

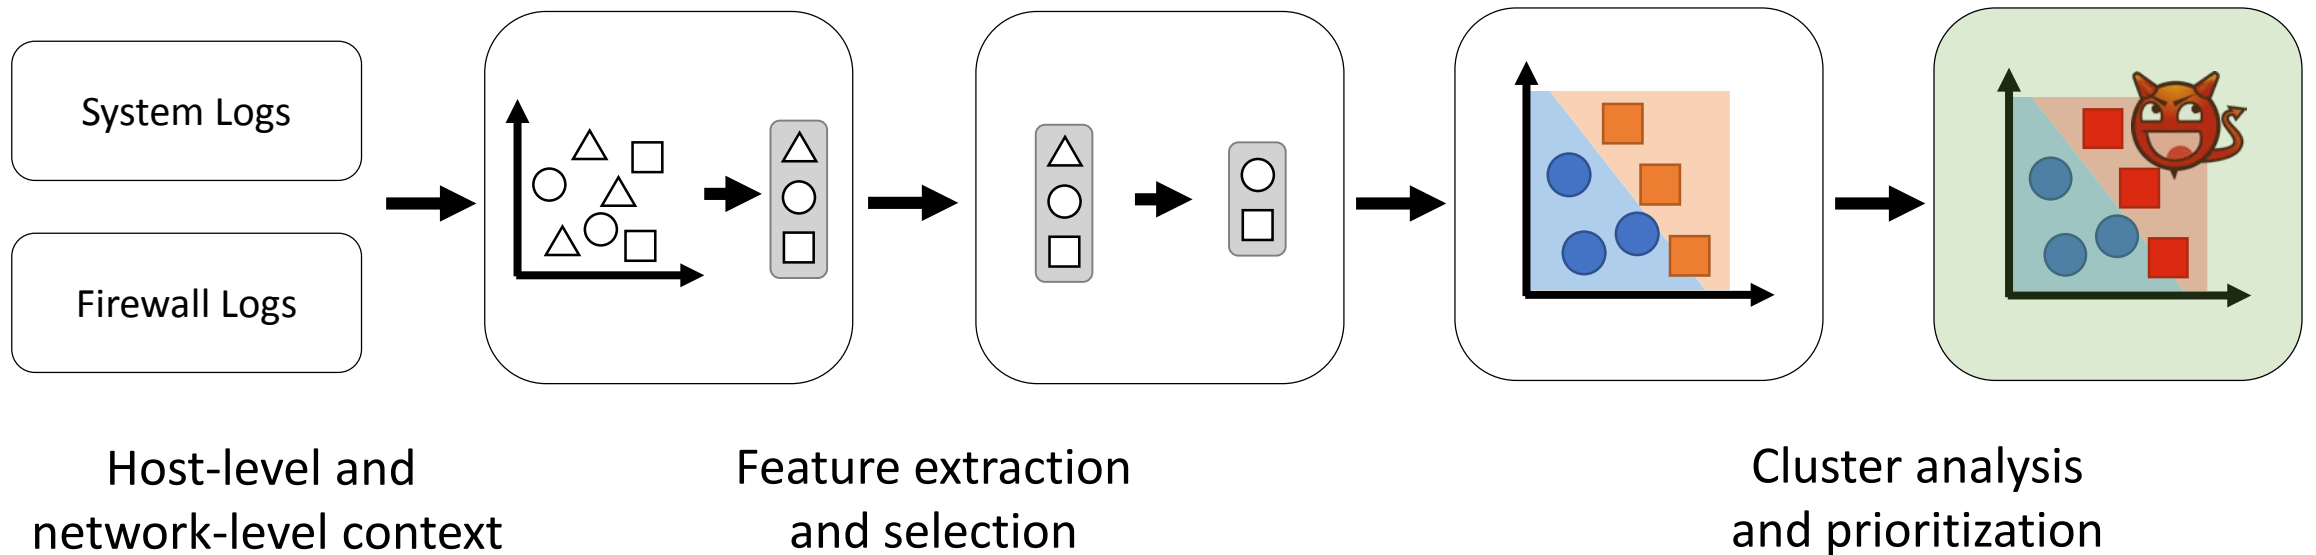
Intrusion Detection Using Monitor Information Fusion

Student: Atul Bohara

P.I.: William H. Sanders


Previous Work [1]

Intrusion detection by combining and clustering diverse monitor data



Previous Work [2]

