White Paper | July 2012

radisys®

Managing Quality of Experience with RAN-Aware Policy Enforcement

By: Chandresh Ruparel, Senior Product Marketing Manager, Radisys

Overview

As mobile traffic continues to grow exponentially, mobile operators require better solutions for reducing cost per bit and increasing revenue per subscriber. A key hurdle is service disruption caused by congested cells, a major source of subscriber frustration and churn. Network capacity upgrades alone are not sufficient to solve the problem—especially in the Radio Access Network (RAN), where the bandwidth is very expensive and therefore limited. Network operators face a huge challenge, as bandwidth-intensive applications, such as gaming, photo sharing, social media and mobile TV, continue to gain popularity.

Policy management is viewed by network operators as one of the effective and viable means to manage network traffic and promote revenue growth. This capability is well-established and deployed in the fixed-line networks, which historically had the greatest data bandwidth demand. However, with advent of smartphones, tablets and other wireless connected devices, mobile networks are increasingly facing a similar challenge.

CONTENTS

Can mobile networks duplicate the policy management solution from fixed-line networks? The migration of the solution to mobile networks is not straightforward as the protocols and element behaviors are different. In addition, in the mobile networks, RAN awareness is very important for effective policy management. For example, blocking the major culprits, like peer-to-peer (P2P) applications, across all cells in the RAN would free up substantial bandwidth, but such a policy will reduce quality of experience (QoE) on non-congested cells and still not provide the assurance that a premium subscriber will get good QoE in congested cells. This does not help average revenue per unit/subscriber (ARPU) growth and instead can cause subscriber churn.

A more effective option is to adopt real-time policy management capable of prioritizing subscribers based on applications and service level ONLY in congested cells. This control mechanism ensures all subscribers get better user experience on non-congested cells, while providing better quality of service and experience for premium users on congested cells.

This paper reviews the market drivers for policy management and how RAN-Aware Traffic Shaping and Policy Management can be very effective in reducing subscriber churn and improving chances for ARPU growth. The paper also highlights how network equipment providers (NEPs) can develop such a solution using best-in-class pre-integrated products from Radisys, Qosmos and 6WIND.

Market Drivers

Mobile Broadband Is Booming!

The usage of mobile data and video services is experiencing impressive growth. According to Allot Communications, video streaming traffic continues to dominate mobile broadband, with a 42 percent share of all global bandwidth (Figure 1).[1] In 2011, YouTube accounted for about 24 percent of global broadband traffic. A large portion of the broadband data upsurge is attributable to social media, where Twitter, YouTube and Facebook traffic was up by between 100 and 200 percent during the second half of 2011, as illustrated in Figure 2.[1] Trends indicate that future growth in mobile broadband traffic will continue to be driven by social media—video in particular.

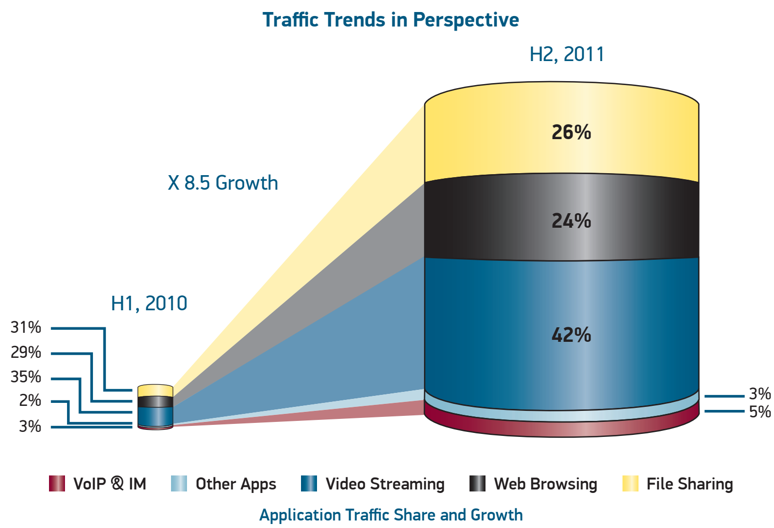

Figure 1. *Mobile Broadband Traffic (Source: Allot Mobile Trends Report H2 2011[1])*

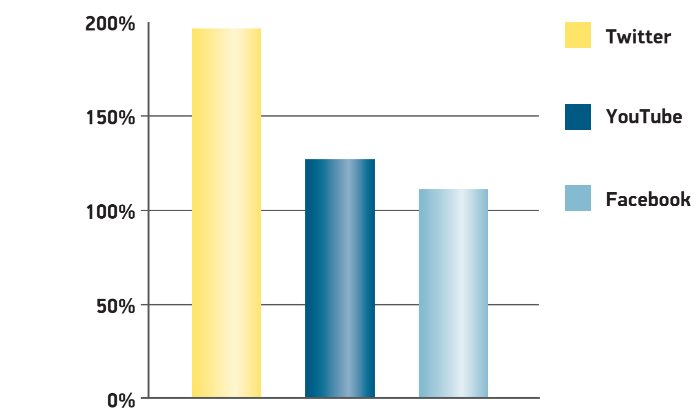

Figure 2. *Social Media and Video Sharing Bandwidth Growth[1]*

Traffic and Revenue Disparity

Revenue isn't increasing as fast as booming broadband demand, creating consternation among mobile operators,. Heavy Reading, an independent research organization, illustrates this trend with Figure 3, where there is a transition from the "Voice Era" to the "Data Era" coinciding with the point when traffic growth begins to outpace revenue growth.[2] In order to deal with this new state of affairs, operators require network equipment that delivers lower cost per bit, hence the move to LTE, and strategies to better monetize network bandwidth, like offering tiered services. The latter requires policy management capable of differentiating subscribers and the applications they use, which is covered in more detail in the next section.

