TECHNICAL NOTE

Scaling Network Flow Analysis

FOR SECURING 10G/100G LINKS

FlowTraq | NetQuest Monitoring Access Solutions

Contents

Introduction 3

The Challenge 3

Solution Overview 4

- Requirements 4
- Instrumentation Architecture 5

Test Results: 100G Peering Link Unsampled Flow Analysis 8

- Test Results 10

High-Capacity Network Flow Monitoring Optimization 12

- 100G Unsampled Flow Generation Optimization 12
- Flow Collection & Storage Optimization 13

Scaling out N x 10G/100G/400G Flow Monitoring 15

Introduction

As Service Providers, Web-scalers and larger Government/Enterprise networks accelerate their migration to 100G links, the industry must respond to the challenges of scaling the subsequent network flow analysis. This technical paper presents solution architecture and test results for a real-world use case to demonstrate the efficacy of the proposed flow analysis solution for securing 10G/100G links.

This paper also discusses various infrastructure optimizations explored during the testing that improve the scalability of 10G/100G unsampled flow generation, collection, and storage.

The Challenge

The expansion of the Internet, the multitude of connected devices and the growing array of services and content continue to fuel network traffic growth. Fortunately, advancements in networking technology and favorable cost-per-bit economics are enabling organizations to upgrade their largest network links from 10G Link Aggregation Groups (LAGs) to 100G, and even 400G technology in some cases, including backbone links, service provider interconnects, Internet access links and growing cloud service provider interconnects.

Tier 1 Internet Service Provider (ISP) peering connections are at the epicenter of this traffic growth. ISPs are embarking on network migrations to 100G peering links to meet this demand but maintaining security visibility at these high speeds is increasingly difficult. Monitoring traffic directly in high-speed switches/routers can adversely impact network performance as traffic grows. Using CPU or NPU-based probes or Network Packet Brokers to inspect traffic with horizontal scaling to expand capacity quickly becomes cost prohibitive. Other approaches, such as traffic sampling to reduce processing burden, limits visibility and increases risk.

Given networking advancements and the drive to 100G/400G technologies is outpacing CPU/NPU processing advancements, security monitoring risks are increasing. As organizations rearchitect their security monitoring infrastructure to support migration to 100G links and beyond, new approaches are required to offload and optimize traffic analysis.

This paper highlights:

- Several key requirements for migrating a security monitoring infrastructure to support highly scalable flow analysis on 10G/100G links.
- An instrumentation architecture that addresses these requirements.
- Test results from a real-world use case validating unsampled flow analysis on a 100G Tier 1 ISP Peering Link.
- Several monitoring infrastructure optimizations explored during the testing to increase the performance and efficacy of the solution.
- A high-capacity flow monitoring architecture for scaling out to support multiple 10G/100G links, and the eventual migration to 400G links.

The focus of this paper is on scaling a security monitoring infrastructure to support high-capacity unsampled flow data generation, collection, and storage.

Solution Overview

Requirements

Unsampled Flow Analysis

NetFlow/IPFIX flow analysis via packet sampling has historically been used for network capacity planning and some security monitoring, e.g. DDoS attack detection. The inability of many networking equipment vendors to support 100% flow visibility is a significant reason why packet sampling is often employed. In recent years, traditional security breaches and volumetric attacks have been augmented with far more sophisticated attacks. A major challenge now is detecting/blocking bad actors while not falsely blocking legitimate users. Another challenge is detecting the new breed of DDoS attacks which attack "low and slow" to evade traditional detection methods.

Modern approaches must handle volumetric cases and source-based mitigation, detecting abnormal/unusual behavior from specific places and devices (not necessarily at volume), detecting distributed attack pattens, and detecting flows with potential "IP Reputation" conflicts. Addressing these challenges requires full unsampled flow visibility and more intelligent monitoring.

Cost-effective High-capacity Flow Generation

The NetFlow/IPFIX flow metering solution must:

- Support unsampled flow metering of traffic at rates up to 100Gbps per port.
- Have large flow cache memory per 100G port to ensure a robust unsampled solution, e.g. 50 million active flows per 100G port.
- Support high-capacity flow cache processing and flow record export to ensure 100% accurate unsampled flow reporting, e.g. at least 5 million flows per second per 100G port.

  Note: Many early implementations of 100G unsampled flow generation involve load balancing traffic and horizontal scaling of 10G flow generation probes due to performance limitations. This approach quickly becomes cost prohibitive and untenable as the number of 100G links grow. The flow generation requirements must be satisfied with technology that provides significant 100G port density and processing capacity to minimize rack space, power/cooling requirements and cost, e.g. up to 16 100G ports per 1RU with modularity to support a pay-as-you-grow deployment model.
- Support high-density 10G links in the same 1RU chassis to provide a graceful migration, e.g. up to 32 x 10G links per 1RU.
- Provide a seamless upgrade path to 400G in the same 1RU chassis.

Cost-effective High-capacity Flow Collection, Storage & Analysis

The flow collection, storage and analysis solution must be optimized for high-capacity flow processing, e.g. a single server shall support unsampled flow records for a 100Gbps link with peak traffic, e.g. 24/7 sustained flow rate of at least 250,000 flows per second. The solution architecture must scale to support millions of flows per second per cluster.

The flow storage architecture must possess high efficiency storage enabling it to keep up to 90 days of flow data at full flow fidelity, e.g. approximately 120TB per 100G link for 90 days. The flow database must also leverage configurable fast memory caching for high performance flow data ingest and analysis, e.g. support high-speed queries for historical flow data up to 48 hours into the past.

