

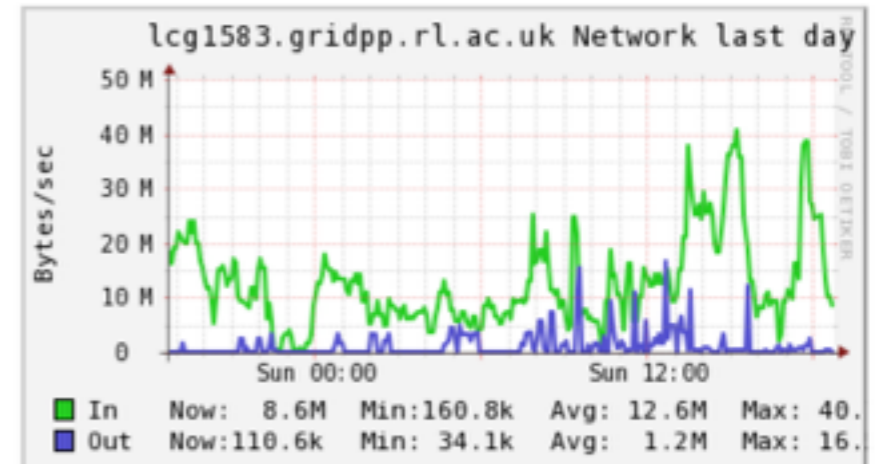
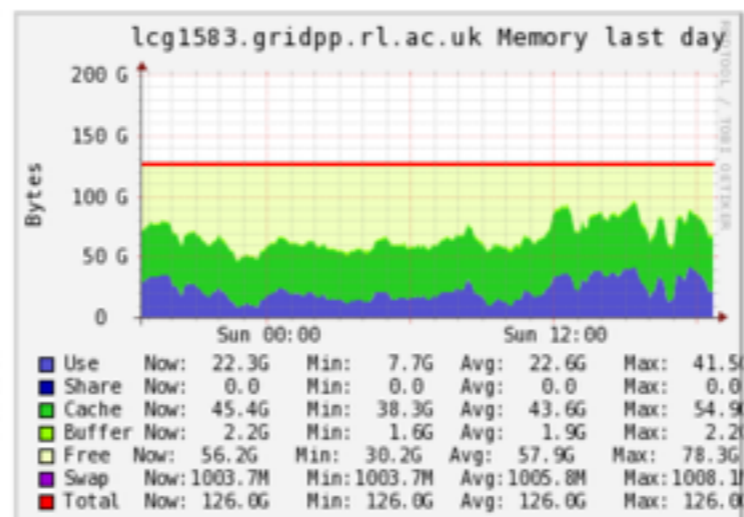

# Monitoring the resource usage of jobs using cAdvisor

Andrew Lahiff,  
HEP SYSMAN, June 2015



# Current monitoring

- Ganglia
  - CPU, memory, network, disk usage for the machine as a whole



- Could be useful to have this information *per job*
- Are there any existing monitoring tools that could provide this information?

# cAdvisor

- cAdvisor - Container Advisor
- open-source software from Google - [github.com/google/cadvisor](https://github.com/google/cadvisor)
- a daemon which collects, aggregates, processes & exports information about running containers & machine as a whole
- works with any type of container using cgroups
  - includes jobs on WNs if the batch system is configured to use cgroups

# cAdvisor

- exposes a web UI showing live statistics
- REST API
  - details of the machine
  - resource usage
  - events (container creation & deletion, OOM events, ...)
- can send data to InfluxDB
  - existing feature request for sending to Elasticsearch



# Installation - cAdvisor

- After installing go & mercurial, can build the executable:

```
go get -d github.com/google/cadvisor
```

```
go get github.com/tools/godep
```

```
cd $GOPATH/src/github.com/google/cadvisor
```

```
$GOPATH/bin/godep go build
```

- Copy the executable to the worker nodes, or make an rpm, ...

# Installation - InfluxDB

- Easy to install

```
rpm -ivh https://s3.amazonaws.com/influxdb/influxdb-latest-1.x86_64.rpm
```

```
service influxdb start
```

- Setup a database using the UI at <http://hostname:8083>

# Installation - Grafana

- Easy to install

```
rpm -ivh https://grafanarel.s3.amazonaws.com/builds/  
grafana-2.0.2-1.x86_64.rpm
```

```
service grafana-server start
```

- Add the InfluxDB database as a data source using the UI at <http://hostname:3000>

## Edit data source

Name	cadvisor	Default	
Type	InfluxDB 0.8.x		

## Http settings

Url	http://vm283.nubes.stfc.ac.uk:8086	Access		direct
Basic Auth	Enable	<input type="checkbox"/>		

## InfluxDB Details

Database	cadvisor		
User	wn	Password	*****

[Update](#) [Cancel](#)

# cAdvisor UI



# UI - overview of whole machine



Host

Docker Containers

Subcontainers

Isolation

CPU

Shares 1281 shares

Allowed Cores 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

Memory

Limit unlimited


Bswap Limit unlimited

Usage

Overview

Processes

Slack	PID	PPID	Start Time	CPU%	MEM%	MEM	Virtual Size	Stack	Running Time	Command	Uptime
gnome08	17544	17519	15:45	85.10	1.33	1.47 Kib	0.23 Kib	8K	00:03:14	bash	✓
gnome08	17562	17569	15:50	85.00	1.33	1.44 Kib	0.23 Kib	8K	00:03:14	bash	✓
gnome08	22608	22768	16:29	86.00	1.33	1.37 Kib	0.23 Kib	8K	00:03:14	bash	✓
gnome08	14585	14502	15:45	82.40	0.41	55.41 Kib	705.53 Kib	8K	00:03:14	bash	✓
gnome08	5114	5071	16:25	22.70	0.90	1.26 Kib	1.62 Kib	8K	00:02:12	bash	✓
gnome08	3544	3497	15:45	25.30	1.00	1.34 Kib	1.89 Kib	8K	00:03:14	bash	✓
gnome08	18562	18561	09:54	25.00	1.00	1.35 Kib	1.91 Kib	8K	00:03:14	bash	✓
gnome08	3427	3443	10:05	22.00	1.00	1.31 Kib	1.89 Kib	8K	00:03:14	bash	✓
gnome08	16428	16664	09:48	22.00	1.00	1.35 Kib	1.92 Kib	8K	00:03:14	bash	✓
gnome08	18242	18241	09:54	21.90	1.00	1.31 Kib	1.88 Kib	8K	00:02:19	bash	✓
gnome08	3478	3762	10:05	21.80	1.00	1.29 Kib	1.89 Kib	8K	00:03:14	bash	✓
gnome08	3438	3372	10:05	21.40	1.00	1.32 Kib	1.89 Kib	8K	00:03:14	bash	✓
gnome08	3371	3335	10:05	21.30	1.00	1.32 Kib	1.89 Kib	8K	00:03:14	bash	✓
gnome08	3541	3443	10:05	21.10	1.00	1.29 Kib	1.91 Kib	8K	00:03:14	bash	✓
gnome08	16421	16638	09:48	21.10	1.00	1.31 Kib	1.92 Kib	8K	00:03:14	bash	✓
gnome08	18175	18111	09:54	21.10	1.00	1.32 Kib	1.92 Kib	8K	00:03:14	bash	✓
gnome08	3709	3694	10:05	21.00	1.00	1.33 Kib	1.92 Kib	8K	00:03:14	bash	✓
gnome08	16401	16344	09:48	21.00	1.00	1.31 Kib	1.89 Kib	8K	00:03:14	bash	✓
gnome08	5116	5104	10:05	21.00	1.00	1.38 Kib	1.90 Kib	8K	00:03:14	bash	✓
gnome08	18162	16262	09:48	20.90	1.00	1.31 Kib	1.91 Kib	8K	00:03:14	bash	✓
gnome08	18128	18063	09:54	20.90	1.00	1.34 Kib	1.92 Kib	8K	00:03:14	bash	✓
gnome08	3479	3462	10:05	20.80	1.00	1.28 Kib	1.92 Kib	8K	00:03:14	bash	✓
gnome08	16420	16392	09:48	20.40	1.00	1.32 Kib	1.92 Kib	8K	00:03:14	bash	✓
gnome08	3708	3612	10:05	20.40	1.00	1.33 Kib	1.89 Kib	8K	00:03:14	bash	✓
gnome08	3490	3416	10:05	20.30	1.00	1.38 Kib	1.91 Kib	8K	00:03:14	bash	✓