Cell Loading Is Uneven

Mobile network congestion is often localized and variable because the loading on a cell will depend on its location, time of day, as well as other factors. For example, cells in densely populated urban areas will be most heavily congested during the work day, whereas residential cells will experience higher demand in the evening hours. In addition, there are often unpredictable increases in traffic through cells, as during major sports, entertainment or political events. Given this variability, policy management should be flexible and dynamic to best respond to changing traffic patterns in the RAN and to deliver improved QoE based on real-time traffic and congestion information.

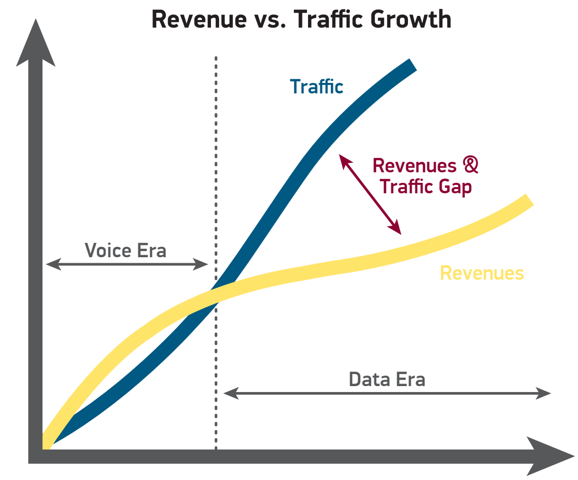

Figure 3. *Traffic Growth Overtaking Revenue Growth*

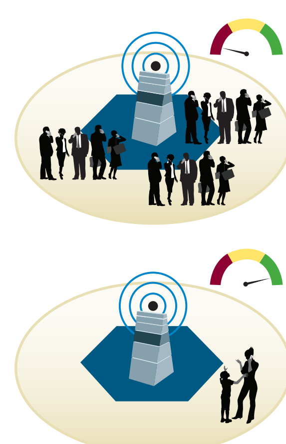

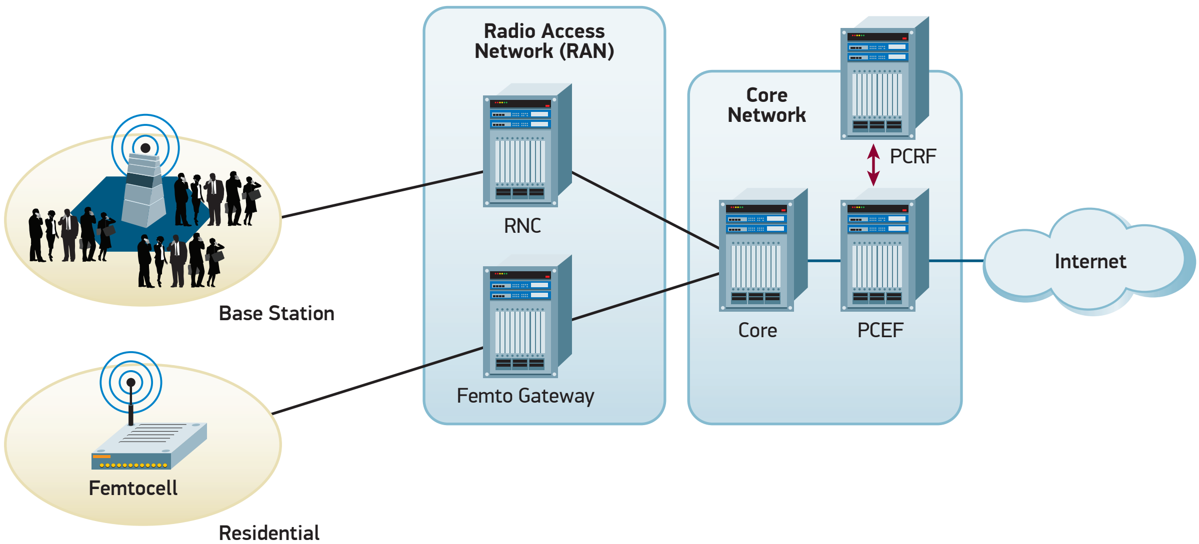

Figure 4. *3G Example: PCRF and PCEF Functions in the Network*

Mobile Policy Management Challenges

Policy management in both wired and wireless networks is achieved by integrating two key functions, which are shown in Figure 4 and described in the following:

- *The Policy and Charging Rule Function (PCRF)* determines policy and charging rules for controlling service data flows and IP bearer resources. It provides the PCEF the applicable policy and charging control (PCC) rules.
- *The Policy and Charging Enforcement Function* (PCEF) is responsible for enforcing policies and charging decisions made by the PCRF by performing functions such as traffic shaping, deep packet inspection (DPI), flow marking and QoS control.

Wired broadband operators, whose core business is moving large amounts of data traffic, deployed policy management early on. This was commonly done by implementing the PCRF as a separate piece of equipment and integrating the PCEF into edge routers, which sniff packets to determine how to handle the flow.

Unlike their wired counterparts, many mobile networks lack sophisticated policy management. Although 3rd Generation Partnership Project (3GPP) specified PCRF and PCEF for policy management, its adoption is limited. One reason is that for decades, mobile networks predominately carried voice and hardly any data; thus, traffic shaping wasn't needed. In fact, the PCEF function was mostly neglected, in that it wasn't incorporated into 3G Core Networks as a standalone network element or integrated into an existing element—analogous to edge routers in wired broadband supporting PCEF functions.

