
CERN

NOTED and BBR at DC24

CERN
IT Department CS Group

LHCOPN LHCONE Meeting #52 - INFN Catania
9th - 11th April 2024

Carmen Misa Moreira
Edoardo Martelli

Outline

NOTED
- Motivation
- Architecture
- Elements
- Modes of operation
- States of execution
- NOTED demo at DC24

Packet pacing
- Motivation
- BBR (Bottleneck Bandwidth and Round-Trip Time)
- Preliminary tests
- Preliminary MSS testing
- BBR demo at DC24

NOTED

Motivation

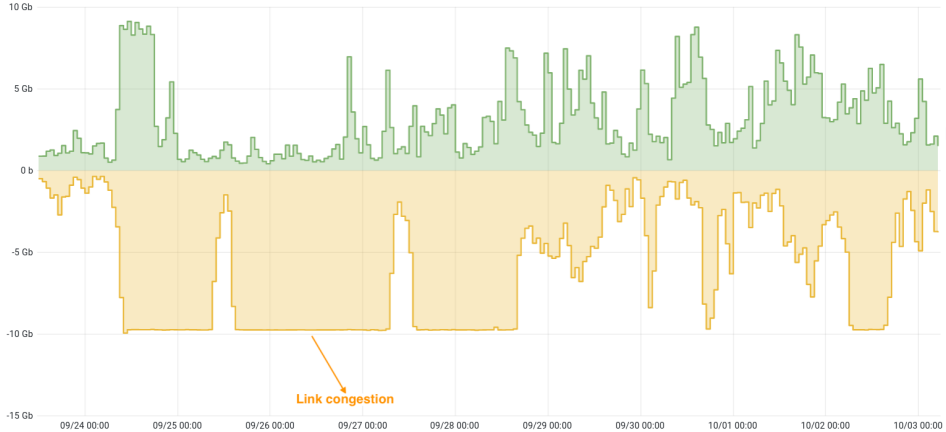

- Large data transfers can saturate network links while alternative paths may be left idle

Architecture

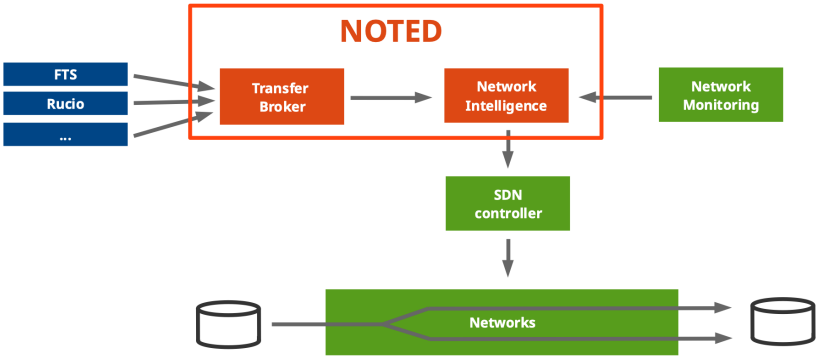

NOTED (Network Optimized Transfer of Experimental Data)

An intelligent network controller to improve the throughput of large data transfers in FTS (File Transfer Services) by handling dynamic circuits.

Elements

FTS (File Transfer Service):

- Analyse data transfers to estimate if any action can be applied to optimise the network utilization → get on-going and queued transfers.

CRIC (Computing Resource Information Catalog):

- Use the CRIC database to get an overview of the network topology → get IPv4/IPv6 addresses, endpoints, rcsite and federation.

FTS
File Transfer Service

elasticsearch

Modes of operation

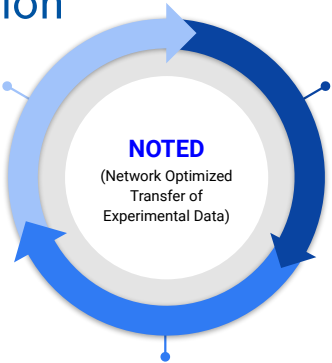

- Much more complex for LHCONE since a single path is shared by multiple sites ~ 100.

States of execution

- Decision-making: NOTED is making the network decision to potentially execute an action or not.
- Running: NOTED is running but there are no transfers in FTS so NOTED is waiting and running until the link-saturation alarm is cleared.
- Monitoring: NOTED is running and there are on-going FTS transfers, but they are below the defined bandwidth threshold that we establish.
- Action: NOTED is running and has triggered an SDN action to provide more bandwidth.
- Stopped: NOTED has stopped because there are no transfers in FTS and the link-saturation alarm has cleared.

