Netflow Collection and Analysis at Tier 1 Internet Peering

FloCon 2017

San Diego, CA January, 2017

Fred Stringer

System Engineer / Architect
AT&T CSO (Chief Security Organization)
TIARE (Threat Intelligence Analysis Response Engineering)

Rethink Possible®

Agenda

- What is Tier 1 Internet Peering
- Threat Analytics Platform
- Automated Data Analysis
- Netflow Collection and processing
- Keeping up:
  - Processing
  - Traceable Records
  - Interaction with Graphical Data Records
- Service Network Evolution
- On-Demand Surveillance

Internet Peering

Peering is a business relationship supported by routing (BGP) policies and procedures creating a network relationship. Can be created at meeting points (public peering) or direct connections (private peering)

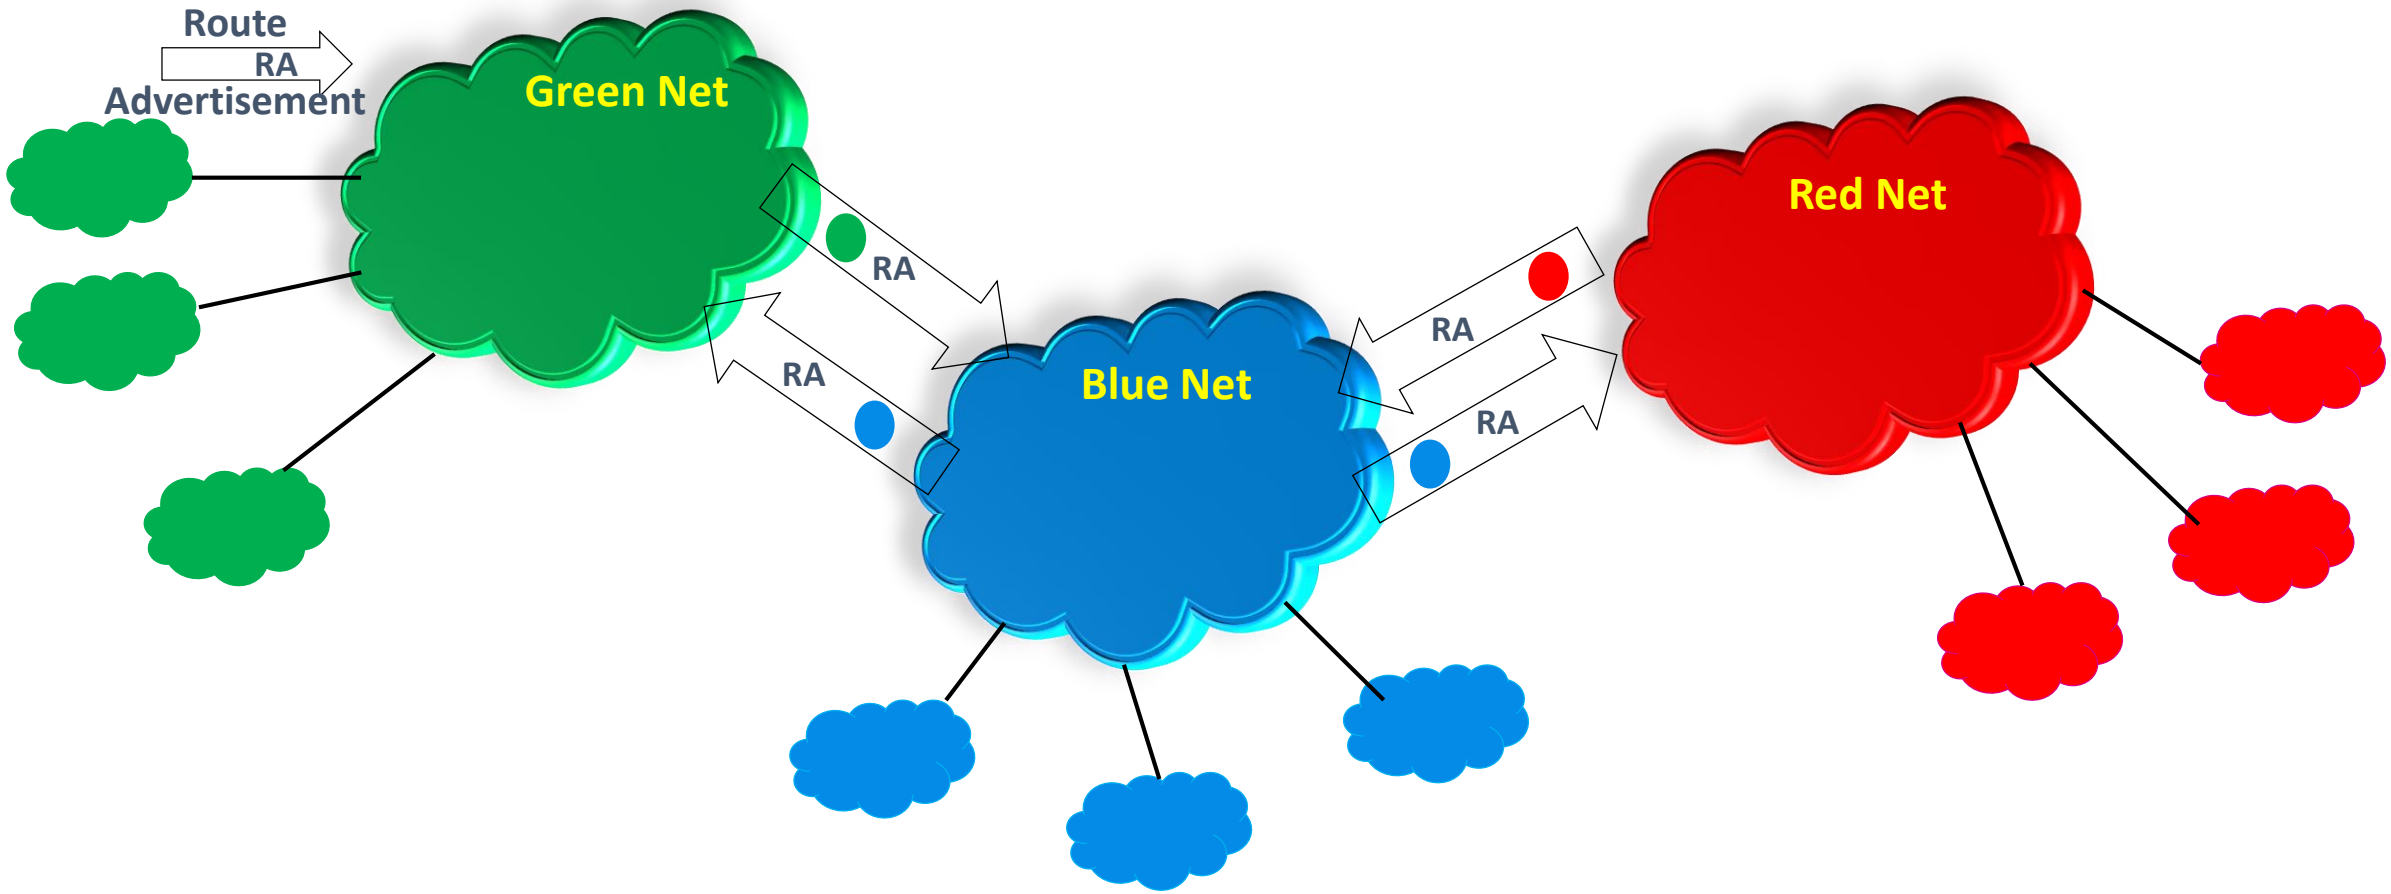

Internet Peering

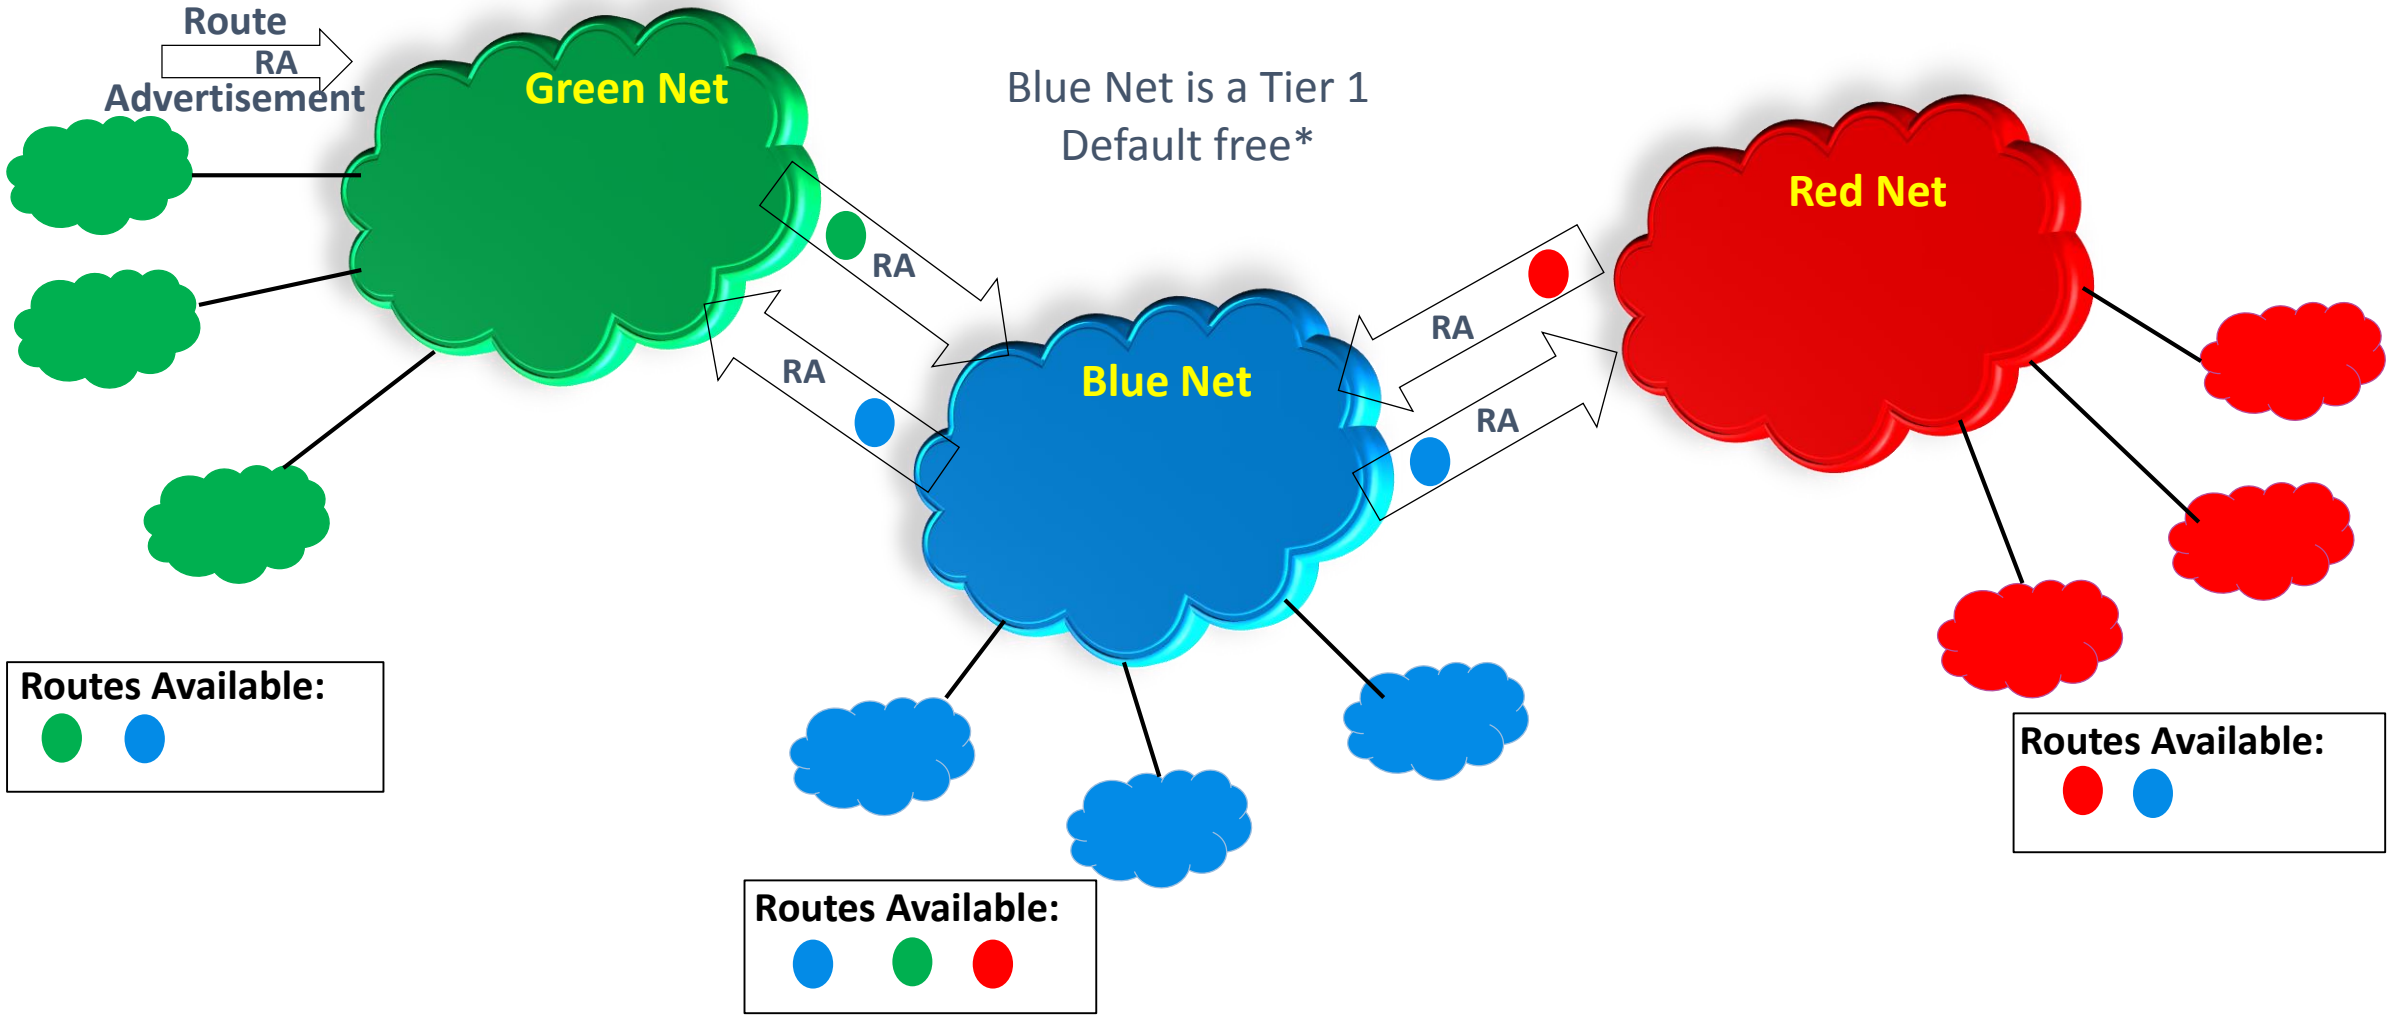

Tier 1 Internet Peering

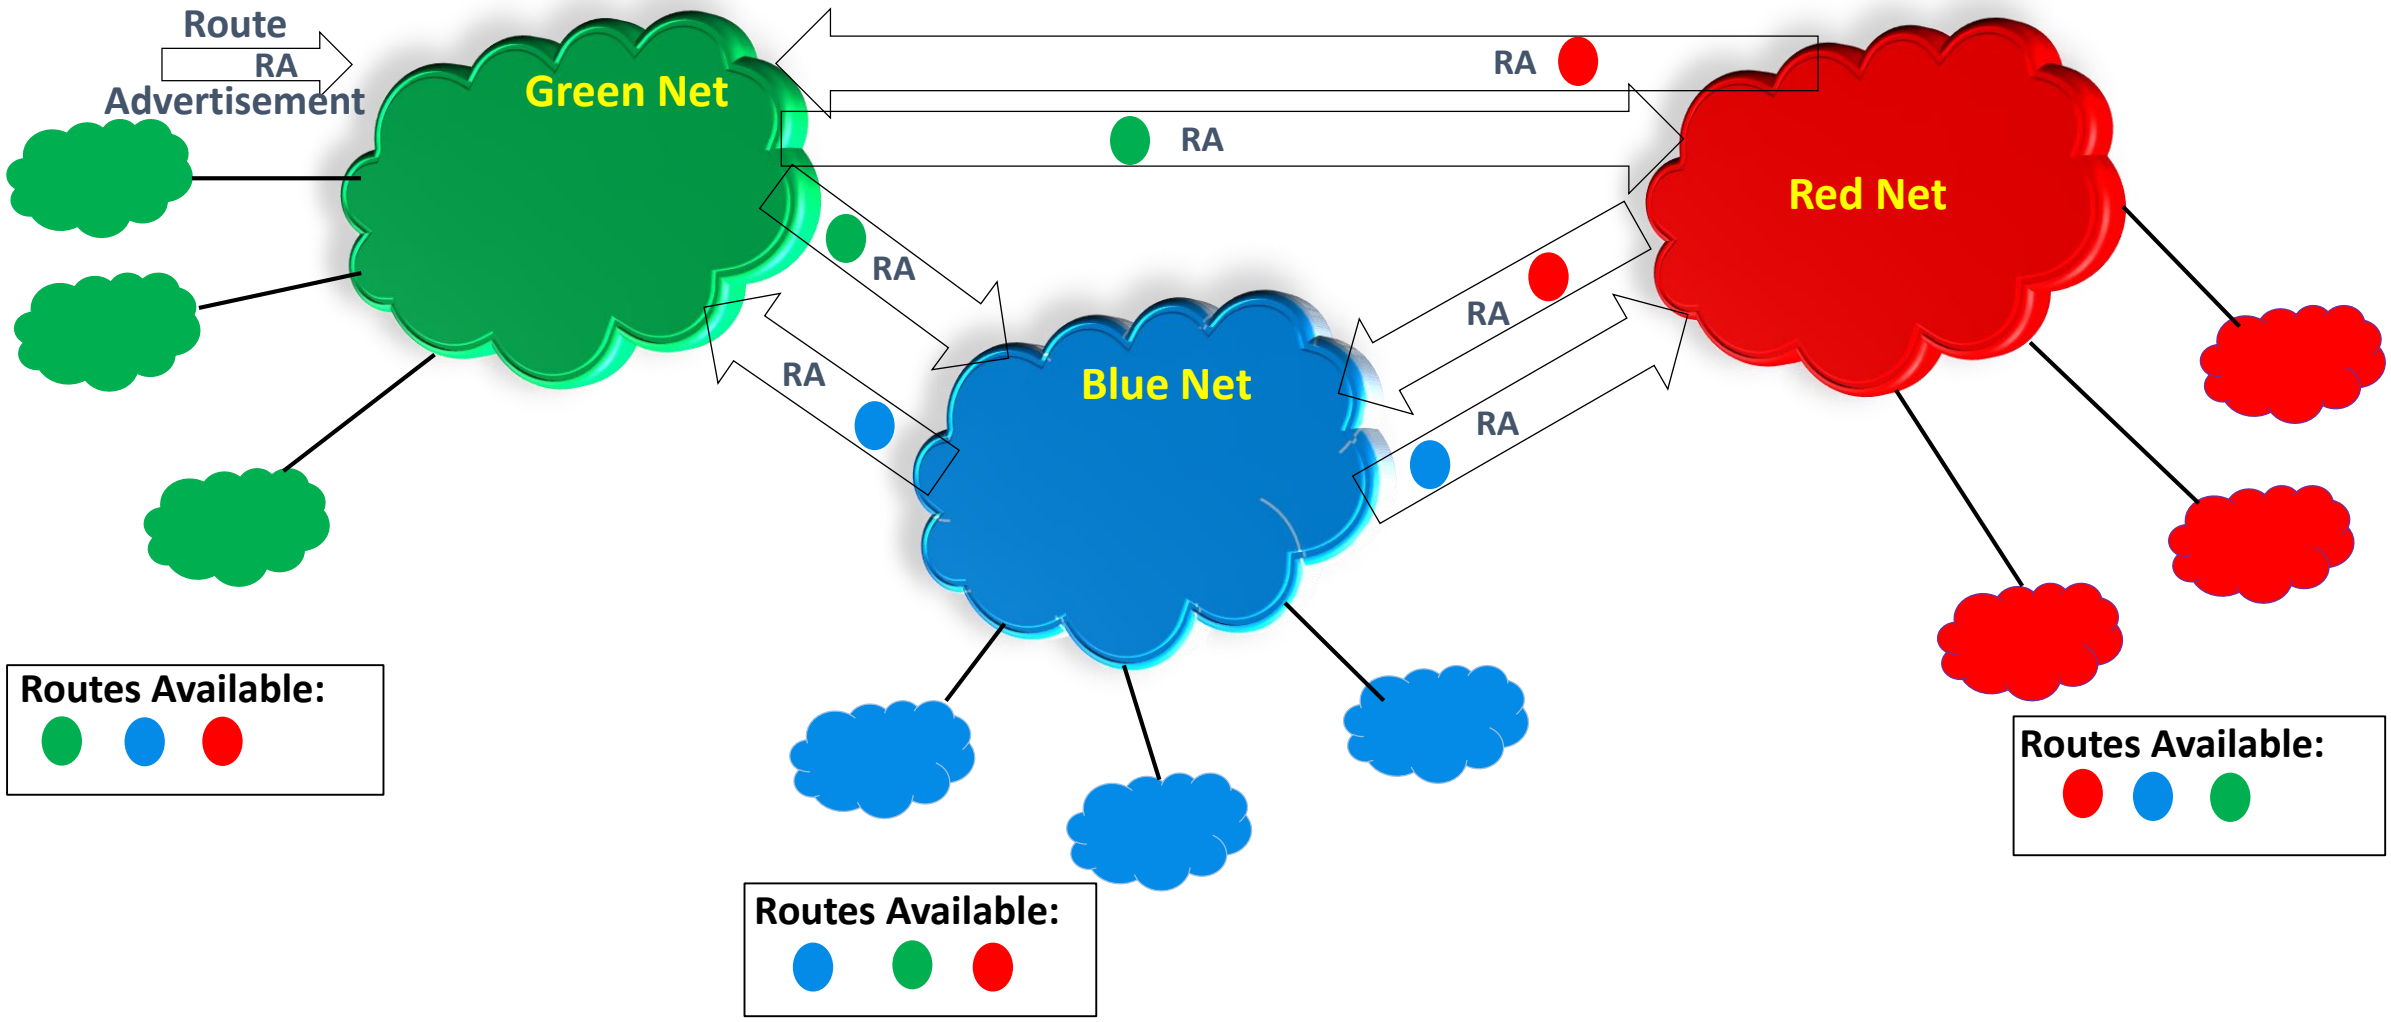

Internet Peering – Always Multiple Connections.

