Raytheon

System of Systems to Provide Quality of Service Monitoring, Management and Response in Cloud Computing Environments

July 16-19, 2012

Paul C. Hershey [1]

Shrisha Rao [2]

Charles B. Silio, Jr. [3]

Akshay Narayan [2]

[1] Raytheon, Intelligence and Information Systems

[2] International Institute of Information Technology Bangalore

[3] University of Maryland, College Park

Raytheon

Agenda

- Problem: Maintain QoS in Presence of Data Overload and Economic Downward Pressure
- Previous Approaches - Issues with Cloud Computing in Complex Systems
- Solution: Apply New 5-Step Procedure to Cloud Computing to Complex System of Systems
- System Model: Mathematical Model for Quality of Service Metrics (Performance, Authentication, Authorization)
- Application Scenario: Distributed Denial of Service Attack on Complex Systems
- Results: Delay, Variation in Delay, and Throughput Performance Metrics Verification
- Conclusions, Present Status, and Path Forward

Problem: Maintain QoS in Presence of Data Overload and Economic Downward Pressure

Raytheon

- Capacity: Dramatic increase in the quantity of data transmitted over DoD, government, and commercial networks threaten QoS
  - Data overload created by evolution of complex, net-centric enterprise systems over which multiple disparate users in dispersed locations share petabytes of data at high speeds
- Economic: Decreasing budgets require a solution beyond increasing processing and bandwidth resources.
  - Sharing resources, as achievable through cloud computing, offers possible solution

"We're going to find ourselves in the not too distant future swimming in sensors and drowning in data,"

Lt. Gen. David A. Deptula, Keynote Address, GEOINT 2009, Oct. 2009.

Capacity and Economic Issues Point to Cloud Computing as Solution

Previous Approaches - Issues with Cloud Computing in Complex Systems

Raytheon

Complex computing systems that use cloud computing are prone to failure and security compromise in five main areas.

1. **Computing Performance**
   - e.g., latency, time delay experienced by a system when processing a request
2. **Cloud Reliability**
   - e.g., network connectivity
3. **Economic Goals**
   - e.g., interoperability between Cloud Providers
4. **Compliance**
   - e.g., digital forensics to discern what happened, learn how to prevent incident, and collect information for future actions
5. **Information Security**
   - e.g., protect the confidentiality and integrity of data and ensure data availability