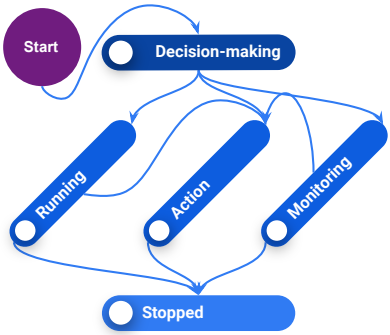

NOTED alarms in MONIT Grafana [Link to the dashboard]

NOTED Alarms

ID	Alarm name	Version	NOTED status	NOTED action	SDN status	Max FTS Throughput [Gb/s]	Interface
184	CH-CERN to CA-TRIUMF	CUSTOM	Running	Spectrum generated an alarm: NOTED is inspecting FTS.	Not provided	0	
187	DE-KIT to CA-TRIUMF	CUSTOM	Action	On-going SDN. FTS throughput [Gb/s]: 5.56	Provided	9.94	
211	CH-CERN to FR-CCIN2P3	CUSTOM	Monitoring	No transfers found in FTS. NOTED is still running until Spectrum clears the alarm.	Not provided		
219	DE-KIT to CA-TRIUMF	CUSTOM	Stopped	The large data transfer is finished.	Released	22.3	
73	ES-ATLAS-T2 to CH-CERN	LHCONE	Decision-making	An action on the link may be required: number of events: 1. Throughput [Gb/s]: 4.12	Not provided		I513-e-rjup1-1_irb.111
83	FR-CCIN2P3 to CH-CERN	LHCONE	Action	On-going SDN. FTS throughput [Gb/s]: 4.94	Provided	7.52	I513-e-rjup1-1_irb.111
84	RO-LCG to CH-CERN	LHCONE	Stopped	The large data transfer is finished.	Released	10.3	I513-e-rjup1-1_irb.111
85	ES-PIC to CH-CERN	LHCONE	Action	On-going SDN. FTS throughput [Gb/s]: 5.94	Provided	12.6	I513-e-rjup1-1_irb.111
107	FR-GRIF to CH-CERN	LHCONE	Monitoring	No transfers found in FTS. NOTED is still running until Spectrum clears the alarm.	Not provided		I513-e-rjup1-1_irb.111
108	IT-INFN-T2 to CH-CERN	LHCONE	Stopped	The large data transfer is finished.	Released	27.9	I513-e-rjup1-1_irb.111
116	UK-SouthGrid to CH-CERN	LHCONE	Running	Spectrum generated an alarm: NOTED is inspecting FTS.	Not provided		I513-e-rjup1-1_irb.111
29	AU-ATLAS to CH-CERN	LHCOPN	Stopped	The large data transfer is finished.	Released	8.79	I513-e-rjup1-1_irb.3530
30	CH-CERN to CA-TRIUMF	LHCOPN	Action	On-going SDN. FTS throughput [Gb/s]: 7.45	Provided	31.5	I513-e-rjup1-1_irb.2126
31	CH-CERN to DE-KIT	LHCOPN	Stopped	The large data transfer is finished.	Released	17.7	I513-e-rjup1-1_irb.3530
32	CH-CERN to DE-KIT	LHCOPN	Monitoring	No transfers found in FTS. NOTED is still running until Spectrum clears the alarm.	Not provided	0	I513-e-rjup1-1_irb.3530
36	NL-T1 to CH-CERN	LHCOPN	Decision-making	An action on the link may be required: number of events: 1. Throughput [Gb/s]: 6.48	Not provided		I513-e-rjup1-1_irb.3530
37	DE-KIT to CH-CERN	LHCOPN	Running	Spectrum generated an alarm: NOTED is inspecting FTS.	Not provided		I513-e-rjup1-1_irb.3530

NOTED demo at DC24 (LHCOPN, LHCONE versions)

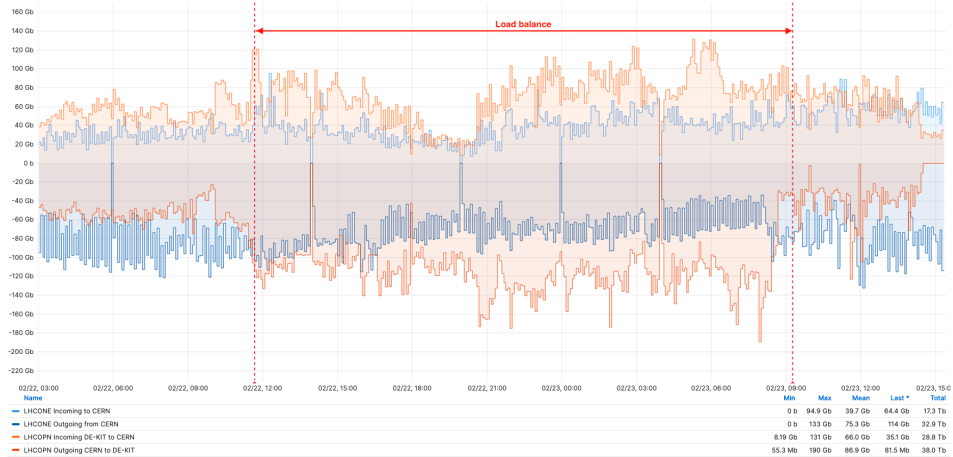

- DE-KIT load balancing between LHCOPN and LHCONE (from 22nd to 23rd of February 2024)

