Shooting for the sky: Testing the limits of condor

HTCondor Week 2015

21 May 2015

Edgar Fajardo
On behalf of OSG Software and Technology

Acknowledgement

Although I am the one presenting. This work is a product of a collaborative effort from:

- HTCondor development team
- GlideInWMS development team
- HTCondor UW CHTC (A Moate), UCSD T2 (Terrence M) and CMS FNAL T1 (Burt H), who provided the hardware for the testbed.
- CMS Computing and WLCG sites provided the worker node resources.
- OSG Software team.

Vanilla Condor on a Slide

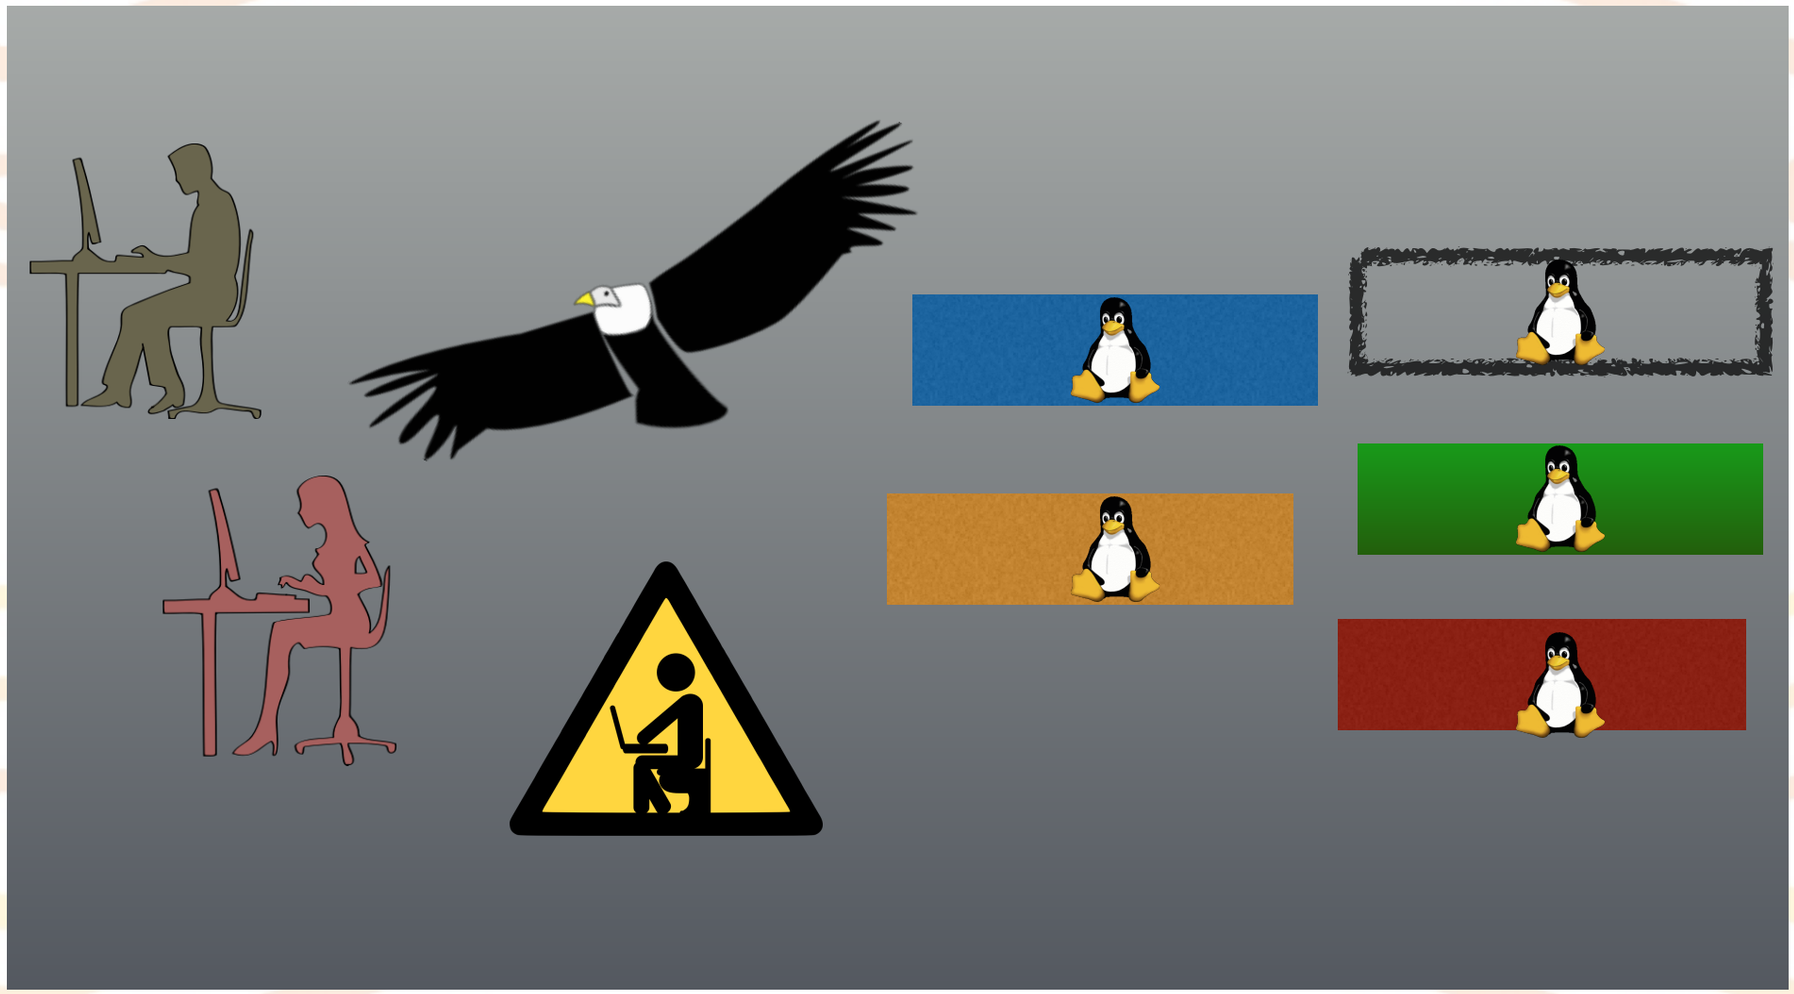

Legend:

Schedd

Central Manger

Worker node

Proudly providing HTC for more than 20 years

GlideInWMS on a slide

Happily working for more than 6 years !!!