Lateral movement detection using distributed data fusion



Ongoing Work

Proactive detection of advanced attacks through fusion



- Hypothesis: events observed in the system, as a result of a multi-stage attack, are correlated. By combining the evidences of different attack stages, we can increase the confidence in the detection of overall attack
- E.g., Fuse the evidence of *C&C* and lateral movement to detect and prevent a possible data exfiltration attack
- Data-driven modelling of attack and defense (system)



Air Force Research Laboratory



Chris Cai
PI: Professor Roy Campbell

Integrity ★ Service ★ Excellence



CRONets: Cloud-Routed Overlay Networks



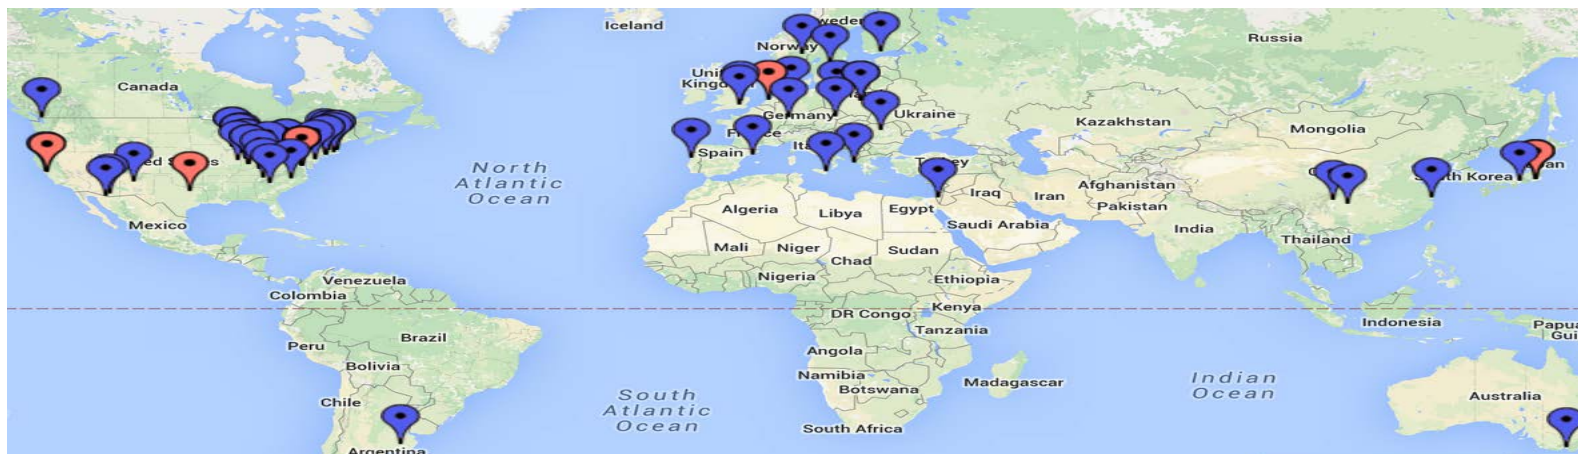
- We aim to understand what level of performance improvement can a user expect to get from leveraging public cloud service to build overlay network, as opposed from other resource providers like ISPs.
- Performance metrics can include throughput, latency, loss rate, etc, corresponding to particular demands of different applications.
- Questions to answer:
 - **Can CRONets provide similar improvements compared to the previous experimental studies, but in a realistic-cloud-setting?**
 - **How can emerging technologies simplify the overlay path selection problem?**



Measurement Testbed



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





Contributions



- Our work will help large companies as well as individual users to best leverage the available commercial cloud network resources to meet their specific network requirement.
- CRONets also has the potential to provide a robust fault-tolerant transmission layer to help application surviving network failures.
- Our work will help cloud provider to better design their inter-datacenter transmission mechanism to be “CRONets-friendly”.



