Introduction to Argus

Network Non-Repudiation and Cyber Security

FloCon 2011
Salt Lake City, UT

Jan 10, 2011

Carter Bullard
CEO/President

QoSient, LLC
150 E 57th Street Suite 12D
New York, New York 10022

carter@qosient.com

Carter Bullard carter@qosient.com

- QoSient - Research and Development Company
  - US DoD, DISA
    - Large Scale Optimization (Operations, Performance, Security)
    - High Performance Network Security Research
    - DARPA CORONET Initial Optical Security Architecture
  - Telecommunications / Media Performance Optimization
  - FBI / CALEA Data Wire-Tapping Working Group
- QoS/Security Network Management - Nortel / Bay
- 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 to Argus

- Discuss the problem space
- Describe Argus design and implementation
- In the context of approaching some real problems
  - Cyber Security
    - Insider Threat protection through Non-Repudiation
  - Degradation of Service
    - Identification
    - Attribution
    - Mitigation

Argus

http://qosient.com/argus

- Argus is a network activity audit system

  Argus was officially started at the CERT-CC as a tool in incident analysis and intrusion research. It was recognized very early that Internet technology had very poor usage accountability, and Argus was a prototype project to demonstrate feasibility of network transactional auditing.

- The first realtime network flow monitor (1989)
- Top 100 security tools used in the Internet today
  - Generates detailed network resource usage logs
  - Source of historical and near realtime data for the complete incident response life cycle
- Designed to provide useful data for network
  - Operations - Service availability and operational status
  - Performance - End-to-end assessment of user traffic
  - Security - Audit / Non-Repudiation

Argus

- Composed of
  - Real-time network flow monitor
  - Network flow data collection system
  - Network flow data processing
  - Audit data repository tools

Argus History

- Georgia Tech (1986)

  Argus was the first data network flow system. Started at Georgia Tech, Argus was used as a real-time network operations and security management tool. Argus monitored the Morris Worm, and was instrumental in discovery for the "Legion of Doom" hacking investigations.

- CERT/SEI/Carnegie Mellon University (1991)

  Argus was officially supported by the CERT as a tool in incident analysis and intrusion research. Used to catalog and annotate any packet file that was provided to the CERT in support of Incident Analysis and Coordination, it was a focal point for research in intrusion analysis and Internet security.

- Argus Open Source (1995 - Present)

  Transitioned into public domain in 1995. Supported by CMU and CERT/SEI at many levels including argus developers mailing list.

  Used now by a large number of educational, commercial and governmental sites for network operations, security and performance management.

Who's using Argus?

- U.S. Government
  - DoD Performance/Security Research - Gargoyle
    - https://software.forge.mil/projects/gargoyle
    - JCTD-Large Data, CORONET, NEMO, JRAE, Millennium Challenge
  - Tactical Network Security Monitoring / Performance Analysis
    - Naval Research Laboratory (NRL), DISA, General Dynamics, IC
- Network Service Providers
  - Operational/Performance Optimization
  - Acceptable Use Policy Verification
- Educational (1000's of sites world-wide)
  - Carnegie Mellon University
  - Stanford University
  - University of Chicago
  - New York University

Enterprise wide near realtime network security audit
Distributed security monitoring
Network security research
Acceptable use policy verification

- ISPs, Enterprises, Corporations, Individuals

Network Situational Awareness

- Argus is designed to be THE network SA sensor
  - Ubiquitously deployable DPI traffic sensor
  - Comprehensive (non-statistical) traffic awareness
  - Provides engineering data, not business intelligence
    - Detailed network transactional performance
    - Network fault identification, discrimination and mitigation
      - Reachability, connectivity, availability, latency, path, flow control etc....
    - Customer gets the primitive data, not just reports/alerts
  - Near realtime and historical capabilities
  - Packet capture replacement
- Supporting a large number of SA applications
  - Advanced Network Functional Assurance (Operations)
    - End-to-End transactional performance tracking (data and control plane)
    - Network component functional assurance (NAT, reachability, encryption)
    - Policy enforcement verification/validation (Access control, path, QoS)
  - Advanced Network Optimization (Security and Performance)
    - Supports network entity and service identification, analysis, planning tracking and control, including baselining, anomaly detection, behavioral analysis and exhaustive forensics

Network Activity Driven Feedback Directed Optimization

Function	Description	
Identify	Discover and Identify comprehensive network behavior	Collect and process network behavioral data
Analyze	Collect and transform data into optimization metrics, establish baselines occurrence probabilities and prioritize events.	
Plan	Establish optimization criteria, both present and future and implement actions, if needed	Provide information and feedback internal and external to the project on the optimization outcomes as events.
Track	Monitor network behavioral indicators to realize an effect.	
Control	Correct for deviations from criteria.	

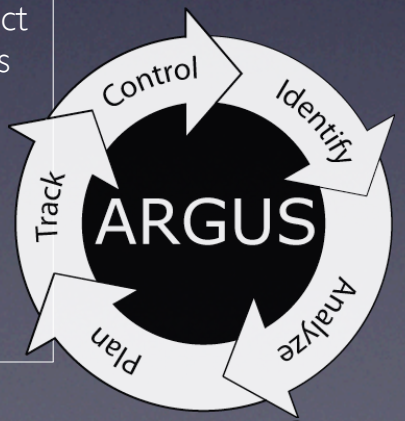

Problem Space

US Cyber Security Focus

- Comprehensive National CyberSecurity Initiative
  - Shifting the US focus from CyberCrime to CyberWarfare
- Strategy and technology focused on new issues
  - Public sector defense, with nation state threats and countermeasures
  - New emphasis on military concepts in Cyber Security
    - Shift from detection to prevention
    - Possible retaliatory mechanisms
- Multi-billion dollar budget will have a significant impact
  - Redefine CyberSecurity for most of the public
  - Compete for best/brightest in security research
  - Determine a new direction for commercial security products

Theoretical Security Threats and Countermeasures

Countermeasures		Threat				
		Unauthorized			Degradation of Service	Repudiation
		Use	Modification	Disclosure		
Authentication	Cryptographic	x		x		
Integrity	Cryptographic		x			
Confidentiality	Cryptographic			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

Non-Repudiation

- Most misunderstood countermeasure *
  - ITU-T Recommendation X.805 security dimension
    - Prevent ability to deny that an activity on the network occurred
- Principal source of true deterrence
  - Non-repudiation provides comprehensive accountability
  - Creates concept that you can get caught
- Argus approach to network non-repudiation
  - Generate data to account for all network activity
    - Comprehensive Network Transactional Audit
    - Mechanism specified by DoD in NCSC-TG-005
      - The Red Book - Trusted Network Interpretation of the Trusted Computer System Evaluation Criteria (1987)
  - Focus on all X.805 Security Planes
    - User, Control and Management network activity

* Crypto-technical redefinition of non-repudiation by Adrian McCullagh in 2000 to apply only to digital signatures has created a great deal of confusion. While you can have repudiation of a signature, it's not the only thing you can repudiate.

Why Non-Repudiation?

- When it exists and structured well, you get
  - Effective information for incident response
    - Fundamental ground truth (if its not there, it didn't happen)
    - Classical forensics support
    - Evidence suitable for criminal and civil complaints
  - Enhanced network situational awareness
    - Network Service Behavioral Baselining
      - Who is really using my DNS servers?
      - What is generating Email in my enterprise?
      - How much data did he transmit last night?
    - Network Policy Enforcement Assurance
      - Are my IPS / IDS / Firewall protections working?
    - Network Fault Attribution
      - Is it an attack? Is it real? Is it a bug? Is it Fred?
  - Enables enhanced analytics, simulation and what if analysis
    - Will this new access control policy, break anything?

Achieving Non-Repudiation

- Comprehensive Activity Accountability
  - Complete Activity Sensing and Reporting
  - Develop Information System with Formal Properties
    - Fundamental ground truth (if its not there, it didn't happen)
- Accurate and Efficient Activity Representation(s)
  - Stored data must represent actual activity
    - Attribute verifiability
      Must be unambiguous with regard to object identification
      Must have a relational algebraic correctness
    - Time synchronization and precision
      Must convey correct order of events
- Fundamental Data Utility
  - Formal and Mature Data Model
  - Useful Data Availability Properties
  - Effective Storage and Retention Strategies