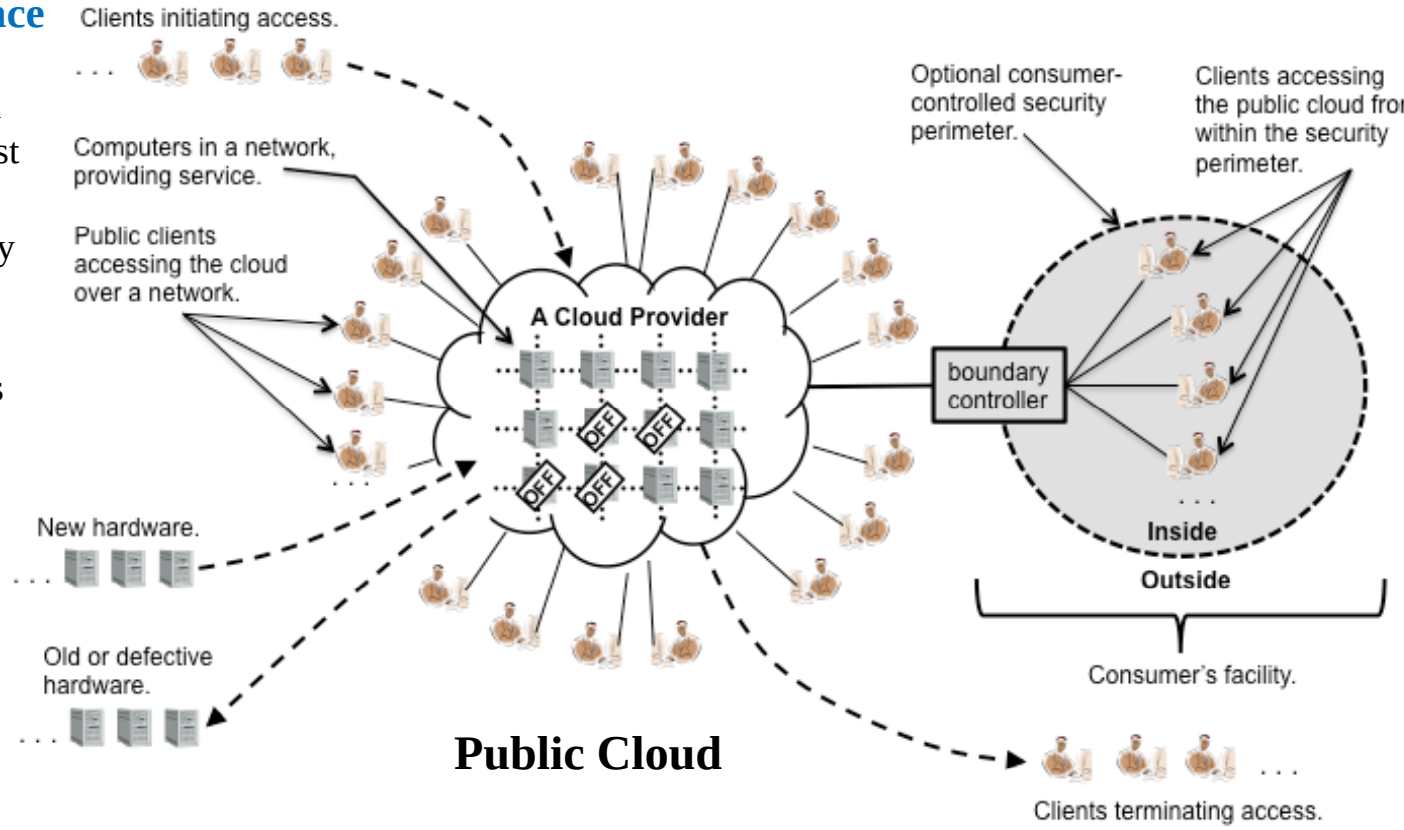

* P. Mell and T. Grance, *The NIST Definition of Cloud Computing. National Institute of Standards and Technology (NIST), US* Dept. of Commerce, Sep. 2011, NIST Special Publication 800-145, http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf.

Previous Approaches are Prone to Failure and Security Compromise

Solution: Apply New 5-Step Procedure to Cloud Computing to Complex System of Systems (SoS)

Raytheon

Designed to overcome the limitations of previous approaches

Step 1: Define a SoS for monitoring, management, and response.

Step 2: Derive framework for Quality of Service (QoS) monitoring, management and response in cloud computing environments.

Step 3: Identify cloud computing metrics.

Step 4: Identify suitable locations within the cloud computing environment for observing and collecting metrics.

Step 5: Identify potential implementation schemes from which to collect and analyze the cloud computing QoS metrics.

New Solution Addresses Performance and Security Deficiencies

Step 1: *Define a SoS for Monitoring, Management, and Response*

Raytheon

- SoS characteristics effective QoS monitoring, management, and response to overcome cloud computing deficiencies
  - Structure
    - Computing Performance
    - Information Security
  - Coupling
    - Cloud Reliability
  - Behavioral
    - Compliance
  - Interoperability
    - Economic

Structural	A SoS has a structure that comprises interdependent systems that integrate to form a higher order system, usually resulting in a hierarchy. This hierarchy can include monitoring and response at the highest-level system down to the smallest sub-component system (i.e., bit-level).
Coupling	The systems that comprise a SoS include coupling with respect to such areas as data, information, functions, state, and algorithm. A loss of any portion of the SoS will degrade the overall performance or capabilities of the higher order system; therefore, the systems are interdependent.
Behavioral	Integration of decisions and actions of systems occurs in the higher order system through governance in contrast to non-SoS where the sharing of information is the basis for collaboration.
Inter-operable	Systems that comprise a SoS interface with one another and interoperate by design in contrast to non-SoS where systems are not designed to do so.

A SoS is Well Suited for Application of Cloud Computing to Complex Systems

Step 1: Representative SoS for Monitoring, Management, and Response

Raytheon

- All domains operate within a Service Oriented Architecture
- Single authority provides Governance as a Service (GaaS) to multiple heterogeneous administrative domains & enables business & collaboration services
- Business as a Service (BaaS) enables end-users who are producing and consuming data using Software as a Service (SaaS) and Infrastructure as a Service (IaaS)