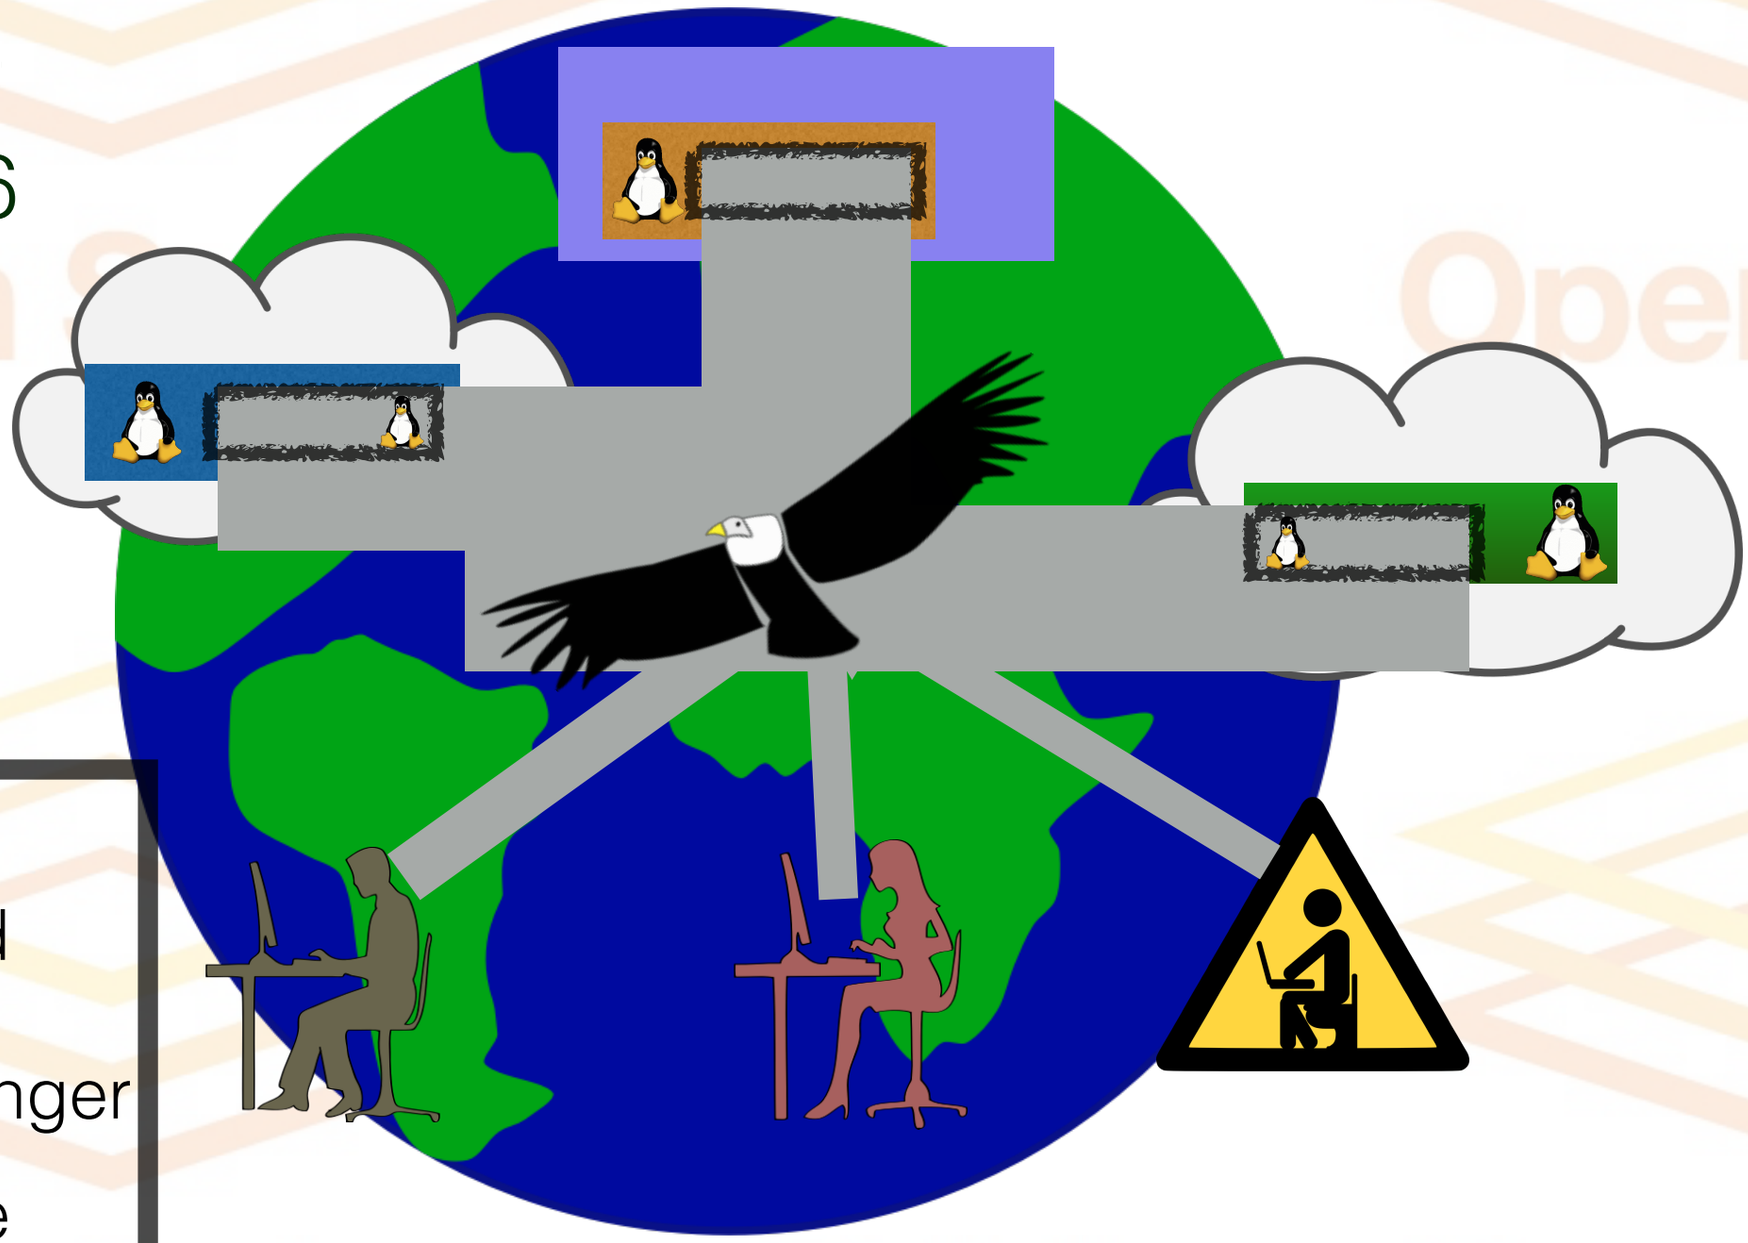

Legend:

Schedd

Central Manger

Worker node

Vanilla HTCondor vs GlideInWMS in the numbers

	Vanilla Condor	GlideinWMS
Number of nodes	O(1k)	O(10k)
# Different types of machines	O(10)	O(100)
Location of Schedds and WN	Private	WAN

The Challenge:

CMS

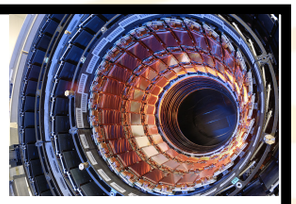

Hi OSG Software folks

For Run II, we would like to have a single global pool of 200,000 running jobs, u think Condor can handle it?

Wow, We would expect so. Nobody has tried that big of a pool

Nobody has found the Higgs before we and ATLAS did

Touche

Btw, u can use some of my WLCG resources to test. About 20k slots

Just 20k? Didn't u say u want a pool of 200k? Nvm we will figure it out.

The Challenge in the numbers

	CMS Requirements for LHC Run II
# Startd	200,000
Autoclusters (types of machines)	600
Number of Schedds	<5
High Availability	YES !!!

Challenge Accepted !!!

How to get 200k slots?

Gathering from a commercial cloud for example*:

$$\frac{\$0.013}{\text{hour} * \text{core}} * \frac{24\ \text{hours}}{\text{day}} * \frac{30\ \text{days}}{\text{month}} * 200\text{k cores} = \frac{\$1,872,000}{\text{month}}$$

Without using your credit card

Our Solution: The UberGlideIn

*At spot pricing and without using the Uber GlideIn

Uber GlideIn

Normal GlideIn

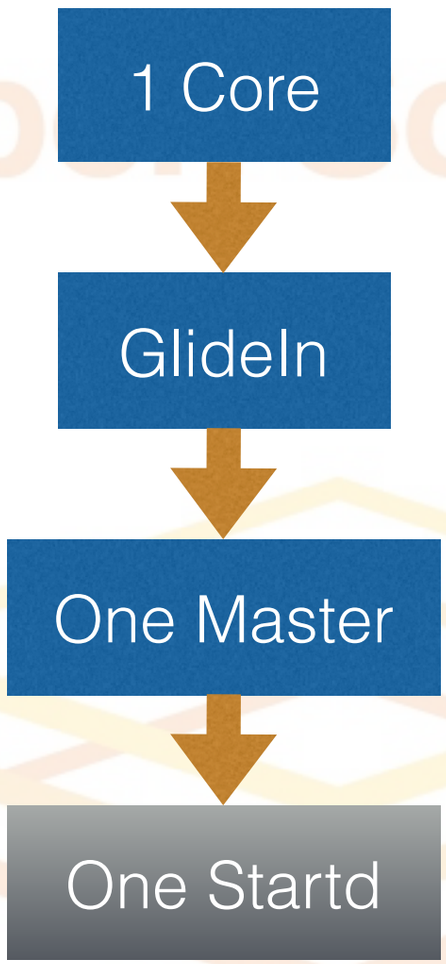

This works because jobs are sleep jobs.

