



# Spectrum Scale Performance Tools Deployment at Nuance

Bob Oesterlin, Sr Storage Engineer  
Nuance HPC Grid  
June 10th, 2016

[Robert.Oesterlin@nuance.com](mailto:Robert.Oesterlin@nuance.com)

---

# Topics

- Quick Overview of Nuance and HPC Grids
- Performance GUI – Experiences and Limitations
- Spectrum Scale Performance Tools – Deployment
- Collector Sizing/Federation
- Dashboards using the Zimon-Grafana Bridge
- What's next

---

# Reinventing the relationship between people and technology

- Defining the next generation of human-computer interaction:  
**Intelligent Systems**
- Deeply invested in creating effortless and natural user experiences
- Best known for rapidly advancing voice-recognition technology

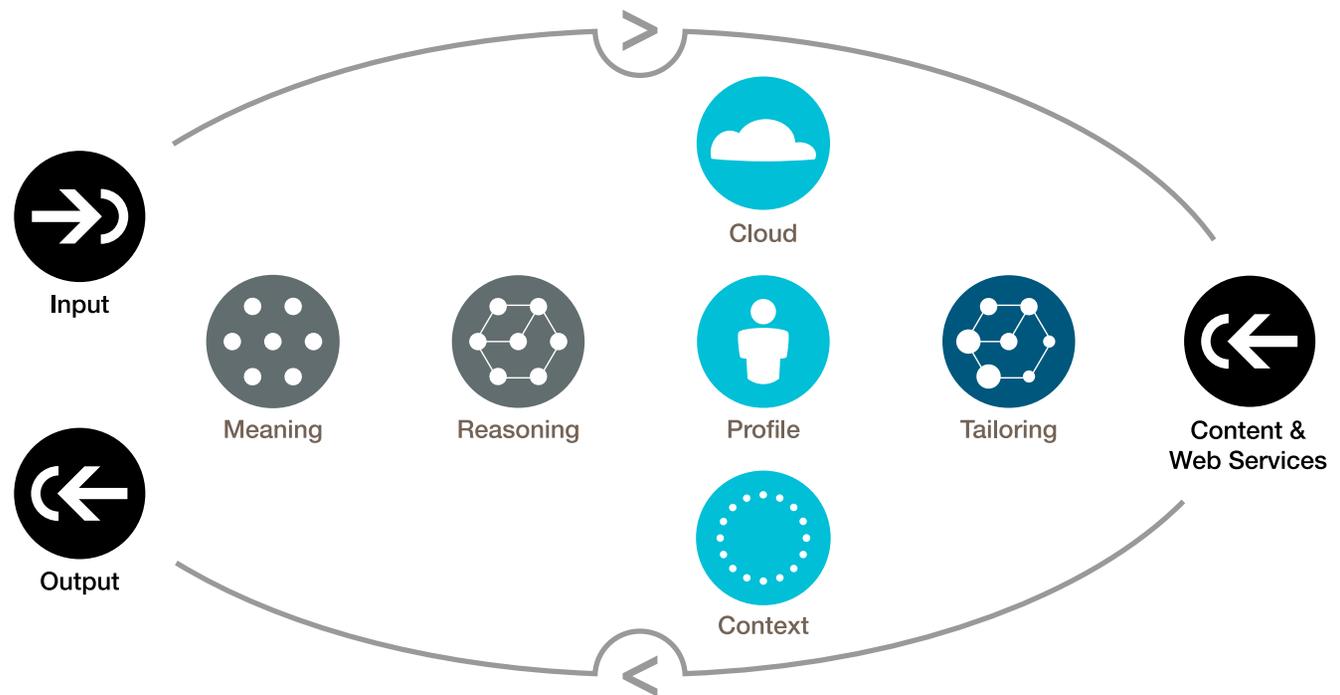


# Nuance Natural Language Framework

The engine that drives Intelligent Systems

“Anything with George Clooney on tonight?”

“Yes, I’ve found three shows, one of which is starting in just a few minutes.”



