

# Updates on Windows TCP

Praveen Balasubramanian

[pravb@microsoft.com](mailto:pravb@microsoft.com)



# Recap and deployment update

- Recap from Chicago IETF
  - TFO (TCP Fast Open) enabled by default in the Edge browser in Creators Update
  - Experimental support for CUBIC
  - RACK and TLP enabled for > 10 msec RTT connections
  - LEDBAT++ being used for internal workloads like crash dump uploads
    - \*in case you missed iccr, the ++ portions were presented there
- Fall Creators Update rolling out worldwide as a free update to Windows 10 users
- Server 2016's 1709 update available for download

# TFO deployment issues

- Fallback algorithm not aggressive enough
- Failure modes impact user experience
- Rolled back on Creators Update after initial ramp due to user reported issues from multiple geographies
  - “Edge should not yet have enabled by default, the tcp fast open. “
  - “Edge and Store Apps can't connect to Google services “
  - “YouTube Website NOT Loading on Edge but it loads on chrome and other browsers”
  - “Can't reach this page... DNS error on Facebook, Google, Youtube... happened after Creators Update”
  - “I can't get to Google sites such as Youtube or Gmail on builds 14977 and 14986 “

# TFNO - The nefarious middlebox

- Built TFNO using WFP (Windows Filtering Platform)
- Supports the following (mis)behaviors
  - drop all TFO segments
  - strip SYN data
  - drop SYN segments with data
  - blackhole TFO connections after they are established
  - blackhole all connections on <src IP, dst IP> when TFO is used
  - blackhole data from the server after a TFO connection is established
  - drop SYN+ACKs that acknowledge SYN data
  - delay all data after a TFO connection is established
  - blackhole the TFO client src IP after a SYN with a TFO option is seen
  - selectively drop SYN segments

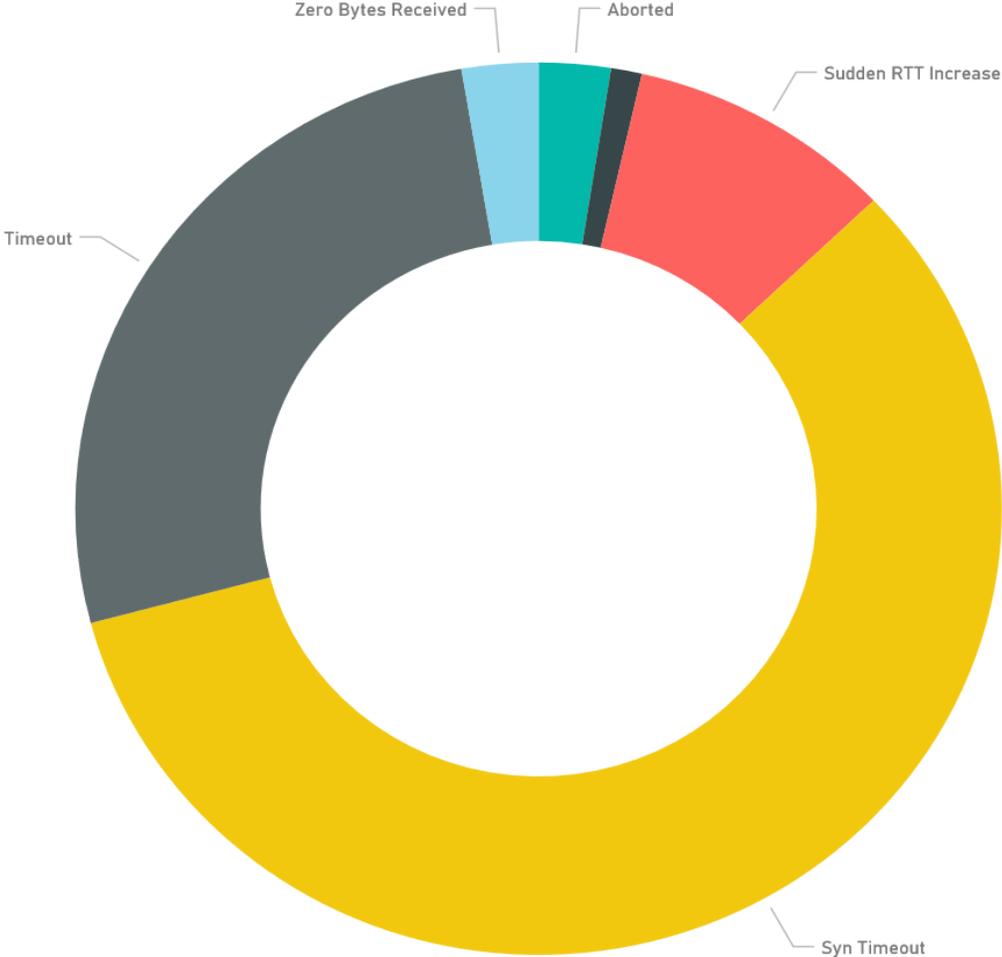
# TCP Fast Open – new fallback algorithm

