

Scalable and Flexible Network Functions using SDN White-Boxes



Vijay Sivaraman, Professor, UNSW

Never Stand Still

**ANZ
SDN**

**AUSTRALIA & NEW ZEALAND
SOFTWARE DEFINED NETWORKING**

Who am I?

- PhD from UCLA in 2000
 - Packet scheduling analysis (Bell Labs)
 - Thesis had more Greek than English!
- Silicon-valley start-up 2000-2003
 - L1-L4 optical switch-router; \$64.5m burnt
 - 100k+ lines of code; NPU-code, L2/L3 protocols, ...
- Gap between academia and industry
 - Depth vs breadth; optimal vs simple; innovative vs practical
- Software Defined Networking (SDN): ray of hope?
 - Software “brain” separated from hardware “brawn”
 - ANZ-SDN Alliance: community to foster SDN ecosystem

$$\alpha_j(t, w) = \frac{d}{dw} P\{A_{ji}(t) \leq w \text{ and } (j, i) \text{ on at } -t\}$$

$$\beta_j(t, w) = \frac{d}{dw} P\{A_{ji}(t) \leq w \text{ and } (j, i) \text{ off at } -t\}$$

Thus, $A(t) = \sum_{j:d_j < t} A_{ji}(t)$ and we can express $\phi_i(w, \lambda)$ as:

$$\phi_i(w, \lambda) = \prod_{j=1}^m \binom{k_j}{l_j} \bigoplus_{j=1}^m \alpha_j(t, \cdot)^{\oplus(l_j)} \oplus \beta_j(t, \cdot)^{\oplus(k_j - l_j)}(\cdot, w) \quad (29)$$

where

$$m = \max\{j : d_j < t\} \text{ and } l_1, l_2, \dots, l_m : \sum_{j=1}^m l_j p_j = \lambda \quad (30)$$



Let's talk Network Functions (NFs)


- Special-purpose operations on packets
 - Service Providers: DPI, load-balancer, ...
 - Enterprises: Firewalls, NAT, ...
- Physical appliances (middle-boxes) are expensive
- Virtual appliances (VNFs on x86) difficult to scale

How can SDN Help?



Hardware: dumb+fast on “flows”

- Abstraction: match + action + counter
- Match: any header field(s)
- Action: out-port, mirror, drop, edit
- Multi-vendor “white-box”
- Scalable: 400 Gbps – 6.4 Tbps
 - 1m+ flow table entries (FIB)
- Cost-effective: \$20-40K



Software: smart+slow on “packets”

- Deep inspection / stateful operations
- Intelligence (e.g. optimization, ML/AI)
- Commodity servers cheap
 - Elastic scale-out, DPDK support
 - Flexible and agile
 - Software dev environments / tools

So how to put the two together?

- NF: Operate on flows when you can and packets when you must
- 75-80% of network traffic is carried in top 1-5% (“long”) flows
 - At UNSW: ~600K concurrent flows at peak hour; ~6k flows transfer > 4MB [1]
 - Corroborates with prior mice/elephant studies in enterprise and carrier networks [2]
- Here’s a simple idea:
 - Operate on mice packets in software (VNF)
 - Operate on elephant flows in hardware (SDN)
- Potential benefits of “SDN Accelerated NF” (SANF):
 - Software VNF load reduces by 4x → lower cost
 - Hardware flow table (FIB) is small → better scalability

[1] S. Madanapalli, M. Lyu, H. Kumar, H. Habibi Gharakheili, V. Sivaraman, “Real-time Detection, Isolation and Monitoring of Elephant Flows using Commodity SDN System”, IEEE/IFIP NOMS workshop on Experience Sessions, Taipei, Taiwan, Apr 2018.

[2] T. Mori et al, “Identifying elephant flows through periodically sampled packets”, ACM IMC, Taormina, Sicily, Italy, 2004.


