



U.S. AIR FORCE



# Spring 2016 Research Update Presentations



# Monitoring Data Fusion to Improve Intrusion Detection

Atul Bohara  
P.I.: William H. Sanders  
ACC Seminar, Jan 27, 2016



# Overview



- **Goal:** protect a real-world networked system against malicious activities
  - E.g., Enterprise network / campus network / cloud data center
  - Prevention techniques are not sufficient
  - Need to rely on security monitoring and detection
- **Data-driven approach to combine information from multiple monitors and detect intrusions**
  - Utilize the vast amount of information generated by security monitors
  - Detect sophisticated attacks



# Recent Progress



- Unsupervised anomaly detection in enterprise system using clustering
  - Identify useful features
    - Represent data in highly useful and concise format
    - Combine across different monitors
  - Unsupervised machine learning
    - Apply clustering and dimensionality reduction techniques to separate normal and anomalous behavior
    - Detect intrusions by analyzing anomalous behavior clusters



# Future Plan



1. Develop techniques to represent and combine data
2. Develop unsupervised intrusion detection techniques
3. Apply them to systems such as campus network, cloud data center



U.S. AIR FORCE



# INTRUSION DETECTION, RESPONSE, AND RECOVERY IN THE CLOUD

**Student:** Uttam Thakore

**P.I.:** William H. Sanders



# Quantitative Methodology for Security Monitor Deployment



- A **cost-effective** methodology for monitor deployment to **meet intrusion detection goals**
  - Uses quantitative metrics to capture monitor utility and cost
  - Uses integer programming to determine optimal monitor deployment based on intrusion detection goals and cost requirements
- Work last semester:
  - Implemented heuristic approach to make solution algorithms scalable
  - Submitted paper to DSN 2016
- Plans for the semester:
  - Exploring collaboration with IBM to apply approach to IBM cloud offering



# Unsupervised Anomaly Detection in Enterprise Systems Using Clustering



- Applying **unsupervised clustering techniques** to network- and host-level security logs to detect malicious behavior
- Work last semester:
  - Devised and implemented initial approach and evaluated on VAST 2011 Mini Challenge 2 data set
  - Submitted paper to HotSoS 2016
- Plans for the semester
  - Apply approach to NCSA security log data



U.S. AIR FORCE

# Adaptive “Learning Responses”



- Deployment and configuration of monitors in response to detected attacker behavior to aid intrusion detection algorithms
  
- Plans for the semester:
  - Investigate existing tools and literature for predictive monitor selection
  - Apply unsupervised learning techniques to NCSA data to identify events that warrant additional monitoring



U.S. AIR FORCE



# A Flexible Fine-Grained Adaptive Framework for Parallel Mobile Hybrid Cloud Applications

Kirill Mechitov


PI: Gul Agha



# A Flexible Fine-Grained Adaptive Framework for Parallel Mobile Hybrid Cloud Applications



- Research Problem
  - Support *hybrid mobile cloud computing* – mobile devices leveraging cloud resources: secure private cloud + public cloud
  - *Dynamic reconfiguration and offloading* can achieve dramatic speedups (over 50x) for compute-intensive tasks such as image processing and/or mobile device energy savings
  - *Security policies* are needed to prevent unauthorized access and leakage of sensitive information from secure devices and private clouds to public clouds
- Status & Plans
  - *Illinois Mobile Cloud Manager (IMCM)* framework for hybrid MCC applications: prototype implemented
  - Optimize for different energy/performance objectives
  - Implement enforcement of actor semantics by the IMCM runtime





U.S. AIR FORCE

# IMCM: Illinois Mobile Cloud Manager



- Code offloading:
  - Automatic
  - Dynamic
  - Fine-grained
  - Parallel
- Supports:
  - Hybrid cloud with multiple cloud spaces
- Provides:
  - Policy-based control by cloud provider, app developer, user