- Limited passive probing
- Probing is limited to “Internet connected” networks
- TFO probing needs multiple connections to same server
- Allow only one probe TFO connection to proceed at a time
- When the probe connection is closed, mark success if all of the following hold
  - no RST in response to SYN
  - no SYN timeout
  - no full data timeout
  - data exchanged in both directions
  - connection wasn't cancelled
  - no sudden RTT increase
- If a successful probe connection succeeded *exercising* cookie, success
- If a network hits fallback, persist & never attempt again
- If a network hits success, stop probing & remember for boot session
- Fall Creators Update – first retail release of TFO without rollback

# TCP Fast Open – Some preliminary data

- Around 26% devices successfully used TFO and did not fallback
- A/B test result - No measurable increase in navigation failures
- Failures are correlated with geography
- Failures are correlated with specific networks
- SYN timeout heuristic makes fallback very aggressive
- Future
  - Fallback only if subsequent SYN (without option) succeeds
  - Increase coverage by experimenting with removing fallback criteria
  - Work with network operators to improve success rates

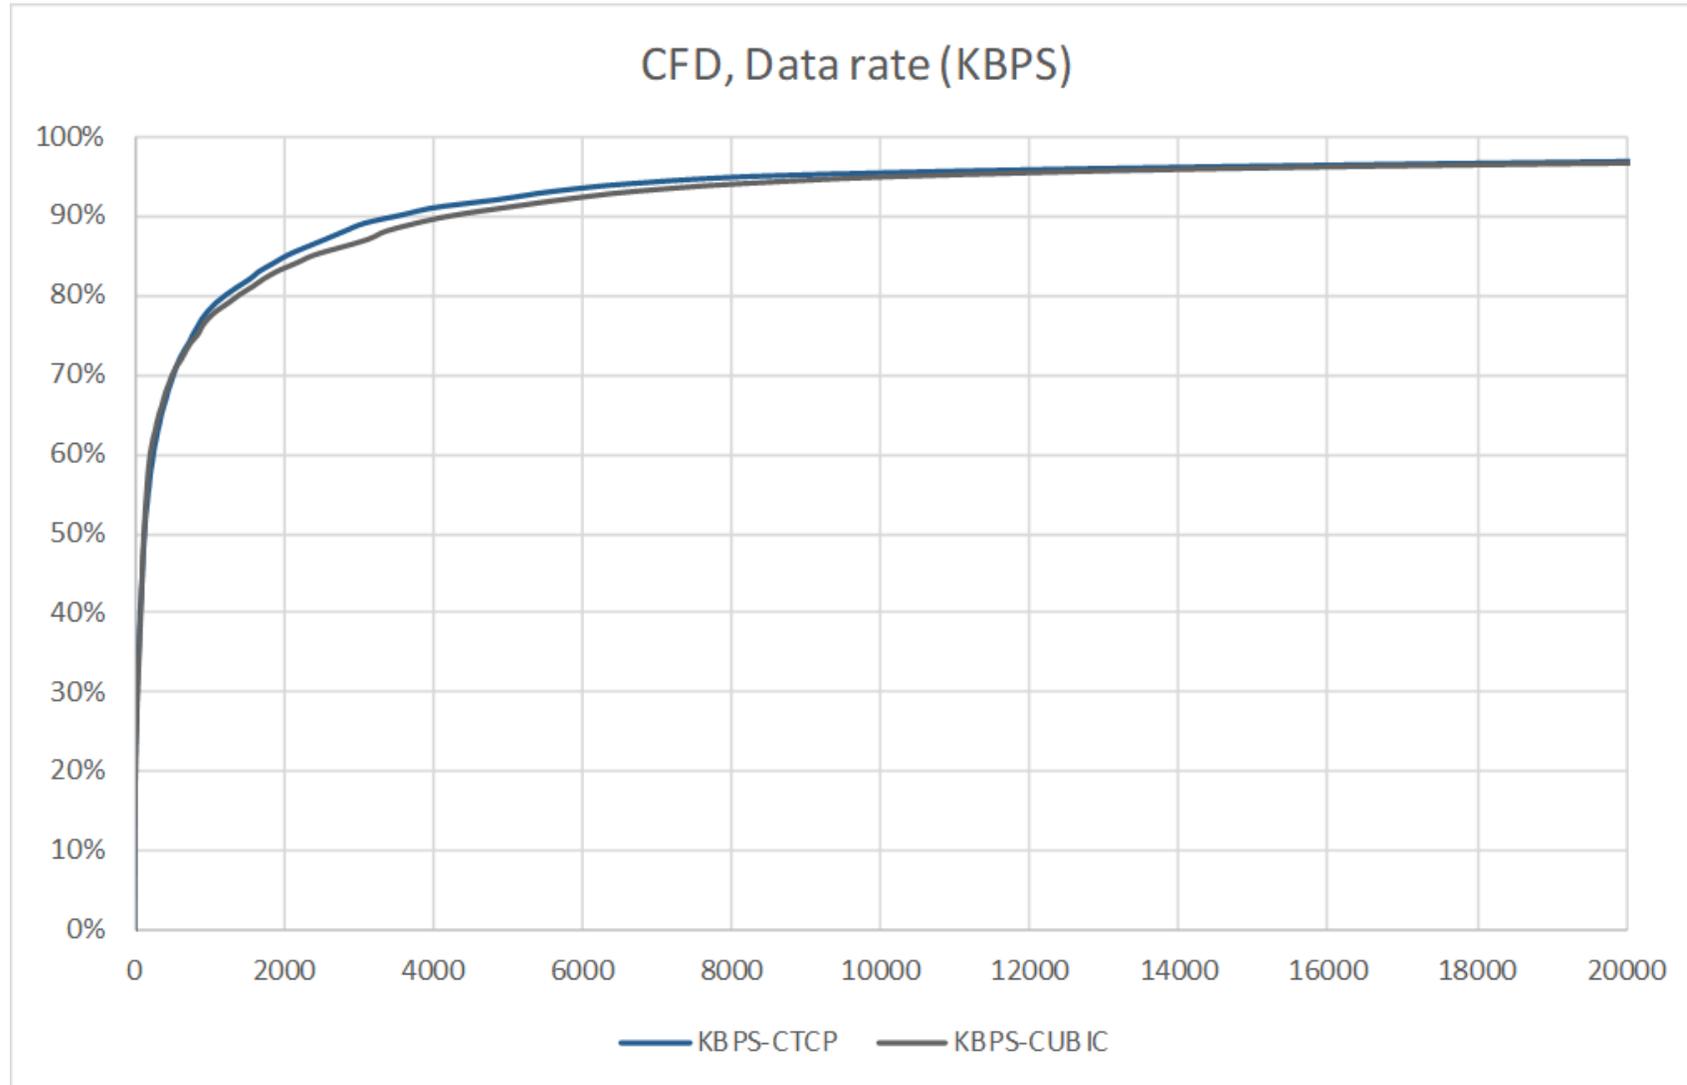
# TCP Fast Open – Probe Failure reasons



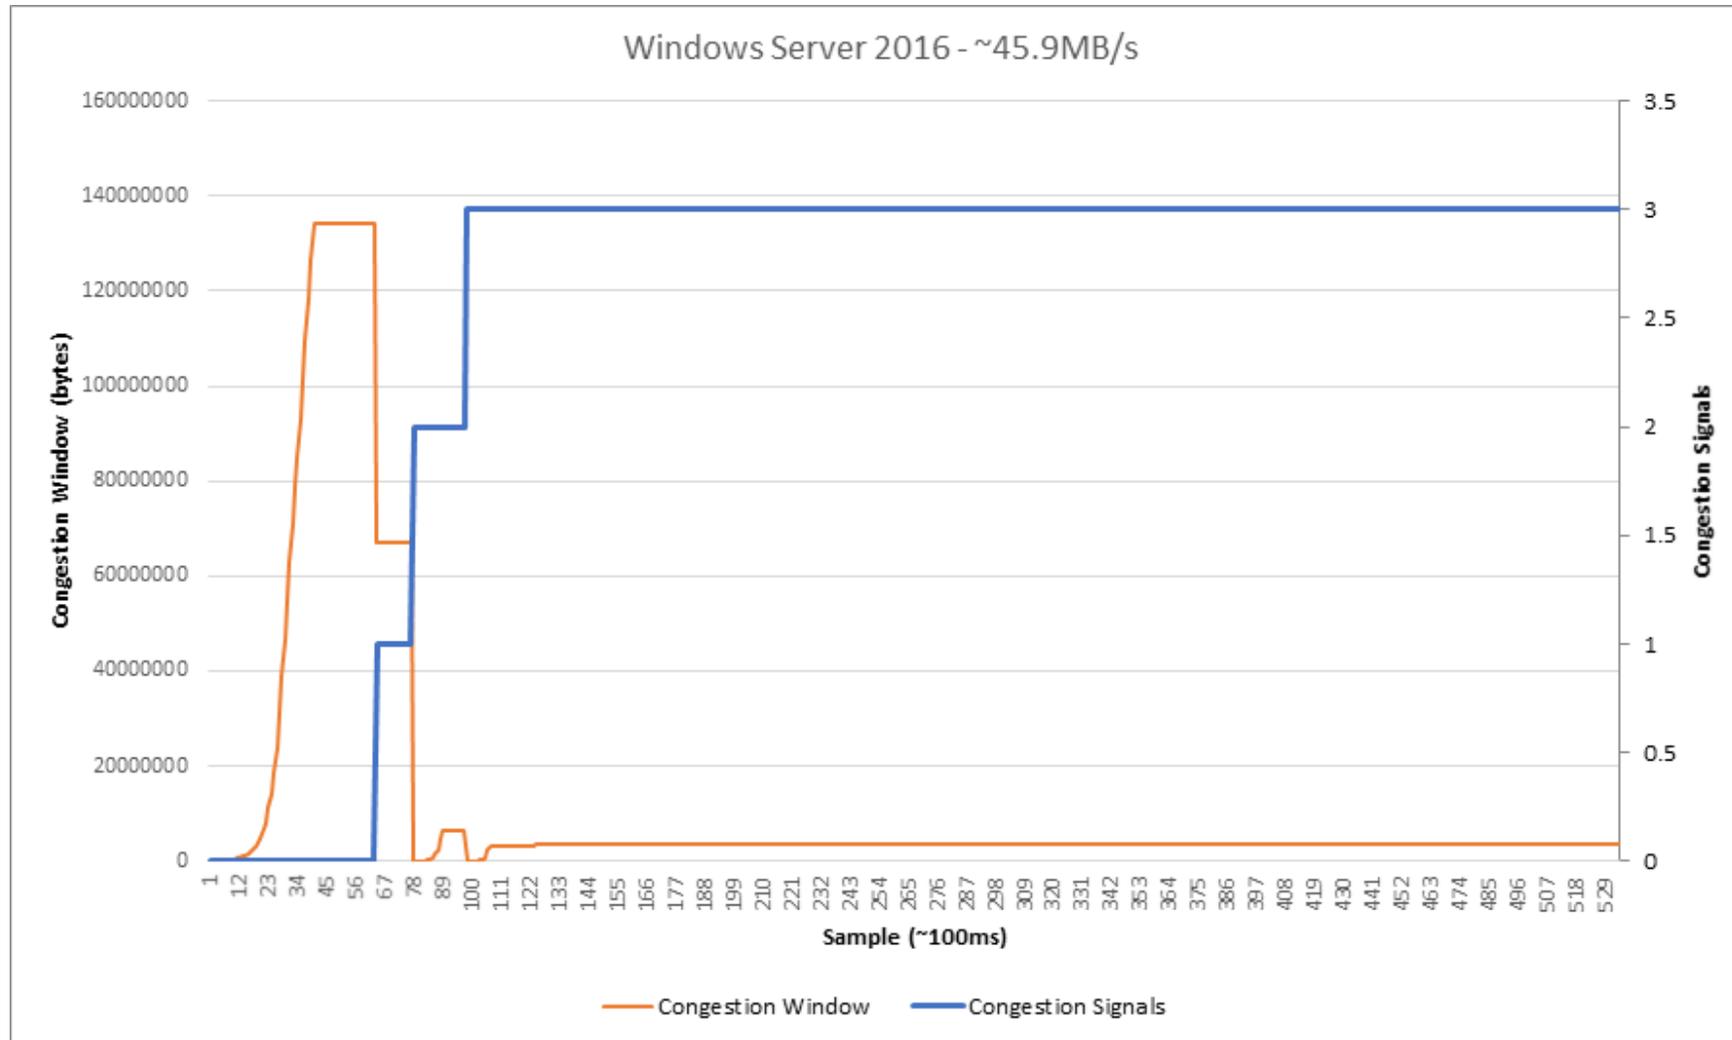
# CUBIC on by default

- Compound TCP sensitive to delay fluctuations
  - Bimodal latency distribution
  - Exacerbated by virtual networking
- Switching to CUBIC default
- Implementation based on draft-ietf-tcpm-cubic
- Fall Creators Update
  - All connections
- Windows Server 1709
  - Internet connections (>10 msec handshake RTT)

