Optimizing TLS for High-Bandwidth Applications in FreeBSD

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

[1]Netflix

[2]Consultant, @encthenet

2015 March 15 / AsiaBSDCon 2015

Why is TLS Necessary

- The videos are encrypted, but metadata maters
- Attacks are ongoing against privacy:
  - Government Spying (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 10.1 and Nginx 1.5
- High Bandwidth – 10Gbps - 40Gbps
- Storage – 10TB - 120TB
- Connections – 10-40k long lived TCP

Classic Webserver

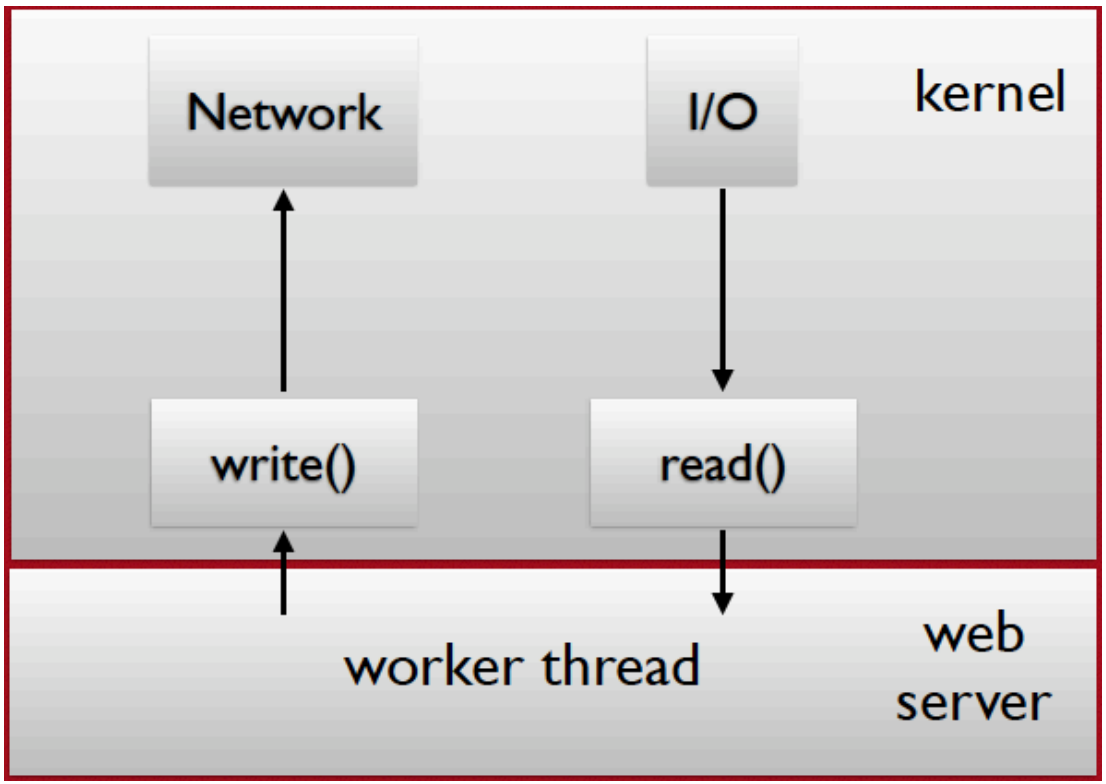

Sendfile Enabled Webserver

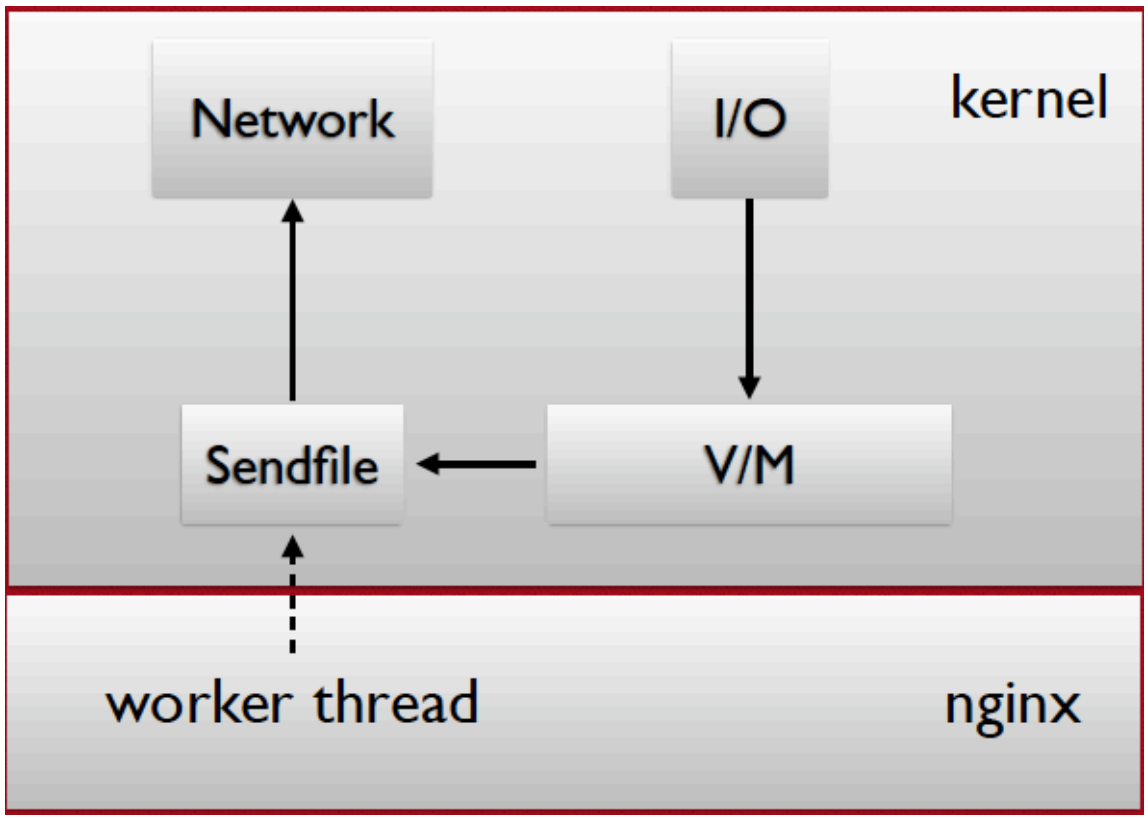

Classic TLS Webserver

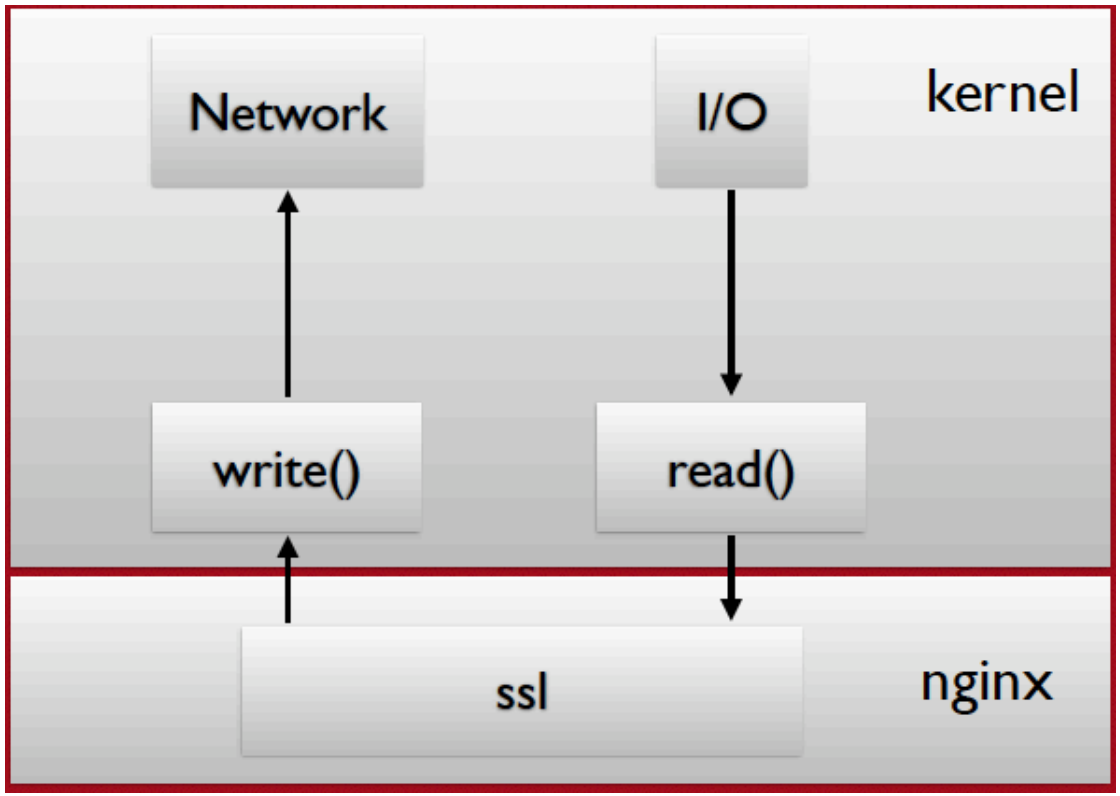

