

# BRILL

Sarah Baker  
Krzysztof Ciba  
Michael Nash  
Peter Sherwood  
Brinick Simmons

SIT Meeting  
ATLAS S/W week  
1st December 2010

# Contents

```
>> find . -name "new"
```

```
>> du {afs,castor,mp3}
```

```
>> cat "state of the nation"
```

```
>> import __future__
```

```
>> find . -name "new"
```

# What's new?

>> `find . -name "new"`

- \* Full Batch node running
- \* Better system monitoring
- \* Refactoring 'n' bugfixes galore  
(18 production tags in 4 months)

# Full Batch Node...

```
>> find . -name "new"
```

1. Jobs run on batch local disk

+

2. Direct file dispatch (Castor/ChainStore/AFS)

=

Mega-AFS space savings (discussed later)

# Better monitoring

>> find . -name "new"

## 1. Disk Space:

AFS and Castor

## 2. Batch nodes:

Machine Status

Number/state/age of jobs

## 3. Runs:

Expectations:

Today's runs going?

Yesterday's runs shut down cleanly?

All jobs in allowed state?

Warning/urgent exceptions

Number of failing/succeeding jobs, invalid packages/jobs

....

# Better monitoring

>> `find . -name "new"`

Run regularly on cron

Produces a series of webpages

Used by RTT shifters

Find problems fast

Shifters can deal with some problems

Dispatch to experts otherwise

```
>> du {afs,castor,mp3}
```



# AFS Usage

>> `du {afs,castor,mp3}`

Until recently

- \* Many runs required > 40GB scratch space
- \* Some runs even needed 100GB
- \* Daily requested ~1.8TB scratch

Now

- \* All runs get 15GB
- \* 15GB = 2-3 times safety margin
- \* Daily requested ~0.5TB scratch

# Castor Usage

>> `du {afs,castor,mp3}`

- \* Files > 20MB go to `atlascerngroupdisk pool`
- \* Until recently we had “Planet Earth syndrome”  
(resources are infinite)
- \* Turns out they are not

# Castor Usage

>> du {afs,castor,mp3}

Until recently

- \* Quota: 10TB, RTT usage: 19TB

- \* ZOMG! Delete, delete, delete (old stuff)

Now

- \* Quota: 15TB, RTT usage: 12TB

- \* Should be able <10TB

# Space questions

```
>> du {afs,castor,mp3}
```

In the past RTT has filled both AFS and Castor....

In the future....

....what if 1 week Castor RTT data too much?

- \* Ship more stuff to AFS?

- \* Keep < 1 week Castor data?

....what if AFS full, and Castor full?

- \* Well, let's not think about that...


**>> cat "state of nation"**

# The Facts - 1 week

>> cat "state of nation"

	RTT	TCT	FCT
Runs	112	77	42
Jobs	22400	2030	1470
AFS (GB)	522	229	78
Castor (GB)	1469	1567	72

