CONFIDENTIAL

Advanced Traffic Steering & Optimization Technologies

Bart Salaets

Solutions Architect EMEA

Agenda

- Recent Evolutions in Traffic Steering
- Flow-based vs Transaction-based Traffic Steering
- Service Chaining & IETF Activities
- TCP Optimization
- Summary

Recent Evolutions in Traffic Steering

Traditional Steering to VAS & Optimization platforms

A router steers all port 80 traffic to VAS platforms

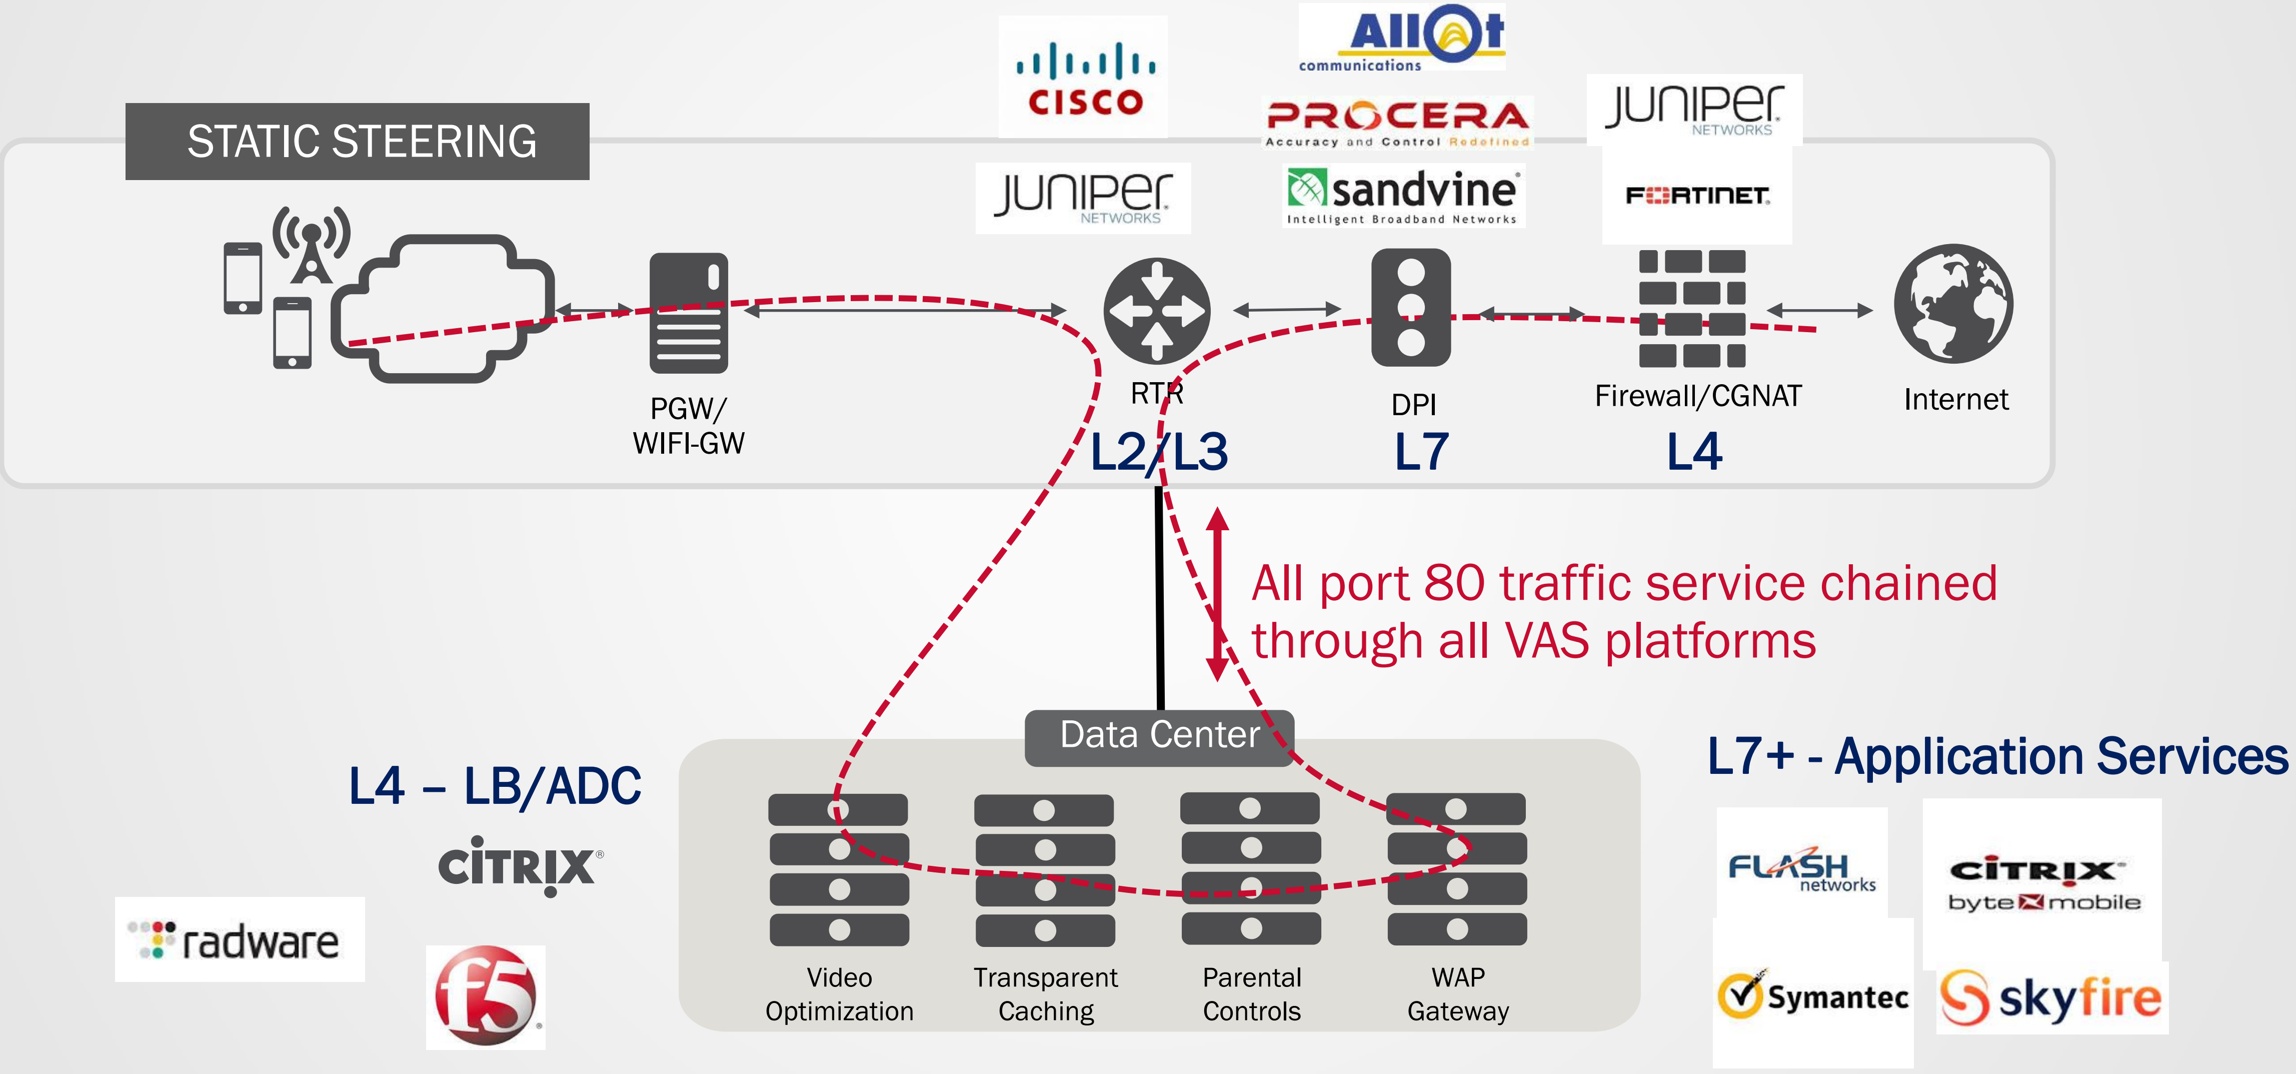

Intelligent traffic steering to VAS platforms

Offloading VAS services & Optimizing infrastructure utilization

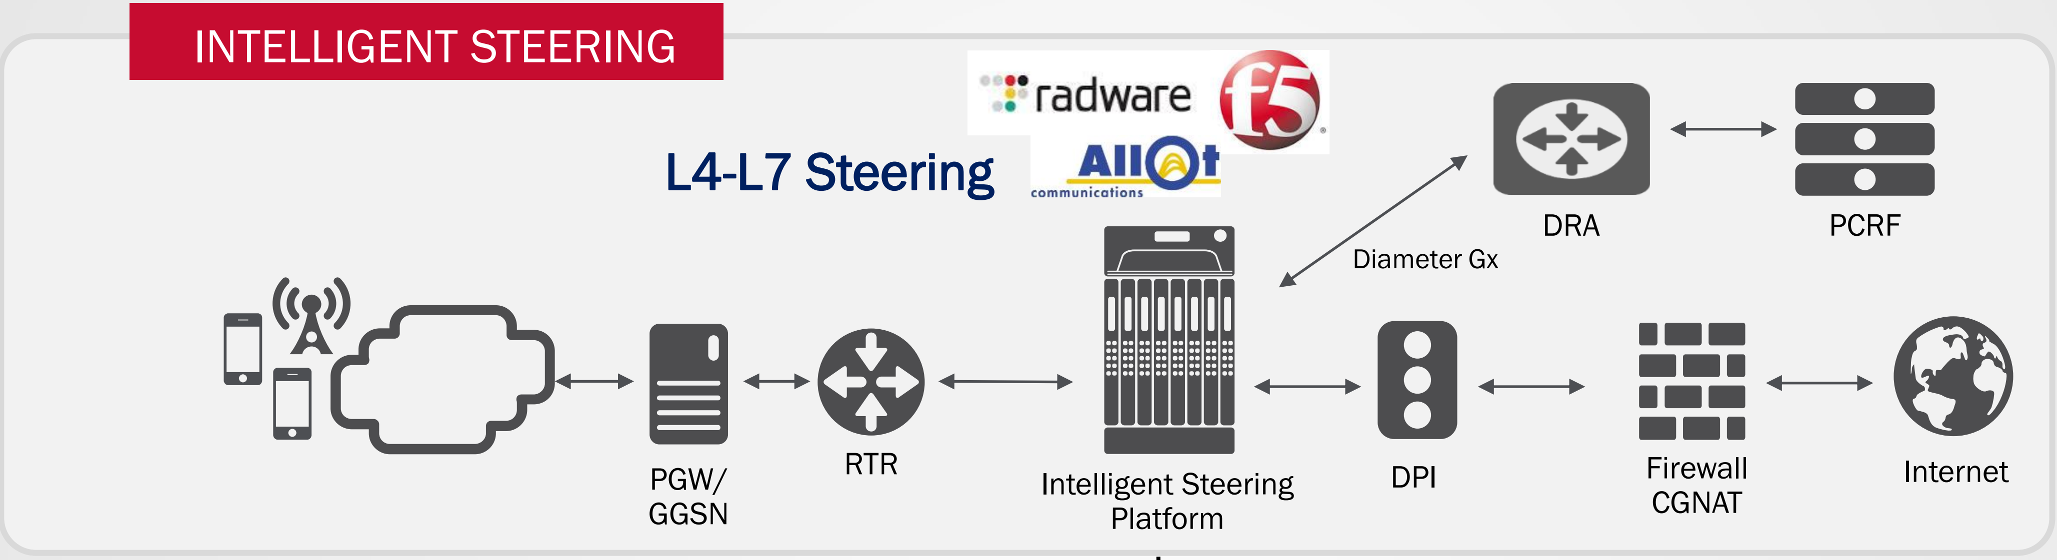

CONTEXT

SUBSCRIBER
DEVICE-TYPE
RAT-TYPE
CONTENT (VIDEO, URI, ...)
CONGESTION

Context-aware & policy-driven steering & intelligent service chaining

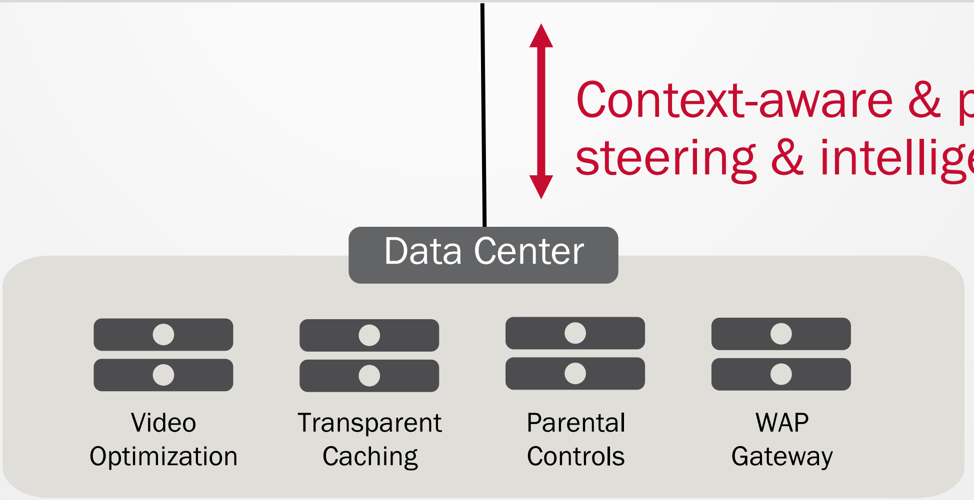

Intelligent traffic steering to VAS platforms

Example : Subscriber and RAT-type based steering / service chaining

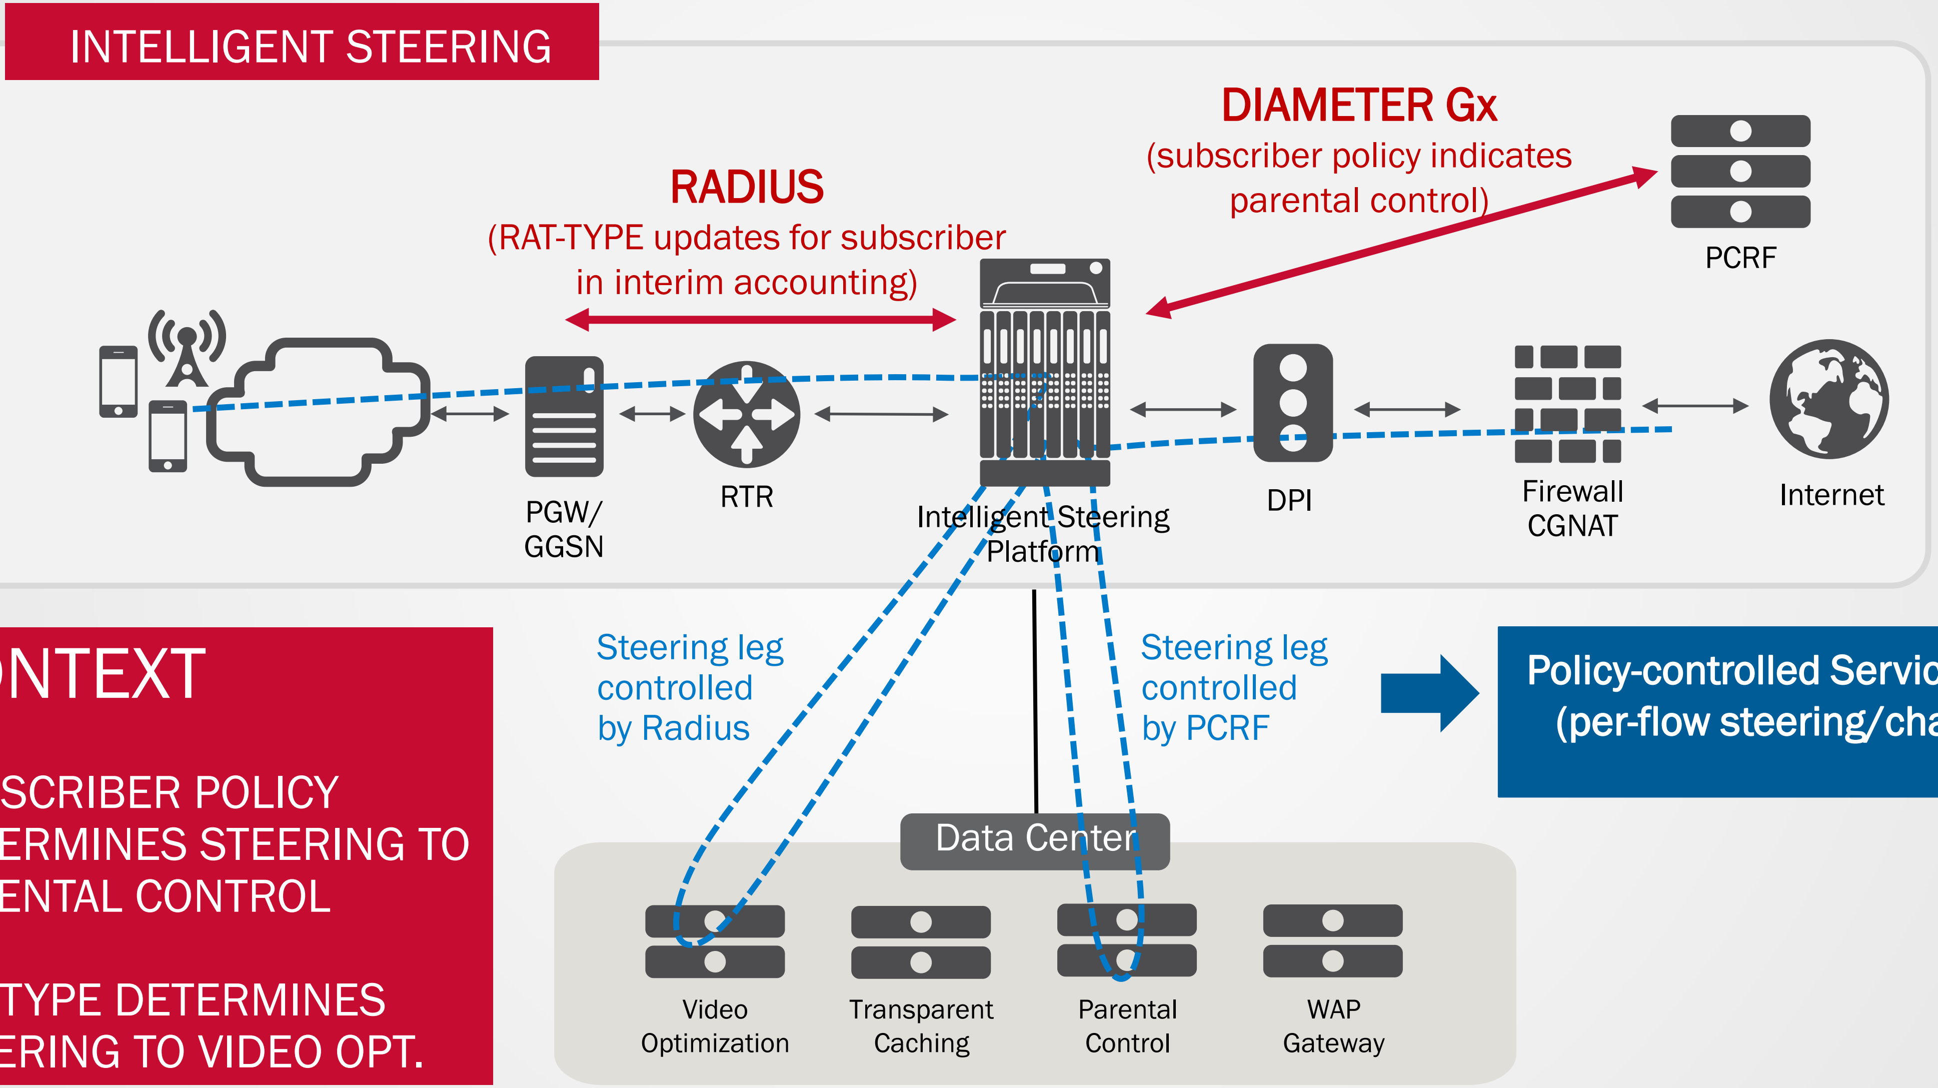

CONTEXT

SUBSCRIBER POLICY DETERMINES STEERING TO PARENTAL CONTROL

RAT-TYPE DETERMINES STEERING TO VIDEO OPT.

Policy-controlled Service Chain (per-flow steering/chaining)

Flow-based vs Transaction-based Traffic Steering

Need for Transaction-based Steering – Video Optimization

- The Service Provider Challenge
  - Video optimization technology is expensive and steering all port 80 traffic to it is not considered economically viable going forward
  - Increasing desire to offload any HTTP traffic that is not carrying video
  - Increasing desire to offload ABR video traffic (as transrating/transcoding no longer needed)
- The Technical Challenge
  - Accurate video detection requires checking both the HTTP request and the response headers
  - If the detection happens at the response level, how can we steer video to video optimizers ‘after-the-facts’ (connection to video server is already established) ?
- The Technical Solution
  - HTTP request-based & response-based steering
  - Per-flow steering is not adequate for this use case (see next slide)

