

The Sky's the Limit:

Preserving Research Data with
Clouds
describing

HEP Legacy Data Project


Network-Enabled Programs Project

funded by



canarie

Canada's Advanced Research and Innovation Network
Le réseau évolué de recherche et d'innovation du Canada



The Sky's the Limit: Preserving Research Data with Clouds

HEP Legacy Data Project team



Wayne Podaima (NRC), Roger Impey (NRC), Ashok Agarwal (UVIC),
Duncan Penfold-Brown (UVIC), Kyle Fransham (UVIC), Ron Desmarias (UVIC), Patrick Armstrong (UVIC),
Andre Charbonneau (NRC), Randall Sobie (UVIC), Colin Leavett-Brown (UVIC)

...and not forgetting our photographer




Ian Gable (UVIC)



The Sky's the Limit: Preserving Research Data with Clouds






The Sky's the Limit: Preserving Research Data with Clouds

The Big Bang and matter/anti-matter asymmetry



- Big Bang produce matter & anti-matter
- But today only matter
- Where did all the anti-matter go?



Three Generations of Matter (Fermions)

	I	II	III	
mass→	2.4 MeV	1.27 GeV	171.2 GeV	0
charge→	$\frac{2}{3}$	$\frac{2}{3}$	$\frac{2}{3}$	0
spin→	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1
name→	u up	c charm	t top	γ photon
	4.8 MeV	104 MeV	4.2 GeV	0
	$-\frac{1}{3}$	$-\frac{1}{3}$	$-\frac{1}{3}$	0
	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1
Quarks	d down	s strange	b bottom	g gluon
	<2.2 eV	<0.17 MeV	<15.5 MeV	91.2 GeV
	0	0	0	0
	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1
	ν_e electron neutrino	ν_μ muon neutrino	ν_τ tau neutrino	Z weak force
	0.511 MeV	105.7 MeV	1.777 GeV	80.4 GeV
	-1	-1	-1	±1
	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1
Leptons	e electron	μ muon	τ tau	W weak force


Bosons (Forces)