# UI - overview of whole machine



## Isolation

Host

Docker Containers

Subcontainers

CPU

Shares 1024 shares

Allowed Cores 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

Memory

Limit unlimited


Bswap Limit unlimited

Usage

Overview

Processes

Slack	PID	PPID	Start Time	CPU%	MEM%	RSS	Virtual Size	Stack	Running Time	Command	Uptime
✓	17544	17519	15:45	85.10	1.33	1.47 K	0.03 K	8K	00:00:14	catmon	✓
✓	17563	17569	15:50	85.00	1.33	1.48 K	0.03 K	8K	00:00:14	catmon	✓
✓	22608	22768	16:29	86.00	1.33	1.37 K	0.03 K	8K	00:00:18	catmon	✓
✓	14585	14502	15:45	82.40	0.41	33.41 K	705.53 K	80K	00:00:18	catmon	✓
✓	5114	3071	16:25	22.70	0.90	1.26 K	0.02 K	80K	00:00:22	catmon	✓
✓	3544	3037	15:45	25.30	1.00	1.34 K	0.02 K	80K	00:00:14	catmon	✓
✓	18562	18561	09:54	25.20	1.00	1.35 K	0.02 K	80K	00:00:14	catmon	✓
✓	3627	3043	16:05	22.00	1.00	1.31 K	0.02 K	80K	00:00:14	catmon	✓
✓	16628	16664	09:48	22.00	1.00	1.35 K	0.02 K	80K	00:00:24	catmon	✓
✓	18262	18261	09:54	21.90	1.00	1.31 K	0.02 K	80K	00:00:27	catmon	✓
✓	3878	3782	10:25	21.80	1.00	1.29 K	0.02 K	80K	00:00:15	catmon	✓
✓	3438	3372	10:05	21.40	1.00	1.32 K	0.02 K	80K	00:00:14	catmon	✓
✓	3371	3335	10:05	21.30	1.00	1.32 K	0.02 K	80K	00:00:14	catmon	✓
✓	3541	3484	10:05	21.10	1.00	1.29 K	0.02 K	80K	00:00:15	catmon	✓
✓	16621	16638	09:48	21.10	1.00	1.31 K	0.02 K	80K	00:00:15	catmon	✓
✓	18175	18111	09:54	21.10	1.00	1.32 K	0.02 K	80K	00:00:15	catmon	✓
✓	3789	3684	10:25	21.00	1.00	1.33 K	0.02 K	80K	00:00:15	catmon	✓
✓	16601	16344	09:48	21.00	1.00	1.31 K	0.02 K	80K	00:00:15	catmon	✓
✓	5116	3104	16:05	21.00	1.00	1.38 K	0.02 K	80K	00:00:13	catmon	✓
✓	18162	16882	09:48	20.90	1.00	1.31 K	0.02 K	80K	00:00:13	catmon	✓
✓	18128	18063	09:54	20.90	1.00	1.34 K	0.02 K	80K	00:00:23	catmon	✓
✓	3979	3852	10:05	20.80	1.00	1.28 K	0.02 K	80K	00:00:12	catmon	✓
✓	16620	16392	09:48	20.40	1.00	1.32 K	0.02 K	80K	00:00:18	catmon	✓
✓	3788	3612	10:25	20.40	1.00	1.33 K	0.02 K	80K	00:00:12	catmon	✓
✓	3980	3916	10:05	20.30	1.00	1.38 K	0.02 K	80K	00:00:14	catmon	✓



# UI - overview of whole machine



Isolation

Host: /

Docker Containers

Subcontainers

Number: [input field]

Resources

CPU

Shares: 1024 shares

Allowed Cores: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

Memory


Limit: unlimited

Bswap Limit: unlimited



Processes

Slack	PID	PPID	Start Time	CPU%	MEM%	MEM	Virtual Size	Stack	Running Time	Container	Group
general	17444	17419	15:45	85.10	1.33	1.87 K	0.23 K	8K	00:00:14	mainfun	✓
general	17443	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17442	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17441	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17440	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17439	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17438	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17437	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17436	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17435	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17434	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17433	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17432	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17431	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17430	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17429	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17428	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17427	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17426	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17425	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17424	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17423	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17422	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17421	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17420	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓
general	17419	17419	15:45	85.00	1.33	1.88 K	0.23 K	8K	00:00:14	mainfun	✓