HTTP Messages Differ from IP Packets & TCP Flows

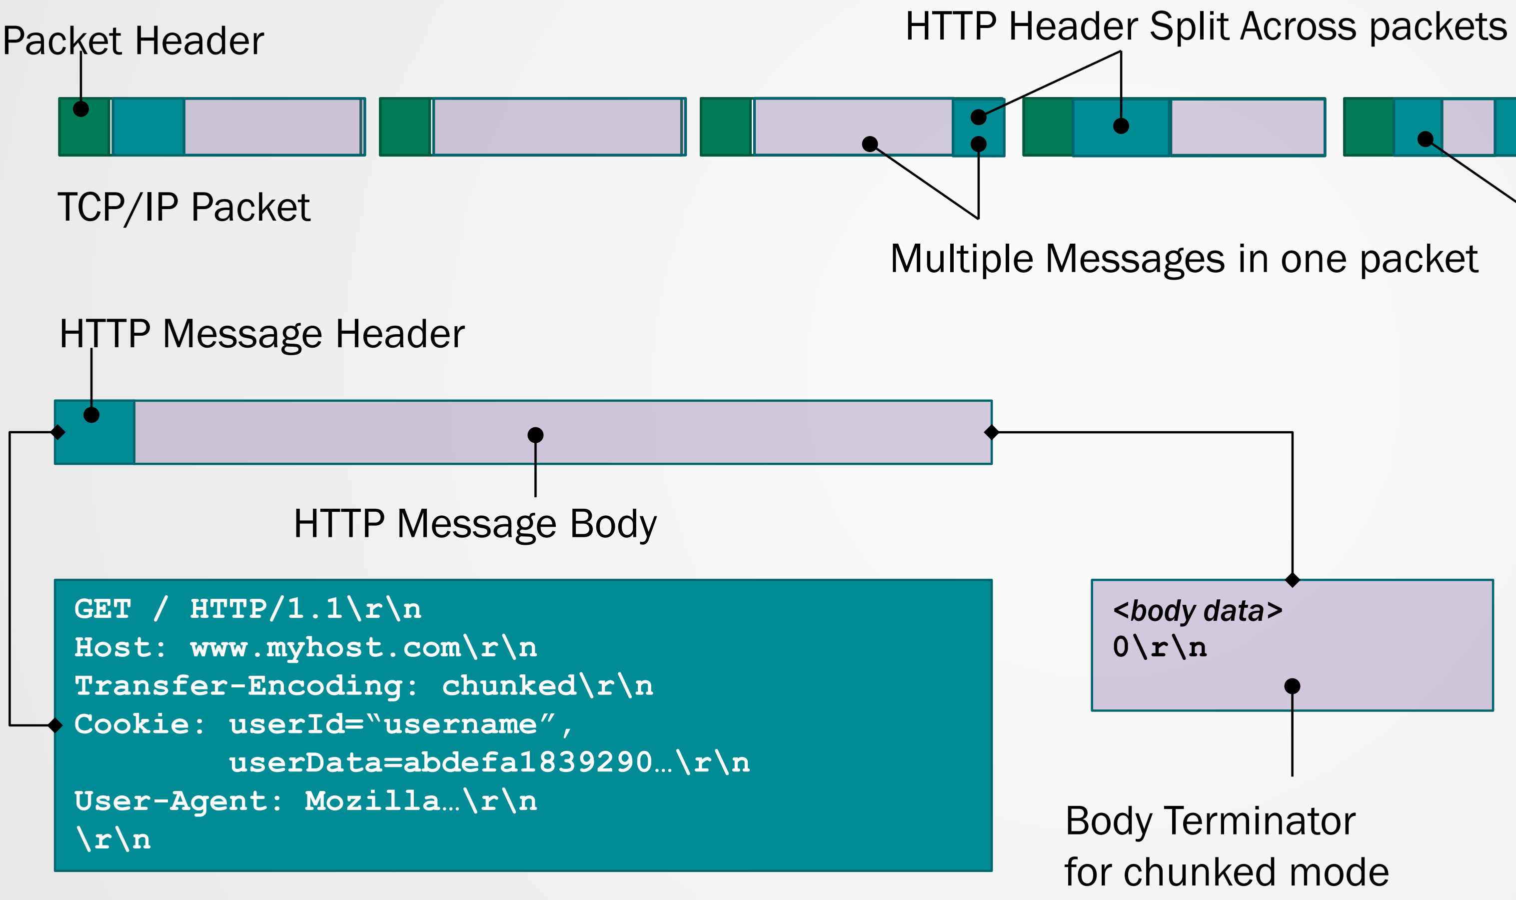

- HTTP message can span multiple packets
- Packets may have multiple HTTP messages
- Delimiting HTTP messages may require inspection of every byte
- Message steering in some cases may cause TCP stream to be split – may lead to chaos in client to end point communication

Steering on HTTP Request & Response

- Steering on request
  - Establish TCP connection with client (full handshake)
  - Accumulate HTTP request message(s) in that TCP connection
  - For each HTTP request message in the TCP connection from the client
    - Parse the HTTP request headers and select VAS based on steering policy
  - Establish new TCP connection with the VAS selected in the steering policy and forward the accumulated HTTP message (in case of service chaining there will be several connections)

- Steering on response
  - Establish TCP connection with client and establish another TCP connection with the server – forward HTTP messages between client and server
  - For each HTTP response message in the TCP connection from the server
    - Parse the HTTP response headers and select VAS based on steering policy
    - But how do we steer to the VAS ? The connection with the server is already established …

