White Paper: **November 2021**

# VMware Tanzu Application Platform with TriggerMesh

## Connect AWS Services to on-premises machine learning

vmware®

## Table of contents

"TriggerMesh SAWS makes it simple to integrate data services on AWS with the rich set of computing platforms available on VMware Tanzu. AWS event sources can provide a vital bridge between cloud-centric designs and the unique network connectivity of on-premises or edge locations."

Evan Anderson
Tech Lead/Co-Founder of Knative
VMware

### VMware Tanzu Application Platform with TriggerMesh

- Tanzu Application Platform is a multi-cloud developer experience built for any K8s distribution.
- Cloud Native Runtimes is the runtime layer for Tanzu Application Platform.
- TriggerMesh Sources for AWS is an open source project that provides cross-cloud triggers from AWS to serverless workloads.
- Knative is an open source serverless framework built on top of Kubernetes.

## Introduction

Modern applications require the infrastructure they are deployed on to react automatically to their needs. That can mean provisioning or deprovisioning resources in response to traffic changes, triggering workloads based on internal or external events, or even just sending a Slack message notifying a developer of a change in their app or its environment. Kubernetes is a powerful abstraction that makes infrastructure dynamic and more attuned to application needs. Yet development teams face a substantial learning curve when getting started, and they must perform additional work to enable deployment of applications with event-driven architectures on the platform.

Cloud Native Runtimes, the runtime layer of VMware Tanzu Application Platform, is a set of capabilities that enables developers to leverage the power of Kubernetes for serverless use cases without first having to master the Kubernetes API. Based on the open source Knative project, Cloud Native Runtimes can be used by itself or in concert with other Tanzu Application Platform capabilities to quickly get modern, cloud native applications with event-based architectures up and running on Kubernetes, regardless of a developer's level of experience with the platform.

TriggerMesh makes it easy for customers to bring Amazon Web Services (AWS) data and events into their cloud native infrastructure and use it not only on AWS, but also in their other public, private, or hybrid clouds as needed. Event-driven applications are built on top of their data and events. With TriggerMesh, customers' data becomes part of the fabric used by developers and DevOps teams to integrate that data across every platform.

## Tanzu Application Platform

Tanzu Application Platform includes foundational elements that enable developers to quickly begin building and testing applications, regardless of their familiarity with Kubernetes. Operations teams can create application scaffolding templates with baked-in security and compliance guardrails, making those considerations mostly invisible to developers. Starting with those templates, developers turn source code into a container and a declarative manifest and get a URL to test their application in minutes. Once the container is built, updating it happens automatically every time there's a new code commit or dependency patch. And connecting to other applications and data—regardless of how they're built or what kind of infrastructure they run on—has never been easier thanks to an internal API management portal.

## TriggerMesh sources for AWS

AWS is the most popular public cloud platform and is used by enterprises as both their primary source of infrastructure and as part of their multi-cloud and hybrid strategies. Working across multiple platforms is the reality for many organizations. Part of the TriggerMesh platform, the open source TriggerMesh Sources for Amazon Web Services (SAWS) integrates AWS events and data with practically any other platform, including Tanzu Application Platform and Cloud Native Runtimes from VMware.

SAWS makes it easy to consume events from AWS services and send them to any system running anywhere—on-premises workloads, SaaS applications, cloud services, serverless functions, and even other public or private clouds. TriggerMesh SAWS supports the most popular AWS services, including CloudWatch, CodeCommit, DynamoDB, Kinesis, Simple Notifications Service (SNS), Simple Queue Service (SQS), and Simple Cloud Storage (S3).

## Tanzu Application Platform image processing with TensorFlow on AWS S3 Object sent via TriggerMesh

TensorFlow is an open source platform for developing, training, and building machine learning (ML) models and applications. In the recent general availability announcement of Cloud Native Runtimes for VMware Tanzu, TriggerMesh and VMware demonstrated an event-driven application architecture for processing images of vehicle license plates with TensorFlow and recording the results in real time using Google Sheets. This represents a real-world use case of automatically recognizing license plates at a loading terminal or perhaps a toll road and using that as a source for automating security or billing as appropriate.

In the demonstration, images of automobiles are uploaded to Amazon S3. This triggers an event that is captured by TriggerMesh and passed to the Tensformation service running on Tanzu Cloud Native Runtimes (shown in Figure 1 below as Step 1), which downloads the image from S3, encodes it for processing, and sends it to the TensorFlow Server (Step 2). The TensorFlow Server processes the image, recognizes whether it has a license plate, and returns the result (Step 3), which is then sent to the Label_analyser, a different service (Step 4), which then returns the license plate number as an event and records it in a Google Sheet (Step 5).

