

Technology Brief Application Note – OpenOnload® 201509

Introduction

The 201509 release of OpenOnload introduces a range of major features to deliver improvements in scalability and performance and allow acceleration of an even wider set of applications:

- **Scalability improvements** avoid hardware restrictions on the number of sockets
- **Higher performance** through improved connection rate
- **Acceleration** of IP_TRANSPARENT sockets provides support for a new class of applications such as haproxy
- **Resilience** through support for Stratus ftServer and hotplug
- **Compatibility** with Linux Teaming

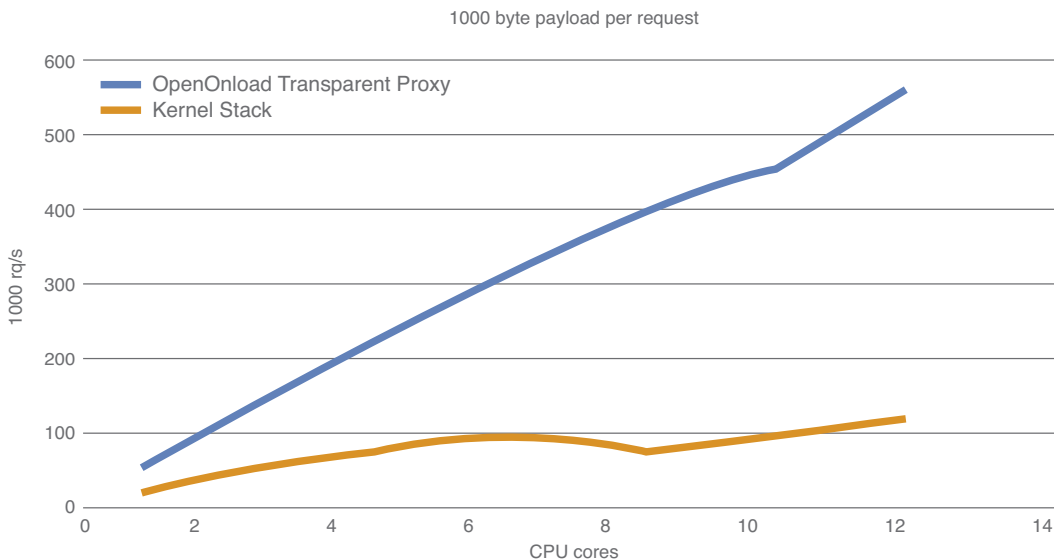
Scalability

The OpenOnload-201509 release provides Scalable Filters, intended to vastly improve scalability and the number of concurrent connections.

Scalable Filters remove the hardware restriction on the number of sockets created by OpenOnload. By using a MAC filter to redirect all traffic arriving at an interface to an OpenOnload stack it avoids the need to use a separate filter for each socket. Any ARP, IGMP or ICMP traffic arriving at the interface is directed to the kernel, allowing the interface to be configured and managed as normal.

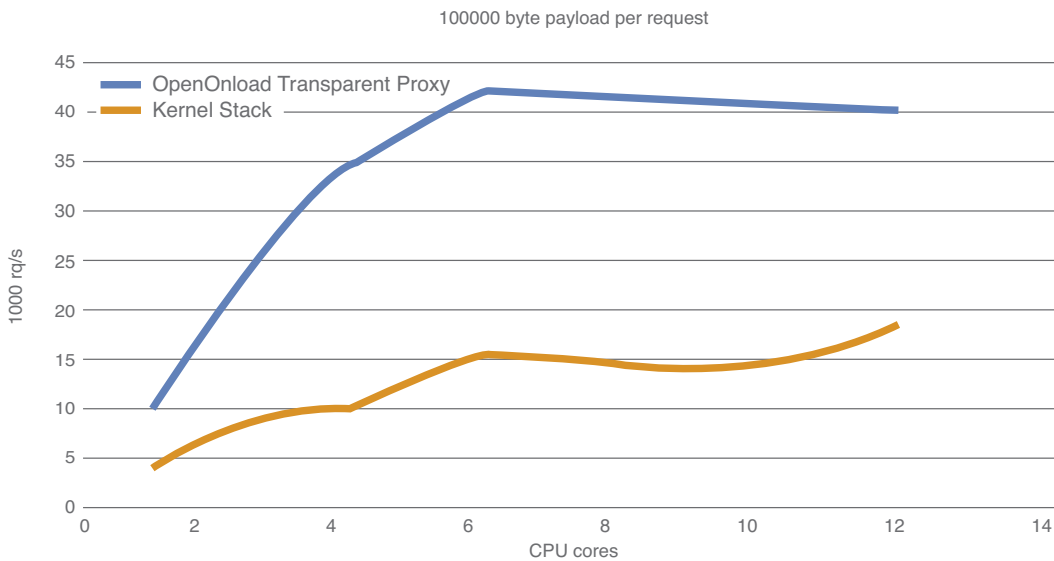
With no hardware limitations this allows the creation of up to 2 million sockets per stack in an accelerated application. Scalable Filters can be used with both TCP listening sockets and active-open TCP sockets using the IP_TRANSPARENT socket option.

Figure 1. Transparent Proxy Performance Onload vs. Kernel Stack



SolarflareTechnologyBrief

Figure 2. Transparent Proxy Performance Onload vs. Kernel Stack



Performance

The new Active Socket Caching feature complements the existing passive socket caching¹, extending the improvement in connection rate to active-open TCP connections. Sockets using the IP_TRANSPARENT socket option that are used for active-open TCP connections can now be cached by the Onload stack when they are closed. Newly created TCP sockets are then created from the cache. This avoids having to enter the kernel on either socket close or socket creation, significantly reducing the cost of these operations, typically giving a four-times increase in connection rate.

Applications

This release brings the acceleration benefits of Onload to new class of application: transparent reverse proxies, and load balancers working in transparent mode, such as haproxy. With the combination of new features such as Scalable Filters, Active Socket Caching, IP_TRANSPARENT socket support and Application Clustering load-balancer applications can run with unprecedented efficiency and scalability:

1. Scalable Filters allow up to 2 million connected sockets per stack as well as making creation of each new socket lightweight.
2. Active Socket Caching can reduce cost of creating sockets even further as resources from closed sockets are reused for new ones.
3. Application Clustering allows each CPU core to work in complete isolation, so achieving linear application scaling over multiple CPU cores.

Resilience

OpenOnload-201509 adds support for hotplug and the Stratus ftServer platform. Stratus ftServer provides complete hardware redundancy, with seamless failover, delivering a resilient platform for applications.

Compatibility

To complement the existing support in Onload for Linux bonding, new support for teaming has now been added. As with bonding, Onload will now seamlessly accelerate interfaces configured with the Linux teaming module and teamd.

¹Passive socket caching is described in [whitepaper SF-114312-CD](#) Nginx acceleration with Solarflare Flareon and Flareon Ultra network adapters.

