AGILITY 2014

SOLUTIONS
FOR AN APPLICATION WORLD

Traffic Steering & Service Chaining

Optimize & Monetize with PEM

Bart Salaets

Solution Architect

Agenda

- F5 Gi LAN Strategy
- Traffic Steering & Service Chaining
  - Recent Evolutions
  - Policy-Based “Per-Flow” and “Per-Transation” Steering
  - Static & Dynamic Service Chaining : Evolution to IETF model
- Policy Enforcement Manager
  - Traffic Classification
  - Policy Actions
- Evolution towards NFV
- Summary

F5 Gi LAN Strategy

F5 in the S/Gi network – A Consolidated Approach

Simplifying the delivery of network services

BEFORE F5

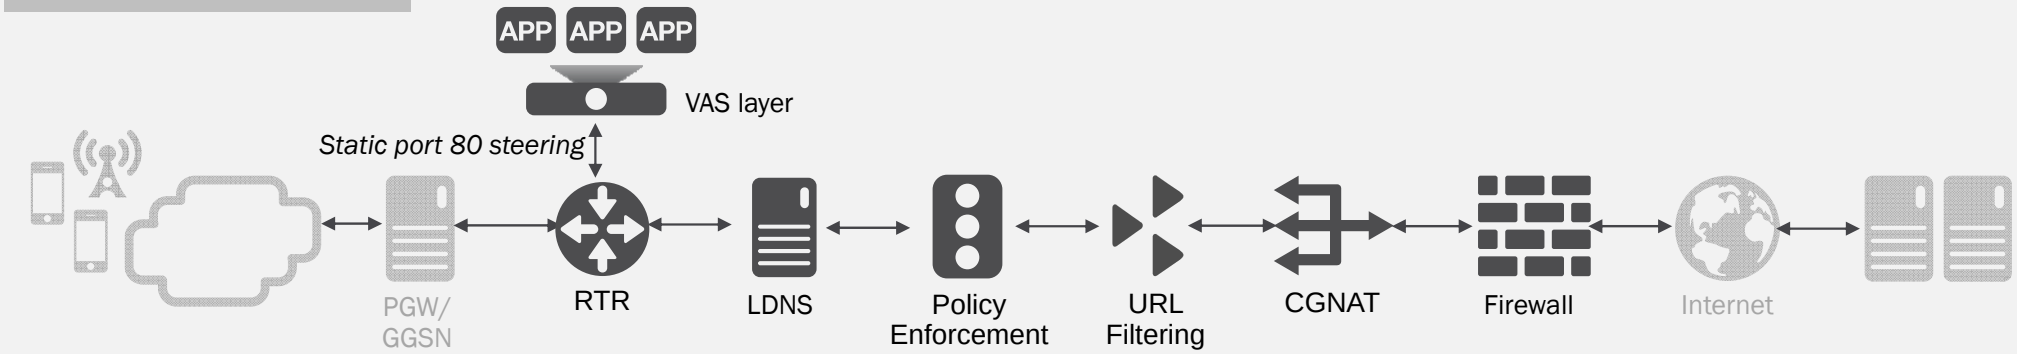

WITH F5

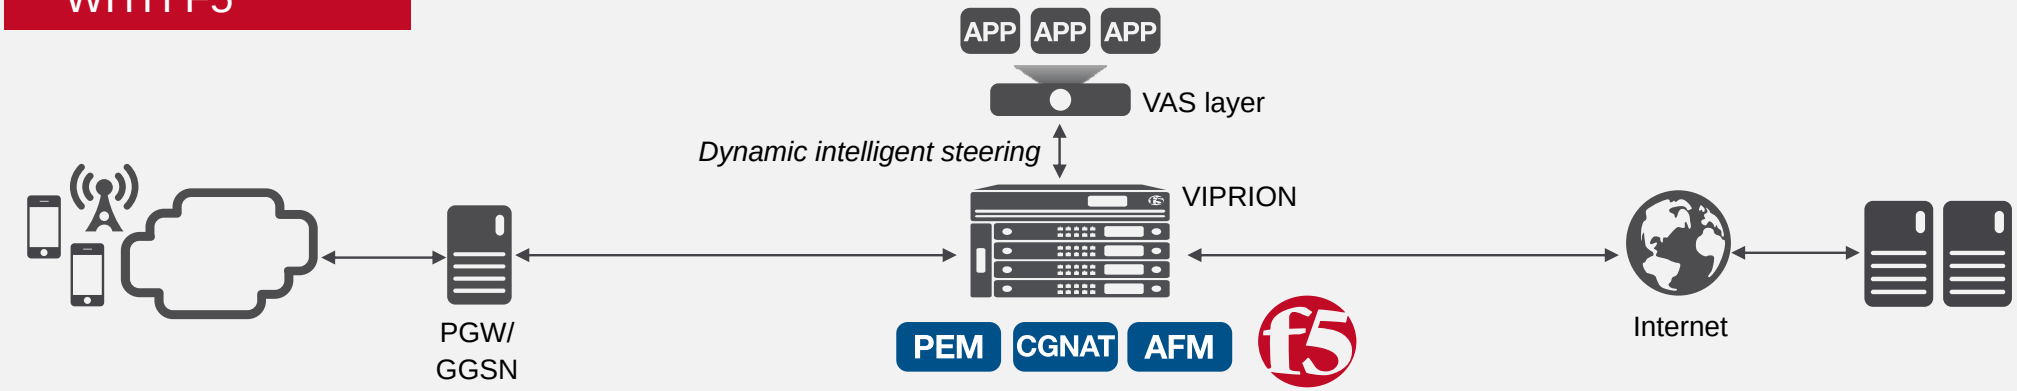

Traffic Steering
&
Service Chaining

Traditional Steering to VAS & Optimization platforms

A router steers all port 80 traffic to VAS platforms

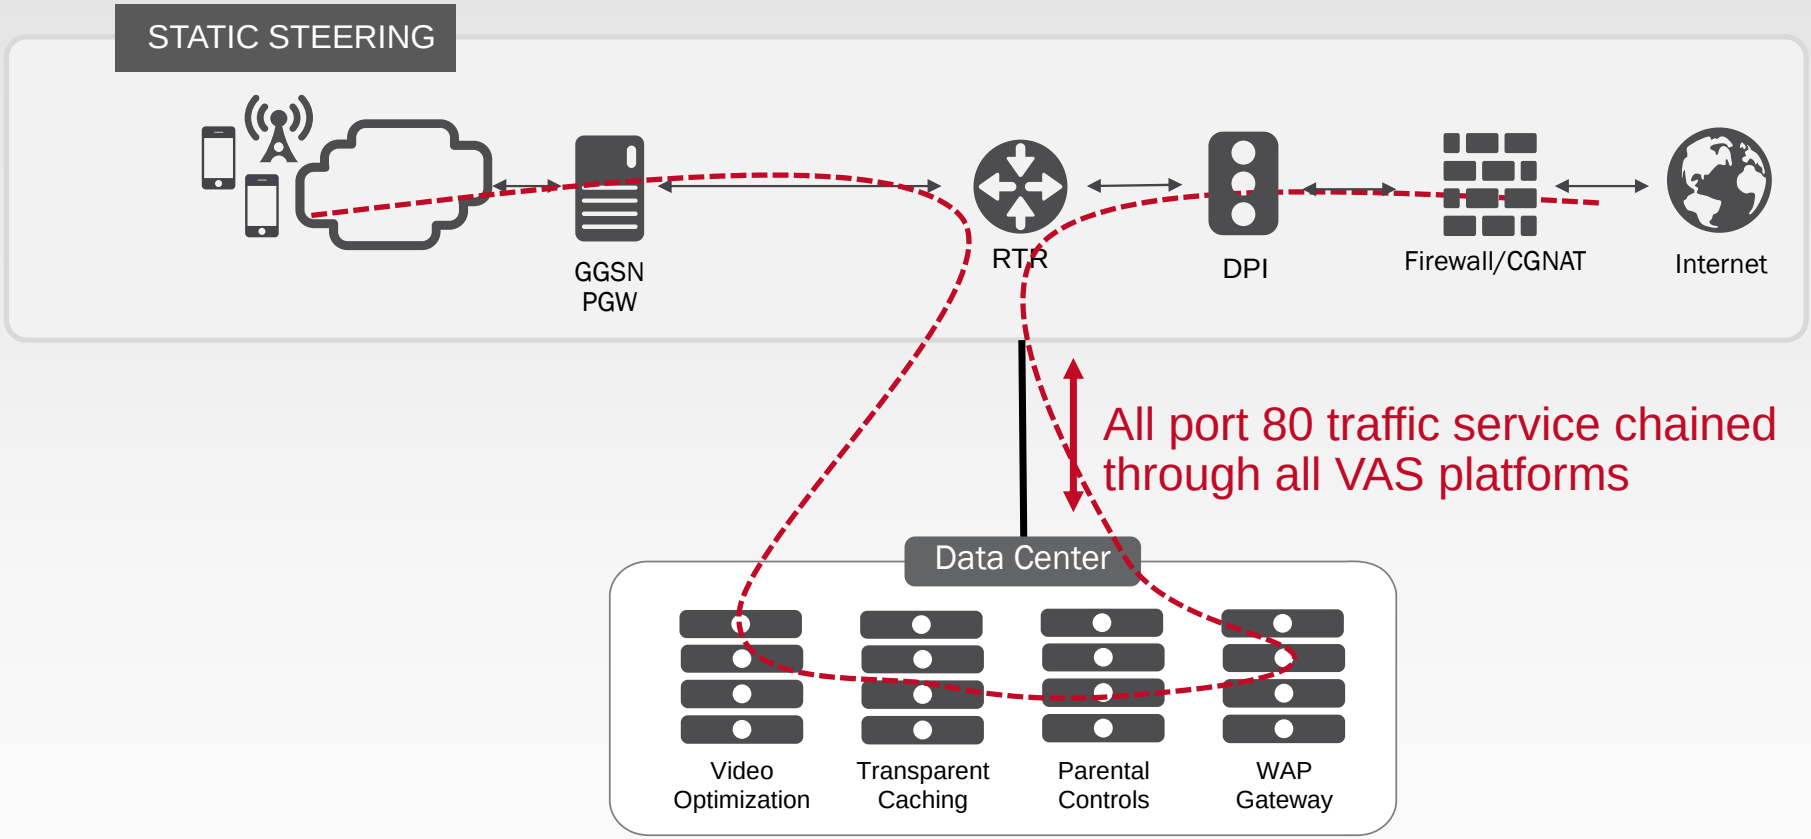

Challenges with the traditional approach

- All VAS platforms need to classify and process all port 80 traffic
  - Waste of resources for pass-through traffic
  - Duplication of resources
- Time-to-market challenge for new services
  - Network integration
  - PCC integration
- Video optimization and transparent caching requires a fresh look
  - Cost benefit or QoE mgmt tool
  - Rise of ABR video (HLS, MPEG-DASH)
  - Increase in SSL / SPDY

Challenges with the traditional approach

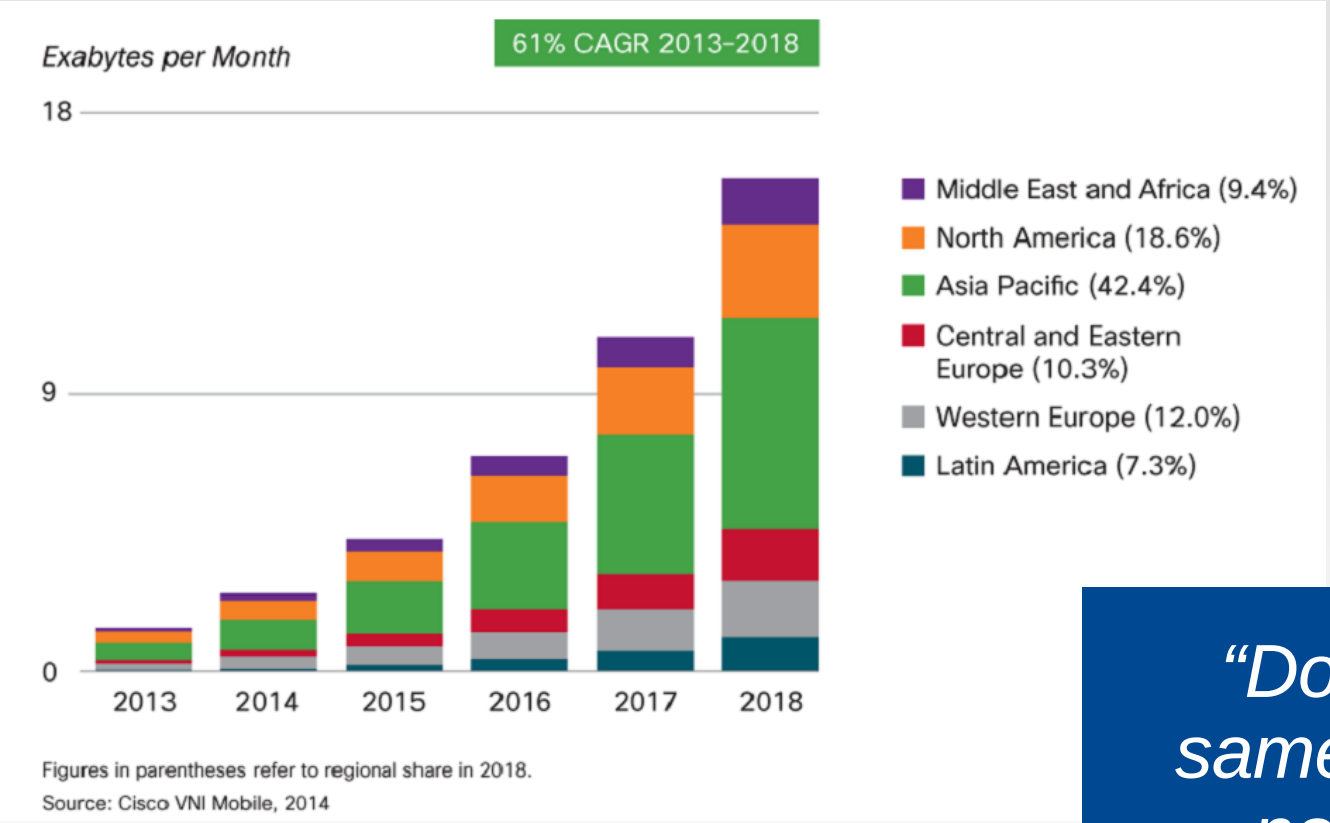

Figures in parentheses refer to regional share in 2018.

