DNS-AS

Done with SDN and Tired of Dealing with Snowflake Network Complexity? Change the Game with a Simple TXT String

Wolfgang Riedel

Principal Engineer, CCIE #13804, VCP #42559

wolfgang@f1-consult.com

Who is Wolfgang Riedel ???

- **Personal:**
  - Location: Erlangen, Germany (between Munich – Frankfurt)
  - Interests: Alpine Snowboarding, High-End Audio, AS51871, Data Center, ZFS/ZOL, Real World LAB, Cybersecurity, High-performance sports cars, Geothermal DC cooling, …
- **Background:**
  - 1985, Started my first company
  - Self-employed as an in-depended consultant in the Networking and IT space
  - 2001, Joined CISCO
    - ✓ SE – RS Germany (2001 – 2006) -> Campus with a DC attached
    - ✓ CSE – DC EMEA (2006 - 2008) -> DC with Campus attached
    - ✓ CE – CTO Office (2008 – 2011)
    - ✓ PE – ARND (2011 – 2013)
    - ✓ PE – CTO Team ENG (2013 - 2014)
    - ✓ PE – Architecture Team ENG (2014 – 2017)
  - Worked with more then 250 customers within several projects over the last +20 years
  - CCIE #13804 (RS), VCP #42559 (3/4/5/6) and pile of CPOC's
  - Individual Contributor: Cat4k, Cat6k, N7k, ASR1k, FC, FCoE, DCB, UCS, N5k, N2k, N1k, PoE FEX, vPC, OTV, LISP (Pioneer Award), SP-DC, OF, SDN, NfV, USP, APIC-EM, AVC
  - 2017, Self-employed as an in-depended consultant in the Networking and IT space, again ;-)
- **Stuff I am currently working on:**
  - DNS-AS (two patents pending)
  - Consulting for some very special customers ;-)

Agenda

1. Introduction – What is DNS-AS
2. What is Network Metadata
3. Network Metadata within DNS RR's
4. How to manage "foreign" domains
5. DNS-AS Operations
6. Actually, what can we do with it?
7. DNS-AS Demo
8. Program Plans & Milestones
9. A Few Conclusions and Q&A, if we have time

Core Message:
Network Metadata

Warning:
A good portion of this session is about DNS and DNS functionalities we use.
This is not about DNS-AS it's just supposed to be a re-fresher for those of us which forgot about it ;-)

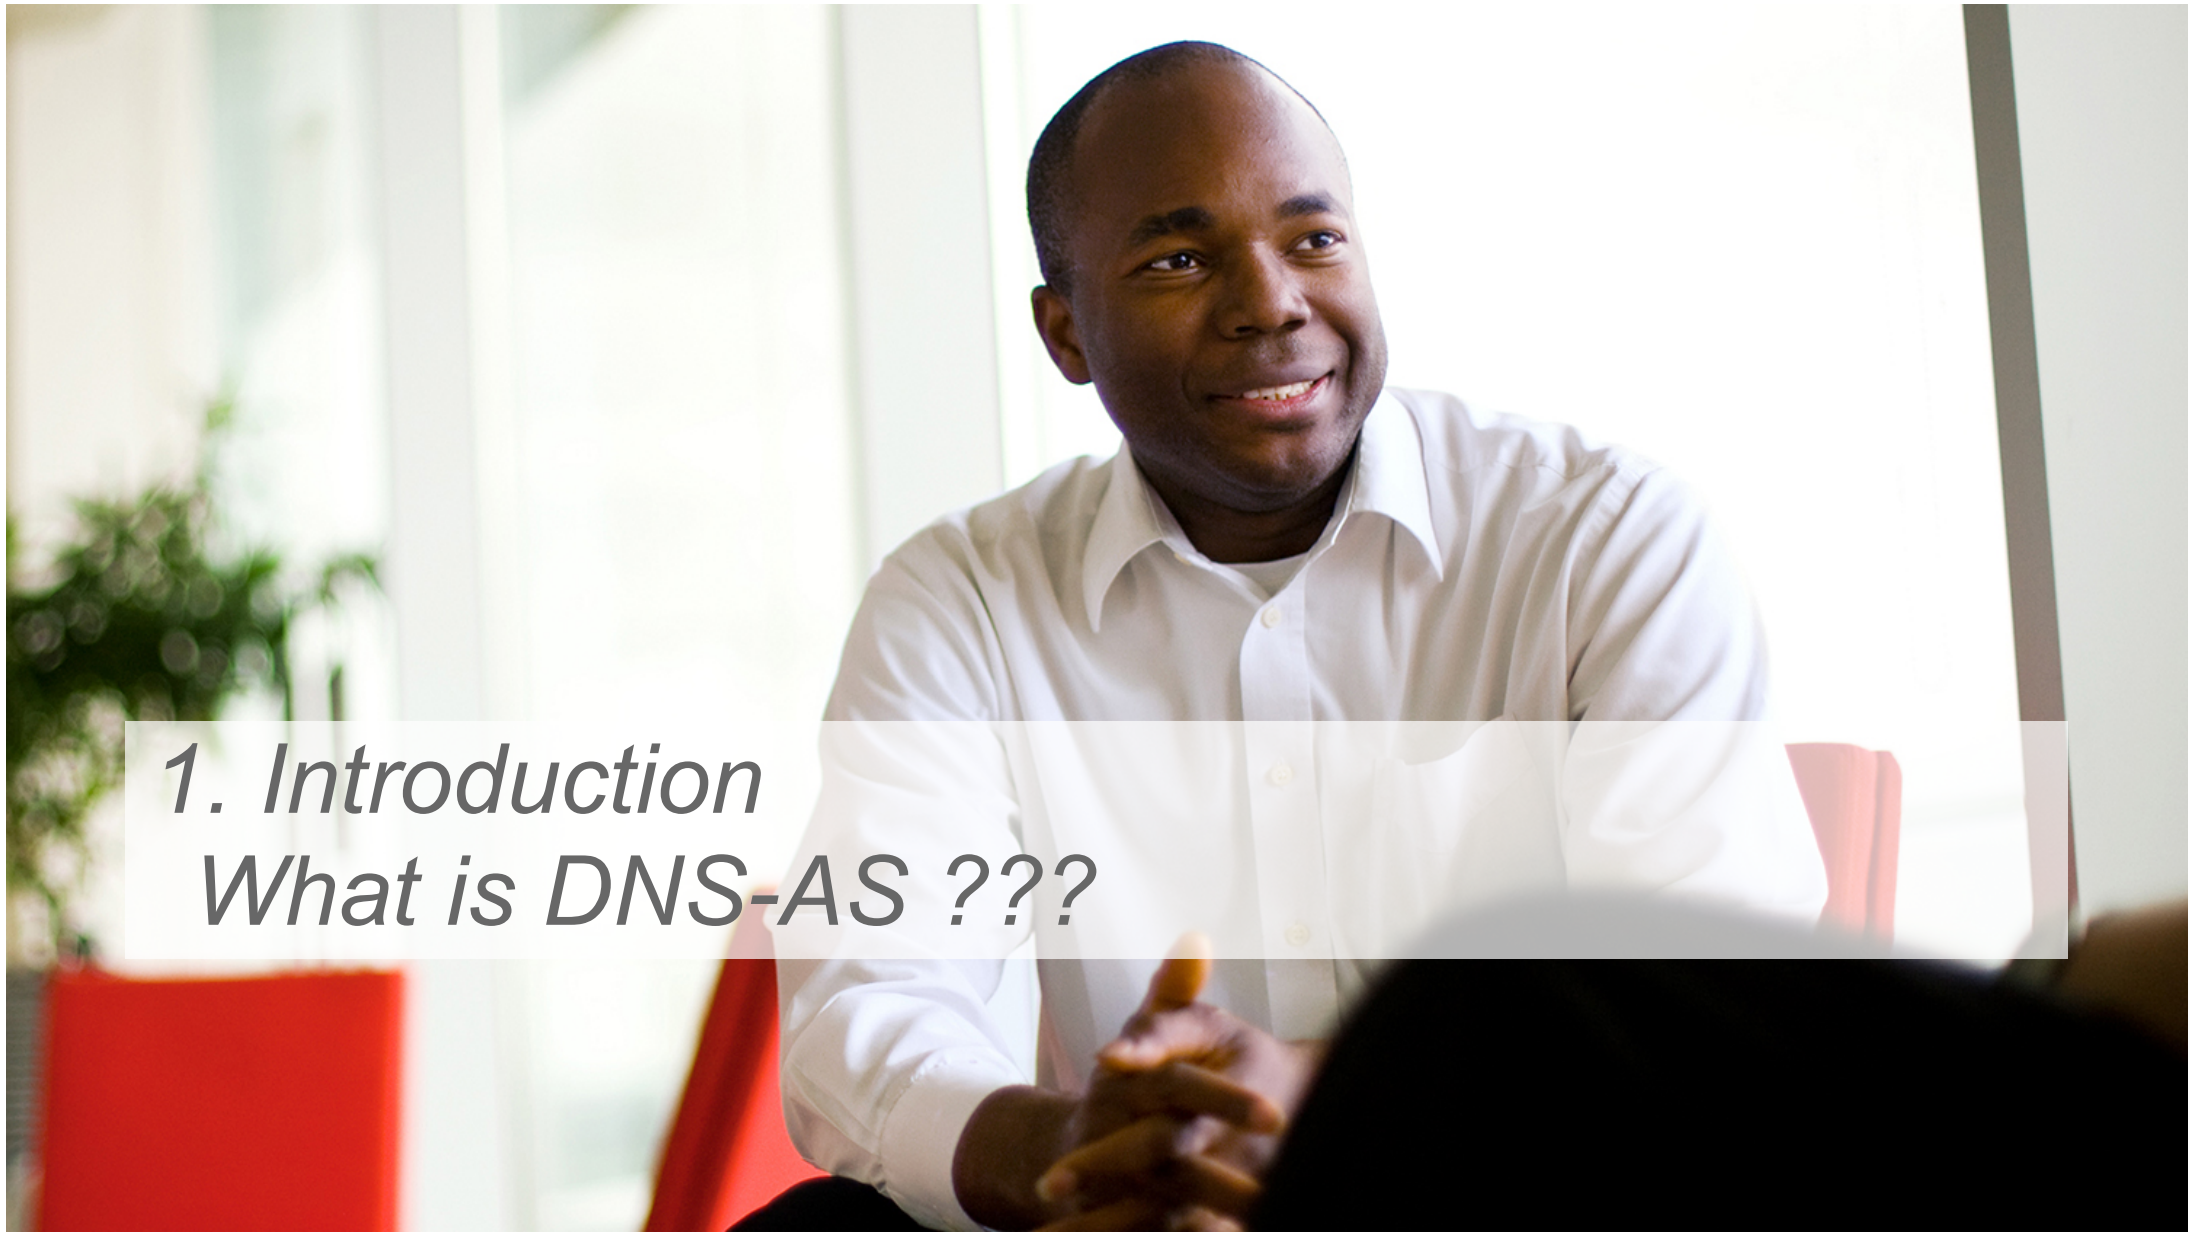

1. Introduction
What is DNS-AS ???

End2End Common Policy Model

The big SDN question in October 2013

Policy Intent (Common Namespace for Business Intent)

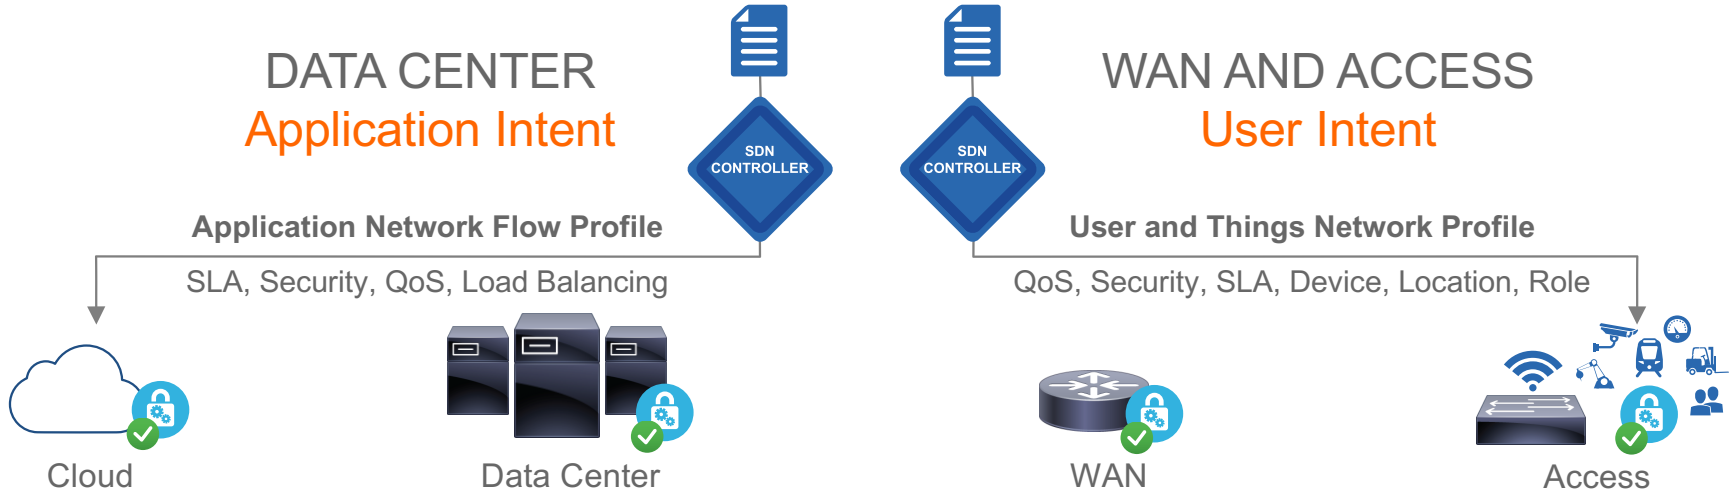

BROWNFIELD AND GREENFIELD | END TO END | POLICY FRAMEWORK: FOCUS ON APPLICATION AND USER ENABLEMENT

DNS-AS – The idea in 17.10.2013

Mike Herbert, Mark Montanez and Wolfgang Riedel @ a Sushi place in SJC

Sorry, no napkin this time…

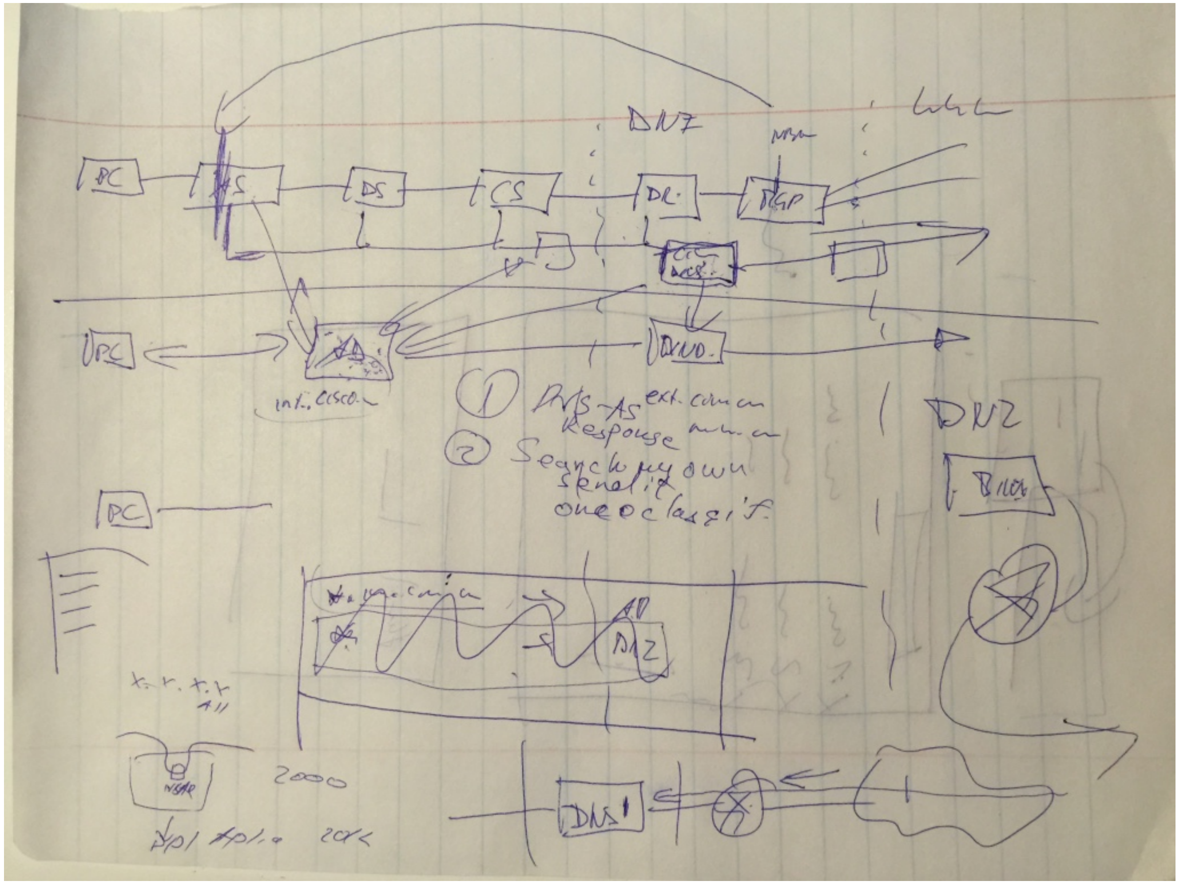

9

DNS-AS - Tenets

Application Visibility

How
can you keep unambiguous visibility if the majority of traffic is encrypted?

Metadata Driven

How
can you holistically program the network so it behaves like a self driving car.

Centralized Control

How
to use DNS as a cross domain application intent policy controller?

DNS-AS - Problem Statement

❑ **Application Visibility**

Today many applications operate in clear text and therefore it is possible to identify these by the use of Deep Packet Inspection (DPI) methods. Tomorrow applications communicate in a confidential way by the use of end2end encryption which renders DPI methods ineffective as a means of application identification and Application Visibility and Control.

❑ **Metadata Driven**

Metadata is information about applications that describes them. Instead of guessing device by device we holistically program the network via metadata, no matter if the traffic is encrypted or not. Suddenly your network behaves like a self driving car.

❑ **Centralized Control**

The Promise of SDN had been “Decoupling Policy from Configuration” in means of Policy Intent Networking. While the industry is busy trying to agree on Cross Domain Policy (NIC, GBP, NEMO) we simply utilize the most scalable and proven controller out there which is already available across all admin domain boundaries.
The DNS infrastructure!

❑ **Control without admin access**

Furthermore customers may no longer own a network at all as everything is up in the cloud and they may just have a small network inside the data-center which needs to take control over network devices spread across the whole internet which may not be under direct administrative control of them. With the proliferation of digitization in the context of IOT and IOE with thousands to millions of devices and sensors it becomes apparent that present controller approaches cannot scale to such exceptional numbers.

DNS-AS - The Burj Khalifa Elevator Pitch

DNS-AS leverages DNS as an Authoritative Source
to publish metadata as a key for common policy across networks,
without the need for a dedicated (SDN) controller.

DNS-AS is a control and data plane separation solution where we leverage the **D**omain **N**ame **S**ystem as an **A**uthoritative **S**ource to publish metadata at large scale as a key for common policy across enterprise and worldwide distributed networks without the need for a dedicated (SDN) controller.

While the application of policies to network devices, applications and services stays local to the device, DNS-AS is able to simplify network operations at large scale without the need of steady reconfiguration of these. Not all network devices have to be capable of supporting DNS-AS which enables phased deployment.

DNS-AS-Client addresses how we enable network elements or applications to retrieve metadata from the DNS Database. We use this metadata to express policy intent and associate this metadata locally and leverage it for local policy enforcement and decision making.

DNS-AS-Proxy will be able to generate metadata in the case an authoritative DNS Server is unable to provide metadata or may not be considered as a trusted as a source.

1.1 Industry Trends
10 minutes on SDN relevance

Industry trends in Networking

Cloud (2008)

OpenFlow Networking (Stanford clean slate) (2011)

Software Defined Networking (2012)

OPEN DAYLIGHT www.opendaylight.org Open Daylight Project (2013)

DevOps, The API Driven Datacenter (2013)

Network Function Virtualization (2013)

Managing Networks through abstractions (2014)

Metadata Driven Networking (2016)

Micro Services (2018)

Managing the network through abstractions

There are two approaches to Control Systems

IMPERATIVE CONTROL

Baggage handlers follow sequences
of simple, basic instructions

DECLARATIVE CONTROL

Air traffic control tells where to
take off from, but not *how* to fly the plane

Managing the network through abstractions

There are two approaches to Control Systems

IMPERATIVE CONTROL

DECLARATIVE CONTROL

It’s 2017 and network admins still enjoy being “masters of complexity”

SDN Controllers – Types

Start with the End in Mind - the RYF-complex (Fragile/Robust)

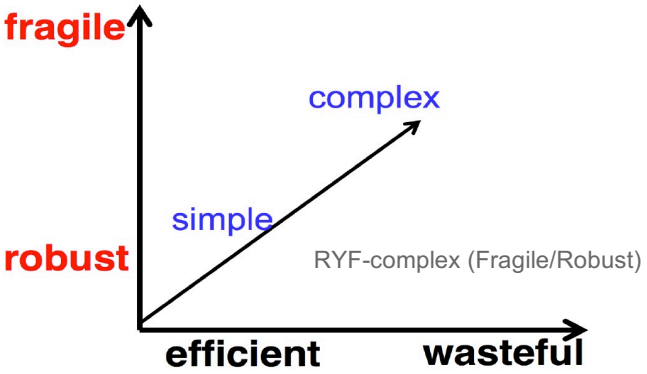

RYF-complex (Fragile/Robust)

Five dimensions of robustness in complex systems
(1) Reliability
(2) Efficiency
(3) Scalability
(4) Modularity
(5) Evolvability

See J. Doyle, et. al.,
"Robustness and the Internet: Theoretical Foundations"

Alderson and Doyle identify four kinds of constraints on system robustness:
(1) Component-level
(2) System-level
(3) Protocols
(4) Emergent constraints

Complex systems science as conflicting constraints
John C. Doyle, HOT and SF networks

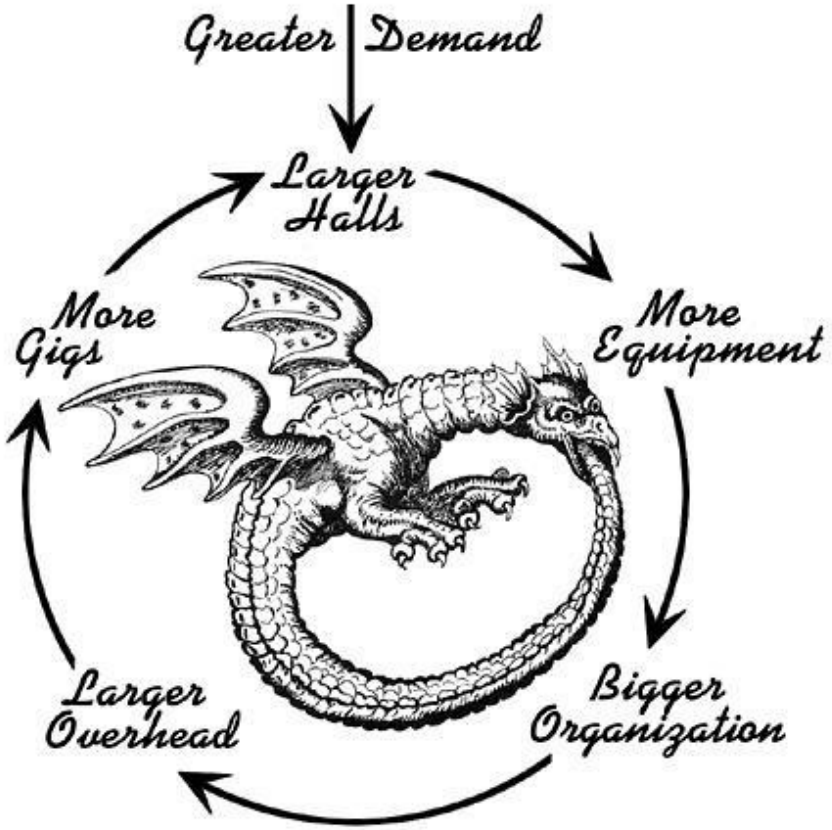

(Giga Exponentia

DNS server as a SDN controller?

