

Experiments with FQ in a Science DMZ environment

Brian Tierney, ESnet
bltierney@es.net

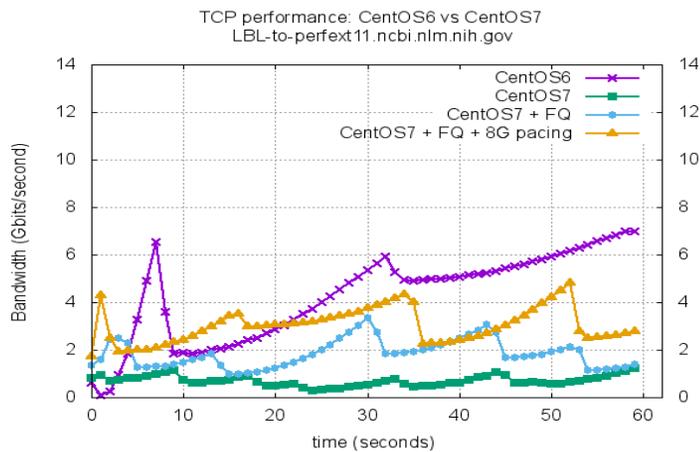
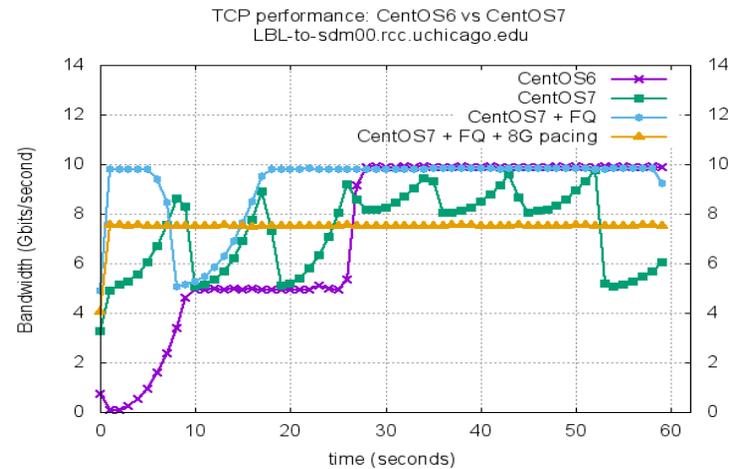
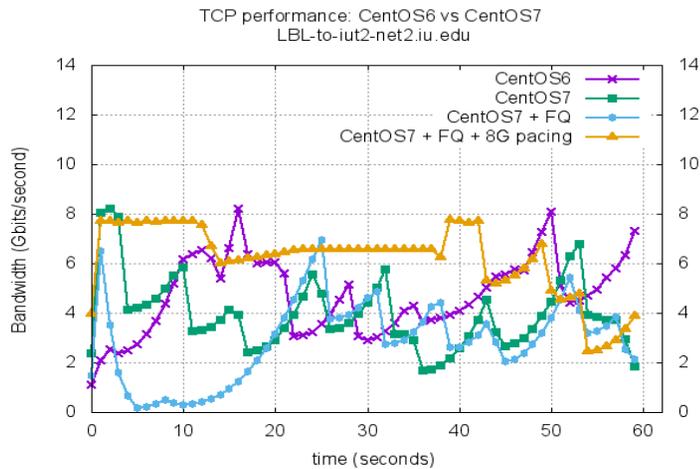
New TCP option: Fair Queuing Scheduler (FQ)

- <http://fasterdata.es.net/host-tuning/linux/fair-queuing-scheduler/>
- Available in Linux kernel 3.11 or higher
 - available in Fedora 20, Debian 8, and Ubuntu 13.10
 - Backported to 3.10.0-327 kernel in v7.2 CentOS/RHEL
- To enable Fair Queuing (which is off by default), do:
 - `tc qdisc add dev $ETH root fq`
- Or to both pace and shape the bandwidth:
 - `tc qdisc add dev $ETH root fq maxrate Ngbits`

Experiment Summary

- Testing from a CentOS 7.2 host to a number of public perfSONAR hosts
- Did both Single stream and 4 stream tests

Observation 1: CentOS 7 +FQ often helps; adding pacing can help even more



CentOS7 + FQ usually has higher overall tput

Adding maxrate of 8G is usually more stable

Parallel Stream Tests

- Most high performance data transfer tools use parallel streams
- Experiment:
 - What is the impact of FQ and pacing on 4 streams