Processes

# UI - overview of whole machine



Isolation

Processes

Host

Docker Containers

Subcontainers

CPU

Shares 1024 shares

Allowed Cores 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

Memory

Limit unlimited

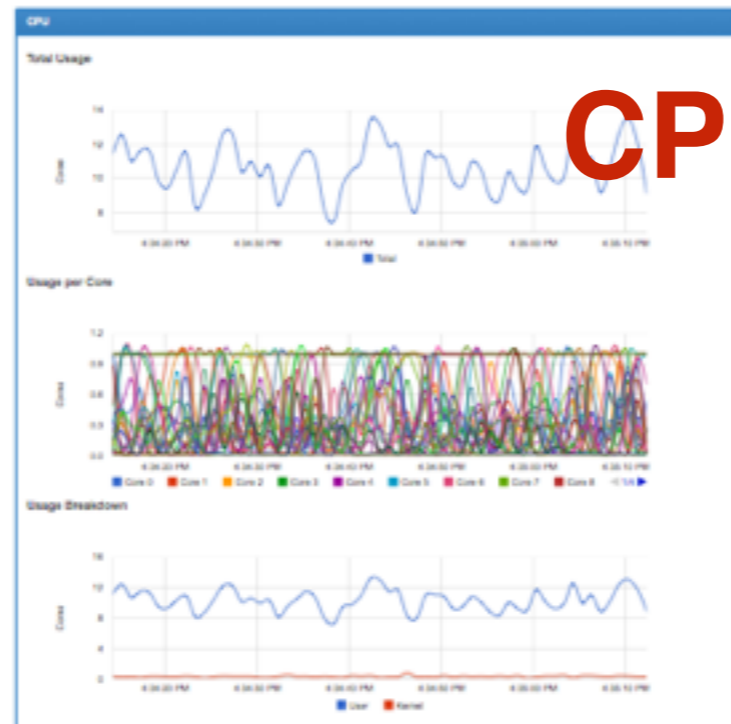
Bswap Limit unlimited

Usage


Overview

Processes

PID	PPID	Start Time	CPU%	MEM%	RSS	Virtual Size	Status	Running Time	Command	Uptime
17944	17933	09:45	89.10	1.33	1.87 K	0.23 K	SS	00:00:14	cat	✓
17945	17933	09:45	89.10	1.33	1.88 K	0.23 K	SS	00:00:14	cat	✓
17946	17933	09:45	89.10	1.33	1.89 K	0.23 K	SS	00:00:14	cat	✓
17947	17933	09:45	89.10	1.33	1.90 K	0.23 K	SS	00:00:14	cat	✓
17948	17933	09:45	89.10	1.33	1.91 K	0.23 K	SS	00:00:14	cat	✓
17949	17933	09:45	89.10	1.33	1.92 K	0.23 K	SS	00:00:14	cat	✓
17950	17933	09:45	89.10	1.33	1.93 K	0.23 K	SS	00:00:14	cat	✓
17951	17933	09:45	89.10	1.33	1.94 K	0.23 K	SS	00:00:14	cat	✓
17952	17933	09:45	89.10	1.33	1.95 K	0.23 K	SS	00:00:14	cat	✓
17953	17933	09:45	89.10	1.33	1.96 K	0.23 K	SS	00:00:14	cat	✓
17954	17933	09:45	89.10	1.33	1.97 K	0.23 K	SS	00:00:14	cat	✓
17955	17933	09:45	89.10	1.33	1.98 K	0.23 K	SS	00:00:14	cat	✓
17956	17933	09:45	89.10	1.33	1.99 K	0.23 K	SS	00:00:14	cat	✓
17957	17933	09:45	89.10	1.33	2.00 K	0.23 K	SS	00:00:14	cat	✓
17958	17933	09:45	89.10	1.33	2.01 K	0.23 K	SS	00:00:14	cat	✓
17959	17933	09:45	89.10	1.33	2.02 K	0.23 K	SS	00:00:14	cat	✓
17960	17933	09:45	89.10	1.33	2.03 K	0.23 K	SS	00:00:14	cat	✓
17961	17933	09:45	89.10	1.33	2.04 K	0.23 K	SS	00:00:14	cat	✓
17962	17933	09:45	89.10	1.33	2.05 K	0.23 K	SS	00:00:14	cat	✓
17963	17933	09:45	89.10	1.33	2.06 K	0.23 K	SS	00:00:14	cat	✓
17964	17933	09:45	89.10	1.33	2.07 K	0.23 K	SS	00:00:14	cat	✓
17965	17933	09:45	89.10	1.33	2.08 K	0.23 K	SS	00:00:14	cat	✓
17966	17933	09:45	89.10	1.33	2.09 K	0.23 K	SS	00:00:14	cat	✓
17967	17933	09:45	89.10	1.33	2.10 K	0.23 K	SS	00:00:14	cat	✓
17968	17933	09:45	89.10	1.33	2.11 K	0.23 K	SS	00:00:14	cat	✓
17969	17933	09:45	89.10	1.33	2.12 K	0.23 K	SS	00:00:14	cat	✓
17970	17933	09:45	89.10	1.33	2.13 K	0.23 K	SS	00:00:14	cat	✓
17971	17933	09:45	89.10	1.33	2.14 K	0.23 K	SS	00:00:14	cat	✓
17972	17933	09:45	89.10	1.33	2.15 K	0.23 K	SS	00:00:14	cat	✓
17973	17933	09:45	89.10	1.33	2.16 K	0.23 K	SS	00:00:14	cat	✓
17974	17933	09:45	89.10	1.33	2.17 K	0.23 K	SS	00:00:14	cat	✓
17975	17933	09:45	89.10	1.33	2.18 K	0.23 K	SS	00:00:14	cat	✓
17976	17933	09:45	89.10	1.33	2.19 K	0.23 K	SS	00:00:14	cat	✓
17977	17933	09:45	89.10	1.33	2.20 K	0.23 K	SS	00:00:14	cat	✓
17978	17933	09:45	89.10	1.33	2.21 K	0.23 K	SS	00:00:14	cat	✓
17979	17933	09:45	89.10	1.33	2.22 K	0.23 K	SS	00:00:14	cat	✓
17980	17933	09:45	89.10	1.33	2.23 K	0.23 K	SS	00:00:14	cat	✓
17981	17933	09:45	89.10	1.33	2.24 K	0.23 K	SS	00:00:14	cat	✓
17982	17933	09:45	89.10	1.33	2.25 K	0.23 K	SS	00:00:14	cat	✓
17983	17933	09:45	89.10	1.33	2.26 K	0.23 K	SS	00:00:14	cat	✓
17984	17933	09:45	89.10	1.33	2.27 K	0.23 K	SS	00:00:14	cat	✓
17985	17933	09:45	89.10	1.33	2.28 K	0.23 K	SS	00:00:14	cat	✓
17986	17933	09:45	89.10	1.33	2.29 K	0.23 K	SS	00:00:14	cat	✓
17987	17933	09:45	89.10	1.33	2.30 K	0.23 K	SS	00:00:14	cat	✓
17988	17933	09:45	89.10	1.33	2.31 K	0.23 K	SS	00:00:14	cat	✓
17989	17933	09:45	89.10	1.33	2.32 K	0.23 K	SS	00:00:14	cat	✓
17990	17933	09:45	89.10	1.33	2.33 K	0.23 K	SS	00:00:14	cat	✓
17991	17933	09:45	89.10	1.33	2.34 K	0.23 K	SS	00:00:14	cat	✓
17992	17933	09:45	89.10	1.33	2.35 K	0.23 K	SS	00:00:14	cat	✓
17993	17933	09:45	89.10	1.33	2.36 K	0.23 K	SS	00:00:14	cat	✓
17994	17933	09:45	89.10	1.33	2.37 K	0.23 K	SS	00:00:14	cat	✓
17995	17933	09:45	89.10	1.33	2.38 K	0.23 K	SS	00:00:14	cat	✓
17996	17933	09:45	89.10	1.33	2.39 K	0.23 K	SS	00:00:14	cat	✓
17997	17933	09:45	89.10	1.33	2.40 K	0.23 K	SS	00:00:14	cat	✓
17998	17933	09:45	89.10	1.33	2.41 K	0.23 K	SS	00:00:14	cat	✓
17999	17933	09:45	89.10	1.33	2.42 K	0.23 K	SS	00:00:14	cat	✓
18000	17933	09:45	89.10	1.33	2.43 K	0.23 K	SS	00:00:14	cat	✓



CPU



# UI - overview of whole machine



Isolation

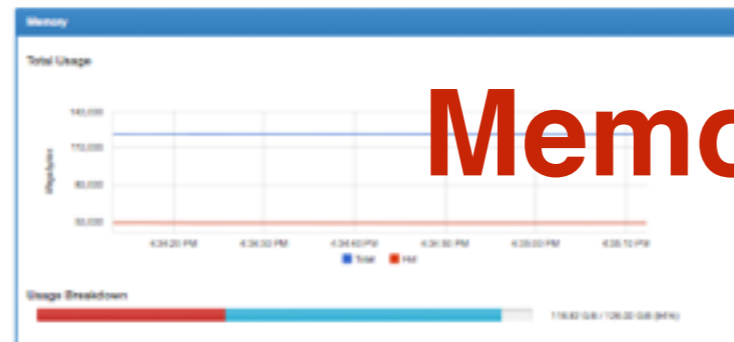

Processes

The screenshot shows the cAdvisor web interface. At the top, there's a navigation bar with 'Host' and 'Docker Containers'. Below that, there's a 'Subcontainers' section with a search bar. The main content area is divided into several sections: 'CPU' (Shares 1024 cores, Allowed Cores 0-31), 'Memory' (Limit unlimited, Swap Limit unlimited), and 'Usage' (Overview with six gauges). At the bottom, there's a 'Processes' table with columns for PID, PPID, Start Time, CPU%, MEM%, RSS, Virtual Size, Status, Running Time, Command, and Group.