It’s a pretty proven and awesome system, right?

Reliability

Using DNS - the most proven, used and scalable system of the Internet, to Distribute Metadata

Efficiency

DNS well proven for it’s efficiency – Light weight & Distributed with Tree Architecture

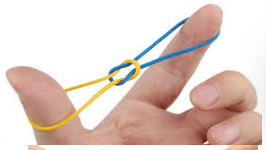

Scalability

DNS is a fully distributed system- scales well for the whole Internet!

Modularity

Decoupled DNS Network Infra and Agent running on Device (No endpoint requirements)

Evolvability

Has the capacity of Adaptive Evolution – Metadata not just limited to Network Devices

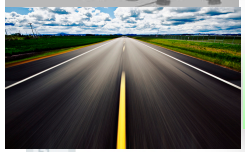

Performance

Hardware Acceleration possible – Potential for applications beyond QoS (security, etc …)

CONSULT GmbH & Co. KG

How About DNS Granularity?

Is DNS granular enough? - IP Address Explosion

Networks continue to grow in size, importance, and complexity, organizations need to implement network services that are secure, scalable and fault tolerant

- One IP per service is the new norm
- IP Address Explosion:
  - VM Sprawl
  - M2M
  - My Own Private Internet
- IPv6 without DNS is impossible to manage
- IPv6 also replaces DSCP in some ISP networks
- DHCP makes the task of network configuration a breeze
- DNS is still key

DNS Security? – Pretty Bad Privacy

Threats: Monitoring and Surveillance (Haya Shulman, irtfopen @ IETF93)

DNS packets:

- Clear text is per se insecure (monitored, collected, logged)
- DNS data is public
- Research
- Operational purposes
- Financial gain: tailored ads
- Intelligence collection
- Censorship

Attackers:

- Eavesdroppers
- DNS/ networks operators
- Third party service providers
- URI dialing for VoIP (looking up phone number)

Question:

Is this really of concern for Enterprise Network?

DNS Data Integrity

Privacy ./. AVC

Security ./. User Experience

End2End Encryption ./. Company Policies

Security Audits?

DNS poisoning or spoofing, or similar vulnerabilities generally requires the attacker to take advantage of poorly configured or vulnerable DNS servers.

Privacy for DNS?

Large effort within research and operations communities to protect DNS

- Proposals for confidentiality
  - DNSCurve / DNSCrypt
  - DNS over TLS
  - Opportunistic encryption
- **Differences**
  - What is protected
  - Channel vs DNS record
  - Adoption requirements
  - Changes to DNS message format
  - Changes to DNS software
  - Key management
  - Required volume support of TCP
- **Infrastructure compatibility, Protocol support**
  - Cloud provider, how to handle third party proxies?
  - Support of basic protocols : TCP, which version?
  - DNS and side channels: timing, sizes, domains dependencies, browsers' prefetching,…
- **DNSSEC for data integrity**
  - Singing DNS resource records using PKI

How About DNS Authenticity DNSSEC

Singing DNS resource records using PKI

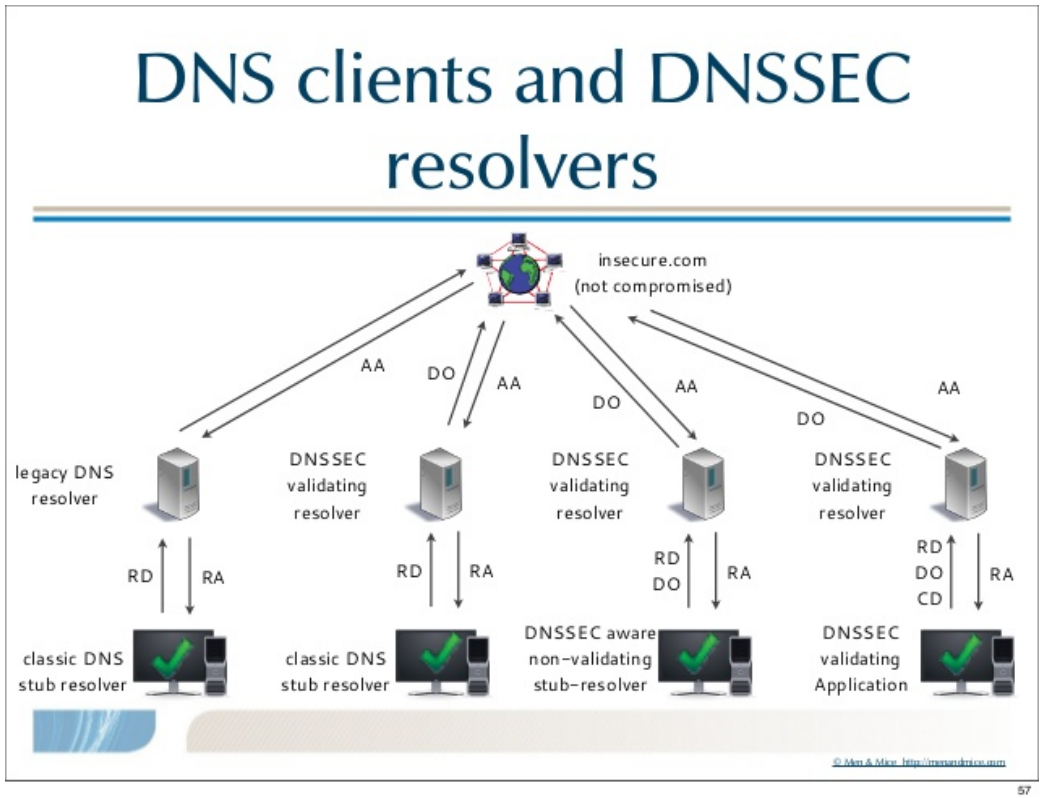

- DNSSEC works by digitally signing each DNS record so that any tampering of that record can be detected.
- The digital signatures, and keys used to create them, are distributed just like any other records in the DNS making DNSSEC backward compatible.
- Keys in each layer in the DNS hierarchy are signed by keys from the preceding layer which effectively vouches for them just like domain names are delegated from one layer to the next.
- This "chain of trust" is used to validate the digital signatures accompanying DNSSEC protected records to detect changes.

Controllers – Types

There's nothing like "the SDN controller"

- SDN Config-Pusher
  - Orchestration (robot micromanaging manual to-do's)
  - NCM (Network Configuration Management)
  - Customers may see or edit any part of the config
  - ✓ Prime Infrastructure, Action Packed, Solarwinds
  - ✓ Puppet, Chef
  - ✓ Openstack
  - ✓ Netconf
- SDN Policy-Compiler
  - Customer is never exposed to nor has access to nor influence over direct snippets of configuration elements.
  - They express their intent and the conversion to machine language is invisible.
  - ✓ Cisco APIC-EM
- SDN Policy-Enabler
  - ✓ Cisco APIC-DC
- SDN Overlay Controller
  - ✓ VMWare: VCS, VCD, NSX
  - ✓ VSM (N1kv), EVP, VTS
  - ✓ Windows Server, Microsoft System Center)
- SDN Open Flow Controller
  - ✓ Primary for research

Cross Controller Domain Policy (NIC, GBP, NEMO)

Cross Domain Application Policy (DNS-AS)

DNS-AS

- leverages **DNS** as an **Authoritative Source** to publish metadata as a key for common policy across networks without the need for a dedicated (SDN) controller.
- https://dns-as.org

NIC

- **Network Intent Composition (Open Daylight)**
- Manage and direct network services and network resources
- describing the "Intent" for network behaviors and network policies.
- Abstracted policy semantics instead of Openflow-like flow rules

GBP

- **Group Based Policy**
- Placing endpoints into groups (EPGs) that share the same semantics
- Defining how these endpoints need to communicate.
- Represent the requirements of the application and then force the infrastructure to figure out how to meet these requirements,
- Rather than defining the policy in terms of the underlying infrastructure.

NEMO

- **Network Modeling Language**
- Allows applications to use intent-based policy to create virtual networks comprised of nodes with policy-controlled flows.
- Intent based policy is prescriptive
- Leaving the details to the network

Span of Control - Cross Controller Architecture

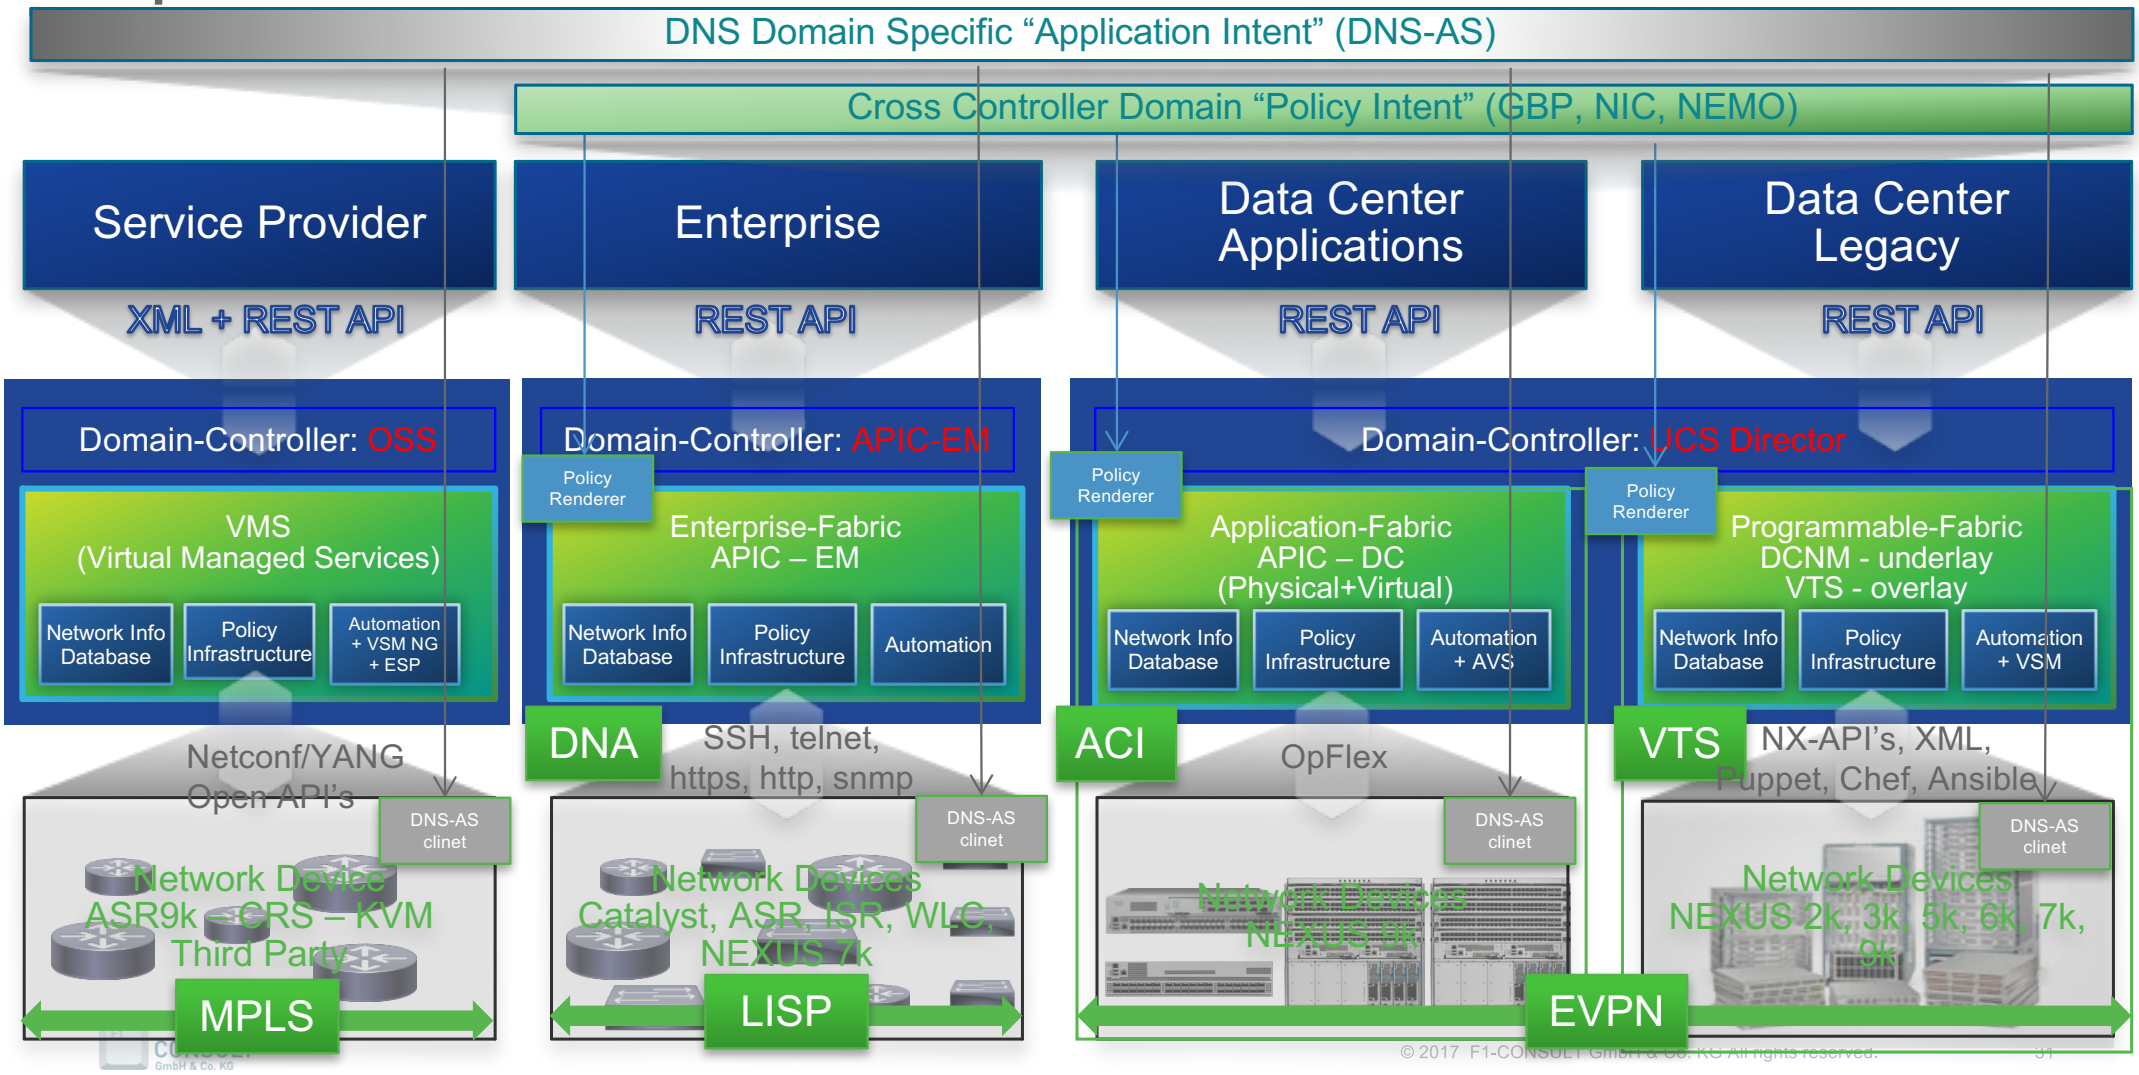

1.2 Application and Protocol challenges

The World After “Snowden”

Growth of Encrypted Network Traffic

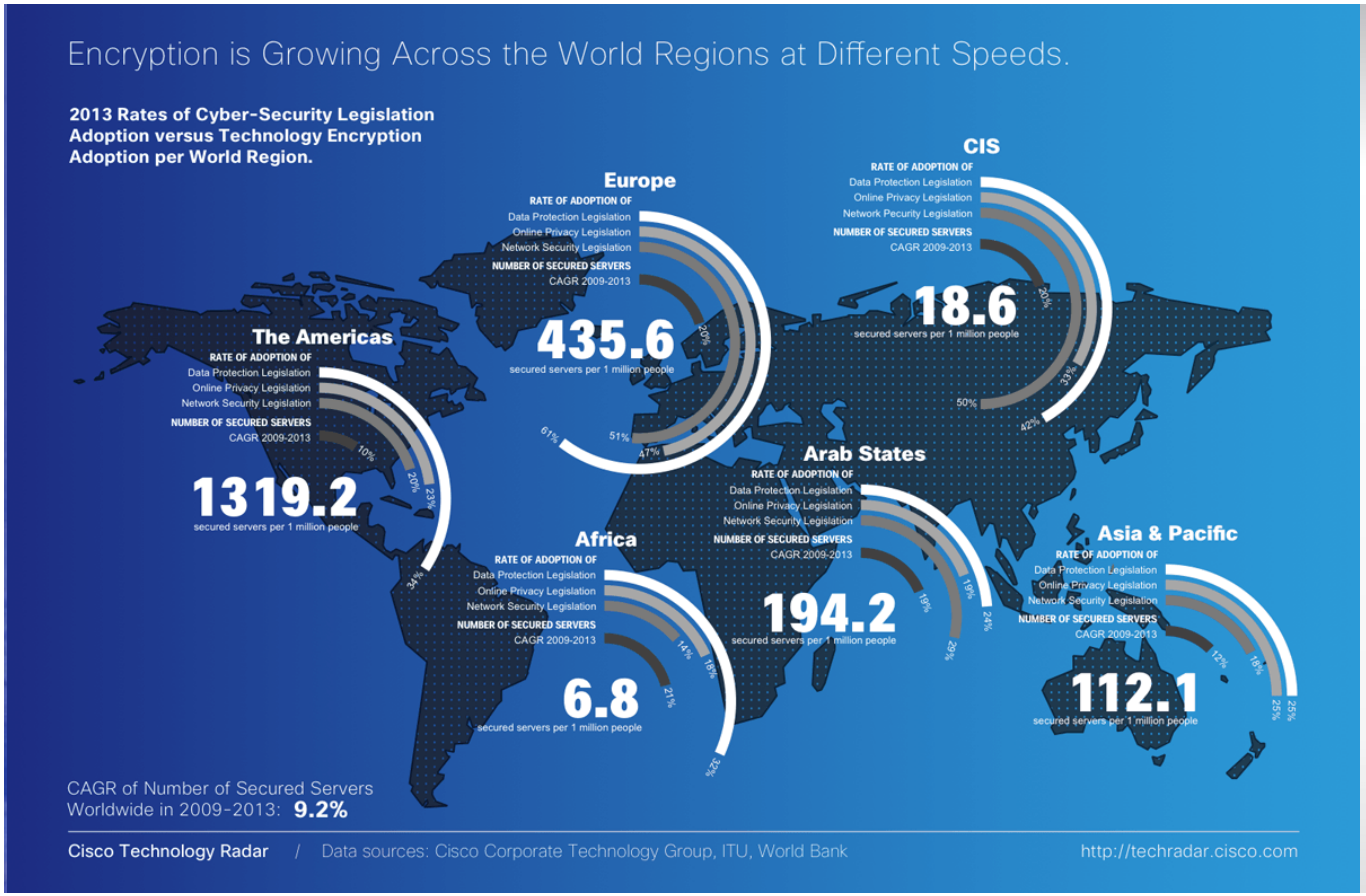

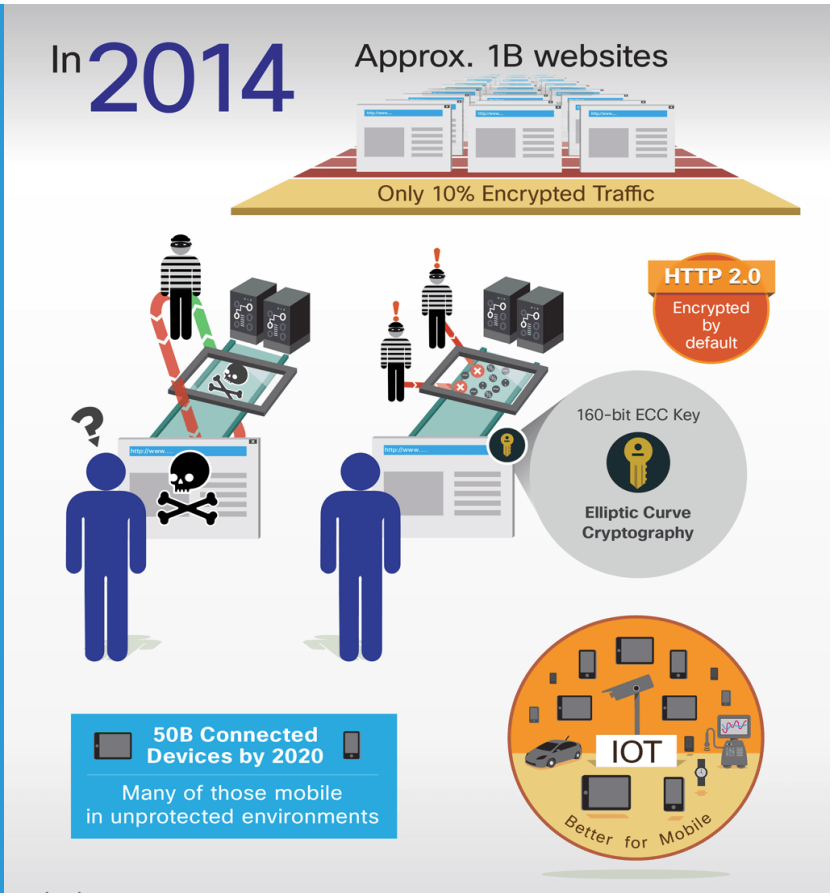

The World After “Snowden”