Source: Cisco VNI Mobile, 2014

"Doing more of the same is economically no longer viable"

Solution : Intelligent traffic steering to VAS platforms

Offloading VAS services & Optimizing infrastructure utilization

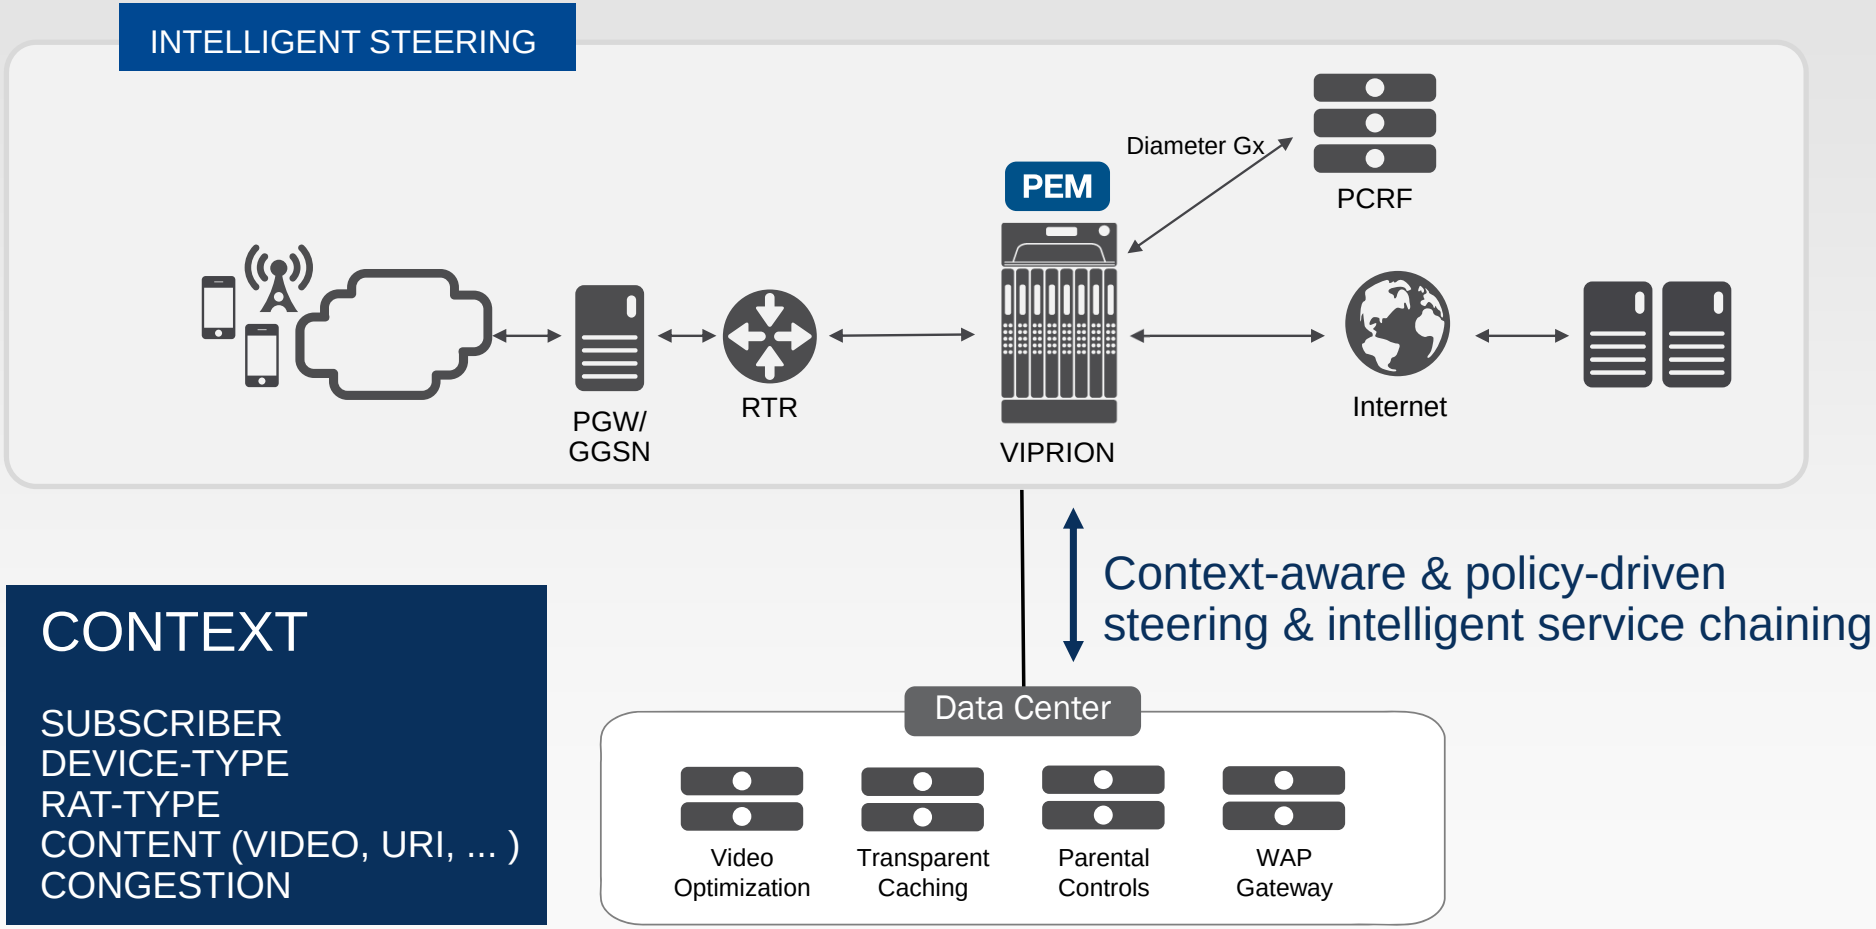

Policy-Based “per-flow” Steering

- Use cases – Steering policy for each flow dependent on ...
  - Subscriber policy
  - RAT-Type (2G / 3G / LTE / Wifi)
  - Location (roaming)
  - Network congestion
  - Device-type (IMEI, HTTP User-Agent)
  - Content-Type (HTTP Content-Type, DPI signature)

  ... or any combination of the above
- Control Plane interactions
  - Diameter Gx (from PCRF)
  - Radius (from GGSN)
  - Custom API (eg for congestion based steering)

Policy-based Flow Steering in Action

Steering to 2 VAS services : Subscriber + RAT-type based

User	Service Policy
John	Video Optimization “LTE bypass” Parental Control “No”
Paul	Video Optimization “Always” Parental Control “Yes”

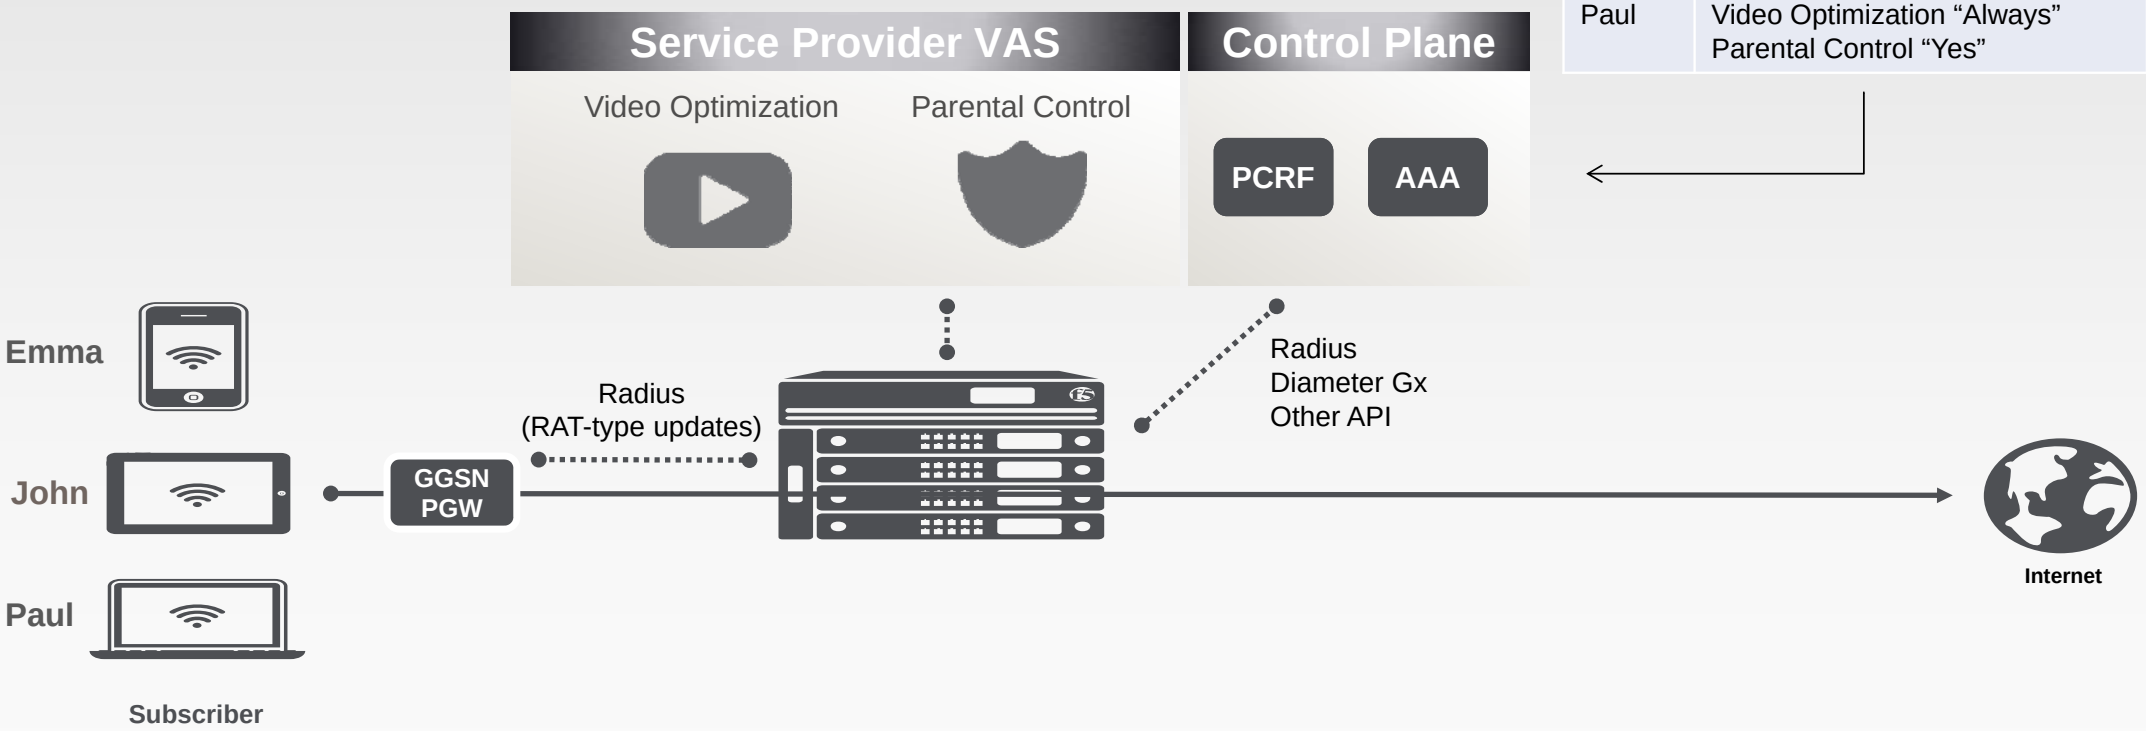

Policy-based Flow Steering in Action

User John : http traffic in LTE

User	Service Policy
John	Video Optimization “LTE bypass” Parental Control “No”
Paul	Video Optimization “Always” Parental Control “Yes”