PID	PPID	Start Time	CPU%	MEM%	RSS	Virtual Size	Status	Running Time	Command	Group
17144	17119	15:45	85.10	1.33	1.87 K	2.23 K	RR	00:00:14	cat	cat
17145	17119	15:45	85.10	1.33	1.88 K	2.27 K	RR	00:00:14	cat	cat
17146	17119	15:45	85.10	1.33	1.89 K	2.31 K	RR	00:00:14	cat	cat
17147	17119	15:45	85.10	1.33	1.90 K	2.35 K	RR	00:00:14	cat	cat
17148	17119	15:45	85.10	1.33	1.91 K	2.39 K	RR	00:00:14	cat	cat
17149	17119	15:45	85.10	1.33	1.92 K	2.43 K	RR	00:00:14	cat	cat
17150	17119	15:45	85.10	1.33	1.93 K	2.47 K	RR	00:00:14	cat	cat
17151	17119	15:45	85.10	1.33	1.94 K	2.51 K	RR	00:00:14	cat	cat
17152	17119	15:45	85.10	1.33	1.95 K	2.55 K	RR	00:00:14	cat	cat
17153	17119	15:45	85.10	1.33	1.96 K	2.59 K	RR	00:00:14	cat	cat
17154	17119	15:45	85.10	1.33	1.97 K	2.63 K	RR	00:00:14	cat	cat
17155	17119	15:45	85.10	1.33	1.98 K	2.67 K	RR	00:00:14	cat	cat
17156	17119	15:45	85.10	1.33	1.99 K	2.71 K	RR	00:00:14	cat	cat
17157	17119	15:45	85.10	1.33	2.00 K	2.75 K	RR	00:00:14	cat	cat
17158	17119	15:45	85.10	1.33	2.01 K	2.79 K	RR	00:00:14	cat	cat
17159	17119	15:45	85.10	1.33	2.02 K	2.83 K	RR	00:00:14	cat	cat
17160	17119	15:45	85.10	1.33	2.03 K	2.87 K	RR	00:00:14	cat	cat
17161	17119	15:45	85.10	1.33	2.04 K	2.91 K	RR	00:00:14	cat	cat
17162	17119	15:45	85.10	1.33	2.05 K	2.95 K	RR	00:00:14	cat	cat
17163	17119	15:45	85.10	1.33	2.06 K	2.99 K	RR	00:00:14	cat	cat
17164	17119	15:45	85.10	1.33	2.07 K	3.03 K	RR	00:00:14	cat	cat
17165	17119	15:45	85.10	1.33	2.08 K	3.07 K	RR	00:00:14	cat	cat
17166	17119	15:45	85.10	1.33	2.09 K	3.11 K	RR	00:00:14	cat	cat
17167	17119	15:45	85.10	1.33	2.10 K	3.15 K	RR	00:00:14	cat	cat
17168	17119	15:45	85.10	1.33	2.11 K	3.19 K	RR	00:00:14	cat	cat
17169	17119	15:45	85.10	1.33	2.12 K	3.23 K	RR	00:00:14	cat	cat
17170	17119	15:45	85.10	1.33	2.13 K	3.27 K	RR	00:00:14	cat	cat
17171	17119	15:45	85.10	1.33	2.14 K	3.31 K	RR	00:00:14	cat	cat
17172	17119	15:45	85.10	1.33	2.15 K	3.35 K	RR	00:00:14	cat	cat
17173	17119	15:45	85.10	1.33	2.16 K	3.39 K	RR	00:00:14	cat	cat
17174	17119	15:45	85.10	1.33	2.17 K	3.43 K	RR	00:00:14	cat	cat
17175	17119	15:45	85.10	1.33	2.18 K	3.47 K	RR	00:00:14	cat	cat
17176	17119	15:45	85.10	1.33	2.19 K	3.51 K	RR	00:00:14	cat	cat
17177	17119	15:45	85.10	1.33	2.20 K	3.55 K	RR	00:00:14	cat	cat
17178	17119	15:45	85.10	1.33	2.21 K	3.59 K	RR	00:00:14	cat	cat
17179	17119	15:45	85.10	1.33	2.22 K	3.63 K	RR	00:00:14	cat	cat
17180	17119	15:45	85.10	1.33	2.23 K	3.67 K	RR	00:00:14	cat	cat
17181	17119	15:45	85.10	1.33	2.24 K	3.71 K	RR	00:00:14	cat	cat
17182	17119	15:45	85.10	1.33	2.25 K	3.75 K	RR	00:00:14	cat	cat
17183	17119	15:45	85.10	1.33	2.26 K	3.79 K	RR	00:00:14	cat	cat
17184	17119	15:45	85.10	1.33	2.27 K	3.83 K	RR	00:00:14	cat	cat
17185	17119	15:45	85.10	1.33	2.28 K	3.87 K	RR	00:00:14	cat	cat
17186	17119	15:45	85.10	1.33	2.29 K	3.91 K	RR	00:00:14	cat	cat
17187	17119	15:45	85.10	1.33	2.30 K	3.95 K	RR	00:00:14	cat	cat
17188	17119	15:45	85.10	1.33	2.31 K	3.99 K	RR	00:00:14	cat	cat
17189	17119	15:45	85.10	1.33	2.32 K	4.03 K	RR	00:00:14	cat	cat
17190	17119	15:45	85.10	1.33	2.33 K	4.07 K	RR	00:00:14	cat	cat
17191	17119	15:45	85.10	1.33	2.34 K	4.11 K	RR	00:00:14	cat	cat
17192	17119	15:45	85.10	1.33	2.35 K	4.15 K	RR	00:00:14	cat	cat
17193	17119	15:45	85.10	1.33	2.36 K	4.19 K	RR	00:00:14	cat	cat
17194	17119	15:45	85.10	1.33	2.37 K	4.23 K	RR	00:00:14	cat	cat
17195	17119	15:45	85.10	1.33	2.38 K	4.27 K	RR	00:00:14	cat	cat
17196	17119	15:45	85.10	1.33	2.39 K	4.31 K	RR	00:00:14	cat	cat
17197	17119	15:45	85.10	1.33	2.40 K	4.35 K	RR	00:00:14	cat	cat
17198	17119	15:45	85.10	1.33	2.41 K	4.39 K	RR	00:00:14	cat	cat
17199	17119	15:45	85.10	1.33	2.42 K	4.43 K	RR	00:00:14	cat	cat
17200	17119	15:45	85.10	1.33	2.43 K	4.47 K	RR	00:00:14	cat	cat



CPU



Memory



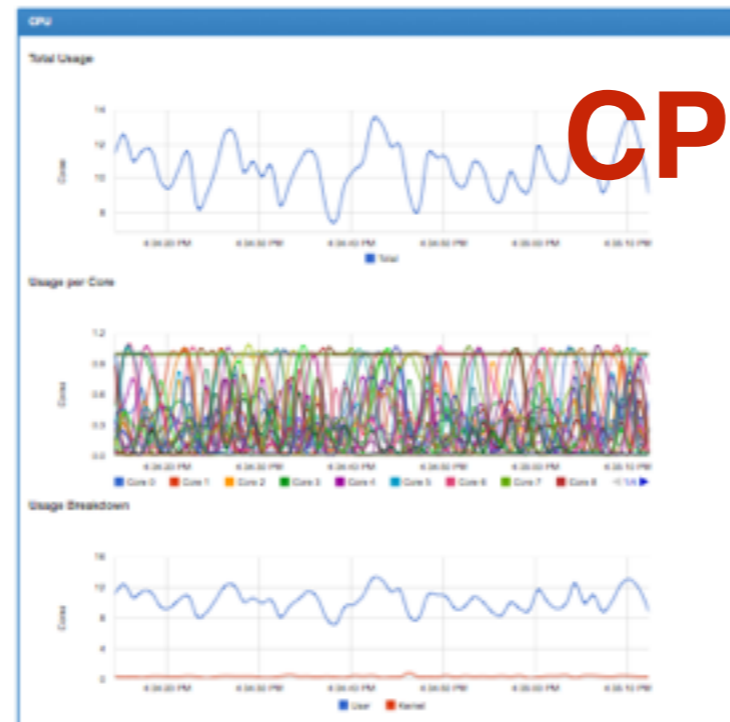
# UI - overview of whole machine



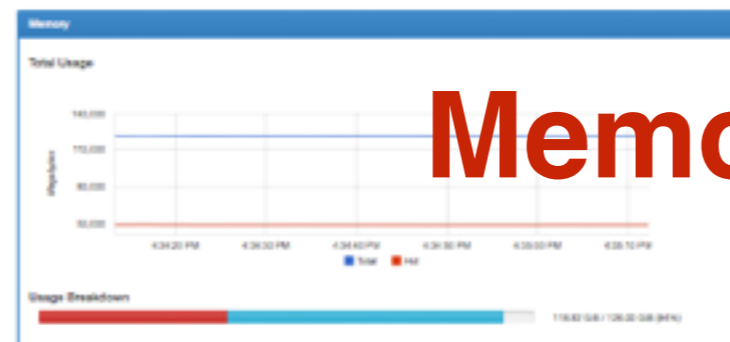
Isolation

Processes

The screenshot shows the cAdvisor web interface. At the top, there's a search bar and navigation links for 'Docker Containers' and 'Subcontainers'. Below that, there are sections for 'CPU', 'Memory', and 'Usage'. The 'Usage' section contains several circular gauges. At the bottom, there is a 'Processes' table with columns for PID, PPID, Start Time, CPU%, MEM%, RSS, Virtual Size, Status, Running Time, Command, and OomKill. The table lists various processes, including 'gnome-shell' and 'gnome-terminal'.




