
Distributed Computing Resources at Duke University

Scalable Computing Support Center

<http://wiki.duke.edu/display/SCSC>

<http://sites.duke.edu/scsc>

scsc@duke.edu

John Pormann, Ph.D.

jbp1@duke.edu

What is the SCSC?

- Scalable Computing Support Center
 - ◆ We connect researchers to hardware, software, educational, and personnel resources, both local and global, to enable novel computational science
 - ◆ We will leverage the parallel computing facilities already in place, help build out the computational infrastructure to handle future work-loads, foster the development of scalable applications, and assist in the training of parallel-aware researchers
 - ◆ We provide expertise in computational science
 - Algorithm design, numerical analysis
 - Parallel and high-performance computing

HPC and HTC

- High Performance Computing (HPC) generally means getting a particular job done in less time (for example, calculations per second).
 - ◆ DSCR
- High Throughput Computing (HTC) means getting lots of work done per large time unit (for example, jobs per month).
 - ◆ Condor
 - ◆ OSG

Duke Shared Cluster Resource

- As of 8/'13, ~460 dedicated machines
 - ◆ 2-16 CPU-cores, 1-512GB
 - ◆ 1 & 10Gbps networking
 - ◆ ~50TB of on-line disk storage
- It uses a “Condo” model
 - ◆ Researchers purchase new machines and add them to the cluster
 - ◆ We guarantee high-priority access to your machines whenever you need them



DSCR/Flexibility - Hardware


- While we would like to provide flexibility in hardware vendors, we have seen great pricing when we “batch” orders and go to one vendor
 - ◆ Dell is currently the preferred vendors
 - ◆ “Blade” form-factor (we can also handle 1U)
 - Machines can go up to 512GB (alt. platforms can get to 1TB)
 - ◆ Intel CPUs, 64-bit
 - Current “sweet-spot” is dual eight-core CPUs
 - ◆ New blades have 10Gbps Ethernet on-board
 - May share a 10Gbps uplink




DSCR/Flexibility - Software

Wolfram
Mathematica[®]

MATLAB[®]
The Language of Technical Computing



Copyright 1984-2004, The MathWorks, Inc.



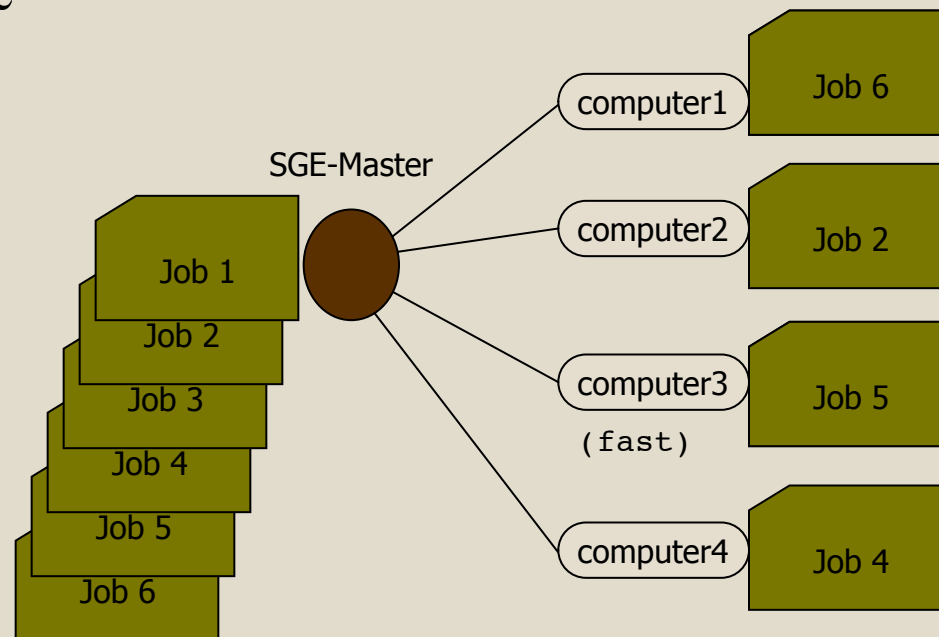
Duke UNIVERSITY




Scalable Computing Support Center
<http://wiki.duke.edu/display/scsc>

DSCR, cont'd

- The DSCR is a “Batch” environment
 - ◆ All jobs go through a queuing system
 - ◆ High-priority jobs launch immediately onto your own machines
 - ◆ Low-priority jobs may wait for an open slot on someone else’s machine




Interesting results ...