UberGlideIn

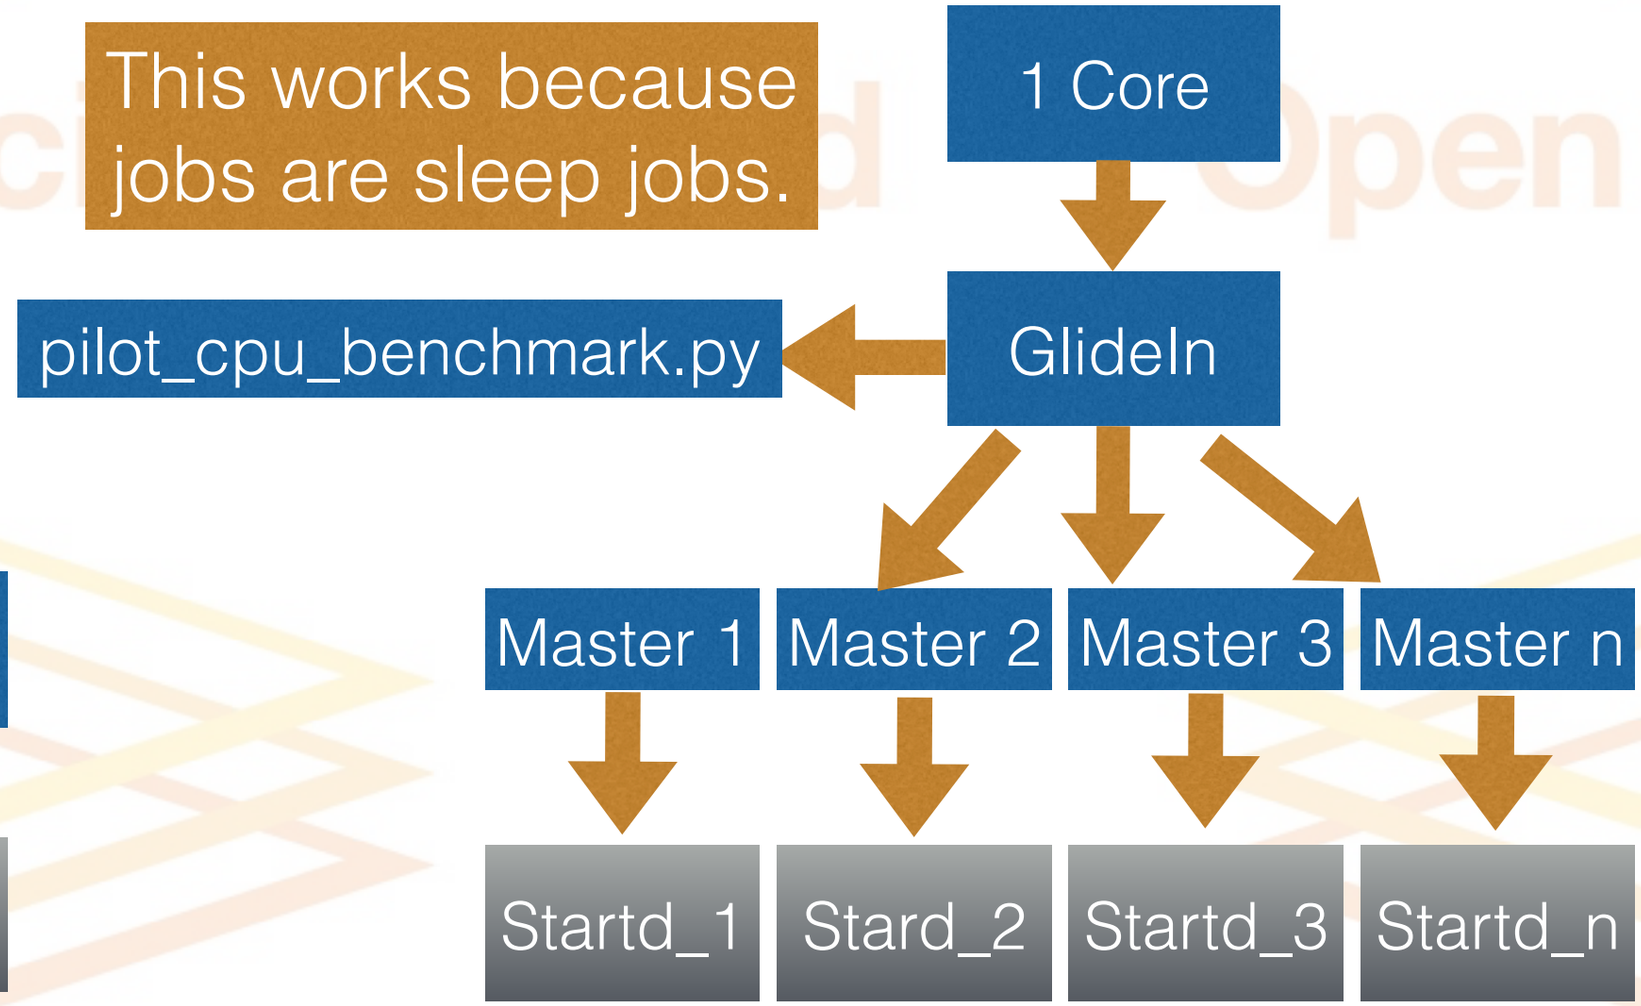

Moreover we wanted them distributed all over

In production, the network latency of having Startd all over the WAN is known to cause problems. See Bradley et all

Does it work?

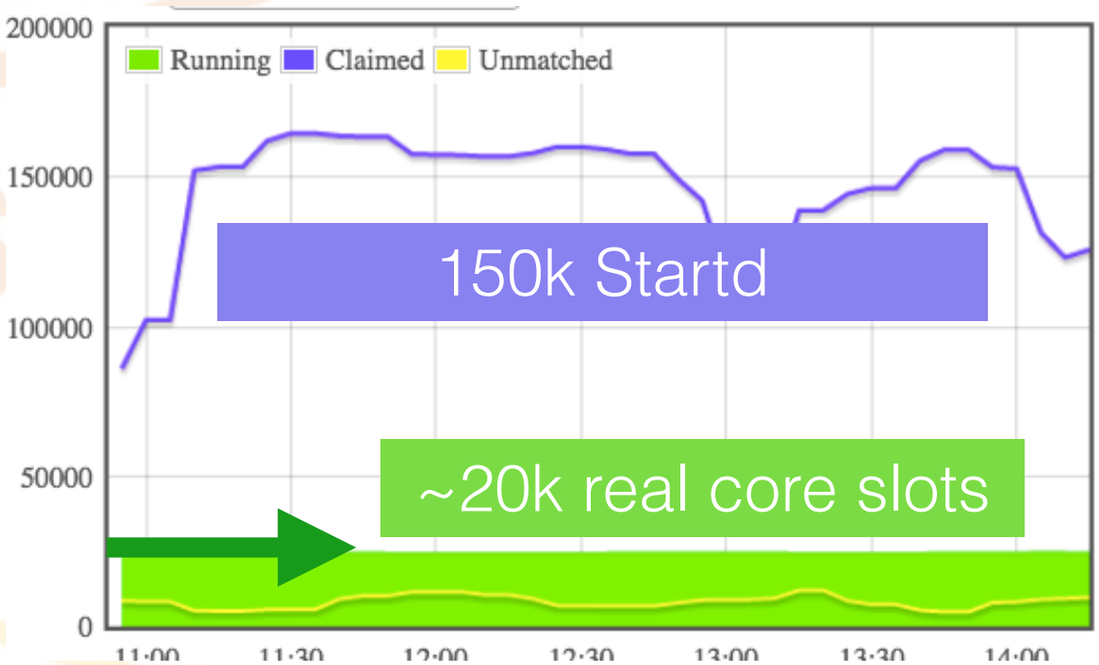

Yes it does !!!

