

Using Docker in BioHPC Cloud

Jaroslav Pillardy
Bioinformatics Facility

Part 4 – hands-on exercises

May 28 2020

Docker container with GUI

The easiest way to create a Docker container capable of running graphical applications (GUI) is to have a ssh server running inside such a container.

1. Start with centos:7 image
`docker1 run -d -t centos:7 /bin/bash`
2. Install X11 subsystem – it contains all graphical libraries needed [takes really long time]
`yum groups mark install "X11"`
`yum groups mark convert "X11"`
`yum groupinstall "X11"`
3. Install ssh server
`yum install openssh-server openssh-clients`
4. Set up ssh encryption keys
`/usr/bin/ssh-keygen -A`
5. Create a password for root so you can connect as root via ssh
`passwd root`

Docker container with GUI

6. Create a startup script for ssh server. Make it executable (chmod a+x /start.sh).

```
----  
#!/bin/bash  
/usr/sbin/sshd -D &  
/bin/bash  
----
```

7. Install software you need

```
yum install eog
```

8. Save image

```
docker1 commit jarekp__biohpc_1 x11  
docker1 save -o x11.tar biohpc_jarekp/x11
```

9. Run image

```
docker1 run -d -t biohpc_jarekp/x11 /start.sh
```

10. Connect and run software, remember to run X11 client on your desktop

```
docker1 inspect jarekp__biohpc_1  
ssh -X root@172.17.0.2  
[run software]
```

Docker container with GUI

11. You may also want to add Development Tools (compilers and libraries need to compile programs from source)

```
yum groups mark install "Development Tools"
```

```
yum groups mark convert "Development Tools"
```

```
yum groupinstall "Development Tools"
```

Exercises

1. Build an “htop example” image following slides 3-6 from Part 3.
2. Set up a Docker web server following slides 16-24 from Part 3.
3. Build UniRep image following slides 25-29 from Part 3.
4. Build MySQL server using slides 31-42 from Part 3.
5. Create a GUI capable container as described on slides 2-3.
6. Two optional exercises are on the following slides.

BioHPC Docker Example – install MAFFT

1. Check software website for instructions.
<https://mafft.cbrc.jp/alignment/software/>
2. Decide what type of image is needed (CentOS, Ubuntu, public minimal or development etc).
3. Follow instructions.
4. Save resulting container for future use.

BioHPC Docker Example – install MAFFT

File Edit View History Bookmarks Tools Help


MAFFT - a multiple sequence al X +

← → ↻ 🏠 🔒 https://mafft.cbrc.jp/alignment/software/ ⋮ 🛡️ ☆

🌐 CBSU 🌐 DFG 🌐 Freeville WTC 🌐 Freeville WU 🌐 Google Maps 🌐 Ithaca NY Cloud Cover 🌐 Blodgett Mills 🌐 jareksastro 🌐 My LastPass Vault 🌐 Trello 🌐 Cloud Cover NOAA 🌐 ZenDesk

MAFFT version 7

Multiple alignment program for amino acid or nucleotide sequences



Download version

- [Mac OS X](#)
- [Windows](#)
- [Linux](#)
- [Source](#)

Online version

- [Alignment](#)
- [mafft --add](#)
- [Merge](#)
- [Phylogeny](#)
- [Rough tree](#)

[Merits / limitations](#)


[Algorithms](#)

[Tips](#)

[Benchmarks](#)

[Feedback](#)

[Follow](#)



Contact email address, kazutaka.katoh@aist.go.jp, is temporarily unavailable from 2018/Feb/7. If you sent an email to this address but have received no response, then please re-send the email to katoh@ifrec.osaka-u.ac.jp.

About


MAFFT is a multiple sequence alignment program for unix-like operating systems. It offers a range of multiple alignment methods, L-INS-i (accurate; for alignment of <~200 sequences), FFT-NS-2 (fast; for alignment of <~30,000 sequences), etc.

Download and Installation

- [Mac OS X](#)
- [Linux](#)
- [Windows](#)
- [Source](#)
- [Changelog](#)

The latest version is 7.450, 2019/Nov.