Now that mobile broadband has taken off, due in large part to faster networks and the advent of smartphones, the need for advanced policy management is clear. Operators have taken an easy route by applying the policies used in wired networks, leading to today's capabilities that tend to be basic, such as applying blanket policies that limit high-bandwidth applications at all times and across all cells. With such a rigid policy, even subscribers on premium plans will experience lower quality of service when cells are congested, leading them to question why they are paying more.

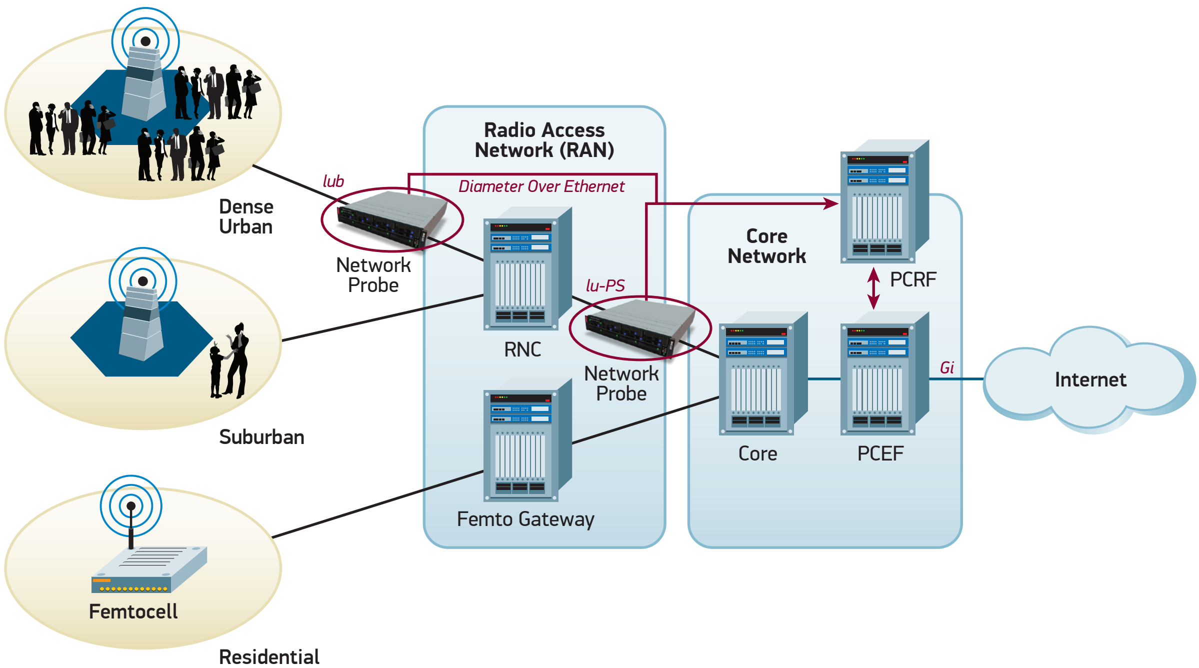

Figure 5. *3G Example: Mobile Networks with RAN-aware Policy Management*

High Performance Policy Management

Network operators face an on-going challenge to maximize their average revenue per user (ARPU) to offset constant increases in CAPEX and OPEX associated with investments in network capacity expansion and the adoption of advanced services. "Network Intelligence" is one key to maximizing ARPU: if operators have real-time knowledge about the traffic characteristics and demand in their network, segmented by application, user and time-of-day, they can deliver customized services that provide high value to specific customers. Subscribers selecting premium rate plans, rather than "best-effort" options, can reliably be given an experience commensurate with the higher costs.

A solution, called RAN-aware traffic shaping and policy management, manages traffic flow during peak periods, leading to improved QoS for most subscribers and ensuring a higher QoE for those on congested cells who are paying a premium. This is achieved by adding another network function—network monitoring/network probe—as shown in Figure 5 (3G example). The network probe adds cell awareness by monitoring the signaling in the RAN; it sniffs IuB/Iu-PS signaling (S1U in LTE networks), maps the user ID and Packet Data Protocol (PDP) sessions, to the cell ID and delivers this information to PCRF over an Ethernet link. The PCRF, leveraging 'RAN awareness,' generates policy and charging control (PCC) rules that are passed to the PCEF, which performs traffic shaping based on real-time network intelligence on traffic flowing through the Gi interface (SGi in LTE networks).

Network Probe Requirements: Since network probes are deployed in multiple locations in the carrier's network, the platform must be cost effective, high performance, carrier grade and short depth to fit most rack sizes. In addition to executing dataplane applications, the platform must be integrated with robust signal monitoring software.

3G/LTE Network Probe

Intel® dramatically improved dataplane performance on its latest generation of server-class processors (codenamed Sandy Bridge), and this significant stride is a major contributor to enabling a more space and cost efficient platform. As shown in the figure below, the combination of the Intel® Data Plane Development Kit (Intel® DPDK) and Intel® microarchitecture improvements are delivering upwards of 165 MPPS of L3 forwarding throughput (64 byte packets) using the dual-socket 8-core Intel® Xeon® processor E5-2658 or approximately four times greater than a 2009 Intel® Xeon® processor-based platform, as shown in Figure 6.

