Troubleshooting Edge Compute Errors

Overview

When you first create a Workload, we recommend reviewing the **Status** column in the Instances section.

At first, the Status of your newly-created Instance should appear as *Starting:*

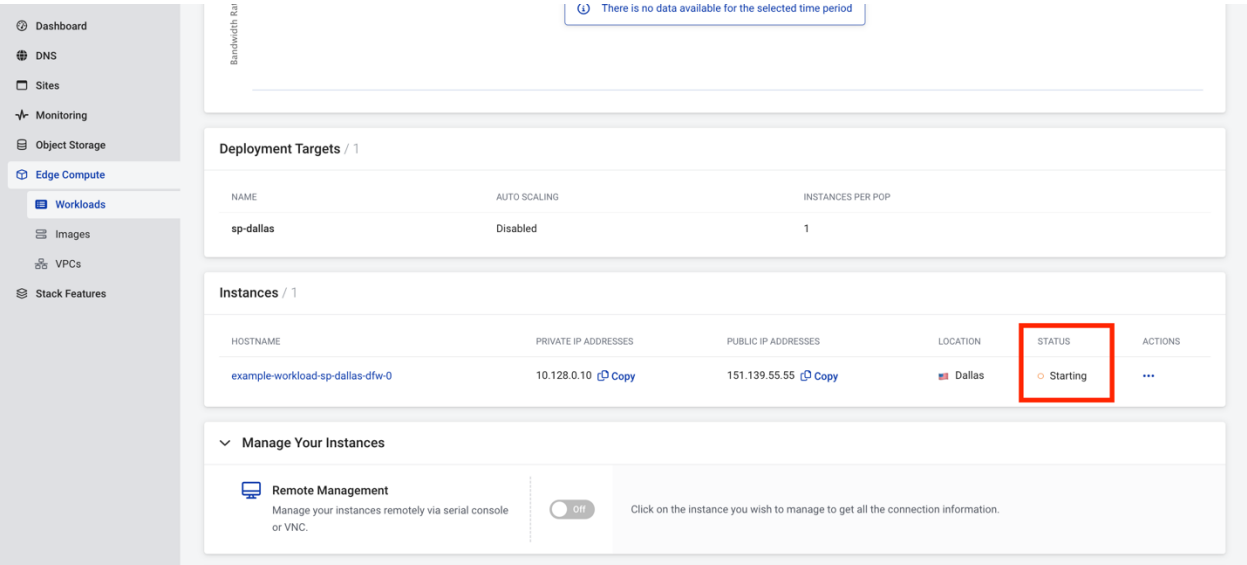

After a few minutes, the Status will update to *Running:*

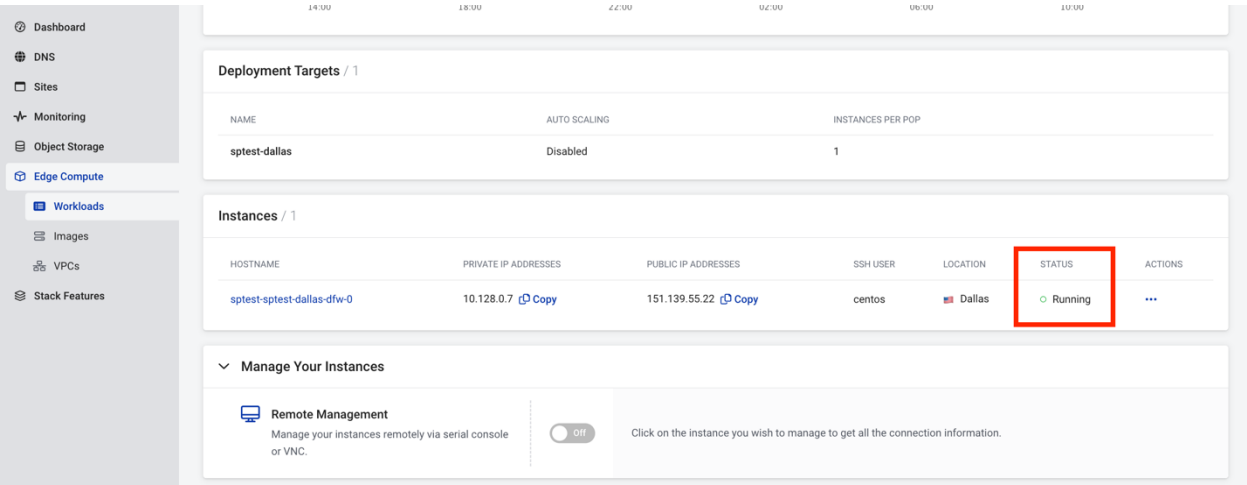

If you notice that the Status does not change from *Starting* to *Running* after a few minutes, there may be an issue with your Workload's configuration.

You can verify whether there is an issue with the Workload's configuration on your end by clicking on the name of the Instance. If there is an issue, a red banner containing an error will appear at the top of the page. In the following example, we are receiving an error because the name of the image is invalid:

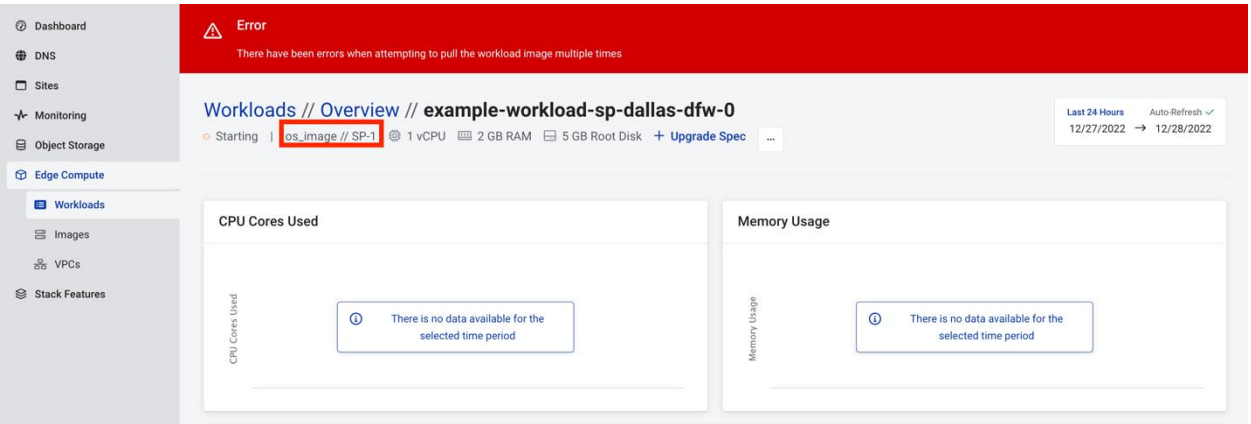

This article will serve as a troubleshooting guide that will explain why these errors occur and how to resolve them.

If one of these errors persists after 7 days and you have not attempted to resolve it, then we will delete the workload on the 7th day.

CrashLoopBackOff

The **CrashLoopBackOff** error occurs when an image or container can't maintain an operational state, resulting in it immediately erroring or exiting. This may happen with container Workloads that don't have permanent operations, such as OS images. Commands need to be provided to the containers to ensure a steady run-state before they can be used.

The cause behind this error varies from Instance to Instance and is caused by the software running within it (or lack thereof).

We recommend consulting your container logs to determine the cause of the container error. Container logs can be accessed via the API, or in the Control Portal.

ImagePullBackOff

The **ImagePullBackOff** error occurs when the incorrect registry credentials have been provided, preventing our platform from being able to pull the specified image from the provided source.

Our platform defaults to using dockerhub if a separate source is not specified. Dockerhub returns the ImagePullBackOff error if:

1. The individual server has exceeded dockerhub's anonymous, IP-based rate limit.
2. The image requested does not exist on dockerhub (may be due to a typo in the image name or path).

We recommend confirming that the image name and registry credentials are correct.