An experimental version, 7.464, supports [full-length MSA of closely-related viral genomes](#). **New!** (2020/Apr/21)

 Version 7.463 had a bug. Please update 7.463 to 7.464. (2020/Apr/21)

Input Format

Fasta format. [example1 \(LSU rRNA\)](#), [example2 \(protein\)](#)

The type of input sequences (amino acid or nucleotide) is automatically recognized.

Usage

```
❯ mafft [arguments] input > output
```

An alias for an accurate option (L-INS-i) for an alignment of up to ~200 sequences × ~2,000 sites:

BioHPC Docker Example – install MAFFT

MAFFT version 7
Multiple alignment program for amino acid or nucleotide sequences

Linux versions

Select either of:

- .rpm
 - [mafft-7.464-gcc_fc6.x86_64.rpm](#) (Built on kernel 2.6) *Experimental*
 - [mafft-7.450-gcc_fc6.x86_64.rpm](#) (Built on kernel 2.6)

Root privileges are necessary.

```
% su -  
Password: *****  
# rpm -Uvh mafft-x.xxx-xxx.xxxx.rpm  
# exit  
% rehash (if necessary)
```

A shell script, mafft, is installed into /usr/bin/ and some binaries are installed into /usr/libexec/mafft/.

[License](#)

- .deb
 - [mafft_7.464-1_amd64.deb](#) (Built on kernel 2.6) *Experimental*
 - [mafft_7.450-1_amd64.deb](#) (Built on kernel 2.6)

Root privileges are necessary.

```
% su -  
Password: *****  
# dpkg -i mafft_x.xxx-x_amd64.deb  
# exit  
% rehash (if necessary)
```

A shell script, mafft, is installed into /usr/bin/ and some binaries are installed into /usr/libexec/mafft/.

[License](#)

- [.deb](#) (maintained by [Debian Med](#))
- [Source](#)

BioHPC Docker Example – install MAFFT

Let's install .rpm package – short Google search reveals it is linked to Red Hat and CentOS, among others.

Let's use basic image – if something goes wrong, we can do it again!

We can download the RPM and place it on our machine in `/workdir/labid`

https://mafft.cbrc.jp/alignment/software/mafft-7.464-gcc_fc6.x86_64.rpm

BioHPC Docker Example – install MAFFT

```
[jarekp@cbsum1c2b014 ~]$ ls -al /workdir/jarekp
total 13924
drwxr-xr-x  2 jarekp root          61 May 22 18:42 .
drwxrwxrwx. 4 root  root          30 May 22 13:26 ..
-rw-r--r--  1 jarekp pmm262 14249143 Apr 15 21:49 mafft-7.464-gcc_fc6.x86_64.rpm
-rw-r--r--  1 jarekp root          6620 May 22 18:15 tophat.help
```

```
[jarekp@cbsum1c2b014 ~]$ docker1 run -d -t centos:7 /bin/bash
```

```
Unable to find image 'centos:latest' locally
```

```
Trying to pull repository dtr.cucloud.net/centos ...
```

```
Trying to pull repository docker.io/library/centos ...
```

```
sha256:989b936d56b1ace20ddf855a301741e52abca38286382cba7f44443210e96d16: Pulling from
```

```
docker.io/library/centos
```

```
469cfcc7a4b3: Pull complete
```

```
Digest: sha256:989b936d56b1ace20ddf855a301741e52abca38286382cba7f44443210e96d16
```

```
Status: Downloaded newer image for docker.io/centos:latest
```

```
c40a7a8f3fd3094b1da391a1a502595f9a5a8362e69ce01f05d765a87a88f85b
```

```
[jarekp@cbsum1c2b014 ~]$ docker1 ps -a
```

CONTAINER ID	IMAGE	COMMAND	CREATED	STATUS	PORTS	NAMES
c40a7a8f3fd3	centos	"/bin/bash"	29 seconds ago	Up 27 seconds		jarekp__biohpc_2

```
[jarekp@cbsum1c2b014 ~]$
```