Example Uses

- Example 1: Load-Balancer
 - Handle first packet of flow in VNF, then insert flow table rule to send to selected output
- Example 2: NAT
 - Determine port mappings in VNF based on first packet of flow; subsequent packets in h/w
- Example 3: Traffic classifier
 - Keep state on flow volume in VNF; Insert flow table rule in h/w on reaching threshold
 - Extract flow telemetry and classify based on behavioral profile
- Example 4: Stateful firewall
 - Determine state on outbound flow based on first packet; insert flow rule for return traffic
- These examples and their complexity/scalability modeled in ^[3]

[3] V. Sivaraman, H. Habibi Gharakheili, M. Lyu, H. Kumar, C. Russell, “Scaling Network Functions via SDN Acceleration”, under review with IEEE Journal on Selected Areas in Communications (JSAC), 2018.

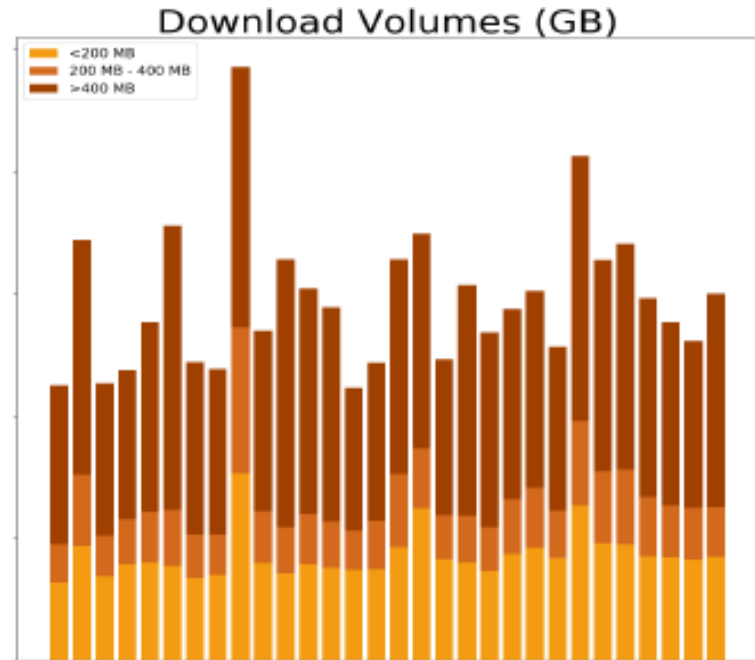
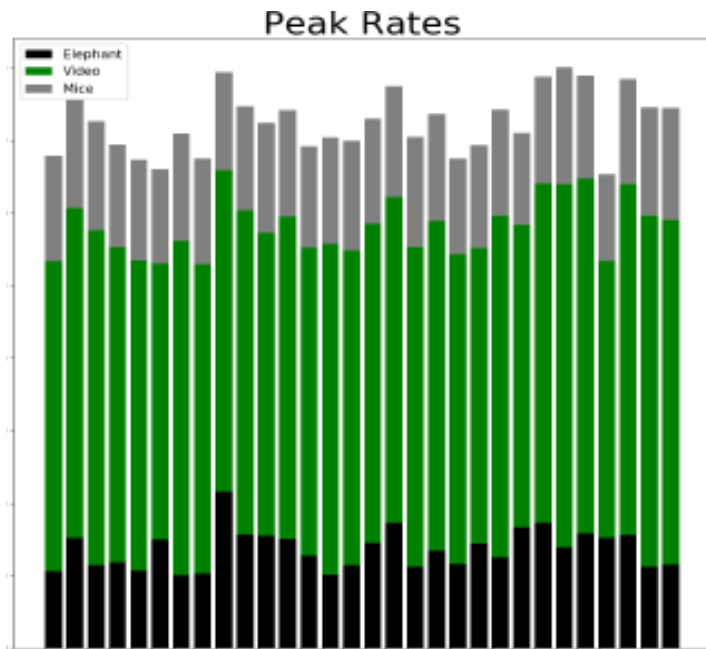
NF Use-Case 1: Traffic Classification

- Motivation: Understand bandwidth usage pattern
 - Break-down of browsing, video, and download sessions over a day
 - Composition of video traffic by provider (Youtube, Netflix, ...)