Protocol Evolution – HTTP/1, SPDY, QUIC, HTTP/2

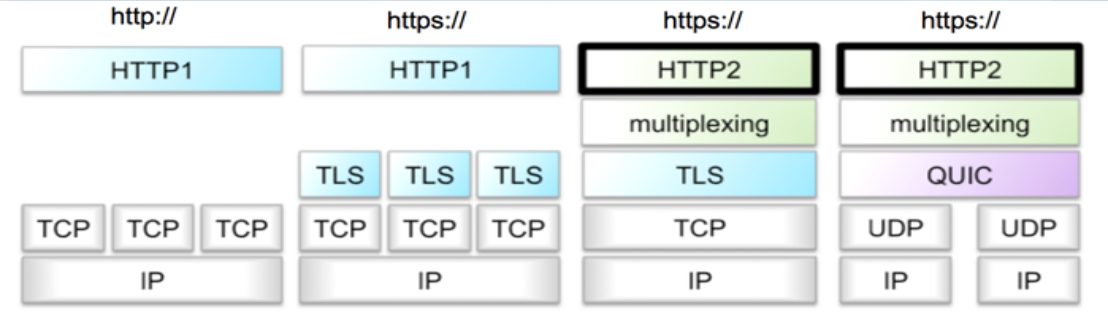

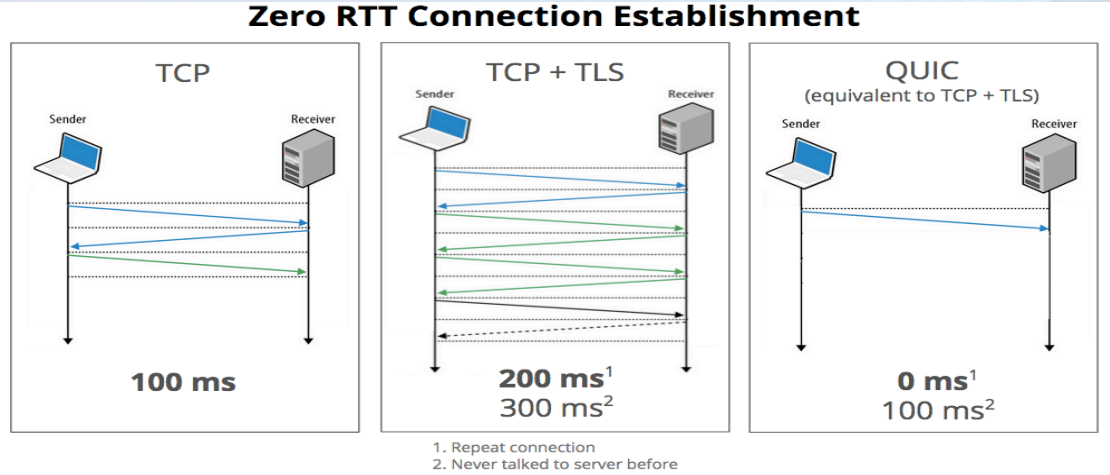

- HTTP/1.0 was pioneered in the late 80’s
- TCP + TLS requires 2 to 3 round trips

- HTTP/2 February 2015 IETF steering group announced completion
- Real performance improvement over TCP
- zero-round-trip connection establishment
- encryption capability by default

- QUIC: bundles streams over the same UDP connection
- If your firewalls block bi-directional UDP traffic, QUIC is blocked also.
- How to differentiate your could delivered QUIC app from an UDP attack?
- How about ICMP to the host

Living in a after “Snowden” world

Google Shame All Websites That Are Unencrypted - [Motherboard]

Google's Eric Schmidt: 'the solution to government surveillance is to encrypt everything'

By Nathan Ingraham on November 21, 2013 02:50 pm Email @NateIngraham

- Google wants everything on the web to be travelling over a secure channel.
- Google Announces 97 Percent of YouTube Traffic is Now Encrypted
- More important is to understand some implications:
  - Prevent content tampering, deny last mile SP to replace, add or filter out advertisement
  - Eliminating the ability of transparent proxies to muck up streaming protocols
  - Prevent last mile SP analytics, monitoring and monetization of user behavior
  - Net-Neutrality, Peering Agreements

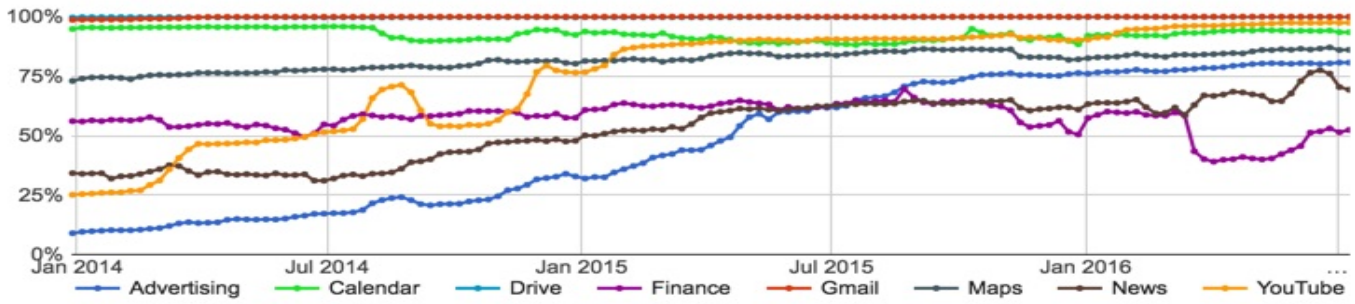

This is an approximate number that represents most of Google traffic for the given product.

- Chrome: “chrome://flags”
- navigate to “mark non-secure as” and selecting “mark non-secure origins as non-secure.”

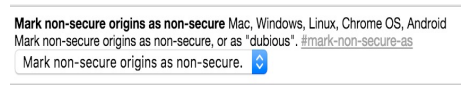

Living in a after “Snowden” world

It becomes harder and harder for us to “guess”

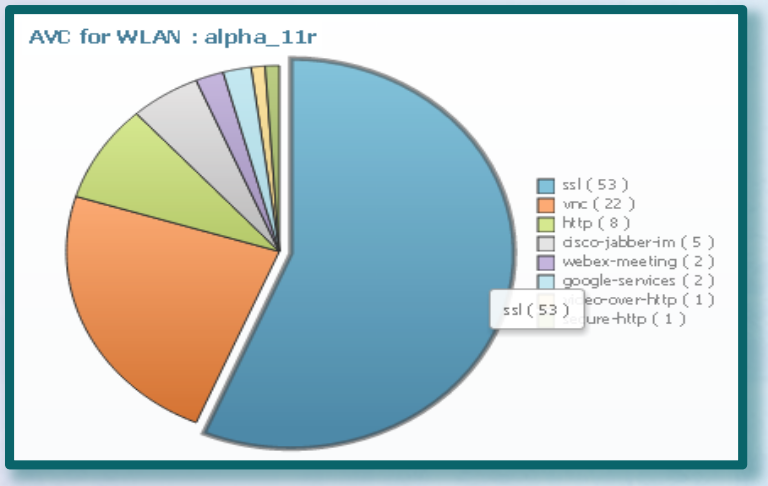

Bottom line: It becomes harder and harder for us to look into into traffic streams in order to “guess” what the apps are based on snooping traffic.

1.3 Evolution of AVC

AVC - Use Cases

Know Applications (includes Growing Encrypted apps) In Your Network Granularly

Needs Support across various PINS - Wireless, UA, WAN/Internet edge, Core, DC, Security

Application Level Reporting

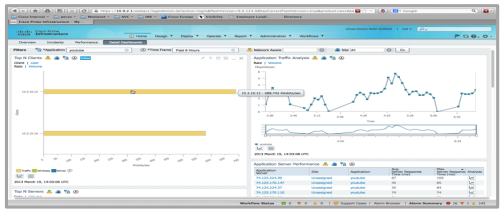

Visibility, Capacity Planning, Reporting on LAN & WAN

Network Data Analytics

Use Application Information to Drive Network Data Analytics – e.g. CMX/ wireless scenarios

Business Level Policy Enforcement

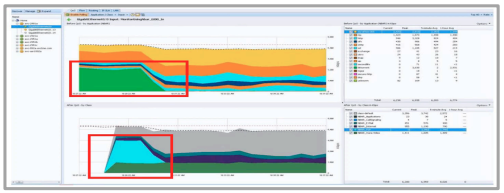

E2E QoS & ACL (any Policy) enforcement – Drop "selectively", Access Marking & Core/WAN Queuing

App-Aware "Domain Based" Routing

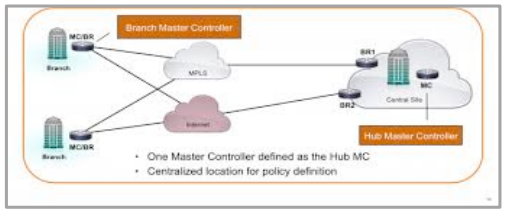

To support cloud apps breakouts to the Internet based on app-aware Routing policies

Application Level Troubleshooting & Easy Fault Isolation

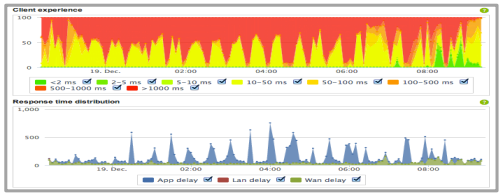

Zoom in on "Where The Problem Is" for business applications – could be ANYWHERE!

Network Readiness for Application Deployment

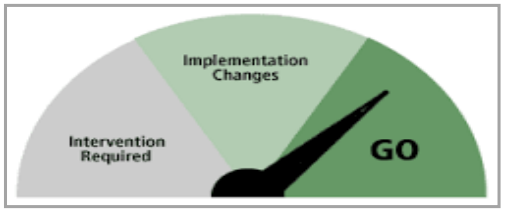

"Readiness Assessments" – Determine readiness for Application Deployment at planned scale

AVC – End to End – How?

Requirements for Future Application Identification:

We need an
Authoritative
Light-Weight
Unambiguous
way to identify applications.

We then need to be able to
link that Application Identity to Organizational Policy
for enforcement, accounting, etc.

How can we do this while addressing the challenges noted?

Network Metadata

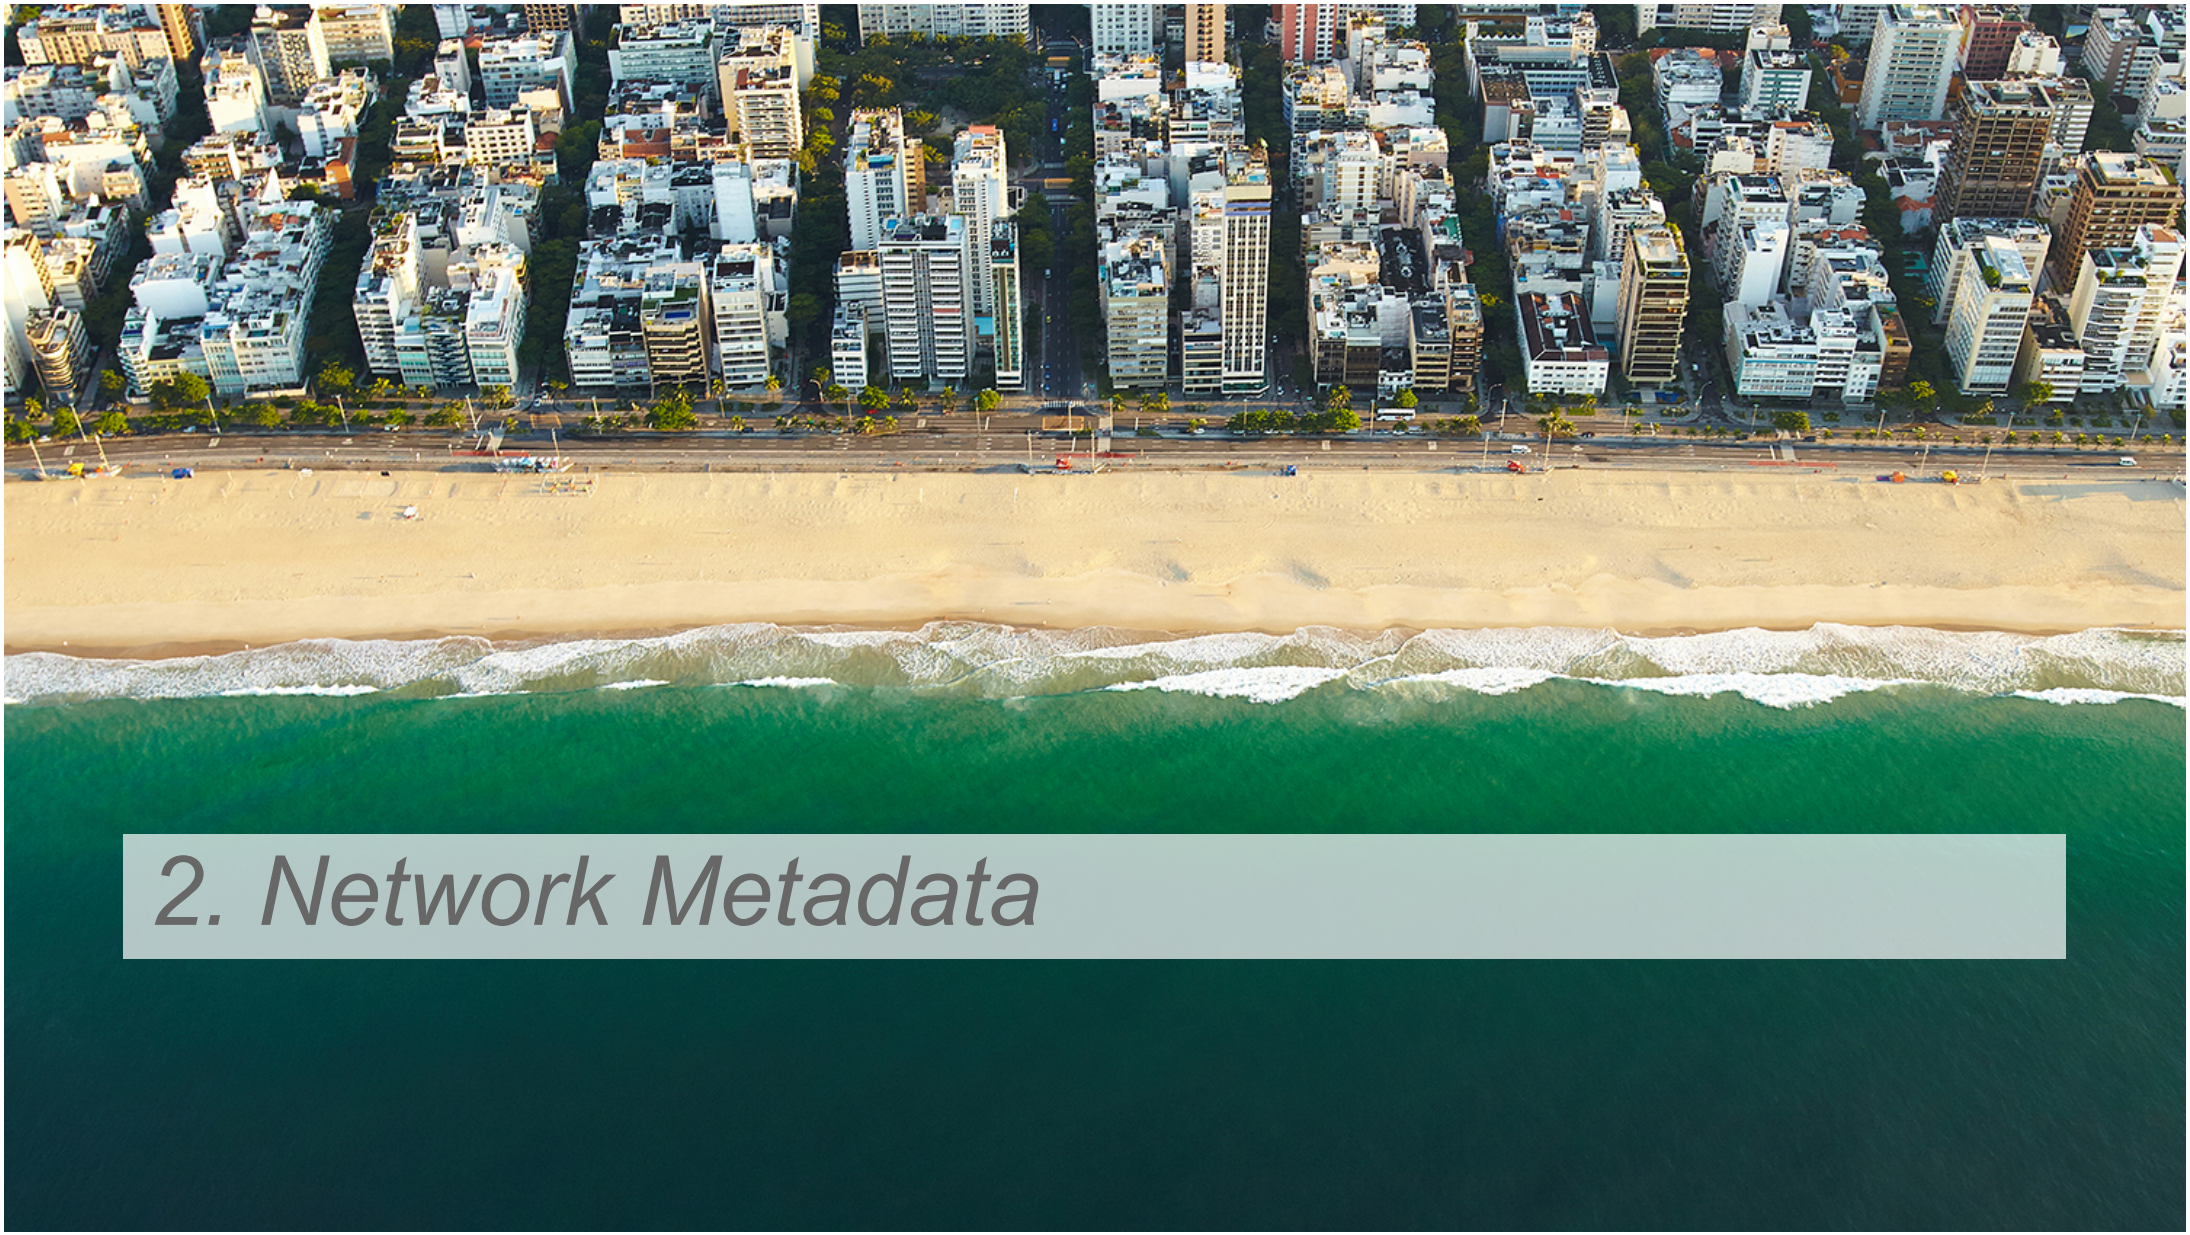

2. Network Metadata

Network Metadata

What is it? Why do we need it?

- literally, “data about the data”
- Identify Enterprise Applications
- Describe what the application **IS**
- Describe what the application **NEEDS**
- No longer any guessing

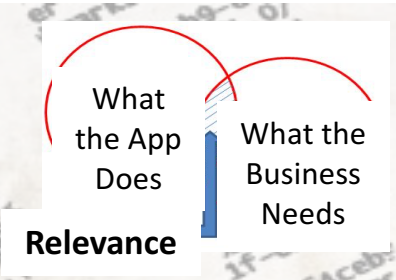

Instead of guessing device by device, we holistically program the network via DNS-AS metadata

Network Metadata – possible sources of truth

Multiple Application ID's out there

- **SEC: Snort Open App ID**
- SourceFire
- FireSIGHT eStreamer Application Protocol
- NBAR
- Meraki
- Simple DNS Matches
- Application Information in IP Flow Information Export (IPFIX)
- **AVC: Global Application ID assignment model**
  http://www.rfc-editor.org/rfc/rfc6759.txt

Application Network Metadata – DNS-AS

RFC6759 Metadata Components

Attributes	Short Name	Comments
Application Name	**app-name**	**custom names are possible, minimum length to be 3 chars**
Application ID	**app-id**	**RFC 6759 based application ID names**
Application Category	app-category	
Application Sub-Category	app-sub-category	
Traffic Class (QoS)	**app-traffic-class**	**RFC 4594 based short names**
Business Relevance	**business**	**[YES|NO|DEFAULT]**
Next Hop	next	NSH - Service Chaining Next Hop
Attributes (tunneled, encrypted, p2p)	tunneled, encrypted, p2p	tunneled, encrypted, p2p
Server Port Range	port-range	to identify an application by ports
IP Protocol Specifier	ip-protocol	
IP Version Specifier	ip-version	
Min/Avg/Max Bandwidth consumption	min-bw, avg-bw, max-bw	
Max. Possible Packet Loss	max-pkt-loss	In %
Max. Possible Jitter	max-jitter	In ms
Max. Possible Latency	max-latency	In ms
Metadata derived from	source	NBAR2, DNS-AS-server, DNS-AS-proxy, RPZ

DNS-AS Application Metadat

RFC1035 Metadata Components within TXT and AVC RTYPEs

TXT RDATA format

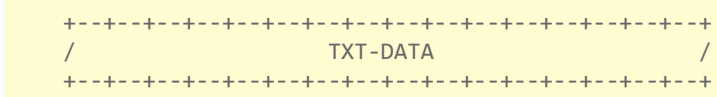

TXT-DATA One or more <character-string>’s

- **depreciated** for DNS-AS
- to be used for backward compatibility reasons
- with not so current DNS servers
- before BIND 9.9.9b1

General DNS-AS TXT record syntax:
"CISCO-CLS=<option>:<val>{|<option>:<val>}*”

AVC RDATA format

AVC-DATA One or more <character-string>’s

- the official **IANA assigned** mnemonic
- preferred RR going forward if the authoritative DNS server already supports this new RR
- Starting with BIND 9.9.9b1 / BIND 9.10.4b2

General DNS-AS TXT record syntax:
"<option>:<val>{|<option>:<val>}*”

- You may have multiple "strings" in a single resource record
- Each "string” may be up to 255 characters in length
- RDATA itself may not exceed 65535 bytes in total
- That 64K limit is a general restriction on DNS records of all types
- Any DNS response which exceeds 512 bytes is slightly undesirable, or use EDNS0
- Responses which exceed 512 bytes will signal truncation and prompt a retry via TCP, optimal to stay within 512 bytes if possible.
- General DNS-AS RR record syntax: ‘<option>:<val>{|<option>:<val>}*’
- Option-value pairs may appear in the same record, separated by a pipe character ‘|’.
- Example for a TXT record with app metadata would be: **"CISCO-CLS=app-name:wolfgang|app-id:CU/67244”**
- Example for a AVC record with app metadata would be: **"app-name:wolfgang|app-id:CU/67244"**

DNS-AS Application Metadata

Metadata Lookup Sequencing with mixed TXT and AVC RTYPEs

Default RDATA Lookup Sequence:

```
1. query for AVC RDATA
        QTYPE=AVC for wolfgang.dns-as.org
        -> "app-name:dns-as-wolfgang|app-class:TD|business:YES|app-id:CU/28203”
        if NODATA or ANCOUNT=0 then goto 2

2. query for RPZ RDATA
        QTYPE=AVC for _avc.wolfgang.dns-as.org
        -> "app-name:dns-as-wolfgang|app-class:TD|business:YES|app-id:CU/28203”
        if NODATA or ANCOUNT=0 then goto 3

3. query for TXT RDATA
        QTYPE=TXT for wolfgang.dns-as.org
        -> "CISCO-CLS=app-name:dns-as-wolfgang|app-class:TD|business:YES|app-
        id:CU/28203”
        if NODATA or ANCOUNT=0 then goto 4

4. no DNS-AS related metadata available
           -> NBAR
```

Override options by trusted-domains:

```
!
avc dns-as client enable
!
avc dns-as client trusted-domains
 domain ^.*f1.*$ AVC RPZ TXT
 domain ^.*cisco.*$ TXT RPZ AVC
 domain *.toocoolforyou.net AVC RPZ TXT
 domain *.blackberry.net TXT
 domain *.dns-as.org AVC
 domain *.nbar2web.org
 domain *.f1-consult.com RPZ
 domain *.f1-consult.de
 domain *.f1-online.net
 domain *.f1v4.net
 domain *.f1v6.net
!
```

We need to accommodate:

- Zones that provide their own AVC information
- Zones who don't provide any AVC information
- Zones whose provided AVC information you want to override locally
- All other DNS lookups passing unimpeded/unaltered

- Query in that sequence and just sent the QTYPES been listed behind the trusted-domain label.
- If there is no QTYPE listed, just follow the default lookup sequence.