- Users have queued up 5000 jobs to run over a weekend
- Someone ran 400 8-CPU jobs (in low-priority mode)
 - ◆ ... completed in about 1 day!
- We've seen a single job use 200-300 CPUs
 - ◆ Many users routinely run 20-CPU jobs
- We've seen 3-month-long jobs run on the DSCR without any problems
 - ◆ We do aim for quarterly maintenance, but not all of them are outages

Virtual Compute Lab

- VCL gives users access to remote desktop machine-images through a web-based reservation system
 - ◆ <https://vcl.oit.duke.edu>
- After reserving your image, you can connect through X11 or RDP
 - ◆ Can reserve multiple seats for classroom use
- And you have ‘root’ on the machine!
 - ◆ For the duration of your reservation
- VCL is now an Apache project:
 - ◆ <http://vcl.apache.org/>





VCL: Virtual Computing Lab

NEW RESERVATION	CURRENT RESERVATIONS	BLOCK ALLOCATIONS	USER PREFERENCES	MANAGE IMAGES
MANAGE COMPUTERS	VIEW TIME TABLE	VIRTUAL HOSTS	STATISTICS	HELP
IMAGE LIST	LOGOUT			

Current Reservations

You currently have the following normal reservations:

	Environment	Starting	Ending	Initially requested
<input type="button" value="Connect!"/> <input type="button" value="End"/> <input type="button" value="Edit"/>	OIT Windows7 Lab Image	Friday, Aug 23rd, 11:13 am	Friday, Aug 23rd, 12:15 pm	Friday, Aug 23rd, 11:13 am

Click the **Connect!** button to get further information about connecting to the reserved system. You must click the button from a web browser running on the same computer from which you will be connecting to the remote computer; otherwise, you may be denied access to the machine.






Duke UNIVERSITY

Scalable Computing Support Center
<http://wiki.duke.edu/display/scsc>


Condor



- Last year, we officially deployed a Condor grid across campus
 - ◆ Mostly Physics-owned machines
 - ◆ Some VMs are contributed nightly from OIT/VCL




- <http://cs.wisc.edu/condor/>



Condor: Opportunistic Computing

- Desktop PCs are idle for half the day
 - ◆ ... or more!



Desktop PCs (and VMs) tend to be active during the day.



But at night, during most of the year, they're idle. So we're only getting half their value (or less).

