WHITE PAPER

Massive Analysis, Correlation and Scalability of IP Metadata (MAXIM®) using the SGI® Altix® UV System

November 2010

Summary

The purpose of this document is to describe a high performance solution for the packet processing, decoding and analysis of massive amounts of IP traffic. Intended fields of application include national cyber security and communications interception.

It is a best-of-breed approach combining:

- Tilera, for low power consumption / advanced multicore packet processing
- Qosmos ixEngine, for real-time traffic decoding and metadata extraction
- SGI Altix UV, for high performance computing power using the Intel® Xeon® processor 7500 series
- Linux environment, with many terabytes of in-memory data for event correlation

Technical Principles

This white paper outlines the value of a combined solution based around the SGI Altix UV shared memory server, the Tilera multicore packet processing, and the Qosmos protocol parsing capability.

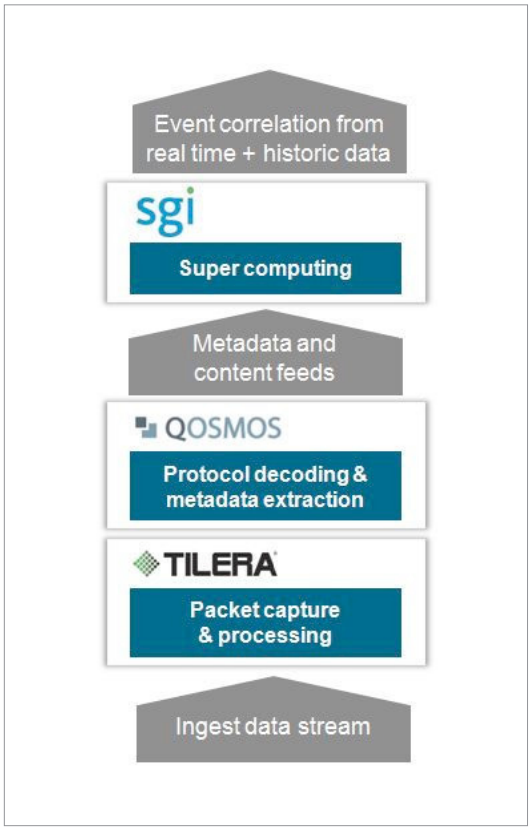

Figure 1: Components for the SGI Solution

This solution is used to create a scalable environment for analyzing packet streams in real time, and matching events described by the metadata produced.

The unique value this combination of technologies provides is a scalable solution that can meet growing processing demands within an open standard processing environment.

The following sections outline the components that when coupled together, enable this unique capability

	Features	Benefits
Tilera Tile Pro 64	Multicore packet processing	-High performance processing -Low power consumption
Qosmos ixEngine	Real-time traffic decoding and metadata extraction	-Completed visibility of traffic -Ready-to-use development tools and libraries: fast time-to-market -500+ protocols and 4000+ communications metadata
SGI Altix UV	Super computing power on an open platform	-Scalable processor, memory, and IO capability to meet growing needs over time -Easy application development in a familiar Intel / Linux environment -Reuse of existing applications