BioHPC Docker Example – install MAFFT

```
[jarekp@cbsum1c2b014 ~]$ docker1 exec -it ee4d845bb163 /bin/bash
root@ee4d845bb163:/workdir# cd /workdir
root@ee4d845bb163:/workdir# ls -al /workdir
total 13924
drwxr-xr-x  2 4965 root          61 May 22 22:42 .
drwxr-xr-x 22 root root        257 May 22 21:52 ..
-rw-r--r--  1 4965 4963 14249143 Apr 16 01:49 mafft-7.464-gcc_fc6.x86_64.rpm
-rw-r--r--  1 4965 root         6620 May 22 22:15 tophat.help

root@ee4d845bb163:/workdir# rpm -Uvh mafft-7.464-gcc_fc6.x86_64.rpm
```

BioHPC Docker Example – install MAFFT

```
[root@31e5b0fcb7b8 workdir]# rpm -Uvh mafft-7.464-gcc_fc6.x86_64.rpm
Preparing... ##### [100%]
Updating / installing...
 1:mafft-7.464-gcc_fc6 ##### [100%]
```

```
[root@31e5b0fcb7b8 workdir]# mafft
```

MAFFT v7.464 (2020/Apr/21)

MBE 30:772-780 (2013), NAR 30:3059-3066 (2002)
<https://mafft.cbrc.jp/alignment/software/>

```
Input file? (fasta format)
@
[...]
```

Remember to save the image! Your container WILL be deleted after your reservation ends.

```
[jarekp@cbsum1c2b014 ~]$ docker1 export -o /home/jarekp/mafft_image.tar 31e5b0fcb7b8
```

BioHPC Docker Example – Install MySQL Database Server

1. Search online for instructions and choose ones best suited for your goal.
2. Pull appropriate image you want to work with
3. Follow instructions to install MySQL
4. Start database server in your container
5. Configure database server so it is accessible for host machine
6. Save the image. Now you have a database server you can move between machines you use! Databases are stored inside the image.

BioHPC Docker Example – install MySQL Database Server

<https://support.rackspace.com/how-to/installing-mysql-server-on-ubuntu/>

The screenshot shows a web browser window with the Rackspace support page. The browser's address bar displays the URL <https://support.rackspace.com/how-to/install-mysql-server-on-the-ubuntu-operating-system/>. The page features a dark navigation bar with the Rackspace logo and links for 'SUPPORT HOME', 'HOW-TO', 'DEVELOPER DOCUMENTATION', and 'BLOGS'. Below this is a 'SUPPORT NETWORK' section with a search bar. The main content area is titled 'Install MySQL Server on the Ubuntu operating system', with a sub-header 'Last updated on: 2019-12-20' and 'Authored by: Jered Heeschen'. The article text explains that MySQL is an open-source relational database and provides a basic installation guide for the Ubuntu operating system. A code block at the bottom shows the terminal commands:

```
sudo apt-get update
sudo apt-get install mysql-server
```

File Edit View History Bookmarks Tools Help

Install MySQL Server on the Ubuntu operating system

https://support.rackspace.com/how-to/install-mysql-server-on-the-ubuntu-operating-system/

CBSU DFG Freeville WTC Freeville WU Google Maps Ithaca NY Cloud Cover Blodgett Mills jareksastro My LastPass Vault Trello Cloud Cover NOAA ZenDesk

SYSTEM STATUS BUY NOW LOG IN

rackspace SUPPORT HOME HOW-TO DEVELOPER DOCUMENTATION BLOGS

SUPPORT NETWORK Search Support documentation...

How-To Home

CLOUD SERVERS

Introduction
FAQ
All Articles

Edit this article

Buy Now

Install MySQL Server on the Ubuntu operating system

Last updated on: 2019-12-20 Authored by: Jered Heeschen

MySQL is an open-source relational database that is free and widely used. It is a good choice if you know that you need a database but don't know much about all the available options.

This article describes a basic installation of a MySQL database server on the Ubuntu operating system. You might need to install other packages to let applications use MySQL, like extensions for PHP. Check your application documentation for details.