Enterprise Border Awareness

Exterior Interior Model

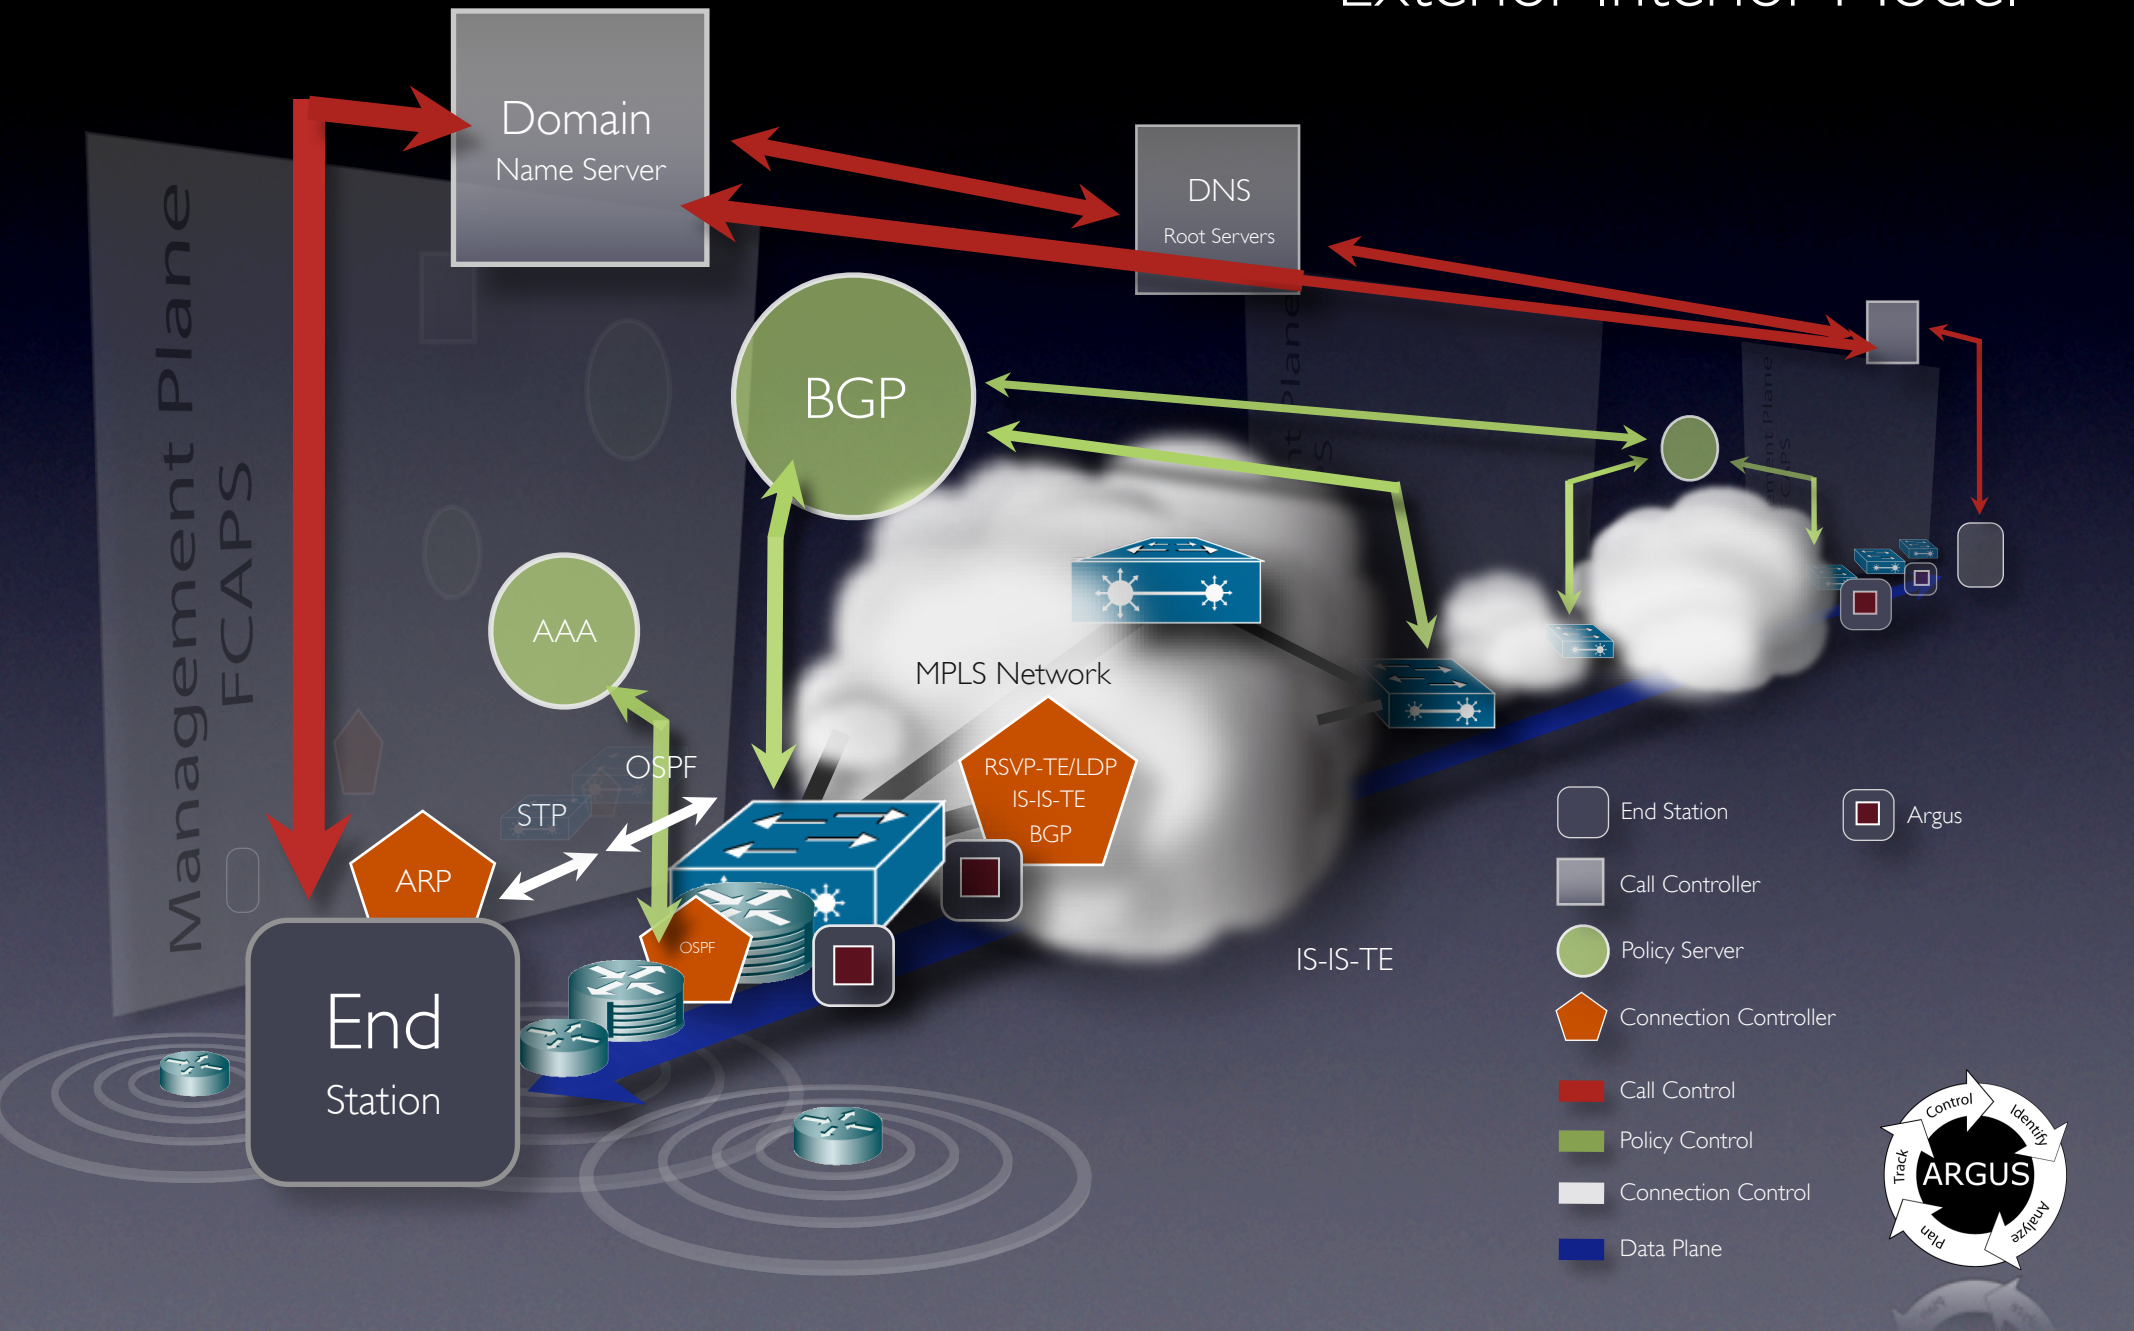

Enterprise Border Monitoring

Internal/External Strategies

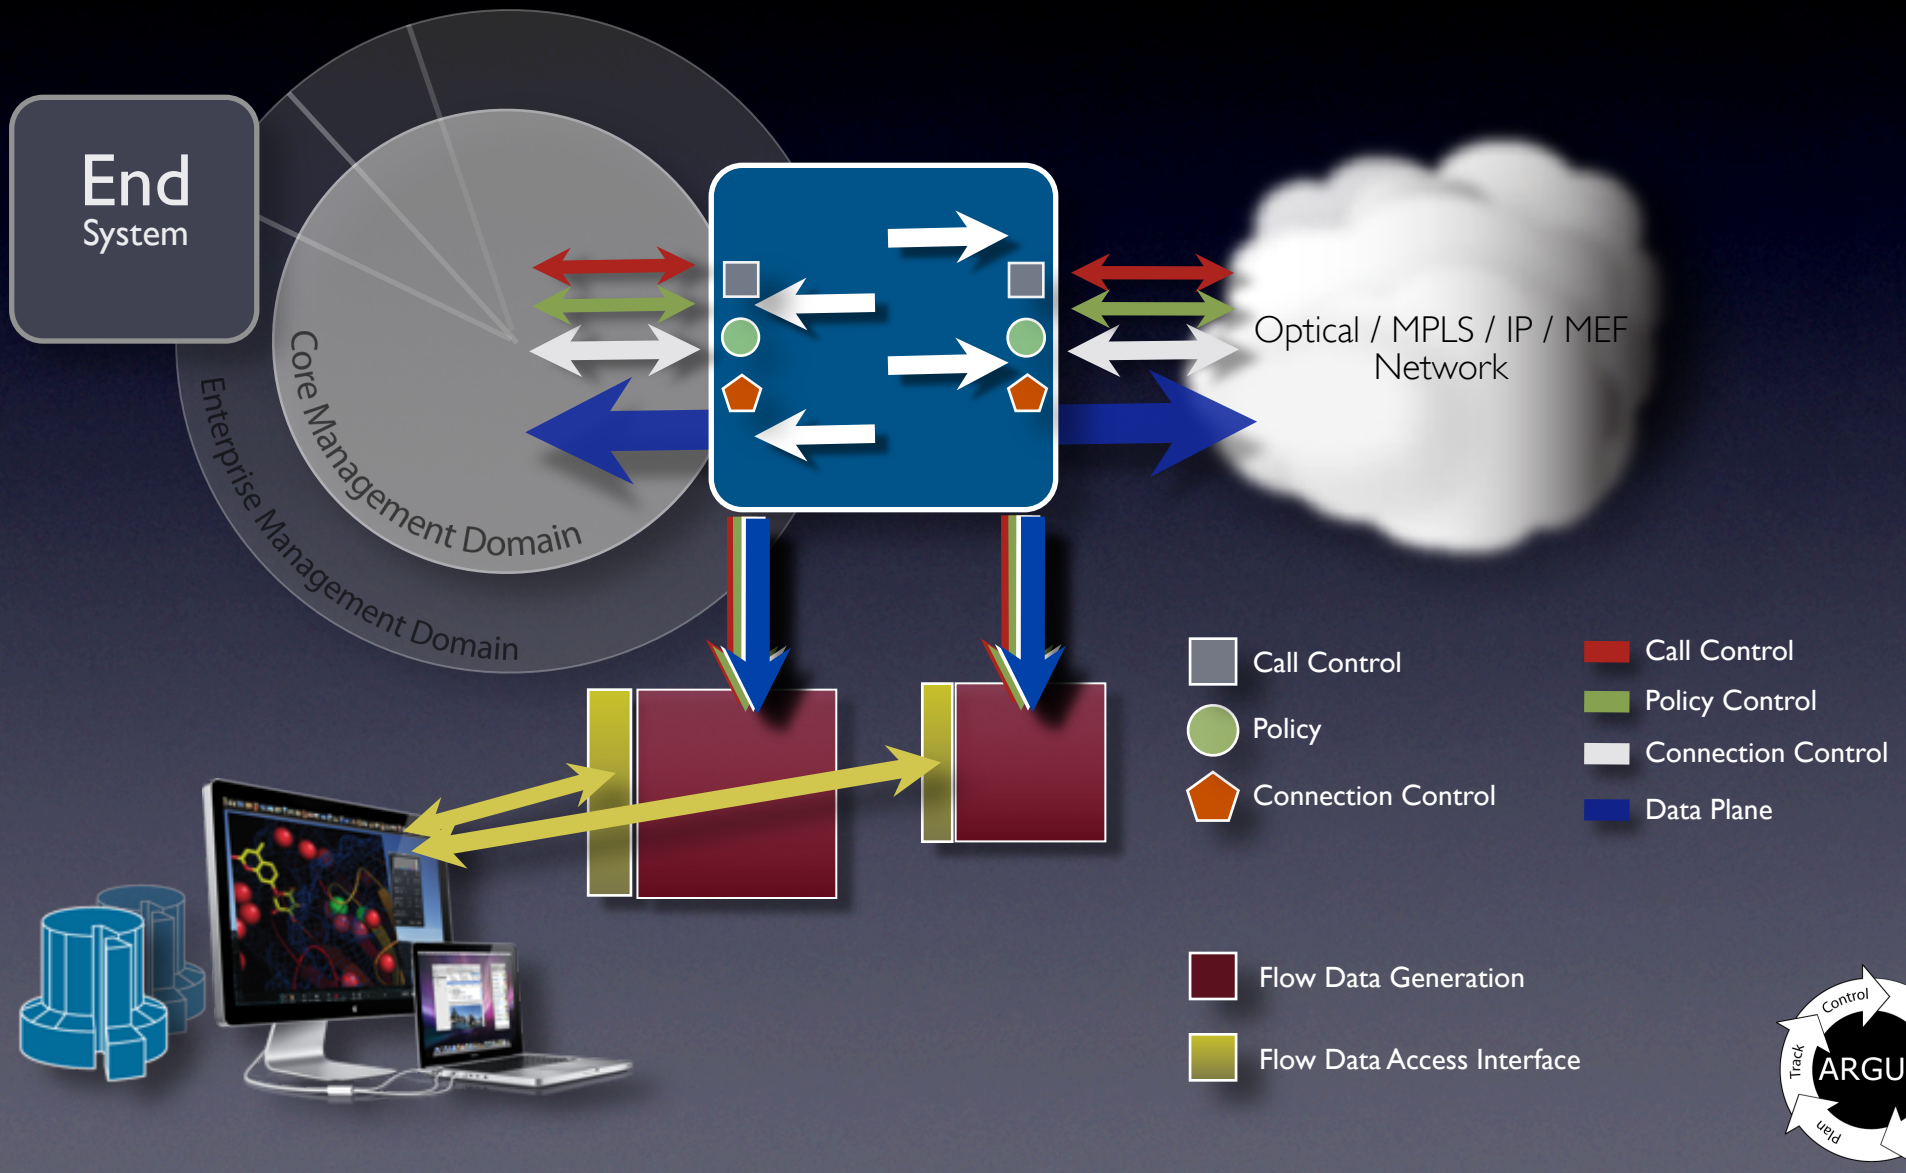

Comprehensive Enterprise Awareness

Approaching the Interior SA

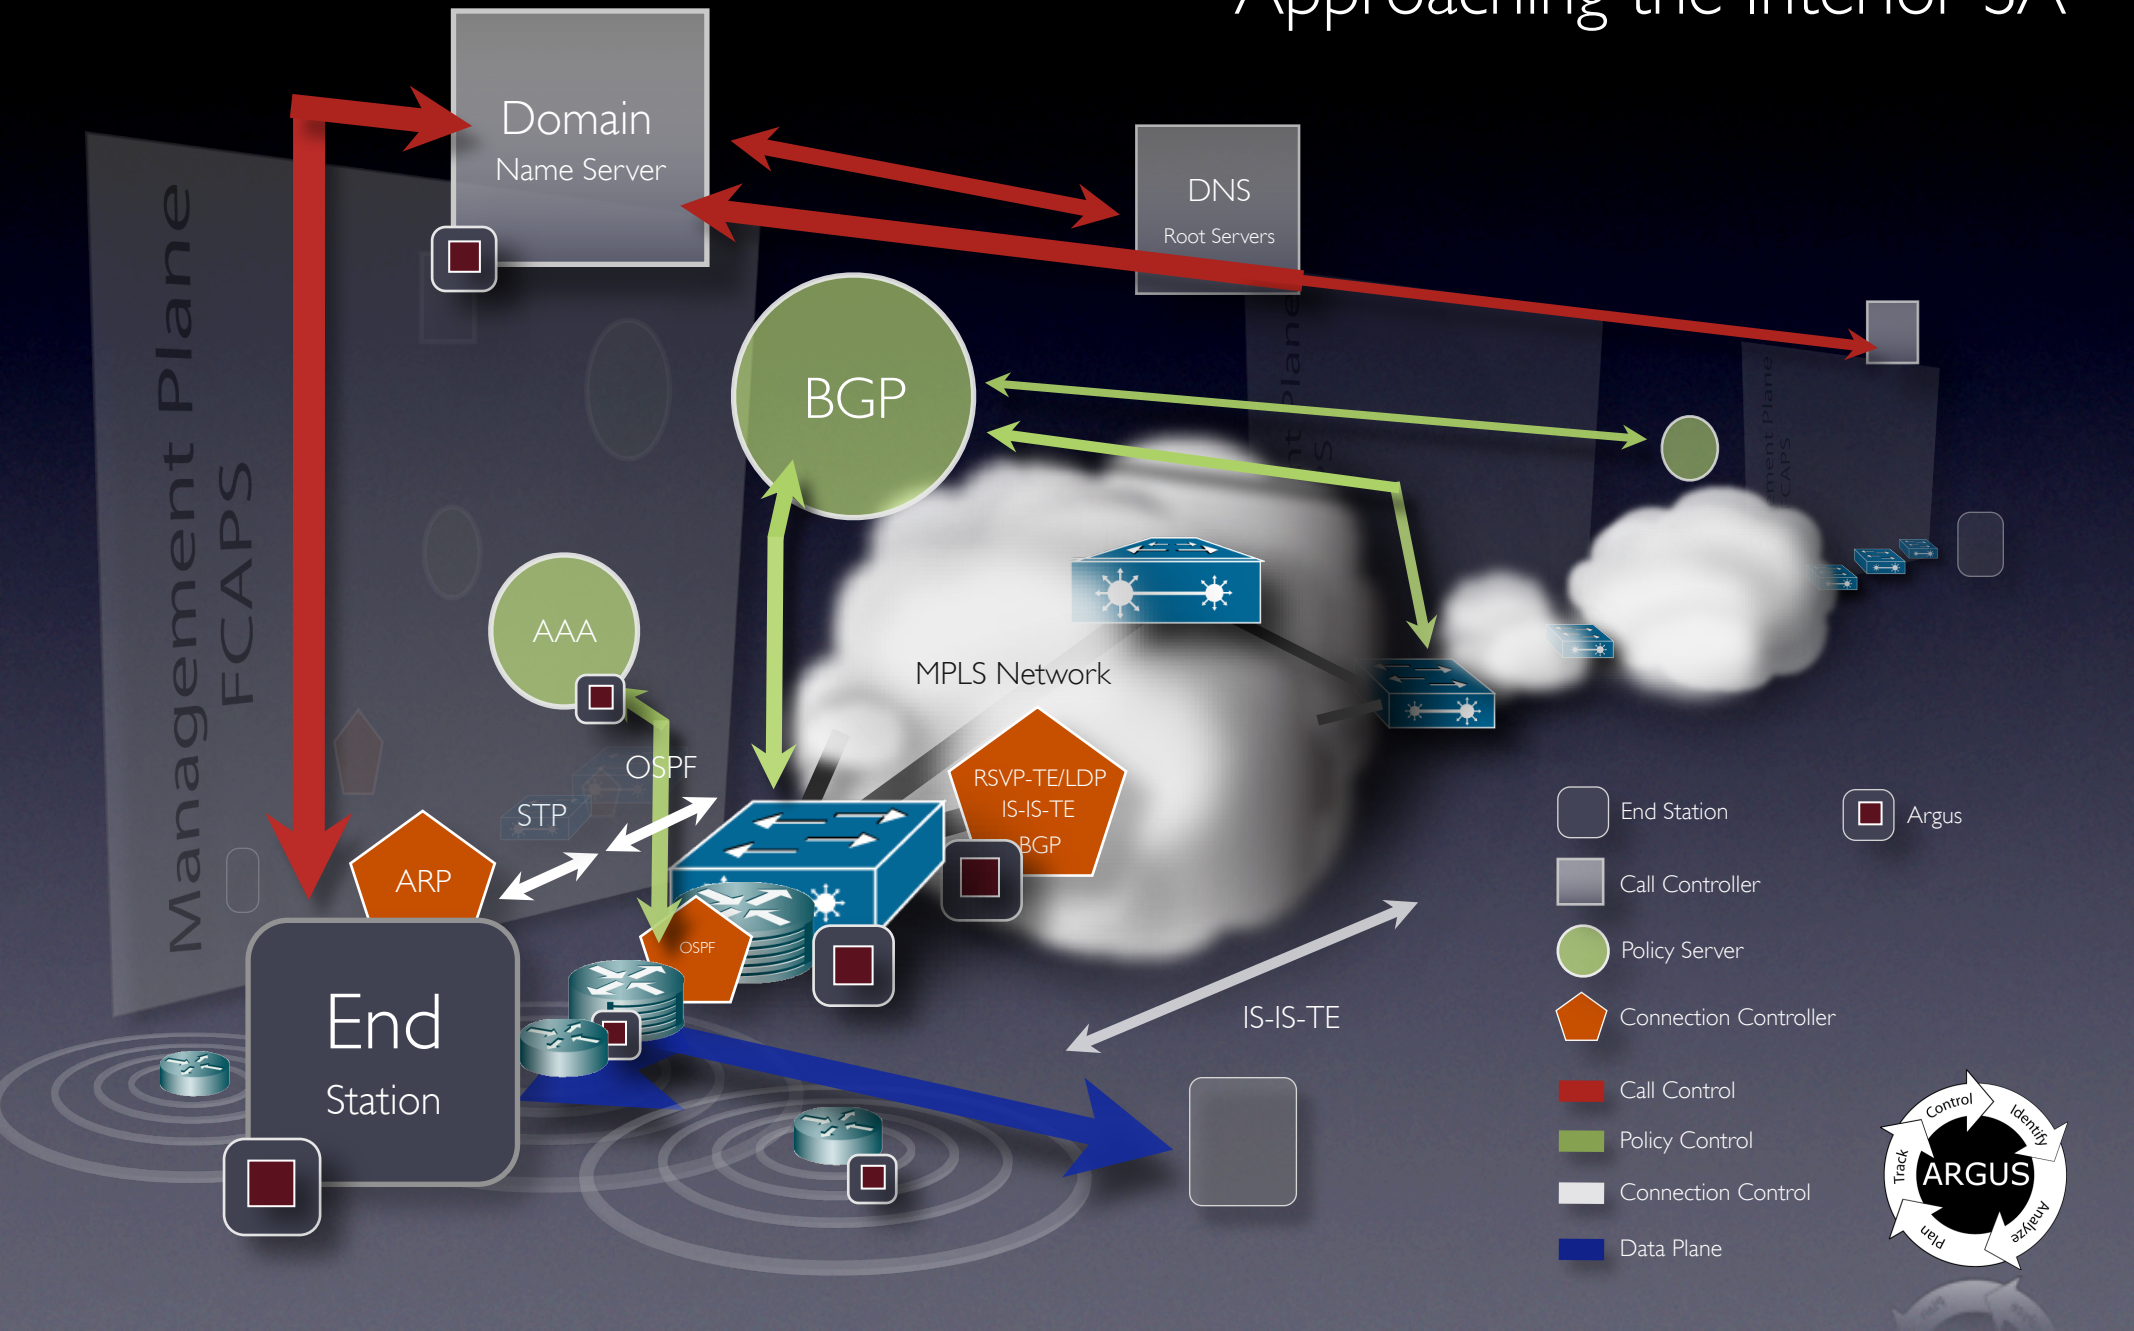

Real world issues

- Non-Repudiation systems must support addressing real world issues
  - Should capture adequate forensics data for incident response
    - Enterprise focused on contemporary security issues
    - Policy enforcement verification validation
- Provide real deterrence
- In a perfect world, you would have a single source for all your network forensics data
  - Support near real-time and historical requirements
  - FISMA continuous network monitoring role

Real world issues

- Incident Response
  - NASA calls. One of your machines attacked a satellite launch
    - Very important military mission
    - Concerned that you may have done it on purpose.
    - Cost the US Gov't $357M
    - 7.5 months ago
    - FBI is coming over in a few minutes
- In a perfect world, you would
  - Review enterprise network activity audit logs as first step
    - Single location for entire enterprises network logs
      - Query for any activity to NASA network or host
      - Pinpoints local hosts involved
      - Now begins the forensics examination
        - Was the attacking machine broken into?
        - If so, (hope so), where did it come from?
    - With multiple internal non-repudiation systems
      - You should be able to identify external / internal attack progression
        - Attack methodologies
        - Identify stepping-stone hosts

Real world issues

- Xerox machines intellectual property loss
  - News story reveals problems with Xerox machines
    - Photocopy machines don't delete copy images
    - Hospitals have lots and lots of Xerox machines
  - What can you do?
  - With single enterprise border non-repudiation system
    - You would know if anyone from the outside ever discovered your Xerox machines in a scan
    - You would know if anything directly accessed your Xerox machines from the outside
  - With non-repudiation system at the Xerox LAN border
    - You would have logs of all network accesses to machine
    - You would know which accesses extracted data rather than presented data to the printer
    - You would have the content visibility needed to identify what images were extracted.

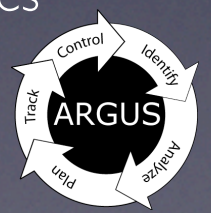

Real world issues

- Intrusion Detection Behavioral Anomalies
  - Access from user X to supercomputer A account
    - Authenticated, acceptable
    - No apparent system log deviations
    - But came from a host outside the normal COI
  - Human analyst noticed the network inconsistency
  - User was on vacation
  - First indication of significant US Gov't problem with Stakkato

Real world issues

- Unintended/Unexpected data exposure
  - Symptom - Poor application performance
    - Database application exhibiting very poor performance
    - Each transaction taking 0.3-0.4 seconds to complete.
    - All software components running on a single machine
    - Absolutely no clues from debugging information
    - Wasn't this bad last week
    - Very, very, very sensitive information
  - Network activity monitoring revealed problem
    - All IPC messaging transmitted into the network
      - Network turned it back around, after it left the domain
    - One software component poorly configured
      - Using server's external name (NAT'ed environment)
    - Very, very, very, very bad

Degradation of Service

- A primary design goal of Argus is DoS identification
  - Argus used in DDoS research papers (1996-2010)
  - CERT Advisory CA-1996-01 UDP Port Denial of Service
  - Many commercial DDoS products are flow data based
- Degradation is an attack on Quality of Service
  - QoS sensitive situational awareness is critical
    - QoS anomaly detection
    - QoS fault management
    - QoS intentional assignments
  - DoS protection really needs to be a part of QoS optimization
    - Can't discriminate QoS degradation when there is poor QoS
- Argus data specifically designed to support:
  - QoS Fault identification/discrimination/mitigation/recovery
    - Pre fault QoS Characterization and Optimization
    - Realtime fault detection and QoS anomaly characterization
    - Post fault recovery, forensics and impact assessments
    - Formal QoS optimization processes

Security and Performance

- Security and performance are tightly coupled concepts
  - Network performance is an asset that needs protection
    - DoD GIG Information availability assurance (DoDD 8500.1)
    - Commercial product delivery dependent on network performance
    - Performance is being specifically attacked
  - Security and performance contribute directly to QoS
  - Security and performance are both optimizations
    - Many times at odds with each other
- Performance awareness data is security awareness data
  - Presence with identifying information is much of the forensics story
- Performance as a leading security indicator
  - Exfiltration and spam generation consume resources
  - Classic "man in the middle" and "traffic diversion" detection
    - Scenarios create measurable end-to-end performance impacts
  - [D]DoS detection is a performance anomaly problem

Degradation of Service (cont)

- QoS Fault Discrimination
  - Traditional QoS fault detection and mitigation
    - End-to-End oriented QoS tracking capability
      - Availability, demands, path, latency and efficiency modifications
      - Host vs Network QoS impact discrimination
      - Distributed sensor strategies provide best "finger pointing" capabilities
    - Historical audit provides baseline analytics for boundary tests
    - Discrimination can involve session dependency analysis
      - Front end network service dependancies
        - ARP, DNS, IP reachability, TCP availability, Service
      - Back end service dependency awareness
  - Discriminating intentional QoS failure
    - Protocol vulnerability exploitations
    - Exclusionary methods for attack designation
      - Flash crowd vs DDoS
      - Indirect attack assessment support