Observations

- 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

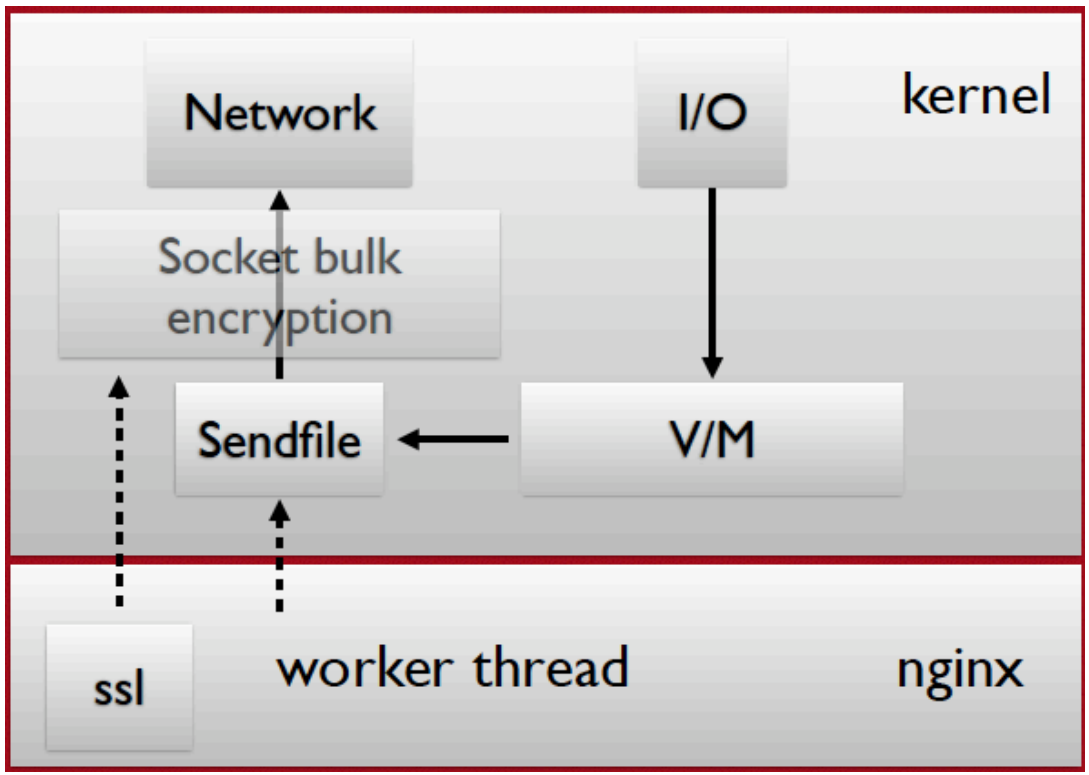

Changes to FreeBSD Open Crypto Framework

- Session Handling
- AES-NI Improvements
- AES-GCM

TLS Handshake

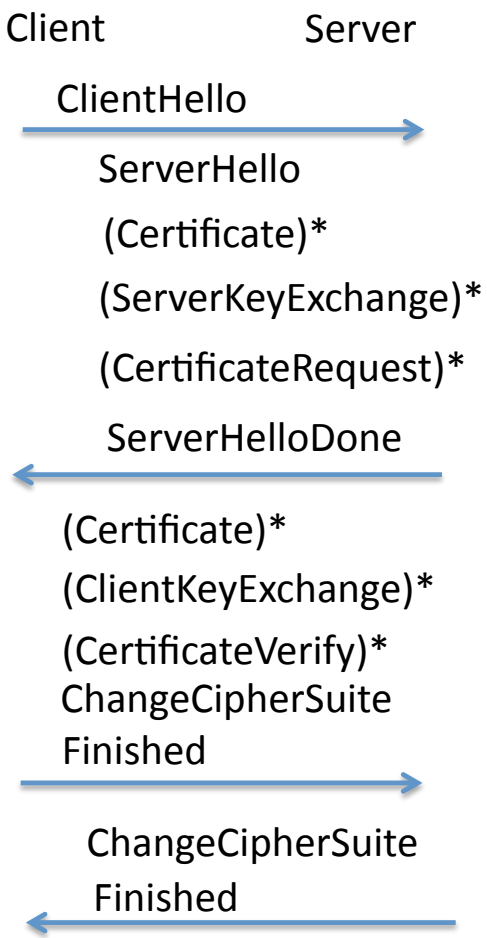

TLS Changes

- 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.

