PROCERA EMPOWERING INTELLIGENCE | NETRONOME

Solution Brief

Building High-performance Deep Packet Inspection Platforms with Netronome and Procera Networks

Netronome NFE-3240

Procera and Netronome have partnered to deliver the industry's most flexible and highest-performance platform solutions for next generation Deep Packet Inspection (DPI) and network flow processing. This integrated solution combines the unparalleled application recognition and metadata extraction capabilities of Procera Networks NAVL DPI engine with the ultra-high performance and functionality of Netronome's Flow Processors, Network Flow Engine (NFE) acceleration cards and 1U/2U appliance reference platforms. Together, this represents an industry leading OEM platform for building and delivering next generation network infrastructure solutions that recognize and identify thousands of today's most popular applications at real time speeds up to 100 Gbps.

Application Recognition and Detection

Procera Networks' Network Application Visibility Library (NAVL) is a next-generation Deep Packet Inspection (DPI) software engine that provides real-time, Layer 7 classification of network traffic. NAVL uses a combination of deep packet and deep flow inspection techniques to accurately identify today's most common applications including mobile, social networking, P2P, instant messaging, file sharing, enterprise and Web 2.0 applications.

Procera's NAVL engine continually evolves to keep up with the rapid development of new applications and changes to existing protocols. Procera's team of DPI engineers receive feedback from tens of thousands of deployments globally to identify the newest applications and trends and constantly refresh plug-in patterns to keep up with perpetual changes in network traffic. This enables Procera to ensure NAVL's signature set guarantees the highest classification ratio possible and the most accurate classification of network traffic.

Packet Processing and Load Balancing up to 100Gbps

Netronome's flow processors dramatically accelerate Procera's NAVL by significantly offloading data plane processing. The NFE-3240 family of PCIe cards for COTs servers and flow processing platform reference designs offer the most complex packet and flow processing with unparalleled performance. These flow processing solutions tightly couple L2-L4 packet processing, L4-L7 flow processing, with the performance and scalability of general-purpose multicore x86 systems. This heterogeneous multicore architecture sets a new performance benchmark for network appliances with multiple layers of workload-specific packet, flow, security and application processing, each with increasing levels of granularity. The NFE-3240 and flow processing platforms enable the acceleration of network and security applications by utilizing high-performance packet processing delivered from 40 networking-optimized flow processor cores. These solutions utilize several techniques to dramatically improve network workloads including packet classification, stateful flow analysis and per flow action processing, Layer 2 switching, Layer 3 routing, IPsec VPN origination/termination, SSL inspection, network address and port translation and dynamic load-balancing of flows across a virtualized PCIe datapath to parallelize application processing.

Netronome and Procera: A unique technology pairing

NAVL quickly and accurately classifies all major network-delivered enterprise and consumer applications and sub-applications, enabling network equipment vendors to differentiate business-critical from non-critical applications. It even distinguishes among the voice, video, chat and file transfer capabilities of most social applications. Due to the incredible speed at which these applications and protocols change in today's dynamic environment, NAVL is optimized to run on x86 multicore processors.

Features and Benefits

- Application Recognition and Detection of today's common applications:
  - Mobile
  - Social Networking
  - P2P
  - Instant Messaging
  - File Sharing
  - Enterprise
  - Web 2.0
- Steady feed of new and updated application signatures ensuring continued accuracy and coverage
- Line-rate throughput for packet classification, flow processing, packet, and capture.

  Support for inline, passive, or L2/L3 modes of operation
- Low Latency
  - <20µs cut-through at switch layer
  - <40µs cut-through at NFE
  - <100µs (inline application processing)
- Optional integrated bypass for inline applications
- Independent per-port pair failover
- 40 microengine cores operating at 1.2GHz, providing over 1600 instructures/packet at 30M pps
- Improved application/server performance:
  - Significantly reduce host CPU utilization by offloading flow action processing to NFE

The fewer cores and on-board memory used by a DPI engine on a multicore x86 appliance, the better. Maintaining a small footprint with high performance helps contain costs for network infrastructure vendors and their customers. By coupling Netronome's flow processors to the NAVL engine, CPU utilization of the engine itself can be reduced to negligible amounts returning all of that processing capability to the customer application.

