



# Combining collectl and ganglia

Roy Dragseth

Team Leader, HPC

The Computer Center

`roy.dragseth@cc.uit.no`



# Collectl

Developed by Mark Seger (HP).

Does collect a LOT of stuff at the system level:

- cpu

- memory

- net - both ethernet and infiniband

- io - both local and lustre

Can log history to local files.



# Ganglia

Can collect metrics from a lot of nodes.

Very light weight.

Scales to thousands of nodes.

Has a very simple extension interface

- python api

- shell script

The standard monitoring tool in Rocks.



# Combining them

Make collectl dump the system state to a file (can do both sexps and lists)

```
--export lexpr --expdir /tmp
```

will dump to /tmp/L

Get the ganglia collector to read this file and publish the desired metrics.



# The result

Global system views and long term history for everything that collectl can collect.

Fine grained system information per system that enables forensics after a system crash.



# Future development

Enable realtime (or at least more frequent) monitoring on a per job basis. Possibly using the collectl socket interface.

Publish hardware counters through ganglia: Cluster wide gigaflop/second or maybe L2 cache misses per second per job.

Improving the web interface.




## More info

Collectl: <http://collectl.sf.net>

Ganglia: <http://www.ganglia.info>

Rocks: <http://www.rocksclusters.org>



# Questions?