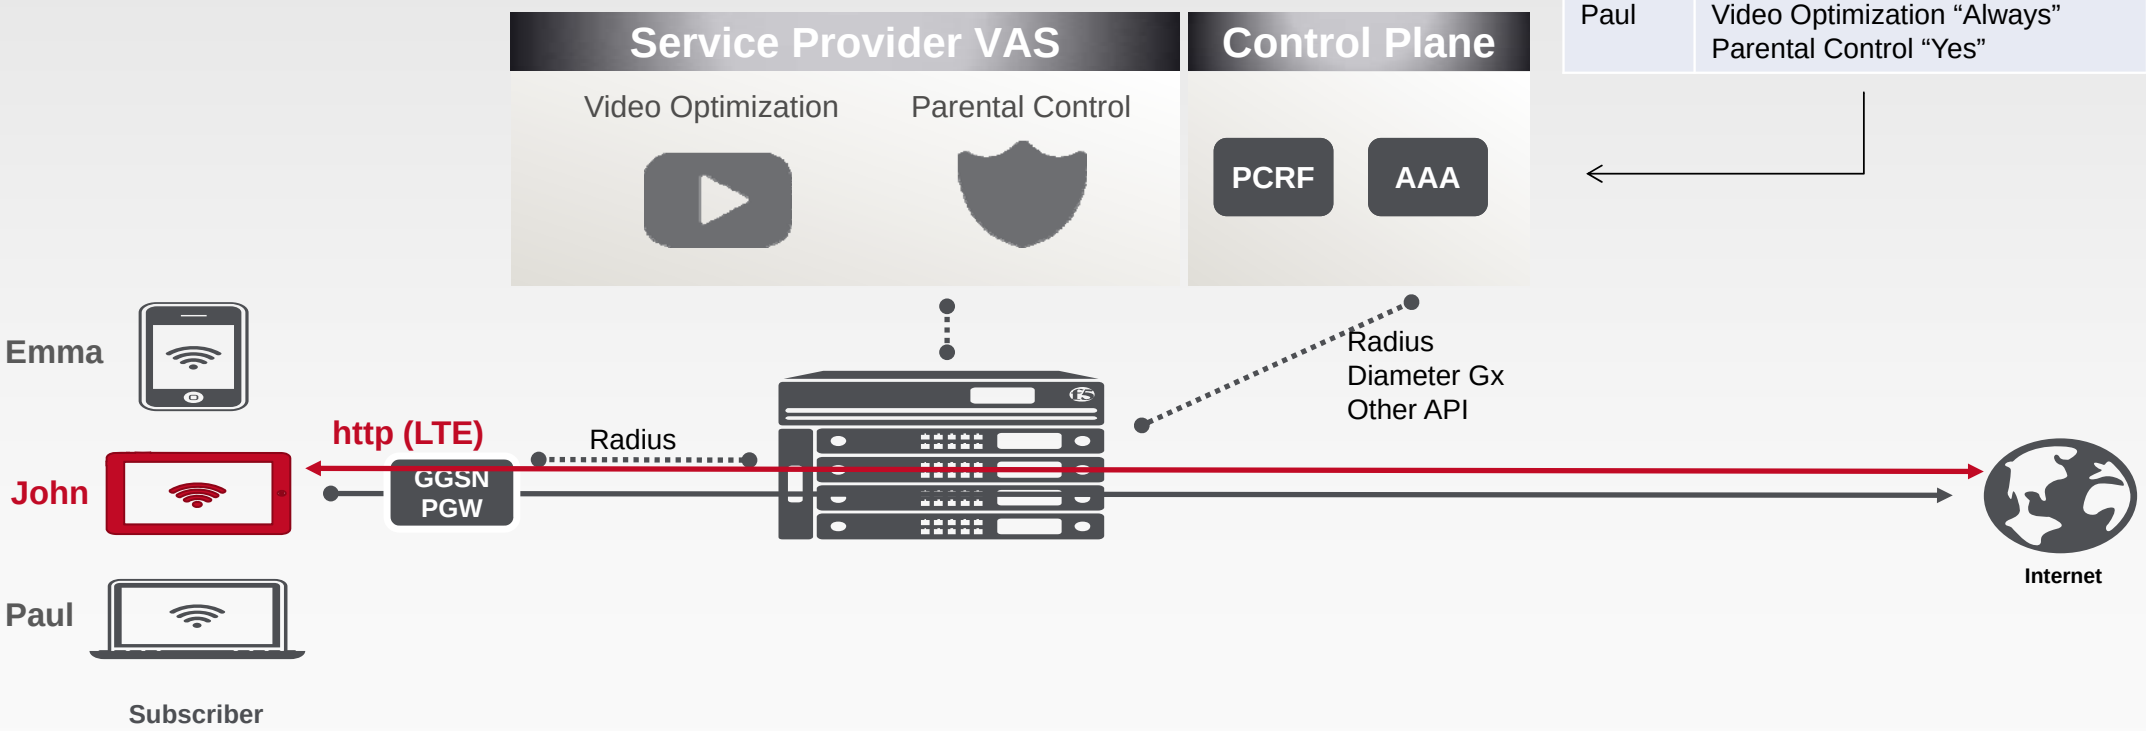

Policy-based Flow Steering in Action

User John : http traffic in 3G

User	Service Policy
John	Video Optimization “LTE bypass” Parental Control “No”
Paul	Video Optimization “Always” Parental Control “Yes”

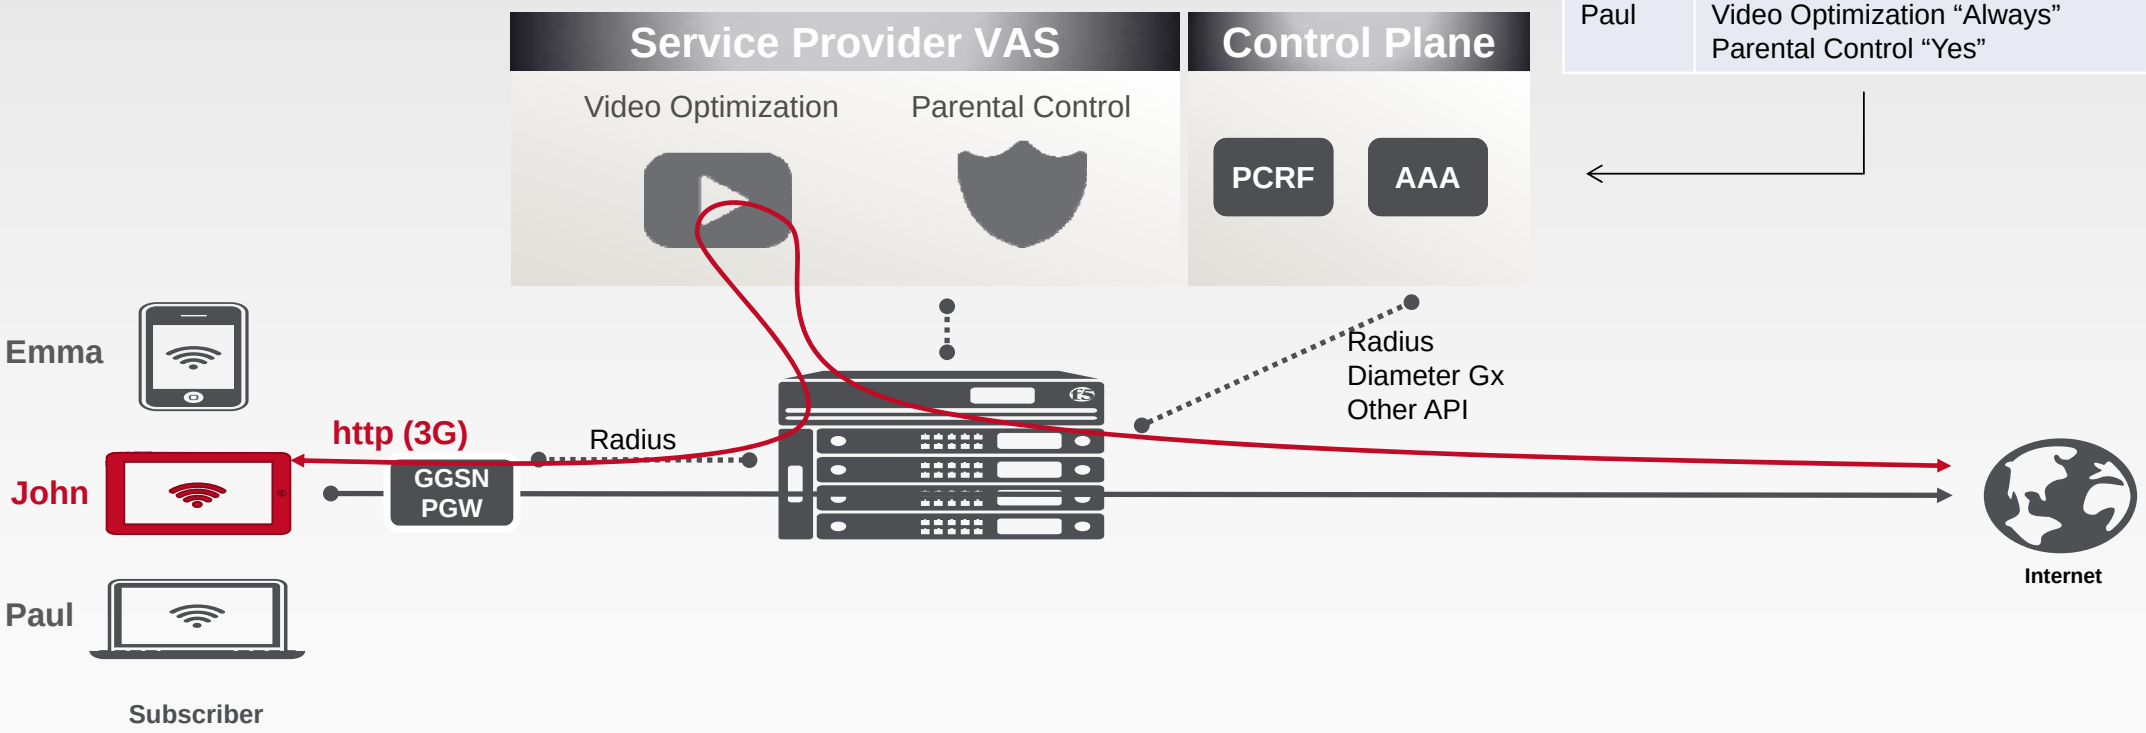

Policy-based Flow Steering in Action

User Paul : http traffic in 3G or LTE

User	Service Policy
John	Video Optimization “LTE bypass” Parental Control “No”
Paul	Video Optimization “Always” Parental Control “Yes”

