

From PC to Cloud or HPC

Introduction

You may have heard of **High Performance Computing (HPC)**, or you may already be using it. This module will discuss the differences between Cloud Computing and HPC, and provide an overview of pros and cons of moving from traditional desktop computing to Cloud or HPC infrastructure.

Videos

The following videos go through most of the content in this module and offer a less in-depth description of the subject than the documentation does.

<https://www.youtube.com/watch?v=MQ-CT05qevM>

<https://www.youtube.com/watch?v=JUhN5EhPSW8>

Conventions

The notation throughout the training documents can be interpreted as follows:

Words in *italics* are used for names and terminology, e.g. name of a software, or name of a computing concept. It may also just emphasise a word in the traditional way.

Quotations are also written in italics and are put in between quotation marks.

Words in **bold** are used to highlight words which identify important concepts of a paragraph, to make it easier for users to skim through the text to find a paragraph which explains a certain idea, concept or technology.



Additional information which is optional to read is displayed in *info boxes* like this one.



Important information is displayed in boxes like this one.



Definition of terms are displayed in boxes of this style.



Possibly specific **prerequisites** for reading a particular section are contained in this type of box at the beginning of a section.

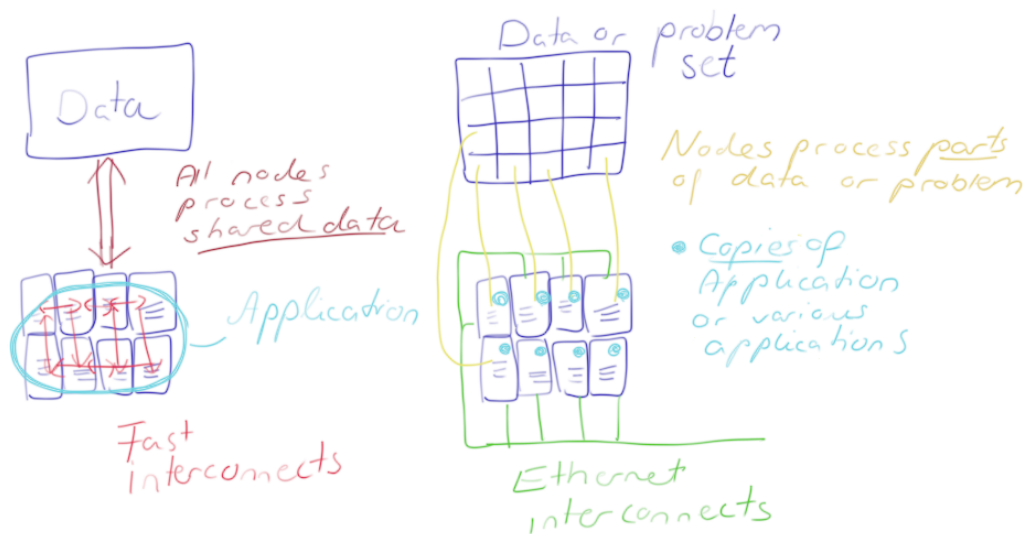
HPC vs Cloud computing

High Performance Computing (short: *HPC*) is not the same as cloud computing. Both technologies differ in a number of ways, and have some similarities as well.

We may refer to both types as “*large scale computing*” – but what is the difference? Both systems target scalability of computing, but in different ways.

HPC targets extremely large sets of data and crunching the information in parallel *while* sharing the data between compute nodes (you can think of a “node” as a computer). The data connection between the nodes has to be very fast (typically, *Infiniband* technology is used), essentially turning the entire grid of nodes into one single “supercomputer”. This requires expensive hardware: nodes with individually high performance, i.e. high processing power and large memory, and very fast network connections between nodes. **One** application can be run **across** a variable number of nodes. We call this **vertical scalability**.

Cloud computing on the other hand targets “embarrassingly parallel problems” (EPP). An embarrassingly parallel problem is one for which little or no effort is required to separate the problem into a number of parallel tasks. This is often the case where there exists no dependency (or communication) between those parallel tasks. A common EPP problem is one in which a very large data set is chopped into pieces which are dispatched to various computers for processing; or, several copies of a smaller data set are distributed across computers to perform different computations on it (e.g. running the application with different parameters). After the processing is finished, the resulting data is re-assembled or the results from all computers summarized. The individual computers don’t have to be super fast, but instead the power lies in having a huge number of computers. *Several* applications (or, *copies* of the same application) run on several nodes. We call this **horizontal scalability**.



i The main point of *horizontal* scalability (in cloud computing) is that data and the same application are replicated across the computers. This can be done to perform calculations, or merely to replicate the same application / data in order to ensure availability.

