


Thomas Stocking
Director of Systems Engineering
GroundWork Open Source, Inc.



Monitoring Automation and Analysis for DevOps

GroundWork Monitor and Big Data



- **Monitoring data and its sources**
- **How industry trends affect the data we collect**
- **Cautions and concerns**
- **Reasons for optimism**
- **Where our customers are**
- **Where we see solutions**
- **GroundWork features and futures**
- **Q & A**



History

ICMP is as old as IP

SNMP is a large body of work and is still the de-facto standard

Using common tools in uncommon ways we have been gathering data for a long time

But things are changing...



Virtualization

- One machine now hosts many
- No such thing as a static configuration
- Everything virtual has to be automated, including monitoring.
- Orchestration layers have to be monitored too



Cloud Transitions

> 60% of businesses in the USA have some level of cloud adoption*

This trend is driven by the need for cost savings and access to innovative technologies

You can't manage what you can't monitor

Having good data is key to future success.

- source:<https://www.forbes.com/sites/louiscolombus/2017/04/23/2017-state-of-cloud-adoption-and-security/#63f594b51848>



Some things to be concerned about with this trend:

Data is expanding exponentially

- Data privacy and security are growing issues

- Cloud reliability is better, but not by all that much

Data standardization and integrity are still lacking



$$P = 1/D$$

Or

**Way more data on all of us
means
Way less privacy for any of us**

And

Private data has better integrity than compromised data

The GDPR is coming, which will help. And hurt.



Speaking of data standardization:

A lot of monitoring metrics ignore the underlying complexity, and need interpretation.

Example: %CPU

Also related:

vCPU measures from the OS or Hypervisor perspective

Does Load Factor give you a better measure?

What does that metric mean?



Brendan Gregg's Blog [home](#)

CPU Utilization is Wrong


09 May 2017

The metric we all use for CPU utilization is deeply misleading, and getting worse every year. What is CPU utilization? How busy your processors are? No, that's not what it measures. Yes, I'm talking about the "%CPU" metric used *everywhere*, by *everyone*. In every performance monitoring product. In top(1).

What you may think 90% CPU utilization means:



What it might really mean:



Source: <http://www.brendangregg.com/blog/2017-05-09/cpu-utilization-is-wrong.html>

What does that metric mean?



```
# perf stat -a -- sleep 10
```

```
Performance counter stats for 'system wide':
```

```
 641398.723351    task-clock (msec)    #    64.116 CPUs utilized    (100.00%)
      379,651      context-switches    #    0.592 K/sec            (100.00%)
      51,546      cpu-migrations      #    0.080 K/sec            (100.00%)
 13,423,039      page-faults         #    0.021 M/sec            (100.00%)
1,433,972,173,374 cycles                #    2.236 GHz              (75.02%)
<not supported> stalled-cycles-frontend
<not supported> stalled-cycles-backend
1,118,336,816,068 instructions          #    0.78  insns per cycle    (75.01%)
 249,644,142,804 branches                # 389.218 M/sec            (75.01%)
 7,791,449,769   branch-misses          #   3.12% of all branches    (75.01%)

10.003794539 seconds time elapsed
```

IPC = 0.78
CPU Width = 4 } 19%

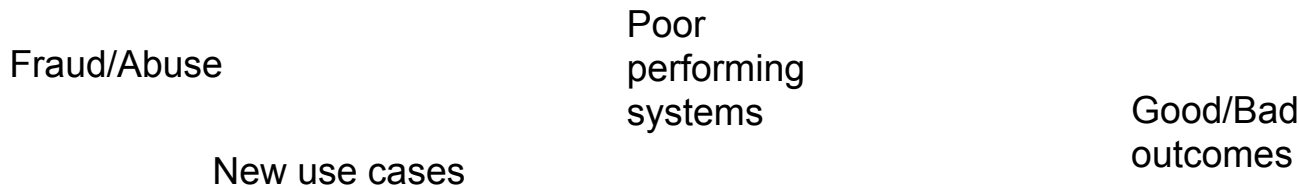


Good data allows good analysis

Lots of businesses integrating data to detect patterns of use and abuse

Simple example: sentiment analysis with Cloudera*

$$S = (\text{positive} - \text{negative}) / (\text{positive} + \text{negative})$$



Source: https://www.cloudera.com/documentation/other/tutorial/CDH5/topics/ht_example_4_sentiment_analysis.html

What can we learn from this?



Our companies and our customers are in transition

There are a lot of people with one foot in the data center and one in the cloud.
Monitoring has to cover both.





Work with standards

Conventions rule

DNS works why? Because it is standard

Name your metrics well

Know your metrics

Publish your rules

Remember documentation? The wiki is your friend.



Megabytes?