Instrumentation Architecture

Given the above challenges, we focus on identifying specific technologies and product designs that optimize high capacity flow generation, collection, and storage. The diagram below highlights an approach we employ which includes:

- A scalable 10G/100G/400G Visibility Fabric that provides low-cost, high-speed packet access
- A Service Node optimized for high-capacity packet processing & flow generation
- A Network Security & Forensics Platform purpose-built for high volume flow processing

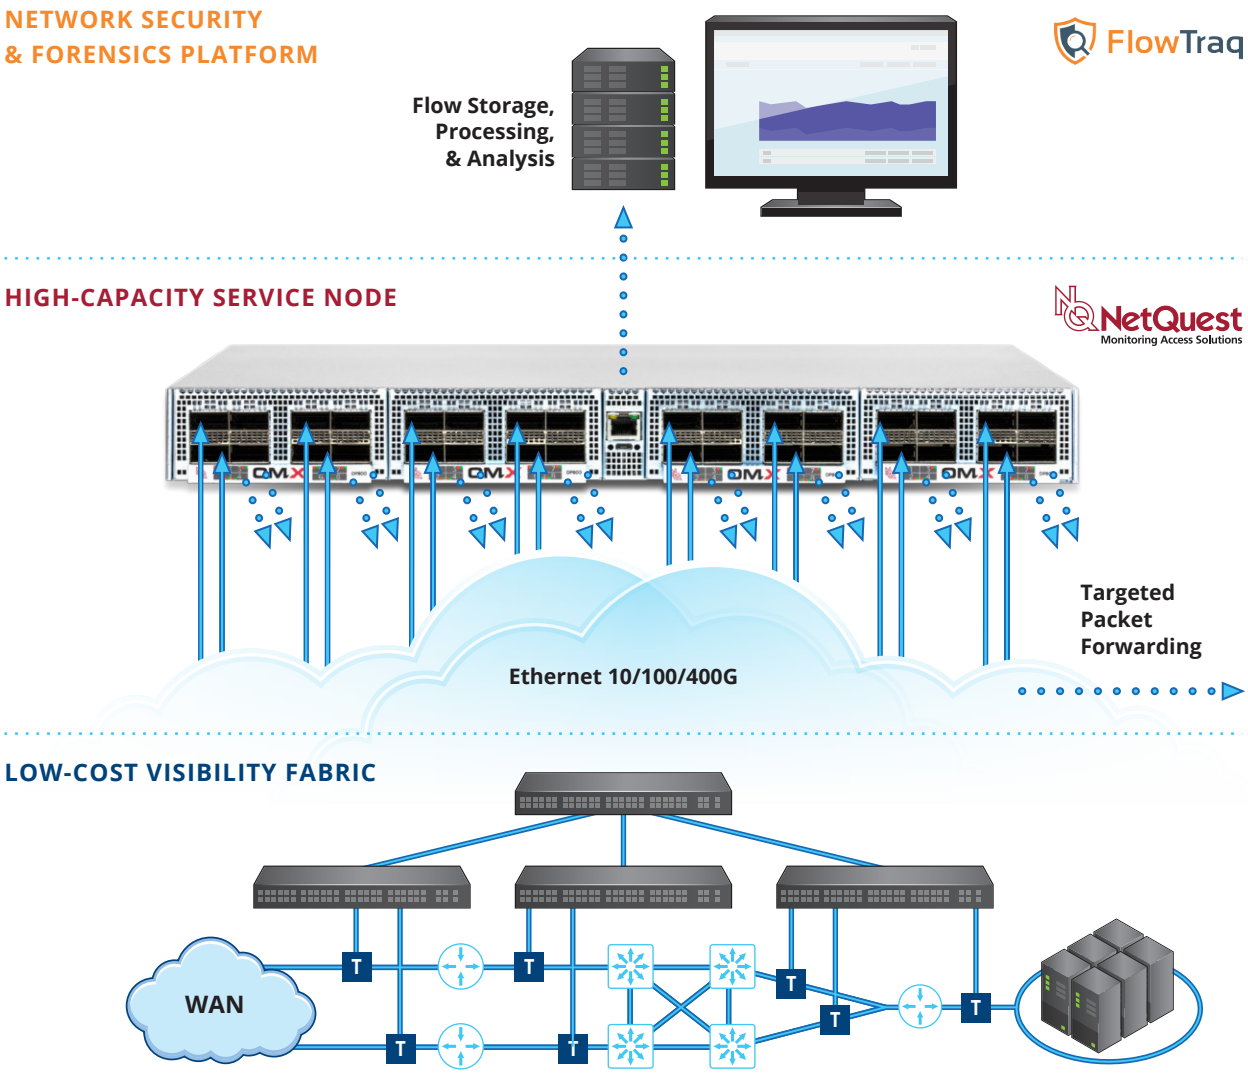

Figure 1: High-Capacity N x 10g/100g/400g Flow Generation & Analysis

The NetQuest OMX3200 Service Node receives network traffic from the Visibility Fabric, performs high-capacity unsampled flow metering, and exports the resulting IPFIX flow records to a cluster of bare metal commercial-off-the-shelf (COTS) servers running FlowTraq software that performs flow collection and security analysis.

Low-cost Visibility Fabric

Several leading switch/router vendors, e.g. Arista, Cisco, have introduced visibility fabric solutions leveraging their merchant silicon switches combined with management software to enable flexible packet access for cost-effective monitoring visibility. When compared to traditional Network Packet Broker offerings, the price-per-port advantage and broader networking technology options of this approach increase as port speeds and feeds migrate to 100G/400G. As a similar alternative, several software defined networking (SDN) vendors have introduced solutions that provide monitoring fabric management software to control 3rd party "white box" switches to achieve cost-effective visibility.

While both approaches solve the cost-effective 100G access problem, advanced packet processing and flow metering needs are typically left to either x86-based "Service Nodes" or legacy Network Packet Brokers, neither of which can cost-effectively scale to N x 100G when performing unsampled flow generation.