The Sky's the Limit: Preserving Research Data with Clouds


SLAC colliding matter with anti-matter






The Sky's the Limit: Preserving Research Data with Clouds

Recording the data with the BaBar detector




- Electron/positron collision
- ~500 TB data captured, 2000 to 2008






Mathematics & Histograms

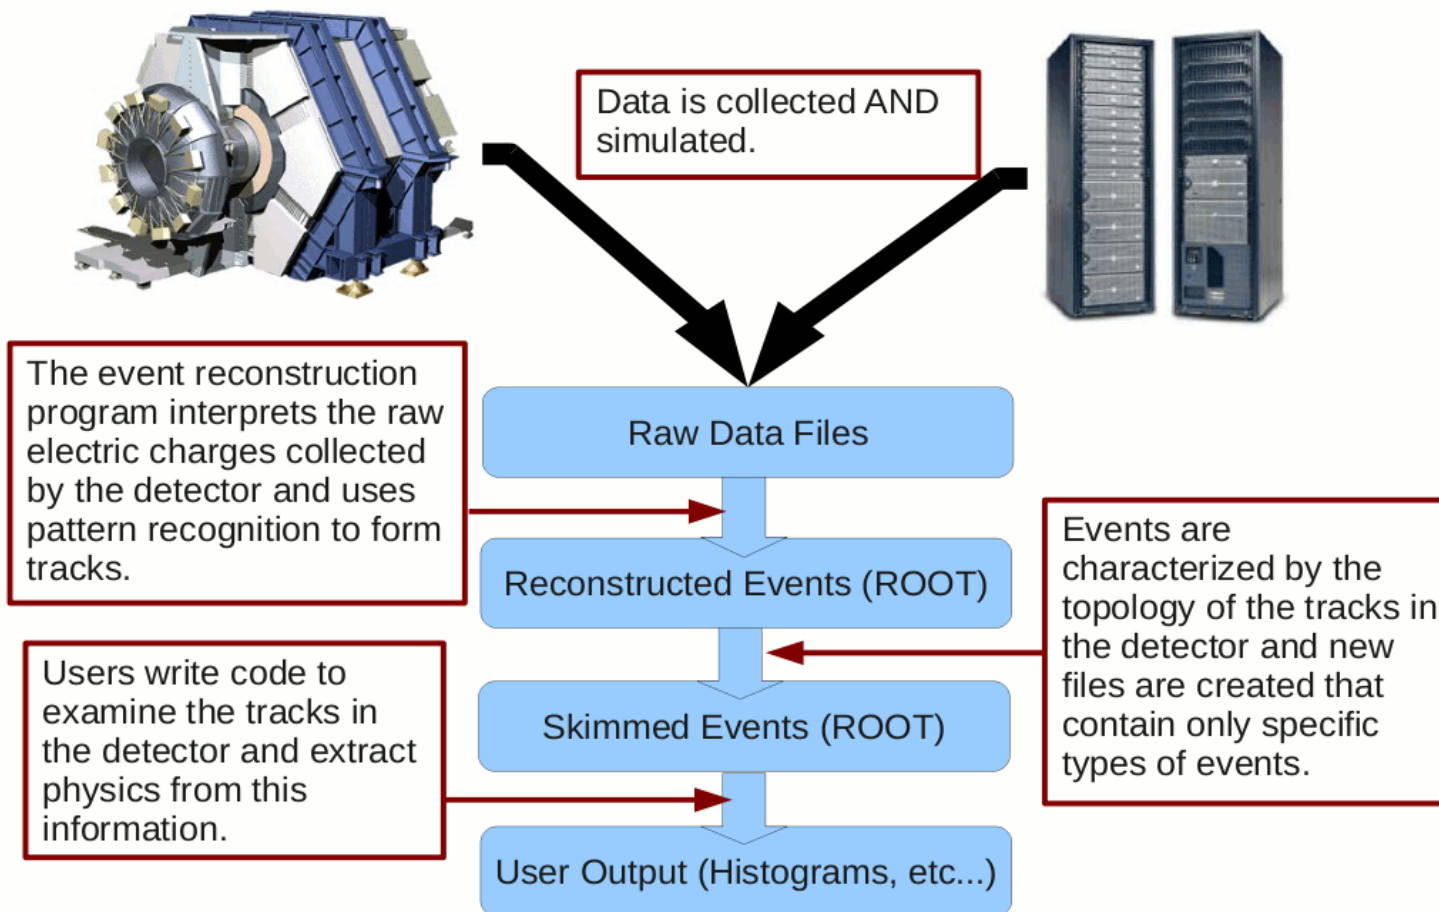
$$C_{ij}^{(2b)} = \Delta_{21}^2 (\lambda_i - \Delta_{31})(\lambda_j - \Delta_{31}) s_{12}^2 c_{11}^2$$






The Sky's the Limit: Preserving Research Data with Clouds

BaBar Software Chain






The Sky's the Limit: Preserving Research Data with Clouds


BaBar is a Complex Application



+




+



Issues:

- 9.5 million lines of C++ & Fortran
- Requires particular OS, compiler, & library levels
- Certification of environment required
- Not easy to share environment
- Sysadmins almost need to become application specialists



The Sky's the Limit: Preserving Research Data with Clouds


What to do about aging environment?



+




+



Goals:


- Keep running for the next 5 to 10 years
- Able to use any available resource
- Easy to adopt (for everyone)




The Sky's the Limit: Preserving Research Data with Clouds

Four pieces needed for our solution:

- Application encapsulation
- Image replication
- User control
- Resource utilization
- Xen, KVM, etc.







The Sky's the Limit: Preserving Research Data with Clouds

Four pieces needed for our solution:

- Application encapsulation
- Image replication
- User control
- Resource utilization
- Xen, KVM, etc.



- Web service to manage VMs
- One interface per cluster
- Nimbus, Eucalyptus, OpenNebula, EC2




The Sky's the Limit: Preserving Research Data with Clouds


Four pieces needed for our solution:

- Application encapsulation
- Image replication
- User control
- Resource utilization
- Xen, KVM, etc.

- Web service to manage VMs
 - One interface per cluster
 - Nimbus, Eucalyptus, OpenNebula, EC2



- Dynamic resources
- Needs good API
- Condor, SGE




The Sky's the Limit: Preserving Research Data with Clouds

Four pieces needed for our solution:


- Application encapsulation
- Image replication
- User control
- Resource utilization
- Xen, KVM, etc.

- Web service to manage VMs
 - One interface per cluster
 - Nimbus, Eucalyptus, OpenNebula, EC2



- Dynamic resources
- Needs good API
- Condor, SGE


- Shares clouds among users
- Controls distributed clouds




The Sky's the Limit: Preserving Research Data with Clouds

How does it work?