NOTED demo at DC24 (LHCOPN, LHCONE versions)

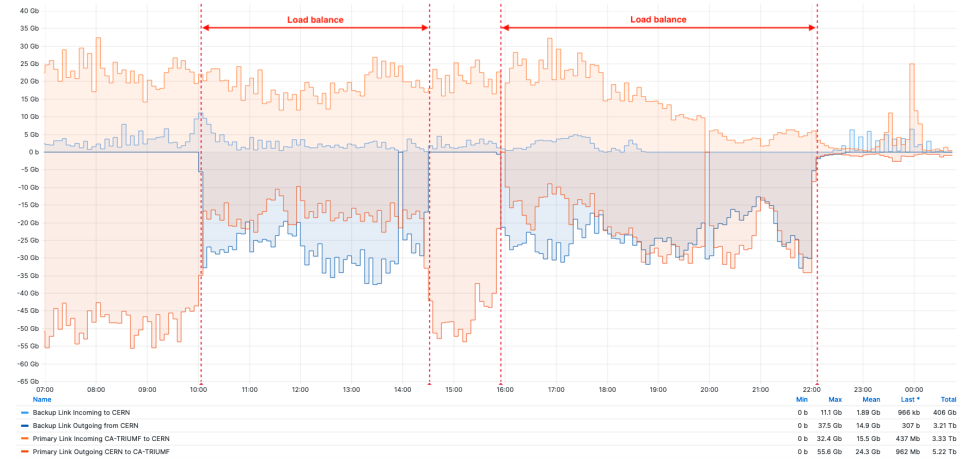

- CA-TRIUMF load balancing between LHCOPN and its backup link (from 21st to 23rd of February 2024)

NOTED demo at DC24 (LHCOPN, LHCONE versions)

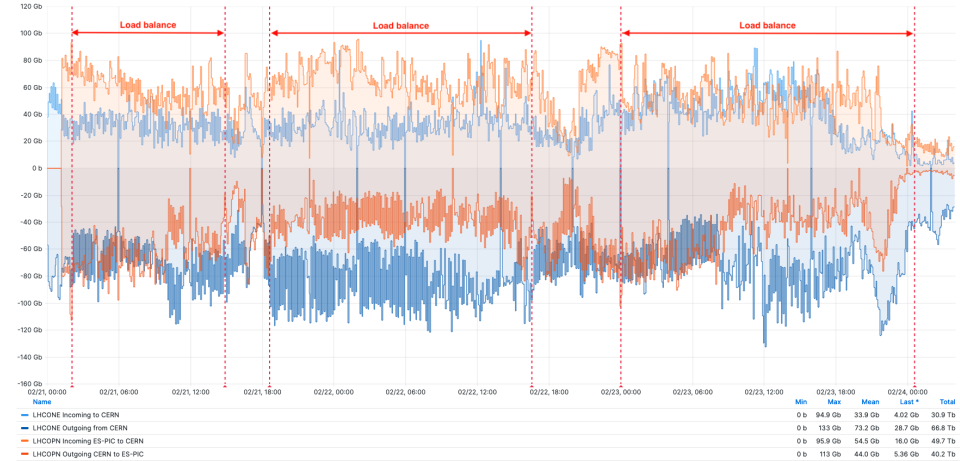

- ES-PIC load balancing between LHCOPN and LHCONE (from 21st to 23rd of February 2024)

DC24 participants

- Monitoring of LHCONE and LHCOPN links at CERN.
- For CA-TRIUMF: load balance with their backup link.
- For ES-PIC and DE-KIT: load balance between LHCOPN and LHCONE.
- Dry-run mode for the rest of Tier 1's.

