## Measuring Broadband America Special Projects: 1Gbps-WG Interim Update and Next Steps

March 18, 2021



Sam Knows Confidential Agenda



## Agenda

- Explore and contrast alternate throughput test methods at Gigabit speeds
- Summary of testing carried out so far
- TCP Cubic 16 Connections
- TCP BBR v1 (BBR v2 not yet supported on mainline Linux kernel)
- BBR v1 Investigation
- Further Testing
- Q&A

Sam Knows Confidential



## **MBA** Testing Completed so far

- 8 parallel connections TCP Cubic
  (For comparison purposes)
- 16 parallel connections TCP Cubic
- 8 parallel connections TCP BBR v1
- Investigation into results seen using BBR v1





4

## MBA Comparisons 8 to 16 thread TCP Cubic

Daily average, split by 7 anonymised ISPs, sample of ~250 units

- Results are derived from tests to the off-net Stackpath servers.
- Daily average; results are used from all tests that ran as scheduled using a 10 second download speed test.
- All tests use 8 parallel connections and TCP (Cubic) before changing to 16 parallel connections.
- Change from 8 to 16 connections on 11/16/2021.



FCC 1Gbps - ISP Representative Cubic 8/16 Thread Download Speed (Multi-Threaded Only) 16 Thread TCP Cubic 8. Nov 9. Nov 10. Nov 11. Nov 12. Nov 13. Nov 14. Nov 15. Nov 16. Nov 17. Nov 18. Nov 19. Nov 20. Nov 21. Nov 22. Nov 23. Nov 24. Nov 25. Nov 26. Nov 27. Nov Date/Time





## MBA Comparisons 8 to 16 thread TCP Cubic,

3 representative units from ISP X, daily average

- Results are derived from tests to the off-net Stackpath servers.
- Daily average; results are used from all tests that ran as scheduled using a 10 second download speed test.
- All tests use 8 parallel connections and TCP (Cubic) before changing to 16 parallel connections.
- Change from 8 to 16 connections on 11/16/2021.





## MBA Comparisons 8 to 16 thread TCP Cubic

*3 representative units from ISP Y, daily average* 

- Results are derived from tests to the off-net Stackpath servers.
- Daily average; results are used from all tests that ran as scheduled using a 10 second download speed test.
- All tests use 8 parallel connections and TCP (Cubic) before changing to 16 parallel connections.
- Change from 8 to 16 connections on 11/16/2021.



7

## MBA Comparisons 8 to 16 thread TCP Cubic,

1 representative unit from ISP Y, scatter plot of all test results

- Results are derived from tests to the off-net Stackpath servers.
- Scatter plot showing all results from all hours of the day are (not just peak) results from a 10 second download speed test.
- All tests use 8 parallel connections and TCP (Cubic) before changing to 16 parallel connections.
- Change from 8 to 16 connections on 11/16/2021.



Sam Knows





## **MBA Comparisons** TCP Cubic to TCP BBR v1

10 representative units from 3 ISPs, daily average

- Results are derived from tests to specifically provisioned offnet Stackpath servers.
- Daily average; results are used from all tests that ran as scheduled using a 10 second download speed test.
- Schedule changes to BBR v1 on 12/02/2020.



# FCC 1Gbps - Representative Units BBR Drop Download Speed (Multi-Threaded Only) 16 Thread TCP Cubic 8 Thread TCP BBR 16. Nov 17. Nov 18. Nov 19. Nov 20. Nov 21. Nov 22. Nov 23. Nov 24. Nov 25. Nov 26. Nov 27. Nov 28. Nov 29. Nov 30. Nov 1. Dec 2. Dec 3. Dec 4. Dec 5. Dec 6. Dec 7. Dec 8. Dec 8 Date/Time