F1 CONSULT GmbH & Co. KG

Network Metadata – AVC Components

Metadata Components for Application Visibility

Important Application Visibility Attributes:

- ✓ Application Name (app-name)
- ✓ Application ID (app-id)

Optional Application Visibility Attributes:

- Attributes (tunneled, encrypted, p2p)
- Server Port Range (to identify an application with ports)
- IP Protocol Specifier
- IP Version Specifier
- Source of Metadata (NBAR2, DNS-AS server etc.)

TXT Example:

```
"CISCO-CLS=app-name:smtp|app-id:IL4/25|server-port:TCP/25,UDP/25”
```

AVC Example:

```
"app-name:smtp|app-id:IL4/25|server-port:TCP/25,UDP/25”
```

Network Metadata – AVC Components

Metadata Components for Application Policy Intent

Important Application Intent Attributes:

- ✓ Traffic Class (`app-class`)
- ✓ Business Relevance (`business`)

Optional Application Intent Attributes:

- Application Category
- Application Sub-Category
- Server Port Range (to identify an application with ports)
- Min/Avg/Max Bandwidth consumption
- Max. Possible Packet Loss (in %)
- Max. Possible Jitter (in ms.)
- Max. Possible Latency (in ms.)

TXT Example:

```
"CISCO-CLS=app-name:smtp|app-class:bulk-data|business:YES|app-id:IL4/25|server-port:TCP/25,UDP/25”
```

AVC Example:

```
"app-name:smtp|app-class:bulk-data|business:YES|app-id:IL4/25|server-port:TCP/25,UDP/25”
```

NBAR and DNS-AS

Different Tools for Different Problems

```
Starting Nmap 5.51 ( http://nmap.org ) at 2016-07-04 18:03 CESTN
map scan report for exchange.toocoolforyou.net (192.168.168.240)
Host is up (0.00042s latency).
Not shown: 976 closed ports
PORT     STATE SERVICE
25/tcp   open  smtp
80/tcp   open  http
110/tcp  open  pop3
135/tcp  open  msrpc
139/tcp  open  netbios-ssn
143/tcp  open  imap
443/tcp  open  https
445/tcp  open  microsoft-ds
587/tcp  open  submission
593/tcp  open  http-rpc-epmap
808/tcp  open  ccproxy-http
993/tcp  open  imaps
995/tcp  open  pop3s
1025/tcp open  NFS-or-IIS
1026/tcp open  LSA-or-nterm
1027/tcp open  IIS
1688/tcp open  nsjtp-data
3389/tcp open  ms-term-serv
5357/tcp open  wsdapi
5666/tcp open  nrpe
6001/tcp open  X11:1
6002/tcp open  X11:2
6003/tcp open  X11:3
6004/tcp open  X11:4
```

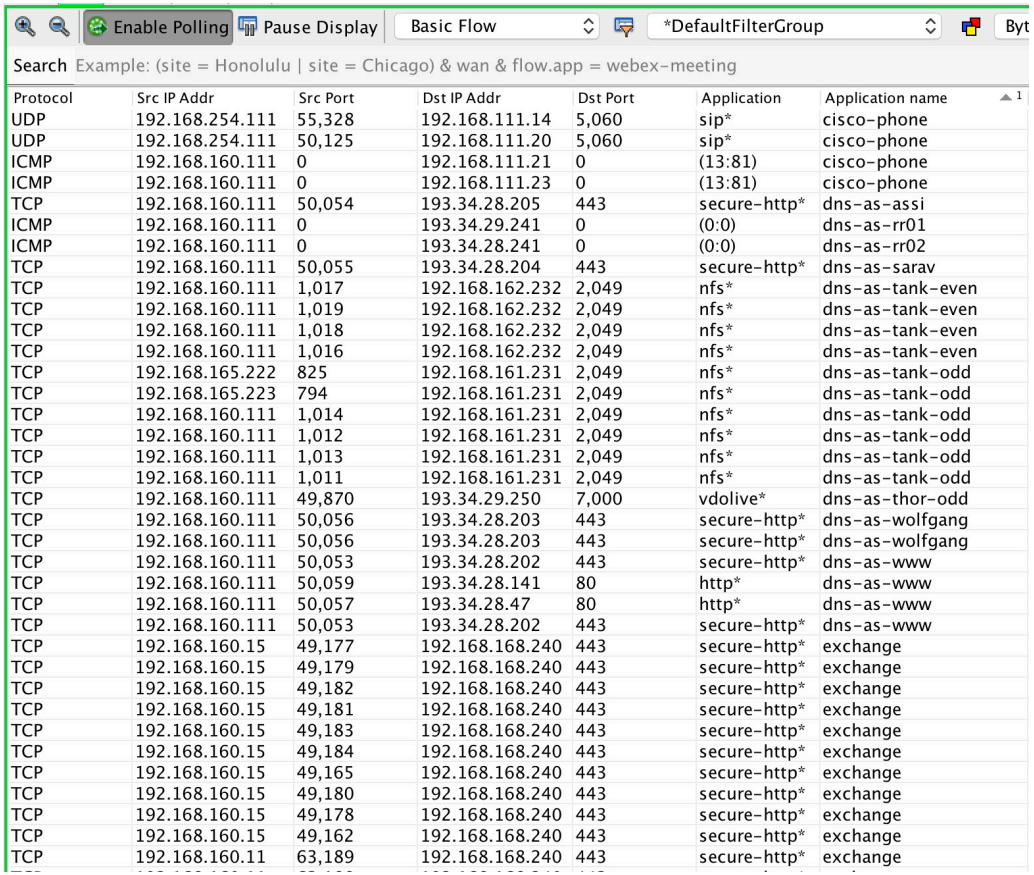

Protocol	Src IP Addr	Src Port	Dst IP Addr	Dst Port	Application	Application name
UDP	192.168.254.111	55,328	192.168.111.14	5,060	sip*	cisco-phone
UDP	192.168.254.111	50,125	192.168.111.20	5,060	sip*	cisco-phone
ICMP	192.168.160.111	0	192.168.111.21	0	(13:81)	cisco-phone
ICMP	192.168.160.111	0	192.168.111.23	0	(13:81)	cisco-phone
TCP	192.168.160.111	50,054	193.34.28.205	443	secure-http*	dns-as-assi
ICMP	192.168.160.111	0	193.34.29.241	0	(0:0)	dns-as-rr01
ICMP	192.168.160.111	0	193.34.28.241	0	(0:0)	dns-as-rr02
TCP	192.168.160.111	50,055	193.34.28.204	443	secure-http*	dns-as-sarav
TCP	192.168.160.111	1,017	192.168.162.232	2,049	nfs*	dns-as-tank-even
TCP	192.168.160.111	1,019	192.168.162.232	2,049	nfs*	dns-as-tank-even
TCP	192.168.160.111	1,018	192.168.162.232	2,049	nfs*	dns-as-tank-even
TCP	192.168.160.111	1,016	192.168.162.232	2,049	nfs*	dns-as-tank-even
TCP	192.168.165.222	825	192.168.161.231	2,049	nfs*	dns-as-tank-odd
TCP	192.168.165.223	794	192.168.161.231	2,049	nfs*	dns-as-tank-odd
TCP	192.168.160.111	1,014	192.168.161.231	2,049	nfs*	dns-as-tank-odd
TCP	192.168.160.111	1,012	192.168.161.231	2,049	nfs*	dns-as-tank-odd
TCP	192.168.160.111	1,013	192.168.161.231	2,049	nfs*	dns-as-tank-odd
TCP	192.168.160.111	1,011	192.168.161.231	2,049	nfs*	dns-as-tank-odd
TCP	192.168.160.111	49,870	193.34.29.250	7,000	vdolive*	dns-as-thor-odd
TCP	192.168.160.111	50,056	193.34.28.203	443	secure-http*	dns-as-wolfgang
TCP	192.168.160.111	50,056	193.34.28.203	443	secure-http*	dns-as-wolfgang
TCP	192.168.160.111	50,053	193.34.28.202	443	secure-http*	dns-as-www
TCP	192.168.160.111	50,059	193.34.28.141	80	http*	dns-as-www
TCP	192.168.160.111	50,057	193.34.28.47	80	http*	dns-as-www
TCP	192.168.160.111	50,053	193.34.28.202	443	secure-http*	dns-as-www
TCP	192.168.160.15	49,177	192.168.168.240	443	secure-http*	exchange
TCP	192.168.160.15	49,179	192.168.168.240	443	secure-http*	exchange
TCP	192.168.160.15	49,182	192.168.168.240	443	secure-http*	exchange
TCP	192.168.160.15	49,181	192.168.168.240	443	secure-http*	exchange
TCP	192.168.160.15	49,183	192.168.168.240	443	secure-http*	exchange
TCP	192.168.160.15	49,184	192.168.168.240	443	secure-http*	exchange
TCP	192.168.160.15	49,165	192.168.168.240	443	secure-http*	exchange
TCP	192.168.160.15	49,180	192.168.168.240	443	secure-http*	exchange
TCP	192.168.160.15	49,178	192.168.168.240	443	secure-http*	exchange
TCP	192.168.160.15	49,162	192.168.168.240	443	secure-http*	exchange
TCP	192.168.160.11	63,189	192.168.168.240	443	secure-http*	exchange

DNS-AS ./. NBAR

Application Name ./. Protocol the Application is delivered over

```
mingla#show avc dns-as client binding-table
```

Protocol name	Vrf	Ip List	Host	Age [min]	Text record	TTL [min]	Time to Expire [min]
dns-as-proxy-srv0	<default>	193.34.28.245	proxy2.f1-online.net	47	app-name:dns-as-proxy-srv02|app-class:BD|business:YES|app-id:CU/28245	188	163
dns-as-proxy-srv0	<default>	193.34.29.245	proxy1.f1-online.net	47	app-name:dns-as-proxy-srv01|app-class:BD|business:YES|app-id:CU/29245	185	163
dns-as-smtp-mx02	<default>	193.34.28.11	mx2.f1-online.net	48	app-name:dns-as-smtp-mx02|app-class:BD|business:YES|app-id:CU/28011	322	300
dns-as-smtp-mx01	<default>	193.34.29.11	mx1.f1-online.net	48	app-name:dns-as-smtp-mx01|app-class:BD|business:YES|app-id:CU/29011	322	300
dns-as-ns02	<default>	193.34.28.244	ns2.f1-online.net	48	app-name:dns-as-ns02|app-class:NC|business:YES|app-id:CU/28244	807	785
dns-as-ns01	<default>	193.34.29.244	ns1.f1-online.net	48	app-name:dns-as-ns01|app-class:NC|business:YES|app-id:CU/29244	808	786
dns-as-ns00	<default>	193.34.28.240	ns0.f1-online.net	48	app-name:dns-as-ns00|app-class:NC|business:YES|app-id:CU/28240	322	300
dns-as-rr02	<default>	193.34.28.241	rr2.f1-online.net	48	app-name:dns-as-rr02|app-class:NC|business:YES|app-id:CU/28241	778	756
dns-as-rr01	<default>	193.34.29.241	rr1.f1-online.net	48	app-name:dns-as-rr01|app-class:NC|business:YES|app-id:CU/29241	808	786
dns-as-wolfgang	<default>	193.34.28.203	wolfgang.dns-as.org	48	app-name:dns-as-wolfgang|app-class:TD|business:YES|app-id:CU/28203	325	299
dns-as-thor-odd	<default>	193.34.29.250	thor-odd.f1-online.net	52	app-name:dns-as-thor-odd|app-class:NC|business:YES|app-id:CU/29250	323	295
dns-as-adc-2	<default>	192.168.168.241	adc2.toocoolforyou.net	113	app-name:dns-as-adc-2|app-class:NC|business:YES|app-id:CU/68241	60	48
dns-as-ntp2	<default>	192.168.168.244	ntp2.toocoolforyou.net	114	app-name:dns-as-ntp2|app-class:NC|business:YES|app-id:CU/68241	60	22
dns-as-ntp1	<default>	192.168.167.244	ntp1.toocoolforyou.net	114	app-name:dns-as-ntp1|app-class:NC|business:YES|app-id:CU/67241	60	22

URL parsing ./. DNS-AS Metadata

A much less expensive way to achieve 80% of the goal

http://username:password@www.dns-as.org:443/path/file.name?query=string#anchor

```
{
        scheme: "http://"
        user: "username",
        password: "password",
        host: "www.dns-as.org",
        port: "8080",
        path: "/path/file.name",
        query: "?query=string",
        fragment: "#anchor”
}
```

- As of today to we need to parse the whole URL to get application specific granularity
- At a fraction of the cost in terms of CPU and Hardware requirements you get similar results
- You get 80% of the goal for 100% consistency
- From a technical feasibility point of view a key enabler for common policy across our product portfolio

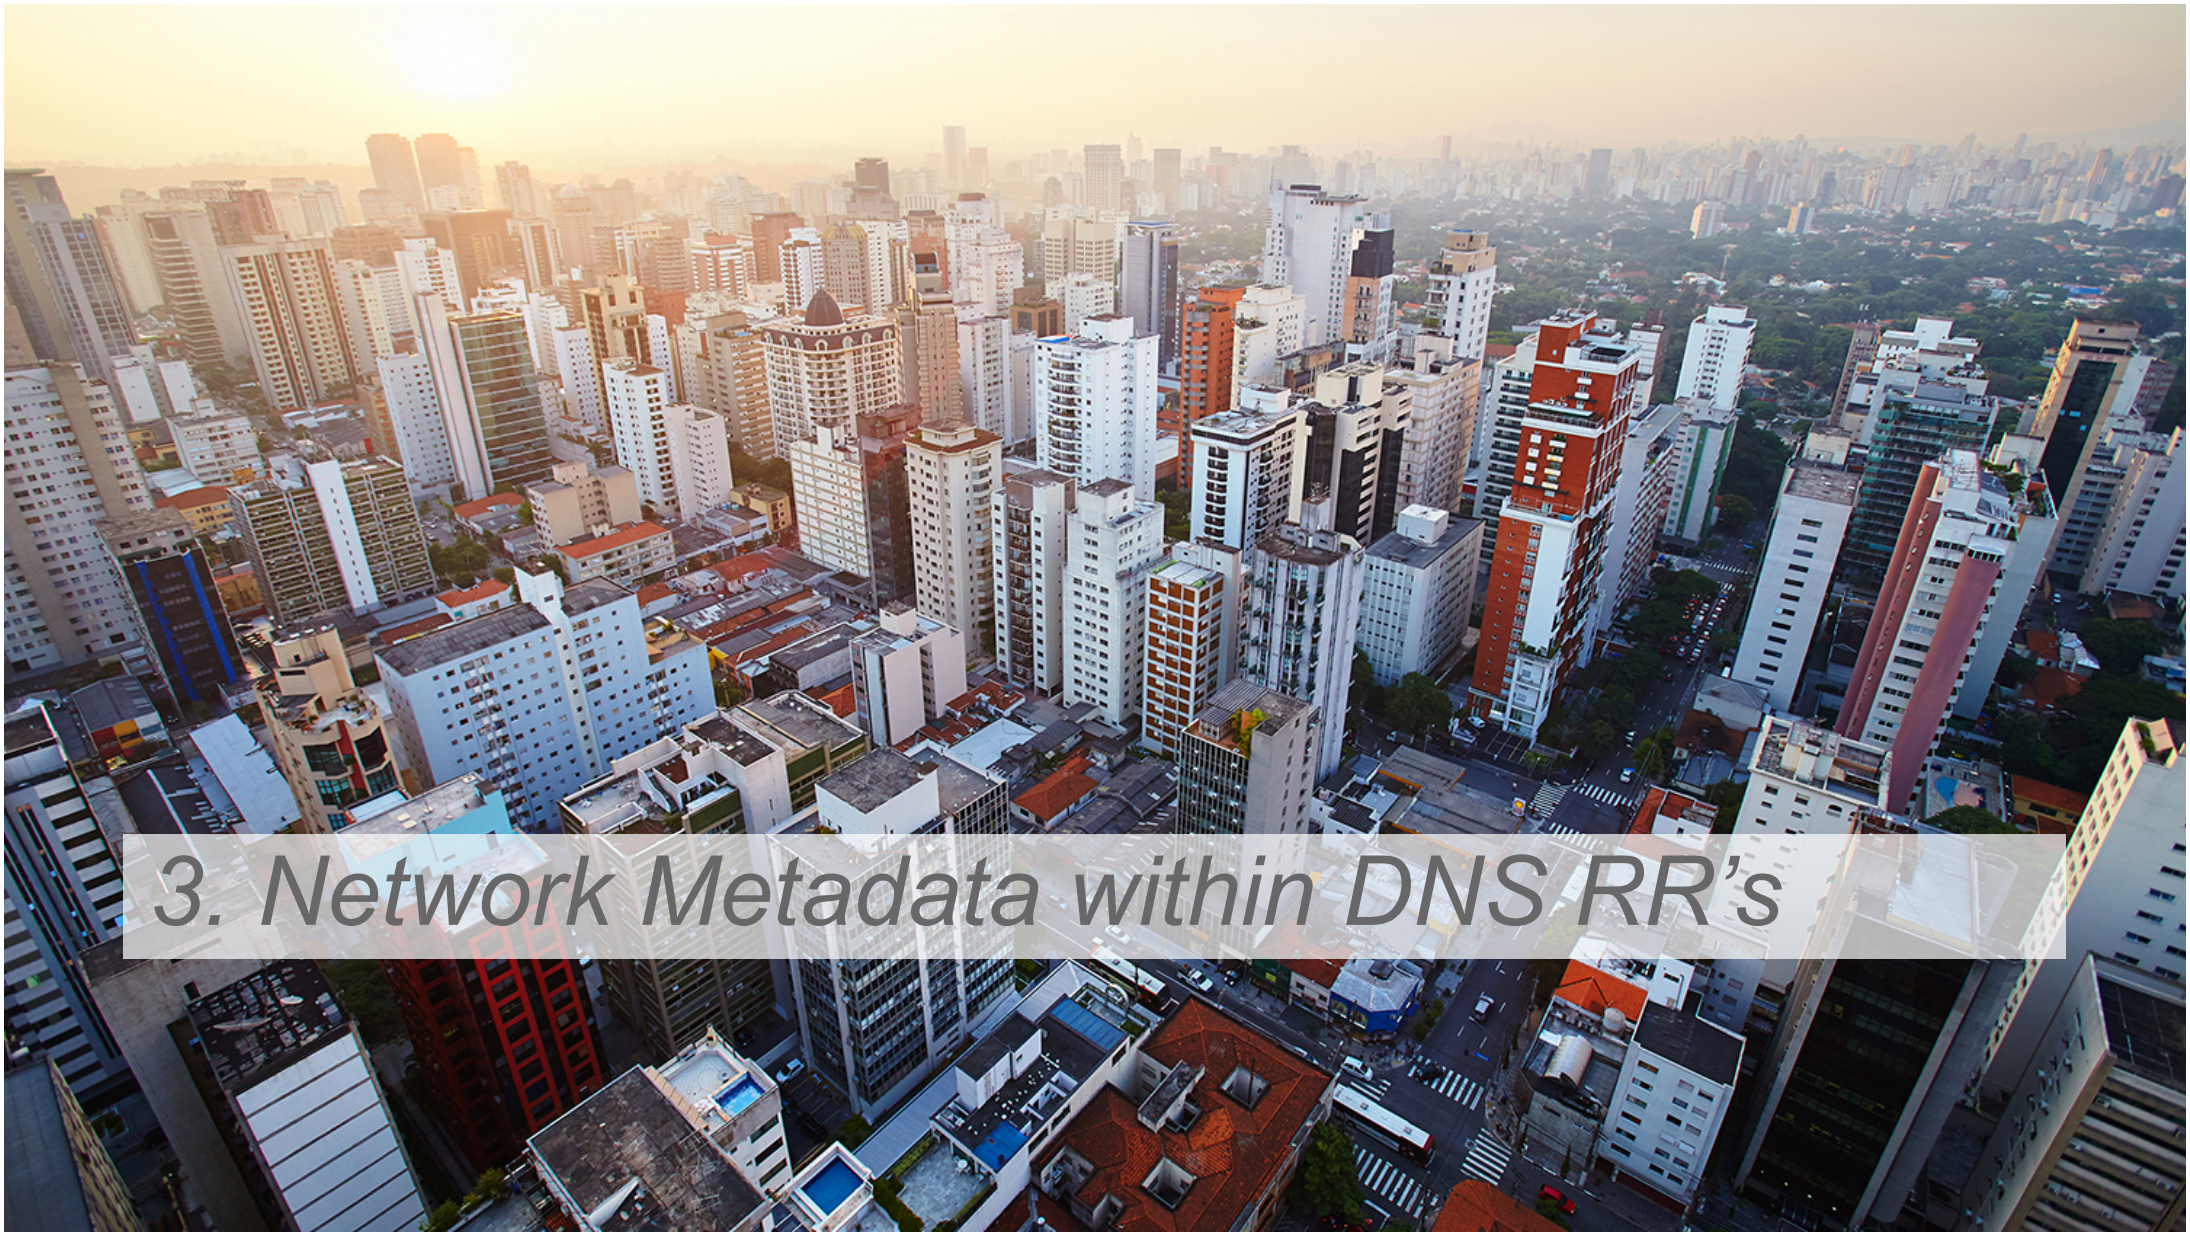

3. Network Metadata within DNS RR's

Network Metadata – How to Generate

https://www.dns-as.org/support/avc-rdata/

Define a TXT record for your Application based on NBAR2 Protocol Pack Taxonomy

Two options:

- **Generate Predefined** use this for well know applications using our best practice defaults
- **Generate Custom** use this for for your own applications using our own values

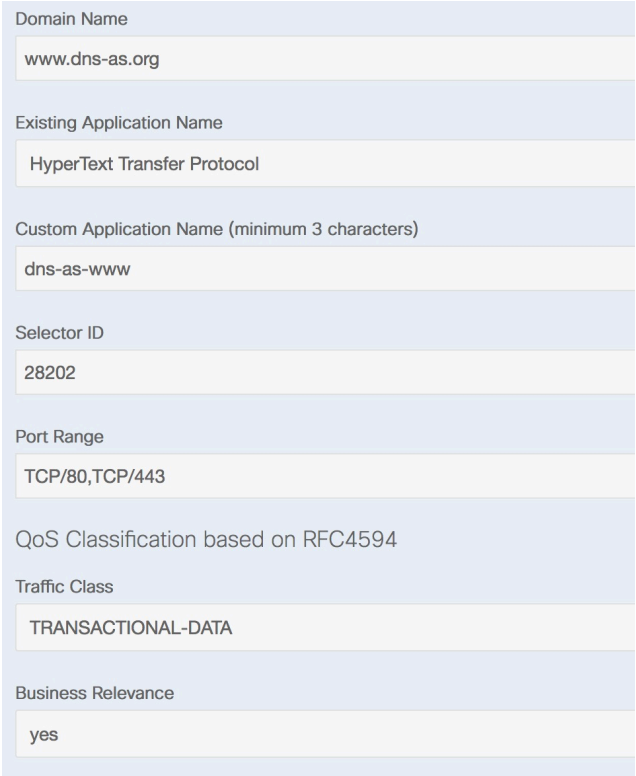

www.dns-as.org IN TXT "CISCO-CLS=app-name:dns-as-www|app-class:TD|business:YES|server-port:TCP/80,TCP/443|app-id:CU/28202"

Network Metadata – BIND

```
$ORIGIN .
$TTL 3600       ; 1 hour
dns-as.org      IN SOA  ns1.f1-online.net. hostmaster.f1-online.net. (
                                2016020101 ; serial ; serial
                                14400      ; refresh (3 hours)
                                3600       ; retry (1 hour)
                                604800    ; expire (2 weeks)
                                3600       ; minimum (1 hour)
                                )
                        NS      ns2.f1-online.net.
                        NS      ns1.f1-online.net.
                        A       193.34.28.202
                        TXT     "CISCO-CLS=app-name:HTTP|app-class:TD"
                        MX      10 mx1.dns-as.org.
                        MX      10 mx2.dns-as.org.
                        TXT     "v=spf1 mx a ip4:193.34.28.0/24 ip4:193.34.29.0/24 ~all“
```

```
$ORIGIN dns-as.org.
assi                    A       193.34.28.205
                        TXT     "CISCO-CLS=app-name:ASSI|app-class:NC“
mail                    A       193.34.28.201
                        A       193.34.29.201
                        TXT     "CISCO-CLS=app-name:MX00|app-class:BD|business=yes"
mx1                     A       193.34.29.201
                        TXT     "CISCO-CLS=app-name:MX01|app-class:BD|business=yes"
mx2                     A       193.34.28.201
                        TXT     "CISCO-CLS=app-name:MX02|app-class:BD|business=yes"
ns1                     A       193.34.29.200
                        TXT     "CISCO-CLS=app-name:DNS-AS|app-class:OAM|business=yes"
ns2                     A       193.34.28.200
                        TXT     "CISCO-CLS=app-name:DNS-AS|app-class:OAM|business=yes"
sarav                   A       193.34.28.204
                        TXT     "CISCO-CLS=app-name:SARAV|app-class:NC"
wolfgang                A       193.34.28.203
                        TXT     "CISCO-CLS=app-name:WOLFGANG|app-class:OAM"
www                     A       193.34.28.202
                        TXT     "CISCO-CLS=app-name:DNS-AS-WWW|app-class:TD"
```

Network Metadata – How to verify

Forward Zone:

```
$ dig TXT +short www.dns-as.org
"CISCO-CLS=app-name:dns-as-www|app-class:TD|business:YES|app-id:CU/28202"

$ dig TXT +short wolfgang.dns-as.org
"CISCO-CLS=app-name:dns-as-wolfgang|app-class:TD|business:YES|app-id:CU/28203"

$ dig TXT +short smtp.cisco.com
"CISCO-CLS=app-name:smtp|app-class:bulk-data|business:YES|app-id:IL4/25|server-port:TCP/25,UDP/25"

$ dig TXT +short inception.toocoolforyou.net
"CISCO-CLS=app-name:dns-as-exchange|app-class:BD|business:YES|app-id:CU/28111”

$ dig TXT +short topic.cisco.com
"CISCO-CLS=app-name:csco-topic|app-class:transactional-data|business:YES|app-id:CU/111|server-port:TCP/80,TCP/443"
```

Reverse Zone:

```
$ dig TXT +short 202.28.34.193.in-addr.arpa.
"CISCO-CLS=app-name:dns-as-www|app-class:TD|business:YES|app-id:CU/28202"

$ dig TXT +short 111.28.34.193.in-addr.arpa.
"CISCO-CLS=app-name:dns-as-exchange|app-class:BD|business:YES|app-id:CU/28111”
```

Network Metadata – Microsoft Active Directory

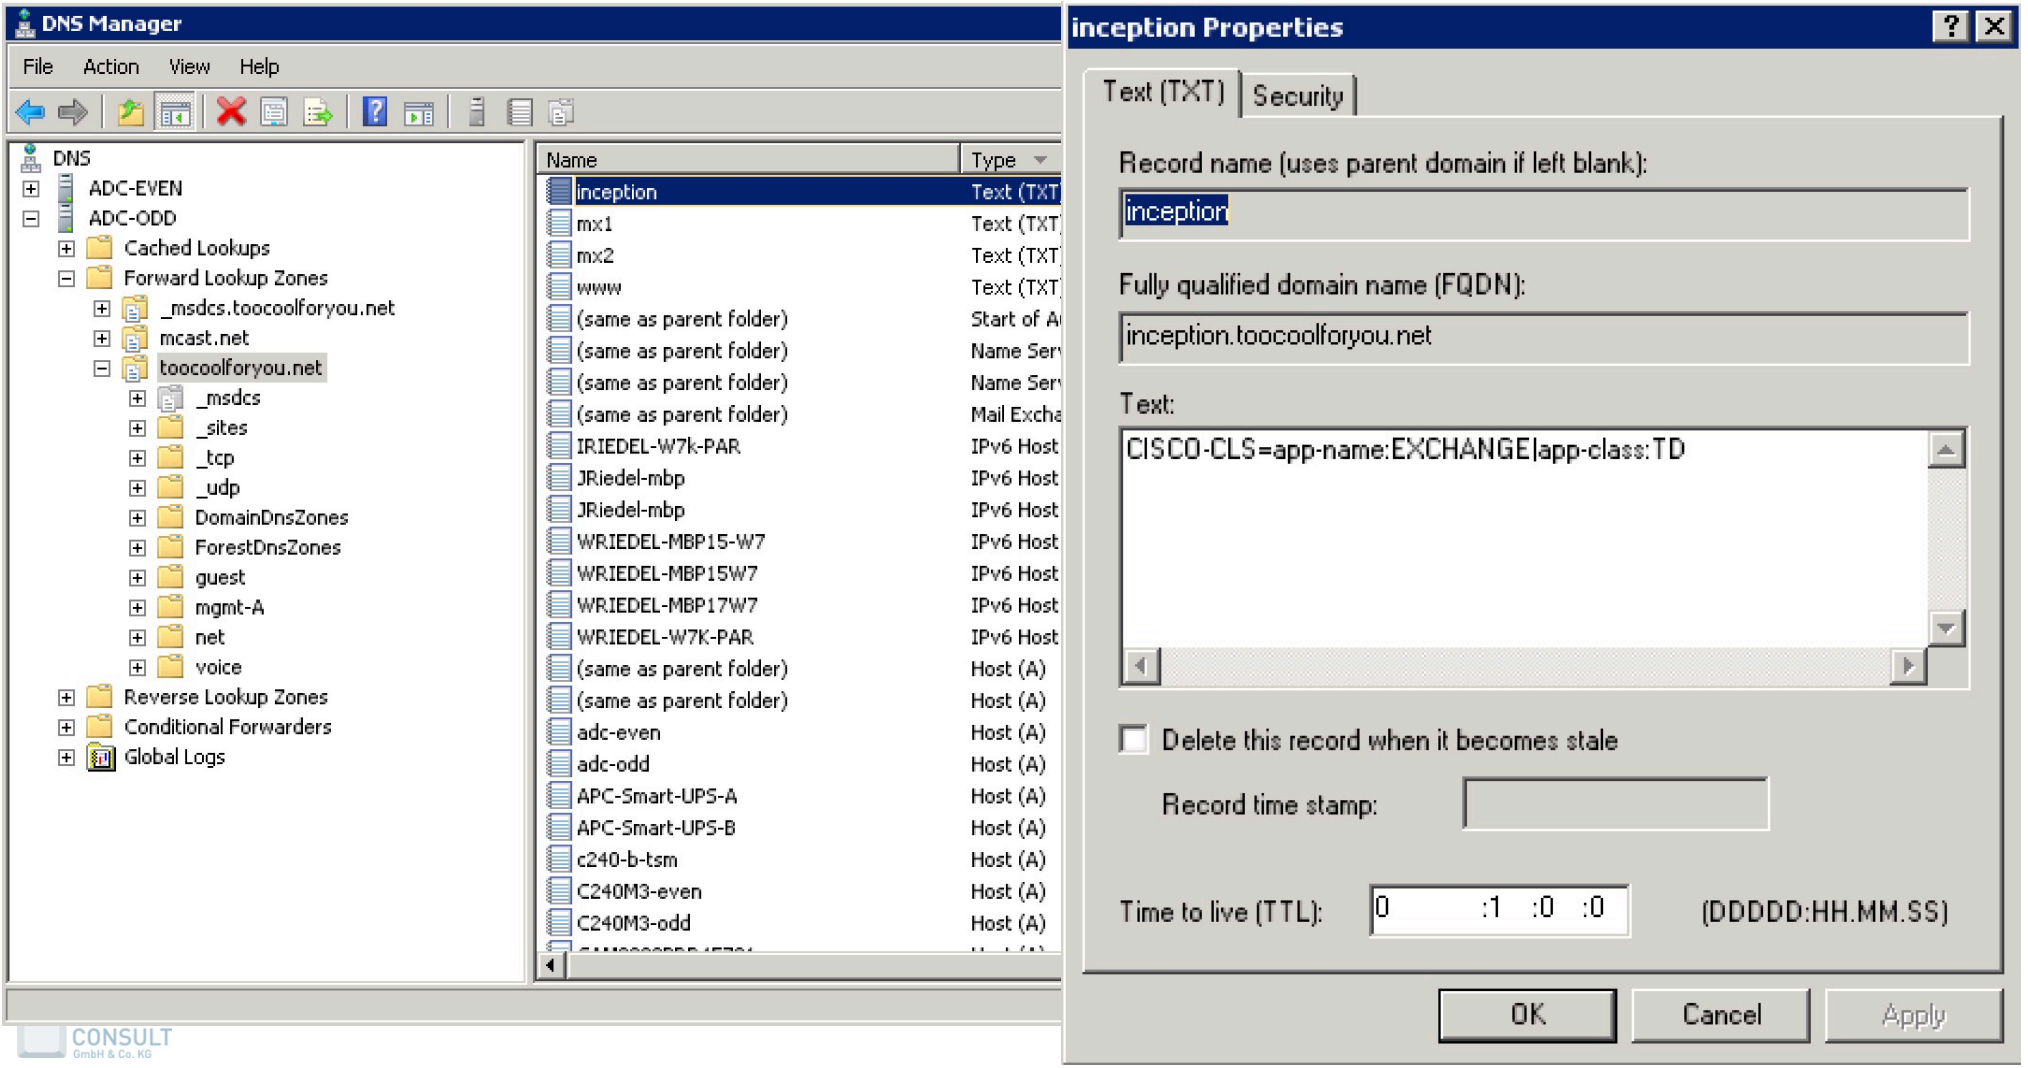

Network Metadata – Abstractions

Microsoft Office 365 with and without DNS-AS

without DNS-AS

*.outlook.com
*.microsoftonline.com
*.microsoftonline-p.com
*.microsoftonline-p.net
*.microsoftonlineimages.com
*.microsoftonlinesupport.net[1]
*.msecnd.net
*.office365.com
*.live.com
*.portal.microsoftonline.com
*.passwordreset.microsoftonline.com
*.msn.com
*.osub.microsoft.com

Ports 80/443
Protocols TCP and HTTPS
Rule must apply to all users
HTTPS/SSL time-out set to 8 hours

In reality, more then 140 entries

A full listing can be found here:
http://www.dns-as.org/support/das-as-cloud-apps/

with DNS-AS

DNS-AS metadata provided by MS:

```
AVC     "app-name:ms-update |app-class:BD|business=yes“
AVC     "app-name:ms-office365-web |app-class:BE|business=yes“
AVC     "app-name:ms-office365-outlook |app-class:BE|business=yes“
AVC     "app-name:ms-office365-live |app-class:MMS|business=yes"
AVC     "app-name:ms-office365-lync |app-class:VO|business=yes“
AVC     “ ... „
```

DNS-AS metadata consumed by customers

```
avc dns-as client trusted-domains
 domain ^.*outlook.*$
 domain ^.*microsoft.*$
 domain ^.*lync.*$
 domain ^.*sway.*$
```

4. How to control “foreign” domains

DNS Firewall Response Policy Zones (RPZ)

BIND Response Policy Zones

- Most modern electronic crime and network abuse relies on the Domain Name System (DNS)
- A DNS firewall can selectively intercept DNS resolution for known-malicious network assets including domain names, IP addresses, and name servers.
- Interception can mean rewriting a DNS response to direct a web browser to a "walled garden", or simply making the malicious network assets invisible and unreachable.
- Policies are applied only on DNS requests that ask for recursion (RD=1) and which either do not request DNSSEC metadata (DO=0) or for which no DNSSEC metadata exists.

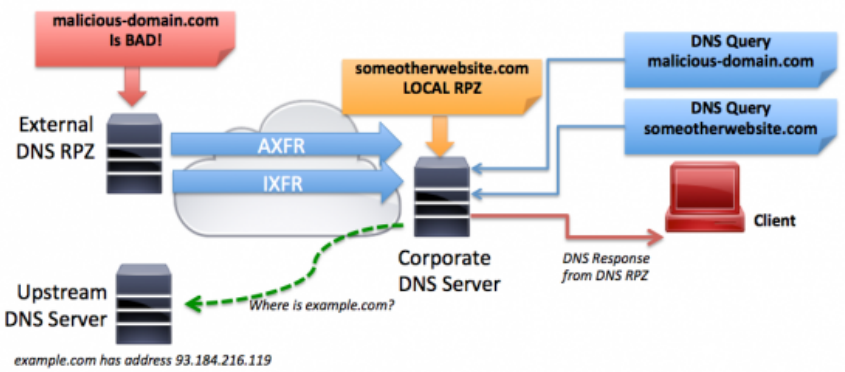

A response policy in DNS RPZ can be matched as follows:

- by the query name (QNAME)
- by an address which would be present in a truthful response
- by the name or address of an authoritative name server responsible for publishing the original response.

A response policy action can be one of the following:

- to synthesize a "domain does not exist" response
- to synthesize a "name exists but there are no records of the requested type" response.
- to replace the response with specified data.
- to exempt the response from further policy processing.

DNS Firewall Response Policy Zones (RPZ)

In a DNS RPZ firewall, the policy rule set is contained in a DNS "zone", which can be transferred using normal "zone transfer" mechanisms. The master copy of your DNS firewall policy can be a DNS "zone file" which you either edit by hand, or which you generate from a database. You can also edit a DNS zone indirectly using DNS dynamic updates (for example, using the "nsupdate" shell level utility.)

RPZ is not a standard DNS feature defined by an IETF RFC. It is, however, an **Open** specification (currently Format 3) whose authors have made it freely available.
It is Copyrighted by ISC but annotated "Distribution of this memo is unlimited, if full attribution is given". However, it must be noted that any specification changes/updates are at the whim of its authors.

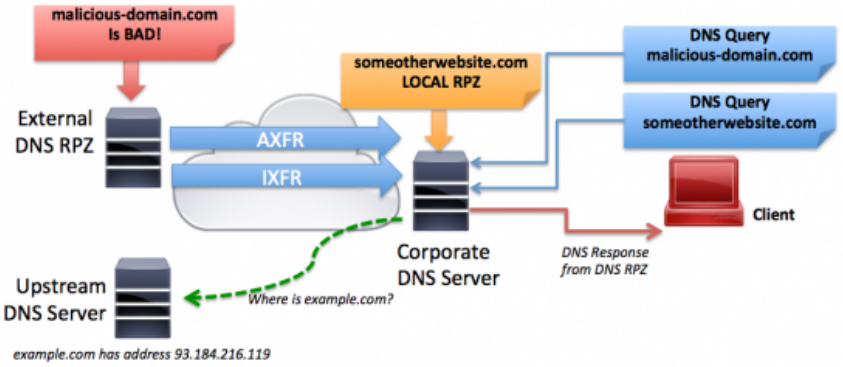

The RPZ specification defines the use of standard zone files whose RR definition invoke Policy Actions by using Policy Triggers in what one may call a Policy Rule Set (though this term is not used in the specification). RPZ is invoked (and its behavior controlled) in BIND9 using a response-policy statement (in named.conf) which is unique to BIND9 and is not defined within the RPZ specification - other implementations will use their own configuration styles and parameter sets. RPZ, by default, does not invoke policy processing on DNSSEC responses (though this can be modified with the break-dnssec parameter). For those familiar with the technology, it is similar to, but more complex than, DNS Black Lists (DNSBL) - a reputational anti-spam technique.

A very helpful configuration guide can be found here: http://www.zytrax.com

DNS Firewall dnsrpz.info

Providers of reputation data	Service	Services Supported
DissectCyber	rpzone.us	
FarsightSecurity	Newly Observed Domains and example	
InternetIdentity	DNS firewall	
SpamHaus	Several of their popular blocklists are available via RPZ. Article Pricing	
SURBL	Data Feed	
ThreatStop	DNS firewall and announcement	
SecurityZones	Provider	Provides product marketing and sales for some of the providers above
Deteque	Provider	Has provided integration consulting for some of the DNS RPZ providers above
OpenDNS		Integrated, Management Overlay, Managed Services

Comparison of DNS blacklists

RPZ - response-policy statement

The response-policy statement controls the behavior of RPZ policy processing

```
response-policy { zone zone-name
        [ policy (given|disabled|passthru|drop|nxdomain|nodata|tcp-only| cname domain-name)
        [ recursive-only yes_or_no ]
        [ max-policy-ttl number ] ; }
        [ max-policy-ttl number ]
        [ break-dnssec yes_or_no ]
        [ min-ns-dots number ]
        [ qname-wait-recurse yes_or_no ] ;
# example
response-policy {zone "dontlike" ; zone "likeless" policy passthru;} recursive-only yes;
```

```
Policy Triggers:
QNAME Trigger on query name.
CLIENT-IP Trigger on DNS client IP.
IP Trigger on query response IP.
NSDNAME Trigger on NS name during delegation.
NS-IP Trigger on NS IP during delegation.
```

```
Policy Actions:
NXDOMAIN Return name does not exist.
NODATA Return name exists but with no answer data.
PASSTHRU Do nothing - normally defines an exception in a range.
TCP-Only Force use of TCP. [not in Format 3]
DROP Causes client timeout. [not in format 3]
Local-Data Response data defined by RR and target-name/left-hand expression.
```

Policy Trigger:
Any Policy Trigger can be used with any Policy Action while the table shows only the most common types used with each Policy Action.

Policy Actions:
Policy Actions define the required outcome or result and are relatively straightforward. They are defined using the RR type and target-name (left-hand-name) of the RR as shown in the table on the next slide:

RPZ - configuration

named.conf

```
options { forward first;
          forwarders {
                208.67.222.222; // opendns.org
                208.67.220.220; // opendns.org
                8.8.8.8; //google-public-dns-a.google.com.
                8.8.4.4; //google-public-dns-b.google.com.
                };
        response-policy { zone "rpz.f1-online.net"; zone "rpz.spamhaus.org"; zone "rpz.surbl.org";
"rpz.ph.surbl.org"; }; };
```

1. response-policy option

2. local RPZ slave zone

3. remote RPZ slave zone's

```
zone "rpz.f1-online.net"    { type slave; file "rpz.f1-online.net.zone"; masters { 193.34.28.244; 193.34.29.244; }; check-names ignore; };
zone "rpz.spamhaus.org"     { type slave; file "dbl.rpz.spamhaus.org.zone"; masters { 199.168.90.51; 199.168.90.52; 199.168.90.53; };
zone "rpz.surbl.org"        { type slave; file "rpz.surbl.org.zone"; masters { 94.228.131.210; 94.228.131.211; }; check-names
zone "rpz.mw.surbl.org"     { type slave; file "rpz.mw.surbl.org.zone"; masters { 94.228.131.210; 94.228.131.211; };
zone "rpz.ph.surbl.org"     { type slave; file "rpz.ph.surbl.org.zone"; masters { 94.228.131.210; 94.228.131.211; };  check
```

rpz.f1-online.net

```
[…]
; return NXDOMAIN for facebook.com
www.facebook.com            666      CNAME .
*.facebook.com              666      CNAME .

; redirect to walled garden IP's
www.badguys.org    666      A        10.10.10.1
*.badguys.org               A        10.10.10.1
rpz.dns-as.org              A        10.0.2.21
wolfgang.cisco.com          A        193.34.28.108

; do not rewrite www.cisco.com (so, PASSTHRU) but add or override DNS-AS metadata
www.cisco.com               CNAME    rpz-passthru.
*.cisco.com                 CNAME    rpz-passthru.
www.cisco.com               TXT      "CISCO-CLS=app-name:HTTP|app-class:TD"
*.cisco.com                 TXT      "CISCO-CLS=app-name:HTTP|app-class:TD"

; rewrite A and add DNS-AS metadata
www.bradreese.com        A           72.163.4.161
www.bradreese.com        TXT         "CISCO-CLS=app-name:HTTP|app-class:SCV"
```

4. local RPZ master zone DNS-AS overrides

passthru for A/AAAA would be great but does not work, today. Working with ISC on this!