Distributed Situational Awareness

Multi-Probe Multi-Site

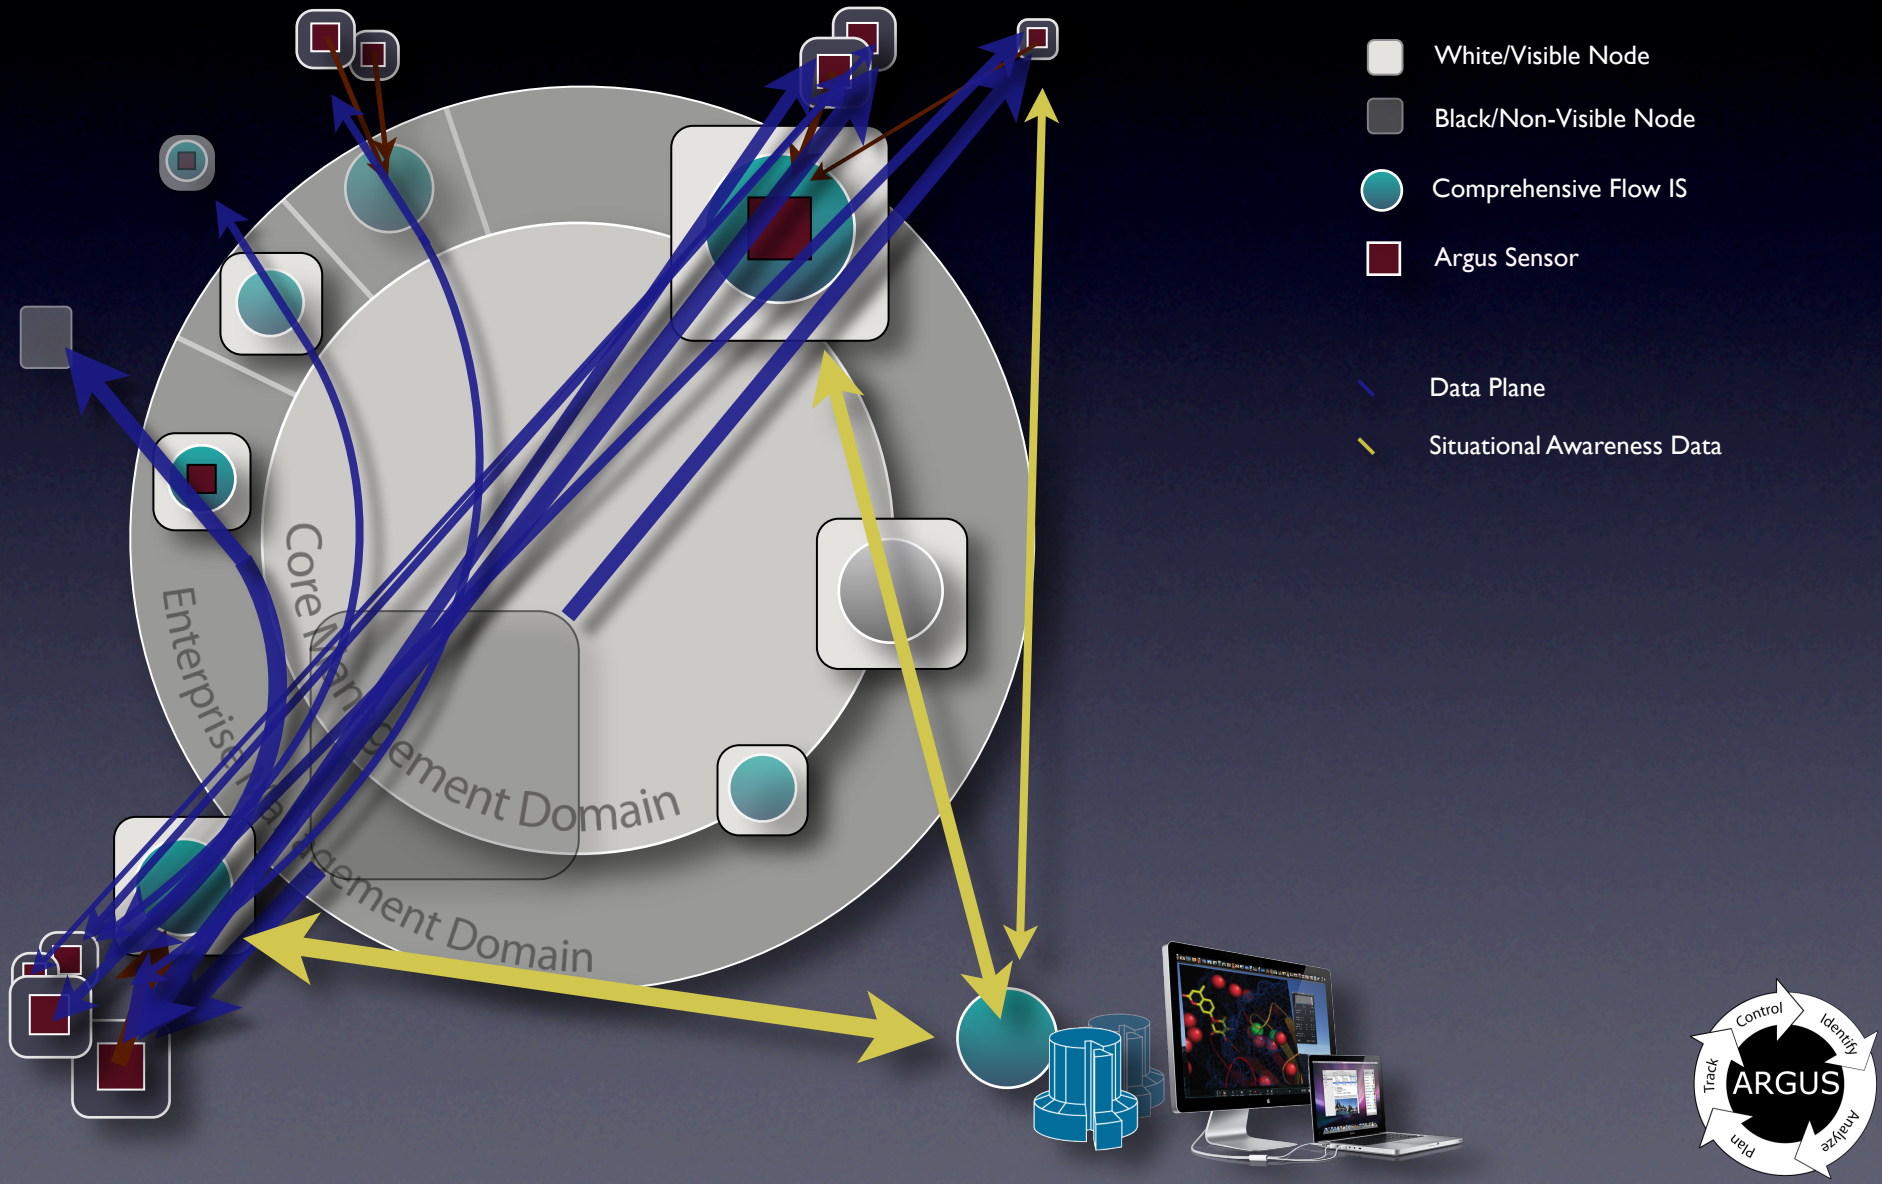

End-to-End QoS Measurement

Network Performance

Black Core Mesh

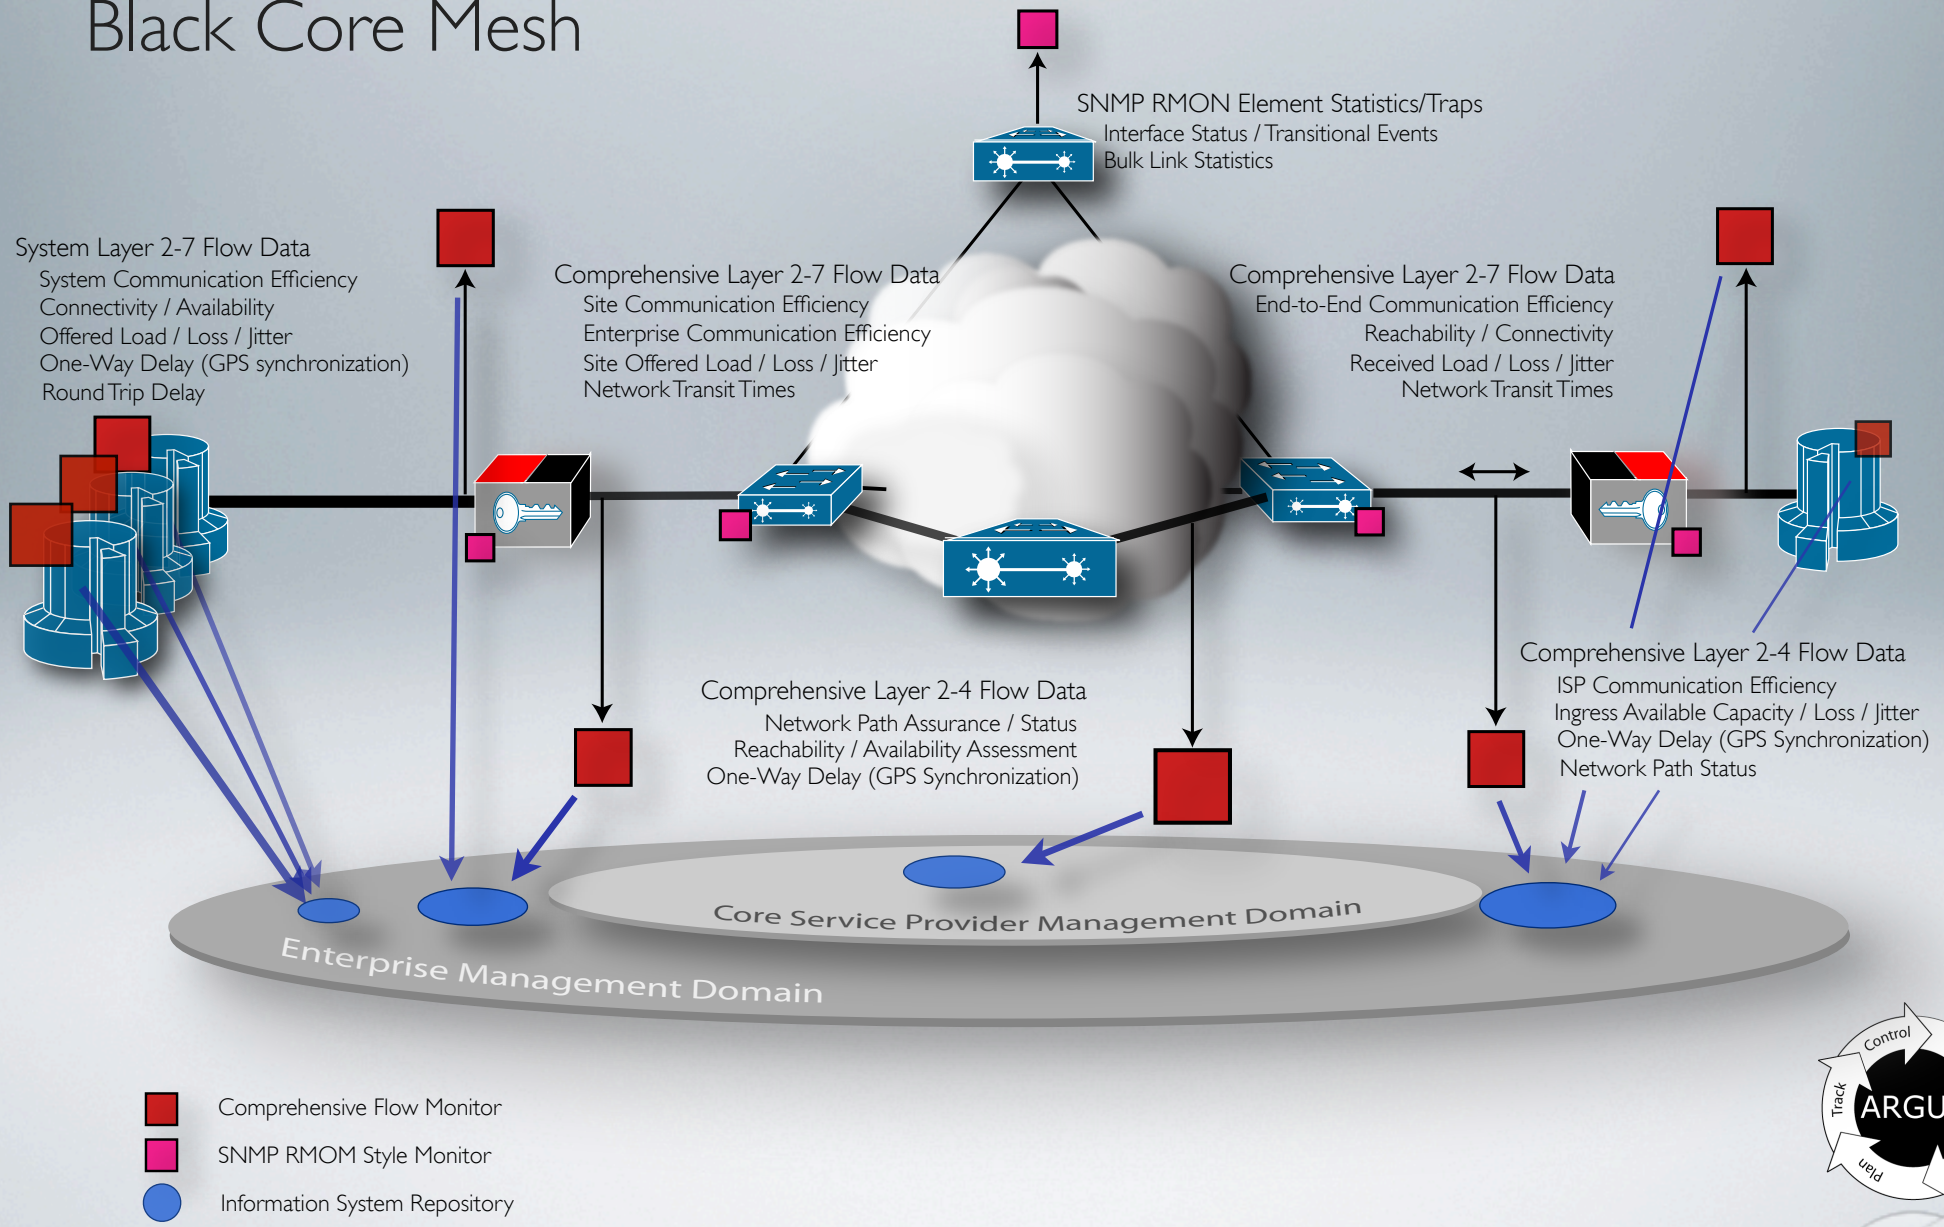

Distributed Situational Awareness

Multi-Probe Multi-Site

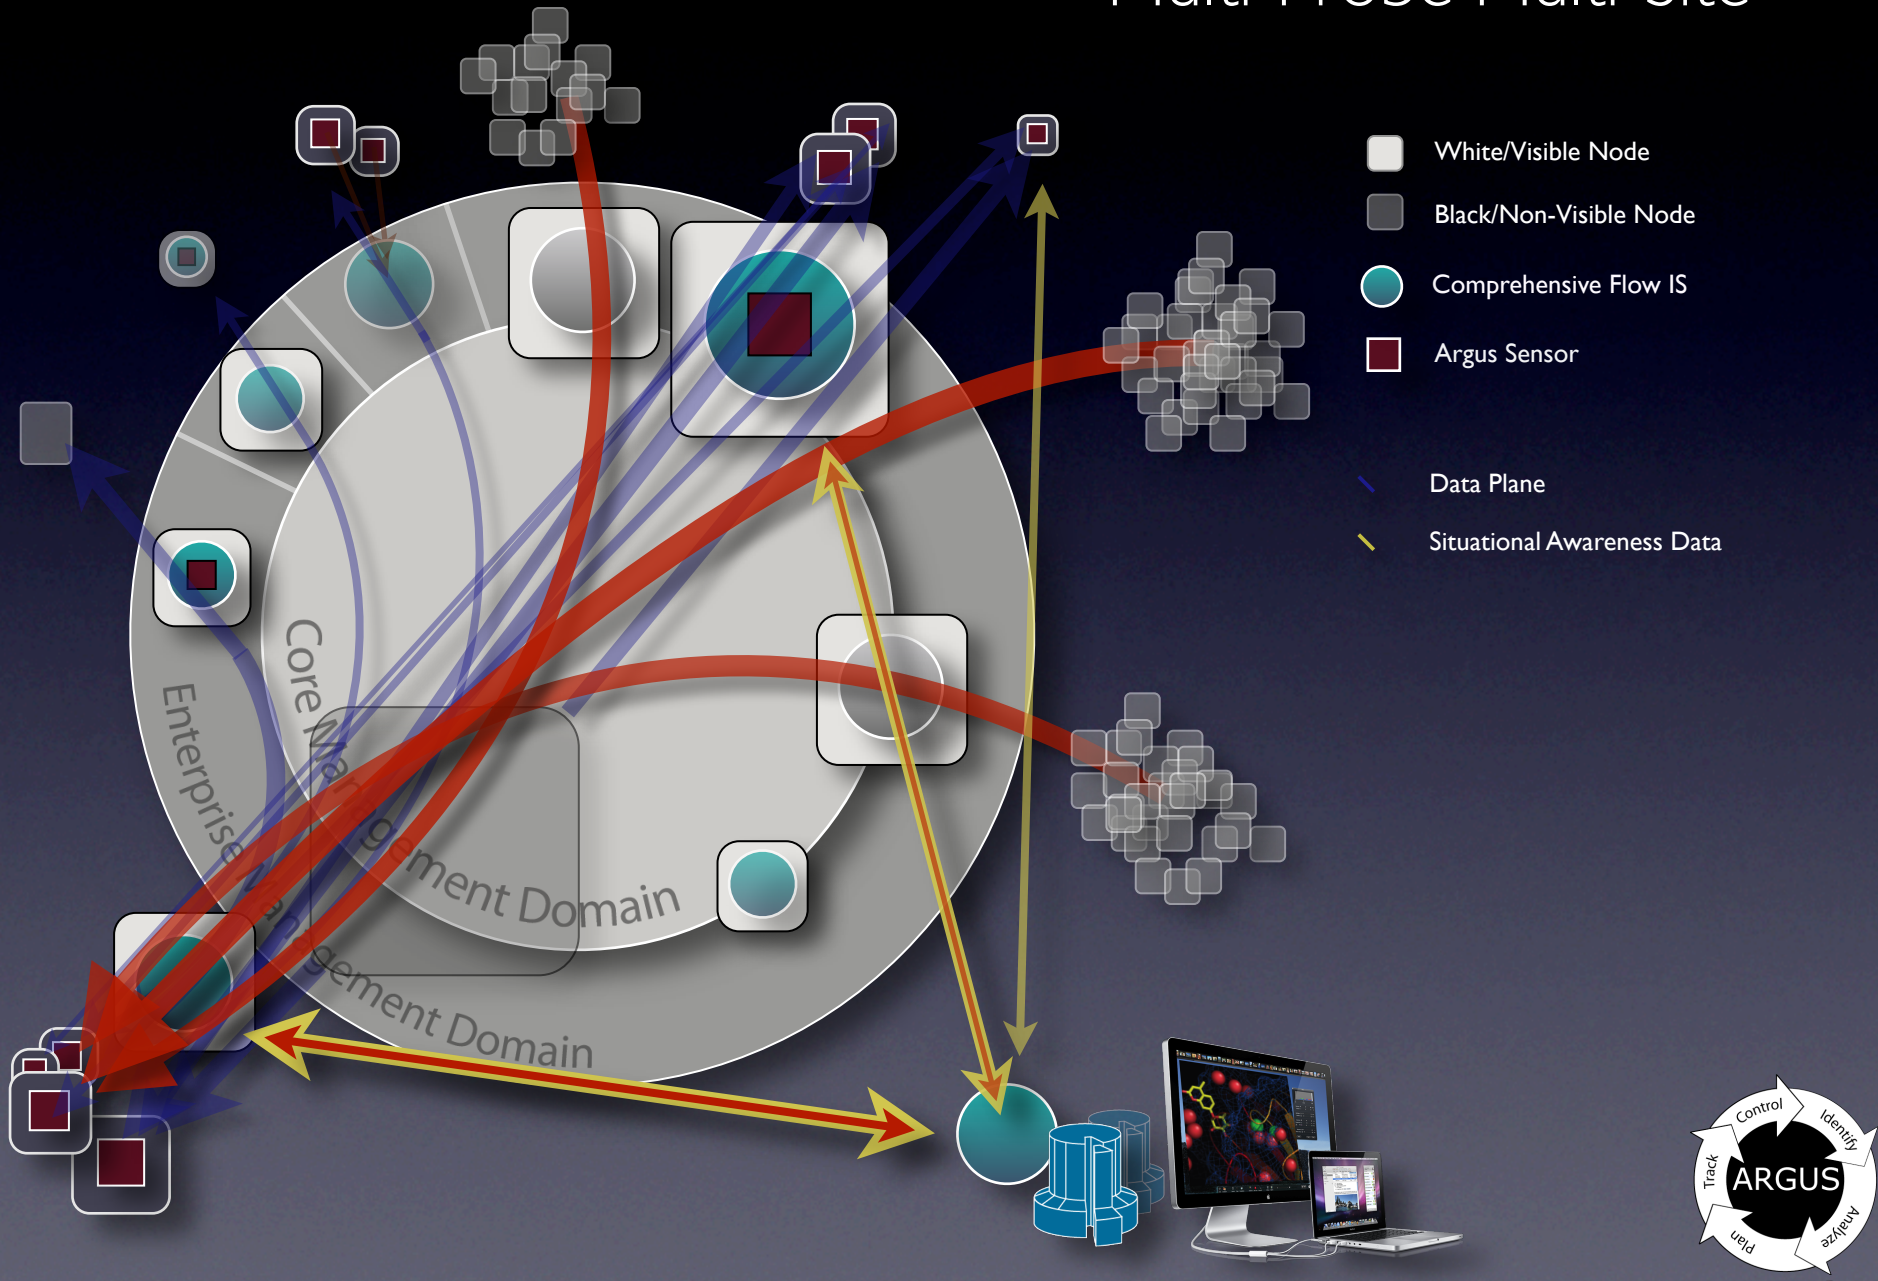

Distributed Situational Awareness

IP Spoofing Scenarios

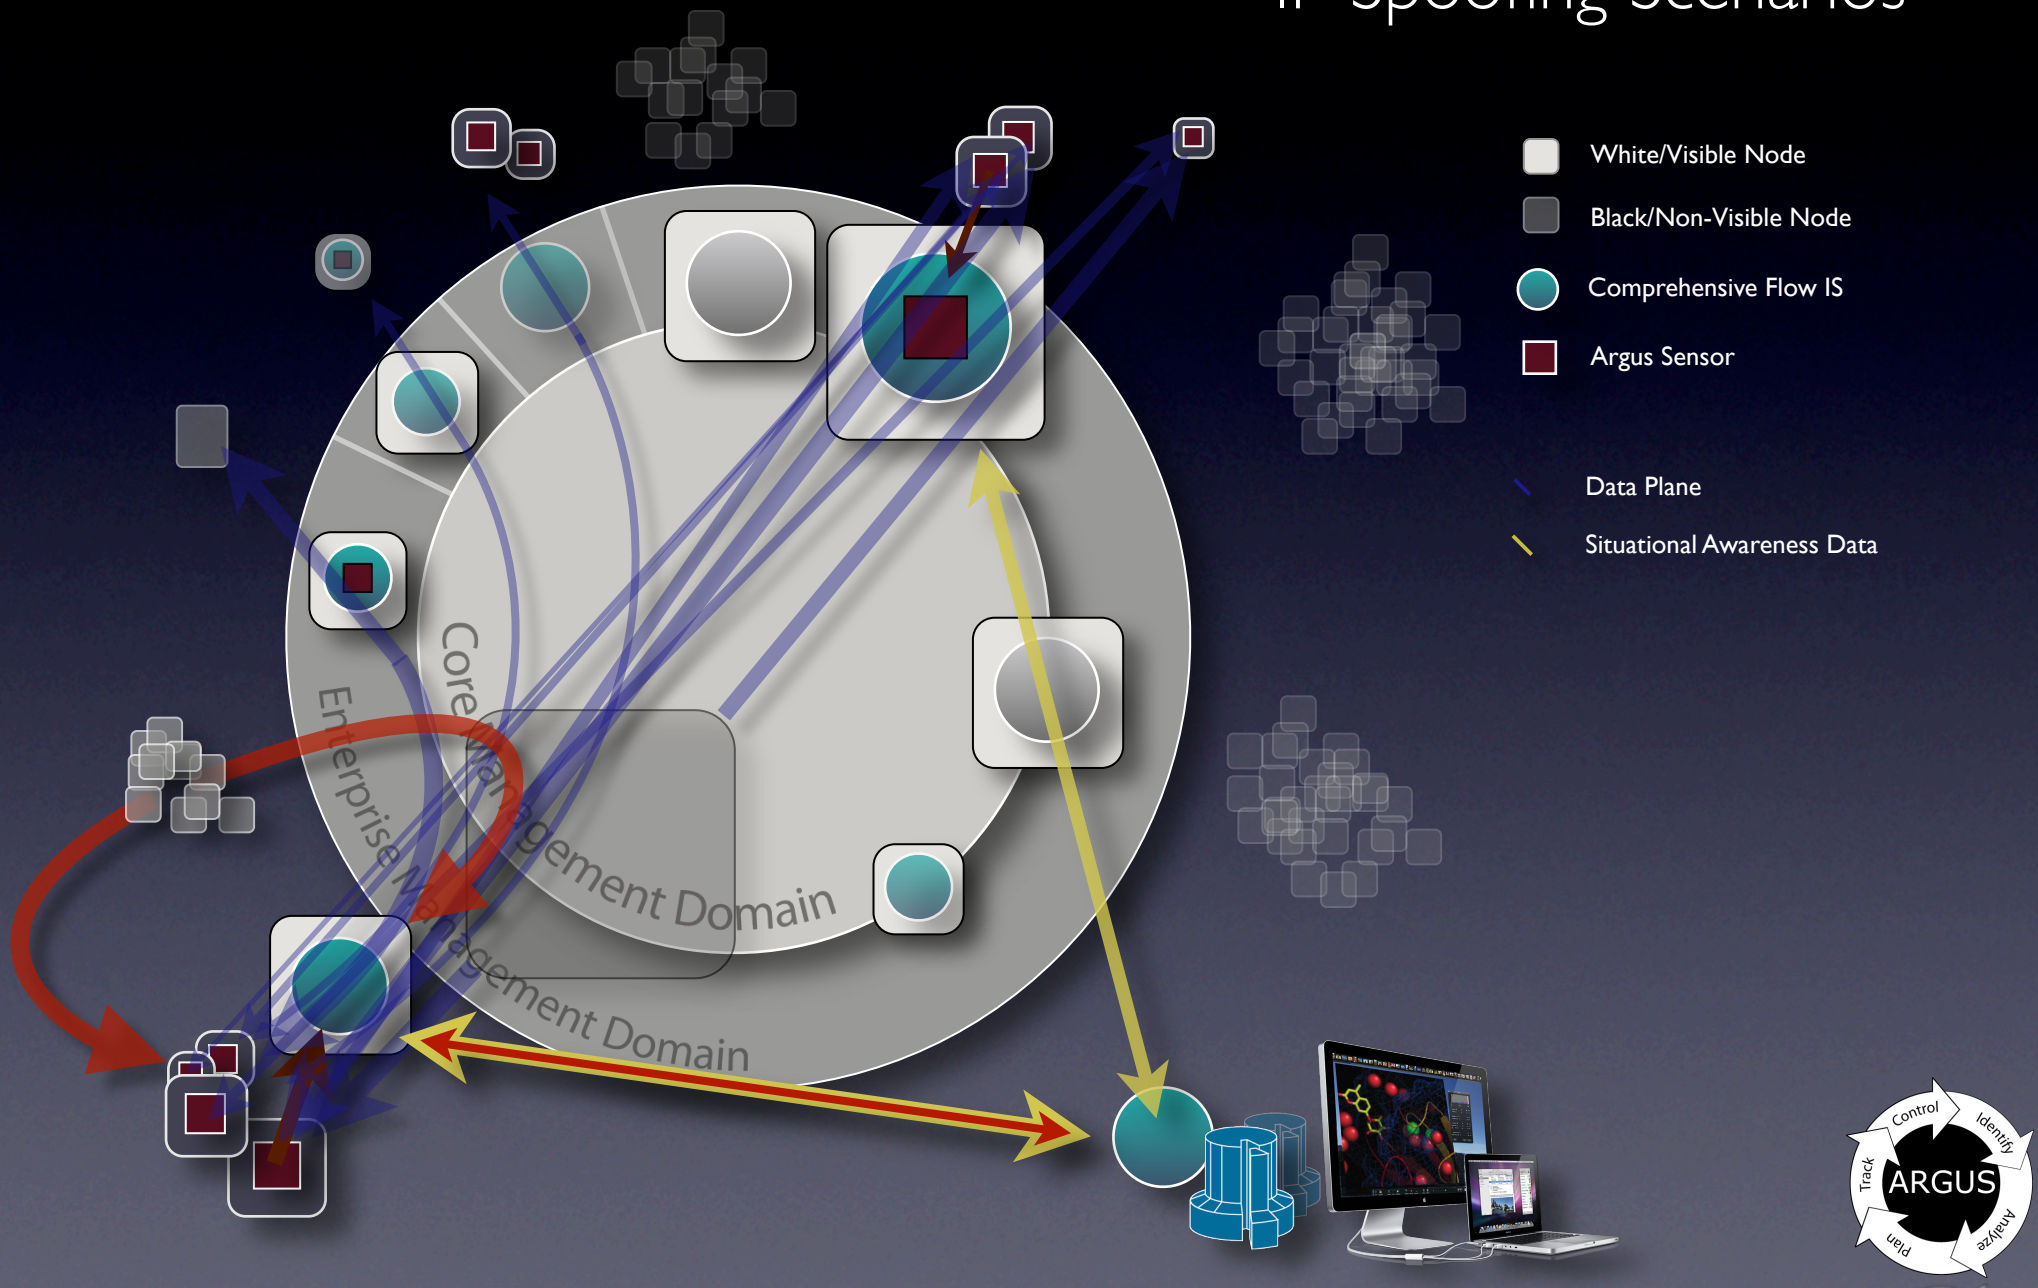

Degradation of Service (cont)

- Methods used to defeat [D]DoS mitigation
  - Mitigation involves denying access from list of exploit IP addresses
  - IP address spoofing
    - Host along attack path emulates [D]DoS traffic
      - Internal host that can "see" the target can forge 100,000's of simultaneous active connections to/from foreign hosts
    - Routing mediated address spoofing
      - BGP modifications allow near local networks to spoof address space
      - Internal modification to locally support foreign address space
        - Static routes can be setup so that "China" is routed to port 23b
        - Control plane attacks (ARP, RIP, OSPF) to advertise "China" is over here
- Result is that you just can't seem to shake the attack
- Distributed sensing detects this scenario
- Net-spatial data and active traceback strategies

