

Implementation of CMSWEB Services Deployment Procedures using HELM

Aroosha Pervaiz,^a Muhammad Imran,^{a,c,*} Valentin Kuznetsov,^b Panos Paparrigopoulos,^a Spyridon Trigazis^a and Andreas Pfeiffer on behalf of the CMS Collaboration^a

^a*CERN,
1211 Geneva 23, Switzerland*

^b*Cornell University
Ithaca, NY 14850, USA*

^c*National Centre for Physics
Islamabad, Pakistan
E-mail: aroosha.pervaiz@cern.ch, muhammad.imran@cern.ch,
vkuznet@protonmail.com, panos.paparrigopoulos@cern.ch,
spyridon.trigazis@cern.ch, andreas.pfeiffer@cern.ch*

The Compact Muon Solenoid (CMS) experiment heavily relies on the CMSWEB cluster to host critical services for its operational needs. Recently, CMSWEB cluster has been migrated from the VM cluster to the Kubernetes (k8s) cluster. The new cluster of CMSWEB in Kubernetes enhances sustainability and reduces the operational cost. In this work, we added new features to the CMSWEB k8s cluster. The new features include the deployment of services using Helm's chart templates and the incorporation of canary releases using Nginx ingress weighted routing that is used to route traffic to multiple versions of the services simultaneously. The usage of Helm simplifies the deployment procedure and no expertise of Kubernetes are needed anymore for service deployment. Helm packages all dependencies, and services are easily deployed, updated and rolled back. Helm enables us to deploy multiple versions of the services to run simultaneously. This feature is very useful for developers to test the new versions of the services by assigning some weight to the new service version and rolling back immediately in case of issues. Using Helm, we can also deploy different application configurations at runtime.

*International Symposium on Grids & Clouds 2022 (ISGC 2022)
21 - 25 March, 2022
Online, Academia Sinica Computing Centre (ASGC), Taipei, Taiwan****

*Speaker

1. Introduction

The CMS [1] is a general-purpose detector at the Large Hadron Collider (LHC) at CERN, Geneva, Switzerland. It has a broad physics program ranging from studying the Standard Model (including the Higgs boson) to searching for extra dimensions and particles that could make up dark matter.

CMSWEB Kubernetes cluster was instantiated when the migration of CMSWEB cluster from VM clusters to Kubernetes cluster was performed. This CMSWEB cluster manages dozens of services in a consistent fashion. The framework for the deployment is written in bash scripts which enables the deployment and management of these services. To put it into perspective, the CMSWEB enables the independent deployment, management and evolution of services, the simplification of integration and regression testing when these services are updated, and also the development and association of external services that allow for the information to be integrated in a clean, organized way.

Each service in CMSWEB has its own development group that upgrades the relevant service. The CMSWEB development life cycle includes all the hosted services that are actively developed. The CMSWEB team is not responsible for maintaining individual services but only provides the web infrastructure and manages service deployment.

