

Network Flows

Past, Present and Future

Carter Bullard

QoSient, LLC

carter@qosient.com

FloCon 2014

Charleston, South Carolina

Jan 13-16, 2014

QoSient



Carter Bullard carter@qosient.com

- QoSient - Research and Development Company
 - US DoD, IC, DARPA, DISA
 - Very Large Scale Optimization (Operations, Performance, Security)
 - High Performance Network Security Research
 - DARPA CORONET Optical Security Architecture
 - Telecommunications / Performance Optimization
 - FBI / CALEA Data Wire-Tapping Working Group
- QoS / Security Network Management - Nortel / Bay
- QoS / Security Product Manager – FORE Systems
- CMU/SEI CERT
 - Network Intrusion Research and Analysis
 - Principal Network Security Incident Coordinator
- NFSnet Core Administrator (SURAnet)
- Standards Efforts
 - Editor of ATM Forum Security Signaling Standards, IETF Working Group(s), Internet2 Security WG, NANOG



Introduction



Network Flow

- Network flow concepts are critical to IT infrastructure
 - Complex flow based packet classification is ubiquitous
 - Integrated into almost all IT network hardware
 - Routers, switches, access control, NAT, encryptors, PEP, wireless access points, cable modems, NICs, Hypervisors
 - Humans seem to work well with this abstract notion
 - Provides opportunity for complex network traffic controls supporting access control, resource allocation and forwarding
 - Supports detail network resource utilization reporting
- It wasn't always this way



Earliest Days

- Complex network functions emerged as networks offered more services
- 1960-1980 - Didn't really have "inter-networking" protocols.
 - Days of single hop dial up networks
 - USEnet introduced formal multi-hop transport
 - USEnet News
 - UUCP
 - Email
- Then there was ethernet, DECnet, Appletalk, HIPPI
- And then there was IP



Earliest Days

- 1985-1990 - Georgia Tech had a particular nasty set of problems
 - Very active student community - ISS founder was a Ga Tech student
 - Legion of Doom breaks into Equifax through Bell South and Ga Tech
 - gatech.edu, managing SURAnet south segments of the NSFnet
 - Devastating ARP storms from all the AppleTalk equipment
 - Morris Worm
- GaTech had a very serious and active networking group - Phil Enslow
 - Developing TCP/IP stacks for various systems
 - Earliest VoIP development
- Realtime Network Activity Monitoring System
 - Tracked Bi-directional per TCP and UDP connections
 - Stored records in binary files
 - Primarily used for debugging and development



CMU SEI CERT

- 1991 - Use of Flow Data In Cyber Security
 - Incident Response Management
 - Case Synopsis
 - Incident metadata
 - Data sharing models with FIRST members
 - Packet data management
 - Summarization, annotation, keyword indexing
 - Large repository management
 - Forensics Data Analysis
 - Incident correlation
 - Penetration analysis
 - Common methods correlation - Mitnick
 - Network Vulnerability Research
 - Loose Source and Strict Source Routing Vulnerabilities
 - DDoS Impact Characterization



