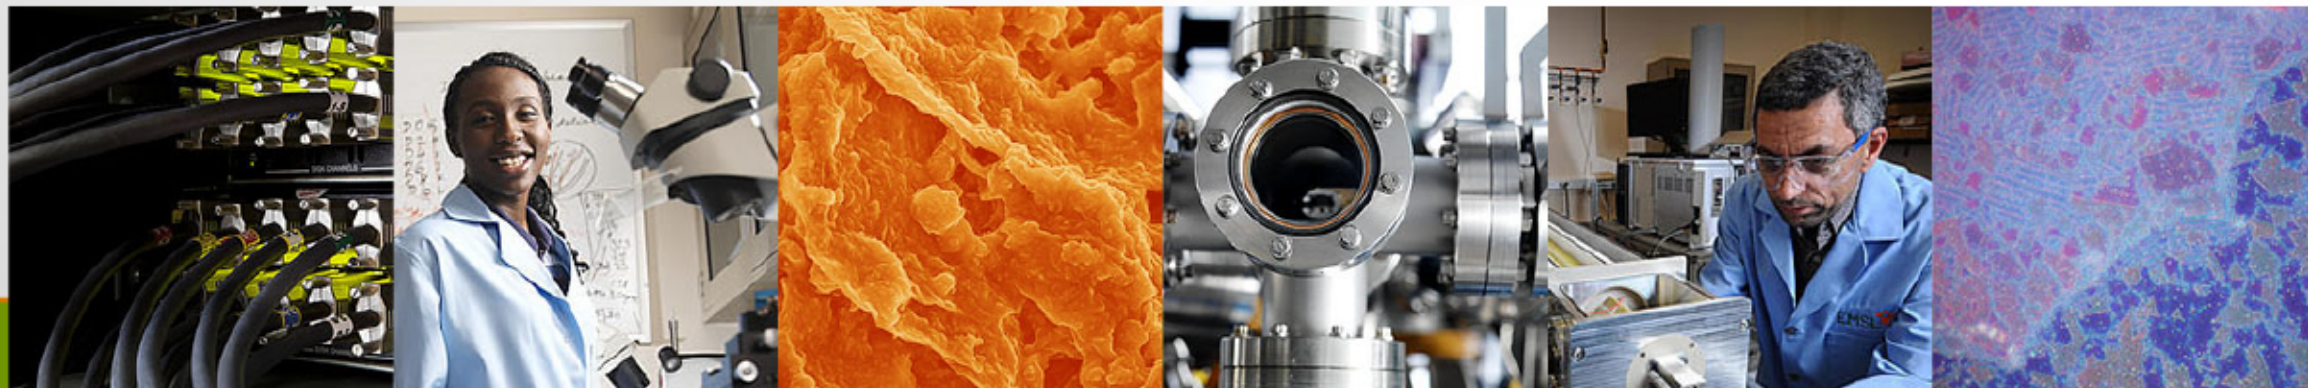


# CView and NWPerf for supercomputer performance collection and display.

Evan Felix - [evan.felix@pnnl.gov](mailto:evan.felix@pnnl.gov)

Ken Schmidt – [kenneth.schmidt@pnnl.gov](mailto:kenneth.schmidt@pnnl.gov)