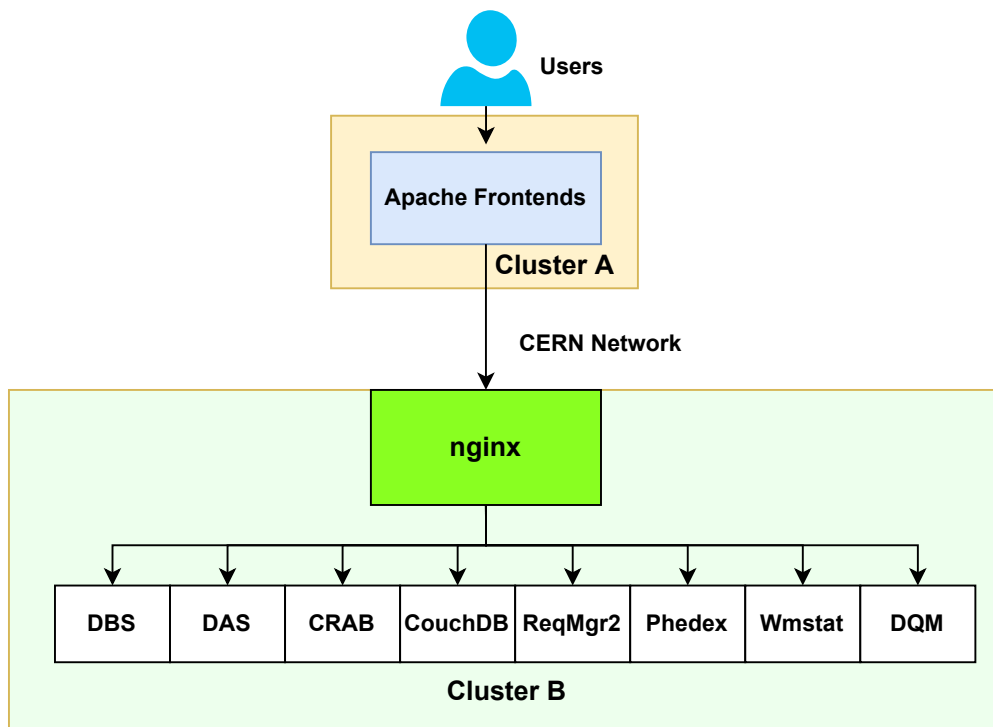


Figure 1: The CMSWEB Kubernetes cluster architecture[2]

When it comes to the architecture of the CMSWEB cluster, it has two layers of services, i.e., frontend and backend services as depicted in Figure 1. The frontend cluster relies on Apache which first provides authentication using X509 certificates and then redirects the request to the backend

cluster. On the backend cluster, the ingress controller has basic redirect rules to the appropriate services and only allows requests from the frontend cluster. It is actually independent and individual backend services that carry out the requested and relevant tasks. The reason for selecting two clusters is related to the structure of Kubernetes, i.e., the mapping of services (ingress => services). Using Apache as an authentication layer requires keeping redirect rules from Apache to other nodes. As the Kubernetes host network does not allow for replicas, the 2 cluster architecture was needed. Plus, it allows independent maintenance of individual clusters. Kubernetes supports namespaces to allow service separation within a cluster. In the Kubernetes cluster, we combined services into namespaces according to the developers' group. More details on these services can be found at [2].

Each of the CMSWEB services are currently deployed through manifest files in yaml using Kubernetes's *kubectl* tool. However, it becomes increasingly tedious to manage all those manifest files because of all the different sorts of objects that Kubernetes has to offer — such as configMaps, services, pods, persistent volumes — in addition to the number of releases you need to manage. Furthermore, it is not possible to run the multiple versions of the services without changing the service manifest files and Nginx ingress rules. It is not easy to rollback, upgrade, and configure services using *kubectl* tool. We also need to share multiple files for deployment configurations with service developers and the developers also need to have knowledge about the deployments' configurations and Kubernetes. This is where Helm charts come to help us. The usage of Helm simplifies the deployment procedure and no expertise of Kubernetes are needed anymore for service deployment. Helm packages all dependencies and services are easily deployed, updated and rolled back. Helm enables us to deploy multiple versions of the services to run simultaneously. This feature is very useful for developers to test the new versions of the services by assigning some weight to the new service version and rolled back immediately in case of issues. Using Helm, we can also deploy different application configurations at run-time.

In this paper, Section 2 gives an overview of the Helm charts and describes how they are used. Section 3 discusses the canary releases using Nginx ingress that is additionally included in those Helm charts. Section 4 provides brief detail about documentation and support. Finally, we conclude in Section 5.

2. HELM Charts

Helm [3] is a combined effort of Deis (which is now a part of Microsoft) and Google. It became a part of the Kubernetes 1.4 release in 2016 and has been easing the life of developers ever since. In April 2020, Helm was classified as a CNCF project in April 2020 which further added to its popularity. It is one of the open-source projects that wish to address the complexities of Kubernetes. Helm, a packet manager for Kubernetes, makes it easier to deploy, upgrade and roll back the Kubernetes deployments by simplifying the processes to a single CLI command.

The package format of Helm is called chart, which is a collection of files describing Kubernetes resources. Helm charts enable the definition and installation of even the most complex Kubernetes applications. These Helm charts can be customized at runtime that also automates the upgrading and rolling back of releases. Helm also maintains the history of previous versions of the releases, so if things go wrong, we can roll back. Helm also allows us to package all the related Kubernetes

resources. Hence, it becomes effortless to package, distribute, and download and install these Helm charts.