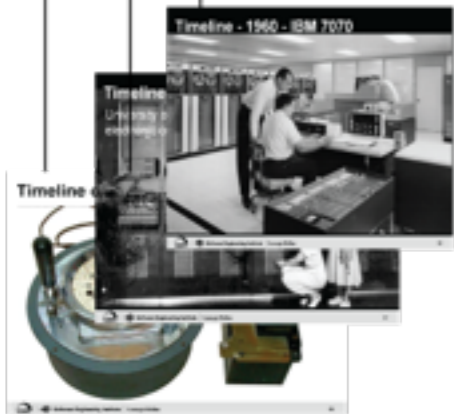
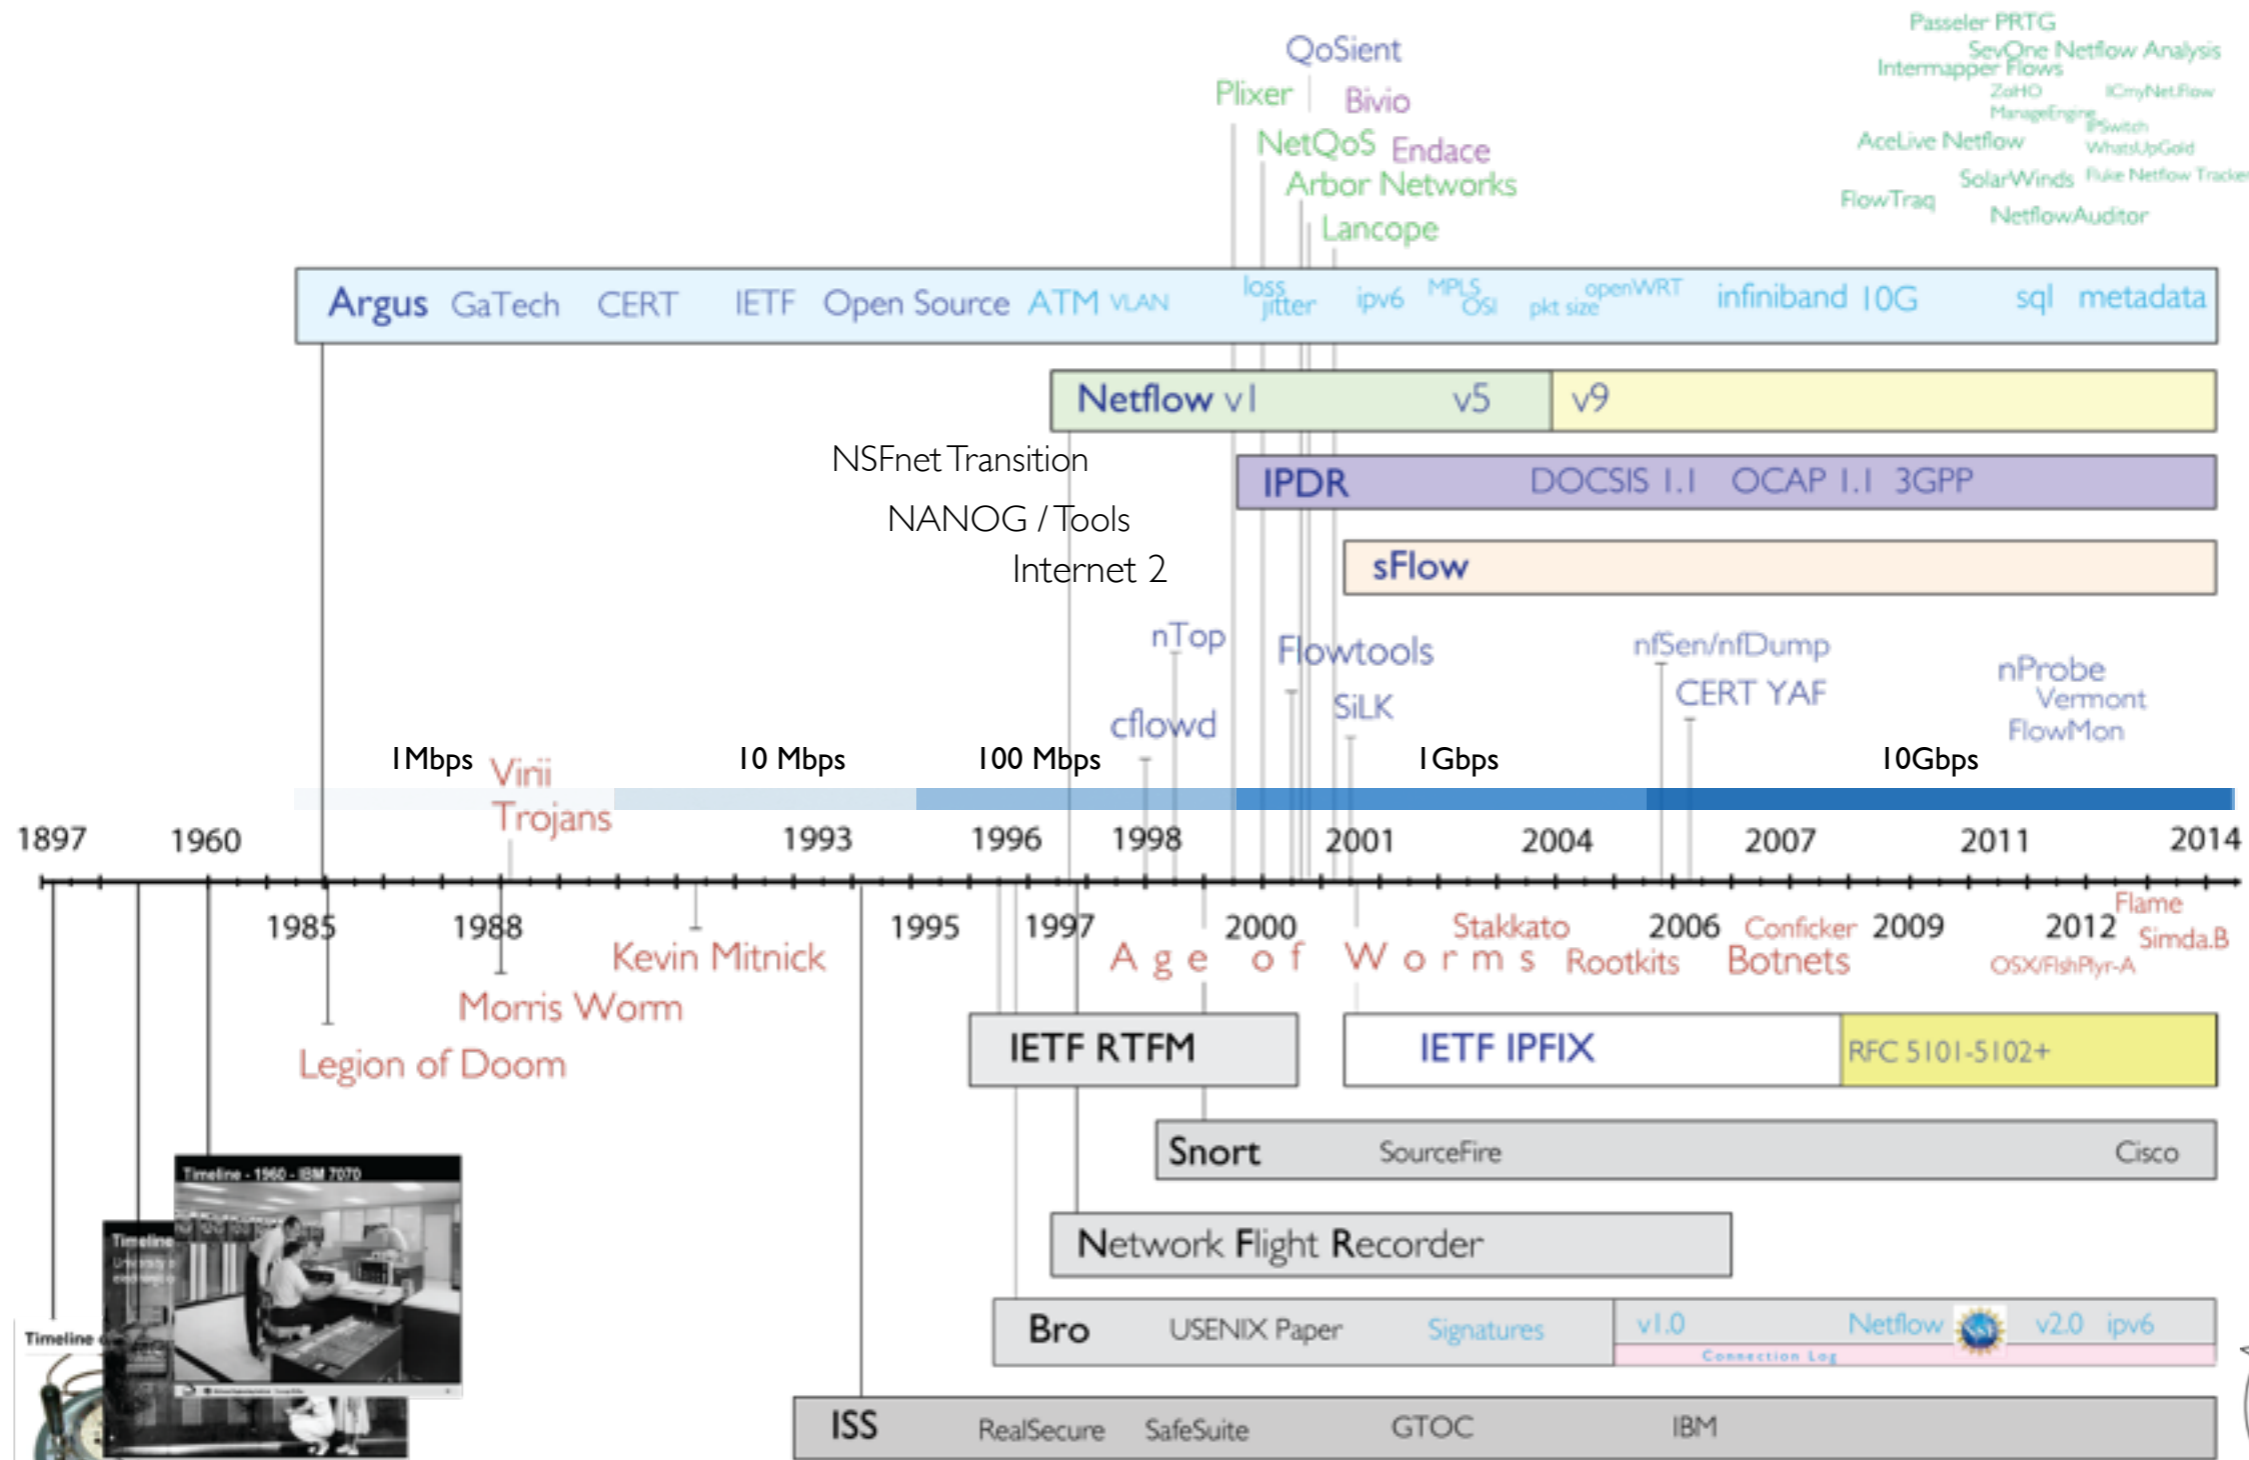
CMU SEI CERT

- CERT Network Flow Outreach
 - FIRST tool suites 1992-1993
 - Internet accountability presentations 1993
 - NANOG / NSFnet transition 1994
 - IETF RMON2 Presentation - 1994
- CERT Advisories 1995-1998
 - Flow Tools for Incident Response
- CMU SEI Flow Monitor Development and Analytics
 - SiLK Netflow v5 and Argus 2.0 - 2003
 - YAF - 2006
 - SiLK drops Argus support - 2008



