

Thirty Billion Metrics a Day:

Large-Scale Performance Metrics with Ganglia

Adam Compton, Quantcast

acompton@quantcast.com

@comptona

usenix

LISA15

November 8–13, 2015 | Washington, D.C.

www.usenix.org/lisa15

#lisa15

~~Thirty~~ Billion Metrics a Day:

Forty

Adam Compton, Quantcast

acompton@quantcast.com

@comptona

usenix

LISA15

November 8–13, 2015 | Washington, D.C.

www.usenix.org/lisa15

#lisa15

- 7,000,000 unique metrics collected every 15 sec
- from 1000s of machines
- 480,000 metrics/second
- 41,000,000,000 (billion) metrics/day

How many machines do we need for this?

Two.

Who am I?




quantcast


- **What is Ganglia?**
- How Ganglia works
- Inputs
- Outputs
- Scaling
- What's next

What is Ganglia?

“Ganglia is a scalable distributed monitoring system for high-performance computing systems”



What does
“monitoring” mean?



What is Ganglia - monitoring vs. metrics

performance metrics: regular, numeric, time-series data

monitoring: scheduled checks (of metrics, processes, endpoints, etc.) to identify anomalous behavior

alerting: getting a human's attention

visualization: charts and graphs


By these definitions,
Ganglia **collects, aggregates, and visualizes**
performance metrics.

What is Ganglia not?


- No monitoring
- No alerting
- No asset database
- No Julianne fries

- What is Ganglia?
- **How Ganglia works**
- Inputs
- Outputs
- Scaling
- What's next


How Ganglia works




How Ganglia works - gmond




How Ganglia works - gmetad



(aside) How RRDs work



How Ganglia works - gweb



How Ganglia works - Upsides

Built-in hierarchy, summarization, and metadata

Web UI is easy to browse and grok

How Ganglia works - Downsides

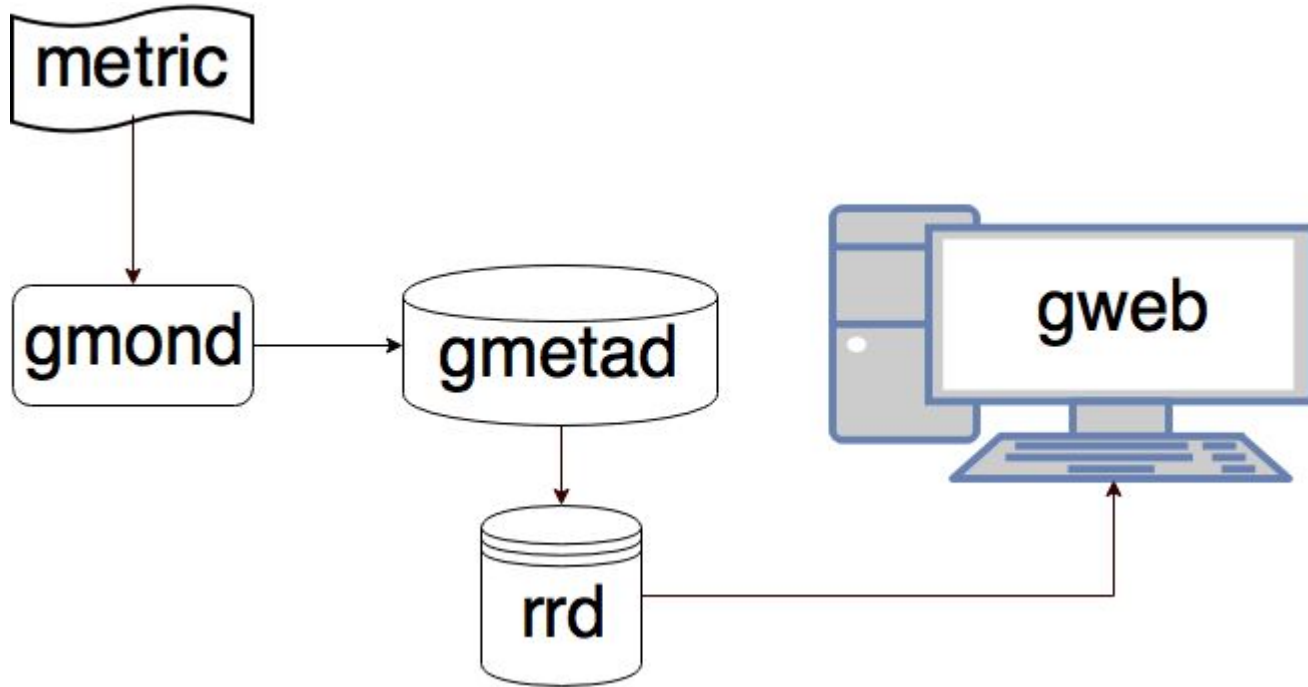
Downsampling loses data :'(

Custom visualizations are hard



- What is Ganglia?
- How Ganglia works
- **Inputs**
- Outputs
- Scaling
- What's next

Life of a Metric



Inputs - How do we feed it?