1. Preparation of cloud infrastructure







The Sky's the Limit: Preserving Research Data with Clouds

How does it work?

2. Preparation of virtual machine images





The Sky's the Limit: Preserving Research Data with Clouds

How does it work?

3. User creates a job script


```
Executable = /home/babar/SP26.0.0/workdir/runJob.sh
Universe   = vanilla
Log        = condor.log
Output     = moose.output
Error      = moose.error
InitialDir = /home/babar/SP26.0.0/workdir
#should_transfer_files = yes
#WhenToTransferOutput = ON_EXIT


Requirements = VMType =?= "bbrSim"
+VMName      = "bbrSim"
+VMLoc       = "http://alto.cloud.nrc.ca/bbrSim"
+VMNetwork   = "public"
+VMCPUArch   = "x86"
+VMStorage   = "1"
+VMCPUCores  = "1"
+VMMem       = "2048"
getenv       = True

Queue
```



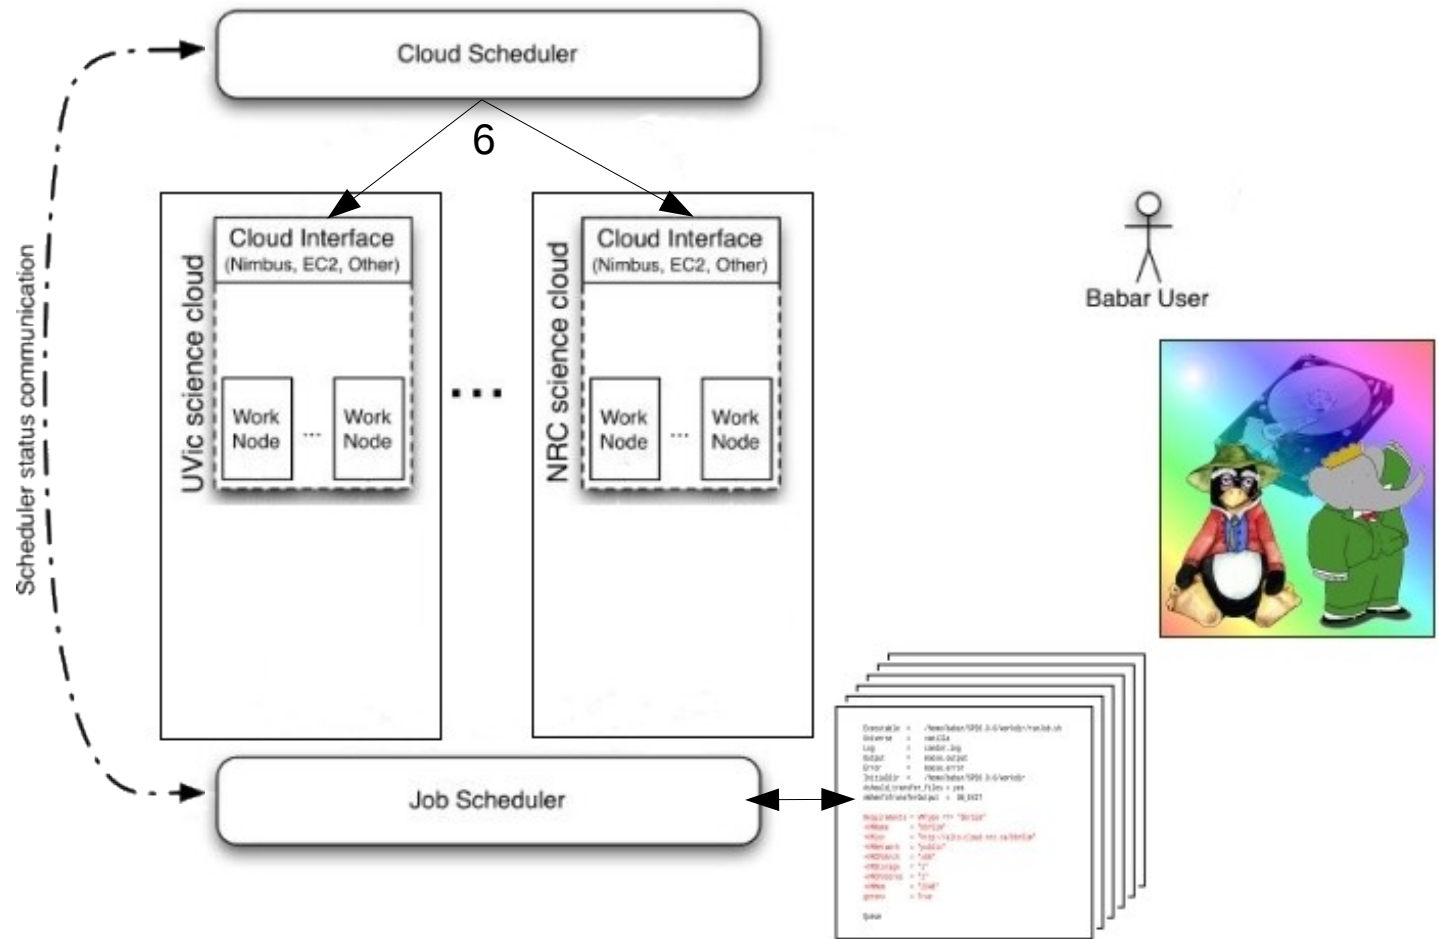

4. User submits jobs to Job Scheduler






How does it work?

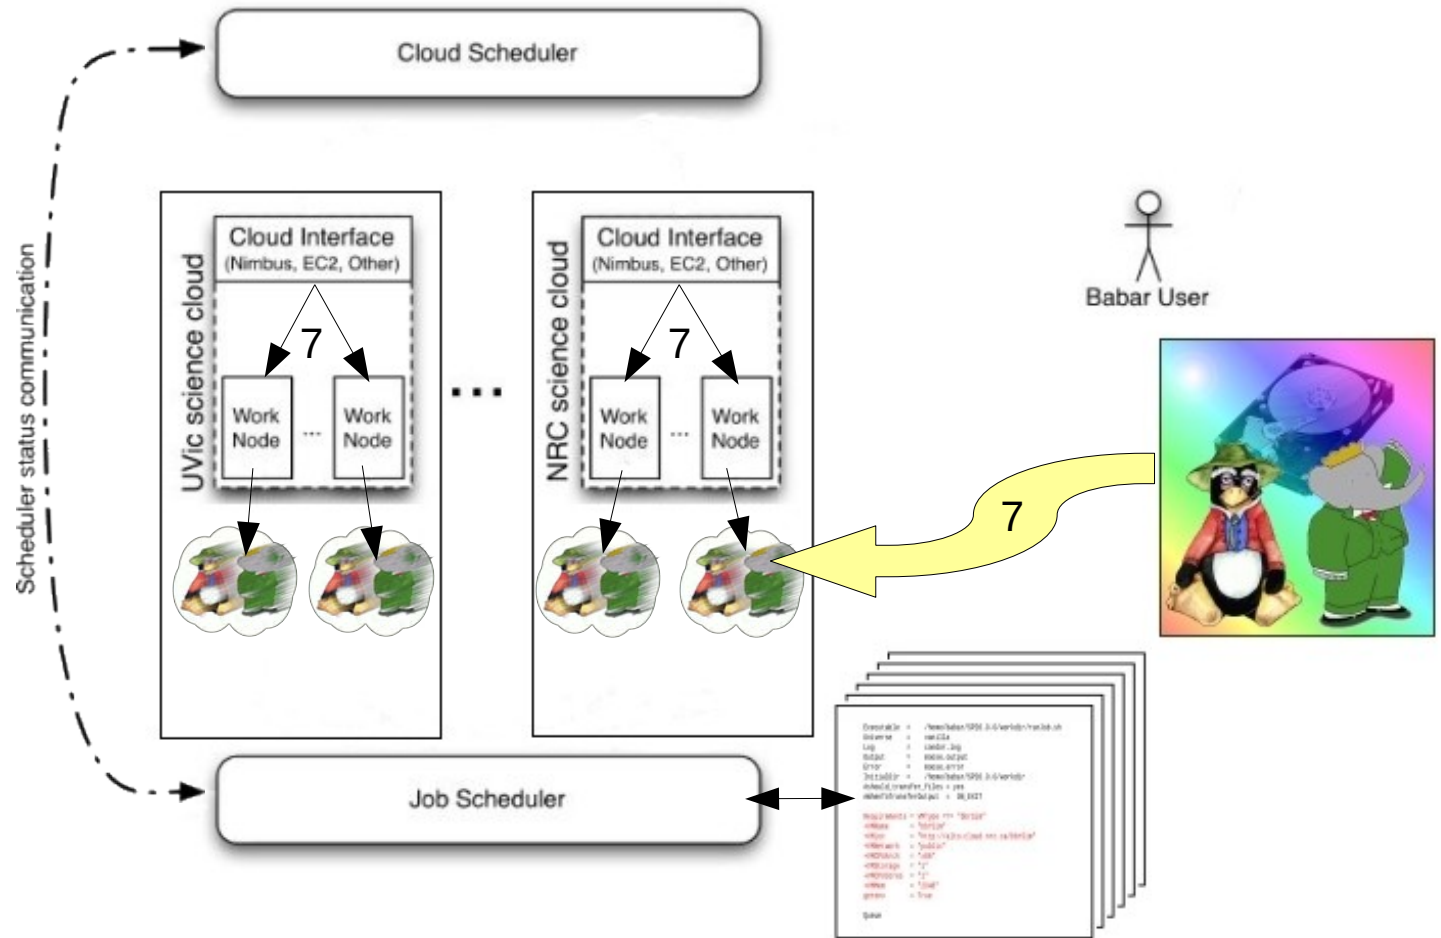
6. Cloud Scheduler requests deployment of VMs