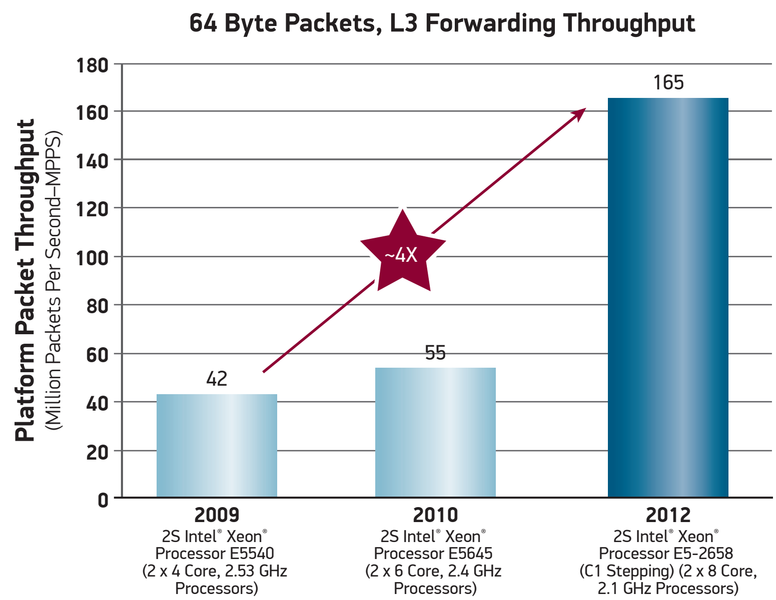

Figure 6. *L3 Forwarding Throughput Benchmark Data for 64 Byte Packets (Source: Intel®)*

Hardware Platform

The Radisys RMS-220 Network Appliance Server is based on dual 8-core Intel® Xeon® E5-2658 (the same configuration as shown on the right side of Figure 6), and is part of the pre-integrated Radisys network probe solution illustrated in Figure 7. The Radisys RMS-220 Network Appliance delivers unmatched combination of power, space, performance and serviceability optimizations. With its maximum front serviceability and configurability, short-depth and carrier-grade design, the RMS-220 can be configured in the field for broad range of applications WITHOUT opening the hood or de-racking the system!

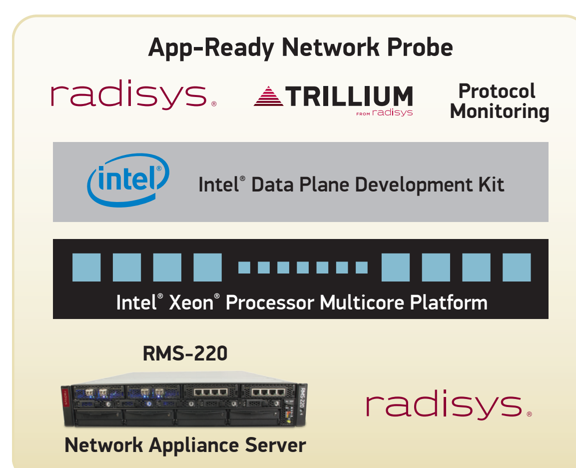

Figure 7. *Key Ingredients of the App-Ready Network Probe*

Key attributes of the RMS-220 include:

- Short depth (20-inches) and carrier grade (NEBS/ETSI compliant)
- Dual 8-core Intel® Xeon® processor E5-2600 series (formerly codenamed Sandy Bridge/Romley platform), and up to 256GB memory via 16 x DIMM slots
- Front I/O modules that can be serviced/configured without de-racking and support a broad range of I/O options. Optimization for packet processing with the Intel® DPDK. Max 240G I/O per server
- 5-10TB capacity of storage/local caching with front-mounted, hot-swappable SAS, SATA or SSD storage with built-in RAID 0/1/5/10
- Highly serviceable (built-in remote/lights-out management, front serviceability) with redundant and hot-swap capabilities
- NEBS and non-NEBS variants available
- Long life support (5-7 years)

Monitoring Software

The network probe provides RAN-awareness via Radisys Trillium signal monitoring software for Iub and Iu-PS, or S1U interfaces. The monitoring solution is based on high-performance, field-hardened, Trillium 3G and LTE protocol stacks, which are used the world over by a majority of the leading telecommunications equipment manufacturers. Critical to the sniffing/monitoring process, the monitoring solution collects data at wireline speed without disrupting the flow of traffic. Furthermore, data is captured at Layers 2-5, allowing the monitoring application to perform comprehensive analysis on each traffic session. For example, an NEP-developed application can analyze the data delivered at each layer and generate information per user, such as throughput, quality level and session type (e.g., YouTube, Bit Torrent, Web browsing). The traffic flowing across the protocol interfaces is characterized via the gathered statistics, including traffic volume, session type and packet errors.

In a 3G network, the probe monitors a wide range of protocols across the Iub interface (between the NodeB and RNC) and the Iu-PS interface (between the RNC and SGSN), as shown in Figure 8. Radisys offers comparable capabilities for 3G WCDMA, as well as LTE networks. The monitoring software can be integrated with the Intel® DPDK. This library helps dramatically increase packet processing performance for comprehensive control and user plane monitoring on Intel® processors.

One Interface

The network probe can be accessible through Trillium Diameter, a highly-interoperable application programming interface (API) used to send information towards the PCRF/PCEF, such as user identity, cell number and PDP IP Address, to the PCRF. Trillium Diameter source code helps NEPs address common concerns such as interoperability, security, quality of service, and legacy interworking through authentication, authorization and accounting (AAA) services.