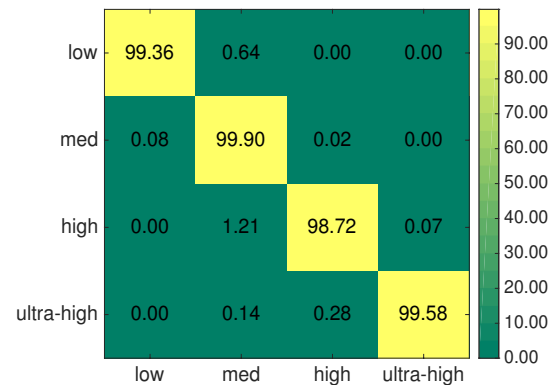
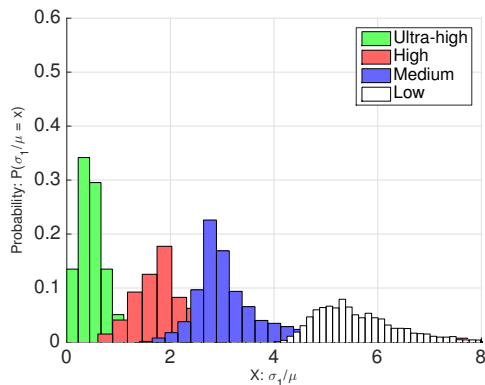
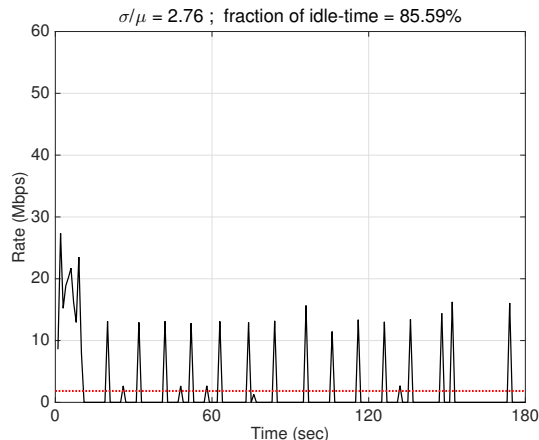
Traffic Classification (contd.)

- Motivation: Manage demand variability (peak provisioning)



DPI vs Flow Classification

- Approach: Classify long flows based on behavioral profile [4],[5]
 - No inspection of packets (except DNS)
 - VNF tracks “short” flows, pushes “long” flows into hardware
 - Extracts statistical profile of long flows and does machine classification

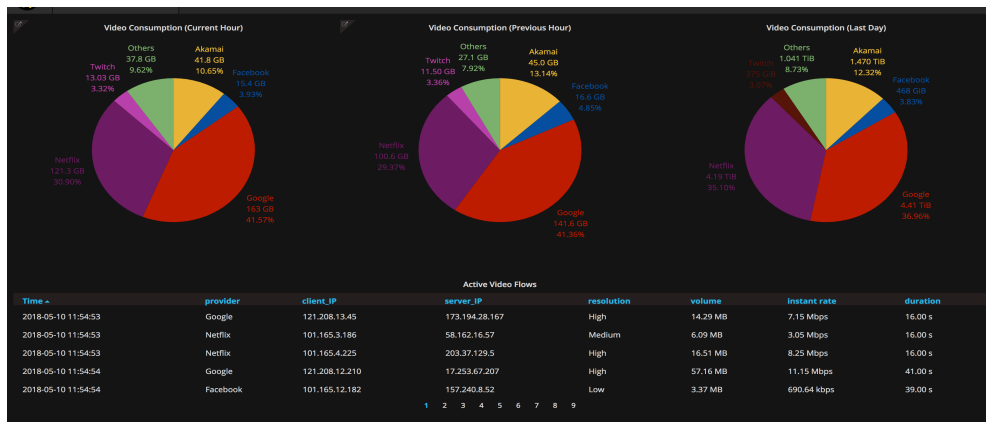


[4] T. Bujlow et al, “Independent comparison of popular DPI tools for traffic classification”, *Computer Networks*, 76(c):75 – 89, Jan 2015.