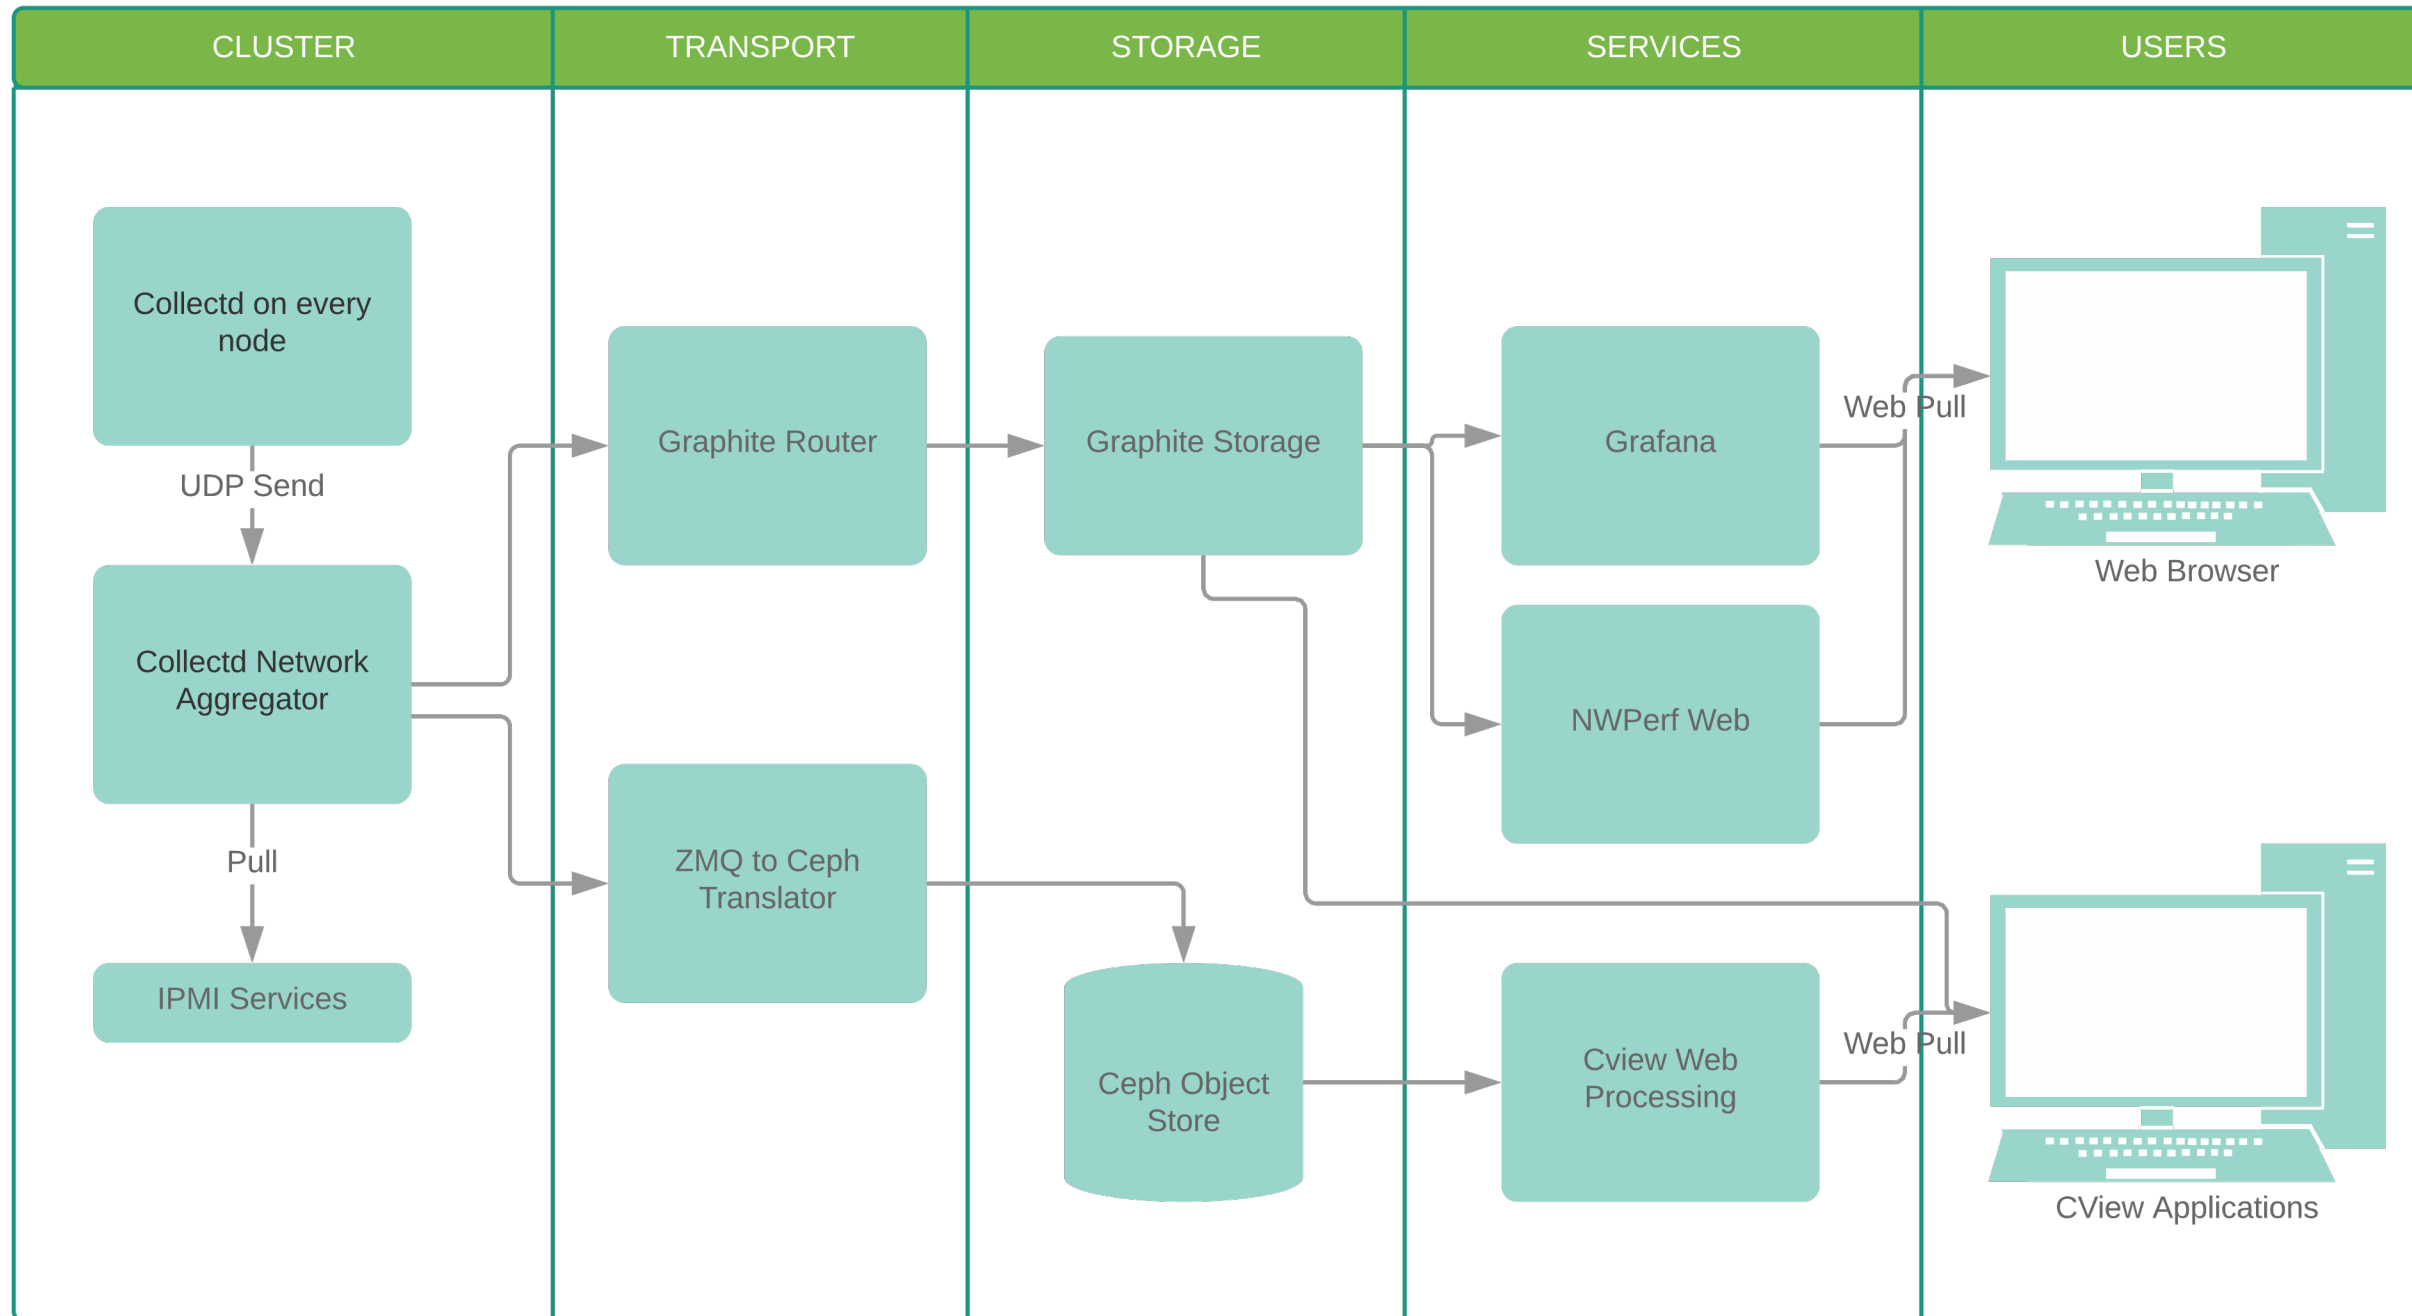


# Why?



- Provide visualization of cluster and job performance metrics for supercomputing sized systems.
  - ▶ System Administrators get an overall 3D View
  - ▶ Users get a 3D View and 2D Web view
- CView is a interactive 3D graphics engine designed for use in displaying large scientific data sets.
- NWPerf is a performance metrics gathering and routing system.

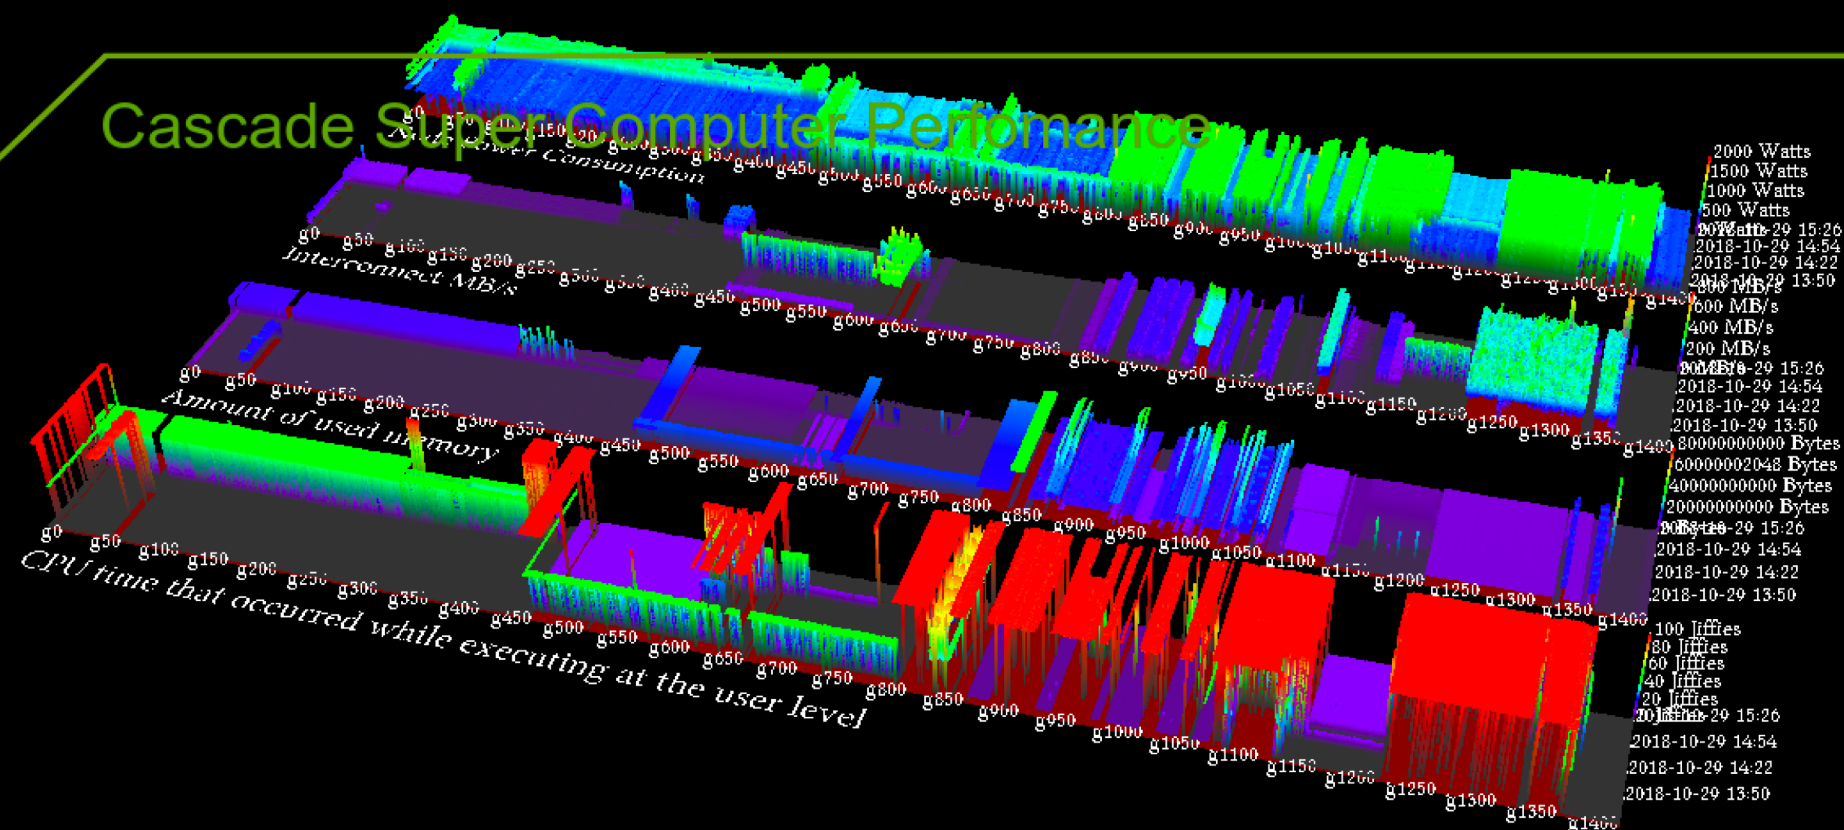
# Architecture Diagram



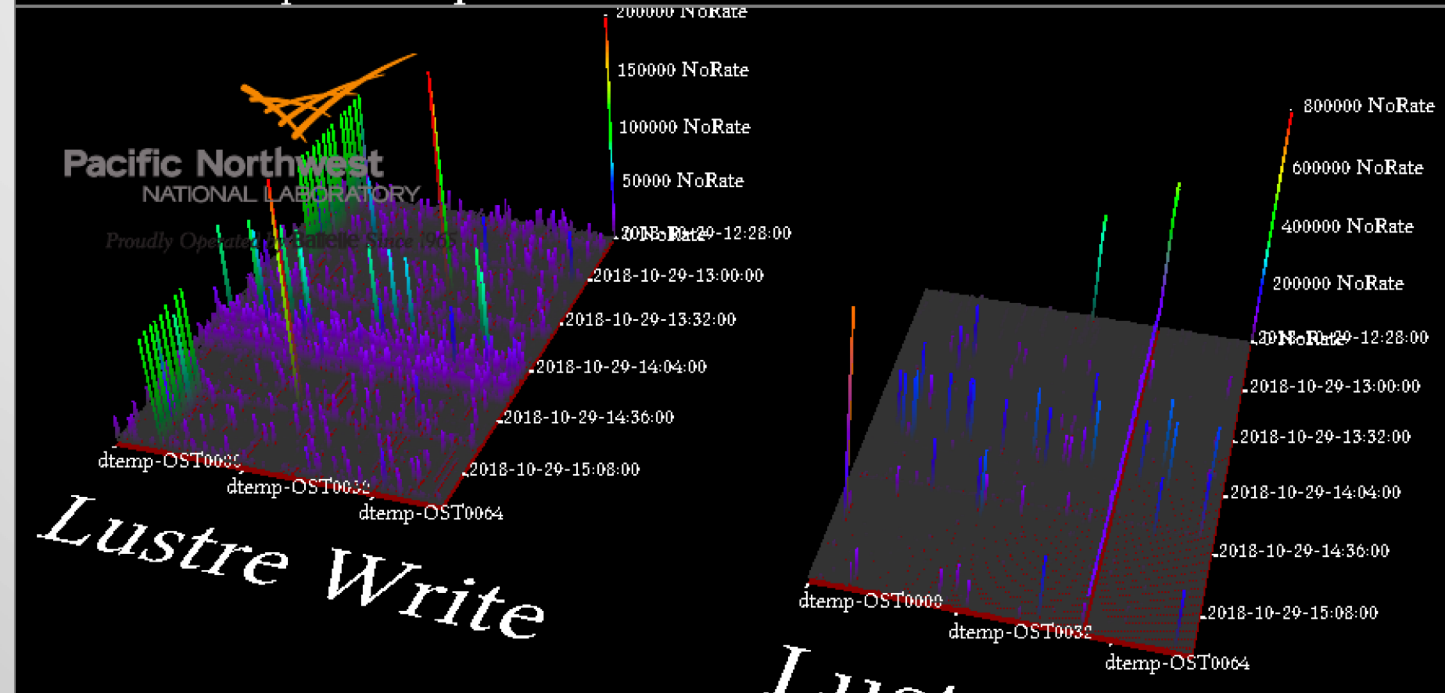


- Interactive 3D Visualization System
- System Views of short or long term data
- Can display multiple types of data,
  - ▶ CView datasets, static or web
  - ▶ Graphite queries
  - ▶ (x,y,val) data similar to gnuplot
  - ▶ Streaming data, such as collectl
- Status dashboards of multiple clusters

