Optimizing TLS for High-Bandwidth Applications in FreeBSD

R. Stewart[1] J. Gurney[2] S. Long[1]

[1]Netflix

[2]Consultant, @encthenet

3 October 2015 / EuroBSDcon 2015

Why is TLS Necessary

AT&T Switching Center, 611 Folsom Street by Remo Corso ⓒⓘⓢ from wikimapia.org

- The videos are encrypted, but metadata maters
- Attacks are ongoing against privacy:
  - Government Spying (Klein, Snowden)
  - Behavioral monitoring: viewing habits, purchases
- and Content Integrity:
  - Comcast — Xfinity WiFi Ad Injection
  - Verizon — User unique cookie inserted

Netflix OpenConnect Appliance

- Intel 64bit Xeon CPUs
- FreeBSD 10-stable and Nginx 1.5
- High Bandwidth – 10Gbps - 40Gbps
- Storage – 10TB - 120TB
- Connections – 10-40k long lived TCP

Classic Webserver

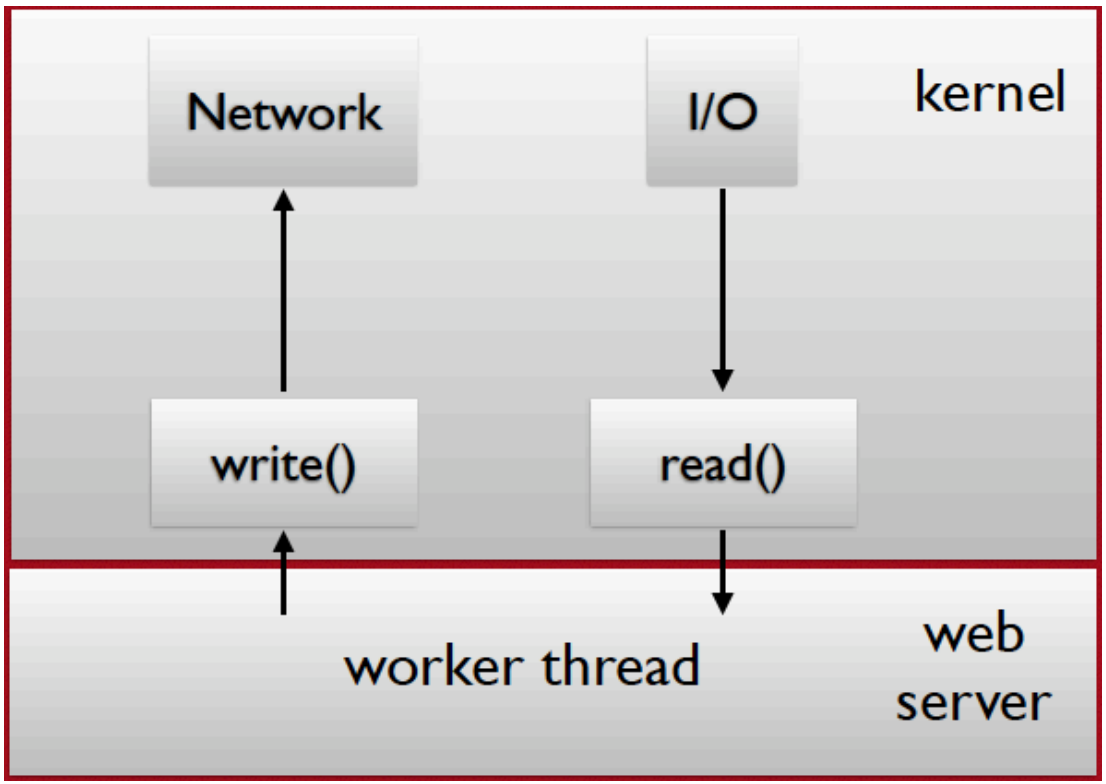

Sendfile Enabled Webserver

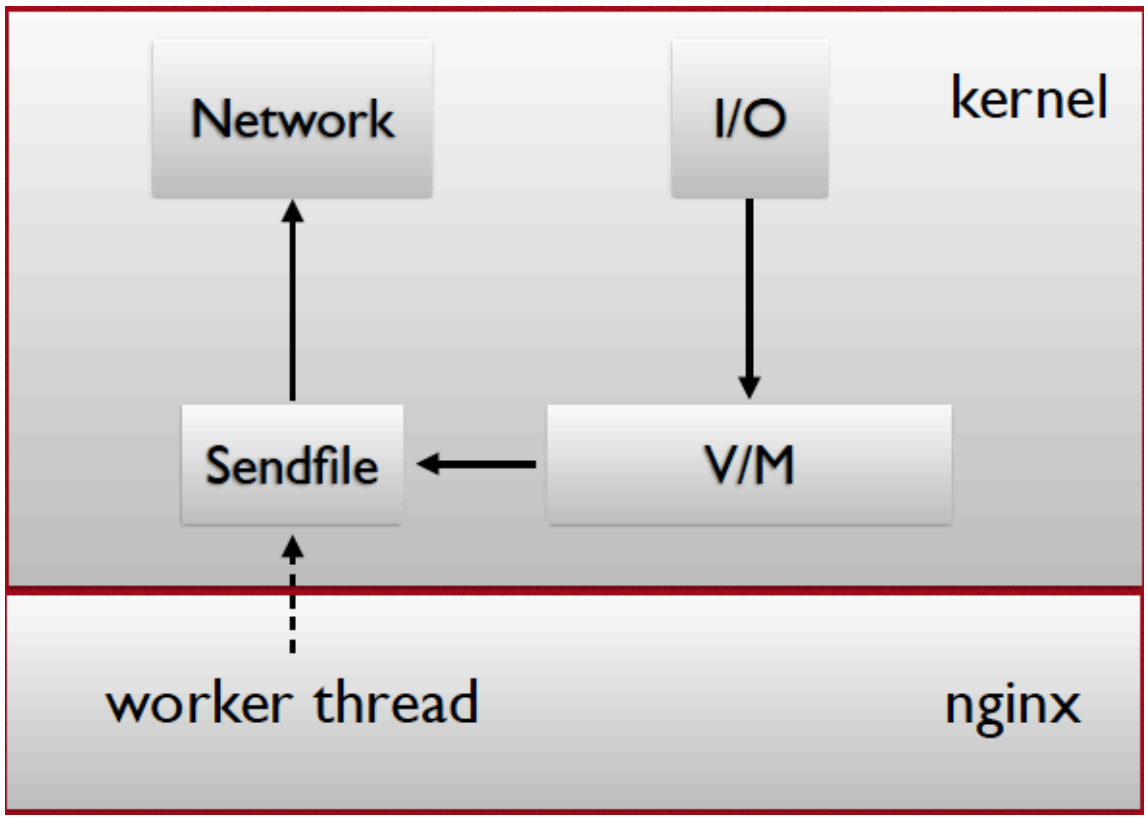

Classic TLS Webserver

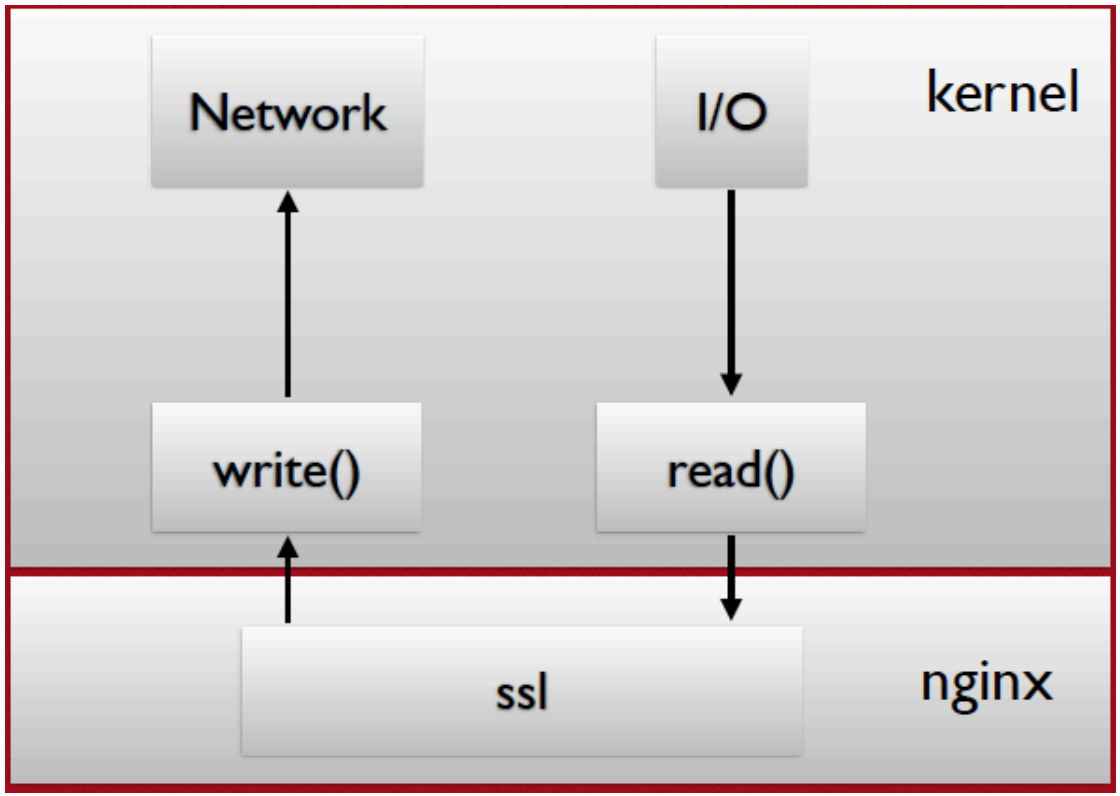

Observations

- sendfile improved performance by 3x
- Async sendfile added an additional 33% improvement
- Adding TLS dropped performance by 2.5-3x
- TLS on traditional, non-sendfile web servers saw a performance drop of 2x

