

Computer Monitoring with Prometheus & Grafana

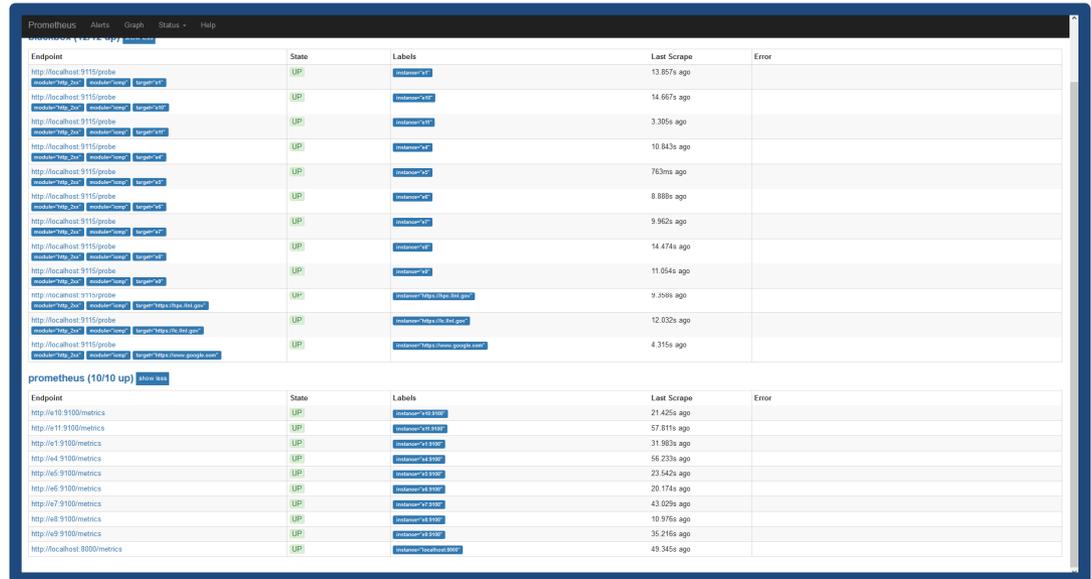
Brad Davis, Enikoe Bihari, Eliana Purcell
HPC Academy 2018

August 14, 2018



Prometheus

- Computer monitoring database software that stores data received from various exporters



Endpoint	State	Labels	Last Scrape	Error
http://localhost:9115/probe	UP	instance="1"	13.857s ago	
http://localhost:9115/probe	UP	instance="2"	14.667s ago	
http://localhost:9115/probe	UP	instance="3"	3.305s ago	
http://localhost:9115/probe	UP	instance="4"	10.843s ago	
http://localhost:9115/probe	UP	instance="5"	763ms ago	
http://localhost:9115/probe	UP	instance="6"	8.888s ago	
http://localhost:9115/probe	UP	instance="7"	9.902s ago	
http://localhost:9115/probe	UP	instance="8"	14.474s ago	
http://localhost:9115/probe	UP	instance="9"	11.054s ago	
http://localhost:9115/probe	UP	instance="10"	9.398s ago	
http://localhost:9115/probe	UP	instance="11"	12.032s ago	
http://localhost:9115/probe	UP	instance="12"	4.315s ago	

Endpoint	State	Labels	Last Scrape	Error
prometheus (10/10 up)				
http://10.9100/metrics	UP	instance="10.9100"	21.455s ago	
http://11.9100/metrics	UP	instance="11.9100"	57.911s ago	
http://1.9100/metrics	UP	instance="1.9100"	31.983s ago	
http://4.9100/metrics	UP	instance="4.9100"	56.273s ago	
http://6.9100/metrics	UP	instance="6.9100"	23.542s ago	
http://6.9100/metrics	UP	instance="6.9100"	20.174s ago	
http://67.9100/metrics	UP	instance="67.9100"	43.029s ago	
http://A8.9100/metrics	UP	instance="A8.9100"	10.976s ago	
http://A9.9100/metrics	UP	instance="A9.9100"	35.216s ago	
http://localhost:9000/metrics	UP	instance="localhost:9000"	49.345s ago	