How does it work?


7. IaaS deploys the requested VMs






How does it work?

9. Cloud Scheduler detects idle VMs



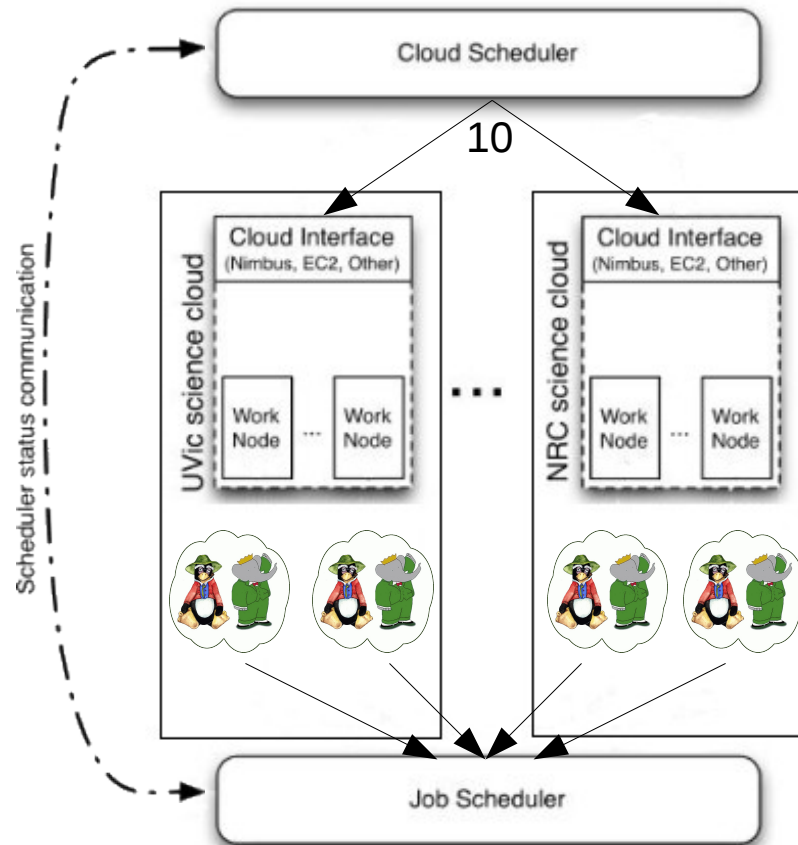
Babar User






How does it work?