Distributed Situational Awareness

Mediation

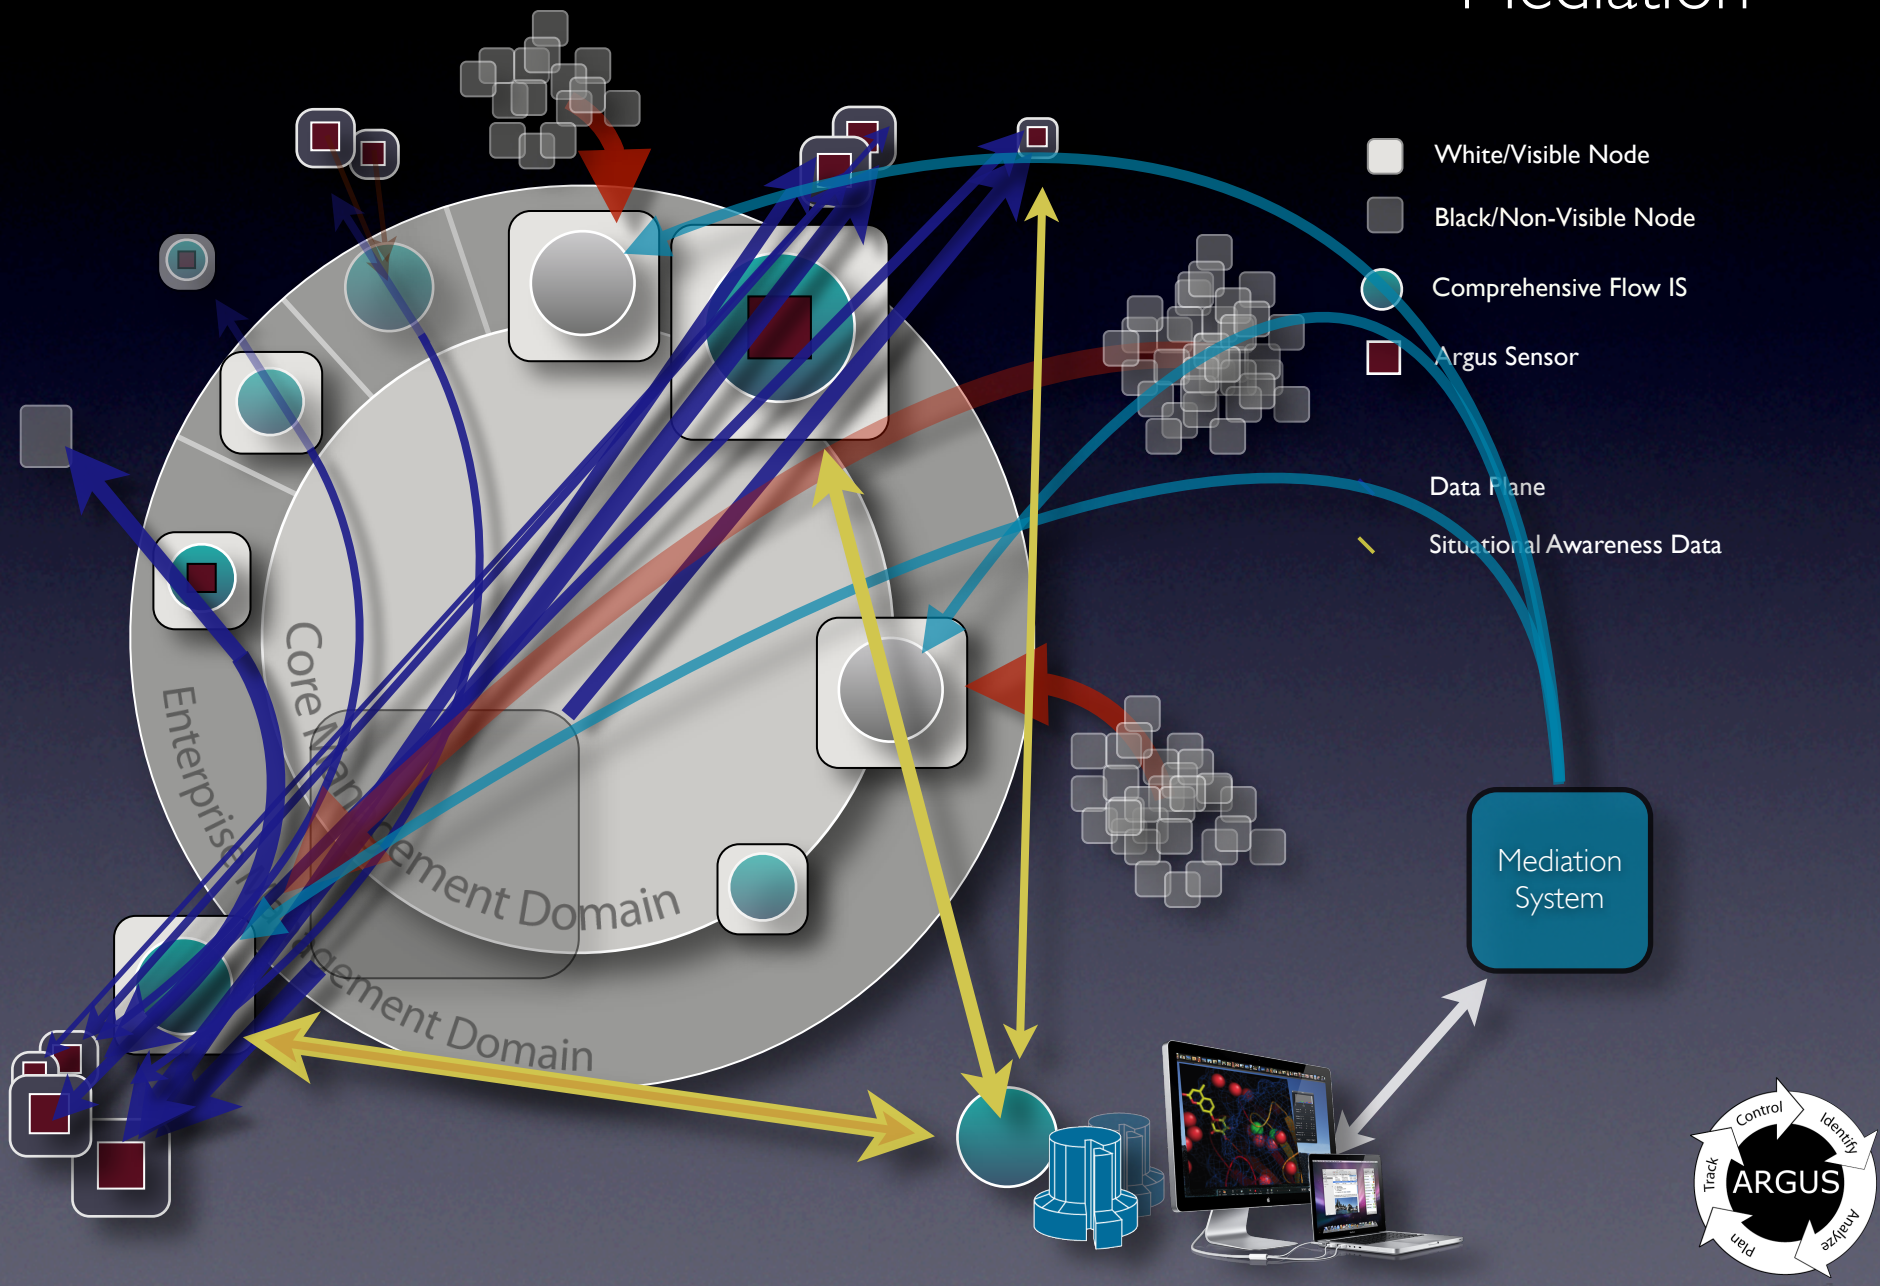

Degradation of Service (cont)

- QoS Fault Mediation
  - Argus can provide information for effective mediation
  - Provide realtime forensics for threat analysis
    - Realize that QoS of critical assets are being affected
    - Provide real-time list of active nodes
    - For web attacks provide recurring URL visits
  - Provide CIDR addresses to block
    - Need to be sensitive to ACL limits of network equipment
    - Need to be cleaver when trying to block 50K IP addresses
  - Provide CIDR addresses to allow
    - Historical Community of Interest (COI) for allowable customers
    - The list of networks active at the initial time of attack
- Argus information to assure mediation worked
  - Network now performing within SLA
  - Track conditions to indicate when to revert, if ever

Building Non-Repudiation Systems

Non-Repudiation Concepts

ITU X.813

Information Technology

Open Systems Interconnection

Security Frameworks in Open Systems: Non-repudiation Framework

"The Non-repudiation service involves the generation, verification and recording of evidence. ..... Disputes cannot be resolved unless the evidence has been previously recorded."

The service provides the following facilities which can be used in the event of an attempted repudiation:

- generation of evidence
- recording of evidence
- verification of generated evidence
- retrieval and re-verification of the evidence

Formal Non-Repudiation Systems

J-STD-025A

WAI/GT/
FuncSpecs
v1.0.1 (2000-06)

- Telephone Billing Records (retrospective)
- J-STD-025A / ETSI TS 101 671 (prospective)
  - Dialed Number Recorder (DNR/Pen Register)
  - Full Audio Interception (Title III/FISA)
- When concepts applied to data networks:
  - Content capture unencrypted (keys)
  - Information Protection Requirements
  - Geo-Location Information
  - Time Constraints
  - Unchanged State of Service

ETSI ES 201 671

Telecommunications Security

Lawful Interception(LI); Handover interface for the lawful interception of telecommunications traffic

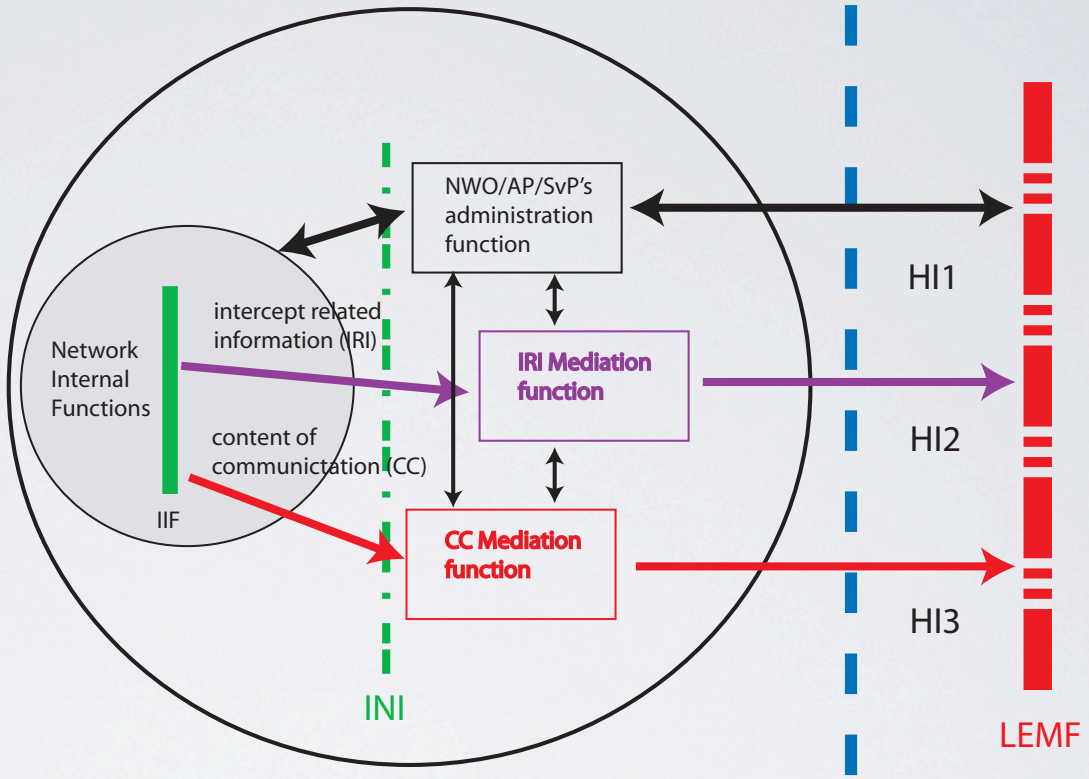

IIF: internal interception function
INI: internal network interface

HI1: administrative infomration
HI2: intercept related information
HI3: content of communication

NOTE 1: Figure 1 shows only a reference configuration, with a logical reprsentation of the entities involved in lawful interception and does not mandate separate physical entities.

NOTE 2: The mediation functions may be transparent.

Functional Block Diagram Showing Handover Interface HI

ETSI ES 201 671 V2.1.1 (2001-09) Telecommunications security; Lawful Interception(LI); Handover interface for the lawful interception o ftelecommunications traffic

X.813 Non-Repudiation Entities

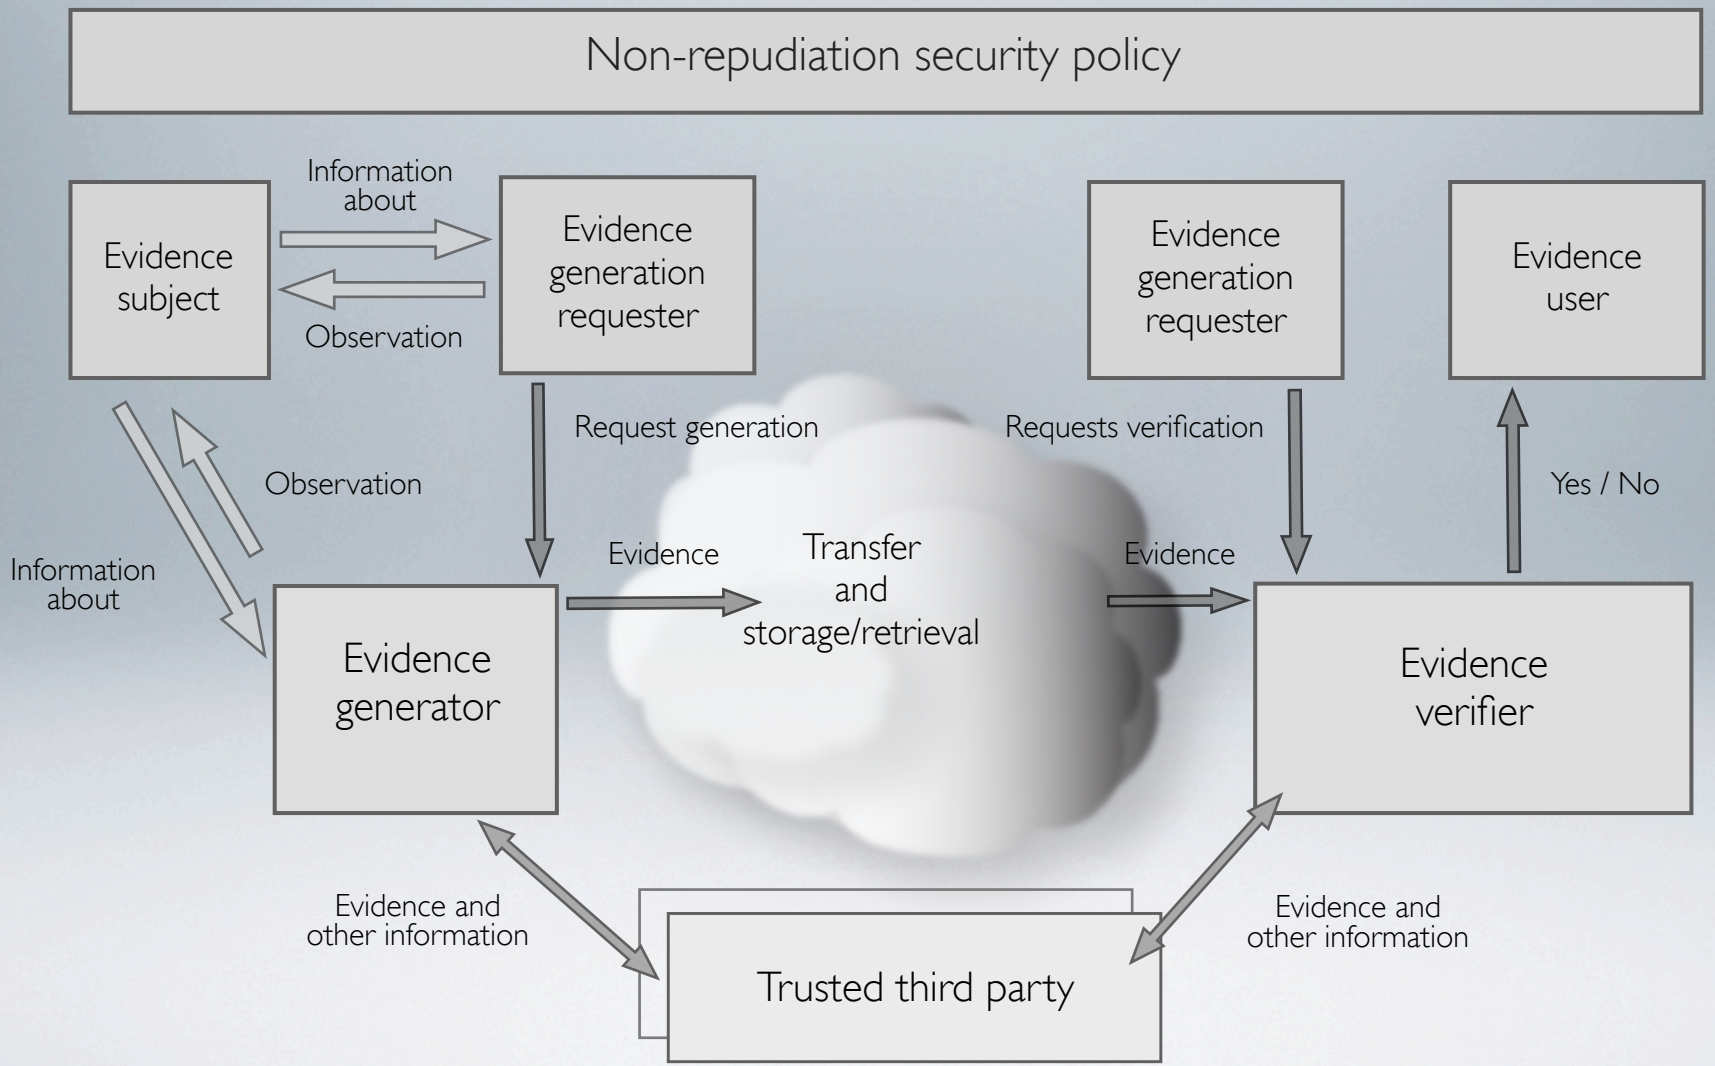

ITU-T Rec. X.813 (1996 E) TISO7760-96/D01

X.813 Non-Repudiation Entities

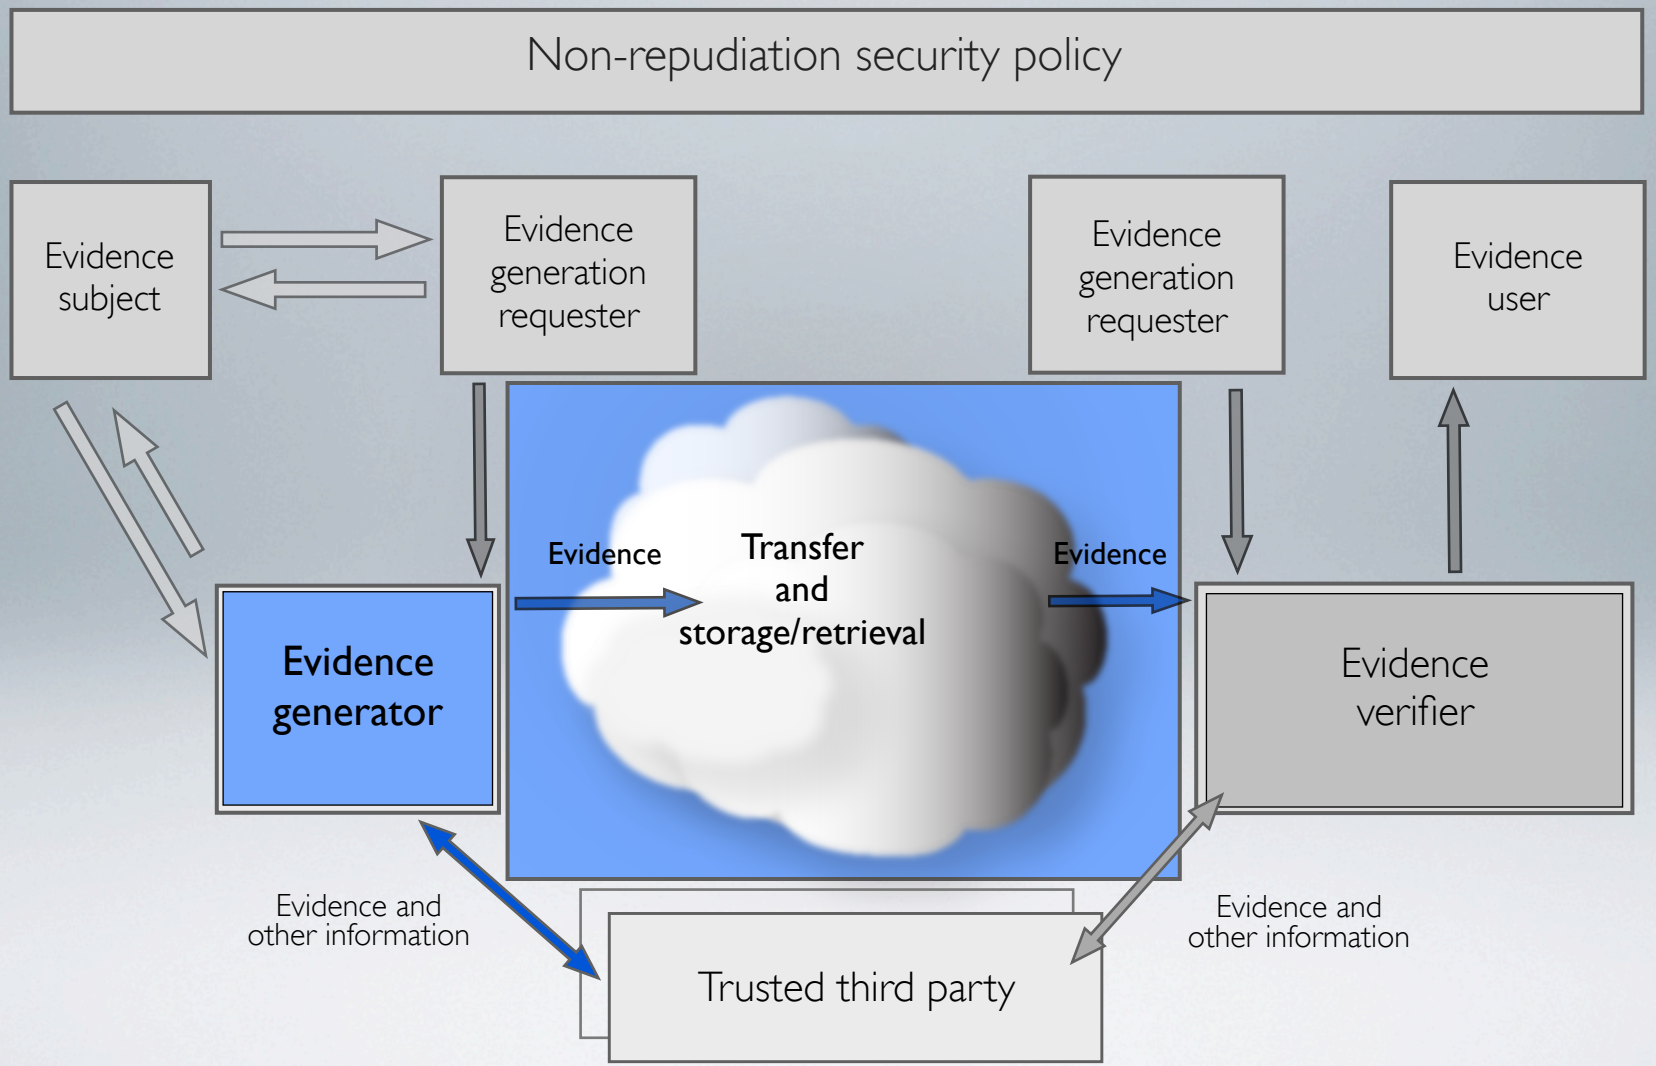

ITU-T Rec. X.813 (1996 E) TISO7760-96/D01

LEAs and Telecommunications

- US Lawful Intercept
  - Pen Register
  - Trap and Trace
  - Content Interception

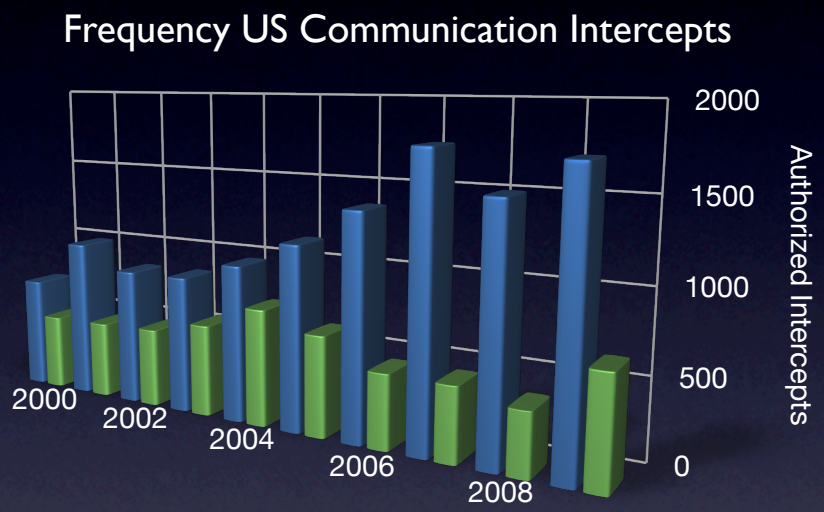

Source http://uscourts.gov/Statistics

- However, the principal interaction of LEAs with the telecommunications industry are subpoenas of telephone billing records. (over 100X number of Lawful Intercepts)

Private-Public Partnership

- Telephone Billing Records, Call Detail Records (CDR), are a by product of Telco network operations and considered Customer Proprietary Network Information (CPNI).
- Society provides privacy protection for CPNI
  - Of course, the customer can have access to the information at anytime
  - No voluntary disclosure by telco, without customer approval
  - Government can gain access through warrants, or trail subpoenas
- CDRs contain no content, but have high security utility
  - Provide an effective and well recognized deterrent against crime
  - Private and Public sectors rely on CDRs for investigative purposes
    - Provides an enhanced Situational Awareness
    - Used by LEAs to demonstrate need for further investigation
- CDRs directly minimizes the use of Lawful Intercept
- Can CDR equivalent strategies be realized in the Internet?
- Is it possible to enable this partnership in the Internet?
- Can the CNCI use this type of partnership for national Cyber Security?