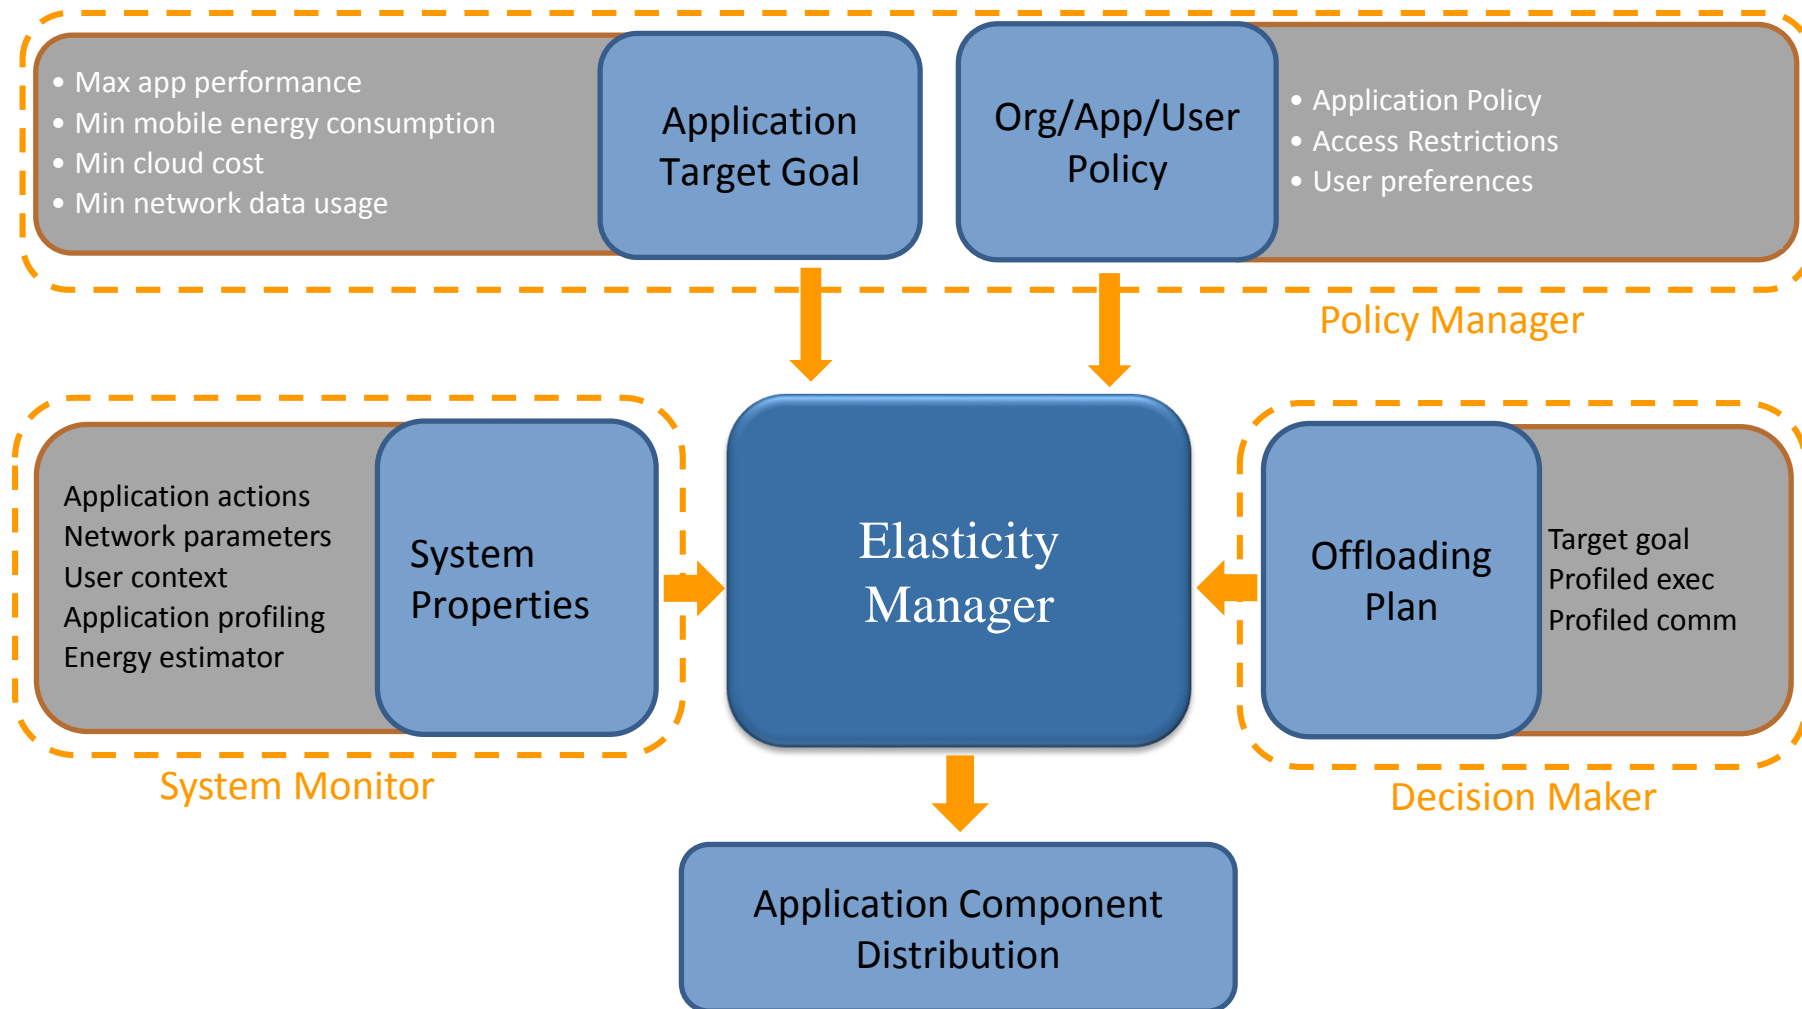




U.S. AIR FORCE



# IMCM framework






U.S. AIR FORCE

# Research Plans



- Enforce actor model
  - Rather than assume well-behaved code or rely on compile-time enforcement
  - E.g., no out-of-band communication
- Energy performance optimization
  - Automate estimation of energy use by application components without user/programmer assist
  - Integrate with current constraint solver for dynamic runtime optimization





U.S. AIR FORCE



# Moving towards a Secure Container Framework

Mohammad Ahmad, Rakesh Bobba,  