System Overview

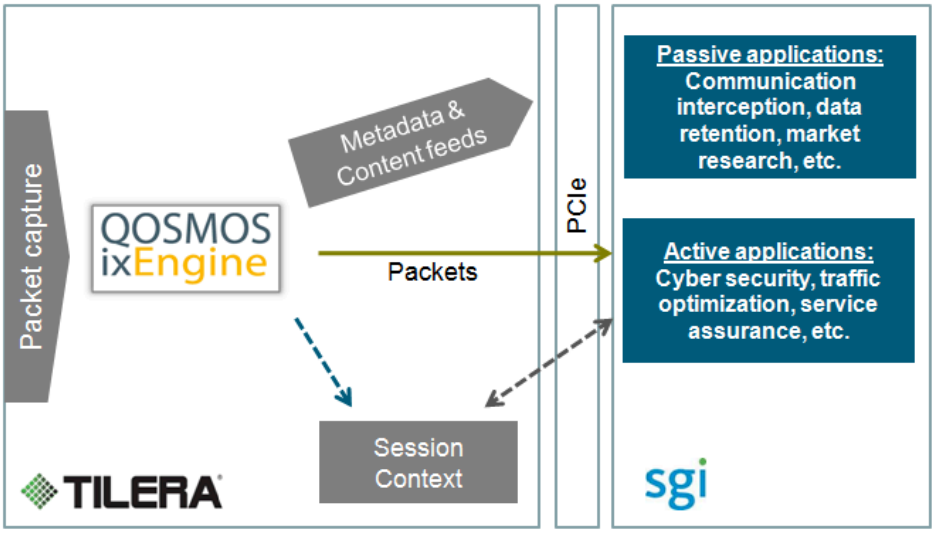

Figure 2: Principles of Interworking

Tilera

The MAXIM® solution takes advantage of Tilera's flexibility and best-in-class performance for Deep Packet Inspection (DPI) applications. By utilizing integrated high speed IOs, coupled with the flexibility of a standard programming model and a scalable architecture, Tilera based systems are able to achieve world class performance on a variety of DPI applications.

High performance packet processing

- Two 10G XAUI Ethernet interfaces
- 64 independent processor tiles, programmable in C/C++
- DMA interface: packets delivered directly to tillers or memory
- Capable of up to 30 Million packets per second

The Network Intelligence architecture utilizes Tilera's DPI capabilities in order to support a high performance, flexible DPI architecture. Within MAXIM®, data flows into the system via multiple multi-Gbps links. Each link is processed by a Tilera card, running the Qosmos ixEngine software. The Tilera card is responsible for executing the ixEngine software, processing the line rate traffic and performing the layer 7 deep packet inspection. The level of processing is configurable, and may involve anything from flow and session level analysis, all the way up to application and user level analysis. While traffic is being inspected, both higher level metadata and raw packets may be sent to the Altix UV system via PCIE. Once at the Altix UV system, higher level analysis may be performed on the traffic, utilizing information from all multi-Gbps links to perform network intelligence across all multi-Gbps links in the system.

Qosmos ixEngine

Qosmos ixEngine is a Deep Packet Inspection (DPI) and Network Intelligence (NI) framework composed of software libraries and tools. Developers use this market-leading DPI and NI technology to accelerate the delivery of applications.

Features

Stateful inspection engine

- Traffic parsing, enabling:
  - Passive mode applications: processing of extracted metadata and content
  - Inline mode applications: smart packet filtering, traffic shaping or content filtering based on session context
- Packet classification at protocol and application levels using layer 7 deep packet inspection. No use of TCP/UDP ports.
- Correlation of data at flow, session, application and user levels.
- Real-time extraction of traffic metadata (caller, type of file downloaded, IMSI, etc.) and content (e.g. email text or VoIP stream based on RTP data).
- Automatic de-capsulation of tunneled / encapsulated traffic.
- Operates on fragmented, duplicated, de-sequenced packets.
- Operates on bidirectional and unidirectional traffic.
- Delivery of structured traffic information over standard APIs.
- Specific ixEngine version optimized for Tilera Tile Pro 64.
- Brings multicore computing to a new level, enabling DPI / Network Intelligence at multi-Gbps in real time.

Protocol plugin library
Qosmos' protocol plugin library includes over 500 protocol and application plugins to classify flows and extract metadata.

Protocol plugin SDK (Software Development Kit)
Allows users to develop custom protocol plugins that integrate in the ixEngine framework.

Protocol watch
Continuous update of ixEngine protocol plugins when new versions of protocols are released (e.g. MSN Messenger V8 to V9).