Network Flow History

Event Timeline



Network Flow Data Adoption

- Adoption was / is painfully slow
- Skeptics in every aspect of the concept of network flow data
 - Performance - Whoa, 1 Mbps !!!! And 100's of flows !!!!
 - Resources Needed - Whoa, toooo much data !!!!!
 - World wide adoption of encryption
 - It doesn't stop any attacks
- Active vs passive monitoring
 - An awful lot of research dollars spent on ping and traceroute
 - Very little insight as to how most national networks are being used
 - US NSFs Internet 2 doesn't collect flow data
 - And they can't tell you who their top 10 talkers are



Network Flow Data Adoption

- Very poor implementations of flow data systems
 - Many expensive systems generated limited results
 - “ Business knowledge “ reports had only limited utility
 - Systems generated reports, but didn't provide data
 - Limited customer customization / extensibility
 - If the 30 reports didn't do it, it didn't get done
 - Some expensive systems generated a lot of really bad data
 - Missing flows, buggy counters, flow records coming out 8 hours later
 - Limited flow attributes i.e. unidirectionality crippled good attempts
 - Very limited actionable outcomes
 - Very bad performance
 - Poor packet processing resulted in waves of statistical flow data
 - Large repositories resulted in poor searching against data
 - Large scale centralized collection techniques made it even worse



Network Flow Data Adoption

- But, flow is the best data for cyber security
 - Enables network activity audit
 - Specified by DoD in NCSC-TG-005
 - The Red Book - Trusted Network Interpretation of the Trusted Computer System Evaluation Criteria (1987)
 - Fundamental theoretical threat countermeasure
 - Can create real deterrence in formal systems
 - Historical intrusion / penetration analysis, long term threat assessments, intrusion impact analysis
 - Complex incident investigations, behavioral baselining, anomaly detection, etc....



Theoretical Security Threats and Countermeasures

Countermeasures		Threat				
		Unauthorized			Degradation of Service	Repudiation
		Use	Modification	Disclosure		
Authentication	Cryptographic	x		x		
Integrity			x			
Confidentiality					x	
Access Control		x	x	x	x	
Non-Repudiation (audit)		x	x	x	x	x

Derived from ITU-T Recommendation X.805
Security Architecture for Systems Providing End-to-End Communications

