DNS for Service Providers

Ivan Ramos

SP Solutions Architect LatAm

Agenda

1.- DNS and SP Trends

2.-Intelligent DNS (DNS in 3G Networks)

3.-Roadmap (DNS for 4G Support)

Key F5 network services

A unified platform and single management framework

- Intelligent DNS
- Intelligent traffic management
- CGNAT and IPv6 migration
- Policy enforcement
- ICSA certified network firewall
- Header enrichment and TCP optimization
- URL filtering

F5 can help

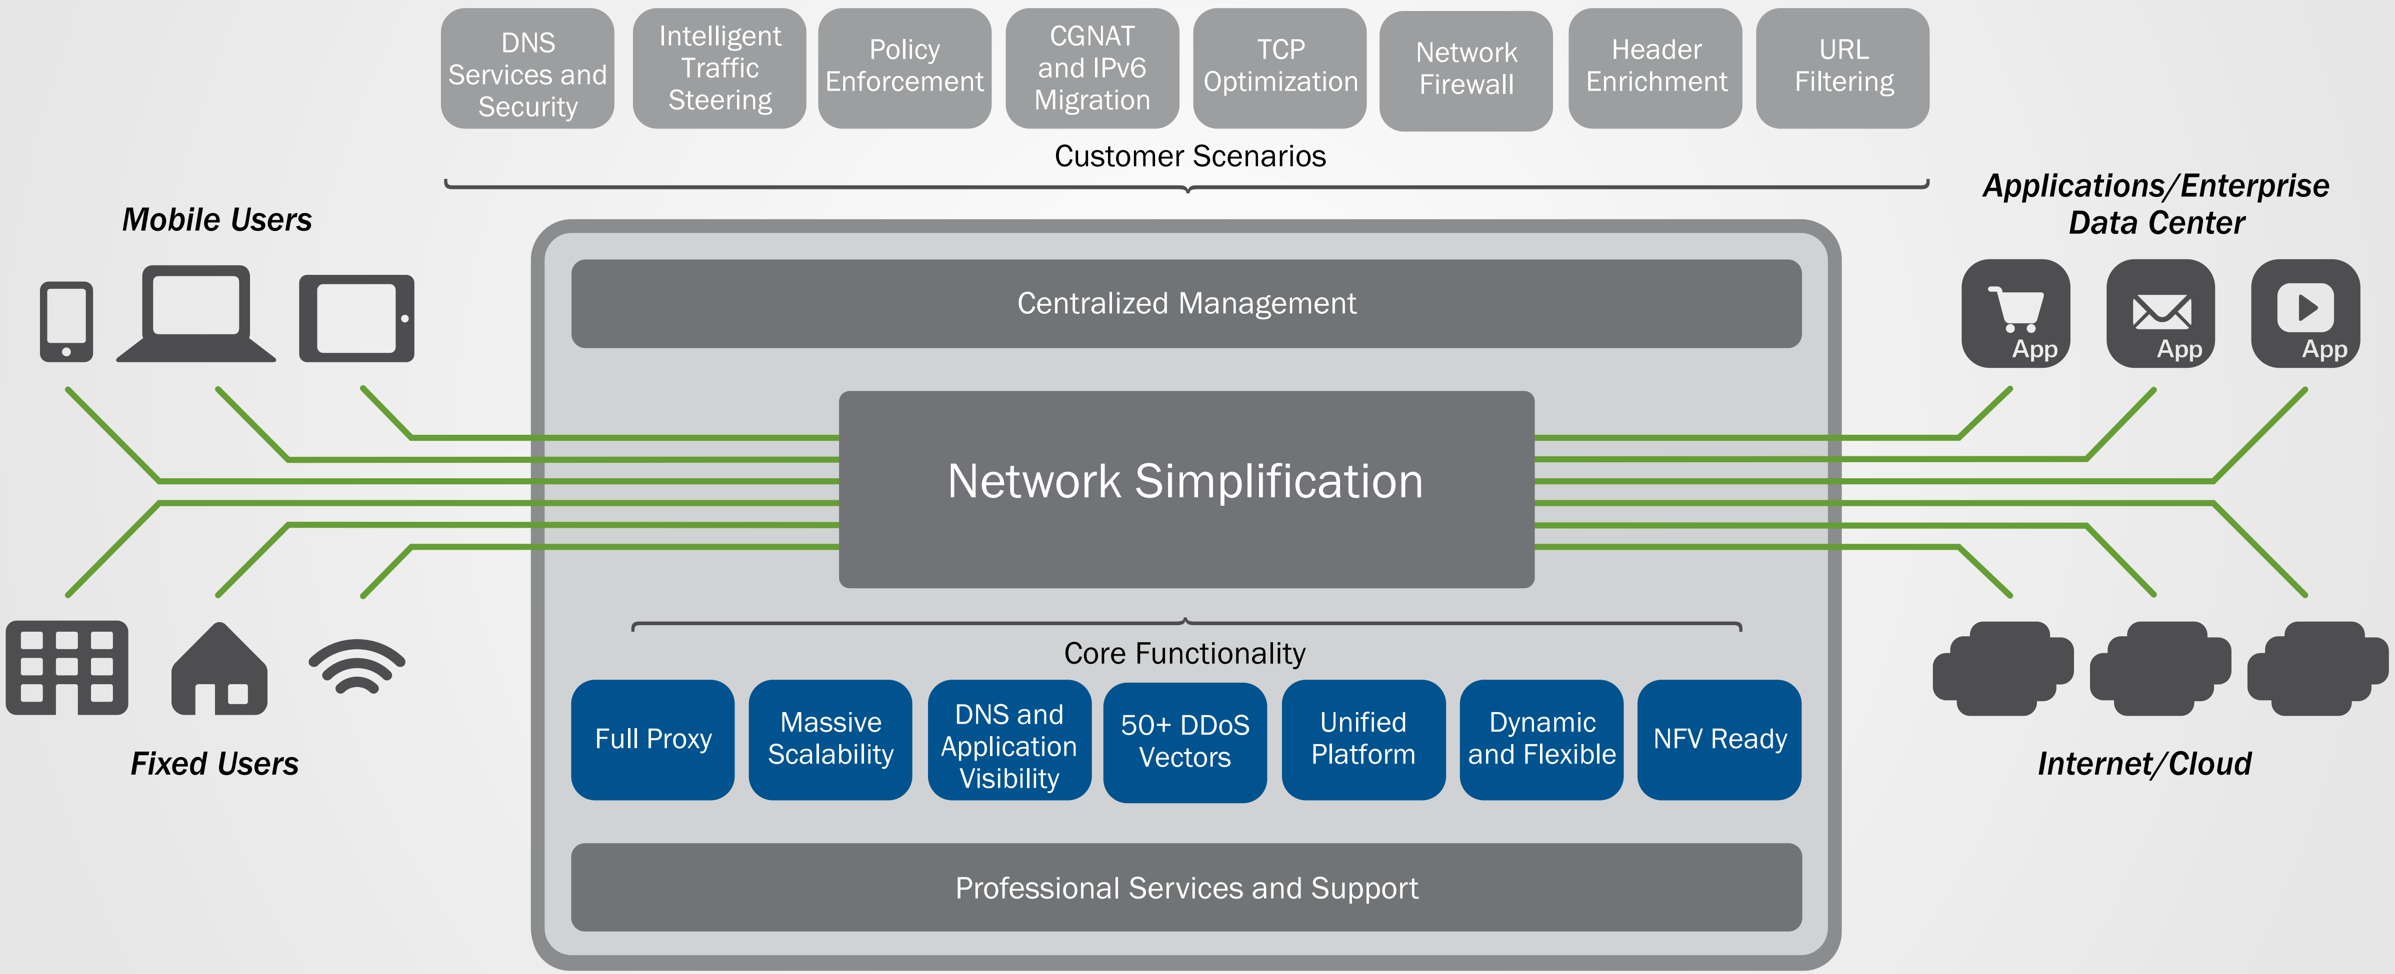

DNS Services and Security
Intelligent Traffic Steering
Policy Enforcement
CGNAT and IPv6 Migration
TCP Optimization
Network Firewall
Header Enrichment
URL Filtering
Customer Scenarios
Mobile Users
Applications/Enterprise Data Center
App
App
App
Centralized Management
Network Simplification
Core Functionality
Full Proxy
Massive Scalability
DNS and Application Visibility
50+ DDoS Vectors
Unified Platform
Dynamic and Flexible
NFV Ready
Professional Services and Support
Fixed Users
Internet/Cloud

DNS and SP Trends

Growth, Innovation, and Pain Points

Explosive data growth

Worldwide mobile data to grow **13 times** by 2018

Total mobile subscriptions to reach **9.1 billion** by 2018

Security attacks

A DDoS attack occurs on the web every **2 minutes**

Attacks over 10 Gbps have increased **nearly 50%**

Network innovation

213 4G LTE networks have launched in **81** countries

Over 50% of leading operators plan to deploy SDN and NFV by 2014

New VAS services

40% of global YouTube traffic is mobile, up from 6% in 2011

Facebook has over 800 million monthly mobile users, up **150%** since 2011

Everything: DNS

- Internet of Things **needs scalable DNS services***
- Combination = 5 to 10 times Internet revolution**
- 10bil devices in 2014 = 77bil mobile apps**
- Ensure really fast connections and responses*

Denial of Service Attacks Against DNS

Application Layer Attacks

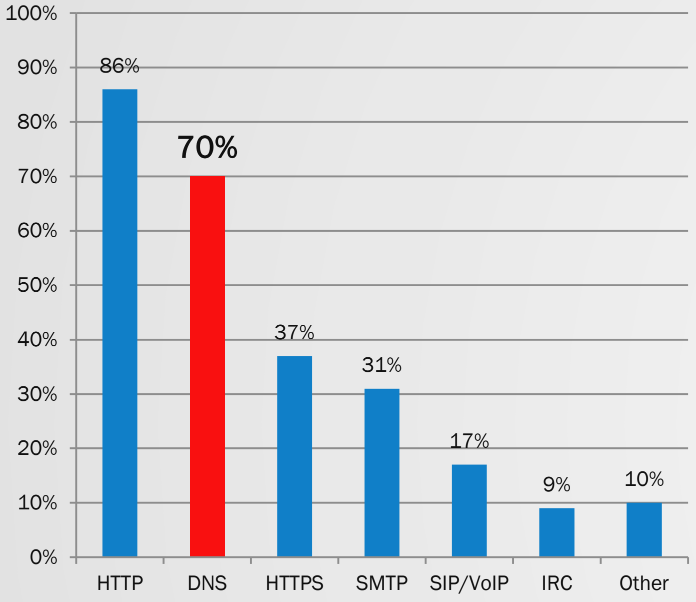

DNS is now the second most targeted protocol after HTTP.

DNS DoS techniques range from:

- Flooding requests to a given host
- Reflection attacks against DNS infrastructure
- Reflect / Amplification attacks
- DNS Cache Poisoning attempts

Traditional DDoS Mitigation

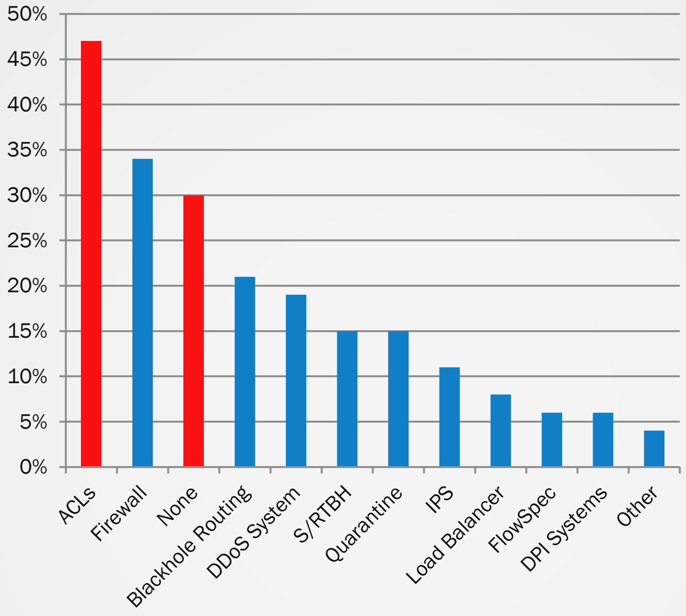

Of the **customers that mitigate DDoS** attacks, many **choose a technique that inhibits the ability of DNS** to do its job

- DNS is based on UDP
- DNS DDoS often uses spoofed sources
- Using an ACL block legitimate clients
- DNS attacks use massive volumes of source addresses, breaking many firewalls.

Attack Sources and Frequency

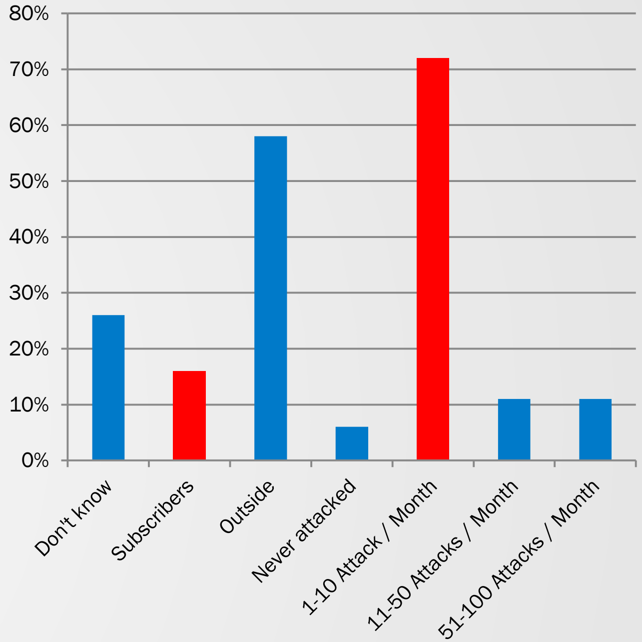

94% of SPs are under at least **one DoS attack per month.**

Of those, 25% of SPs don't know if the cause was their own subscribers, while 16% reported that a combination of bad actors and malware on their own network was root cause.

Greater DNS Threats and Outages. Access to Malicious Sites.

PROBLEMS

Greater threats and volume loads on DNS infrastructure. *26% increase in DDoS attacks, 40% increase in UDP/DNS attacks**

$$$

App and site outage affects business viability and revenue. *$27mil. avg. loss for 24hr. outage from DDoS***

SPs desire to block subscriber access to IPs with malware and viruses

CURRENT SOLUTION/ PROJECT

Buy more DNS servers to handle the volumes. Buy better DNS load balancing

Route DNS DDoS volumes to external cloud scrubbing service but unaware of what's dropped

Cloud Services

Some DNS Firewall (domain filtering) services only offer one list/database to choose

*"Organizations should invest in protecting their DNS infrastructure." Gartner****

Intelligent DNS

F5 DNS Optimization

Control and Data Plane Management

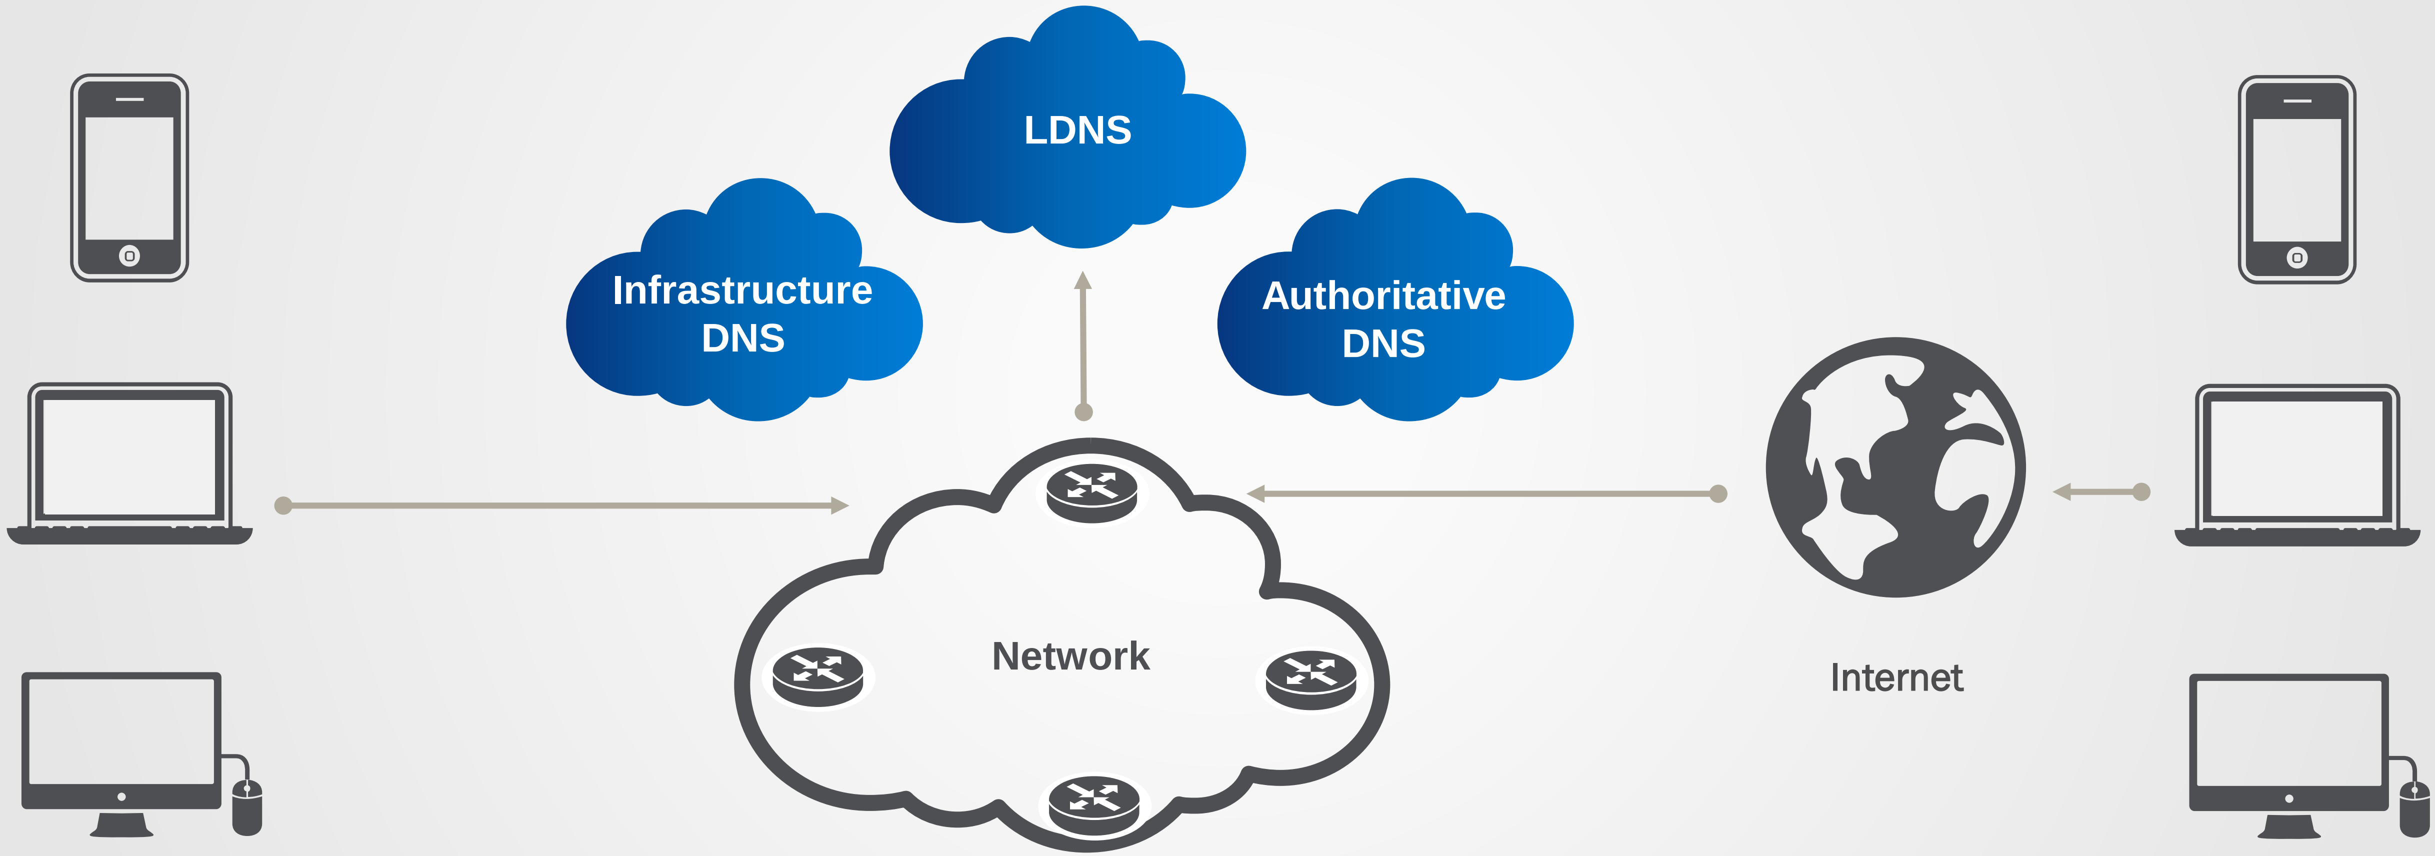

- Improve end user Quality of Experience
- (QoE) Increased flexibility and automation
- Performance for the highest service demands

F5 Intelligent DNS and Global Service Optimization

LDNS

- Faster DNS for 3G and 4G LTE
- Enhanced perf. through transparent cache
- Caching resolver for server consolidation
- Mitigate DNS threats by blocking access to malicious IPs

Authoritative

- Robust, scalable portal and service access
- Exponential DNS performance and DDoS security protection
- Optimize global service delivery

Infrastructure

- Intelligent DNS for Evolved Packet Core
- Proactively monitor for service-level adherence
- Enhanced subscriber experience