What Are CDRs Used For?

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

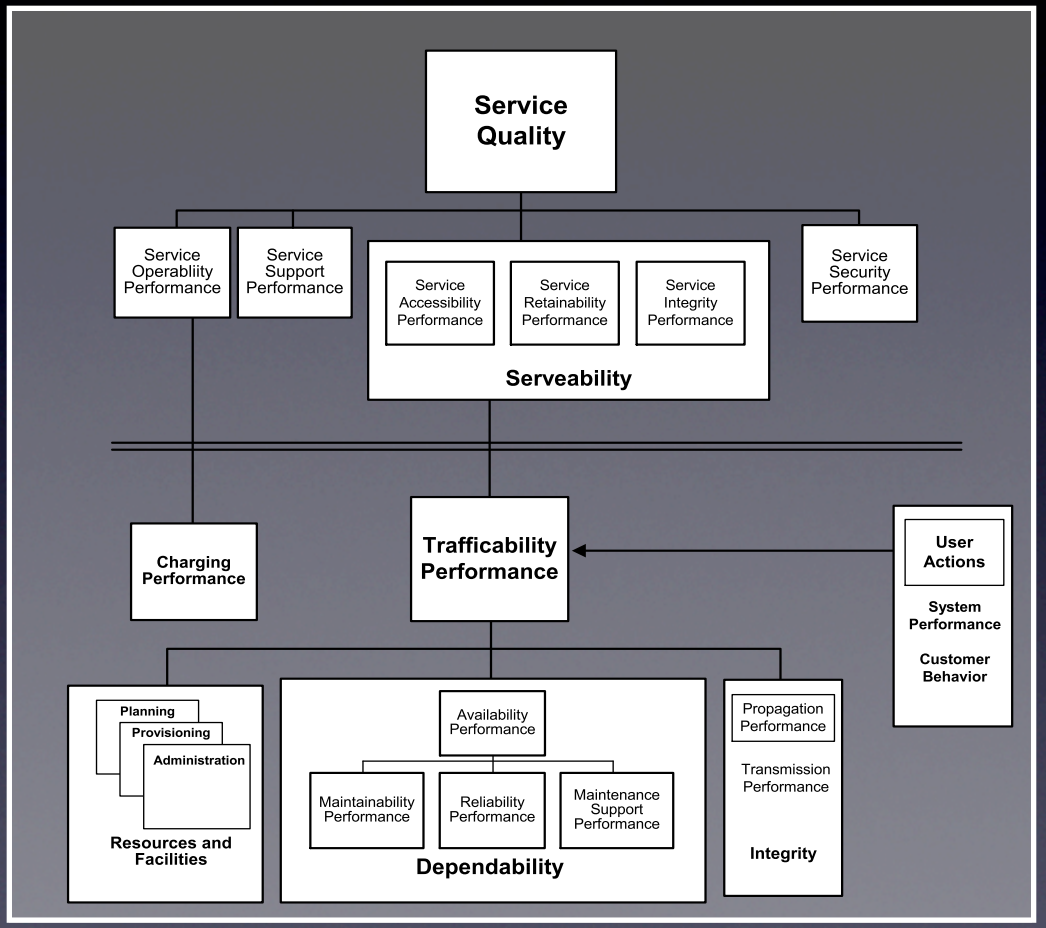

From ITU-T Recommendation E.800 Quality of Service, Network Management and Traffic Engineering

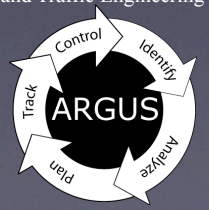

Network Auditing

- Specified by DoD in NCSC-TG-005
  - The Red Book - Trusted Network Interpretation of the Trusted Computer System Evaluation Criteria (1987)
- Goal to provide Non-Repudiation
  - Comprehensive audits accounting for all network use
  - Creates **real deterrence** in formal systems
    - Fear of getting caught is extremely powerful
  - Utility comes from the quality of collected information
- Internet network transaction auditing is emerging
  - Started at the CMU CERT-CC in early 1990's - Argus
  - Directly modeled after the PSTN CDR
  - Aspects of IP network auditing are being standardized

Achieving Non-Repudiation

- Comprehensive Activity Accountability
  - Complete Activity Sensing and Reporting
  - Develop Information System with Formal Properties
    - Fundamental ground truth (if its not there, it didn't happen)
- Accurate and Efficient Activity Representation(s)
  - Stored data must represent actual activity
    - Attribute verifiability
      - Must be unambiguous with regard to object identification
      - Must have a relational algebraic correctness
    - Time synchronization and precision
      - Must convey correct order of events
- Fundamental Data Utility
  - Formal and Mature Data Model
  - Useful Data Availability Properties
  - Effective Storage and Retention Strategies

Comprehensive Accountability

- Account for all network activity
  - Because any network activity can be associated with a cyber-security activity
    - Generally, if you aren't looking 'there', 'there' is where they will be
    - Hidden variables enable the adversary
  - Observation scope must be relevant
    - Utility of collected information should be very high
    - Using PSTN as guide, ISP can collect anything, but share nothing.
- Argus approach to network non-repudiation
  - Generate data to account for all network activity
    - Comprehensive Network Transactional Audit
    - Mechanism specified by DoD in NCSC-TG-005
      - The Red Book - Trusted Network Interpretation of the Trusted Computer System Evaluation Criteria (1987)
  - Focus on all X.805 Security Planes
    - User, Control and Management network activity

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.
- Available IP Flow Information
  - Argus
    - Control and Data Plane network forensics auditing
    - Archive, file, stream formats. (Binary, SQL, CSV, XML)
  - YAF/SiLK - CERT-CC (IP data only)
    - Designed for Cyber security forensics analysis
    - IETF IPFIX stream formats. Binary file format.
  - IPDR - Billing and Usage Accountability (IP data only)
    - ATIS, ANSI, CableLabs, SCTE, 3GPP, Java CP, ITU/NGN
    - File and stream formats (XML).
  - Netflow, JFlow, Sflow (IP data only)
    - Integrated network vendor flow information - statistical/sampled
    - Used primarily for router operations, network management