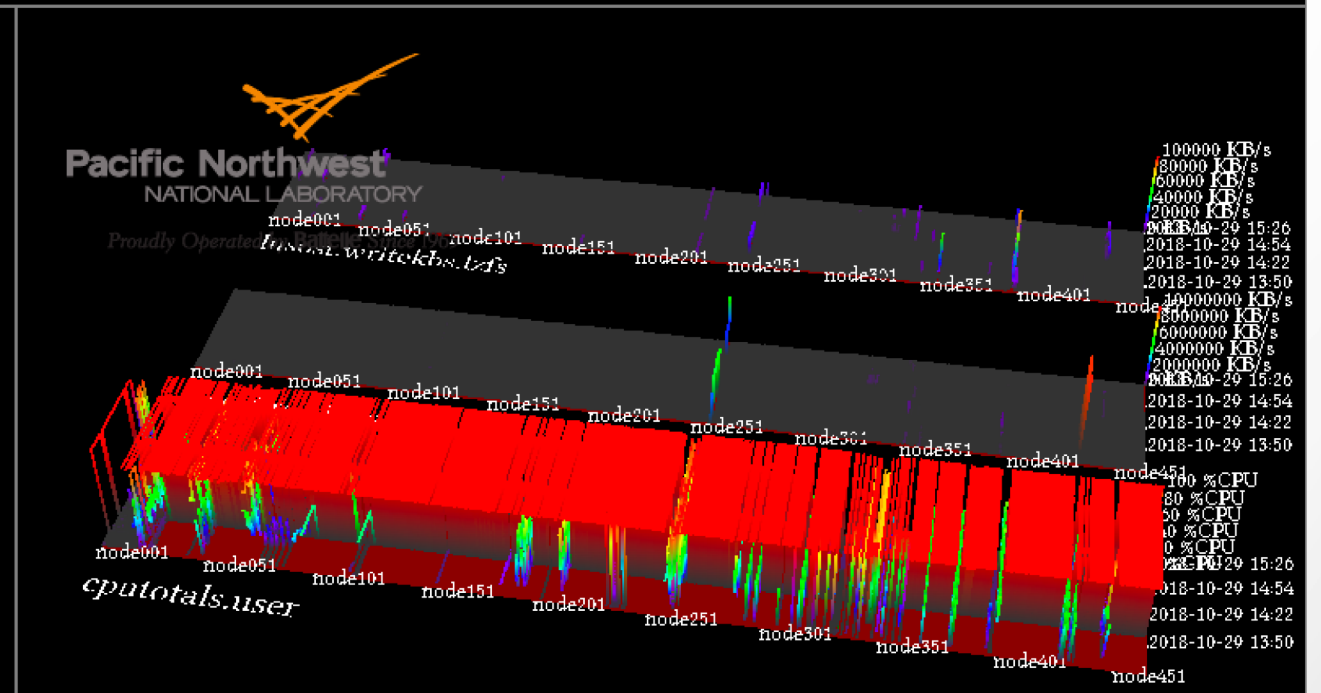
# Cascade Super Computer Performance



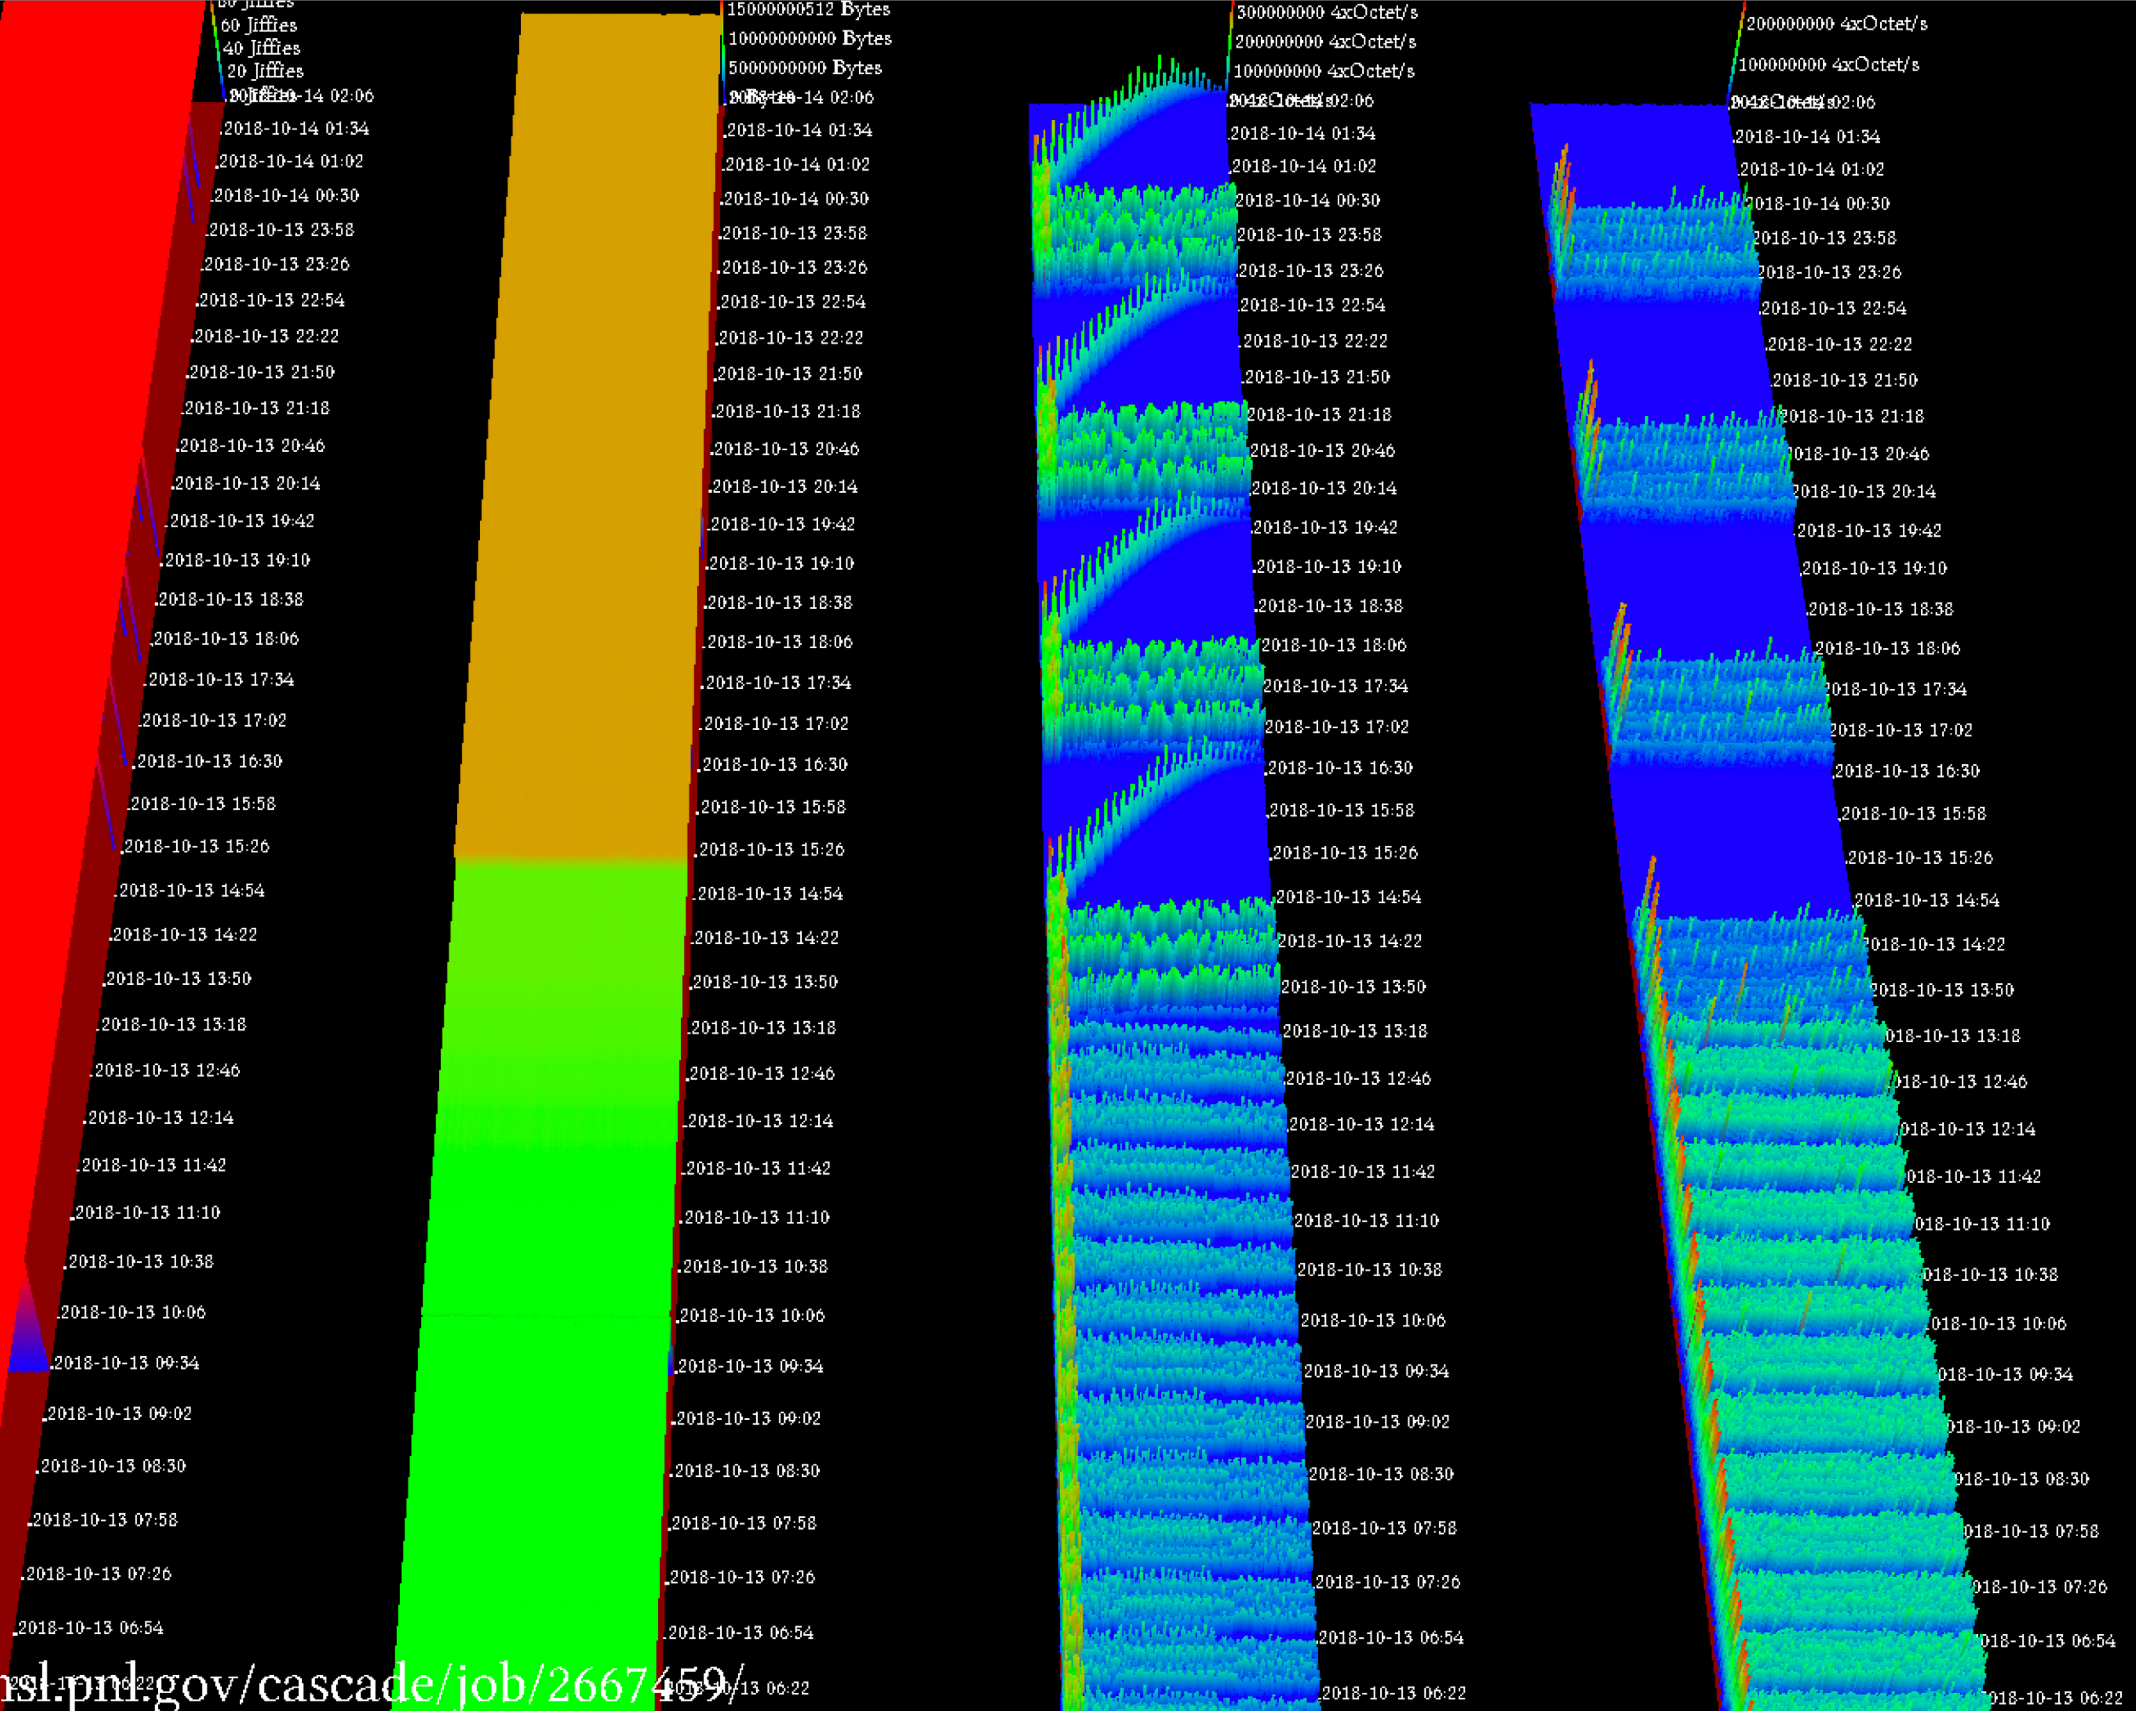
## Cascade Supercomputer

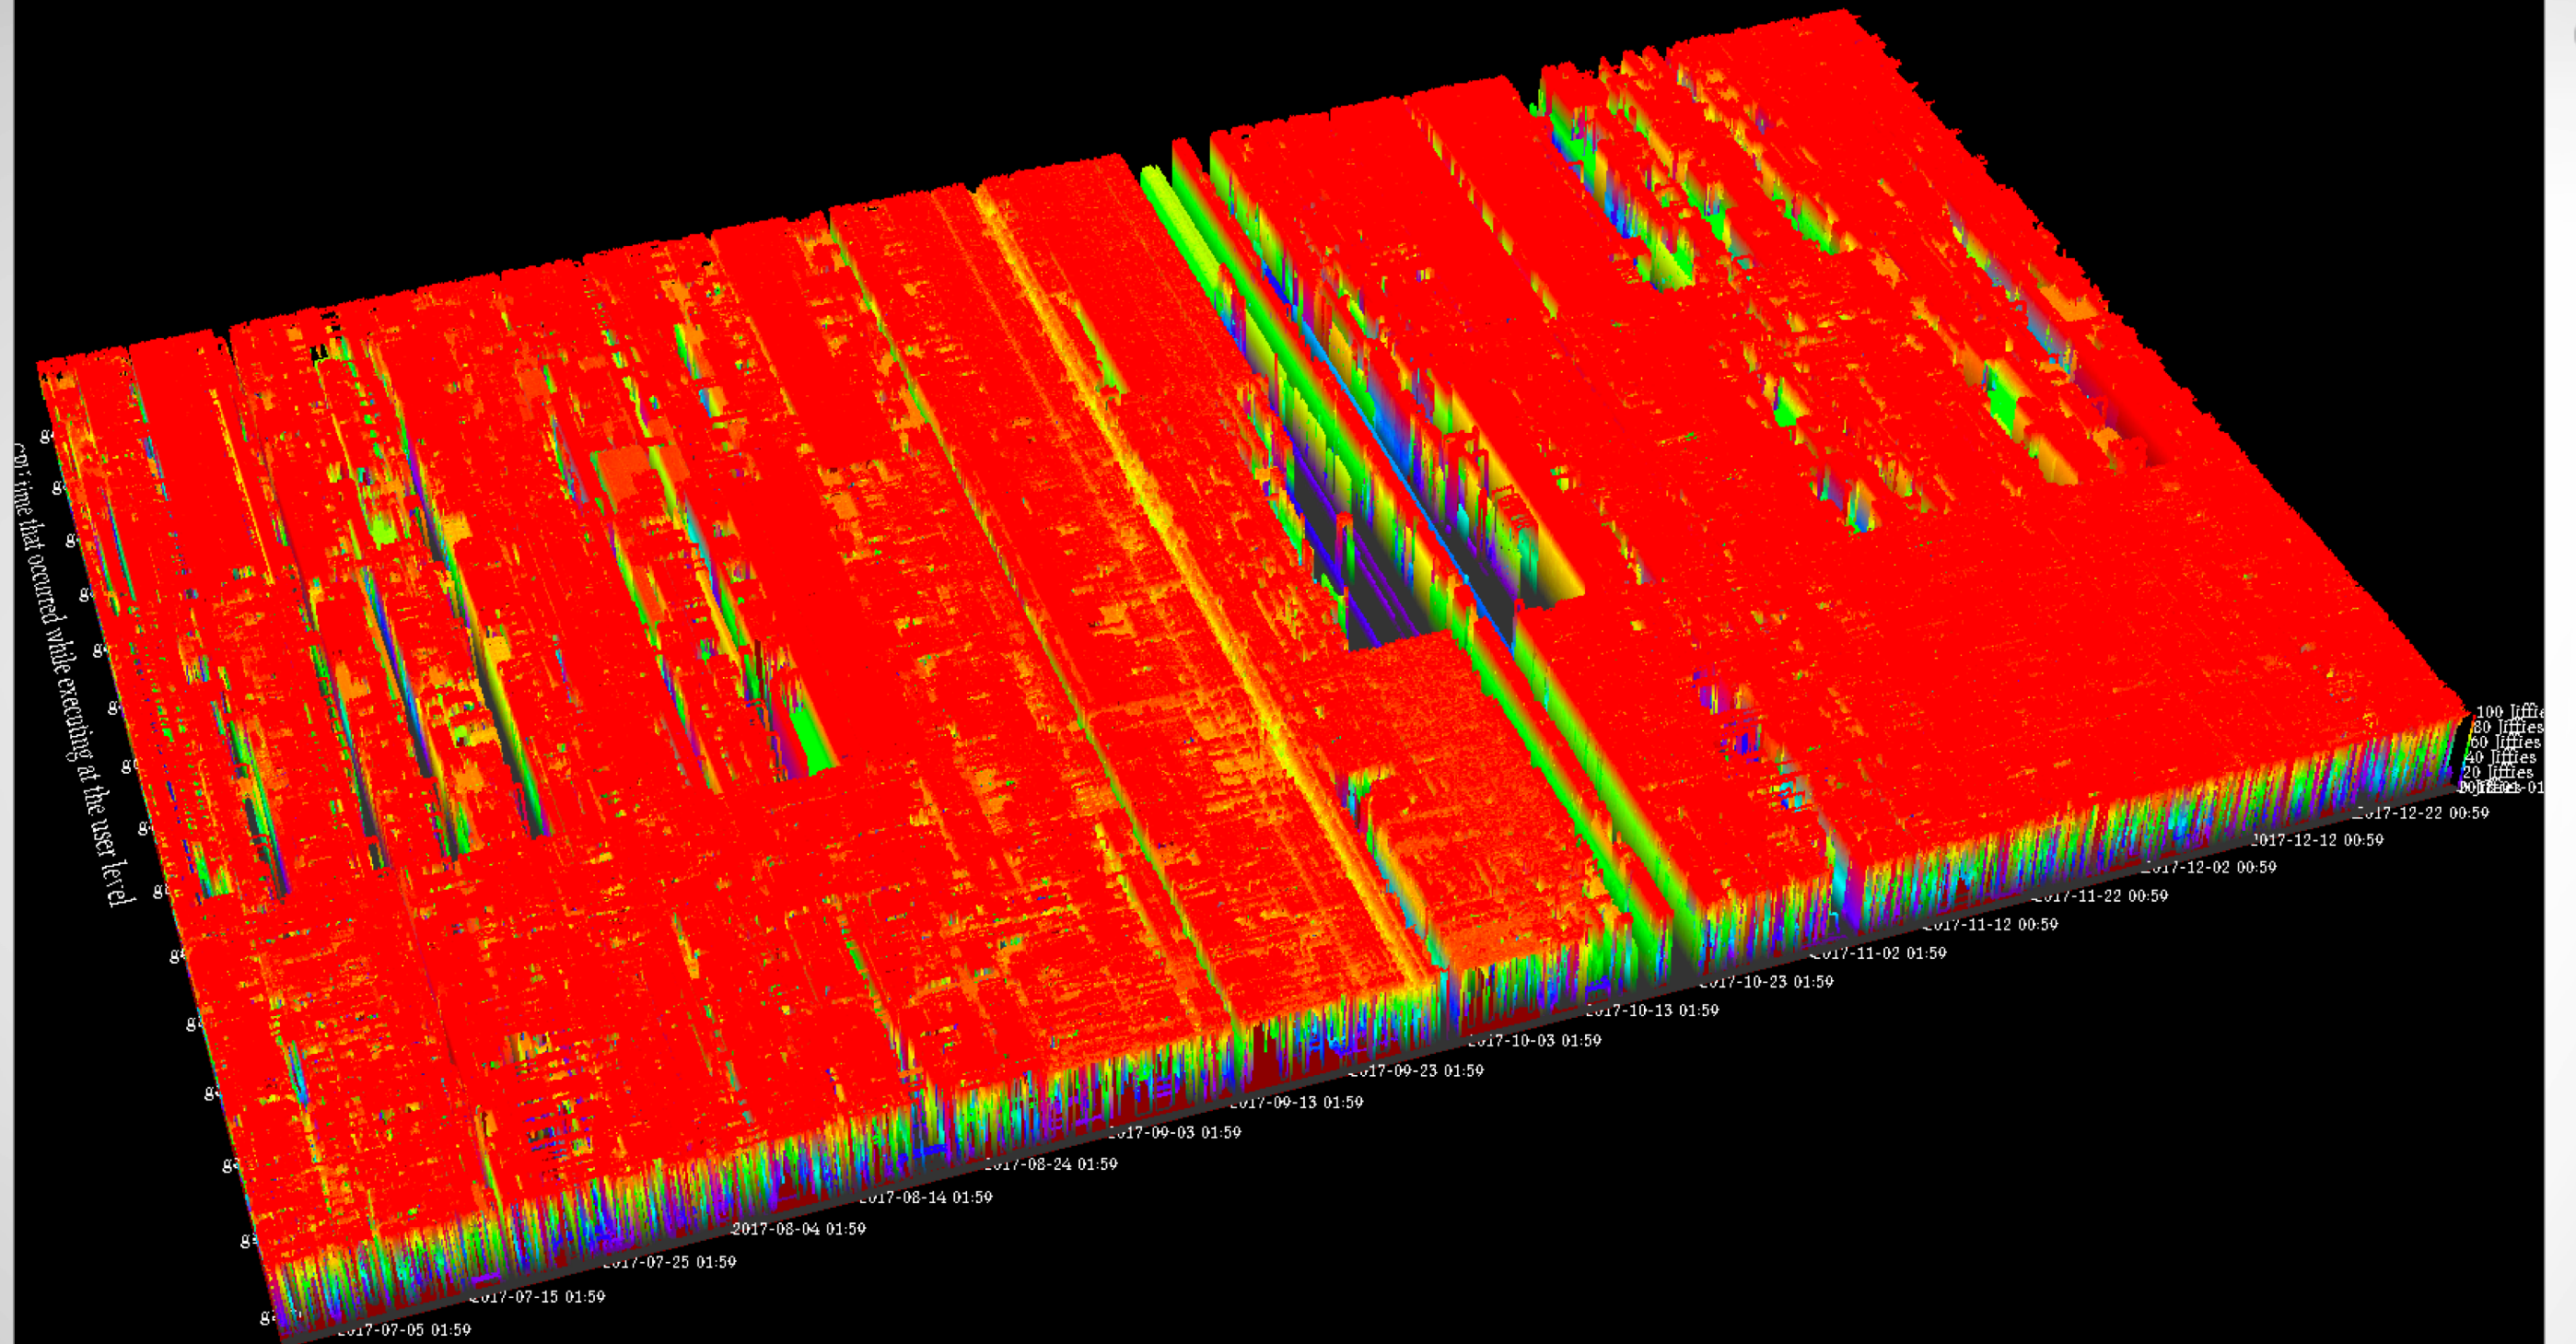


Cascade Lustre Performance