A,	
4	
/	
Dec 9. Dec	



## **MBA** BBR v1 - further investigation

- Following the observed drop in results on BBR v1 we undertook a further investigation to establish the cause.
- We found that when using BBR v1, we saw high amounts of TCP retransmissions (sometimes over 15% of the entire transfer).
- We believe we are hitting the "high packet loss rates if bottleneck queue < 1.5\*BDP" issue documented with BBRv1 (see page 3 of https://datatracker.ietf.org/ meeting/104/materials/slides-104-iccrg-an-update-onbbr-00).
- In other words, we believe that this very high packet loss is caused by the server sending data at such a high speed that a bottleneck queue on a network device somewhere on the path is filling up, and BBR v1 is not reacting to the ensuing packet loss.
- BBR v2 resolves this, but is not yet supported in mainline Linux.



## MBA BBR v1 TCP\_INFO Logs

## CUBIC TCP\_INFO Logs (8 parallel TCP connections, for times 5000ms - 5200ms in a test):

CONN	TIME	SSTH	CWND	RETRANS	MSS	RTT	BYTES
<snip></snip>							
CONNO:	5100	23	24	109	1398	3675	5553834
CONN1:	5098	50	63	84	1398	3836	6933660
CONN2:	5096	62	63	74	1398	3367	9018078
CONN3:	5094	31	32	80	1398	3616	6587515
CONN4:	5092	27	28	85	1398	3182	8096237
CONN5:	5090	75	81	73	1398	3551	10174224
CONN6:	5088	49	50	106	1398	3569	6804485
CONN7:	5086	24	37	105	1398	3939	594052 <sup>°</sup>
CONNO:	5200	23	32	109	1398	4488	5630165
CONN1:	5198	50	69	84	1398	4098	7139166
CONN2:	5196	62	70	74	1398	4178	9223584
CONN3:	5194	31	40	80	1398	4181	6704947
CONN4:	5192	27	34	85	1398	4088	818431 <sup>°</sup>
CONN5:	5190	75	87	73	1398	4188	1043844
CONN6:	5188	37	39	109	1398	4242	6921917
CONN7:	5186	24	44	105	1398	4692	6034467
CONNO:	5300	25	26	111	1398	3334	572411
<snip></snip>							

## BBR TCP\_INFO Logs (8 parallel TCP connections, for times 5000ms - 5200ms in a test):

CONN	TIME	SSTH	CWND	RETRANS	MSS	RTT	BYTES
<snip></snip>							
CONNO:	5100	150	396	1933	1398	25127	53431454
CONN1:	5098	134	118	479	1398	23140	25127652
CONN2:	5096	74	176	1558	1398	29312	38750828
CONN3:	5094	120	612	2318	1398	23047	59448978
CONN4:	5092	116	124	1205	1398	25229	42994624
CONN5:	5089	116	38	335	1398	23278	19495110
CONN6:	5087	126	168	1816	1398	25191	61373720
CONN7:	5085	108	401	3314	1398	22677	102900034
CONNO:	5200	150	154	2326	1398	22198	55528606
CONN1:	5198	134	60	589	1398	23205	26008392
CONN2:	5196	74	144	1779	1398	22611	39799404
CONN3:	5194	120	208	2704	1398	21524	61546130
CONN4:	5192	116	129	1345	1398	23660	44043200
CONN5:	5189	116	36	374	1398	22478	19495110
CONN6:	5187	126	164	1985	1398	20984	62422296
CONN7:	5185	108	364	3718	1398	22441	104997186
CONNO:	5300	150	154	2365	1398	5012	58674334
CONN1:	5298	134	134	589	1398	5518	26830416
<snip></snip>							



## Further Testing

The experimentation phase of the investigation is ongoing.

The next phase of the workshop is to continue on to these tests:

- Using multiple servers in parallel
- SK UDP speedtest
- UDPST

There will be another TWG meeting to review findings when all experiments are completed.



![](_page_12_Picture_0.jpeg)

Sam Knows Confidential

![](_page_12_Picture_2.jpeg)

![](_page_13_Picture_1.jpeg)