CPU



Memory



Network



# UI - overview of whole machine



**Isolation**

Host

Docker Containers

Subcontainers

Resources

CPU

Shares 1024 shares

Allowed Cores 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

Memory

Limit unlimited

Bswap Limit unlimited

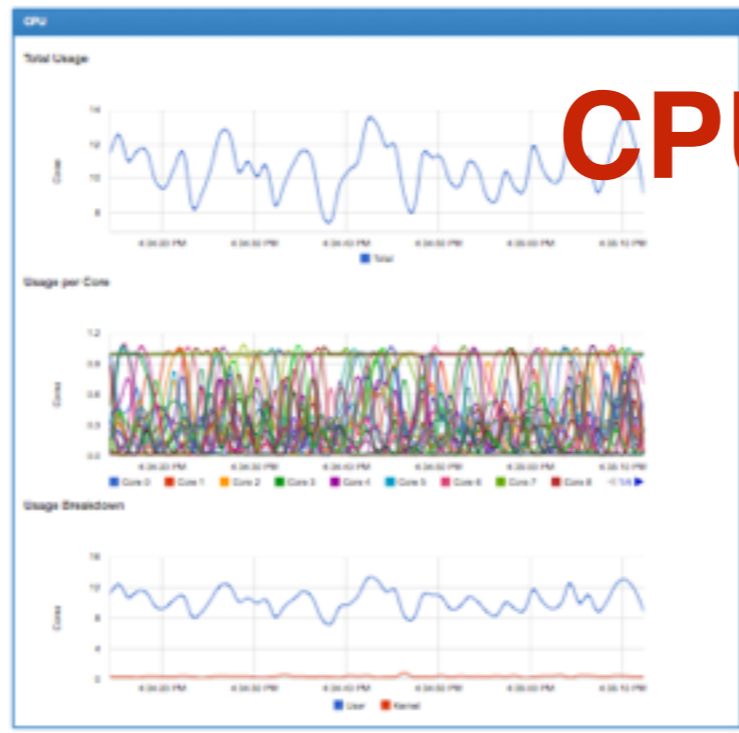
Usage

Overview

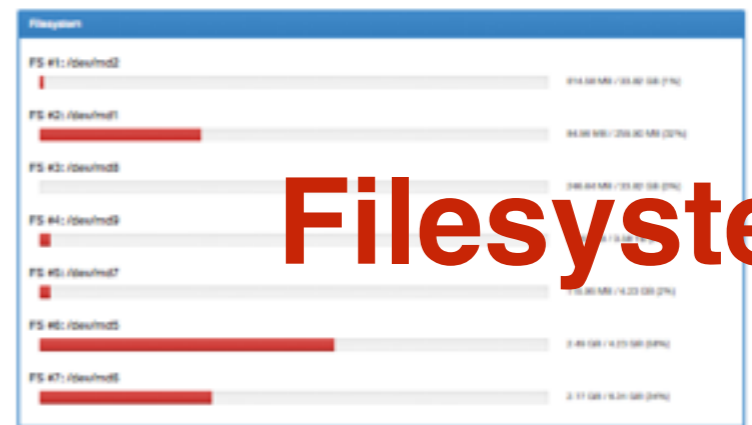
Processes

PID	PPID	Start Time	CPU%	MEM%	RSS	Virtual Size	Status	Running Time	Command	Group
17144	17133	09:45	89.10	1.33	1.87 K	0.03 K	SS	00:00:14	cat	
17145	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17146	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17147	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17148	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17149	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17150	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17151	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17152	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17153	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17154	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17155	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17156	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17157	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17158	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17159	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17160	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17161	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17162	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17163	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17164	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17165	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17166	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17167	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17168	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17169	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17170	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17171	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17172	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17173	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17174	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17175	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17176	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17177	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17178	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17179	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17180	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17181	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17182	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17183	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17184	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17185	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17186	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17187	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17188	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17189	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17190	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17191	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17192	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17193	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17194	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17195	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17196	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17197	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17198	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17199	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	
17200	17133	09:45	89.00	1.33	1.88 K	0.03 K	SS	00:00:14	cat	