NetQuest OMX3200 High-capacity Service Node

The NetQuest OMX3200 High-capacity Service Node is a multi-terabit traffic flow visibility platform supporting metadata generation and advanced packet processing in a compact modular 1RU hardware platform.

The OMX3200 receives network traffic from either a Visibility Fabric (as shown in Figure 1) or direct from full-duplex (FDX) Network TAPs as show above. Each of the four OMX3200 modules support high-speed packet processing on 8 configurable QSFP28 ports (100G or 4x10G) with a total port density of up to 32x100G in 1RU. Each module supports the following services:

- Unsampled Network Metadata Generation transforming packets to flow-based metadata (IPFIX) for highly scalable N x 10G/100G Flow Analysis

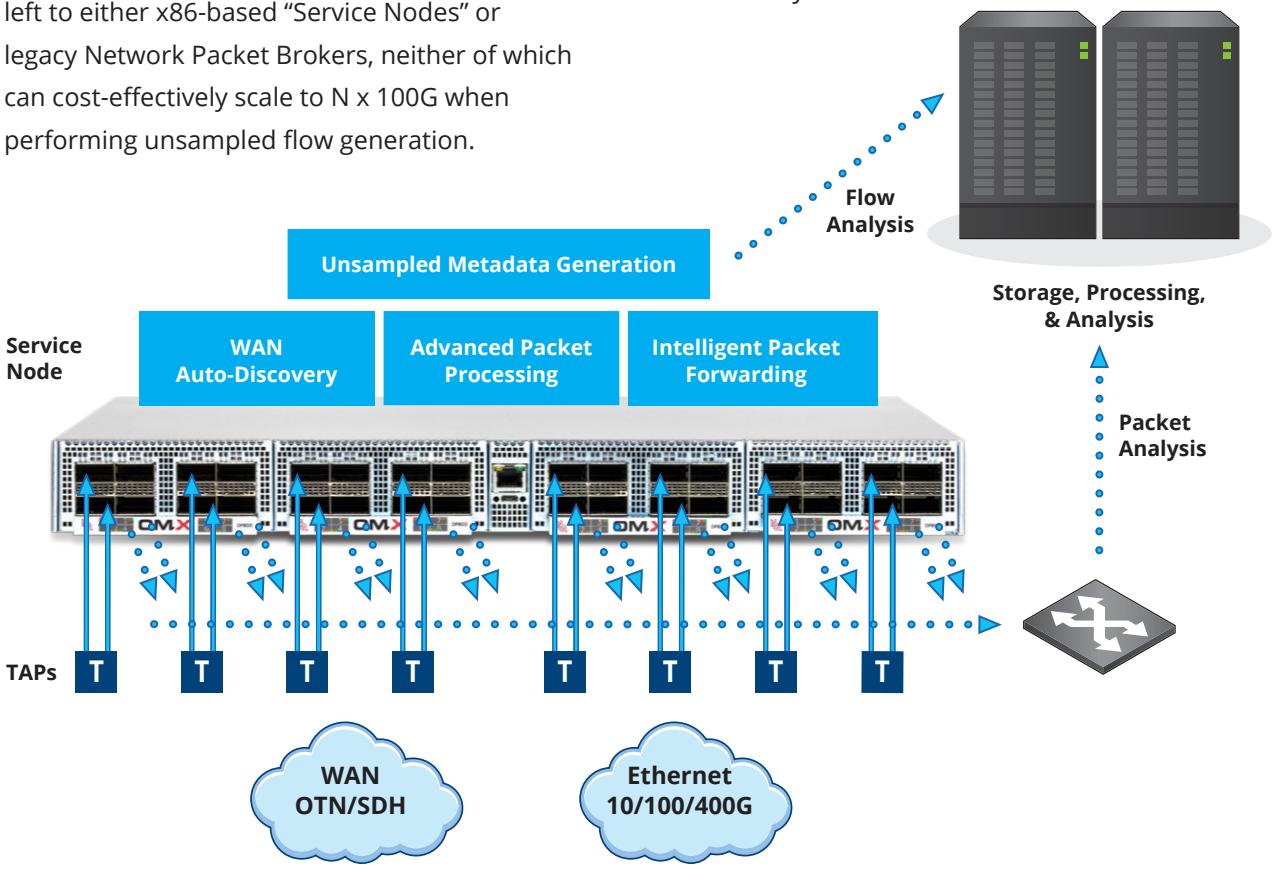

Figure 2: Netquest OMX3200 High-Capacity Service Node

- Advanced Packet Processing providing 10G/100G line-rate Intelligent Packet Forwarding including Header/Tunnel Stripping, Flow-based Packet Shunting, Packet Filtering with DPI Pattern Matching, Port Labeling and MAC Tagging to offload tool packet processing enabling scalable, targeted Packet Analysis
- WAN Auto-Discovery providing visibility within Optical Transport Network OTN & SONET/SDH links

FlowTraq Network Security and Forensics Platform

The FlowTraq Network Security and Forensics Platform gives security professionals and cyber hunters the complete visibility they need to monitor large-scale complex networks. This highly scalable solution helps detect and alert on suspicious activity. This helps prevent DDoS and brute force attacks, malware, zero-day threats, malicious botnets, new viruses, and other threat vectors. FlowTraq performs interactive network traffic analysis enabling network security engineers to detect network traffic anomalies, data breaches, risks to servers storing intellectual property, nation state-backed intrusions and social engineering attacks.

FlowTraq is optimized to collect and store large volumes of unsampled flow data to analyze network traffic. Its powerful network behavioral engine "learns" the various patterns of behavior on your network. Then, when any server, device, or specialized system starts to behave in a way outside these normal patterns, FlowTraq will generate real-time alerts.