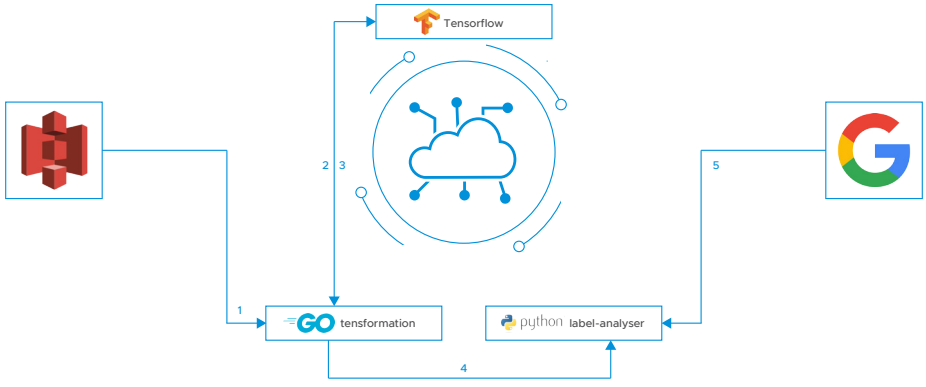


**Figure 1:** Image processing of an S3 Object on a different service, with results published to a Google Sheet.

In this real-world example, the TriggerMesh cloud native integration platform ties these services and events together, running on top of the Cloud Native Runtimes. The Amazon events from S3 are generated when an image is uploaded into the S3 bucket and TriggerMesh's SAWS source creates an S3:ObjectCreated CloudEvent and sends it to the Tensformation service. The Tensformation service is running as a Cloud Native Runtimes serverless application, allowing it to automatically scale up and down as the number of uploaded S3 images increases and decreases accordingly. The Tensformation service downloads the image from S3, Base64 encodes it and pushes it to the TensorFlow tf-inference-server service, which is also running as a serverless application under Cloud Native Runtimes for VMware Tanzu.

```
apiVersion: sources.triggermesh.io/v1alpha1
kind: AWSS3Source
metadata:
  name: my-bucket
spec:
  # This should point to your aws s3 bucket IE: arn:aws:s3:us-west-2:925906438773:demobkt-triggermesh
  # The default ARN from AWS needs the region and your account ID added to it as above:
  # https://docs.triggermesh.io/sources/awss3/#amazon-resource-name-arn
  arn: ""
  credentials:
    accessKeyID:
      valueFromSecret:
        key: AWS_ACCESS_KEY
        name: aws
    secretAccessKey:
      valueFromSecret:
        key: AWS_SECRET_KEY
        name: aws
  eventTypes:
  - s3:ObjectCreated:*
  sink:
    ref:
      apiVersion: eventing.knative.dev/v1
      kind: Broker
      name: default
```

**Figure 2:** Configuration of AWS S3 event source from TriggerMesh.

The TensorFlow tf-inference-server service is running the Automatic Number Plate Recognition model as a container image provided by VMware Engineer Myles Gray. The TensorFlow Server is resource intensive, but Cloud Native Runtimes ensures that the balance between the number of Tensformation application instances consuming S3 images and the tf-inference-server application instances applying the ML model to those images rapidly scales to handle the workload. When all of the S3 images have been processed, Cloud Native Runtimes frees the computing resources. The results of the TensorFlow ML model are sent to the Label-analyzer serverless application as a CloudEvent, which in turn publishes the identified license plate number to a Google Sheet.

Although this example is quite technical, it provides a template for processing events streamed from AWS and incorporating those into an AI/ML process running in a Kubernetes cluster with Tanzu Application Platform. The automatic scaling provided by Cloud Native Runtimes allows the application to map to actual traffic patterns, only consuming the resources needed as demand arises.

"Our partnership with VMware opens the door to multicloud event-driven integration within the Tanzu platform. It brings together several years of open source work to make applications more scalable in response to changes in the computing environment. Customers will be able to use their favorite cloud and leverage their familiar VMware environment at the same time."

Sébastien Goasguen
CTO and Co-Founder
TriggerMesh

The destination for the processed image events could be anywhere that the data is needed using the data sinks provided by TriggerMesh, like Elasticsearch, Splunk, Google Storage, and others.

## Benefits of running event-driven Kubernetes workloads on Tanzu Application Platform with TriggerMesh

Leveraging the combined capabilities of Knative and TriggerMesh event sources packaged in Cloud Native Runtimes, teams can build event-driven applications that automatically scale within their VMware infrastructure. Those event-driven applications can be described using a declarative API familiar to Kubernetes operators. Events coming from AWS services are readily available, making it simple to react to changes in AWS services and to perform actions within VMware Tanzu.

## Summary

Cloud Native Runtimes for VMware Tanzu is generally available now and comes with the TriggerMesh Sources for AWS included. Tanzu Application Platform is available today in beta. As demonstrated in this brief, this enables teams to react to changes in their AWS S3 buckets and dynamically scale containerized workloads running within VMware Tanzu Advanced.