Packet pacing

Motivation

- Improve performance of data transfer:
  - being more resilient to packet loss
  - better fitting in buffer-constrained networks
  - reducing load on sending hosts
  - more effectively sharing available bandwidth

BBR (Bottleneck Bandwidth and Round-Trip Time)

- TCP congestion control algorithm developed by Google and standardised within the IETF
- Seeks to achieve maximum throughput with minimum congestion, while also keeping queues short which minimises delay
- In traditional loss-based algorithms (CUBIC, RENO,), the sending rate is given by the size of the congestion window, and the sender node may send packets in bursts, up to the maximum rate of the sender's interface. Thus, traditional algorithms rely on routers' buffers to absorb packet bursts. BBR uses packet pacing to set the sending rate to the estimated bottleneck bandwidth. The pacing technique spaces out or paces packets at the sender node, spreading them over time.
- BBRv3 vs BBRv1: better coexistence with RENO/CUBIC and lower loss rate

BBRv3 is not available yet on any of the WLCG recommended Linux distributions

Testbed used for preliminary tests

- PerfSONAR hosts located at CERN, ESnet and JISC
- Run 100 tests with iperf3 and perfSONAR scheduler for IPv4/IPv6 by using a Python script
  - Single stream
  - Duration: 60 sec
  - For different congestion protocol: CUBIC, RENO, BBRv1 and BBRv3
- The data is exported to Elastic Search and MONIT Grafana

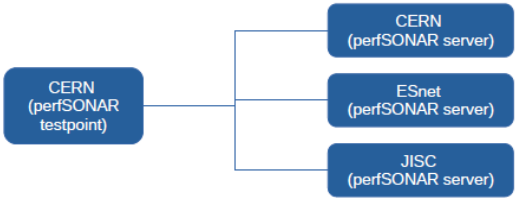

NOTE: the perfSONAR servers are not properly tune (yet)

Monitoring system: Grafana

- Throughput, latency, losss, RTT and hops measurements

Preliminary tests: summary for IPv4

- MTU = 9000
- For CERN hosts: we don't see an improvement by using BBR → maybe we reach the CPU limit or is due to the proximity of the 2 servers, they are on the same rack
- For ESnet hosts: by using BBR we get 1 Gb/s more throughput in comparison with CUBIC and 2 Gb/s more in comparison with RENO
- For JISC hosts: by using BBR we get 4 Gb/s more throughput in comparison with CUBIC and 11 Gb/s more in comparison with RENO

Source	Summary IPv4															
psb02-gva.cern.ch					Throughput		Retransmits		RTT		Hops		Latency		Loss	
Destination	Protocol	IP version	MTU	Streams	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg
pse01-gva.cern.ch	CUBIC	IPv4	9000	1	22.0 Gb/s	18.5 Gb/s	116732	48233	-	-	-	-	-	-	-	-
pse01-gva.cern.ch	RENO	IPv4	9000	1	22.4 Gb/s	18.5 Gb/s	19956	10069	-	-	-	-	-	-	-	-
pse01-gva.cern.ch	BBRv1	IPv4	9000	1	22.6 Gb/s	19.0 Gb/s	128566	25317	-	-	-	-	0.550 ms	0.043 ms	0.003	0
pse01-gva.cern.ch	BBRv3	IPv4	9000	1	21.9 Gb/s	13.9 Gb/s	3342	270	-	-	-	-	0.820 ms	0.163 ms	0.003	0
lbl-dev-dtn.es.net	CUBIC	IPv4	9000	1	9.13 Gb/s	7.67 gb/s	-	-	-	-	-	-	-	-	-	-
lbl-dev-dtn.es.net	RENO	IPv4	9000	1	9.48 Gb/s	6.76 Gb/s	-	-	-	-	-	-	-	-	-	-
lbl-dev-dtn.es.net	BBRv1	IPv4	9000	1	9.01 Gb/s	8.83 Gb/s	-	-	404 ms	161 ms	18	17.9	76.1 ms	75.1 ms	0.002	0
lbl-dev-dtn.es.net	BBRv3	IPv4	9000	1	9.00 Gb/s	8.87 Gb/s	-	-	712 ms	197 ms	18	17.7	76.8 ms	75.2 ms	0	0
ps-london-bw.perf.ja.net	CUBIC	IPv4	9000	1	36.9 Gb/s	24.0 Gb/s	-	-	325 ms	39.3 ms	16	15.8	-	-	-	-
ps-london-bw.perf.ja.net	RENO	IPv4	9000	1	35.8 Gb/s	17.0 Gb/s	-	-	392 ms	46.4 ms	17	16.1	-	-	-	-
ps-london-bw.perf.ja.net	BBRv1	IPv4	9000	1	36.7 Gb/s	28.0 Gb/s	-	-	274 ms	33.2 ms	16	15.5	9.30 ms	8.80 ms	0	0
ps-london-bw.perf.ja.net	BBRv3	IPv4	9000	1	32.2 Gb/s	27.0 Gb/s	-	-	377 ms	70.5 ms	16	15.6	9.38 ms	8.89 ms	0	0

