

# HTCondor & Cloud stuff...



# Game plan for today:

Introduction to Condor

Condor use patterns

StarCluster: On-demand clusters, in “The Cloud”!

Course logistics update

Thus far, we've largely talked about Hadoop...

M/R and Hadoop ecosystem is great, but:

... what if you've already got programs written?

... what if your data can't go on HDFS?

... what if your cluster doesn't include Hadoop?

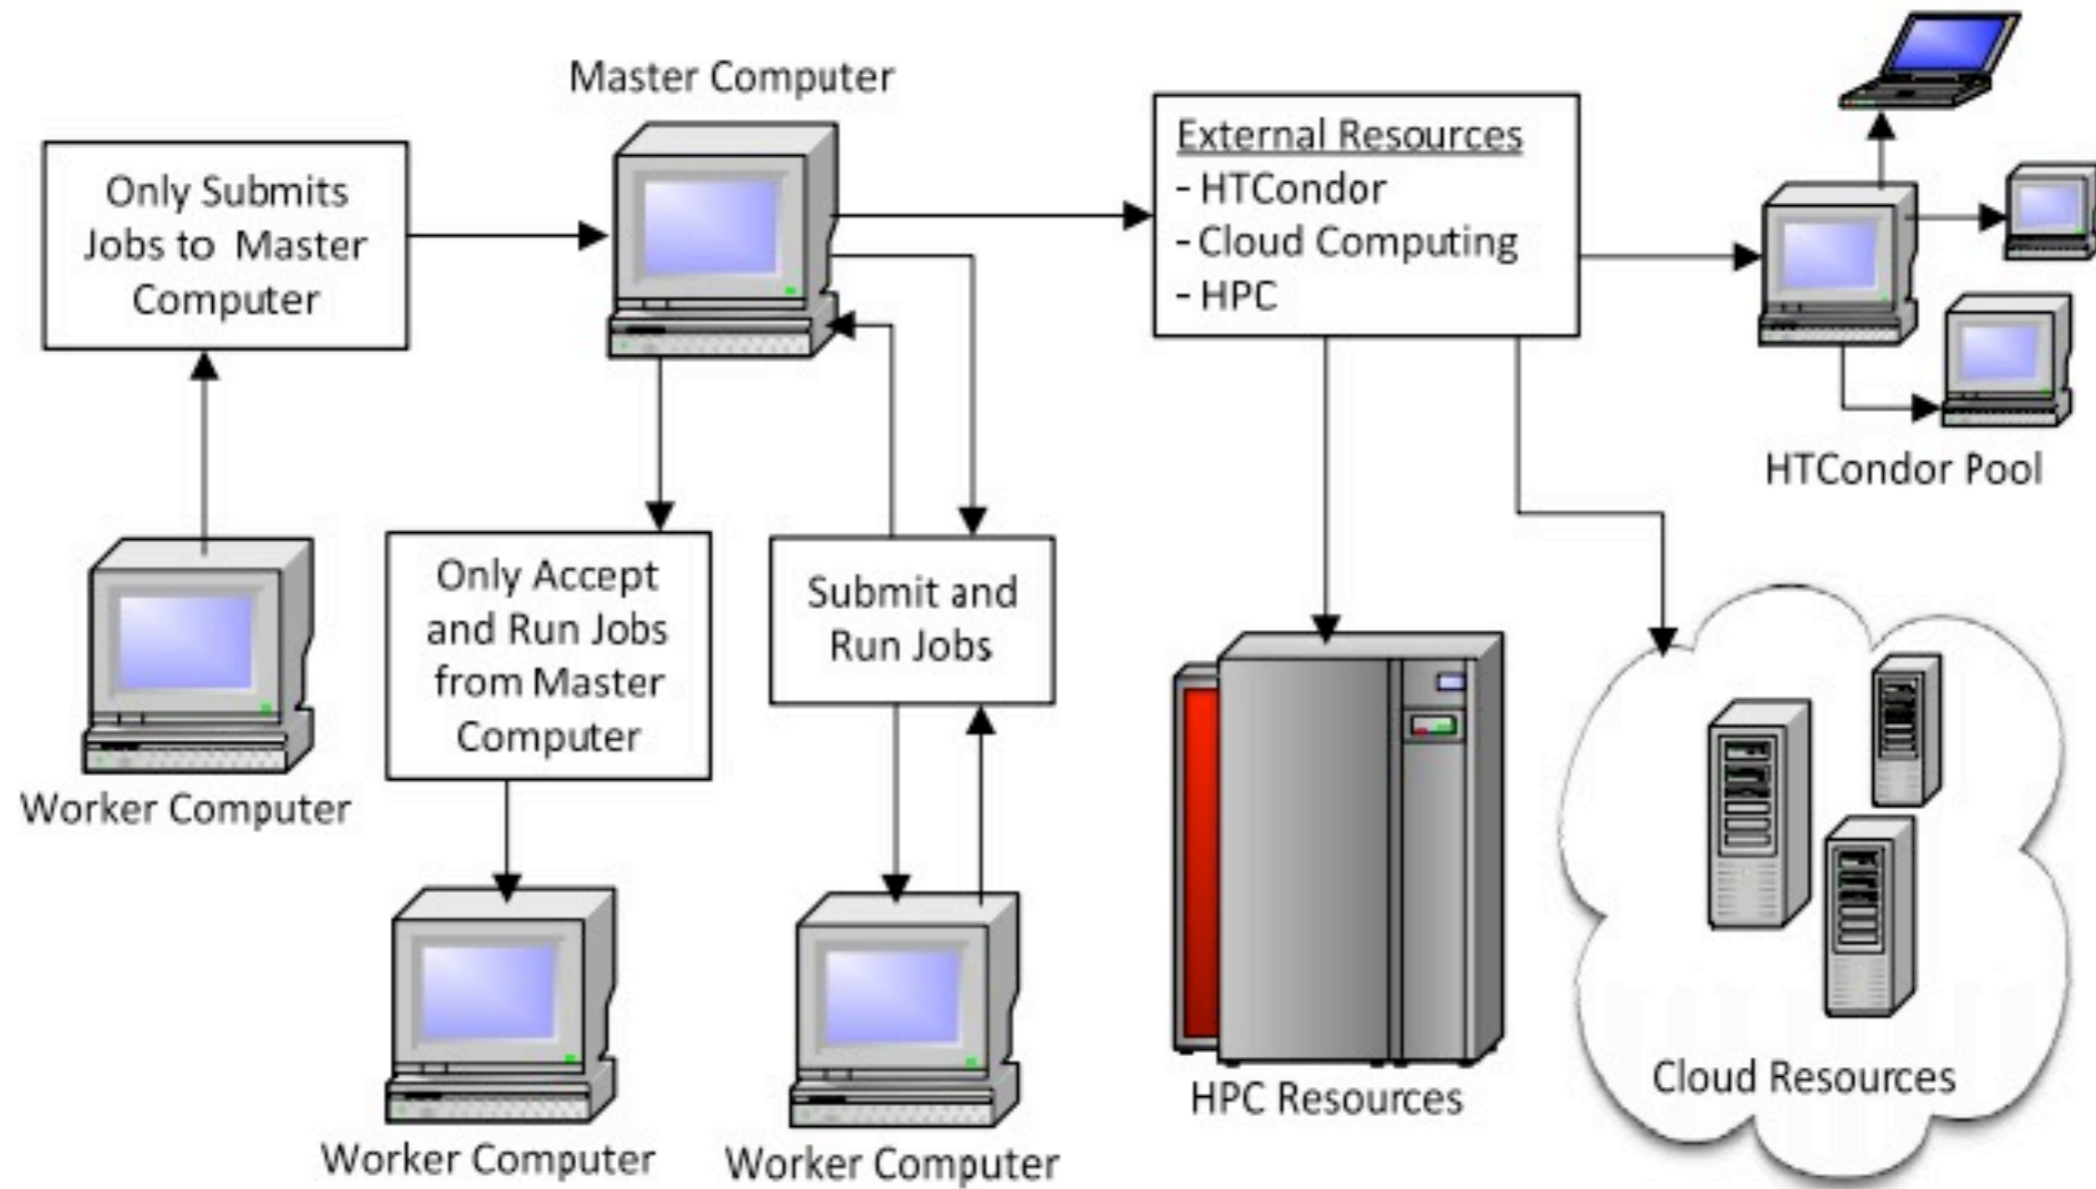
There are other paradigms!