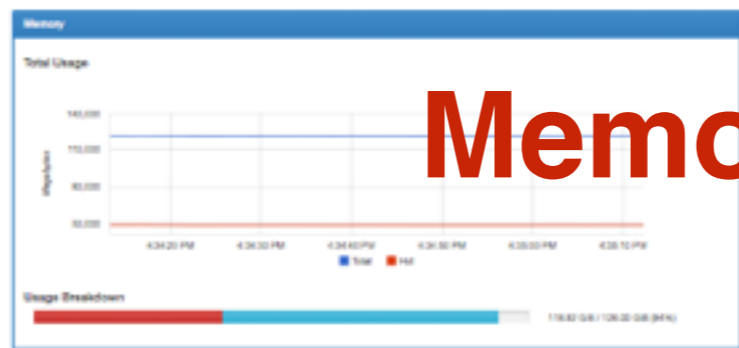
**CPU**



**Filesystem**



**Memory**



**Processes**

**Network**



# Running jobs

## /htcondor

root / htcondor

### Subcontainers

/htcondor/condor\_pool\_condor\_slot1\_1@lcg1583.gridpp.rl.ac.uk

/htcondor/condor\_pool\_condor\_slot1\_2@lcg1583.gridpp.rl.ac.uk

/htcondor/condor\_pool\_condor\_slot1\_3@lcg1583.gridpp.rl.ac.uk

/htcondor/condor\_pool\_condor\_slot1\_4@lcg1583.gridpp.rl.ac.uk

### Isolation

CPU

Shares 1024 shares

Memory


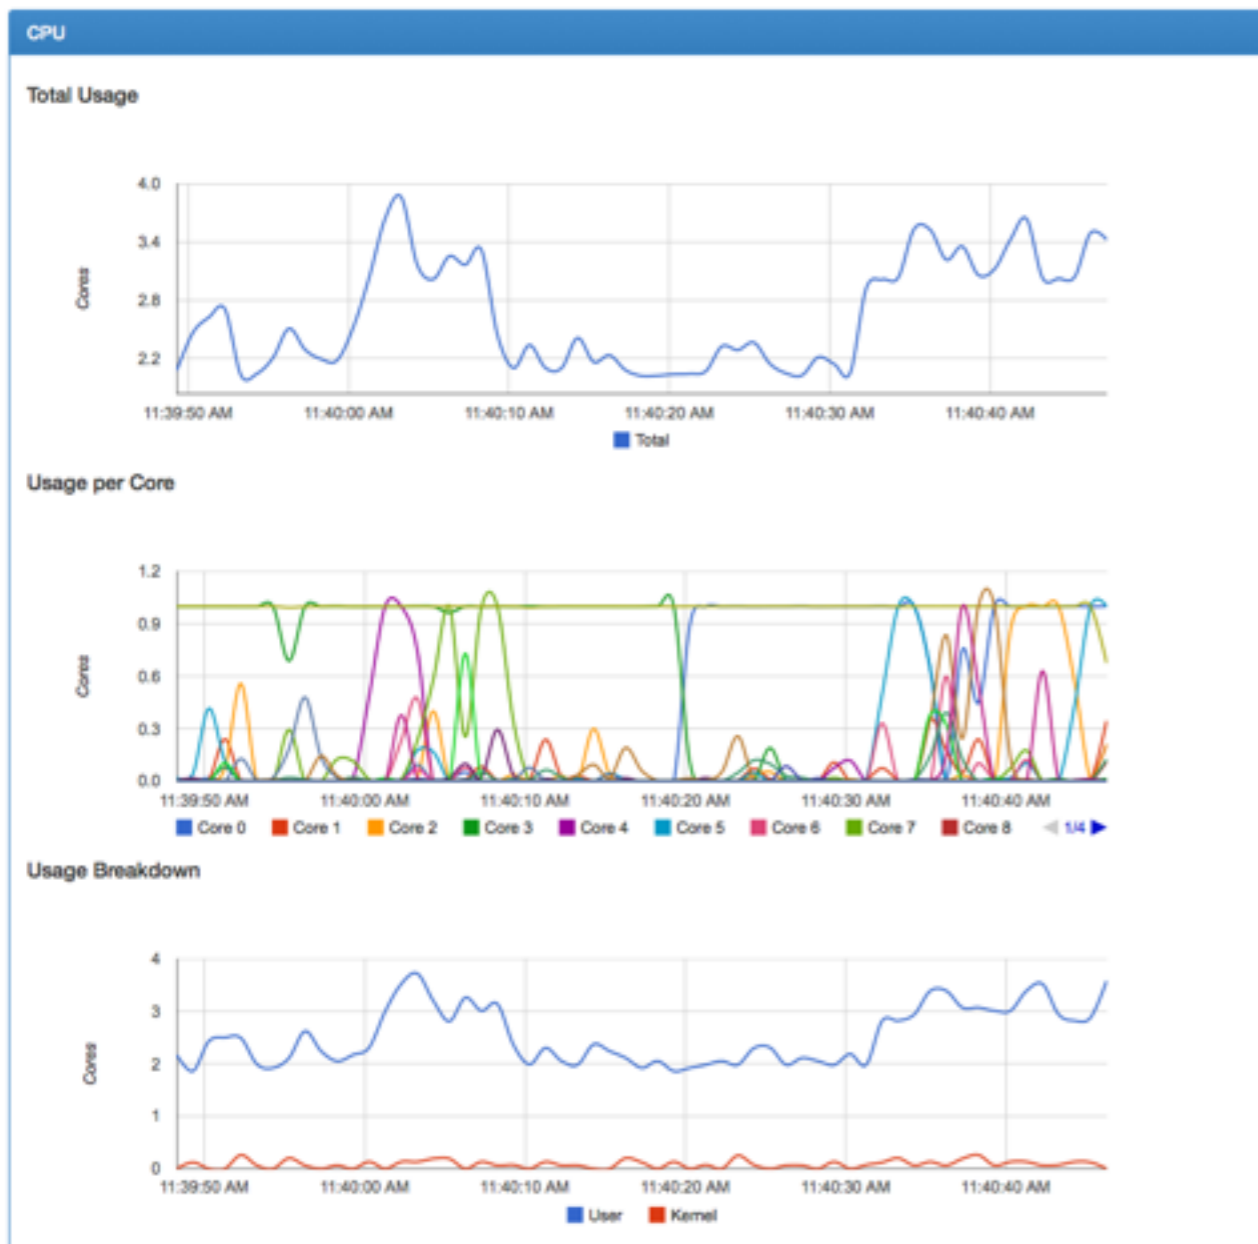
Limit 126.00 GB

Swap Limit 151.20 GB



# Individual jobs

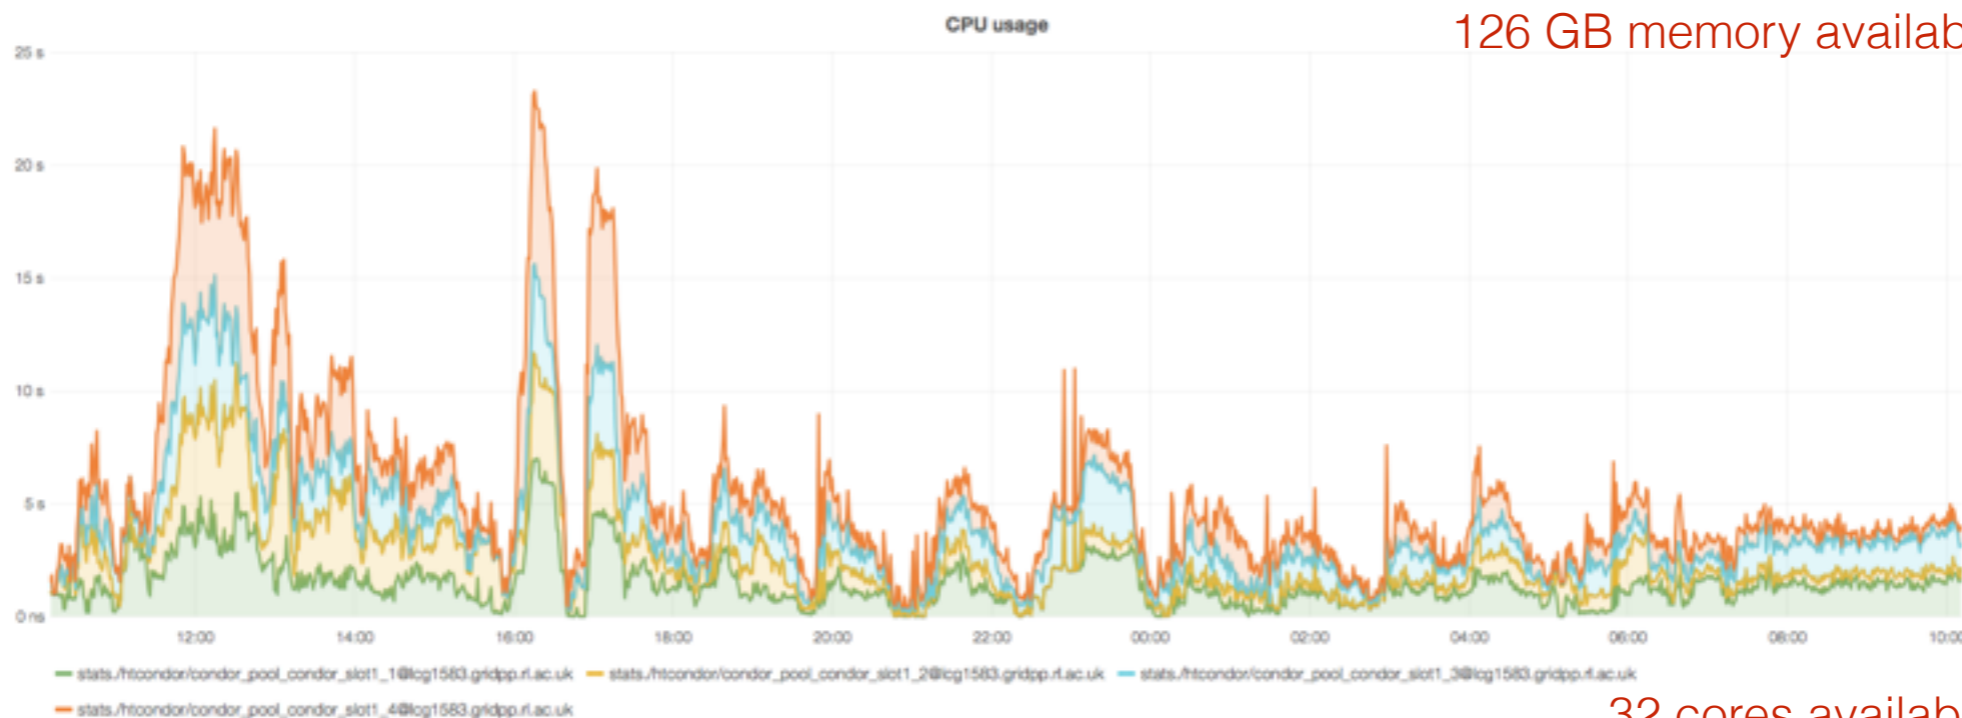
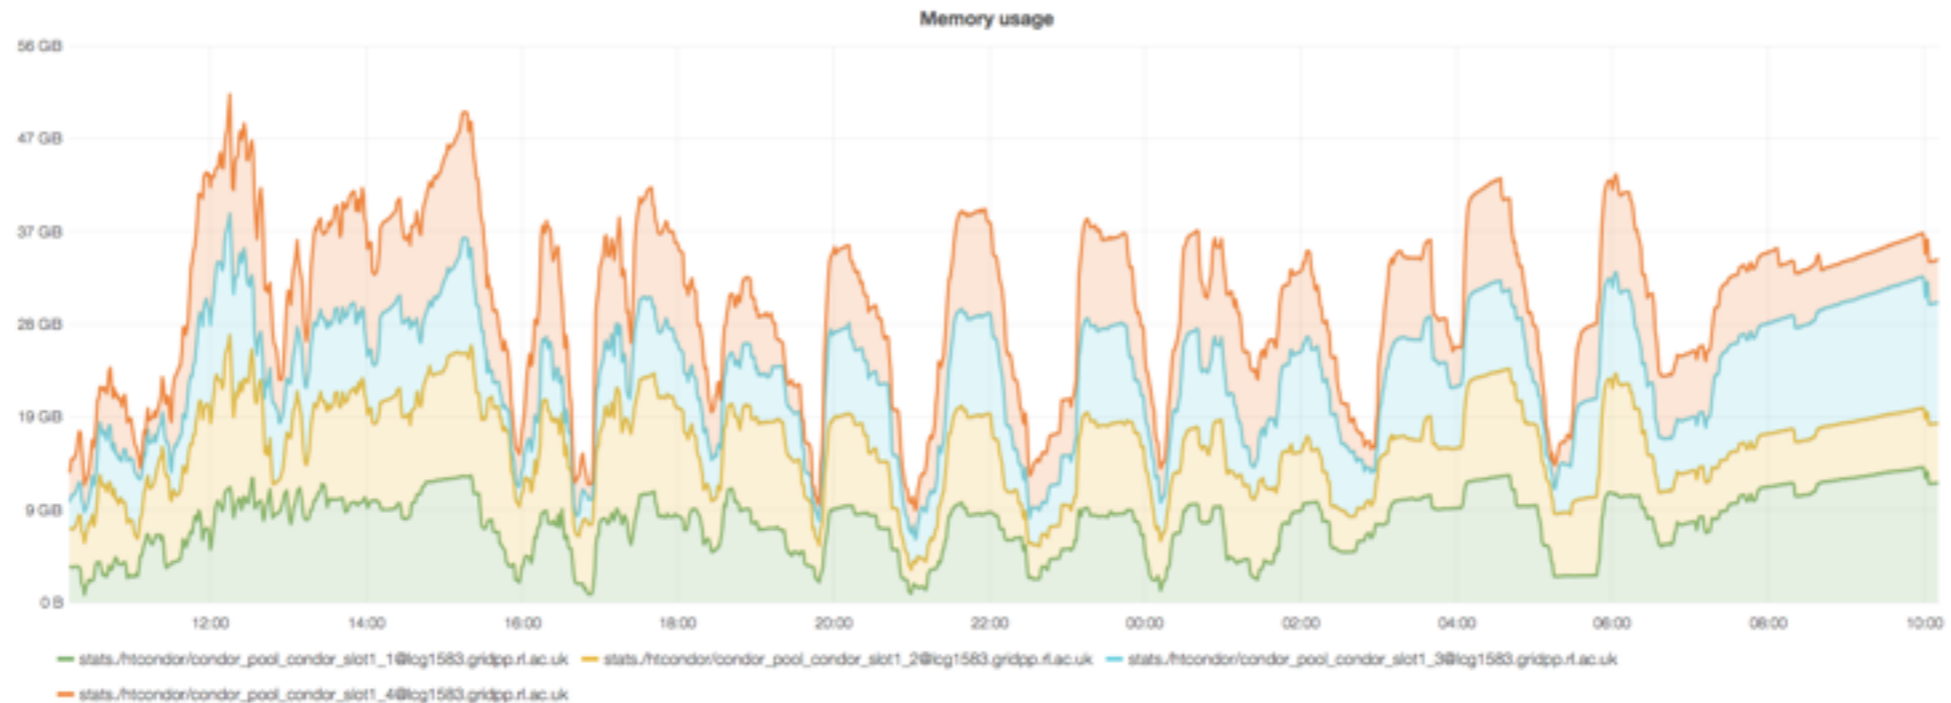
Example: multi-core CMS job



# Grafana

# CPU & memory

4 multi-core CMS jobs on a 32-core WN




126 GB memory available to HTCondor

32 cores available to HTCondor

# Memory

4 multi-core CMS jobs on a 32-core WN

Memory usage




	min	max	avg	current