Comprehensive Enterprise Awareness

Approaching the Interior SA

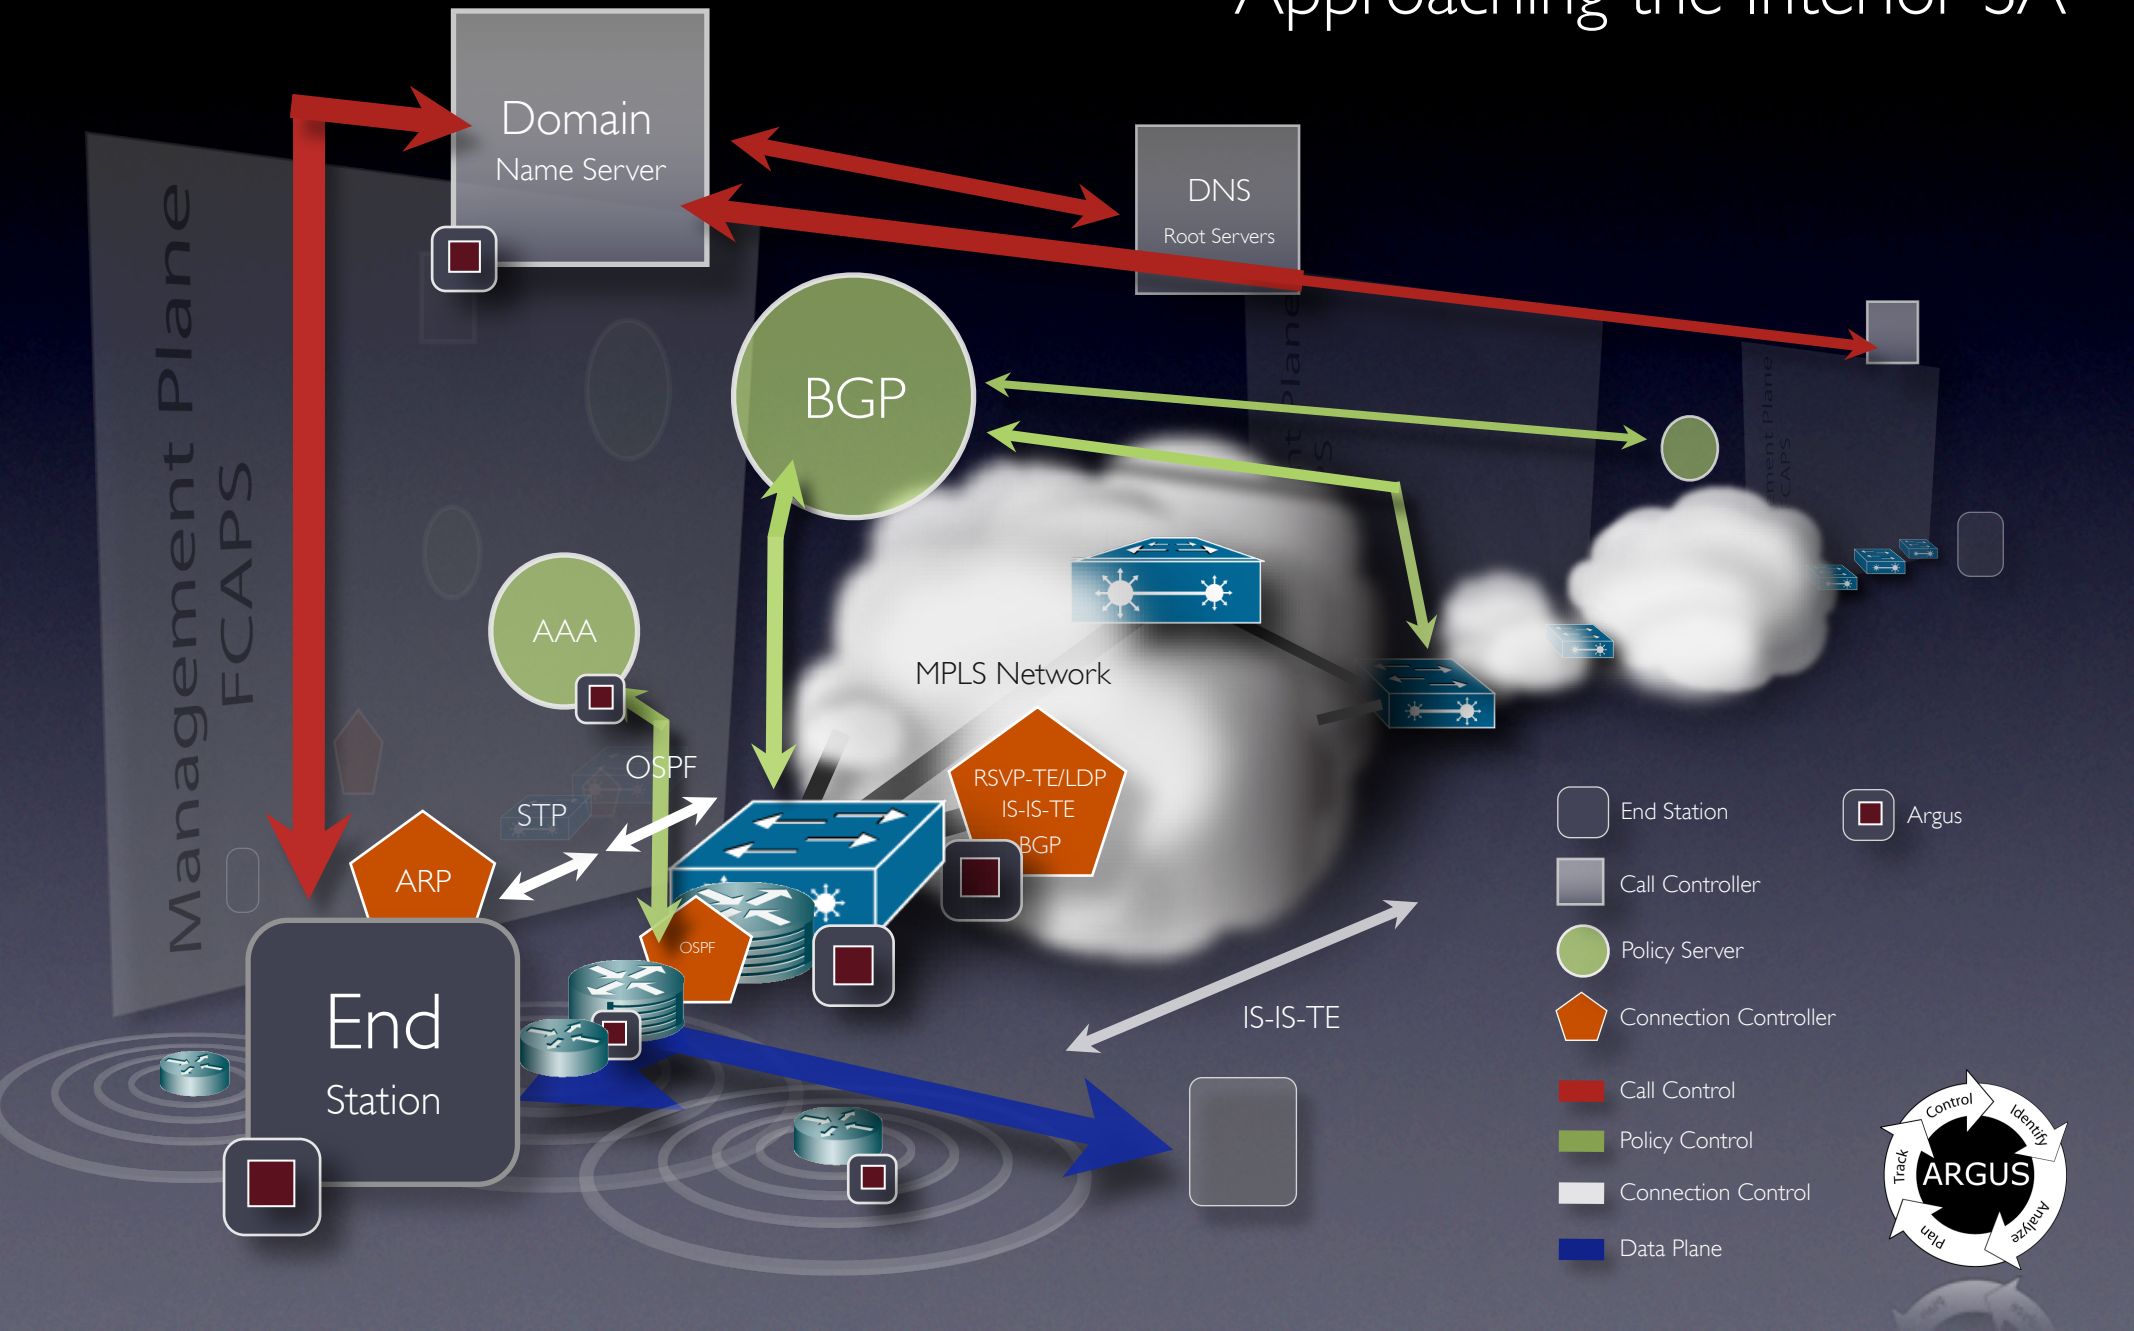

X.813 Non-Repudiation Entities

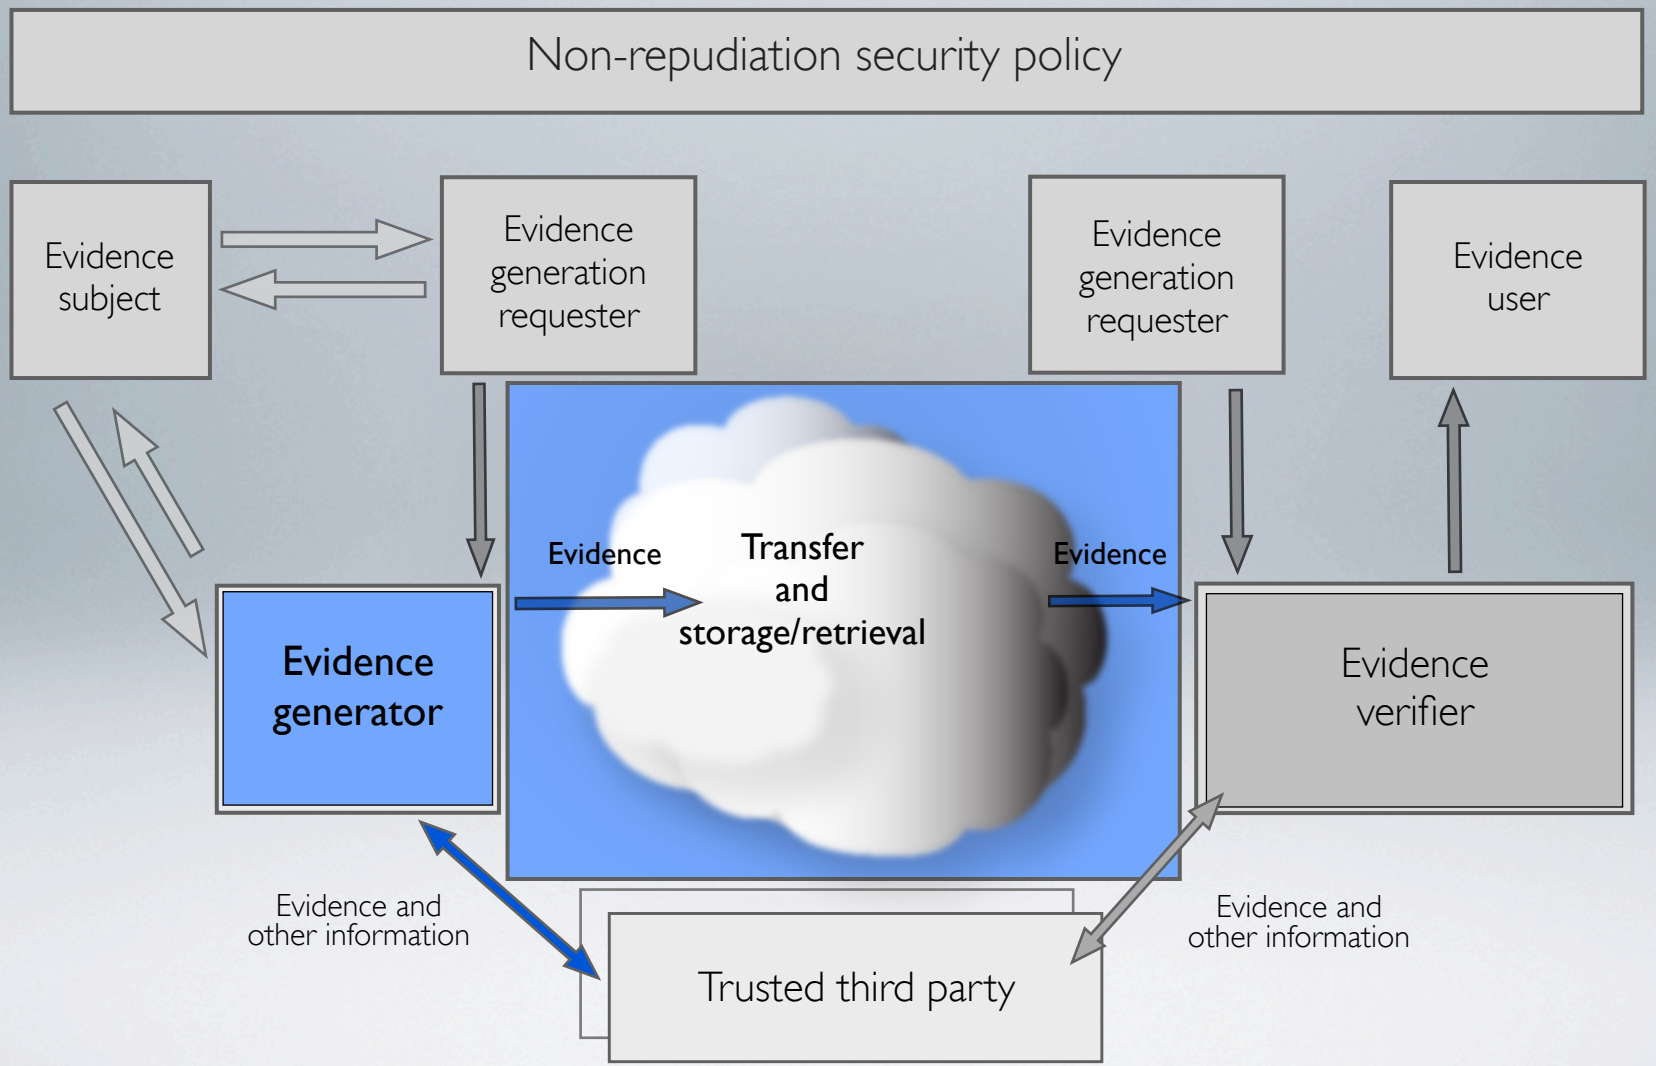

ITU-T Rec. X.813 (1996 E) TISO7760-96/D01

Data Generation and Collection

Argus System Design

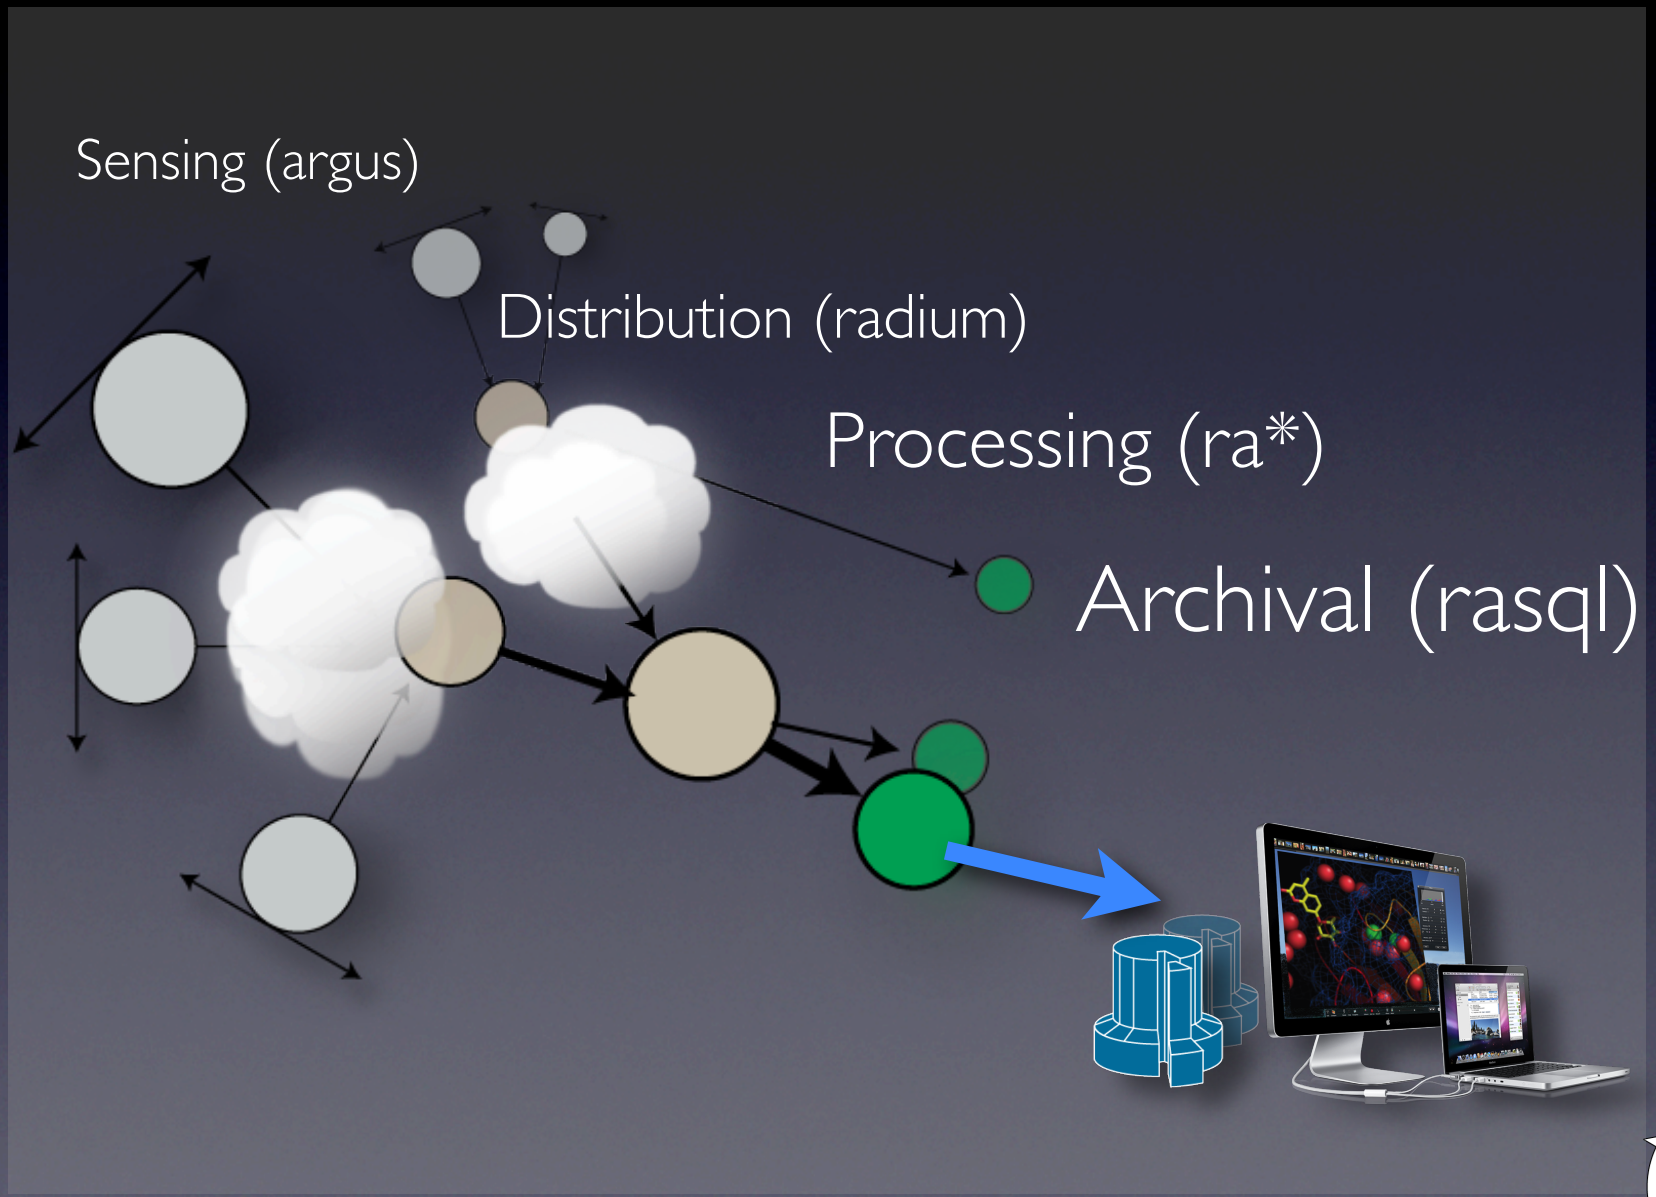

Argus Sensor Design

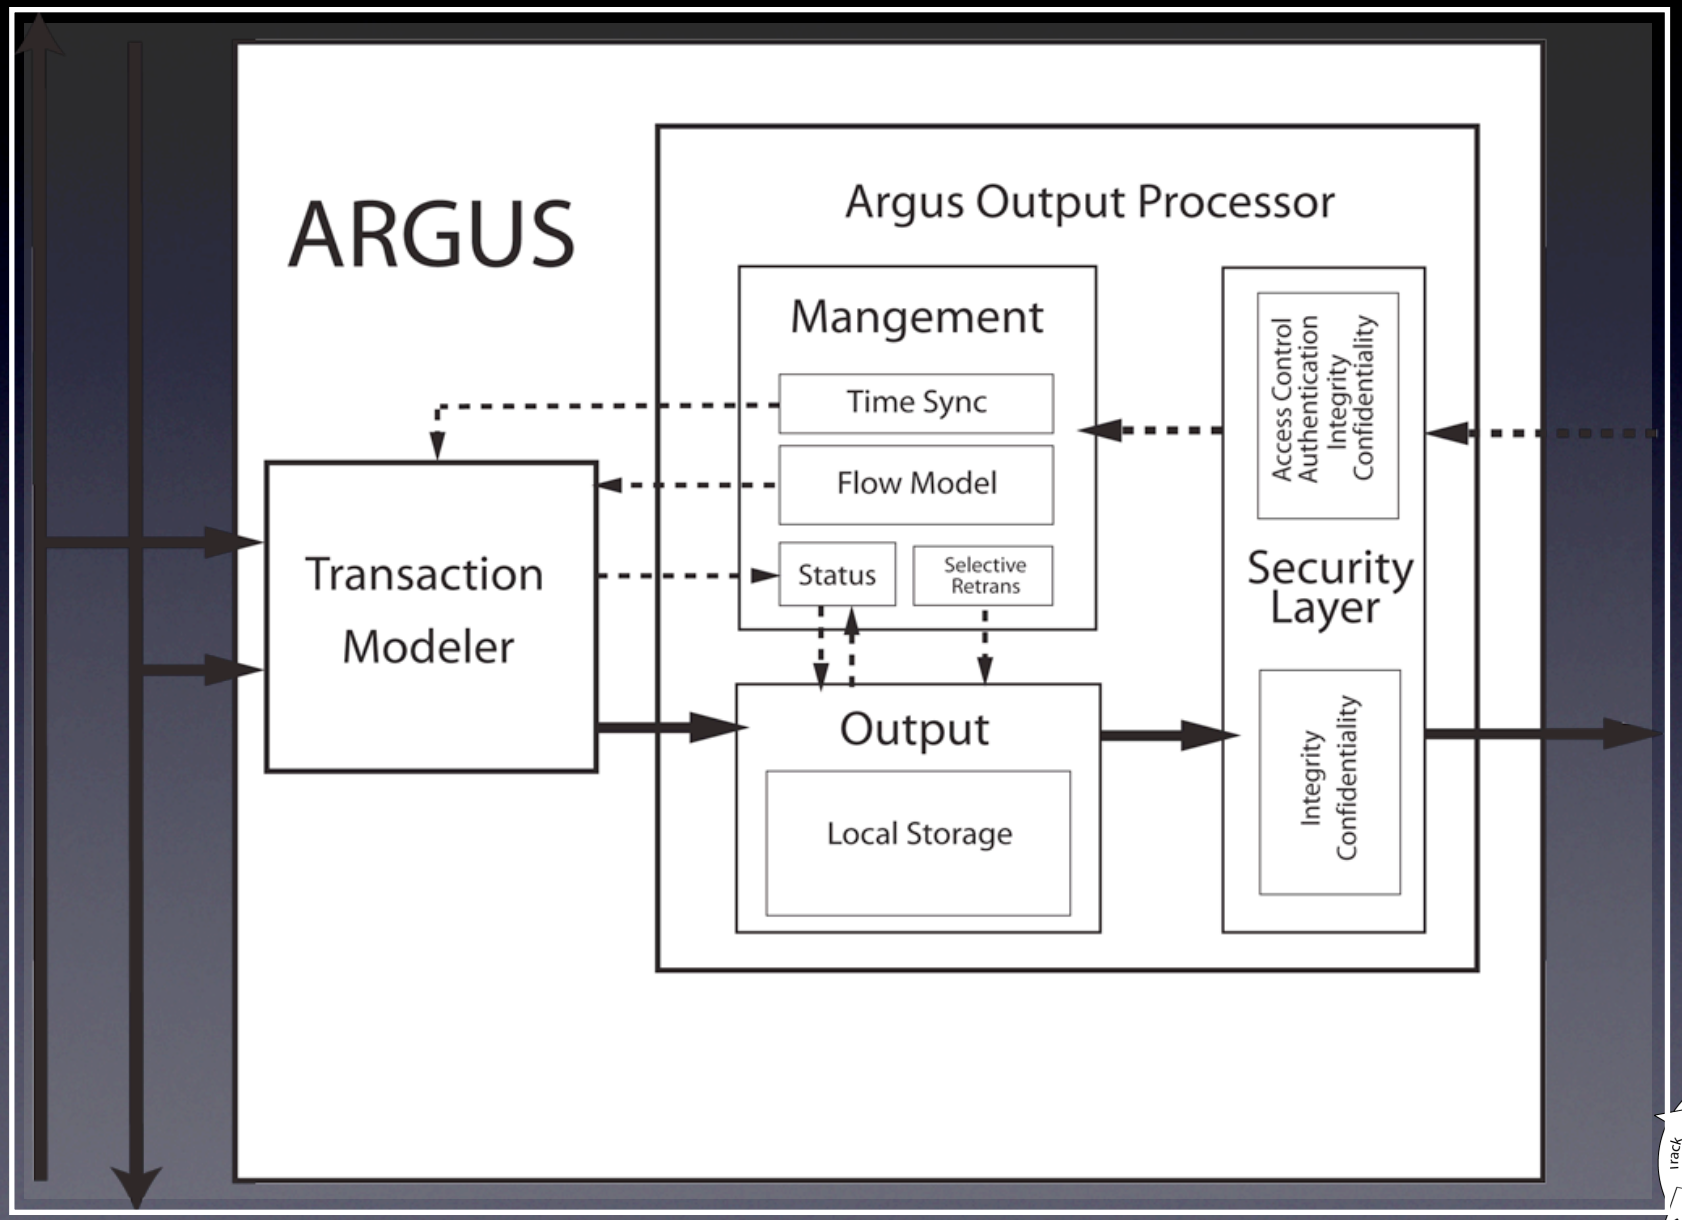

Argus Sensor Design

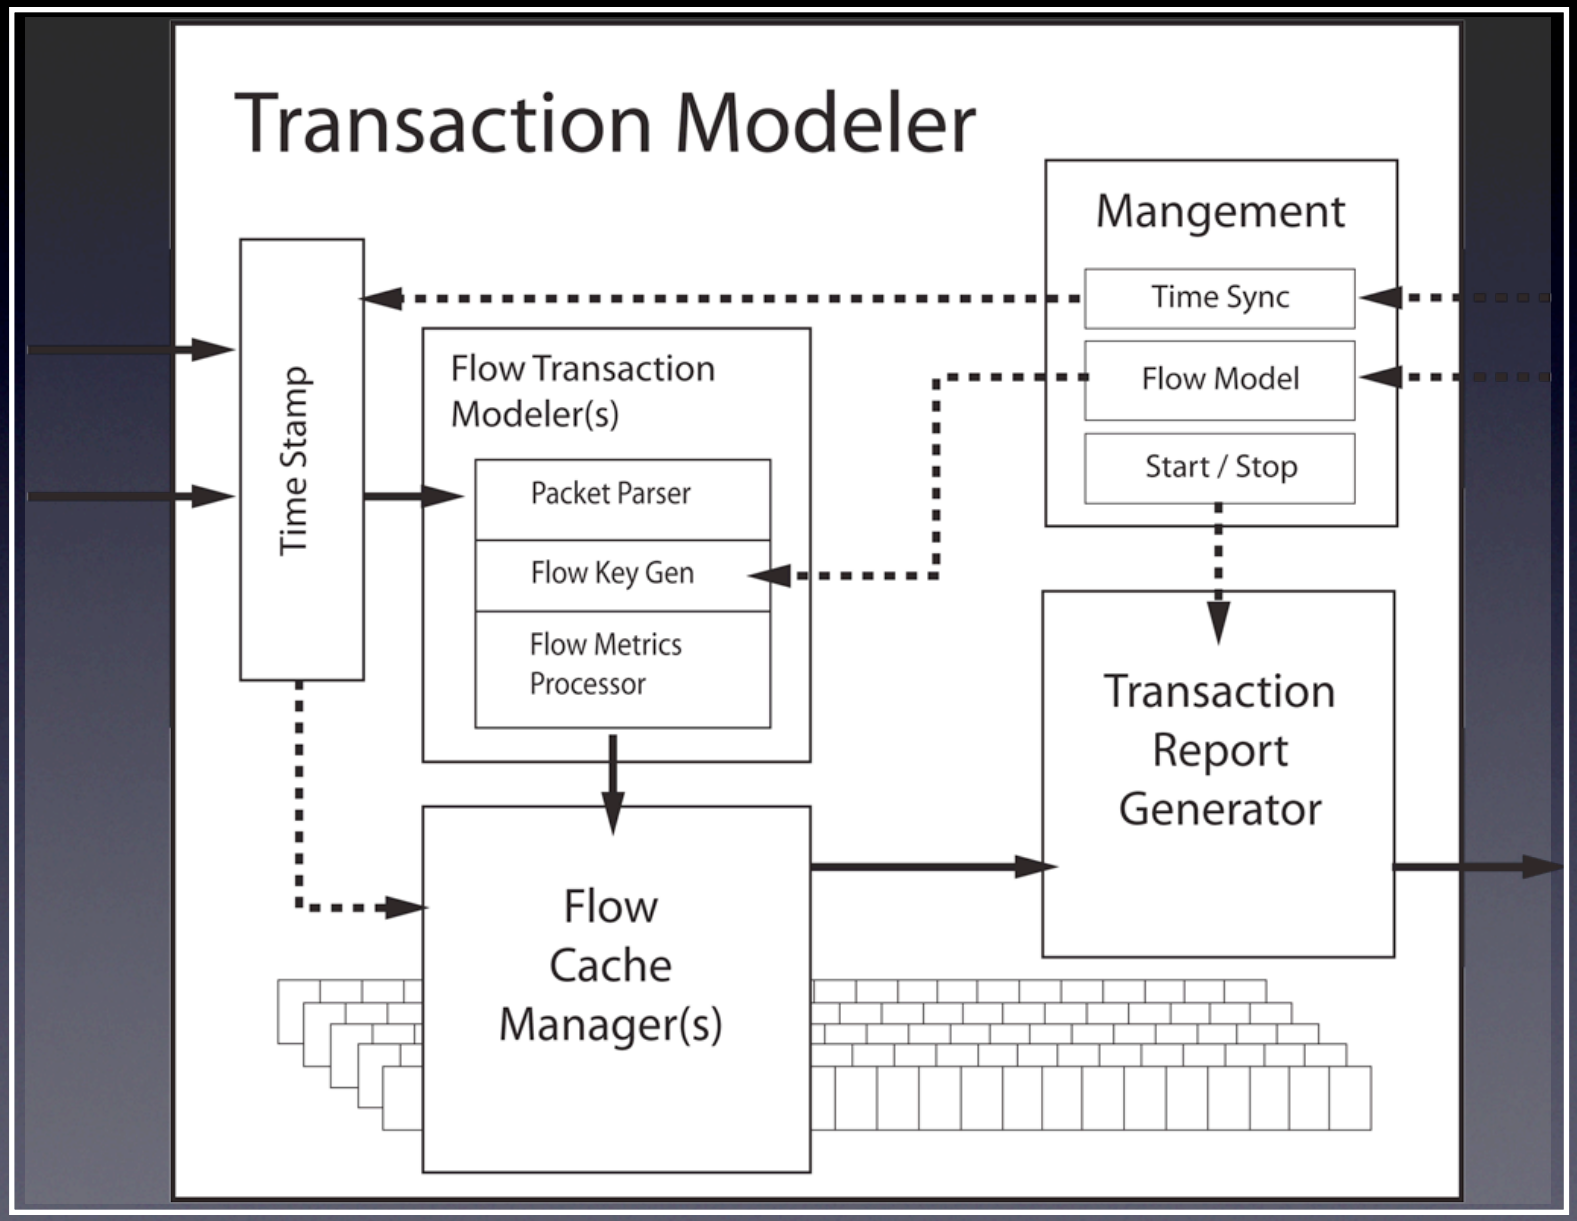

Radium Distribution

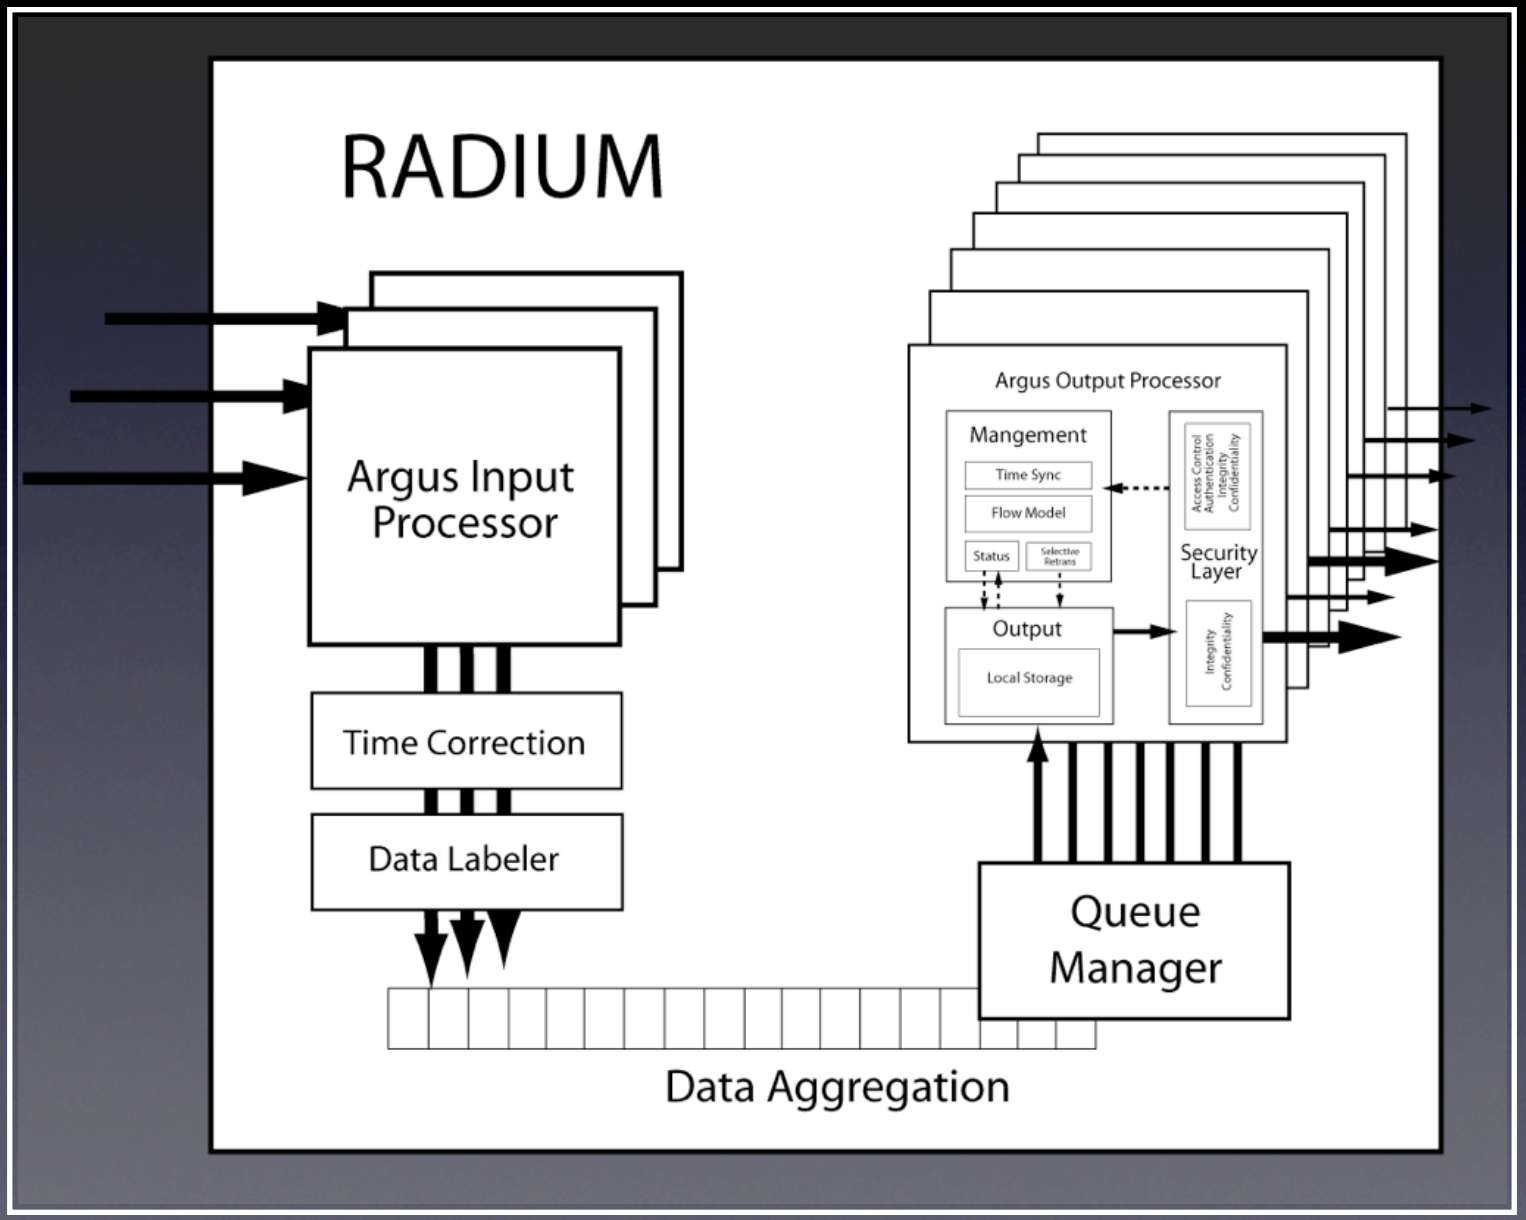

Argus Processing Design

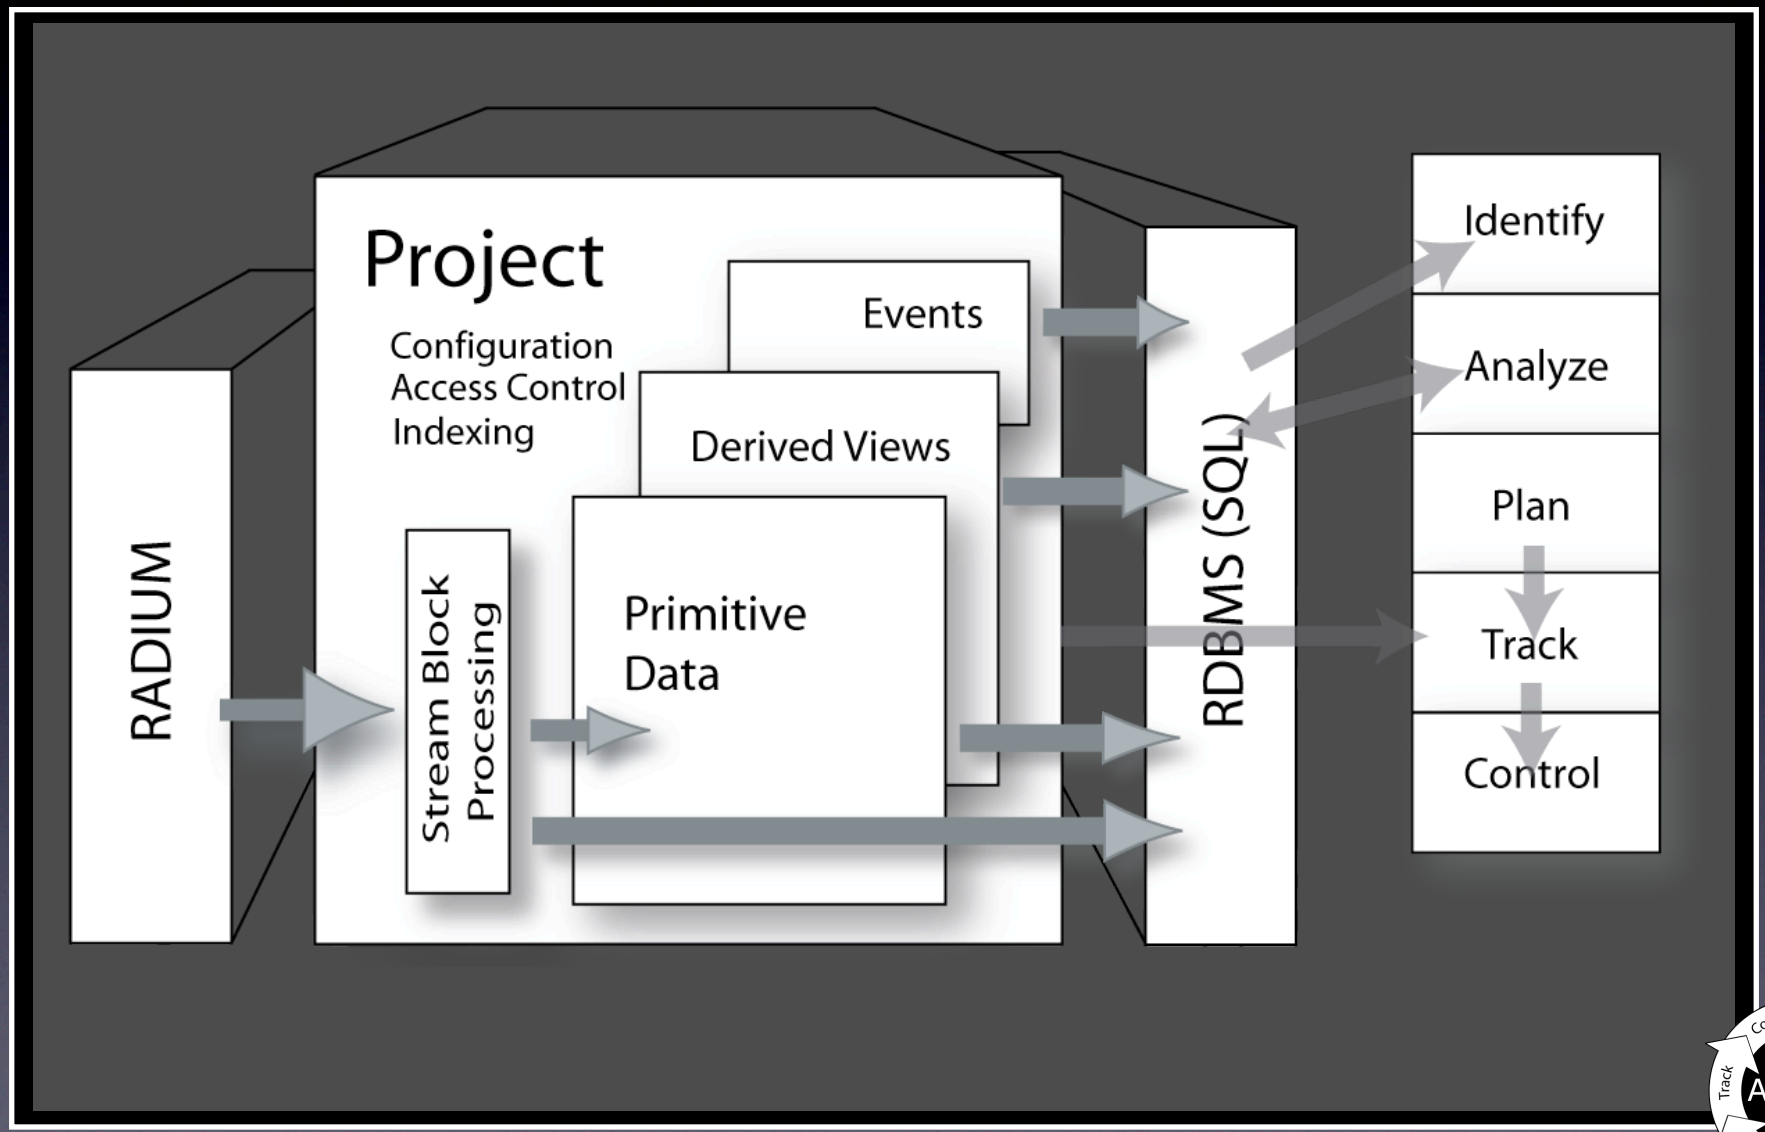

Stream Block Processor Design

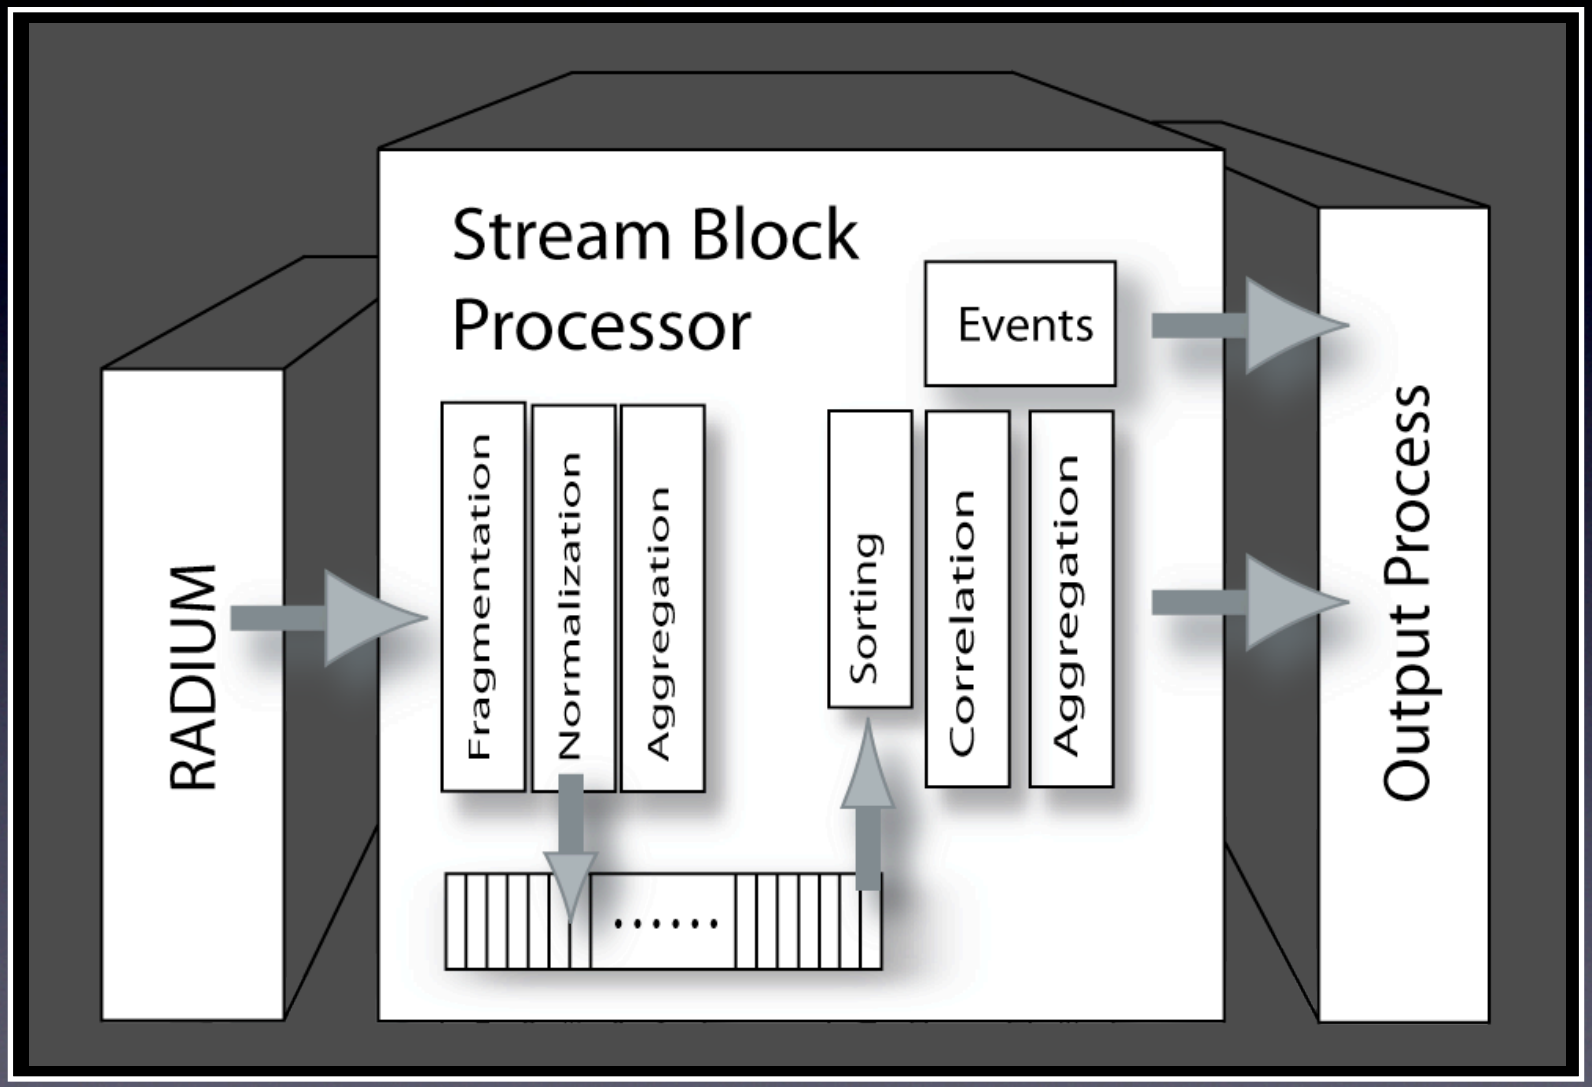

Data Collection

All ra* programs can read data from any Argus data source, files, stream, encrypted, and/or compressed, and can write current file structure.

Making an argus data repository needs just a little bit more.

- File Distribution
- Radium Distribution
- Argus Repository Establishment
  - cron
  - rasplit/rastream
  - rasqlinsert/rasql

Data Collection

Argus reading from packet files or network and writing directly to disk

Argus reading from the network and writing directly to network based client

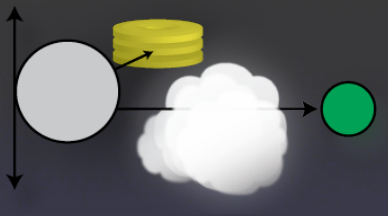

Argus reading from the network and writing directly to disk and network based client

Argus reading from the network and writing directly to a network Radium, writing to a client

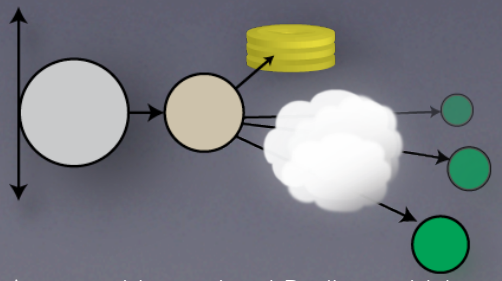

Argus writing to local Radium which is writing directly to disk and to network based clients

Many Argi writing directly to a Radium based distribution network, which is providing data to a set of clients.

Data Collection

- Local Generation and Storage
  - Basis for argus-2.0 argusarchive.sh
  - Direct argus support for renaming files
  - Normally cron mediated
  - Issues with time and record spans
  - System designer has most control !!!

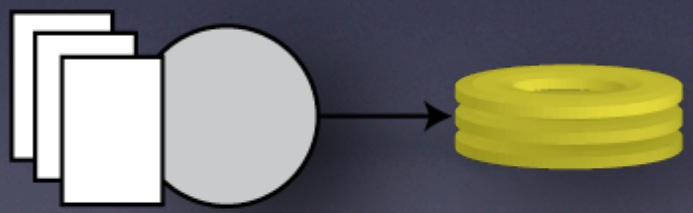

Argus reading from packet files or
network and writing directly to disk

Data Collection

- Local Generation Remote Collection
  - Most high performance systems use this strategy
    - Provides explicit scalability and performance capabilities
    - Relieves argus from physical device blocking
    - Network interfaces generally faster than local storage devices
  - Introduces network transport issues
    - Reliability, connection vs. connection-less, unicast vs multicast, congestion avoidance, access control and confidentiality

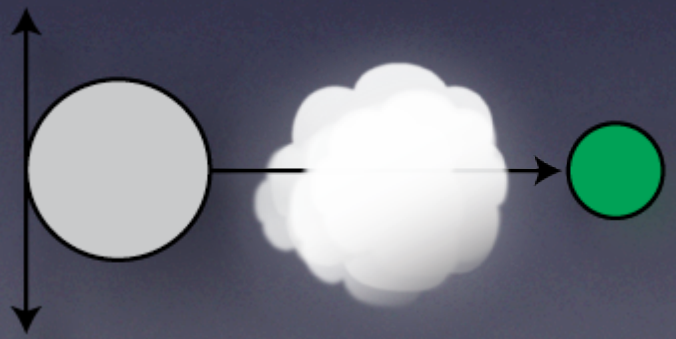

Argus reading from the network and writing directly to network based client

Data Collection

- Local Storage and Remote Collection
  - Used when data reliability is most critical
    - Local storage provides explicit data recovery
    - File collection provides additional distribution flexibility