Sibin Mohan, Roy Campbell



# Background



- Container benefits
  - Startup on the order of milliseconds
  - Packaging dependencies & portability
- Container usage
  - Platform as a Service Clouds
  - Openshift, DotCloud
- Cross container side-channel attacks shown on public clouds [1]

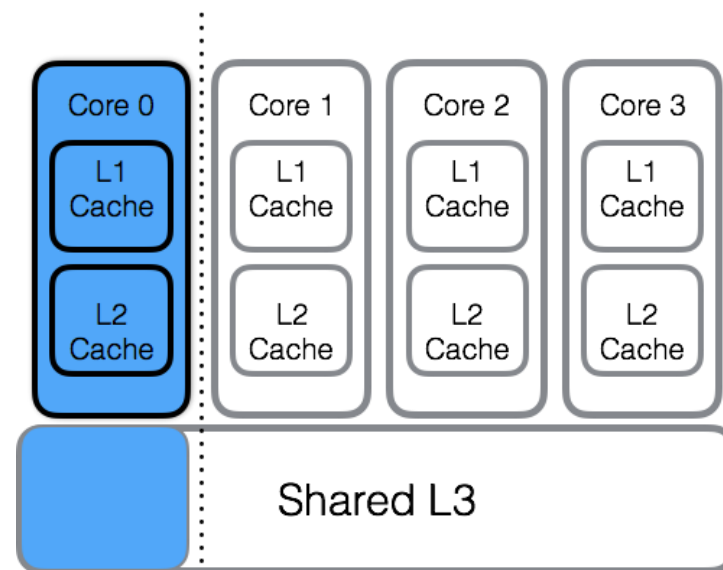


U.S. AIR FORCE

# Secure Container Framework



- Phase 1 – Defenses against cache based side-channels
  - Scheduling-based defenses
    - Cache flushing
  - Incorporate hardware support
    - Intel Cache Allocation Technology to isolate parts of the LLC