Will vary with  
job success/  
failure



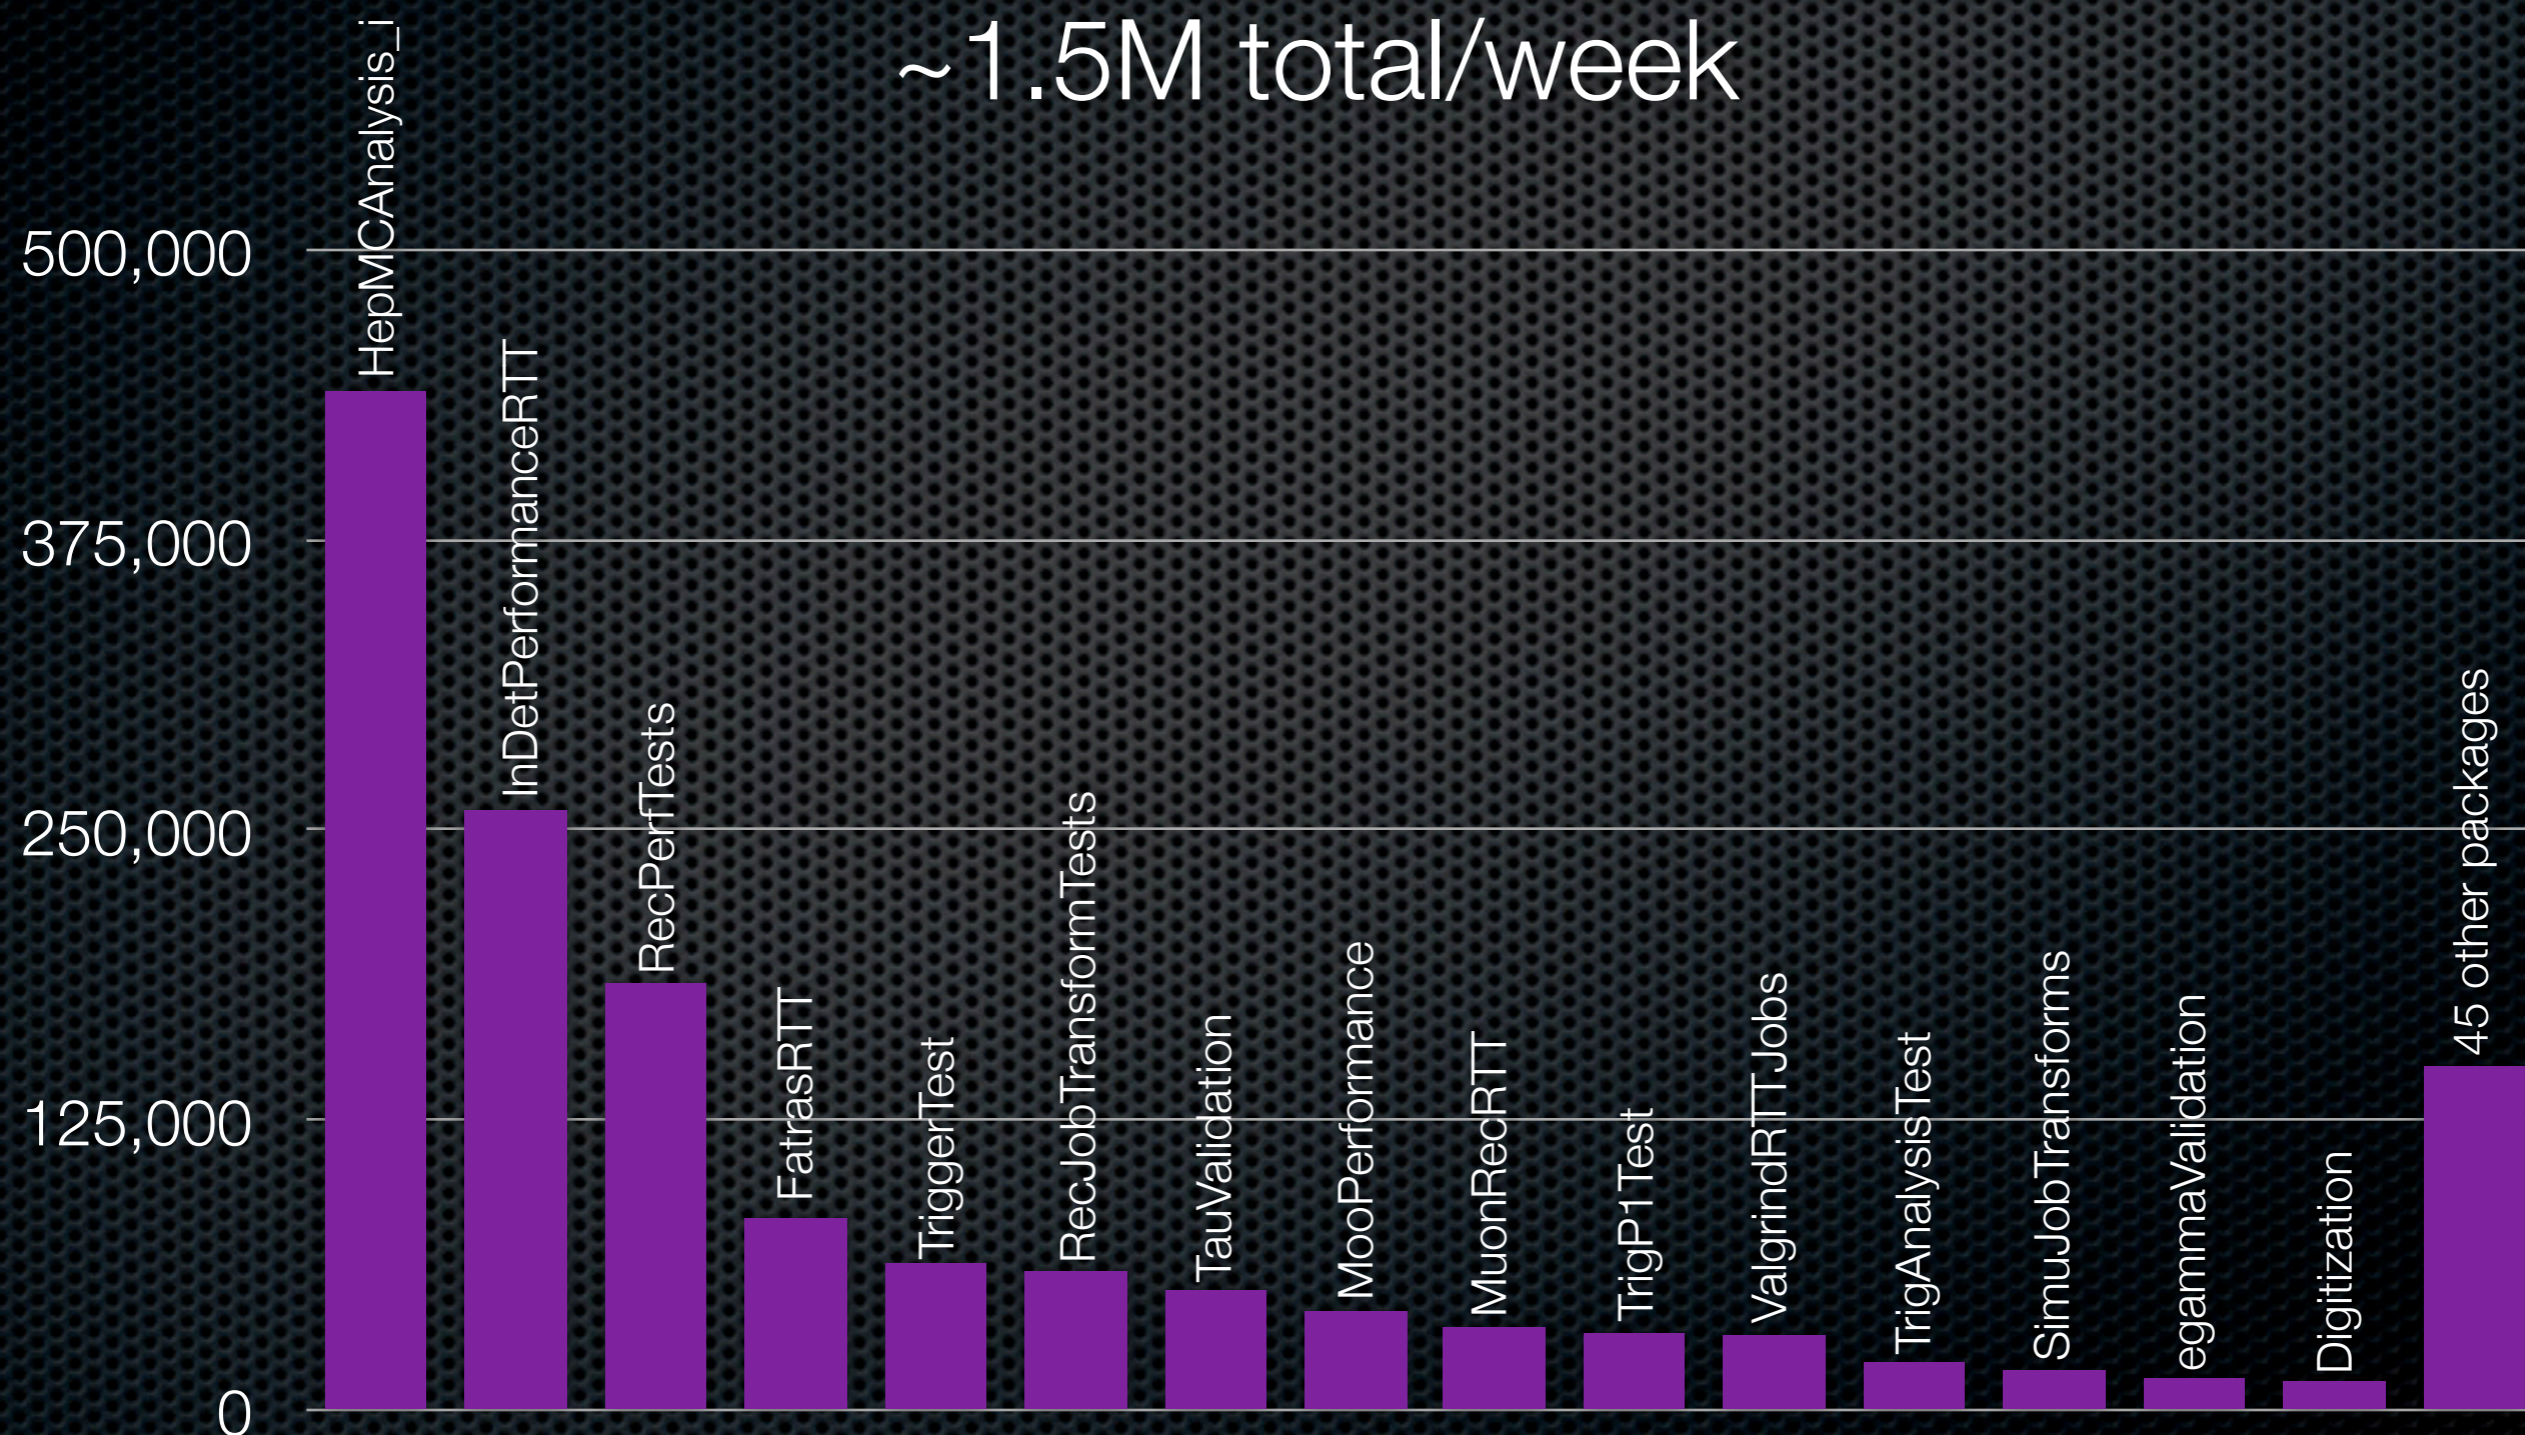
Total AFS use : 829GB  
Total CASTOR use : 3108GB