PNNL Institutional Computing: Olympus







- User Focused
  - ▶ No software to install
  - ▶ User friendly interface
  - ▶ Specifically designed around jobs
  - ▶ Jobs performance automatically stored
- Originally based on NWPerf
- Migrated to Graphite backend

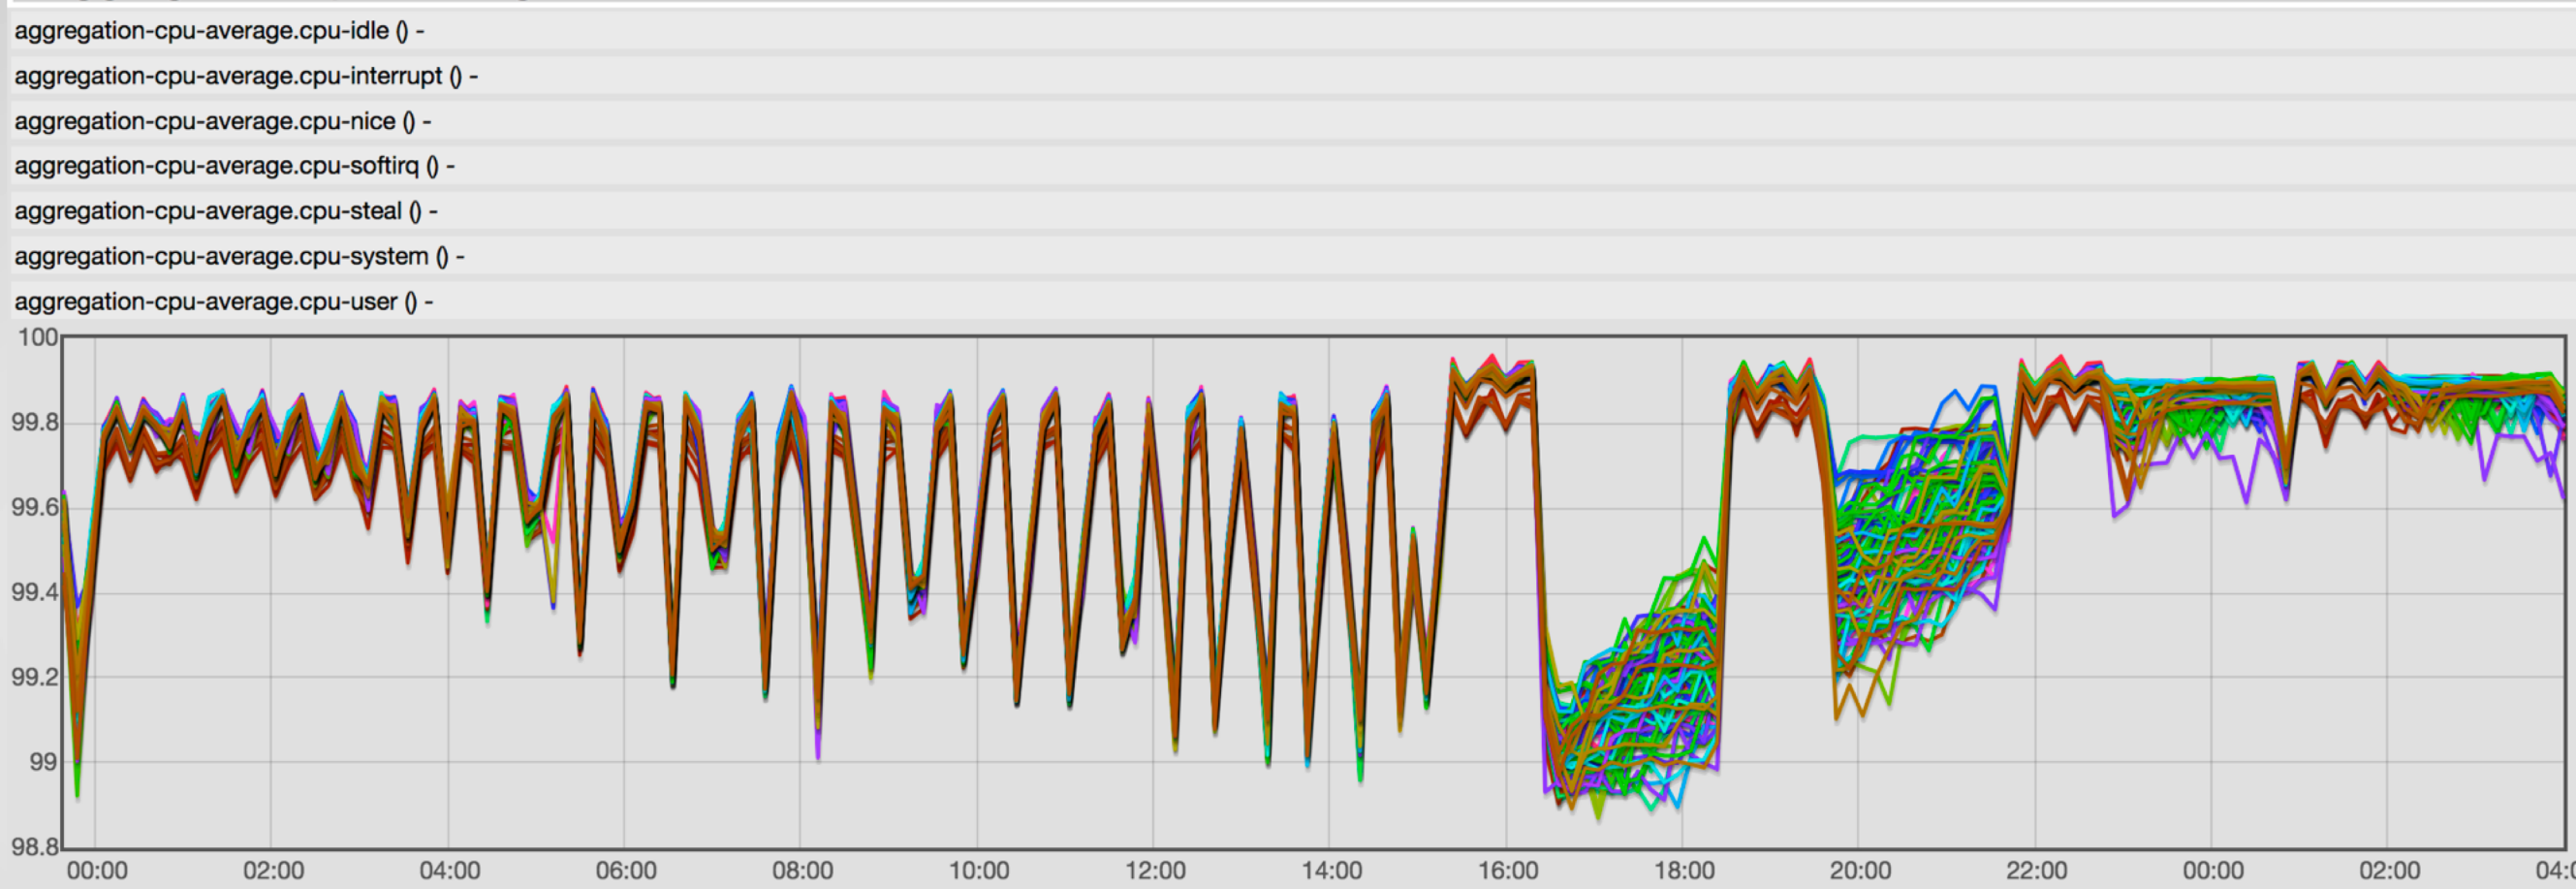


EMSL NwPerf HPC Monitoring

Start Date: After Nov 3 2018  
 End Date: Before Nov 10 2018  
 Start Date: After Nov 3 2018  
 End Date: [blank]