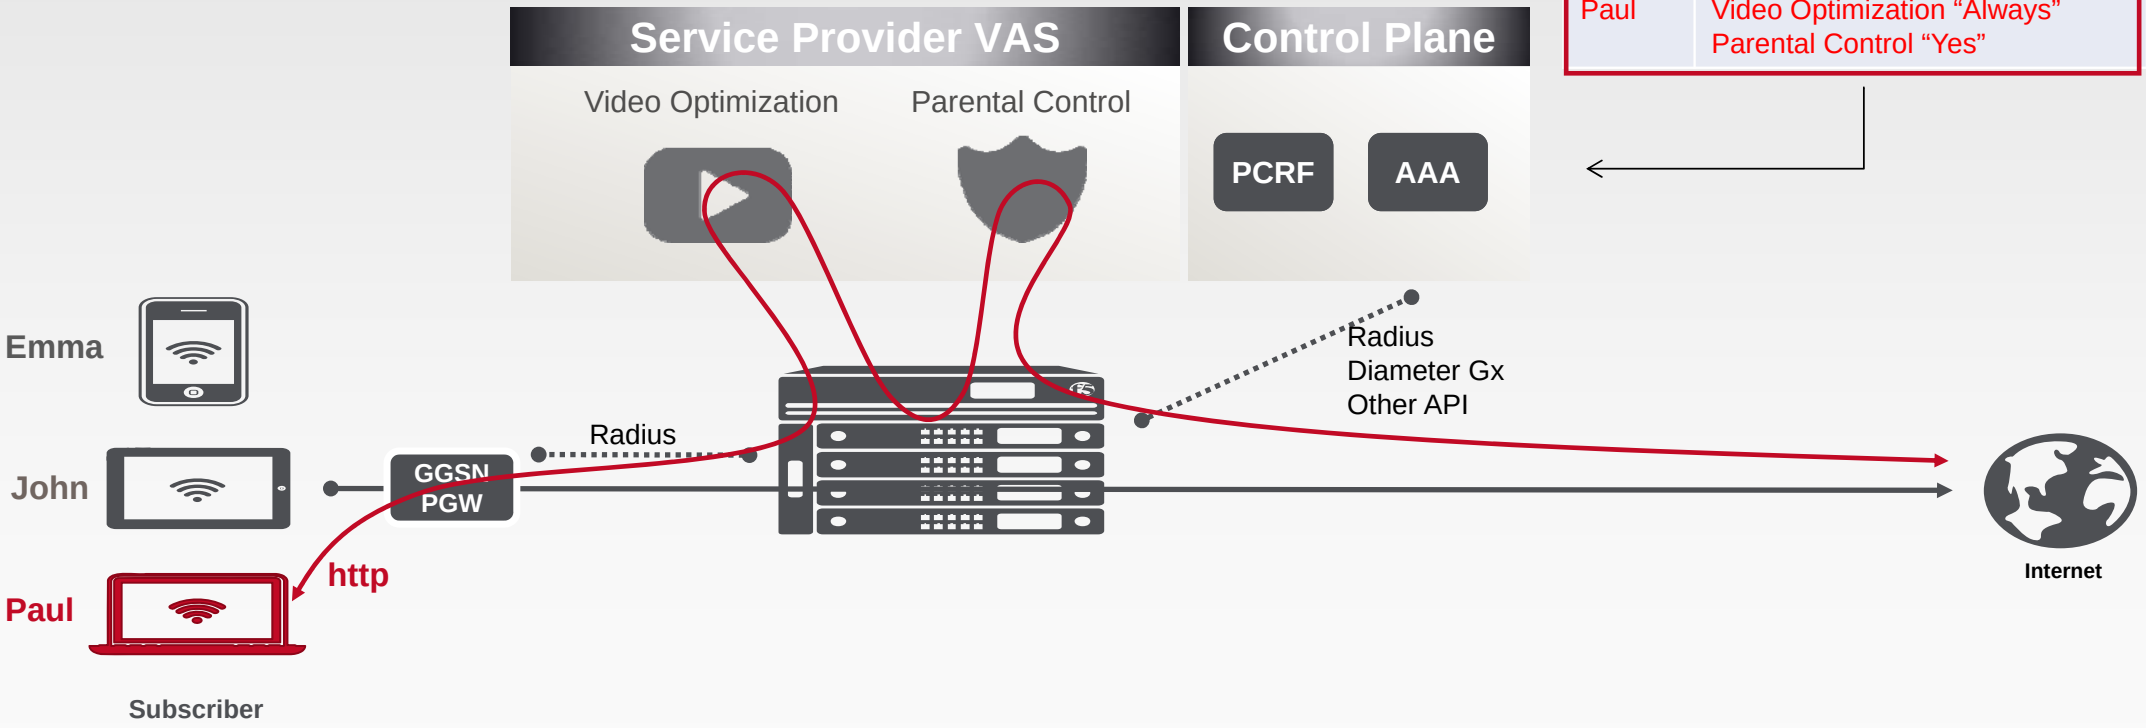

Policy-Based Flow Steering & Service Chaining
Summary

POLICY-BASED STEERING + LOAD BALANCING

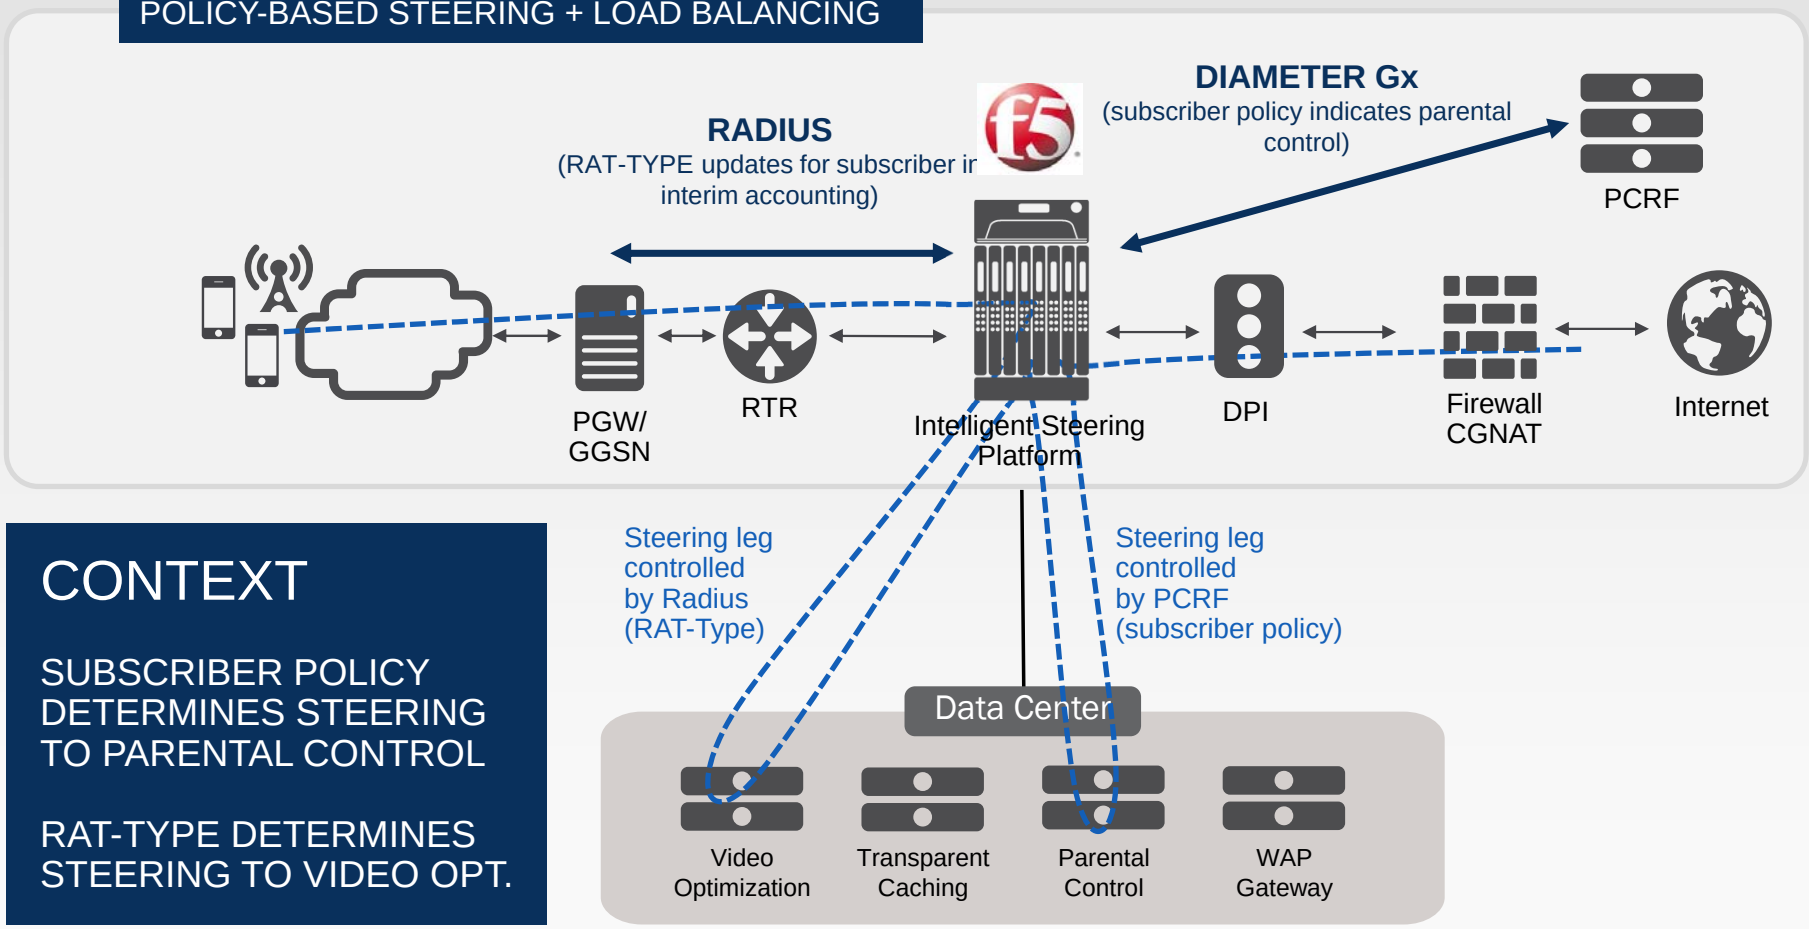

CONTEXT

SUBSCRIBER POLICY DETERMINES STEERING TO PARENTAL CONTROL

RAT-TYPE DETERMINES STEERING TO VIDEO OPT.

Policy-based “per-transaction” Steering

- Video Optimization – A changing use case
  - Increasing desire to offload any HTTP traffic that is not carrying video
  - Increasing desire to offload ABR video traffic
  - Possibly further refined with per-subscriber (for QoE differentiation) and RAT-type based (for LTE offload) steering policies
- 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
  - Requires HTTP request-based & response-based (per-transaction) steering

HTTP Messages Differ from IP Packets & TCP Flows

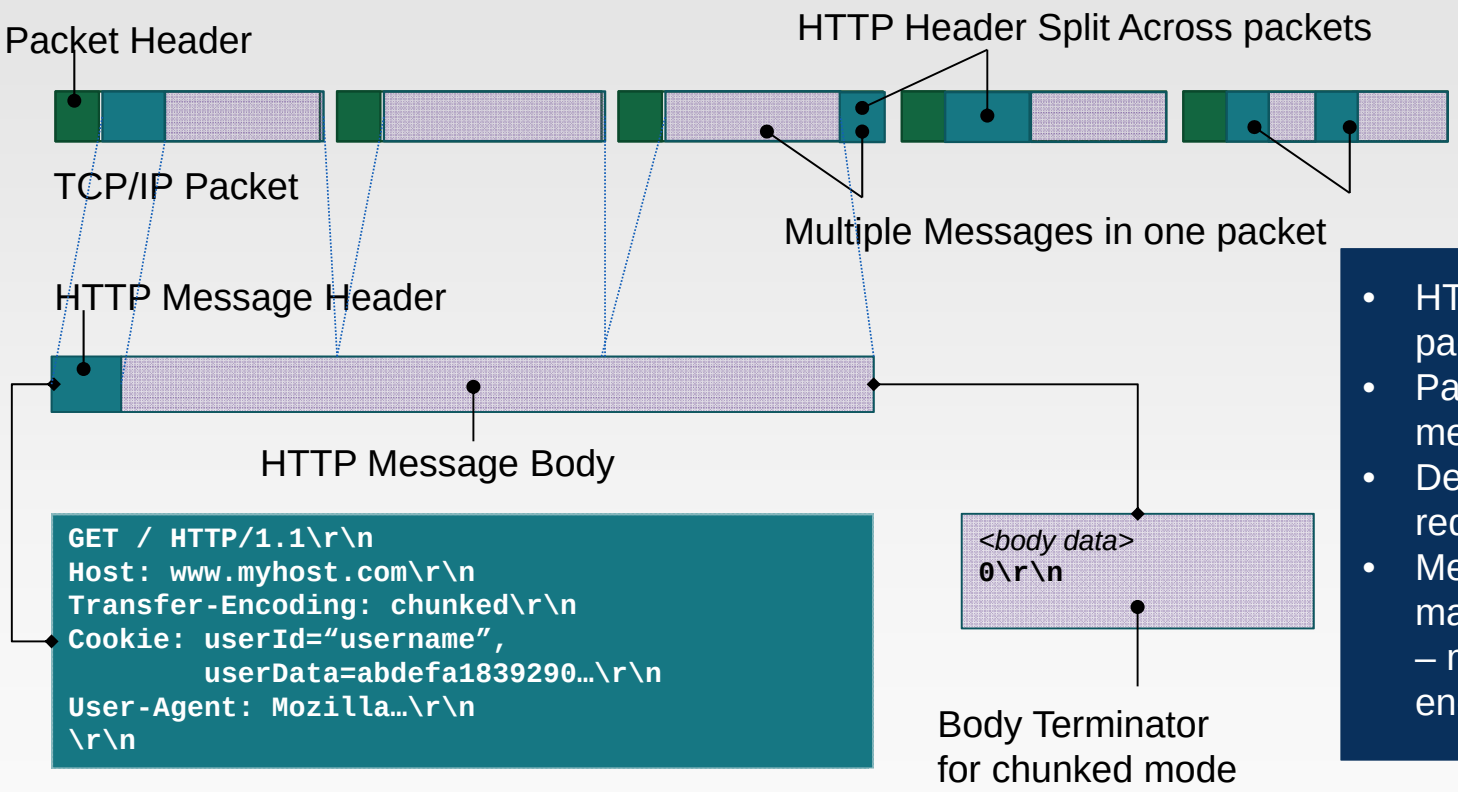

- HTTP message can span multiple packets
- Packets may have multiple HTTP messages
- Delimiting HTTP messages may require inspection of every byte
- Message steering in some cases may cause TCP stream to be split – may lead to chaos in client to end point communication

Steering on HTTP Request

- Establish TCP connection with client (full handshake)
- For each HTTP request message within that TCP connection from client
  - Parse the HTTP request headers
  - If steering policy is dependent on value(s) of one or more of the HTTP headers, then determine the nexthop (VAS endpoint) for this HTTP message according to this steering policy
- Establish new TCP connection with the VAS selected in the steering policy and forward the HTTP message between client and selected VAS
- In case of service chaining (multiple VAS endpoints) there will be several TCP connections being set up over which the HTTP message will be forwarded

Steering on HTTP Response

- Establish TCP connection with client (full handshake)
- Establish another TCP connection with the server (full handshake)
- Forward HTTP message from client to server over the full proxy
- For each HTTP response message within that TCP connection from server
  - Parse the HTTP response headers
  - If steering policy is dependent on value(s) of one or more of the HTTP headers, then determine the nexthop (VAS endpoint) for this HTTP message according to this steering policy
  - But how do we steer to the VAS ? The connection with the server is already established ...

