

100G Monitoring: More Bandwidth, New Challenges Josh Joiner, Director of Solutions Engineering June 2017

Challenges for Financial Trading Networks



100G Deployment Accelerating



100G port shipments forecasted to pass 40G in 2017, 1G in 2018



Networks Confidential and Proprietary. Not to be distributed without express written consent of cPacket Networks, Inc.

© 2017 cPacket Networks

100G By the Numbers – Capturing to Disk

$$\frac{10^{11}bits}{1\,second} * \frac{1byte}{8\,bits} = 12,500,000,000\ bytes/second$$
(1)
$$\frac{10^{12}bytes}{12,500,000,000\ bytes/second} = 80\ seconds/terabyte$$
(2)
$$\frac{86,400\ seconds}{1\ day} * \frac{1\ day}{80\frac{seconds}{terabyte}} = 1,080\ terabytes = 1.08\ petabytes$$
(3)

At 100 Gbps, 1TB is filled every 80 seconds and...

→ 1.08 petabytes is required to store 24 hours of data



© 2017 cPacket Networks Confidential and Proprietary. Not to be distributed without express written consent of cPacket Networks,

100G By the Numbers – Burst Size

(1)

(2)

 $\frac{10^{11}bits}{1\,second} * \frac{1byte}{8\,bits} = 12,500,000,000\ bytes/second$

 $\frac{12,500,000,000 \text{ bytes}}{\text{second}} * \frac{1 \text{ second}}{1000 \text{ milliseconds}} = 12,500,000 \text{ bytes/millisecond}$

At 100 Gbps, a one millisecond spike is 12.5 megabytes...

this is more data than all the memory in some switches

How much market data is dropped due to microbursts?



© 2017 cPacket Networks Confidential and Proprietary. Not to be distributed without express written consent of cP

Putting line-rate Performance at the Wire

- Packet brokers are oversubscribed by design
 - Providing multiple monitoring points is good but not when they create choke points
- Capture NICs and write to disk reduce confidence
 - Garbage-in, garbage-out can't make good decisions with bad data
- Deploy purpose-built hardware monitoring at the wire:
 - Performance and accuracy
 - Scalability
 - Distributed



cPacket's Solution vs. Traditional NPM Solutions





Always On Millisecond Accuracy

Spike detection for market data feeds

Challenges

- Markets = Event Driven
- Spikes = Lost Data
- Spikes Are Invisible

Solution - cBurst Predictive high-resolution monitoring that identifies problems before end-users are impacted





cBurst Finds Problems Before Users Impacted



Proactive alerts identify problems in real-time

233.200.79.128:62128 (CTS 1/Network A)

Allowing you to identify the **specific** problem feed

High Resolution Analytics

Identify problems your other tools can't see

Challenge

 Low Resolution (>1 Sec) monitoring hides your network problems

Solution: Millisecond Resolution Identifies problems that your other monitoring tools can't detect

High Resolution Analytics Definition Matters

Order of magnitude difference between millisecond and second

Market Feed Gap Detection

Sequence anomaly detection in real time

Challenges

- Traffic Bursts = Gaps
- NPBs = Gaps or False Pos
- Full Coverage = \$\$\$\$

Solution: cPacket MFGD

Real-time detection of sequence anomalies (gaps) performed in hardware at line rate