Helm also allows for the same chart to be deployed multiple times across multiple environments with different values, unlike Kubernetes. Pertaining to all the benefits that Helm offers, we converted our Kubernetes manifest files of CMSWEB services into Helm charts and stored them in a separate Helm repository. We believe that it will reduce the complexity of modification and upgrading of services. We will also have the entire history of our previous versions of the releases which will make it easier to track back. By creating a repository for the Helm charts of the services, the users will be easily able to install all the services locally. They can also easily update the Helm repository and leverage it to install their desired services.

2.1 Repository

Helm charts can be hosted at an online repository. In order to host our Helm charts, and make them available to the CMSWEB team, they are being served at two locations: [4] and [5]. One can fetch these repositories, push Helm charts to these repositories, and update their locally cached repositories to match the changes in the aforementioned repositories.

2.2 Contact Command and Control servers

Helm uses a packaging format known as chart. A chart is a collection of interrelated files inside a directory. The name of the directory reflects the name of the chart. A typical chart contains a *Chart.yaml* file and a *templates* directory. This *template* directory then contains the interrelated deployments for the same chart. The *Chart.yaml* and *templates* directory is a required prerequisite for a directory to be deemed as Helm chart. It is shown in Figure 2.

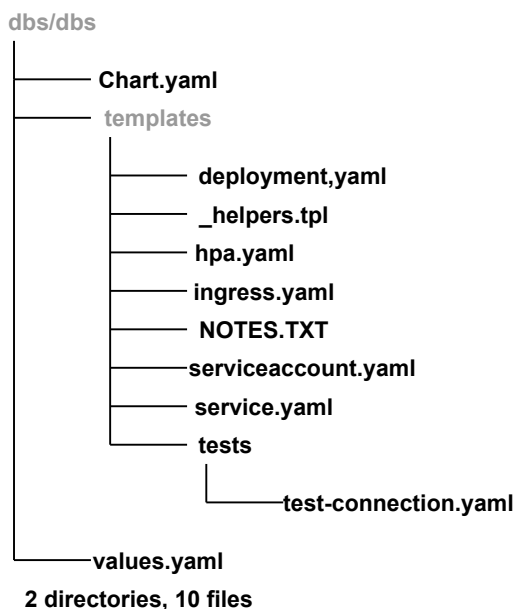


Figure 2: A typical Helm chart resources.

The brief description of each file is shown below:

- Chart.yaml - A YAML file containing information about the chart.
- A templates directory: templates/ - A directory of templates with Kubernetes manifest files or that will generate valid Kubernetes manifest files when combined with values.yaml.
- LICENSE - A plain text file containing the license for the chart
- README.md - A human-readable README file
- values.yaml - The default configuration values for this chart
- templates/NOTES.txt - A plain text file containing short usage notes
- template/_helpers.tpl – template helpers that you can re-use throughout the chart
- We can add more manifest files in the templates directory corresponding to our needs for more Kubernetes objects.

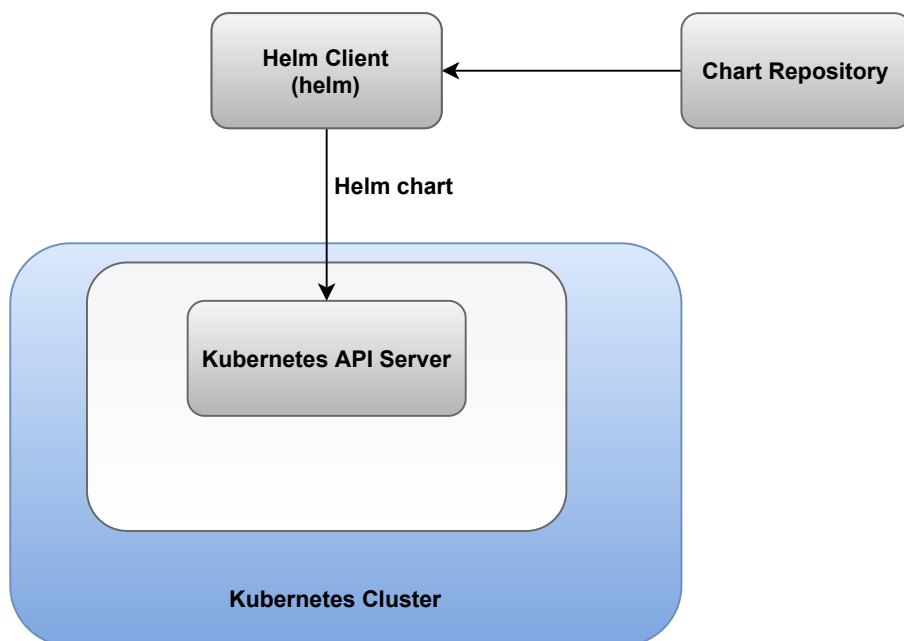


Figure 3: Helm client connects and installs the helm chart to the Kubernetes cluster by interacting with Kubernetes API Server.