Key capabilities of the FlowTraq architecture include:

- **Unrestricted Visibility** – FlowTraq gives the operator infinite control over how to filter the data, and how to view the data.
- **Big Data Database** – built for modern parallel architectures, FlowTraq makes optimal use of all compute/IO resources so you do not need to aggregate your valuable flow data and reduce visibility.
- **Triple-Split Storage Architecture** – the most recent flow data is quickly and directly available from the RAM database, longer queries are serviced from a local SSD database, while the furthest history is stored on archive that can be provisioned on spinning RAID, or remote SAN/NAS solutions.
- **Long-Term Trends & Baselines** – builds baselines of normal traffic and stores long-term trends for all interfaces and network traffic groups to quickly evaluate peering relationships and link utilization over years of collected data.
- **Traffic & Security Alerts** – detects a complete range of network anomalies, from simple traffic threshold alerts, deviations from baselines, to Distributed Denial of Service events and security violations.
- **Distributed Load Balanced Clusters** – patented flow analysis clustering technology allows you to run over multiple servers, combining storage and processing power to handle unlimited quantities of flow data.

Test Results: 100G Peering Link Unsampled Flow Analysis

This section presents the results of using unsampled IPFIX to monitor and secure a Tier 1 ISP 100G peering connection. The tests use a 100G full-duplex (FDX) traffic profile based on the characteristics of a live production 100G peering connection between two Tier 1 ISPs. The results demonstrate the following:

- OMX3200 IPFIX IPv4/IPv6 flow record interworking with FlowTraq collection and analysis.
- OMX3200 IPFIX flow metering and export for a 100G FDX peering link at average & peak traffic levels using a single module with 100% unsampled performance.
- FlowTraq IPFIX record collection and analysis for a 100G FDX peering link at average & peak traffic levels using a single commercial-of-the-shelf server with 100% unsampled performance.
- OMX3200 100G FDX line-rate Advanced Packet Processing and Intelligent Packet Forwarding to packet-based tools in parallel, with no impact on IPFIX generation processing.

OMX3200 High-capacity Service Node

Single OMX3200 chassis, 1RU, with single 8-port module installed

FlowTraq Bare-metal Server Configuration

- FlowTraq installed on Dell R740 Dual Intel Xeon Platinum 8280 2.7GHz, 56 Core/112 Thread,
- 384GB 2933MHz RAM, 4x 400GB SSD SAS 12Gbps
- Flow ingest: 10G SFP+ Intel XL710 NIC
- Management: 1GBaseT LOM

OMX3200 module Port 1 & Port 2	OMX3200 module Port 8
100G FDX ingress, peering link Tx & Rx traffic Packets also hairpinned to 100G egress on same ports	10G egress, IPFIX export

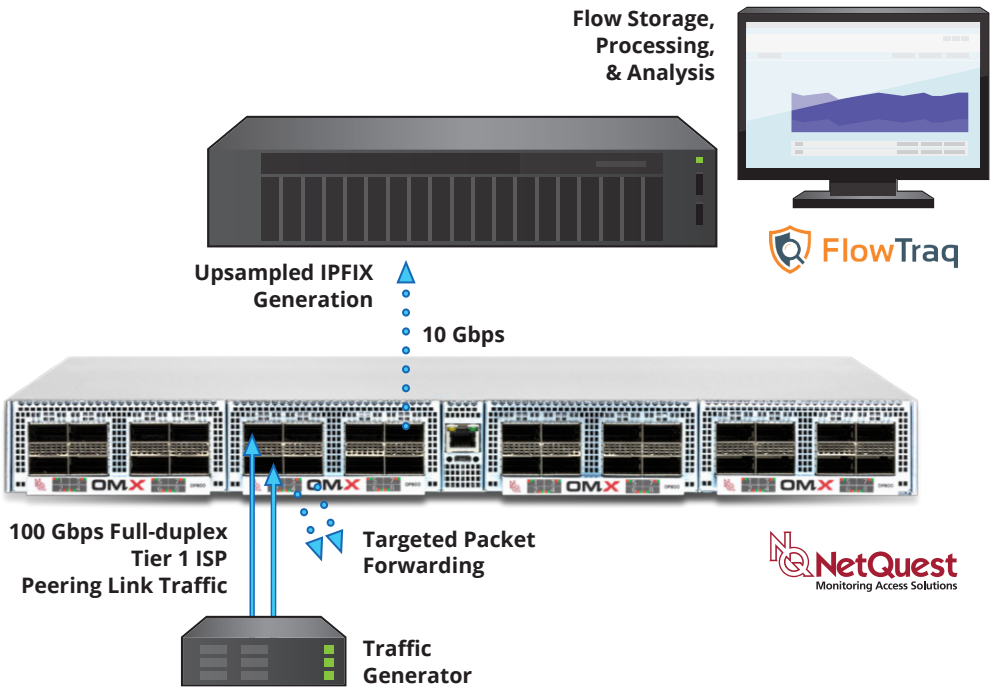

Figure 3: Tier 1 Isp 100g Peering Link Flow Analysis Test Environment

Tier 1 ISP 100G Peering Link Traffic Profile

The traffic profiles below represent an averaging of traffic characteristics measured across multiple 100G peering links in a major US city in 2020. The table below shows the average traffic characteristics measured on the 100G peering links:

100G Tier 1 ISP Peering Link – Average Traffic Profile					
Rx			**Tx**		
Avg Pkt Rate	Avg Traffic Rate	Avg Pkt Size	Avg Pkt Rate	Avg Traffic Rate	Avg Pkt Size
5–6 MPPS	55–60 Gbps	1350–1400	3–4 MPPS	7–8 Gbps	210–240

