

# SONiC-**Event Driven Data Plane Telemetry**

#### Aviad Raveh, VP Business Development Technology Mellanox







# PLATINUM



## **Telemetry Data Sources**





#### **Statistical Sampling**







## **Event driven telemetry indications**

#### Statistical Telemetry









#### **Events**





# **Telemetry Mechanisms Functional Comparison**

Network Scenario	Statistical Sampling	Event Driven	Flow Driven
Physical Layer Fault			
Application Distribution			
Network utilization			
Flow path			
Forwarding Black Holes			
Incasts			
Micro bursts			
E2E application latency			
DOS Attack			
Policy Engine Hits			
Dropped Packets Data			







## Network debug is like application debug

#### **Network Debug**









#### **Application Debug**

top - Tasks Cpu(s)	12:16:01 : 141 tot : 0.8%u	up 1 al, s, 0	day 1 r	, 23:3 unning y, 0	31, 3 3, 139 .0%ni,	2 use 9 slee , 98.5	cs, ep: 54:	, lo; ing, id, (	ad ave 0 st 0.0%wa	opped, , 0.0%hi	0, 0.00, 0.00 1 zombie , 0.0%si, 0.0%st
Nem:	1021108	k tot	al,	983	128k 1	used,		3798	SOk ir	ee, 134	600k buffers
FID to	2046968 cenice:	2545	4		UR	usea,		204691	SSK II	ee, 599	b/2K cached
PID	USER	PR	NI	VIRT	RES	SHR	3	SCPU	\$ MEM	TIME+	COMMAND
25454	root	20	0	397m	107m	13m	\$	2.3	10.8	26:13.86	skype
26211	root	20	0	2680	1140	884	R	0.7	0.1	0:00.34	top
31031	root	20	0	5544	736	528	\$	0.3	0.1	0:13.15	udisks-daemon
1	root	20	0	2872	1400	1200	5	0.0	0.1	0:00.89	init
2	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kthreadd

top

logfile

Nov	17	09-17-02	WS	recheck for undates: Bunning /usr/local/bin/check for undates
Nov	17	09:17:02	MIC	check for undates. Starting
Alexa		00.17.02	WY-S	check for updates. Updating heldligt
NOV		09:17:02	WS	check_for_updates: Opdating notatist
Nov		09:17:11	WS	<pre>check_for_updates: Running 'apt-get update'</pre>
Nov		09:17:11	WS	check for updates: Error exit called
Nov	17	09:24:18	WS	<pre>nagios3: Auto-save of retention data completed successfully.</pre>
Nov	17	10:05:58	WS	kernel: [ 6185.581784] br0: port 1(eth0) entering forwarding state
Nov	17	10:05:58	WS	kernel: [ 6185.637674] device eth0 left promiscuous mode
Nov		10:05:58	WS	kernel: [ 6185.637681] br0: port 1(eth0) entering disabled state
Nov		10:05:58	WS	kernel: [ 6185.639415] ADDRCONF(NETDEV_UP): eth0: link is not ready
Nov	17	10:06:02	WS	kernel: [ 6188.955834] device eth0 entered promiscuous mode
Nov	17	10:06:02	WS	kernel: [ 6189.097303] r8169 0000:03:00.0: eth0: link down
Nov		10:06:02	WS	kernel: [ 6189.097309] r8169 0000:03:00.0: eth0: link down

gdb

		· · ·		****	000 000	reakpoint gdb) run tarting p reakpoint	: 1 at God program: / : 1, main for	3848426: /home/gar () at he r(i=0;i<1	file he y/hello llo10.c 0;i++)	llo10.c, 10 :6	line (
	••••							******			
		-	-	- 18 C	- 25			-		-	
			10 m m m m m		19						
		and the second	and the first of the		1.12			المالية المالية			
63 - F			- C - C - C - C - C - C - C - C - C - C		- 62		* ***		1. S. S. S. S.		
- C - L			* *	•	- 33			+	- 12 E		
100 C	***	******	*****	- C	æ 🗌	****					











Prediction



#### Remediation











# **SONiC and Open Fabric Monitor**

- Lets make OFM Our Fabric Monitor and not One Feature Missing
- All nodes speak the same telemetry language
- SONIC community defines recommended integration with set of open tools (e.g. InfluxDB, Grafana)
- Lesson learned repository sharing debug scenarios with SONiC as the network NOS, from problem to solution









# **Call For Action:**

- Add event driven infrastructure to SONiC
  - Unified agent on top of SAI that retrieves dataplane events and streams it to a collector
- Define the streaming data format
  - Not the transport
  - Data streamed from SONiC should be consumed by a unified database
  - Standard format (e.g. JSON)









# Open. Together.



OCP Global Summit | March 14–15, 2019