10. Cloud Scheduler requests shutdown and destruction of VMs



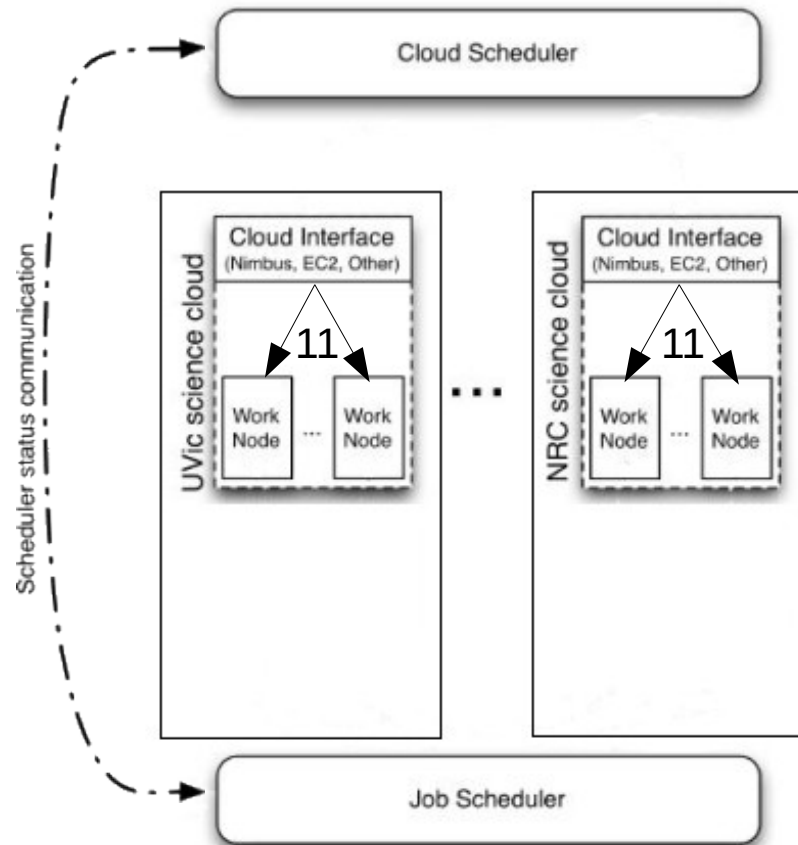
Babar User






How does it work?

11. IaaS destroys the requested VMs



Babar User







The Sky's the Limit: Preserving Research Data with Clouds

Does it work?

1. Single site test in Ottawa






The Sky's the Limit: Preserving Research Data with Clouds

Does it work?

2. Using multiple clouds

Initial Distribution Test: 300 Jobs to Clouds at NRC and UVic






The Sky's the Limit: Preserving Research Data with Clouds

Next Steps

- Incorporate Amazon EC2 within our distributed cloud environment
 - Awarded research grant of US \$3500 of time on Amazon EC2
- Commence BaBar simulation production; minimum 3 clouds
- Augment/improve image and data management; functionality and performance
 - Awarded US \$5000 Google summer of code
- Greater emphasis on BaBar analysis
- Reliability and scalability improvements
- Work with WestGrid to incorporate cloud capabilities on their UVIC installation.



The Sky's the Limit: Preserving Research Data with Clouds


Wider Application




- Jobs are embarrassingly parallel much like HEP.
- Six different surveys, each require a different processing environment, which require:
 - A specific version of a Linux distribution
 - A specific compiler version
 - Specific libraries




NGVS VM



PAndAS VM



SASSy VM




The Sky's the Limit: Preserving Research Data with Clouds

More Information:

- <http://cloudscheduler.org>
- Code on GitHub:
 - <http://github.com/hep-gc/cloud-scheduler>
 - <http://wiki.github.com/hep-gc/cloud-scheduler/cloud-scheduler-test-drive>
- <http://heprc.phys.uvic.ca/>

And Twitter:

- <http://twitter.com/cloudscheduler>



The Sky's the Limit: Preserving Research Data with Clouds

Test Drive Cloud Scheduler

Publicly available pre-configured EC2 AMI ready to go:

```
#create a security group
```

```
$ ec2addgrp cloudscheduler -d "Used for Cloud Scheduler"
```

```
$ ec2auth cloudscheduler -P icmp -t "-1:-1"
```

```
$ ec2auth cloudscheduler -P tcp -p 22
```

```
$ ec2auth cloudscheduler -P tcp -p 40000-40050
```

```
$ ec2auth cloudscheduler -P udp -p 40000-40050
```

```
$ ec2auth cloudscheduler -P tcp -p 9618
```

```
$ ec2auth cloudscheduler -P udp -p 9618
```

```
#boot the cloud scheduler/condor VM
```

```
$ ec2run ami-f9ff1190 -k ec2-keypair -g cloudscheduler
```

```
$ ssh -i ~/.ec2/id_rsa-ec2-keypair \
```

```
root@ec2-75-101-197-134.compute-1.amazonaws.com
```

```
[root@ec2-75-101-197-134 ~]# cat README
```



The Sky's the Limit: Preserving Research Data with Clouds

Summary:

- Babar and Clouds are very good friends
- We are creating a unique system to utilize clouds for research computing
- Wider applicability: Others starting to make use of this solution