OCA Appliance: No TLS

19-20Gbps

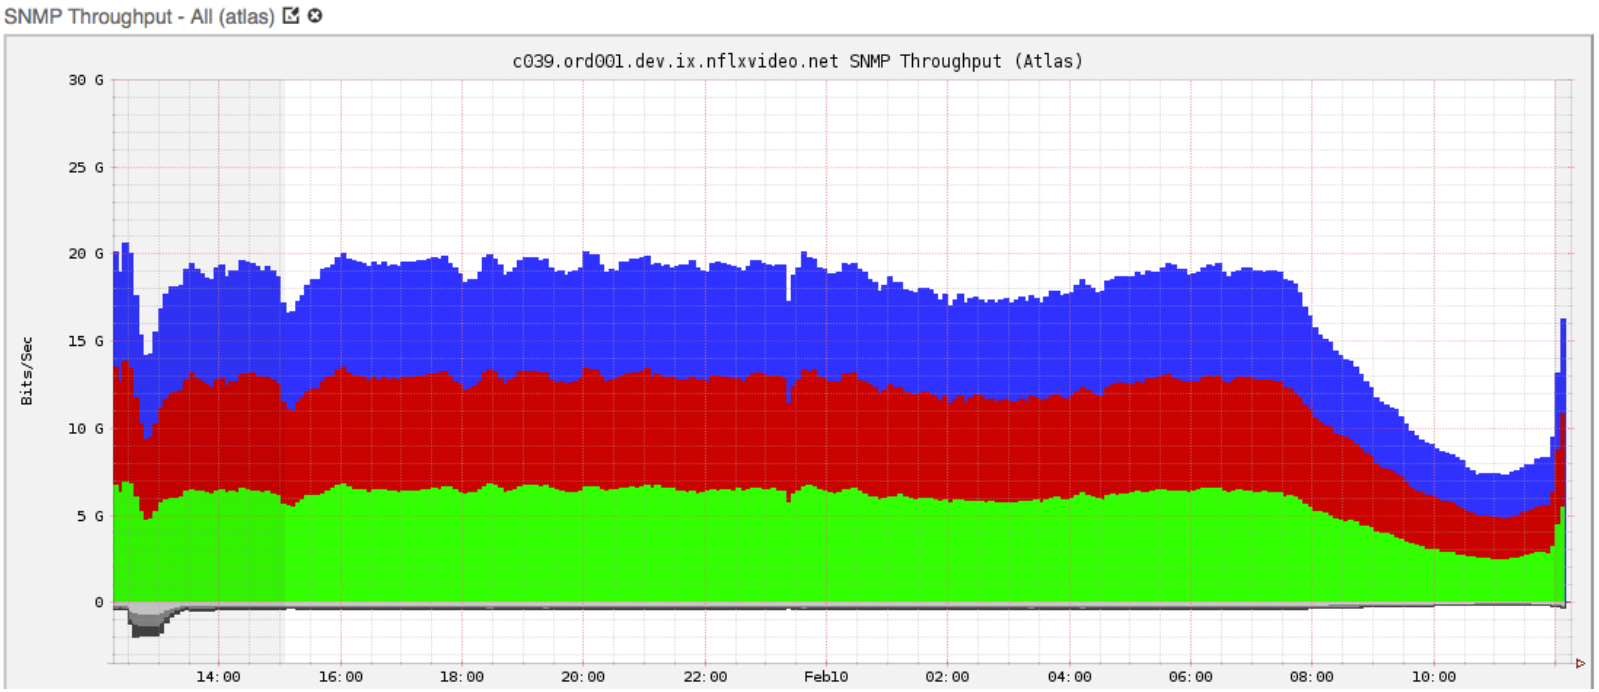

OCA Appliance: OpenSSL

8.5Gbps

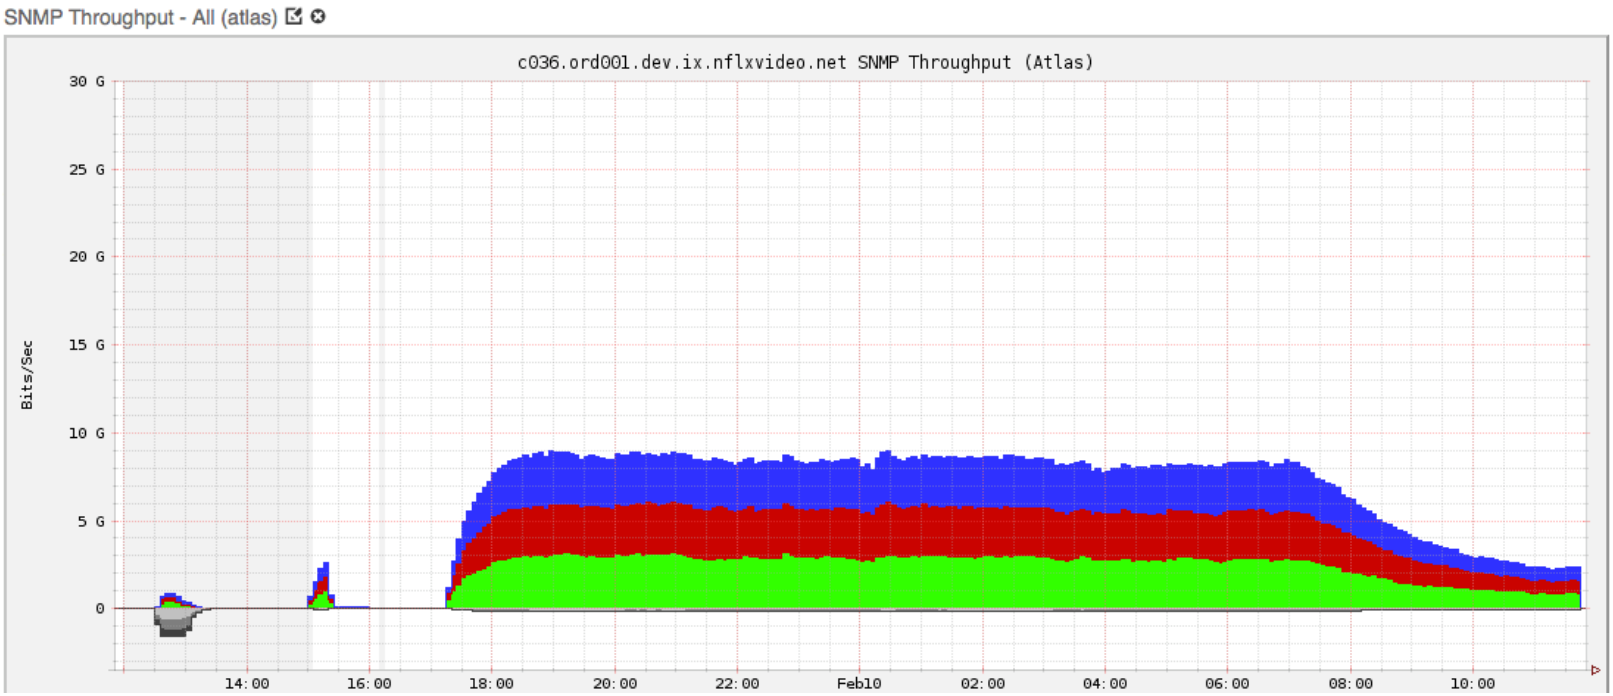

OCA Appliance: TLS w/ sendfile

9Gbps

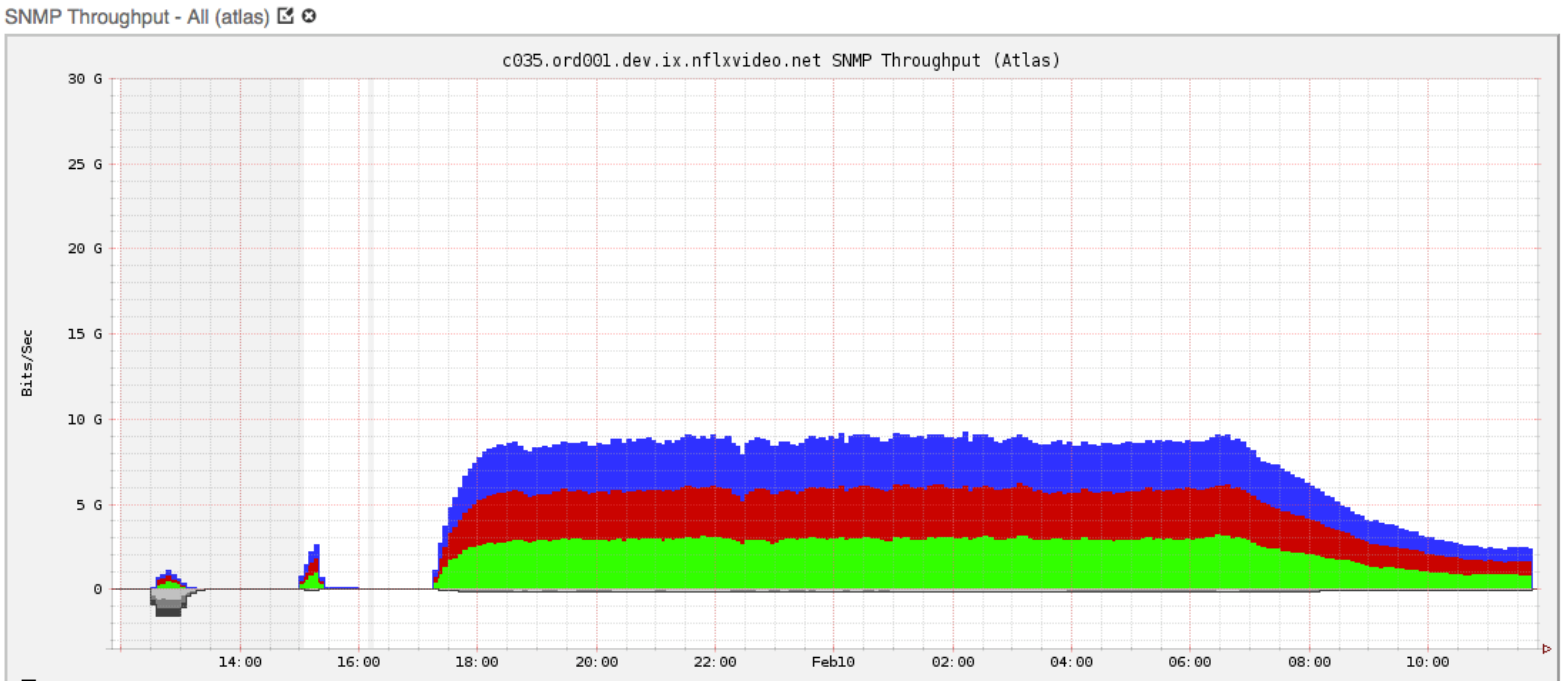

Summary

- Using TLS in the kernel only provides moderate gains.
- BUT...

Areas for future improvement

- In kernel SHA code has not been optimized which is used with most ciphers except AES-GCM
- Combine SHA256 and AES-CBC to hide the latency of AES-CBC
- Extra copy is made instead of doing direct to buffer processing
- Optimize FPU save/restores
- Pipelining: Process multiple streams at once for AES-CBC and SHA-256
- Support decryption, and automatic framing/deframing of data

Questions?