Architecture

By: The Cloud Platform Team at Statistics Canada

Introduction

This document represents a high-level overview of how Drupal should be architected in the cloud to support any of the Government of Canada procured cloud service providers (AWS, Azure, and GCP), as well as upcoming ones such as IBM and Oracle. This would also work in an on-premise environment with the appropriate infrastructure.

A key mandate is to follow the Open Source Directive as given by the Treasury Board Secretariat (C.2.3.8) which states where possible, use open standards and open source software first. Additionally, where possible expose all functionality as services (RESTFul) and leverage microservices via a containerized approach (C2.3.10). We are leveraging a microservices design pattern utilizing immutable images through containerization running on Kubernetes with a platform that has been built and open sourced by Statistics Canada. The platform will be discussed briefly to provide context but the bulk of the document discusses how Drupal is installed on top of it.

We are also following the DevSecOps approach documented by the United States Department of Defense (DoD), who has adopted a modern cloud infrastructure and development workflow as part of their "Platform One" team.

Kubernetes

> The base of the platform - Kubernetes is the first graduate of the CNCF (Cloud Native Computing Foundation).

Kubernetes orchestrates the computing, networking, and storage infrastructure on behalf of user workloads. It assigns workloads and resources to a series of nearly identically-configured machines.

Kukbernetes supports worloads running anywhere, from IoT devices, to private cloud and all the way to public cloud. This is possible due to Kubernetes' pluggable architecture, which defines interfaces that are then implemented for the different environments. Kubernetes provides an Infrastructure as Code environment defined through declarative configuration. Because Kubernetes abastracts away the implementation of the computing environment, application dependencies such as storage, networking, etc., applications do not have to concern themselves with these differences.

Kubernetes is backed by a huge (tens ouf thousands) and vibrant growing community, consisting of end users, business, vendors and large cloud providers.

Key Points

The Statistics Canada architecture brings many benefits to the Government of Canada:

- Support for hybrid workloads (Linux and Windows), deployed using the same methodology
- Abstraction of underlying hardware ("cattle rather than pets") enabling an automated, highly-available and scaleable infrastructure for microservices

- Declarative configuration enabling Infrastructure as Code allowing for deployment automation, reproducibility and re-use
- Constructs to support advanced deployment patterns (blue/green, canary, etc.) enabling zero-downtime deployments
- Platform-level tools for traffic handling (routing, error recovery, encyption, etc.), monitoring, observability and logging, secrets management; avoiding duplication across applications in the environment

Kubernetes is supported across all cloud service providers (fully managed and self managed), preventing vendor lock-in. Managed offerings are available from Google, IBM, Azure, Digital Ocean, Amazon, Oracle and more. The choice whether to roll your own, using a managed service or a Platform as a Service (PAaS offering) is up to the organization to decide based on their requirements and risks. The Statistics Canada platform stays as close to Open Source version and tools as possible in order to remain compatible with the different Kubernetes offerings (raw, managed, platform, etc.). Other offerings, such as OpenShift, Pivotal and VMWare PKS, are more opiniated in hwo to use them, provinding more guard rails and being more locked down out of the box.

Government

Kubernetes is being actively investigated and/or used by many departments across the Government of Canada. Departments are starting to collaborate more and work together towards a common, well-vetted solution and this is why we have have Open Sourced our platform on the GC Accelerators hoping to foster this collaboration and form a community of practice.

Provided below is the Terraform (Infrastructure as Code) necessarily to install the Azure Kubernetes Service Infrastructure as well as configure with optional platform components (RBAC, Service Mesh, Policies, etc).

- Terraform for Kubernetes Infrastructure: https://github.com/canada-ca-terraform-modules/terraform-kubernetes-aks
- Terraform for Kubernetes Platform: https://github.com/canada-ca-terraform-modules/terraform-kubernetes-aks-platform

To highlight that this solution can run on any cloud service provider, we are currently working on support for the IBM Kubernetes Services:

- Terraform for Kubernetes Infrastructure: https://github.com/canada-ca-terraform-modules/terraform-kubernetes-iks
- Terraform for Kubernetes Platform: https://github.com/canada-ca-terraform-modules/terraform-kubernetes-iks-platform

IBM's managed Kubernetes offering is entirely managed, reducing operational requirements. This is similar to Google Cloud Run, and is where other other cloud providers (such as Azure) is moving towards.

