B2HOST

About

Documentation about how to use the B2HOST (Service Hosting Framework).

Modified: 22 February 2016

Synopsis

The EUDAT Collaborative Data Infrastructure offers a Service Hosting Framework called B2HOST that allows communities to deploy and operate their own applications and data-oriented services on machines next to the data storage location. Reasons for such services next to the data can be:

- that the volume of the data is too large to be transferred efficiently on demand to third party data processing and analysis facilities
- licensing restrictions that prevent even the smallest volume of data from being copied to a third party which provides the compute facilities.

In both cases, the use of B2HOST allows for the data to remain local, with a (community-specific) service interfacing between the data and external clients.

Resource providers within EUDAT offer service hosting capabilities in tandem with their storage service. These provide access to resources such as bare-metal or virtualized machines with basic execution system platforms (operating systems with a selection of software, tools and libraries). The only allowable use of these resources is for EUDAT communities to deploy and operate data-oriented services hosted at specific data centers. Community service managers can request the appropriate resources through B2HOST, as described in this document. Community service managers can also offer resources and join B2HOST.

How to access B2HOST

Community data managers can check the available service hosting environments and the conditions via the Resource Coordination Tool (RCT). The RCT provides an overview of the available hosting environments in EUDAT and of the options and conditions for requesting and using machines. The requests via the RCT are centrally dispatched and the requestors are notified when the requested platforms are ready for service deployment. Community service managers can then login to the provided machines, deploy and run their services. The EUDAT Terms of Use apply.

In more detail, in order to request the use of a hosting infrastructure, users can make a new request in the RCT. Important information that should be specified within the request is as follows:

- **Provider:** Target infrastructure for B2HOST.
- **Name:** Name for the virtual host or service to be provisioned. No spaces or special characters are allowed, as this will be used as part of the host name
- **Description:** What service(s) will be deployed on the machine? Is it a production or test installation?
- **Ports:** What ports should be opened on the firewall?
- **Contact:** Name, organization, e-mail and telephone number of person responsible for the service(s) being deployed
- **Characteristics:** Required characteristics of the service (CPU/Cores, Disk space, RAM size, Other)

- **Operating System:** Initial OS patching policy and interval. Operating system flavour, including RedHat/Centos/Debian and Windows. In order to use Redhat, Windows or other licensed OS, you have to bring your own licences
- **Preferred hosting platform (if applicable):** The options vary between providers and include OpenNebula, OpenStack and VMware
- **Other:** Any other relevant information

After a successful request, you will receive an e-mail with your credentials and other needed information like host name and IP address of the system you have been granted access to.

For more information about the use of the RCT please check the corresponding documentation in the EUDAT website.

Available hosting environments

Communities can offer service-hosting resources as discussed in the user documentation. The following EUDAT providers are currently offering service hosting capabilities: CSC, JUELICH, RZG, SURFsara. We discuss the offering in turn below.

B2HOST at CSC

OpenStack

CSC offers virtual cluster resources based on the OpenStack cloud service for data intensive computing (Pouta). There is no guaranteed QoS.

Pouta is the main production IaaS cloud at CSC. The Pouta service allows customers to run virtual machines connected to the Internet. It provides an easy to use web interface and a programmable API for managing virtual machines, networks and storage.

This service is targeted to high performance computing (HPC), so the scheduling of virtual machines does not over-commit resources. This allows customers to run virtual machines with exclusive access to up to 16 cores. The same principle applies to RAM and other resources available to the virtual machines. This should provide more predictable performance characteristics compared to general purpose IaaS platforms.

VM hosting environment

When requesting the use of a virtual machine, you can choose between the following characteristics:

- CPU/Cores
- OS:
  - Centos 6.x
  - Other
- Hard disk size:
  - 20 GB
  - Other (max 100 GB)
- RAM size:
  - 2 GB
  - 4 GB
  - 8 GB
  - 16 GB
  - Other

For more information, see https://research.csc.fi/pouta-flavours

The virtual machines run on part of the Taito supercluster. The nodes are HP ProLiant SL230s servers with two Intel Xeon 2.6 GHz E5-2670 CPUs (16 cores per server) and 64 GB of RAM each. The nodes are connected using a fast FDR InfiniBand fabric.

Virtual machines can be given external IP addresses and accessed directly from the Internet. This provides an easy way to access virtual machines from anywhere on the Internet, but customers must also take care to secure their machines. The virtual machines do not have access to any other part of the CSC infrastructure, other than what is already visible to the Internet. Application data and software can be uploaded either via the Internet or copied from CSC's existing shared storage or applications.

CSC account

In order to deploy at CSC, you first need to apply for a CSC account.

Students in Finnish universities or polytechnics can apply for a regular academic user account. The following is valid for a member or a close collaborator of a Finnish research group who works abroad and the requirements for international use are satisfied:

The right for international use of a researcher working abroad is valid for a fixed period provided that the information given in the application form is valid and the grounds for international use are effective. The project manager (i.e., the principal investigator) is obliged to inform CSC of any changes in the international use and the information given in the application. The use rights will be checked regularly.