Steering on Response – Call flows

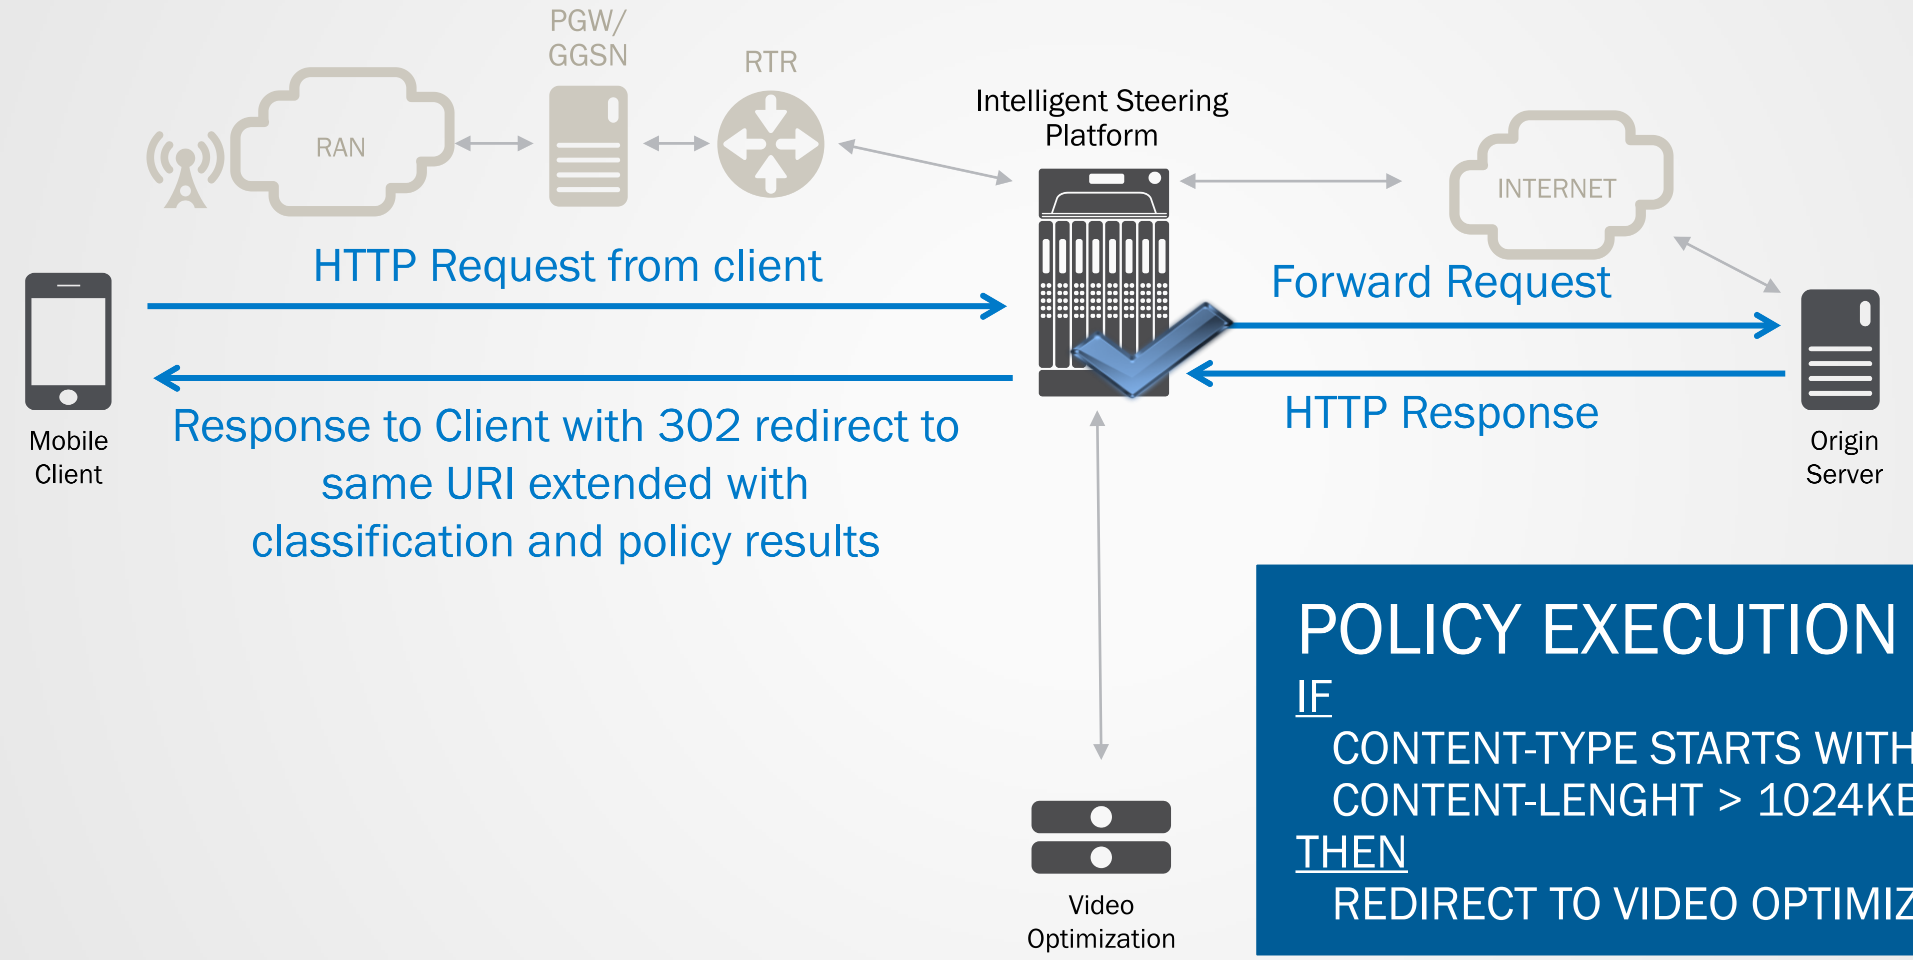

POLICY EXECUTION

<u>IF</u>

CONTENT-TYPE STARTS WITH “VIDEO/”
CONTENT-LENGHT > 1024KB

<u>THEN</u>

REDIRECT TO VIDEO OPTIMIZATION

Steering on Response – After the HTTP redirect

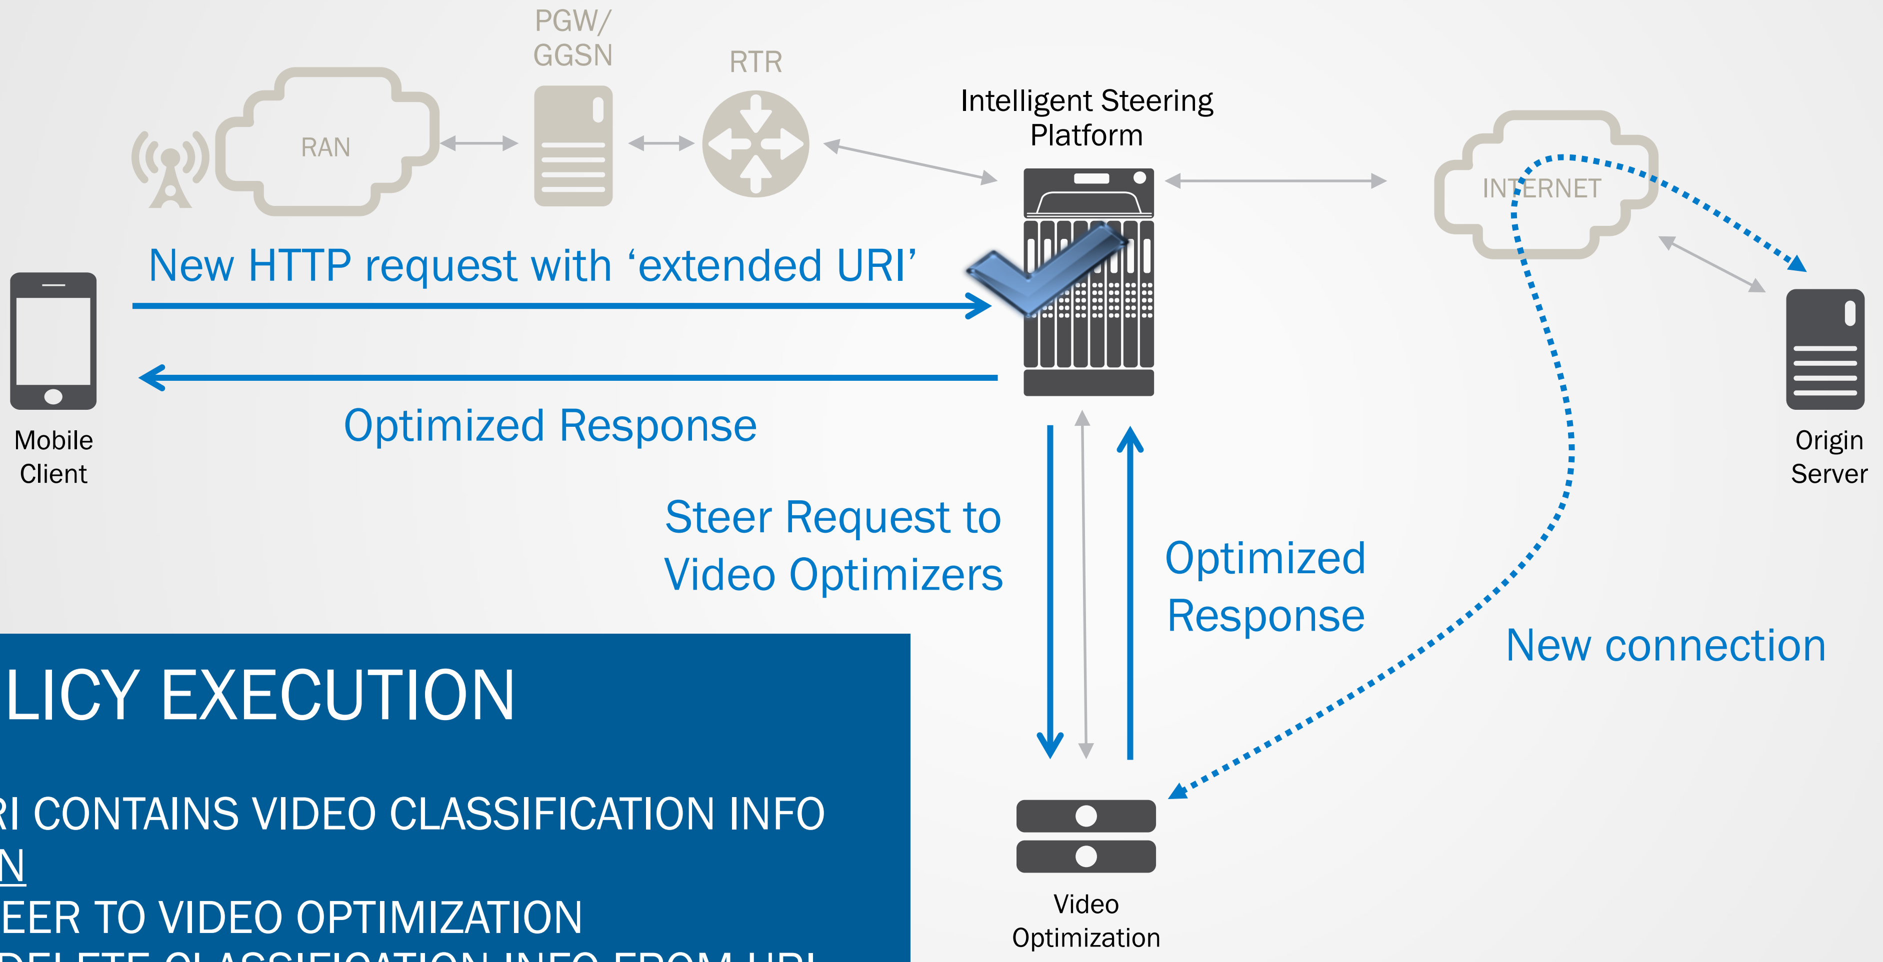

POLICY EXECUTION

IF
URI CONTAINS VIDEO CLASSIFICATION INFO
THEN
STEER TO VIDEO OPTIMIZATION
& DELETE CLASSIFICATION INFO FROM URI

Service Chaining & IETF Activities

IETF – Service Chaining Working Group

- IP networks rely more and more on the combination of advanced functions
  - Besides basic routing and forwarding functions
- Goal : Enforce service-inferred forwarding for traffic traversing a given domain
  - Differentiated by the set of Service Functions to be invoked
  - Service-inferred forwarding is policy-based. Policies may be:
    - Subscriber-aware
    - Based on flow characteristics
    - TE-oriented (e.g., optimize network resource usage)
    - Combination of the above
- Several Service Function Chaining (SFC) IETF drafts available

IETF – Service Function Chaining Examples

- SFC ingress : Policy classification will determine service chain SFC-ID – pointing to a sequence of service functions (SFs)
  - All Service Functions may be policy controlled via a control plane
  - Meta-data can be added to the packets (to convey the SFC-ID to the SFs)
  - Service Functions can be physical or virtual (NFV)
- Packet forwarding between SFs can be plain IP, SDN, overlay networks, …