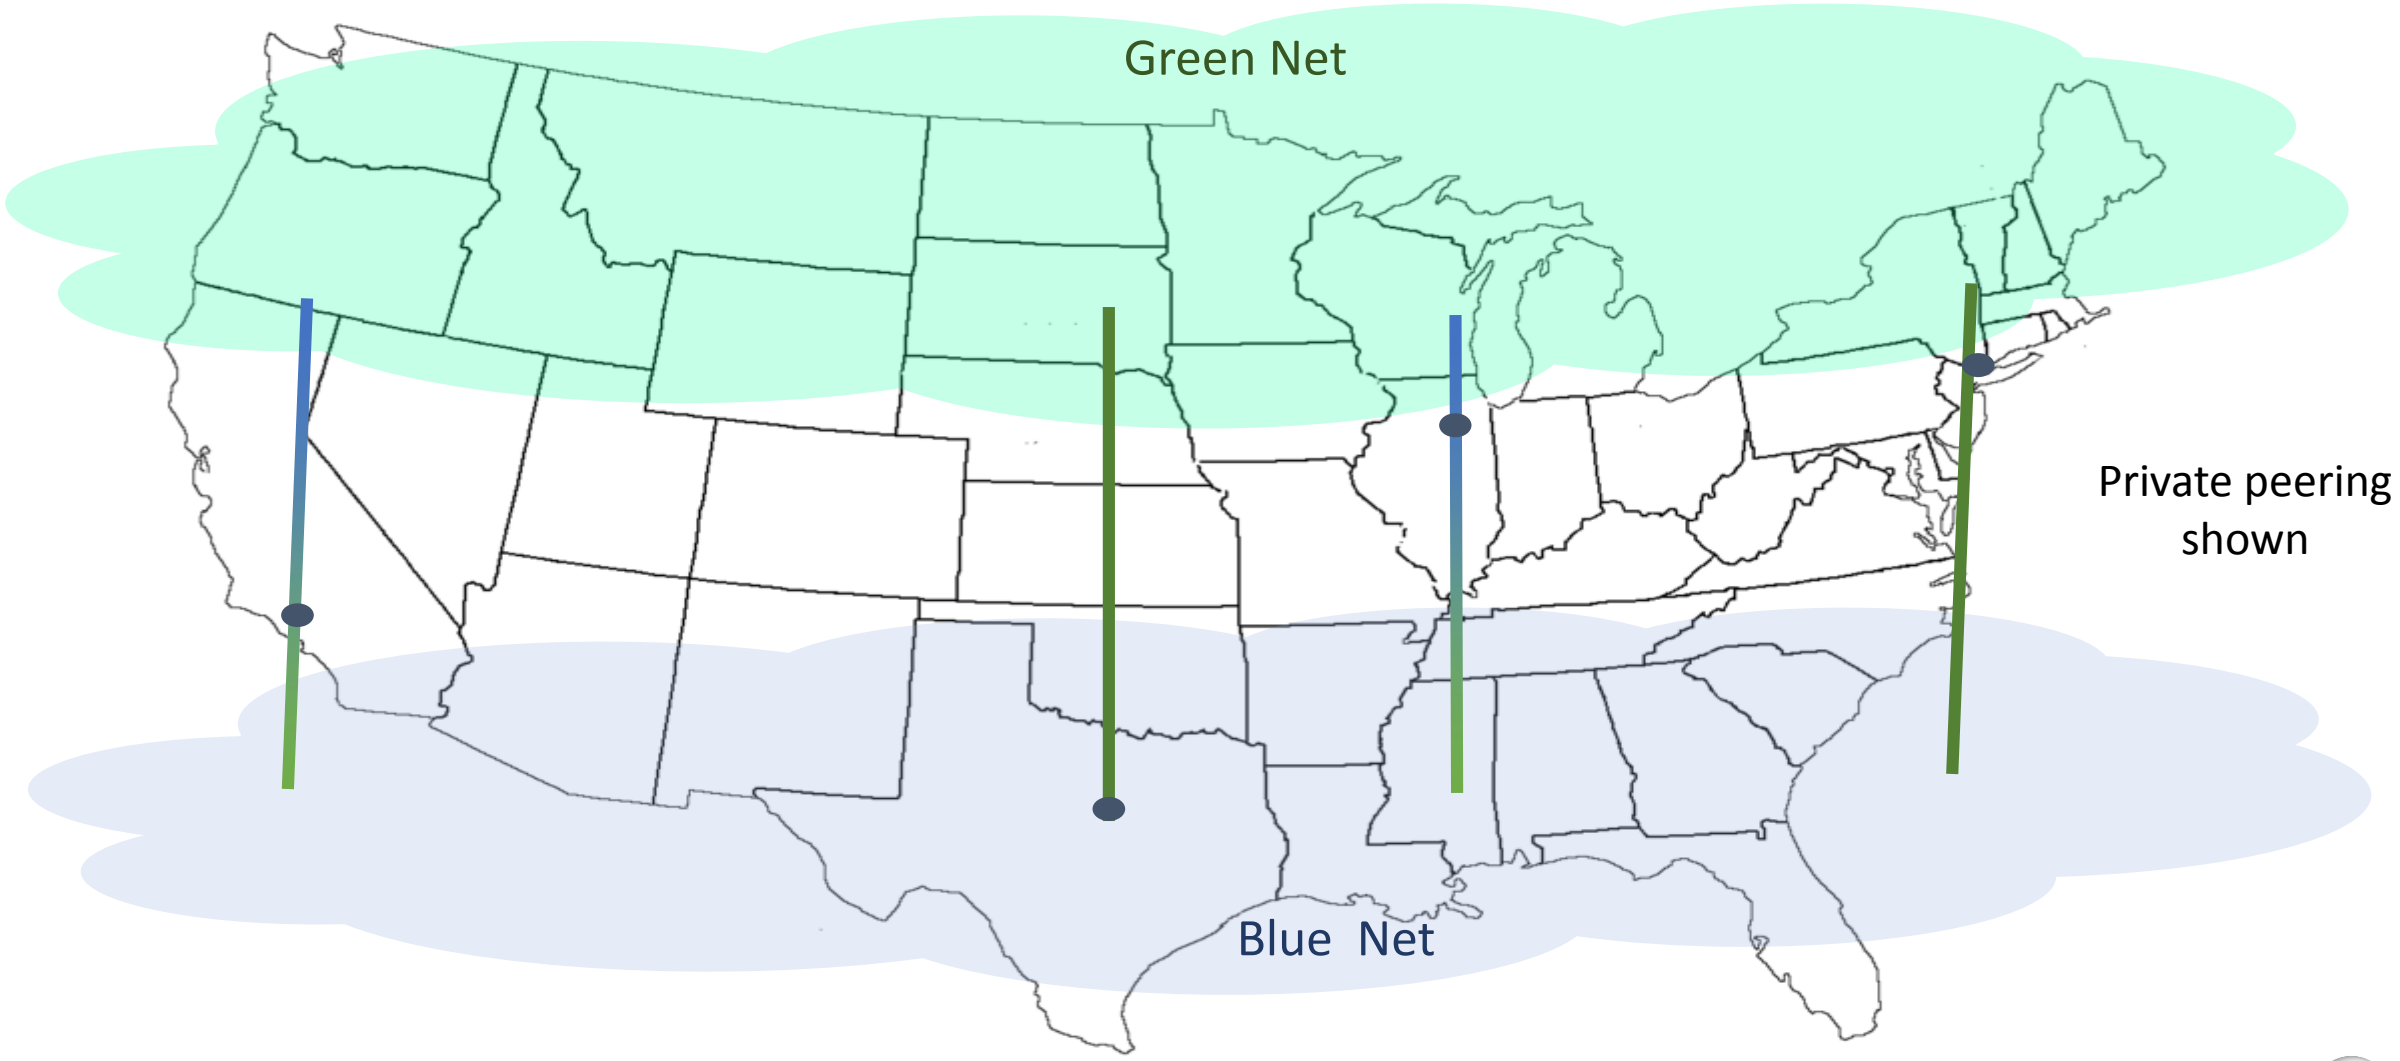

Peering – Hot Potato Traffic Routing – Asymmetrical Traffic Patterns

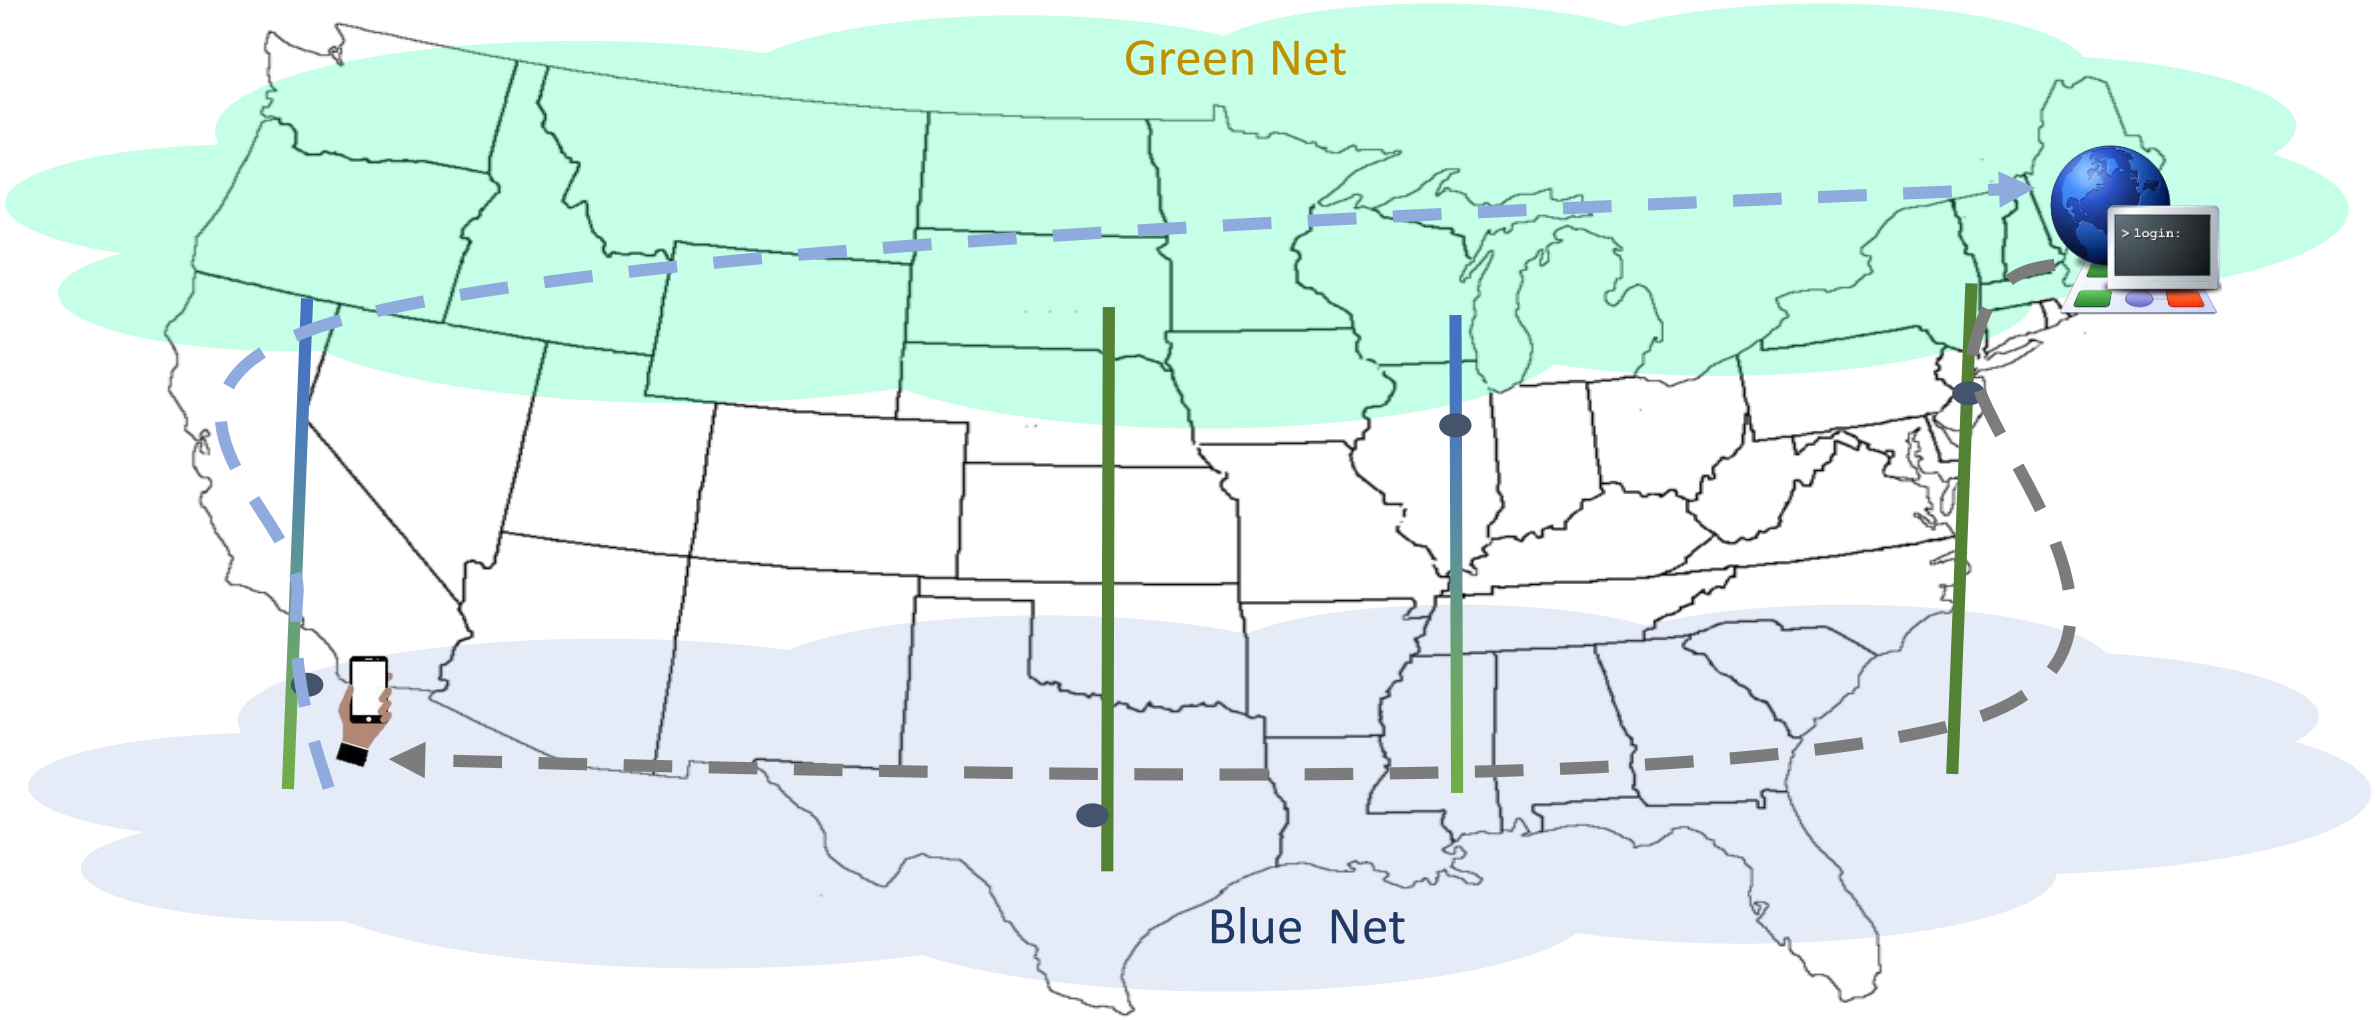

AT&T (ASN 7018) Tier 1 Internet Peering.