[5] A. Dainotti et al., “Issues and future directions in traffic classification”, *IEEE Network*, 26(1): 35–40, Jan 2012.

TeleScope: Flow Classification at Line Speed

- Fully implemented and operational today [6]
 - Uses NoviFlow SDN switch, Ryu Controller, NFF-GO DPDK, Weka/SciKit ML, ...
 - Real-time UI, logging of every long flow, analysis/insights each midnight
 - Operating on 10 Gbps mirror traffic at UNSW for past 6 months



[6] H. Habibi Gharakheili, M. Lyu, Y. Wang, H. Kumar, V. Sivaraman, “iTeleScope: Intelligent Video Telemetry and Classification in Real-time using Software Defined Networking”, under review with IEEE/ACM Transactions on Networking, 2018.

TeleScope: Commercial offering



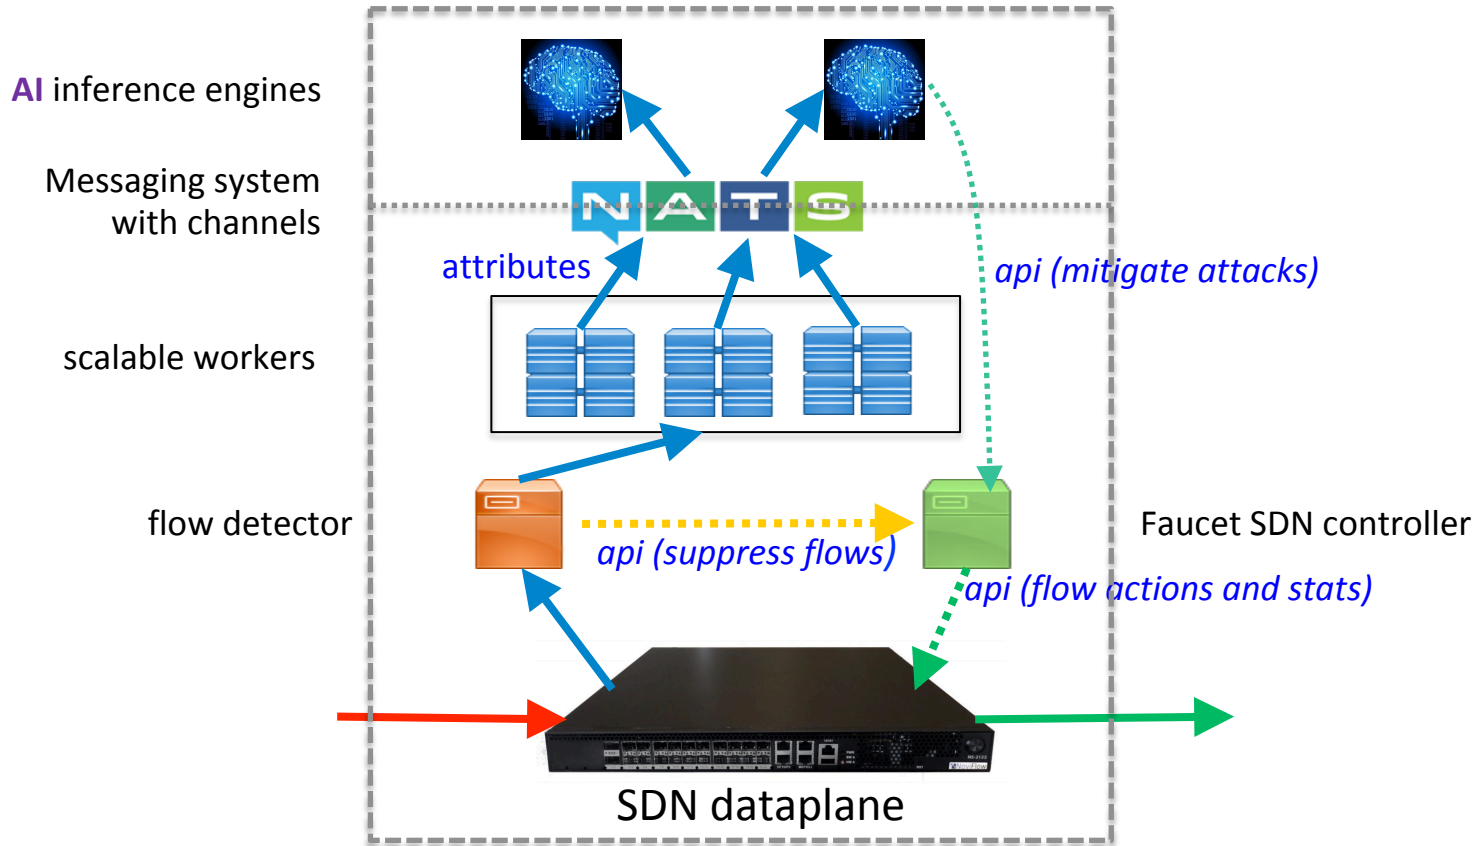
- IP licensed to Canopus Networks
 - In commercial trials with multiple Tier-1 ISPs
 - Showing insights into video consumption (real-time / hourly / daily)
 - Identifies 300+ video content providers, including rate, resolution, duration
 - Showing insights into large downloads (real-time / hourly daily)
 - Storage, CDNs, games (Fortnite), ...
 - Developing signatures for interactive video, streaming music, gaming, VR, ...
- Partnering with white-box BNG vendors for in-line b/w management
 - Framework for customer experience optimization based on theory of “marginal utility”
 - Open, flexible, and verifiable framework that addresses net neutrality concerns [7]