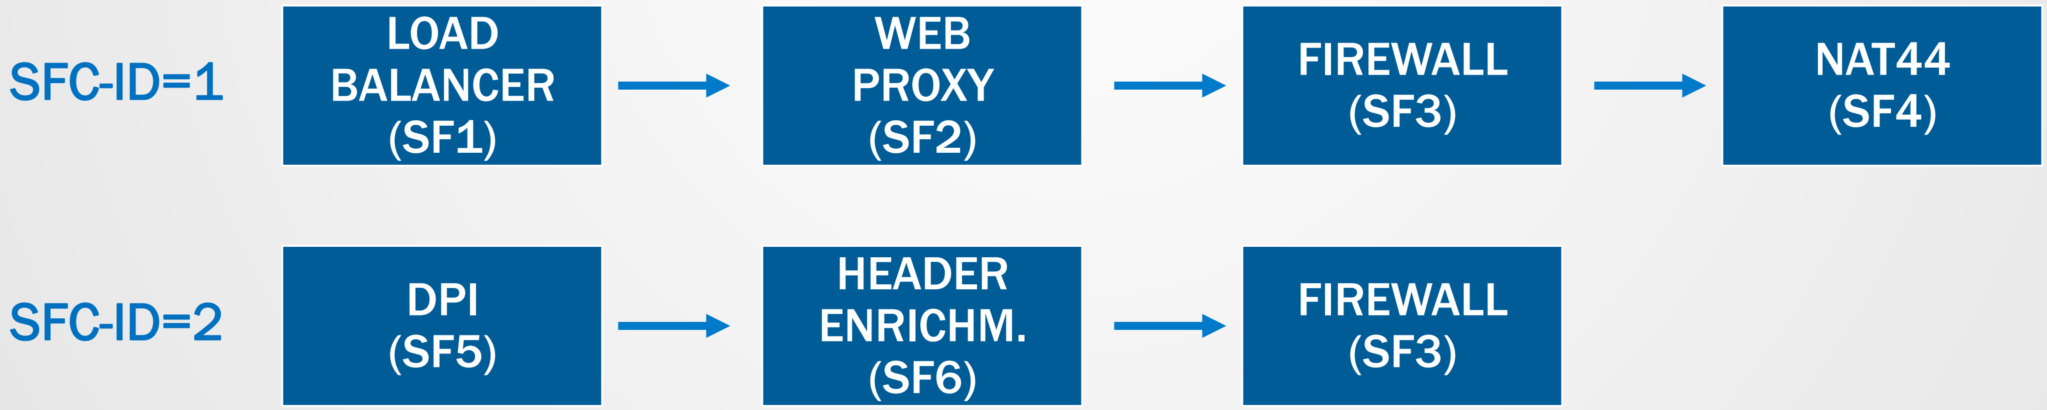

Static & Dynamic Service Chaining – Today

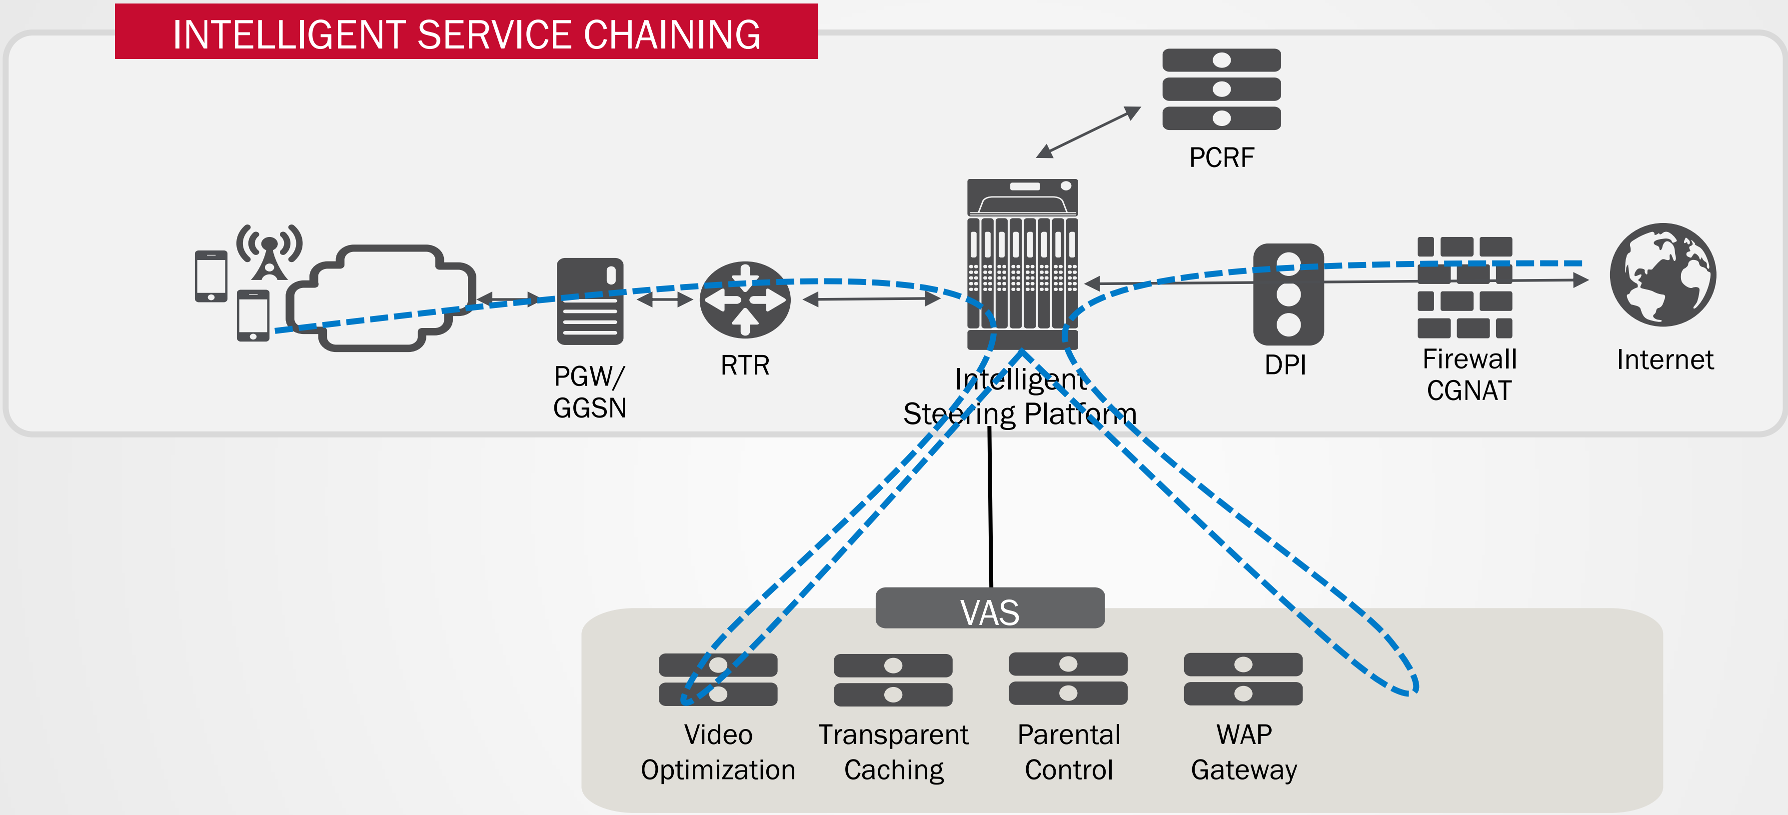

STATIC SERVICE CHAINING

INTELLIGENT STEERING POLICY DEFINES A FIXED SFC (E.G. VAS1-VAS4)

DYNAMIC SERVICE CHAINING

INTELLIGENT STEERING POLICY PER VAS LEG TO FULLY CONTROL THE SERVICE CHAIN ORDER BASED ON STATIC OR DYNAMIC PARAMETERS

Static Service Chaining

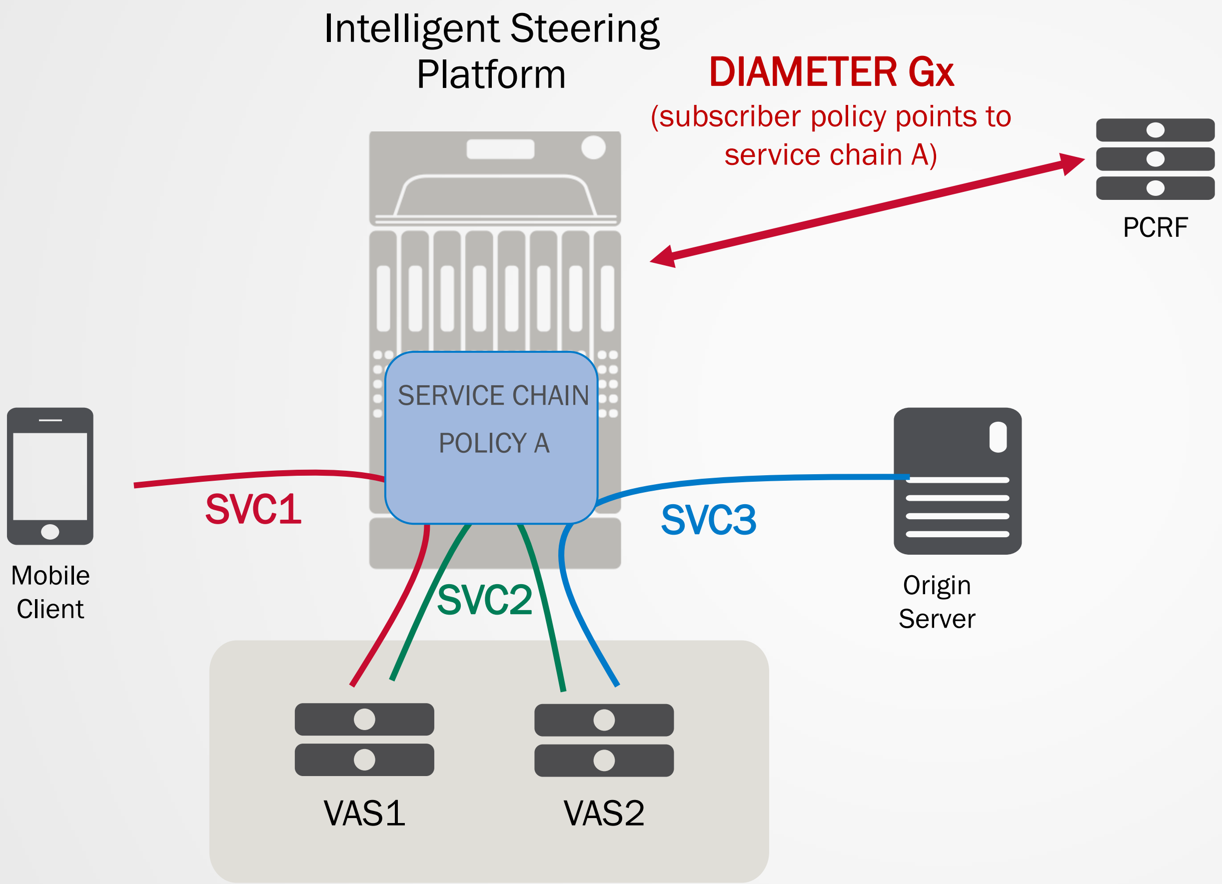

SERVICE CHAIN POLICY A

```
SVC1 {
    FROM SUBSCRIBER
    TO VAS1
}
SVC2 {
    FROM VAS1
    TO VAS2
}
SVC3 {
    FROM VAS2
    TO INTERNET
}
```

Dynamic Service Chaining

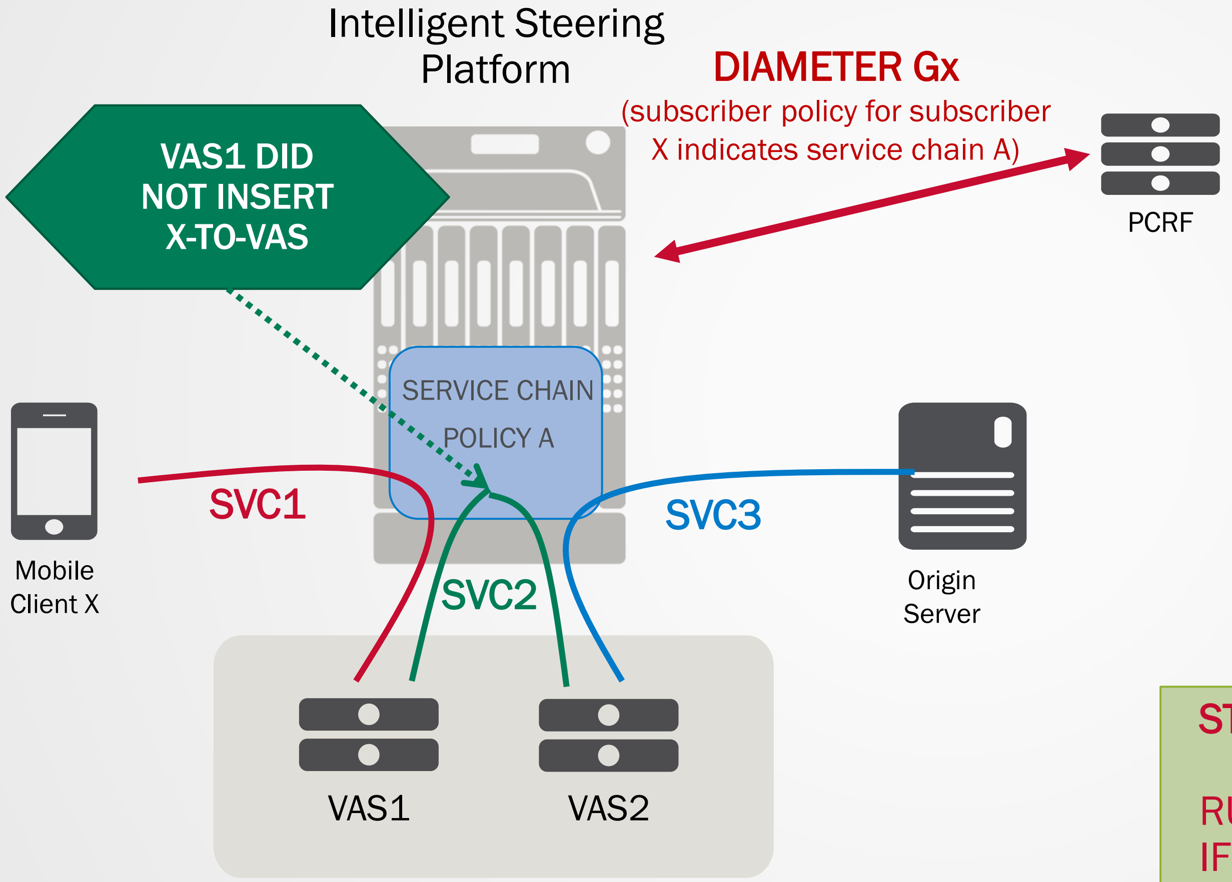

SERVICE CHAIN POLICY A

```
SVC1 {
    FROM SUBSCRIBER
    TO VAS1
}

SVC2 {
    FROM VAS1
    TO VAS2
    STEERING POLICY P1
}

SVC3 {
    FROM VAS2
    TO INTERNET
}
```