- 20 – 25 peers all Tier 1.
  - Some recent mergers have provided small decrease in number of Tier 1 ISPs in the last few years.
- ~ 10 different geographic locations in US.
  - 40% of the peers are connected at all locations another
  - 25% connected to all but 1 locations. 30% connected to about half the locations.
- Connection bandwidth ranges from 10GE to 300GE per connection.
  - All connections are Ethernet LAG bundles – hash is FLOW SAFE
  - Average is ~50GE per connection per peer per location.

Netflow Metering for Threat Analytics – 2nd Generation

- In-line probes, COTS w/ requirements push to vendors and Internal development.
- IPFIX, no elements – but one element over-write
  - ingressinterface (EID 10) aka ifindex
- Records from every relationship but not every 10GE.
  - Metered from >50% of 10GE circuits in the bundles.
- Records metered for 100% of packets. 1:1 sampling.

Threat Analytics (TA) Platform Functions & Technology

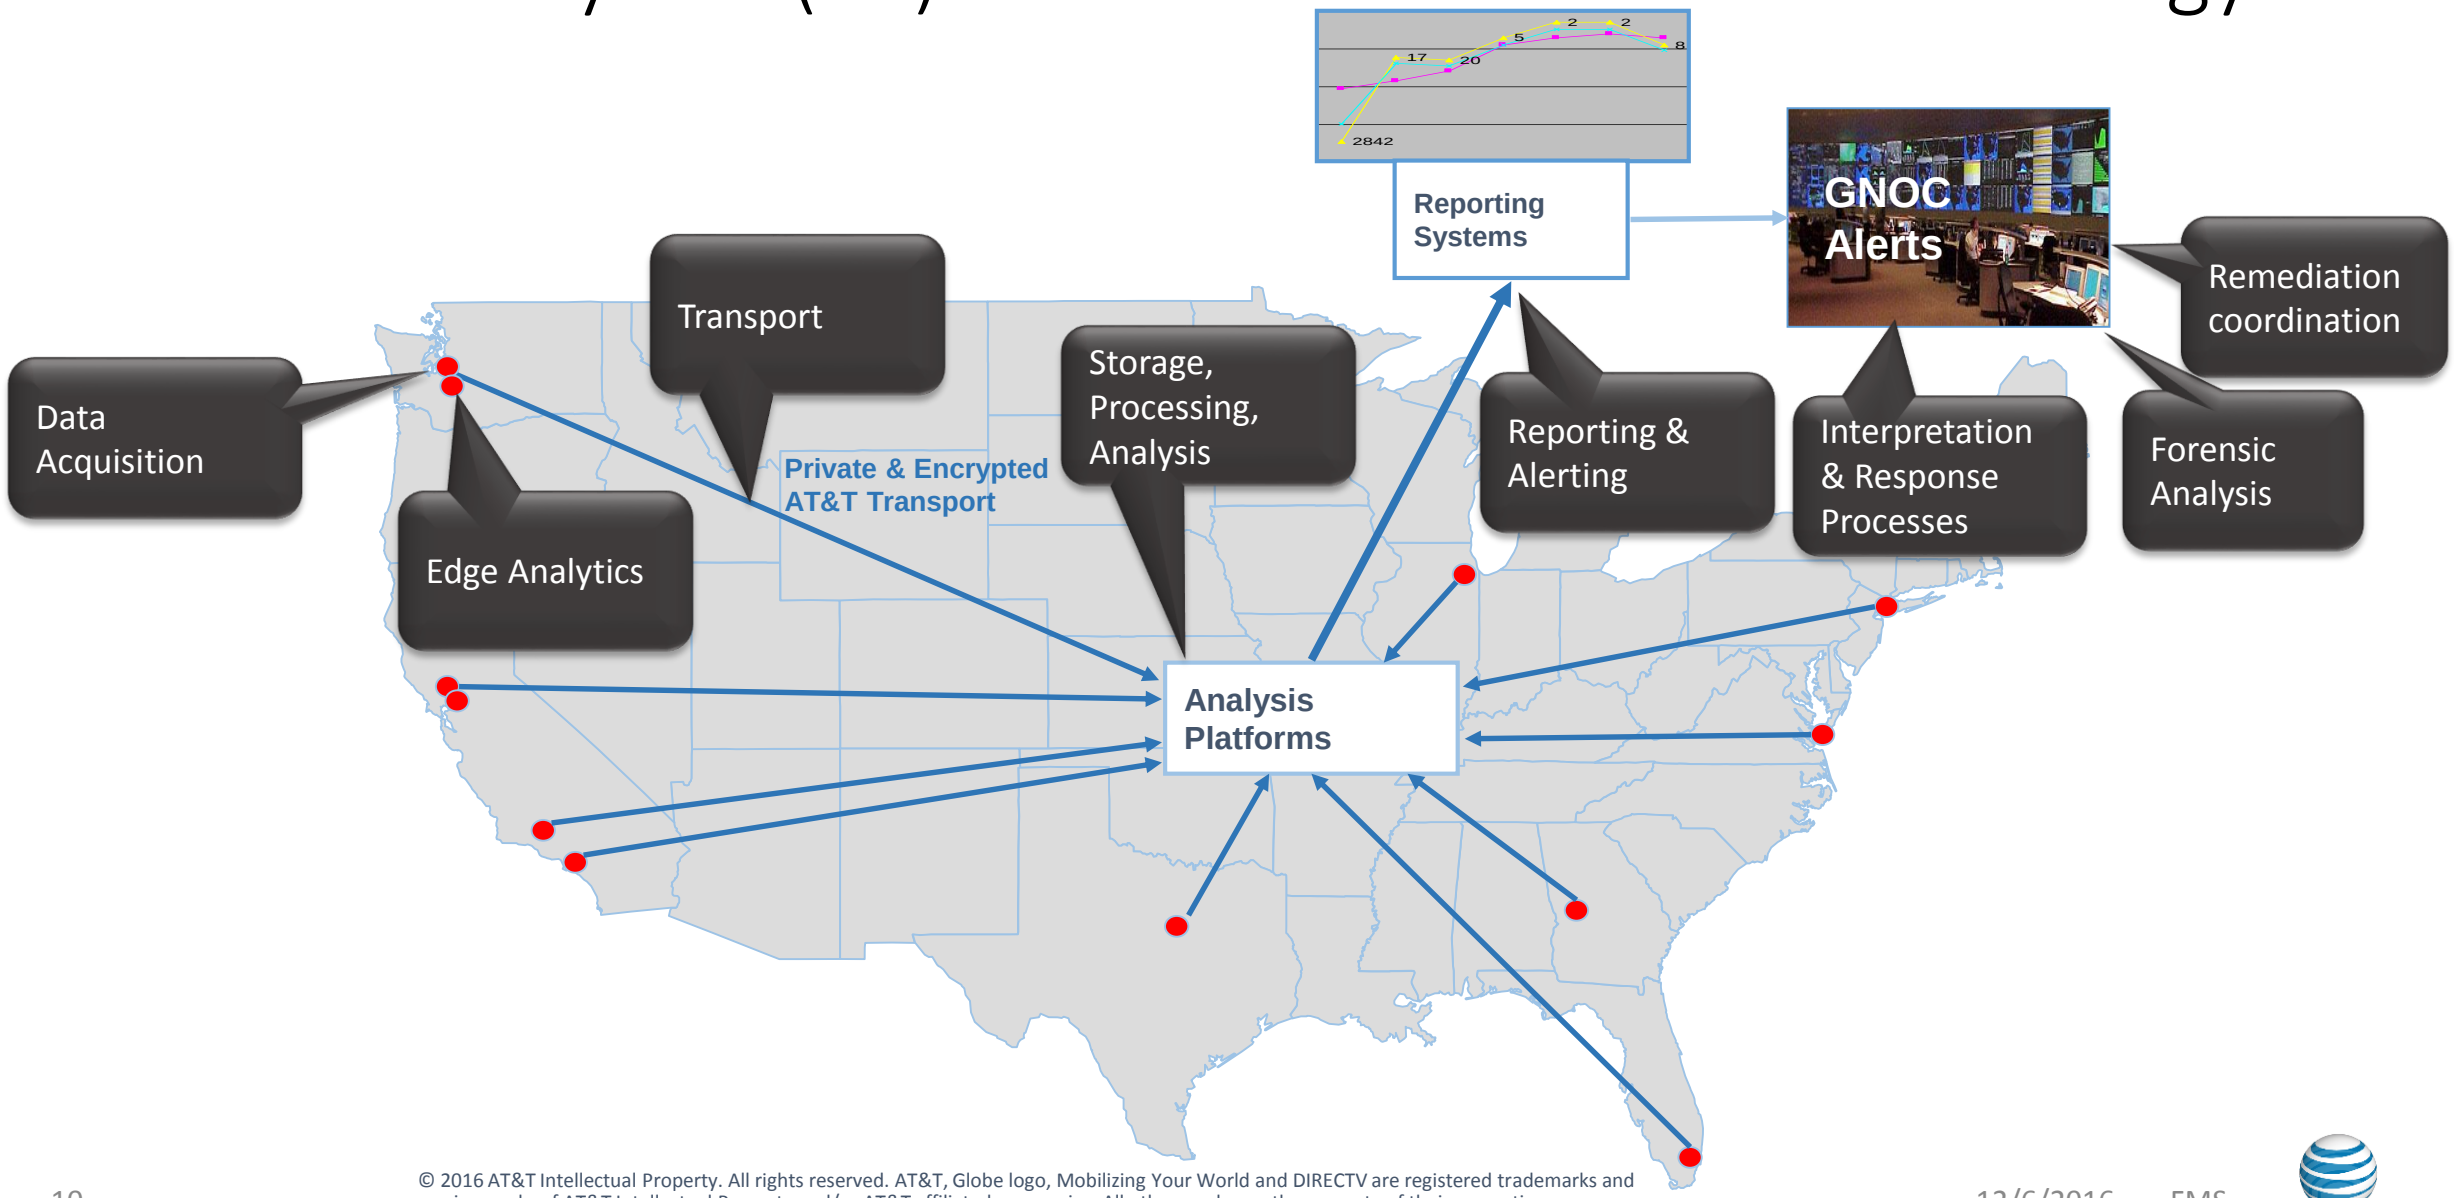

TA Processing Architecture Evolution

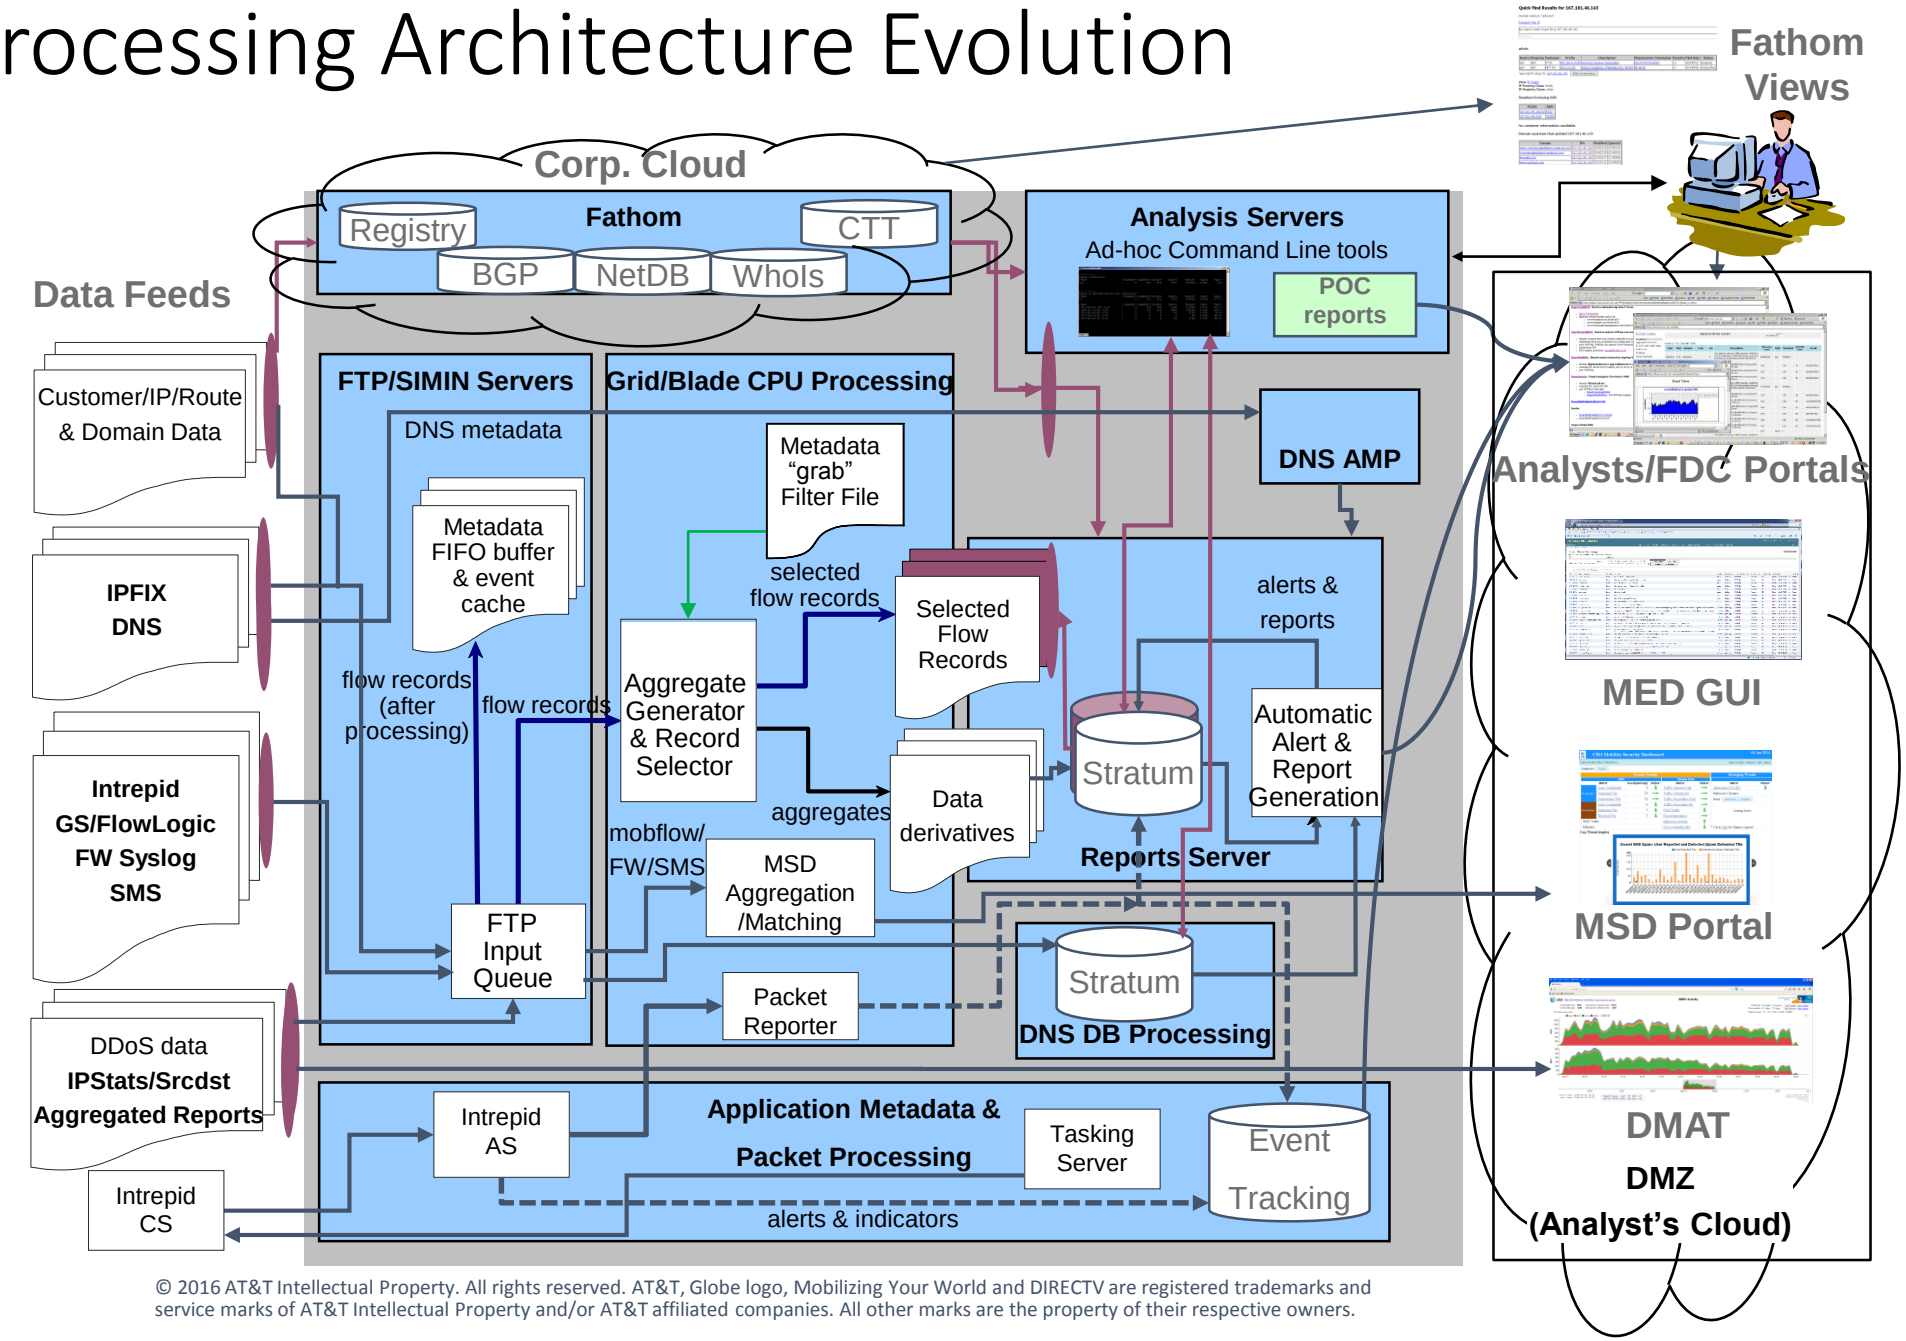

Automated Analysis Functions

- **Volumetric Alerting**
  - Port, protocol, address block anomalies
  - DDoS Attacks, other otherwise undetected events
- **Scan Detection**
  - Source address making many attempts to connect to many destination addresses or ports
- **Worm propagation (derived from scan detection):**
  - Alarming on rapid increase in the number of sources of a particular scan type (per circuit)
- **Scan Volume Alarm (derived from scan detection)**
  - Increases in scan probes or scan packets per protocol/port

Automated Analysis Functions, 2/2

- **APT Event Detection**
  - Track changes in suspected APT malware domains
- **Exfiltration Event Reporting**
  - Reports activity to/from known APT drop servers