Steering on Response – Sequence of events

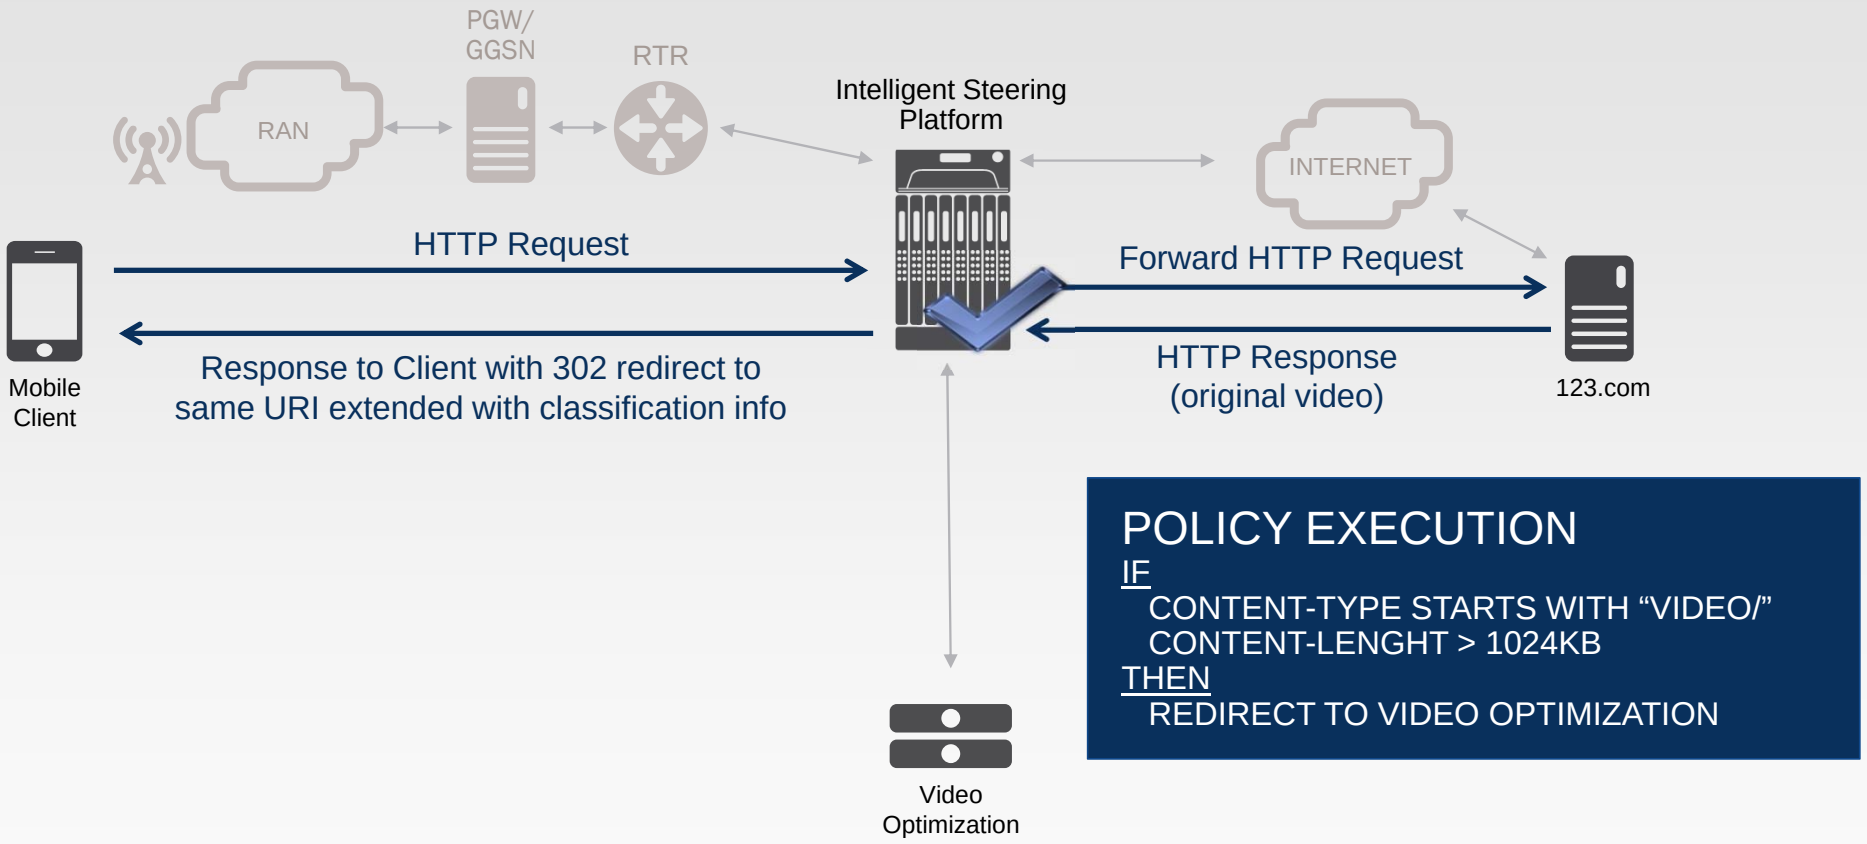

Steering on Response – After the HTTP redirect

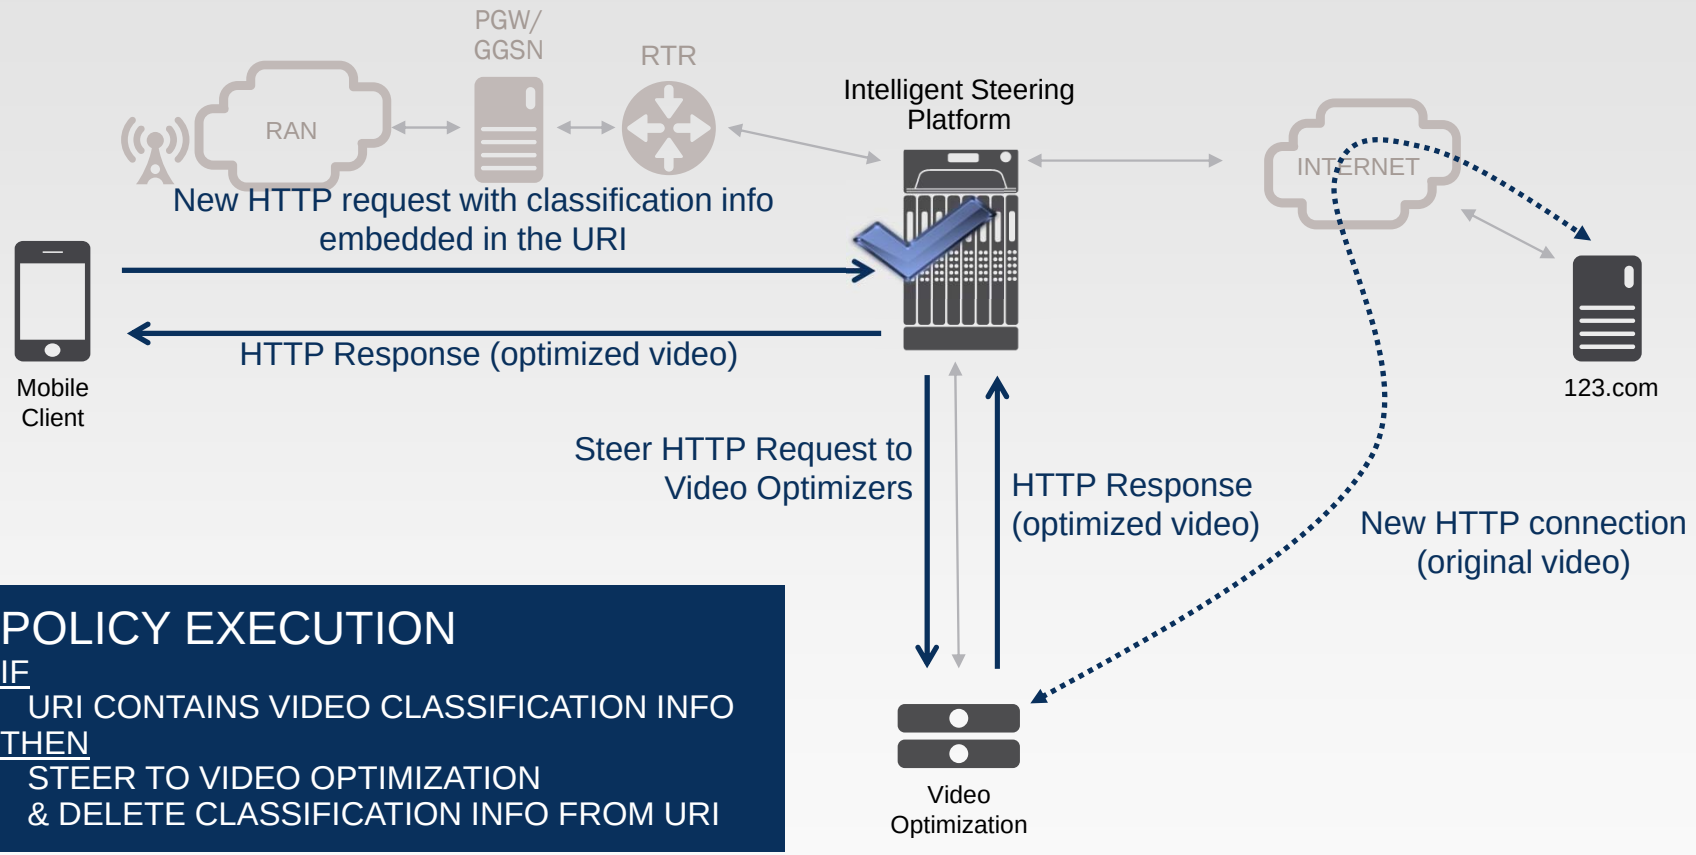

POLICY EXECUTION

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

Steering on Response – ICAP variant with Skyfire

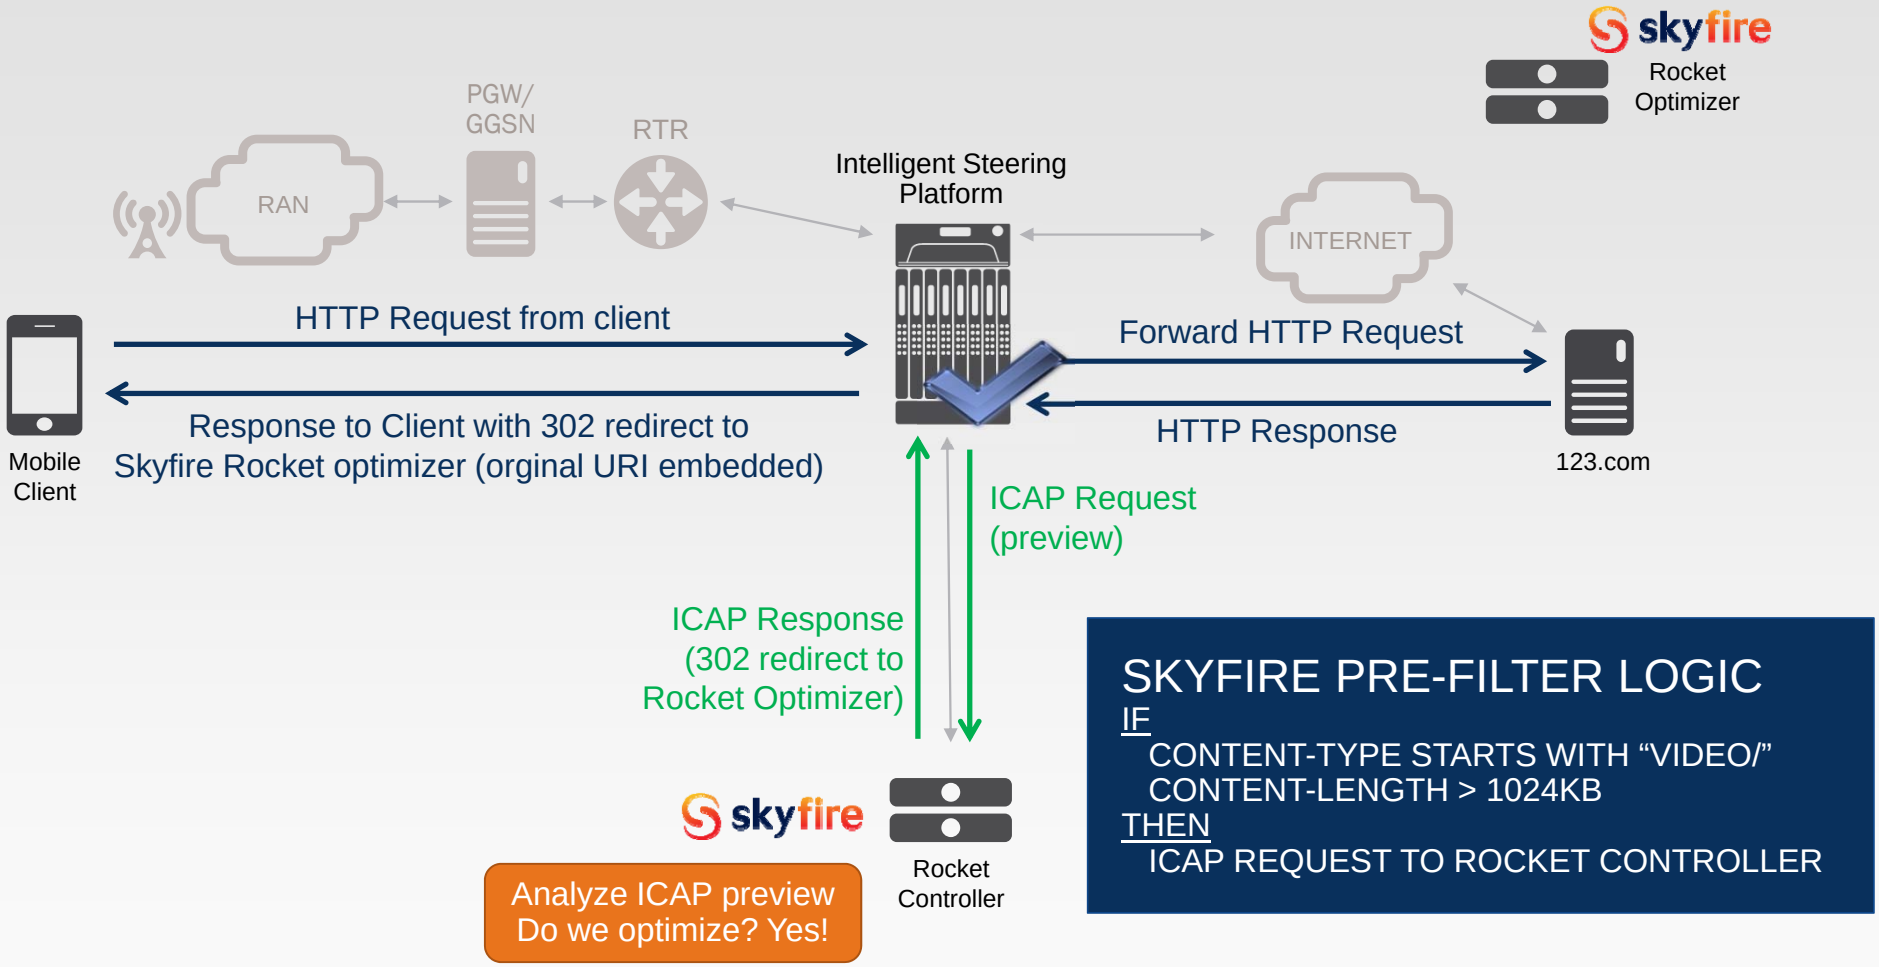

Steering on Response – After the HTTP redirect

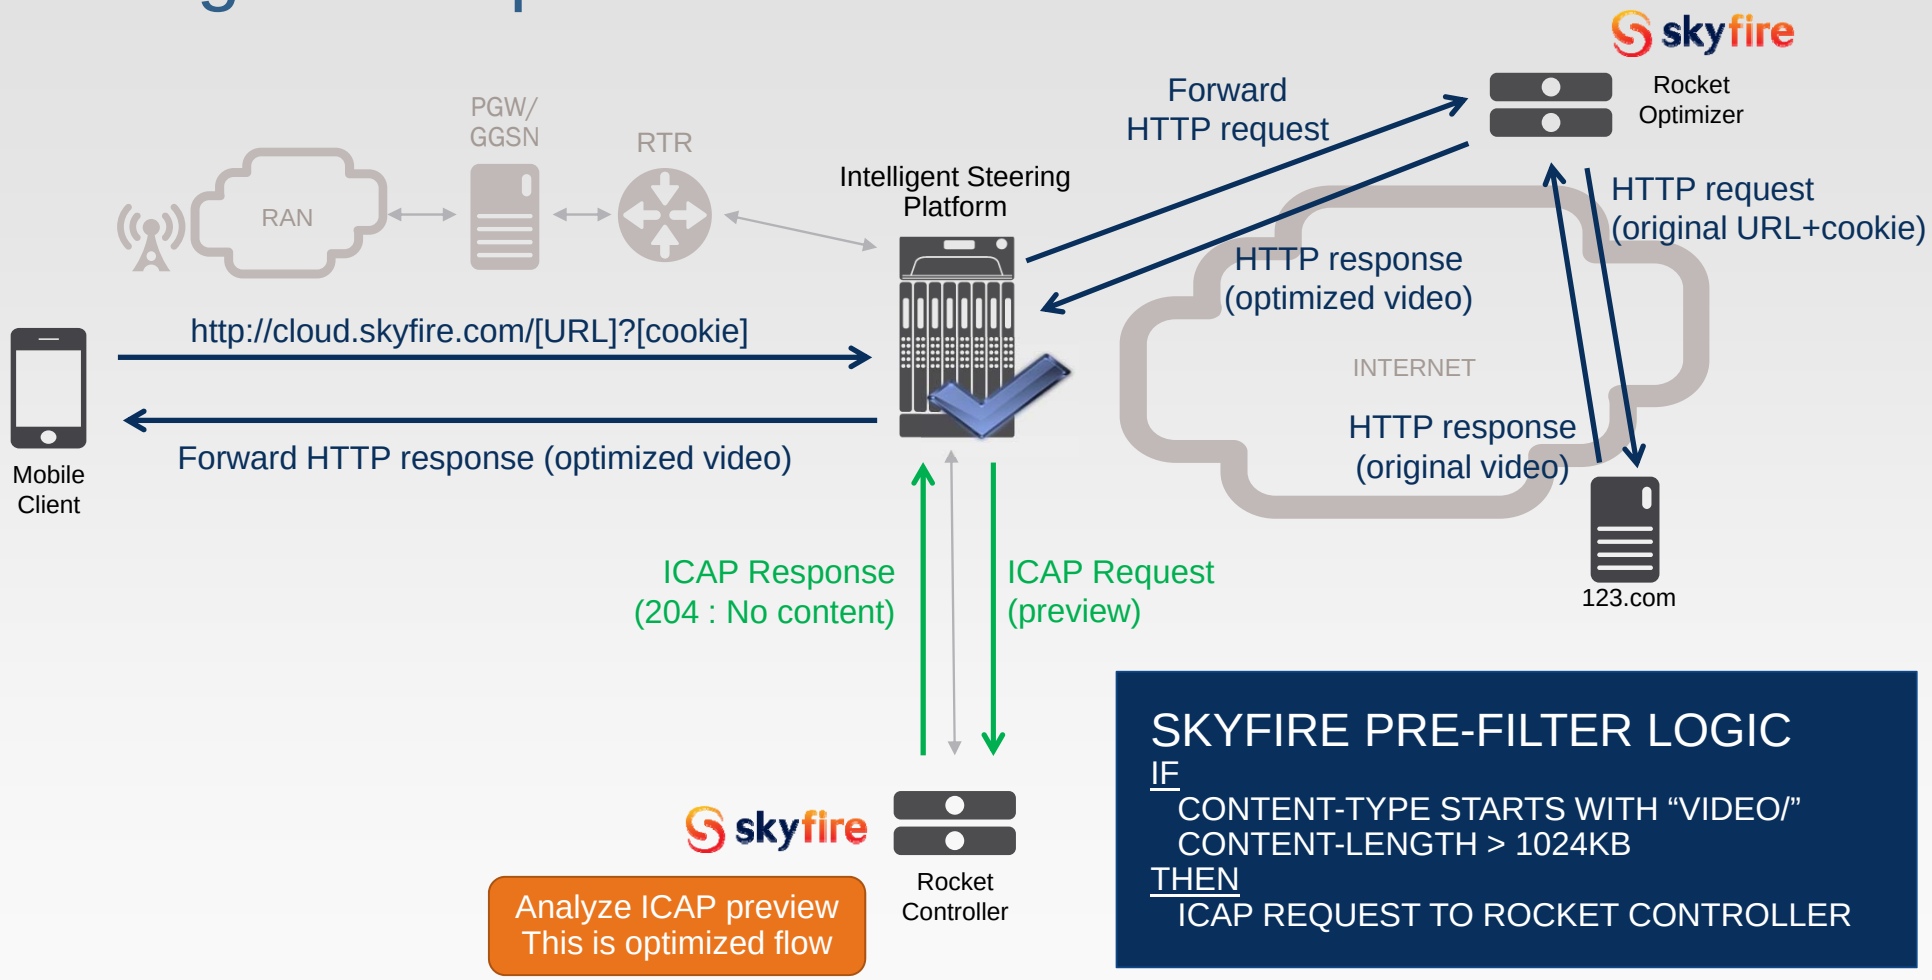

SKYFIRE PRE-FILTER LOGIC

IF
CONTENT-TYPE STARTS WITH “VIDEO/”
CONTENT-LENGTH > 1024KB
THEN
ICAP REQUEST TO ROCKET CONTROLLER

IETF – Service Chaining Working Group

- IP networks rely more and more on the combination of advanced functions