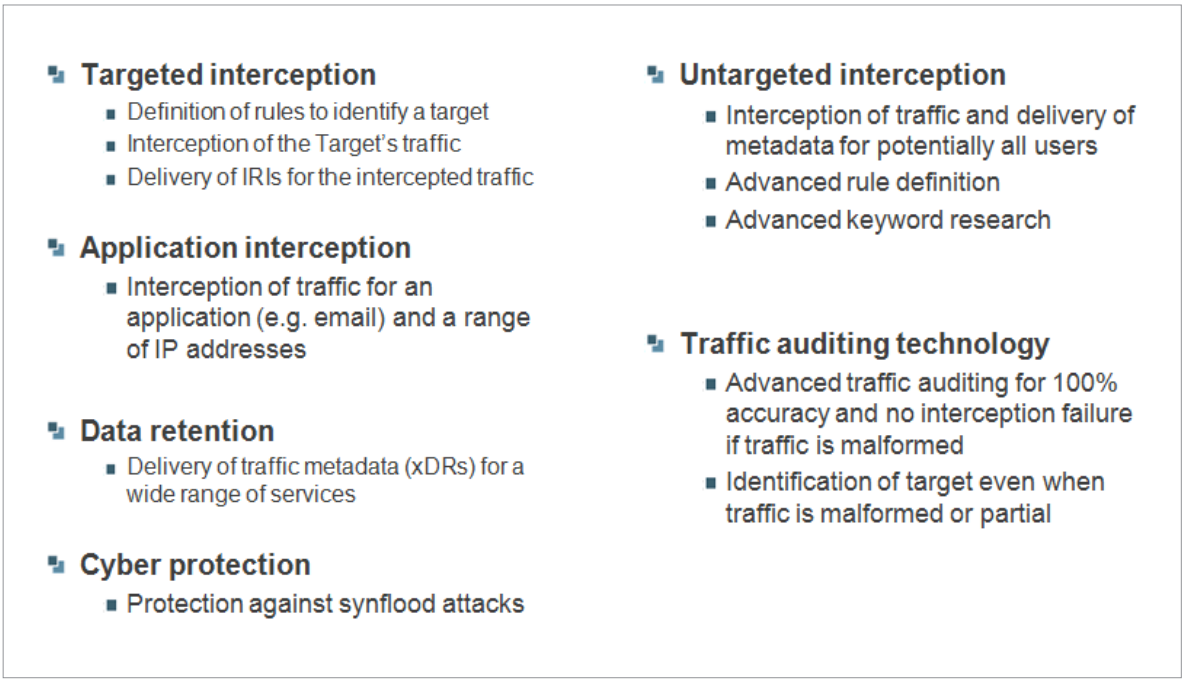

Figure 3: Qosmos ixEngine capabilities for interception

SGI Altix UV System

The Altix UV system is the latest generation of SGI scalable shared memory platforms. It is based on the Intel Xeon 7500 series processor. This processor has 4, 6, and 8 core variants. Altix UV runs standard versions of the Linux operating system from Novell and RedHat, ensuring binary compatibility with the x86 applications eco-system, which is unrivalled in the industry.

The Altix UV is built around a 2 socket blade, as the base unit of compute. These blades are coupled together using the SGI NUMAlink® 5 interconnect. What's unique about this interconnect, is that it allows the operating system to expand to manage the entire set of blades under one instance of the OS. This means a single OS can manage up to 2048 cores, 16TB of memory, as well as connecting PCIe devices to each of the processing blades. This ability to grow over time to these levels of capability is unique to Altix UV.