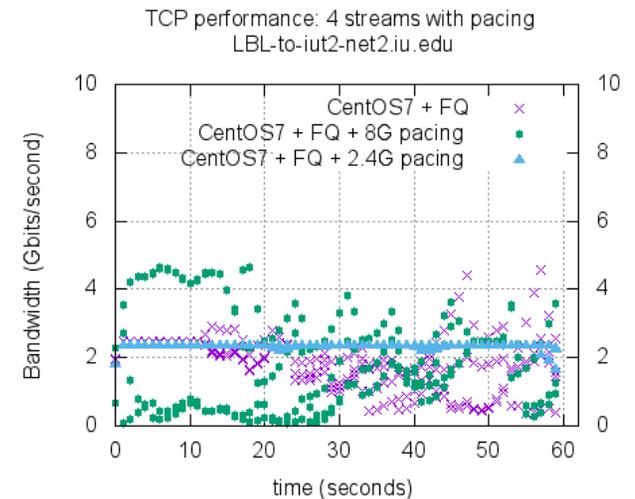
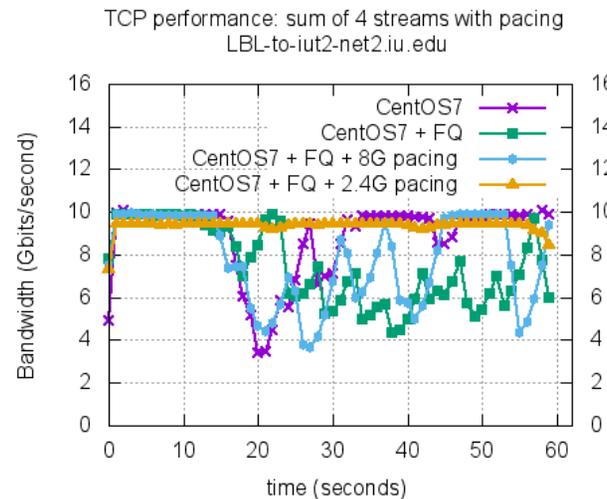
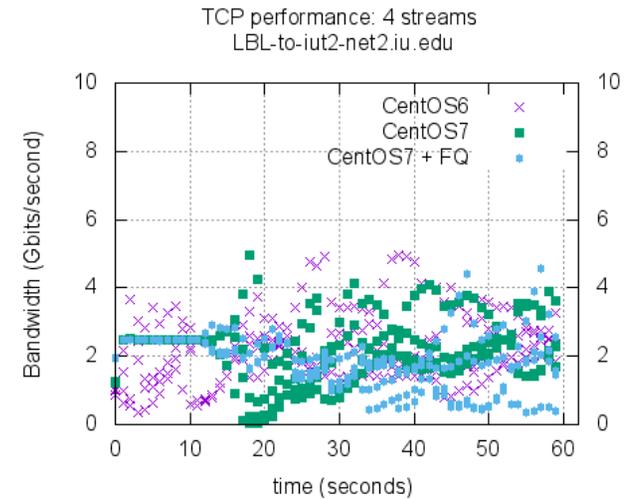
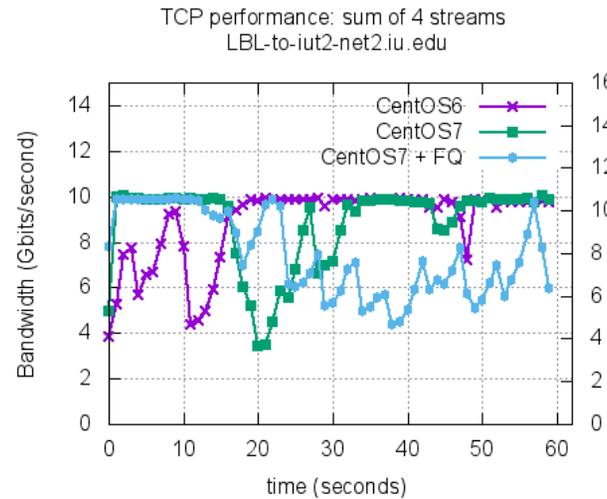
Parallel Stream Test 1