A + TXT works today

5. DNS-AS Operations

BIND and DNS

What Constitutes an Authoritative Source

The BIND software distribution has three parts:

- Domain Name Resolver
- Domain Name Authority server
- Tools

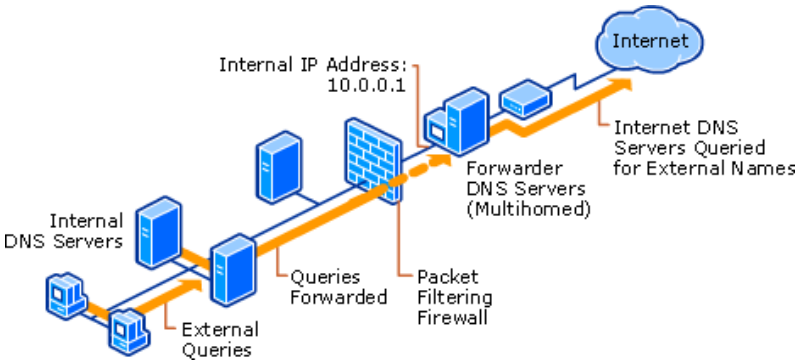

Domain Name Authority server

- An authoritative DNS server answers requests from resolvers, using information about the domain names it is authoritative for
- There can just be ONE ZONE being authoritative per domain

Domain Name Resolver

- A resolver is a program that resolves questions about names by sending those questions to appropriate servers and responding appropriately to the servers' replies.
- In the most common application, a web browser uses a local stub resolver library on the same computer to look up names in the DNS. That stub resolver is part of the operating system.
- The stub resolver usually will forward queries to a caching resolver, a server or group of servers on the network dedicated to DNS services. Those resolvers will send queries to one or multiple authoritative servers in order to find the IP address for that DNS name.

76

The DNS-AS Acronym Decoder Ring

Split-DNS

An enterprise typically has different authoritative servers for internal and external clients, and publish some zones on the internal servers only.

- ✓ Internal zones, managed from an Active Directory
- ✓ External zones, managed from a single ‘master’ BIND system (DMZ)
- ✓ Caching recursive resolvers for “external” domains (DMZ)

Response Policy Zones

RPZ is a BIND mechanism to selectively override foreign zones we are not authoritative for

DNS-AS-RR

A DNS TXT record inside a forward or reverse ZONE file

TXT "CISCO-CLS=app-name:HTTP|app-class:TD”

DNS-AS-client (Enterprise: client -> application server)

A client side Network Element running a DNS stub resolver for resolving DNS-AS-RR by using the client DNS request as a trigger for a forward lookup with a fallback to a reverse lookup

DNS-AS-client (Datacenter: application server -> client)

An application server side Network Element running a DNS stub resolver for resolving DNS-AS-RR by using the application IP as a trigger for a reverse lookup

DNS-AS-proxy

Inserts metadata (DNS-AS-RR) in case not being provided by a northbound DNS server or in the case we don’t trust a specific domain (malware, porn,…)

DNS-AS-edge

Internet facing Border Routers running two DNS-AS functions

- ✓ DNS-AS-client (even if running a DNS-AS-proxy on the same box) derives it's DNS-AS-RR from a southbound DNS Server (DMZ)
- ✓ DNS-AS-proxy (ensures that the southbound DNS servers (DMZ BIND) have meaningful DNS-AS-RR for external domains)

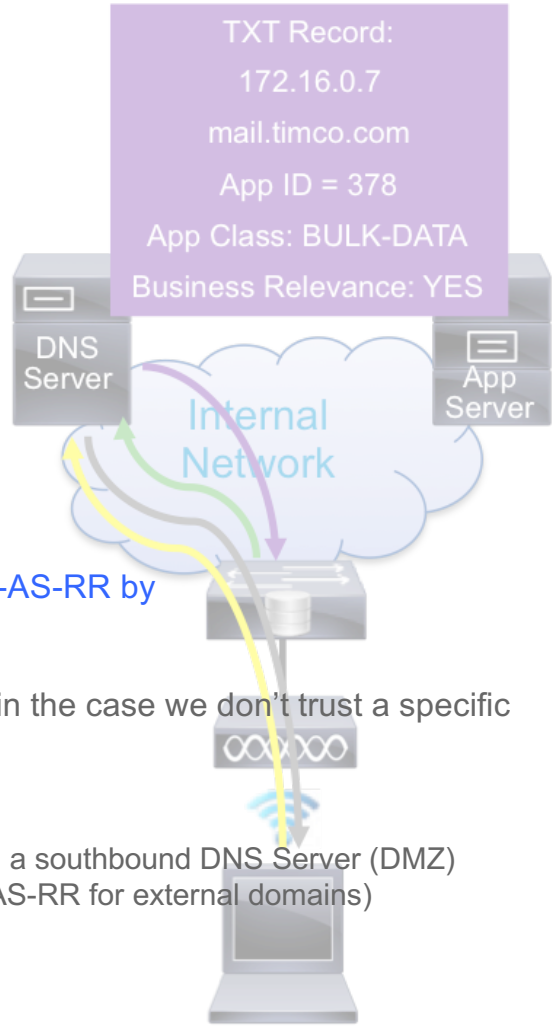

Enterprise DNS Deployment

Internal Namespace
INTRANET - Full Trust
MS Active Directory Integrated

Microsoft Active Directory

Active Directory
Domain
Windows 2012
Server

authoritative for internal resolvers
adc0.toocoolforyou.net (192.168.168.240)
adc1.toocoolforyou.net (192.168.167.244)
adc2.toocoolforyou.net (192.168.168.244)

DMZ Hybrid Namespace
EXTRANET DMZ - Medium Trust
My authoritative Named
Recursive Resolvers (RPZ)

ISC Internet Systems Consortium

ISC BIND
DNS Server
Fedora 24
BIND 9.10.4

Authoritative Hidden Master
dnssec-enable (TXT + AVC)
ns0.f1-online.net (193.34.28.240)
Authoritative Public Master
dnssec-enable (TXT + AVC)
ns1.f1-online.net (193.34.29.244)
ns2.f1-online.net (193.34.28.244)

Recursive Caching Resolvers (RPZ)
dnssec-validation (TXT + AVC + RPZ)
rr1.f1-online.net (193.34.29.241)
rr2.f1-online.net (193.34.28.241)

Public Namespace
INTERNET - No Trust
External authoritative Named
ROOT Server

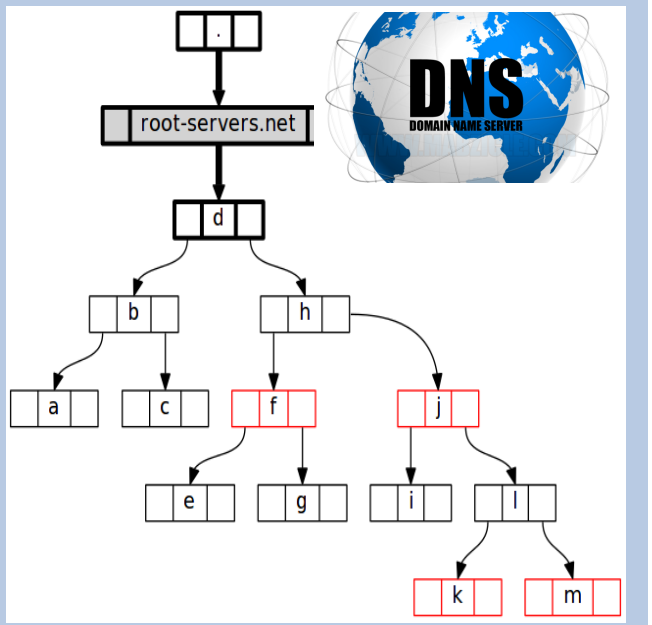

Recursive Lookup
Recursive Lookup
Recursive Lookup
Authoritative Lookup

In an Enterprise - DNS lives in multiple places

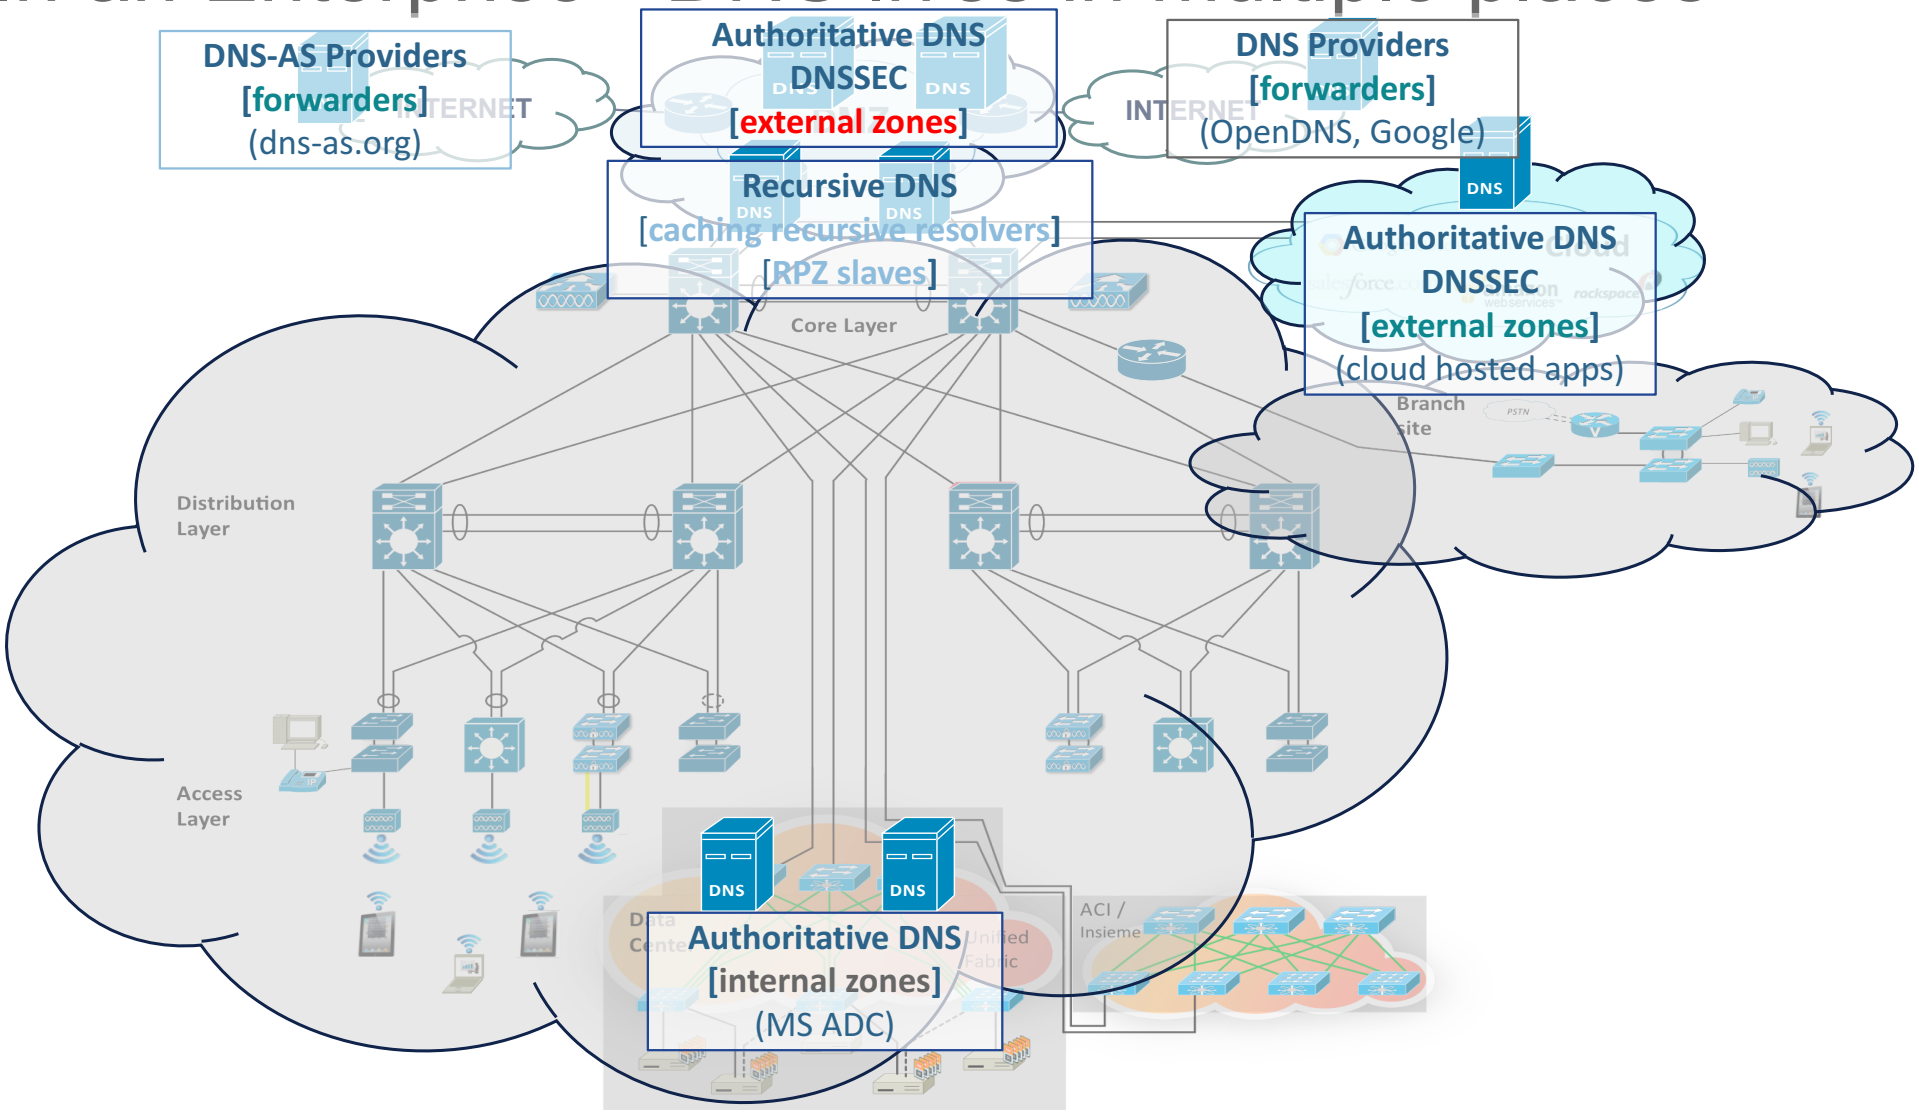

DNS-AS-Client - Operations

DNS-AS Client (APs, Switches, Routers)

C3PL Policy Enforcement
based on AVC Binding Table
SRC-IP: 192.168.160.10
DST-IP: 193.34.28.202
"CISCO-CLS=app-name:HTTP|app-class:TD"

| query
| query

e A | 193.34.28.202
e TXT | "CISCO-CLS=app-name:HTTP|app-class:TD"

DNS-AS
Client
192.168.254.100

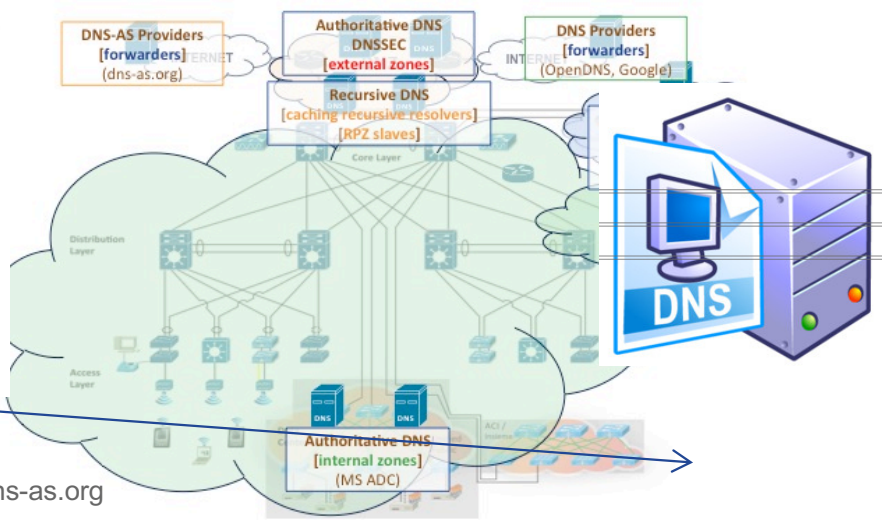

User

192.168.160.10

DNS
snooping

standard query | type A | www.dns-as.org

standard query response | type A | 193.34.28.202


6. Actually, what can we do with it?

DNS-AS Use Case Matrix

Everywhere you want to match on Metadata

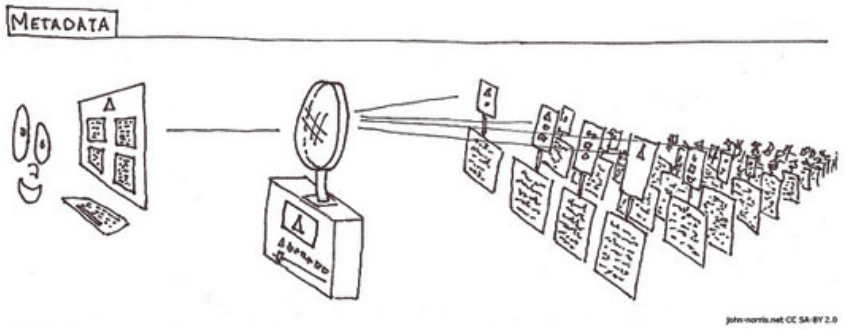

- Reporting via FNF even if encrypted
- Easy QoS
- Troubleshooting
- SPAN
- Martian ACL’s
- IPSLA
- Domain Based Routing
- ZBF (Zone Based Firewalls)
- NSH (Service Chaining)

DNS-AS Use Case Matrix

DNS-AS <metadata> as a variable to match within C3PL MQC

1) QoS

```
class-map match-all NETWORK-CONTROL
 match protocol attribute traffic-class network-control
 match protocol attribute business-relevance business-relevant
 match protocol <metadata>
```

2) Zone Based Firewalls

```
class-map type inspect match-all class-in-ssh
 match access-group name ACL-IPv4-ssh-in
 match protocol ssh
 match protocol <metadata>
```

3) Security ACL's

```
ip access-list extended ACL-IPv4-Minecraft-in
 remark ----- minecraft.f1-online.net ------
 permit tcp any host 193.34.29.143 eq 25565
 permit protocol <metadata>

ip access-list standard ACL-IPv4-NMS
 remark ----- NOC DMZ
 permit aaa.bb.ccc.ddd
 permit protocol <metadata>
 remark ---- deny everything else --------
 deny   any log
```

4) Object Group

```
object-group service port-proxy-server
 tcp eq 8080
 match protocol <metadata>
```

5) Domain Based Routing

```
track 104 match protocol <metadata>
ip route 192.168.168.0 255.255.255.0 192.168.252.114 111 track 104
```

Easy QoS Integration

DNS-AS Shortcuts for Cisco's (RFC 4594-Based) 12-Class QoS Model

APPLICATION CLASS (RFC 4594)	APPLICATION CLASS long DNS-AS-RR (LONG)	APPLICATION CLASS short DNS-AS-RR(SHORT)	BUSINESS-RELEVANCE DNS-AS-RR(SHORT)	DSCP	COS 802.1e	WMM	QUEUING & DROPPING	APPLICATION EXAMPLES
VoIP Telephony	app-class:VOIP-TELEPHONY	app-class:VO	business:yes	EF			Priority Queue (PQ)	Cisco IP Phones (G.711, G.729)
Broadcast Video	app-class:BROADCAST-VIDEO	app-class:BV	business:yes	CS5			(Optional) PQ	Cisco IP Video Surveillance / Cisco Enterprise TV
Real-Time Interactive	app-class:REALTIME-INTERACTIVE	app-class:RTI	business:yes	CS4			(Optional) PQ	Cisco TelePresence
Multimedia Conferencing	app-class:MULTIMEDIA-CONFERENCING	app-class:MMC	business:yes	AF4			BW Queue + DSCP WRED	Cisco Jabber, Cisco WebEx
Multimedia Streaming	app-class:MULTIMEDIA-STREAMING	app-class:MMS	business:yes	AF3			BW Queue + DSCP WRED	Cisco Digital Media System (VoDs)
Network Control	app-class:NETWORK-CONTROL	app-class:NC	business:yes	CS6			BW Queue	EIGRP, OSPF, BGP, ISIS, HSRP, IKE
Signaling	app-class:SIGNALING	app-class:CS	business:yes	CS3			BW Queue	SCCP, SIP, H.323
Ops / Admin / Mgmt	app-class:OPS-ADMIN-MGMT	app-class:OAM	business:yes	CS2			BW Queue	SNMP, SSH, Syslog
Transactional Data	app-class:TRANSACTIONAL-DATA	app-class:TD	business:yes	AF2			BW Queue + DSCP WRED	ERP Apps, CRM Apps, Database Apps
Bulk Data	app-class:BULK-DATA	app-class:BD	business:yes	AF1			BW Queue + DSCP WRED	E-mail, FTP, Backup Apps, Content Distribution
Best Effort	app-class:BEST-EFFORD	app-class:BE	business:default	DF	0		Default Queue + RED	Default Class
Scavenger	app-class:SCAVENGER	app-class:SCV	business:no	CS1	0		Min BW Queue (Deferential)	YouTube, Netflix, iTunes, BitTorrent, Xbox Live

91

Easy QoS Integration

```
class-map match-all VOICE
   match protocol attribute traffic-class voip-telephony
   match protocol attribute business-relevance business-relevant
class-map match-all BROADCAST-VIDEO
   match protocol attribute traffic-class broadcast-video
   match protocol attribute business-relevance business-relevant
class-map match-all INTERACTIVE-VIDEO
   match protocol attribute traffic-class real-time-interactive
   match protocol attribute business-relevance business-relevant
class-map match-all MULTIMEDIA-CONFERENCING
   match protocol attribute traffic-class multimedia-conferencing
   match protocol attribute business-relevance business-relevant
class-map match-all MULTIMEDIA-STREAMING
   match protocol attribute traffic-class multimedia-streaming
   match protocol attribute business-relevance business-relevant
 class-map match-all SIGNALING
   match protocol attribute traffic-class signaling
   match protocol attribute business-relevance business-relevant
class-map match-all NETWORK-CONTROL
   match protocol attribute traffic-class network-control
   match protocol attribute business-relevance business-relevant
class-map match-all NETWORK-MANAGEMENT
   match protocol attribute traffic-class ops-admin-mgmt
   match protocol attribute business-relevance business-relevant
class-map match-all TRANSACTIONAL-DATA
   match protocol attribute traffic-class transactional-data
   match protocol attribute business-relevance business-relevant
class-map match-all BULK-DATA
   match protocol attribute traffic-class bulk-data
   match protocol attribute business-relevance business-relevant
class-map match-all SCAVENGER
   match protocol attribute business-relevance business-irrelevant
```