In Kernel TLS Webserver

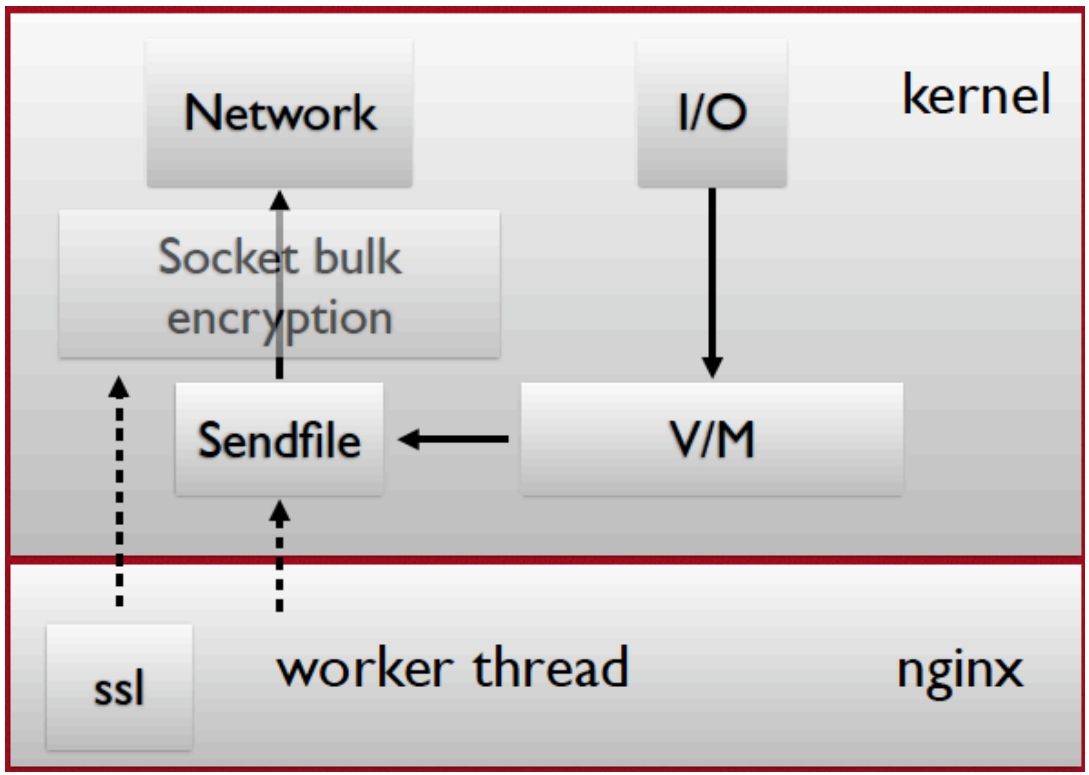

TLS Handshake

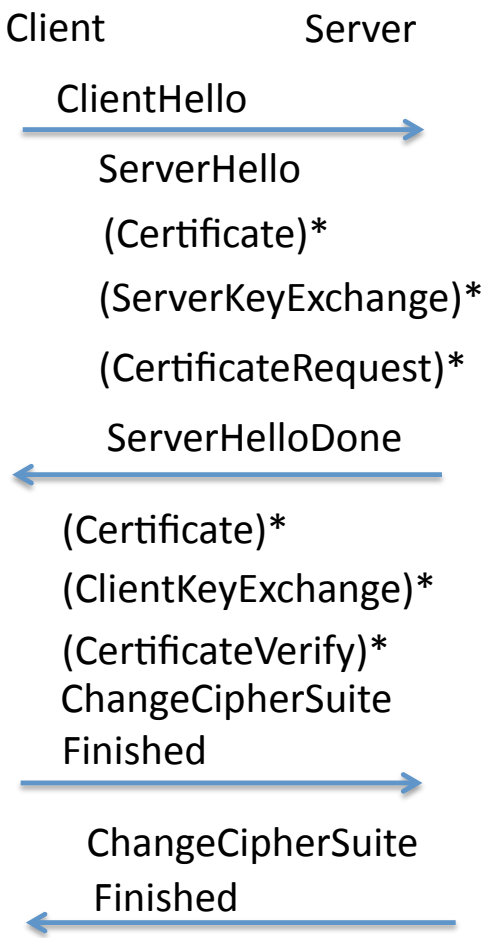

TLS Changes

- Only do encryption in kernel. Certificates, key exchange and decryption in user land.
- Unsupported ciphers work by using traditional, non-sendfile path.
- Two socket options added: Cipher Supported and Cipher Keys.