What is a megabyte? It depends...

Decimal: 1000k

Hexadecimal: 1024k

Make no assumptions!



Data storage methods are changing..

Optimize on write

- Not a lot of data
- Not a lot of storage
- Just save the data you need

Optimize on read

- Store a lot of data
- Store it really fast
- Check the details later



Version 7 (2014)

- Log Bridge
- Cloud Hub

New Release: 7.2.0 (2017)

- GrafBridge: Grafana integration with influxDB
- Grafana dashboards
- Grafana graphs in Status viewer
- Cloudera Connector for Cloud Hub

What's GroundWork doing about this?



Cloud Hub

Add a Connection



Docker

+ ADD

openstack™

OpenStack

+ ADD



Red Hat Enterprise
Virtualization

+ ADD

vmware®

VMware vSphere

+ ADD



Amazon AWS

+ ADD

NetApp

NetApp

+ ADD

cloudera®
Ask Bigger Questions

Cloudera

+ ADD

icinga

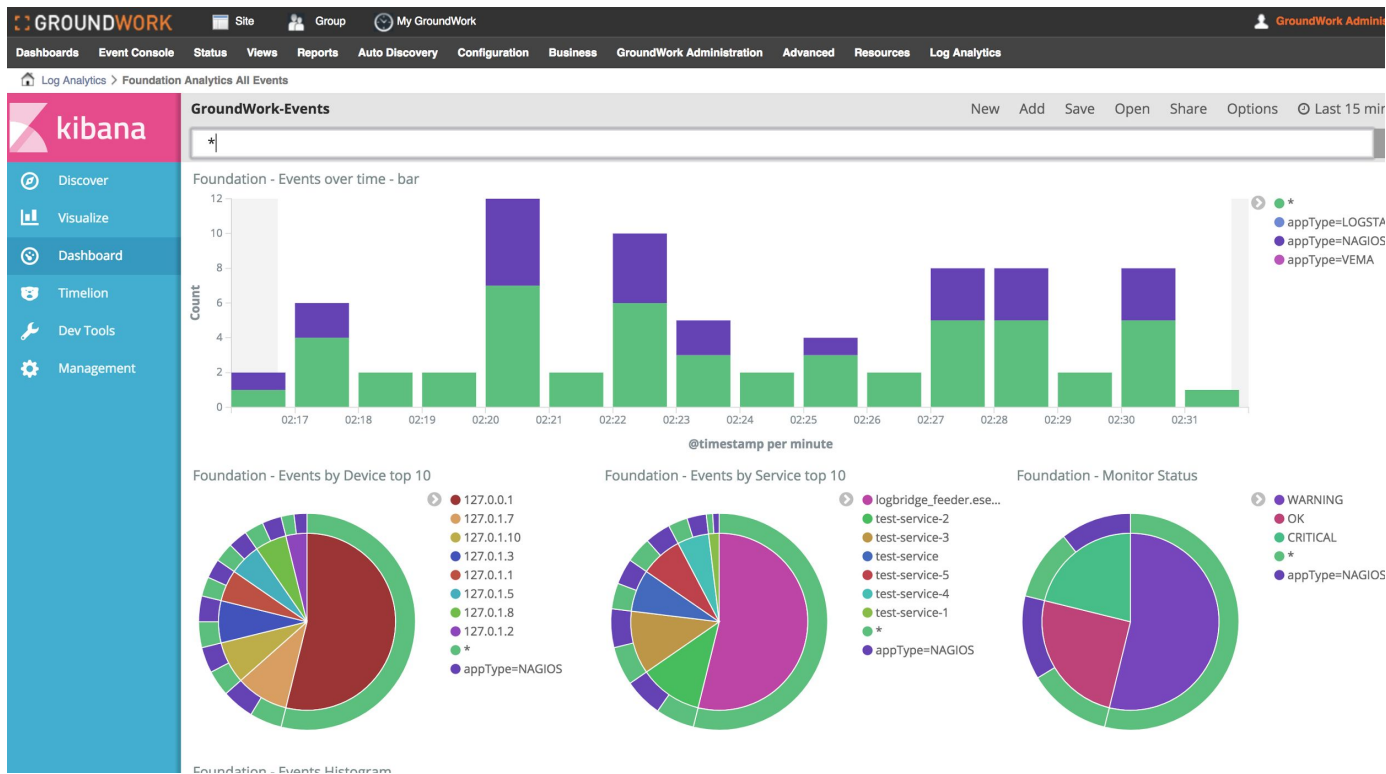
Icinga 2 Monitoring

+ ADD

What's GroundWork doing about this?