"CISCO-CLS=app-name:WOLFGANG|app-class:NC"

magically allows "wolfgang.dns-as.org" to sneak underneath class-map NETWORK-CONTROL With ZERO configuration

```
policy-map MARKING
   class VOICE
     set dscp ef
   class BROADCAST-VIDEO
     set dscp cs5
   class INTERACTIVE-VIDEO
     set dscp cs4
   class MULTIMEDIA-CONFERENCING
     set dscp af41
   class MULTIMEDIA-STREAMING
     set dscp af31
   class SIGNALING
     set dscp cs3
   class NETWORK-CONTROL
     set dscp cs6
   class NETWORK-MANAGEMENT
     set dscp cs2
   class TRANSACTIONAL-DATA
     set dscp af21
   class BULK-DATA
     set dscp af11
   class SCAVENGER
     set dscp cs1
   class class-default
     set dscp default
```

DNS-AS Metadata:

```
www.dns-as.org          TXT "CISCO-CLS=app-name:HTTP|app-class:TD"
wolfgang.dns-as.org     TXT "CISCO-CLS=app-name:WOLFGANG|app-class:NC"
```

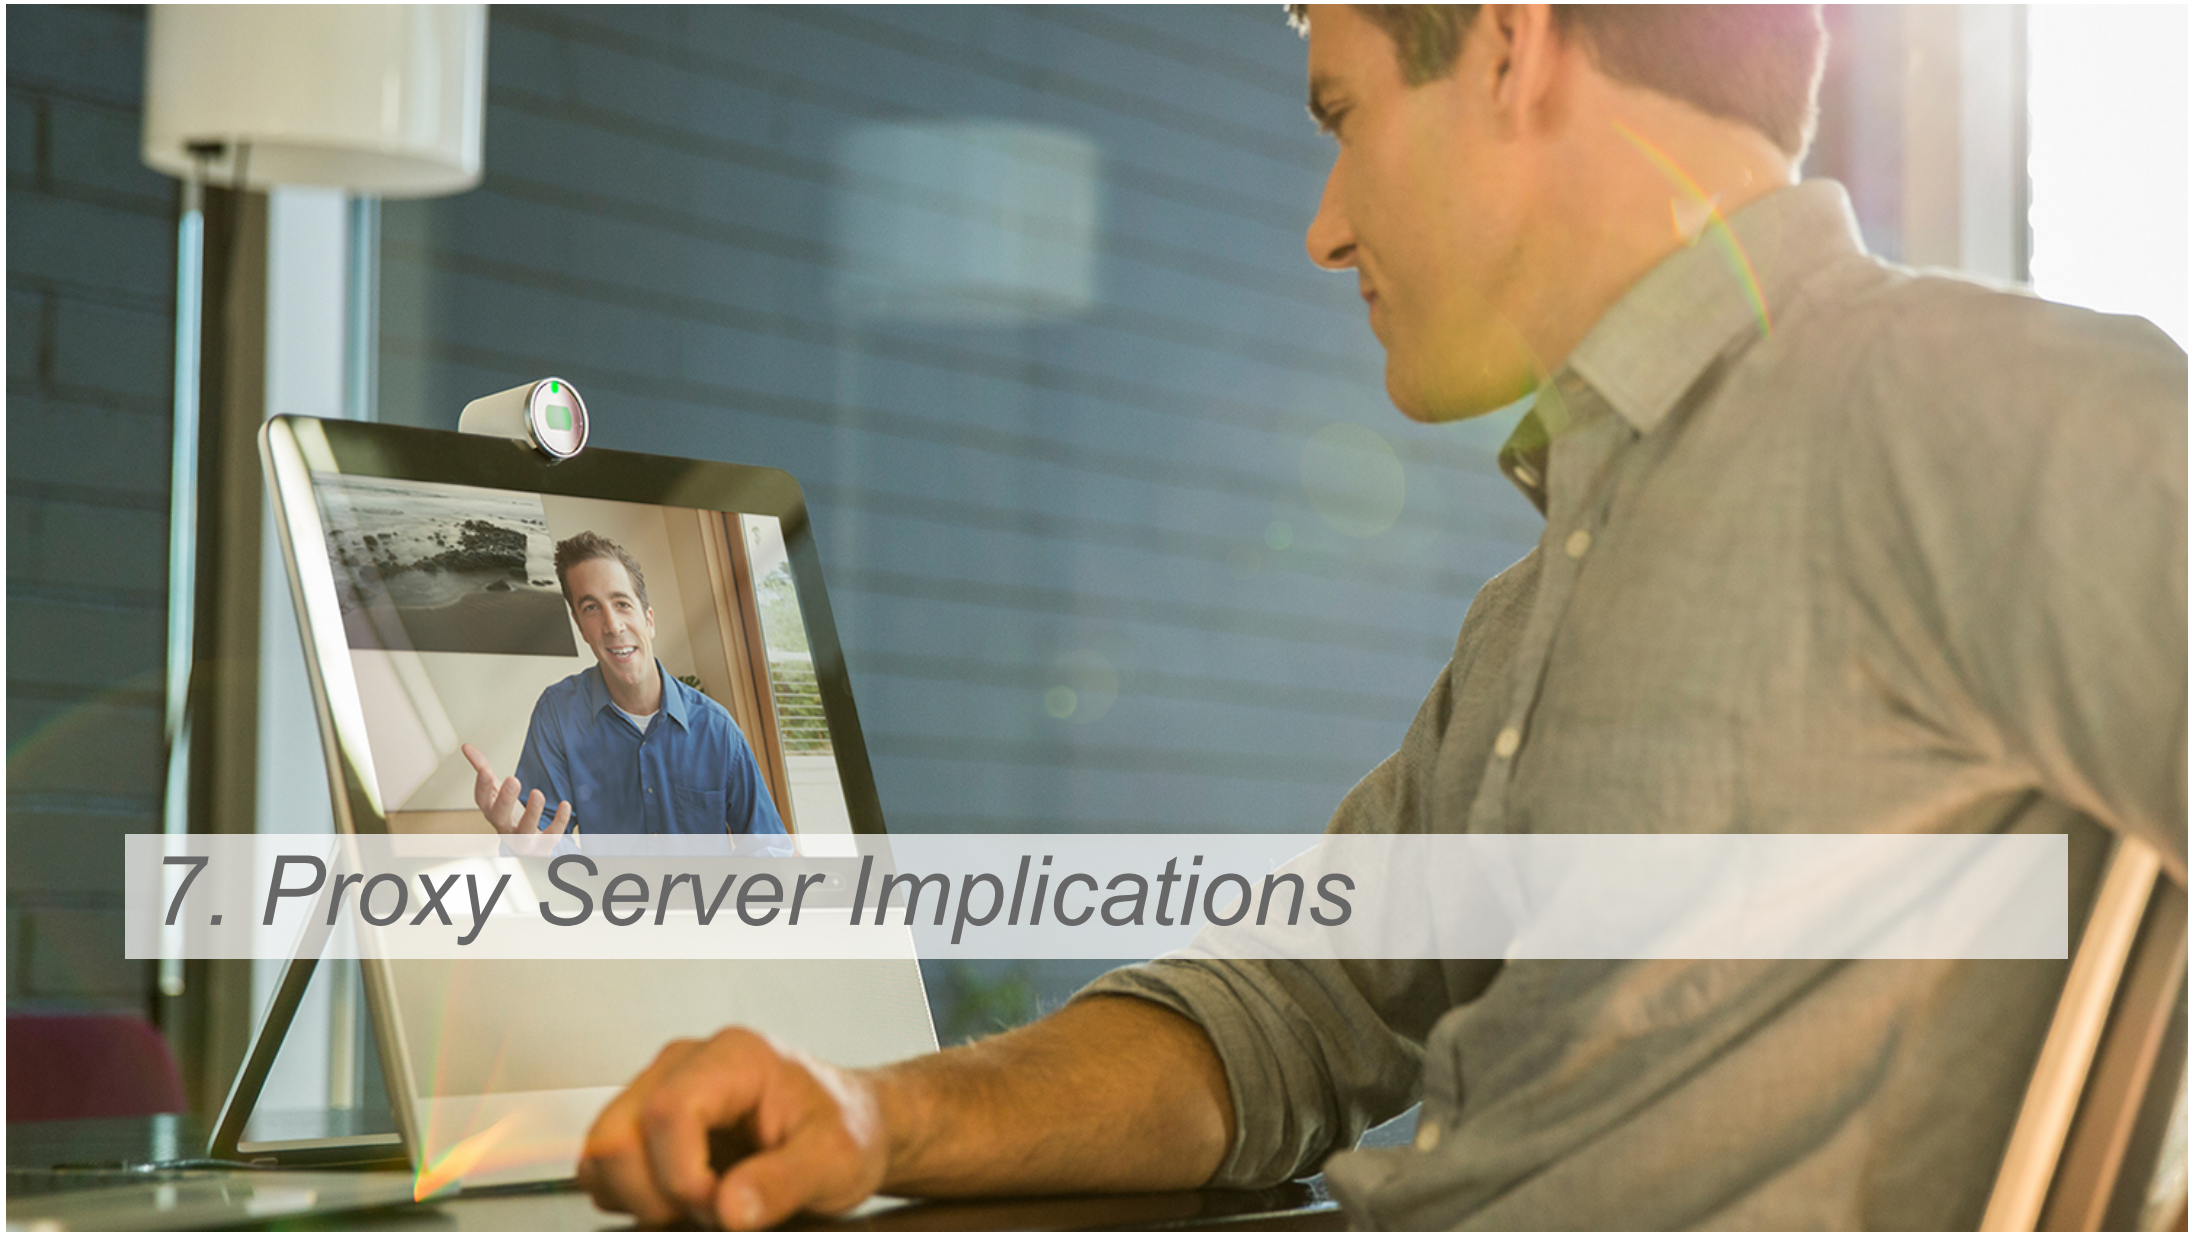

7. Proxy Server Implications

DNS-AS and Proxy Servers

Don’t tunnel If you care about quality of experience

```
// proxy.pac JavaScript
// Wolfgang Riedel wolfgang@cisco.com

localdomain  = "*.toocoolforyou.net";
dmz_odd_net   = "193.34.29.0";
dmz_odd_mask  = "255.255.255.0";
dmz_even_net  = "193.34.28.0";
dmz_even_mask = "255.255.255.0";

function FindProxyForURL(url,host)
{
    // If the hostname matches, send direct.
    if (shExpMatch(host, "*.toocoolforyou.net"))
        return "DIRECT";

   if (dnsDomainIs(host, "localhost") ||
        dnsDomainIs(host,localdomain))
   return "DIRECT";

   if (
        shExpMatch(host, "*.local") ||
        isPlainHostName(host) ||
        isInNet(dnsResolve(host), "10.0.0.0", "255.0.0.0") ||
        isInNet(dnsResolve(host), "172.16.0.0",  "255.240.0.0") ||
        isInNet(dnsResolve(host), "192.168.0.0",  "255.255.0.0") ||
        isInNet(dnsResolve(host), "127.0.0.0", "255.255.255.0") ||
        isPlainHostName(host) ||
        !isResolvable(host)
      )
   return "DIRECT";
```

What to classify?

- Outer (proxy traffic)
- Inner (application traffic)

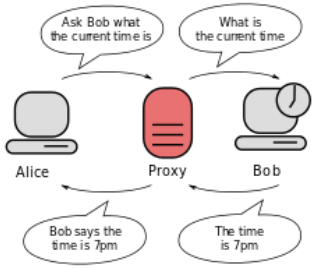

Typically exclude intranet traffic
Code Snippets: findproxyforurl.com

```
   if (url.substring(0, 5) == "http:") {
        return "PROXY proxy.f1-online.net:8080; DIRECT";
    }
   else if (url.substring(0, 4) == "ftp:") {
        return "PROXY proxy.f1-online.net:2121; DIRECT";
    }
   else if (url.substring(0, 6) == "https:") {
        return "DIRECT";
    }
   else if (url.substring(0, 7) == "gopher:") {
        return "DIRECT";
    }
   else {
        return "DIRECT";
   }
}
```

9. DNS-AS – Switches (no NBAR)

Catalyst 4k / Catalyst 2k

DNS-AS Classification & Marking Policy Example (Part 1 of 3)

```
!
class-map match-all VOICE
 match protocol attribute traffic-class voip-telephony
 match protocol attribute business-relevance business-relevant
class-map match-all BROADCAST-VIDEO
 match protocol attribute traffic-class broadcast-video
 match protocol attribute business-relevance business-relevant
class-map match-all REAL-TIME-INTERACTIVE
 match protocol attribute traffic-class real-time-interactive
 match protocol attribute business-relevance business-relevant
class-map match-all MULTIMEDIA-CONFERENCING
 match protocol attribute traffic-class multimedia-conferencing
 match protocol attribute business-relevance business-relevant
class-map match-all MULTIMEDIA-STREAMING
 match protocol attribute traffic-class multimedia-streaming
 match protocol attribute business-relevance business-relevant
class-map match-all SIGNALING
 match protocol attribute traffic-class signaling
 match protocol attribute business-relevance business-relevant
cla
 ma
 ma
cla
 ma
 ma
cla
 ma
 ma
cla
 ma
 match protocol attribute business-relevance business-relevant
class-map match-all SCAVENGER
 match protocol attribute business-relevance business-irrelevant
!
```

```
!
policy-map INGRESS-MARKING
 class-map match-all AUTOQOS_VOIP_VIDEO
  match cos  4
 class-map match-all AUTOQOS_VOIP_VOICE
  match cos  5
 class-map match-all AUTOQOS_VOIP_SIG
  match cos  3
!
```

Same ‘holy grail’ classification policy as on other router/switch platforms

Same ‘holy grail’ marking policy as on other router/switch platforms

Small extension of the trust boundary for voice and video

```
!
policy-map INGRESS-MARKING
 class VOICE
  set dscp ef
 class BROADCAST-VIDEO
  set dscp cs5
 class REAL-TIME-INTERACTIVE
  set dscp cs4
 class MULTIMEDIA-CONFERENCING
  set dscp af41
 class MULTIMEDIA-STREAMING
  set dscp af31
 class SIGNALING
  set dscp cs3
 class NETWORK-CONTROL
  set dscp cs6
 class NETWORK-MANAGEMENT
  set dscp cs2
 class TRANSACTIONAL-DATA
  set dscp af21
 class BULK-DATA
  set dscp af11
 class SCAVENGER
  set dscp cs1
  class class-default
  set dscp default
!
```

Catalyst 4k / Catalyst 2k

DNS-AS Classification & Marking Policy Example (Part 2 of 3)

```
!
interface GigabitEthernet2/14
 description IP-Phone
 switchport access vlan 165
 switchport mode access
 switchport voice vlan 111
 switchport port-security maximum 3
 switchport port-security violation restrict
 switchport port-security aging time 2
 switchport port-security aging type inactivit
 switchport port-security
 load-interval 30
 power inline police
 power efficient-ethernet auto
 auto qos voip cisco-phone
 storm-control broadcast level 10.00
 storm-control action trap
 qos trust device cisco-phone
 spanning-tree portfast edge
 spanning-tree bpduguard enable
 service-policy input INGRESS-MARKING
 service-policy output EGRESS-QUEUEING-1P7Q1T
!
```

In case trust boundary is extended to cisco-phone

Allow DSCP marking through the ingress policy-map

```
!
policy-map INGRESS-MARKING
 class AUTOQOS_VOIP_VOICE
  set dscp ef
  police cir 128000 bc 8000 conform-action
transmit  exceed-action set-dscp-transmit cs1
violate-action set-cos-transmit 1
 class AUTOQOS_VOIP_VIDEO
  set dscp af41
  police cir 10000000 bc 8000 conform-action
transmit  exceed-action set-dscp-transmit cs1
violate-action set-cos-transmit 1
 class AUTOQOS_VOIP_SIG
  set dscp cs3
  police cir 32000 bc 8000 conform-action
transmit  exceed-action set-dscp-transmit cs1
violate-action set-cos-transmit 1
!
```

Catalyst 4k / Catalyst 2k

DNS-AS Classification & Marking Policy Example (Part 3 of 3)

Configures basic DNS info

```
!
ip domain round-robin
ip domain-list toocoolforyou.net
ip domain-lookup source-interface Loopback0
ip domain-name toocoolforyou.net
ip name-server 192.168.167.244
ip name-server 192.168.168.244
!
```

DNS-AS snooping capability enabled by service-policy input

```
!
interface range TenGigabitEthernet2/1-40
 service-policy input INGRESS-MARKING
 service-policy output EGRESS-QUEUEING-1P7Q1T
!
```

```
!
avc dns-as client enable
!
avc dns-as client trusted-domains
domain ^.*f1.*$
domain ^.*cisco.*$
domain *.toocoolforyou.net
domain *.dns-as.org
domain *.nbar2web.org
domain *.f1v4.net
domain *.f1v6.net
!
```

Enables DNS-AS client

Whitelisted domains for which metadata may be queried and used for policy-purposes

10. DNS-AS – Routers (with NBAR)

ASR1k / ISR4k / CSR1kv

DNS-AS Classification & Marking Policy Example (Part 1 of 2)

```
!
class-map match-all VOICE
 match protocol attribute traffic-class voip-telephony
 match protocol attribute business-relevance business-relevant
class-map match-all BROADCAST-VIDEO
 match protocol attribute traffic-class broadcast-video
 match protocol attribute business-relevance business-relevant
class-map match-all REAL-TIME-INTERACTIVE
 match protocol attribute traffic-class real-time-interactive
 match protocol attribute business-relevance business-relevant
class-map match-all MULTIMEDIA-CONFERENCING
 match protocol attribute traffic-class multimedia-conferencing
 match protocol attribute business-relevance business-relevant
class-map match-all MULTIMEDIA-STREAMING
 match protocol attribute traffic-class multimedia-streaming
 match protocol attribute business-relevance business-relevant
class-map match-all SIGNALING
 match protocol attribute traffic-class signaling
 match protocol attribute business-relevance business-relevant
class-map match-all NETWORK-CONTROL
 match protocol attribute traffic-class network-control
 match protocol attribute business-relevance business-relevant
class-map match-all NETWORK-MANAGEMENT
 match protocol attribute traffic-class ops-admin-mgmt
 match protocol attribute business-relevance business-relevant
class-map match-all TRANSACTIONAL-DATA
 match protocol attribute traffic-class transactional-data
 match protocol attribute business-relevance business-relevant
class-map match-all BULK-DATA
 match protocol attribute traffic-class bulk-data
 match protocol attribute business-relevance business-relevant
class-map match-all SCAVENGER
 match protocol attribute business-relevance business-irrelevant
!
```

Same ‘holy grail’ classification policy as on other router/switch platforms

Same ‘holy grail’ marking policy as on other router/switch platforms

```
!
policy-map INGRESS-MARKING
 class VOICE
  set dscp ef
 class BROADCAST-VIDEO
  set dscp cs5
 class REAL-TIME-INTERACTIVE
  set dscp cs4
 class MULTIMEDIA-CONFERENCING
  set dscp af41
 class MULTIMEDIA-STREAMING
  set dscp af31
 class SIGNALING
  set dscp cs3
 class NETWORK-CONTROL
  set dscp cs6
 class NETWORK-MANAGEMENT
  set dscp cs2
 class TRANSACTIONAL-DATA
  set dscp af21
 class BULK-DATA
  set dscp af11
 class SCAVENGER
  set dscp cs1
  class class-default
  set dscp default
!
```

ASR1k / ISR4k / CSR1kv

DNS-AS Classification & Marking Policy Example (Part 2 of 2)

Configures basic DNS info

```
!
ip domain round-robin
ip domain-list toocoolforyou.net
ip domain-lookup source-interface Loopback0
ip domain-name toocoolforyou.net
ip name-server 192.168.167.244
ip name-server 192.168.168.244
!
```

DNS-AS snooping combined with NBAR

```
interface GigaEthernet0/0/0
 ip nbar protocol-discovery
 service-policy input ingress-MARKING
 service-policy output egress-hqos-95000
```

```
!
avc dns-as client enable
!
avc dns-as client trusted-domains
domain ^.*f1.*$
domain ^.*cisco.*$
domain *.toocoolforyou.net
domain *.dns-as.org
domain *.nbar2web.org
domain *.f1v4.net
domain *.f1v6.net
!
```

Enables DNS-AS client

Whitelisted domains for which metadata may be queried and used for policy-purposes

DNS-AS snooping without NBAR

```
interface GigabitEthernet0/0/0
 avc dns-as learning
 service-policy input ingress-MARKING
 service-policy output egress-hqos-95000
```