Identify Gaps in Market Data Feeds

Deploy atErritica lapoi set inlanfmastructure

Identify Gaps in Market Data Feeds

		The Wireshark Network Ana	lyzer				<u>×</u>		
		Elle Edit View Go Capture	Analyze Statistics Hel	р — <u>—</u> .					0
1 2 3 4	5 6 7 8 9		📙 × 🕲 📇 🛛	3 🗢 🛸 轮 🚡 🖞	Ł	📑 🔍 Q, Q, 🖤 M 🖬 M 🚻 💥			
de: O Live Timeline Edit			💌 💠 Expression 🗞 Clear 🖋 Apply				n alerts: 🗹	user alerts: 🧭	search: 🗹
30:04 # 321	, 47,	40, 50	ource	Destination	Protocol	Info			
		-	7.183.142.87	204.252.103.16	TCP	1013 > 22 [FIN, ACK] Seq=3084 Ack=644 Win=	-		
		115 60 53.550000 2	04.252.103.16	207.183.142.87	TCP	22 > 1013 [ACK] Seq=644 Ack=3085 Win=16384		0.90%	_
		116 60 53.550000 2	04.252.103.16	207.183.142.87	TCP	22 > 1013 [FIN, ACK] Seq=644 Ack=3085 Win=		0.0070	
		117 54 53.550000 2	07.183.142.87	204.252.103.16	TCP	1013 > 22 [ACK] Seq=3085 Ack=645 Win=32256			
Client 1	Time	118 342 53.920000 2	04.252.103.79	255.255.255.255	BOOTP	[Packet size limited during capture]		rm	1
		119 240 54.210000 0	0000000.00609/3960/1	00000000.TTTTTTTTTTTTT	NMPI	[Packet size limited during capture]			
		120 189 54.250000 0	0:20:41:92:04:51	03:00:00:00:00:01	SMB	[Packet size limited during capture]			
		121 60 54.630000 0	07 192 142 97	204 252 102 2	BOD	Cont. Root = 65535708:00:40:08:50:56 Cost			
		123 66 54 710000 2	04 252 102 2	207 183 142 87	POP	Response: +OK 2 3467			
Client 2 2017-05-13 12:55:18 1 Frame 122 (60 bytes on wire, 60 bytes captured) Ethernet II, Src: 00:c0:4f:c7:eb:c0 (00:c0:4f:c7:eb:c0), Dst: 00:00:0c:36:00:19 (00:00:0c:36:00 Internet Protocol, Src: 207.183.142.87), Dst: 204.252.102.2 (204.252.102.2) Internet Protocol, Src: 207.183.142.87 (207.183.142.87), Dst: 204.252.102.2 (204.252.102.2)							5,227		_
		Source port: 22587 (225 Destination port: 110 (
2017-05-13 16:22:55 Sequence number: 29 (relative sequence number) [Next sequence number: 35 (relative sequence number)] Acknowledgement number: 134 (relative ack number)									
		010, 1011, 10	49	0 45 006 0 7 cc fcu.@.@. 4. 0 50 18 f.X;.nj }x=ST AT.	E. W P.			2	
		Sequence number (tcp.seq), 4 byt	es] [] F	: 3632 D: 3	632 M: 0	-		

Access parte de points ten biet on "proof"

© 2017 cPacket Networks Confidential and Proprietary. Not to be distributed without express written consent of cPacket Networks, Inc.

Timestamping for MiFID II The Industry's Most Accurate Timestamp

Challenges

- MiFID II Compliance
- Contamination From Queuing & Buffering
- Necessary for Diagnostics

Solution:

cPacket nanosecond timestamp is most accurate in the industry

cPacket's Built-in Timestamp Significantly Exceeds MiFID II (RTS-25) Minimum Requirements

MiFID II RTS-25 Requirements

Type of Trading Activity	Divergence from UTC	Granularity Required Under MiFID II	cPacket Timestamp Accuracy		
Activity using high frequency algorithmic trading technique	100 μs	1 μs or better			
Activity on voice trading systems	1 sec	1 sec or better	+ 1 nc		
Activity on request for quote systems where the response requires human intervention and algorithmic trading	1 sec	1 sec or better	using PPS		
Activity of concluding negotiated transactions	1 sec	1 sec or better			
Any other trading activity	1 msec	1 msec or better			

Not STAC Benchmarks

Thank You

Please email questions to: josh@cpacket.com

Confidential and Proprietary. Not to be distributed without express written consent of cPacket Networks, Inc.

© 2017 cPacket Networks

Ideas

- 100G is deploying
 - packet brokering and NICs are not the answer
 - Gilder-Moore
- 100G math for spikes
 - Acceleration-deceleration

Financial Trading Networks need 100% of time capacity

NEED GRAPHS WITH 100GBPS

Confidential and Proprietary. Not to be distributed without express written consent of cPacket Networks, Inc.