STEERING POLICY P1

```
RULE R1 {
IF (HDR $X-TO-VAS == "INTERNET") {
        STEER INTERNET
    }
}
```

Dynamic Service Chaining

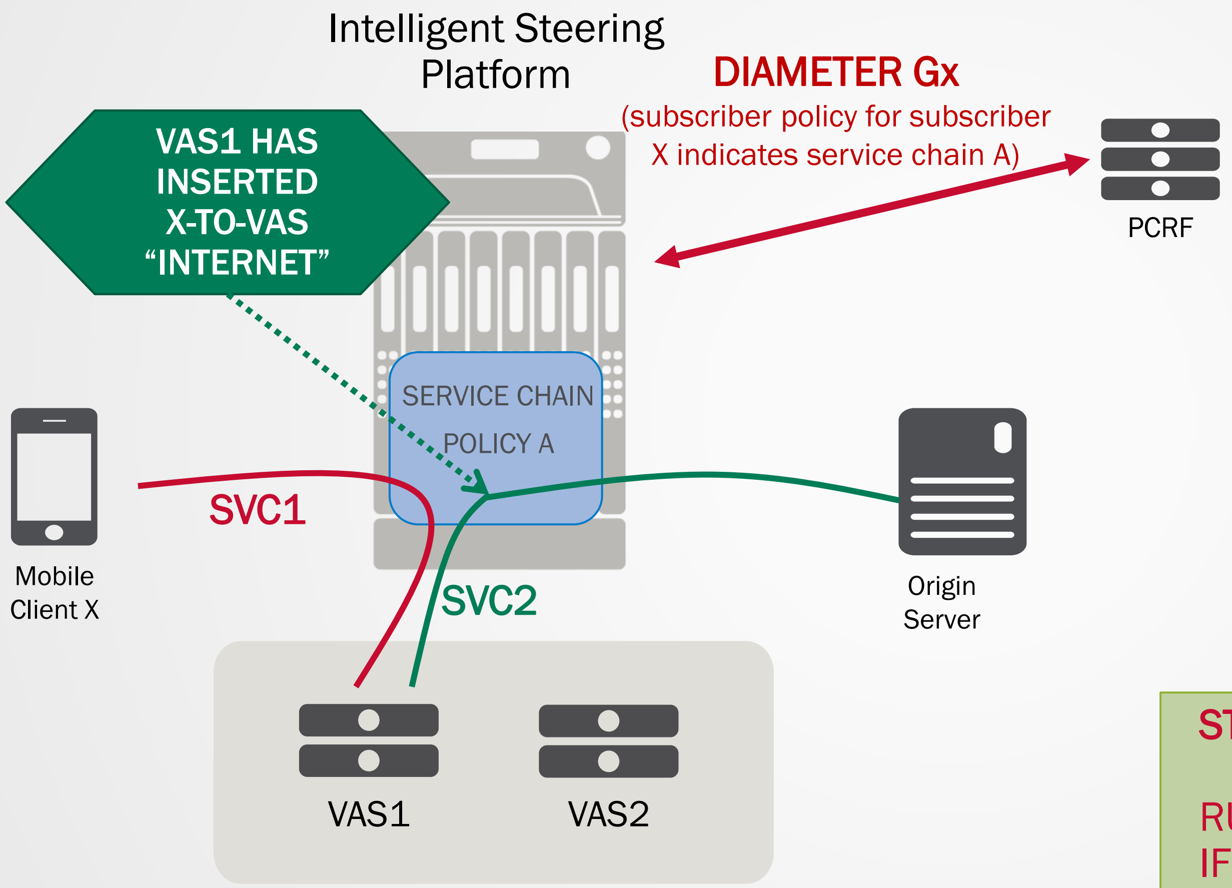

SERVICE CHAIN POLICY A

```
SVC1 {
    FROM SUBSCRIBER
    TO VAS1
}

SVC2 {
    FROM VAS1
    TO VAS2
    STEERING POLICY P1
}

SVC3 {
    FROM VAS2
    TO INTERNET
}
```

STEERING POLICY P1

```
RULE R1 {
IF (HDR $X-TO-VAS == "INTERNET")  {
        STEER INTERNET
    }
}
```

Service Chaining – Packet Forwarding

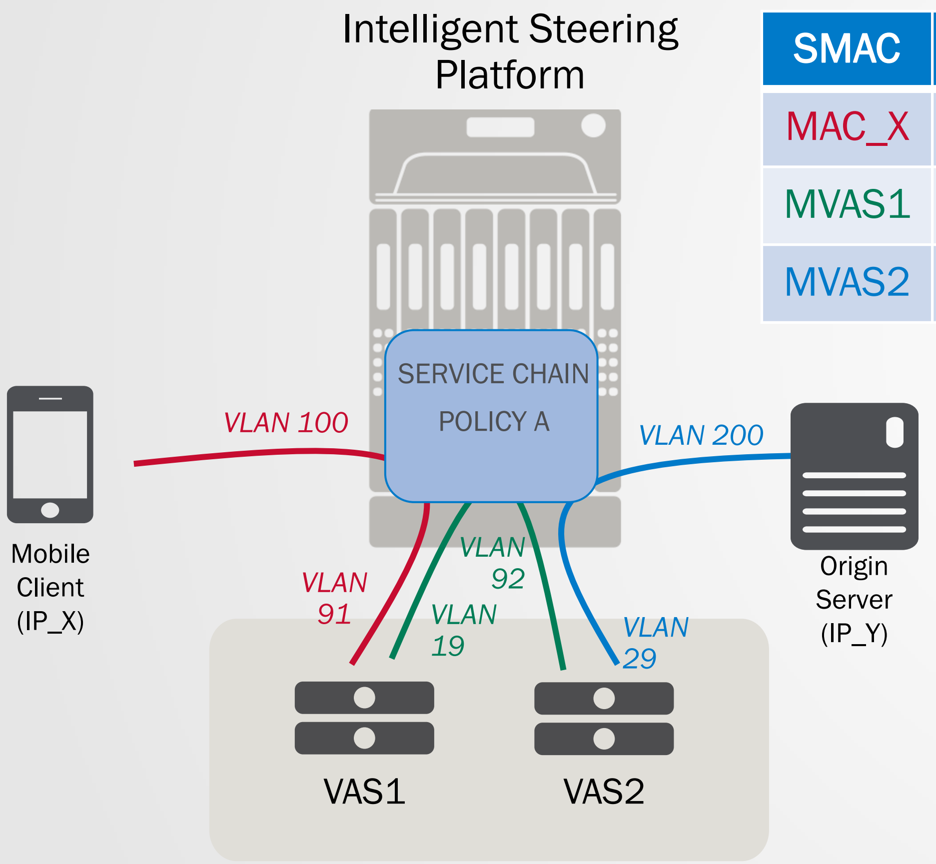

CONNECTION TABLE

IN					OUT				
SMAC	DMAC	SIP	DIP	VLAN	SMAC	DMAC	SIP	DIP	VLAN
MAC_X	M100	IP_X	IP_Y	100	M91	MVAS1	IP_X	IP_Y	91
MVAS1	M19	IP_X	IP_Y	19	M92	MVAS2	IP_X	IP_Y	92
MVAS2	M29	IP_X	IP_Y	29	M200	MAC_Y	IP_X	IP_Y	200

CONNECTION-ORIENTED FORWARDING

- Intelligent steering platform tracks the source MAC address and VLAN of incoming connections in the connection table
- Return traffic from endpoints and/or pools is sent back to the MAC address (on the VLAN) that transmitted the request

Service Chaining – Today and Future

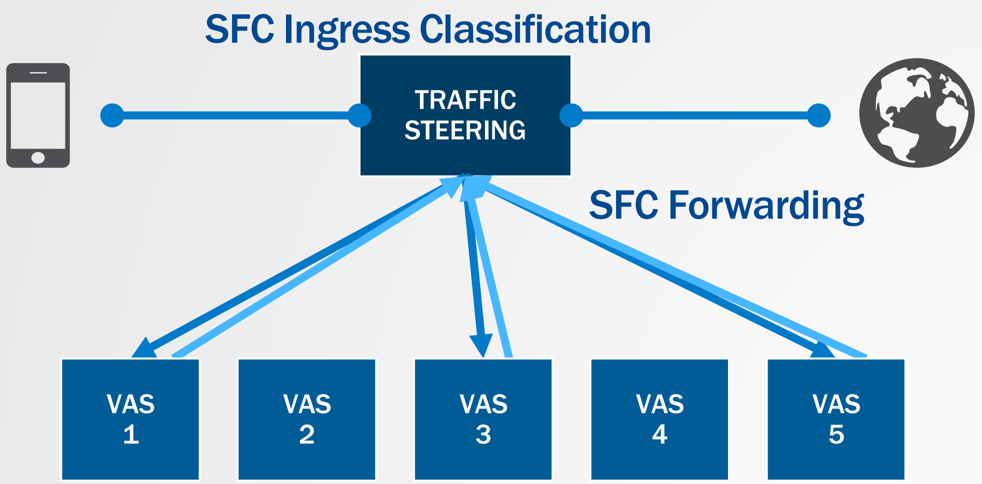

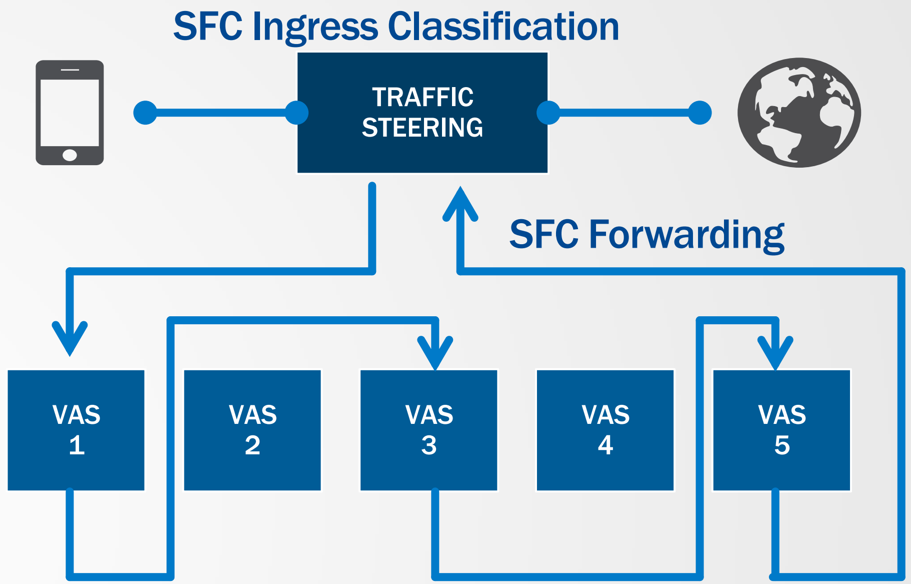

- Available today – TCP & HTTP proxy technology
- Flexible use of ‘steering headers’ towards VAS platforms (HTTP headers, DSCP, …)
- Works with ICAP as well (control plane steer)
- Practical model for few VAS services

- Discussed in several IETF drafts
- Requires all vendors to agree on same standard (packet header for metadata)
- How to leverage SDN/NFV and overlay networking (VXLAN, NVGRE) technology
- Scales to many VAS services