 Primary Security Countermeasure
 Secondary Security Countermeasure



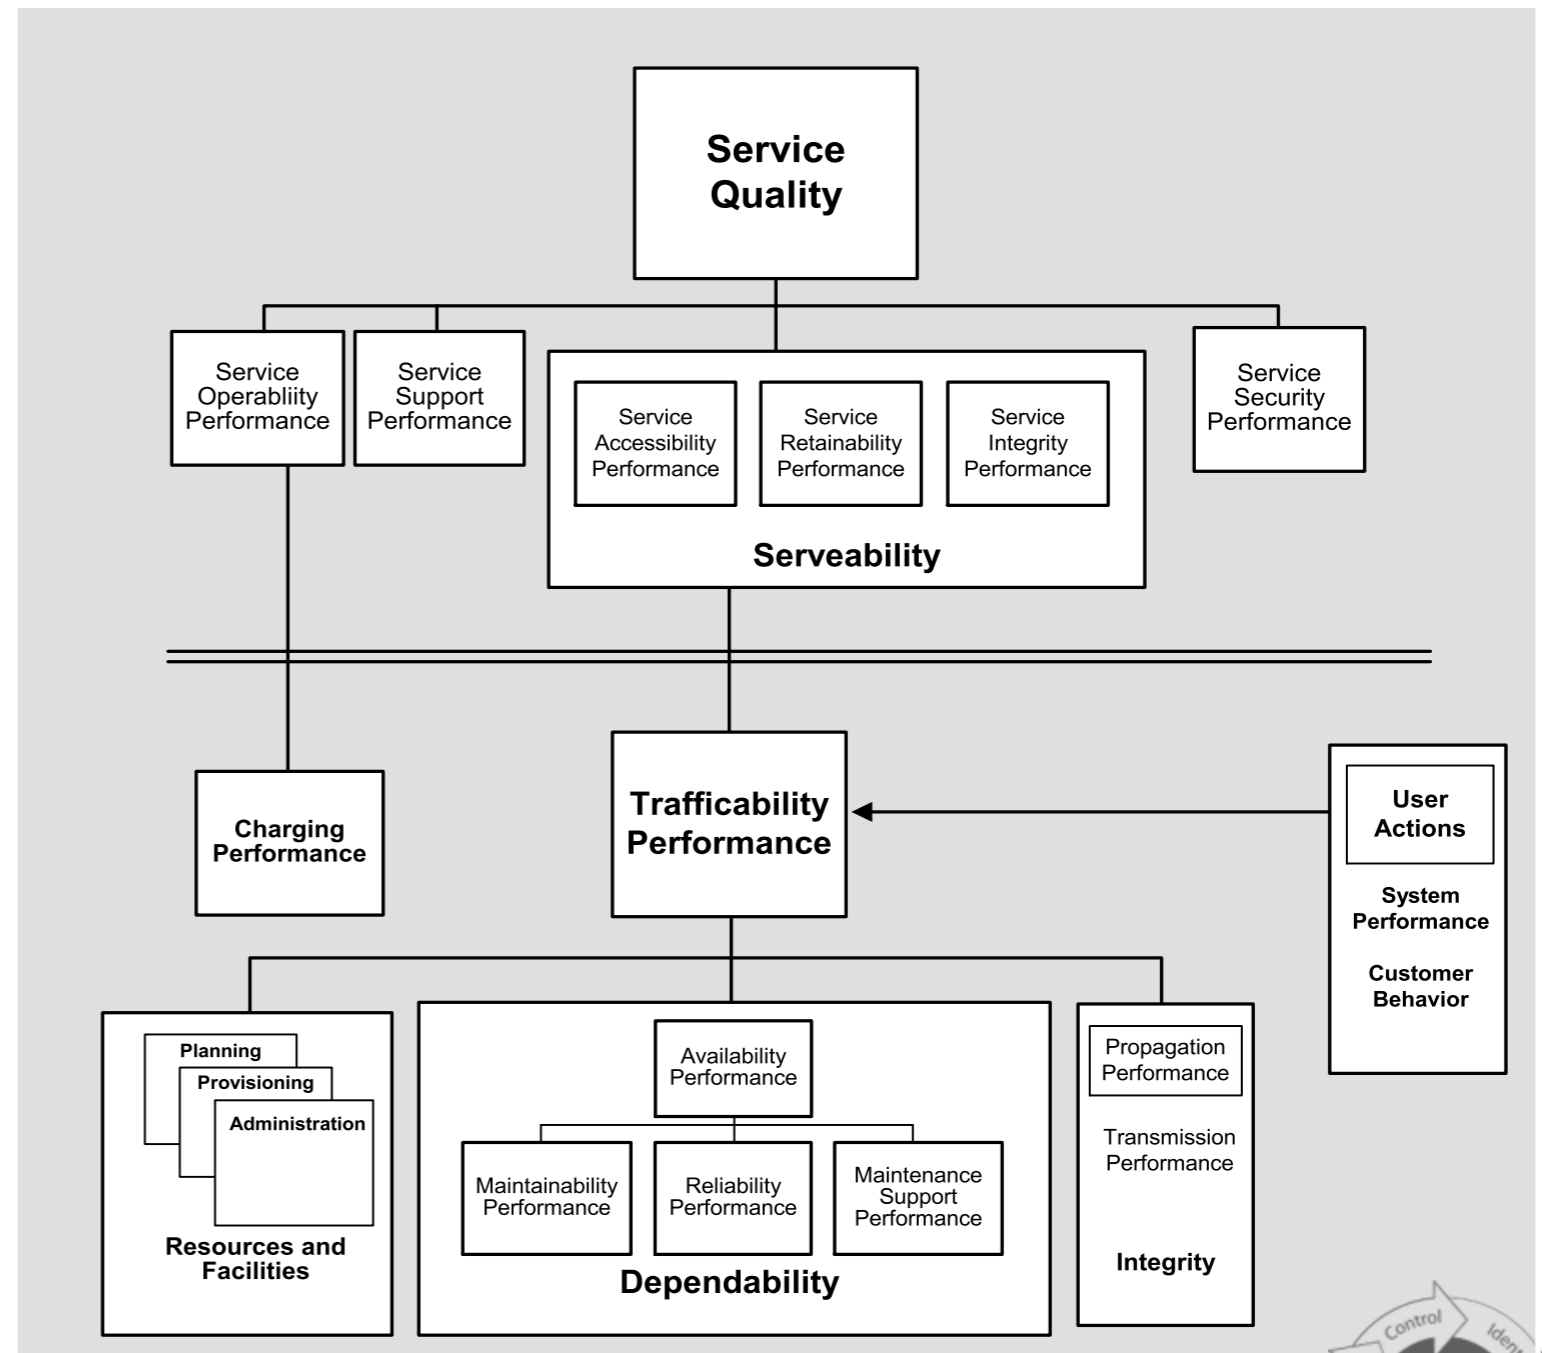
Network Flow Data Adoption

- End-to-end performance optimization
 - Bi-directional flow data with performance metrics
 - Connectivity, Availability, Rate, Load, Loss
 - Host and Network Demand, PktSize, Pkt Arrival, Jitter
 - Control plane performance; ARP, DHCP, DNS, Routing
 - Network fault identification, ICMP Tracking
- Complex operations problems
 - Differentiate between network and end system error
- Assist in technology development
 - 40-100 Gbps transport modeling
 - E2E QoS utilization modeling (ATM, VoIP, 4K Video)



Telephony CDR Utility

- Billing
- Traffic Engineering
- Quality Assurance
- Network Management
- Maintenance
- Marketing
- Product Development
- Security
 - Fraud Detection
 - Forensics Analysis
 - Incident Response
 - Non-Repudiation / Audit



From ITU-T Recommendation E.800 Quality of Service, Network Management and Traffic Engineering



Now

QoSient
ARGUS



Who's Using Network Flow?

- Educational Sites (10,000's of sites world-wide)
 - Carnegie Mellon University, Stanford University, Purdue, University of Chicago, Columbia University, University of Ga, American University, etc.....
 - Enterprise wide near realtime network security auditing
 - Distributed Security Monitoring
 - Network forensics security research
- U.S. Government / Nation States
 - DISA has added flow data generation to procurement criteria
 - DHS, DoD, i.e. Naval Research Laboratory
 - Developing their own security technology
 - Embedded in a lot of networking and security equipment
 - Many (inter)national networks collect flow based usage data for many purposes.
- ISPs, Network Service Providers , Enterprises, Corporations, Individuals
- Technology Developers



IP Network Flow Information

- All types contain IP addresses, network service identifiers, starting time, duration and some usage metrics, such as number of bytes transmitted.
- More advanced types are transactional, convey network status and treatment information, service identification, performance data, geo-spatial and net-spatial information, control plane information, and extended service content.
- IPDR - Billing and Usage Accountability
 - ATIS, ANSI, CableLabs, SCTE, 3GPP, Java CP, ITU/NGN
 - File and stream formats (XML).
- Netflow, JFlow, Sflow, Qflow, Rflow, cFlowd, NetStream
 - Integrated and standalone vendor flow information
- Argus - Ops, Performance and Security Management
 - L2/L3/L4/L5 control and data plane network auditing
 - Archive, file, stream formats. (Binary, SQL, CSV, XML)
- YAF/SiLK - Cyber Security Forensics
 - IETF IPFIX stream formats. Binary file format.