We also recommend verifying the specified image sources, image paths and image names. If these look correct, we suggest utilizing private pull credentials to bypass dockerhub's anonymous rate limit.

InvalidImageName

The InvalidImageName error occurs when there is either a misspelled or non-existent image.

We recommend confirming that the name of the image is correct and publicly available from one of the following repositories:

- registry.centos.org
- Red Hat Ecosystem Catalog
- Docker: Accelerated, Containerized Application Development
- Quay

Using the API

You can use the API to retrieve more details about your Instance and any errors associated with it. Using our first example, where we provided the incorrect image name, the response to the Get a workload instance call would appear as follows:


```
{
  "instance": {
    "stackId": "a3c1310f-4279-4eb3-9c08-b97b5d6f7ed1",
    "workloadId": "56ca88a5-a5af-4ffe-ab8f-1ddf527b4ce3",
    "id": "7cdbd752-616c-41a4-98d1-23a16c4d85d2",
    "name": "example-workload-fakedallas-dfw-0",
    "metadata": {
      "labels": {
        "workload.platform.stackpath.net/deployment-scope": "dfw",
        "workload.platform.stackpath.net/instance-geometry": "SP-1",
        "workload.platform.stackpath.net/stack-id": "a3c1310f-4279-4eb3-9c08-b97b5d6f7ed1",
        "workload.platform.stackpath.net/stack-slug": "edgecompute-20190205-3cb7de",
        "workload.platform.stackpath.net/target-name": "fakedallas",
        "workload.platform.stackpath.net/workload-id": "56ca88a5-a5af-4ffe-ab8f-
1ddf527b4ce3",
```

```
    "workload.platform.stackpath.net/workload-slug": "example-workload"
   },
   "version": ""
  },
  "location": {
   "city": "Dallas",
   "cityCode": "DFW",
   "subdivision": "Texas",
   "subdivisionCode": "TX",
   "country": "United States",
   "countryCode": "US",
   "continent": "North America",
   "latitude": 32.78014,
   "longitude": -96.800453
  },
  "phase": "STARTING",
  "ipAddress": "10.128.0.10",
  "externalIpAddress": "151.139.55.55",
  "createdAt": "2022-12-28T20:45:58Z",
  "networkInterfaces": [
   {
    "network": "default",
    "ipAddress": "10.128.0.10",
    "ipAddressAliases": [
     "151.139.55.55"
    ],
    "gateway": "10.128.0.1"
   }
  ],
  "resources": {
   "requests": {
    "cpu": "1",
    "ephemeral-storage": "5Gi",
    "memory": "2Gi"
   },
   "limits": {
    "cpu": "1",
    "ephemeral-storage": "5Gi",
    "memory": "2Gi"
   }
  },
  "containers": {
   "container-0": {
    "image": "fake_image",
```

```
      "command": [],
      "resources": {
        "requests": {
          "cpu": "1",
          "ephemeral-storage": "5Gi",
          "memory": "2Gi"
        },
        "limits": {
          "cpu": "1",
          "ephemeral-storage": "5Gi",
          "memory": "2Gi"
        }
      },
      "volumeMounts": []
    }
  },
  "containerStatuses": [
    {
      "name": "container-0",
      "phase": "STARTING",
      "waiting": {
        "reason": "ImagePullBackOff",
        "message": "Back-off pulling image \"fake_image\""
      },
      "ready": false,
      "restartCount": 0,
      "containerId": ""
    }
  ],
  "virtualMachineStatuses": [],
  "reason": "",
  "message": ""
 }
}
```


If you take a closer look at **containerStatuses**, you will see the error preventing the Container from running along with a message. In this case, we are receiving the **ImagePullBackOff** error, meaning we provided an incorrect image name.


```
"containerStatuses": [
    {
      "name": "container-0",
      "phase": "STARTING",
      "waiting": {
        "reason": "ImagePullBackOff",
```

```
      "message": "Back-off pulling image \"fake_image\""
    },
```


If you need further assistance with creating a new Workload, please refer to our Create and Manage Virtual Machines, Containers and Workloads guide, or reach out to your seller.