Using RStudio with Memory Machine CE

Summary

Memory Machine CE is a software product from MemVerge that is used to deploy containerized applications in public clouds. Usually, the containers are used to run batch jobs, that is, jobs that run without an attached terminal — results are written to a file (or files) and retrieved when the job completes. RStudio is an integrated development environment (IDE) for the R programming language, which means that an RStudio user engages in an interactive session. This document describes how to use Memory Machine CE to deploy a container that supports an interactive browser-based RStudio session. The container can run on a Spot Instance in which case Memory Machine CE seamlessly migrates the job to a new virtual machine if the Spot Instance is reclaimed.

Memory Machine CE's support for RStudio sessions includes a gateway, which is a reverse proxy that ensures that the RStudio session appears to clients as an IP address that never changes even when the RStudio session moves to a different virtual machine.

Architecture

As shown in the figure, the gateway has two sides: the client side and the server side. The gateway represents all back-end servers as a single IP address on the client side. Multiple servers, hosting the same or different services, can connect to the same gateway. For example, RStudio servers and Jupyter Notebook servers can run simultaneously in a server farm. On the client side, each service, for example, an RStudio session, is represented as an IP address:Port combination.

In the case of a Spot reclaim event, the RStudio session moves to a new virtual machine instance (either Spot or On-demand depending on your policy). The new virtual machine has a different IP address from the reclaimed Spot instance, but to the client this change is transparent — the client session continues without breaking the connection.

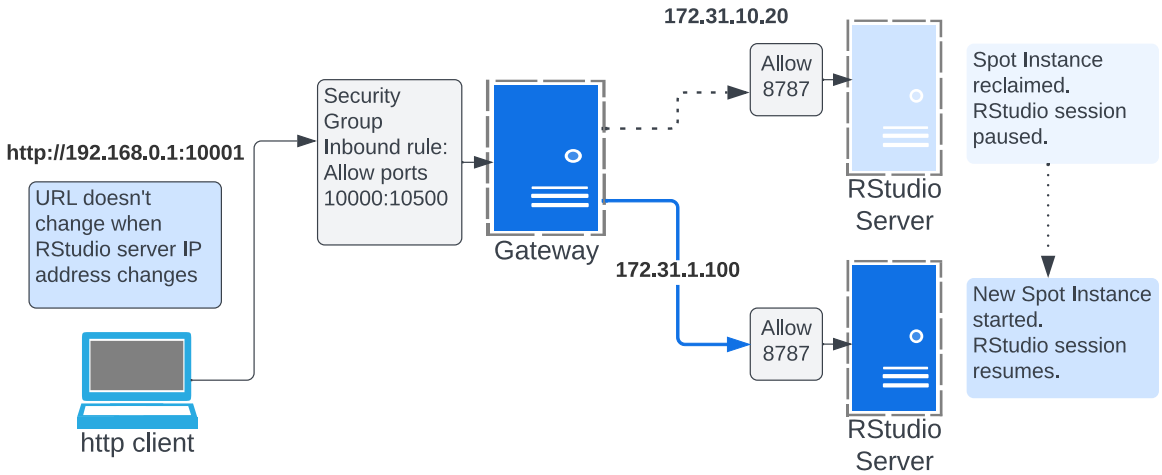

Figure 1. RStudio Session Connected to Gateway

Operation

To the OpCenter, the RStudio session is a job, identified by a Job ID, that continues running until it is manually canceled. This job can run on a Spot Instance or an On-demand Instance. The gateway is a specialized server, started by the OpCenter, that runs on an On-demand Instance and continues running until it is manually canceled. The gateway is identified by a gateway ID.

The gateway and the RStudio server connect via a TCP port (the default port number for RStudio is 8787). The gateway selects the first available port from the range of ports configured on the client side of the gateway to identify the RStudio server to clients. In the example shown, the port 10001 is used by the gateway to identify the RStudio server to clients. When the client opens an http session on the client side of the gateway, the client uses port 10001 to indicate which RStudio server to connect to.

Prerequisites

To use RStudio, and the gateway, with Memory Machine CE, you need the following:

- Memory Machine CE Carmel 2.0 release or later
- Running instance of OpCenter with valid license

Create Inbound Firewall Rules

To access RStudio running on a container, a port on the container is *published*, that is, a port on the container host is mapped to a port on the container. This enables a browser to establish a TCP connection with the container. To open a port on the container host, create an inbound firewall rule by following these steps.