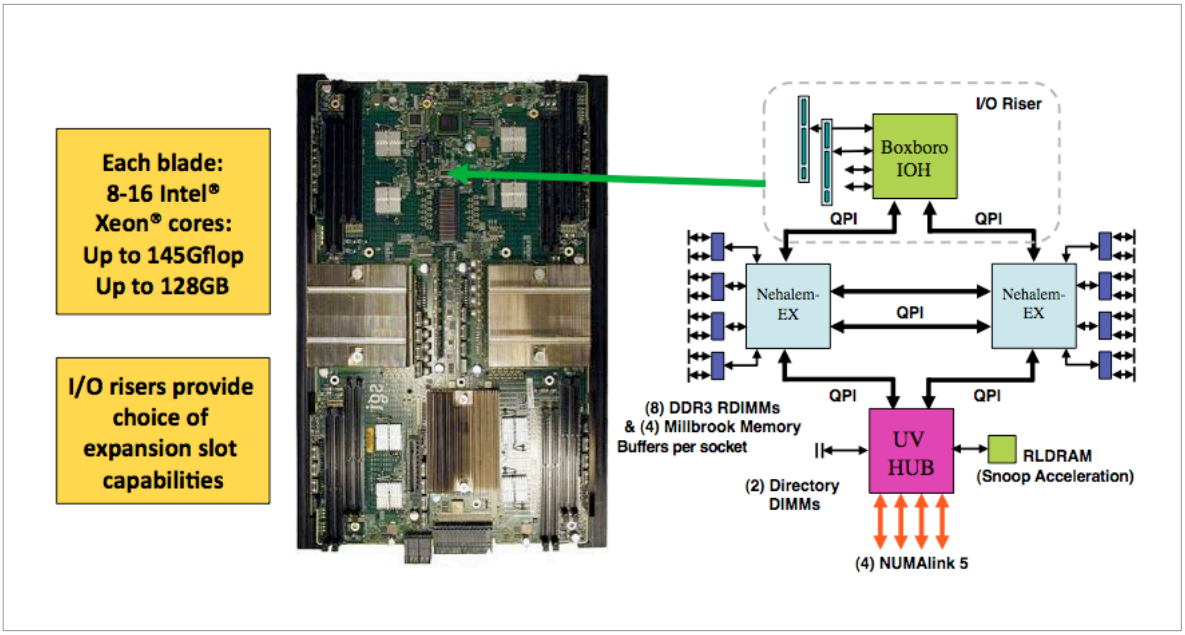
Each blade:
8-16 Intel® Xeon® cores:
Up to 145Gflop
Up to 128GB
I/O risers provide choice of expansion slot capabilities
I/O Riser
Boxboro IOH
QPI
QPI
Nehalem-EX
QPI
Nehalem-EX
QPI
QPI
(8) DDR3 RDIMMs & (4) Millbrook Memory Buffers per socket
UV HUB
RLDRAM (Snoop Acceleration)
(2) Directory DIMMs
(4) NUMAlink 5

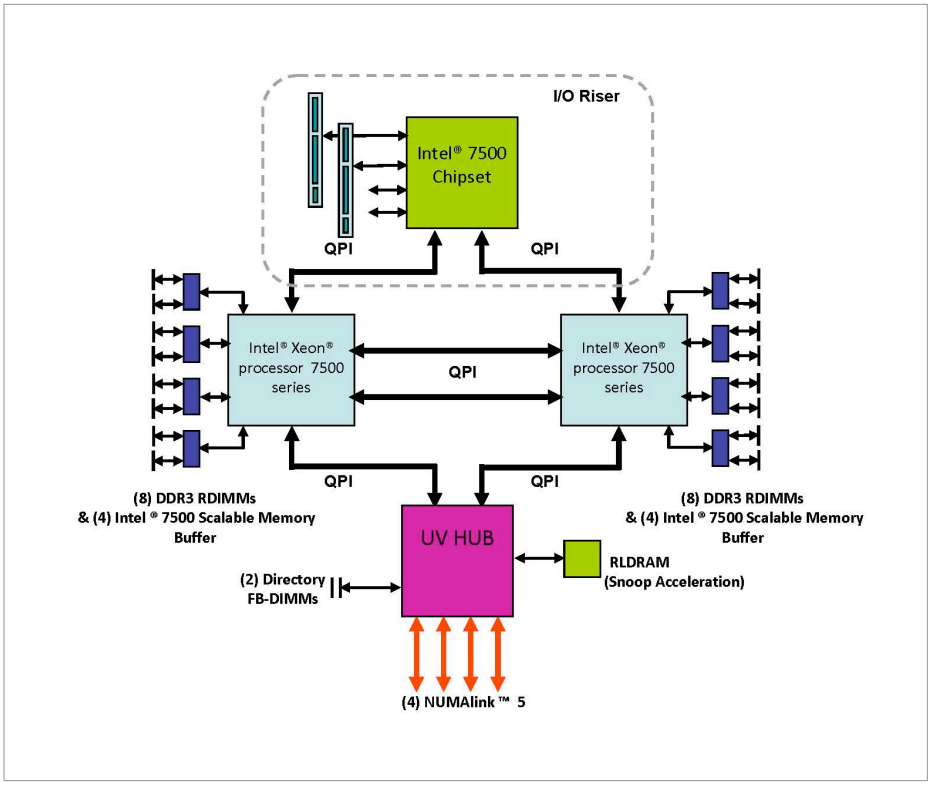
I/O Riser
Intel® 7500 Chipset
QPI
QPI
Intel® Xeon® processor 7500 series
QPI
Intel® Xeon® processor 7500 series
QPI
QPI
(8) DDR3 RDIMMs & (4) Intel ® 7500 Scalable Memory Buffer
UV HUB
RLDRAM (Snoop Acceleration)
(2) Directory FB-DIMMs
(4) NUMAlink ™ 5
(8) DDR3 RDIMMs & (4) Intel ® 7500 Scalable Memory Buffer