Left side:
sum of 4 streams

Right side:
tput of each stream

Streams appear to be much better balanced with FQ, pacing to 2.4 performed best

Parallel Stream Testing



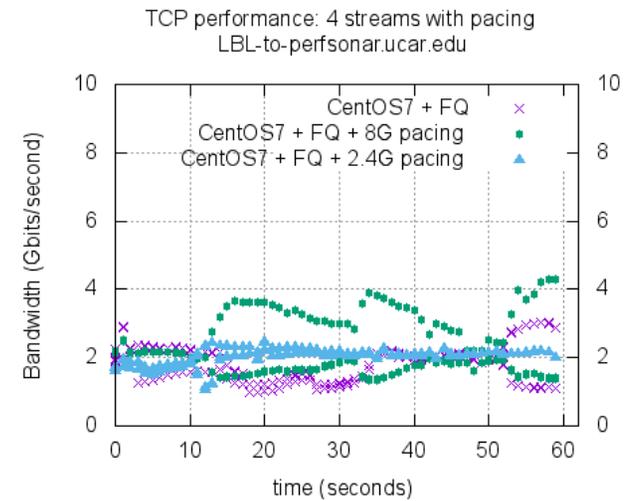
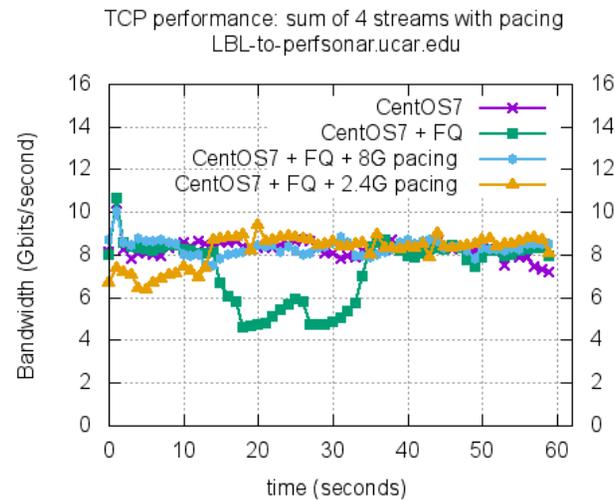
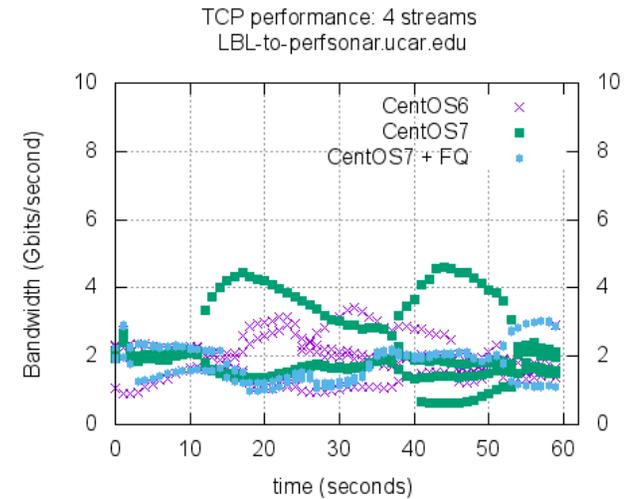
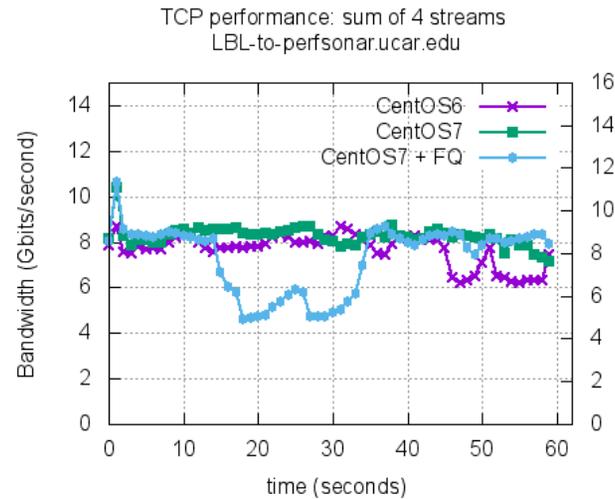
Parallel Stream Test 2

Left side:
sum of 4 streams

Right side:
tput of each stream

Streams appear to be much better balanced with FQ

Parallel Stream Testing



But...