Early Results on OCA Appliance

No TLS
19-20Gbps

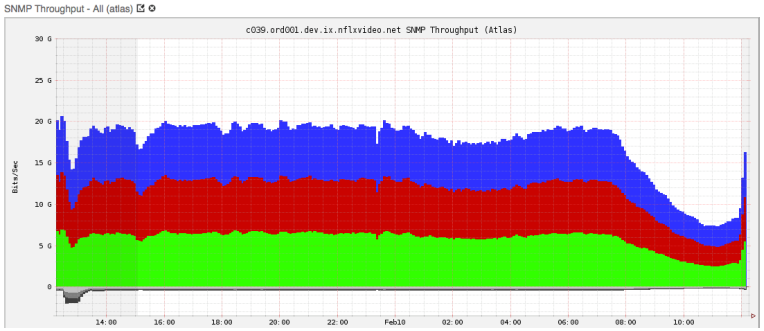

OpenSSL
8.5Gbps

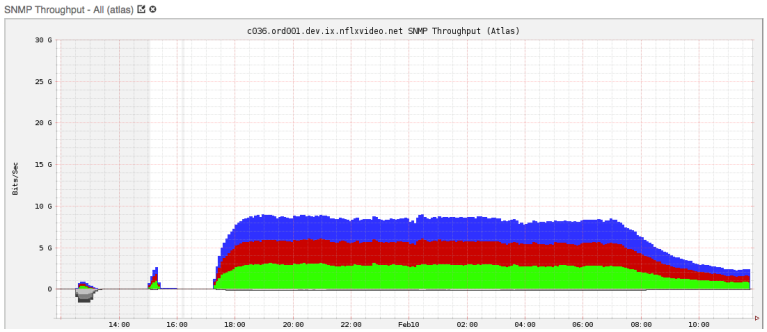

TLS w/ sendfile
9Gbps

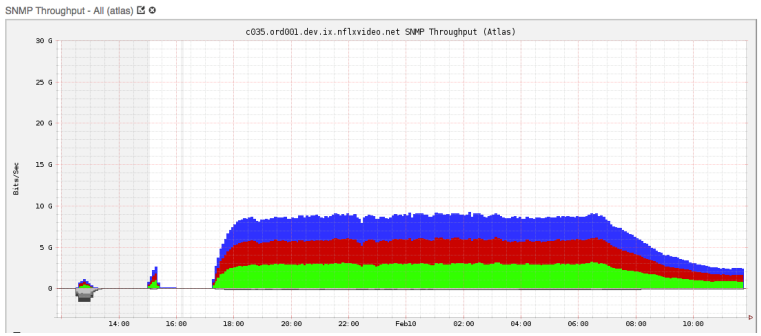

Latest Results

Benchmarking on production machines has advantages and disadvantages

- Real Work Loads
- Unpredictable loading

Rev D Disk Appliance 1 hour

OpenSSL
6-12Gbps

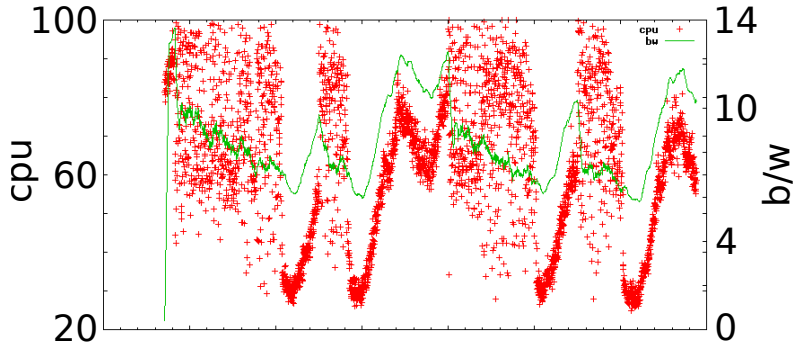

OCF
16-17Gbps