Simplified and Consolidated DNS

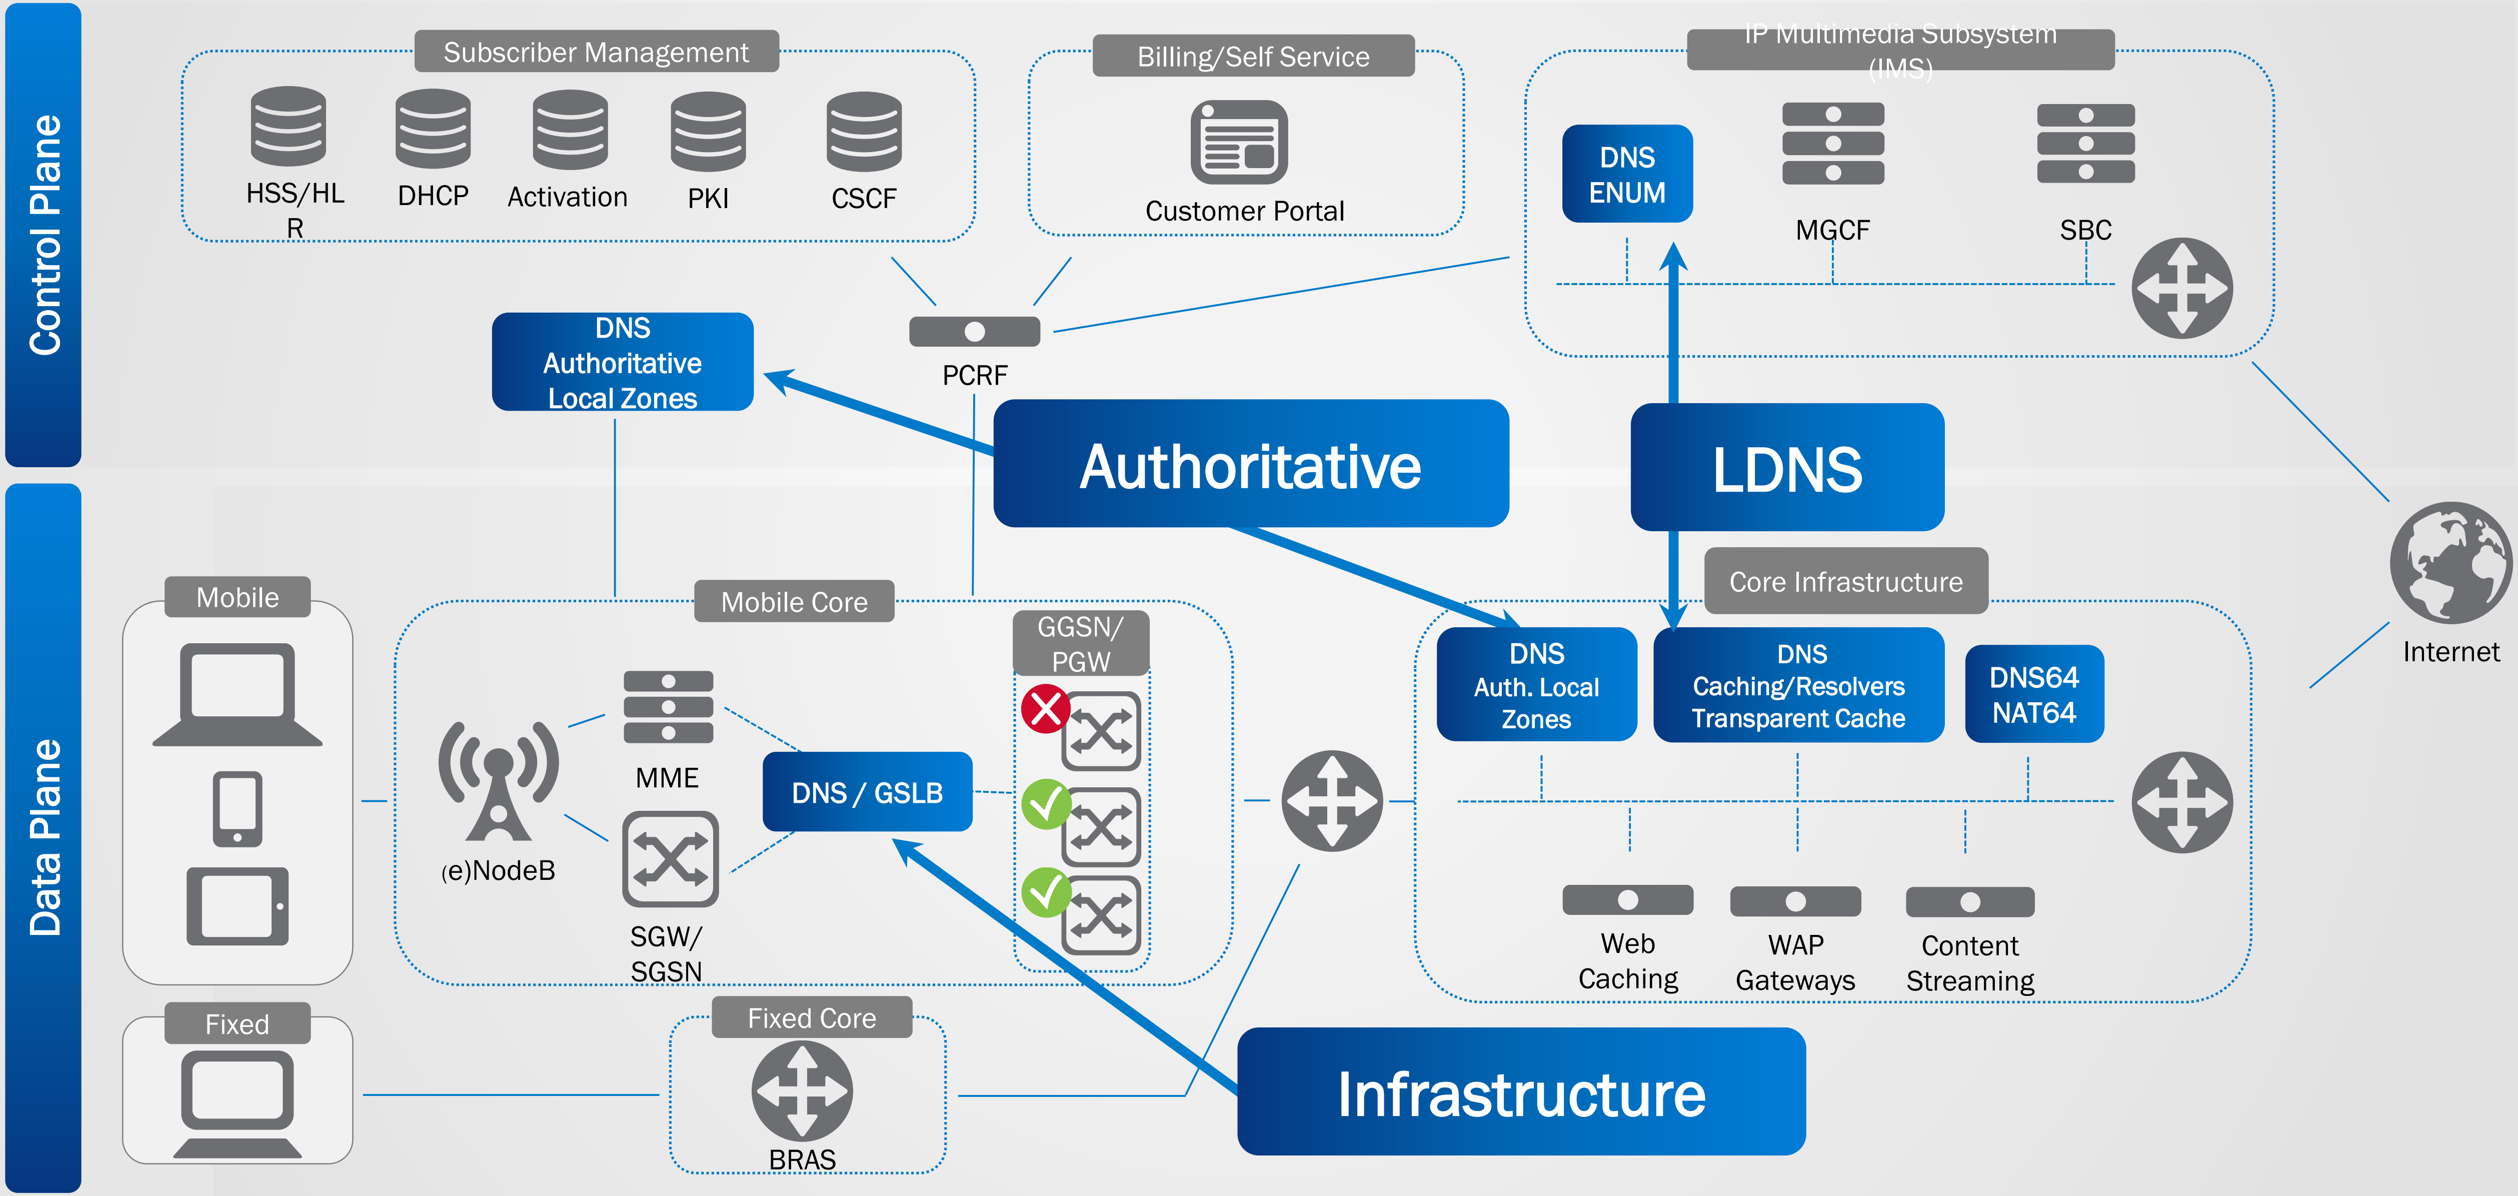

LDNS: Faster DNS Responses with BIG-IP

The Business Case

- SPs add DNS servers to accommodate growth
- Effective DNS responsive with load balancing
- Easily deploy F5 leading DNS Delivery solution
- Low barrier to entry
  – Works with existing servers, and policies

The F5 Advantage

- Ensure a consistent experience for subscribers
  - Health monitors with DNS/Service delivery
- Protect existing infrastructure with firewall services
- Take advantage of F5 capabilities and services
  - Same framework, same topology, greater scalability

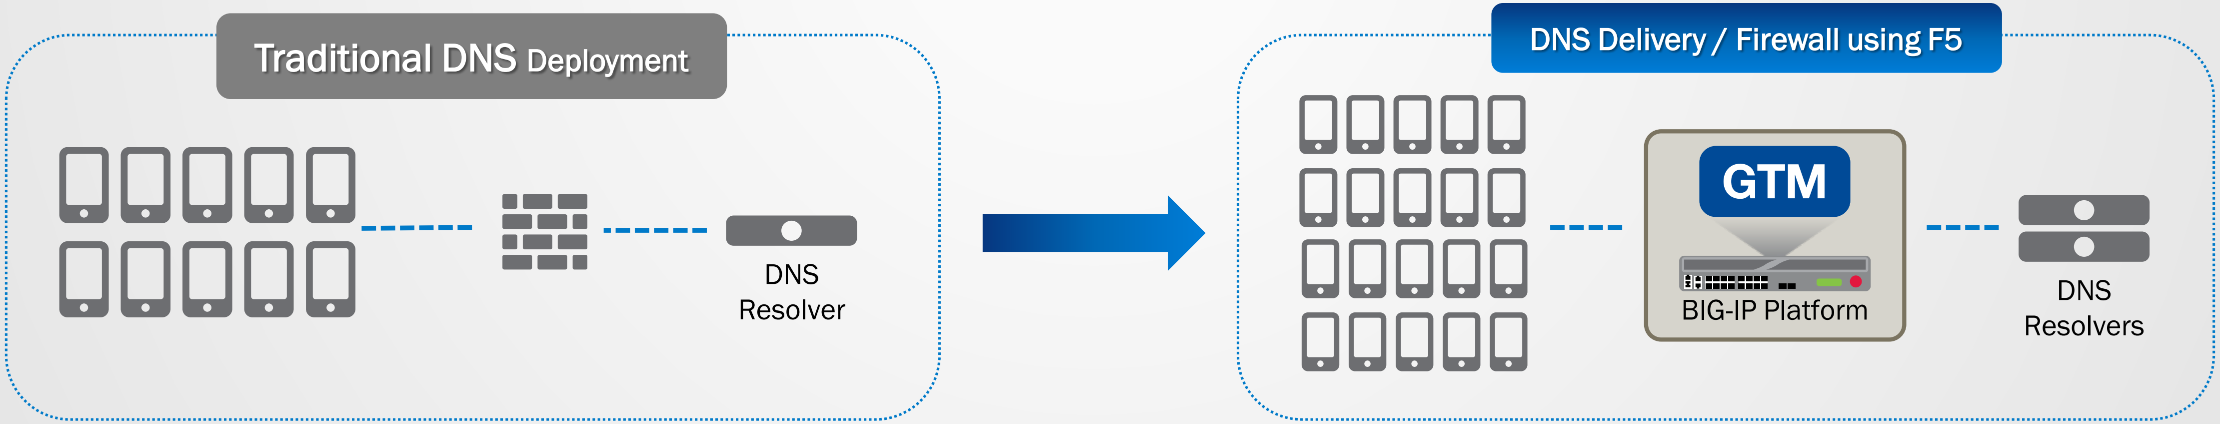

LDNS: Scale with Transparent Cache

The Business Case

- Need to **decrease DNS latency and offload DNS resolvers**
- **Implement transparent DNS caches** close to the subscriber
- **Deliver DNS scale** without impacting service

The F5 Advantage

- **Scale DNS transparent caches** as demand increases. **Offloads existing DNS infrastructure**
- Provides a **simple upgrade path** to a full caching resolver
  - **Eliminate the need** for centralized DNS

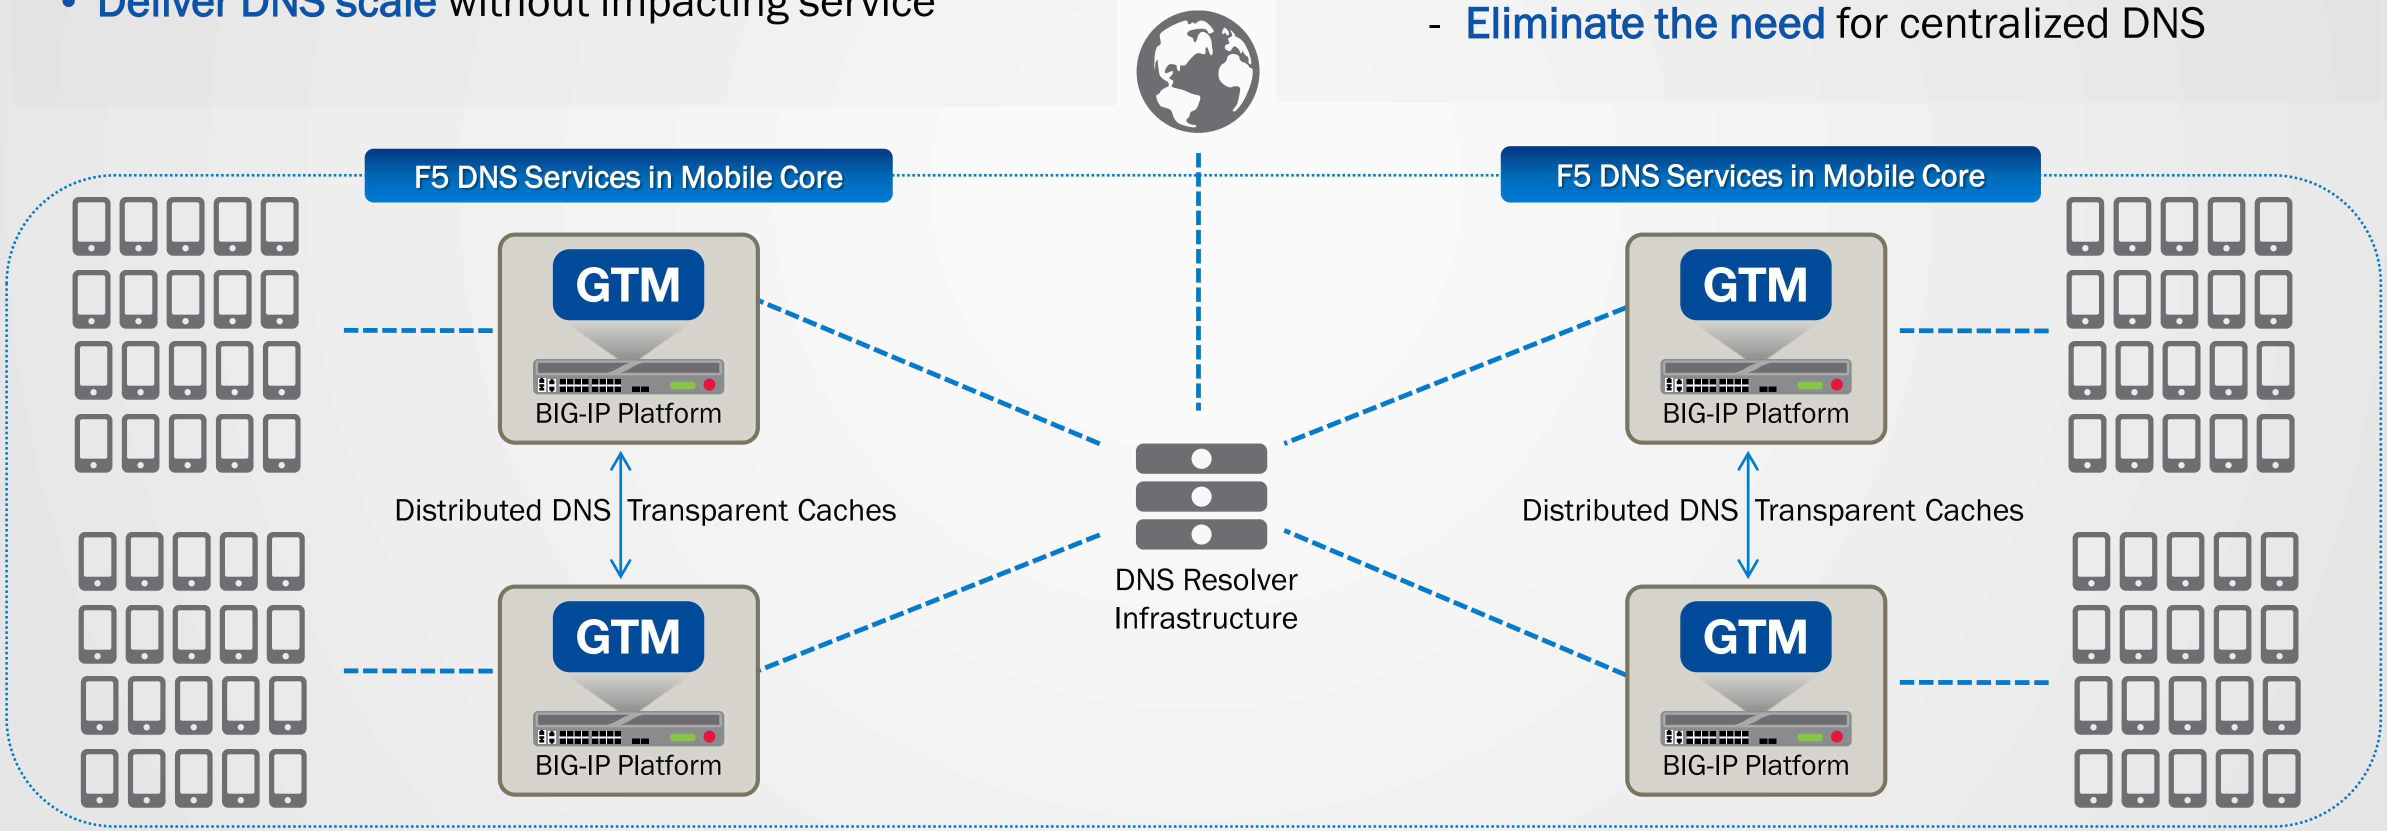

Optimized DNS Resolving and Cache Zone Forwarding

FASTER WEB BROWSING

- DNS Caching passes queries to the Resolver when response isn't cached
- Resolver uses root hints to kick off process

FASTEST WEB BROWSING

- Requests for specific zones sent to specific recursive name server
- Zone not listed, then Resolver follows root hints

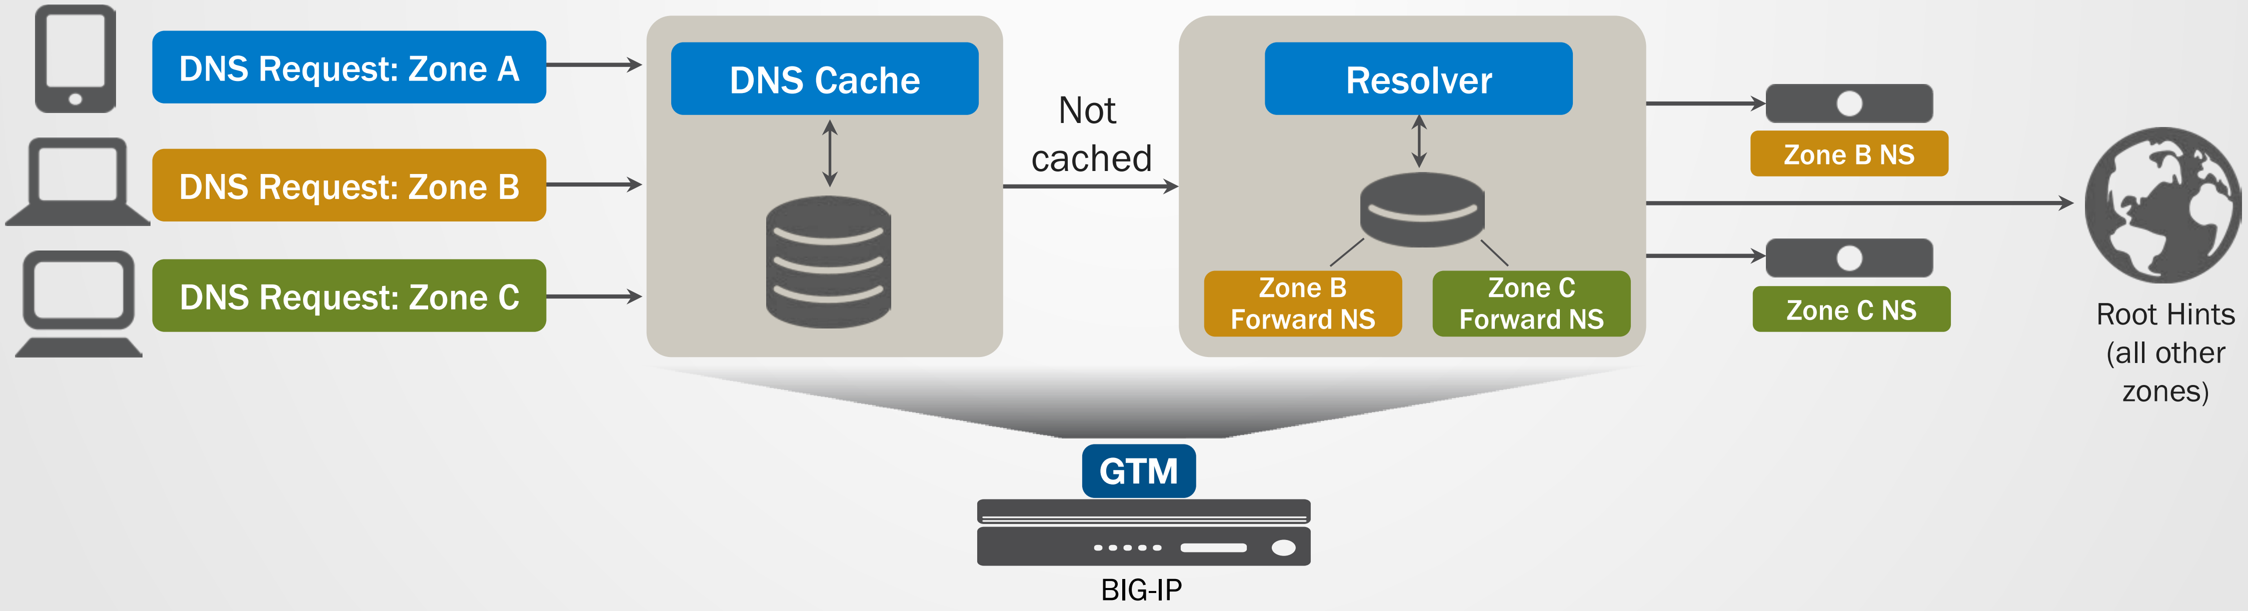

LDNS: Mitigate Malicious Communication

Open Service DNS Query Filtering by Reputation

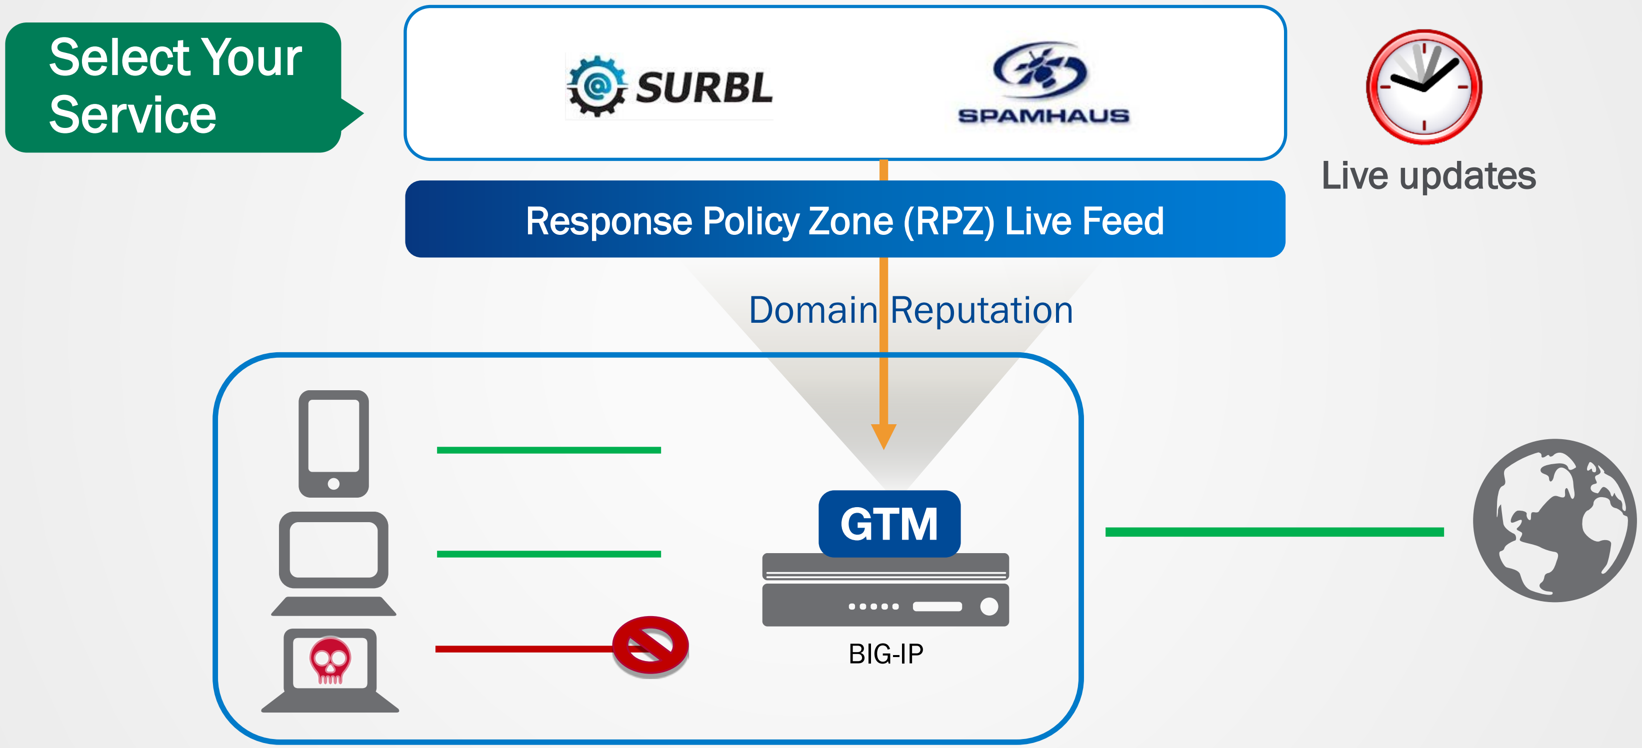

Mitigate DNS threats by blocking access to malicious IPs. Reduce malware and virus infections.

Prevent malware and sites hosting malicious content from ever communicating with a client.

Inhibit the threat at the earliest opportunity. Internet activity starts with a DNS request.