Preliminary tests: summary for IPv6

- MTU = 9000
- For CERN hosts: by using BBR we get 2 Gb/s more throughput in comparison with CUBIC and 1 Gb/s more in comparison with RENO
- For ESnet hosts: by using BBR we get 1 Gb/s more throughput in comparison with RENO and no gaining in comparison with CUBIC
- For JISC hosts: by using BBR we get 5 Gb/s more throughput in comparison with CUBIC and 12 Gb/s more in comparison with RENO

Source	Summary IPv6															
psb02-gva.cern.ch					Throughput		Retransmits		RTT		Hops		Latency		Loss	
Destination	Protocol	IP version	MTU	Streams	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg
pse01-gva.cern.ch	CUBIC	IPv6	9000	1	18.6 Gb/s	14.9 Gb/s	128067	90135	-	-	-	-	-	-	-	-
pse01-gva.cern.ch	RENO	IPv6	9000	1	19.1 Gb/s	15.1 Gb/s	81991	44902	-	-	-	-	-	-	-	-
pse01-gva.cern.ch	BBRv1	IPv6	9000	1	22.9 Gb/s	16.7 Gb/s	521182	245265	-	-	-	-	0.460 ms	0.0729 ms	0	0
pse01-gva.cern.ch	BBRv3	IPv6	9000	1	23.1 Gb/s	16.9 Gb/s	10973	3003	-	-	-	-	0.780 ms	0.173 ms	0	0
lbl-dev-dtn.es.net	CUBIC	IPv6	9000	1	9.47 Gb/s	8.40 Gb/s	-	-	-	-	-	-	-	-	-	-
lbl-dev-dtn.es.net	RENO	IPv6	9000	1	9.48 Gb/s	7.67 Gb/s	-	-	-	-	-	-	-	-	-	-
lbl-dev-dtn.es.net	BBRv1	IPv6	9000	1	8.98 Gb/s	8.82 Gb/s	-	-	411 ms	176 ms	11	10.9	76.1 ms	75.3 ms	0.012	0
lbl-dev-dtn.es.net	BBRv3	IPv6	9000	1	8.97 Gb/s	8.81 Gb/s	-	-	396 ms	165 ms	11	10.8	76.0 ms	75.1 ms	0.002	0
ps-london-bw.perf.ja.net	CUBIC	IPv6	9000	1	35.8 Gb/s	25.3 Gb/s	-	-	385 ms	62.8 ms	17	16.8	-	-	-	-
ps-london-bw.perf.ja.net	RENO	IPv6	9000	1	36.3 Gb/s	18.6 Gb/s	-	-	182 ms	34.3 ms	17	15	-	-	-	-
ps-london-bw.perf.ja.net	BBRv1	IPv6	9000	1	36.9 Gb/s	30.7 Gb/s	-	-	329 ms	62.6 ms	17	16.6	9.27 ms	8.78 ms	0.007	0
ps-london-bw.perf.ja.net	BBRv3	IPv6	9000	1	34.4 Gb/s	28.1 Gb/s	-	-	298 ms	44.1 ms	17	16.5	9.62 ms	8.87 ms	0.002	0