- **Botnet Controller Detection (flow-based & DNS metadata DB)**
  - Reports on suspect bot activity based on correlated flow characteristics

Automated Flow Analysis Functions

- **Protocol/Port Volumetrics:**
  - Alarming in flow, packet, byte share anomalies
  - Measurements relative to:
    - 255 IP protocols,
    - 65536 TCP ports,
    - 65536 UDP ports,
    - 255 ICMP types, and
    - 65636 IP address blocks (A.B.x.x)
- **Scan Detection:**
  - Address scans: one-to-many IP addresses on recurring protocol/port
  - Reconnaissance Index
  - Basis of Scan Volume Alarms & Worm Alarms
  - Seed for botnet analysis

Flood – AT&T Service Network Security Analysis

Processes over 567 billion flow records per day

Records represent approximately 15 PBytes of traffic / day

Detects and characterizes approximately 500 anomalies / day

Keeping Up:

- Automated aggregation and processing of anomalies to create ACTIONABLE report.
- Integration with workflow support system
- Interaction / action with graphical representation of data

Processing efficiency

As we moved to a Big (Huge) Data platform the linkage to processing on a per circuit basis was eliminated.

Still wanted source(s) traceability of flows.

Upon collection of flow records, the collection server maps it's server id, the input port, and the ifindex written in the record by the probe and over writes that with a network wide unique ID. Therefore without any other data or index we can determine exactly where a flow record was metered.

Note:
Ifindex (SNMP Index) in most routers is very stable today. Typically survives reboot and restart.
We pull and dif the ifindex table from all the probes periodically, after every restart and after every software upgrade.
The map file in the Collector is updated as needed when the ifindex table changes.