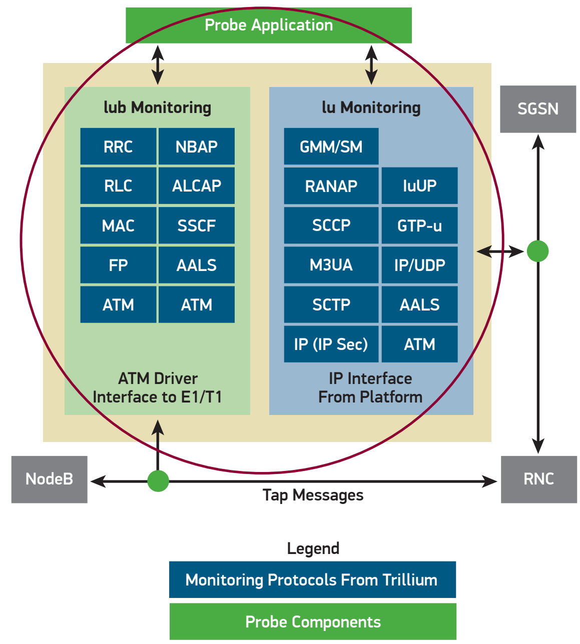

Figure 8. *UMTS Network Probe Software Architecture*

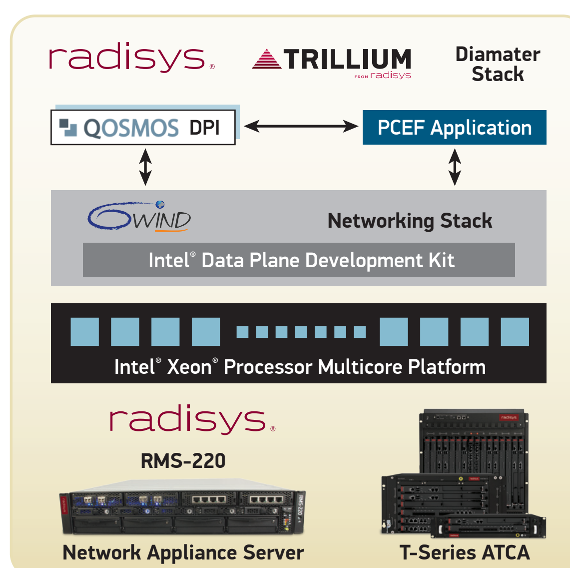

Figure 9. *Key Ingredients of the RAN-aware PCEF*

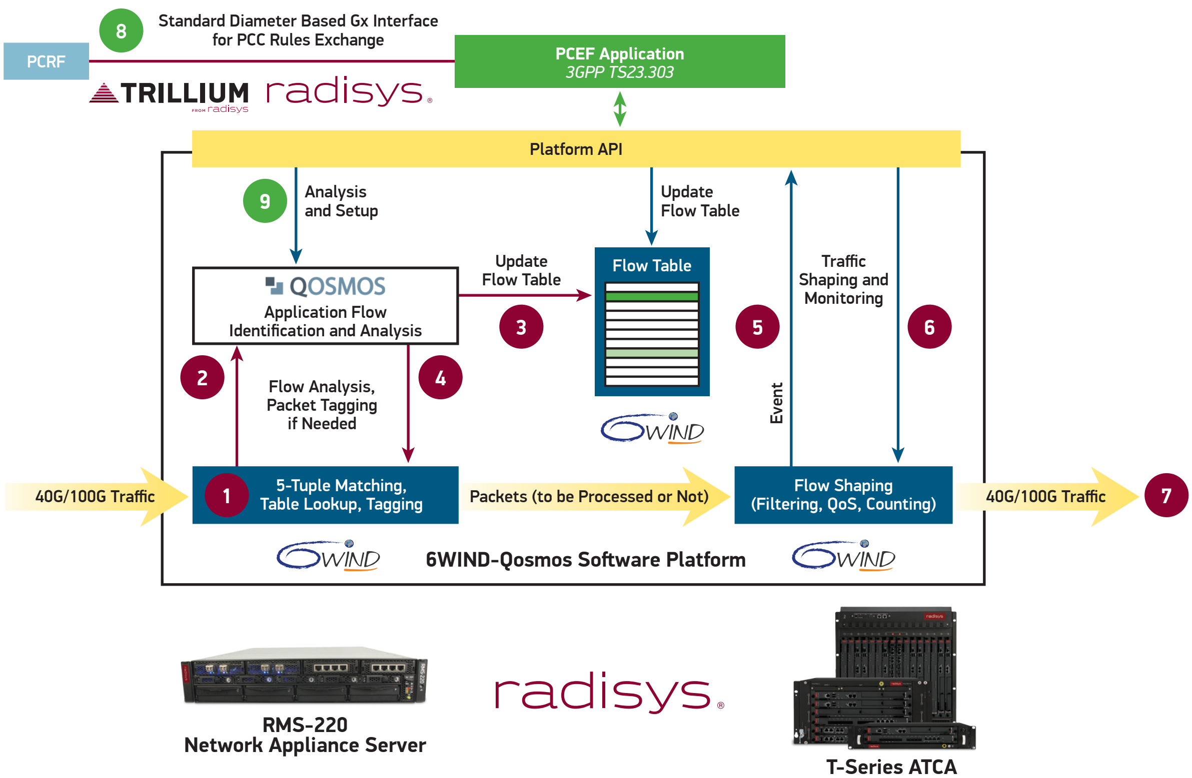

Figure 10. *PCEF Framework — High Level Architecture*

PCEF High Level Framework

The PCEF must provide high data plane performance and support state-of-the-art packet classification, DPI and traffic shaping capabilities. To address these requirements, the application-ready PCEF solution, presented in this whitepaper, was designed to enable NEPs to quickly develop a cost-effective, high-performance solution (Figure 9). It includes:

- Carrier-grade hardware platforms such as RMS-220 Network Appliance and high density ATCA platforms combined with high-performance Trillium Diameter software from Radisys
- DPI and metadata extraction engine from Qosmos
- High-performance packet processing software from 6WIND
- Multi-core processors and software libraries from Intel®

PCEF Operation in Detail

Since the PCEF is in the data path, the following events, illustrated in Figure 10, must occur at wire speed.

