


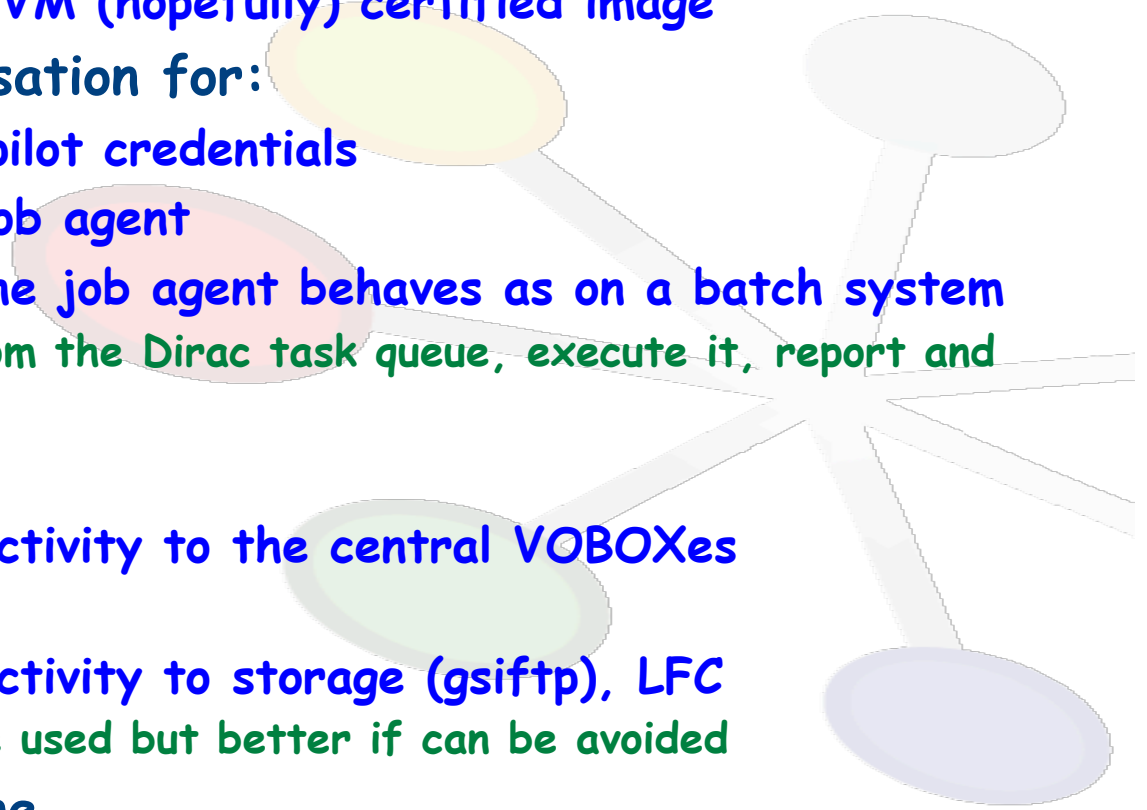


LHCb and virtualisation





- Aim: be able to replace a Dirac pilot with a customised CERNVM virtual machine
 - Why CERNVM: it is there, our applications and Dirac are installed on CVMFS
 - Using the CERNVM (hopefully) certified image
- Using contextualisation for:
 - Installing Grid pilot credentials
 - Start a Dirac job agent
 - From then on the job agent behaves as on a batch system
 - ★ Gets a job from the Dirac task queue, execute it, report and upload data
- Requirements:
 - Outbound connectivity to the central VOBOXes
 - ★ Mandatory
 - Outbound connectivity to storage (gsiftp), LFC
 - ★ Proxies can be used but better if can be avoided
- Starting a machine
 - Generic cloud interface (EC2?) would be OK