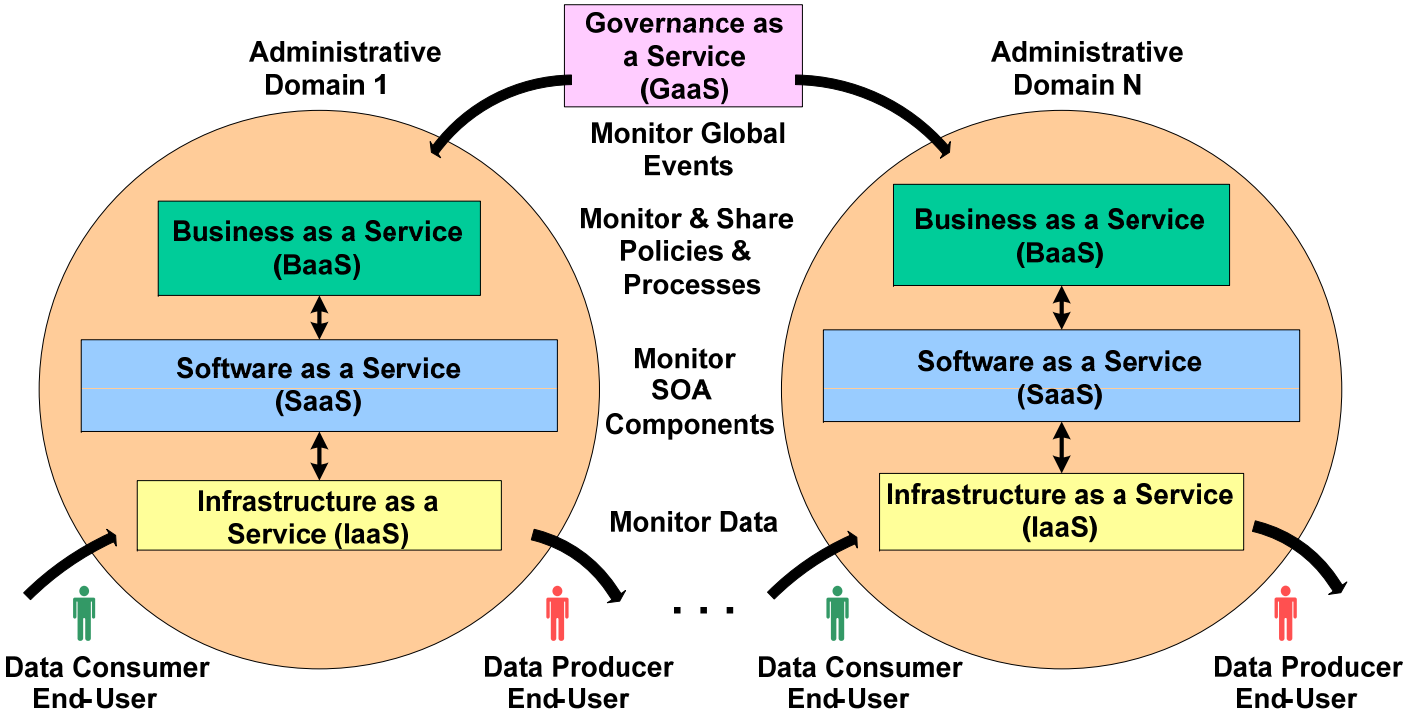

Representative SoS Includes IaaS, SaaS, BaaS, and GaaS

Step 2. *Derive Framework for Cloud Computing Environment QoS Monitoring, Management & Response*