The table below shows the peak traffic characteristics measured on the 100G peering links:

100G Tier 1 ISP Peering Link – Peak Traffic Profile					
Rx			**Tx**		
Peak Pkt Rate	Peak Traffic Rate	Avg Pkt Size	Peak Pkt Rate	Peak Traffic Rate	Avg Pkt Size
7–9 MPPS	95–98 Gbps	1350–1400	5–7 MPPS	10–14 Gbps	210–240

Other interesting 100G peering link traffic characteristics measured:

Avg IPFIX Flows Per Second	Avg Packets Per Flow	Avg Flow Duration	Avg Long Flow Duration
125–250 KFPS	90–110 packets	3–4 sec	1000–1200 sec

We used Cisco TRex traffic generator software with Mellanox ConnectX-5 NICs to recreate stateful 100G FDX session traffic resembling the above Tier 1 ISP Peering Link traffic. This includes modeling the asymmetric nature of the traffic, the average packet size in each direction, the packet rate, and the inclusion of long flows. The application mix consists of the following:

- BitTorrent, BitTorrent Handshake, DNS, Ebay, Facebook SSL, Google, Gmail, Google Maps, ESPN, HTTP, Netflix, SMTP, Generic TCP, Generic UDP, Youtube-SSL

Traffic Profile 2 – Tier 1 ISP Peering Link (100G FDX – Peak Load)

A total of 30,441,565 sessions were generated to model the PEAK peering link load while IPFIX generation performance statistics were recorded at 10 second intervals. The OMX3200 metered, processed, and forwarded 100% of the packets. FlowTraq collected/stored 100% of the flow records with < 80% Input Buffer utilization and < 2% Write Buffer utilization, showing FlowTraq supports 100G FDX links at Peak Load (but is approaching input flow record processing capacity).

P1–Avg Pkt Size	P1–Avg Pkt Size	Avg Flow Duration	Avg Pkts/Flow	Long Flow %	Avg Long Flow Duration
1379	185	3.54 sec	62	0.90	73 sec

MPPS	P1 Gbps	P2 Gbps	Export KFPS	Export Mbps
1.61	10.12	1.48	38.60	35.72
3.22	22.98	2.18	85.67	51.33
4.56	34.63	2.72	107.41	62.45
6.11	47.20	3.12	133.43	88.09
7.53	59.24	3.51	149.45	89.28
8.02	61.72	3.61	154.01	91.75
8.03	61.73	3.62	142.50	92.47
8.03	61.73	3.62	156.31	92.92
8.03	61.73	3.62	156.35	93.03
8.03	61.73	3.62	143.25	92.89
8.03	61.73	3.62	156.10	92.74
8.03	61.73	3.62	156.35	92.83
7.85	60.72	3.33	142.54	91.29
6.06	48.70	1.98	100.91	49.98
4.49	36.13	1.27	69.65	41.01
3.16	24.50	0.80	49.61	24.79
1.74	12.86	0.46	7.73	4.58
0.35	1.39	0.07	6.07	2.81
0.01	0.00	0.00	1.77	1.02
0.00	0.00	0.00	0.45	0.08
0.00	0.00	0.00	0.00	0.00

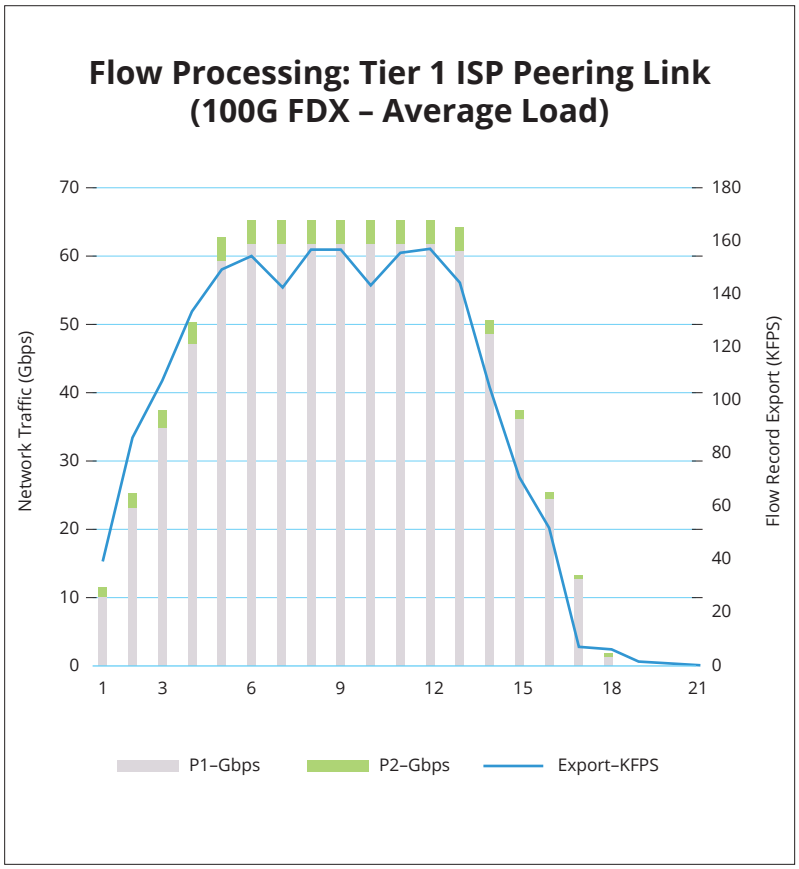

Test Results

Traffic Profile 1 – Tier 1 ISP Peering Link (100G FDX – Average Load)

A total of 19,371,944 sessions were generated to model the AVERAGE peering link load while IPFIX generation performance statistics were recorded at 10 second intervals. The OMX3200 metered, processed, and forwarded 100% of the packets. FlowTraq collected/stored 100% of the flow records, more than 150,000 flows per second (FPS), with < 1% Input Buffer utilization and < 1% Write Buffer utilization, indicating ample available capacity.