Advisor: Professor Iyer, Professor Kalbarczyk

KEYWHAN CHUNG



Security as a Signaling Game



- Continued work w/ Dr. Kamhoua & Dr. Kwiat at AFRL
- An approach on modeling the **decision making process** for security under limited observation on the environment as a **signaling game**, and studying the effectiveness of the optimized decisions
- Simulation results had shown:
 - That the signaling game can reason the decisions of the attacker
 - Worst case scenarios for the defender
 - Promising evaluation results compared to the common approach
- Further steps:
 - Comparison with more advanced mitigation methods or other attack models
 - Deployment to a real system w/ real monitors and responses



Attack on Computing Infrastructures through Targeted Alteration of ICS



- A study on seeking the possibility of utilizing the relatively weak security of the ICS systems to attack a well hardened computing infrastructure that requires advanced environmental control
- Studied the cooling system for Blue Waters
 - Campus / Building / Cooling cabinet level
 - Interdependency between the systems
- Studied Blue Waters failures related to the cooling system
 - Three failure scenarios with possibility of the attacker replaying through alteration in the monitoring / control system
- Further steps:
 - Formulation of the attack model
 - Mitigation methods (Bro IDS, etc.)

Intel VT-x on QEMU

Lavin Devnani

PROJECT GOALS

- ▶ Extend QEMU (Quick Emulator) to emulate Intel VT-x instruction set
- ▶ Run a hypervisor + guest OS in emulated operating system
- ▶ Support future security and reliability projects

QEMU + VT-x

HOST OS


KVM

GUEST OS

```
lvml@lvml-desktop:~$ cat /proc/cpuinfo
cat: /proc/info: No such file or directory
lvml@lvml-desktop:~$ cat /proc/c
lvml@lvml-desktop:~$ cat /proc/c
cgroups  cmdline  consoles  cputime  crypto
lvml@lvml-desktop:~$ cat /proc/cputime
lvml@lvml-desktop:~$ cat /proc/cputime
processor      : 0
vendor_id     : GenuineIntel
cpu family    : 6
model        : 69
model name    : Intel(R) Core(TM) i7-4790 CPU @ 3.60GHz
stepping     : 3
microcode    : 0x1c
cpu MHz      : 3720.093
cache size   : 8192 KB
physical id  : 0
siblings     : 8
core id      : 0
cpu cores    : 4
apicid       : 0
local apicid : 0
fpu          : yes
fpu_exception : yes
cpuid level  : 13
wp           : yes
flags        : fpu vme de pse tsc nsr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush dts_
l mmx fxsr sse sse2 ss ht tm pbe syscall nx pdpe1gb rdtscp lm constant_tsc arch_perfmon pebs bts rep_good
nopl_xtopology nonstop_tsc aperfnperf eagerfpu pni pclmulqdq dtes64 monitor ds_cpl vmx smx est tm2 sse3
fma cx16 xtpr pdcm pcid sse4_1 sse4_2 x2apic movbe popcnt tsc_deadline_timer amd_
ahf ln_ahb ida arat epb pln pts dtherm tpr_shadow vmmi flexpriority ept vpid fsgsbase tsc_adjust bmi1 avx
2 smep bmi2 erms invpcid xsaveopt
bugs         :
bogomips    : 7195.43
clflush_size : 64
cache_alignm : 64
address sizes : 39 bits physical, 48 bits virtual
power management:
processor    : 1
vendor_id   : GenuineIntel
cpu family  : 6
```

Future Applications

- ▶ Taint analysis of VT-x
- ▶ Taint analysis + Symbolic execution
- ▶ Profiling existing hypervisors
- ▶ Prototyping new hypervisors
- ▶ Extension of VMX functionality




Cloud Security Certifications: A Comparison to Improve Cloud Service Provider Security

Carlo Di Giulio (cdigiul2@illinois.edu)

Masooda Bashir (mn@illinois.edu)

09/21/2016

Previous Steps



Goals:

- Security & Privacy in Cloud Environments
- Evaluation of cloud vendors
- Market trends

3 Pillars:

- Laws and Regulations
- Cloud services
- Privacy and security policies

- Focus on the first and third pillar
- Security and privacy certifications and standards
- FedRAMP, ISO27001


Contribution

- ☁ Evaluation of the impact and relevance of Privacy and Security certifications for Cloud Services
- ☁ Deeper understanding of vendors' commitment in promoting information assurance
- ☁ Suggestion of improvements to current standards and guidelines




Current Status, Accomplishments


☁ ISO27001:2005 and 2013



☁ FedRAMP Moderate and High baseline
(DoD Lev 2-4)




☁ AICPA SOC2 (TSPC 2014 and 2016)



☁ BSI Cloud Computing Compliance Control
Criteria (C5)



Bundesamt
für Sicherheit in der
Informationstechnik



Secure Containers

Konstantin Evchenko, Read Sprabery, Abhilash Raj*,
Sibin Mohan, Rakesh Bobba*, Roy H. Campbell