  - Besides basic routing and forwarding functions (optimization, proxy, DPI, FW)
- The goal is to 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 thereof
- Several Service Function Chaining (SFC) IETF drafts available

Sounds familiar ?

IETF – Service Function Chaining

- 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)
- Packet forwarding between SFs can be plain IP, SDN, overlay networks, ...

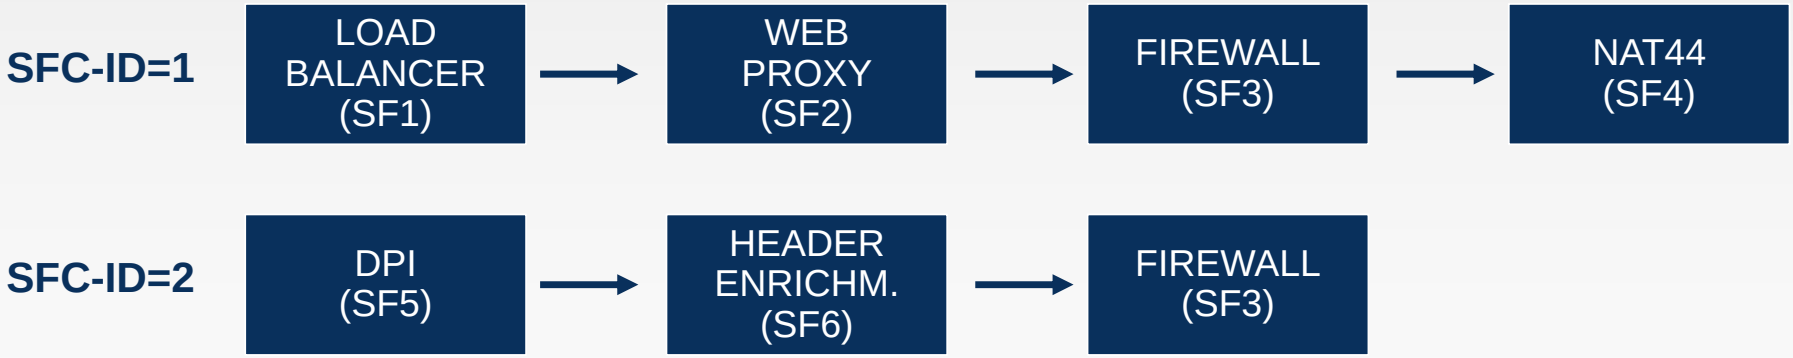

Static & Dynamic Service Chaining – Today with F5

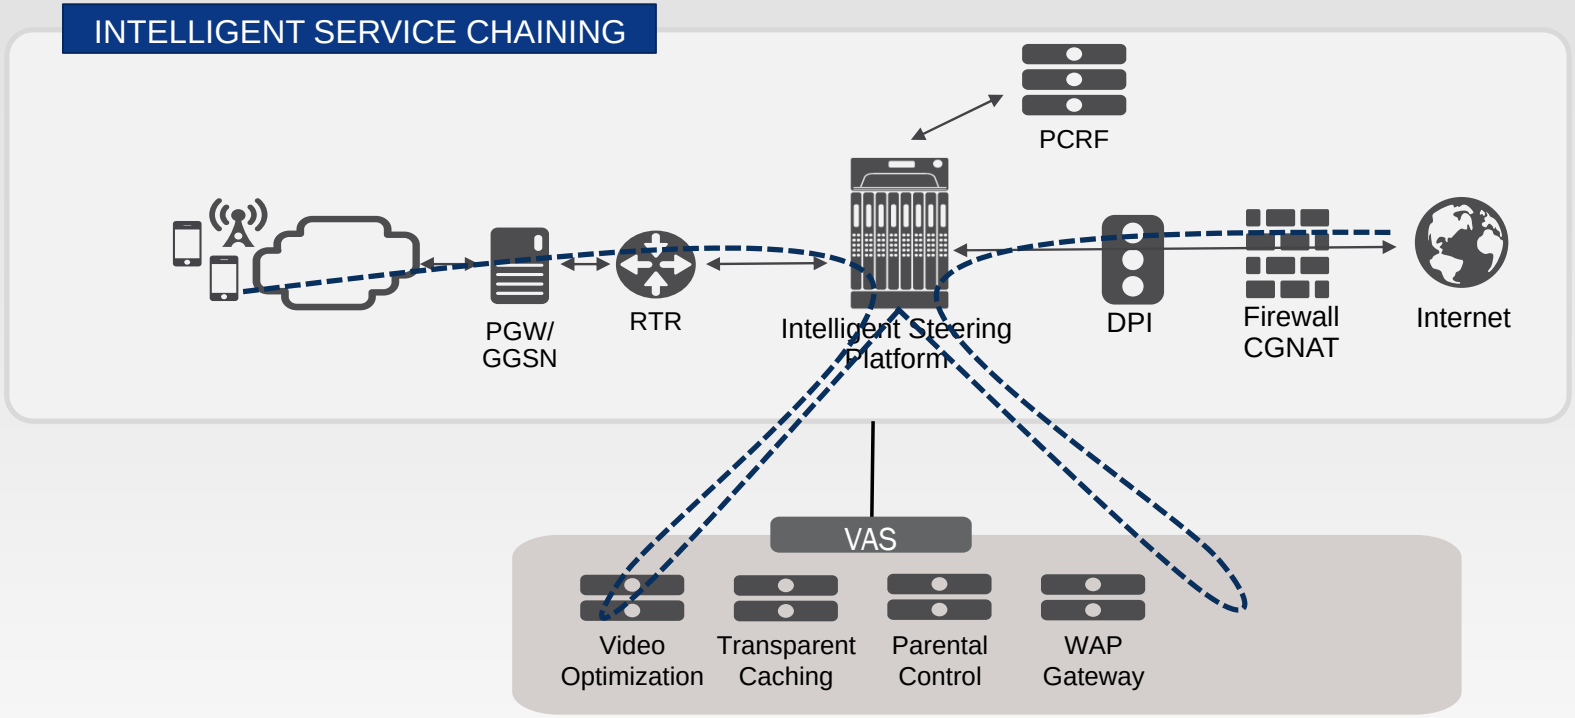

STATIC SERVICE CHAINING

INTELLIGENT STEERING POLICY DEFINES A FIXED SERVICE FUNCTION CHAIN (E.G. VAS1-VAS4)