- Log in to your Cloud Service Provider (CSP)'s *Management Console* and follow the steps to create an inbound rule. The steps are different for each CSP. The steps shown here are for AWS (firewall rules are called *security groups* in AWS).
    - Open the *EC2 Dashboard* and click on *Security Groups*.
    - On the upper right, click on *Create security group*. The *Create security group* screen opens.
    - Enter a name and a description for the security group.
    - In the *Inbound rules* box, click on *Add rule*.
    - In the *Port range* box, enter 8787 (this is the default port number for using http to connect to RStudio).
    - In the *Source* box, enter 0.0.0.0/0 (this allows access from any host).

    - Scroll to the bottom of the page and click on *Create security group*.
    - After the security group is created, it appears in the table of *Security Groups*. Note the entry in the column titled *Security group ID*. It has the format *sg-xxxx*.
- Repeat these steps to create an inbound rule to allow access to a range of ports on the gateway. For example, if the gateway assigns ports in the range from 10000 to 10500, create an inbound rule to allow access to all ports in the range 10000 to 10500. Note the *Security group ID* for this rule.

Start the Gateway

- Log in to OpCenter.
- Enter **`float gateway create -h`** to show the options that are available when starting a new gateway. For most deployments, the default options are sufficient except for **`portRange`** and **`securityGroup`**, which must be provided. The port range must be between 10000 and 65535 (inclusive).
- Create a gateway by entering the following command:

  `float gateway create --portRange <minport>:<maxport> --securityGroup <sg-yyyy>`

  where **`<minport>`** is the lowest port number in the range, **`<maxport>`** is the highest port number in the range, and **`<sg-yyyy>`** is the ID for the Security Group that allows access to all the ports in the range. For example:

```
float gateway create --portRange 10000:10500 --securityGroup sg-0fbb6a83983183364
```

- Show all running gateways by entering **`float gateway list`**. For example:

```
float gateway list
+-------+------+----------+--------------+------------+--------+-----+-------+
|   ID  |STATUS| CONFIG   |   PUBLICIP   |  PORTRANGE | START  | JOBS|  COST |
+-------+------+----------+--------------+------------+--------+-----+-------+
| g-Qeh | Ready| 2Cores4GB| 35.175.116.18| 10000-10500|16:10:09| 1   | 14.20 |
+-------+------+----------+--------------+------------+--------+-----+-------+
(edited for clarity)
```

- To delete a gateway, enter **`float gateway destroy -g <gw_id>`** where **`<gw_id>`** is the ID of the gateway to delete.

Download the Memory Machine CE RStudio Script

Memory Machine CE includes a fully-functional CLI, but using a wrapper script simplifies starting and maintaining RStudio sessions.

- Use a browser to download the script here.
- Make the script executable by entering:

```
chmod +x mmce_rstudio.sh
```

- Add the path to **`mmce_rstudio.sh`** to your PATH environment variable.
- To see what options are available to use with **`mmce_rstudio.sh`**, enter the following (default values are also displayed where applicable):

```
mmce_rstudio.sh -h
```

Regardless of the options used, the format must always be the following:

```
mmce_rstudio.sh <option_1> <option_2>...<option_n> <action>
```

where **<action>** is one of **start|list|stop|migrate**.

Start an RStudio Session and Connect to Gateway

- Log in to OpCenter.
- Check that the RStudio image is in the OpCenter repository. Enter the following:

```
float image ls
+------------+----------------------------+--------+-------------+
|      NAME  |                URI         |  TAGS  | ACCESS USER |
+------------+----------------------------+--------+-------------+
| rstudio    | docker.io/memverge/rstudio | latest | opcenter    |
+------------+----------------------------+--------+-------------+
[edited for clarity]
```

If the URI **docker.io/memverge/rstudio** is missing, upgrade the OpCenter software.

- Start an RStudio job using the default values and connect to a gateway (using port 8787) by entering:

```
mmce_rstudio.sh -A <OpCenter_ip_address> -s <sg-xxxx> -g <gw_id> start
```

where **<OpCenter_ip_address>** is the IP address of the OpCenter, **<sg-xxxx>** is the ID of the security group to allow access to port 8787, and **<gw_id>** is the ID of the gateway to connect to.

The RStudio environment can be customized by overriding the default values. For example, use the **-t** option to change the RStudio login password.