LDNS: Cache/Resolving and DNSSEC Validation

Subscriber – Need Scalable, Secure DNS

- Subscribers utilize DNS when selecting services
- Some name servers sign DNSSEC responses
- Subscribers need quick responses with lower latency
- Offload DNS services to F5 high performance platform for consolidation

Consolidated DNS Services

- DNSSEC validation guarantees authenticity of DNSSEC responses
- Client-side resolver knows the IP address received is authentic when DNSSEC validation is used
- Rapid cached and resolver responses closest to client
- Lower DNS latency leads to lower subscriber churn

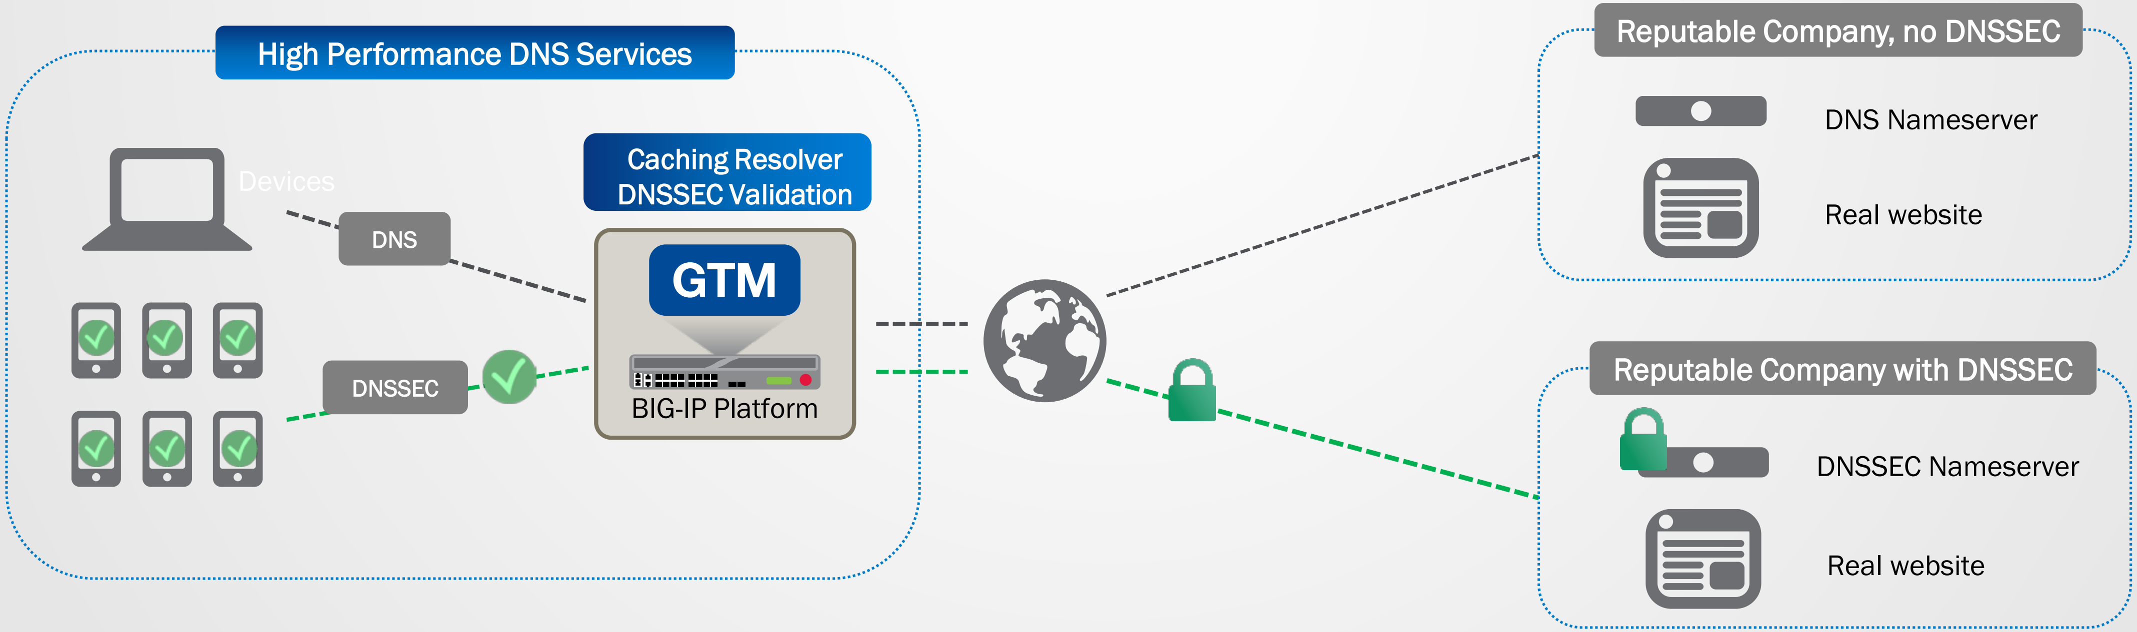

Infrastructure: Automatically Monitor Packet Gateways For Availability

Problem: Manually Remove Packet Gateways

- Many SPs don't monitor the PGW/GGSN from DNS
- SGSN/MME selects an APN by DNS lookup (apn.provider.com)
- DNS responds with the available PGW/GGSN
- Manually remove PGW from record list given to mobile unit

Solution: Automatically Monitor, Remove and Reload

- Higher availability of services
- Closer mapping of network capacity to required load
- Reduced overhead through overprovisioning
- Allows for capacity to be added or removed automatically

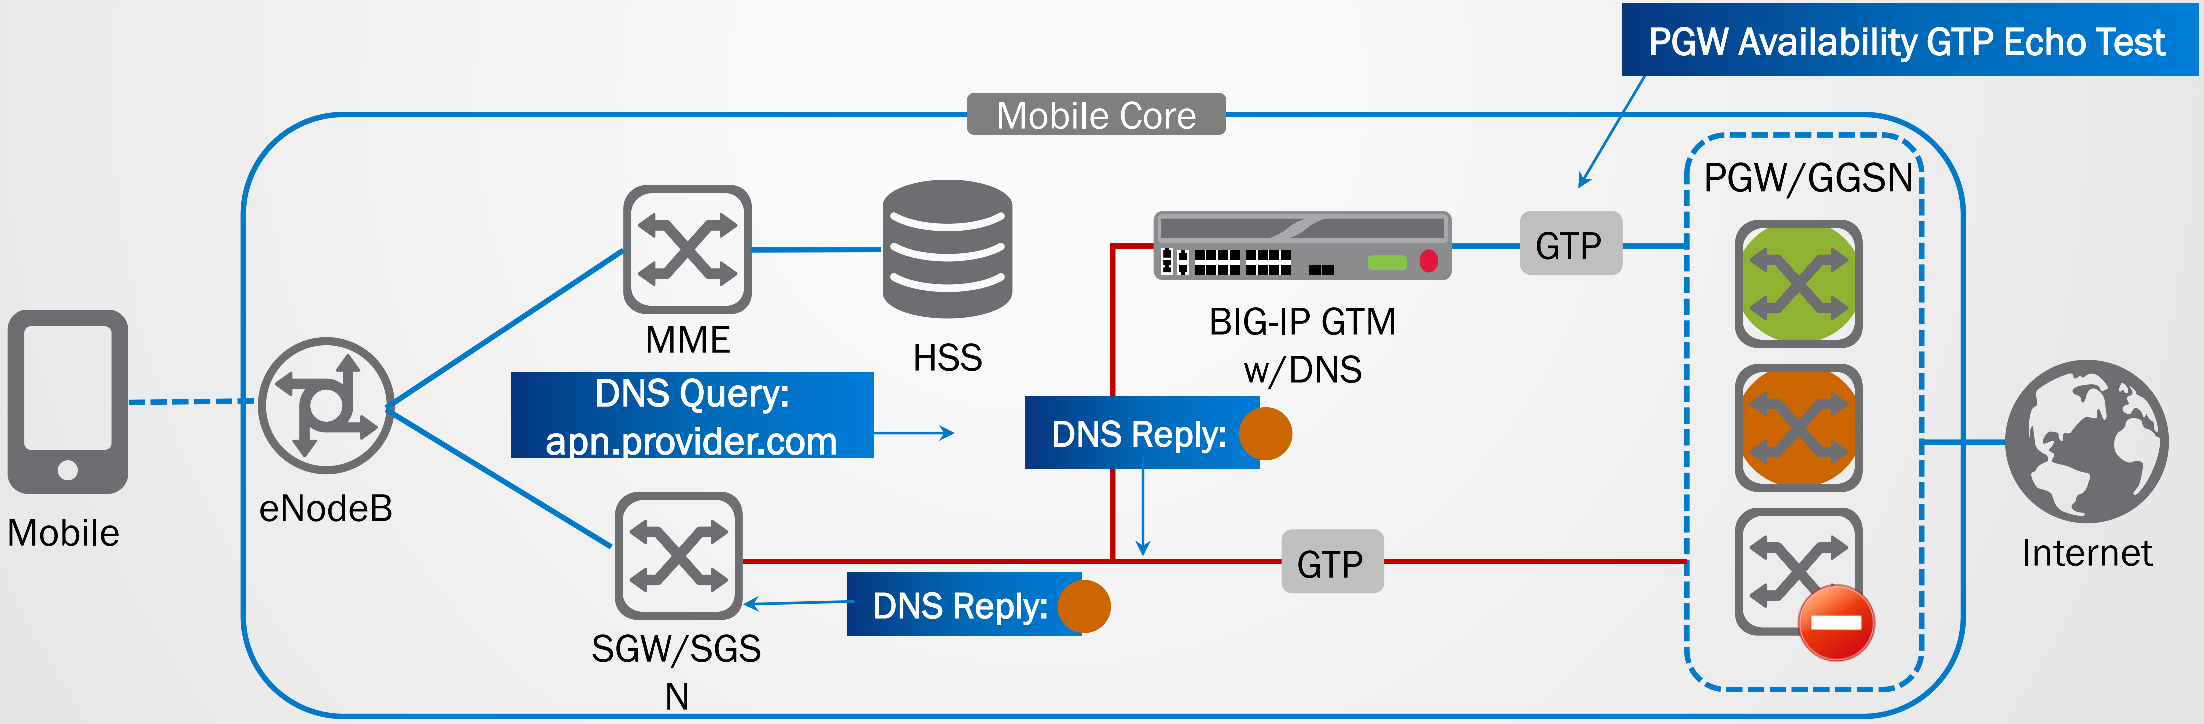

Infrastructure: Proactive Traffic Management

Assure Availability through DNS Rate Shaping: iRules

Protect Critical Infrastructure

- Subscribers can disrupt DNS infrastructure:
  - Maliciously behavior
  - Unintentional / unknowingly via bots and malware
- Protect critical infrastructure from abuse with BIG-IP:
  - Ensure DNS availability for others
  - Mitigate DNS DDoS attacks

Per-client IP Rate Limiting

- Monitor each individual subscriber with specific RPS quota
- Excess queries logged or dropped when quota exceeded
- Support for multiple rate limits to adjust to specific classes of service
- Keep DNS infrastructure responding to legitimate traffic

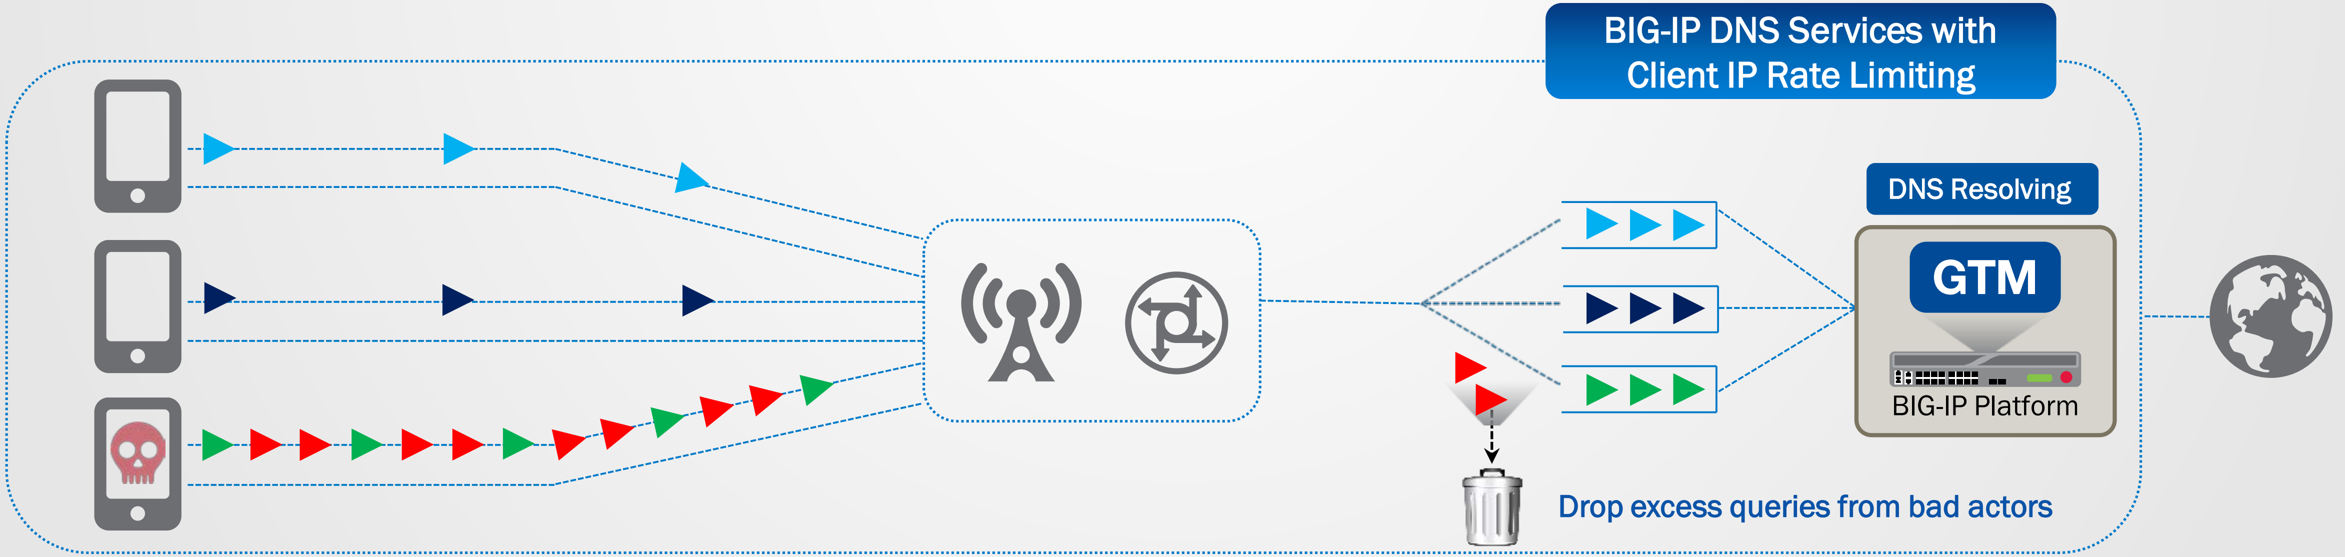

Authoritative: Optimize Service Delivery Globally

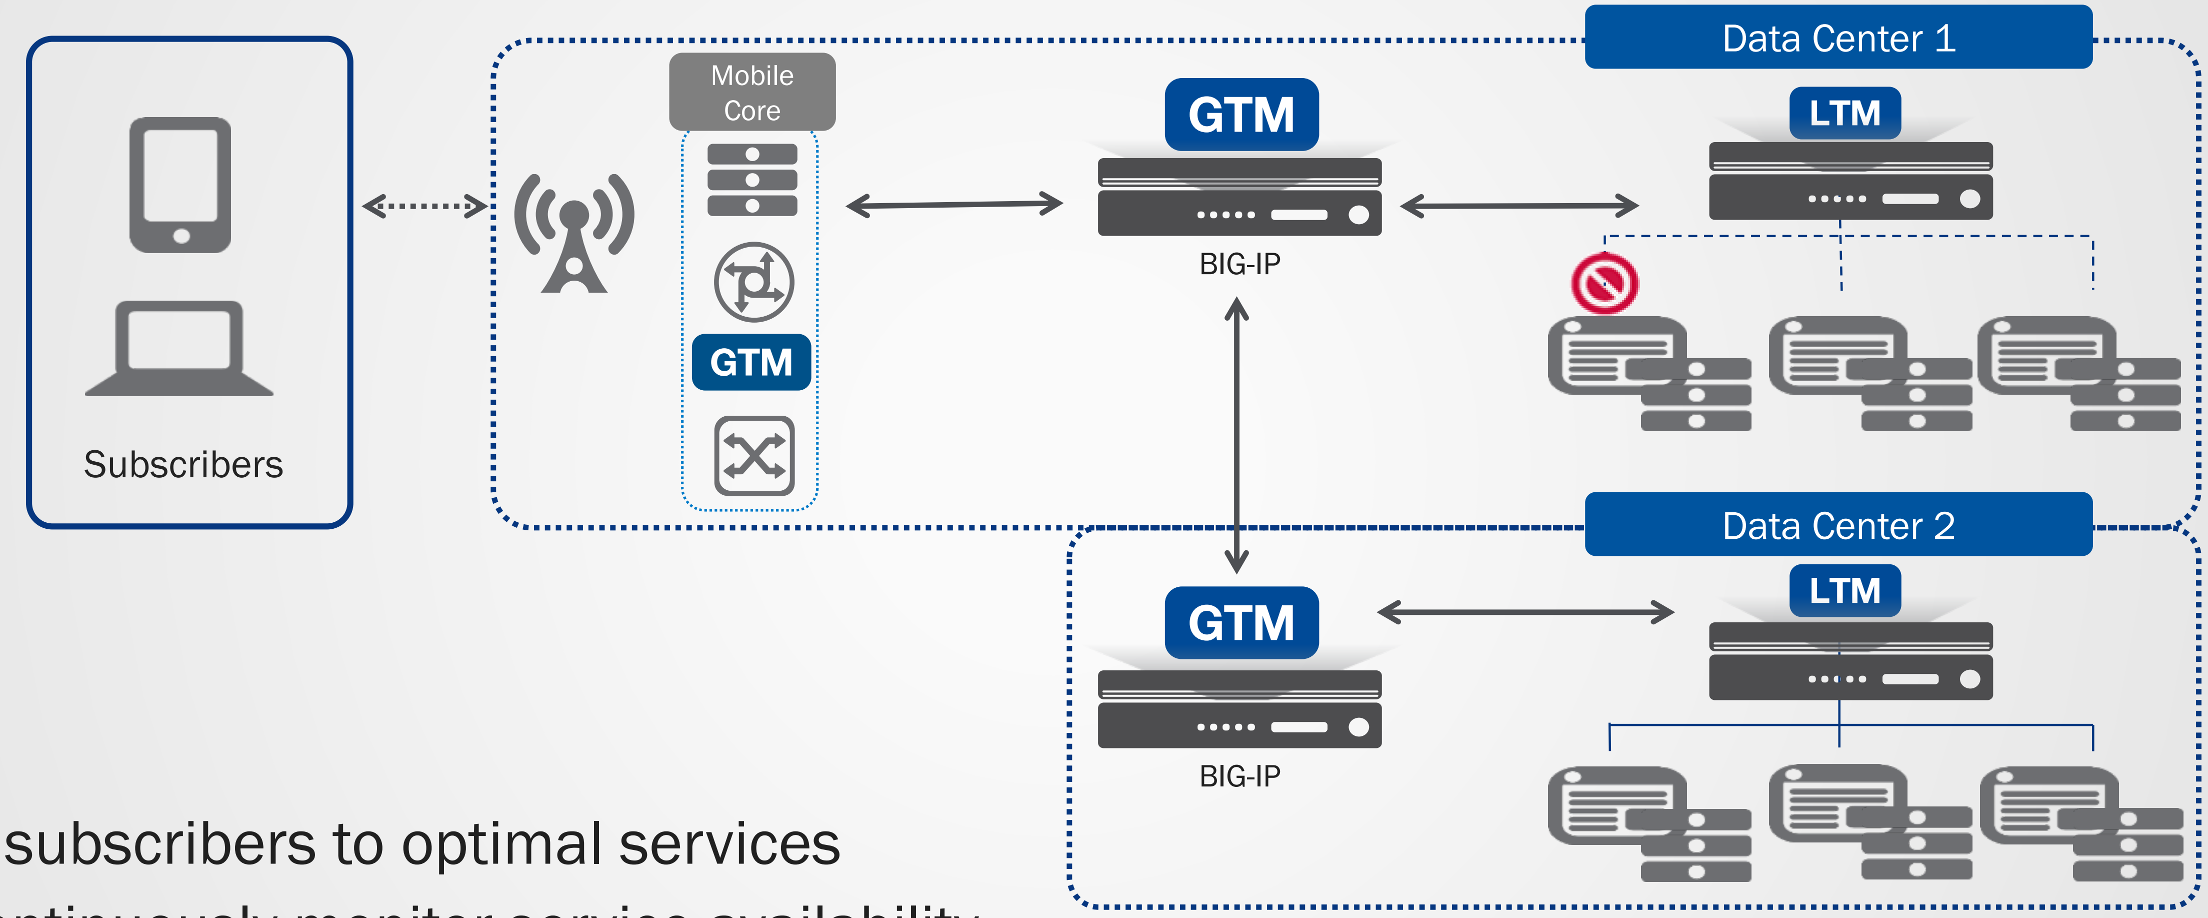

Direct subscribers to optimal services

- Continuously monitor service availability
- Route based on business logic to available services
- Enable LDNS caching and resolving for faster responses

Authoritative: Scaling the F5 DNS Solution through IP Anycast

Solution Overview

- **Expand your DNS infrastructure** to support multiple scalable name servers using a single IP address with IP Anycast
- **Advantage in mitigating against DNS DDoS** as traffic is sent to the closest authoritative DNS server, spreading the load

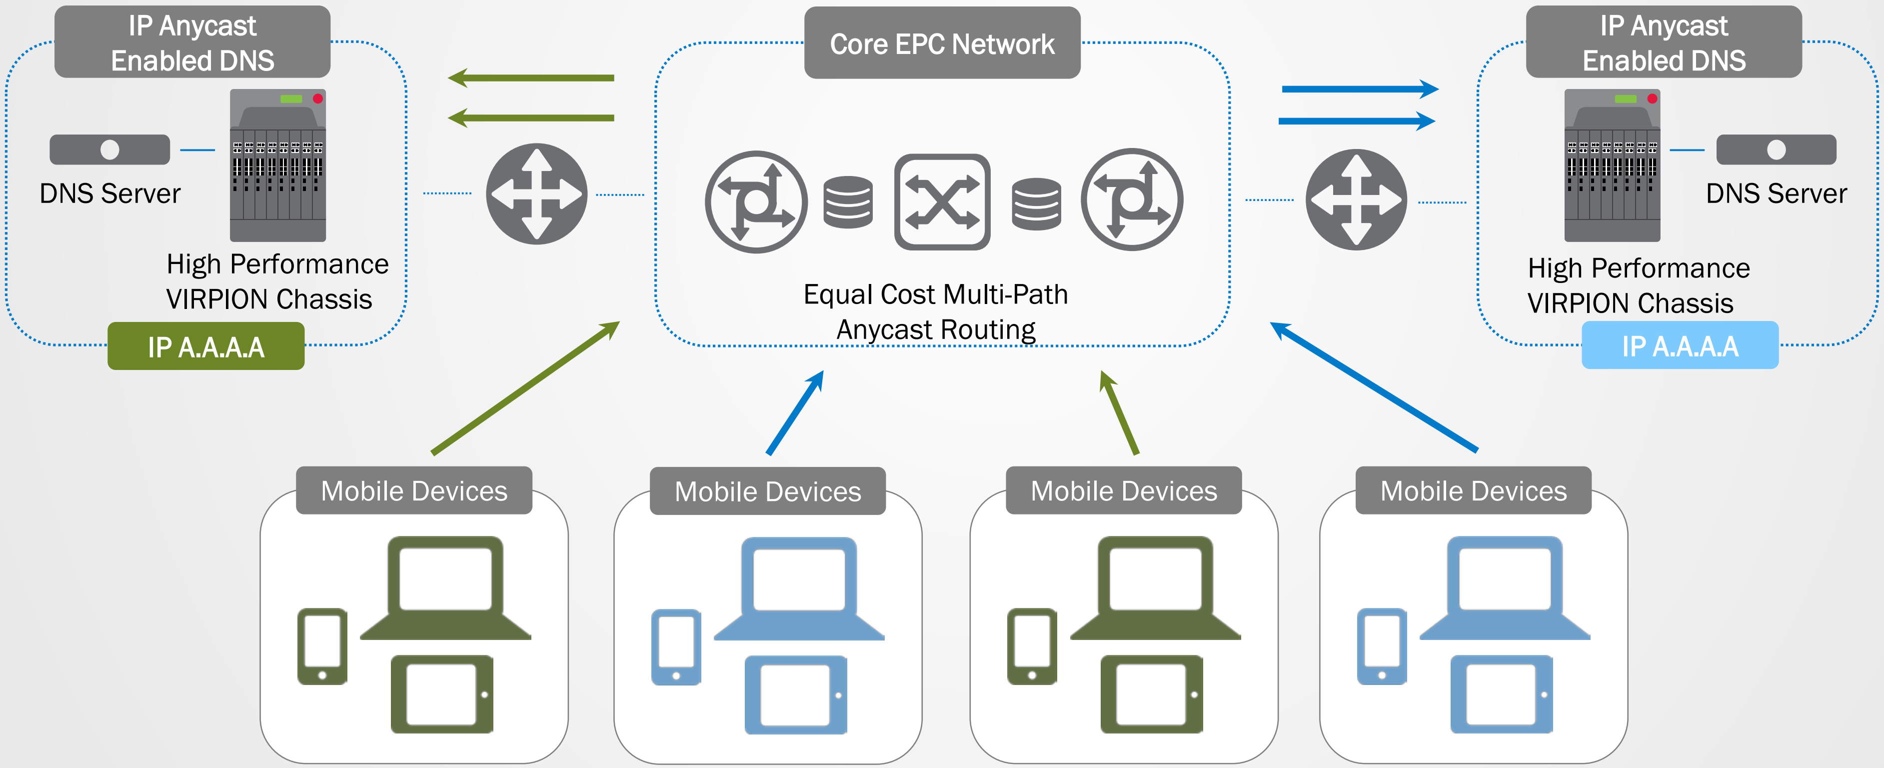

Solutions for an application world.

CONFIDENTIAL

TCP Optimization for Mobile Networks

Ivan Ramos

SP Solutions Architect LatAm

Agenda

- Market Trends – Network & Content Optimization
- TCP in Mobile Networks
- F5 TCP Express – The Full-proxy
- Deployment Models
- Test Result
- Summary

Market Trends
Network and Content Optimization

Mobile Has Unique Challenges

Why is the web so slow on my mobile device?