Graphical Data with actions - DMAT

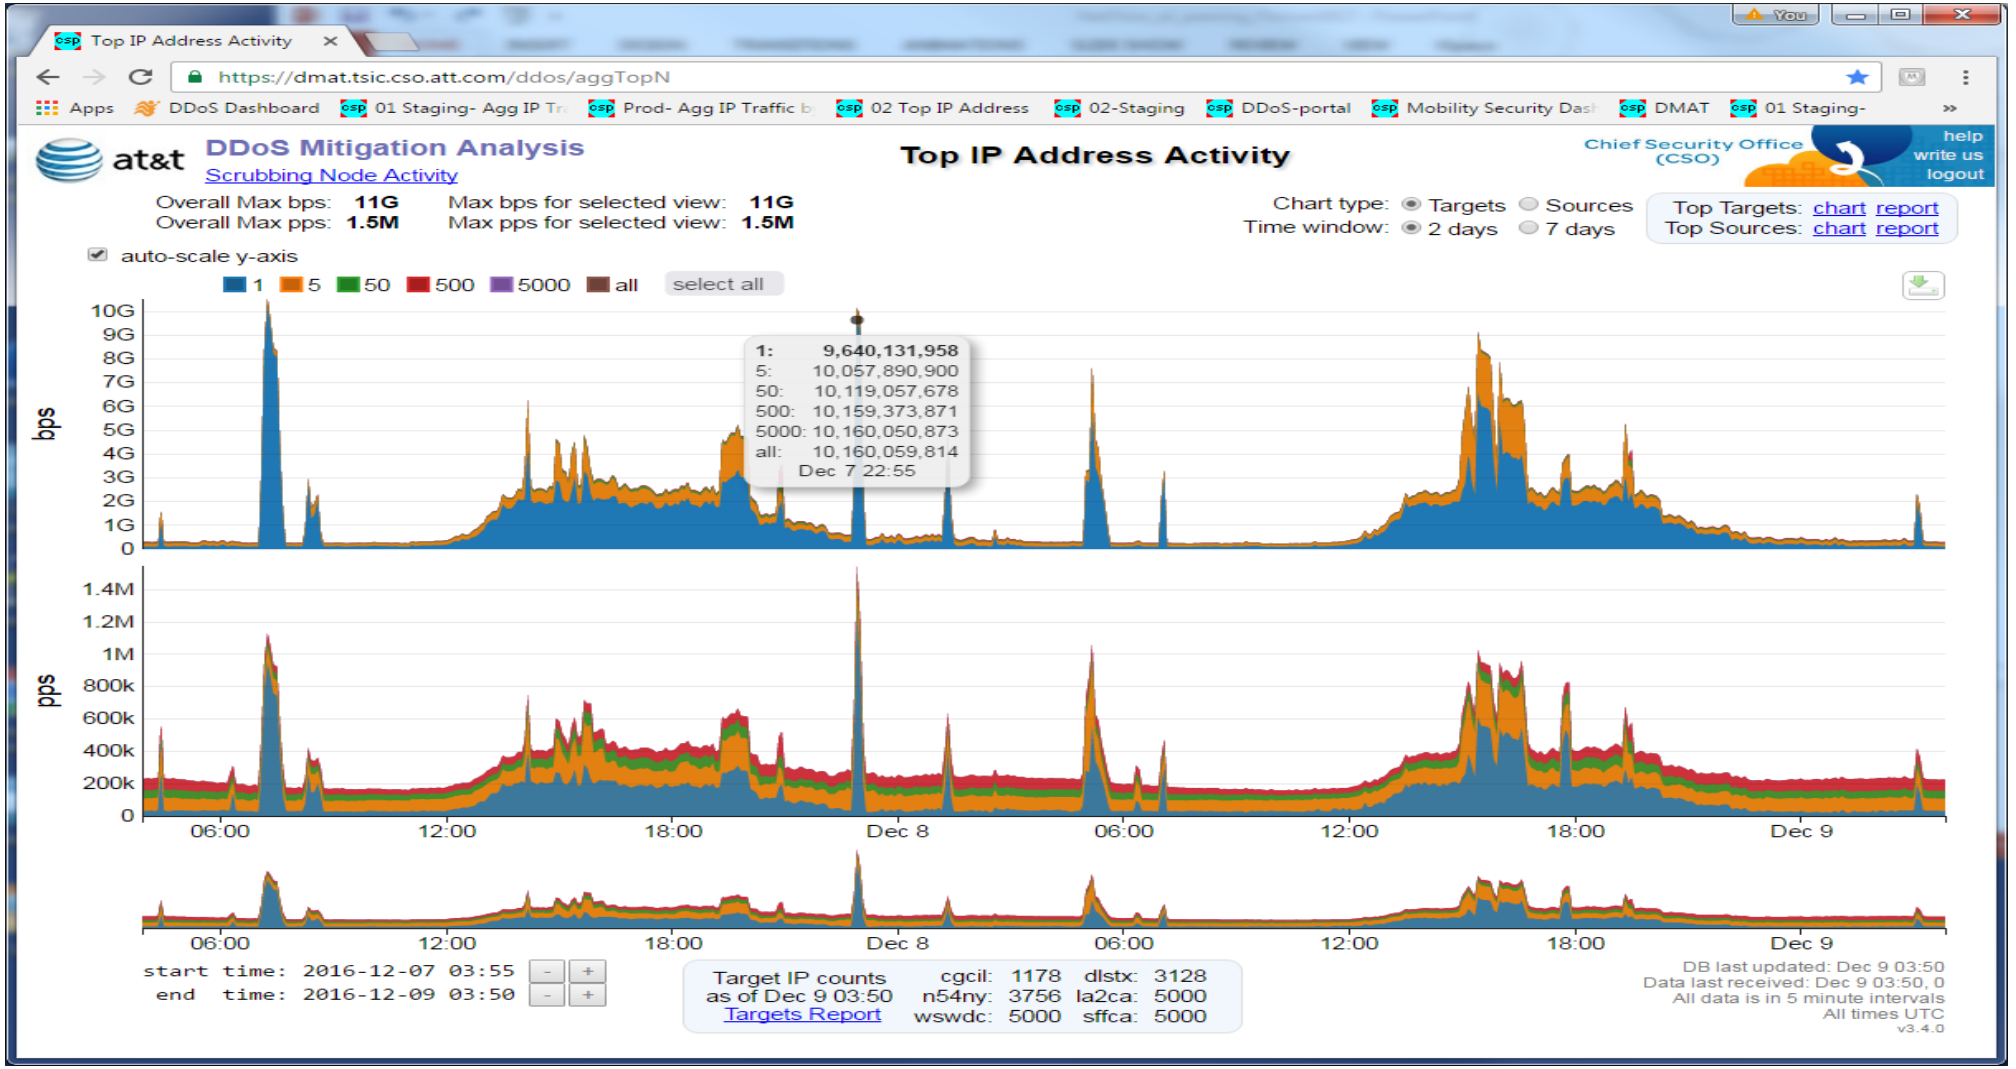

DMAT 2

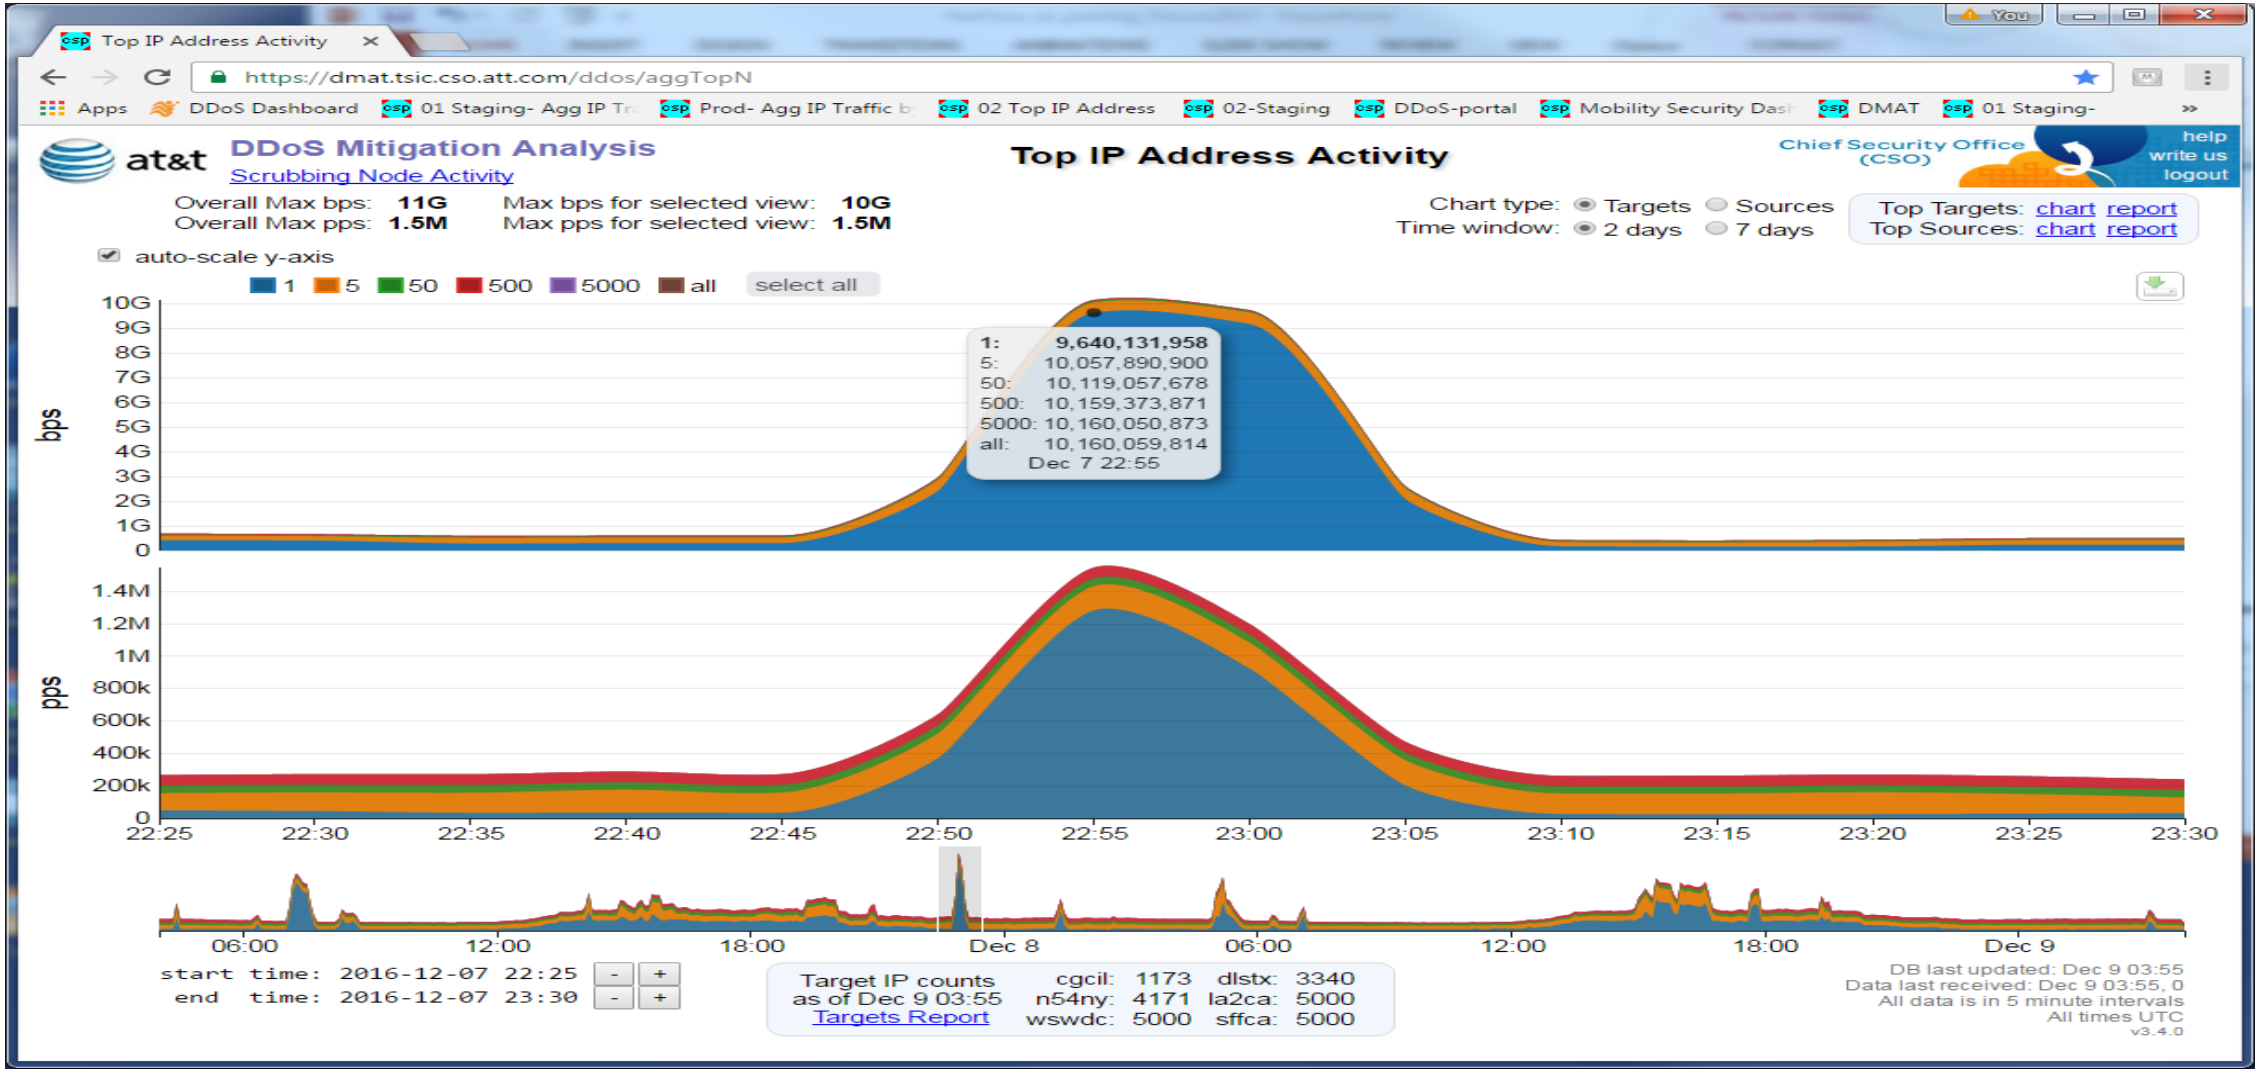

DMAT 3

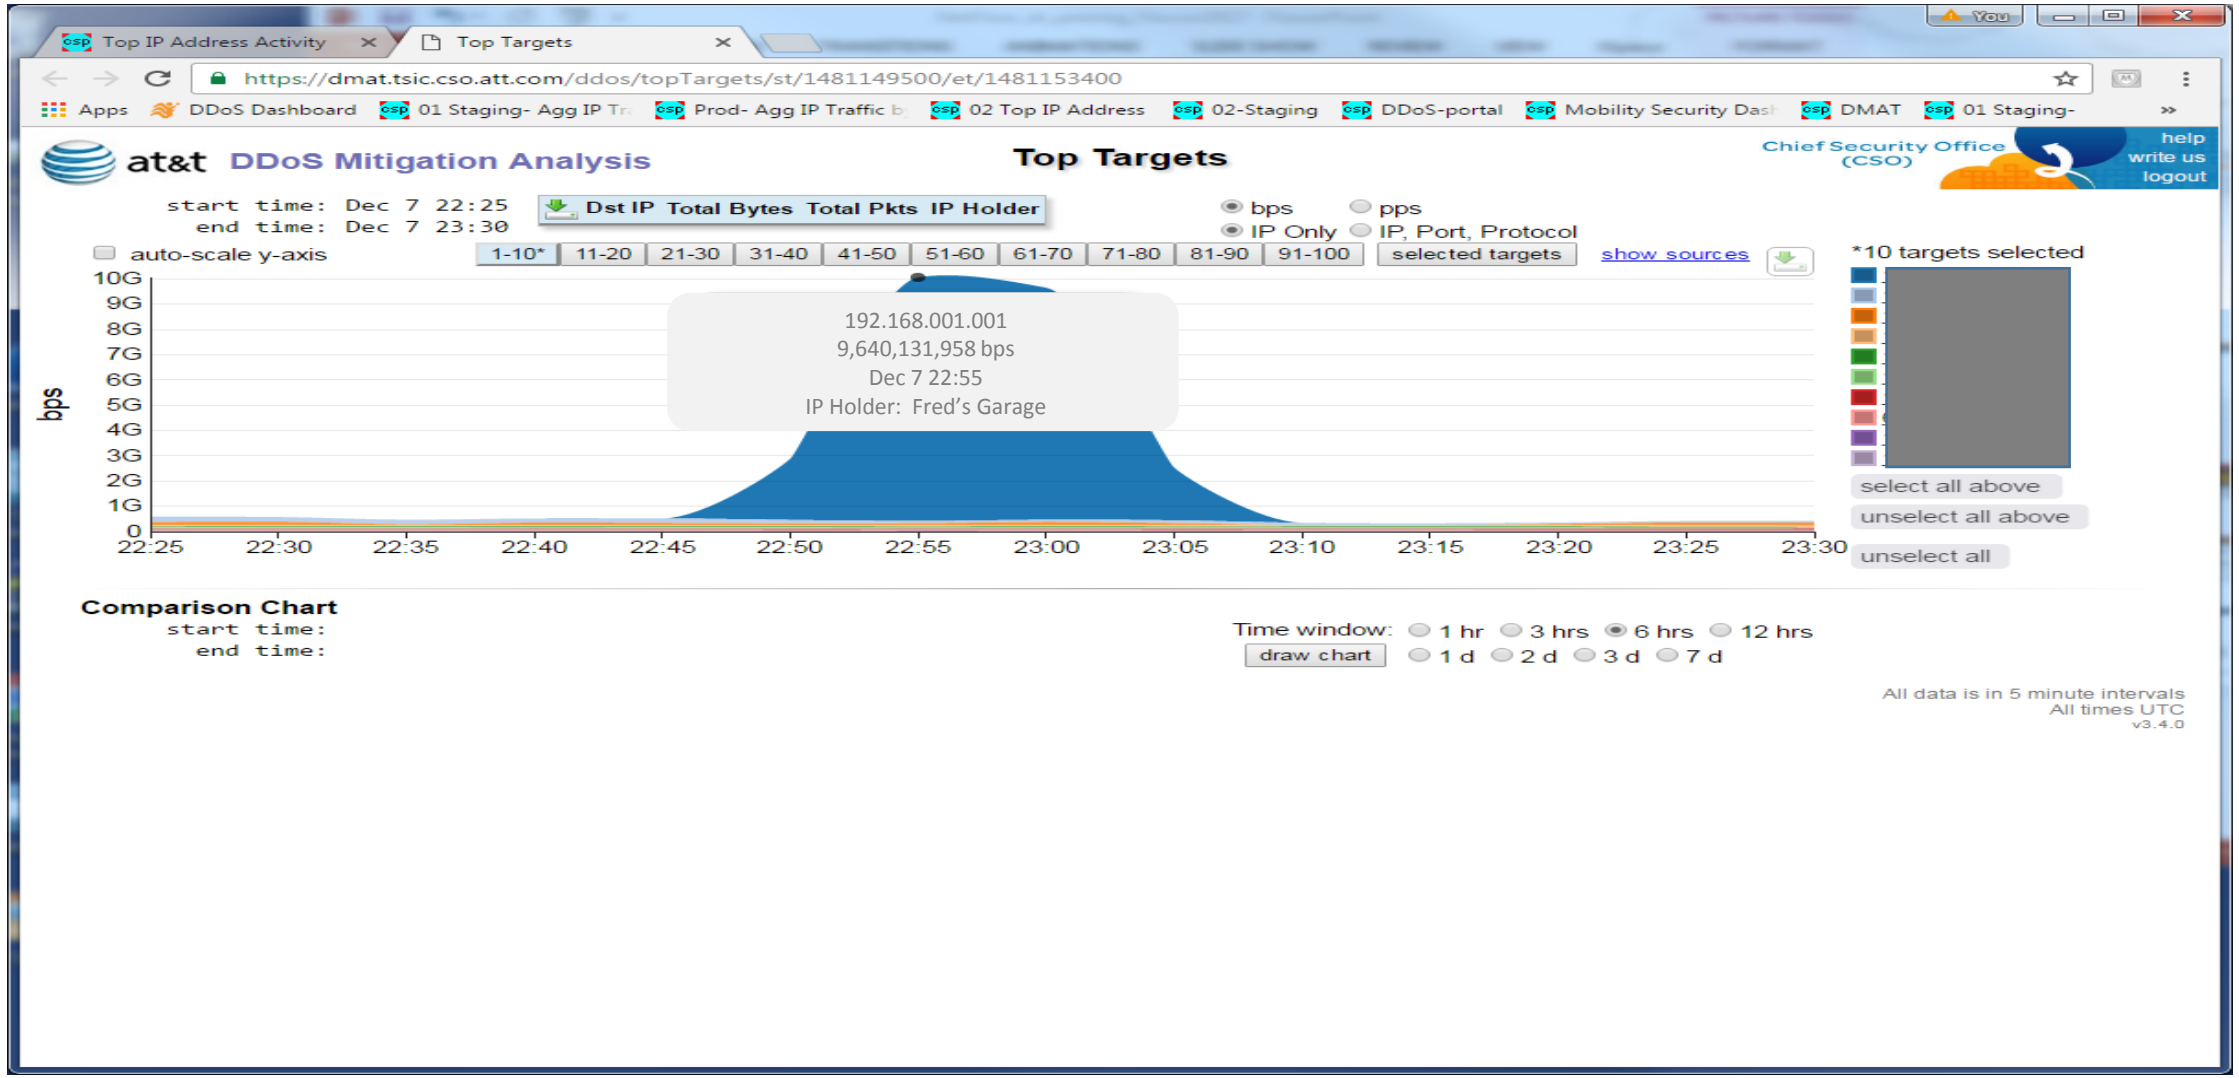

DMAT 4

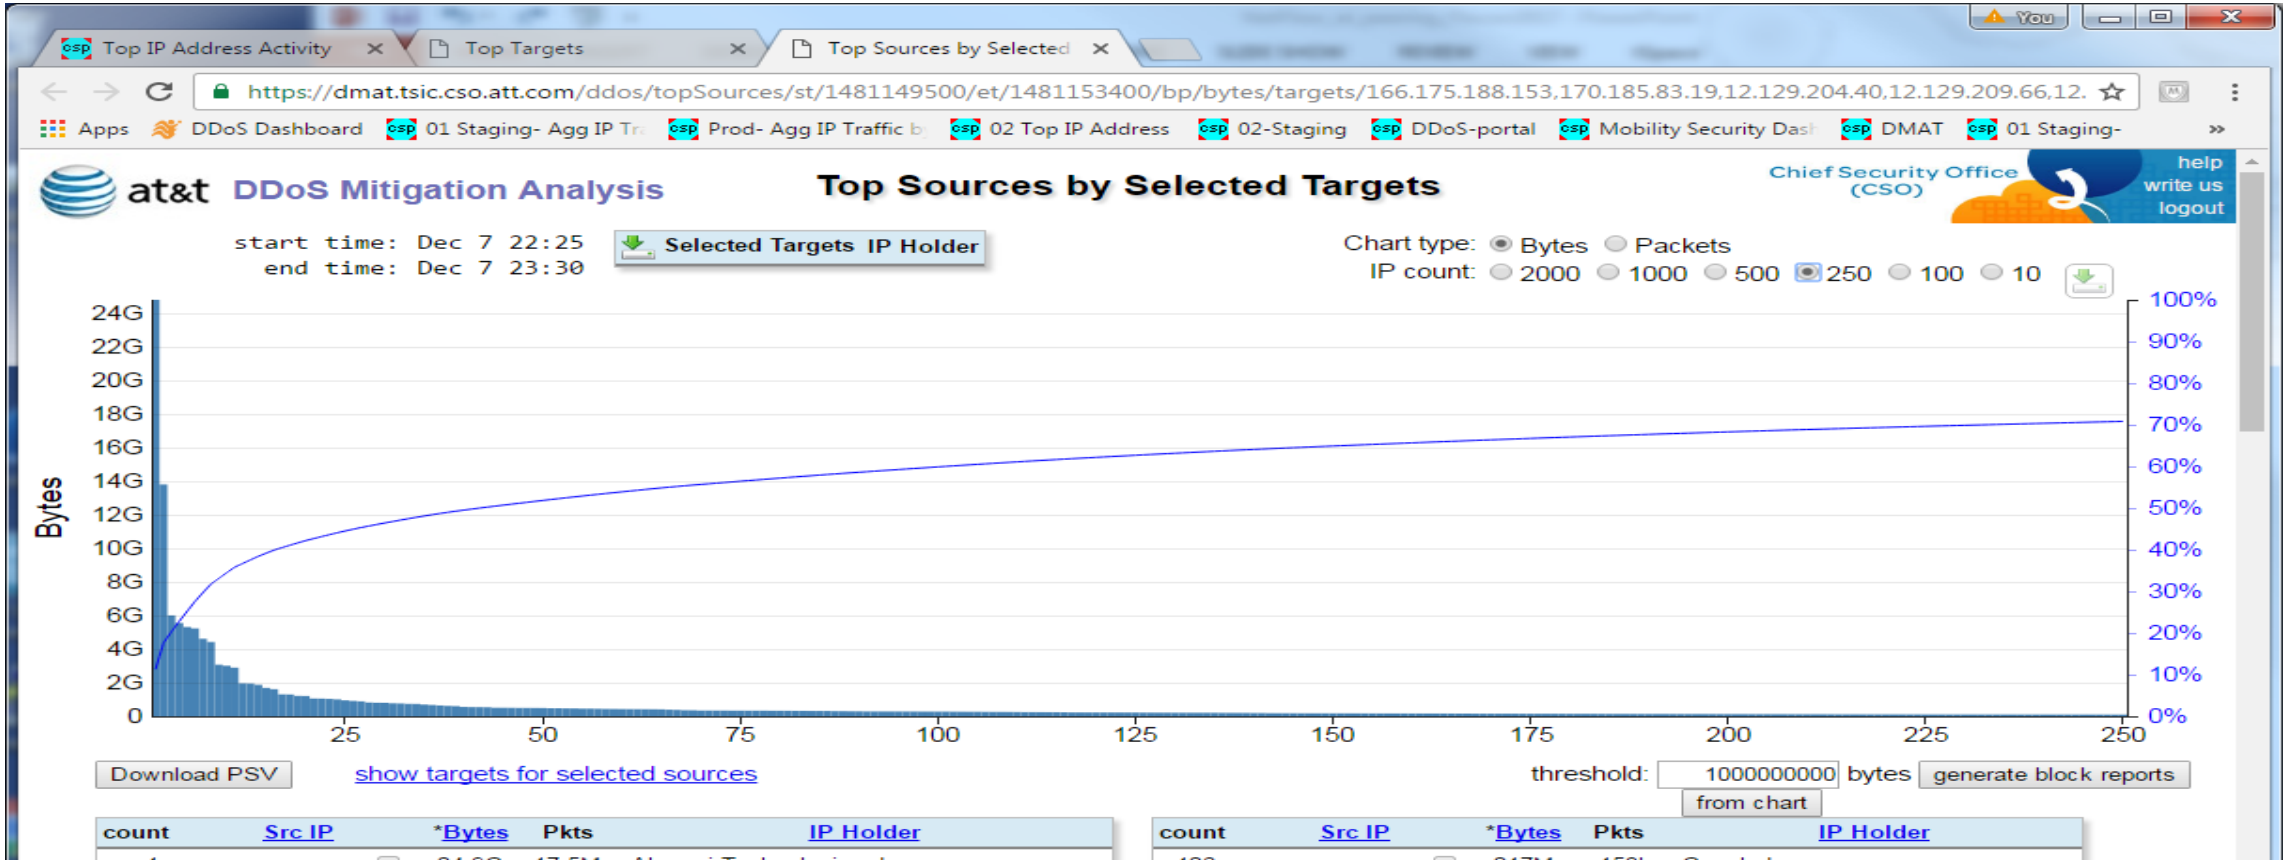

Network Evolution

- The NFV all Ethernet network.

Threat Analysis Transformation Vision

Today

Traffic access
- Passive Probes - static

Traffic Selection
- Dedicated Routers
- Home grown “tasking system” for filter implementation

Metadata Generation
- Dedicated probes (a probe / 10GE)
- Dedicated Appliances

Data Files
- “Collector” application, file server on dedicated servers and storage in SNRC/CO space.

Data Transport
- Private IP Network

Data Analysis
- Centralized – dedicated data center

Future

Traffic access
- Taps, Agents and Probes in Service Network Elements (SNEs)

Traffic Selection
- NFV (Network Function Virtualization) filters
- SDN & Common Orchestration with Service Network
- All the time, sampled, event driven and on-demand

Metadata Generation
- Virtual Probes in SNEs
- NFV Multipoint Probes in SNEs (described later)
- All the time, sampled, event driven and on-demand metadata generation.

Data Files
- NFV Collectors in SNEs and VM Collectors in private space

Data Transport
- AVPN with Orchestration

Data Analysis
- Streaming point of collection analytics, data at rest analytics and chained analysis functions.

Terms and abbreviations, as used in this presentation

Probe	a device or function which monitors traffic and generates a data record about that traffic. An IP Flow generator (IPFIX metering and export process) is a probe.