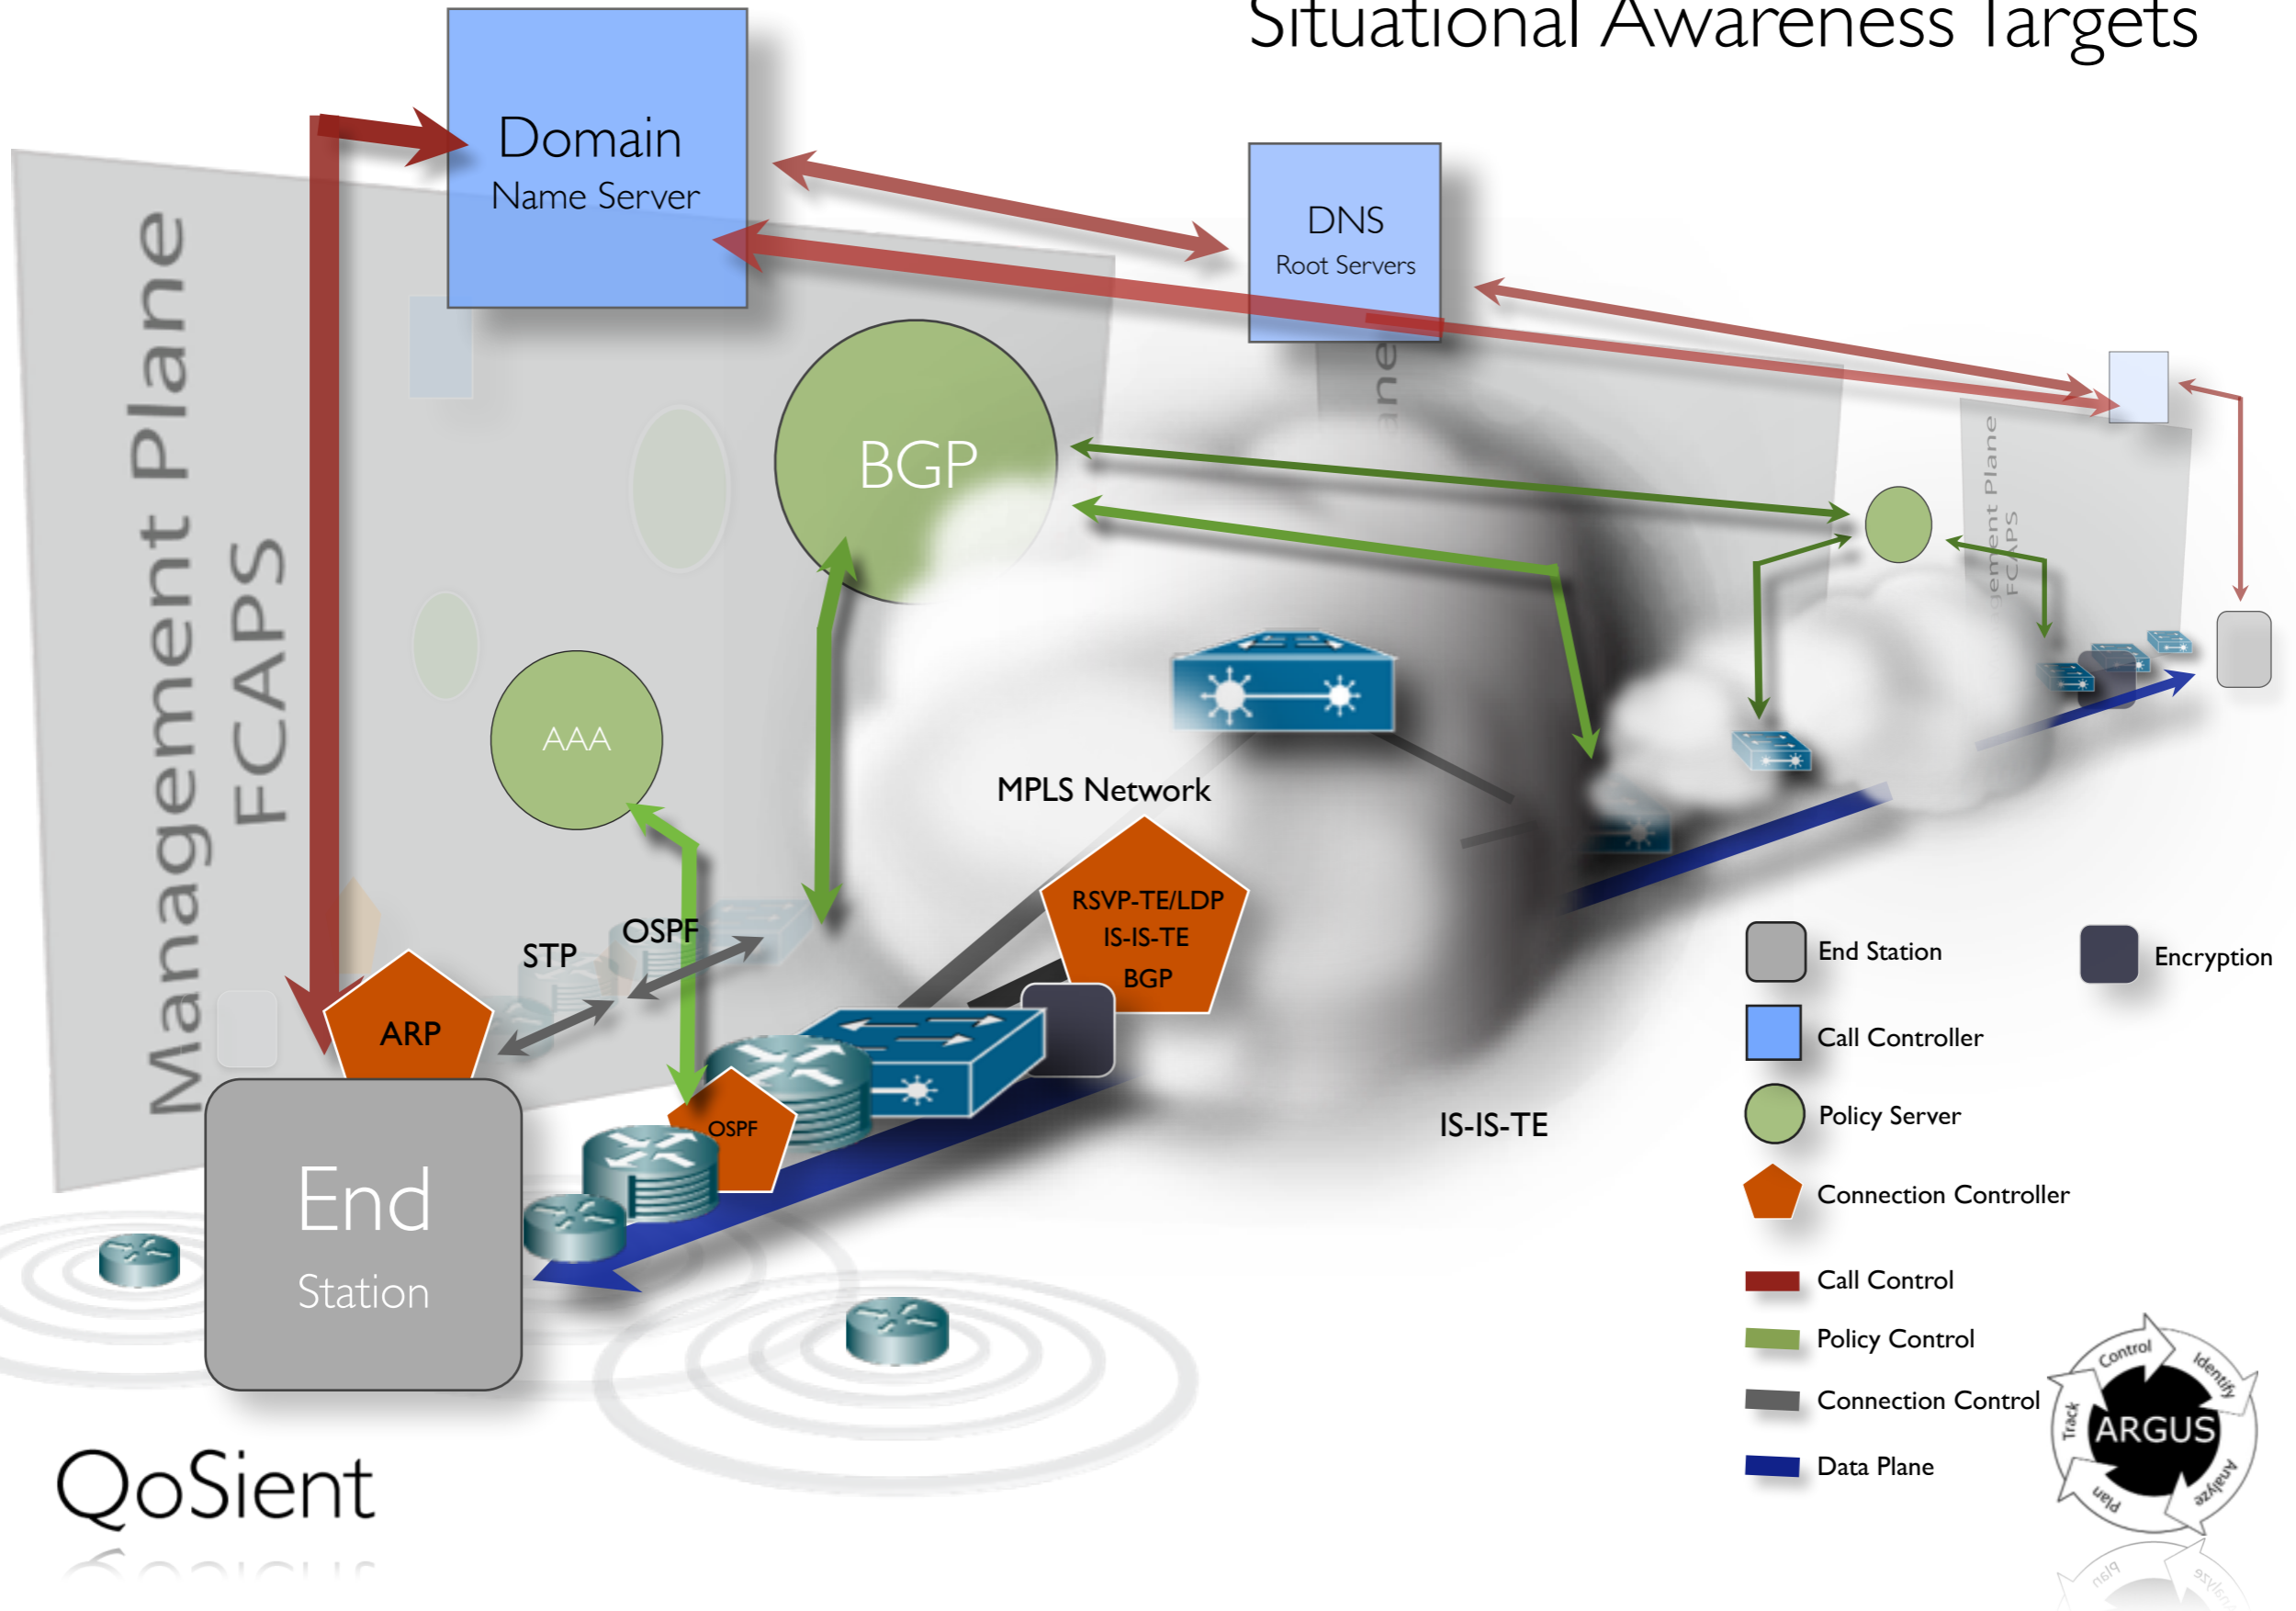
IPDR Adoption Status by Standard Bodies and Industry Forums