Mobile Device

- TCP stacks are different on different mobile OS
- JavaScript parsing and execution is relatively slow on mobile devices

Mobile Network

- Higher packet loss rate
- High network latency: 300ms via 3G vs <50ms on LTE
- Connections are made ad-hoc and frequently dropped to preserve spectrum and battery life

Internet

- Low packet loss rate
- Low latency (except for intercontinental traffic)

Application

- Different TCP stacks being used on servers, some of which are not optimal for mobile networks
- SPDY / HTTP2.0

Network and Content Optimization Technologies

Technology	What it provides	Original goals	New goals / trends
Video and Image Optimization	• Video trans-rating • Video trans-coding • Video pacing • Image optimization	Save cost and enhance mobile video experience	User QoE differentiation
Web Optimization	• Content re-ordering • Content inlining • JavaScript minification	Enhance mobile browsing experience	No longer relevant
TCP Optimization	• TCP proxy with profile tuned to the radio characteristics	Maximize utilization of radio assets	Enhance QoE for all users
Transparent Caching	• Caches internet content locally	Cost savings	Enhance QoE for all users

A Changing Environment

SSL / SPDY INCREASE

- In Europe, SSL traffic (HTTPS and SPDY) on mobile networks is currently reaching around 50% of total Internet traffic
- Top web sites such as Google, Facebook, and Twitter use SPDY
- HTTP 2.0 being standardized in IETF with browsers requiring TLS encryption when setting up HTTP 2.0 connections

RISE OF ADAPTIVE BIT RATE VIDEO STREAMING

- Top video sites such as YouTube, Netflix, Hulu, and BBC iPlayer have all embraced ABR video technology
- Video is encoded at different bit rates, client dynamically chooses or changes appropriate bit rate based on network conditions
- ABR video can be "optimized" using bandwidth control techniques

SPDY – Load Web Pages Faster

- HTTP inefficient and outdated
  - HTTP protocol inefficiencies have a negative impact on mobile web browsing experience
  - Due to higher latencies in mobile networks
- SPDY: New app layer protocol developed by Google
  - Overcomes inherent inefficiencies with HTTP
  - Improved performance (~ 20-50%)
  - Good for low bandwidth / high latency mobile networks
  - Forms the basis for HTTP 2.0 in IETF

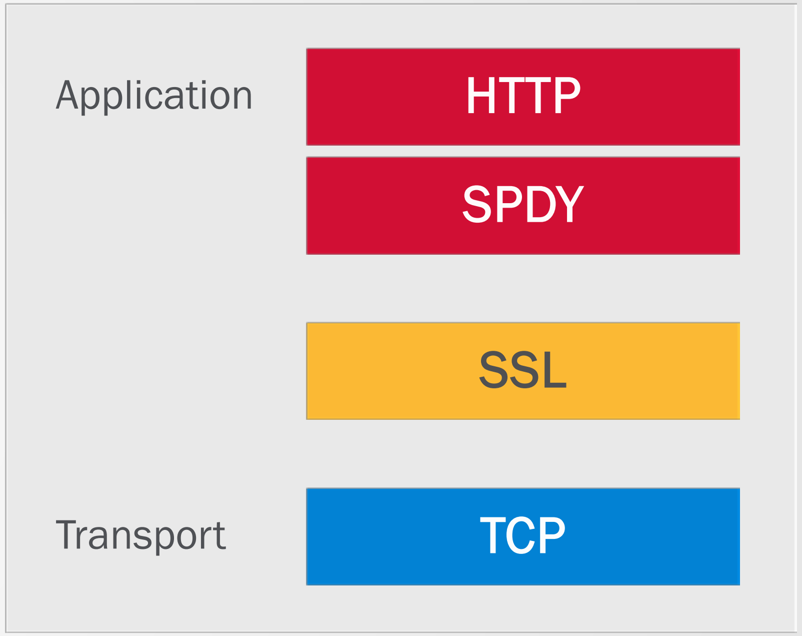

SPDY – For a Better Web Experience

- Multiplexed bi-directional streams within TCP connection
- Fewer network connections required
- HTTP header compression

Reduction in page load times (SPDY vs HTTPS)

	Google News	Google Sites	Google Drive	Google Maps
Median	-43%	-27%	-23%	-24%
5th percentile (fast connections)	-32%	-30%	-15%	-20%
95th percentile (slow connections)	-44%	-33%	-36%	-28%

Source: Google internal testing

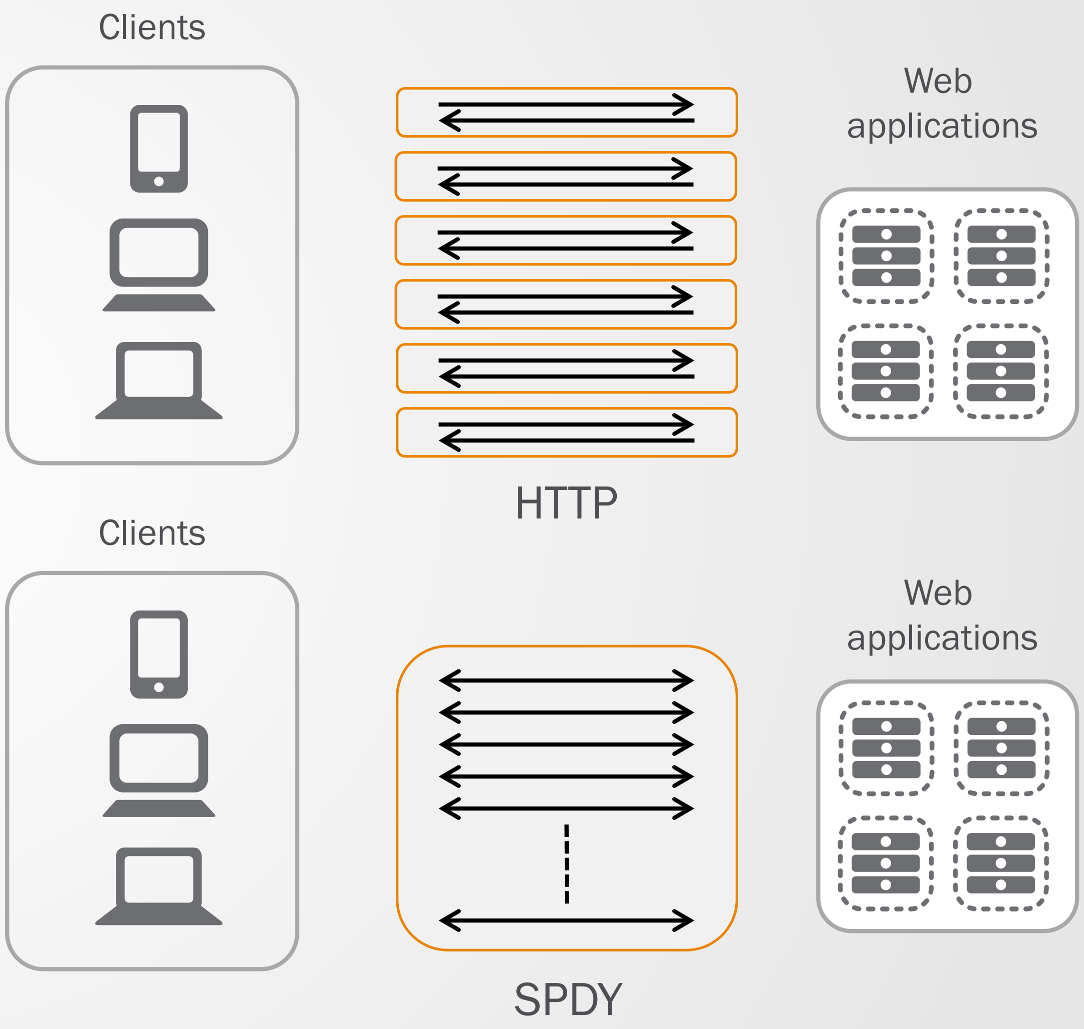

Impact on Optimization Technologies

Video and Web Optimization

Increase of video encryption and ABR video are reducing the benefits of this technology

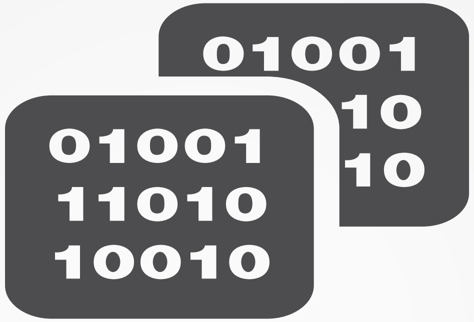

Transparent Caching

Increase of SSL and SPDY on the web are reducing the benefits of this technology

TCP Optimization and Bandwidth Control

Will continue to provide benefits to majority of traffic as > 90% of all traffic rides on top of TCP (including SSL/SPDY)

TCP in Mobile Networks

Specifics of Mobile Broadband Networks (2G/3G/LTE)

- Bandwidth / delay characteristics
  - Relatively low bandwidth compared to VDSL/FTTH (improving with LTE though)
  - Relatively high latency in 2.5G/3G networks : 100ms+ delay
  - Conclusion : high bandwidth-delay product (BDP) – TCP slow start challenges
- Random packet loss (not congestion related) is common in radio networks
  - 3G network designers have implemented link layer retransmission protocols such as RLC (ARQ) to mitigate the 'random' packet loss
  - These techniques while reducing packet loss probability to less than 1% also introduce increased "delay" and "delay variability" which may have an adverse effect on standard TCP stacks
  - Result : RLC techniques contributing to high (and variable) BDP in mobile networks

Specifics of Mobile Broadband Networks (2G/3G/LTE)

- Bufferbloat issues
  - Most TCP stacks use loss-based congestion control – sender will only slow down when packet loss is observed
  - The exceptionally large buffers in cellular networks – to accommodate bursty traffic and absorb channel variability - along with link layer retransmission conceals packet losses from TCP senders
  - Result : the TCP sender continues to increase its sending rate even if it has already exceeded the bottleneck link capacity since all of the overshoot packets are absorbed by the buffers. This results in up to several seconds of round trip delays.
- Mobility / inter-RAT handovers : problematic for TCP stacks as it requires dynamic changes
- Bandwidth oscillations : allocate/deallocate resources to users that want bandwidth at the same time

Average vs Peak LTE Speeds

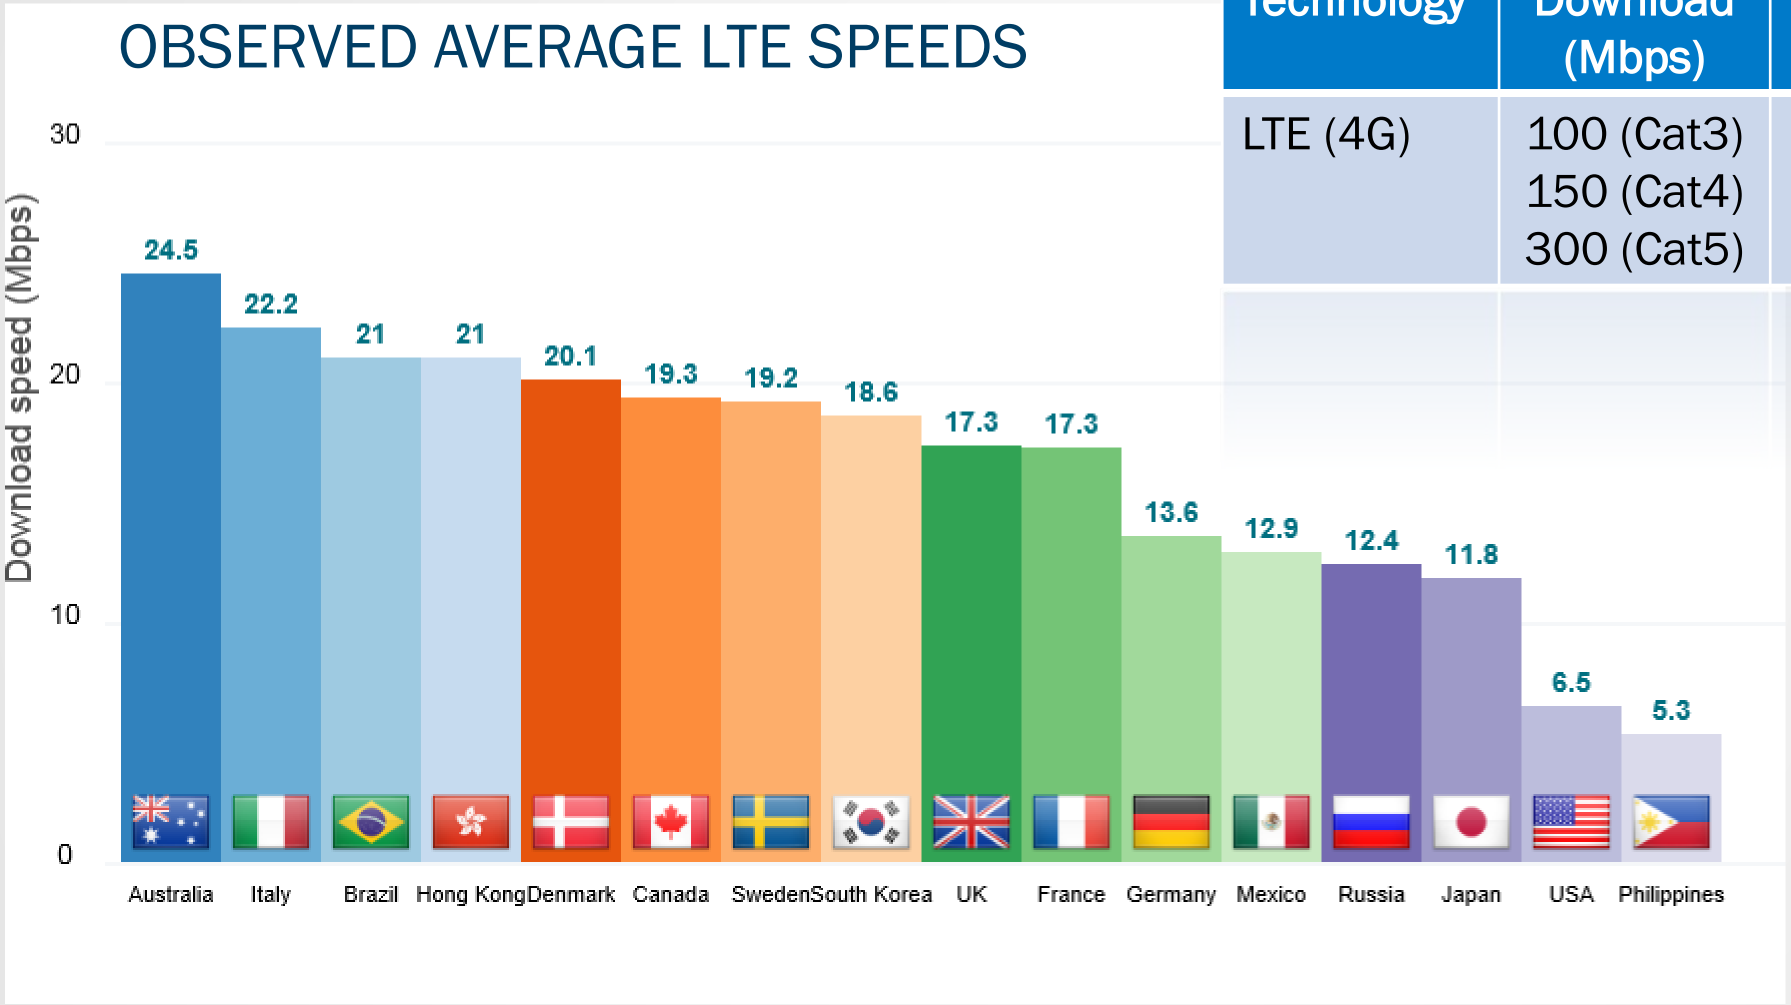

THEORETICAL PEAK SPEEDS

Technology	Download (Mbps)	Upload (Mbps)
LTE (4G)	100 (Cat3) 150 (Cat4) 300 (Cat5)	50 (Cat3) 50 (Cat4) 75 (Cat5)

http://opensignal.com/reports/state-of-lte-q1-2014/

 CONFIDENTIAL

Average vs Peak 3G Speeds

THEORETICAL PEAK SPEEDS

Technology	Download (Mbps)	Upload (Mbps)
HSPA+ (3G)	21 42 84	5.8 11.5 22

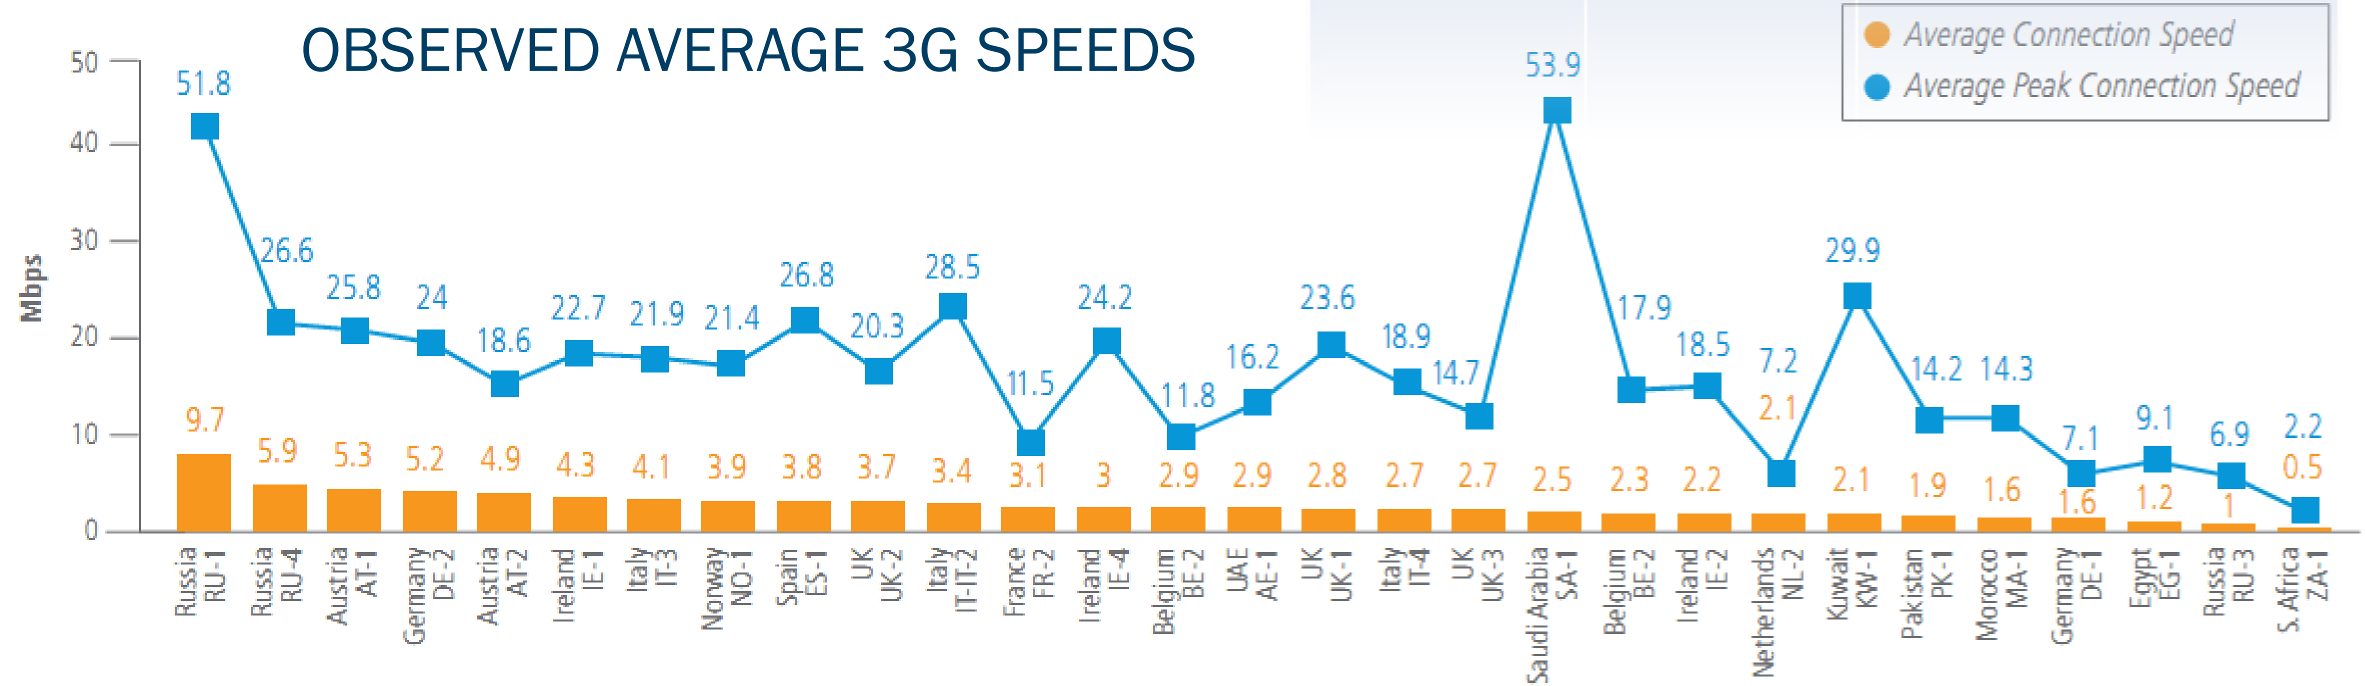

NOTE: The average and average peak connection speeds presented above are based on end-user connections from those mobile networks to the Akamai Intelligent Platform, and are not necessarily representative of a single provider's full set of service offerings or capabilities.