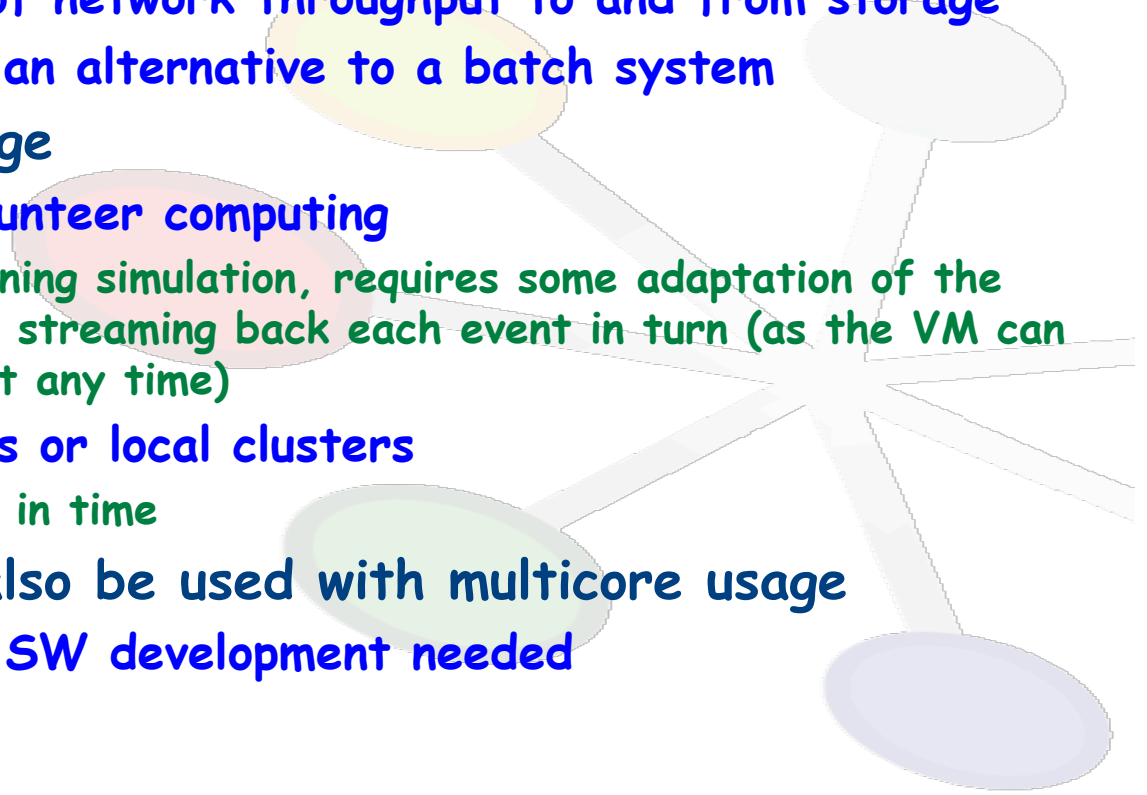
Shutting down the VM

- Do we need to shutdown the VM is not necessary?
- Ideally the VM could run as long as necessary and compatible with pledges
- What if the WM needs to be shut down?
 - Migrate the VM to another hardware
- What if the share is exhausted?
 - Define a protocol with a site control service that gives authorisation to start another job, or gives a deadline for shutdown
 - ☆ Up to the agent to see if it can run more jobs
 - Clouds may allow just the number of VMs corresponding to the pledge (à la EC2: you get what you pay for)
- What if the site wants to shutdown the VM?
 - Alternatively allow sites to send a signal to the pilot for a scheduled shutdown (delay 24 hours)
 - ☆ If the machine is not shutdown, the site is allowed to kill it



Usage of CVMDirac

- EC2 Amazon cloud (OK but someone should give me her credit card number...)
- LXCLLOUD or similar institutional cloud
 - Better control of network throughput to and from storage
 - Can be used as an alternative to a batch system
- Opportunistic usage
 - E.g. BOINC volunteer computing
 - ☆ Mainly for running simulation, requires some adaptation of the application for streaming back each event in turn (as the VM can be shutdown at any time)
 - Idle office CPUs or local clusters
 - ☆ Less limitation in time
- CVMDirac could also be used with multicore usage
 - Still some Core SW development needed





- Primary usage: application deployment
 - Already in use at several sites
 - ☆ PIC, NIKHEF, CERN, RAL (in order of appearance)
 - ☆ In the pipeline for IN2P3, CNAF, SARA, GRIDKA
 - ☆ Some Tier2s
 - Very positive experience
 - ☆ Probably needs for tuning on our side as well
- Potential usage: Dirac usage on WNs
 - Just a matter of changing the pilots' code
- ConditionsDB
 - Snapshots of the ConditionsDB are created as XML files and distributed with Sw
 - Therefore they are available from CVMFS where it's used