P1–Avg Pkt Size	P1–Avg Pkt Size	Avg Flow Duration	Avg Pkts/Flow	Long Flow %	Avg Long Flow Duration
1379	185	3.52 sec	62	0.82	73 sec

MPPS	P1 Gbps	P2 Gbps	Export KFPS	Export Mbps
2.41	15.45	2.25	54.85	50.83
4.95	34.92	3.41	134.05	80.47
7.06	53.22	4.23	149.71	91.02
9.31	72.13	4.83	203.99	137.52
11.55	90.60	5.42	233.91	139.41
12.59	96.99	5.68	241.07	143.86
12.62	97.00	5.69	223.77	145.38
12.62	96.99	5.68	245.52	145.82
12.62	97.00	5.69	245.63	146
12.62	97.01	5.69	245.59	145.98
12.62	97.01	5.69	224.95	145.92
12.62	97.00	5.69	245.42	145.69
12.62	97.00	5.69	225.28	145.93
12.62	97.00	5.69	245.55	145.88
12.62	97.01	5.69	245.56	145.95
12.62	97.01	5.69	225.03	145.82
12.62	97.01	5.69	245.61	145.91
12.62	97.00	5.68	225.13	145.94
12.62	97.00	5.69	245.42	145.81
12.62	97.01	5.68	245.59	146
12.62	97.01	5.69	225.09	145.88
12.62	97.00	5.69	245.53	145.88
12.62	97.00	5.69	225.15	145.86
12.62	97.00	5.69	245.48	145.84
12.62	97.00	5.69	245.59	145.95
12.46	96.21	5.44	224.94	145.62
9.93	79.37	3.27	169.06	84.95
7.43	59.34	2.13	101.15	65.02
5.14	41.05	1.35	89.73	47.62
3.08	22.86	0.79	16.98	7.23
0.85	3.94	0.20	9.35	5.42
0.02	0.01	0.01	3.54	1.79
0.00	0.00	0.00	1.09	0.38

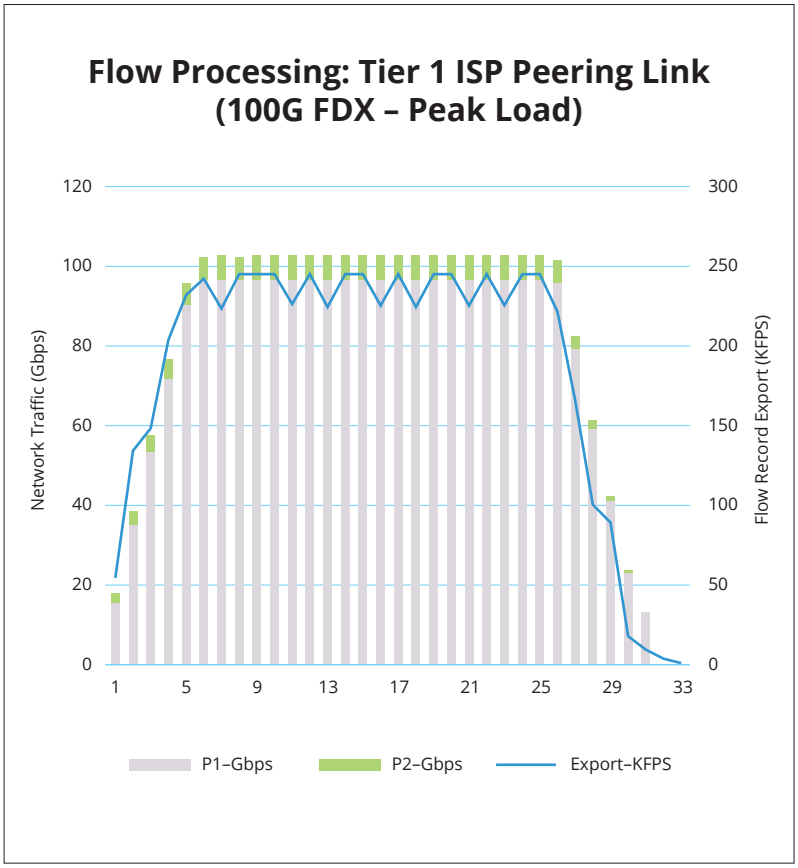

High-Capacity Network Flow Monitoring Optimization

This section presents several monitoring infrastructure optimizations explored during the testing to increase the performance and efficacy of generating unsampled flow records on 100G links and collecting/storing the resulting IPFIX flow records.

100G Unsampled Flow Generation Optimization

Flow Metering Optimization

The OMX3200 High-capacity Service Node supports four modules, and each module supports IPFIX flow generation on 4 x 100G ingress interfaces (up to 16 x 100G per 1RU chassis). Each module uses Field Programmable Gate Array (FPGA) state machines and high-speed memory that support deterministic, high-capacity unsampled IPFIX packet processing performance. The Tier 1 ISP 100G Peering Link Traffic Profile statistics above show a peak traffic rate of Rx: 7–9 MPPS + Tx: 5–7 MPPS = 16 MPPS per 100G peering link. The OMX3200 has sufficient processing capacity to support two 100G FDX peering links with 100% unsampled performance (4 x 100G per module).

Note: The OMX3200 also supports configurable line-rate Microburst Detection on all 100G ports. Each module provides high-precision Framerate (MPPS) and Bandwidth (Gbps) metrics for granular traffic monitoring.