- In some cases the results are less clear....

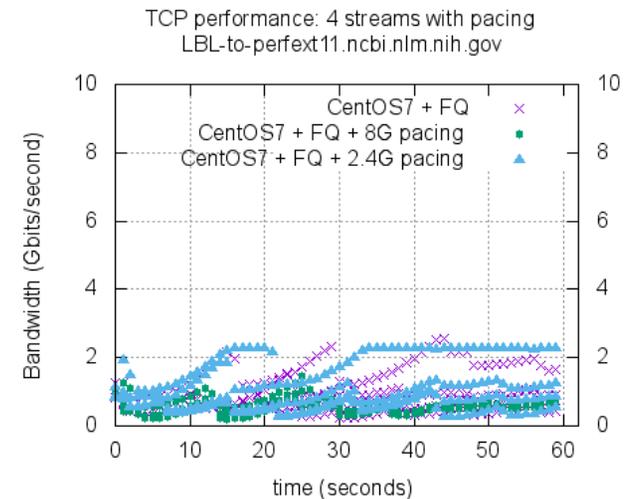
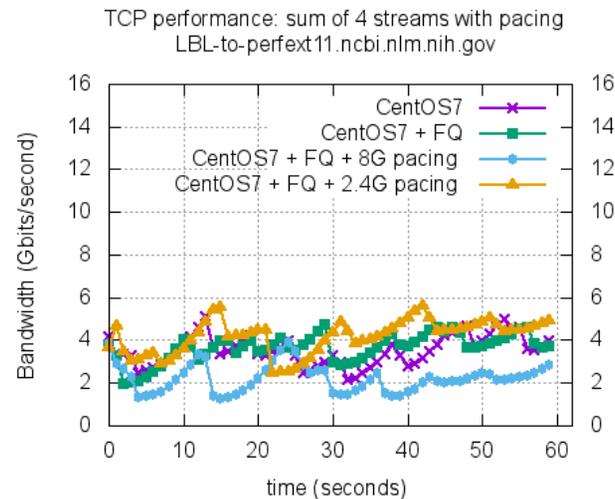
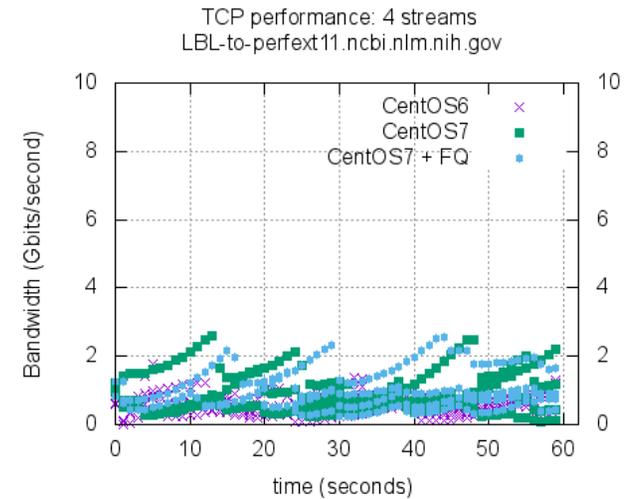
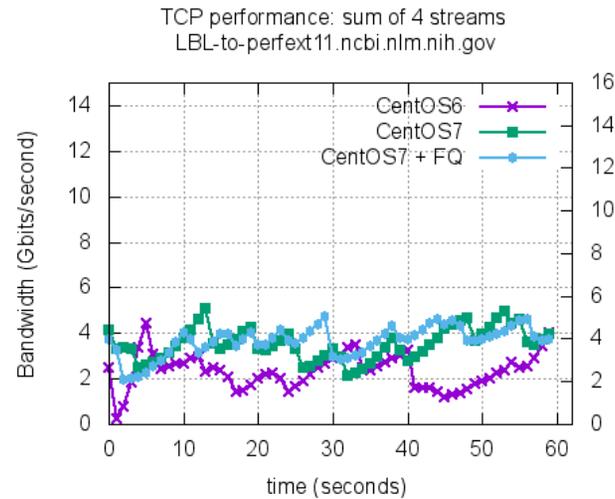
Parallel Stream Test 3