HTCondor is an example of a “queue”-based clustering system.

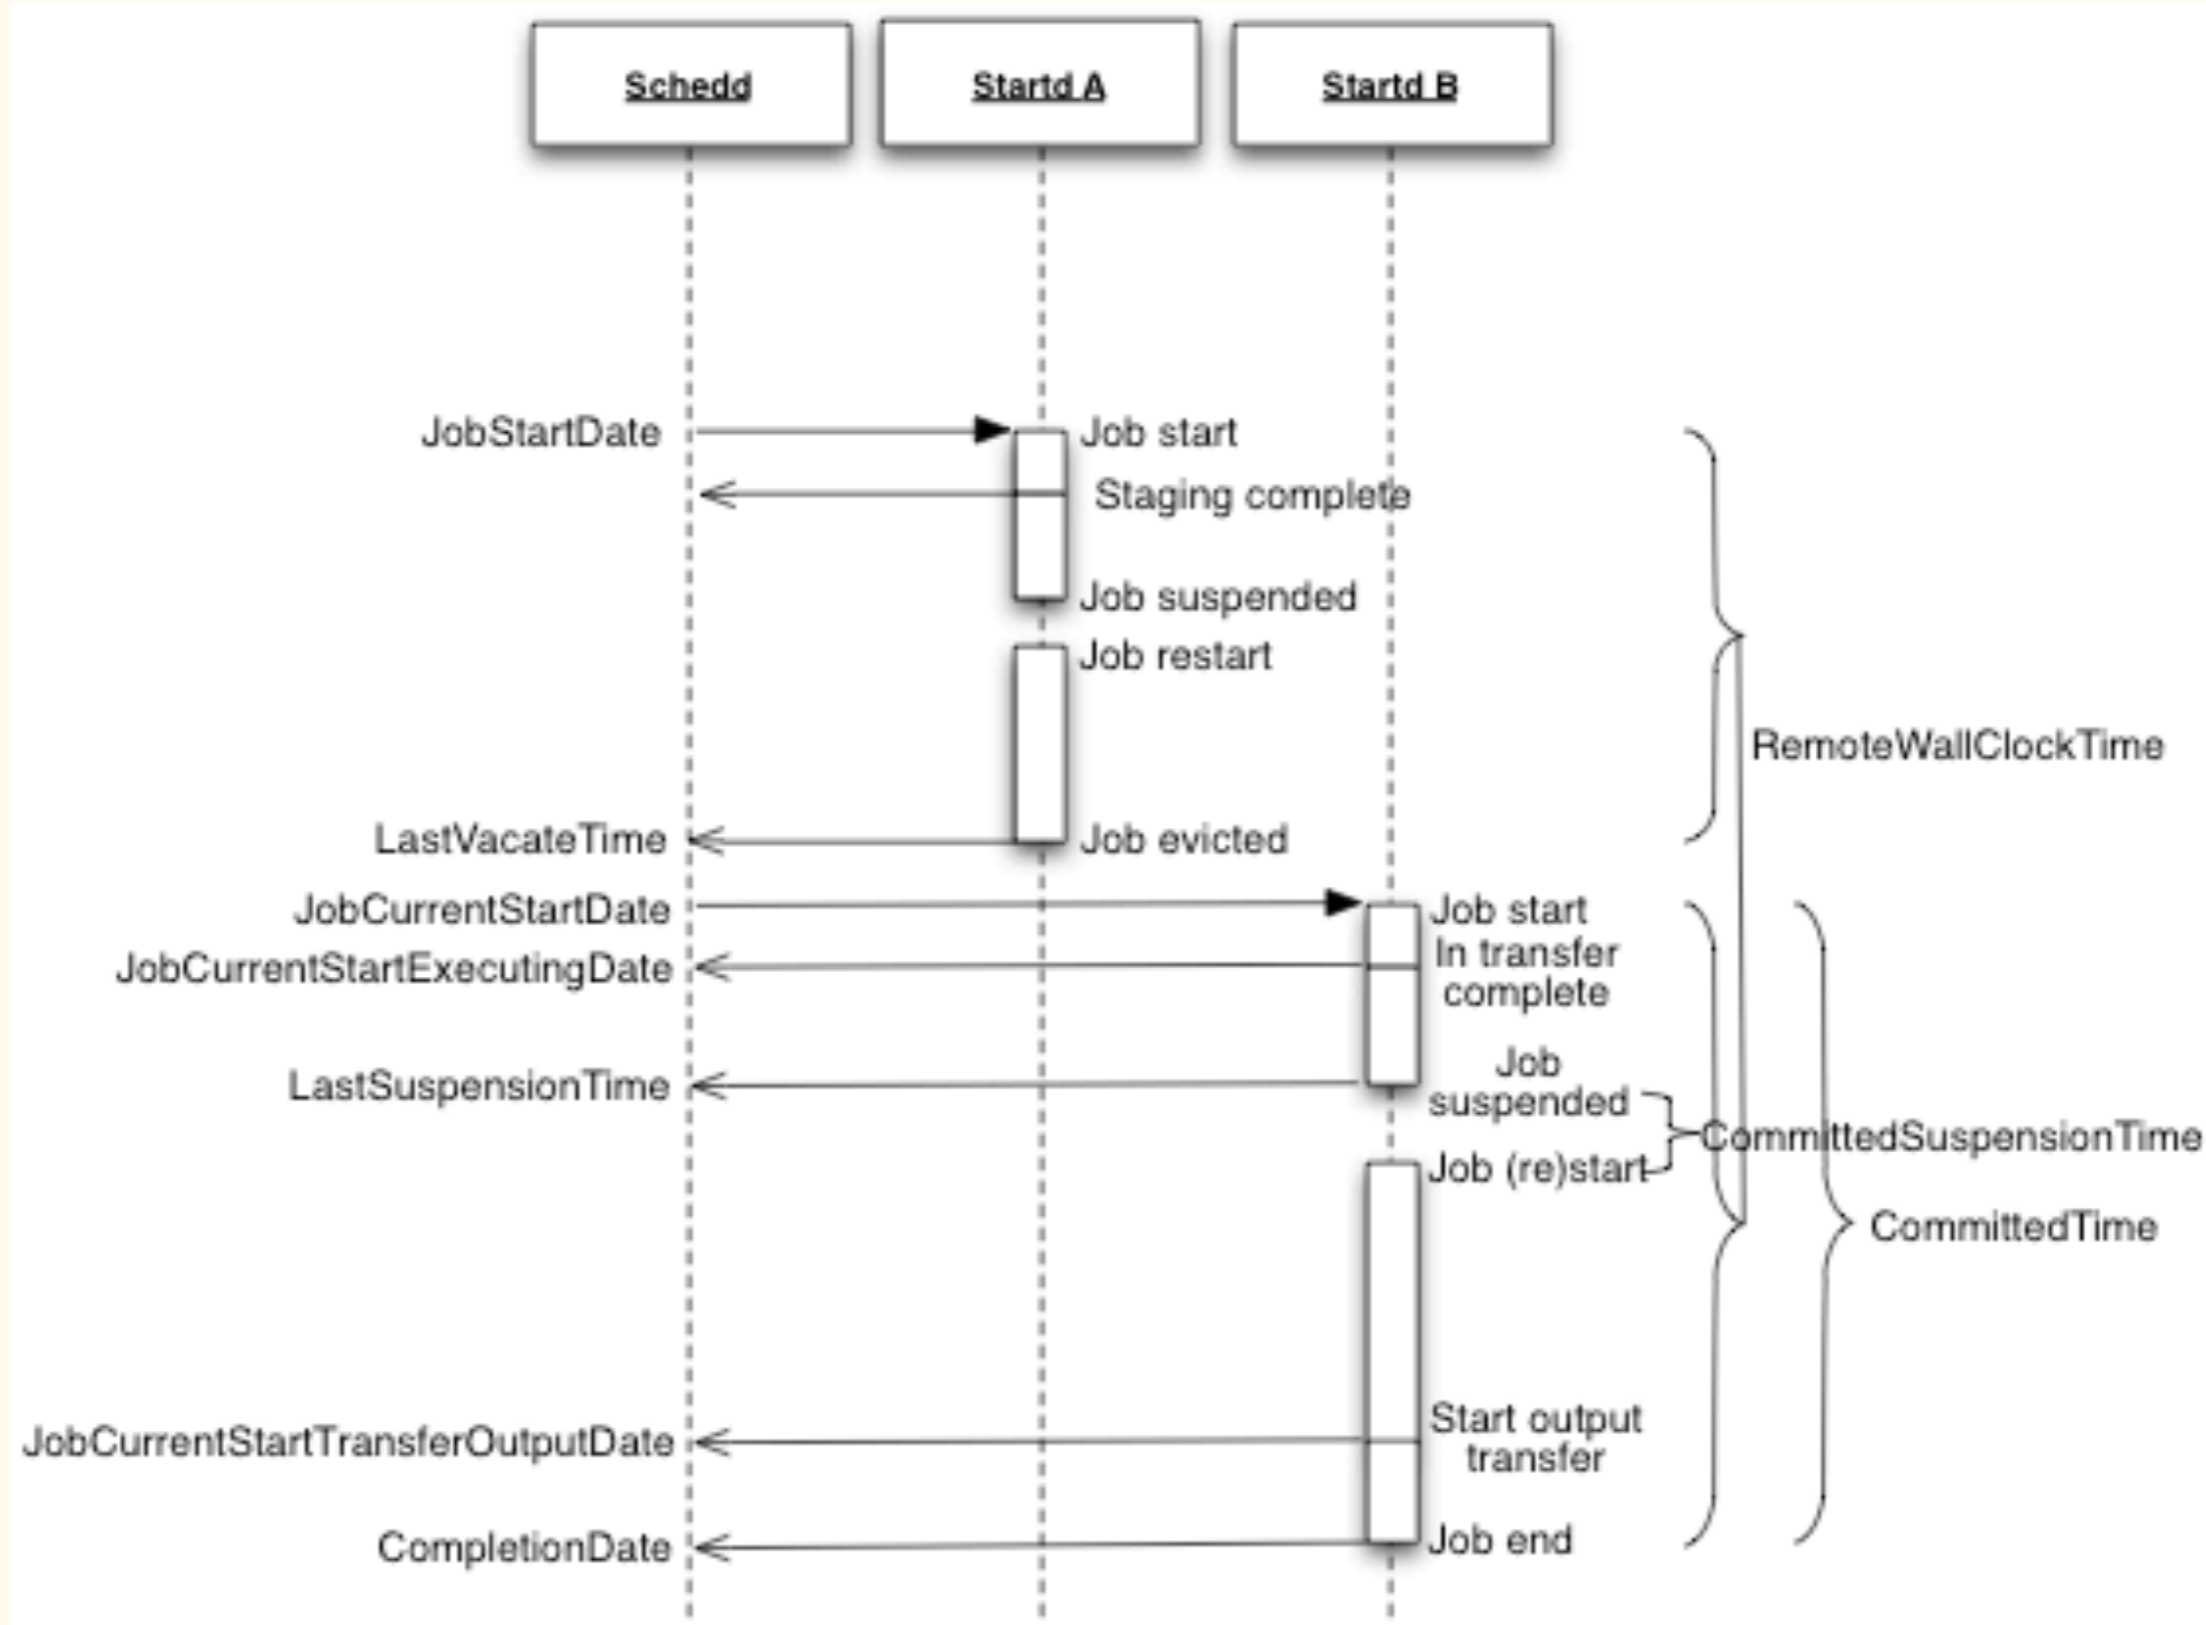
Users submit jobs, which are processed in order by worker machines.

Jobs specify what to do, where to do it, etc.

Others include SGE, OpenGrid, etc.







HTCondor was designed to work in a very heterogeneous environment...

... and so provides lots of flexibility in its configuration!

Condor is built around “slots” and “ads”

Compute nodes provide some number of “slots”...

... and advertise those slots capabilities:

CPU type, memory size, other capabilities, etc.

Jobs then submit requests specifying their runtime requirements: “> 1 gig of RAM”, “2 CPUs”, etc.

Condor figures out which slot should service which jobs!