Flow Cache Optimization

The OMX3200 High-capacity Service Node supports configurable Active Flow and Flow Inactivity timers to balance flow reporting frequency and flow cache utilization. Given each module supports a very large 200M+ Active Flows Cache, the customer has full flexibility to adjust reporting frequency as desired. In the tests above we selected Active Flow Timer = 60 seconds and Flow Inactivity Timer = 30 seconds.

Flow Export Optimization

Unsampled flow generation on 100G links creates significant performance challenges for many Flow Collectors. The OMX3200 High-capacity Service Node supports several features to help the scalability of Flow Collectors.

Export by Observation Domain vs. Observation Point

The IPFIX standard defines Observation Domain as the largest set of Observation Points (interfaces) for which flow information can be aggregated by a Metering Process. An OMX3200 module can be configured for Export Mode as follows:

- **Export by Observation Domain** – each module is an Observation Domain and flow records for its up to four 100G ports are aggregated within IPFIX packets.
- **Export by Observation Point** – each 100G port is a unique Observation Domain and flow records for each 100G port are segregated within a given IPFIX packet.

Export by Observation Point yields modest performance improvement with some Flow Collectors, including the FlowTraq Network Security and Forensics Platform.

Export MTU Size

The OMX3200 High-capacity Service Node supports a configurable export Maximum Transmission Unit (MTU) from 1500 bytes to jumbo frames. Larger frame sizes provide a small benefit with export bandwidth efficiency. However, they also impact Flow Collector receive-

side processing. Some collectors may have better performance with fewer/larger IPFIX packets; in this testing the FlowTraq Network Security and Forensics Platform had slightly better performance with smaller IPFIX packets, i.e. MTU = 1500. Note that the Flow Collector 10G NIC ring buffer should be configured with sufficient regular and jumbo buffers based on the MTU setting to avoid Rx overruns/drops.

Export Load Balancing

The OMX3200 High-capacity Service Node supports IPFIX export load balancing to accommodate flexible scaling of flow collectors. Flow records for each metered 100G port can be sent to a Destination Group with up to 4 Flow Collectors, where each collector is defined by an IP Address/Port pair. The OMX3200 supports flow-based load balancing to ensure a given collector will receive ALL flow records for a given flow.

The FlowTraq Network Security and Forensics Platform is both highly performant and flexible in collecting flow records. We tested exporting flow records to up to 4 listening ports within FlowTraq to see if the aggregate collection performance increased. We found that FlowTraq performs best when all flow records are sent to a single listening port.

Flow Collection & Storage Optimization

Linux OS/Kernel Optimizations

The FlowTraq Network Security and Forensics Platform runs on the Linux operating system. Given the real-time nature of FlowTraq processing, several OS/Kernel optimizations were found to improve performance at scale.

Disable the Linux IO Scheduler

This must be done for all drives that store flow data. Use the "lsblk" command to verify the drive being used. The command to disable the scheduler is:

```
echo "noop" > /sys/block/sdb/queue/
scheduler
```

Simple trick to hit them all:

```
for DISK in `lsblk -n -d | awk '{print $1}' `
do
echo "noop" > /sys/block/${DISK}/queue/
scheduler
done
```

To ensure this persists across a reboot, insert the above in the '/etc/rc.local' file.

Disable "Transparent Hugepages" in the Kernel

The kernel assigns very big memory pages for continuous regions of memory. Since FlowTraq writing is very linear, it will often elect to stick large memory maps of the database files in a "huge page". But this means that when only a single very small update happens to an older flow, the entire huge page must be faulted into memory (lots of IOPS), and then written back (lots of IOPS) again.

```
echo never > /sys/kernel/mm/transparent_
hugepage/enabled

echo never > /sys/kernel/mm/transparent_
hugepage/defrag
```

To ensure this persists across a reboot, insert the above in the '/etc/rc.local' file.

One thing to note is the current running system will have many hugepages assigned to IO blocks, so it will take some time (running overnight) before the improvements will be noticed, as the huge pages are cycled out.

Disable Access Time Updating of Files on ext4 Filesystem

Edit /etc/fstab and insert “noatime” as shown, and then reboot so the mounts take effect.

```
/dev/sda1 / ext4 errors=remount-ro,noatime
0 1
```

Intel Server Hyperthreading

The FlowTraq Network Security and Forensics Platform ran on a dual-socket Dell/Intel server, with two CPUs that each have 28 cores/56 threads. We compared performance with hyperthreading enabled and disabled, and found a modest improvement in FlowTraq processing with hyperthreading enabled.

FlowTraq Application Optimizations

Input Buffer Size

The FlowTraq Network Security and Forensics Platform provides a configurable input buffer size to accommodate high packet rates, especially bursty packet rates. We encountered a “99% Flow Input Buffer” FlowTraq warning during initial peak traffic testing and we often detected dropped sessions between the OMX3200 and FlowTraq when this occurred. Given our high flow rates, we increased the buffer size from the 16MB default to 256MB in the /opt/flowtraq.conf file and there was a significant increase in loss-free performance.

Ingest Thread Processing

The FlowTraq Network Security and Forensics Platform provides a configurable number of threads that process ingested packets. When there is a high number of cores, the resulting threads processing ingested packets contend with each when writing information to session tables. By default, FlowTraq allocates “processthreads” equal to 50% of the cores. Given our large 56 core/112 thread server, we reduced the number of threads and compared performance. We found that allocating processthreads = 1 in the /opt/flowtraq.conf file results in the best FlowTraq ingest performance for the server used in these tests.

Intermediate Buffers

The FlowTraq Network Security and Forensics Platform provides a configurable intermediate buffer size to accommodate high packet rates, especially bursty packet rates. Given our high flow rates, we increased “’bufferslots” from the 16K default to 32K in the /opt/flowtraq.conf file and there was a minor increase in loss-free performance.