2.3 Helm Working Scenario

Helm is an executable that is implemented in two components:

- Helm Client: It is a command line client for the end users of Helm. It is responsible for development of local charts, managing charts, interfacing with the Helm library, and managing releases. It interacts with the Kubernetes API server to install the helm chart to the Kubernetes cluster.

- Helm Library: It provides logic for executing Helm operations. It is a standalone library which has all the logic needed for Helm operations to execute.

Helm client connects and installs the helm chart to the Kubernetes cluster by interacting with Kubernetes API Server as shown in Figure 3.

2.4 Updating and Maintaining the Helm Charts

The Helm charts are updated, maintained and updated through CI/CD pipeline. If a Helm chart has been created or updated the chart and needs to be uploaded to our GitHub repository, you just need to upload the updated code to the *helm* directory in the repository. It will automatically add the release at [4].

2.5 Helm Versioning System

Helm uses semantic versioning to name the charts through which these charts are identified. As depicted in Figure 4, the semantic versioning has three components, MAJOR version, MINOR version, PATCH version. It is in the format MAJOR.MINOR.PATCH.

Given this format, we increment:

1. MAJOR version is changed when we make changes to API that are incompatible to the previous versions.
2. MINOR version is changed when such functionality is added that is backwards compatibility.
3. PATCH version is changed when we fix backwards compatible bug fixes.

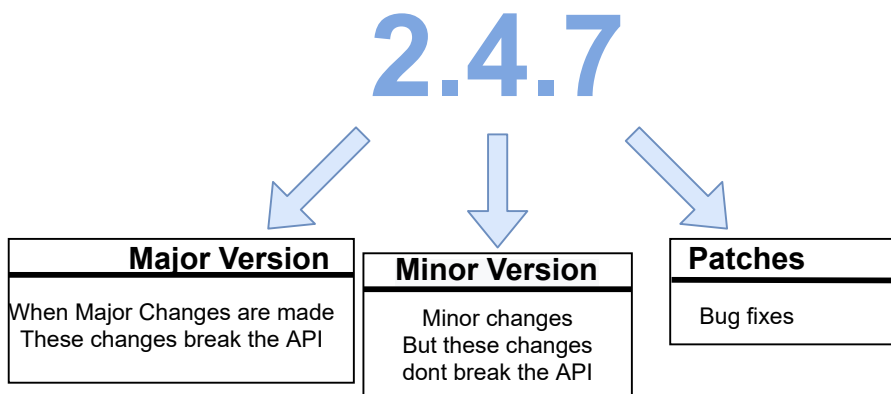


Figure 4: Helm Versioning System uses semantic versioning to name its charts. It comprises of three parts: major, minor, and patch version.

2.6 Usage of Helm charts in Different Environments

Currently, there are three environments available for CMSWEB services, namely, production, pre-production and development environments. The Kubernetes manifest files needed to be commented at some parts, and uncommented on others, to be deemed suitable to work in these different

environments. So, we devised a way through Helm, to do this all with just one variable. The users need to set the environment during the installation of the Helm chart through CLI, and this environment can be upgraded, rolled back and reset using CLI. More detail on this implementation can be found at [6] and [7].

3. Canary Releases Using Nginx Ingress

Our production cluster hosts all the major services required by the CMS for operational needs. Therefore, we cannot risk the unavailability of our Kubernetes cluster, or any particular service. In order to make our system more scalable, we used and implemented the concept of *canary release* provided by Nginx ingress. Using this technique, we can redirect some percentage of the traffic to the new, upgraded releases without having to modify or delete the existing release. This mechanism is shown in Figure 5