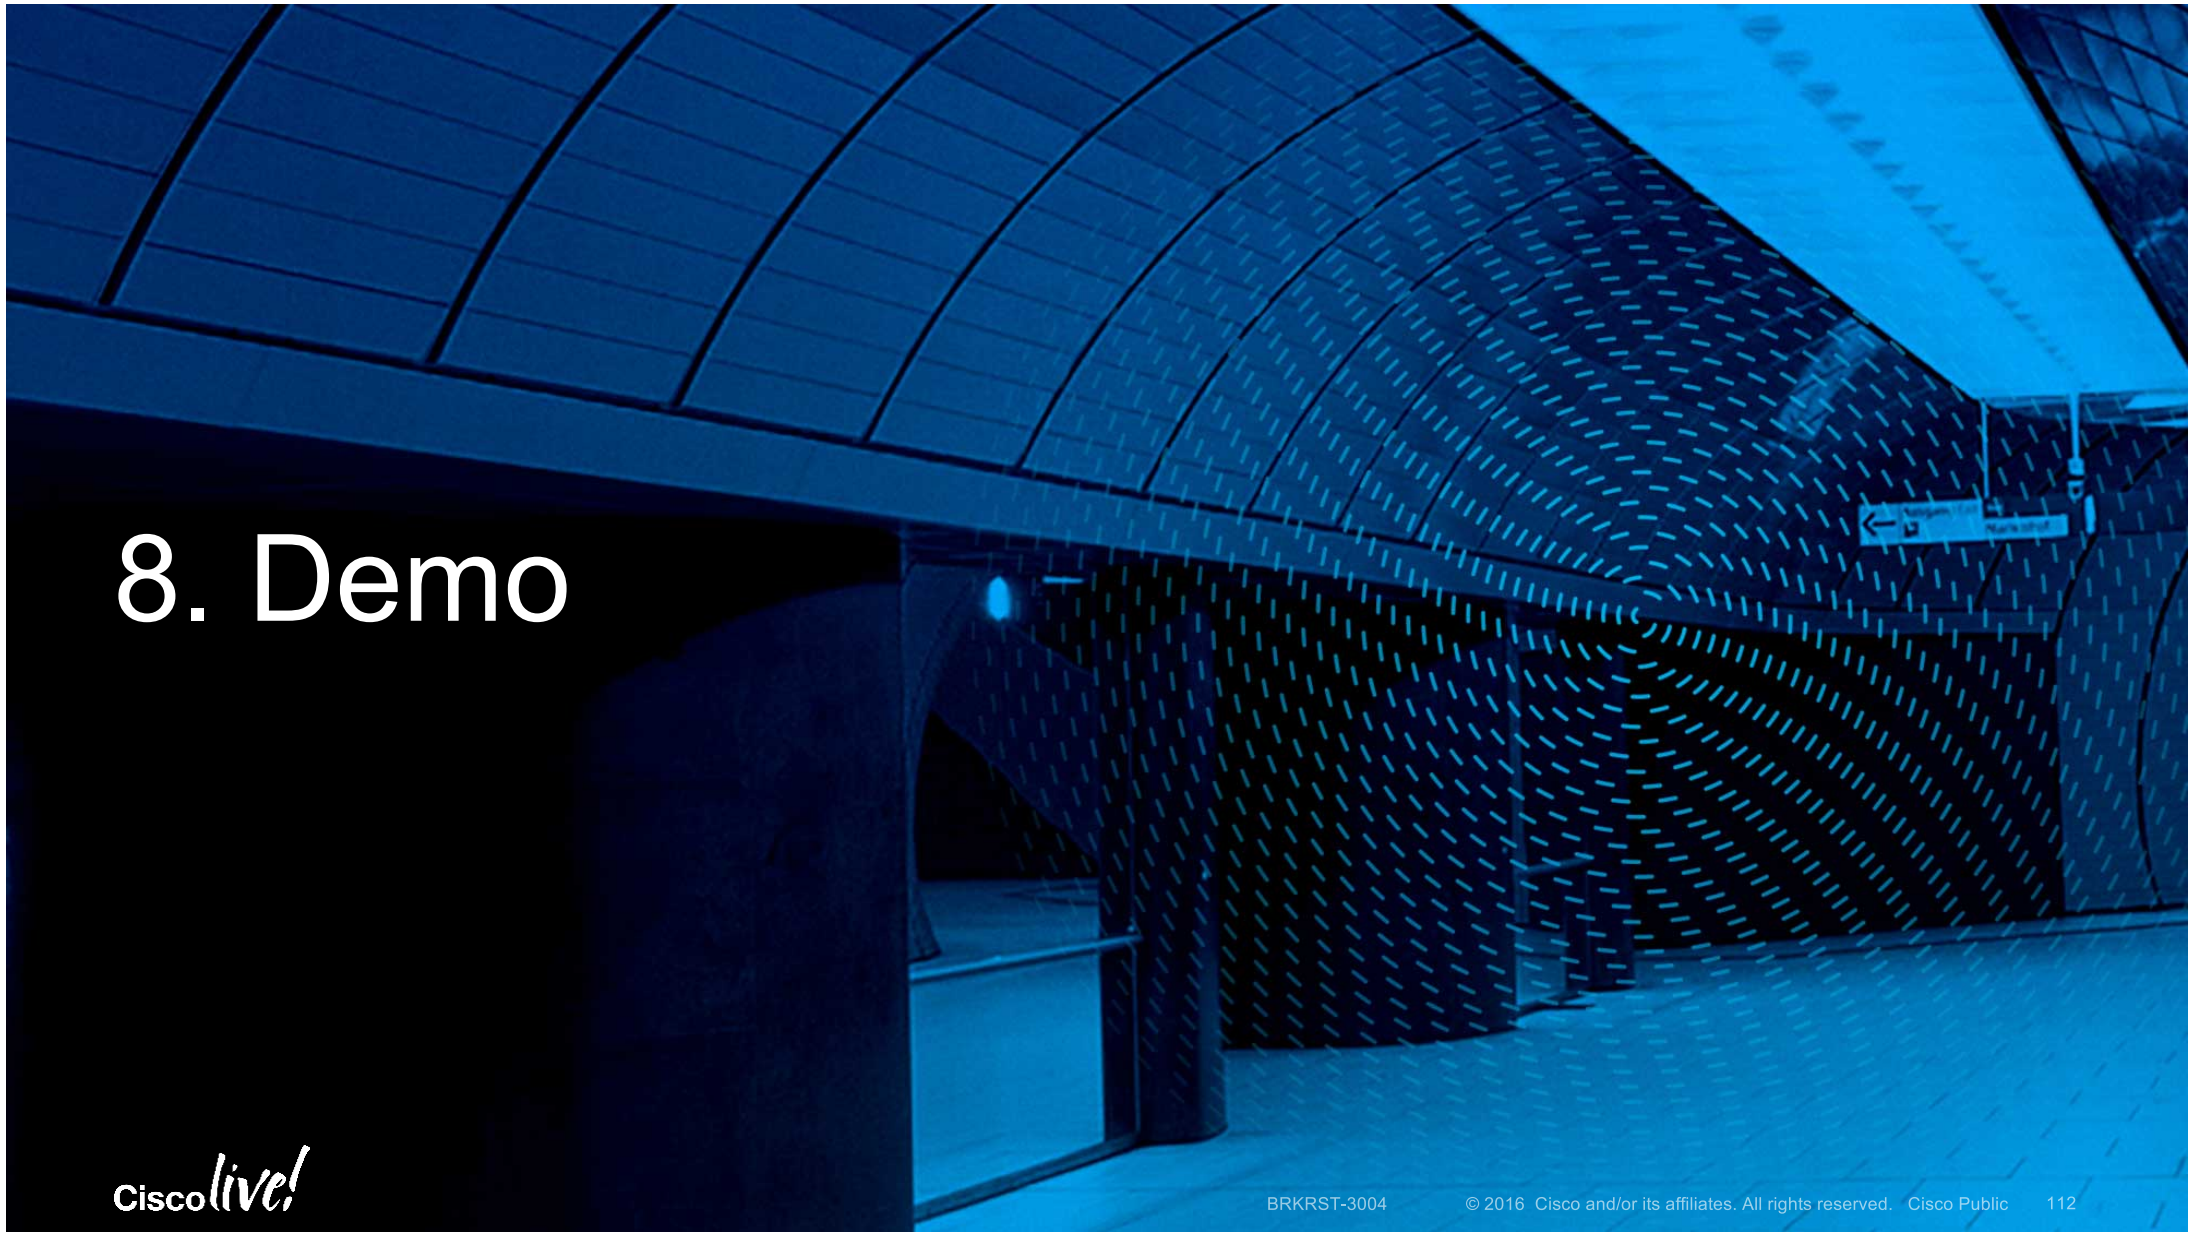

8. Demo

DNS-AS Visualization

DNS-AS Binding table into Prime Infrastructure and LiveAction

```
stealth-odd#show avc dns-as client binding-table

 -------------------------------------------------------------------------------------------------------------------------------------------------
 |                  |           |                |                          |       |                                                 |      | Time to |
 | Protocol name    | Vrf       | Ip List        | Host                     | Age   | Text record                                     | TTL  | Expire  |
 |                  |           |                |                          | [min] |                                                 | [min]| [min]   |
 -------------------------------------------------------------------------------------------------------------------------------------------------
 | DNS-RR2          |<default>  |193.34.28.241   |rr2.f1-online.net         |4136   |app-name:DNS-RR2|app-class:NC|business:yes       |2879  |919      |
 | WWW0-PROXY2      |<default>  |193.34.28.245   |proxy2.f1-online.net      |4129   |app-name:WWW0-PROXY2|app-class:TD|business:yes   |2874  |<1       |
 | WWW0             |<default>  |193.34.29.161   |www.dns-as.org            |1767   |app-name:WWW0|app-class:TD                       |2879  |1112     |
 | DNS-RR1          |<default>  |193.34.29.241   |rr1.f1-online.net         |1235   |app-name:DNS-RR1|app-class:NC|business:yes       |2187  |950      |
 | N-BORDER         |<default>  |193.34.28.50    |border.dns-as.org         |733    |app-name:N-BORDER|app-class:TD|business:yes      |2879  |2145     |
 | N-CONNECT        |<default>  |193.34.29.50    |connect.dns-as.org        |511    |app-name:N-CONNECT|app-class:TD|business:yes     |2879  |2367     |
 -------------------------------------------------------------------------------------------------------------------------------------------------
```

```
stealth-even#show avc dns-as client binding-table
 -------------------------------------------------------------------------------------------------------------------------------------------------
 |                  |           |                |                          |       |                                                 |      | Time to |
 | Protocol name    | Vrf       | Ip List        | Host                     | Age   | Text record                                     | TTL  | Expire  |
 |                  |           |                |                          | [min] |                                                 | [min]| [min]   |
 -------------------------------------------------------------------------------------------------------------------------------------------------
 | WWW0-PROXY2      |<default>  |193.34.28.245   |proxy2.f1-online.net      |4035   |app-name:WWW0-PROXY2|app-class:TD|business:yes   |1561  |<1       |
 | WWW0             |<default>  |193.34.28.47    |www.dns-as.org            |3560   |app-name:WWW0|app-class:TD|business:yes          |400   |37       |
 | VPN-GW-odd       |<default>  |193.34.31.242   |vpn-gw-odd.f1-online.net  |3542   |app-name:VPN-GW-odd|app-class:BD|business:yes    |1297  |723      |
 | N-BORDER         |<default>  |193.34.28.153   |border.dns-as.org         |868    |app-name:N-BORDER|app-class:TD|business:yes      |802   |764      |
 | MX00             |<default>  |193.34.29.140,  |mail.dns-as.org           |430    |app-name:MX00|app-class:BD|business:yes          |2880  |2437     |
 |                  |           |193.34.28.140   |                          |       |                                                 |      |         |
 -------------------------------------------------------------------------------------------------------------------------------------------------
```

DNS-AS & PI Visualization per https app

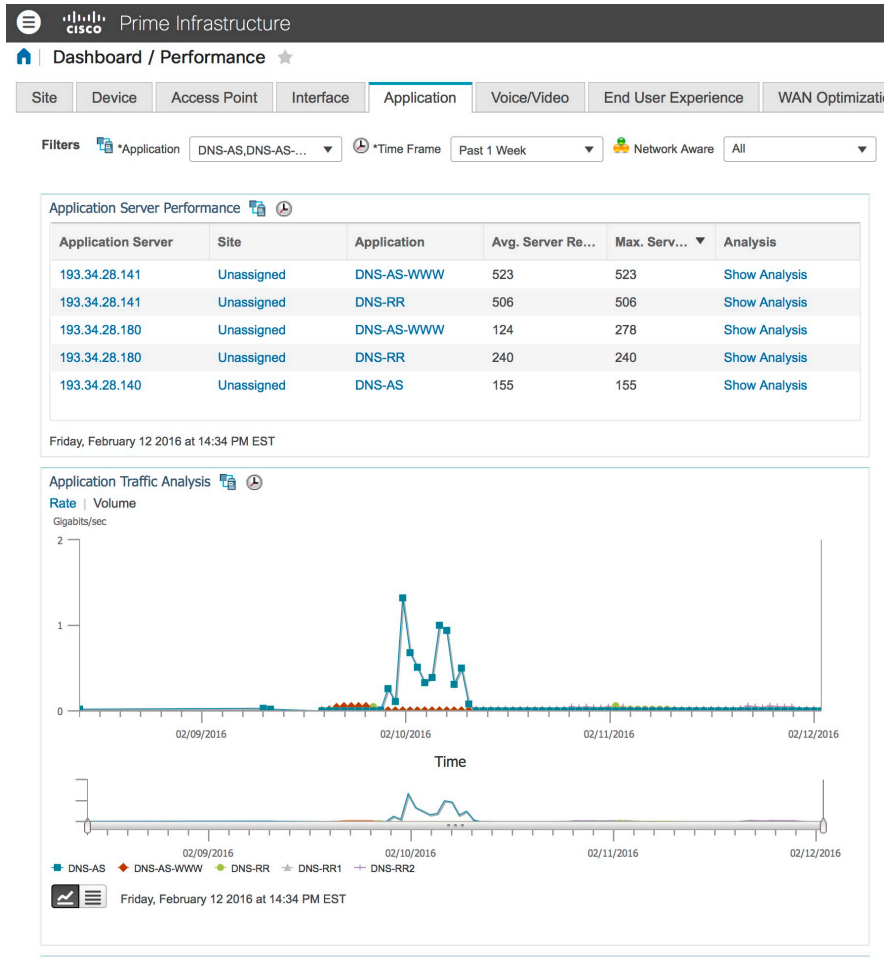

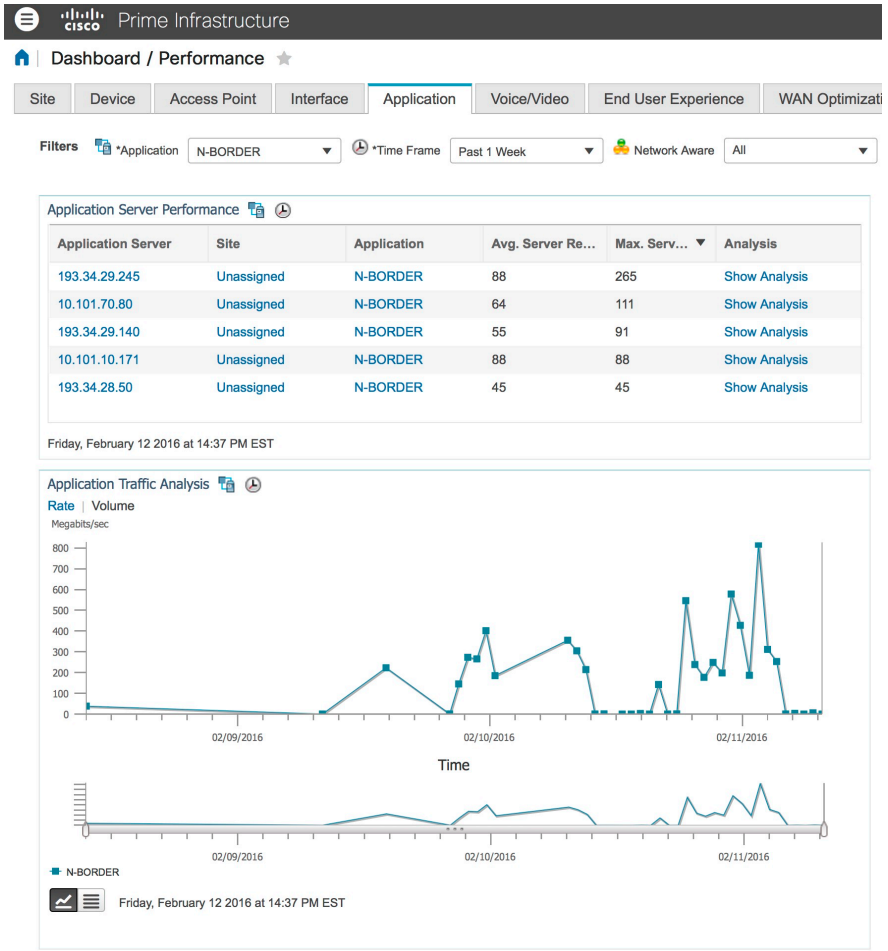

DNS-AS & LiveAction Visualization per https app

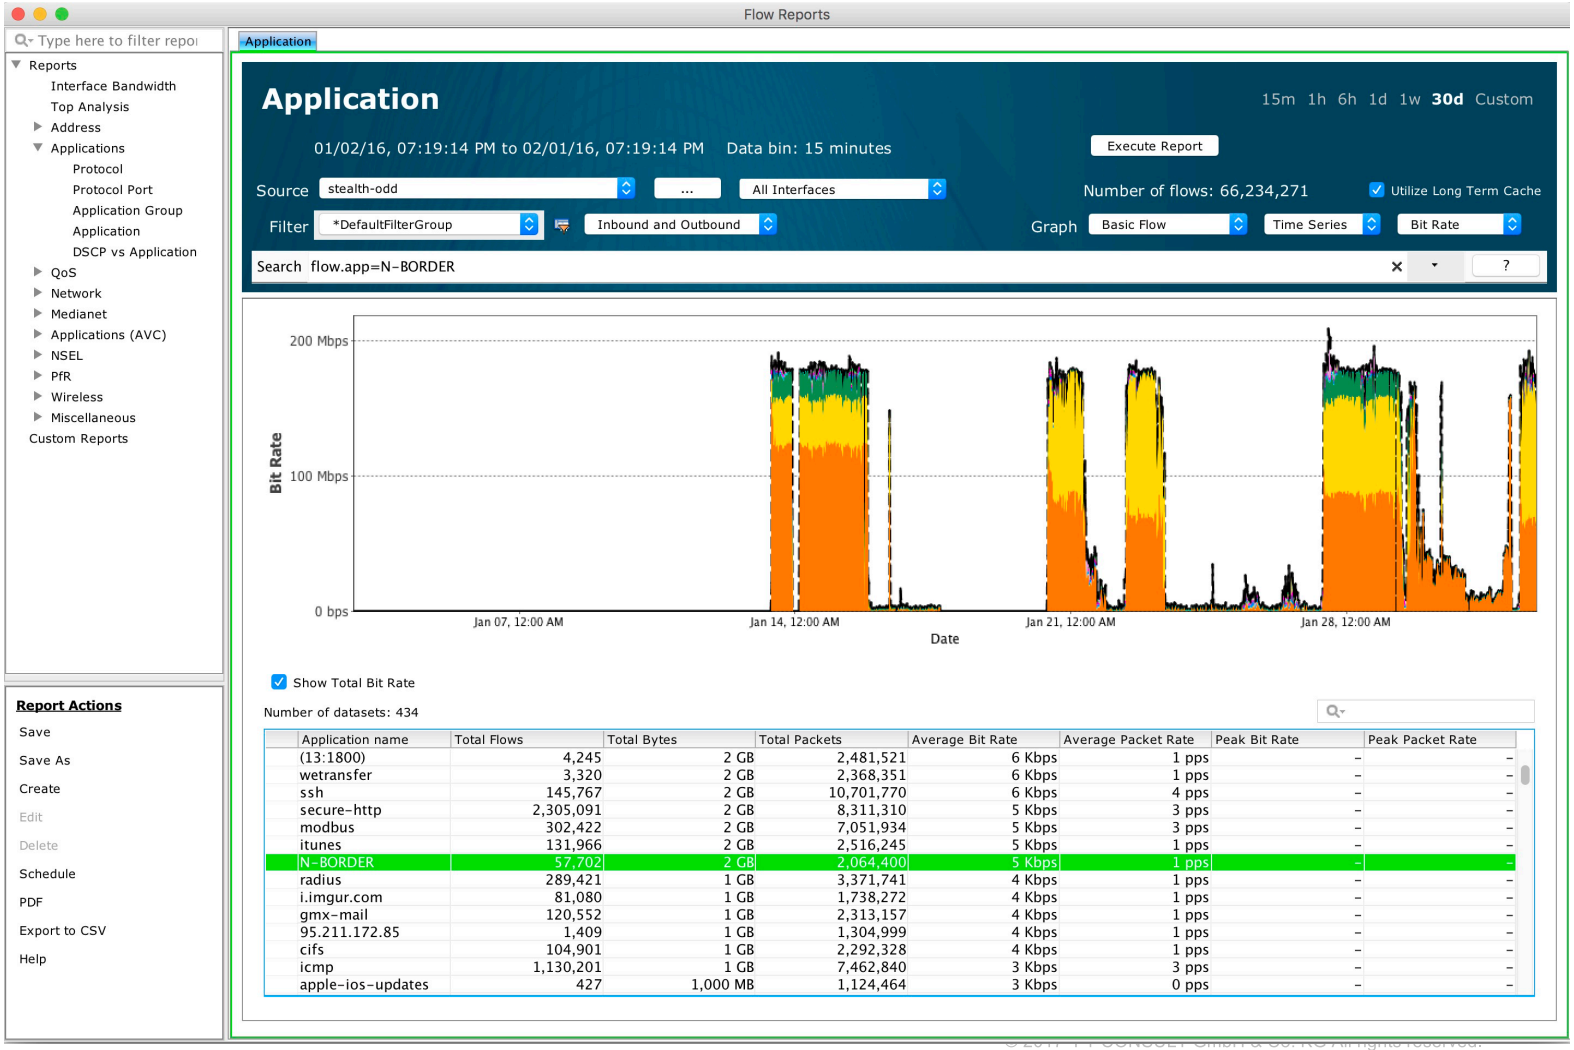

11. Conclusion and Open Discussion

We have come a Mile… but still a Way to Go!

Stages in the Application Assurance Lifecycle

Blindfolded ☹

Some Light…

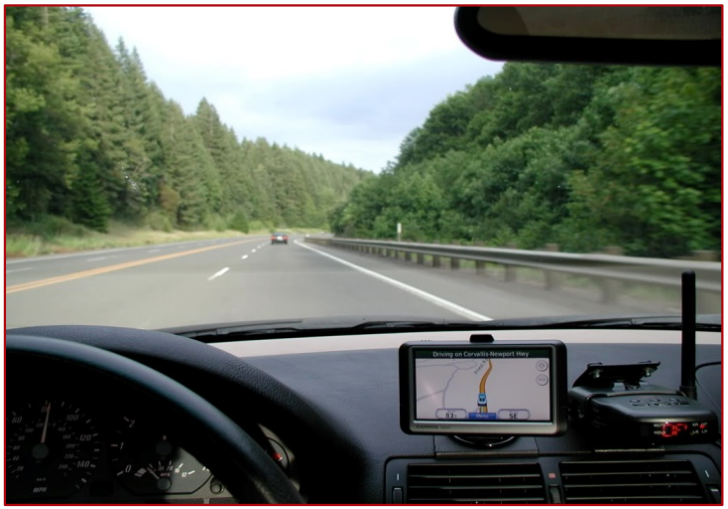

Clear View ☺

DNS - Summary

DNS, as it’s today already gives us a bunch of options

- Don’t fix what’s not fundamentally broken, don’t develop a new protocol and controller for every new use case, utilize what we already use today
- We can assume that DNS really scales well, right ;-)
- Incremental steps
- RPZ allows us to fix others shortcomings (forward and reverse)
- How about DNS Security?
  - OK, don’t let me get started on that one ;-)
  - Follow Best Practice's
  - If DNS is screwed we have a much bigger problem
  - VRF’s
  - Autonomic Networking (self-managed PKI + ACP)
  - DNSSEC
  - MACSEC
  - BIND-CHROOT, SE-linux
  - Split DNS: MS AD, DMZ RR’s, DMZ AS
  - Did I already mention, follow Best Practice’s

Summary - Why DNS-AS ?

- Done with SDN?
- Tired of Dealing with Snowflake Network Complexity?
- More info: http://dns-as.org
- Why would I want to make a best guess if I can know?
- As more CPU cycles you could free up by using DNS-AS more you have left for running DPI
- DPI will have a hard time working with encrypted traffic
- DPI can never work at wire rate and as more throughput we need as less feasible DPI methods become
- DPI as all other current methods just work if you have direct admin control over the box
- Emerging protocols like SPDY, HTTP/2, QUIC makes it impossible to have a clear AVC view
- DNS-AS is single point of administration without the need for having admin control over the network's in between.
- It's all about METADATA

Questions:

✓ **Do you like the idea?**

✓ **Are you interested to help?**

✓ **IETF related work?**

✓ **IETF DNS-AS AVC RDATA implementation guide**

✓ **DNS-AS Proxy?**

✓ **LINUX nftables implementation?**

wolfgang@dns-as.org


F1
Our Network for your Network