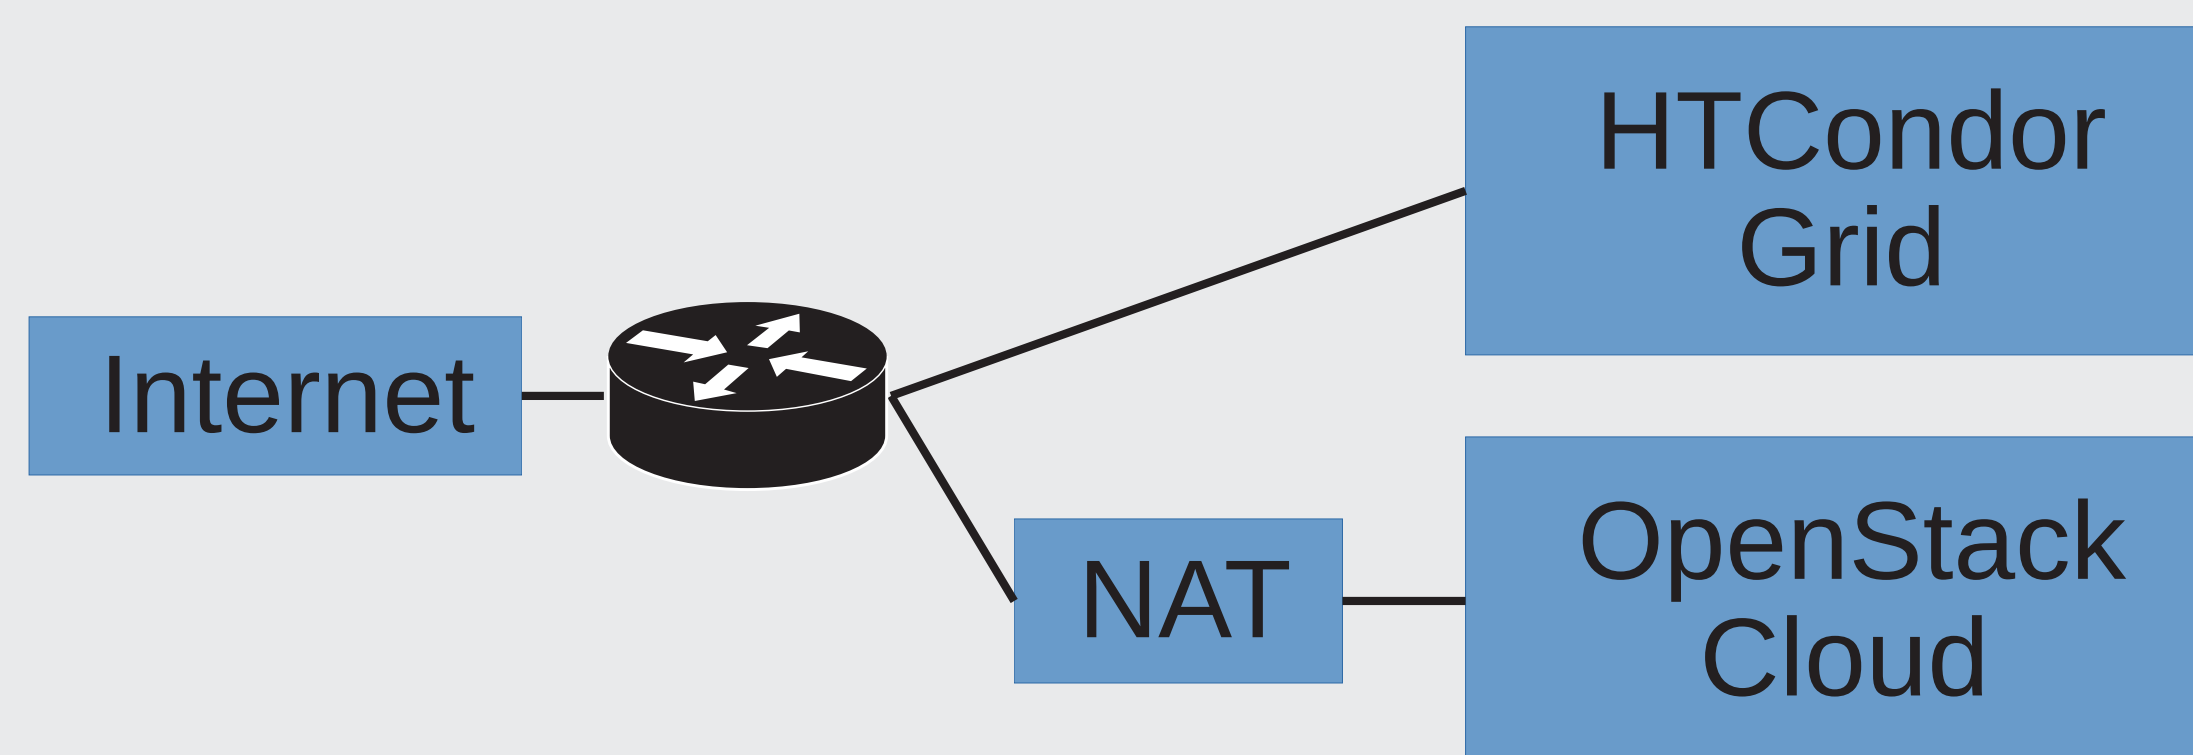


Backfilling An OpenStack Cloud With Grid Jobs

D. Whitehouse & S. Fayer
lcg-site-admin@imperial.ac.uk

Introduction

- Problem: OpenStack Cloud resources were underutilised.
- Few long term instances, mainly projects, development work and other ad-hoc use.
- Cloud is colocated with a busy Tier-2 Site.



- Use spare capacity (when available) to provide resources for WLCG workloads.

Requirements

- Images must be up to date and reflect bare-metal worker nodes in terms of security patches & package versions.
- Instances must work with our HTCondor batch system and contain appropriate credentials & configuration to request and execute jobs.
- Instances must start up and stop quickly in order to release resources when other cloud projects require them.


Backfill Algorithm

- Algorithm script queries each hypervisor to determine how many 8-core worker VMs (*Target*) can be started in the available space, taking into account free vCPUs and RAM.
- The calculation must also consider and avoid fragmented resources.
- 90% of spare capacity used for ephemeral backfill nodes, with 10% left spare as headroom for other cloud users.
- Idle & shutdown instances are regularly removed to ensure resource availability.

Implementation

- Separate OpenStack backfill project segregates resource use.
- Image builds are carried out nightly using Jenkins and uploaded to meet our security patching standards.
- Oz and custom scripts install & configure HTCondor alongside a standard WLCG environment within the image.
- Instances automatically start draining shortly after they have booted, ensuring only one set of jobs runs before the instance terminates.

Backfill Monitoring



- Used in production for 18 months with key metrics sent to Ganglia.
- *Running instances* tracks *Target instances* well, demonstrating good use of spare capacity.
- The number of *Running instances* is sometimes greater than *Target instances* when cloud projects are waiting for HTCondor jobs to finish.
- No issues were observed by cloud projects.