Statistics Canada

Statistics Canada leverages Kubernetes currently on the Azure Kubernetes Managed Service for the bulk of its platform services and all net new applications are first considered whether they can run on our managed cloud native platform.

Statistics Canada is looking at levearing Kubernetes for a majority of its workloads. The platform that we are sharing was initially developed at Statistics Canada, and enables a quick onboarding of new development teams, reduced application scope via platform-level pluggable services such as:

- Artifactory/X-Ray: Providing package management and CVE vulnerability scanning
- Istio: Service mesh, providing better traffic management and observability, as well as automatic mutually-verified TLS encryption
- Elastic: Provides both cluster-level and application-level log aggregation and searching
- Prometheus/Grafana: Provides real-time cluster and application metrics
- Velero: Provides Kubernetes state and data backups
- Hashicorp Vault: Provides static and dynamic secret management
- Cert-manager: Using Let's Encrypt, providing automatic certificate issuance and renewal
- Open Policy Agent: Provides enforcement of defined policies (for example, using images from authorized sources only)

Canadian Digital Services

The Canadian Digital Services (CDS) develops all of their applications to run in a Kubernetes-based environment on Google Cloud. They are actively involved in promoting Kubernetes, DevOps and other modern workflows across the Government of Canada. They are launching tools suchs as GC Notify, their Security Goals compliance reporting tool, and work closely with GC Tools and GC Connex.

Others

Many others across the Government of Canada, municipalities and internal Government agencies are actively and/or investigating Kubernetes. This includes Shared Services Canada (SSC) through their investigation into a managed cloud platform powered by OpenShift (Kubernetes) for use by other departments. The Communication Security Establishment (CSE) is leveraging Kubernetes in their on-premise computing environment, running cloud native technologies in a restricted environment. The Royal Canadian Mounted Police kick started their Kubernetes adoption through the Report a Cyber Crime application developed by CDS.

To highlight a few more:

- The City of Montreal
- The City of Ottawa
- UK Home Office

United States Air Force

The United States Department of Defense (DoD) has Open Sourced all the documentation associated with their modernization initiative for the operation and deployment of infrastructure and applications in environments of all data classifications. The work of the DoD is not limited to just the technology but also the culture changes their undertook to move to a DevSecOps model. The "Platform One" team is responsible for the Cloud One platform and supporting the business and weapon systems.

The DoD was looking to avoid vendor lock-in and empower development teams and enable application deployment into secure environments at a rapid pace. They have defined requirements that enable flexibility while maintaining a strict security posture through their "hardening" process. Their Cloud One system is based

on Kubernetes, Istio, and other Cloud Native Computing Foundation technologies, aligning with the decisions made by Statistics Canada in the development of their architecture.

Lead by the Chief Software Officer, a combination of both Microsoft and Amazon Web Services' cloud platforms has allowed the Air Force to operate at heightened speeds, providing access to cloud capabilities to airmen within days to enable software development on the cloud or leveraging artificial intelligence (AI).

Shared Services Canada and the Royal Canadian Mounted Police are in active talks with the DoD to better understand their environment.

Enterprises

Kubernetes' use extends beyond the government and is used by many large and small organizations. At the hyperscale level, we have Google, GitHub, Reddit, Shopify using Kubernetes. In research and data science, organizations such as CERN use Kubernetes for analyzing data from the Large Hadron Collider. Additionally, large regulated organizations in the financial industry are also embracing Kubernetes, such Bloomberg, Capital One, MasterCard and American Express. Finally, Kubernetes and cloud native is playing role in the rollout of 5G technologies around the world, such as the large contract with AT&T.

Drupal on Kubernetes

A managed Drupal Platform as a Service is a strong candidate to take advantage of what the Statistics Canada platform offers. The design enables a quick onboarding of new workloads through the repeatabile deployment methodology provided by Kubernetes.

Components

The components are individually described below. The components perform the same function in both the content staging and public environments.

Kubernetes

Recommendation: Kubernetes

Kubernetes is the basis of this Drupal platform and is further discussed above.

The whole Drupal application stack can be easily installed in a distributed fashion in minutes using our Helm chart, The chart facilitates a managed service workflow (rolling updates, cronjobs, health checks, auto-scaling, etc.) without user intervention.