# RTT - 1 week

>> cat "state of nation"

Number of keep-files on AFS

~ 1.5M total/week

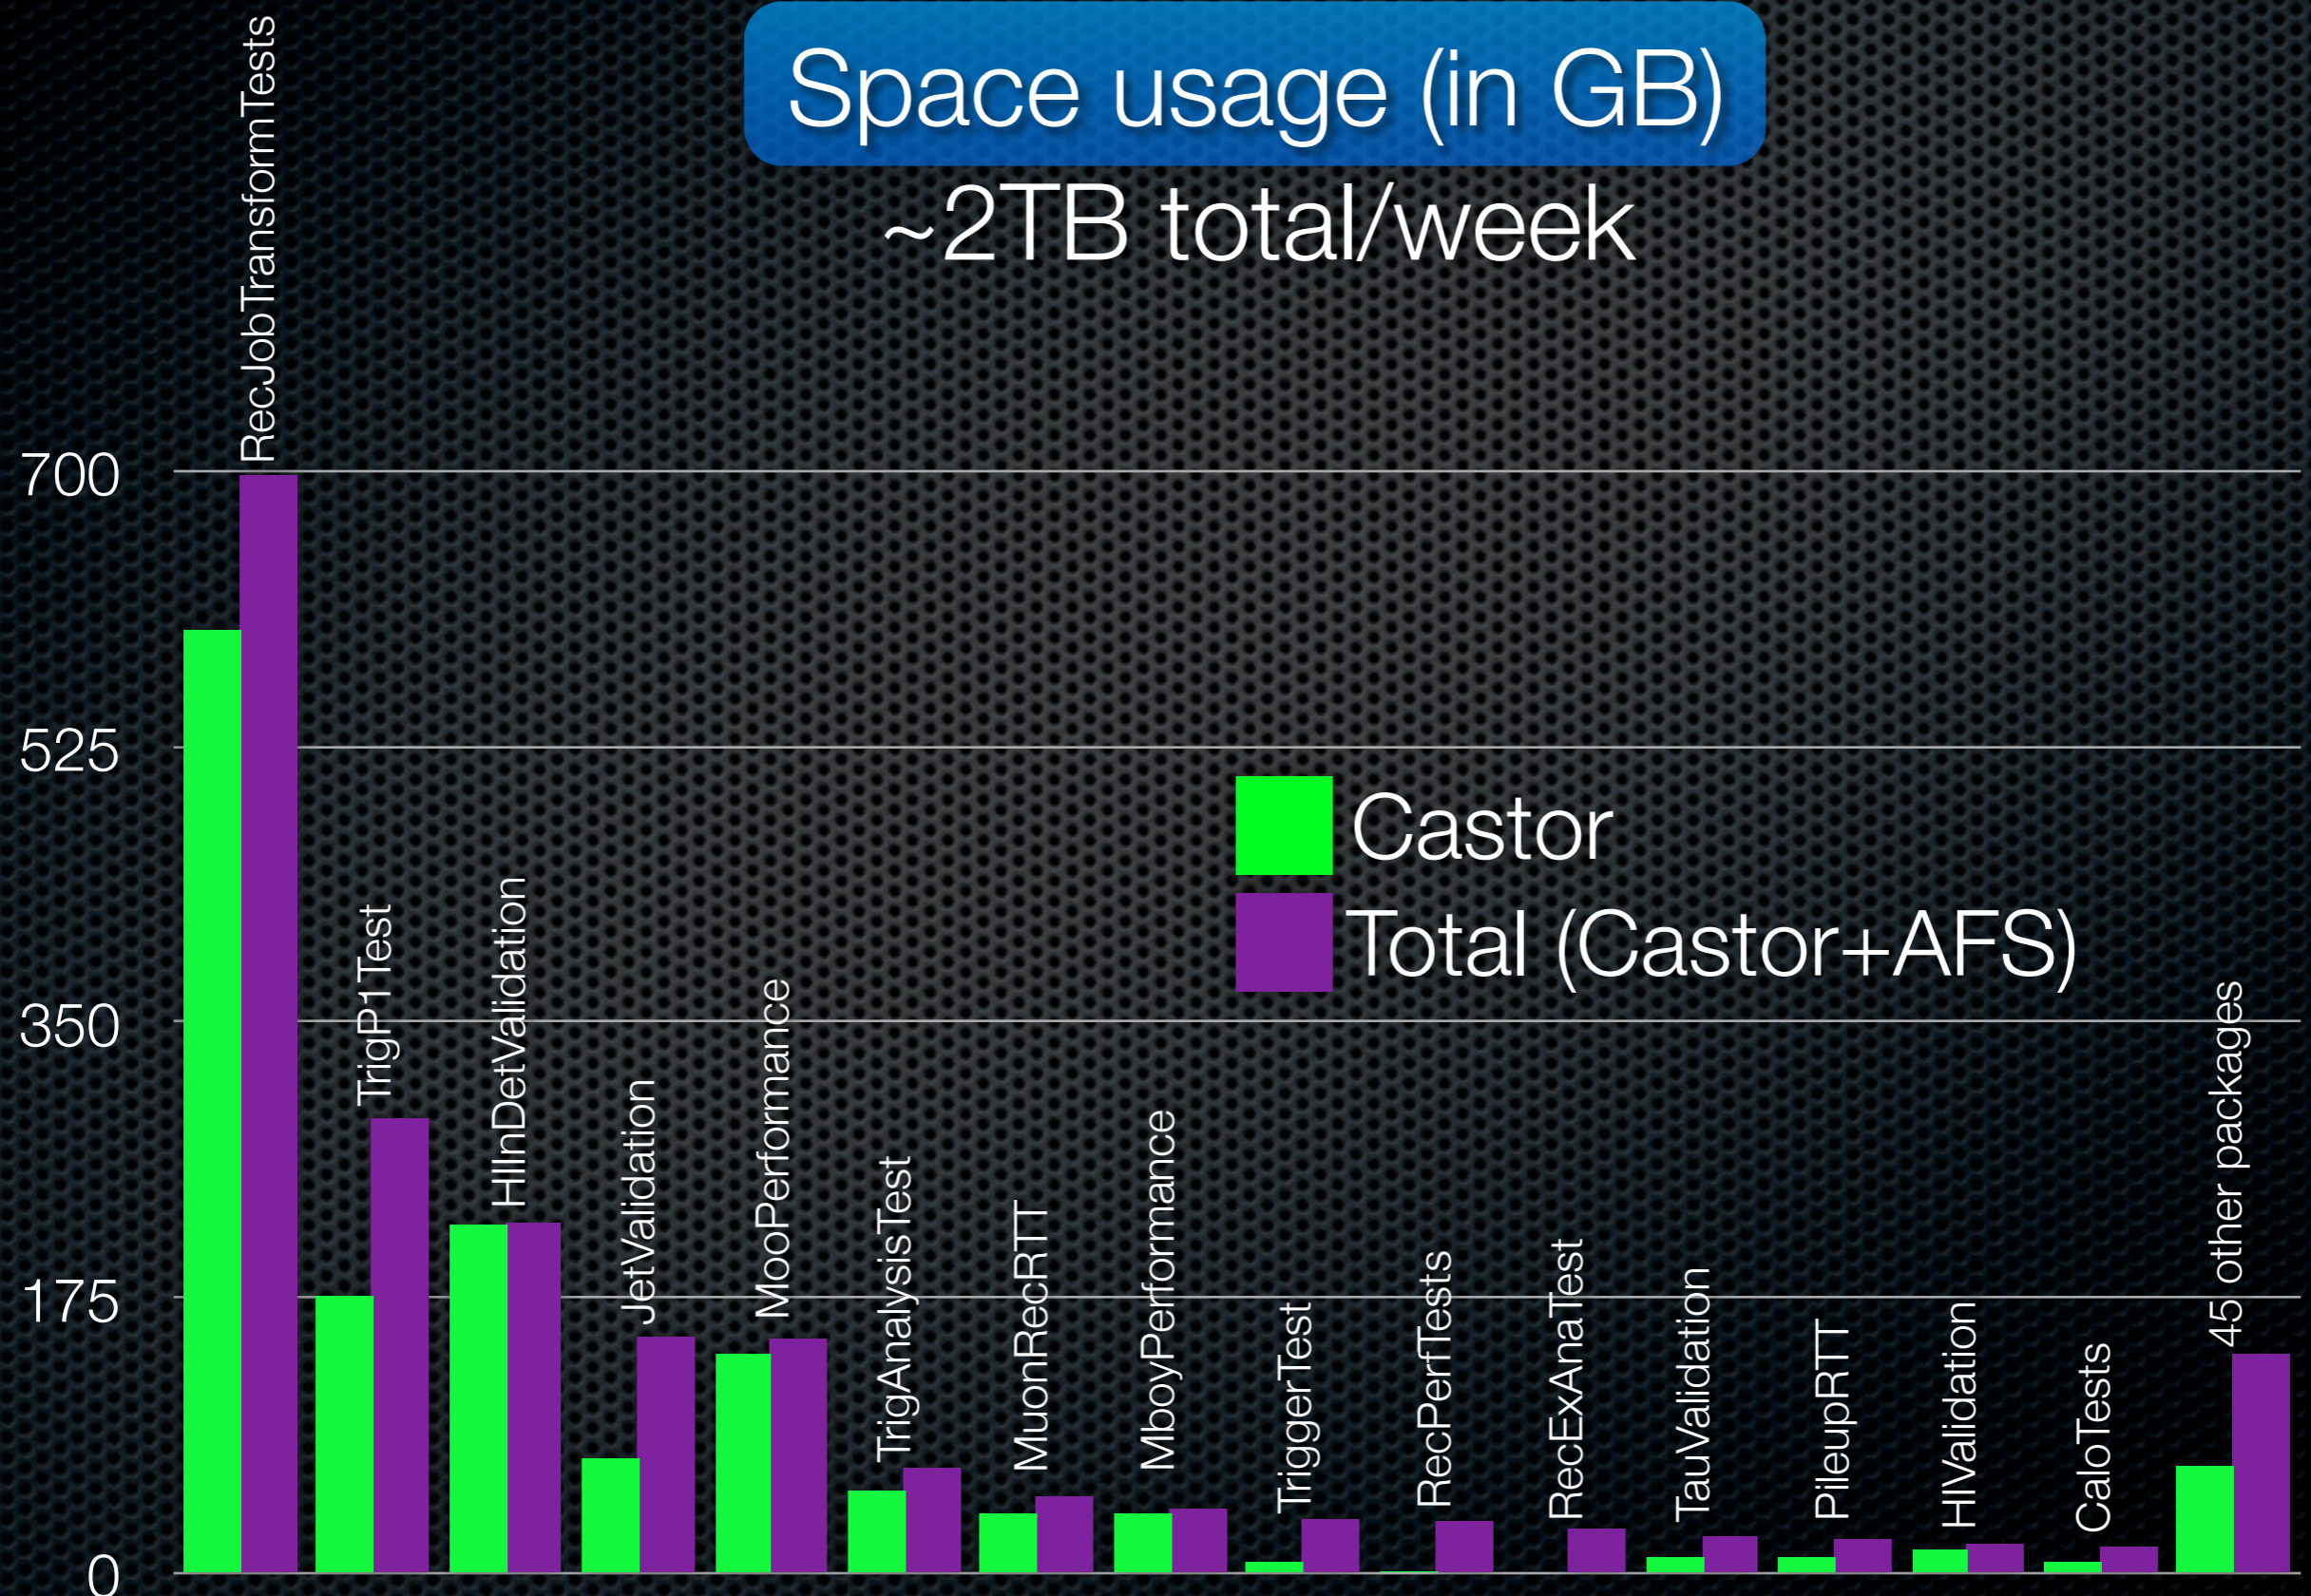


# RTT - 1 week

>> cat "state of nation"

Space usage (in GB)