Well most of the time it works:

PIC Admin

Hi OSG Software folks

About those tests you are running at our site. Just FYI u brought down our firewall, while using > 10k ports

Ohh dude, sorry about that. We will fix it

Mental note: Talk to the HTCondor dev team to reduce the long lived TCP connections from the WN to the outside (Collector, Negotiator …)

Now we have the resources, lets test:

For ~ 3 months

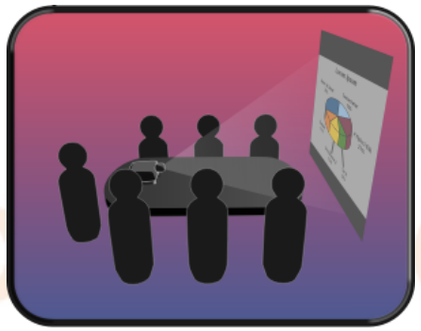

Report/
Brainstorm

Fix

Test

Ganglia features was key. THAnKS !!!

```
GANGLIAD_REQUIREMENTS = MyType =!= "Machine"
```

Did it work? YES !!!

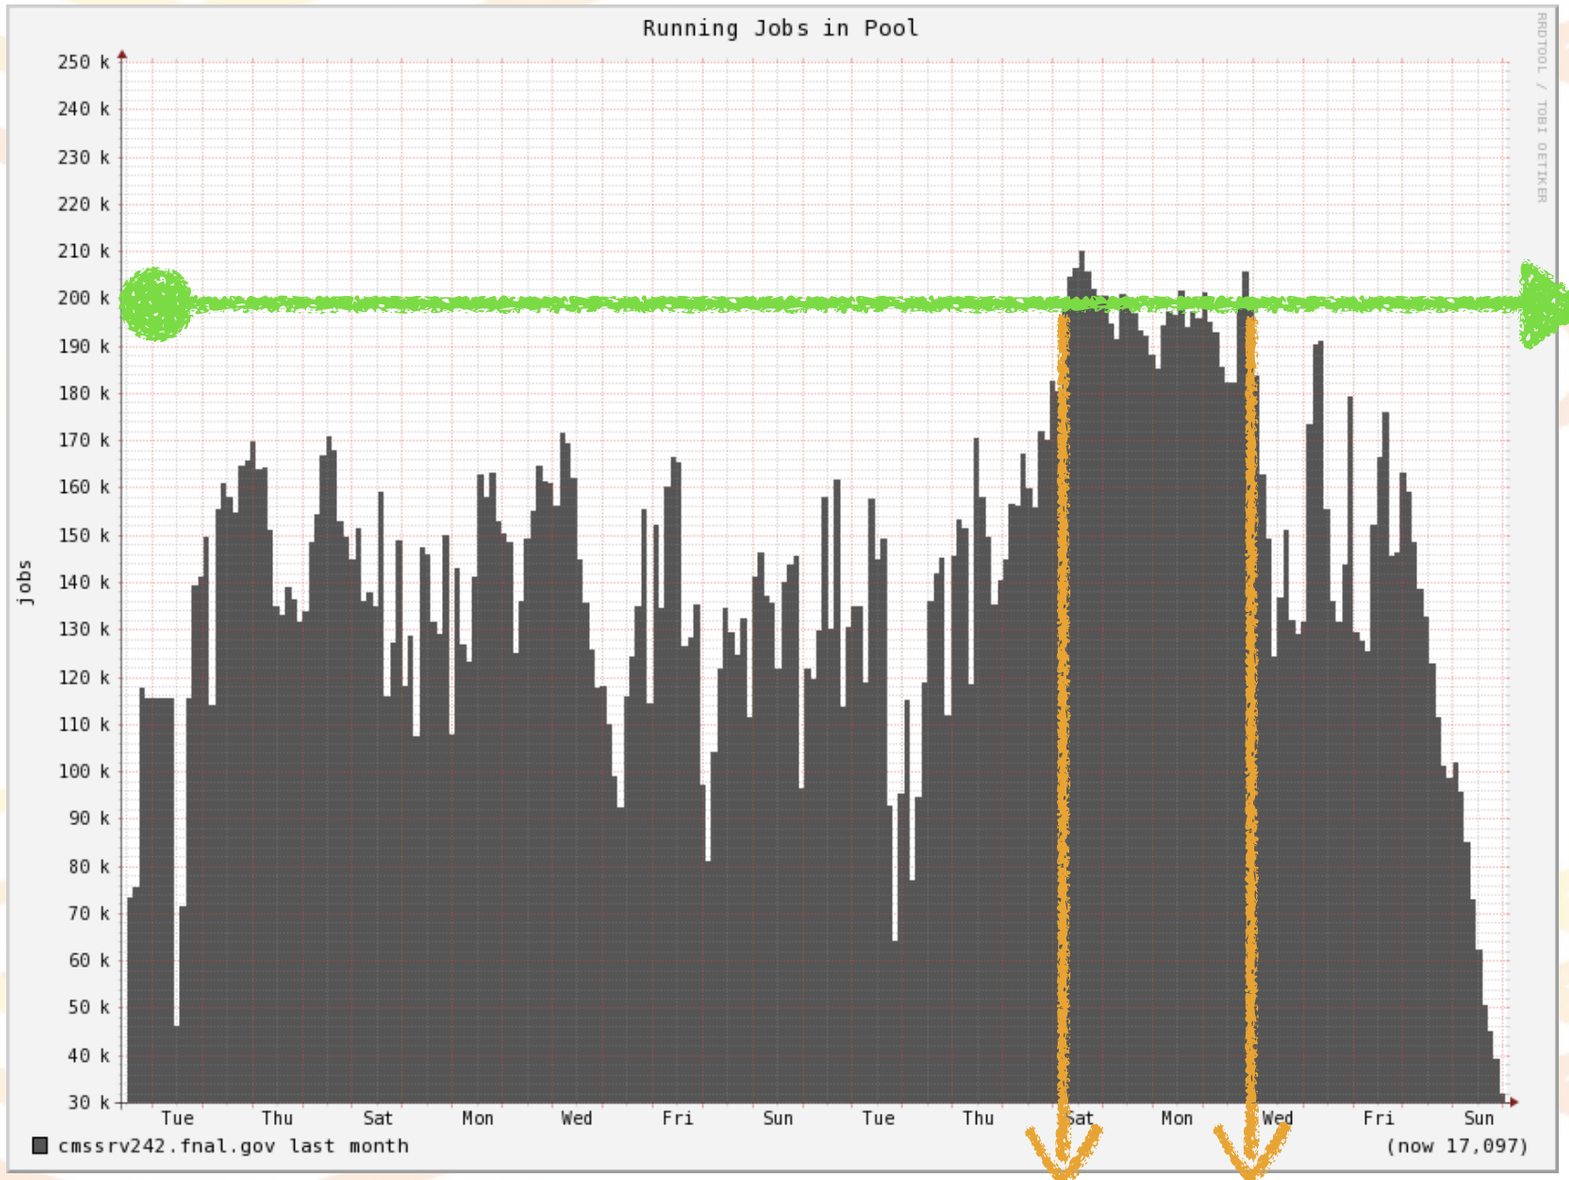

One full weekend