- Helm chart: https://github.com/drupalwxt/helm-drupal

Ingress controller

Recommendation: Istio

The ingress controller is responsible for accepting external HTTPS connections and routing them to backend applications based on configuration defined in Kubernetes Ingress objects. Routing can be done by domain and/or path.

Nginx

Recommendation: Nginx

Nginx is an open source web server that can also be used a reverse proxy, HTTP cache, and load balancer. Due to its root in performance optimization under scale, Nginx often outperforms similarly popular web servers and is built to offer low memory usage, and high concurrency.

> Note: It should be noted that Nginx in this model addresses the cache requirements that are needed in Drupal.

Web (PHP-FPM)

Recommendation: PHP-FPM

Drupal runs in the PHP runtime environment. PHP-FPM is the process manager organized as a master process managing pools of individual worker processes. Its architecture shares design similarities with event-driven web servers such as Nginx and allows for PHP scripts to use as much of the server's available resources as necessary without additional overhead that comes from running them inside of web server processes. The PHP-FPM master process dynamically creates and terminates worker processes (within configurable limits) as traffic to PHP scripts increases and decreases. Processing scripts in this way allows for much higher processing performance, improved security, and better stability. The primary performance benefits from using PHP-FPM are more efficient PHP handling and ability to use opcode caching.

Database

Recommendation: MySQL or PostgreSQL

Drupal maintains its state in a database and while supports several types only MySQL or PostgreSQL should be considered. Personally, we would highly recommend PostgreSQL based on the experience we had building / launching the Open Data portal but in the end, both run quite well with minimal operational concerns.

> Our recommendation would be to use a managed database offering from the cloud providers for a production environment. Coupled with a managed file service, this removes all stateful components from the cluster enabling the best application experience possible.

Stateful Assets

Drupal stores generated CSS/JS assets and uploaded content (images, videos, etc.) in a file storage. As the architecture is designed to be distributed, this present some design considerations for us.

Azure Files (CIFS / NFS)

Fully managed file shares in the cloud that are accessible via Server Message Block (SMB) protocol (also known as Common Internet File System or CIFS). Support is provided for dynamically creating and using a persistent volume with Azure Files in the Azure Kubernetes Service.

For more information on Azure Files, please see Azure Files for applications in AKS.

> Note: This is currently our recommended choice as it results in a simpler installation in Azure then relying on an S3 compatible object store discussed below. Similar storage solutions exist with the other cloud providers.

S3 Compatible Object Store (deprecated)

Recommendation: Minio

To support S3 style object storage we will use a Drupal module (s3fs) to store all of these stateful assets inside an object store. The object store can be Amazon S3 itself or another compatible store (e.g., Minio) which can work with both Azure and Google cloud storage.

> This was originally what we inteded to support for managed files. Unfortunately, the integration with Drupal added more operational concerns than desired.

Installation Profile

[https://github.com/drupalwxt/wxt]

Installation profiles provide site features and functions for a specific type of site as a single download containing Drupal core, contributed modules, themes, and predefined configuration. They make it possible to quickly set up a complex, use-specific site. The Canada installation profile is a sub-profile off of WxT which is a sub-profile of Lightning.

Note: The WxT profile is very light and only adds the GCWeb theme and plugins support.

Composer Project

Project for managing your dependencies via Composer.

https://github.com/drupalwxt/site-wxt

Continuous Integration

Currently we are running all of our CI builds using Travis CI but can be easily switched to any of the other options such as CircleCI, VSTS, Jenkins, etc.

- https://travis-ci.org/drupalwxt/wxt
- https://travis-ci.org/drupalwxt/site-wxt

> Note: We are currently moving our CI builds from Travis CI to run on GitHub Actions.

Docker Registry

https://hub.docker.com/r/drupalwxt/site-wxt

Currently we are storing all of our containers in Docker Hub.

These containers are also stored and scanned against Statistics Canada's internal docker registry namely Artifactory which coupled with XRay scans all of our containers against the known CVE's in the wild. Shoud any critical vulnerability be found we ensure that our public containers stored in Docker Hub recieve these fixes / mitigations as well.

Resources

Here are some resources that we recommend should you wish to further explore Kubernetes, Cloud Native and DevSecOps principals.

- Cloud Native Platform for Government (In Development)
- United States Department of Defense
- Digital Academy Course on Cloud Native Development