[7] V. Sivaraman, H. Habibi Gharakheili, S. Madanapalli, H. Kumar, C. Russell, “Competing on User Experience in a Post-Neutral World”, under preparation for submission to ACM SOSR, 2019.

NF Use-Case 2: Cyber-Security

- Appliance based security solutions are:
 - Expensive → 10G/100G service may be cheaper than the security appliance!
 - Bundled → can't combine “best-of-breed” or get “second opinion”
 - Performance-limited → features have to be turned off for line-rate performance
- Our goal: build platform (**Nozzle**) to augment current solutions:
 - Leverages white-box SDN hardware and commodity server-class compute
 - is **scalable** (modular expansion of capacity as needed)
 - is **flexible** (best-of-breed components)
 - is **low-cost** (commodity components)
 - set of “**workers**” extract attributes ← **SDN**
 - set of “**inference engines**” make deductions ← **AI**

What Nozzle looks like



So what are “attributes”?

- **DNS attributes**

- query & response volume, solicited vs. unsolicited
- #RRs, TTL, sub-domain length, udp payload size, ...

- **Host attributes**

- #TCP-connections, volume, connection-rate, ...

- **Flow attributes**

- long-flow data volume, 5-tuple end-points, SSL/TLS cipher-suites, ...

- **Intra-flow profile**

- inter-packet gaps, packet size distribution, ...

- **Other signaling attributes**

- NTP, DHCP, SNMP, SSDP, SIP, etc.


A deeper look at the DNS Attribute

- DNS:
 - comprises < 0.1% of traffic volume; often allowed through firewalls
 - wealth of information on external services accessed & internal servers
 - used for DoS, data exfiltration, reflection/amplification attacks
- DNS worker in **Nozzle**
 - SDN flow-rule + Elasticsearch; Intel DPDK with RUST/Golang
 - publishes aggregated DNS records to NATS every second



Analysis of one day's DNS Traffic at UNSW

- Less than 10 Mbps (port 53)
- Extracted 4 attributes per-host:
 - Outbound queries / inbound queries
 - # external servers, # external clients, fraction of active time
- Built simple clustering-based ML classifier:
 - UNSW: 40 name servers; 14 resolvers; 368 public-servers
- Result: automated, dynamic, daily inventory of key assets
 - DNS resolvers, DNS name servers, external-facing web-servers