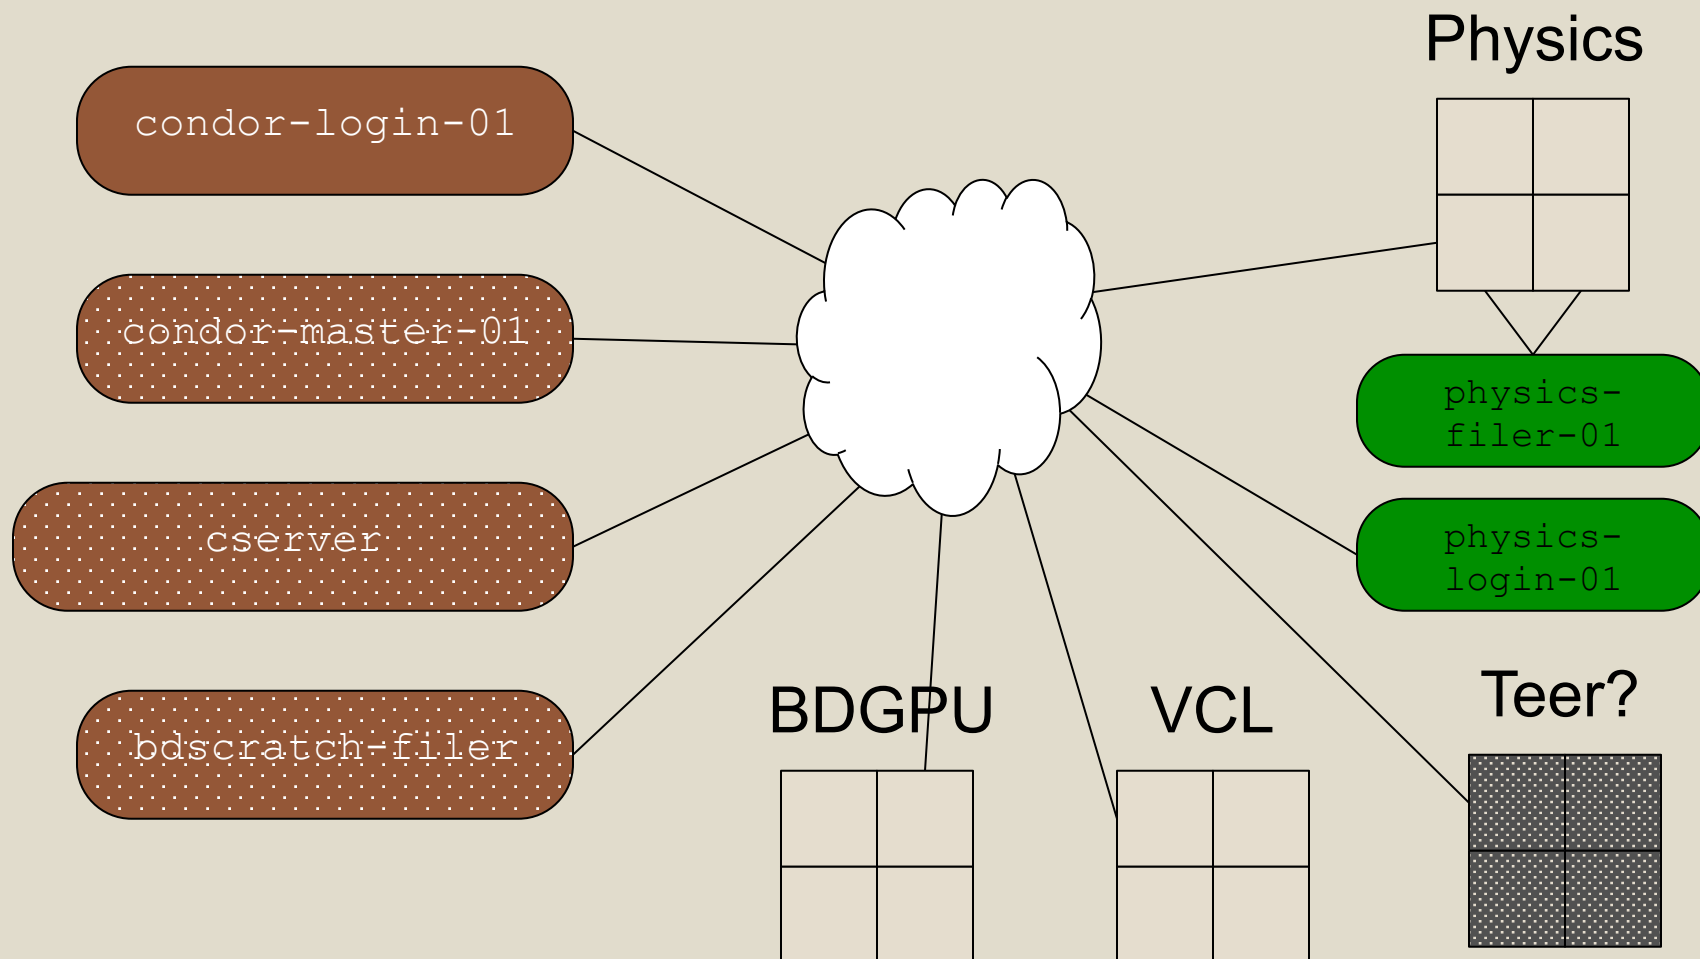
Condor, cont'd

- Condor allows (embraces?) more heterogeneity than the DSCR
 - ◆ This potentially means more work for end-users to make use of the resource
 - What machines/-types/“-sizes” can your job run on?
 - What input/output files does your job need?
 - How much time do you need?