TCP Optimization

TCP Proxy – Optimizing both sides of the TCP connection

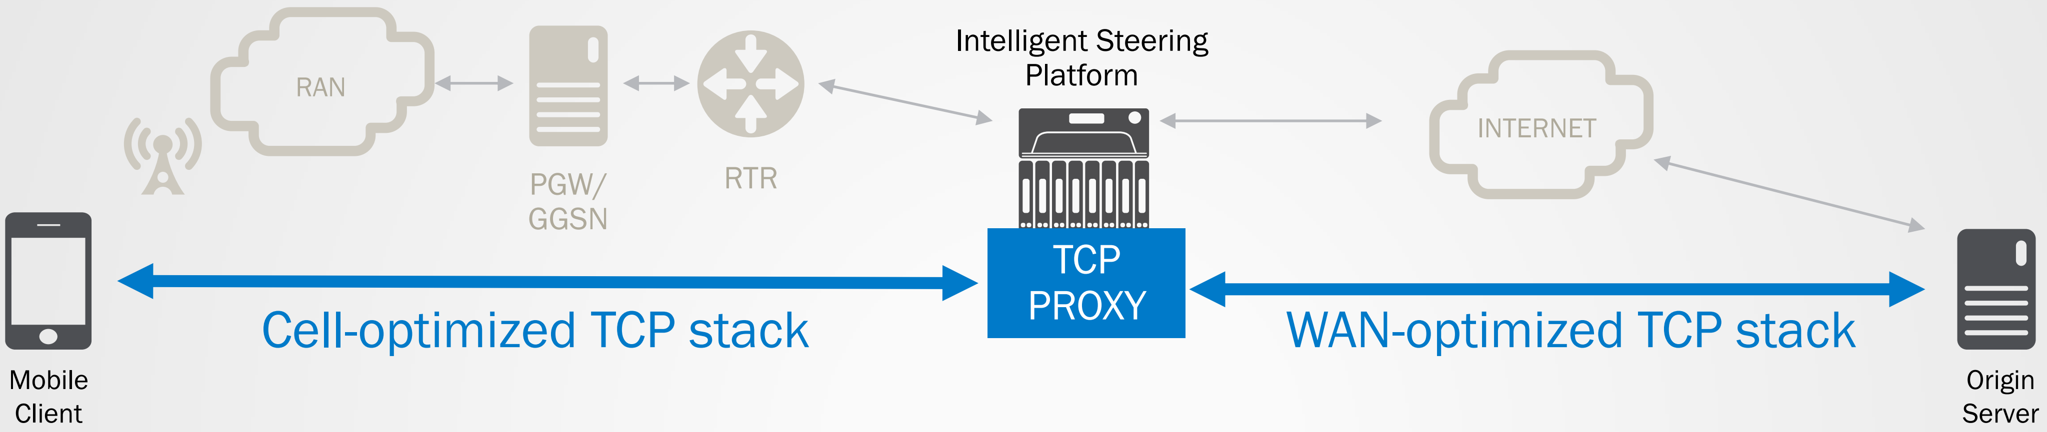

- TCP proxy approach allows for adequate TCP options & window scaling parameters to be negotiated separately with the client and the server, optimized for the access technology
  - Window scaling
  - Selective ACK
  - Congestion control mechanisms, Nagle algorithm, etc.

- Patent pending optimizations to deal with packet loss & delay specific to cellular networks
  - Remove the effect of the first few percent of packet loss on congestion control typical for 2G/3G
  - Avoiding the buffer bloat problems in LTE networks

Ideal TCP stacks would result in …

High Goodput

Minimal Buffer Bloat

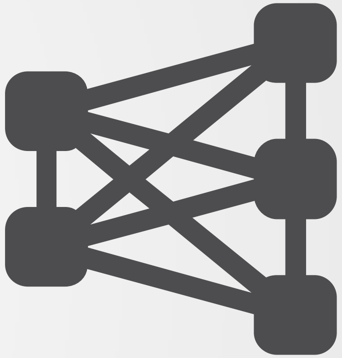

Flow Fairness

HOW DO WE ACHIEVE THIS IN 2G, 3G AND 4G NETWORKS ?

Impact of Latency : Web Page Load Times

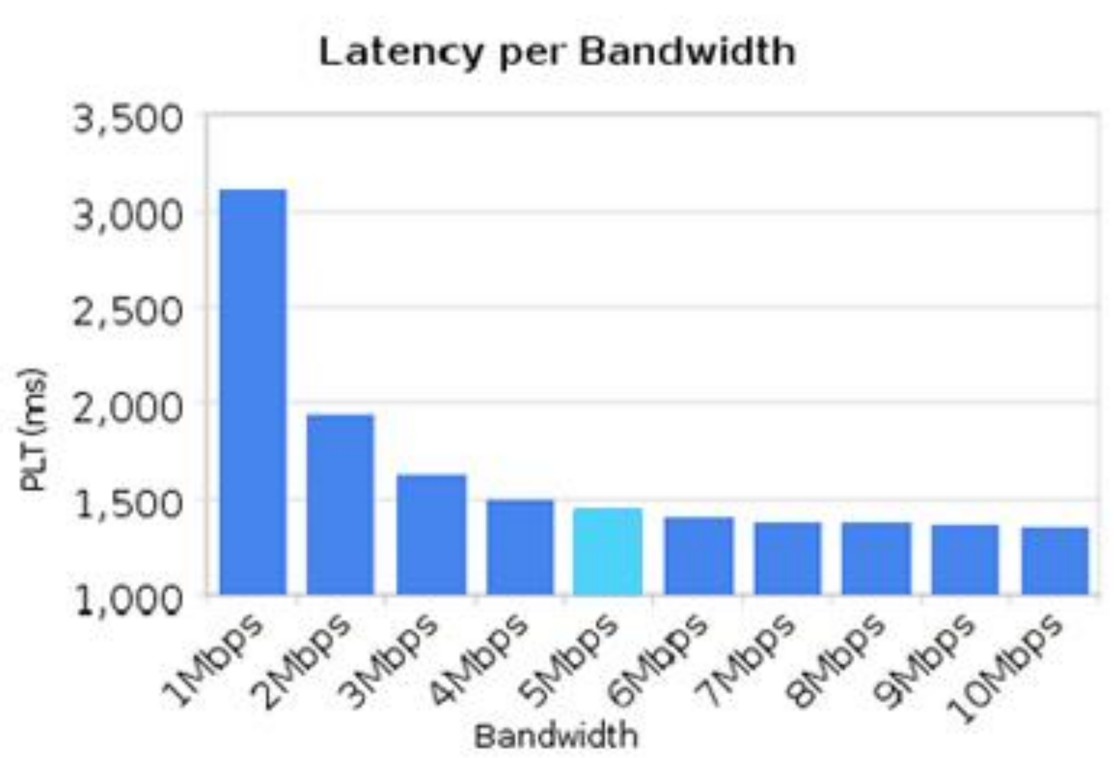

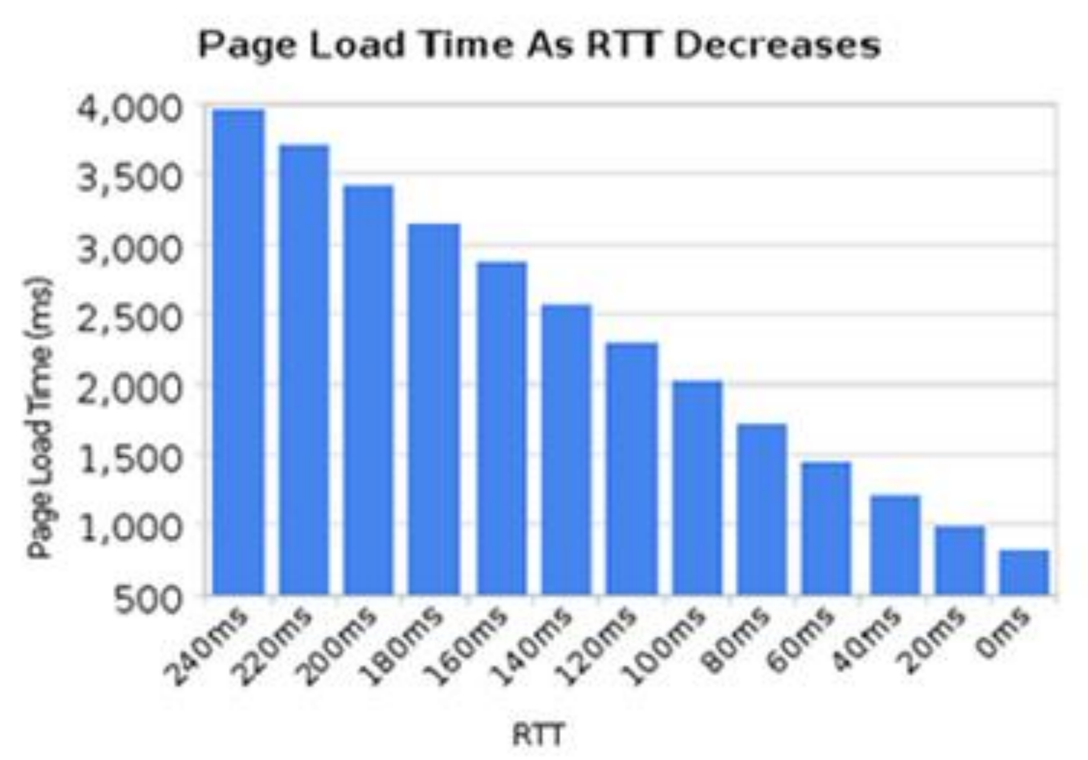

Bandwidth doesn't matter (much) - Google

Slide courtesy of Ilya Grigorik @ Google:
http://www.igvita.com/slides/2012/webperf-crash-course.pdf

Impact of Packet Loss : Throughput Degradation

- TCP is designed to probe the network to figure out available capacity
- TCP slow start is a feature, not a bug