- *Step 1.* 6WINDGate packet processing software computes 5-tuple and generates an entry in the flow table for each new 5 tuple (hash function, table lookup, tagging)
- *Step 2.* First packets of a flow are passed to the Qosmos DPI engine that identifies the application
- *Step 3.* Qosmos DPI engine updates the 6WINDGate flow table with application name and metadata, accessible by the PCEF application (written by the NEP) via a platform API
- *Steps 4-7.* Analyzed packets go back to the 6WINDGate fast path in the data plane that performs traffic shaping and QoS management as per the instructions received from PCEF application and sends out the packets. Subsequent packets for this resolved flow no longer need to go through the Qosmos DPI engine. Overall performance of the solution is optimized.

The following activities are not in the data path.

- The PCEF is made aware of the RAN traffic composition via the Trillium Diameter API to the PCRF.
- THE PCEF application requests additional flow information from the Qosmos DPI engine for in-depth analysis.

PCEF Solution Components

Hardware Platforms

PCEF solutions based on Radisys carrier-grade hardware platforms are highly scalable, as illustrated in Figure 11. Based on various Intel® DPDK, 6WINDGate and Qosmos ixEngine benchmarks (when used in clusters), the Radisys RMS-220 network appliance provides an optimal platform for 10-100G DPI throughput for 3G/LTE deployments. 6U and 14U ATCA platforms handle up to 300G and 500G of traffic, respectively. All three hardware platforms are based on high performance Intel® Xeon® processors and the Intel® DPDK. The platforms also have the capability to augment application-specific performance with additional CPU/NPUs.

Trillium Diameter Software

Diameter was defined by the Internet Engineering Task Force (IETF) as a follow on to widely-deployed, legacy access-control protocols, including RADIUS (Remote Authentication Dial in User Service). RADIUS was commonly used for most dial-up and initial broadband internet service provider (ISP) networks; however, the evolving and growing complexity of network architectures, like IP Multimedia Subsystem (IMS) and the services delivered over those networks, posed their own, higher demands on the AAA framework to provide stable, fault-tolerant and scalable protocols. Now, Diameter is being widely used across all-IP (Internet Protocol) networks and in next generation telecom networks.

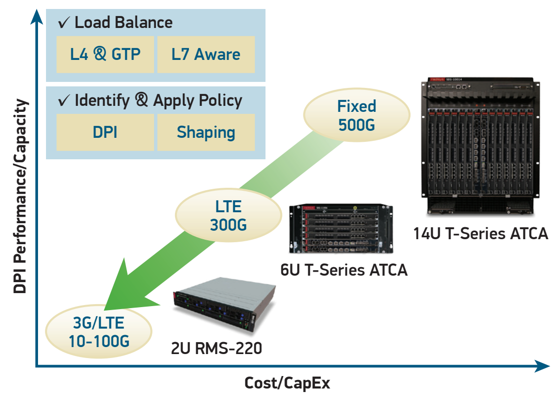

Figure 11. *Radisys Hardware Platform Options for PCEF*

Trillium protocol software from Radisys is synonymous with high quality and superior performance. Trillium Diameter software leverages this strong heritage to simplify the Next Generation Network (NGN) and facilitate the deployment of innovative, revenue-generating services. Trillium Diameter software is standards-based, interoperability-tested, multi-threaded and optimized to meet core network performance demands (e.g., deployments on multi-core CPUs).

Deep Packet Inspection

One of the major tasks of the PCEF is to perform flow control per instructions from the PCRF. This means the PCEF must be able to identify the application types (e.g., email, video streaming, and web browsing) associated with every flow. This is typically accomplished using DPI, which enables the PCEF to look deep into a packet's payload in order to classify the application, as well as to capture appropriate application-specific content from flow. With this information, the PCRF can perform networking functions such as forwarding, security, usage analysis, traffic management, resource allocation, QoS and billing.

The pre-integrated PCEF solution includes the 3rd generation Qosmos DPI engine, called ixEngine, which can be embedded into core network elements (e.g., compute blade) or standalone PCEF boxes. The network intelligence provided by Qosmos DPI enables these functions to be performed with much more sophistication than in traditional networking equipment, where decisions are based only on the packet header. Operators now have access to a wealth of information as a result of data mining, profiling and analytics, all of which are key to the creation of personalized service packages based on a deeper understanding of customers' needs.

Designed with developers in mind, ixEngine accelerates the product development cycle by offering the capabilities listed in Figure 12. Ready-to-use software libraries reduce the effort and lower the risks associated with developing and maintaining a highly complex technology internally.

Table 1 describes the features of the Qosmos ixEngine in more detail.

Packet Processing

The 6WINDGate software has been specifically designed for high-performance networking applications. As a drop-in replacement for the standard Linux networking stack, it maintains compatibility with Linux APIs so existing applications can easily be reused in new systems.