“It was not all a bed of roses”, maybe it was…

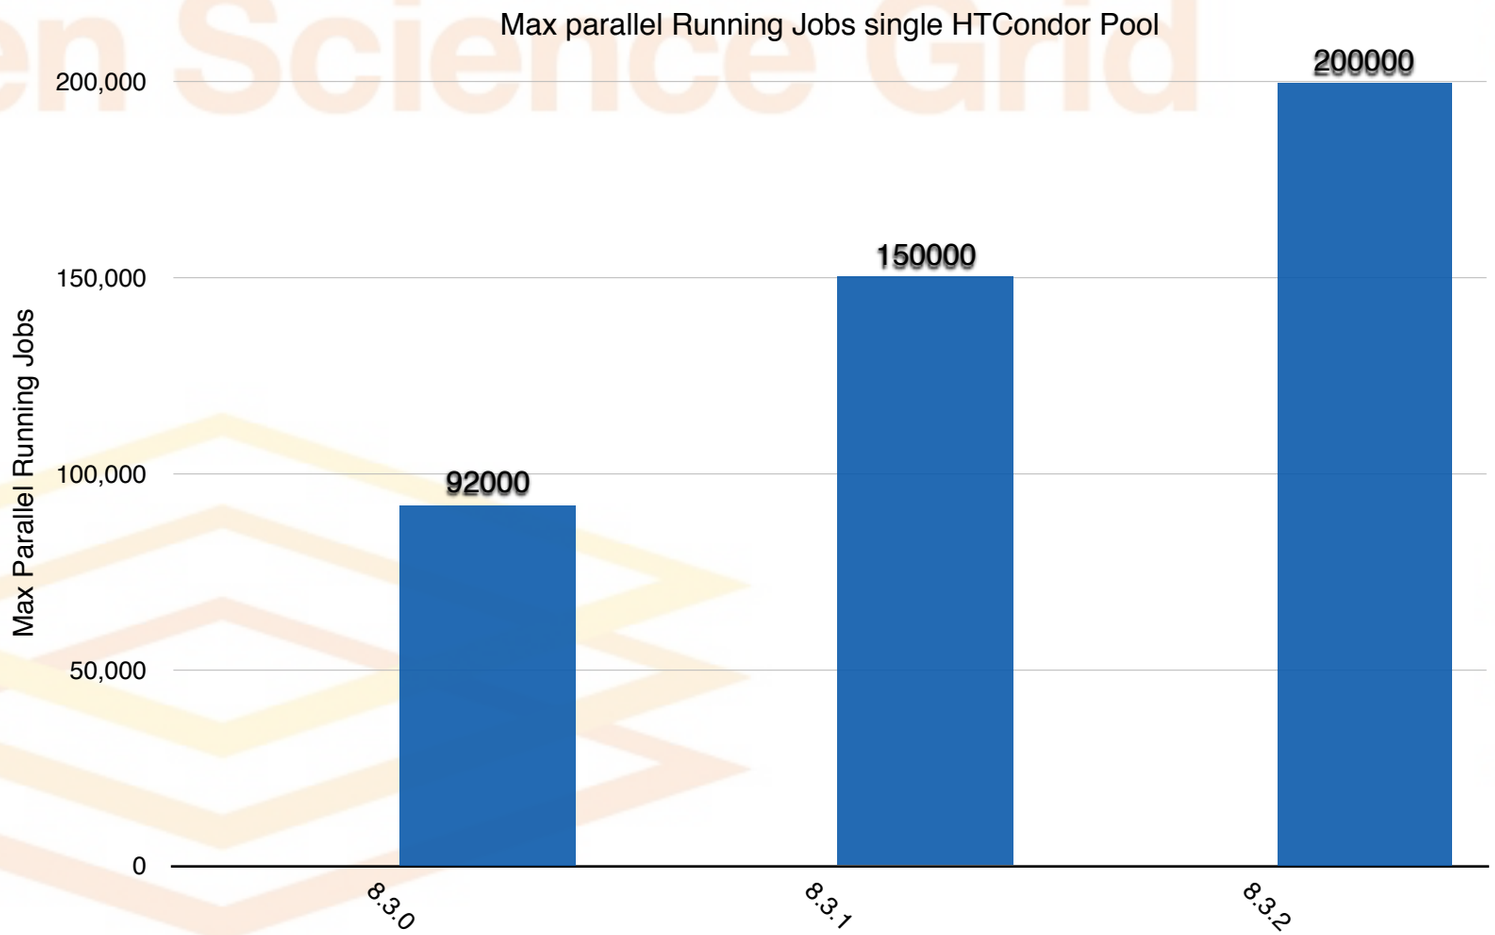

HTCondor Improvements

For more details see Todd's talk on What's new in HTCondor 2015?

- Non blocking GSI authentication at the Collector
- Shared port at the worker node. In general reduce # of long lived TCP connections.
- Removed file locking at the Schedd
- Reduced incoming TCP connections at the Schedd
- Batched resources request from the Collector

Scale Improvements throughout "History"

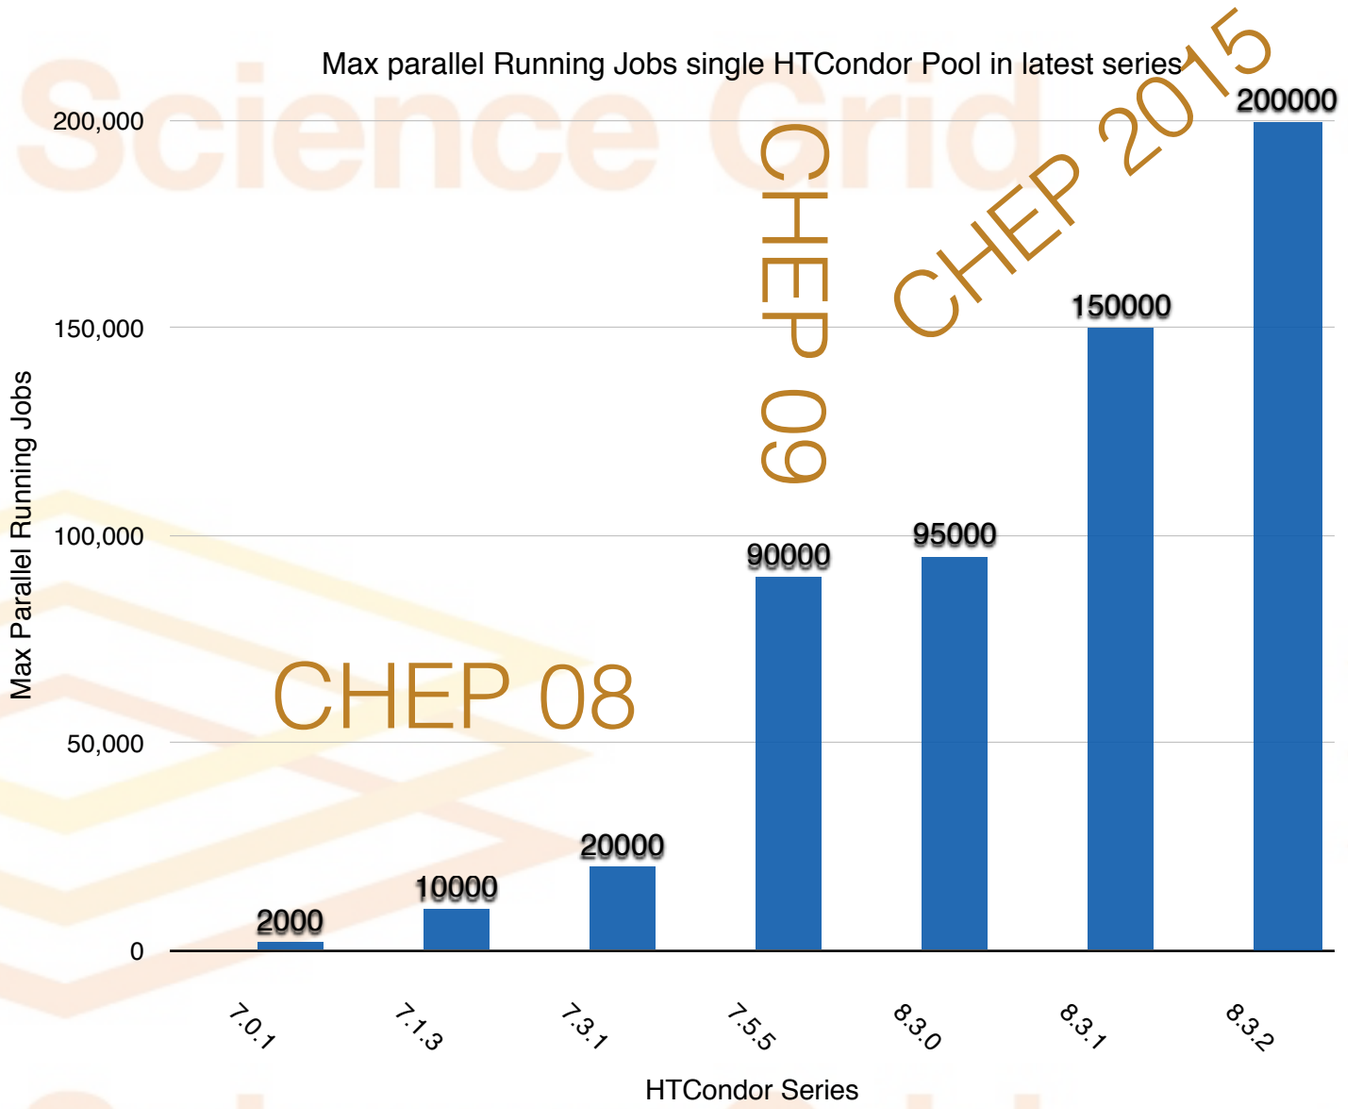

Ahh, One more thing …

Brian B

Hi Edgar,

Since u and Jeff are working on the scaling tests, what about we scale test our new rockstar: The HTCondor CE.