Multipoint probe –	please see slide 6
TAP	A device/function which duplicates a bit stream or packet/cell/frame stream and forwards to at least one additional destination. A span port in an Ethernet switch and an optical splitter are both taps.
NFV	Network Function Virtualization – a software implementation which can run on common compute platforms providing functions typically found previously in Network Element dedicated platforms. Synonymous with VNF (Virtual[ized] Network Function
sNFV	Security Network Function Virtualization – applications like IP filtering, Firewall filters (ACLs), DPI, NAC, intrusion detection implemented in virtual machines and not dedicated appliances
NE	Network Element – a device, typically dedicated hardware and software for network transport, connections, traffic management and network management. An IP Router and an MPLS switch would be considered Network Elements.
Service Network	A network infrastructure operated to deliver commercial services to customers. In 2015 and later architectures all the adjunct function like OAM, capacity management, etc. are integrated in a virtual platform. Generic term to cover all AT&T Service offering infrastructures especially wireline IP and mobility.
SDN	Software Defined Networking - A network infrastructure implemented, configured and provisioned via rules, policy and Service Orders by software driven systems.
vSelectiveTap	An NFV or combination of NFVs which selects specific traffic by a data field in IP headers, frame, GTP header, etc. and forwards a copy of the matching traffic to a tap destination, which could be another VNF.