CSC reserves the right to change this policy and limit the international use if necessary.

- Requirements for international use
- Applying instructions

The applications are handled by the CSC Resource Allocation Group, which assembles every three to five weeks. The applicant has the right of appeal.

You can access your resources with your CSC account username via ssh or other methods, depending on the hosting platform; this is clarified on an e-mail.

Once you have a HAKA account you can request an account for Pouta through SUI. Detailed instructions are available in our new computing environment user guide. If you are not able to create an account, perhaps because you are not an academic user, please contact: cloud-support@csc.fi.

B2HOST at JUELICH

OpenStack

As part of the B2HOST service, Jülich Supercomputing Centre provides a KVM virtual cluster with an OpenStack front-end. Root access is provided for service deployers. The estimated set up time is between 5 minutes and 2 days, depending on whether the user just wants to get a running server of if they want to have direct IaaS access. Local monitoring is available. There is no guaranteed QoS.

VM hosting environment

When requesting the use of a virtual machine, you can choose the following characteristics:

- CPU/Cores
  - up to 16
- OS:
  - Ubuntu 14.04 (64 bit)
  - CentOS 7 (64 bit)
  - Other
- Hard disk size:
  - 20+ GB
- RAM size:
  - Up to 16 GB

JUELICH account

If the user chooses direct IaaS access, they can inject their SSH keys by themselves. If the user requested a server, they will be asked to provide a SSH key as part of the RCT ticket discussed above.

After creation of VM you could access it with the following command:

```
    ssh <your-username>@<ip-of-your-vm>
```

Please make sure that your VM configuration is reboot safe.

B2HOST at RZG

VMware vSphere

RZG provides an extensible service hosting framework based on a VMware vSphere cluster. Root access is provided for service deployers. The estimated set-up time is of 1 to 2 days. The terms of use can be found here: http://www.rzg.mpg.de/userspace/forms/regulations

VM hosting environment

By default, VMs are configured with 20GB hard disk space, 1GB RAM and 1 vCPU running SuSE Linux Enterprise Server (SLES) 11 SP4 x86_64. Also available are SLES 12 and Scientific Linux 7. If required, the amount of RAM and virtual CPUs can be adjusted.

The extensible VM hosting framework currently (May 2014) consists of 32 Sandy Bridge cores and 128GB RAM in total. Attached disk storage capacity of the order of several TBs is available.

The purpose of the hosting environment is to provide VMs for test environments as well as for production services. The VMs are secured by a firewall. Limited amount of ports for inbound traffic can be opened, if the purpose is sufficiently justified (why the service or application requires this).

A hostname of the form *testservice1-eudat.esc.rzg.mpg.de* will be assigned to the VM. Please contact us if you want to point your own domain at the machine.

The operating system of the VMs will be periodically updated. You will be informed via E-Mail if the VM needs to be rebooted (usually a few times per year). You are responsible for keeping you services up to date. Please make sure that your VM configuration is reboot safe.

RZG account

If this is your first service that you want to deploy at RZG you will need to apply for an RZG account in the following link:

https://www.rzg.mpg.de/secure/registrieren/antrag.php?deisa=1&inst=RZG&projekt=EUDAT&lang=en

Please contact eudat-support@rzg.mpg.de for assistance.

With your RZG account, you can log into the gateway machine:

```
   ssh <your-account-name>@con01.rzg.mpg.de
```

From there, you can log into the individual VMs as root:

```
   ssh root@<address-of-your-vm>
```

(authentication via SSH keys automatically created on login node)

B2HOST at SURFsara

OpenNebula

SURFsara offers KVM virtualization with an OpenNebula front-end. Root access is provided for service deployers. The estimated set up time is immediate (deployer does set up).

SURFsara provides a BiG-Grid sponsored HPC Cloud environment. This cloud can be accessed via an extended OpenNebula cloud management software layer. This infrastructure is offered as a "self-service" IaaS cloud; a cloud user configures his or her own virtual machines.

The cloud user has a free choice of operating system. Commercial operating systems (for instance Microsoft Windows) are allowed. Please note that the cloud user must provide the software licenses the VM(s) may require.

The SURFsara HPC Cloud does not provide pre-built virtual machines. The infrastructure is provided and you are free to instantiate any and all virtual environment(s) you require. You can get an overview of the actual usage of the resources from the following links:

- Compute Nodes
- Services

VM hosting environment

The HPC Cloud currently consists of a number of services plus 30 dedicated compute nodes. 19 nodes of these nodes each have:

- four 8-core Westmere E7 CPUs (Intel(R) Xeon(R) CPU E7-4830), totaling 32 cores per compute node
- 8GB RAM per core, totaling 256GB RAM per compute node

The remaining 11 nodes each have:

- four 8-core Westmere E5 CPUs (Intel(R) Xeon(R) CPU E5-4650), totaling 32 cores per compute node
- 8GB RAM per core, totaling 256GB RAM per compute node

In total the cluster has 960 HPC cores with 7680GB of memory, several lighter nodes and a high-memory node with 40 cores and 2TB of memory. The compute nodes are interconnected using a non-blocking 10Gbps network backbone around an Arista DCS-7500 switch. Each node is attached to this network backbone with four bonded 10Gbps connections. A 400TB DDN NAS provides the data storage.