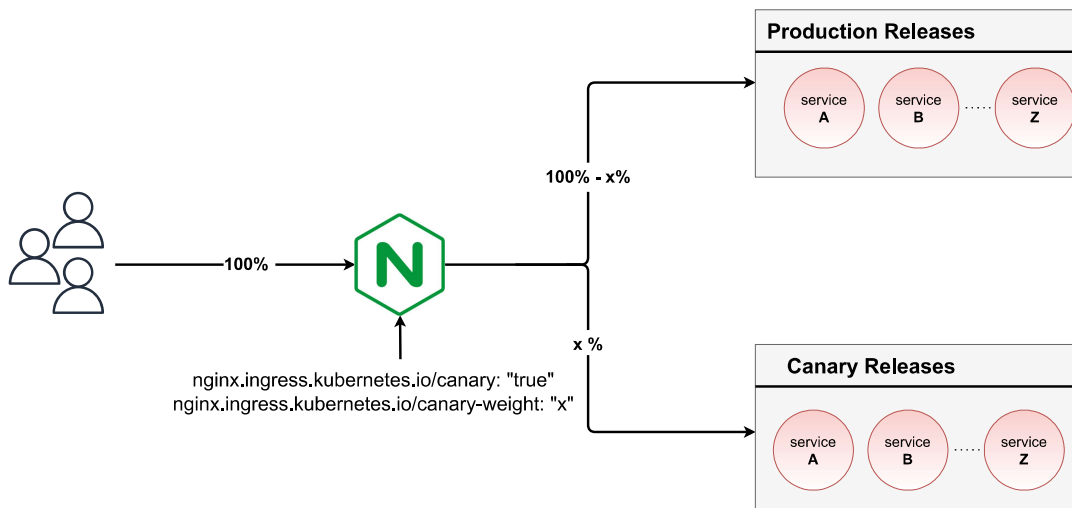


Figure 5: A depiction of how ingress divides the traffic between canary and production releases. As it can be seen, we assign a value to X variable; X% of the traffic goes to the canary releases, while the rest of the traffic stays assigned to the production releases.

Canary releases are a way to release updates to a shifting percentage of the total user base. If there is a need to deploy a new code, or update the image, this newer release can be installed as a canary release. In case of any bug, it will only affect a small percentage of the users before we can detect and rollback.

4. Documentation and Support

Documentation plays an important part in familiarizing others of the work you have done and increases the usability of one’s code by fourfold. For this reason, we have documented the entire

procedure along with details of how my work can be used. This can be found at [6]. If there is a need to familiarize oneself further with how Helm works, the official Helm documentation can be found at [8]. The code for all these Helm charts and the Nginx functionality can be found at the repository [9].

5. Conclusion

In this paper, we presented our work regarding addition of new features to the CMSWEB Kubernetes cluster. The new features include the deployment of services using Helm's chart templates and the incorporation of canary releases using Nginx ingress weighted routing that is used to route traffic to multiple versions of the services simultaneously.

The usage of Helm simplified the deployment procedure and no expertise of Kubernetes is needed anymore for service deployment. Helm packages all dependencies and services are easily deployed, updated and rolled back. Helm enables us to deploy multiple versions of the services to run simultaneously. This feature is very useful for developers to test the new versions of the services by assigning some weight to the new service version and rolled back immediately in case of issues.

References

- [1] R. Adolphi et al., *The CMS experiment at the CERN LHC*, *Jinst* **803** (2008) S08004.
- [2] M. Imran, V. Kuznetsov, K.M. Dziejziniewicz-Wojcik, A. Pfeiffer, P. Paparrigopoulos, S. Trigazis et al., *Migration of cmsweb cluster at cern to kubernetes: a comprehensive study*, *Cluster Computing* **24** (2021) 3085.
- [3] "Helm." <https://helm.sh/>.
- [4] "CMSKubernetes repository for releases." <https://github.com/dmwm/CMSKubernetes/releases>.
- [5] "CERN Registry." <https://registry.cern.ch/chartrepo/CMSWEB/index.yaml>.
- [6] "CMSWeb Documentation ." <https://cms-http-group.docs.cern.ch/Helm/deploy-services/>.
- [7] "CMSWEB New Features." <https://indico.cern.ch/event/1075972/contributions/4525202/attachments/2309994/3954006/Helm%20CMSWEB.pdf>.
- [8] "Helm Documentation ." <https://Helm.sh/docs/>.
- [9] "Set of instructions and examples to deploy CMS data-services to Kubernetes cluster.." <https://github.com/dmwm/CMSKubernetes>.