Sounds good, Which levels are u looking for?

About the size of UNL sleeper pool ~16k parallel running jobs?

Are u kidding me? That is twice as much of what any OSG site needs?

Step back: What is a CE?

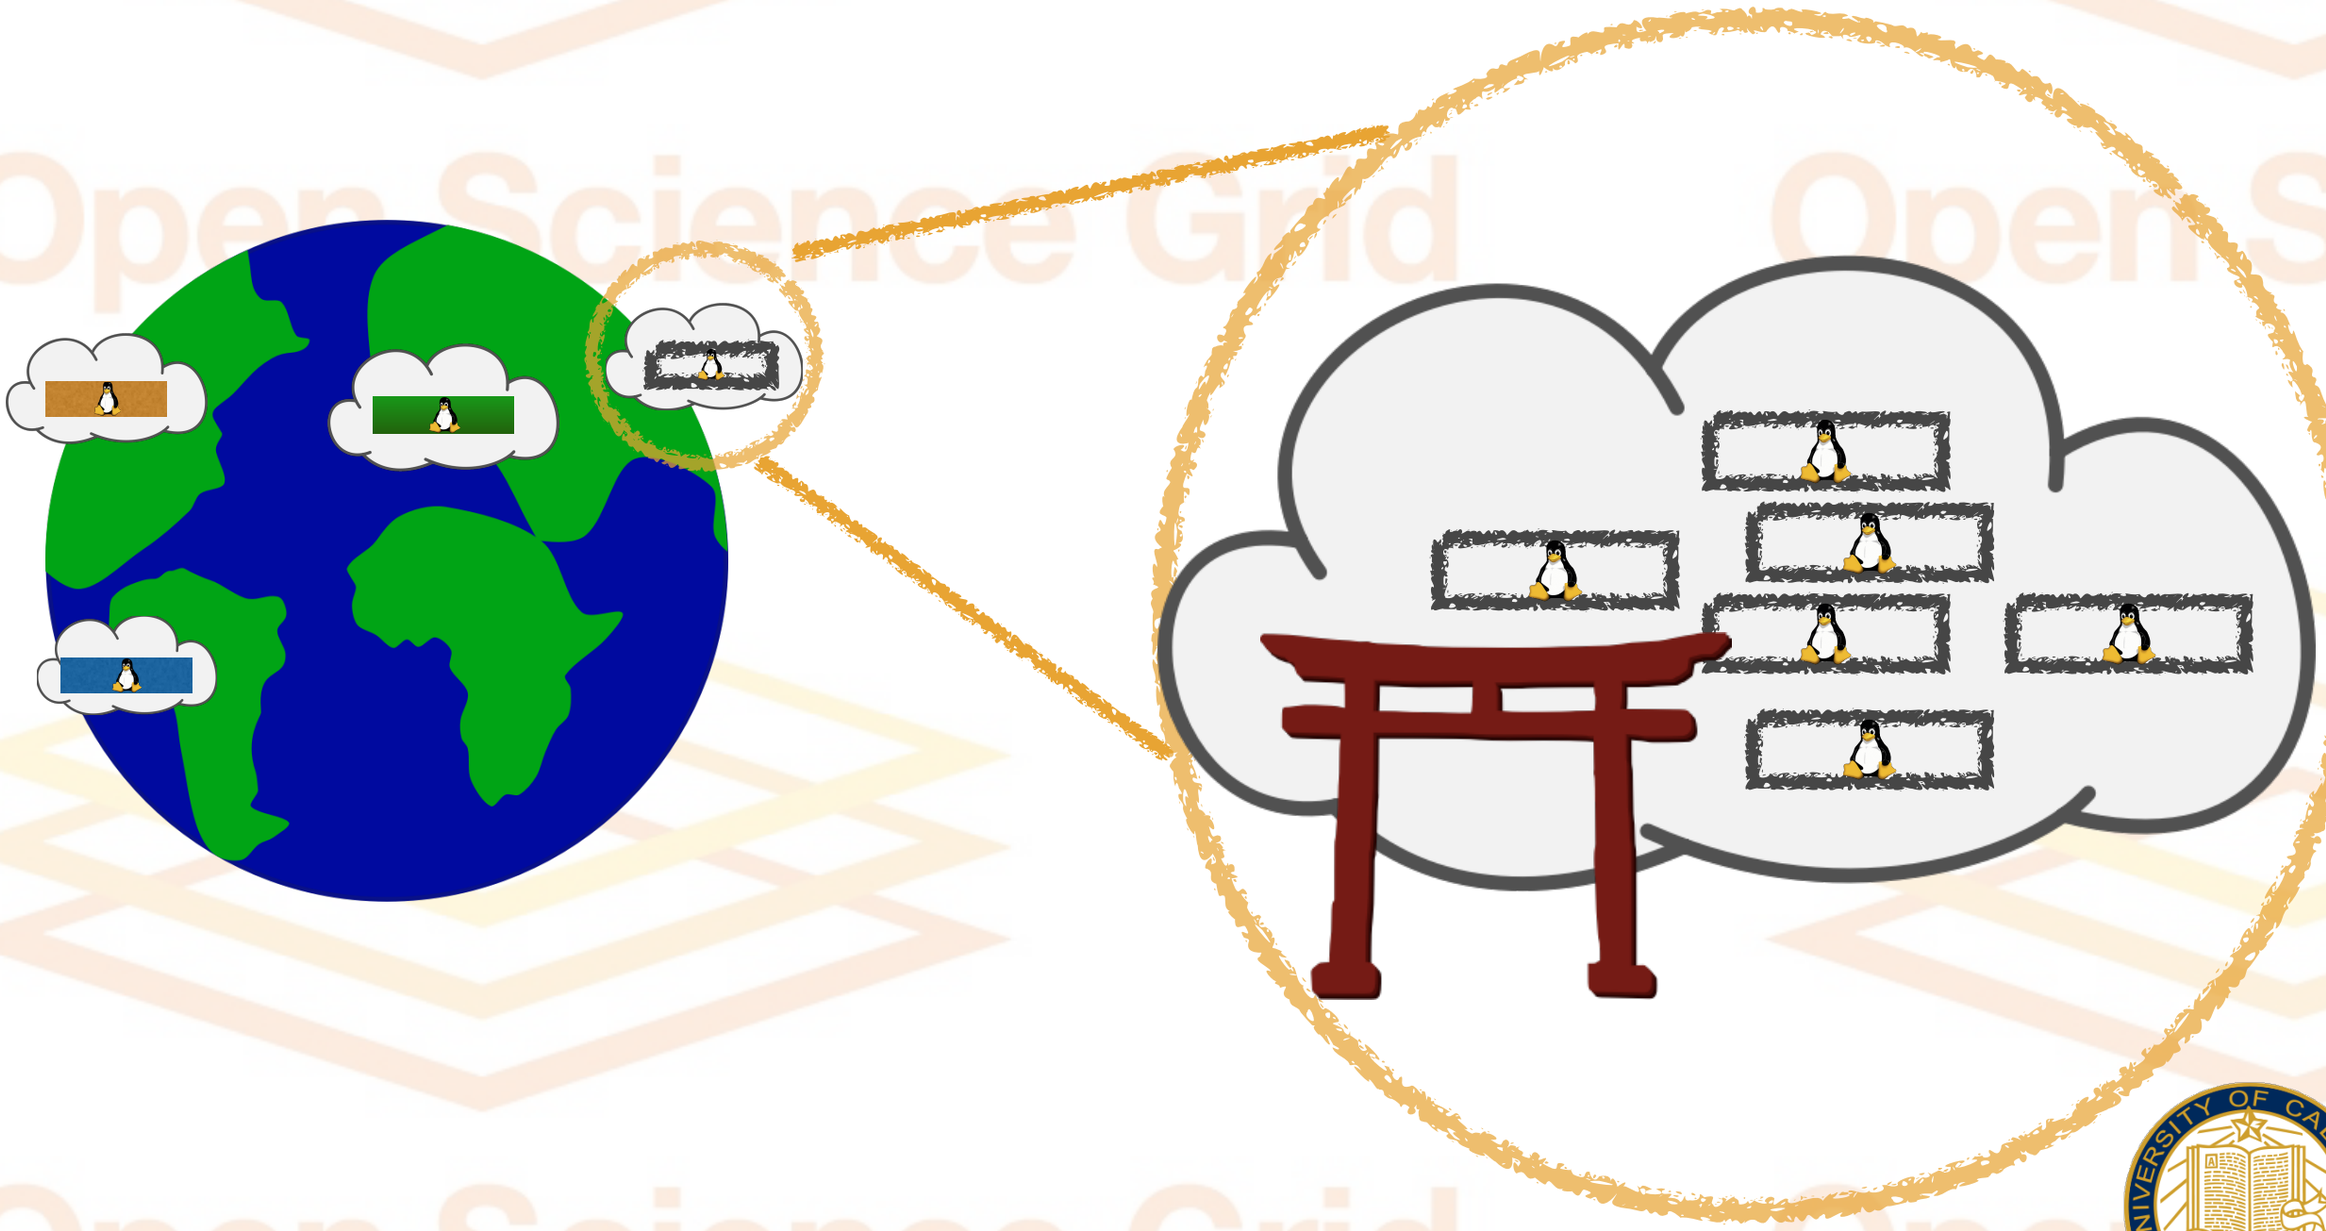

HTCondor-CE in a slide

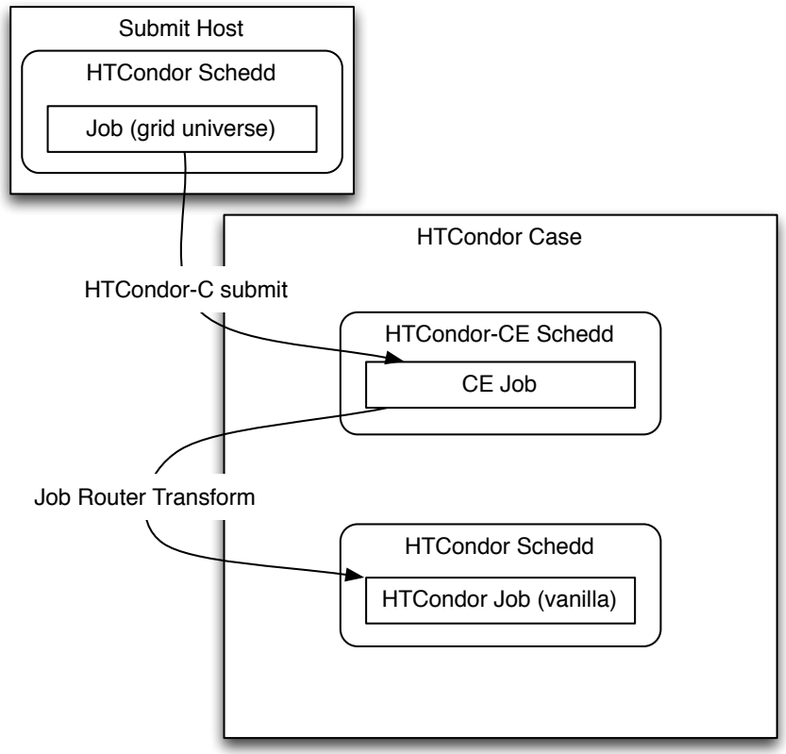

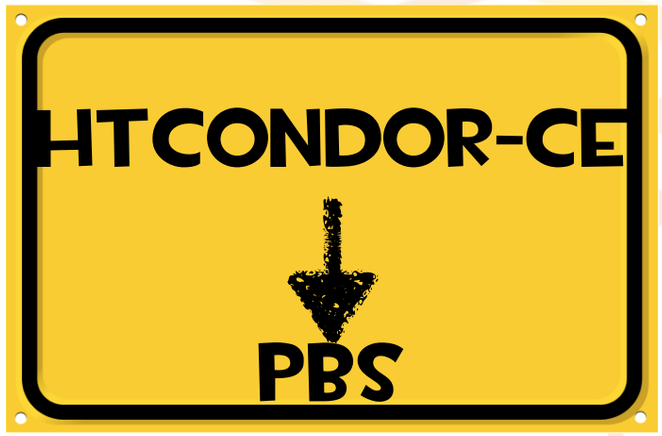

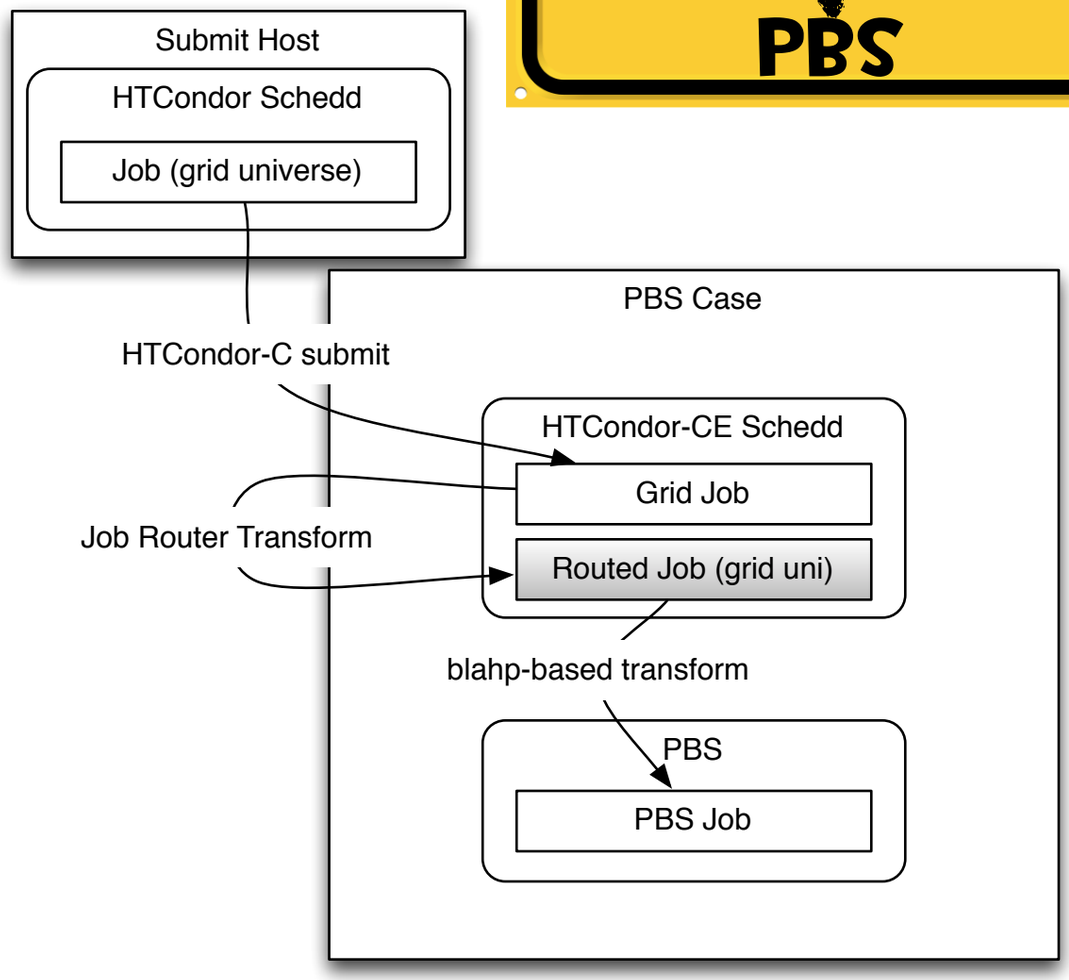