Raytheon

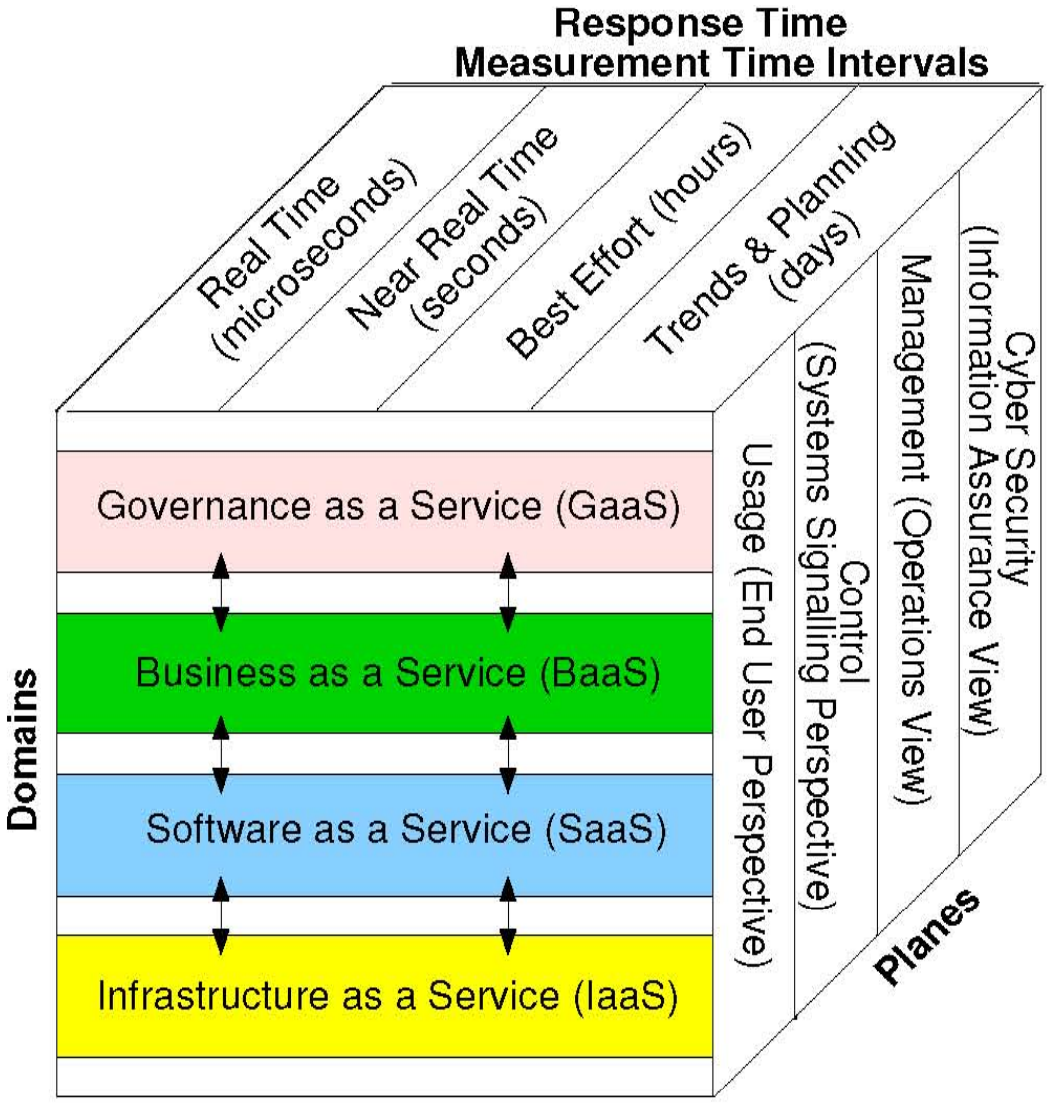

- Enterprise Monitoring, Management, and Response Architecture for Cloud Computing (EMMRA CC)
  - Detect and respond to CC events at enterprise-level
  - Applicable to data and voice
- Response Time
  - Services-based requirements
- EMMRA CC Domains
  - Similar techniques to monitor & manage in CC environment
- EMMRA CC Planes
  - Across-domain view of CC events

Multi-dimensional Reference Architecture Provides Broad Enterprise Coverage

Step 3. *Identify Metrics for Performance and Information Security*

Raytheon

Category	Metric
Performance	Delay
	Delay Variation
	Throughput
	Information Overhead
Security	Authentication
	Authorization
	Non-repudiation
	Integrity
	Information Availability
	Certification & Accreditation
	Physical Security

Key: Focus of This Paper

- Use standardized metrics for DDoS detection
  - Voice and Data
  - Enable sharing across informational domain boundaries
- Organize metrics into categories
  - Refine, focus, and group based on end user needs
- Determine Measurable DDoS Attack Thresholds
  - Simulate, test, and conduct correlation and analysis of historical data

Standard Metrics and Categories with Measurable Thresholds

Step 4. *Identify locations at which to Observe Performance and Information Security Events*

Raytheon

- User communities
  - End-users
  - Help desk
  - Operations
  - Engineering
- System components
  - Workstations
  - Computing services
  - Network
  - Transport

