

Introduction to NCI Cloud Computing

General Guide to NCI and other Cloud environments

NCI operates an Openstack based private cloud computing environment which is available to NCI users and projects who have a dynamic or long running service requirement. At its core Cloud computing is the delivery of computing services including servers, storage, networking, software, analytics and intelligence over the Internet (or "the cloud") to offer faster innovation, flexible resources and economies of scale. Typically, you only pay for cloud services you use, helping you utilise NCI services budget more effectively, operate the infrastructure you need more efficiently and scale across NCI Cloud and HPC environments as your research needs change.

NCI's Nirin Cloud environment provides NCI users with the capability to compose a dynamic and scalable compute environment, built using Virtual Machines (VMs) for compute, Volume and Object storage to hold data, Private and Internet access for networking, and access to NCI gdata filesystems.

This guide covers how to request access to Nirin; log in to the dashboard; set up or import an SSH key pair; create, start and stop Virtual Machines; create and connect additional volume storage; and allocate and attach externally accessible IP addresses. The intended audience is for users who are beginning their use of NCI Cloud services.

Types of Cloud computing

Most cloud computing services fall into four broad categories; infrastructure as a service (IaaS), platform as a service (PaaS), serverless, and software as a service (SaaS). NCI Cloud is used to provide many of the data services and capabilities like the Australian Research Environment (ARE) which uses NCI's IaaS (Nirin Cloud) to build a PaaS (ARE) to deliver SaaS.

| | |
|---|---|
| Infrastructure as a service (IaaS) | The most basic category of cloud computing services. With IaaS, you rent IT infrastructure – virtual machines (VMs), storage, networks, operating systems – from NCI on a pay-as-you-go basis. |
| Platform as a service (PaaS) | Platform as a service refers to cloud computing services that supply an on-demand environment for developing, testing, delivering and managing software applications. PaaS is designed to make it easier for developers to quickly create web or mobile apps, without worrying about setting up or managing the underlying infrastructure of servers, storage, network and databases needed for development. |
| Serverless computing | Overlapping with PaaS, serverless computing focuses on building app functionality without spending time continually managing the servers and infrastructure required to do so. NCI handles the setup, capacity planning, and server management for you. Serverless architectures are highly scalable and event-driven, only using resources when a specific function or trigger occurs. |
| Software as a service (SaaS) | Software as a service is a method for delivering software applications over the Internet, on demand and typically on a subscription basis. With SaaS, cloud providers host and manage the software application and underlying infrastructure, and handle any maintenance, such as software upgrades and security patching. Users connect to the application over the Internet, usually with a web browser on their phone, tablet or PC. |

NCI Cloud Terminology

| Term | Definition |
|----------------------------------|--|
| Virtual Machine (VM) or Instance | A Virtual Machine or Instance is a virtual computer which has a number of processing cores called Virtual CPUs (vCPUs), Memory, operating system storage, virtual disks and virtual networks. Clouds allow many Virtual Machines to share the same underlying physical infrastructure allowing efficient and flexible use of shared resources |
| Flavour | The core resource allocations for a virtual machine - a flavour defines the number of vCPUs, memory and local storage allocated to the virtual machine when it is created. |
| Volume | An external and persistent block storage device, which can be attached to a virtual machine - effectively an external hard drive. Note a virtual machine has local storage which is tightly coupled to the virtual machine's life cycle, and will be deleted along with the virtual machine; volume storage is by default separate from the virtual machine and will persist even if the virtual machine is deleted. |
| OpenStack | OpenStack is an open source cloud operating system that controls large pools of compute, storage, and networking resources throughout a datacenter, all managed and provisioned through APIs with common authentication mechanisms. Beyond standard infrastructure-as-a-service(IaaS) functionality, additional components provide orchestration, fault management and service management amongst other services to ensure high availability of user applications |
| Hypervisor | A hypervisor or virtual machine monitor (VMM) is a piece of software that allows physical devices to share their resources among virtual machines (VMs) running on top of that physical hardware. The hypervisor creates, runs and manages VMs |
| Floating IP | An externally accessible IP address which can be attached to a virtual machine to allow external access. |

Openstack provides a number of underlying capabilities which are coordinated using a Web facing Dashboard which can be accessed via <https://cloud.nci.org.au>

| Capability | Services |
|------------------|---|
| Dashboard | Provides a front end for managing the other resources and capabilities. |
| Identity Service | Controls access to the cloud Dashboard and manages credentials across other Capabilities |
| Compute | Manages the creation, execution and deletion of Virtual Machines (VMs). |
| Block Storage | Provides block storage capabilities similar to a conventional disk drive. Block Storage presents Volumes to the Virtual Machines which presents as disk devices to the operating system of the Virtual Machine. |
| Networking | Provides IP networking between the Virtual Machines within a Project (using Private addresses) and external access to the Internet (using Floating IP addresses) |
| Image Service | Provides access to operating system Images and Snapshots used to boot a Virtual Machine. |
| Object storage | Provides an S3 compatible storage service for object data. |