The data plane is split into two layers, where the lower layer is called the fast path; this processes the vast majority of the incoming packets outside the OS environment, thereby avoiding OS overheads, such as preemptions, threads, timers and locking, that degrade overall performance. The upper layer is the 6WIND-optimized networking stack, which performs the necessary management, signaling and control functions, while also handling those rare packets that require complex processing. The 6WINDGate fast path runs in a Linux userspace. environment on dedicated cores within the Intel® processor. Other cores run Linux, the 6WINDGate control plane and the application software, as shown in Figure 13.

ixEngine Software Development Kit

Key Features:

- Real-time, stateful analysis of IP flows
- Protocol and application Identification
- Extraction of metadata and content from flows
- Correlation of flows at user, application and service level

Figure 12. *Qosmos ixEngine Capabilities*

Decoding Engine	• Identification of protocol and applications contained in IP Flows using stateful inspection and heuristic analysis. Qosmos' protocol plug-in suite includes over 1,000 protocol and application plug-ins to classify flows and extract metadata (5000 metadata available to date). • Extraction of metadata from protocols (e.g., volume, delay, jitter etc.) and applications (e.g., name of file streamed on a website).
Protocol Updates	• Continuous watch of protocols and applications, and update of ixEngine when new versions appear that impact classification and metadata extraction.
Protocol Services	• Instant access and continuous updates to 1000+ protocols and 5000+ communications. • On-demand development of new protocol plug-ins. • Protocol change reports for notification of changes in protocols.
Protocol Plug-in SDK	• Ready-to-use DPI development tools and libraries
	• Allows users to develop custom protocol plug-ins for their ixEngine.
Supported Operating System	• Linux Standard Base 3

Table 1. *Features of the Qosmos ixEngine*

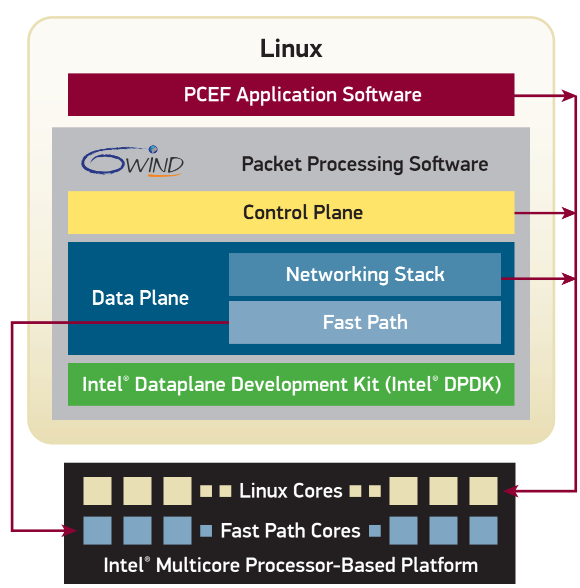

Figure 13. *6WINDGate Packet Processing Software Architecture*

Typically, 6WINDGate delivers ten times the packet processing performance of a standard Linux networking stack, with performance increasing linearly based on the number of processor cores configured to run the fast path. The IP forwarding performance of 6WINDGate running on a dual socket Intel® Xeon® Processor E5 family platform is approximately 14 million packets per second (MPPS) per core or 10 Gbps of network traffic.

Contributing to this impressive throughput is the Intel® DPDK, which is essentially a streamlined driver that gives software "bare metal" access to an Intel® processor-based platform. The Intel® DPDK enables 6WIND's networking software to extract maximum processing performance from Intel® processors and Intel® network interface cards (NICs). The Intel® DPDK is pre-integrated within the 6WINDGate solution, and its presence is transparent to the application software that calls 6WINDGate, which is a major time-to-market advantage for OEMs preferring not to perform the integration themselves.

6WINDGate includes a comprehensive set of networking protocols for routing, switching, security, mobility and protocol termination, minimizing the need to integrate protocols from multiple software suppliers.

Available protocols include:

- VLAN, Mobile IP, link aggregation, GRE, PPP, L2TP, GTP, MPLS, ROHC
- IP forwarding, IP tunneling, routing, virtual routing
- IPv4, IPv6, IP multicast, IPsec, IKE, firewall, NAT
- UDP, TCP, SCTP, flow inspection, SSL termination, TCP termination

By offering one-stop shopping for networking protocols, 6WIND minimizes development time and schedule risk.

High Performance Multi-core Processors

The Radisys RMS-220 network appliance and various ATCA compute blades are equipped with the Intel® Xeon® processor E5-2600 product family, which enables platforms to offer solutions ranging from eight-core/single socket up to 16-core/two-socket.

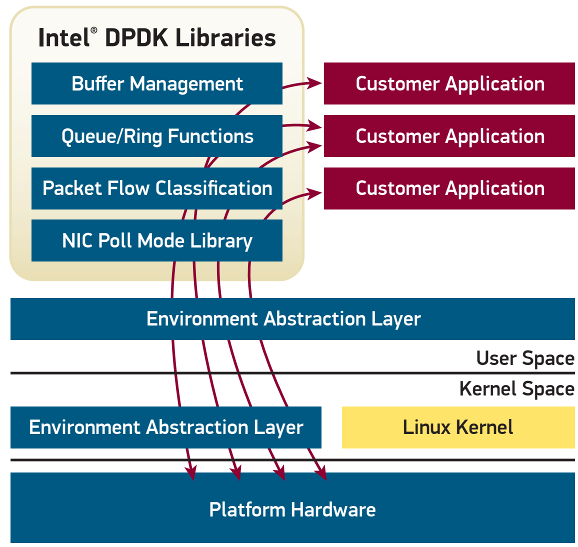

Figure 14. *Major Development Kit Components*

These processors provide significant performance improvement over previous-generation Intel® Xeon® processors. For example, the eight-core Intel® Xeon® processor E5-2658Δ delivers nearly 68 percent performance gain over the previous-generation six-core Intel® Xeon® processor E5645Δ,3 within a similar platform thermal profile. Increased memory capacity is provided with four memory channels per processor and up to 1600 MHz DDR3 memory.