# To run a condor job, first write a jobfile...

```
executable = run_monkeys.py
universe = vanilla
arguments = -brachiate=True -num_bananas=$(Process)

output = /g/steven/monkeys/out/$(Process).out
error = /g/steven/monkeys/err/$(Process).err
initialdir = /g/steven/monkey_data/

queue 10
```

# To run a condor job, first write a jobfile...

```
executable = run_monkeys.py
universe = vanilla
arguments = -brachiate=True -num_bananas=$(Process)

output = /g/steven/monkeys/out/$(Process).out
error = /g/steven/monkeys/err/$(Process).err
initialdir = /g/steven/monkey_data/

queue 10
```

## ... and submit it:

```
$ condor_submit monkey_job.condor
Submitting job(s).....
10 job(s) submitted to cluster 4.
```

# Jobfiles can contain all sorts of other information:

```
request_memory = 1 GB
```

```
Requirements = (Arch == "INTEL" && OpSys == "LINUX") || \  
                (Arch == "X86_64" && OpSys == "LINUX")
```

```
Requirements = (OpSysName == "Ubuntu")
```

```
request_GPUs = 1
```

# They can also provide hints to the scheduler:

```
TimeTaken = 60
```

```
+MaxExecutionTime = $(TimeTaken) # Short jobs get priority
```

```
+JobPrio = -$(TimeTaken)
```

```
Rank = memory # Use amount of RAM to break ties
```

# Useful Condor commands:

```
$ condor_q # list local Condor jobs
```

```
$ condor_q - global # list all Condor jobs
```

```
$ condor_q - global - submitter YOURNAME # list your jobs
```

```
$ condor_status # query the Condor system for other data
```

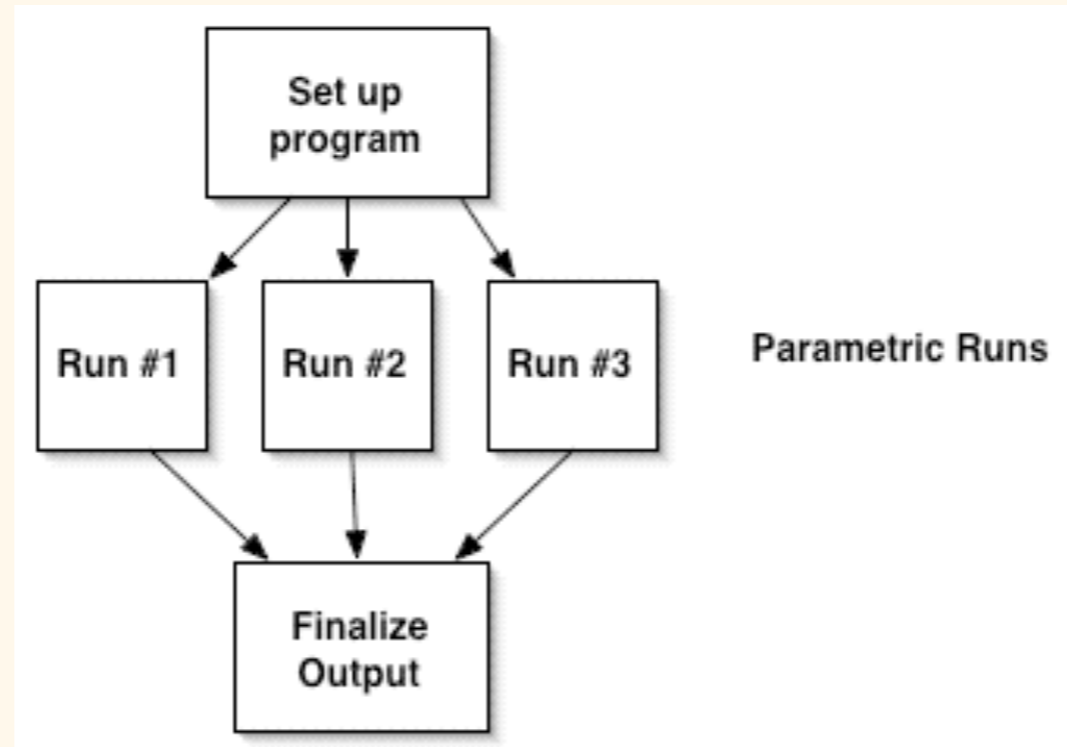
```
$ condor_status HOSTNAME # view slot info on HOSTNAME
```

```
# view stats on all hosts with >= 20 CPU cores
```

```
$ condor_status -format "%s " Name -format "%d " Cpus \  
-format "%d\n" Memory -constraint "Cpus >= 20"
```

```
$ condor_submit -interactive # get shell on a compute node
```

# Users can specify a DAG of related condor jobs:



```
Job Setup setup.submit
```

```
Job WorkerNode_1 run1.submit
```

```
Job WorkerNode_2 run2.submit
```

```
Job WorkerNode_3 run3.submit
```

```
Job Finalize finalize.submit
```

```
PARENT Setup CHILD WorkerNode_1 WorkerNode_2 WorkerNode_3
```

```
PARENT WorkerNode_1 WorkerNode_2 WorkerNode_3 CHILD Finalize
```

```
$ condor_submit_dag myjob.dag
```

# Other (potentially) useful Condor features:

- The “java” universe
- The “parallel” universe
- File transfer
  - Not often necessary, when there is shared storage...
- Checkpointing
  - Only sometimes possible...



# Condor use patterns:

Scenario: Process lots of input data in parallel

```
$ run_monkeys.py -brachiate True -num_bananas 4 -input my_file.txt
```

Solution: shard data, key shards by process ID:

```
executable = run_monkeys.py
```

```
universe = vanilla
```

```
initialdir = /g/steven/monkey_data/
```

```
arguments = -brachiate True -num_bananas 1 -input $(Process).txt
```

```
output = /g/steven/monkeys/out/$(Process).out
```

```
error = /g/steven/monkeys/err/$(Process).err
```

```
queue 100 # or however many shards there are
```

Use DAG to consolidate output?

# Condor use patterns:

Scenario: Grid-searching over parameter space

```
$ run_monkeys.py -brachiate True -num_bananas 1 -input my_input.txt
```

1. Programmatically generate submit file, queue once
2. Pass process ID in as argument, queue many

# 1. Programmatically generate submit file, queue once

```
executable = run_monkeys.py
```

```
universe = vanilla
```

```
initialdir = /g/steven/monkey_data/
```

```
arguments = -brachiate True -num_bananas 1 -input my_input.txt
```

```
output = /g/steven/monkeys/out/brachiate.True.num.1.out
```

```
error = /g/steven/monkeys/err/brachiate.True.num.1.err
```

```
queue
```

```
arguments = -brachiate True -num_bananas 2 -input my_input.txt
```

```
output = /g/steven/monkeys/out/brachiate.True.num.2.out
```

```
error = /g/steven/monkeys/err/brachiate.True.num.2.err
```

```
queue
```

```
...
```

```
arguments = -brachiate False -num_bananas 1 -input my_input.txt
```

```
output = /g/steven/monkeys/out/brachiate.False.num.1.out
```

```
error = /g/steven/monkeys/err/brachiate.False.num.1.err
```

```
queue
```