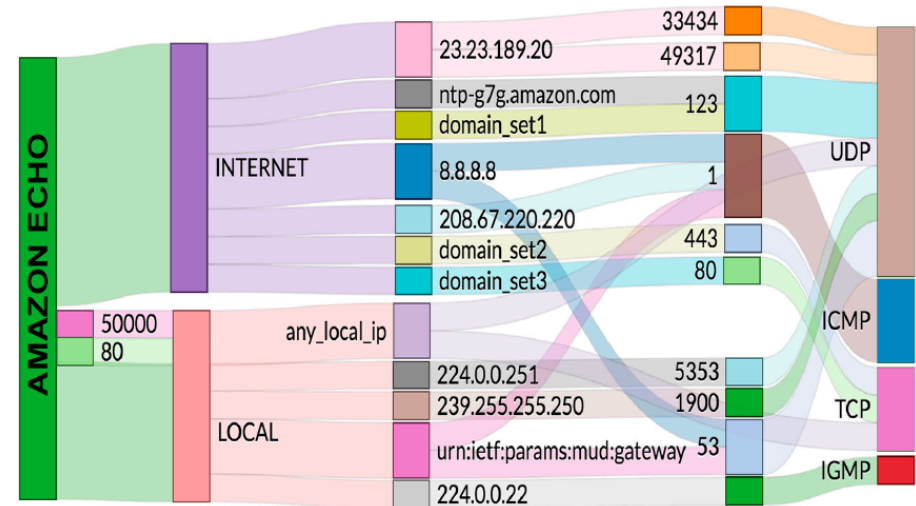
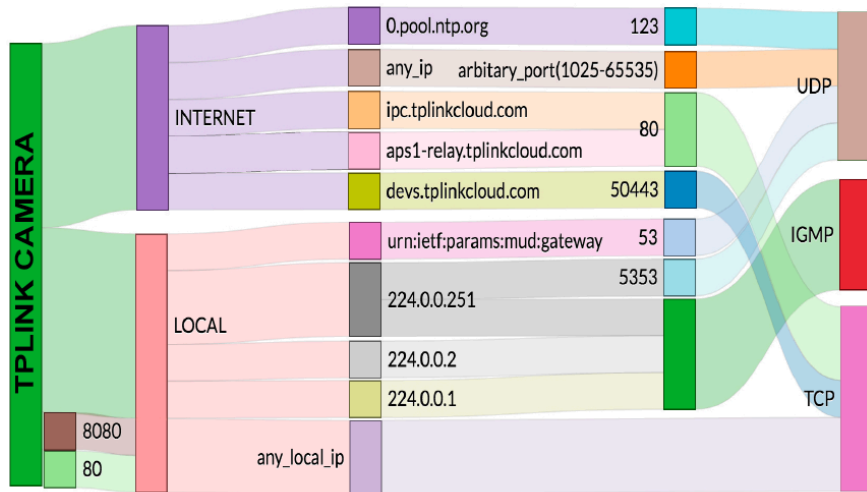
[8] M. Lyu, H. Habibi Gharakheili, C. Russell and V. Sivaraman, "Mapping an Enterprise Network By Analysing DNS Traffic", under review at ACM CoNEXT, Dec 2018.

DNS Attacks and Threats

- Evidence of attack vectors (what and how):
 - DNS scan of network
 - Recursive resolvers used for reflection attacks
 - Public-server DoS-sed with DNS responses
- DNS exfiltration possible and happening
- ML built using attributes: [9]
 - length, entropy, numerical, capital, dots, longest label, reputation
- Our machine detects exfiltration (e.g. cnr.io) in real-time

[9] J. Ahmed, H. Habibi, C. Russell, V. Sivaraman, "Real-Time Detection of DNS Exfiltration and Tunnelling from Enterprise Networks", under review at IFIP IM, 2018.