# Progress



- Built a loadable kernel module
  - Plugs into the Linux scheduler routine
  - Return probes (kretprobes)
- Currently adapting relevant benchmark suites



# CRONets: Cloud-Routed Overlay Networks

Chris Cai

PI: Professor Roy Campbell



U.S. AIR FORCE

## Phurti: Application and Network-Aware Flow Scheduling for Multi-Tenant MapReduce Clusters



- Phurti: Application and Network-Aware Flow Scheduling for Multi-Tenant MapReduce Clusters (*Chris Cai, Shayan Saeed, Indranil Gupta, Roy Campbell, Franck Le*) has been accepted at IC2E 2016
- Will be presented at the conference in April



U.S. AIR FORCE

## CRONets: Cloud-Routed Overlay Networks



- We aim to understand what level of performance improvement can a user expect to get from leveraging public cloud service to build overlay network, as opposed from other resource providers like ISPs.
- Performance metrics can include throughput, latency, loss rate, etc, corresponding to particular demands of different applications.



U.S. AIR FORCE

## Measurement Testbed



- We used PlanetLab nodes as clients and Eclipse mirrors as servers. We used IBM Softlayer as cloud provider to provide overlay nodes.
- Blue labels indicate locations of PlanetLab nodes. Red labels indicate locations of overlay nodes.





U.S. AIR FORCE

## Ongoing work



- Investigating
  - How persistent can a user expect the improvement to be over a certain period, say, a week?
  - What types of network connections can expect the greatest improvement?
  - How many overlay nodes are needed to achieve the best performance?
  - How to automatically choose the best overlay path?



U.S. AIR FORCE



# Monitoring through Inference

Imani Palmer

P.I. Roy Campbell



U.S. AIR FORCE



- Current research:
  - Determine the latest inferencing methods
  - Analyzed current methods
  - Defined a framework
- Plans for the year:
  - Build an inference engine
  - Define a set of policies for intrusion detection/VMI introspection case
  - Show a demonstration of the monitoring system cases



# Research Update

Mainak Ghosh

PI: Indranil Gupta



# Project Status



- **Morphus and Parqua - Completed**
  - Supports reconfiguration in two popular NoSQL databases – MongoDB and Cassandra.
  - Reconfiguration involves changing table level configuration parameters like shard key which affects a lot of data at once.
  - Morphus was accepted as a conference (ICAC) and journal (IEEE TETC) publication. Parqua accepted as a short paper to ICCAC.
- **Getafix – Ongoing**
  - Real-time analytics system like Druid batch temporal data by time segments. They support aggregation queries like COUNT over a time interval.
  - For supporting high query throughput, segments are replicated. Current replication strategies are naïve which do not account for popularity. This leads to poor disk utilization.
  - In Getafix, we first propose an algorithm which provably gives the lowest replication factor required to maintain best query throughput.
  - We design and implement a new adaptive replication scheme in Druid which considers segment popularity.
  - Currently working on the implementation.



# A GAME THEORETIC APPROACH FOR SECURITY

Keywhan Chung

Advisor: Professor Iyer, Professor Kalbarczyk



# Status



- Continued work with Dr. Kamhoua & Dr. Kwiat at AFRL
- Game Theory with Learning for Cyber Security Monitoring
  - Application of **Q-Learning** models for decision making under a *released assumption/restriction* of the attack model
  - Accepted & Presented at HASE 2016
- Signaling Game
  - Decision making is based on *limited (and inaccurate) information*
  - Signaling game derives the optimal decision **given a possibly corrupted message** (observation).
  - Attack model: SlowDoS



# Progress



## Completed

- Preliminary Analysis of Web traffic @NCSA
- Measurements on the victim web server under attack
- Signaling Game simulator

## In Progress

- Formulation of the reward model and justification
- Simulation based evaluation (accuracy)
- Experiment on an actual web application (timeliness, accuracy)



U.S. AIR FORCE



# Research update

Zak Estrada

PI: Ravishankar K. Iyer