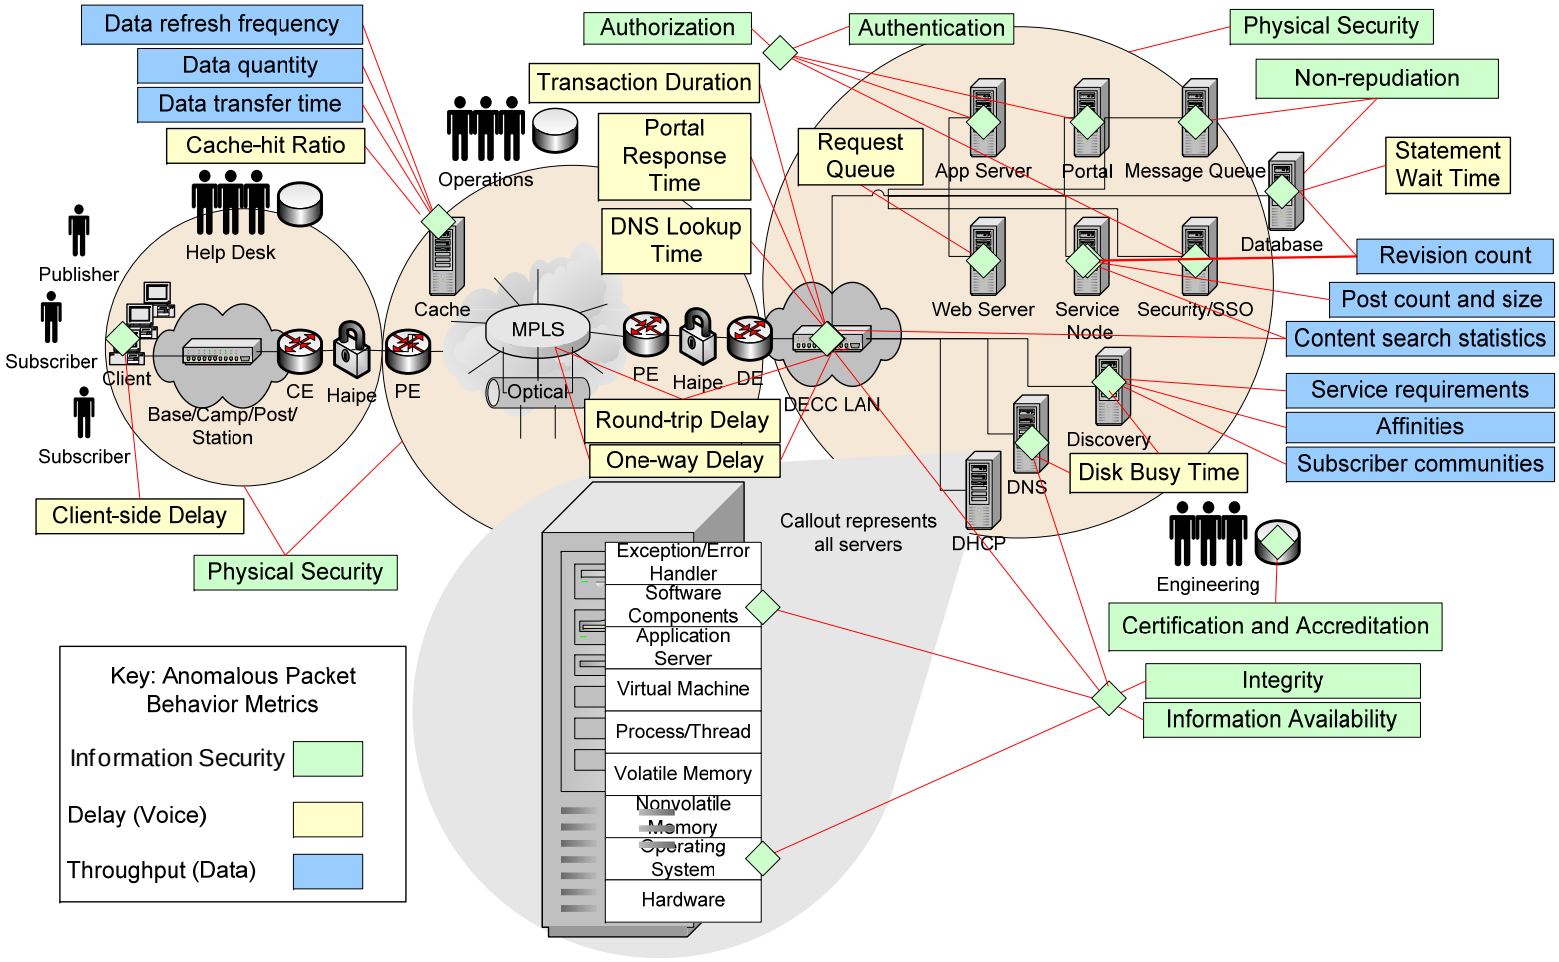

Metrics Detection Locations for Performance [Voice (Delay) & Data (Throughput)] and Information Security