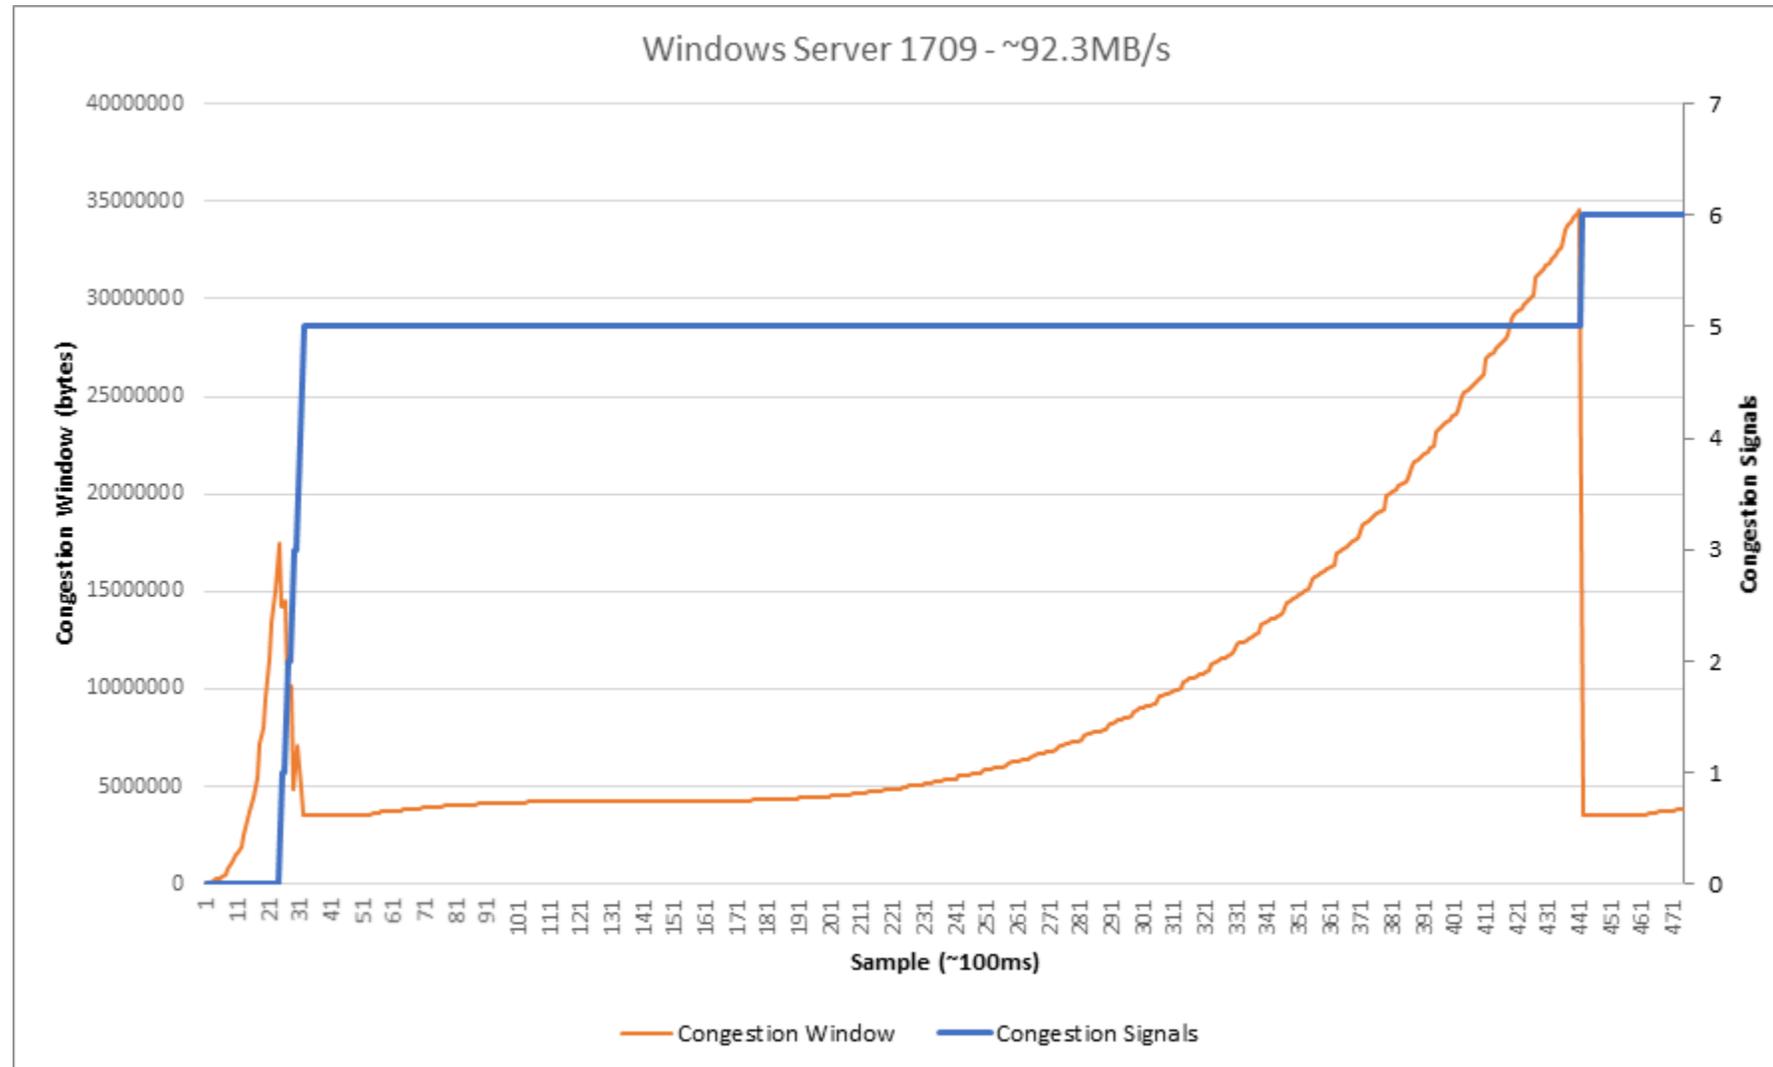
# Marginal improvement for client uploads



# CTCP (Azure US west to east)



# CUBIC (Azure US west to east)



# Q&A