~2TB total/week



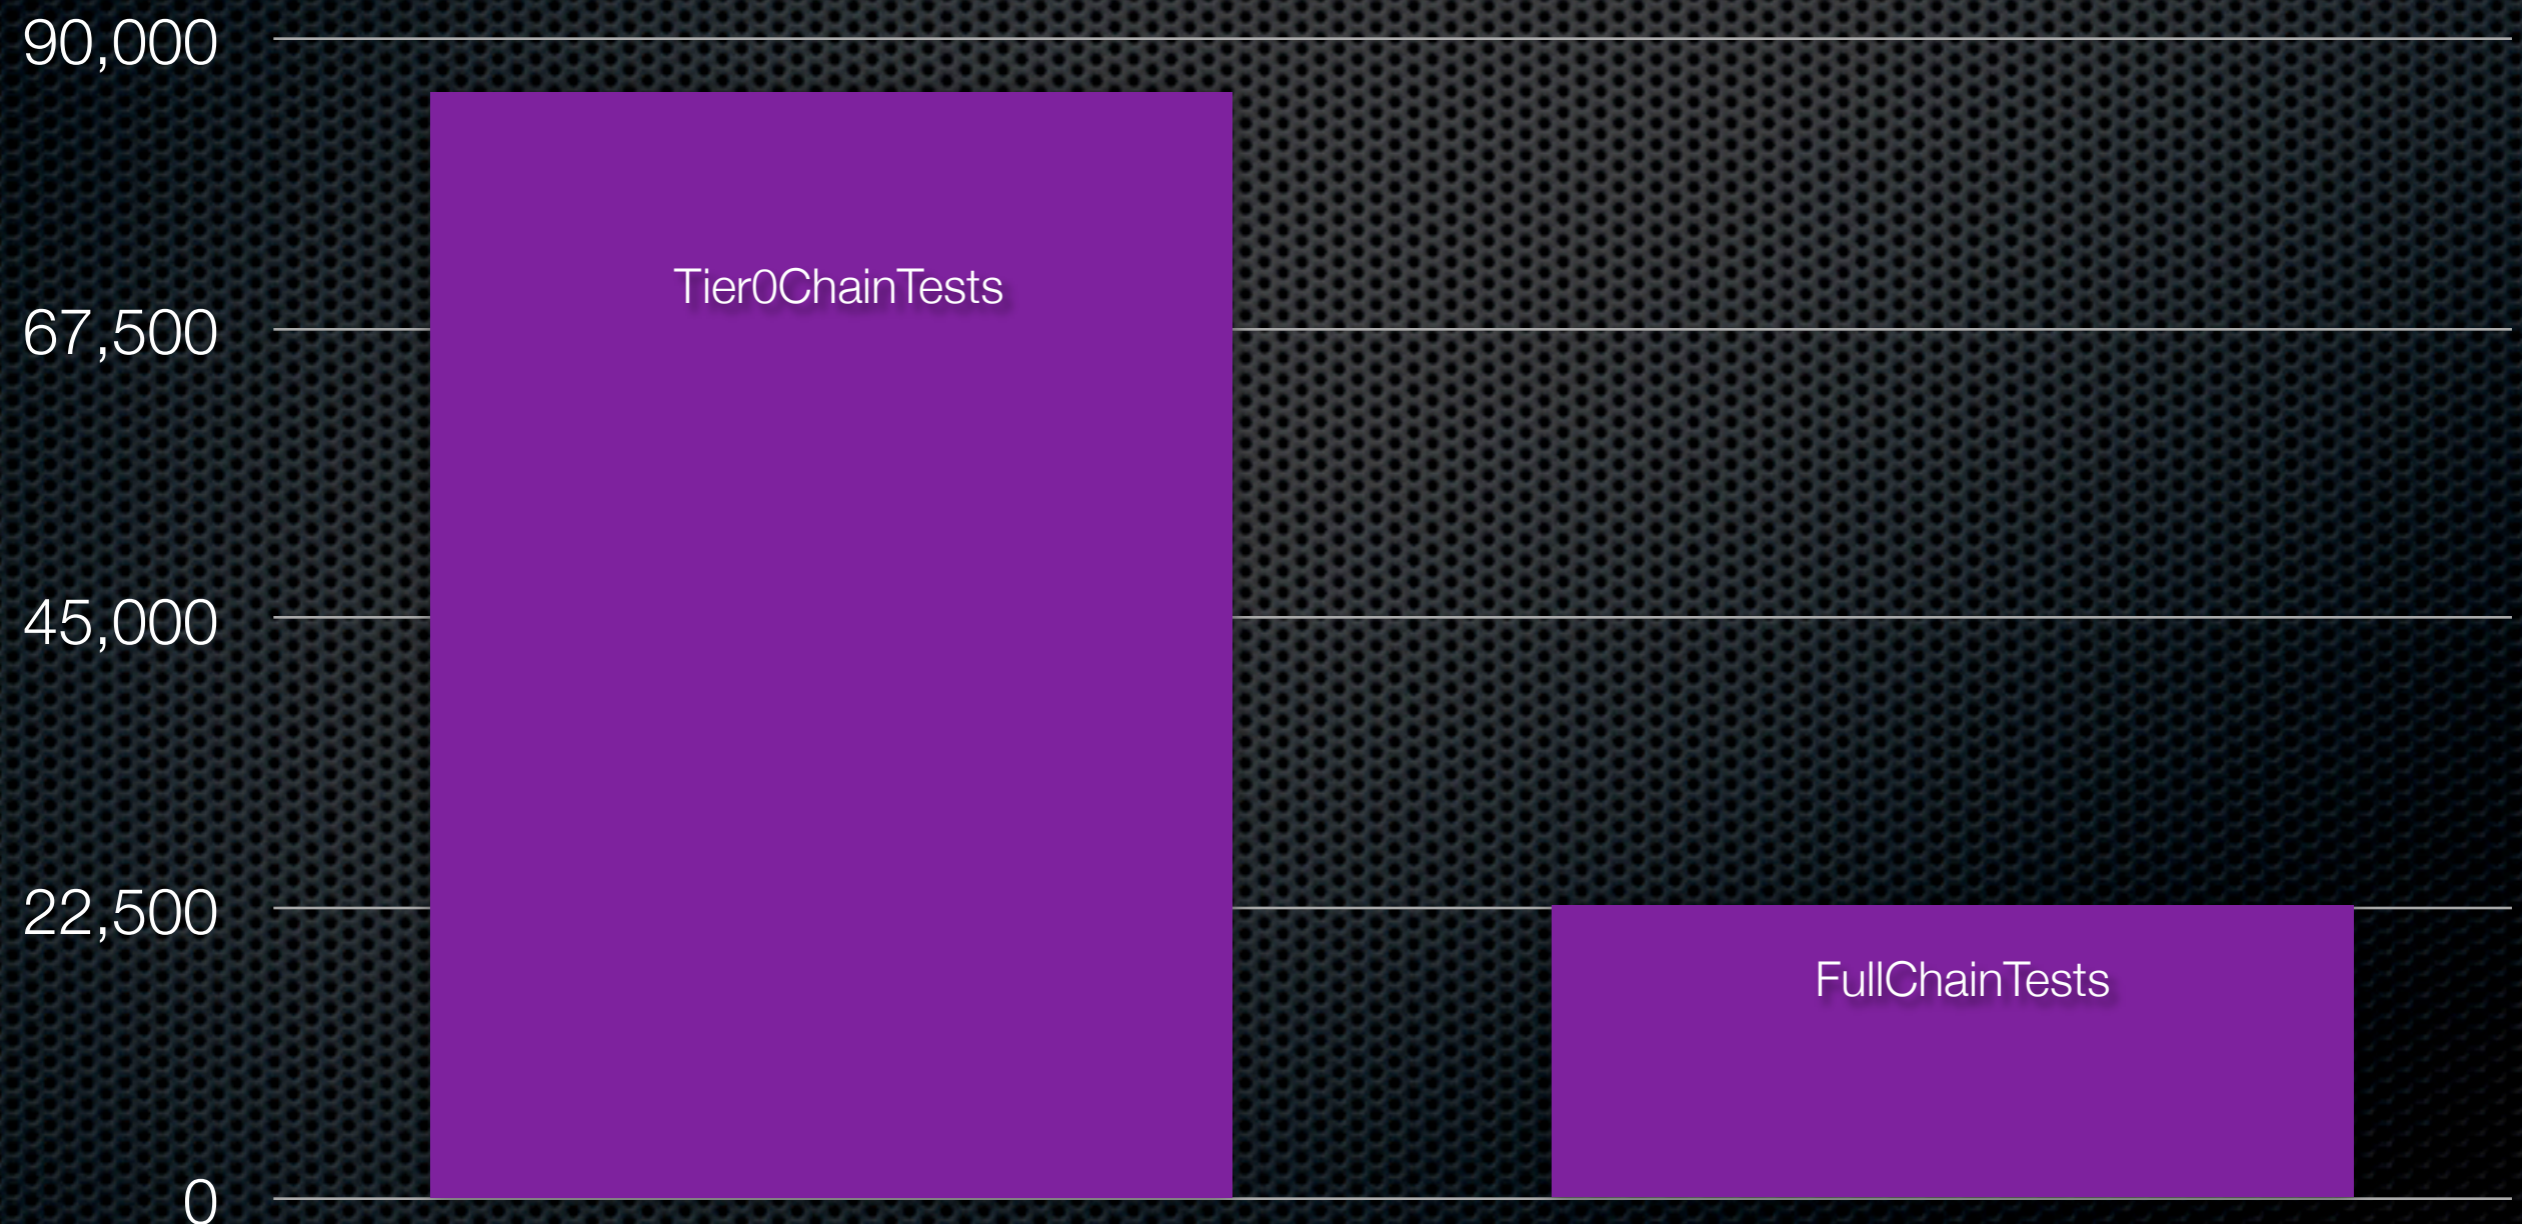


# TCT/FCT - 1 week

>> cat "state of nation"