stats./htcondor/condor_pool_condor_slot1_1@lcg1583.gridpp.rl.ac.uk	307 MiB	27.9 GiB	11.7 GiB	308 MiB
stats./htcondor/condor_pool_condor_slot1_2@lcg1583.gridpp.rl.ac.uk	21 MiB	23.1 GiB	14.2 GiB	9.0 GiB
stats./htcondor/condor_pool_condor_slot1_3@lcg1583.gridpp.rl.ac.uk	18 MiB	31.3 GiB	15.2 GiB	8.3 GiB
stats./htcondor/condor_pool_condor_slot1_4@lcg1583.gridpp.rl.ac.uk	22 MiB	20.1 GiB	12.8 GiB	8.2 GiB

# Memory

4 multi-core CMS jobs on a 32-core WN

Memory usage

**CMS glidein which is stuck & not doing anything**




	min	max	avg	current
stats./htcondor/condor_pool_condor_slot1_1@lcg1583.gridpp.rl.ac.uk	307 MiB	27.9 GiB	11.7 GiB	308 MiB
stats./htcondor/condor_pool_condor_slot1_2@lcg1583.gridpp.rl.ac.uk	21 MiB	23.1 GiB	14.2 GiB	9.0 GiB
stats./htcondor/condor_pool_condor_slot1_3@lcg1583.gridpp.rl.ac.uk	18 MiB	31.3 GiB	15.2 GiB	8.3 GiB
stats./htcondor/condor_pool_condor_slot1_4@lcg1583.gridpp.rl.ac.uk	22 MiB	20.1 GiB	12.8 GiB	8.2 GiB