- ATIS:
 - Approved the Requirements and Architecture – ATIS-0300075
 - Referred by the IIF and IPTV OSS architectures
- ANSI adopted IPDR/SP as an ANS:
 - ATIS-0300075.1
 - ANSI/SCTE 23-3 2005 DOCSIS 1.1 Part 3: Operations Support System Interface
- CableLabs adopted IPDR/SP as a mandatory part of:
 - DOCSIS® 2.0 and 3.0
 - OpenCable™ OCP 1.1
 - Receiver Metrics Gathering Specification (OC-SP-Metrics-I01-061229)
- SCTE - SCTE 135-4 2007 DOCSIS 3.0 Part 4: OSS I
- 3GPP adopted IPDR/SP as optional implementation of the Bx
- JCP adopted IPDR/SP as part of JSR#190
- Harmonization at ITU as part of the NGN



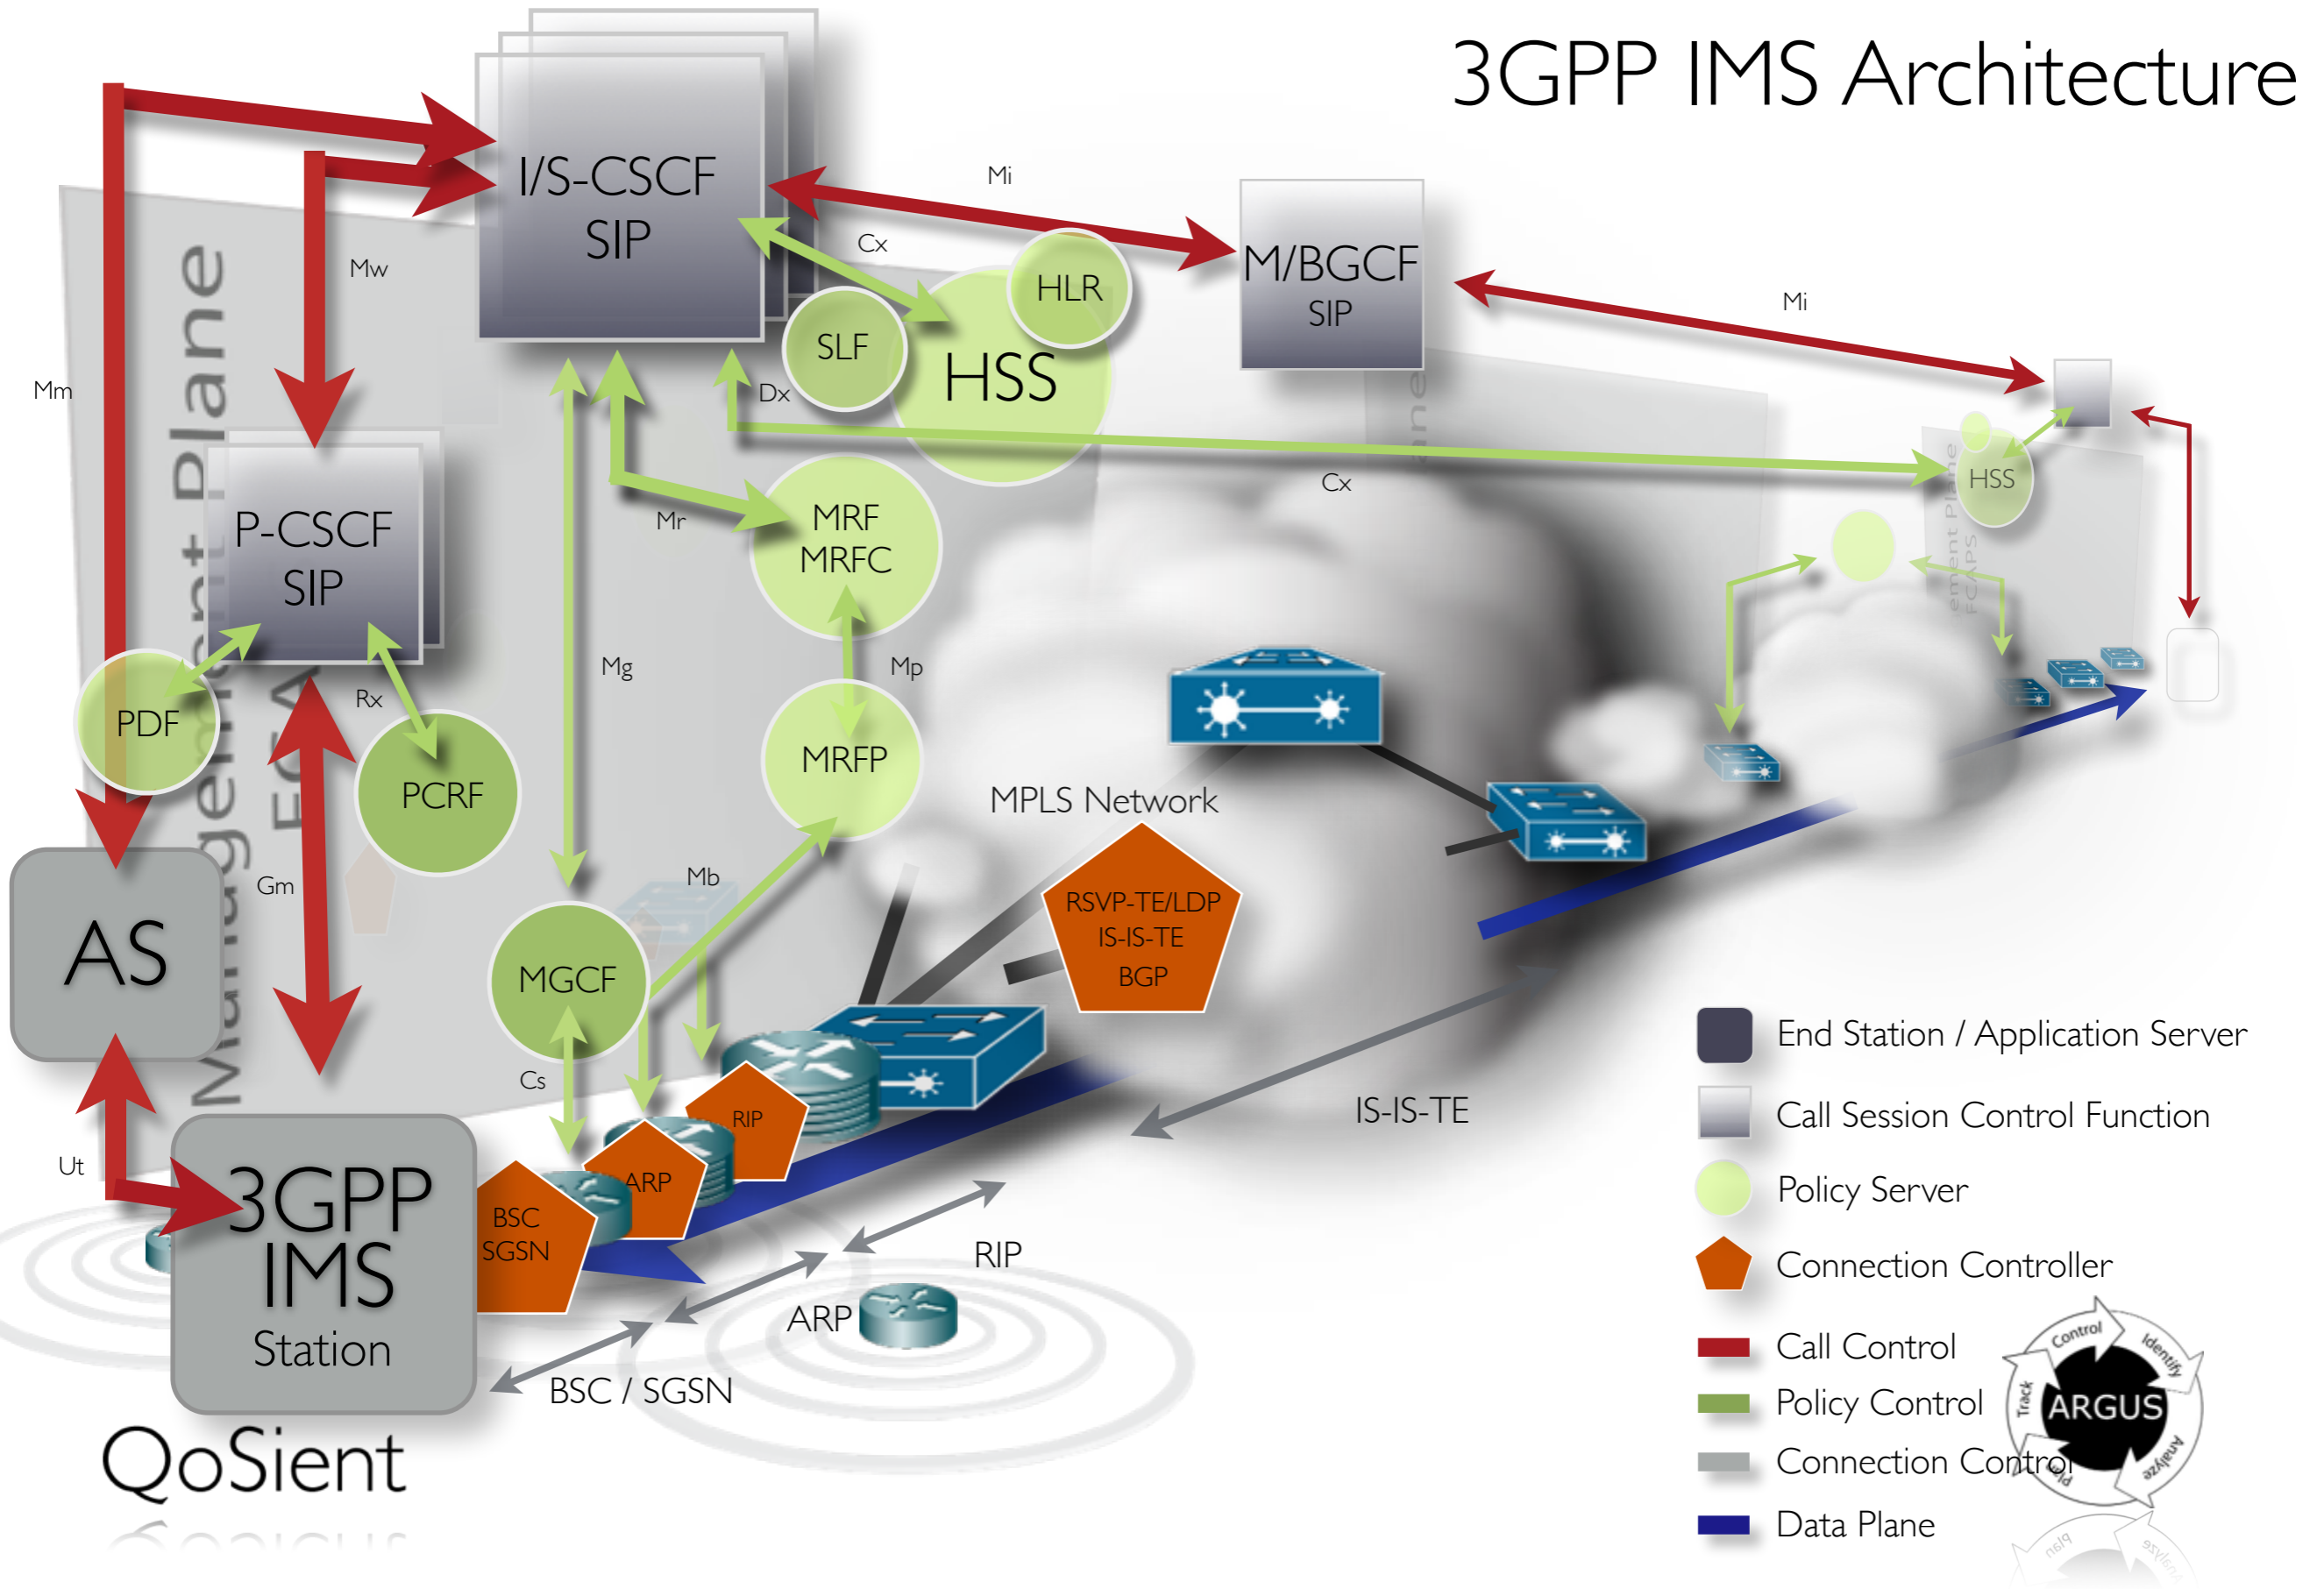
Standard Internet Architecture

Situational Awareness Targets



Mobile User Data Networks

3GPP IMS Architecture



QoSient

Present

- Network Flows have changed a bit
 - Summarization of network activity
 - Multiple global and local identifiers
 - Multi layer network service (what)
 - Time (when)
 - Distributed observation domains (where)
 - Multiple descriptions of activity
 - Protocol and behavioral conformance
 - Network engineering information
 - Metadata Support



Future

QoSient
ARGUS



This Future

- Advances in Technology
 - Ultra Performance Switching
 - Very High Performance Messaging - ZeroMQ
 - Near Field Communications
- Changes in Technological Use
 - Distributed Processing, Clouds, SDN, NFV
 - White Box Infrastructure
 - Much Much Much Larger Data