In contrast, with *vertical* scalability (in HPC) there is only *one* instance of the application; replicating it does not improve performance. Instead, the application itself works in a distributed way over multiple instances: one single application uses hundreds or even thousands of cores and accesses the data on a storage entity that is attached via the fast network to all the nodes.

Bernhard Schott, CTO of VCODYNE and formerly project manager on distributed complex systems at Platform Computing, describes the difference between horizontal and vertical scaling in terms of a schoolyard.

“If you have 200 school kids and want each of them to pick a piece of paper off the floor, that’s a perfectly parallel problem [an EPP] that scales really well, like in the cloud. If you want to coordinate those children to perform together in the same ballet, you have a whole new set of problems, and it doesn’t scale well”

So HPC and Cloud Computing try to achieve a different type of scalability. To achieve their aim, both techniques use their own optimized hardware. Depending on the requirements of your research application, one or the other may be the better solution.

i Some providers also offer HPC systems in the Cloud. HPC requires specialized hardware, so the provider must have such a specialized system as part of their infrastructure. Usage of *HPC in the Cloud* then works

just as using HPC systems which your University may provide, only that the HPC System is located at the Cloud providers infrastructure instead of at your Universities data center.

Amazon, for example, offers *HPC in the Cloud*, as does eResearch South Australia.

When to use HPC

HPC is the right choice if you have an application that pushes on *all* levels of performance:

- computing
- fast interconnects, and
- high-performance storage

This computing environment is not found in a typical Cloud and is unique to HPC.

Optimized HPC libraries (the result of years of research) may be required for your application, working closely with the hardware. The hardware can even be specialized “tuned” hardware which is not available on the simpler cloud computers.

Some applications rely on a technology called **MPI (Message Passing Interface)** which explicitly handles communication between computing nodes within the program code. Such applications may not run efficiently in the Cloud because inter-node communication is slower. Also, some Clouds do not support a technique called *broadcast* or *multicast* in any form. So if your *MPI* implementation uses this, it cannot run in the Cloud. As a rule of thumb, if your application is using *MPI*, then HPC is probably the right choice.

Some applications require very **fast interconnects** (e.g., *Infiniband* or *High-Performance Ethernet*) which require communication that bypasses the OS kernel. This makes the use of the Cloud very difficult because most virtualization schemes do not support this “*kernel bypass*”. If high-performance inter-connections are not available, the application runs slowly and get no performance gain when adding more nodes. So if your application demands high-performance inter-connections between the compute nodes, you should choose HPC over the Cloud.

Other specialised hardware which your application may benefit from are **performance accelerators** or *SIMD* units (Single-Instruction Multiple Data processors), available from *Nvidia* and *AMD/ATI*. This is not found on typical Cloud infrastructure — if your

application makes use of such performance accelerators, you should choose an HPC infrastructure which provides this.

Some HPC solutions already offer a set of pre-installed **software packages**, sometimes including the license, ready for you to use. When using the Cloud, you have to install specialized software yourself, and pay for the license.

When to use the Cloud

Using virtual machines (VM) in the Cloud is the right choice if you want to process one data set with a variety of parameters, or when you can split the data set into several pieces for parallel processing.

In other words, you have an **Embarassingly Parallel Problem (EEP)**. The application does not rely on fast shared memory or storage so it can be distributed into many independent processing units. With the Cloud, you just use the resources you actually need at the time. You can't do that in HPC, you reserve a large number of resources and use them (maybe pay for it) for the whole time.

One often cited example is digital rendering, in which many non-interacting jobs can be distributed across a large number of nodes. These applications often work well with standard Ethernet and do not require a specialized interconnect for high performance.

In HPC, we may need to wait for the resources to become available, in the Cloud they are typically **instantly available**. With Cloud computing, you can get a whole lot of resources for a time, without the need of big capital investments. With HPC on the other hand, you have a specialized machine with its limits, which can only do so much.

Another reason to choose the Cloud can be the possibility of **software choice**: users can design virtual machines to suit their needs, including choice of OS and installed applications. In HPC, installing specialized software can be more problematic. On the other hand, on HPC systems, specialized software may already be pre-installed and ready to use for you, sometimes including the license to use it. When using the Cloud, you *can*, but also *have to* install the software yourself and pay for the license. You have to evaluate this depending on which software you require.