- gmond
- gmetrics
- snmp2ganglia
- json2gmetrics

Inputs - json2gmetrics

Translates a structured JSON block
into metrics

```
'[{ "name": "mood", "value": "happy", "type": "string"},  
  { "name": "sad_people", "value": 0, "type": "uint8" }  
' | json2gmetrics
```

Why so many inputs?

Less work to
report = more
metrics

More metrics =
more useful



Classes of metrics

- High-level operating system health
- In-depth system performance
- Infrastructure metrics
- Application performance
- String “metrics”
- Anything!

Classes of metrics - Monitoring checks

Privileged checks run locally and report to Ganglia




Monitoring interprets metrics as check results




- What is Ganglia?
- How Ganglia works
- Inputs
- **Outputs**
- Scaling
- What's next


Visualizations - trendlines




Visualizations - CPU/Memory per-user




Visualizations - Red background for failures



Visualizations - running variance + bounds



Visualizations - Holt Winters




Monitoring


Ganglia metrics are a huge component (> 50%) of our monitoring world

- Simple, static thresholds
- Historical averages
- Smart “raw data” graphs


Monitoring - Static Thresholds



Monitoring - Historical Averages



Monitoring - Smart Graphs



Tracking Deployments

Before deployment, identify
key metrics and expected
change

After deployment, verify the
changes are as expected



Troubleshooting

- All metrics for a host at once
- Same metric across all hosts
- Easy to make aggregate graphs



Capacity Planning

- Ganglia can store long term data (years)
- Simple to see trends over long spans of time




- What is Ganglia?
- How Ganglia works
- Inputs
- Outputs
- **Scaling**
- What's next

Scaling

Performance metrics must be:

- Realtime
- Highly-available
- Actionable
- Comprehensive
- “At scale”




Scaling

“at scale”:

functionality dictated by business
requirements, not the other way around

Scaling Ganglia - sharding



Scaling Ganglia - gmond

- gmetad -> every gmond? Too slow
- gmetad -> master gmond? SPoF
- gmetad -> multicast gmond? Just right

Scaling Ganglia - tmpfs and mmap

- Metrics -> tmpfs
- tmpfs -> disk
- disk -> tmpfs (on reboot)

Note: Turn off mmap in RRDtool!

Scaling Ganglia - patches

Several other bottlenecks in the Ganglia code

Some fixes [have been published](#), others not yet

- What is Ganglia?
- How Ganglia works
- Inputs
- Outputs
- Scaling
- **What's next**




What's next - Full resolution


High-resolution metrics downsample too quickly

Why not store all data forever?

What's next - Visualizations



VS.



What's next - still Ganglia

Why not move away from Ganglia entirely?



Takeaways

- Ganglia is powerful, reliable, and scales incredibly well
- Don't let your metrics system decide what data you can collect
- You are probably ignoring useful metrics right now

Thank you!

acompton@quantcast.com
@comptona

Image Sources

Slide 7: <http://shop.oreilly.com/product/0636920025573.do>

Slide 16: <https://www.flickr.com/photos/internetarchivebookimages/14782490515/>

Slide 24: <https://www.flickr.com/photos/thomashawk/9135191382/>

Slide 26: <https://pixabay.com/en/approved-button-check-green-round-151676/>

Slide 37: <https://www.flickr.com/photos/naturesdawn/2758645685/>

Slide 38: <https://www.flickr.com/photos/kristiand/3223044657/>

Slide 39: <https://www.flickr.com/photos/kajisagook/3640015021/>

Slide 42: [https://commons.wikimedia.org/wiki/File:Scale_\(PSF\).png](https://commons.wikimedia.org/wiki/File:Scale_(PSF).png)

Slide 52: <https://commons.wikimedia.org/wiki/File:EndOfRoadSign-Margate.jpg>