Preliminary MSS testing: summary for IPv4

- BBRv1 and BBRv3 for MSS = 1500, 2048, 4096 and 8192

Source	Summary MSS IPv4															
psb02-gva.cern.ch					Throughput		Retransmits		RTT		Hops		Latency		Loss	
Destination	Protocol	IP version	MSS	Streams	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg
pse01-gva.cern.ch	BBRv1	IPv4	1500	1	23.5 Gb/s	19.2 Gb/s	188548	29957	-	-	-	-	0.680 ms	0.160 ms	0	0
pse01-gva.cern.ch	BBRv1	IPv4	2048	1	22.5 Gb/s	18.8 Gb/s	137269	20329	-	-	-	-	0.440 ms	0.101 ms	0.007	0
pse01-gva.cern.ch	BBRv1	IPv4	4096	1	22.6 Gb/s	19.1 Gb/s	156898	25848	-	-	-	-	0.620 ms	0.0763 ms	0.005	0
pse01-gva.cern.ch	BBRv1	IPv4	8192	1	23.1 Gb/s	19.2 Gb/s	155002	30203	-	-	-	-	0.450 ms	0.0813 ms	0.007	0
pse01-gva.cern.ch	BBRv3	IPv4	1500	1	22.7 Gb/s	13.9 Gb/s	2083	175	-	-	-	-	1.03 ms	0.126 ms	0	0
pse01-gva.cern.ch	BBRv3	IPv4	2048	1	23.7 Gb/s	14.1 Gb/s	1772	205	-	-	-	-	0.580 ms	0.132 ms	0.003	0
pse01-gva.cern.ch	BBRv3	IPv4	4096	1	21.1 Gb/s	14.1 Gb/s	2086	204	-	-	-	-	0.630 ms	0.116 ms	0.018	0
pse01-gva.cern.ch	BBRv3	IPv4	8192	1	21.0 Gb/s	13.9 Gb/s	1335	151	-	-	-	-	0.530 ms	0.156 ms	0	0
lbl-dev-dtn.es.net	BBRv1	IPv4	1500	1	8.98 Gb/s	8.86 Gb/s	-	-	385 ms	154 ms	18	18	78.9 ms	75.3 ms	0	0
lbl-dev-dtn.es.net	BBRv1	IPv4	2048	1	9.00 Gb/s	8.87 Gb/s	-	-	460 ms	158 ms	18	18	77.8 ms	75.1 ms	0.013	0
lbl-dev-dtn.es.net	BBRv1	IPv4	4096	1	9.00 Gb/s	8.87 Gb/s	-	-	370 ms	154 ms	18	18	76.1 ms	75.2 ms	0	0
lbl-dev-dtn.es.net	BBRv1	IPv4	8192	1	8.97 Gb/s	8.86 Gb/s	-	-	628 ms	158 ms	18	18	79.7 ms	75.3 ms	0	0
lbl-dev-dtn.es.net	BBRv3	IPv4	1500	1	8.97 Gb/s	8.82 Gb/s	-	-	393 ms	154 ms	18	18	75.9 ms	75.1 ms	0.042	0
lbl-dev-dtn.es.net	BBRv3	IPv4	2048	1	8.90 Gb/s	8.76 Gb/s	-	-	425 ms	155 ms	18	18	76.1 ms	75.1 ms	0	0
lbl-dev-dtn.es.net	BBRv3	IPv4	4096	1	8.97 Gb/s	8.90 Gb/s	-	-	420 ms	155 ms	18	18	76.4 ms	75.2 ms	0.007	0
lbl-dev-dtn.es.net	BBRv3	IPv4	8192	1	8.99 Gb/s	8.91 Gb/s	-	-	197 ms	152 ms	18	18	76.2 ms	75.1 ms	0	0
ps-london-bw.perf.ja.net	BBRv1	IPv4	1500	1	5.37 Gb/s	3.97 Gb/s	-	-	77 ms	25.6 ms	16	16	9.30 ms	8.76 ms	0	0
ps-london-bw.perf.ja.net	BBRv1	IPv4	2048	1	14.7 Gb/s	9.91 Gb/s	-	-	419 ms	33.5 ms	16	16	9 ms	8.55 ms	0	0
ps-london-bw.perf.ja.net	BBRv1	IPv4	4096	1	28.7 Gb/s	23.2 Gb/s	-	-	161 ms	28.8 ms	16	16	9.17 ms	8.51 ms	0	0
ps-london-bw.perf.ja.net	BBRv1	IPv4	8192	1	35.2 Gb/s	27.6 Gb/s	-	-	157 ms	28.3 ms	16	15.9	9.08 ms	8.47 ms	0	0
ps-london-bw.perf.ja.net	BBRv3	IPv4	1500	1	7.82 Gb/s	3.91 Gb/s	-	-	84.3 ms	28.5 ms	16	16	9.14 ms	8.52 ms	0	0
ps-london-bw.perf.ja.net	BBRv3	IPv4	2048	1	12.1 Gb/s	8.52 Gb/s	-	-	388 ms	38.0 ms	16	16	9.18 ms	8.66 ms	0	0
ps-london-bw.perf.ja.net	BBRv3	IPv4	4096	1	26.6 Gb/s	19.8 Gb/s	-	-	82.9 ms	25.3 ms	16	16	9.87 ms	8.80 ms	0	0
ps-london-bw.perf.ja.net	BBRv3	IPv4	8192	1	28.8 Gb/s	25.2 Gb/s	-	-	245 ms	34.5 ms	16	16	9.41 ms	8.68 ms	0	0

Preliminary MSS testing: summary for IPv4

- BBRv1 and BBRv3 for MSS = 1500, 2048, 4096 and 8192
- For CERN hosts: we don't see any improvements by setting different MSS → why?
- For ESnet hosts: we don't see any improvements by setting different MSS → why?
- For JISC hosts:
  - For MSS = 4096 we can get 13 Gb/s more throughput in comparison with MSS = 2048 and 19 Gb/s more in comparison with MSS = 1500
  - For MSS = 8192 we can get 17 Gb/s more throughput in comparison with MSS = 2048 and 24 Gb/s more in comparison with MSS = 1500