Install MySQL

Install the MySQL server by using the Ubuntu operating system package manager:

```
sudo apt-get update
sudo apt-get install mysql-server
```

Feedback

BioHPC Docker Example – install MySQL Database Server

Note: many tutorials suggest using sudo command before actual commands. This is only necessary if you are NOT root user, and therefore should NOT be used in Docker containers – you ARE root there already. Just skip this prefix:

`sudo apt-get update` => `apt-get update`

```
[jarekp@cbsum1c2b014 ~]$ docker1 run -d -t docker.io/biohpc/ubuntu dev /bin/bash
c8cb131bd2deab36221db92531627609b32666eccbd9237c7ed4973028cd921e
[jarekp@cbsum1c2b014 ~]$ docker1 ps -a
CONTAINER ID   IMAGE                                COMMAND                  CREATED          STATUS          PORTS          NAMES
c8cb131bd2de   docker.io/biohpc/ubuntu dev      "/bin/bash"            9 seconds ago   Up 7 seconds   jarekp__biohpc_1
[jarekp@cbsum1c2b014 ~]$ docker1 exec -it c8cb131bd2de /bin/bash
root@c8cb131bd2de:/workdir# apt-get update
Hit:1 http://archive.ubuntu.com/ubuntu xenial InRelease
Get:2 http://archive.ubuntu.com/ubuntu xenial-updates InRelease [109 kB]

[...]

Fetched 4024 kB in 8s (500 kB/s)
Reading package lists... Done

root@c8cb131bd2de:/#
```

BioHPC Docker Example – install MySQL Database Server

```
root@c8cb131bd2de:/# apt-get install mysql-server
```

```
Reading package lists... Done
```

```
Building dependency tree
```

```
Reading state information... Done
```

```
The following packages were automatically installed and are no longer required:
```

```
[...]
```

```
Configuring mysql-server-5.7
```

```
-----
```

```
While not mandatory, it is highly recommended that you set a password for the MySQL administrative "root" user.
```

```
If this field is left blank, the password will not be changed.
```

```
New password for the MySQL "root" user:
```

Decide what password you want for MySQL admin account



```
Repeat password for the MySQL "root" user:
```

```
Unpacking mysql-server-5.7 (5.7.22-0ubuntu0.16.04.1) ...
```

```
[...]
```


BioHPC Docker Example – install MySQL Database Server

```
Setting up mysql-server (5.7.22-0ubuntu0.16.04.1) ...  
Processing triggers for libc-bin (2.23-0ubuntu9) ...  
Processing triggers for systemd (229-4ubuntu17) ...  
root@c8cb131bd2de:/#
```

```
root@c8cb131bd2de:/# systemctl start mysql ←  
Failed to connect to bus: No such file or directory  
root@c8cb131bd2de:/#
```

```
root@c8cb131bd2de:/# /etc/init.d/mysql start ←  
* Starting MySQL database server mysqld  
No directory, logging in with HOME=/  
  
[ OK ]  
root@c8cb131bd2de:/#
```

Something is wrong! Container cannot start MySQL in a server way.

Google search reveals containers have this problem and workaround is to start the service manually.

You will have to start the service manually each time the container is run. Not a problem here. You can also create a script and run it when starting this container.

BioHPC Docker Example – install MySQL Database Server

```
root@c8cb131bd2de:/# mysql -u root -p
```

```
Enter password:
```

```
Welcome to the MySQL monitor.  Commands end with ; or \g.
```

```
Your MySQL connection id is 4
```

```
Server version: 5.7.22-0ubuntu0.16.04.1 (Ubuntu)
```

```
Copyright (c) 2000, 2018, Oracle and/or its affiliates. All rights reserved.
```

```
Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.
```

```
mysql> show databases;
```

```
+-----+  
| Database          |  
+-----+  
| information_schema |  
| mysql             |  
| performance_schema |  
| sys               |  
+-----+
```

```
4 rows in set (0.00 sec)
```

```
mysql> exit
```

```
Bye
```

```
root@c8cb131bd2de:/#
```

BioHPC Docker Example – install MySQL Database Server

Now we have running MySQL server in a container and we can use it INSIDE container. We would like to have MySQL server that can be accessed from the host, this way is more convenient and allows for software or pipelines to use it too.