This Future

- Evolving Needs in Technology
 - Cost and Complexity Reduction
 - Minimize network components
 - Pushing it all to the edge
 - Dynamic Adaptive Computing - or Not !!!
 - Data Sharing
- Getting Around to the Real Problems
 - The Insider Threat
 - Establishing Deterrence



Advances in Technology

Network Virtualization

Implementing an OpenFlow Switch on the NetFPGA platform

Jad Naous
Stanford University
California, USA
jnaous@stanford.edu

David Erickson
Stanford University
California, USA
derickso@stanford.edu

G. Adam Covington
Stanford University
California, USA
gcoving@stanford.edu

Guido Appenzeller
Stanford University
California, USA
appenz@cs.stanford.edu

Nick McKeown
Stanford University
California, USA
nickm@stanford.edu

ABSTRACT

We describe the implementation of an OpenFlow Switch on the NetFPGA platform. OpenFlow is a way to deploy experimental or new protocols in networks that carry production traffic. An OpenFlow network consists of simple flow-based switches in the datapath, with a remote controller to manage several switches. In practice, OpenFlow is most often added as a feature to an existing Ethernet switch, IPv4 router or wireless access point. An OpenFlow-enabled device has an internal flow-table and a standardized interface to add and remove flow entries remotely.

Our implementation of OpenFlow on the NetFPGA is one of several reference implementations we have implemented on different platforms. Our simple OpenFlow implementa-

Keywords

Computer networks, Flow switching, NetFPGA, OpenFlow, Packet switching, Programmable networks

1. INTRODUCTION

Today it has become extremely difficult to innovate in the computer networks that we use everyday in our schools, businesses, and homes. Current implementations of mainstream network devices, such as Ethernet switches and IP routers, are typically closed platforms that can not be easily modified or extended. The processing and routing of packets is restricted to the functionality supported by the vendor. And even if it were possible to reprogram net-



Advances in Technology

Network Virtualization

- SDN - Software Defined Networks
 - Open Network Foundation (ONF), OpenFlow, OpenDaylight, VxLan, NVGRE
 - Cisco ACI fabric, VMware NSX, Microsoft HyperV, Juniper Contrail
 - Use of SDNs to provide security
 - Opportunity to build-in strong attachment authentication and authorization
 - Employ ATM security strategies to insert security mechanisms into traffic path
 - Security Challenges of SDNs
 - Dynamism is the antithesis of security
 - New dials for the bad guys to turn
 - SDN developers see NAT as a good thing
- NFV - Network Functional Virtualization
 - Virtualize components of managed services
 - Authentication, Authorization, Provisioning, Billing, Resource Allocation



Advances in Technology

SDNs and Network Flow Data

- Primary interest in SDNs today is advanced monitoring
- SDN capabilities driven by merchant silicon
 - Hardware support for flow based usage monitoring
 - Netflow v9 style metrics, complex flow object identifiers with wildcarding
 - Fixed metrics strategy. Packet and byte accumulations, with limited flow state - unidirectional
 - Results in complex multi-dimensional aggregated flow record types
 - Chips include packet classification (DPI) and filtering logic
 - Lot of horsepower to do a lot of interesting flow data operations.
 - Much of the cycles are needed for SDN function, not much left for additional DPI functions.
- Most SDN switches support 48 10Gbps ports and some 40Gbps
- Per port flow cache limitations will push for 'wildcard' flow entries
 - Compel network designers to support aggregated rule sets, and overlay networks
 - Most detailed data plane monitoring will be on demand



Advances in Technology

SDNs and Network Flow Data

- SDN networks are flow oriented networks
- Flow data systems are/were used in the development of SDNs
 - Argus used in Stanford OpenFlow development
 - Used for flow modeling to understand flow demand of operational networks
 - Flexible flow models for complex traffic
 - Provide flow per second demands
 - Flow duration behaviors
- SDNs add flow oriented control to the network operations and optimization paradigm
 - Network flow data now supports all phases of the optimization cycle.
 - Identification, Analysis, Planning, Tracking and Control
 - This enables feedback-directed optimization for SDNs, using flow data.
 - Requires a common flow data model and usable optimization metrics.



Advances in Technology

SDNs and Network Flow Data

- Flow data systems are needed to monitor SDN health
 - Build out operational verification and validation
 - Verify end-to-end network function in the presence of dynamic change
 - Continuous reachability debugging and troubleshooting - Loop detection
 - Verify that change is the intended result of the SDN control plane
- SDN data plane complexity will dramatically increase
 - Multi-dimensional overlay network strategies
 - Supporting multiple stacks simultaneously, i.e. doesn't have to be IP
 - Expect Active Networking strategies, where packet strategies are changed per hop
 - Very strong opportunity to use protocol incompatibility for separation
 - Native Infiniband, Fiber Channel and Ethernet
 - OSI, IPv4, and IPv6
- Resulting in complex non-IP object identifiers



Focus on CyberSecurity

- Insider Threat
- New Threat Methodology
- Active Defensive Response
- Dynamic Defense



New Public Private Partnership

- With enterprises generating and collecting IP network flow data, for their own Cyber Security purposes, we have a key part of the puzzle.
- CDR data equivalents can be realized for the Internet
 - Can IP network flow data minimize the need for content capture?
 - Enterprises are effectively identifying, analyzing, and responding to CyberSecurity incidents using some network flow strategies.
 - Question is can LEAs get the same level of utility
- Can Society accept the similarities of IP network flow data and Telco CDRs, and give IP network flow data equivalent considerations?
 - Public debate and legislation can address this issue.



Communications Metadata

The New York Times **U.S.** Search All NYTimes.com 

WORLD U.S. N.Y. / REGION BUSINESS TECHNOLOGY SCIENCE HEALTH SPORTS OPINION ARTS STYLE TRAVEL JOBS REAL ESTATE AUTOS

POLITICS EDUCATION TEXAS

Judge Upholds N.S.A.'s Bulk Collection of Data on Calls

By ADAM LIPTAK and MICHAEL S. SCHMIDT
Published: December 27, 2013

WASHINGTON — A federal judge on Friday [ruled](#) that a National Security Agency program that collects enormous troves of phone records is legal, making the latest contribution to an extraordinary debate among courts and a presidential review group about how to balance security and privacy in the era of big data.

Related in Opinion

Editorial: This Week, Mass Surveillance Wins (December 28, 2013)


Letters: Should New Limits Be Put on N.S.A. Surveillance? (December 27, 2013)

Connect With Us on Twitter
Follow @NYTNational for





In just 11 days, the two judges and the presidential panel reached the opposite of consensus on every significant question before them, including the intelligence value of the program, the privacy interests at stake and how the Constitution figures in the analysis.


The latest decision, from Judge William H. Pauley III in New York, could not have been


 FACEBOOK


 TWITTER

 GOOGLE+


 SAVE

 EMAIL

 SHARE

 PRINT

 SINGLE PAGE

 REPRINTS



Search All NYTimes.com

Go



MOST EMAILED

RECOMMENDED FOR YOU

We don't have any personalized recommendations for you at this time. Please try again later.



QoSient

ARGUS

Going Dark

- Changes in technology and billing models in the traditional PSTN are driving some telcos to consider stopping CDR collection and retention.
- Because there are no current statutes or regulations to compel telcos to collect and retain CDRs, assuring CDR availability may be difficult.
- Should we recognize this as a national security vulnerability?
- The CNCI strategy may need to consider more than just data network security issues.