Step 5. *Identify Potential Implementation Schemes*

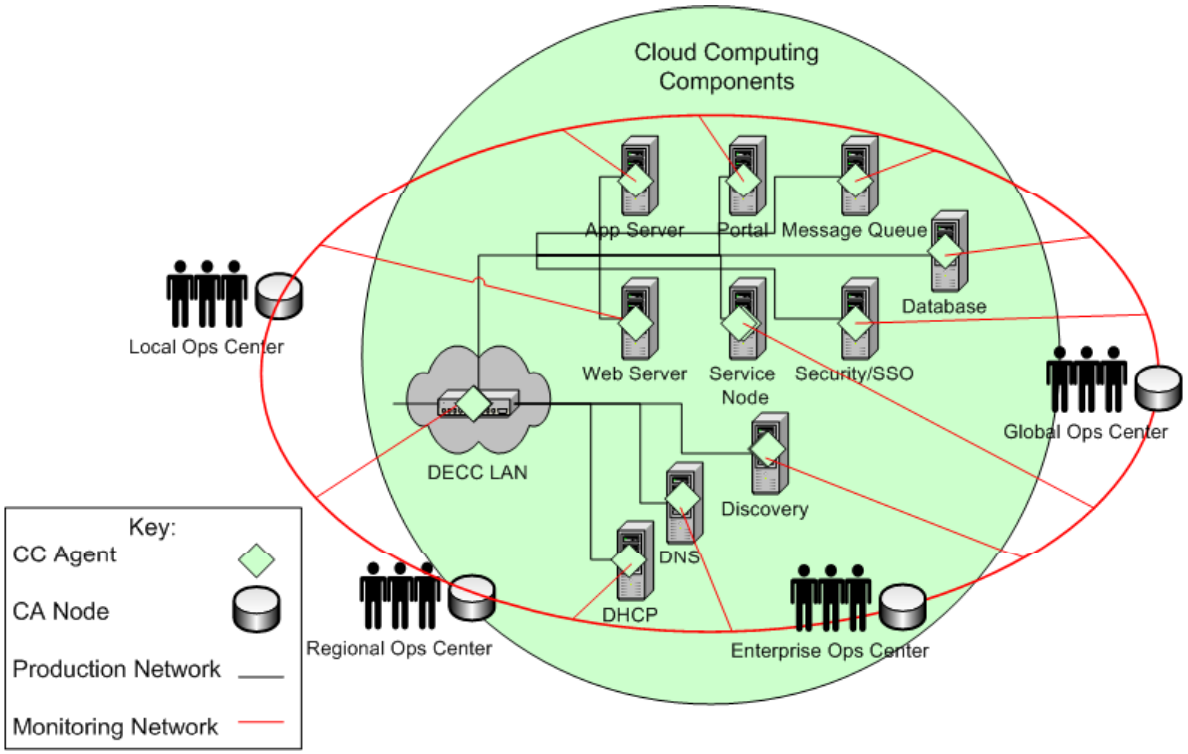

- Embed EMMRA Cloud Computing (CC) agents within multiple diverse cloud computing components
- Continuously monitor enterprise system for QoS metrics
- Agents communicate over an out-of-band (OOB) monitoring network to EMMRA Cloud Collection and Analysis (CA) nodes
- CA nodes are located at local, regional, enterprise and global operations centers

EMMRA CC Agents & CA Nodes Enable Monitoring, Management and Response

System Model: Mathematical Model for QoS Performance Metrics

Raytheon

- Delay
  - SoS view from top level domain (i.e., GaaS) perceives ***delay*** as sum of delays in lower domain levels of cloud.
  - $D_{SoS} = p_1 D_G + p_2 D_B + p_3 D_S + p_4 D_I$

    Where:
    - Each *pi* parameter is dependent on the infrastructure component used.
    - *Dj* is the delay experienced in each layer j in EMMRA,

      Where the specific letter for j is the EMMRA domain (i.e., GaaS, BaaS, SaaS, IaaS)

- Throughput
  - Defined at EMMRA domain level as number of transactions completed per unit time.
  - Visualized at different levels.
    - At GaaS level: order of few days
    - At lower levels: multiplicative in nature.
      - Function of throughput at a lower level:

$$T_I = n \times TransactionThroughput$$

$$T_S = m \times T_I$$

$$T_B = q \times T_S$$

Where m, n and q are numbers of transactions at the lower domain needed to complete the transaction at the higher domain.

Mathematical Models Derived for QoS Performance Metrics (Delay and Throughput)

System Model: Mathematical Model for QoS Information Security Metrics (Authentication)

Raytheon

- Focus on Information Security as a SoS functional requirement comprising authentication and authorization using certificates and accreditation
- *Authentication* metric *is the logical conjunction at* each domain level in EMMRA
  - User's access to the system ceases at level authentication fails.
  - SoS view of authentication is a logical AND of the authentications at various levels in EMMRA (i.e., a top down metric)
    - Lower level EMMRA components have to be kept secure from the end user.
    - User at the top level can obtain service from the bottom levels, but, is not authorized to access the components directly.
    - Only specific personnel are allowed access to the lower level components (*viz, administrators).*
    - Hence in order to obtain access to lower level components the user needs to be authenticated at the top level.

$$A_{SoS} = A_G \wedge A_B \wedge A_S \wedge A_I$$

Authentication Metric is the Logical Conjunction at Each Domain Level in EMMRA

System Model: Mathematical Model for QoS Information Security Metrics (Authorization)

Raytheon

- *Authorization* metric is a bottom-up metric and is applicable at each EMMRA domain level.
  - User access to the service at any layer of EMMRA is subject to authorization.
  - Authorization is such that the least privilege is granted sufficient to accomplish the operation.
  - Authorization is applicable at each level in EMMRA Cloud.
    - e.g., in a banking application, an administrator is not authorized to access account details of the customer of the bank.
  - Authorization at the IaaS level can be represented as

$$Auth_I = \min \left\{ \bigcap_{i \in \text{Set of actions}} p_i \right\}$$

    where *pi* is the permission to perform action *i* at the IaaS level.
  - Similarly, authorization is defined for rest of the domain levels in EMMRA Cloud.

* SoS view of authorization can be obtained using methods such as linear logic.