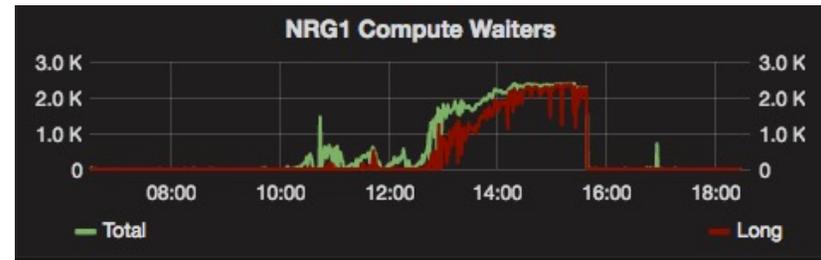
---

# Nuance HPC Grids

- Supports the Worldwide Nuance R&D Community
- Approximately 2000 users
- ~7500 TB per day of data processed
  - 85% Read, 15% Write
- ~20,000,000 jobs processed per month
- Over 6 PB of Elastic Storage across multiple clusters
- On-premise Object storage, 4+ PB
- VMs for casual access/job submission

# Performance data collection - legacy

- Large number of locally written tools
- Collectl for system stats (CPU, disk, network, etc)
- Periodic mmpmon collections feeds local database
- Scripts to track RPC waiters
- Dashboards based around Grafana



---

# SS Performance Sensors (aka “zimon”)

- Part of all releases since 4.1.1
- Integrated with Spectrum Scale
- Wide variety of metrics, both system and GPFS
- New metrics being added (RPC waiters)
- Integrates with Spectrum Control

---

# IBM Performance GUI

- Provides access to all key SS performance metrics
- Early beta participant
- “Fairly” easy deployment
- RH 7 dependency proved to be a challenge; current grids are all RH 6.6 based
- Table provide good overview of overall performance
- Better for my Ops team than Engineering
- Graphs – problematic in large clusters

# IBM Performance GUI



---

# Sensor Deployment - Problems

- Using the default sensor configuration in large cluster is a bad idea
- Deployment with federated (multiple) collectors
- Which sensors drive the GUI?

---

# Collector Sizing and deployment

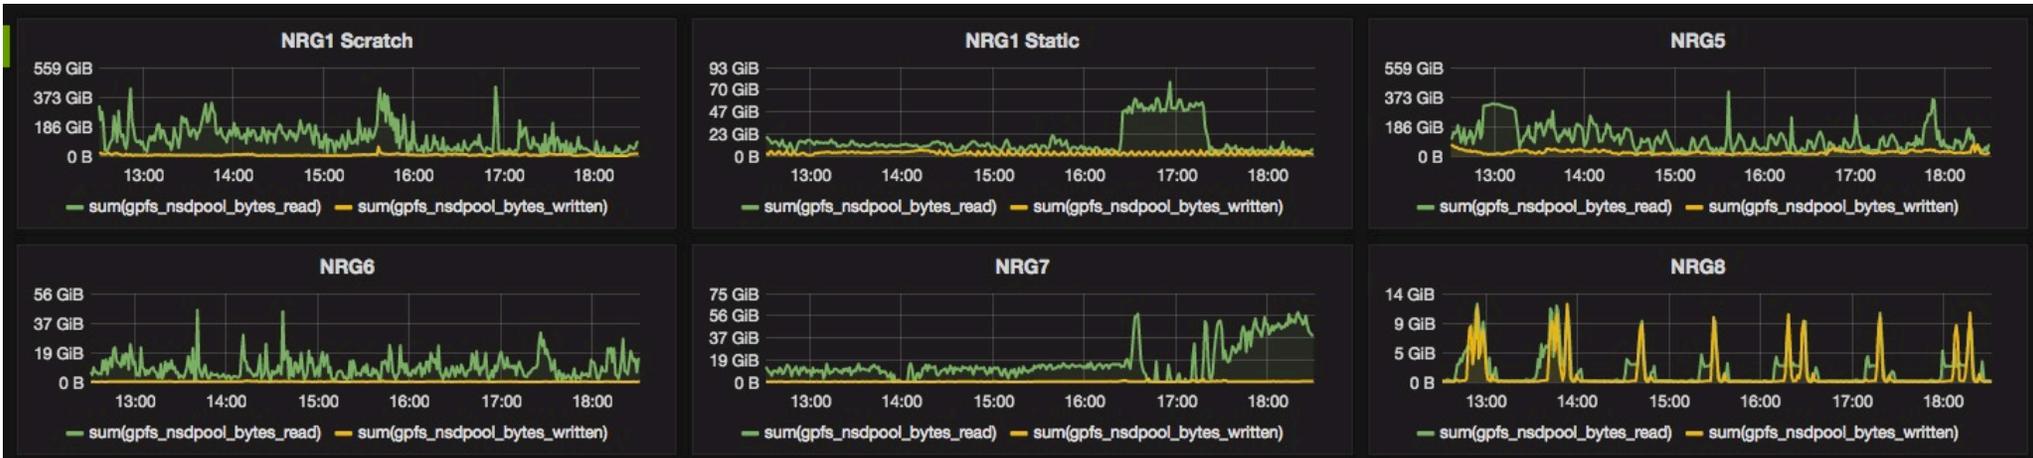
- Default configuration is perfect for small environments
- Collector memory requirements grow quickly
- Difficult to retain large numbers of frequently collected metrics
- Keep an eye on scaling:
  - 500 NSDs \* 500 nodes \* 16 metrics = 4 million!
- Example:
  - 500 nodes, 500 NSDs, 16 file systems, 7 days of 1/min data = 66GB collector memory