DYNAMIC SERVICE CHAINING

INTELLIGENT STEERING POLICY HAS BUILT-IN CONDITIONAL CHECKS PER VAS LEG TO DETERMINE NEXT-HOP IN THE SERVICE CHAIN

Intelligent Service Chaining – Today and Future

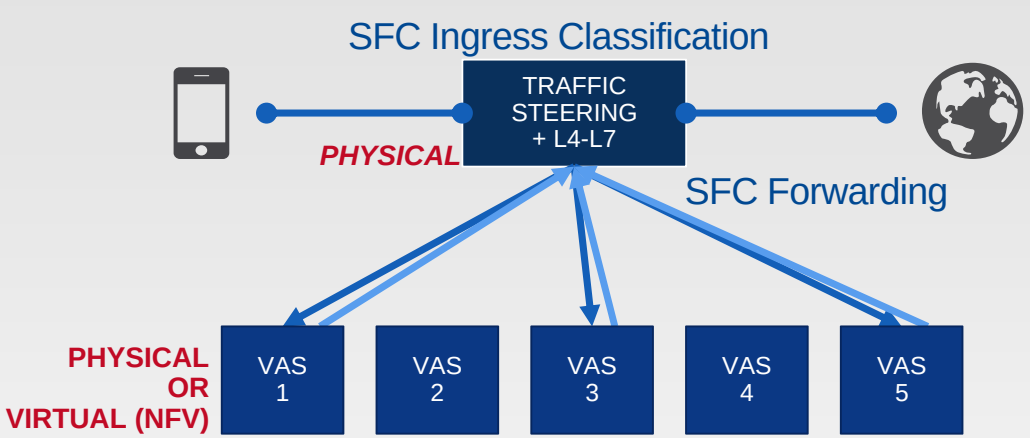

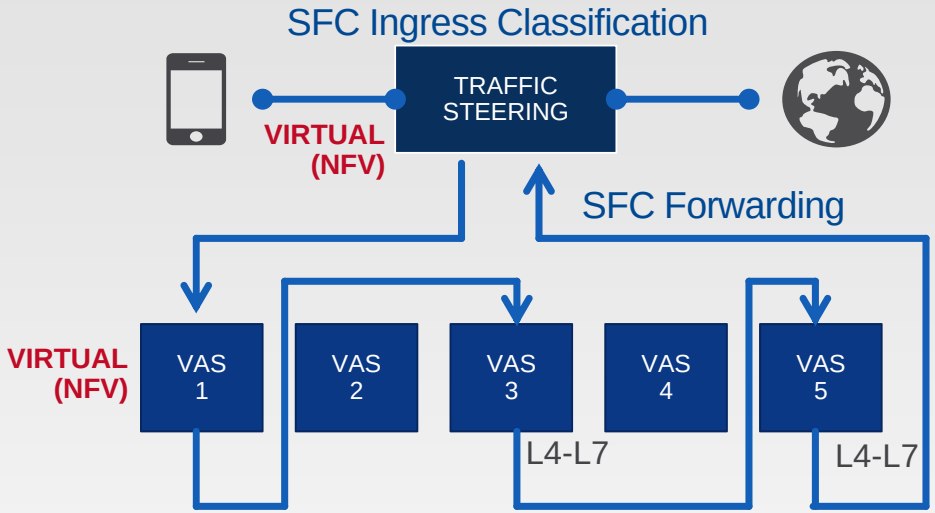

- Available today – F5 TCP & HTTP proxy
- Flexible use of "steering headers" towards VAS platforms (HTTP headers, DSCP, ...)
- Forwarding based on “connection entries”
- VAS health check built in (load balancing)
- Control plane steering possible (eg ICAP)

- IETF drafts
- Requires all vendors to agree on same standard (packet header for metadata)
- How to leverage SDN for forwarding ?
- How to do VAS health-checks ?
- How to do control plane steering ?


Policy
Enforcement
Manager
(PEM)

Policy Enforcement Manager – Policy Definition

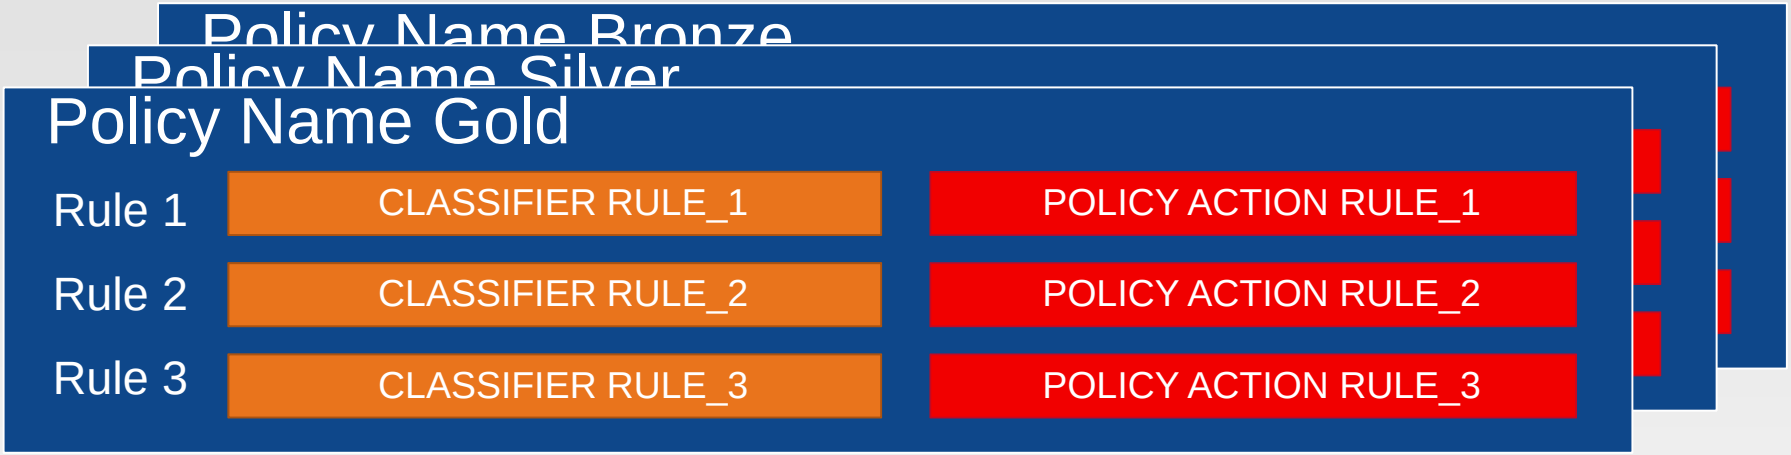

POLICY TYPE	SUBSCRIBER TYPE	POLICY ASSIGNMENT	ANALYTICS & CHARGING

• Global Policy
• Unknown Subscriber Policy
• Subscriber Policy
	• Static subscriber
	• Dynamic subscriber
	  • Radius
	  • DHCP
	  • Unknown IP SA
		• Diameter Gx
		  • Predefined
		  • Dynamic (gate, QoS)
		• Radius
		• Custom
			• Syslog
			• IPFIX
			• Radius
			• Gy
			• Gx Usage Monitoring

Classification & Policy Actions

APPLICATION CLASSIF.

- Application Category (eg. P2P)
- Application (eg. bittorrent)
- Some applications are using F5 signatures, other applications rely on third party DPI signature engine

URL CLASSIF.

- URL Category (eg. Gambling)
- URL database from third party
- Ability to create custom DB
- Used for HTTP and HTTPS (SNI check)

FLOW CLASSIF.

- DSCP
- Protocol (TCP/UDP)
- IP source address range & port
- IP destination address range & port
- Incoming VLAN

CUSTOM CLASSIF.

- irule / TCL script
- Examples
  - Other fields in the traffic flow (ip header, http header, ...)
  - Other fields stored in the PEM sessionDB for that subscriber (RAT-type, roaming, tower-id)

POLICY ACTIONS

REPORTING	HTTP REDIRECT	HTTP HDR ENR.	QOS MARKING
QUOTA MGMT	STEERING (NH)	STEERING (ICAP)	BW CONTROL
	SERVICE CHAIN	GATE (FWD)	CUSTOM / TCL

PEM – Wide range of use cases

Per-subscriber Application & URL Bandwidth Control & Filtering

- TCP-friendly rate limiter
- Separate up/down rates
- Highly scalable solution
- TCP Optimization as a bonus

Subscriber Application Analytics

- Subscriber ID / Rate Plan
- Charging rules
- Application Usage Reporting

Intelligent Traffic Steering & Service Chaining to VAS

- Steer traffic based on subscriber profile to Value Added Services & Optimization Services
- Intelligent Service Chaining

Online Charging (Gy)

- Flexible rating group definitions based on applications and/or URI
- Redirect or block upon quota expiration

URL Filtering & Parental Control

- Government lists
- Per-subscriber parental control opt-in/opt-out service
- For HTTP & HTTPS

OTT Identification & Monetization

- Per-subscriber OTT application detection
- Per-OTT bandwidth, marking and charging rules