    - Scheduled transport
  - Reduces ultimate sensor performance
    - Argus itself is doing a lot of work
    - Packet processing is really the ultimate limit

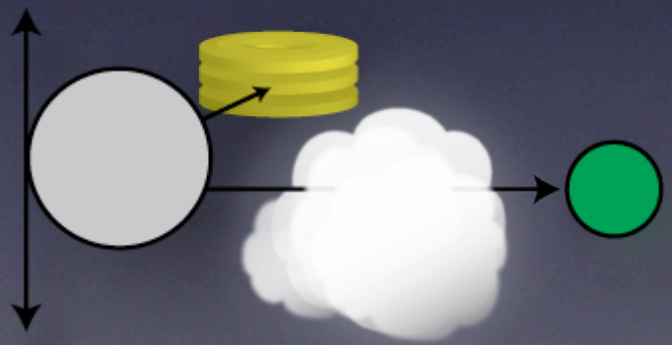

Argus reading from the network and writing directly to disk and network based client

Data Collection

- Radium
  - Primary argus data distribution technology
  - Radium is a ra* program with an argus output processor.
    - Read from many sources
    - Write to many clients
    - Serve up argus data files
    - Process/transform data
    - Configuration is combo of argus() and ra()
- Supports very complex data flow machine architectures.

Data Collection

- Local Generation Remote Distribution
  - Most prevalent strategy used in argus-3.0
    - Provides explicit scalability and performance capabilities
    - Provides most stable collection architecture from client perspective
    - Single point of attachment for complete enterprise
  - Least reliable of 'advanced' strategies
    - Radium failure interrupts continuous stream collection, with no opportunity for recovery

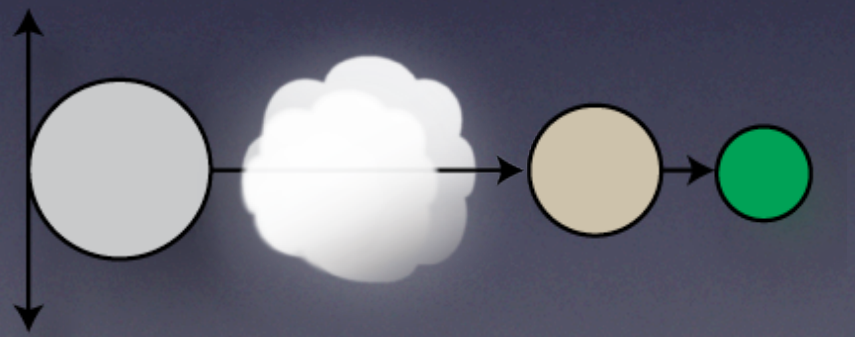

Argus reading from the network and writing directly to a network Radium, writing to a client

Data Collection

- Local Distribution and Storage
  - Best methodology
    - Provides explicit scalability and performance capabilities
    - Provides most reliable collection architecture
    - Multiple points of attachment, multiple points of control
  - Most expensive strategy at data generation
    - Radium deals with device and remote client requests for data which does come with a processor and memory cost

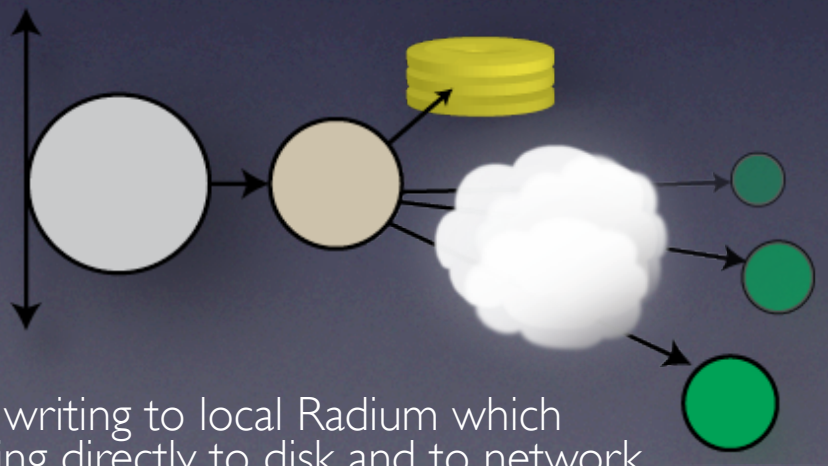

Argus writing to local Radium which is writing directly to disk and to network based clients

Data Collection

- Complex data flow machine architectures
  - Architecture of choice for scalability
    - Provides explicit scalability and performance capabilities
    - Provides most parallelism
    - Multiple points of attachment, multiple points of control
  - Can get a little complex
    - Merging of multiple flows, multiple times, introduces complex data duplication issues, and allows for complex, incompatible data schemas to co-exist

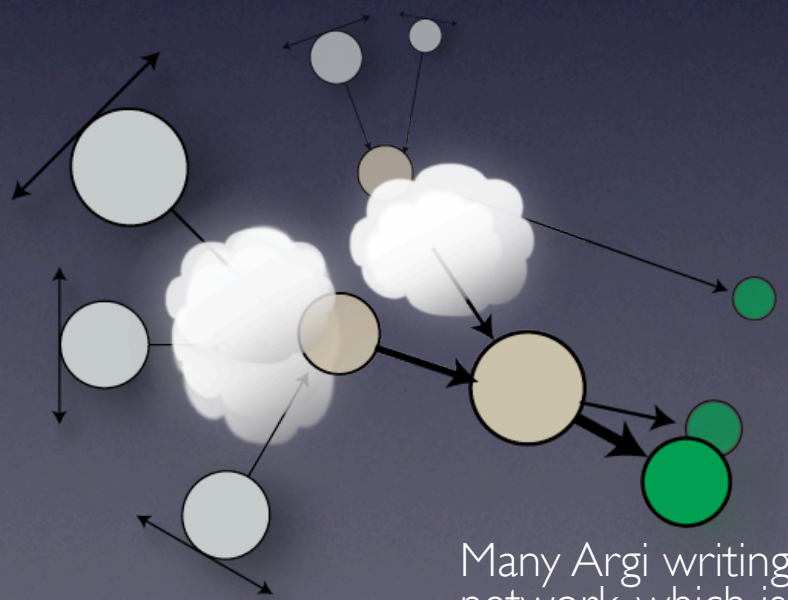

Many Argi writing directly to a Radium based distribution network, which is providing data to a set of clients.

Comprehensive Network Monitoring

Network Activity

Black Core Mesh

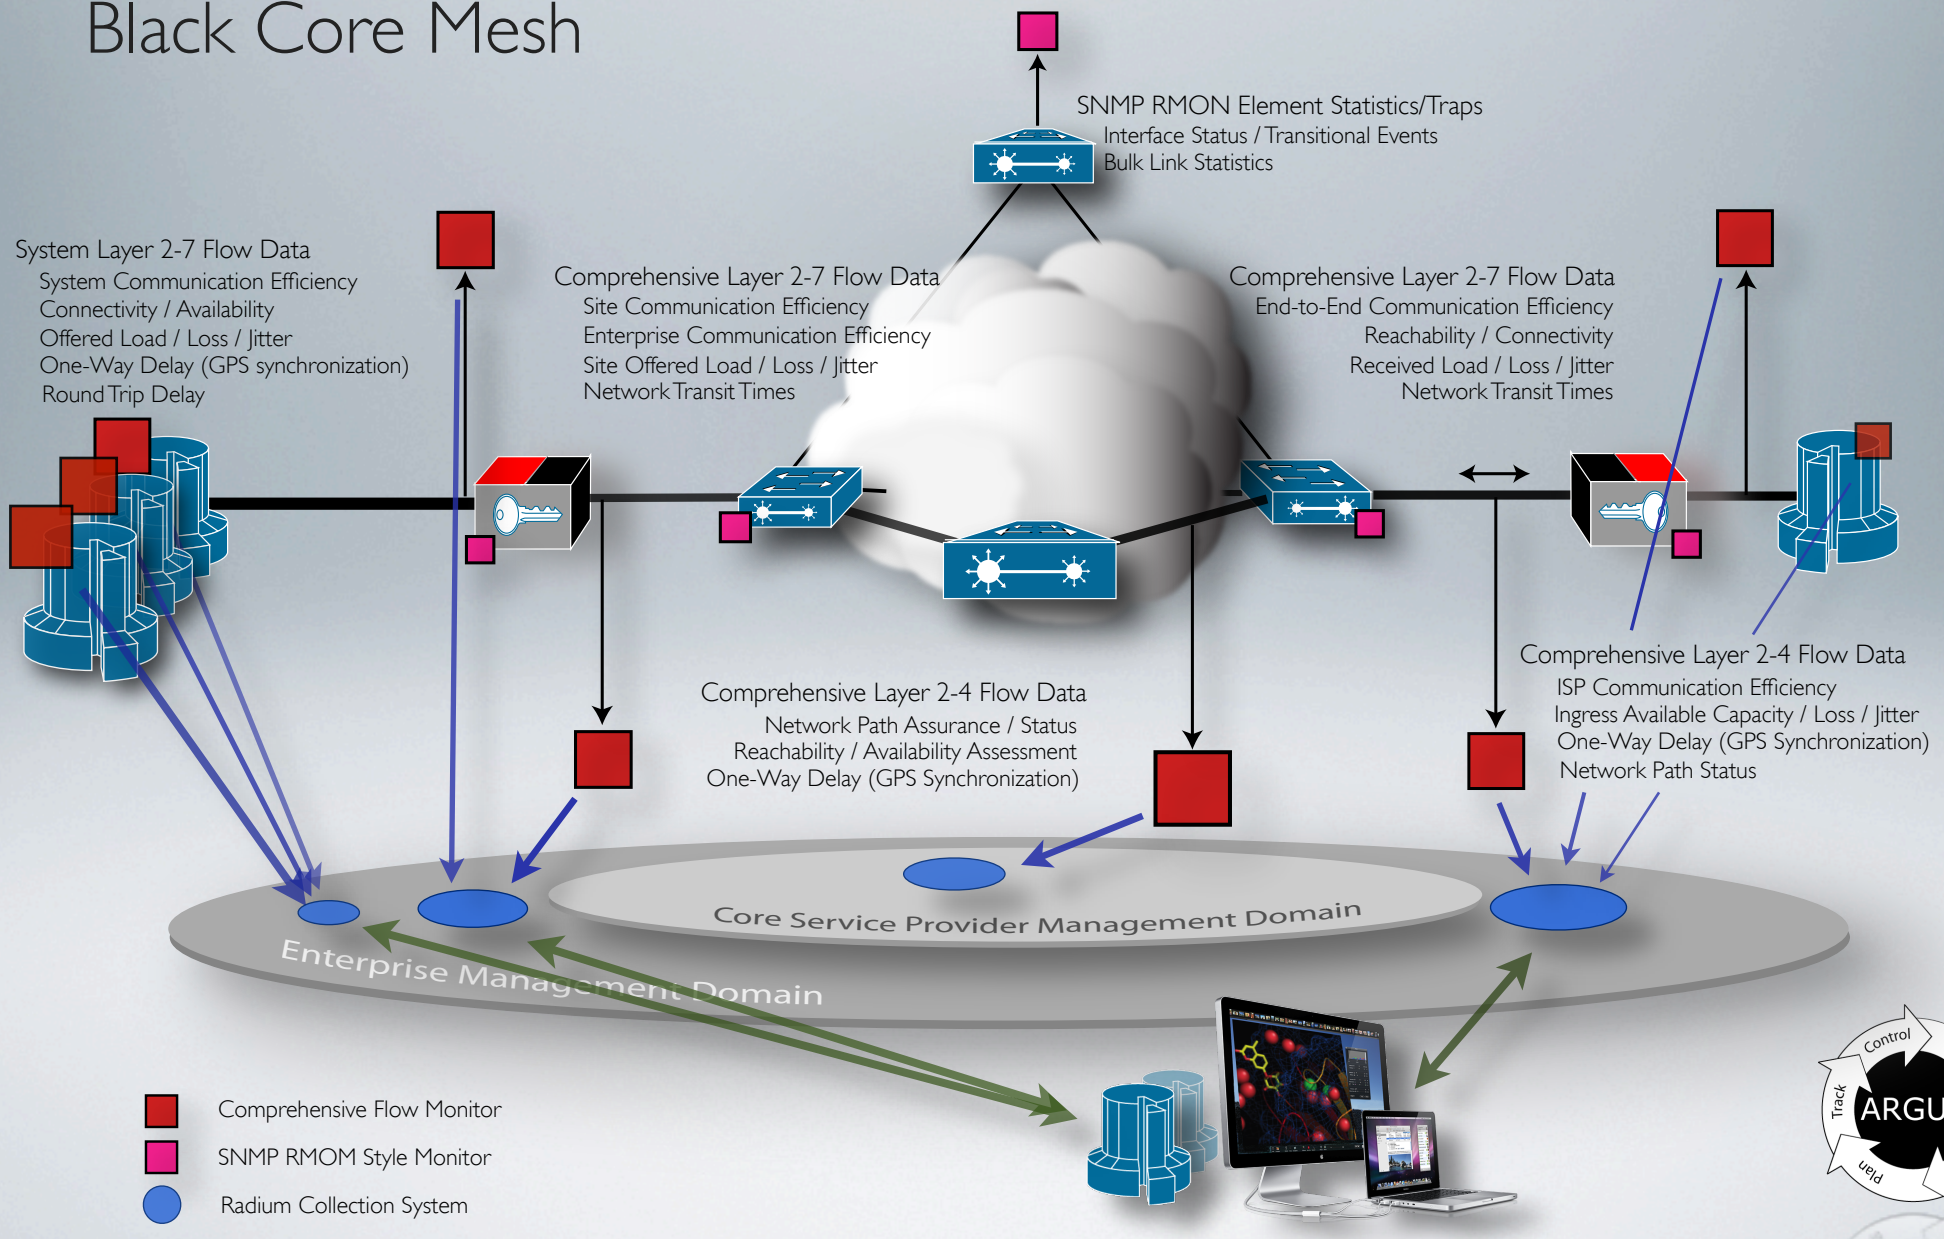

Argus Repositories

- Argus Repository Establishment
  - Formal Ingest/Disposition
- Repository Function
  - Primitive Data Repository
    - General Archive
    - Access Control
    - Retention Policies
    - Modification Policy (Compression)
  - Derived Data Repositories