Terms and abbreviations pg2, example graphics

IPFIX (netflow) Probe example

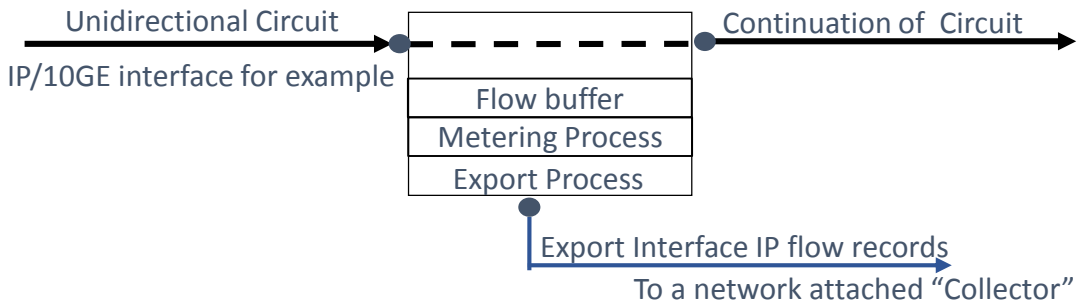

TAP, passive optical Splitter example

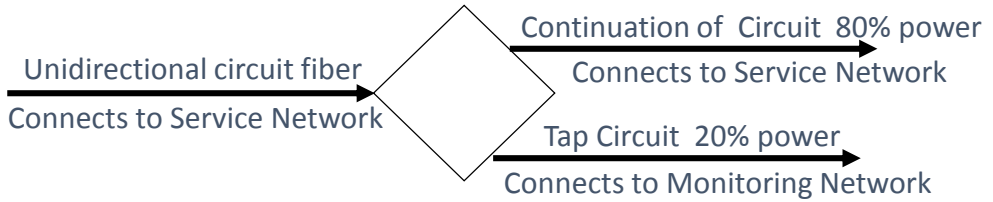

Splitters are available in different power splits 50/50 and 80/20 are common.
Often implemented as plug-in devices in LGX (Lighguide Cross Connect) frames.

Multipoint Probe example

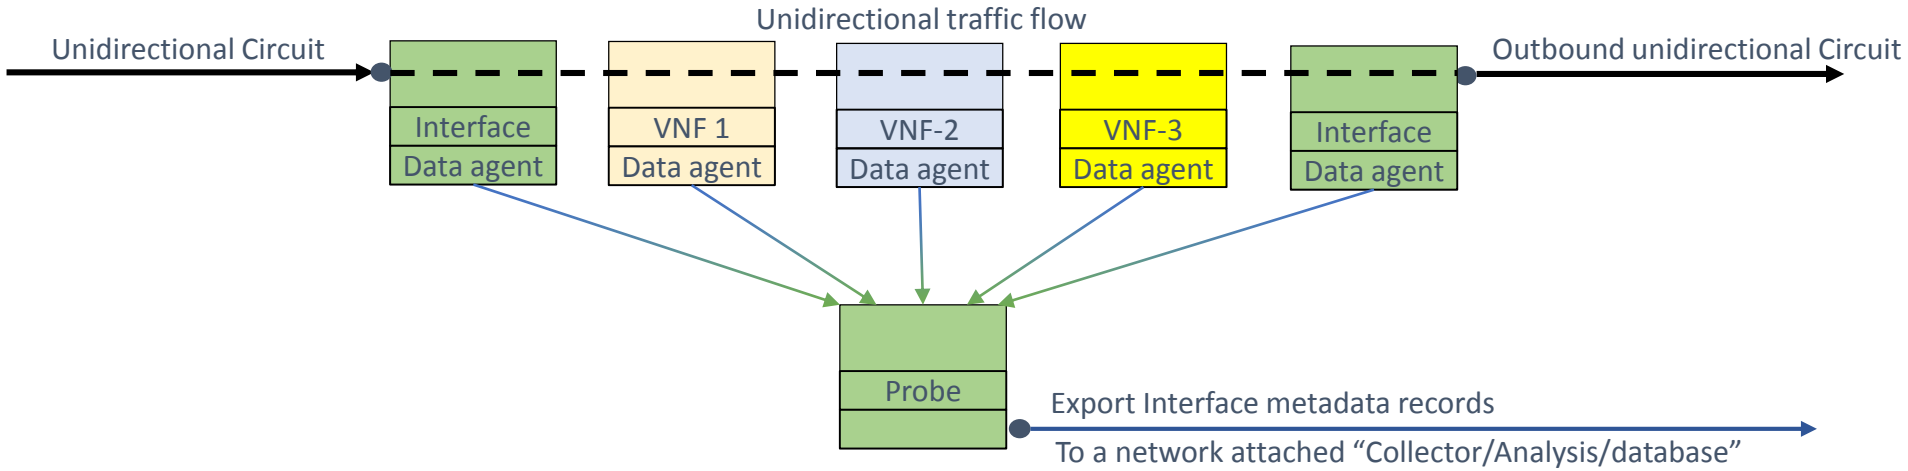

The NFV based network ENABLES collection of metadata records with fields from processes (functions – services) along the data flow.
For example inbound interface data, user authentication, filter matches, before and after NAT etc. all in one record with no correlation from many different taps and probes.

Surveillance Plan – enabled by NFV based network w/Orchestration

- Traffic monitoring, traffic selection and data analysis for Threat Detection and Analysis (TA) is analogous to street crime monitoring and planning police patrols.
- High crime – high threat probability (from risk analysis and history) – High risk/ moderate value
  - Crime – constant monitoring of area camera, regular patrols.
  - TA – constant circuit connection, full time metadata on a sample, full time anomaly detection
    - TA where – peering edge (current today) and hosting edge (future)
    - TA how – full time vTap, vNetflow probe, vDNS selector exported to TA network.
- High Value Asset – lower threat probability but high value, therefore moderate risk
  - Crime – constant monitoring of specific camera, patrols defined by time of most risk.
  - TA – selective monitoring (by IP address or userID), sampled metadata, exception reporting.
    - TA where – high risk/high values customer's edge.
    - TA how – on demand vSelectiveTap, tap traffic connected to network attached probe and security appliances which are typically virtual in a SNE or TA VM.

Surveillance Plan – enabled by NFV based network w/Orchestration, page 2

- No known risk – high threat probability (defined by Intelligence) –
  - Crime – investigation, surveillance and follow-up. Investigate source
  - TA – the same as crime – investigate, monitor, generate metadata, investigate threat.
    - TA where – anywhere in the Service Network depending on intelligence.
    - TA how – vSelectiveTap, vProbes, copy of traffic to security tools, map flow patterns, traffic pattern match.
- Everywhere else – unknown risk / unknown value – You can say it is moderate threat with no analysis or history. Unknown has no comfort to it. It is what analysts will lose sleep over.
  - Crime – sampled monitoring of area cameras, patrols defined to be hard to predict.
  - TA – selective monitoring (by IP address or userID), sampled metadata, exception reporting.
    - TA where – anywhere that is not being surveyed in another category – but overlap is fine.
    - TA how – on demand vSelectiveTaps defined by pseudo random algorithm, tap traffic connected to network attached probe and security appliances which are typically virtual in a SNE or TA VM.

Surveillance Plan – enabled by NFV based network w/Orchestration, page 3

- Completing the analogy
- Active Crime (reported, burglar alarm) and Active Attack (Customer report, alarms, alerts, traffic pattern anomaly).
  - Crime – respond, mitigate damage, escalate resources until no more are needed, capture perpetrator
  - TA – exactly the same as Crime.
    - TA where – pick a point and expand or contract as appropriate.
    - TA how – every and any tool in the kit.

AT&T ThreatTraq

Weekly Cyber Threat Report

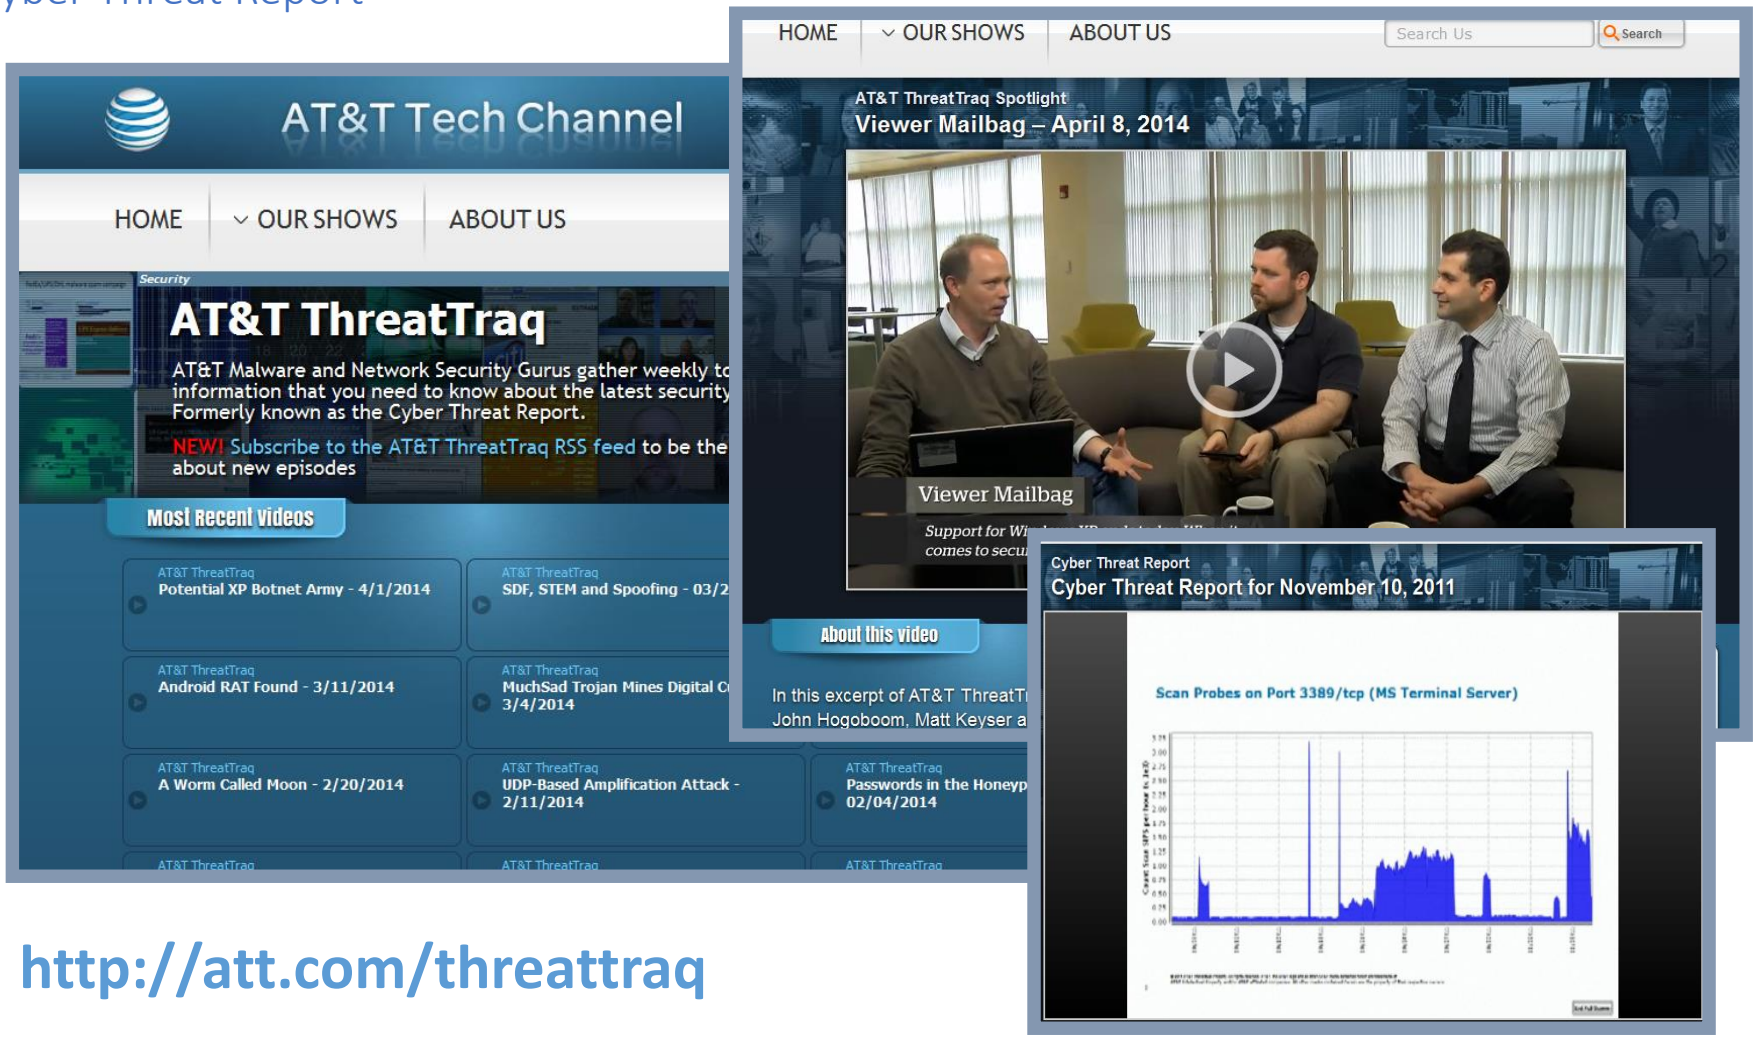

http://att.com/threattraq