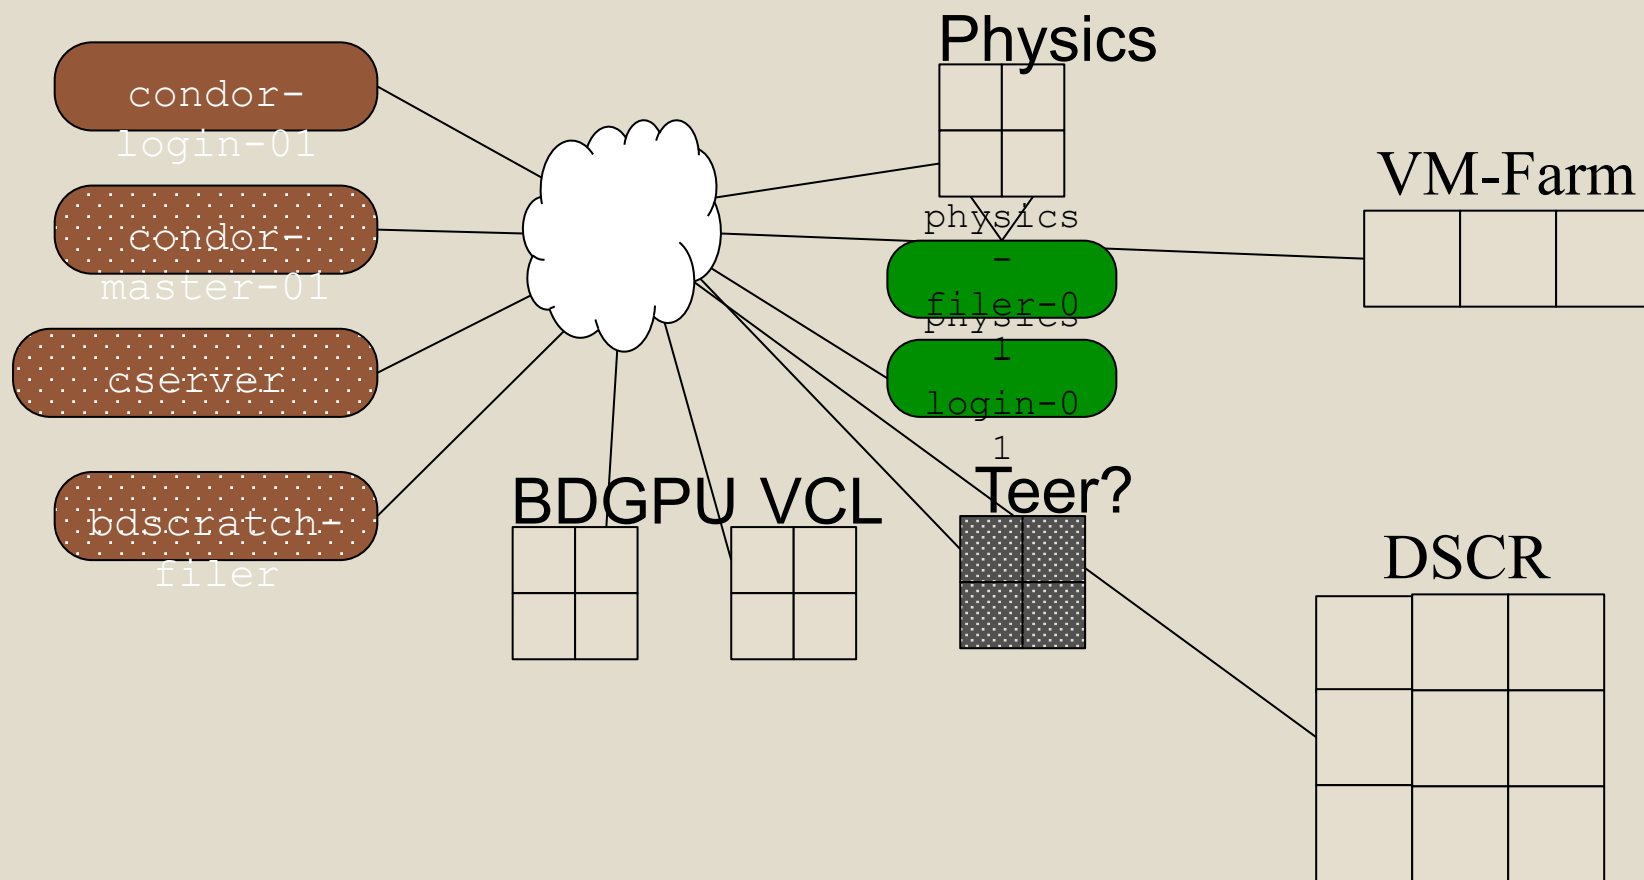
- But potentially gives access to a much larger set of resources
 - ◆ Especially with connection to OSG!



Duke Condor Architecture



Duke Condor Architecture (Future)



Make your job Condor-Ready

Must run in the background:

- No interactive input
- No GUI/Window Clicks
- Can Use STDIN, STDOUT, and STDERR through files instead of actual input devices
- Similar to Linux command:

```
$ ./myprogram <input.txt >output.txt
```

Really – this is making it “Batch-ready”