Another use of “attributes” – IoT



- Each IoT device has its “signature” of network behaviour
- Enforce network behaviour policy (MUD) using SDN

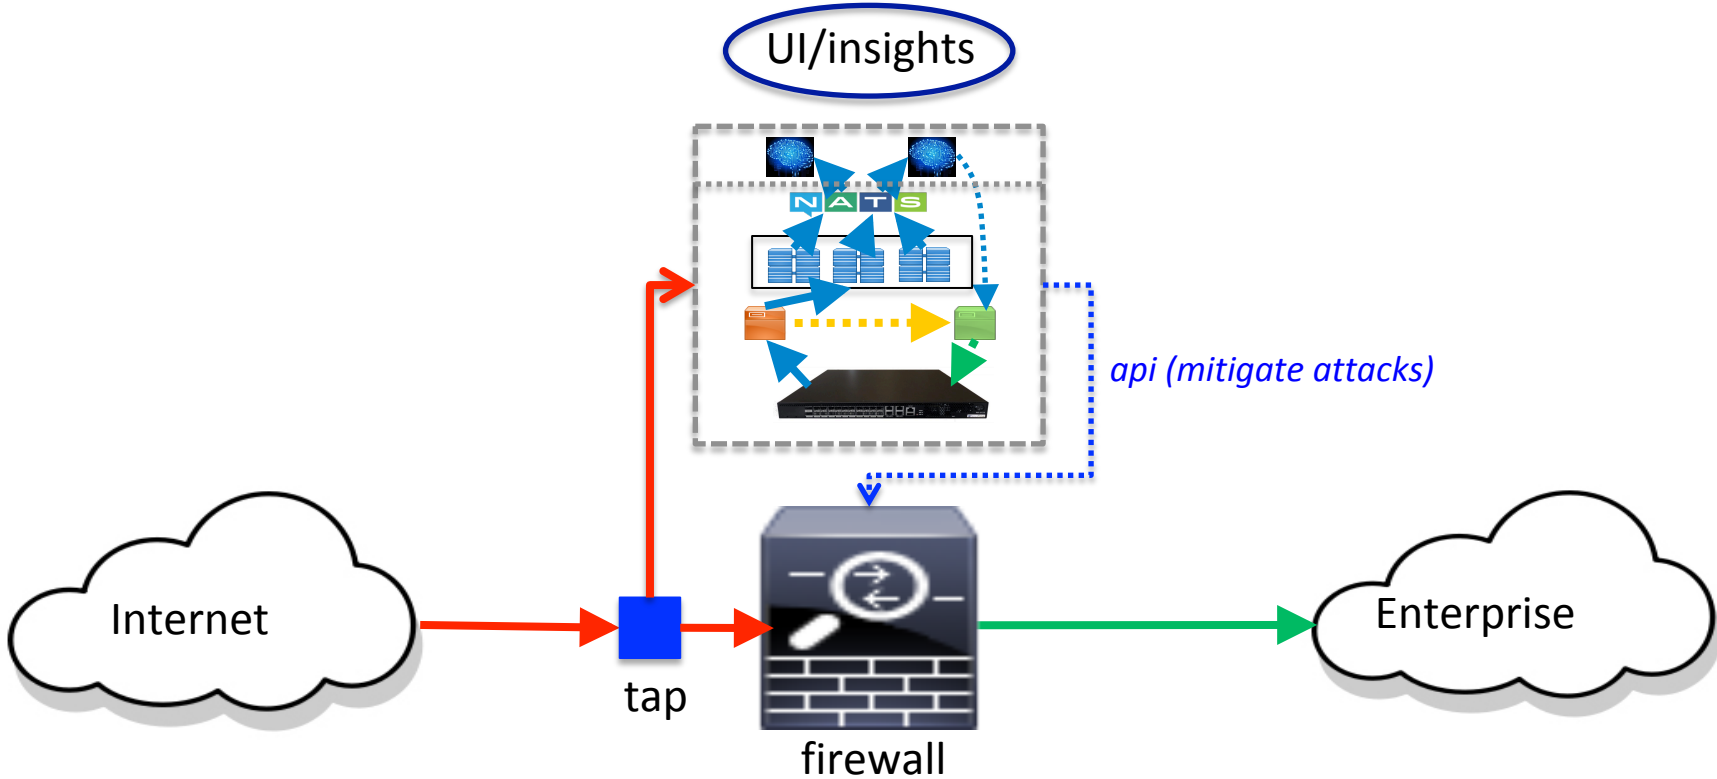
[10] A. Hamza, H. Habibi Gharakheili and V. Sivaraman, “Detecting Volumetric Attacks on IoT Devices via SDN-Based monitoring of MUD Activity”, submitted to ACM CoNEXT, Dec 2018.