Intel® Dataplane Development Kit (Intel® DPDK)

The Intel® DPDK provides Intel® architecture-optimized libraries that allow developers to focus on their applications. The Intel® DPDK provides non-GPL source code libraries to support exceptional data plane performance and ease software development, while minimizing development time. This allows the developers to make additions and modifications to the Intel® DPDK, as required, to meet their individual system needs.

Designed to accelerate packet processing performance, the Intel® DPDK contains a growing number of libraries (Figure 14), whose source code is available for developers to use and/or modify in a production network element, with examples showing usage cases and performance. Developers can build applications with the libraries using “run-to-completion” and/or “pipeline” models that enable the equipment manufacturer’s application to maintain complete control. The following list provides a brief description of key software components.

- *The Environment Abstraction Layer (EAL)* provides access to low-level resources (hardware, memory space, logical cores, etc.) through a generic interface that hides the environment specifics from the applications and libraries.
- *The Memory Pool Manager* allocates NUMA-aware pools of objects in memory. The pools are created in huge-page memory space to increase performance by reducing translation lookaside buffer (TLB) misses, and a ring is used to store free objects. It also provides an alignment helper to ensure objects are distributed evenly across all DRAM channels, thus balancing memory bandwidth utilization across the channels.
- *The Buffer Manager* significantly reduces the amount of time the system spends allocating and de-allocating buffers. The Intel® DPDK pre-allocates fixed size buffers, which are stored in memory pools for fast, efficient cache-aligned memory allocation and de-allocation from NUMA-aware memory pools. Each core is provided a dedicated buffer cache to the memory pools, which is replenished as required. This provides a method for quick access and release of buffers without locks.
- *The Queue Manager* implements safe lockless queues instead of using spinlocks that allow different software components to process packets, while avoiding unnecessary wait times.
- *The Ring Manager* provides a lockless implementation for single- or multi-producer/consumer enqueue/de-queue operations, supporting bulk operations to reduce overhead for efficient passing of events, data and packet buffers.
- *Flow Classification* is an efficient mechanism for generating a hash (based on tuple information) used to combine packets into flows, which enables faster processing and greater throughput.
- *Poll Mode Drivers* greatly speed up the packet pipeline for 1 GbE and 10 GbE Ethernet controllers by receiving and transmitting packets without the use of asynchronous, interrupt-based signaling mechanisms, which have a lot of overhead.

So in summary, the proposed PCEF Application Ready Platform brings together “best-in-class” technologies:

- Radisys Trillium provides multi-threaded Diameter stack with unparalleled performance to handle IuB/Iu-PS/SU1 and Gx/SGx interface message exchange
- Qosmos provides comprehensive DPI and metadata engine, along with instant access to 1000+ protocols and 5000+ comms metadata, parsing of tunneling protocols, data filtering, visibility at ALL levels (flow, session, IP application and subscriber)
- 6WIND provides a high performance packet processing engine, integrated with the Intel® DPDK. It offers a comprehensive set of networking protocols and support for high availability frameworks
- Radisys platforms offer the broadest range of Intel®-based carrier-grade options, including the RMS-220 Network Appliance and ATCA platforms

The benefits to NEPs of a pre-integrated network probe and PCEF solutions include:

- Faster time-to-market with proven interoperability between 6WIND, Qosmos and Radisys components
- Ability to focus on core business by leveraging best-in-class flow management, DPI, protocol interface/monitoring and network appliance
- Maximizing system performance through already-developed optimizations
- Further reductions in development time, cost and risk given common APIs, ready-to-deploy advanced policy control and rules methodologies, and easy interoperability with a PCRF using the standard Diameter protocol interface

Summary

Improving the Quality of Experience

Utilizing a RAN-Aware PCEF to perform adaptive traffic shaping and policy management, network operators can manage QoS for all users, prioritize premium customers within a congested cell, and offer tiered services to increase revenue. Pre-integrated solutions, featuring technology from Radisys, Intel®, Qosmos and 6WIND, enable network equipment providers to reduce their development time on the order of years and significantly mitigate risk in delivering this solution. With mobile data traffic expected to increase to incredible levels, the RAN-Aware PCEF is an essential tool for network operators.

References

[Δ] Intel® processor numbers are not a measure of performance. Processor numbers differentiate features within each processor series, not across different processor sequences. See http://www.intel.com/products/processor_number for details.

[1] Source: Allot Communications, "Mobile Trends Report H2, 2011."

[2] Source: Heavy Reading

[3] Intel® Xeon® processor E5-2658 benchmarking results collected by Intel® Corporation, September 2011. The Intel® Xeon® processor E5645 benchmarking results collected by Intel® Corporation, June 2010.

Platform configurations:

- Intel® Xeon® processor E5-2658 (2 sockets) at 2.1 GHz, 20 MB L3 Cache, 90W; Intel® C604 chipset; 16 x 4 GB RDIMM DDR3-133 MHz
- Intel® Xeon® processor E5645 (2 sockets) at 2.4 GHz, 12 MB Last Level Cache, 80W; Intel® 5520 chipset, 12 x 4 GB RDIMM DDR3-1333 MHz

Software configurations:

- Intel® Xeon® processor E5-2658: OS: Red Hat 6.1 Beta, kernel version 2.6.32-122.el6.x86_64; Compiler: Intel® C/C+ 12.1RC1; Benchmark CPU2006

radisys.

Corporate Headquarters

5435 NE Dawson Creek Drive
Hillsboro, OR 97124 USA
503-615-1100 | Fax 503-615-1121
Toll-Free: 800-950-0044
www.radisys.com | info@radisys.com