If the job is accepted, information about the job is displayed. Note the entries labeled *id:*, *publicIP:*, and *gatewayPort:*. For example:

```
mmce_rstudiogw.sh -A 18.206.202.17 -s sg-0b23d4d7482ddb825 -g g-QehqwJldgogOFzeBKGG9P start
float submit -i rstudio --cpu 2 --mem 4 -e RSTUDIO_USER=rstudio -e
 RSTUDIO_PASS=Welcome123! --securityGroup sg-0b23d4d7482ddb825  --publish
 8787:8787 --gateway g-QehqwJldgogOFzeBKGG9P --targetPort 8787  --extraOptions
 "--irmap-scan-path /home/rstudio/" -f
id: Rl6MmFMienMtyIwhbpQoR
name: rstudio-
user: admin
imageID: docker.io/memverge/rstudio:latest
status: Submitted
submitTime: "2023-03-20T21:52:39Z"
duration: 14s
cost: 0.0000 USD
inputArgs: -i rstudio --force -m 4 -c 2 --extraOptions
 '--irmap-scan-path /home/rstudio/' --gateway g-QehqwJldgogOFzeBKGG9P
 --targetPort 8787 --securityGroup sg-0b23d4d7482ddb825 --env
 RSTUDIO_USER=rstudio --env RSTUDIO_PASS=Welcome123! --publish 8787:8787
extraOptions: --irmap-scan-path /home/rstudio/
vmPolicy:
```

```
    policy: spotFirst
    retryLimit: 3
    retryInterval: 10m0s
envs:
    - RSTUDIO_USER=rstudio
    - RSTUDIO_PASS=Welcome123!
publishes:
    - 8787:8787
securityGroups:
    - sg-0b23d4d7482ddb825
gateway:
    gatewayID: g-QehqwJldgogOFzeBKGG9P
    publicIP: 35.175.116.18
    gatewayPort: "10001"
    targetPort: "8787"
```

- Check the status of the submitted job by entering the following (all RStudio jobs are shown):

```
mmce_rstudio.sh -A <OpCenter_ip_address> list
```

The first column (labeled *ID*) shows the IDs associated with RStudio jobs. Look for the ID that matches your job. When the *STATUS* changes from *Initializing* to *Executing*, you can connect to the RStudio session.

- Use a browser to open an http session on the RStudio host by entering:

```
http://<publicIP>:<gatewayPort>
```

where **<publicIP>** is the client-side IP address of the gateway and **<gatewayPort>** is the port that the gateway uses to identify the RStudio server. These are the values noted previously. You can confirm what **gatewayPort** your job is connected to by entering:

float gateway info -g <gw_id> and matching it with the information displayed by **float squeue**.

Loading the RStudio image can take about ten minutes (the uncompressed image is over a GB). If the RStudio host refuses the http connection, wait a few minutes and retry.

- At the RStudio prompt, log in with the default username **rstudio** and default password **Welcome123!**. If you overrode the defaults, use the username and password that you specified when submitting the RStudio job.

Migrate an RStudio Job

There is no difference between running an RStudio job on an On-demand instance and running on a Spot Instance. For either instance type, the RStudio job can be manually migrated to a new instance of any type.

If the RStudio job runs on a Spot Instance, the Spot Instance can be reclaimed by the CSP. If this happens, OpCenter automatically migrates the RStudio job to a new instance and resumes execution — without losing any work in progress. Effectively, the RStudio session continues uninterrupted.

To manually migrate an RStudio job, enter the following:

```
mmce_rstudio.sh -A <OpCenter_ip_address> -j <job_id> migrate
```

where `<job_id>` is the ID associated with the RStudio job to migrate. You can also specify the number of virtual CPUs and memory required for the new host.

Regardless of whether the RStudio job migrated manually or automatically, the IP address of the new instance is different. The ID associated with the RStudio job **does not** change, neither does the client-side gateway IP address and port number. The RStudio job remains connected to the gateway. Your browser session is not interrupted and you do not have to log in again.

Stop an RStudio Job

The host running the RStudio job continues to run indefinitely. To stop the RStudio job and remove the host, enter the following:

```
mmce_rstudio.sh -A <OpCenter_ip_address> -j <job_id> stop
```

where `<job_id>` is the ID of the RStudio job.

Managing Gateways

Gateways can be created (`float gateway create`), destroyed (`float gateway destroy`), or modified (`float gateway modify`). Use the `float gateway modify` command to add or remove a security group that applies to a particular gateway.

An Rstudio job can be disconnected from one gateway (use the `float gateway disconnect` command) and connected to another gateway (use the `float gateway connect` command).

Enter any command with the `-h` option to display the information that is required for each command.

Publication History

Version	Publication Date
3	April 5, 2023
2	March 23, 2023
1	February 6, 2023