Number of keep-files on AFS

~0.11M total/week

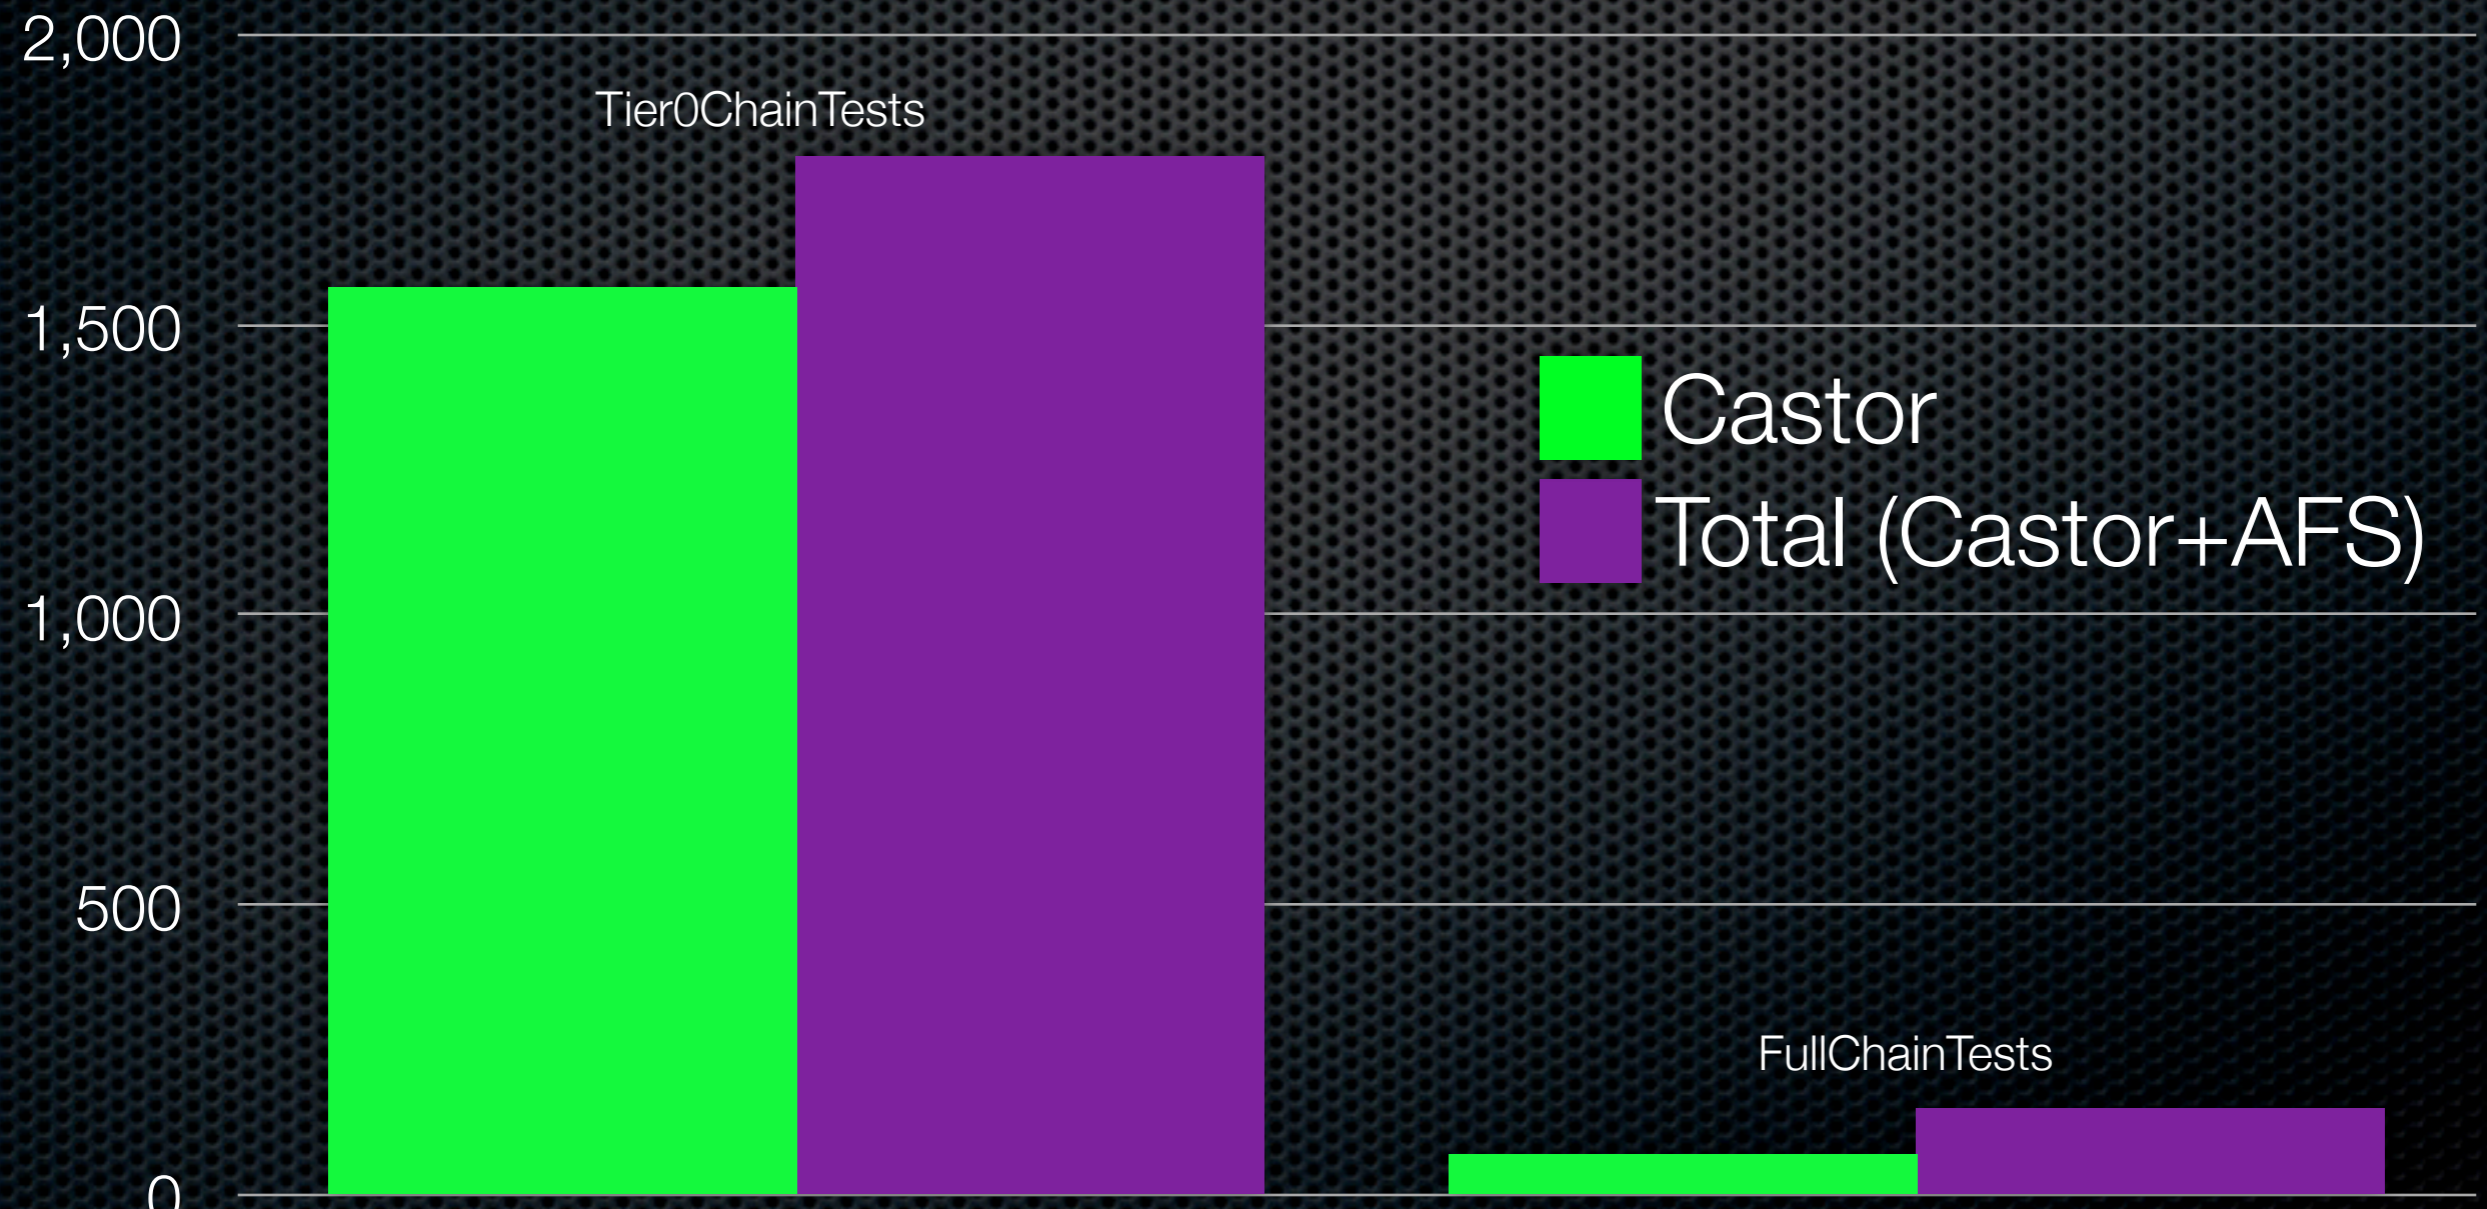


# TCT/FCT - 1 week

>> cat "state of nation"

Space usage (in GB)

~2TB total/week



**>> import future**

# What's coming?

>> `import __future__`

AtlasSetup support [~1 week]

Migration to SLC5 [~1 month]

Launch nodes:

- All 6 are SLC4
- 4 SLC5 nodes available and tested

Web server:

- One SLC4 machine
- IT want us to move to an SLC5 virtual machine
- Will require some nifty coordination to avoid excessive disruption

Improved web site [~1 month]

Behind the scenes:

- More logical database table heirarchy

On the website:

- Better error reporting

**>> echo “One last thing”**

(I've always wanted to say this....)