http://www.akamai.com/dl/akamai/q2_2013_soti_infographic.pdf

Worldwide Average Speeds : 2G vs 3G vs LTE

OBSERVED AVERAGE 2G/3G/LTE SPEEDS

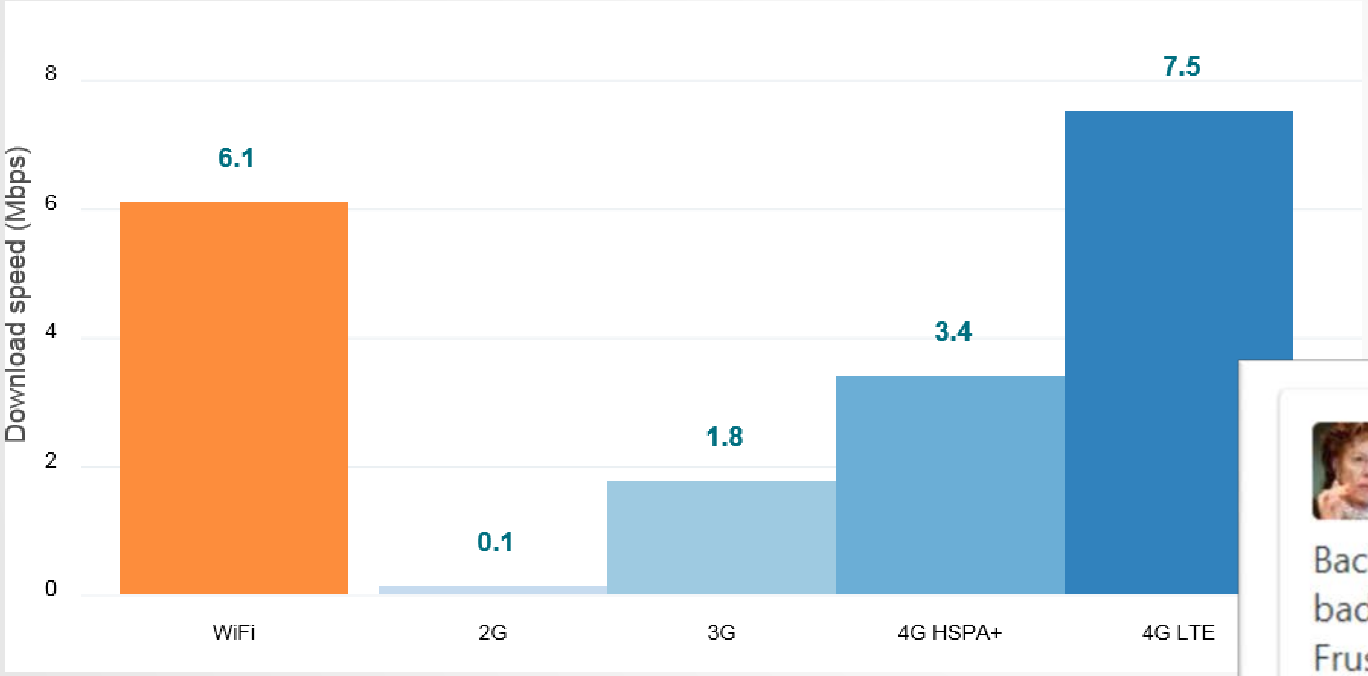

http://opensignal.com/reports/state-of-lte-q1-2014/

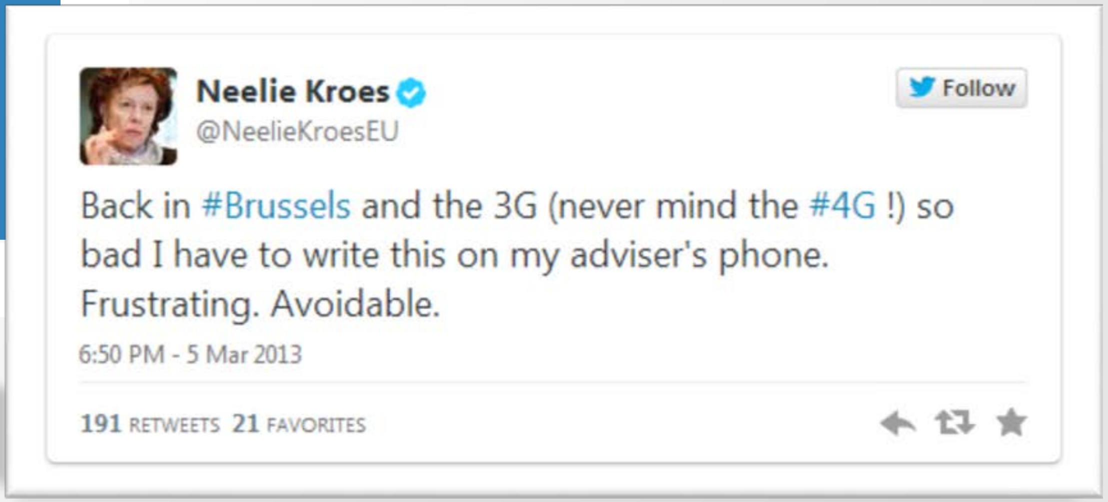

Latency in 3G/4G Networks – Examples from USA

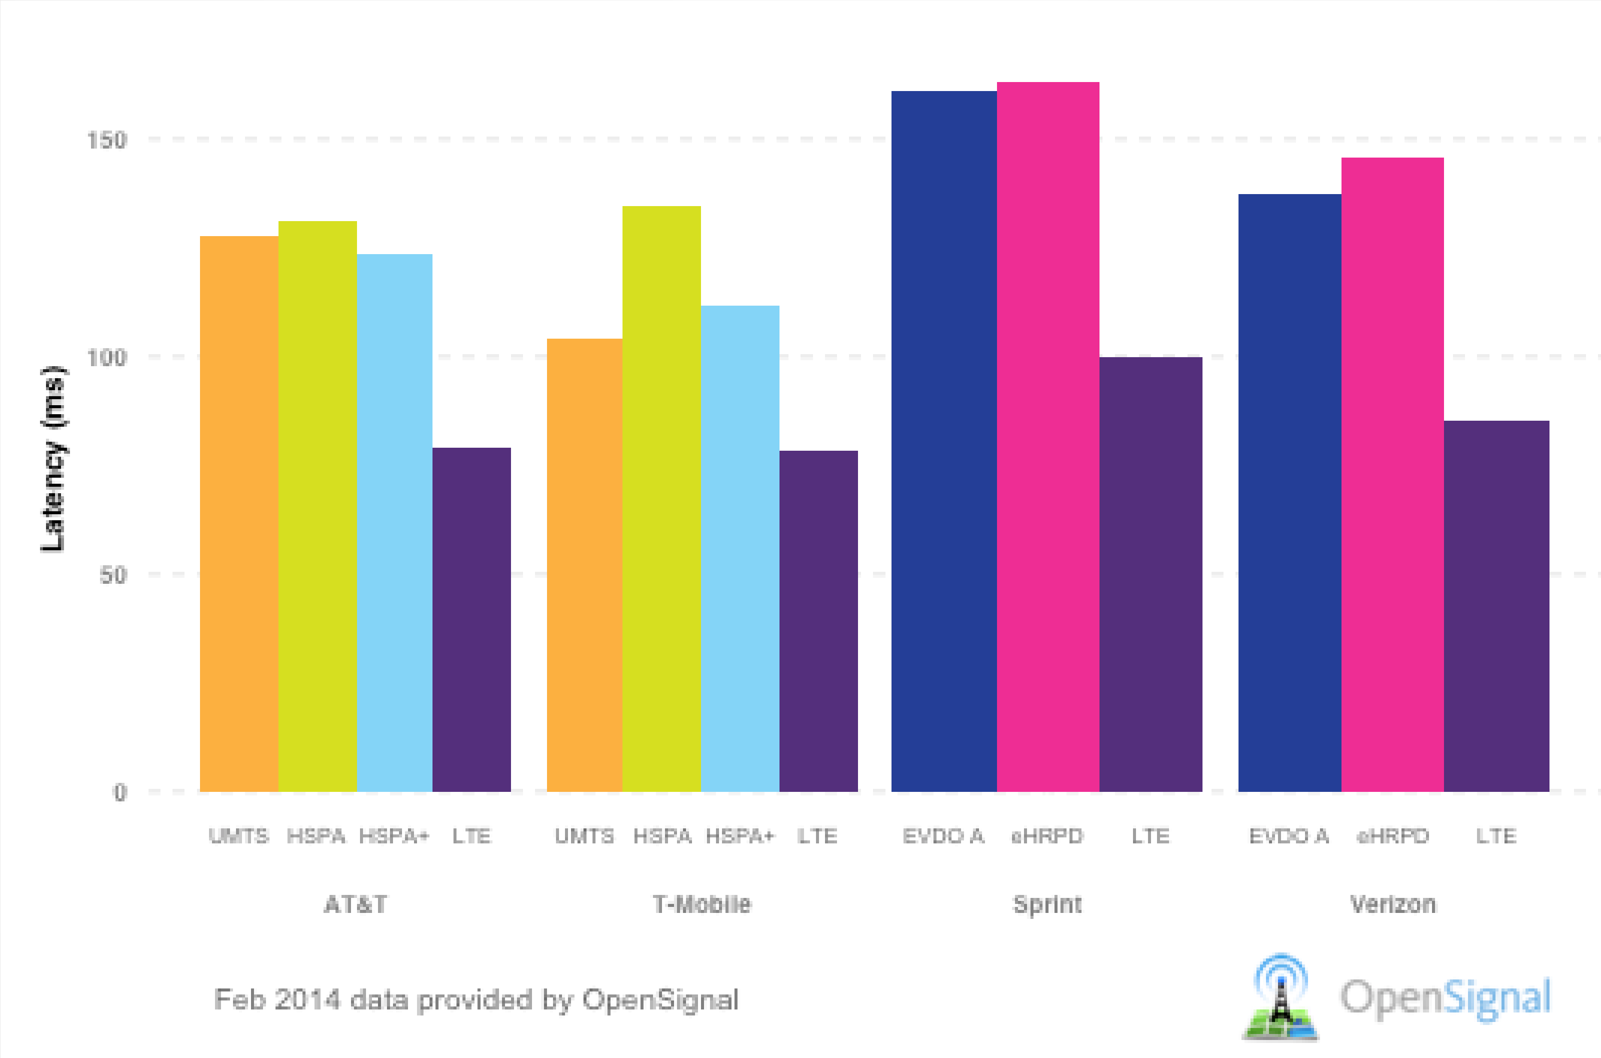

Source : http://www.fiercewireless.com/special-reports/3g4g-wireless-network-latency-comparing-verizon-att-sprint-and-t-mobile-feb

Some stats from US LTE networks

- TCP dominates the data set
  - 95.3% of all flows are TCP
  - 97.2% of all bytes are TCP
- Within TCP
  - HTTP : 76.6% (bytes), 50.1% (flows)
  - HTTPS : 14.8% (bytes), 42.1% (flows)
- TCP flows / payload sizes
  - Top 0.6% of flows ranked by payload sizes (each with >1MB data) accounts for 61.7% of total downlink bytes
  - Top 5% downlink flows
    - At least 89.5KB of data
    - 80.3% is HTTP
    - 74.4% : video or audio

TCP ACCOUNTS FOR 95% OF ALL MOBILE DATA TRAFFIC

Source: http://www-personal.umich.edu/~hjx/file/sigcomm13.pdf

Impact of Latency – Web Page Load Times

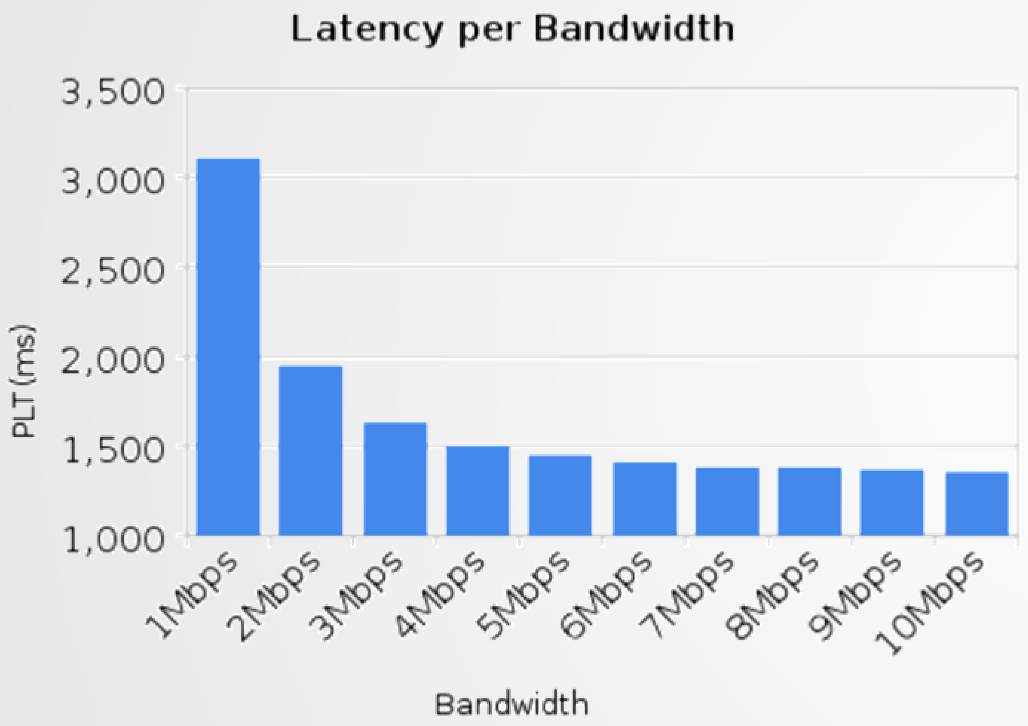

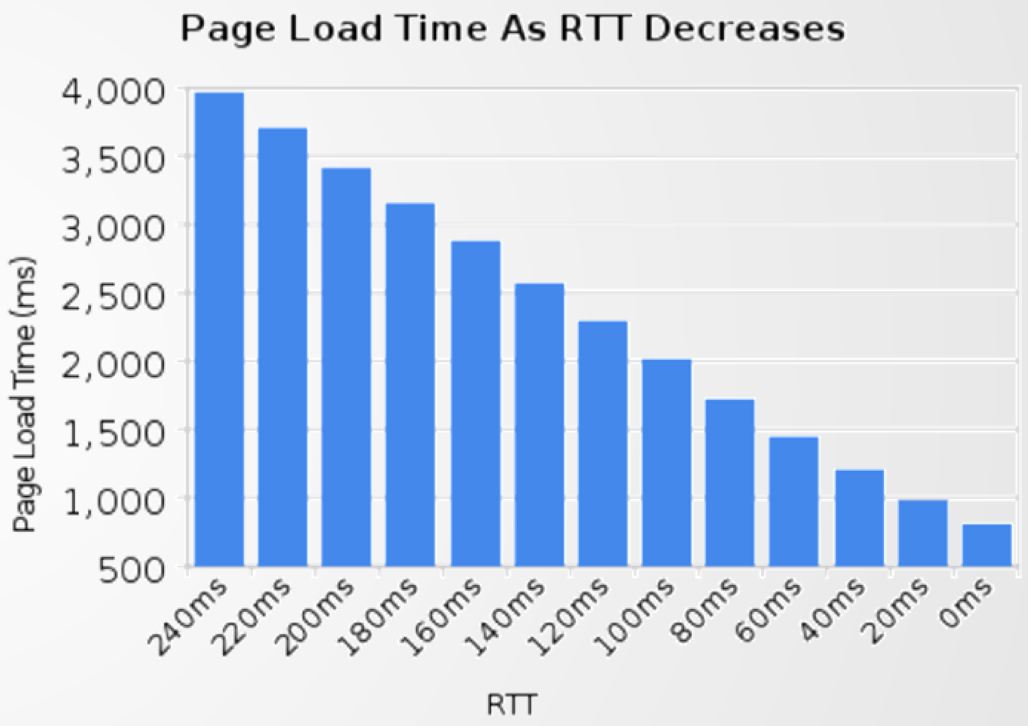

Source: Ilya Grigorik, Google

Impact of Packet Loss – Throughput Degradation

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

Avg HTTP response size 16 kB (3 round trips)

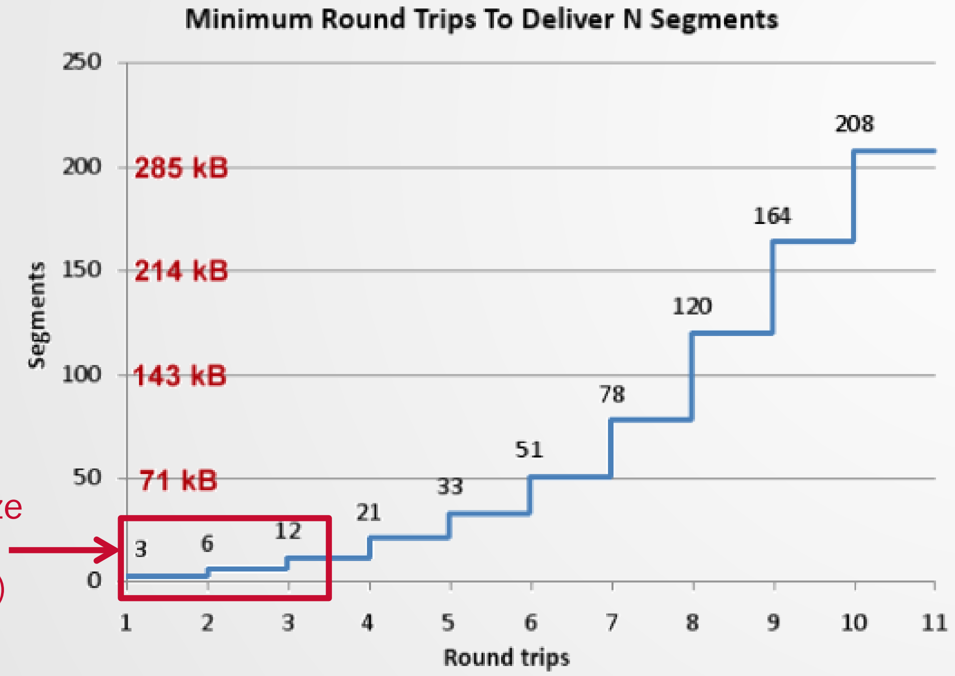

In mobile networks packet loss does not necessarily imply congestion

Packet Loss

Exponential growth

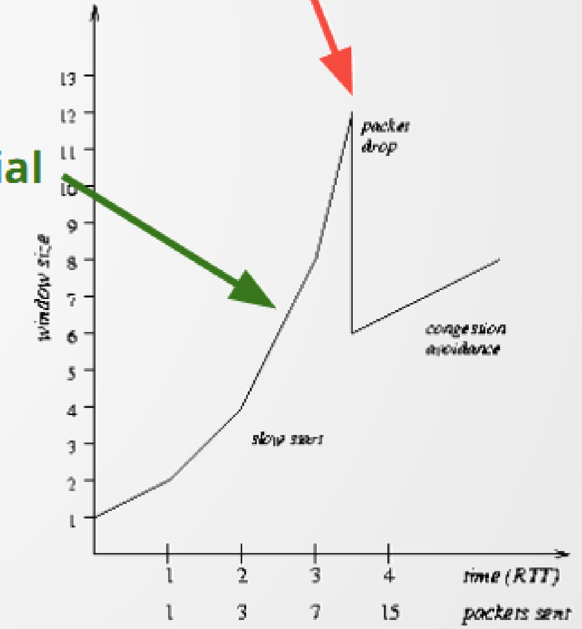

Source: Ilya Grigorik, Google

Conclusions

- Mobile networks have a large BDP
  - Tune your TCP buffers accordingly
- Mobile networks can exhibit random packet loss
  - Choose a TCP congestion control algorithm/technique that takes this into account (don't get into slow start upon random packet loss)
- Mobile networks can suffer from buffer bloat issues
  - Choose a TCP congestion control algorithm that does not rely solely on packet loss
  - Enable TCP rate shaping to ensure 'smoother' delivery packets (less strain on buffers)
- Mobile networks have relatively high latency
  - Tune your settings to increase performance and web page load times (window size, initial congestion window, ...)
- Real life mobile performance is very 'variable' – room for market differentiation !

F5 TCP Express – The Full Proxy

“Especially the heart of TCP, namely flow control and retransmission mechanisms, may cause problems over wireless interfaces. These problems originate mainly because the basic TCP assumes that all packet losses are due to network congestion, not bit errors. When this assumption is combined with the rough flow control scheme of TCP, the performance of TCP transmissions over wireless networks can be severely degraded.