Automatic detection of IoT assets

IoT Device Category	Hubs		Cameras					Switches & Triggers				Air quality sensors		Healthcare devices			Light Bulbs	Electronics			Non-IoT devices		
	Smart Things	Amazon Echo	Netatmo Welcome	TP-Link Day Night Cloud camera	Samsung SmartCam	Dropcam	Insteon Camera	Withings Smart Baby Monitor	Belkin Wemo switch	TP-Link Smart plug	iHome	Belkin wemo motion sensor	NEST Protect smoke alarm	Netatmo weather station	Withings Smart scale	Blipcare Blood Pressure meter	Withings Aura smart sleep sensor	LiFX Smart Bulb	Tribby Speaker	PIX-STAR Photo-frame	HP Printer	Laptop	Smart Phone
Sleep time	1	1	1	2	1	1	2	1	1	3	2	1	1	1	1	5	1	1	1	2	1	1	1
Active volume	1	1	1	2	4	1	3	1	5	2	2	2	5	2	3	4	2	1	1	2	1	2	1
Avg. Pckt size	1	2	4	1	4	2	2	1	4	2	1	3	3	3	2	2	1	2	3	1	5	4	4
Mean. rate	1	1	2	1	2	1	1	1	3	1	1	3	2	1	1	1	1	1	1	1	3	2	2
Peak / Mean rate	1	1	2	1	1	2	1	1	1	1	1	1	1	1	1	2	1	2	1	1	2	2	2
Active time	1	1	1	1	1	1	1	1	1	1	2	1	1	1	3	1	1	1	1	1	1	1	1
No. of servers	1	2	1	2	2	1	2	1	1	1	1	1	1	2	1	1	1	1	1	1	3	3	3
No. of protocols	1	4	4	5	4	1	3	1	1	1	1	1	1	1	1	1	1	4	2	1	2	2	2
Unique DNS req.	1	1	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4	4	4
DNS interval	2	1	1	2	1	1	1	2	3	1	4	3	1	1	4	5	1	1	2	2	1	1	1
NTP interval	2	1	1	4	4	1	1	1	1	3	1	1	1	1	1	1	1	1	1	1	1	1	1

[11] A. Sivanathan, D. Sherratt, H. Habibi Gharakheili, A. Radford, C. Wijenayake, A. Vishwanath and V. Sivaraman, "Characterizing and Classifying IoT Traffic in Smart Cities and Campuses", IEEE Infocom SmartCity17 Workshop on Smart Cities and Urban Computing, Atlanta, GA, USA, May 2017.

Potential model of trial: Augment current appliance



Conclusions

- SDN-Accelerated-NF has a strong value proposition
 - Commodity hardware → lower cost
 - Multi-vendor → choice of port speeds and densities
 - Modular and elastic → scalable in performance
 - Unbundled → more innovative and agile “state-of-art” software (ML/AI)
- Looking for operators with curiosity and appetite
 - Co-develop and test innovative solutions (journey/learnings are important)
 - Gain new insights that would not emerge otherwise
 - Develop skills in emerging areas (software/ML) in work-force
 - Support Australian innovation and education eco-system

Sponsors

- SDN Research and Education at UNSW has been supported by:
 - Australian Research Council (ARC)
 - Google
 - Optus
 - Telstra
 - AARNet
 - Cisco
 - HP
 - NoviFlow
 - Open Networking Foundation (ONF)