Job ID	Account	User	Submit Time	Start Time	End Time	Run Time	Node Count
2745332	emslc50576		2018-11-04 10:00:06	2018-11-04 10:00:07	2018-11-04 11:01:26		3
2745333	emslc50576		2018-11-04 10:00:20	2018-11-04 10:00:20	2018-11-04 10:56:30		3
2745334	emslc50576		2018-11-04 10:01:11	2018-11-04 10:01:11	2018-11-04 13:22:59		3
2745335	emslc50576		2018-11-04 10:01:16	2018-11-04 10:01:17	2018-11-04 13:03:06		3
2745336	emslc50576		2018-11-04 10:01:23	2018-11-04 10:01:24	2018-11-04 12:48:55		3
2745337	emslc50576		2018-11-04 10:01:35	2018-11-04 10:01:35	2018-11-04 13:14:21		3
2745338	emslc50576		2018-11-04 10:20:14	2018-11-04 10:20:15	2018-11-04 14:55:51		6
2745339	emslc50576		2018-11-04 10:49:53	2018-11-04 10:49:53	2018-11-04 19:56:14		6
2745340	emslg50557		2018-11-04 11:57:07	2018-11-04 11:57:07	2018-11-06 09:53:17	1 day, 21:56:10	14
2745341	mscfhpcc		2018-11-04 14:42:08	2018-11-04 14:42:08	2018-11-04 15:14:04		4
2745342	emsls50086		2018-11-04 15:07:57	2018-11-04 15:07:57	2018-11-04 15:21:39		2
2745343	emsls50086		2018-11-04 15:12:34	2018-11-04 15:12:36	2018-11-04 15:35:48		4
2745344	mscfhpcc		2018-11-04 15:15:50	2018-11-04 15:15:50	2018-11-04 15:15:55		4
2745345	emslc50576		2018-11-04 15:16:10	2018-11-04 15:16:11	2018-11-04 18:08:17		6
2745346	emsls50086		2018-11-04 15:17:50	2018-11-04 15:17:50	2018-11-04 16:01:48		4
2745347	mscfhpcc		2018-11-04 15:19:49	2018-11-04 15:19:49	2018-11-04 15:20:35		4
2745348	mscfhpcc		2018-11-04 15:22:33	2018-11-04 15:22:36	2018-11-04 15:23:21		4
2745349	emsls50086		2018-11-04 15:24:31	2018-11-04 15:24:31	2018-11-04 16:24:41		8
2745350	mscfhpcc		2018-11-04 15:30:16	2018-11-04 15:30:17	2018-11-04 15:31:05		4
2745351	mscfhpcc		2018-11-04 15:31:24	2018-11-04 15:31:25	2018-11-04 17:20:11		4
2745352	emslc50583		2018-11-04 17:18:05	2018-11-04 17:18:05	2018-11-04 17:19:03		1
2745353	emslc50583		2018-11-04 17:23:25	2018-11-04 17:23:26	2018-11-04 17:24:41		1
2745354	emslc50583		2018-11-04 17:26:47	2018-11-04 17:26:47	2018-11-04 17:27:57		1
2745355	emslc50583		2018-11-04 17:30:22	2018-11-04 17:30:23	2018-11-04 17:31:33		1
2745356	emslc50583		2018-11-04 17:48:39	2018-11-04 17:48:40	2018-11-04 17:49:15		1
2745357	emslc50576		2018-11-04 18:08:37	2018-11-04 18:08:37	2018-11-05 09:22:09	15:13:32	6
2745358	emslc50576		2018-11-04 18:18:45	2018-11-04 18:18:46	2018-11-05 07:50:55	13:32:09	6
2745359	emslc50583		2018-11-04 18:23:34	2018-11-04 18:23:36	2018-11-04 18:23:41		1
2745360	emslc50576		2018-11-04 18:24:40	2018-11-04 18:24:41	2018-11-05 01:48:16	7:23:35	6
2745361	emslc50583		2018-11-04 18:25:45	2018-11-04 18:25:45	2018-11-04 18:26:03		1
2745362	emslc50583		2018-11-04 18:36:52	2018-11-04 18:36:53	2018-11-04 18:37:03		1
2745363	emslc50620		2018-11-04 18:47:20	2018-11-04 18:47:21	2018-11-04 18:47:26		1
2745364	emslc50620		2018-11-04 18:52:14	2018-11-04 18:52:15	2018-11-04 18:52:24		1

**Job ID:** 2667459      **User:**      **Start Time:** 10/12/2018, 9:33:49 PM  
**Account:** emsls50086      **Group:**      **End Time:** 10/14/2018, 2:07:56 AM  
**Number of Hosts:** 128      **Run Time:** 1 days 4:34:7



**Average:** 76.75  
**Sum:** 9824

- Select All
- g94: 99.04
- g95: 99.13
- g96: 0.000
- g97: 99.37
- g98: 99.28
- g99: 99.08
- g100: 99.26
- g101: 99.28
- g102: 0.000
- g103: 99.32
- g104: 99.24
- g105: 99.28
- g106: 99.23
- g108: 0.000
- g109: 0.000
- g110: 99.30
- g111: 99.23
- g112: 0.000
- g113: 99.33
- g115: 99.21
- g116: 99.20
- g117: 99.37
- g118: 99.23
- g119: 99.29
- g120: 99.29
- g121: 99.33
- g122: 0.000
- g123: 99.07
- g129: 99.21
- g130: 99.14

aggregation-cpu-average.cpu-wait () -

▶ interface-eth1

▶ interface-eth0

▼ infiniband-mlx4\_0\_1

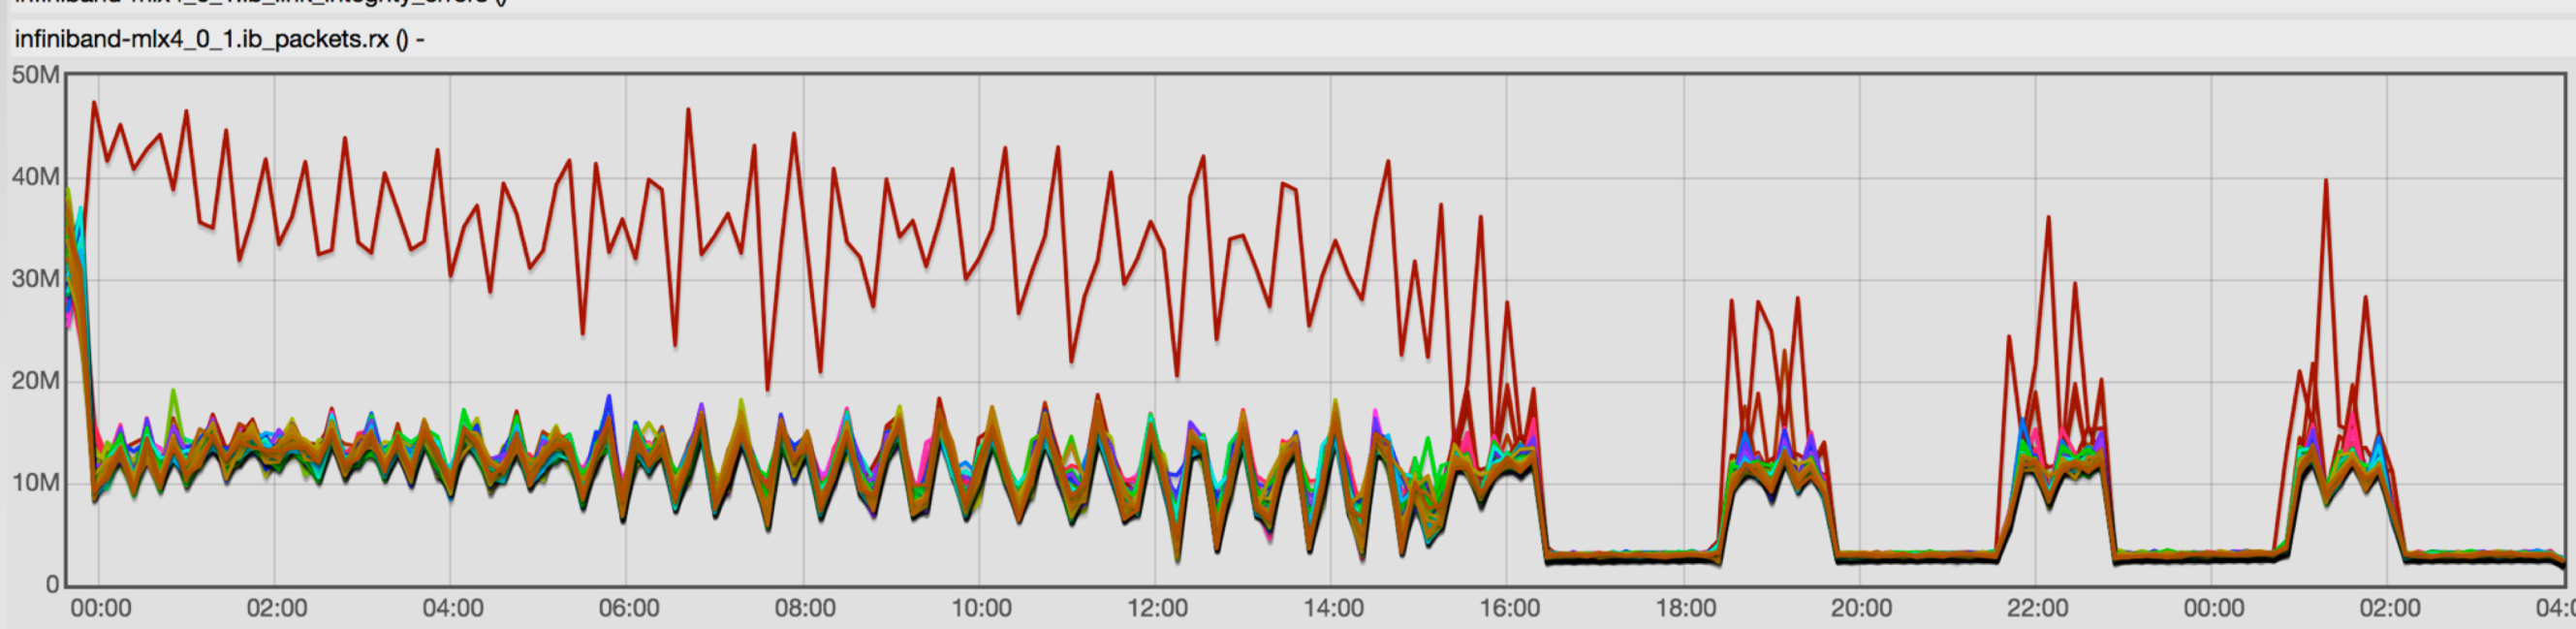
infiniband-mlx4\_0\_1.ib\_buffer\_overruns () -