With this capability to scale, as a workload grows over time, the resources of the system can be increased to match the workload requirements, while still running as a singe system. This makes the programming of the system more flexible to cover growing demands. For example, it may be the case that over time, more network connections are needed to provide data ingest to the system. These can be added to existing nodes in the system, or to new nodes if additional processing and memory capability is needed. Adding blades increasing the memory available, so more historic data can remain resident directly in DRAM to ensure fast access for algorithms working on data correlation. Up to 16TB of DRAM is available directly to a user application, and with this extreme capacity, the latency to read any part of the memory is still under a micro-second.

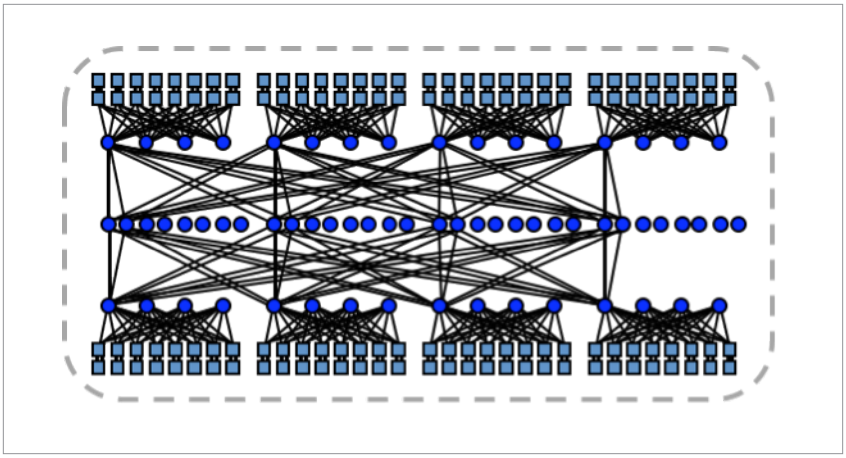

Putting It All Together

With these technologies coming together, different design points can be realized. For example, 8 Tilera PCIe cards, supporting a total of 16 10GE ingest streams, would provide the ability to extract metadata using the ixEngine, and feed it into the system at a rate of 8GB/s. This would fill memory at a rate one Terabyte every 2 minutes. A single Altix UV system memory of 16TB could capture up to 30 minutes of real time meta-data for event correlation. The Altix UV systems can grow even larger using the Altix UV with Jolt™ (up to 32K sockets, and 8PB of RAM). The intelligent communications accelerator technology embedded in NUMAlink enables direct user level program access to Petabytes of in-memory data.

Summary

The combination of SGI Altix UV, along with Tilera multicore packet processing and the Qosmos protocol parsing and extraction of metadata, enables unique processing of data streams and matching information from multiple data streams, across large periods of time, due the large capability to store data within a shared memory environment.

The combined solution provides a best-of-breed approach for environments which need to process massive amounts IP traffic, such as national cyber security and communications interception.

For more information, see:

http://www.sgi.com/products/servers/altix/uv/
http://www.qosmos.com/products/ixengine
http://www.tilera.com/products/processors/TILEPRO64

Corporate Office
46600 Landing Parkway
Fremont, CA 94538
tel 510.933.8300
fax 408.321.0293
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5185.2860
Europe +44 118.912.7500
Asia Pacific +61 2.9448.1463