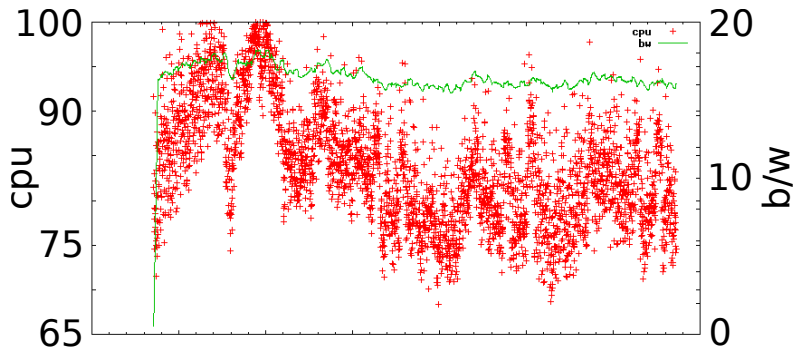

BoringSSL
17-18Gbps

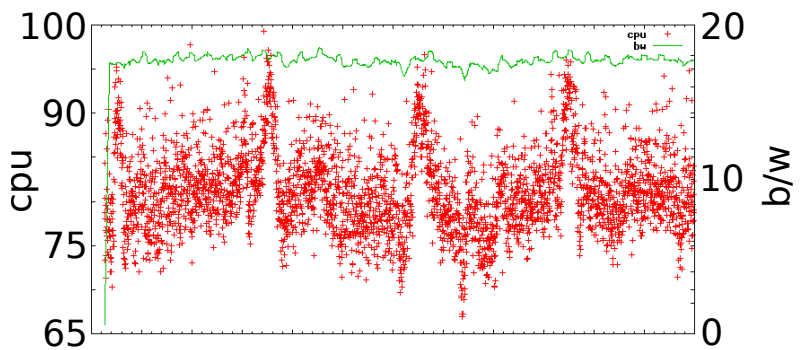

Flash Cache Appliance

OpenSSL
7-16Gbps

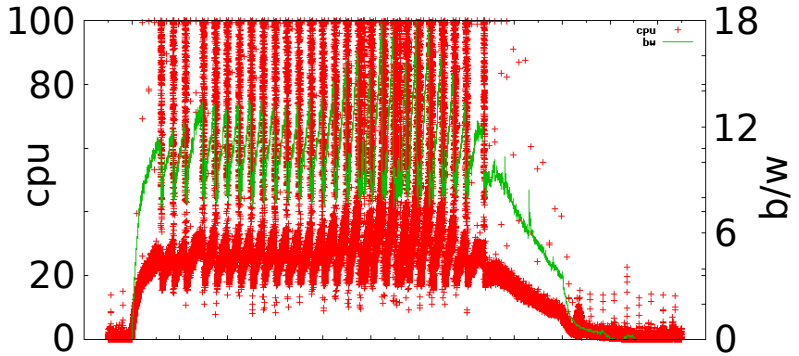

OCF
25Gbps

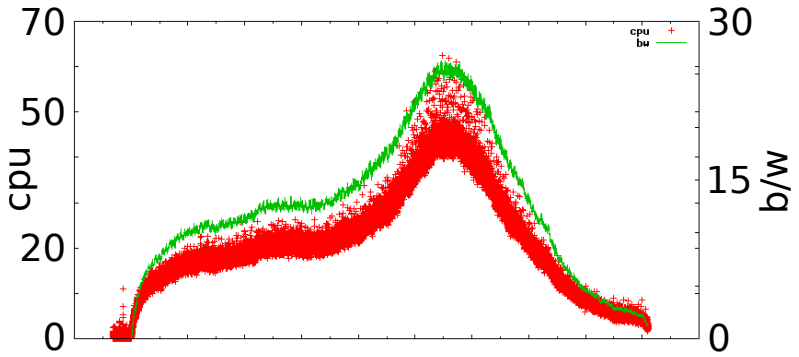

BoringSSL
26Gbps

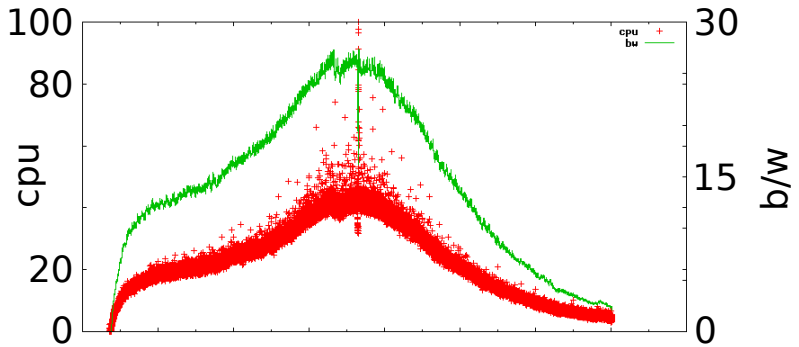

Future

- Support decryption, and automatic framing/deframing of data
- Investigate possibility of integrating into mainline FreeBSD

Questions?