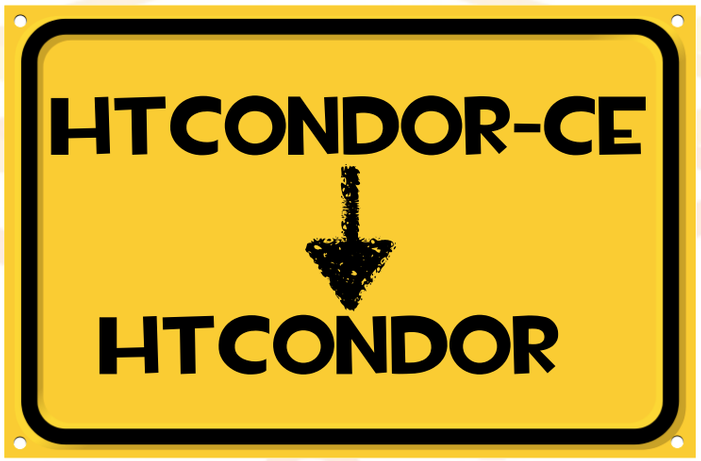

What is HTCondor-CE?

Condor + Configuration

For more details see OSG AHM 2015 or CHEP 2015 talks.

Did we make it?
YES !!!

After some development cycles:

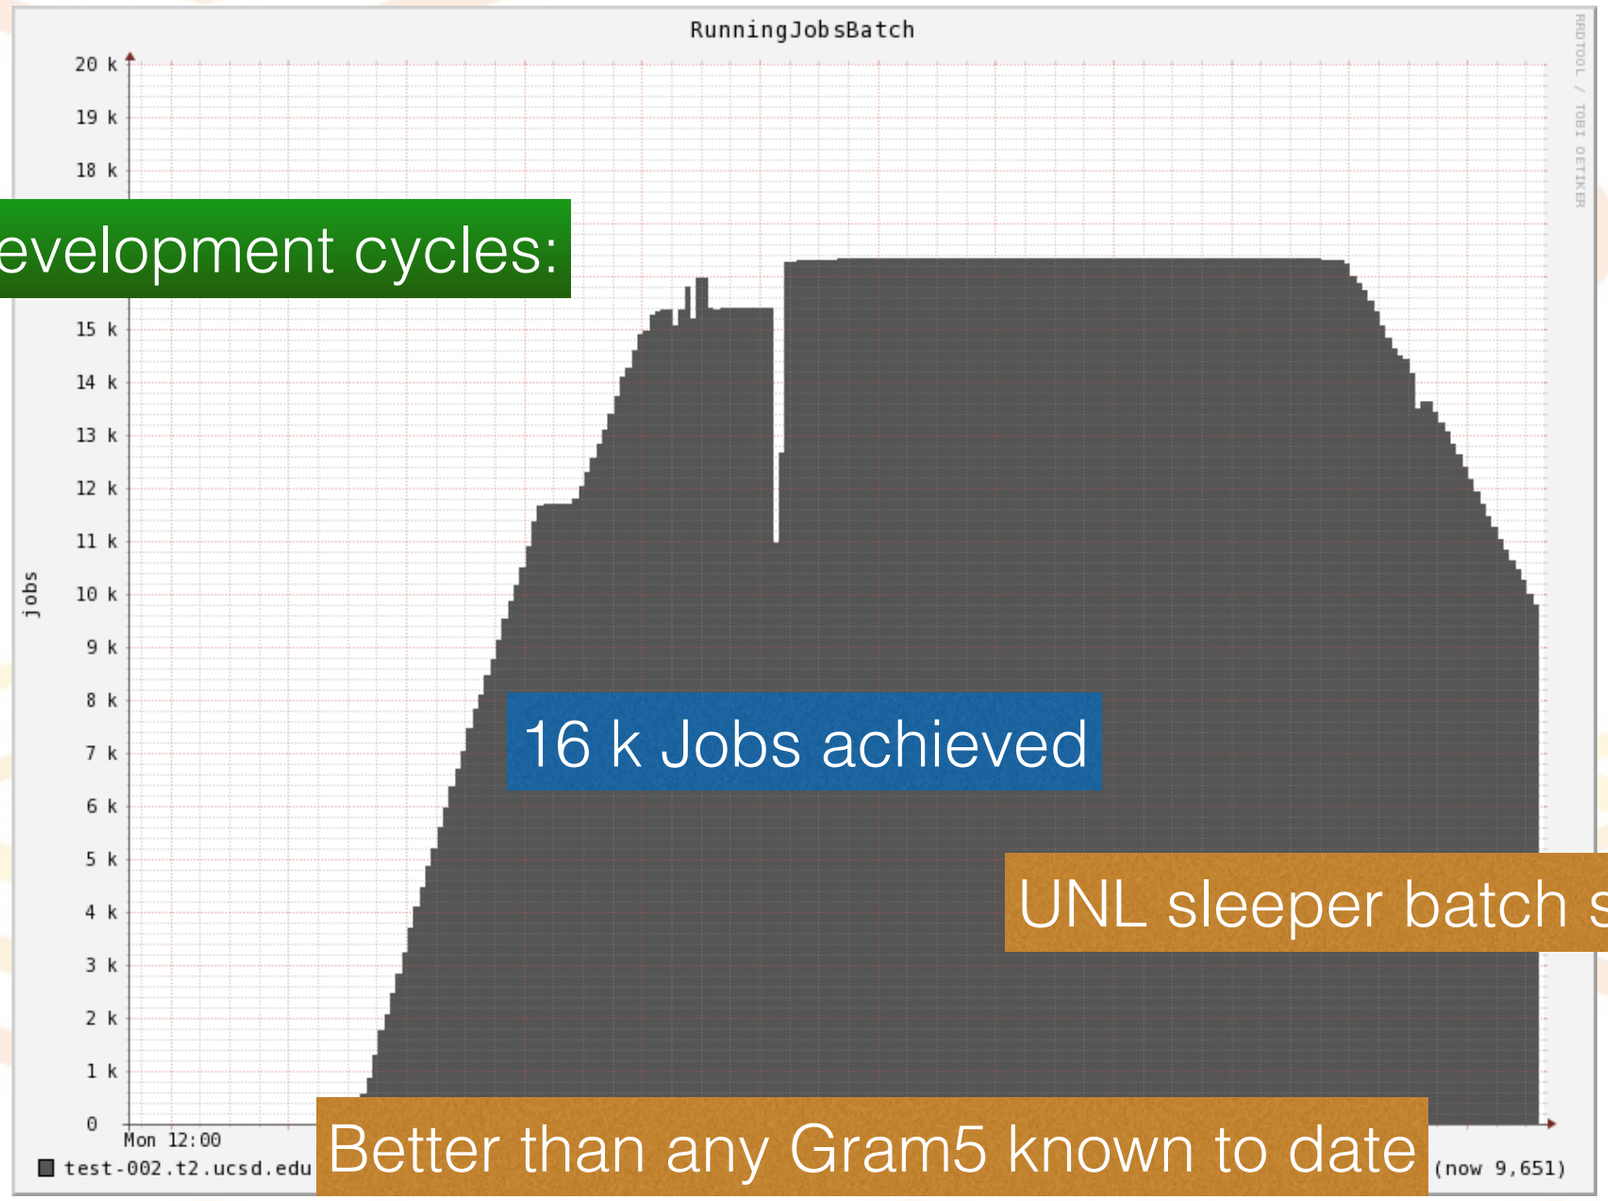

HTCondor CE In the Numbers:

	HTCondor-CE	GRAM 5
Best max running jobs	16k*	10k
Network Port usage (per running job	2	4
StartUp Rate	70 jobs/min	55 jobs/min*

*Disclaimer: This tests were done on different hardware with 5 years in between them.

Conclusions

- "The OSG Software team, in conjunction with HTCondor and GlideinWMS development teams have collaborated to push the scalability limits of a single HTCondor pool "
- The HTCondor-CE is ready to rock and roll

Questions?

Contact us at:

1-900-scale-masters

Just Kidding

Contact us:

osg-software@opensciencegrid.org