## 2. Pass process ID in as argument, queue many

```
executable = run_monkeys.py
```

```
universe = vanilla
```

```
initialdir = /g/steven/monkey_data/
```

```
arguments = -setting $(Process) -input my_input.txt
```

```
output = /g/steven/monkeys/out/$(Process).out
```

```
error = /g/steven/monkeys/err/$(Process).err
```

```
queue 200
```

```
true 1
```

```
true 2
```

```
true 3
```

```
true 4
```

```
...
```

```
true 100
```

```
false 1
```

```
false 2
```

```
...
```

```
should_brachiate = parameters[args.setting][0]
```

```
num_bananas = parameters[args.setting][1]
```

# Game plan for today:

Introduction to Condor

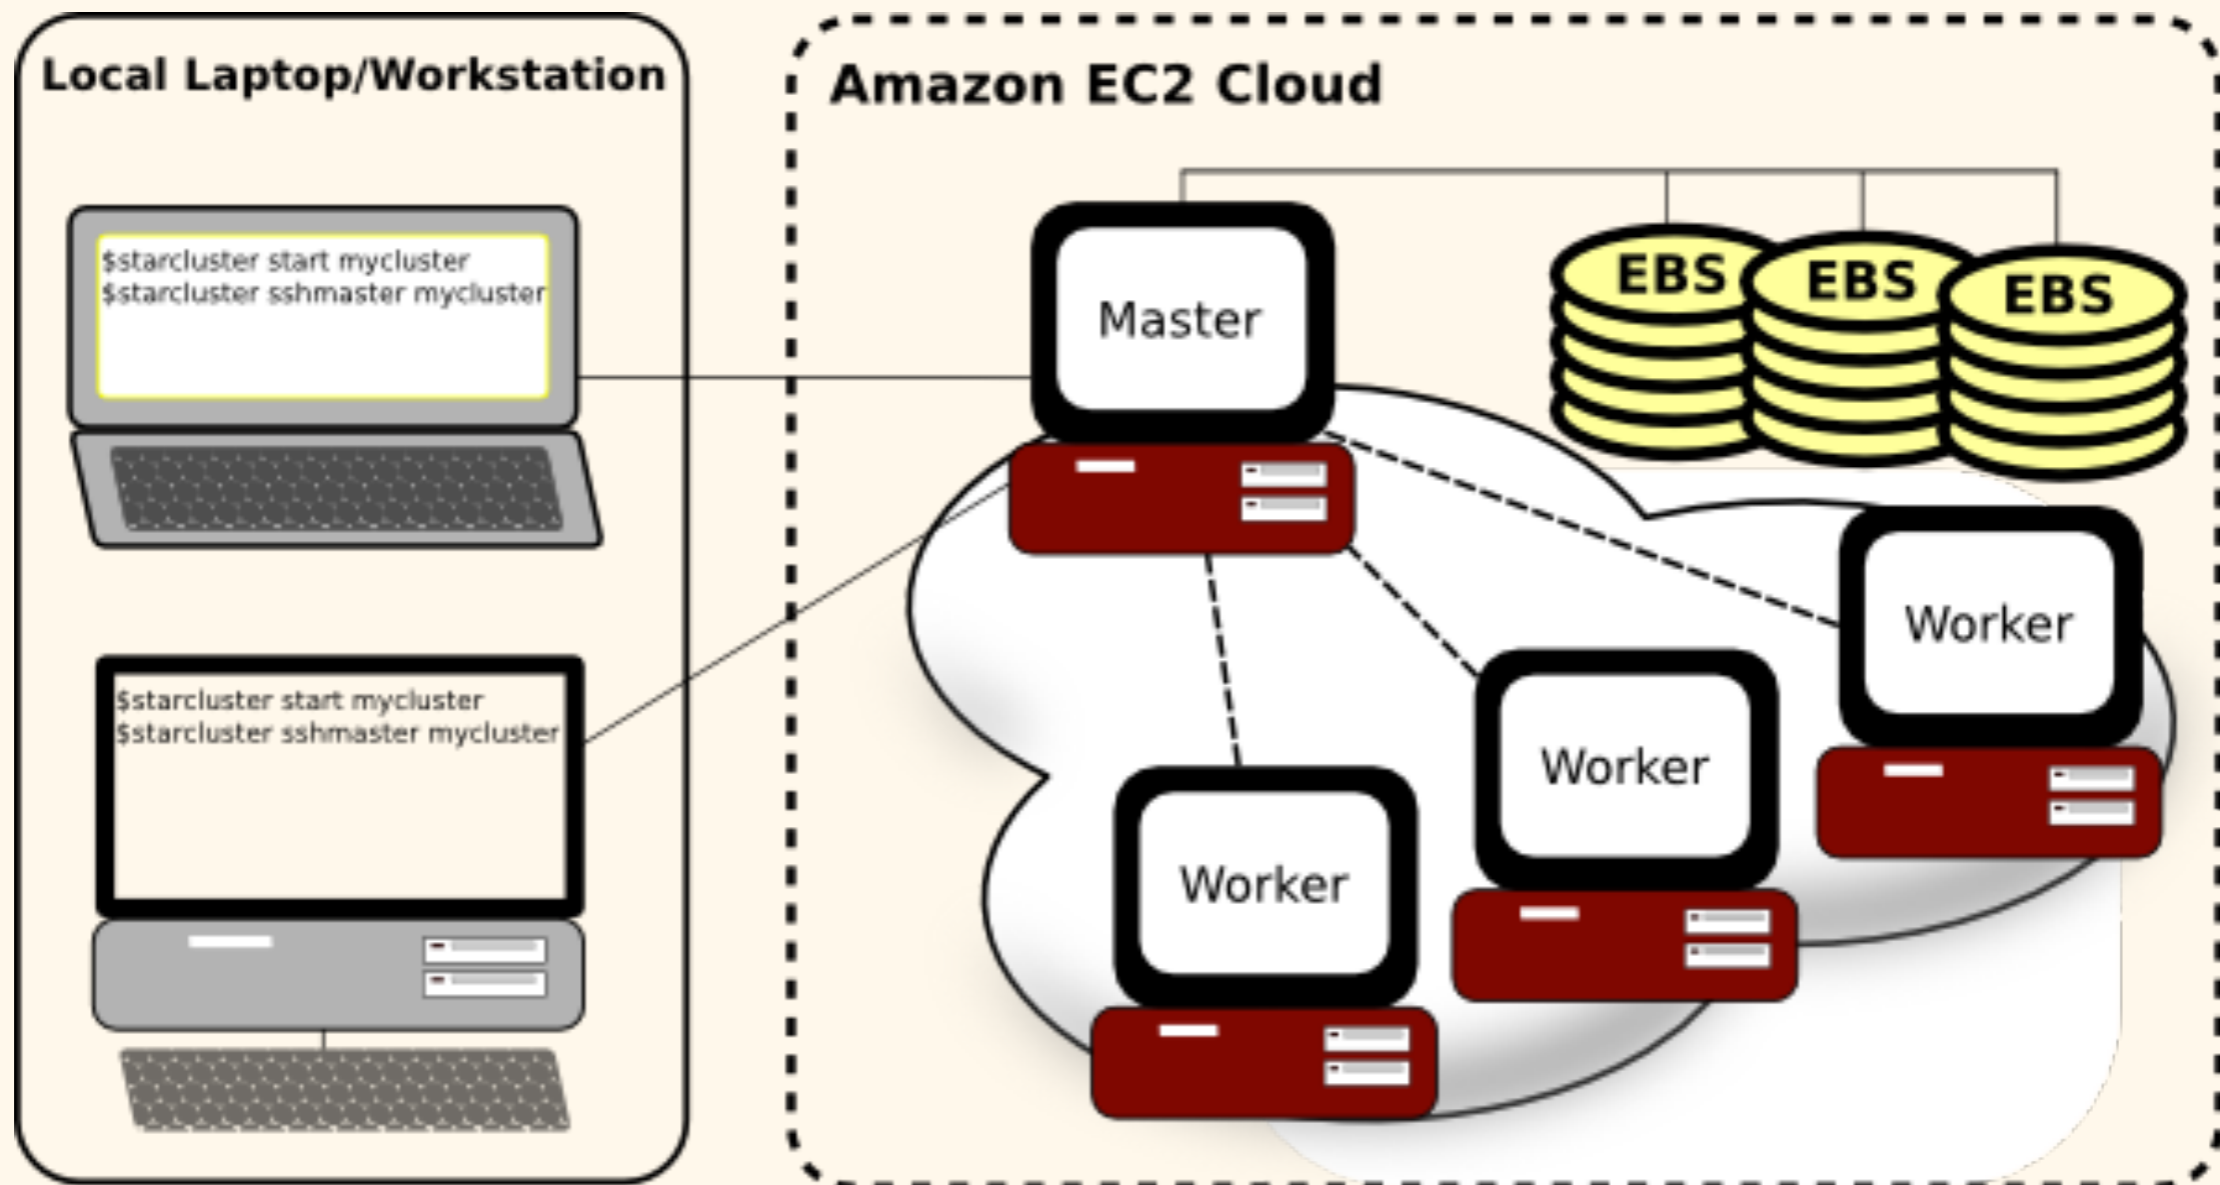
Condor use patterns

StarCluster: On-demand clusters, in “The Cloud”!