infiniband-mlx4\_0\_1.ib\_constraint\_errors.rx () -

infiniband-mlx4\_0\_1.ib\_constraint\_errors.tx () -

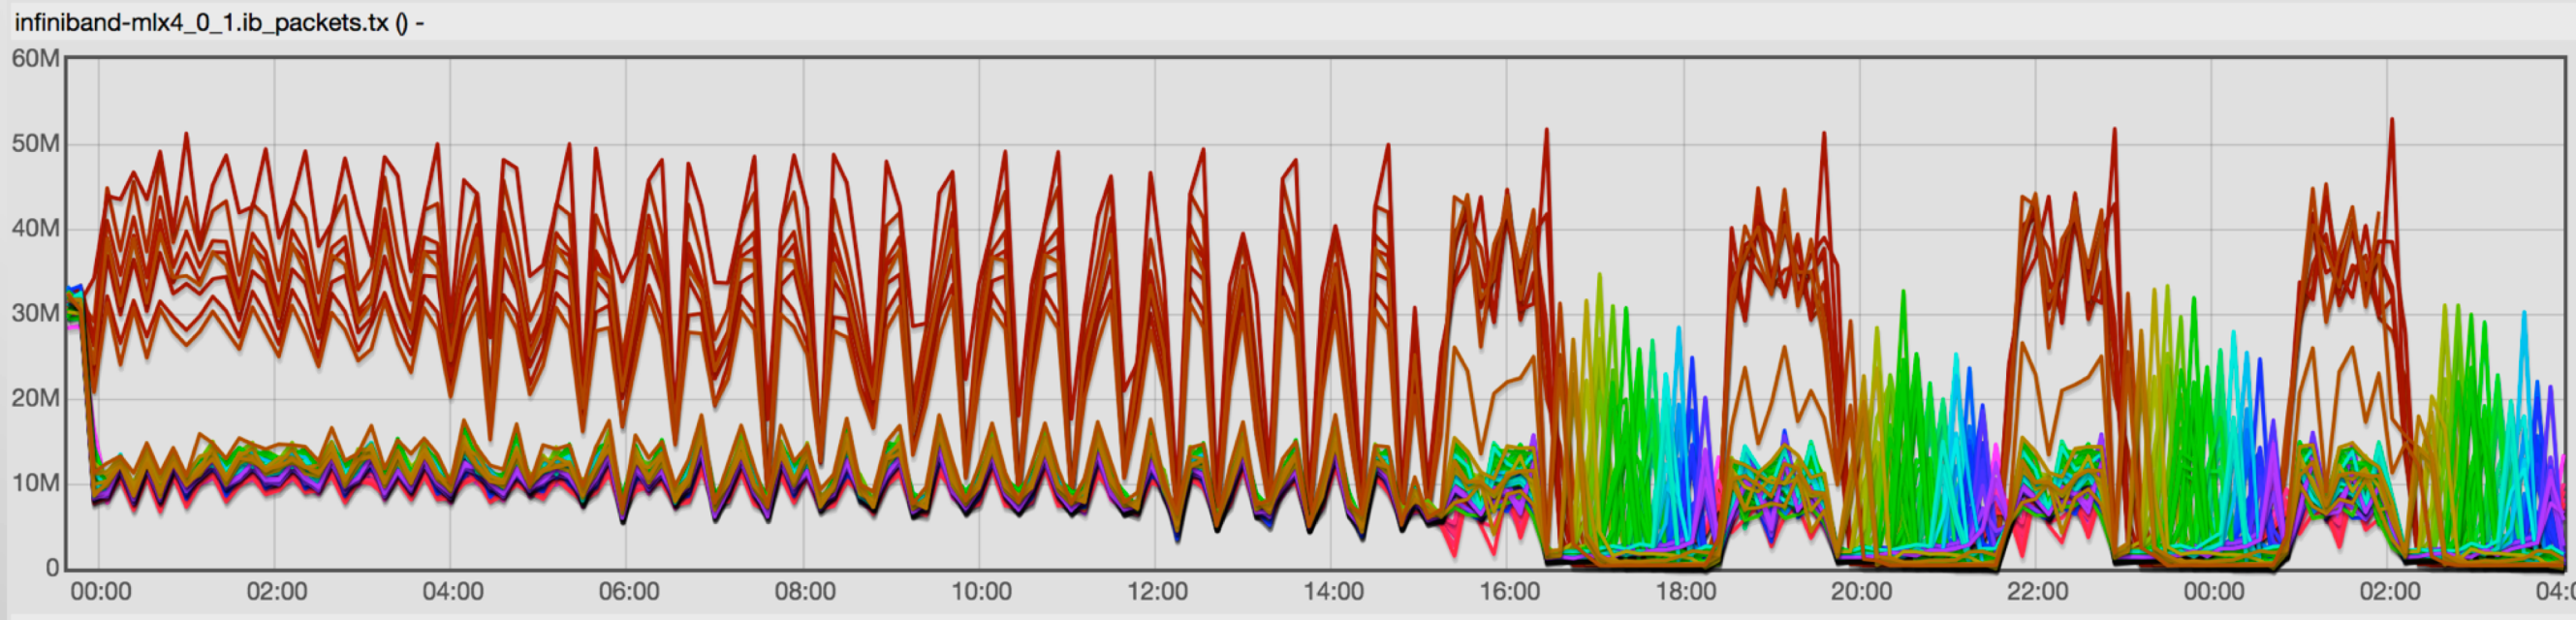
infiniband-mlx4\_0\_1.ib\_data.rx () -

**Job ID:** 2667459    **User:**    **Start Time:** 10/12/2018, 9:33:49 PM  
**Account:** emsls50086    **Group:**    **End Time:** 10/14/2018, 2:07:56 AM  
**Number of Hosts:** 128    **Run Time:** 1 days 4:34:7