Argus Repositories

- Native File System
  - Simplicity
  - Performance
  - Compatibility
- Relational Database System (RDBMS)
  - Extensive Data Handling Capabilities
  - Complex Management Strategies
  - Performance Issues


Control
Identify
Analyze
Plan
Track
ARGUS

US Cyber Security Focus

- Comprehensive National CyberSecurity Initiative
  - Shifting the US focus from CyberCrime to CyberWarfare
- Strategy and technology focused on new issues
  - Public sector defense, with nation state threats and countermeasures
  - New emphasis on military concepts in Cyber Security
    - Shift from detection to prevention
    - Possible retaliatory mechanisms
- Multi-billion dollar budget will have a significant impact
  - Redefine CyberSecurity for most of the public
  - Compete for best/brightest in security research
  - Determine a new direction for commercial security products

US Cyber Strategy Issues

- Cyber Crime still represents 99% of the cyber problem
- Change in focus may create strategies and technologies that are inappropriate for addressing Cyber Crime.
  - Example: many CNCI initiatives involve enhanced monitoring
    - To support advanced intrusion detection and prevention.
    - Sharing of network monitoring data for enhanced situational awareness.
    - In the pubic sector's .gov, .mil and classified networks, where there is no expectation of privacy. Enhanced monitoring is a very good thing.
    - In the private sector, however, any level of enhanced monitoring is perceived by the public as wiretapping.
- Can the CNCI produce a surveillance strategy that represents an acceptable privacy strategy?
- An old public-private partnership may be able to help

Distributed Situational Awareness

Multi-Probe Multi-Site

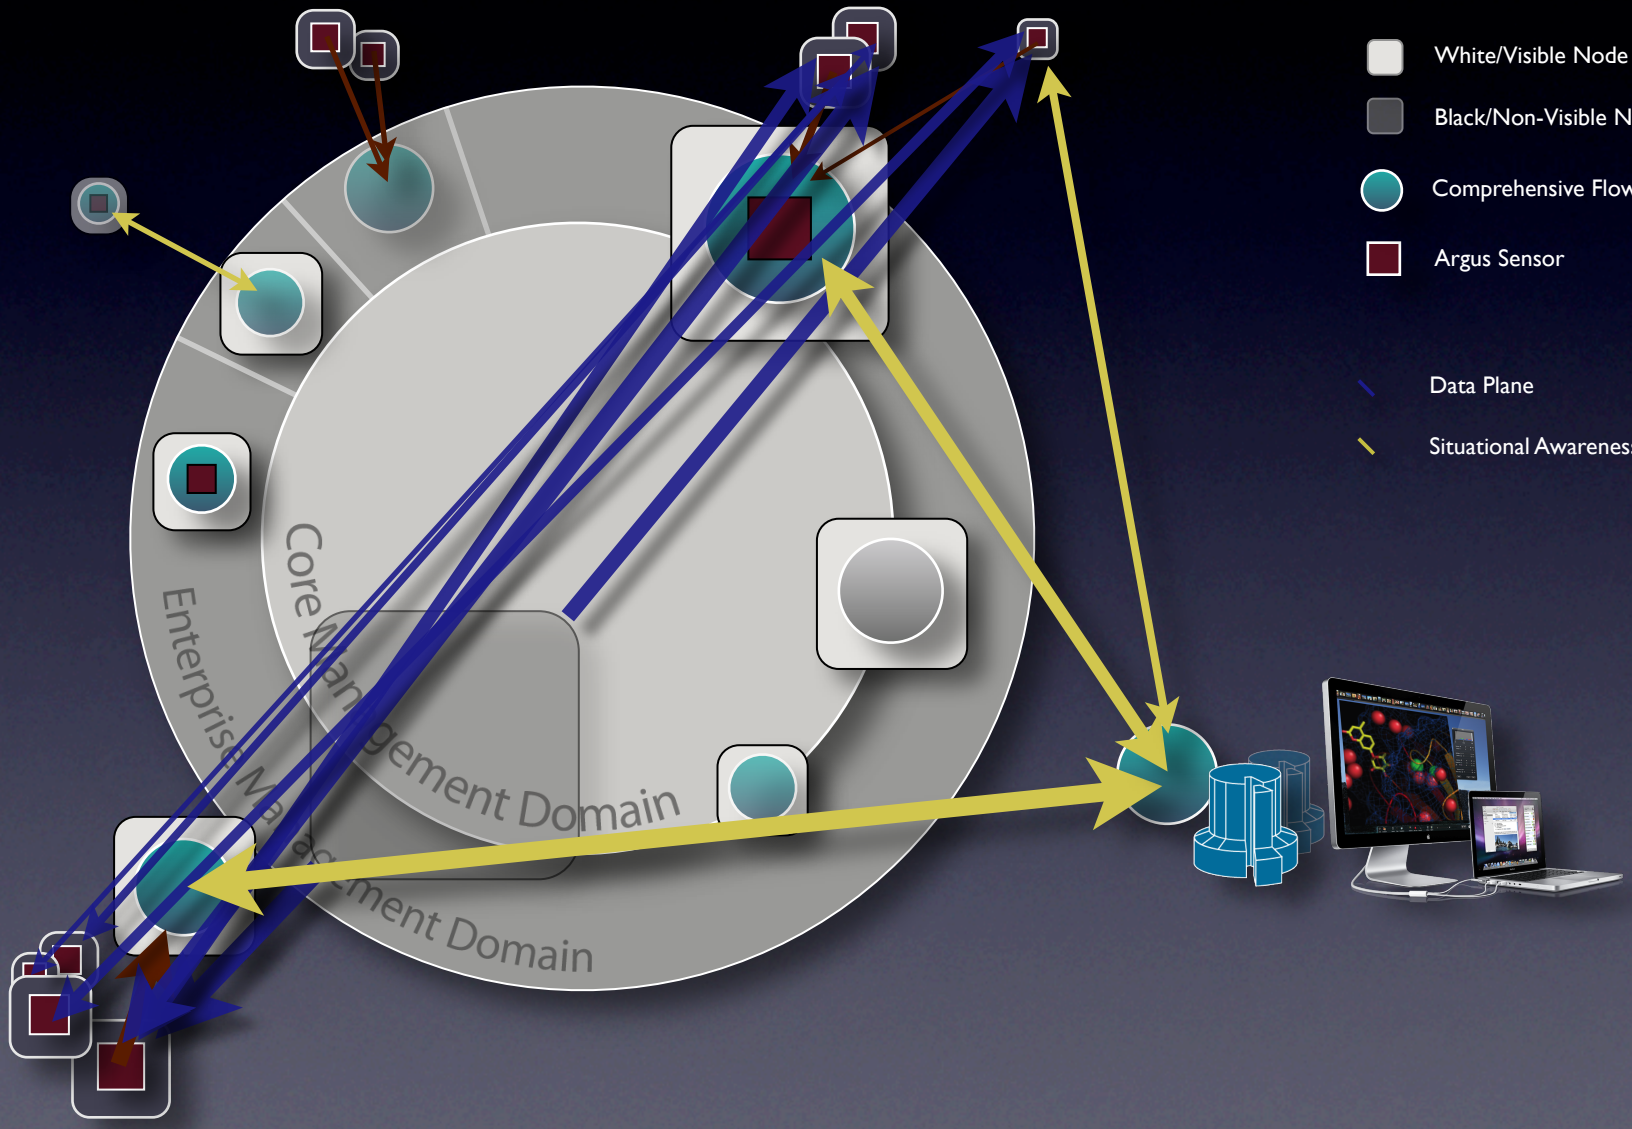

Private-Public Partnership

- 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 IP flow audit 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.

New Public-Private Partnership?

- The private sector is generating and collecting its own IP network flow data for most of the same reasons that the PSTN processes CDRs.
- Society has learned how to effectively use IP network flow data for its benefit, giving up some aspects of privacy in order to achieve a higher level of general privacy protection through minimizing Lawful Intercept.
- The private sector actively contributes to national Cyber Security through controlled sharing of its own network session data.
- Adoption of this public-private partnership enables a historically recognizable deterrence to crime.

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.

Questions?

- For more information please visit

  http://qosient.com/argus

- Contact me directly via email

  carter@qosient.com

Thanks for your interest in Argus

Backup

Next Generation Architectures

3GPP / IMS / TISPAN

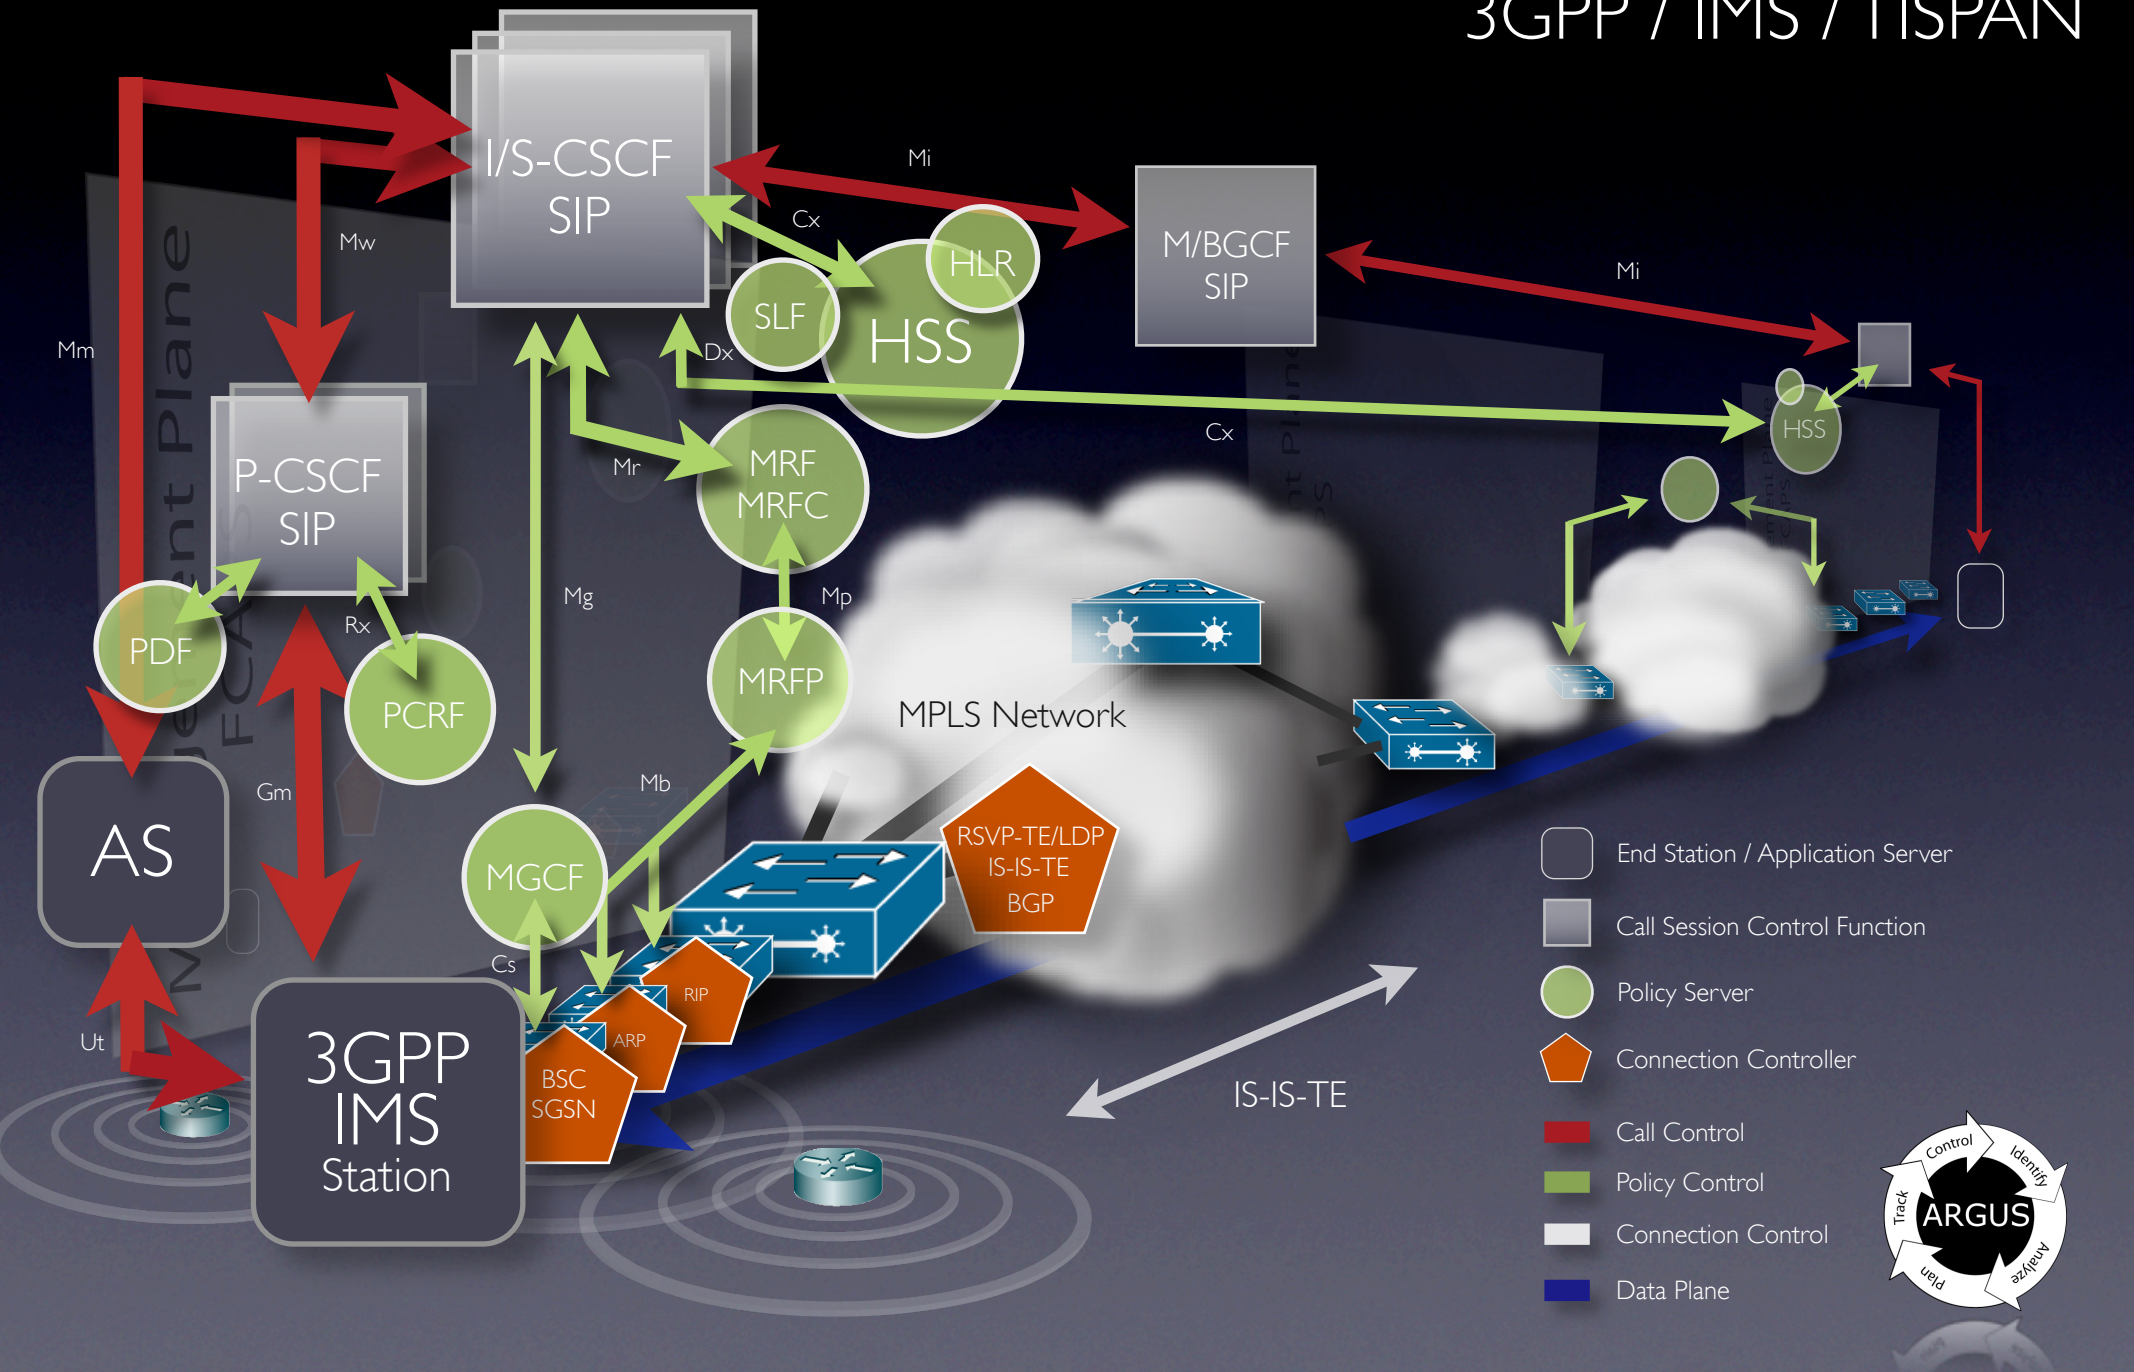

Argus Approach

- Use formal network activity audit models
  - Network service oriented
    - Initiation, status, termination state indications for complex flows
    - Bidirectional realtime network flow traffic monitoring
    - For all services - ARP, DHCP, DNS, TCP, UDP, VoIP, P2P....
  - Convey as much about the traffic as possible
    - Deep packet inspection strategies to extract traffic semantics
    - Tunnel identifiers (MPLS, VLAN, Ethernet, IPnIP, GRE, RTP, Teredo)
    - Reachability, connectivity, availability, rate, load, latency, loss, jitter, packet size distribution
    - Security issue reporting (protocol issues i.e. fragmentation overlap attack, tunnel hopping)
    - Controlled content capture
- Flexible transport, collection and storage strategies
- Data processing tools
  - Aggregation, analysis, anonymization, filtering, graphing, … , zipping.
  - Native OS archive management, MySQL database, 3rd party integration
- Realizable audit data resource requirements
  - CMU generates ~100-300 GB/day of 'primitive' data
  - Naval Research Lab retains ~ 50-100 GB/day
  - QoSient.com manages 120 MB/day

Network Situational Awareness

- Argus is designed to be THE network SA sensor
  - Ubiquitously deployable DPI traffic sensor
  - Comprehensive (non-statistical) traffic awareness
  - Provides engineering data, not business intelligence
    - Detailed network transactional performance
    - Network fault identification, discrimination and mitigation
      - Reachability, connectivity, availability, latency, path, flow control etc....
    - Customer gets the primitive data, not just reports/alerts
  - Near realtime and historical capabilities
  - Packet capture replacement
- Supporting a large number of SA applications
  - Advanced Network Functional Assurance (Operations)
    - End-to-End transactional performance tracking (data and control plane)
    - Network component functional assurance (NAT, reachability, encryption)
    - Policy enforcement verification/validation (Access control, path, QoS)
  - Advanced Network Optimization (Security and Performance)
    - Supports network entity and service identification, analysis, planning tracking and control, including baselining, anomaly detection, behavioral analysis and exhaustive forensics

Benefits of Argus

- Argus provides excellent data
  - Data drives many applications
    - Advanced network activity metrics
    - Rich security information model
  - Real time access
- Deployable throughout the infrastructure
  - High Performance - 10 Gbps using Endace, Bivio, etc....
  - End Systems - Unix, Linux, Mac OS X, Windows (Cygwin), AIX, IRIX, HPUX, Ericsson ViPR
  - Router Systems - VxWorks, DIY Routers, OpenWRT
- Rich data collection, management and processing
  - Key differentiator from commercial offerings
    - Many advanced sites want their own data analysis
    - All want data processing and reporting extensibility

Where are we headed?

- Distributed Network Auditing
  - Very Large Scale Situational Awareness
    - Auditing system scalability using cloud architectures
  - Complete end-to-end capability
  - Automated Attribution
  - New security mechanisms
- Sensor Improvements
  - Higher performance - multi-core
  - More Control Plane Auditing
    - OSPF, BGP, , SIP ...
  - Wireless
- Audit system applications
  - Real-time situational awareness
  - Security forensics tools

Carnegie Mellon®

Flow Monitoring Infrastructure

- Argus is the predominant tool for network flow monitoring/policy enforcement
- Probes at key points on network
  - Border
  - Core
  - Wireless network
  - Ad-hoc on edge routers (moved as necessary)

Carnegie Mellon®

Flow Monitoring Infrastructure

- Success stories:
  - Forensic examination of compromised machine traffic
    - Determining size and scope
    - Correlating with other events
  - Auditing correct router ACLs
    - Examine real time flows on both sides of the router
  - User consultations regarding bandwidth usage
    - Reports of machine traffic can be generated
  - Configuration issues with VPN infrastructure
    - Examining flows identified source of problem