Course logistics update

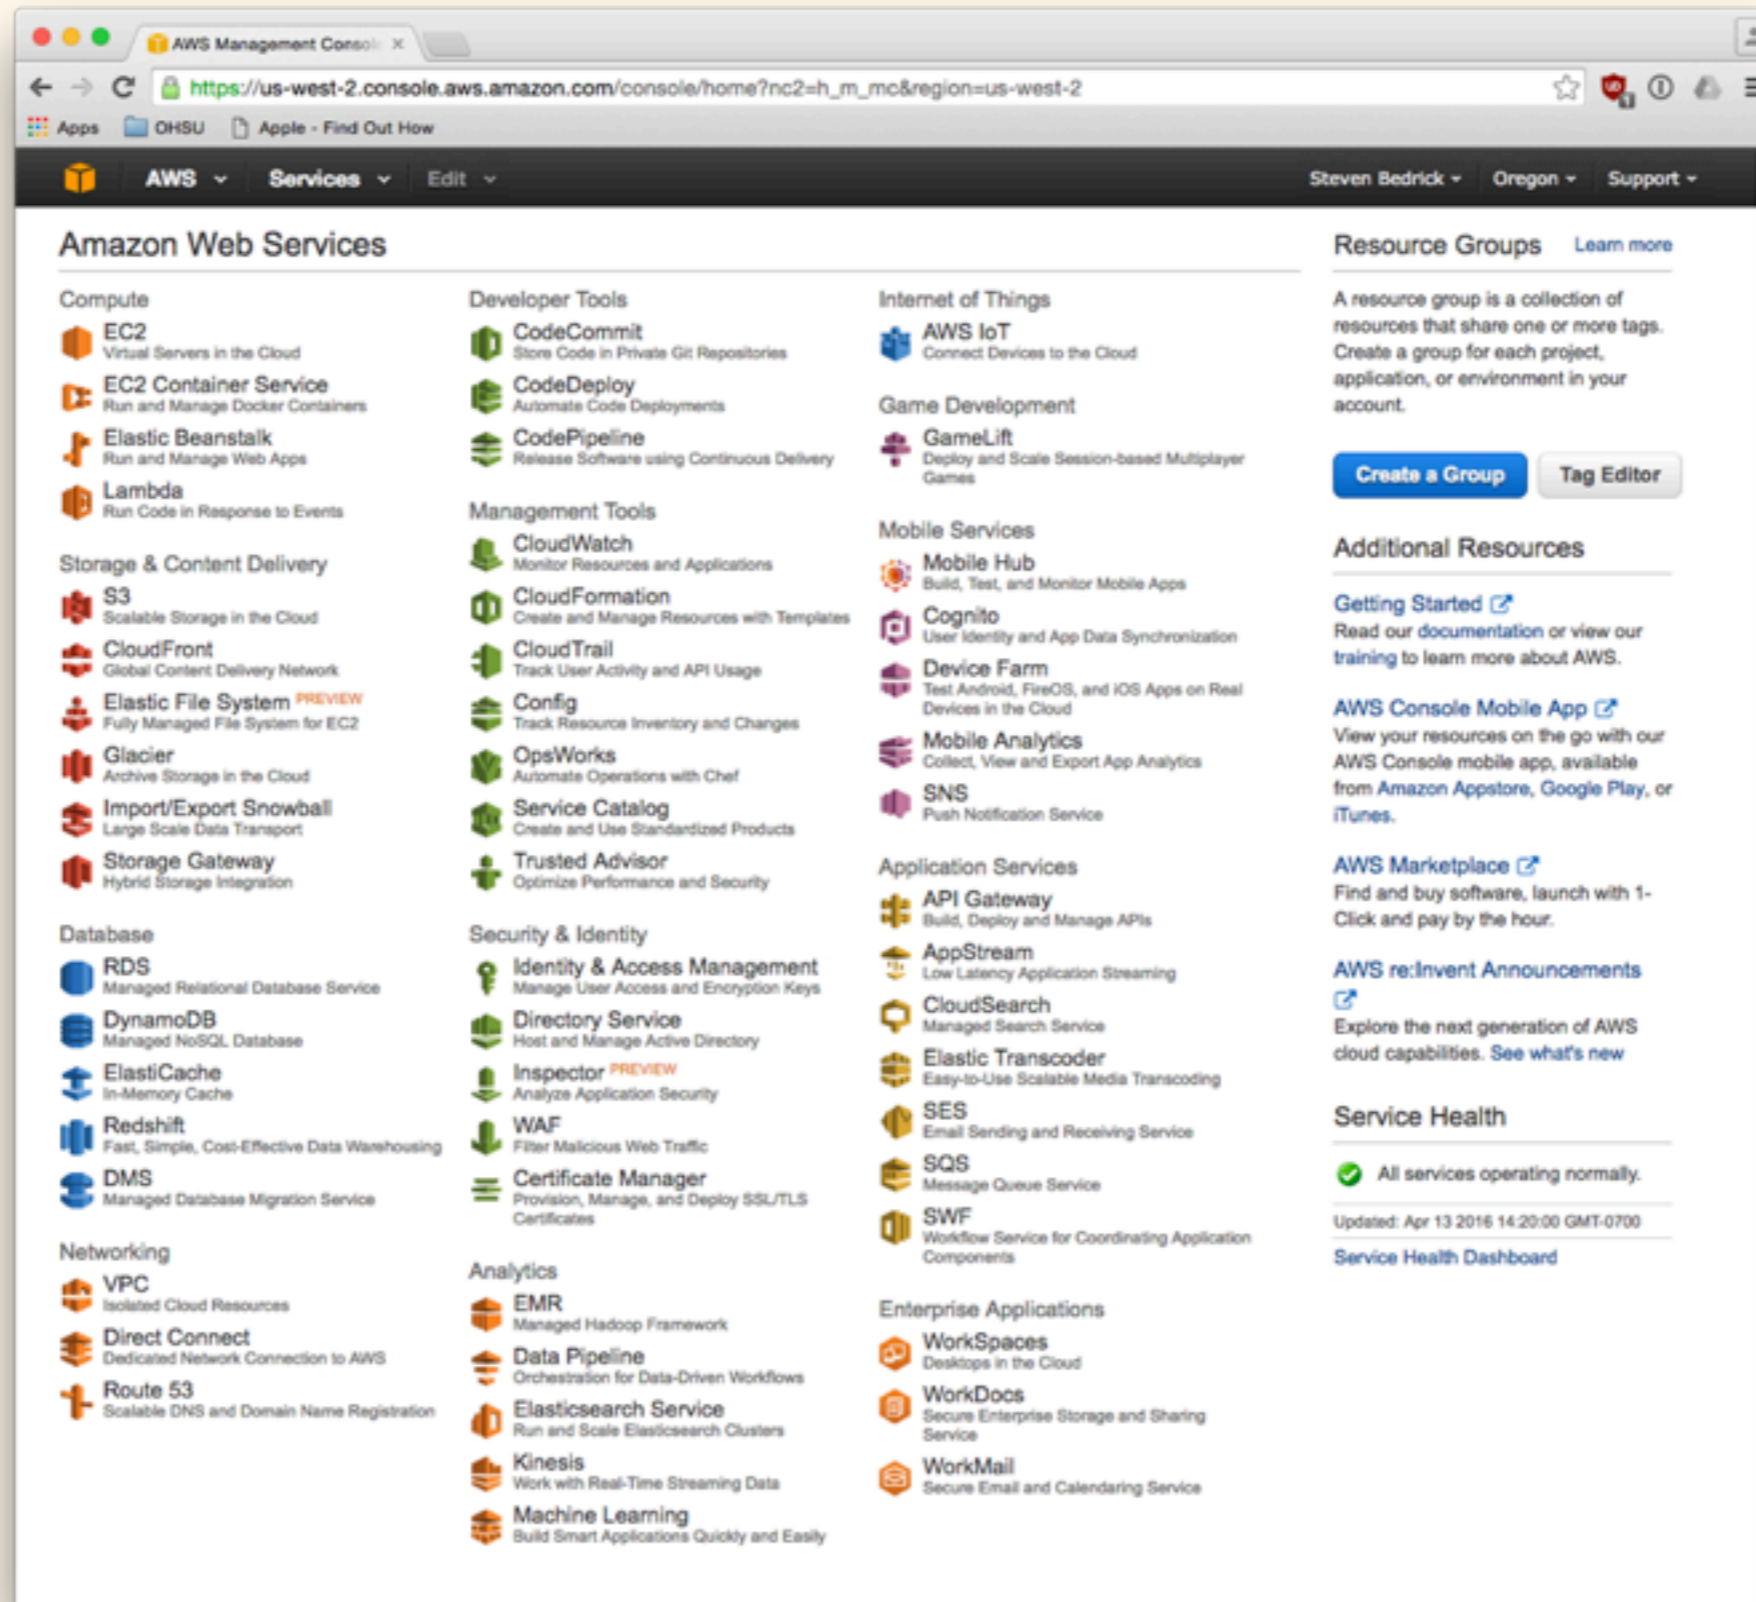
# StarCluster:

A Python library to automatically setup and configure clusters on AWS!





# AWS is Amazon.com's suite of cloud services:



# Key ideas:

EC2 “instances” are virtual machines...

Instances can be of many sizes:

Compute Optimized

## C4

C4 instances are the latest generation of Compute-optimized instances, featuring the highest performing processors and the lowest price/compute performance in EC2.

**Features:**

- High frequency Intel Xeon E5-2666 v3 (Haswell) processors optimized specifically for EC2
- EBS-optimized by default and at no additional cost
- Ability to control processor C-state and P-state configuration on the c4.8xlarge instance type
- Support for [Enhanced Networking](#) and Clustering

Model	vCPU	Mem (GiB)	Storage	Dedicated EBS Throughput (Mbps)
c4.large	2	3.75	EBS-Only	500
c4.xlarge	4	7.5	EBS-Only	750
c4.2xlarge	8	15	EBS-Only	1,000
c4.4xlarge	16	30	EBS-Only	2,000
c4.8xlarge	36	60	EBS-Only	4,000

<https://aws.amazon.com/ec2/instance-types/>

Payment by cpu time, varies by instance size

# Key ideas:

EBS = “Elastic Block Volumes”

Virtual storage, can be of (almost) any size or type.