EUDAT receives funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 654065.

The cloud cluster runs CentOS6 with the QEMU-KVM hypervisor under libvirt. Access to the cloud infrastructure is provided through the OpenNebula Sunstone web interface, or the OpenNebula XML-RPC interface.
A cloud user is able to initiate a virtual machine with a maximum of 32 CPUs and about 240GB RAM. All cloud users are completely separated from each other using a combination of VLANs and firewalls. An IPFilter based network firewall must be set on all internet connections. By default, two incoming internet ports are available, while others can be manually added using the integrated network filter setup feature of the OpenNebula interface. The deployer has full freedom, but also responsibility.

SURFsara account

For obtaining an HPC Cloud account please check the online guide.

Security considerations

The providers are running their hosting environments in different DMZ networks which are secured by firewalls. By default, most of the ports stay closed to the Internet. If your service needs some port to be open or if it also needs connection to other machines outside the provider's network (for example, connection to one external ldap server on port 636), you should specify it like already mentioned in the RCT request. Once the service is running, some sites may then want to run some kind of security port scan (for example a Nessus port scan) so please send an e-mail to the support address of the provider specifying the kind of service and the port where it is listening.

Once the security test has run, the results of the scan are sent to the site administrator.

- If the port scan is successful, that means that **the service doesn't present any (known) vulnerability or exploit**, and the port will be open.
- If the port scan fails, the site administrator will send a list of suggestions for correcting the vulnerabilities to you as the service deployer. After you have made the pertinent changes, please inform again the provider so the port scan can run again.

Here follow some examples of secure configuration parameters checked by the Nessus port scan:

Plugin ID: 10107 - HTTP Server Type and Version

Synopsis: A web server is running on the remote host.

Solution: You can set the directive 'ServerTokens Prod' to limit the information emanating from the server in its response headers.

In Apache:

```
# Nessus Plugin 10107
ServerTokens ProductOnly
ServerSignature Off
```

Plugin ID: 11213 - HTTP TRACE / TRACK Methods Allowed

Synopsis: Debugging functions are enabled on the remote web server.

Solution: To disable these methods, add the following lines for each virtual host in your configuration file :

```
RewriteEngine on


RewriteCond %{REQUEST_METHOD} ^(TRACE|TRACK)


RewriteRule .* - [F]
```

Alternatively, note that Apache versions 1.3.34, 2.0.55, and 2.2 support disabling the TRACE method natively via the 'TraceEnable' directive.

In Apache:

```
# Nessus Plugin 11213
TraceEnable Off


# And for good measure in <Location />:
Options SymLinksIfOwnerMatch     # required for RewriteEngine
RewriteEngine on
RewriteCond %{REQUEST_METHOD} ^(TRACE|TRACK)
# For the RewriteRule to work, mod_rewrite must be loaded (-> a2enmod rewrite).
```

Plugin ID: 10884 - Network Time Protocol (NTP) Server Detection

Synopsis: An NTP server is listening on the remote host. An NTP (Network Time Protocol) server is listening on this port. It provides information about the current date and time of the remote system and may provide system information.

Solution: When possible deactivate the NTP service and use the JUELICH central time service ntp.fz-juelich.de. Example of use in the crontab:

```
crontab -l


00 2,8,14,20 * * * (/usr/sbin/sntp -P no -r ntp.fz-juelich.de)
```

Support

In case of questions, you can always contact the providers at the different support addresses:

Site	Site Support Address
CSC	eudat-support@csc.fi
JUELICH	eudat-support@fz-juelich.de
RZG	eudat-support@rzg.mpg.de
SURFsara	eudat-support@surfsara.nl

If you have comments on this page, please submit them though the EUDAT ticketing system.

Document Data

Version: 1.2

Authors:

Cristina Manzano, c.manzano@fz-juelich.de

Florian Kaiser, florian.kaiser@rzg.mpg.de

Hans van Piggelen, hans.vanpiggelen@surfsara.nl

Pietari Hyvärinen, pietari.hyvarinen@csc.fi

Sander Apweiler, sa.apweiler@fz-juelich.de

Editors:

Kostas Kavoussanakis, kavousan@epcc.ed.ac.uk

Johannes Reetz, johannes.reetz@rzg.mpg.de

Carl Johan Håkansson, cjhak@kth.se

Read more

EUDAT receives funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 654065.