Header Enrichment & WAP offload

- HTTP HE for content-based charging
- WAP GW bypass/offload and replacement

Content Injection / Toolbars

- Java-script based content injection
- Targeted advertisements

Lightweight BRAS/BNG

- DHCP-based BNG model for wifi and wireline deployments
- Radius AAA client

Evolution towards NFV

Step 1 : Consolidate SGi Functions & Virtualize VAS Layer

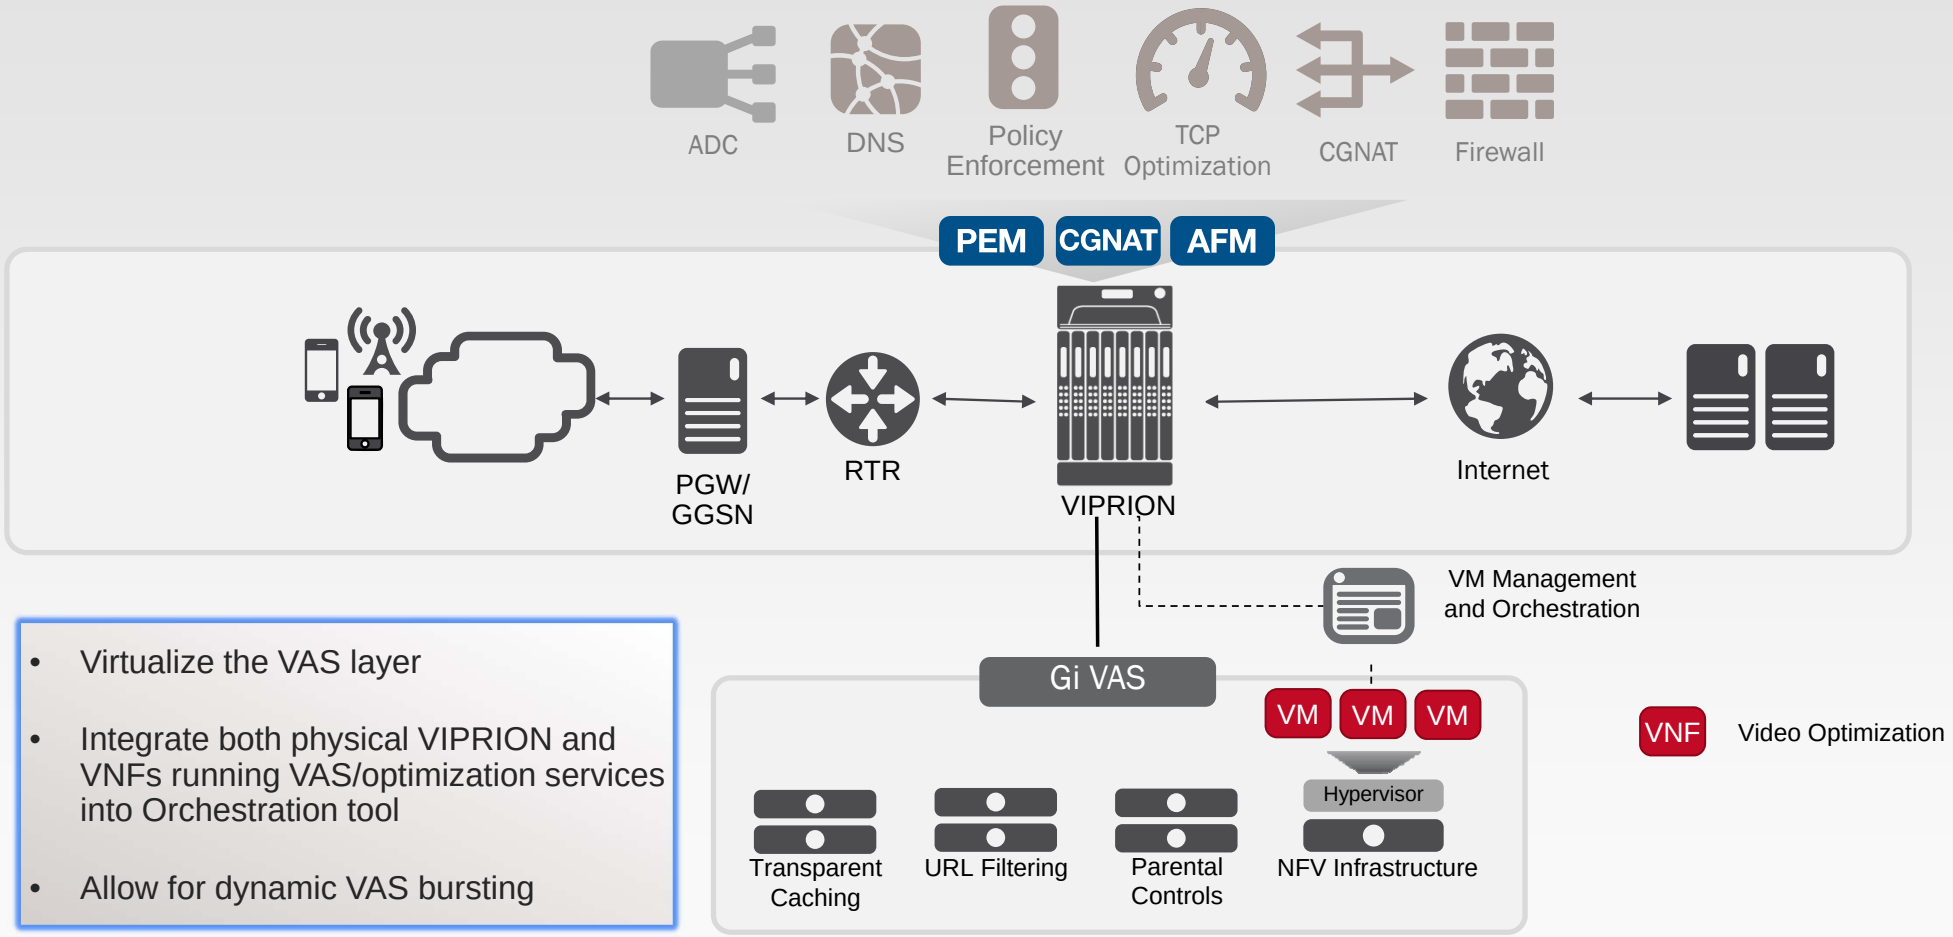

- Virtualize the VAS layer
- Integrate both physical VIPRION and VNFs running VAS/optimization services into Orchestration tool
- Allow for dynamic VAS bursting

Step 2 : Virtualize SGi + VAS layer (NFV)

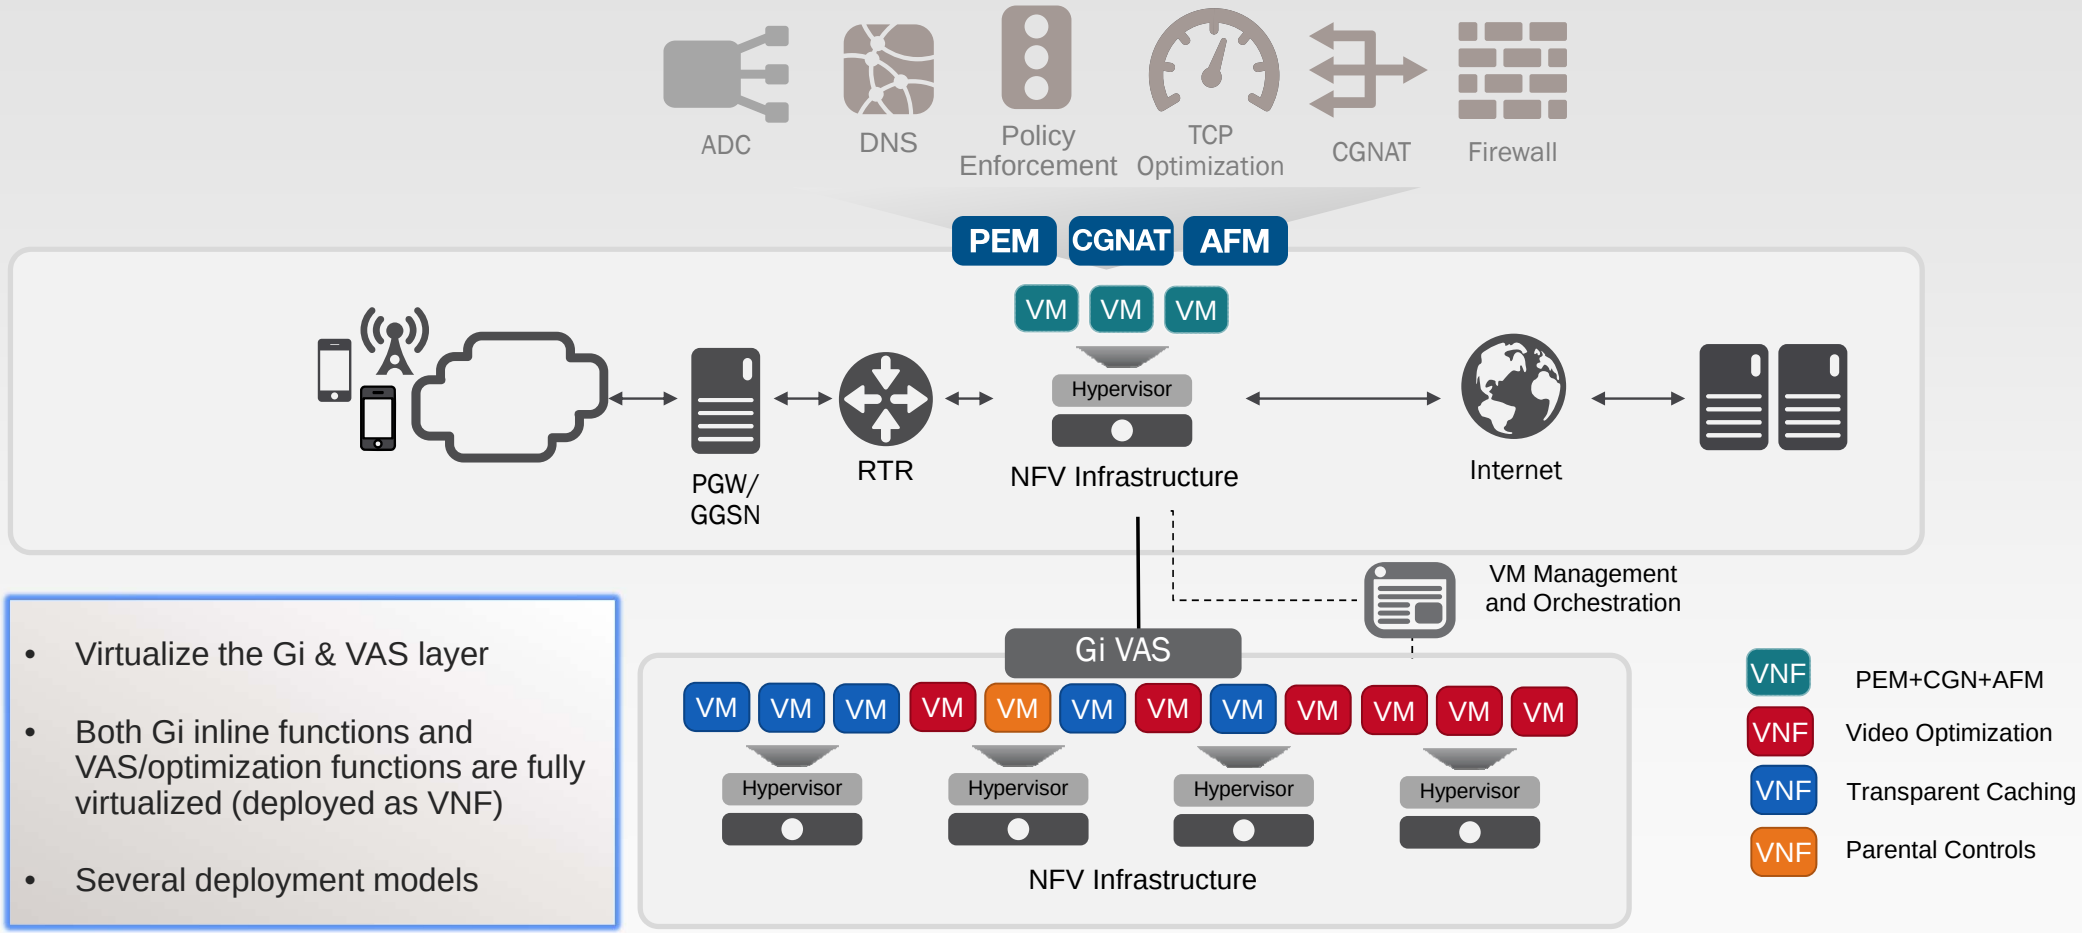

- Virtualize the Gi & VAS layer
- Both Gi inline functions and VAS/optimization functions are fully virtualized (deployed as VNF)
- Several deployment models

Virtualizing SGi functions – VNF mapping alternatives

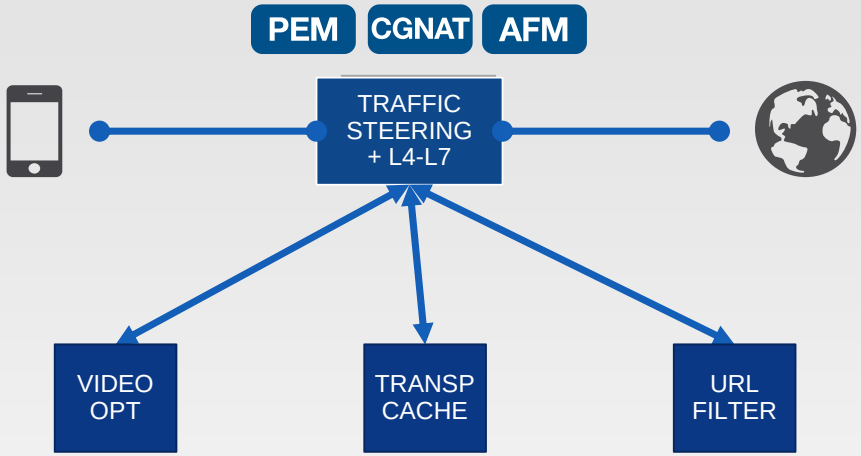

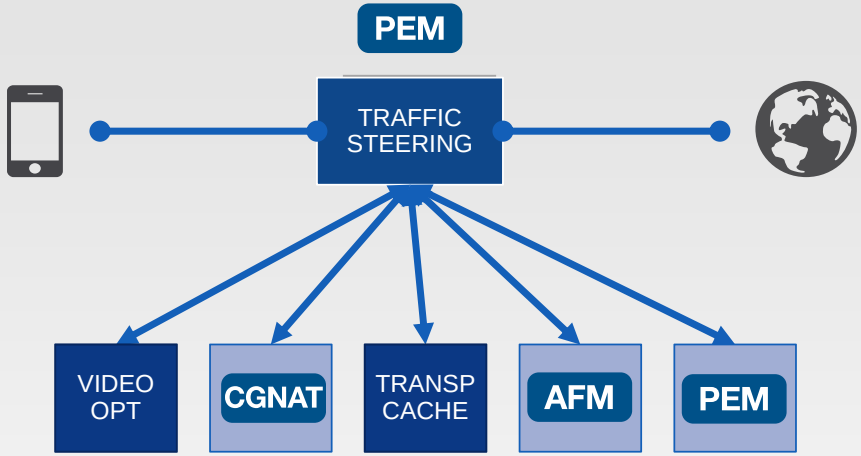

- Maintain the L4-L7 consolidation model as deployed in the physical world in the NFV world
- So you run PEM, CGN and AFM on the same VE (or on separate VEs but all inline in the data path)

- Only maintain PEM inline in the data path for intelligent traffic steering
- Treat functions such as CGNAT and security/firewall functions as VAS services (which are only used for selected flows)

Summary

Summary – Traffic Steering & Service Chaining

- Market is demanding advanced steering capabilities for VAS offload
  - F5 PEM platform supports context-aware intelligent steering and service chaining
- Market is demanding an evolution path towards IETF-based service chaining
  - F5 PEM platform provides a stepwise approach
- Market is demanding a solution that supports a wide variety of use cases
  - F5 PEM platform supports a wide range of classification criteria coupled with strong policy action mechanisms
- Market is demanding a solution that can migrate from physical to virtual
  - F5 PEM platform provides such a migration path, offering multiple alternative deployment options for NFV

Solutions for an application world.