Authorization is a Bottom-up Metric Applicable at Each Domain Level in EMMRA

Application Scenario: Distributed Denial of Service (DDoS) Attack on Complex Systems

Raytheon

- Apply the new approach presented here to monitor, manage, and respond to QoS in the presence of DDoS attacks in cloud computing environment as follows:
  1. Use the SoS, framework, and metrics defined in Steps 1, 2, an 3.
  2. Use step 4 to identify the locations at which to observe those metrics.
  3. Use Step 5 to deploy EMMRA CC agents at those locations.

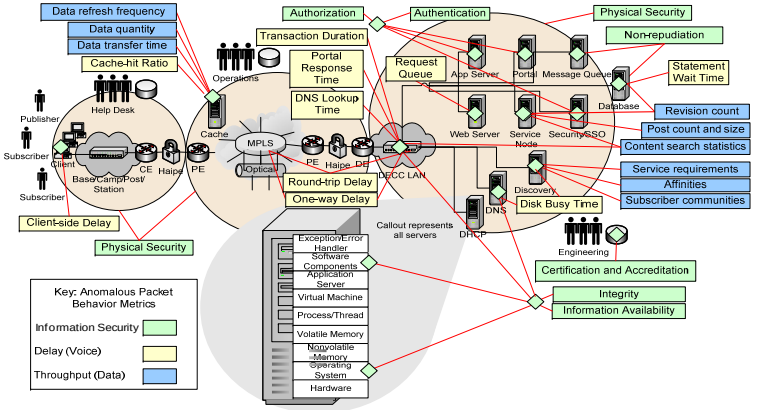

- Rationale:
  1. A*uthentication can be monitored at the Apps,* Portal, and Security/SSO servers (e.g., EMMRA CC agents can monitor Security Assertion Markup Language (SAML) authentication assertions at Security/SSO server.
  2. EMMRA CC agents can monitor and respond to *Authorization events from the Apps, Portal, and Security/SSO* servers where they can access info. such as need-to-know determination required to grant resource authorization.
  3. EMMRA CC agents distributed within the engineering project control and development-tracking database can provide the relevant information to support ongoing certification and accreditation.
- Use Case: Security monitoring and response for a financial/banking application.
  1. Apply SoS and framework
     - Complete one transaction at the business domain
     - Policies established and enforced at GaaS domain require that multiple sub-transactions occur at the AaaS and SaaS domains that are distributed to end-users through the IaaS domain.
  2. Cyber Security Plane monitors across all EMMRA domains to detect and enable proactive response to DDoS security events
     - Apply within all EMMRA domains to prevent transactions that could cause potentially devastating consequences

EMMRA CC Enables Proactive Detection and Response for Security Events on Financial/Banking SoS

Results: Delay, Variation in Delay, and Throughput Performance Metrics Verification

Raytheon

- Performance metrics were measured & recorded at diverse time granularities using a prototype transaction processing application.
- Assumptions
  - QoS thresholds can be changed for different application scenarios (i.e., need not be fixed a priori for all applications to be deployed on a cloud).
- Observations
  - Within an Complex SoS, Delay metrics are additive
  - Both Variation in Delay over time and Throughput are indicators of the overall system performance.
  - Well-established QoS monitoring guidelines and frameworks exist for IaaS and SaaS cloud deployments.
- Actions
  - QoS thresholds were fixed (e.g., throughput per second and delay per millisecond) for the application scenario to be verified
  - Prototype transaction processing application monitored EMMRA service domains for QoS breach.
  - If a QoS breach was observed, then a response action (RA) (i.e., an automated action to rectify the breach) was initiated
  - Experiments establish a method to correlate the IaaS/SaaS QoS breach events to the Baas and GaaS EMMRA domains
  - Correlation provided a SoS view of the QoS monitoring and management in a cloud environment

Results verified EMMRA Cloud Approach Provides a SoS View of QoS in a Cloud Environment