Last but not least, your requirements may be a simple case of needing easy access to computing infrastructure (already connected to the Internet) which you cannot easily get at your research organisation. You may not even need to use a specialized application which runs across several computers, and only need access to one or a few computers that you can use for your research or teaching. Using the cloud will save you the need to invest in local facilities and configuring access from various remote locations. In such cases, the Cloud is the right choice for you: You get

compute and storage infrastructure quickly and easily, and the virtual machines already are accessible over the Internet.

There are a number of reasons why the Cloud may be the right choice. See also Module 3 for a discussion of common use cases.

The Cloud - Pros and Cons

There are a number of advantages and drawbacks to both local computing (the traditional computing model you are familiar with, which is using your local office computers) and to cloud computing. Depending on your needs and concerns, you may take the pros and cons into account when deciding whether to move your applications and data to the cloud.

Pros and Cons of *local* computing

Pros:

- Access to the computer is **fast and easy**.
- Users feel **more comfortable** having their data and processing *on-premises* — it removes safety and data ownership concerns. But: is this justified? Security measures for example are not necessarily better *on-premises* than at the cloud provider's data center.

Cons:

- **Cost:** Need to invest in facilities and maintenance.
- **Limited Access:** Access to existing facilities may be limited and on high demand.

Pros and Cons of *cloud* computing

Pros:

- **Cost:** Cloud computing can be viewed as *pay-as-you-go* computing: You pay only when you need the service (NeCTAR services are even *free* to you). This is usually much **cheaper** than building and maintaining an *on-premises* physical server farm.

Shifting large infrastructure spendings to an operational expense has many advantages, including

- almost unlimited storage
- instant availability and

- the ability to scale up (and down) rapidly.

While NeCTAR services are free to you, but you must request an allocation of a certain amount of maximum resources you will be using.

- **Individual setup:** Users can configure their own server (e.g. choose the operating system and install software) and run it in a cloud whenever they need the computational resources.
- **Access independence:** via the Internet, cloud computing can be done from anywhere (office, home, on conference travel, business trips etc) and with a variety of devices (laptop, smart phone, tablet..). There is no need to install (and maintain) the research applications on each device.
- **Large computing capacity** can be accessed quickly, and only for the time you need it.
- **“Elasticity”** (Flexibility and scalability): Typically, a cloud consists of a dynamically assigned group of *virtual machines* that can scale up quickly at your request. This gives the users the ability to scale up or scale down technological infrastructure resources as required at the time.
- **Resource sharing** is possible across many users. Multiple users can work on the same data simultaneously, which avoids having to wait for it to be emailed.
- **Security** is often as good as or better than other traditional systems (more about this later). This is often mistakenly perceived as believed to be a risk of cloud computing.
- **Saving electricity** through shared infrastructure. This does not only save costs, it is good for the environment!

Cons:

- Requires the **Internet** to access. If the Internet drops out, you lose access. However your services (e.g. your data analysis) will still be running—you just cannot access it.
- **Indirect access control:** The ISPs, telecommunication and media companies control your access. Putting your faith in the cloud means you’re also putting all your faith in continued, uninterrupted access. ISPs may even charge more for higher bandwidth demands.
- **Service outage:** When there are problems at the cloud service provider, it can take out all your services. However usually such outages are for just hours.

- **Concerns about ownership:** Who owns the data you store online: you, or the company storing it? There is also a difference between data which is uploaded, and the data created in the cloud itself (providers could have stronger claims on the latter). Policies of cloud providers vary. NeCTAR never lays any claim on any of your data!
- **Service charge** is based upon usage, which may come at a significant cost with some cloud providers. However NeCTAR services are free to you as an Australian Researcher.

Summary

Hooray!

You should now have a good idea about the difference between Cloud and HPC.

If you have found that the Cloud is a great choice for your research IT infrastructure, you will enjoy the rest of this course.

You may have found that HPC is the right choice for you, in which case you can go on to find out more about your options to use HPC computing. Contact your local IT department to find out more. You may also be interested to take a look at the [Intersect course](#) “Introduction to Unix for HPC”. You may still benefit from attending the rest of this course, in which we will learn more about the Research Cloud and how to use it for your research. Maybe you will be able to identify use cases in your future research, in which case you will be all prepared!

In summary, this module has discussed:

- The difference between Cloud and HPC
- When to use HPC
- When to use Cloud Computing
- Pros and cons of the Cloud

You may now continue with [Module 5](#).