Left side:
sum of 4 streams

Right side:
tput of each stream

FQ + pacing give best throughput, but streams are not very well balanced

Parallel Stream Testing



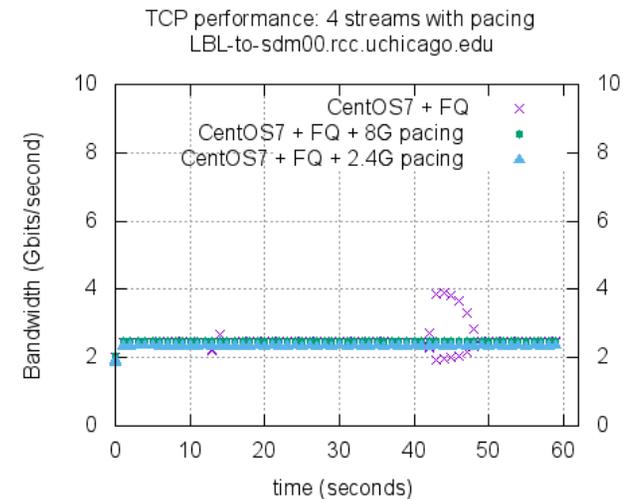
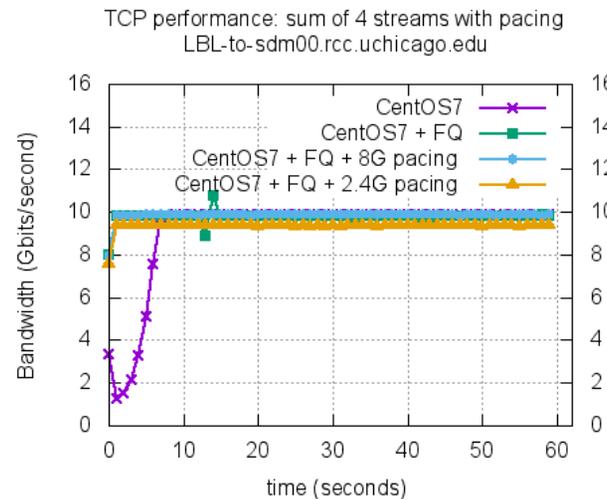
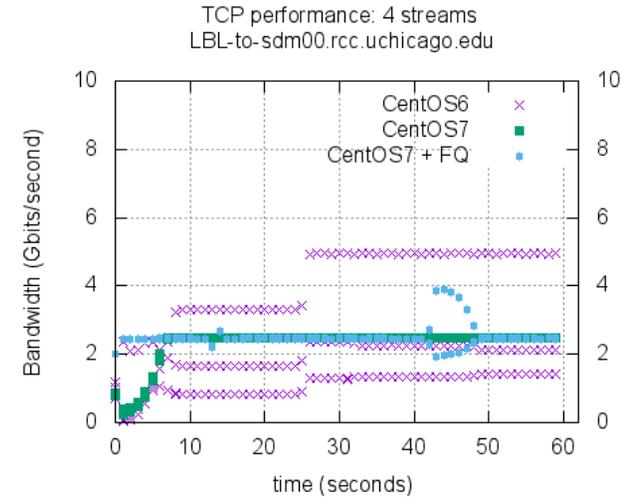
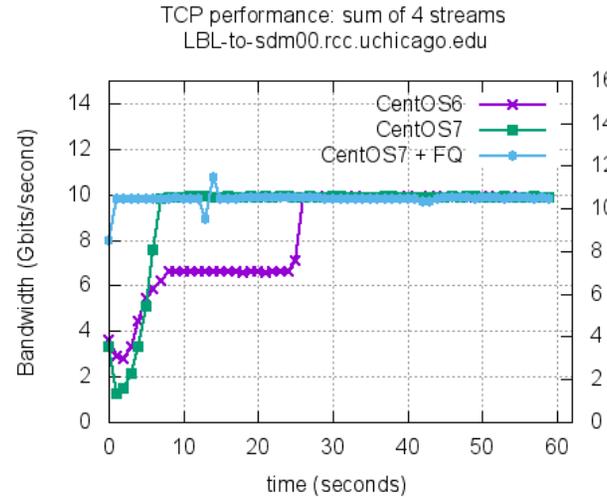
Parallel Stream Test 4

Left side:
sum of 4 streams

Right side:
tput of each stream

Streams appear to be much better balanced with FQ, but no clear advantage from pacing

Parallel Stream Testing



Disclaimers

- Result here seem to be true in general, but its easy to find counter examples