Through a set of open APIs, the NAVL engine and the Netronome flow processors work in tandem to provide unparalled classification accuracy with industry leading performance. In the flow processors, a stateful flow table is maintained that allows per-flow action processing. Initially, all packets of a flow are sent to the NAVL engine for application classification through a variety of techniques including:

- Surgical Pattern Matching
- Deep Protocol Dissection
- Semantic and Conversational Awareness
- Behavioral Analysis
- Flow registration and association

Once the flow has been identified, through simple API calls, the actual action handling of the traffic is offloaded to the Netronome flow processors. On a per-flow or per-application basis, a wide range of actions can be applied to the packets of a flow(s). Traffic can be:

- Actively or passively dropped
- Cut-through the appliance from ingress to egress physical interface
- Redirected from the core(s) NAVL is using to different x86 destination
- Load-balanced across a set of x86 cores inside the appliance
- Load-balanced across a set of egress interfaces
- L2 switched or L3 routed
- Encrypt/decrypt
- Inserted into a tunnel (IPsec, SSL, IP in IP, GRE)
- QoS applied
- Rate limited
- Translated via NAPT
- Add/translate packet fields: VLAN, MPLS, IP, VxLAN, DSCP

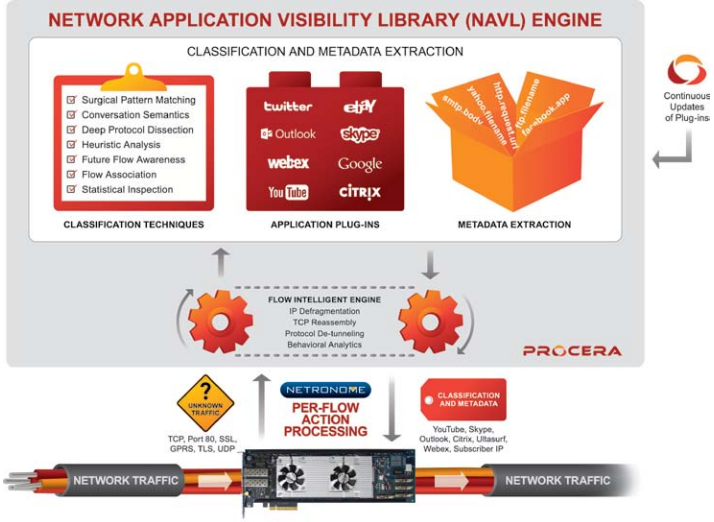

For many applications, NAVL only needs to see the beginning of each flow, and once classified all of the data handling for the flow is handled in the flow processor increasing PCIe throughput and reducing x86 CPU utilization.

Conclusion

In the quest for more intelligent applications operating at ever higher throughputs, finding tools that speed your applications as well as speed your time to market are rare finds. The powerful combination of Procera's NAVL DPI engine and Netronome's flow processing are that extraordinary powerful technology pairing. Providing industry-leading visibility into network flows at incredible throughputs.

Netronome's Flow Processors, Acceleration Cards and Reference Platforms

www.proceranetworks.com
info@proceranetworks.com

Corporate Offices
Procera Networks, Inc.
4121 Clipper Court
Fremont, CA 94538
P. +1 510-230-2777
F. +1 510-656-1355

Canadian Headquarters
Procera Networks
#302 - 1353 Ellis Street
Kelowna, BC V1Y 1Z9, Canada
P. +1 250-448-1925
F. +1 250-412-3558

European Headquarters
Procera Networks
Birger Svenssons Väg 28D
432 40 Varberg, Sweden
P. +46 (0)340-48 38 00
F. +46 (0)340-48 38 28

Asia/Pacific Headquarters
Procera Networks, Pte. Ltd.
Penthouse #44-01, Suntec Tower Three
8 Tamasek Boulevard, Singapore 038988
Phone: +65 6829 2220
Fax: +65 6829 2206

Copyright © 2013 Procera Networks. All rights reserved. All other trademarks are property of their respective owners. SB_Netronome_Q2_2013_5_1