In mobile networks packet loss does not necessarily imply congestion

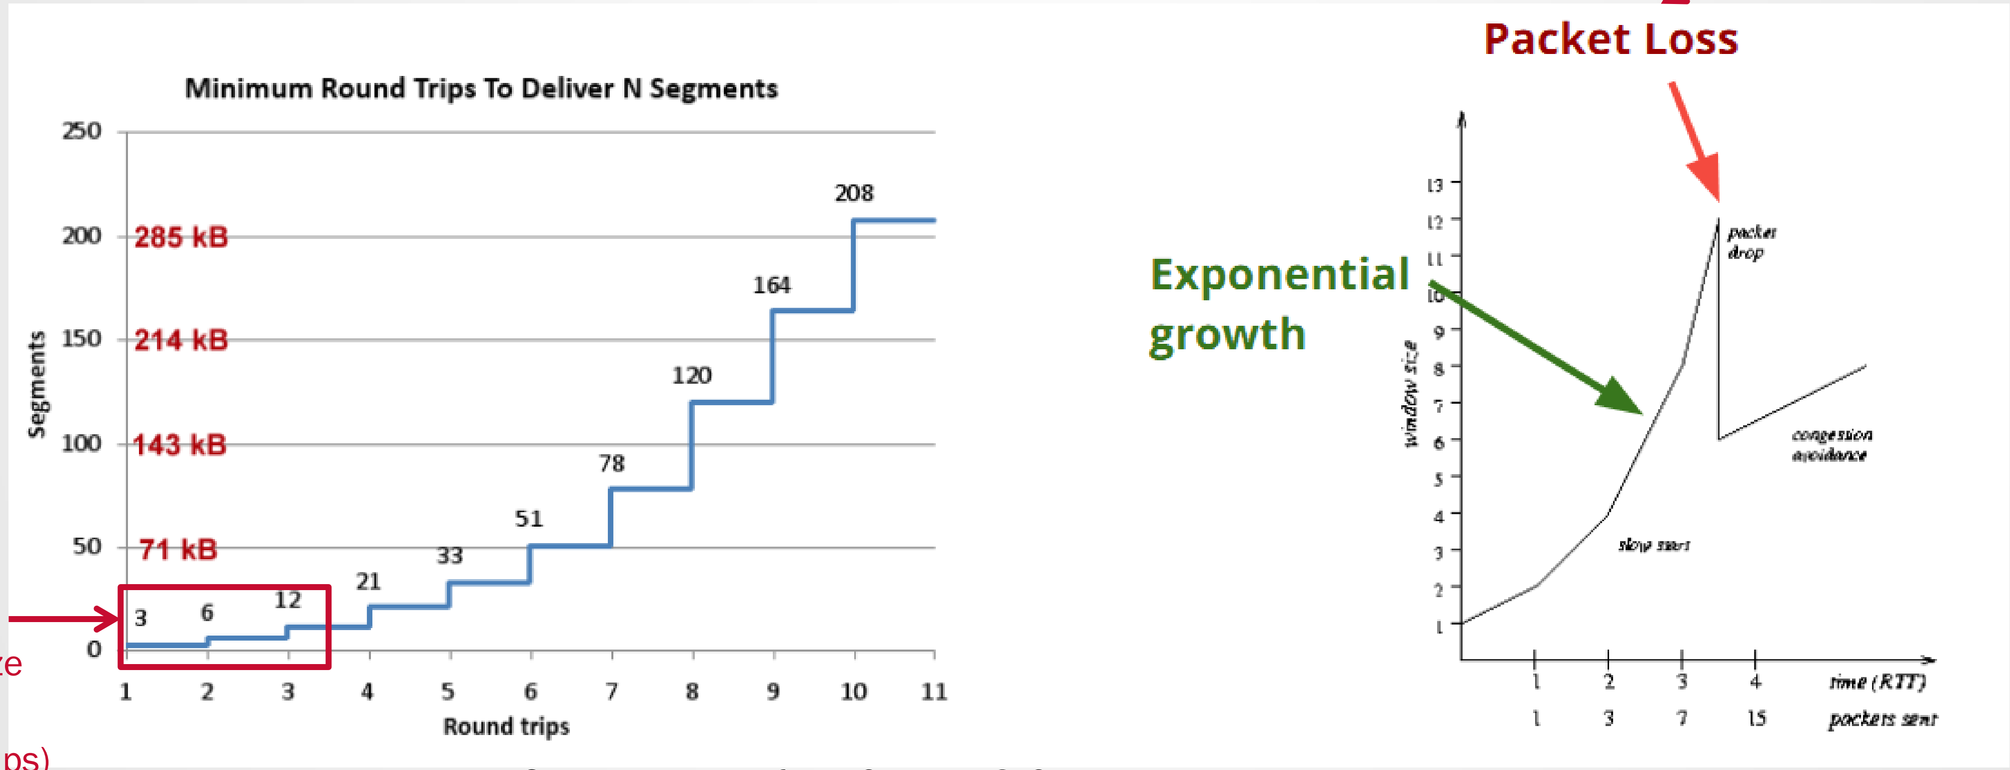

Avg HTTP response size 16kB

(so 3 round trips)

Slide courtesy of Ilya Grigorik @ Google:
http://www.igvita.com/slides/2012/webperf-crash-course.pdf

TCP Congestion Control Algorithms in 3G and LTE

TCP Woodside

- F5 created algorithm.
- Hybrid loss and latency based algorithm.
- Minimizes buffer bloat by constantly monitoring network buffering.

TCP Vegas

- Emphasizes packet delay rather than packet loss
- Detects congestion based on increasing RTT values of packets.

TCP Illinois

- Targeted at high speed long distance networks
- Loss-delay based algorithm.
- Primary congestion of packet loss determines direction of window size change.
- Secondary congestion of queuing delay determines the pace of window size changes.

H-TCP

- Targeted for high speed networks with high latency.
- Loss-based algorithm.

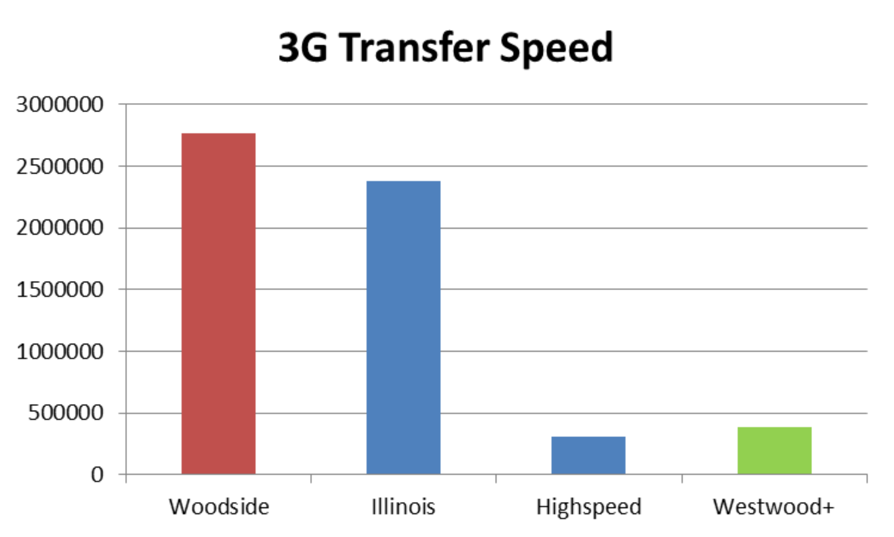

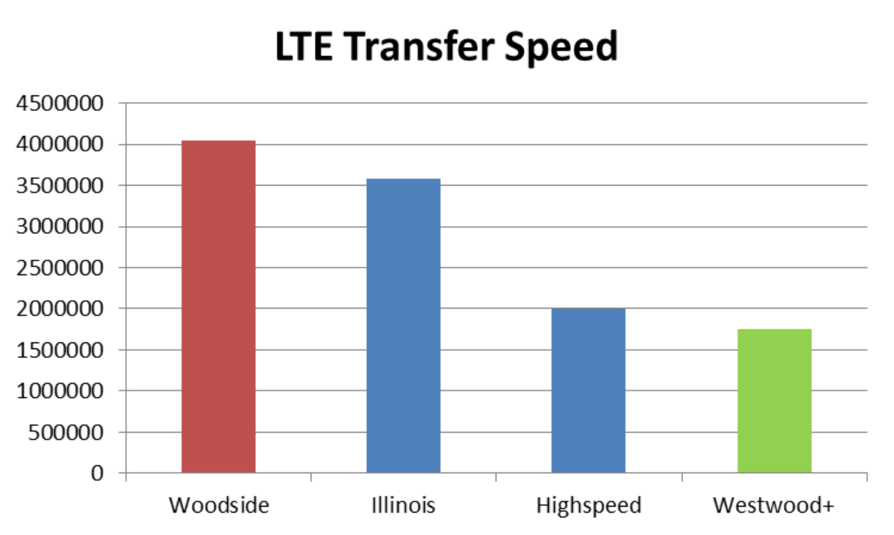

Reducing Web Page Load Times with TCP Optimization

Real life test results – MNO in APAC

Case 1 – 100 * 64KB images

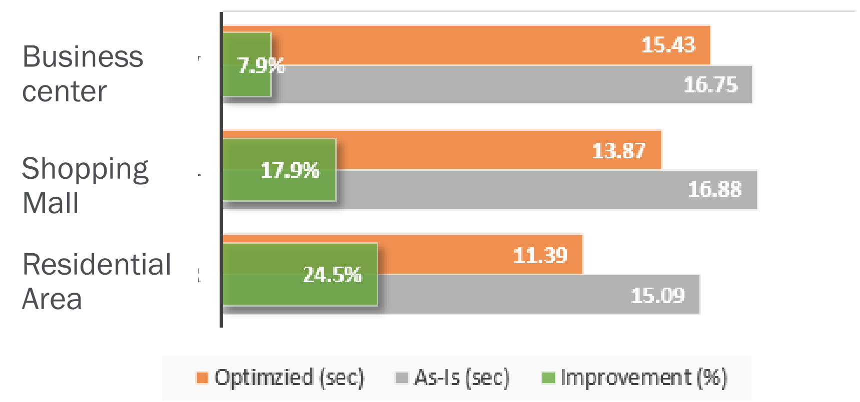

Case 2 – 1 * 10MB image

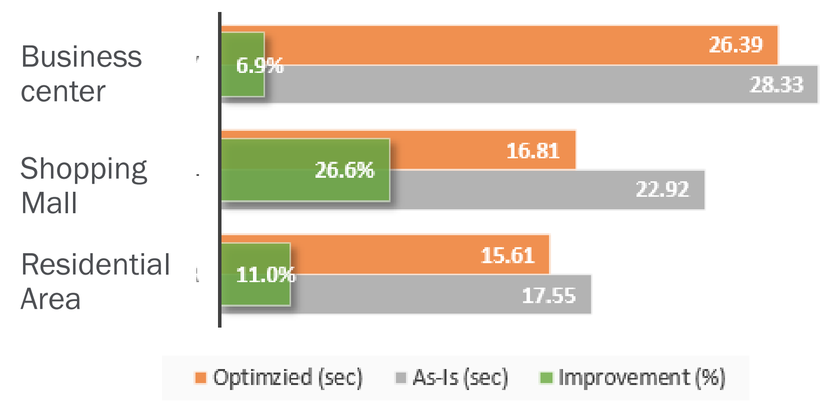

Case 3 – Regular website 1

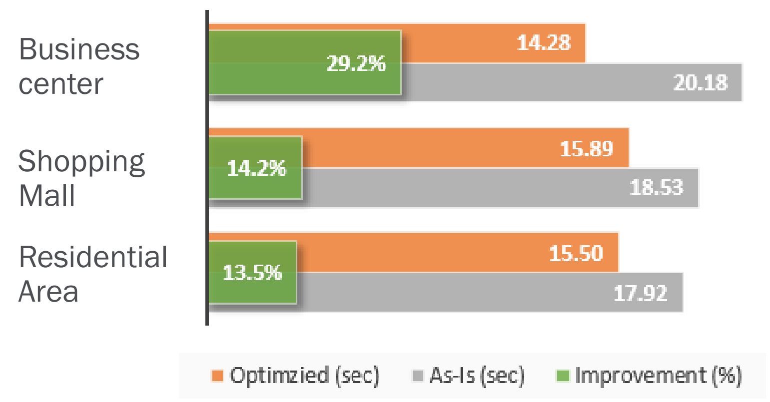

Case 4 – Regular website 2

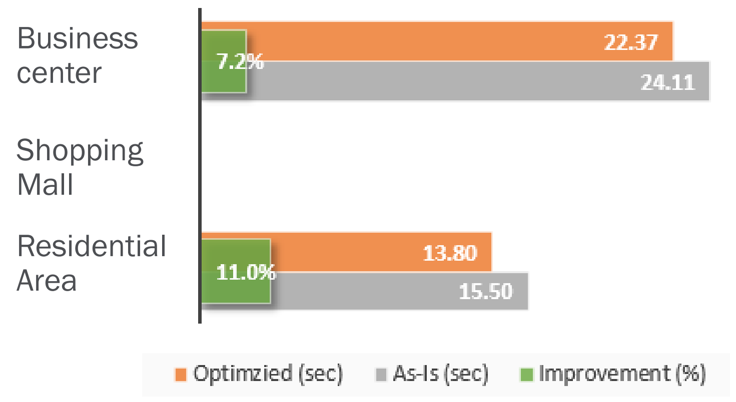

Summary

Traffic Optimization with TCP & HTTP Proxy

- Allows for Policy-based Intelligent Traffic Steering
  - Offloading & cost optimizing the VAS infrastructure

- Allows for Static and Dynamic Service Chaining Today
  - Avoiding to pipe all traffic through all VAS platforms in sequence

- Allows for Enhancing the Mobile Subscriber's Quality of Experience
  - Advanced TCP optimization techiques increases the "goodput" and user experience over the 2G, 3G and LTE radio infrastructure

CONFIDENTIAL