Preliminary MSS testing: summary for IPv6

- BBRv1 and BBRv3 for MSS = 1500, 2048, 4096 and 8192

Source	Summary MSS IPv6															
psb02-gva.cern.ch					Throughput		Retransmits		RTT		Hops		Latency		Loss	
Destination	Protocol	IP version	MSS	Streams	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg
pse01-gva.cern.ch	BBRv1	IPv6	1500	1	22.7 Gb/s	16.8 Gb/s	506007	225346	-	-	-	-	0.580 ms	0.160 ms	0	0
pse01-gva.cern.ch	BBRv1	IPv6	2048	1	22.8 Gb/s	16.9 Gb/s	864335	254022	-	-	-	-	0.380 ms	0.0848 ms	0.012	0
pse01-gva.cern.ch	BBRv1	IPv6	4096	1	22.6 Gb/s	16.9 Gb/s	565427	259222	-	-	-	-	0.550 ms	0.123 ms	0.002	0
pse01-gva.cern.ch	BBRv1	IPv6	8192	1	23.2 Gb/s	17.2 Gb/s	559422	262979	-	-	-	-	0.470 ms	0.0968 ms	0.007	0
pse01-gva.cern.ch	BBRv3	IPv6	1500	1	26.0 Gb/s	17.3 Gb/s	11993	2614	-	-	-	-	0.470 ms	0.0986 ms	0.003	0
pse01-gva.cern.ch	BBRv3	IPv6	2048	1	25.8 Gb/s	17.3 Gb/s	9855	2965	-	-	-	-	0.440 ms	0.129 ms	0.003	0
pse01-gva.cern.ch	BBRv3	IPv6	4096	1	25.0 Gb/s	16.7 Gb/s	10958	2954	-	-	-	-	0.640 ms	0.120 ms	0.007	0
pse01-gva.cern.ch	BBRv3	IPv6	8192	1	23.6 Gb/s	17.2 Gb/s	9483	2851	-	-	-	-	0.440 ms	0.0928 ms	0.02	0
lbl-dev-dtn.es.net	BBRv1	IPv6	1500	1	8.98 Gb/s	8.83 Gb/s	-	-	336 ms	163 ms	11	10	76.1 ms	75.2 ms	0	0
lbl-dev-dtn.es.net	BBRv1	IPv6	2048	1	9.31 Gb/s	8.84 Gb/s	-	-	402 ms	174 ms	11	9.98	76.5 ms	75.0 ms	0.037	0
lbl-dev-dtn.es.net	BBRv1	IPv6	4096	1	8.99 Gb/s	8.87 Gb/s	-	-	435 ms	169 ms	11	11	76.8 ms	75.2 ms	0	0
lbl-dev-dtn.es.net	BBRv1	IPv6	8192	1	8.97 Gb/s	8.85 Gb/s	-	-	385 ms	167 ms	11	10.9	76.9 ms	75.1 ms	0.002	0
lbl-dev-dtn.es.net	BBRv3	IPv6	1500	1	8.95 Gb/s	8.74 Gb/s	-	-	405 ms	160 ms	11	11	76.3 ms	75.3 ms	0	0
lbl-dev-dtn.es.net	BBRv3	IPv6	2048	1	8.9 Gb/s	8.80 Gb/s	-	-	378 ms	163 ms	11	11	77.1 ms	75.3 ms	0	0
lbl-dev-dtn.es.net	BBRv3	IPv6	4096	1	8.97 Gb/s	8.89 Gb/s	-	-	374 ms	164 ms	11	10.9	76.3 ms	75.2 ms	0	0
lbl-dev-dtn.es.net	BBRv3	IPv6	8192	1	8.99 Gb/s	8.91 Gb/s	-	-	383 ms	155 ms	11	11	76.0 ms	75.2 ms	0.002	0
ps-london-bw.perf.ja.net	BBRv1	IPv6	1500	1	4.54 Gb/s	3.65 Gb/s	-	-	337 ms	69.5 ms	17	15.7	9.07 ms	8.56 ms	0	0
ps-london-bw.perf.ja.net	BBRv1	IPv6	2048	1	14.9 Gb/s	8.80 Gb/s	-	-	377 ms	82.4 ms	17	16	9.02 ms	8.49 ms	0	0
ps-london-bw.perf.ja.net	BBRv1	IPv6	4096	1	27.8 Gb/s	21.1 Gb/s	-	-	344 ms	64.7 ms	17	17	9.32 ms	8.54 ms	0	0
ps-london-bw.perf.ja.net	BBRv1	IPv6	8192	1	35.4 Gb/s	28.3 Gb/s	-	-	227 ms	48 ms	17	16.8	9.41 ms	8.58 ms	0.002	0
ps-london-bw.perf.ja.net	BBRv3	IPv6	1500	1	7.03 Gb/s	3.86 Gb/s	-	-	207 ms	37.7 ms	17	16.9	9.32 ms	8.72 ms	0	0
ps-london-bw.perf.ja.net	BBRv3	IPv6	2048	1	11.3 Gb/s	8.18 Gb/s	-	-	229 ms	45.6 ms	17	17	9.43 ms	8.91 ms	0	0
ps-london-bw.perf.ja.net	BBRv3	IPv6	4096	1	27.3 Gb/s	18.3 Gb/s	-	-	249 ms	40.8 ms	17	17	9.49 ms	8.94 ms	0	0
ps-london-bw.perf.ja.net	BBRv3	IPv6	8192	1	31.1 Gb/s	25.5 Gb/s	-	-	237 ms	33.1 ms	17	16.9	9.54 ms	8.60 ms	0.002	0