However by default MySQL is only accessible locally, we need to change MySQL configuration to do so. Google “allow MySQL remote network connections”. There are two suggestions:

1. Comment out line “bind” in [/etc/mysql/mysql.conf.d/mysqld.cnf](#)

```
bind-address = 127.0.0.1 => #bind-address = 127.0.0.1
```

After that we need to restart our MySQL server:

```
/etc/init.d/mysqld restart
```

2. Add remote access privileges to our “root” account

BioHPC Docker Example – install MySQL Database Server

2. Add remote access privileges to our “root” account

check if root has restricted access

```
SELECT host FROM mysql.user WHERE User = 'root';
```

then fix it

```
CREATE USER 'root'@'%' IDENTIFIED BY 'our_root_password';  
GRANT ALL PRIVILEGES ON *.* TO 'root'@'%';
```

The above commands need to be executed inside “mysql -u root -p”

All this info can be found using Google, commands can be copied. Some thinking is still necessary 😊.

BioHPC Docker Example – install MySQL Database Server

```
mysql> SELECT host FROM mysql.user WHERE User = 'root';
```

```
+-----+  
| host   |  
+-----+  
| localhost |  
+-----+
```

```
1 row in set (0.00 sec)
```

```
mysql> CREATE USER 'root'@'%' IDENTIFIED BY 'docker';
```

```
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> GRANT ALL PRIVILEGES ON *.* TO 'root'@'%';
```

```
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> SELECT host FROM mysql.user WHERE User = 'root';
```

```
+-----+  
| host   |  
+-----+  
| %      |  
| localhost |  
+-----+
```

```
2 rows in set (0.00 sec)
```

```
mysql>
```

BioHPC Docker Example – install MySQL Database Server

Before we can connect to our MySQL database remotely, we need to figure out what is our container IP address.

Docker maintains an internal network inside host usually 172.17.0.*.

Use command “`docker1 inspect conatinerid`” to find out more.

```
[jarekp@cbsum1c2b014 ~]$ docker1 inspect c8cb131bd2de
```

```
[
  {
    "Id": "c8cb131bd2deab36221db92531627609b32666eccbd9237c7ed4973028cd921e",
    "Created": "2018-05-23T15:20:40.129886637Z",
    "Path": "/bin/bash",
    "Args": [],
    "State": {
      "Status": "running",
      "Running": true,
      "Paused": false,
      "Restarting": false,
      "OOMKilled": false,
      "Dead": false,
      "Error": "",
      "ExitCode": 0
    },
    "Image": "alpine:latest",
    "NetworkSettings": {
      "Bridge": "docker0",
      "Gateway": "172.17.0.1",
      "GlobalIPv6Address": "",
      "GlobalIPv6PrefixLen": 0,
      "IPAddress": "172.17.0.2",
      "IPPrefixLen": 16,
      "IPv6Gateway": "",
      "MacAddress": "02:42:ac:11:00:02",
      "Networks": {
        "docker0": {
          "IPAMConfig": {
            "IPv4Address": "172.17.0.2",
            "IPv4PrefixLen": 16,
            "IPv6Address": "",
            "IPv6PrefixLen": 0
          },
          "Links": null,
          "MacAddress": "02:42:ac:11:00:02",
          "Name": "docker0",
          "NetworkID": "17217002",
          "Parent": "docker0",
          "PortMapping": null,
          "Scope": "local"
        }
      }
    }
  }
]
```

Here it is!



BioHPC Docker Example – install MySQL Database Server

```
[jarekp@cbsum1c2b014 ~]$ mysql -u root -p -h 172.17.0.2
```

```
Enter password:
```

```
Welcome to the MariaDB monitor.  Commands end with ; or \g.
```

```
Your MySQL connection id is 8
```

```
Server version: 5.7.22-0ubuntu0.16.04.1 (Ubuntu)
```

```
Copyright (c) 2000, 2017, Oracle, MariaDB Corporation Ab and others.
```

```
Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.
```

```
MySQL [(none)]> show databases;
```

```
+-----+
| Database          |
+-----+
| information_schema |
| mysql              |
| performance_schema |
| sys                |
+-----+
```

```
4 rows in set (0.00 sec)
```

```
MySQL [(none)]>
```

BioHPC Docker Example – install MySQL Database Server

Now we can save the image for future use.

Each time we need to start it in the background,

connect to it with “docker1 exec” start MySQL with “/etc/init.d/mysql start”

and then exit container and use the database.

```
[jarekp@cbsum1c2b014 ~]$ docker1 export -o /home/jarekp/mysql_docker.tar c8cb131bd2de
```