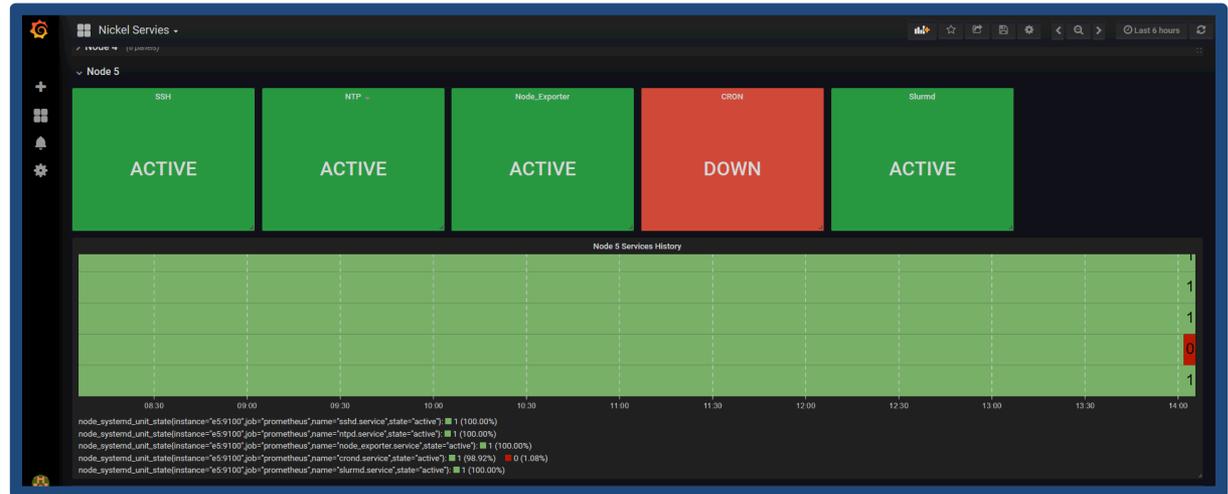


Prometheus Configuration

```
1 global:
2   scrape_interval: 15s
3
4 rule_files:
5   - alert.rules.yml
6
7 alerting:
8   alertmanagers:
9     - static_configs:
10       - targets:
11         - localhost:9093
12
13 scrape_configs:
14   - job_name: 'prometheus'
15     scrape_interval: 60s
16     scrape_timeout: 30s
17     static_configs:
18       - targets: ['localhost:8000', 'e1:9100', 'e4:9100', 'e5:9100', 'e6:9100', 'e7:9100', 'e8:9100', 'e9:9100', 'e10:9100', 'e11:9100']
19
20   - job_name: 'blackbox'
21     metrics_path: /probe
22     params:
23       module: [http_2xx, icmp]
24     static_configs:
25       - targets:
26         - https://www.google.com
27         - https://hpc.llnl.gov
28         - https://lc.llnl.gov
29         - e1
30         - e4
31         - e5
32         - e6
33         - e7
34         - e8
35         - e9
36         - e10
37         - e11
38     relabel_configs:
39       - source_labels: [__address__]
40         target_label: __param_target
41       - source_labels: [__param_target]
42         target_label: instance
43       - target_label: __address__
44         replacement: localhost:9115
45
46 ~
```

Exporters

- Scrape various data from the nodes
 - Node Exporter
 - IPMI Exporter
 - Blackbox Exporter



Grafana



- Grafana is a versatile UI
- Displays the data from Prometheus on dashboards



Main Dashboard

The screenshot shows a dashboard titled "Main Overview" with a dark theme. At the top, it displays the date "07-31-2018" and the time "4:09:52 PM". Below this, there is a "Cluster Status" section with 12 green boxes, each labeled "UP". Underneath is an "External Ping" section with three green boxes labeled "UP". At the bottom, there is a section for "HPC Clusters" with a list of clusters: #Cluster X, #Cluster Y, #Cluster Z, and Nickel Cluster, each with a star icon to its right.

The screenshot shows a dashboard titled "Cluster Nodes Ready" with a dark theme. It features a "Nickel Cluster" section with 11 green boxes, each labeled "UP". Below this, there are sections for "Radon Cluster" and "Boron Cluster", both labeled "(0 panels)".

Queries

RECEIVED
SENT

Graph General **Metrics** Axes Legend Display Alert Time range

Data Source: Nickel Prometheus

A `sum(irrate(node_network_receive_bytes_total{instance=~"$node"}[5m]))`
Legend format: RECEIVED Min step: 15s Resolution: 1/2 Format as: Time series Instant

B `- sum(irrate(node_network_transmit_bytes_total{instance=~"$node"}[5m]))`
Legend format: SENT Min step: 15s Resolution: 1/2 Format as: Time series Instant

