Graph.Build RDF Data Synthesizer Usage Instructions

Usage instructions

These usage instructions will help ensure the successful configuration and running of the Graph.Build RDF Data Synthesizer.

Using the 1-Click Quick Create Stack setup

Using a cloudformation template, which can be found here, we can utilise a Quick Create Stack to launch the Graph.Build RDF Data Synthesizer with 1-Click.

1. Launch the product via 1-Click Quick Create Stack (more information on the Cloudformation template can be found below)

https://us-east-1.console.aws.amazon.com/cloudformation/home?region=us-east-1#/stacks/create/review?templateURL=https://data-lens-cloudformation.s3.us-east-1.amazonaws.com/RDF-Data-Synthesizer-Templates/rdfizer-only-stack.yml&stackName=RDF-Data-Synthesizer

2. Once deployed, open the SSH tunnel by copying the command found in the Outputs tab on your AWS console
    - This is derived from `ssh -i "<private-key-file>" ec2-user@<bastion-host-dns-name>`
3. Then the application can be accessed at the link specified by the API endpoint key
4. Use the `/config` endpoint to check the application's health

Manually setting up your stack

1. Create an ECS Cluster with an EC2 instance of the desired size (usually t2.medium or t2.large) with an EBS volume of 30GB attached (that is the minimum EBS volume size allowed with ECS)
2. Place the instance into a public subnet in a VPC with an internet gateway
3. Create a security group to allow access to the EC2 instance from your personal ip
4. Create a task definition to run the RDF Data Synthesizer as a container

Note that this is the most basic initial setup. For improved security we recommend creating a service with an application loadbalancer attached that uses your cluster. The load balancer should have a SSL/TLS certificate attached to it. Amazon Cognito can be integrated with the Load balancer to authenticate user requests against an AWS user pool. The task can then be securely run by the service.

Upgrade instructions

To preserve your data during upgrade, please ensure the volume is not deleted. Note: when deleting a stack created by the 1-Click deployment, the data volume will not be deleted

CloudFormation Delivery Information

- All sensitive information is saved within your environment and nothing is shared with Graph.Build
- IAM roles and Policies: The template includes roles for allowing the ECS tasks to access metered usage, upload and down from S3, Download images from ECR, upload logs to Cloudwatch, Create ECS Clusters, Services and Tasks, Create and write to an EBS volume.

- Network Configuration: The template creates a VPC with 2 public subnets. A loadbalancer sits across the 2 subnets and forwards requests to a single EC2 instance that is running the UI and UI backend as part of an ECS cluster. Requests are routed by port to the relevant container.
- AWS resources pricing breakdown:

Resource	**$ hourly**	**$ month**
t2.medium	0.0464	33.87
EBS Volume 30GB	0.0032	2.4
Detailed monitoring	0.0032	2.4
Application LoadBalancer	0.0225	16.43

Monitoring and assessing application functions

1. Navigate to your Amazon EC2 console and verify that you're in the correct region.
2. Choose Instances and select your launched instance.
3. Select the server to display your metadata page and choose the Status checks tab at the bottom of the page to review if your status checks passed or failed.

User Documentation

For further instructions please contact support@graph.build to get access to our user documentation at https://docs.graph.build.