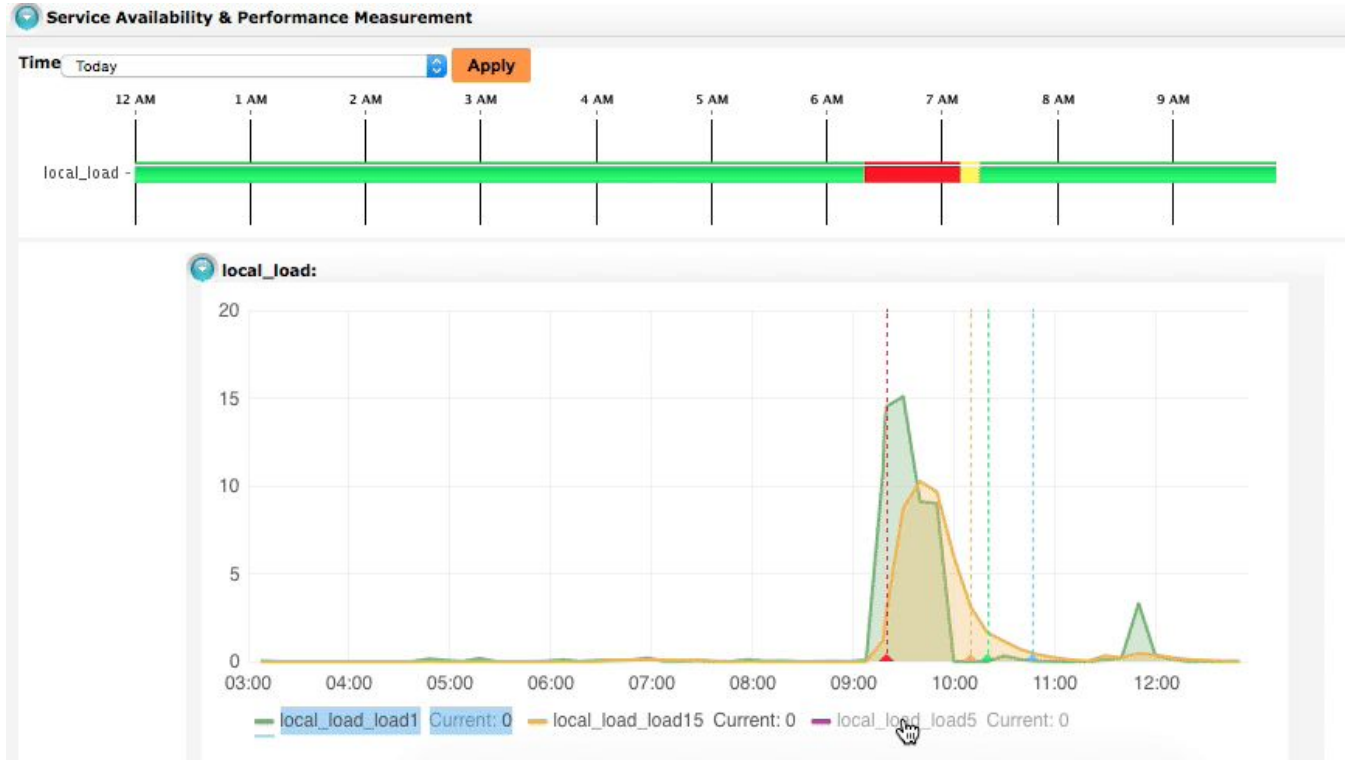
Log Bridge



What's GroundWork doing about this?



GrafBridge



Other updates in 7.2.0



- Enhanced performance of the foundation database at scale
- Upgrades to underlying packages like PHP, Nagios, Postgres
- Diagnostic dashboards for GroundWork itself (influxdb data source)



GroundWork plans for 7.2 Enhancements to take advantage of new technologies:

In the cloud:

- Additional Cloud Hub connectors
 - Azure
 - Google Cloud
 - InfluxDB
- Enhancements to current connectors
 - VMware
 - AWS

On the box:

- Agent integrations
 - Beats
 - Telegraf
 - User-configurable alerts and thresholds
- Agent updates
 - GDMA for windows update
 - GDMA service auto-discovery




From our partner RealStuff:

- SLA Widget integration (RStools)
 - User-configurable dashboards
- Report server
 - Jasper integration



From our Partner NeDi:

- Asset management with NoDi



What's on the horizon?



GroundWork 8.0

Design considerations:

Data centric - we will collect data, store and analyze for operations

Unified Monitoring Infrastructure will be included:

- Alerting
- Reporting
- Dashboards
- Analysis

Modular

- We will bundle less, integrate more





Your questions?

Contact us:

sales@gwos.com

<https://www.gwos.com>

Download Trial:

<https://www.gwos.com/download>

Thomas Stocking

tstocking@gwos.com