Preliminary MSS testing: summary for IPv6

- BBRv1 and BBRv3 for MSS = 1500, 2048, 4096 and 8192
- For CERN hosts: we don't see any improvements by setting different MSS → why?
- For ESnet hosts: we don't see any improvements by setting different MSS → why?
- For JISC hosts:
  - For MSS = 4096 we can get 12 Gb/s more throughput in comparison with MSS = 2048 and 17 Gb/s more in comparison with MSS = 1500
  - For MSS = 8192 we can get 19 Gb/s more throughput in comparison with MSS = 2048 and 24 Gb/s more in comparison with MSS = 1500

BBR demo at DC24

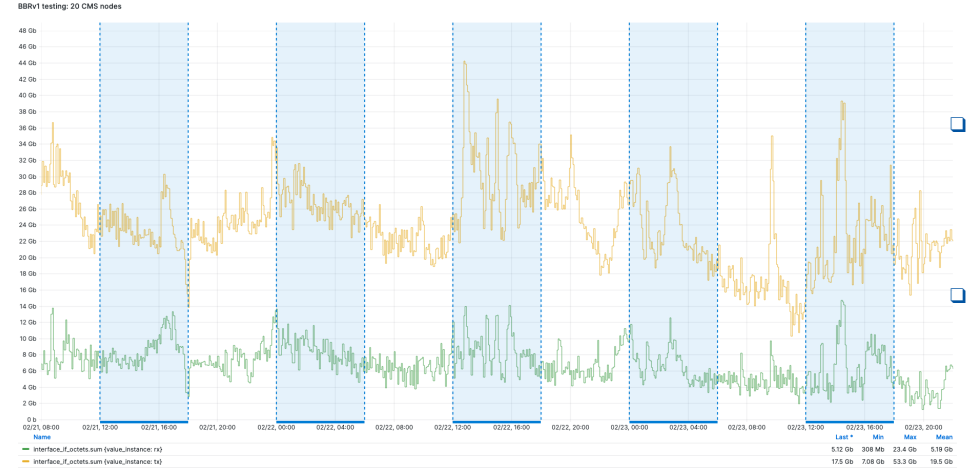

- During WLCG DC24 20 CMS nodes swapped every 6 hours between CUBIC and BBRv1 congestion protocol
- No evidence of gain nor loss using BBRv1

BBR demo at DC24

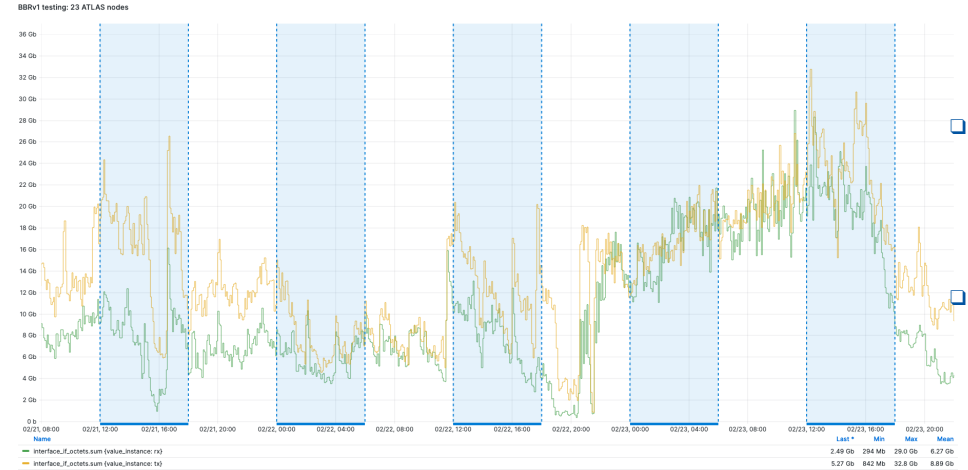

- During WLCG DC24 23 ATLAS nodes swapped every 6 hours between CUBIC and BBRv1 congestion protocol
- No evidence of gain nor loss using BBRv1

Thanks for your attention!


CERN