A. Delay recorded in 10 sample transactions

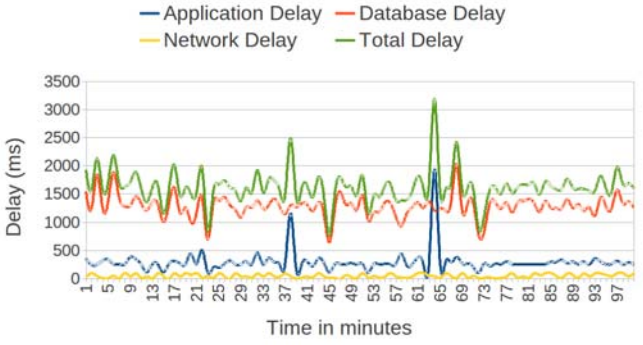

B. Variation in delay recorded over time second

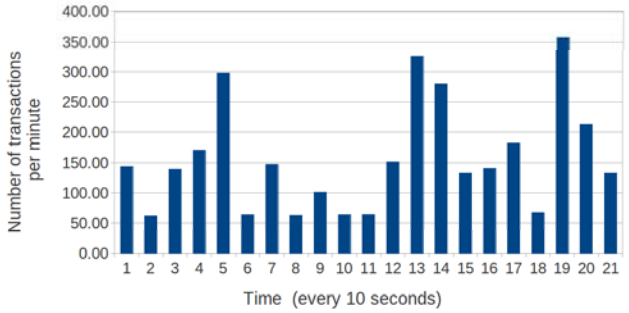

C. Throughput: Number of transactions per minute

Conclusions, Present Status, and Path Forward

- EMMRA CC enables cloud computing service providers and operations centers to meet committed customer QoS levels
  - Uses a trusted QoS metric collection and analysis implementation scheme
  - Extends traditional monitoring, management and response for IaaS and SaaS to complete SOA-stack that includes business logic (BaaS) and governance (GaaS).
- Present Status:
  - EMMRA Architecture is mature and well vetted
  - EMMRA CC performance metrics verified using a prototype transaction processing application

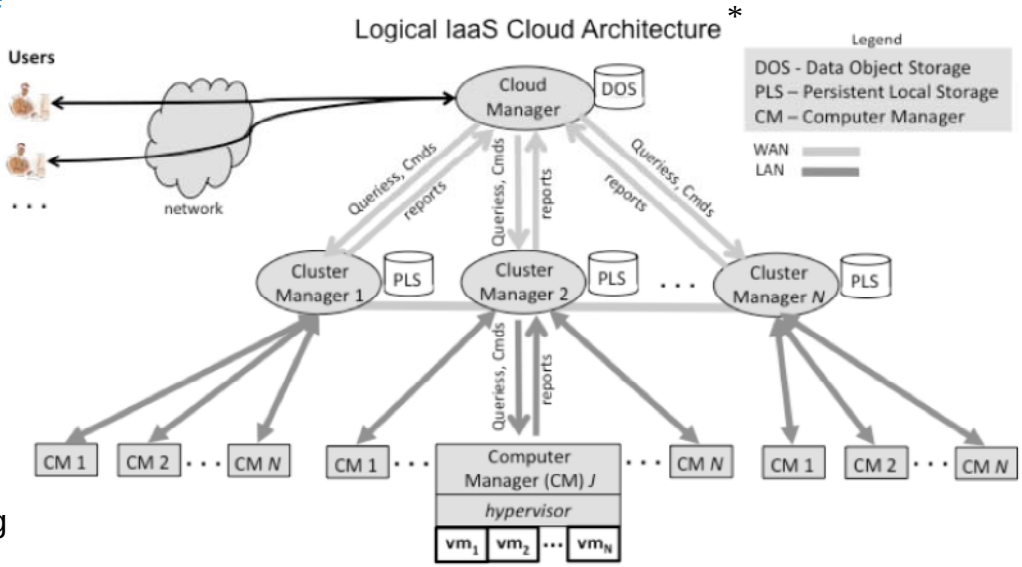

SaaS Cloud Architecture *

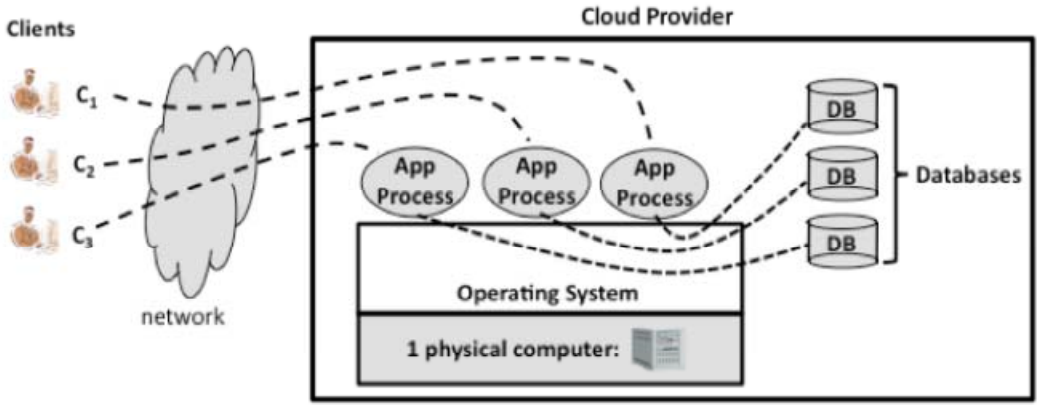

- Next steps
  - Conduct full simulation with diverse scenarios for all EMMRA domains to quantify the effectiveness of this approach
  - Include operations center response time to restore QoS in the presence of anomalous enterprise events.
  - Implement prototype EMMRA Cloud system for single domain (IaaS or SaaS)

New EMMRA Cloud Procedure Enables Operators/Analysts to Effectively Monitor, Manage and Respond within a Complex SoS