F5 TMOS Full-Proxy Architecture

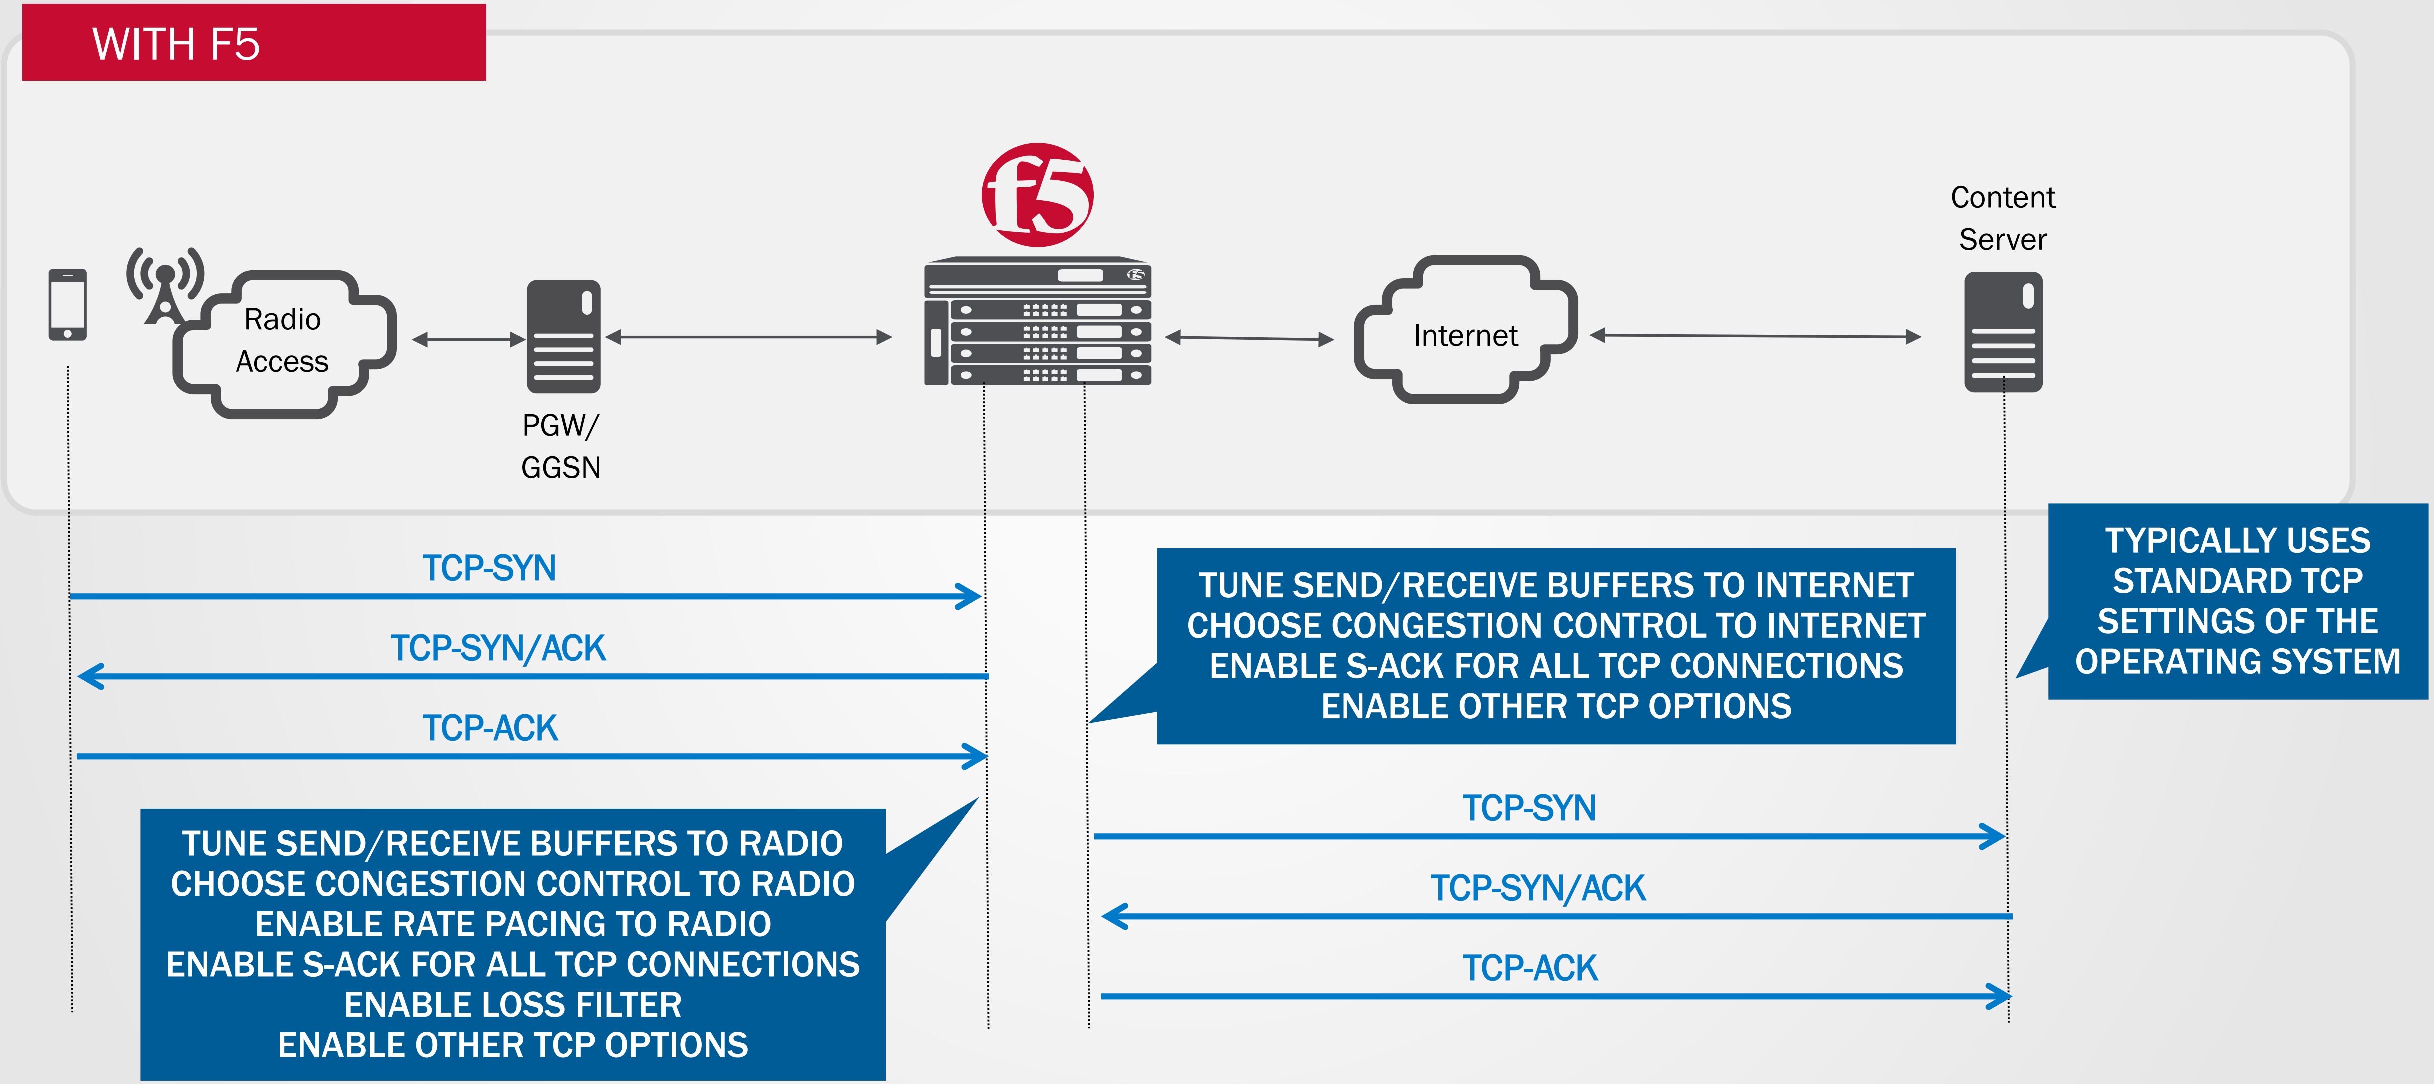

F5 TMOS Full-Proxy Architecture – TCP Profile Settings

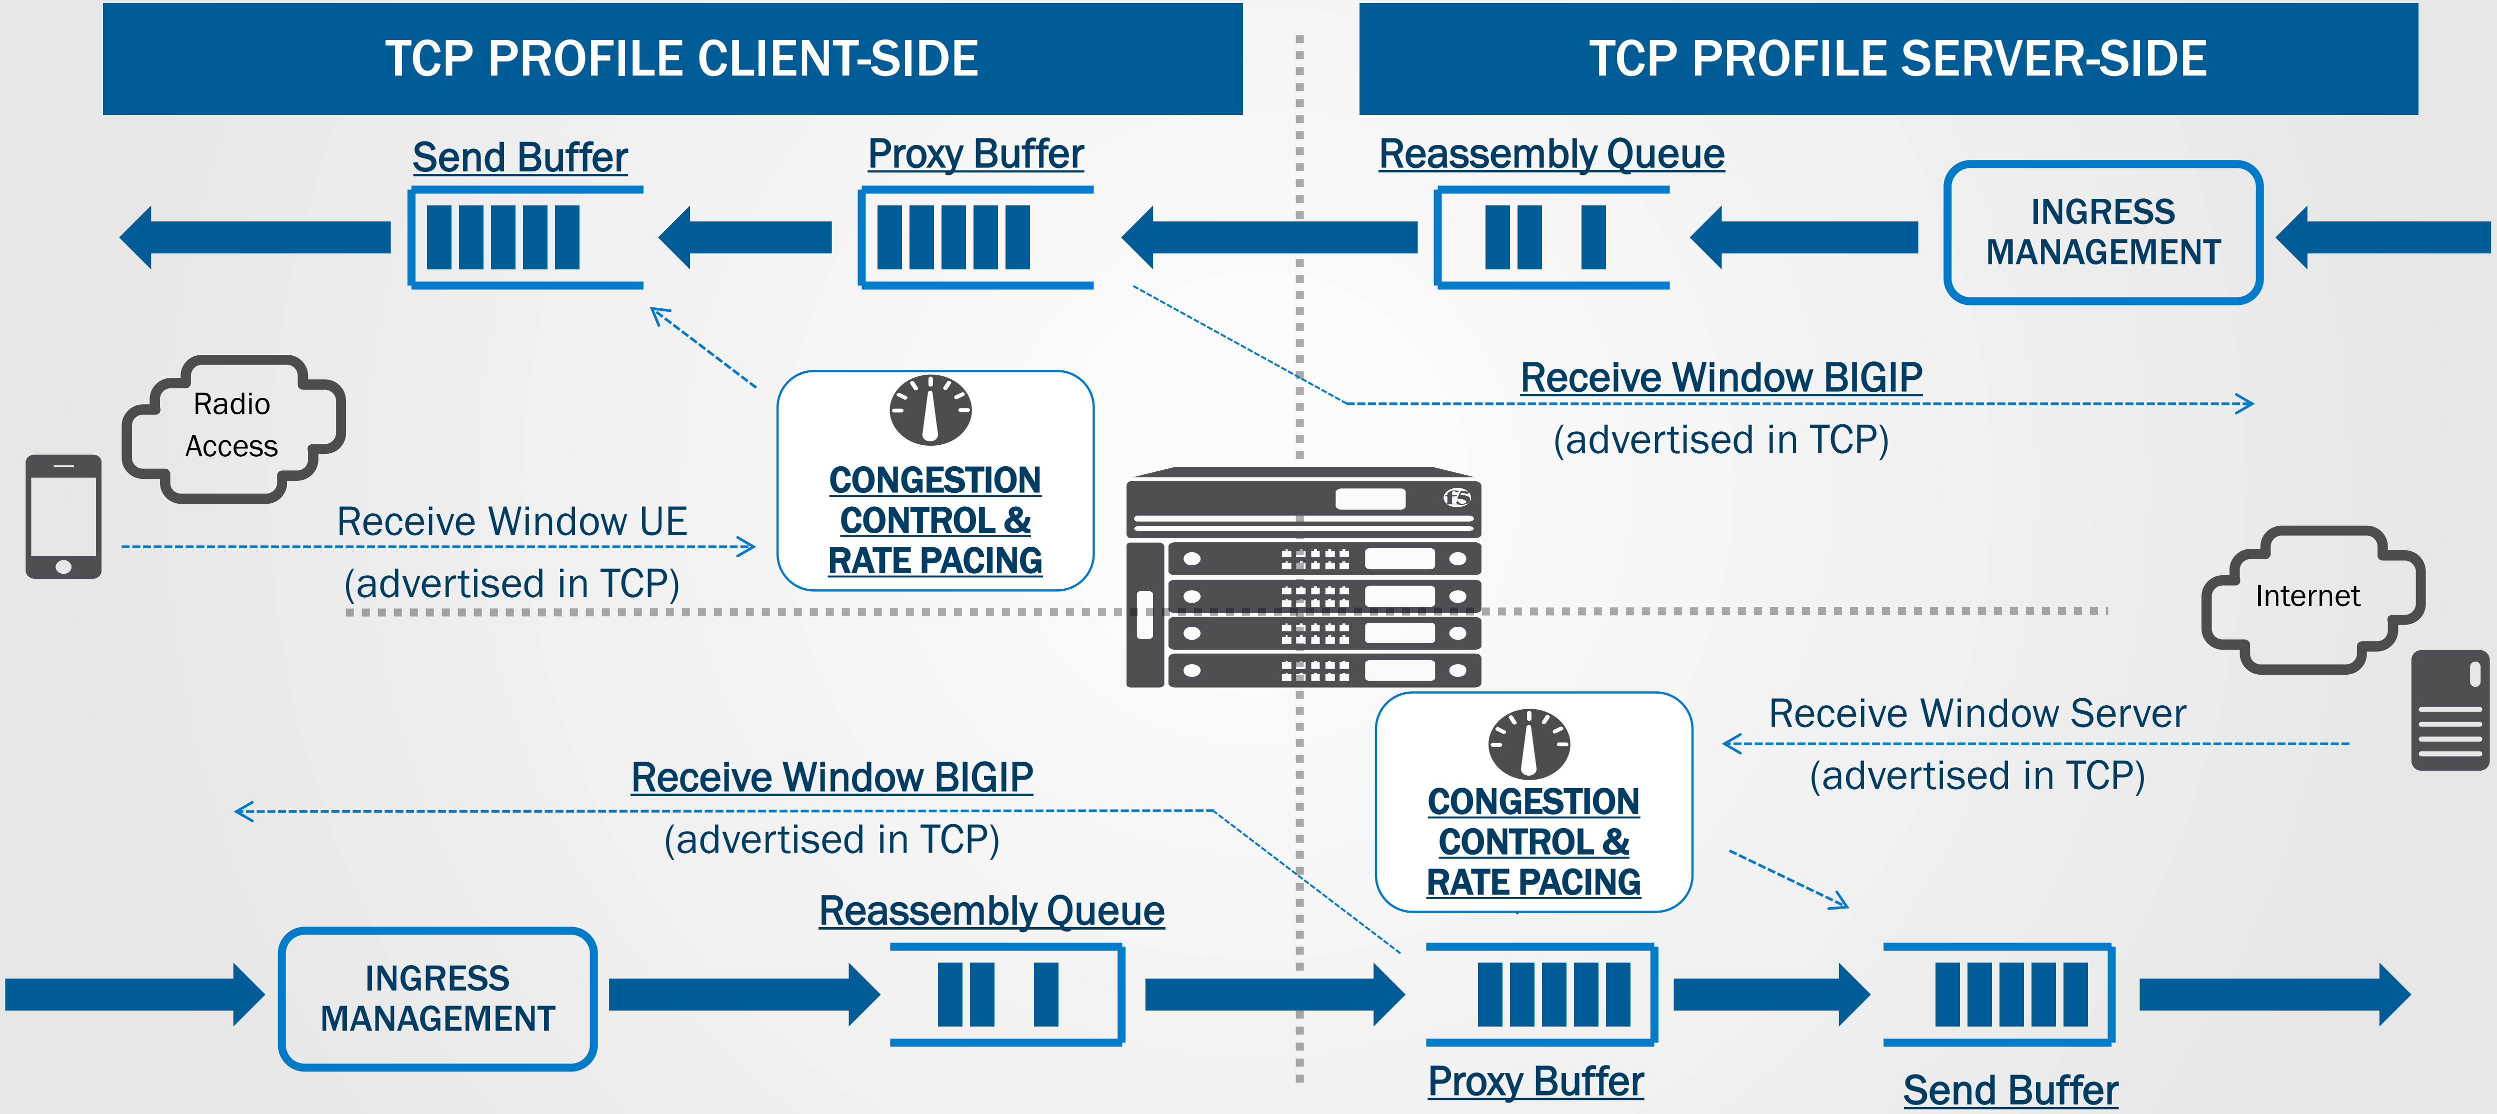

F5 TCP Profile Settings - Snapshot

Memory Management

Proxy Buffer High	49152 bytes
Proxy Buffer Low	32768 bytes
Receive Window	65535 bytes
Send Buffer	65535 bytes

Connection Setup

Deferred Accept	
Proxy Maximum Segment	
Proxy Options	
Verified Accept	

Data Transfer

Acknowledge on Push	Enabled
Delayed Acks	Enabled
Initial Receive Window Size	0 MSS units
Max Segment Size (MSS)	1460 bytes
Nagle's Algorithm	

Congestion Control

Appropriate Byte Counting (RFC 3465)	Enabled
Congestion Metrics Cache	Enabled
Congestion Control	High Speed
Delay Window Control	
Explicit Congestion Notification	
Initial Congestion Window Size	0 MSS units
Packet Loss Ignore Burst	0 packet count
Packet Loss Ignore Rate	0 packets lost per million
Rate Pace	
Slow Start	Enabled
Timestamps Extension for High Performance (RFC 1323)	Enabled

The New TCP Express – TMOS Full Proxy Architecture

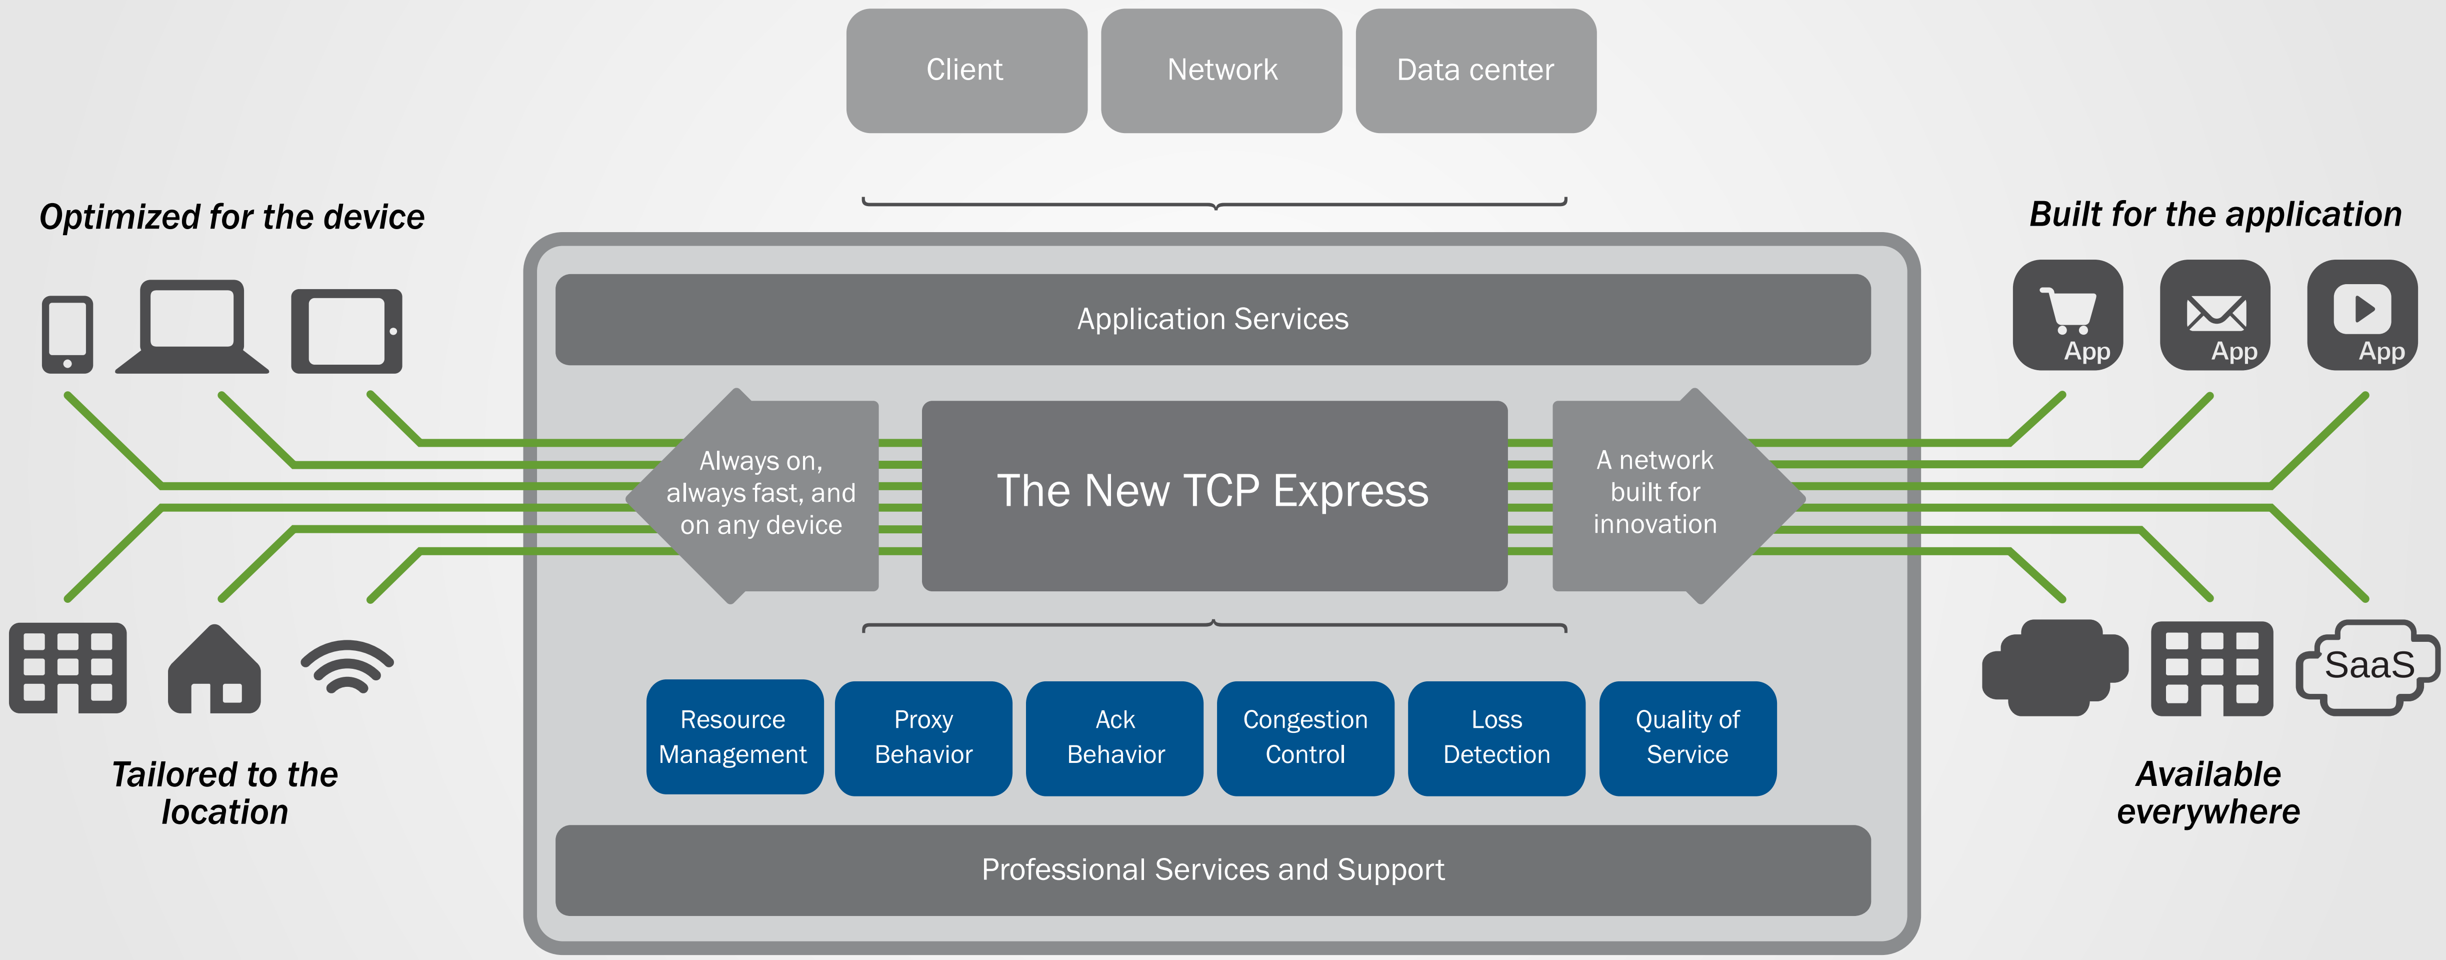

The New TCP Express

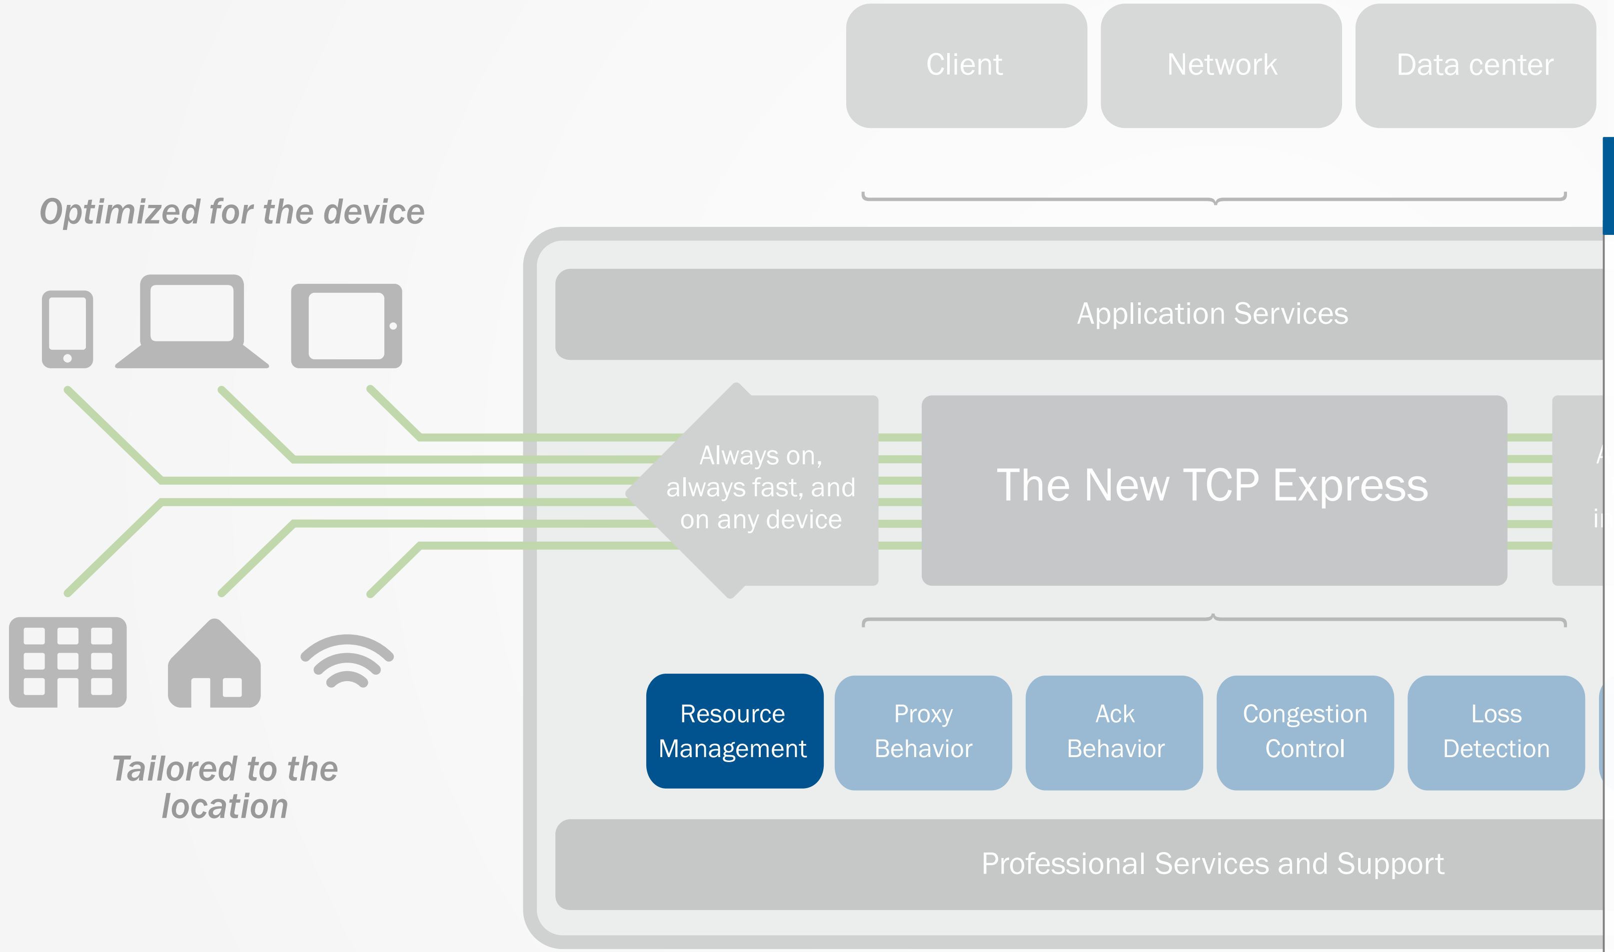

RESOURCE MANAGEMENT

- Drop packets under pressure*
- Timer management
- Memory management

*New in 11.5

The New TCP Express

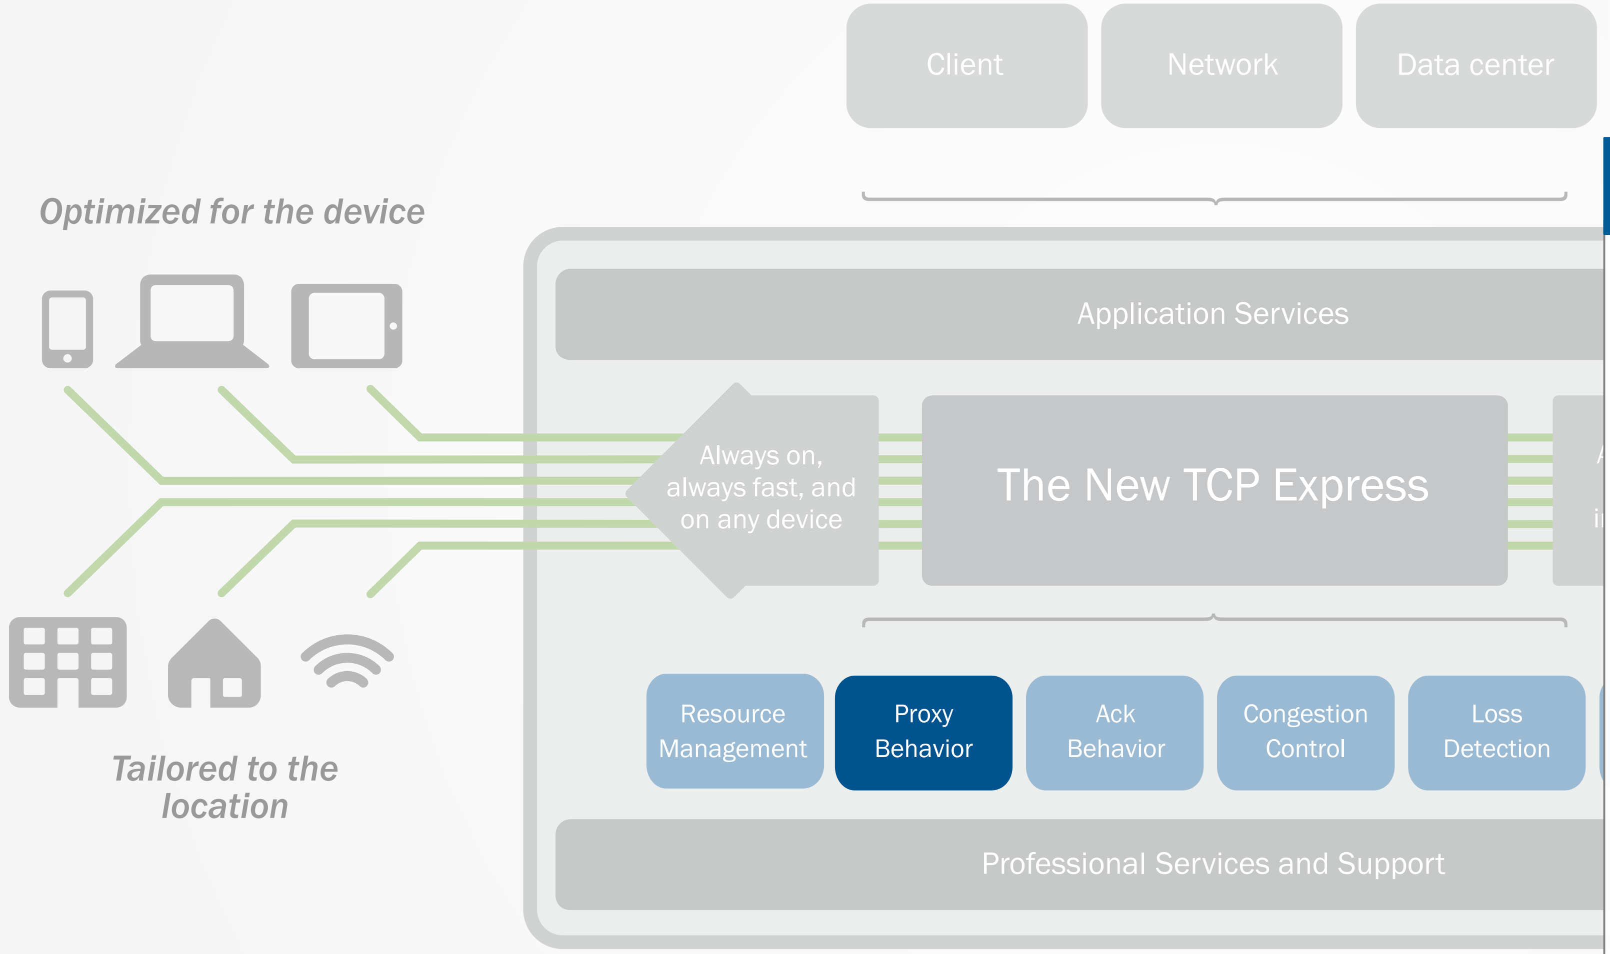

PROXY BEHAVIOR

- **Multi-Path TCP (MPTCP)***
- **Maximum Segment Size (MSS)***
- **Full proxy**

*New in 11.5

Multipath TCP
Mobility

The User chooses which path to transmit
Also independently
The Server chooses which path to transmit on

What's New

- The ability to connect and maintain a continuous connection to the internet over multiple wired and wireless connections
- Shim on the TCP which allows other TCP connections to join in Parallel
- Needed at client and Server (Ios 7 and siri use it)

Use Case

- Device initially connects to site over mobile network.
- Device comes in range of wifi, associates with and connects over Wifi
- Congestion control favors high bandwidth Wifi path
- Device disconnects from Wifi but continues to use 3G network
- Internal to External to internal as long as the app is going to the same BIG-IP

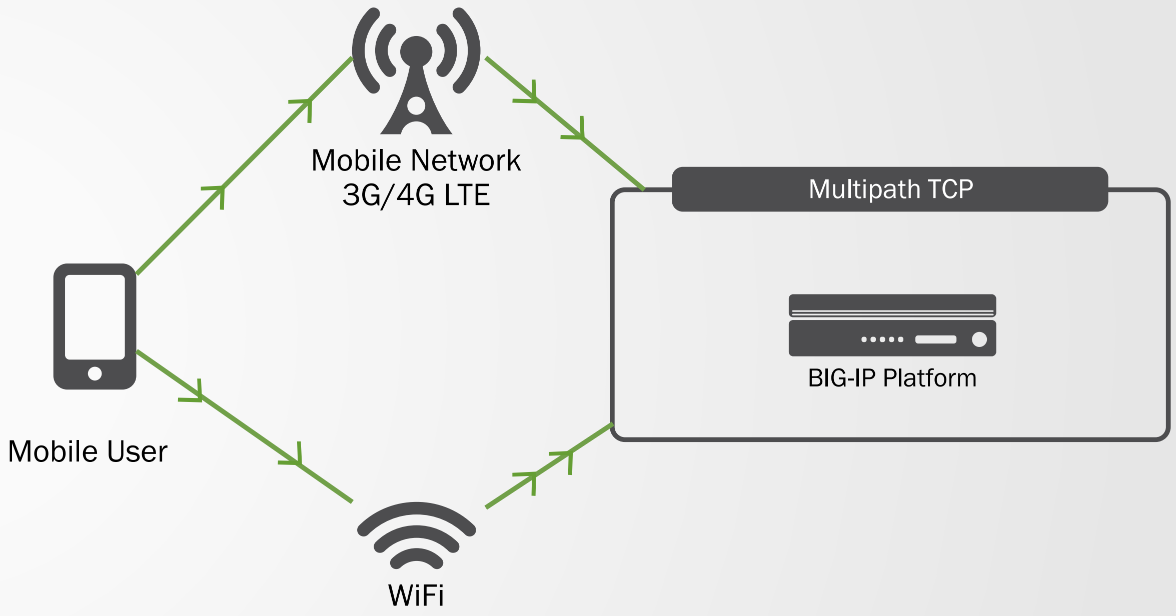

It's arguably the first and most important change to the low-level architecture of the internet to reflect the fact that our connections to it are more mobile and wireless than ever.

