

vultr_inference

Manages Vultr Serverless Inference subscriptions for deploying and running machine learning models in the cloud.

Contents

01	Introduction	3
02	Example Usage	3
03	Argument Reference	3
04	Attributes Reference	3
05	Import	4

vultr_inference

Introduction

Provides a Vultr Serverless Inference resource. This can be used to create, read, modify, and delete managed inference subscriptions on your Vultr account.

Example Usage

Create a new inference subscription:

HCL

```
resource "vultr_inference" "my_inference_subscription" {  
  label = "my_inference_label"  
}
```

Argument Reference

The following arguments are supported:

- `label` - (Required) A label for the inference subscription.

Attributes Reference

The following attributes are exported:

- `id` - The ID of the inference subscription.
- `date_created` - The date the inference subscription was added to your Vultr account.
- `label` - The inference subscription's label.

- `api_key` - The inference subscription's API key for accessing the Vultr Inference API.

Import

Inference subscriptions can be imported using the subscription's `ID`, e.g.

```
terraform import vultr_inference.my_inference_subscription b6a859c5-b299-49dd-8888-b1abbc517d08
```



VULTR