# Memory

Multiple VOs

Memory usage

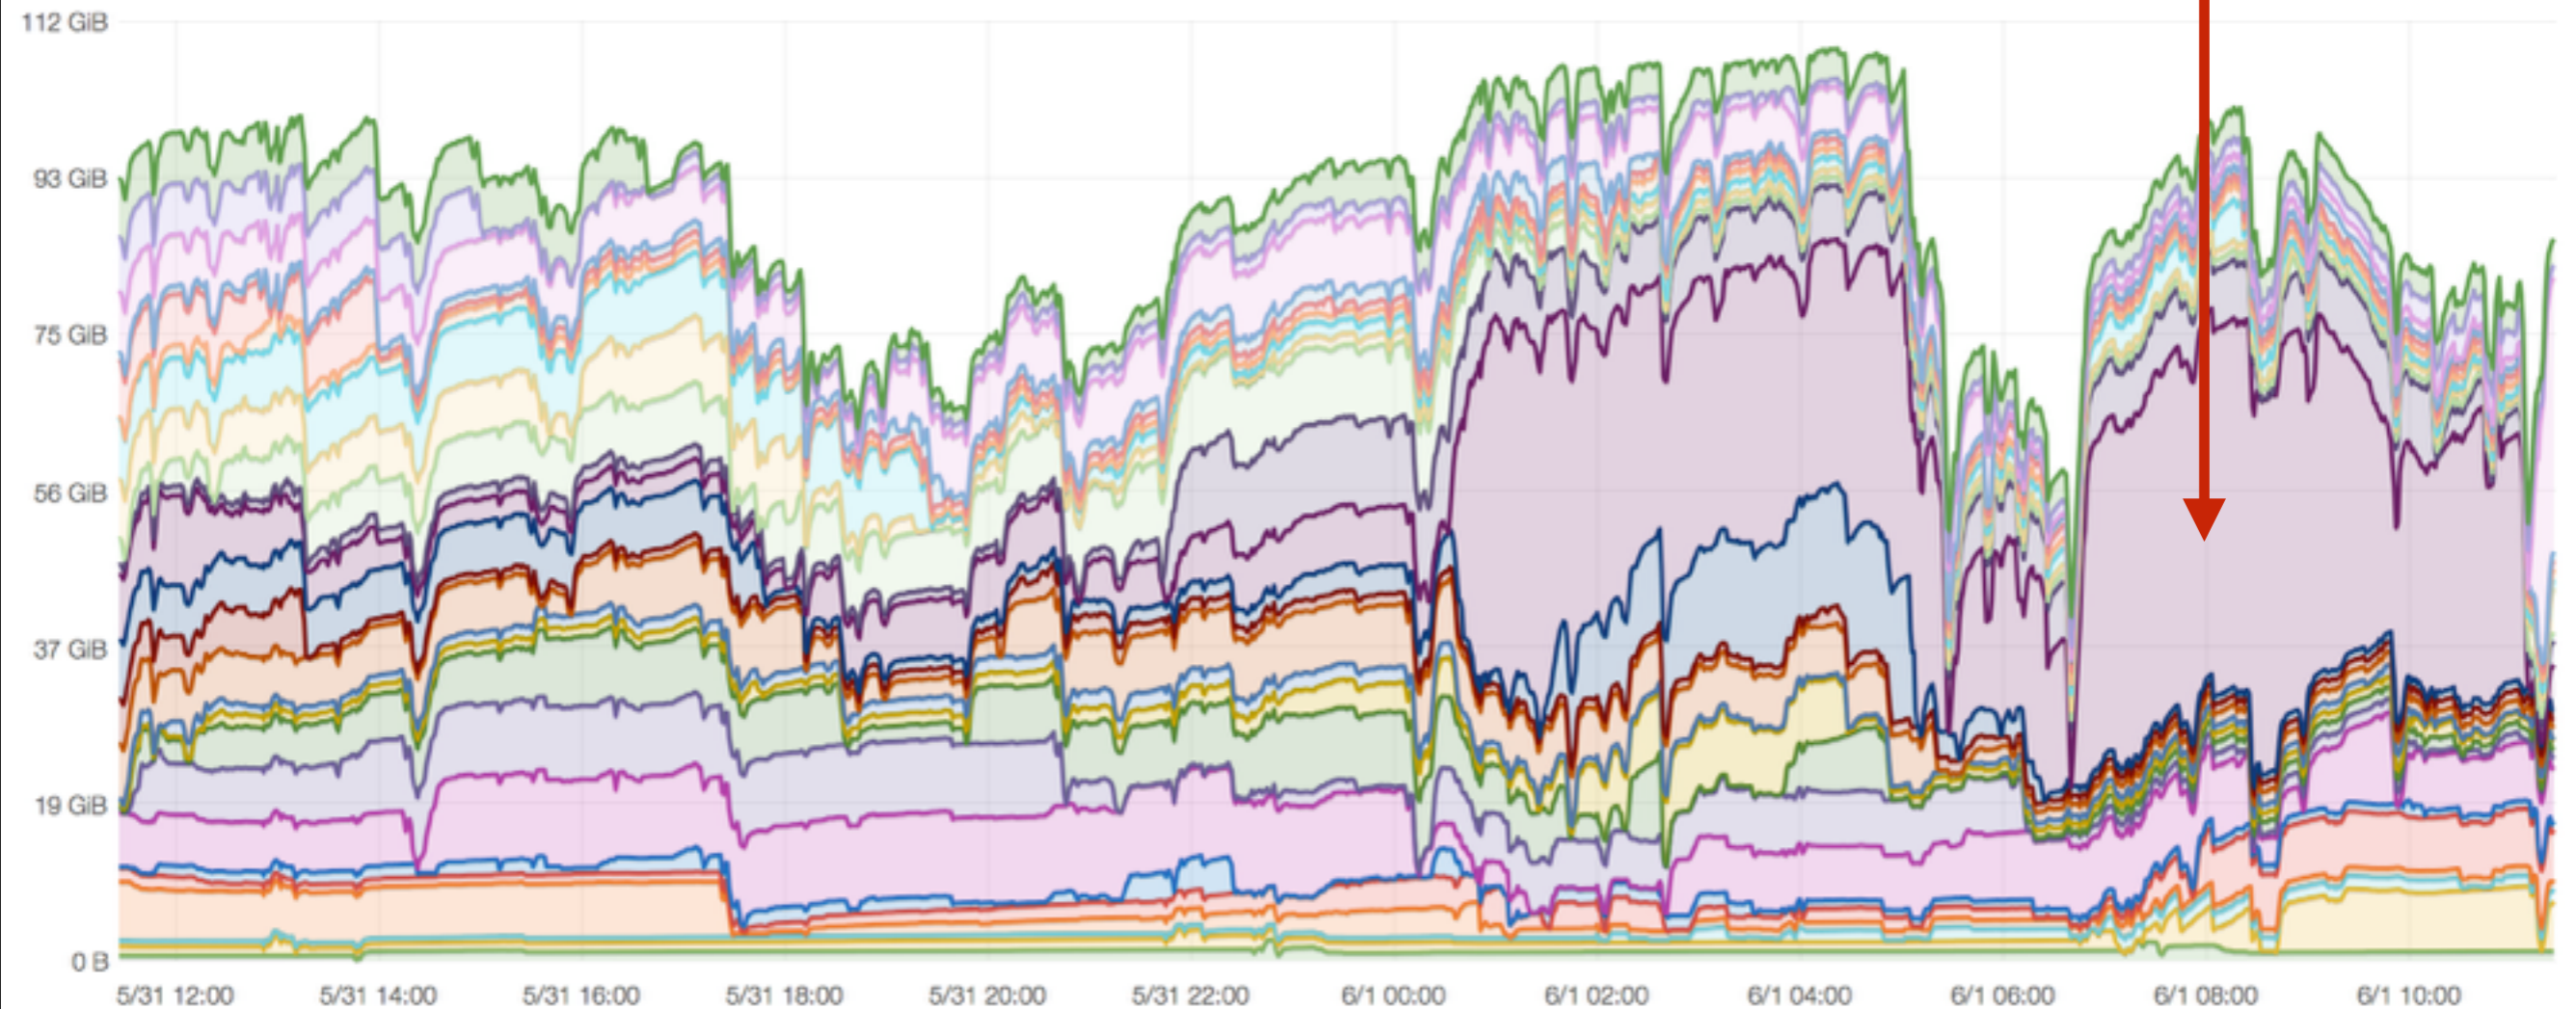


# Memory

Multiple VOs


**ATLAS multicore job using ~ 30GB memory  
(requested 16 GB)**

Memory usage



# Network

- Network throughput per job not available with HTCondor
- Network usage per container is available with e.g. Docker





# Issues

- Not easy to tell when jobs start/finish from Grafana plots - labels are the container names (contain slot numbers, not job ids)
- By default collects data every second - but can change this easily
  - DB grows quickly
  - Quite slow to make plots in Grafana

# Future

- cAdvisor future plans include
  - blkio stats
  - advice on the performance of a container, e.g.
    - when it is being negatively affected by another
    - when it is not receiving the resources it requires