The New TCP Express

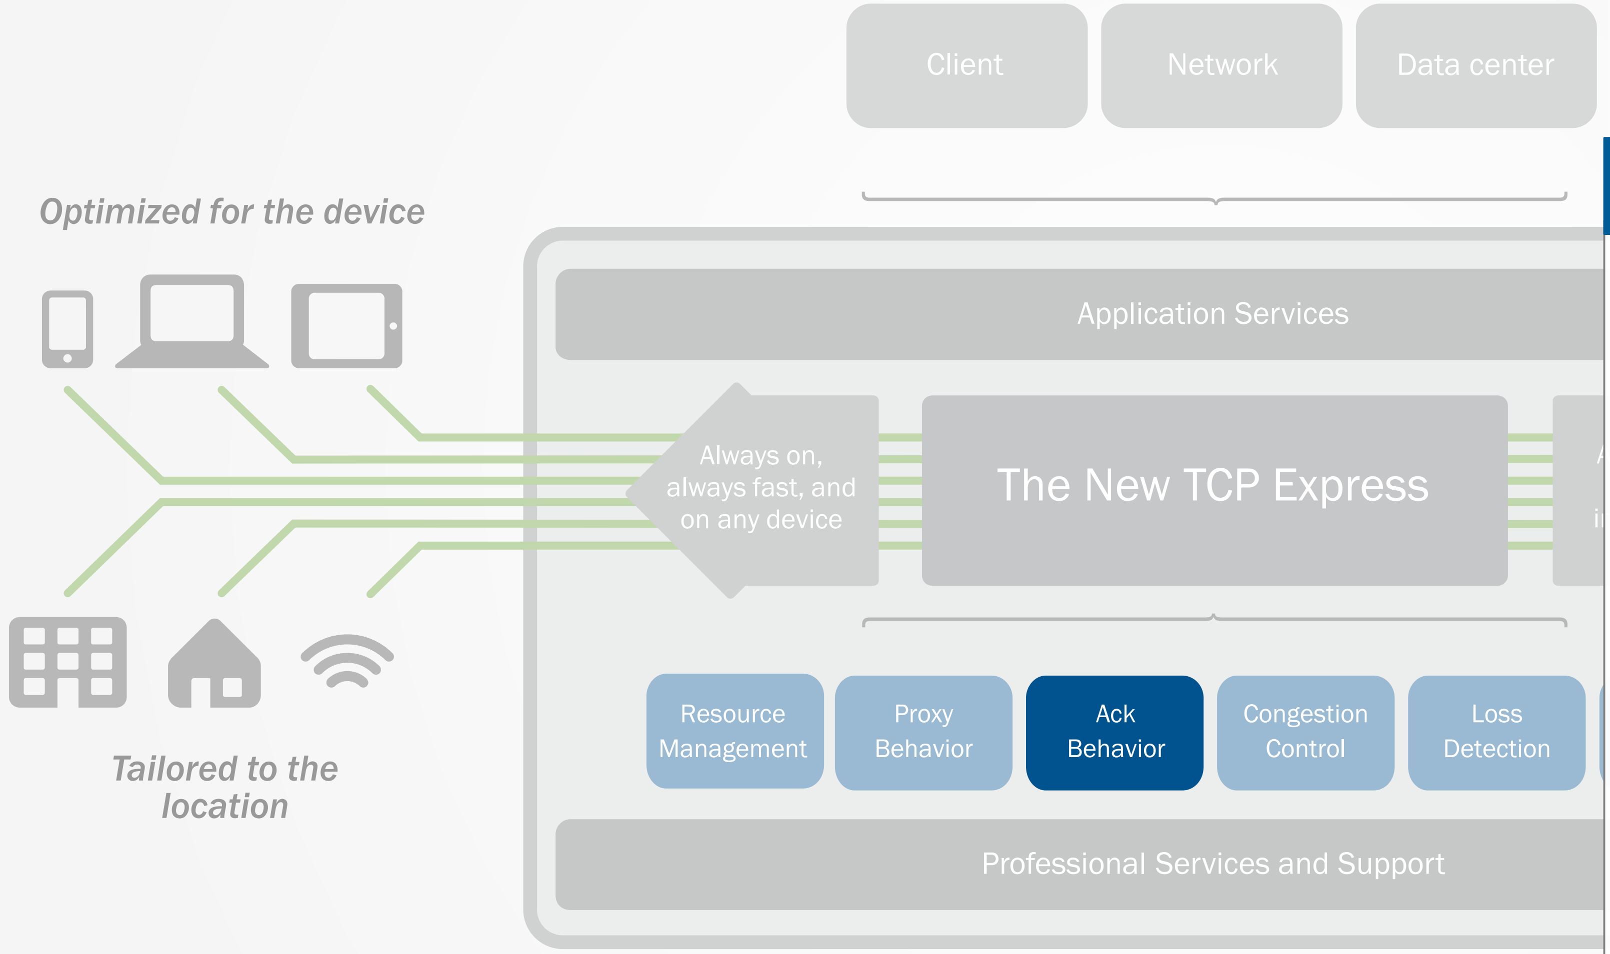

ACKNOWLEDGEMENT

- Delayed ACK
- Selective ACK
- Nagle’s Algorithm

The New TCP Express

CONGESTION CONTROL

- Mobile Optimized Profile
- New Algorithms
  - Woodside
  - Vegas
  - Illinois
  - H-TCP
- Initial Congestion Window Size

*New in 11.5

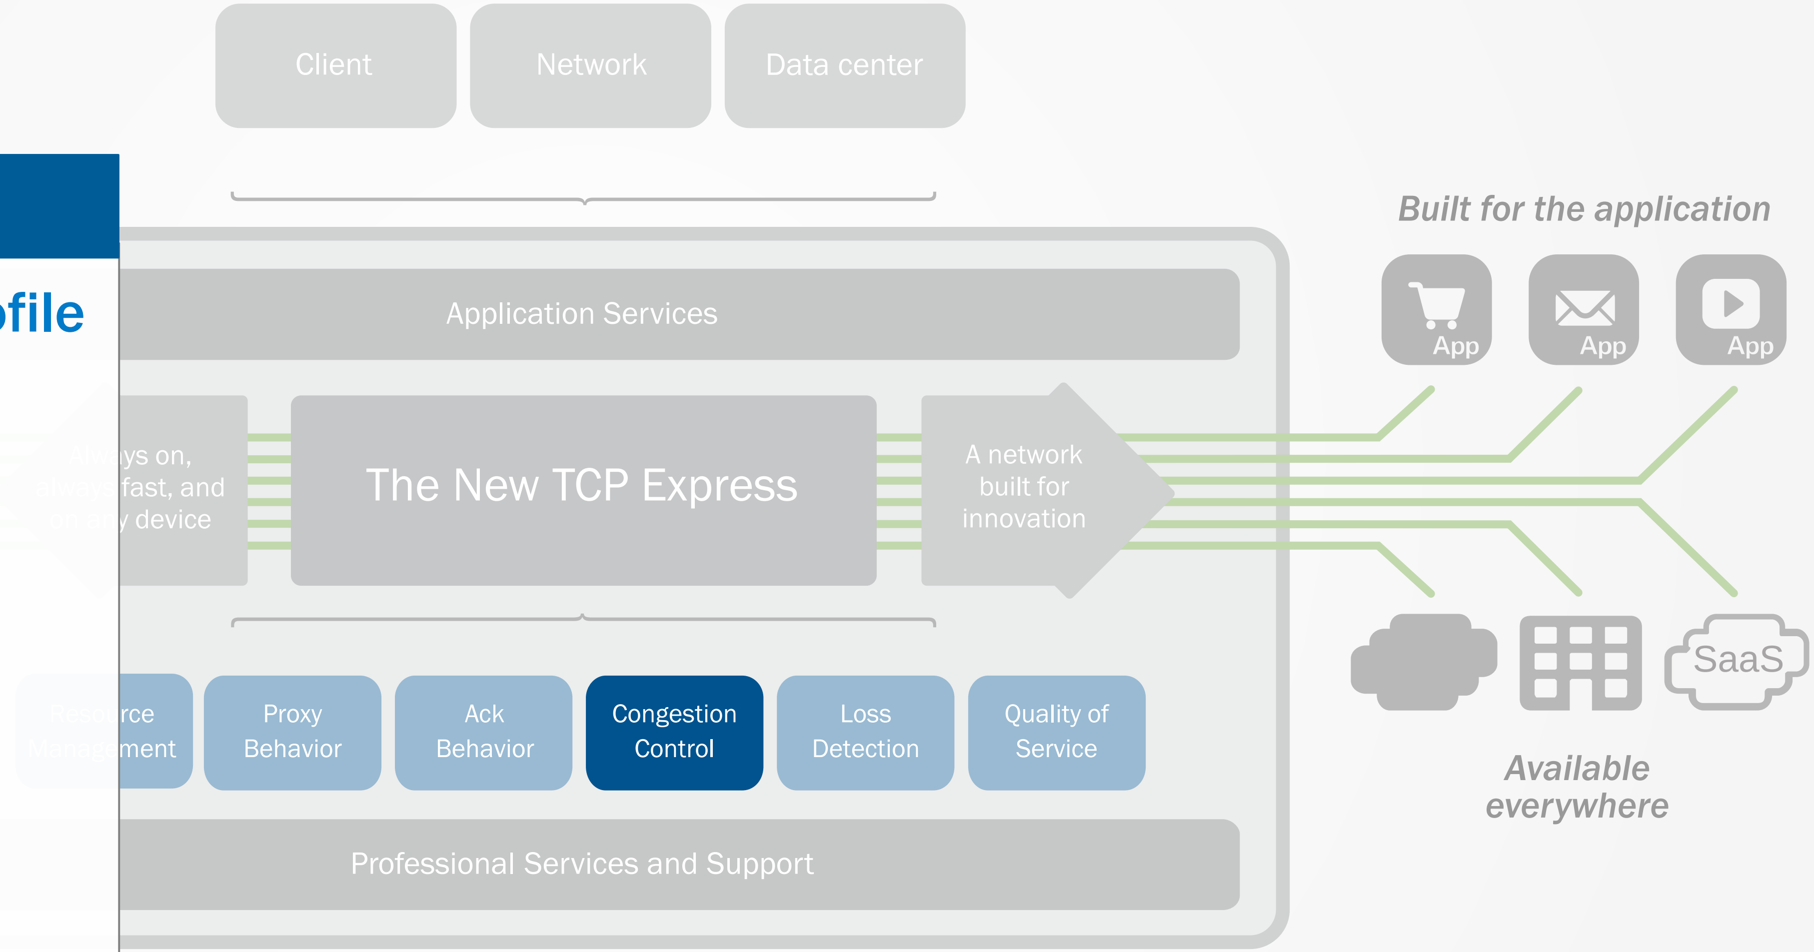

TCP Congestion Control Algorithms

- Loss-based algorithms
  - Reno, New Reno, High-Speed, Scalable, BIC, CUBIC
- Latency-based algorithms
  - Vegas
- Bandwidth-estimating algorithms
  - Westwood, Westwood+
- Hybrid delay/loss algorithms
  - Illinois, Woodside (F5)

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.

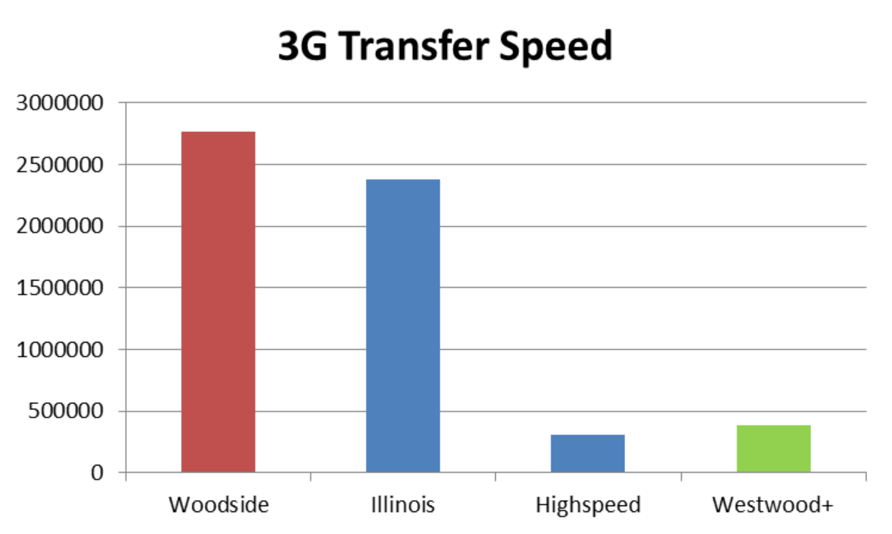

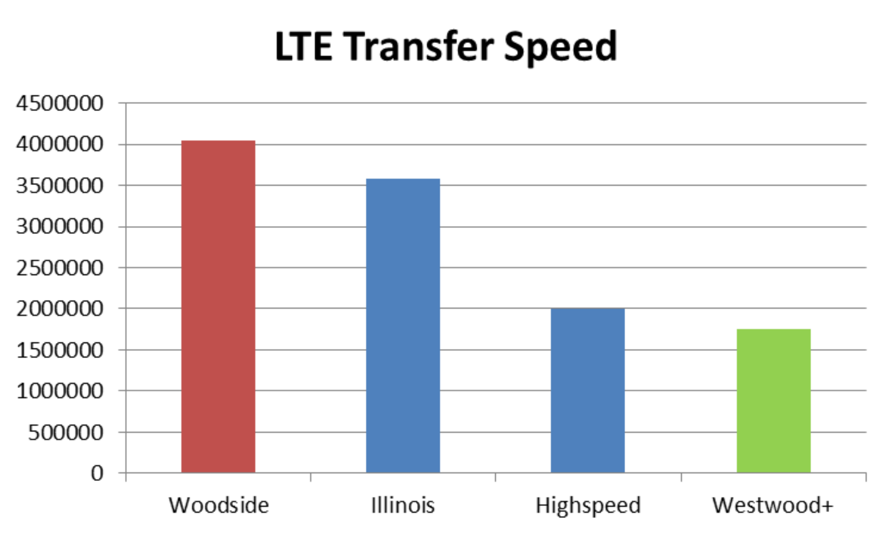

The New TCP Express

LOSS DETECTION

- Rate Pacing
- Forward Error Correction (FEC)
- Retransmission Timeout

*New in 11.5

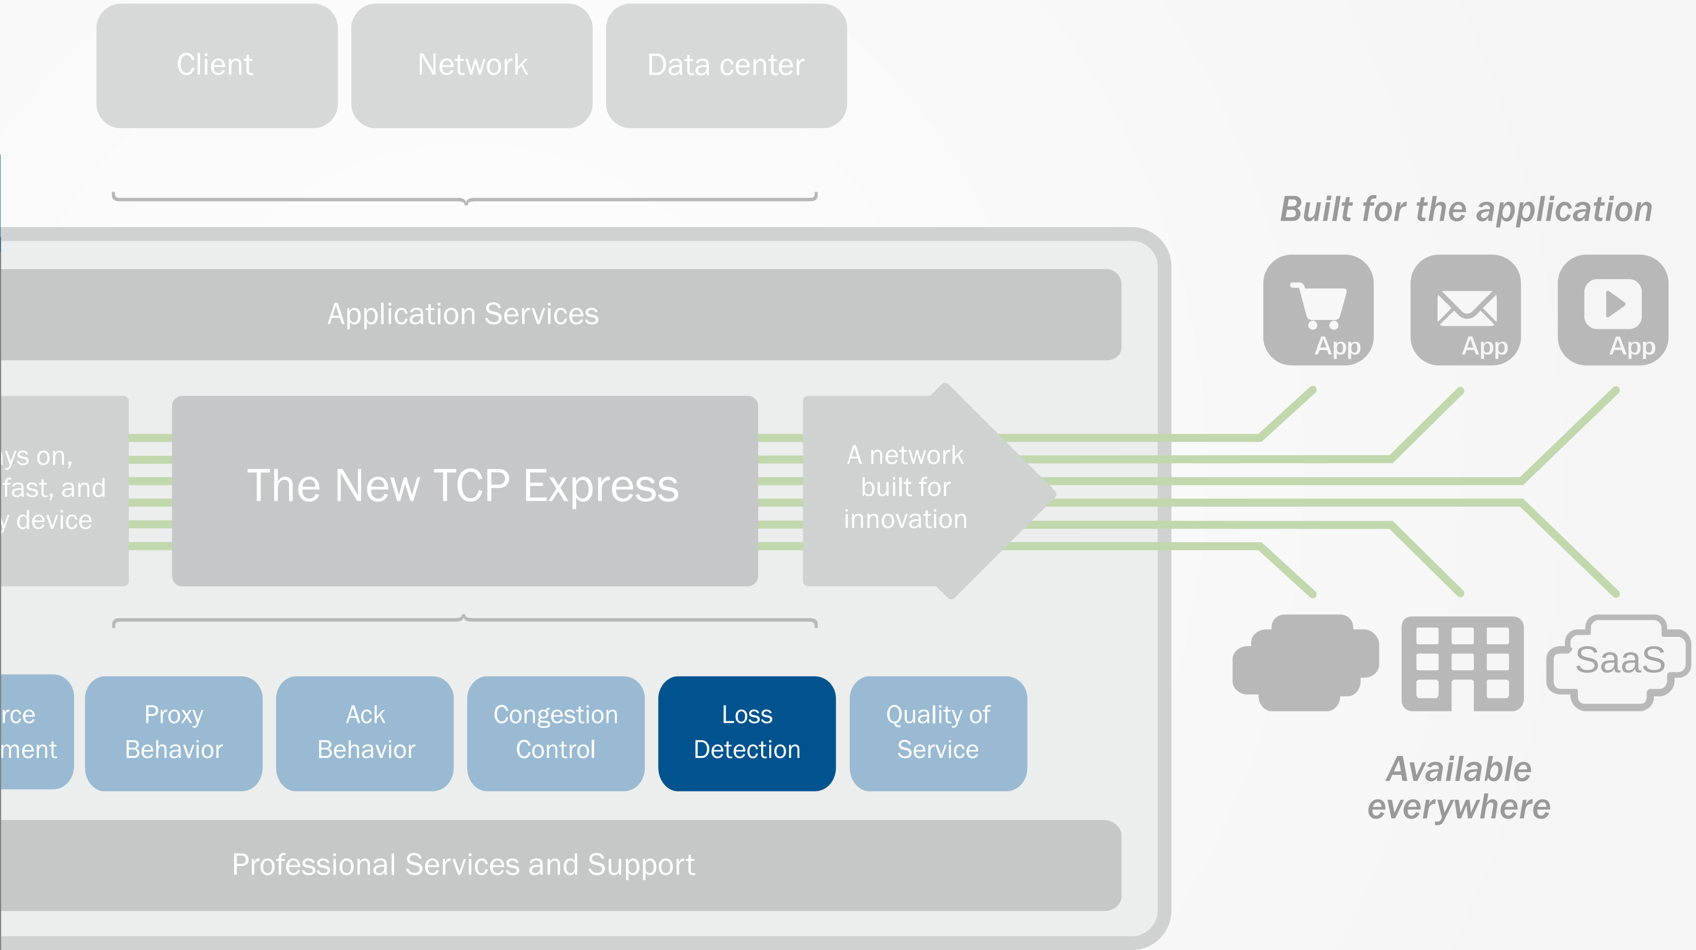

Mobile Optimization – Rate-based TCP

- TCP Express with Rate Pacing
  - Rate Pacing prevents bursts
  - Transmission is paced smoothly by the stack
  - Speed of transmission determined by congestion control
  - Minimal overruns even in high BDP networks
- Benefit
  - Improve the user experience by altering how packets are sent based on feedback received from client.