---

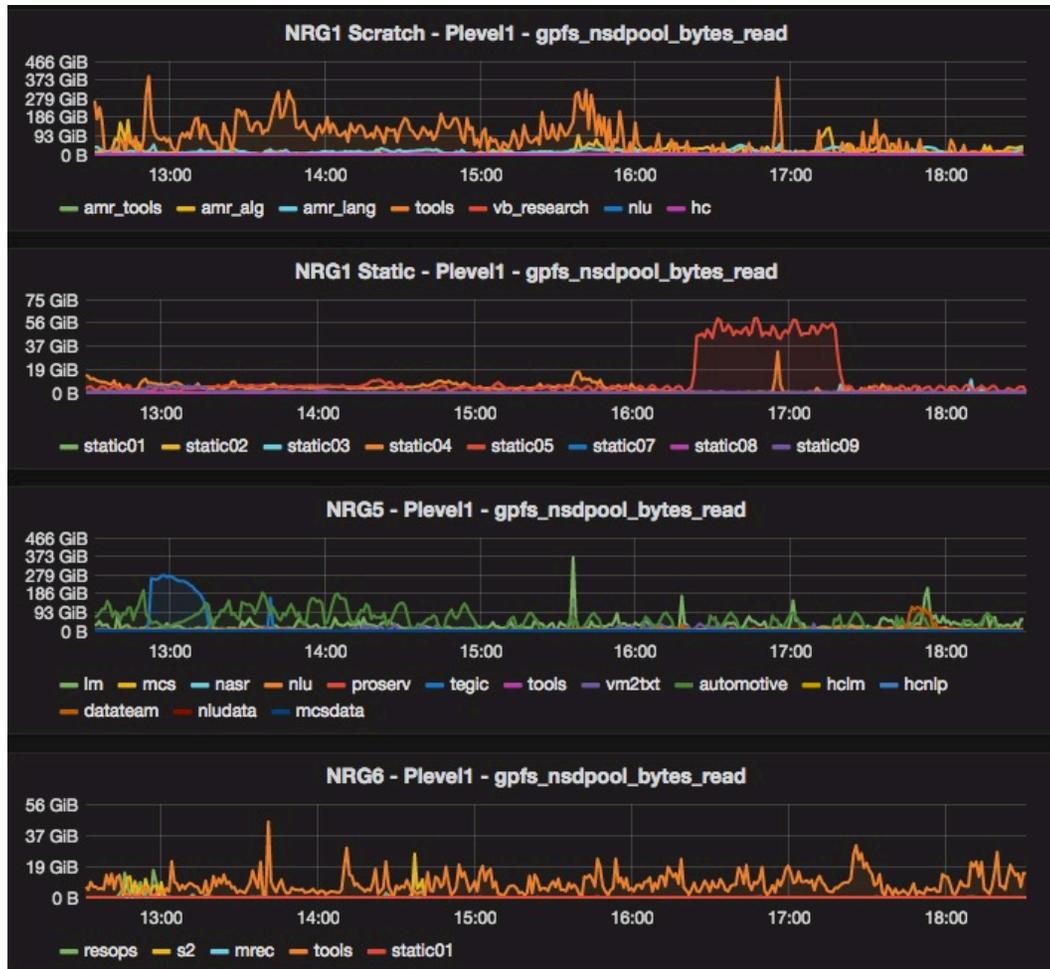
# Grafana with Zimon

- IBM GUI is a great start – but limited, especially on larger clusters
- Zimon-grafana bridge code by Metin Feridun @ IBM ZRL
  - Provides Open TSDB Interface to IBM zimon data
  - Simple python script, runs on collector node, lightweight
  - All collected zimon metrics are available
  - Easy to construct complex/custom dashboards
- Distribution...
  - IBM Developerworks?

# Sample Grafana Dashboards



# Sample Grafana Dashboards



# Sample Grafana Dashboards



---

# What's Next

- SS 4.2.1 Upgrade
  - RPC Waiter metrics in zimon
  - Cloud Tiering
- Consolidation of Grids
  - Combine Compute/NSD Clusters
  - Consistent deployment architecture
- Move from CNFS to CES



**Thank you**