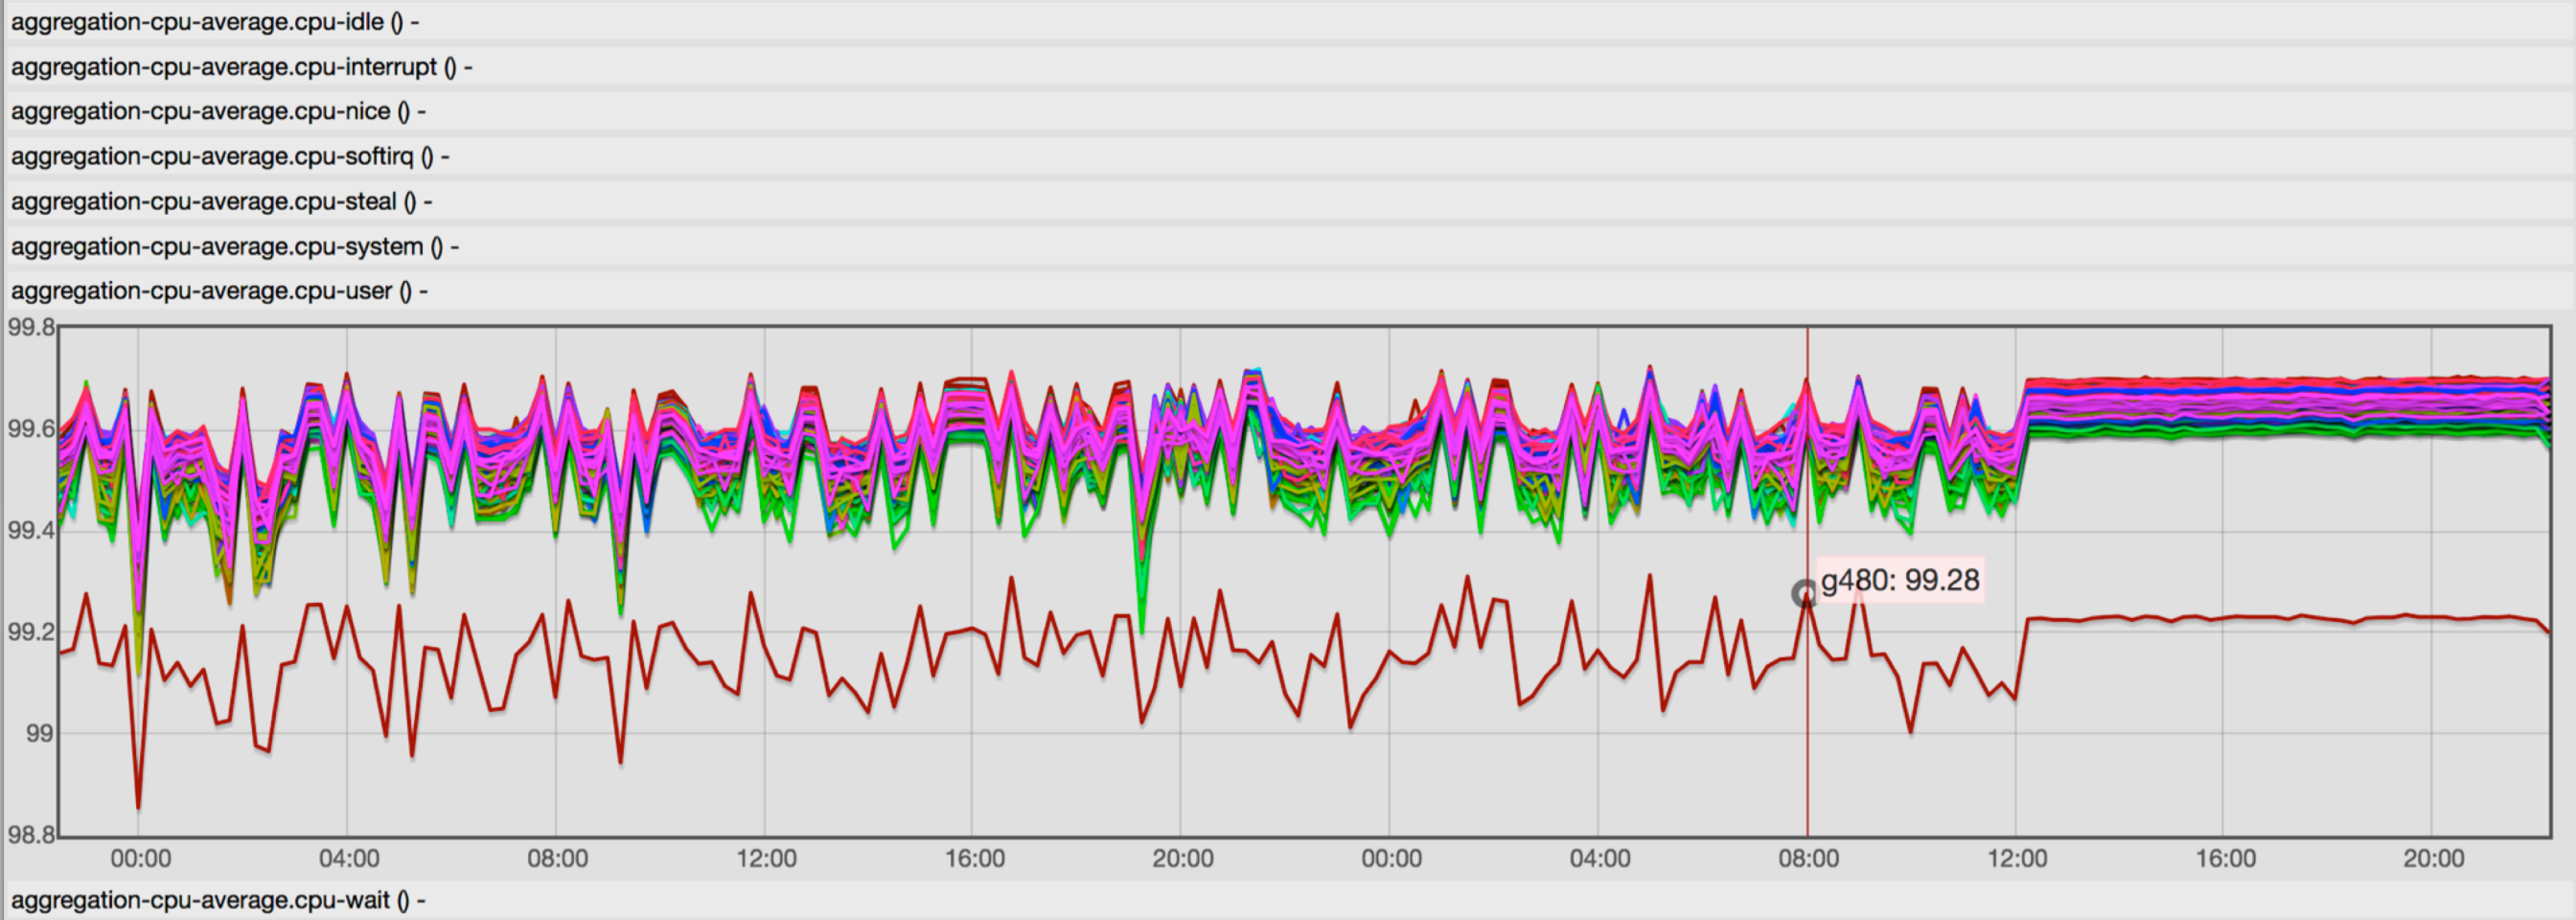
Average: 2.376e+6  
Sum: 3.041e+8

- Select All
- g0: 3.078e+6
- g1: 3.044e+6
- g3: 0.000
- g4: 2.950e+6
- g5: 2.868e+6
- g6: 0.000
- g7: 2.971e+6
- g8: 2.839e+6
- g9: 2.887e+6
- g10: 2.957e+6
- g11: 0.000
- g12: 2.907e+6
- g13: 2.960e+6
- g14: 0.000
- g15: 3.123e+6
- g16: 3.136e+6
- g17: 3.060e+6
- g18: 2.939e+6
- g19: 2.978e+6
- g20: 3.042e+6
- g21: 3.044e+6
- g22: 0.000
- g23: 3.056e+6
- g24: 3.033e+6
- g25: 2.865e+6
- g26: 2.927e+6
- g27: 3.083e+6
- g28: 3.045e+6
- g29: 0.000
- g30: 2.984e+6



infiniband-mlx4\_0\_1.ib\_remote\_phys\_errors () -  
infiniband-mlx4\_0\_1.ib\_switch\_relay\_errors () -

**Job ID:** 2082835      **User:**      **Start Time:** 3/1/2018, 8:14:48 PM  
**Account:** emslc50024      **Group:**      **End Time:** 3/3/2018, 8:14:48 PM  
**Number of Hosts:** 256      **Run Time:** 2 days 0:0:0



**Average:** 99.63  
**Sum:** 2.551e+4

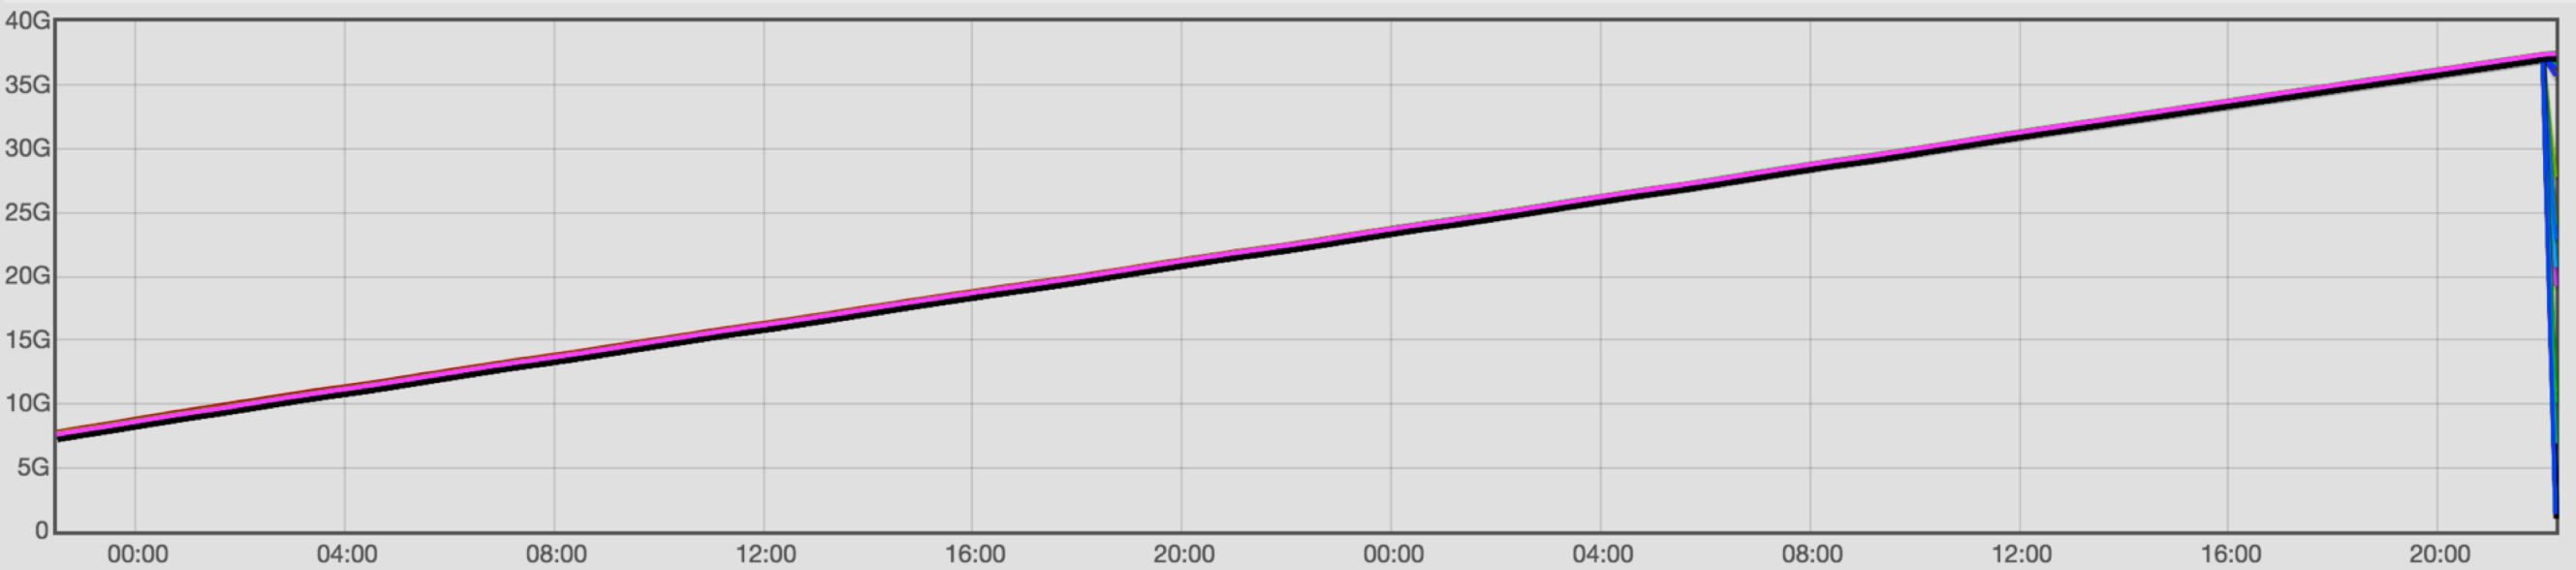
- Select All
- g480 : 99.26
- g481 : 99.68
- g482 : 99.68
- g483 : 99.69
- g484 : 99.67
- g486 : 99.67
- g487 : 99.67
- g488 : 99.67
- g490 : 99.66
- g491 : 99.67
- g492 : 99.67
- g493 : 99.66
- g495 : 99.61
- g496 : 99.64
- g498 : 99.65
- g499 : 99.67
- g500 : 99.61
- g501 : 99.67
- g502 : 99.66
- g505 : 99.65
- g507 : 99.64
- g508 : 99.64
- g511 : 99.61
- g512 : 99.64
- g513 : 99.65
- g514 : 99.63
- g516 : 99.63
- g517 : 99.64
- g518 : 99.59
- g519 : 99.62

- aggregation-cpu-average.cpu-wait () -
- ▶ interface-eth1
- ▶ interface-eth0
- ▶ infiniband-mlx4\_0\_1
- ▶ df-home
- ▶ df-boot
- ▶ df-run-user

**Job ID:** 2082835    **User:**    **Start Time:** 3/1/2018, 8:14:48 PM  
**Account:** emslc50024    **Group:**    **End Time:** 3/3/2018, 8:14:48 PM  
**Number of Hosts:** 256    **Run Time:** 2 days 0:0:0

▼ memory

- memory.memory-buffered () -
- memory.memory-cached () -
- memory.memory-free () -
- memory.memory-slab\_recl () -
- memory.memory-slab\_unrecl () -
- memory.memory-used () -



Average: 2.881e+10

Sum: 7.376e+12

- Select All
- g480: 2.884e+10
- g481: 2.869e+10
- g482: 2.869e+10
- g483: 2.869e+10
- g484: 2.883e+10
- g486: 2.883e+10
- g487: 2.884e+10
- g488: 2.883e+10
- g490: 2.883e+10
- g491: 2.884e+10
- g492: 2.883e+10
- g493: 2.881e+10
- g495: 2.880e+10
- g496: 2.885e+10
- g498: 2.881e+10
- g499: 2.886e+10
- g500: 2.883e+10
- g501: 2.883e+10
- g502: 2.883e+10
- g505: 2.882e+10
- g507: 2.882e+10
- g508: 2.883e+10
- g511: 2.879e+10
- g512: 2.881e+10
- g513: 2.881e+10
- g514: 2.880e+10
- g516: 2.882e+10
- g517: 2.881e+10
- g518: 2.882e+10
- g519: 2.882e+10

- ▶ df-dev-shm
- ▶ interface-lo
- ▶ df-opt
- ▶ interface-enp1s0f1
- ▶ interface-enp1s0f0
- ▶ interface-ib0



# Questions?

NWPerf – <https://github.com/EMSL-MSL/NWPerf>

Cview – <https://github.com/EMSL-MSL/cview>

Collectd – <https://collectd.org/>

Graphite – <https://graphiteapp.org/>

Grafana – <https://grafana.com/>

# Thanks!



[www.emsl.pnnl.gov](http://www.emsl.pnnl.gov)



U.S. DEPARTMENT OF  
**ENERGY**