C Add Query

Nickel RAM

Node: None

E1 Free RAM -

62.7 GB

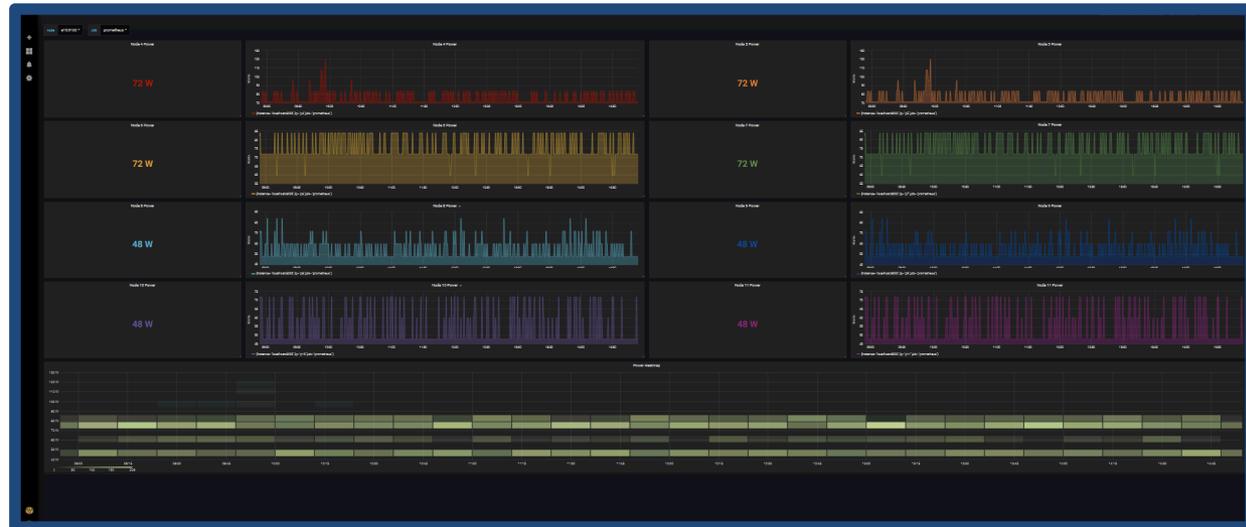
Singlestat General **Metrics** Options Value Mappings Time range

Data Source: default

A `node_memory_free_bytes{instance="*19180*"}[5m]`
Legend format: Legend format Min step: 5m Resolution: 1/1 Format as: Time series Instant

Add Query

Drilldowns



Alertmanager Configuration

```
1 global:
2   smtp_smarthost: 'localhost:25'
3   smtp_from: 'alertmanager@smtp.llnl.gov'
4   smtp_require_tls: false
5 route:
6   group_by: ['alertname', 'instance', 'severity']
7   group_wait: 30s
8   group_interval: 5m
9   repeat_interval: 3h
10  receiver: nickel
11
12 routes:
13   - match:
14     severity: critical
15     receiver: TeslaMan
16
17 inhibit_rules:
18 - source_match:
19   severity: 'critical'
20 target_match:
21   severity: 'page'
22
23 receivers:
24 - name: 'nickel'
25   email_configs:
26     - to: 'bihari2@llnl.gov, davis282@llnl.gov, purcell8@llnl.gov'
27
28 - name: 'TeslaMan'
29   email_configs:
30     - to: 'dixon30@llnl.gov'
31
32
```

```
41 rule_files:
40 - alert.rules.yml
39
38 alerting:
37 alertmanagers:
36 - static_configs:
35   - targets:
34     - localhost:9093
```

- Alertmanager groups and sends alerts when certain metrics reach a threshold
- Sends alerts over email and Slack to different groups depending on the type of alert

```
1 groups:
2 - name: RAM
3   rules:
4 - alert: low_ram
5   expr: sum(node_memory_MemFree_bytes) / sum(node_memory_MemTotal_bytes) < 0.20
6   for: 1m
7   labels:
8     severity: "page"
9   annotations:
10    summary: "RAM usage has surpassed 80%"
11
12 - alert: very_low_ram
13   expr: sum(node_memory_MemFree_bytes) / sum(node_memory_MemTotal_bytes) < 0.05
14   for: 1m
15   labels:
16     severity: "critical"
17   annotations:
18    summary: "RAM usage has surpassed 95%"
```

Alert Examples

Prometheus Alerts Graph Status Help

Alerts

Show annotations

low_ram (0 active)

```
alert: low_ram
expr: sum(node_memory_MemFree_bytes)
    / sum(node_memory_MemTotal_bytes) < 0.2
for: 1m
labels:
  severity: page
annotations:
  summary: RAM usage has surpassed 80%
```

very_low_ram (0 active)

```
alert: very_low_ram
expr: sum(node_memory_MemFree_bytes)
    / sum(node_memory_MemTotal_bytes) < 0.05
for: 1m
labels:
  severity: critical
annotations:
  summary: RAM usage has surpassed 95%
```

Mon 7/30/2018 2:56 PM

AlertManager <alertmanager@smtp.llnl.gov>

[FIRING:1] low_ram page

To: Bihari, Enikoe; Davis, Bradley Taylor; purcell8@llnl.gov.localdomain

If there are problems with how this message is displayed, click here to view it in a web browser.

1 alert for alertname=low_ram severity=page

[View in AlertManager](#)

[1] Firing

Labels
alertname = low_ram
severity = page

Annotations
summary = RAM usage has surpassed 80%

[Source](#)

Sent by **AlertManager**

Scalability With Ansible

- Automated the installation and configuration of Prometheus, Grafana, etc. with Ansible
- Allows for scalability for future use



Questions?





Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.