Scaling out N x 10G/100G/400G Flow Monitoring

Addressing today's security monitoring and digital forensics challenges requires full unsampled flow visibility and precision monitoring. Scaling a flow monitoring architecture for multiple 10G and 100G links, and the eventual migration to 400G links, requires cost-effective, high-capacity unsampled flow generation using a state-of-the-art FPGA implementation, and pairing this with a flow collection, storage and analysis architecture built for scale. The architecture below has been proven to provide a cost-effective growth path from high-density 10G unsampled flow monitoring to 100G as your network traffic grows.

The FlowTraq Network Security and Forensics Platform gives security professionals and cyber hunters the complete visibility required to monitor large-scale complex networks. This paper demonstrates the platform scales to more than 250,000 flows per second per server. Further, the FlowTraq design supports scaling out with individual bare-metal Worker Nodes performing flow collection and processing, and a single Portal Node providing cluster orchestration and powerful, centralized flow analysis for millions of flows.

The NetQuest OMX3200 Service Node provides unsampled flow generation and performs advanced packet processing and targeting in parallel to also support packet analysis. The OMX3200 compact modular pay-as-you-grow design supports up to four packet processing modules, and up to 32 x 10G, 16 x 100G and 8 x 400G* links per 1RU system. The flexible design allows a graceful migration from high-density 10G link unsampled metering to 100G links by simply soft-reconfiguring a module.

* OMX3200 Service Node 2 x 400G module in development

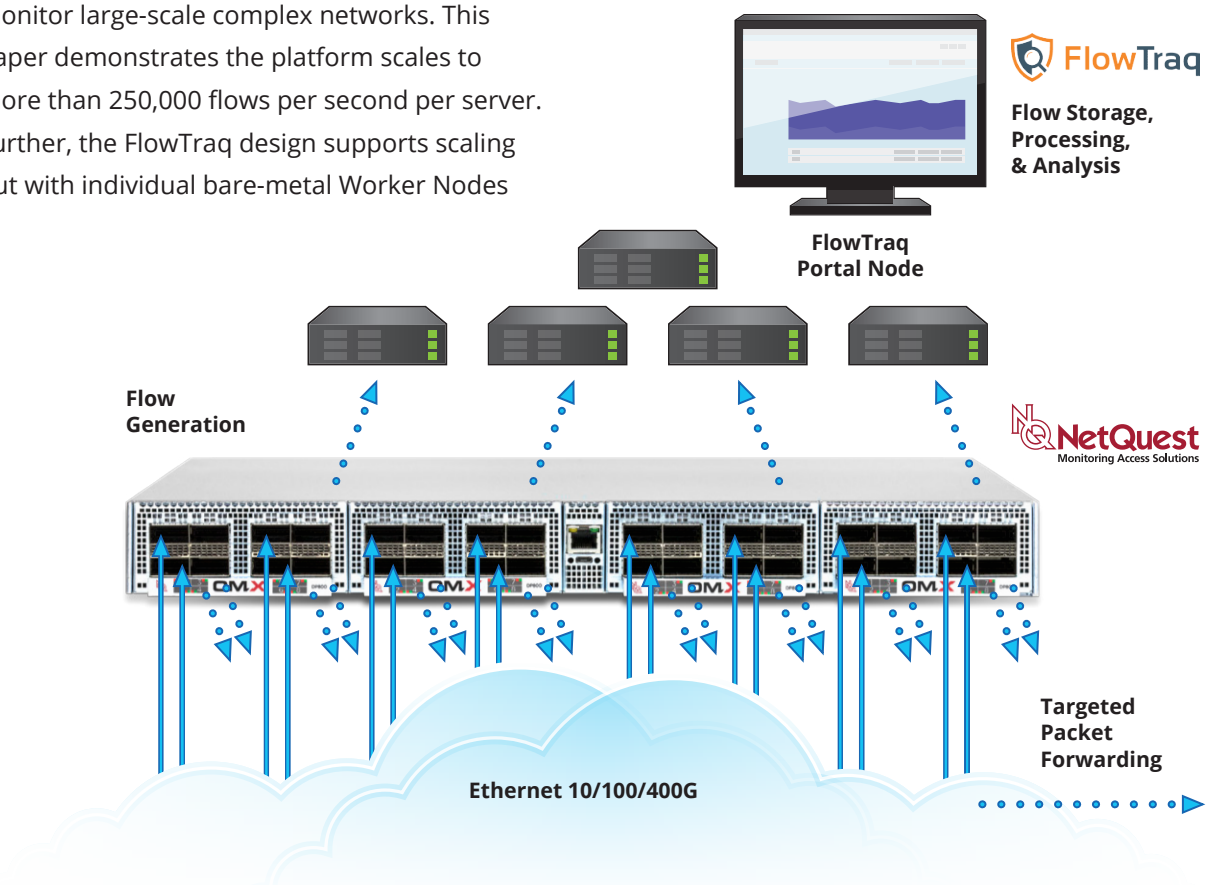

Figure 4: High-Capacity N x10g/100g/400g Unsampled Flow Analysis Architecture

About NetQuest Corporation

NetQuest has been a longstanding and trusted supplier of Cyber Surveillance Appliances to government agencies. With the introduction of the OMX3200, we have built upon our many years of network monitoring experience to offer an optimized network visibility solution ideal for cyber intelligence applications serving government, service providers and large enterprises.

For more information about the NetQuest OMX3200 Service Node visit: **www.netquestcorp.com**

About FlowTraq

FlowTraq®, a Riverbed Company, provides software and services for high performance network monitoring, analytics, security and forensics to detect a range of network behaviors, including distributed denial of service (DDoS), brute force attacks, botnets, worms, network scans and other network traffic anomalies.

To request a free trial of FlowTraq, visit:
www.flowtraq.com/product/free-trial