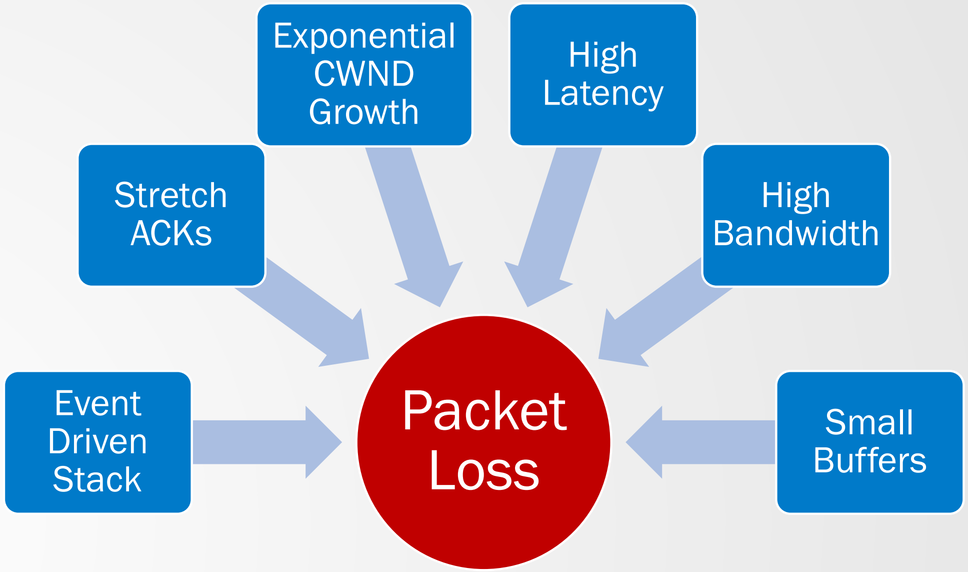

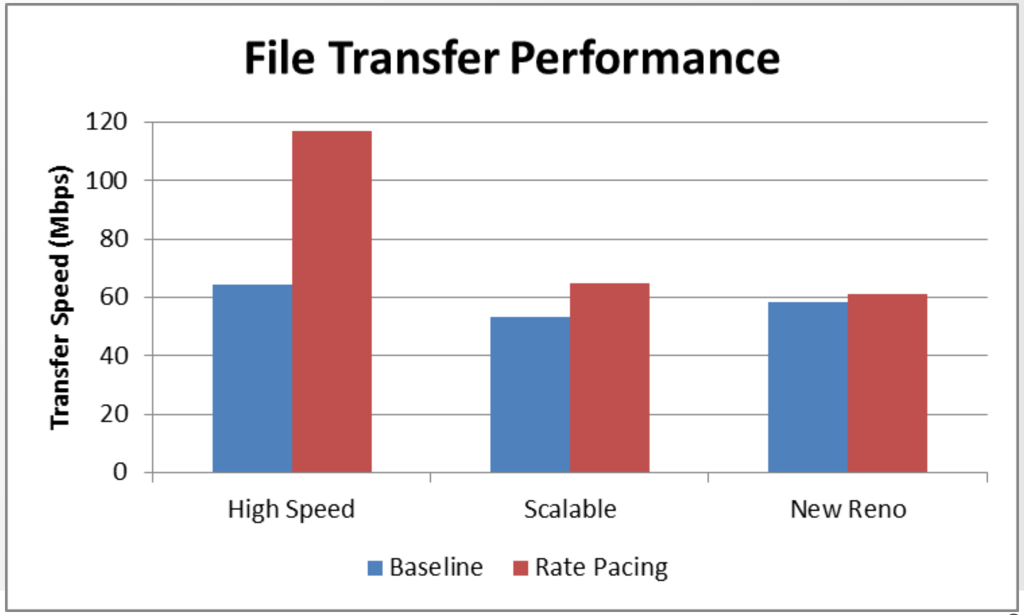

The New TCP Express

QUALITY OF SERVICE

- **TCP-aware bandwidth controller**

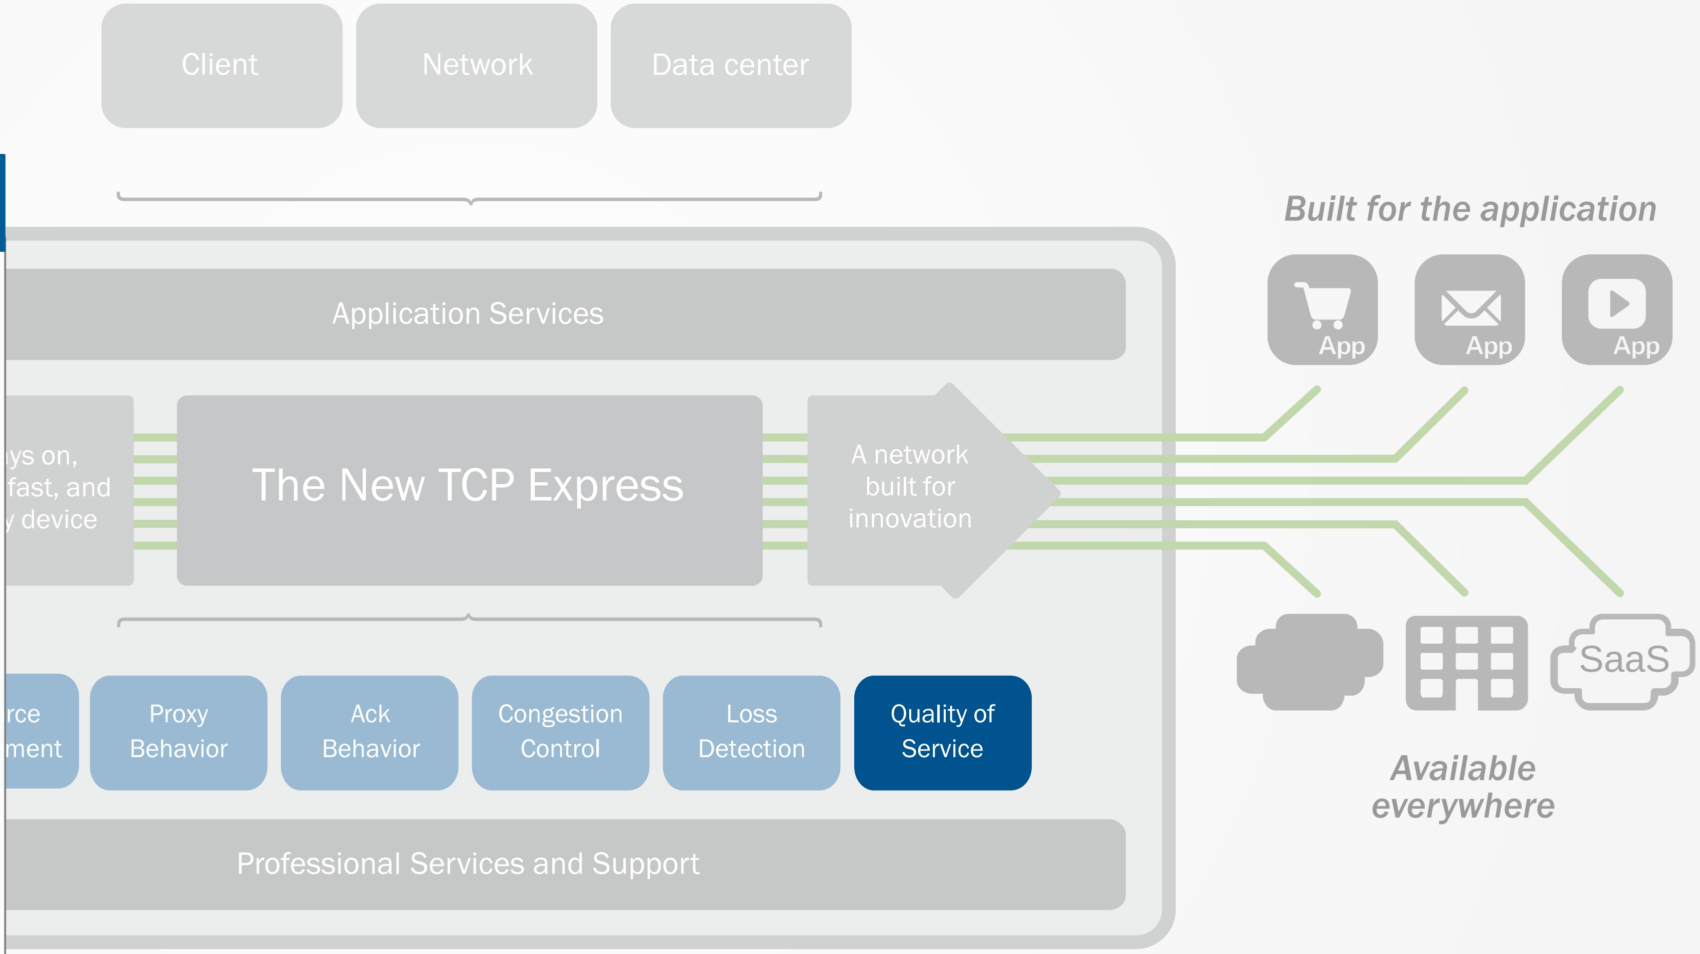

The New TCP Express

QUALITY OF SERVICE

- ToS
- QoS
- MD5 Signature

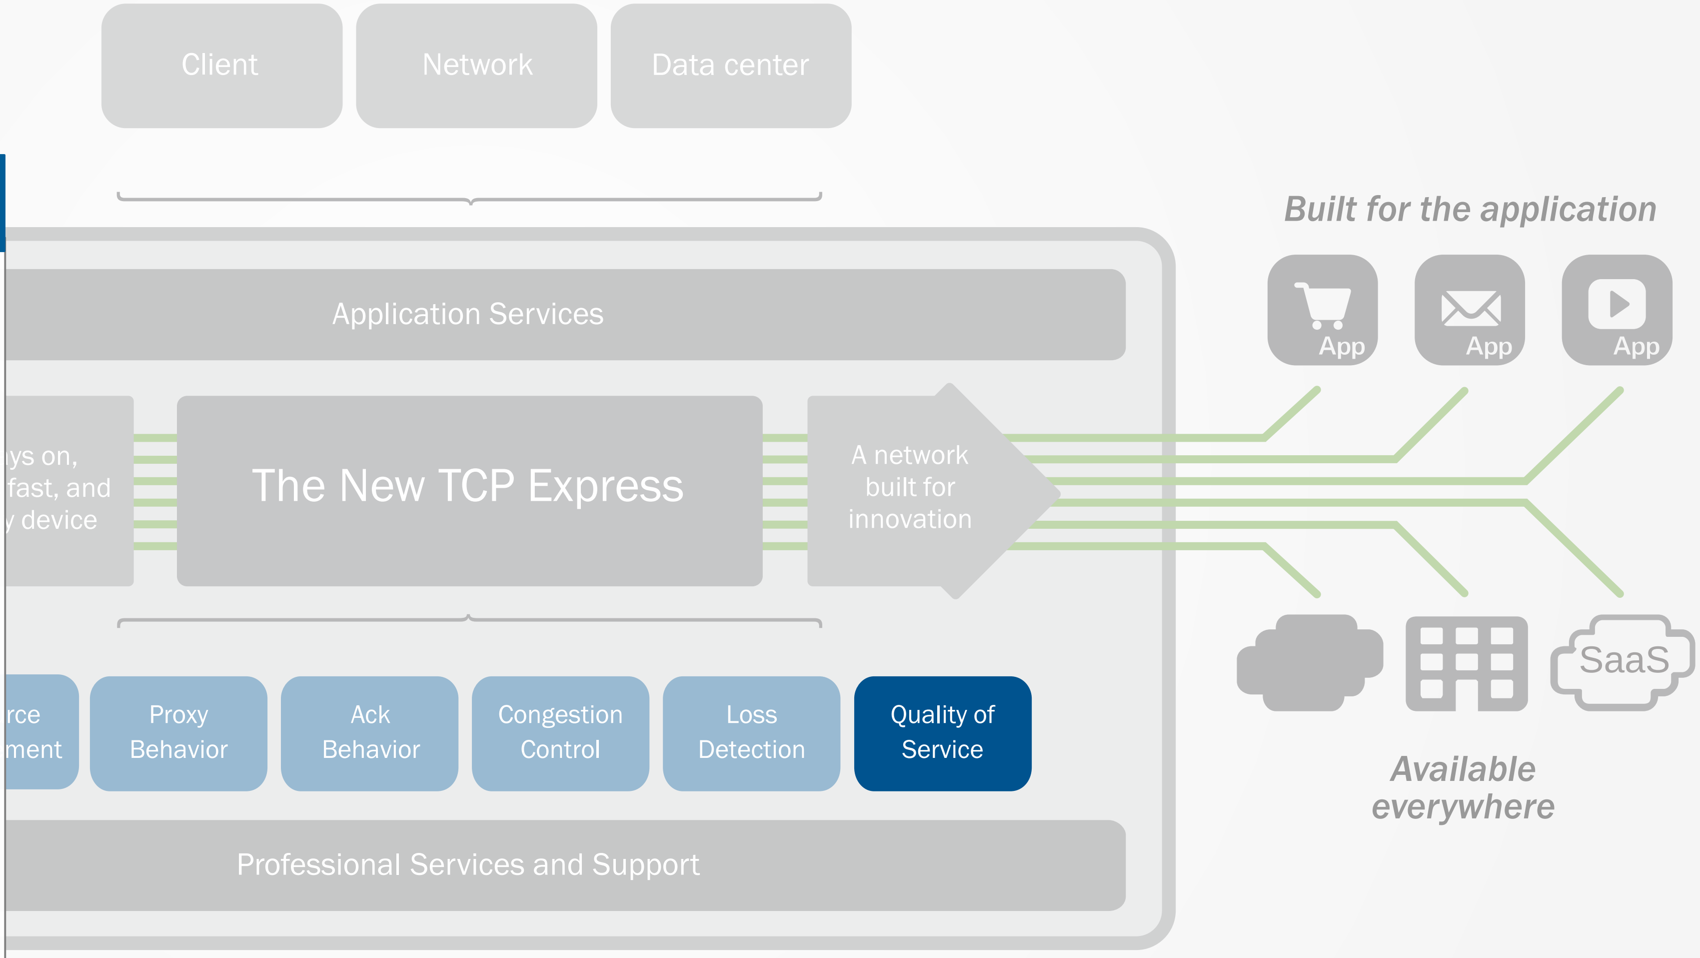

Deployment Models

Traditional Optimization Architecture

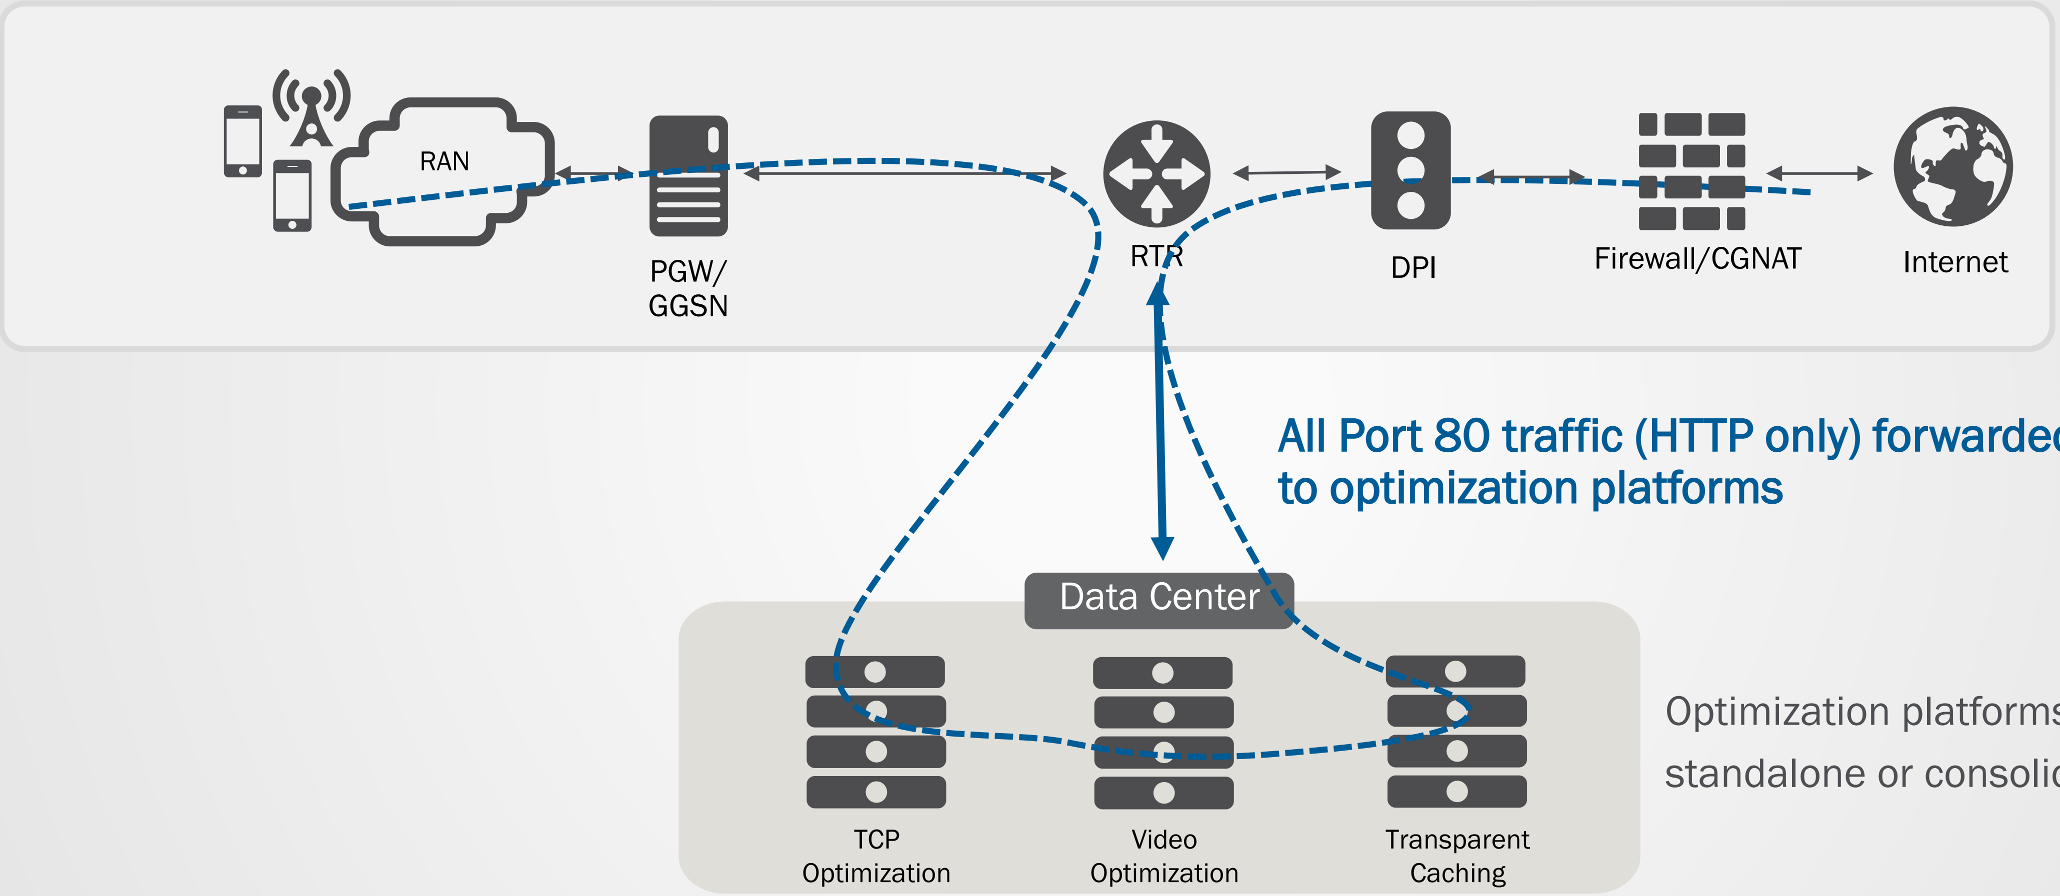

All Port 80 traffic (HTTP only) forwarded to optimization platforms

Optimization platforms can be standalone or consolidated

Next-Generation Optimization Architecture

Inline TCP optimization with intelligent steering consolidated

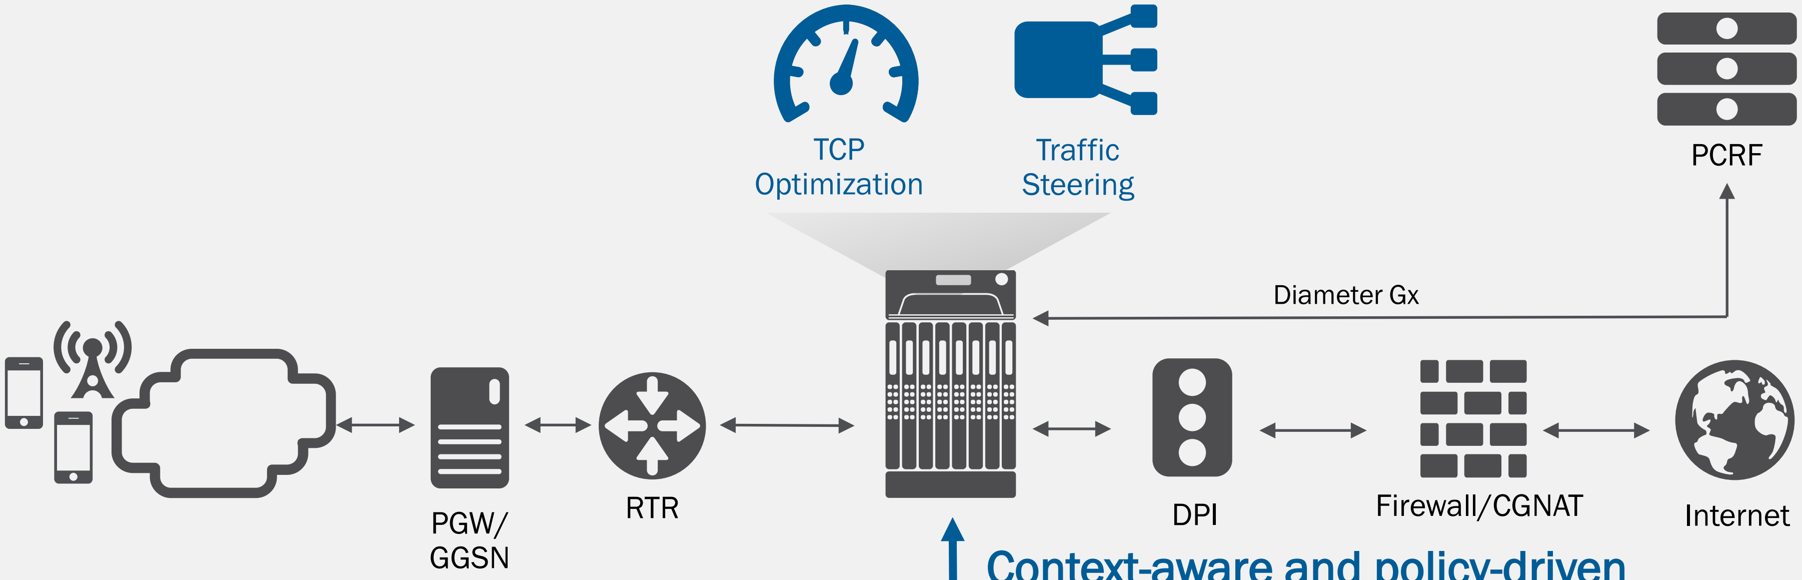

Context-aware and policy-driven traffic steering and service chaining

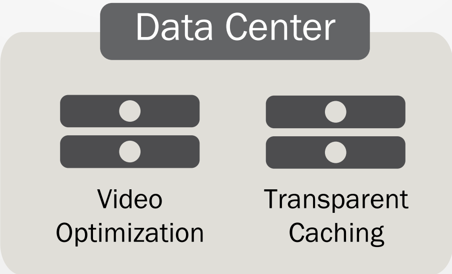

Content optimization platforms

CONTEXT-AWARE STEERING

- Subscriber
- Device-type
- RAT-type
- Content (Video, URI, ...)
- Congestion

Next-Generation Consolidated Gi LAN Architecture

All L4-L7 functionality on a single platform on the Gi LAN

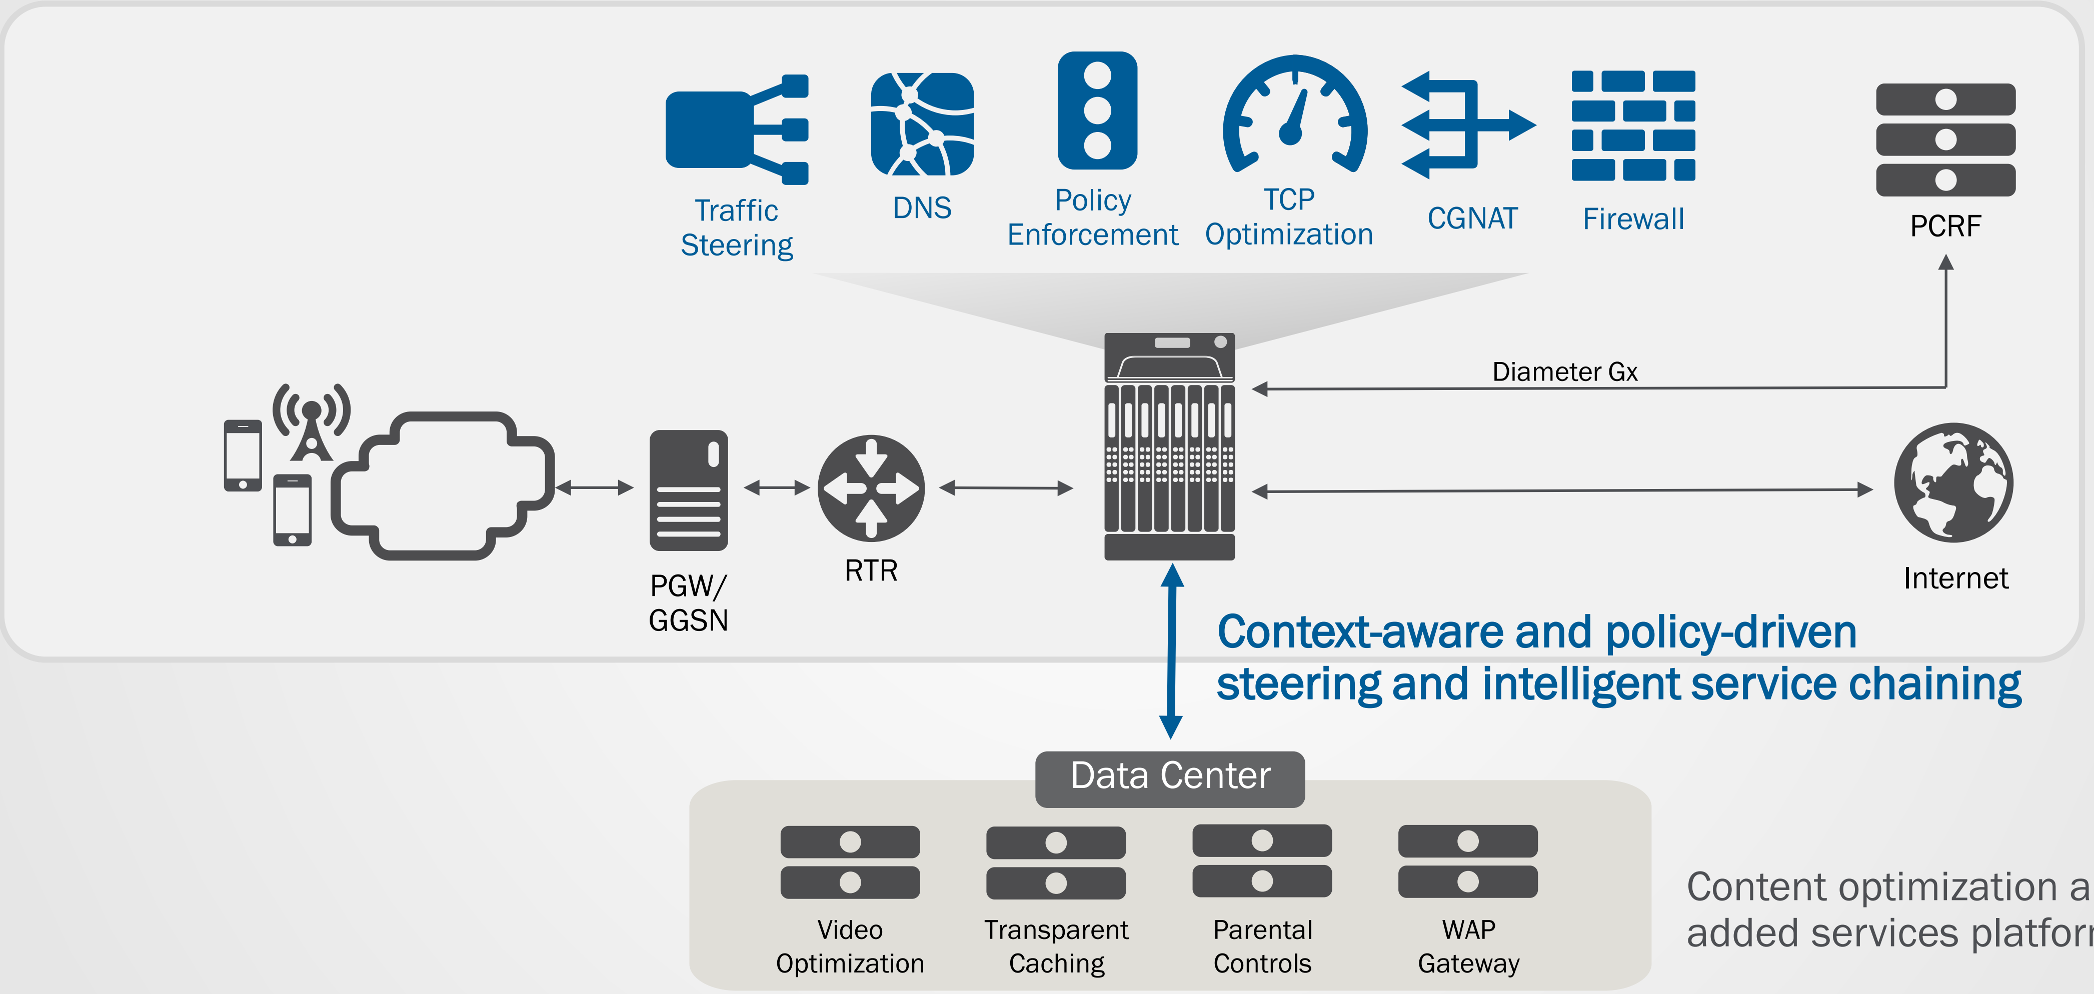

Content optimization and value added services platforms

Solutions for an application world.