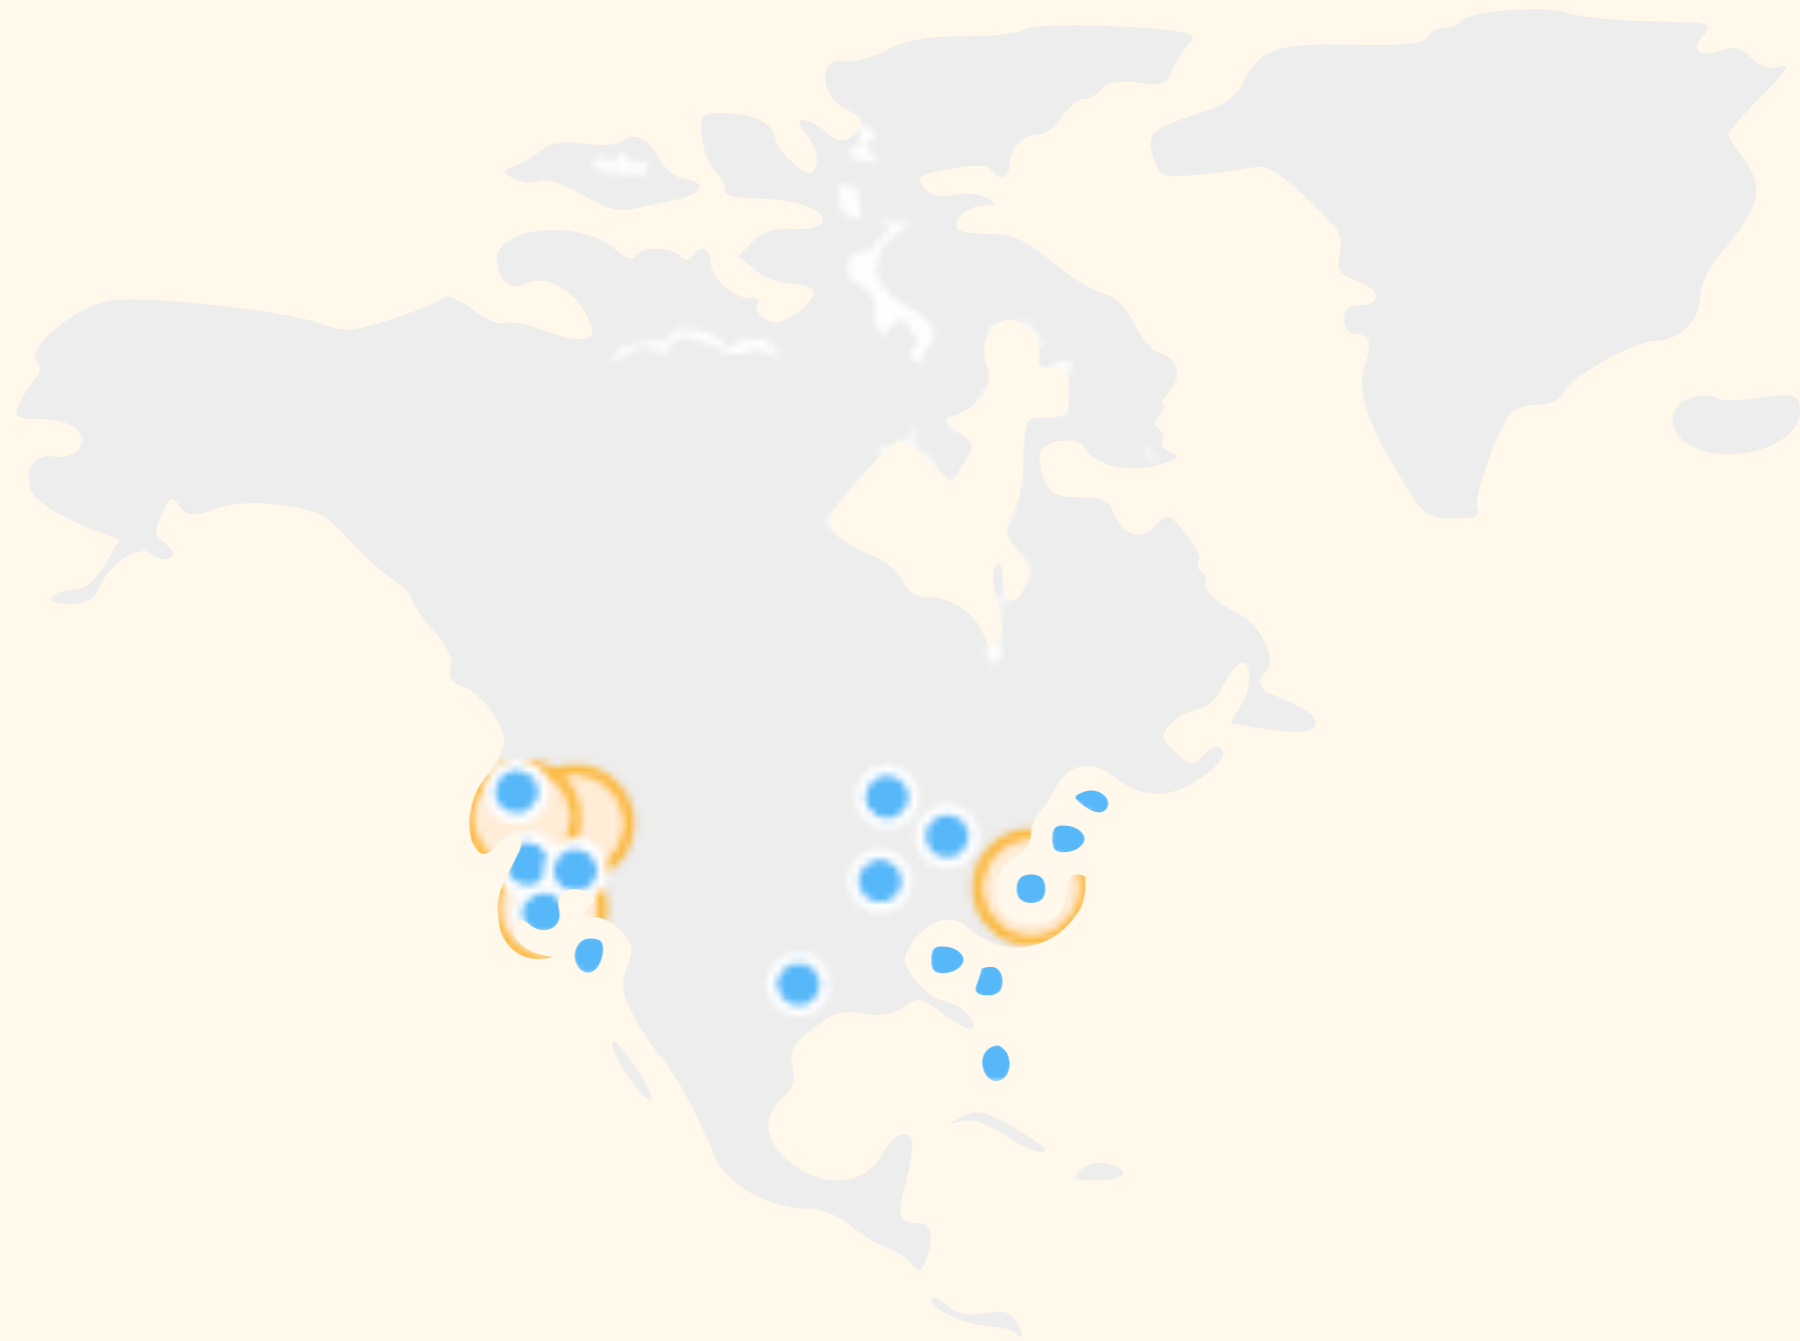
Payment based on space used and amount of I/O.

S3 = “Simple Storage System”

“Bucket” of data in the sky

Payment based on space used and number of requests





Setting up an EC2 instance is easy...

... configuring many instances gets old fast.

Enter StarCluster!