University of Illinois at Urbana-Champaign

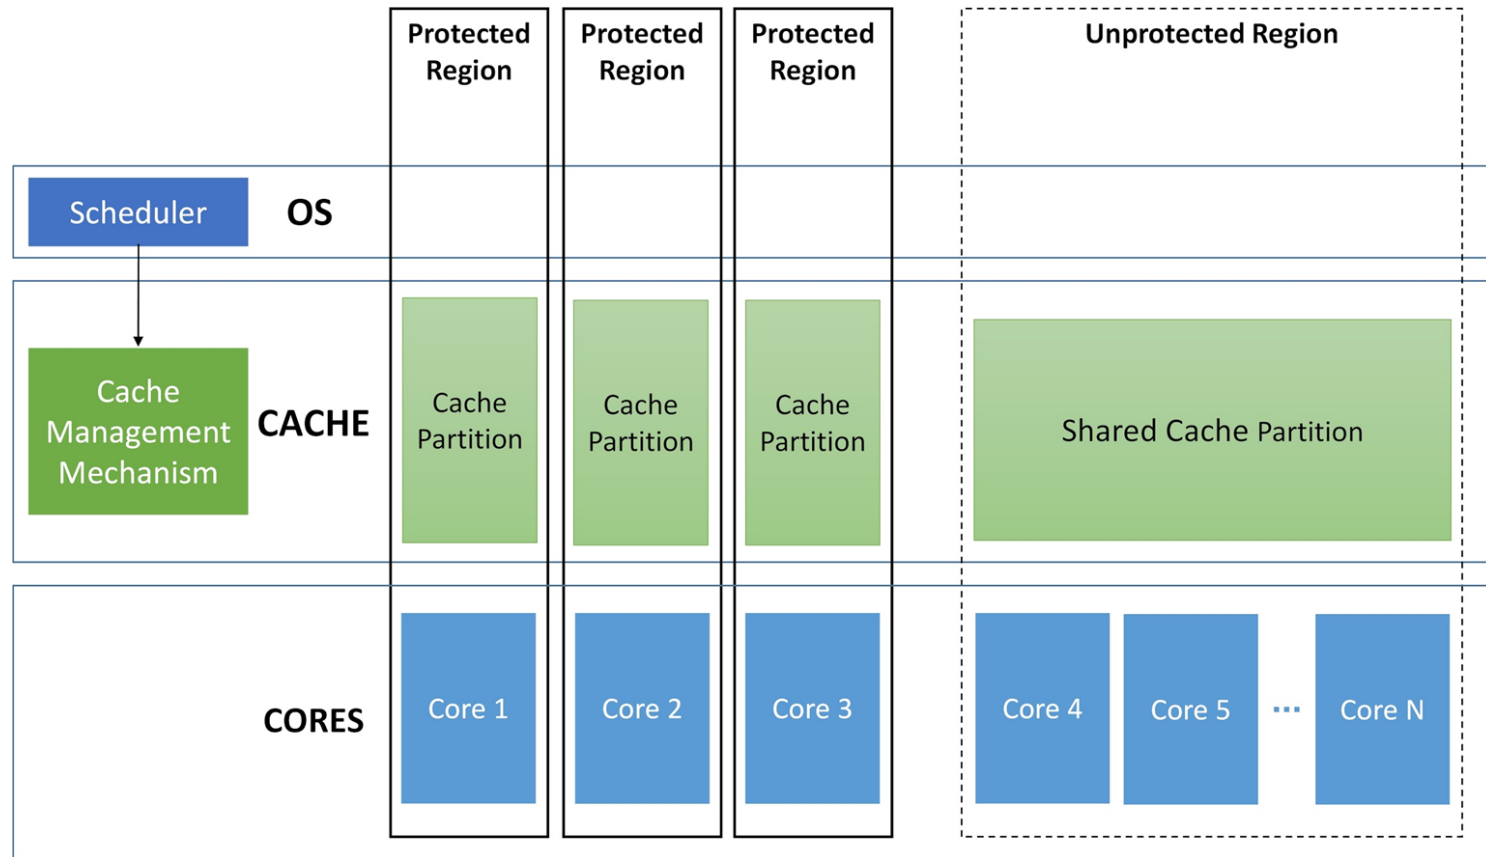
*Oregon State University

Motivation

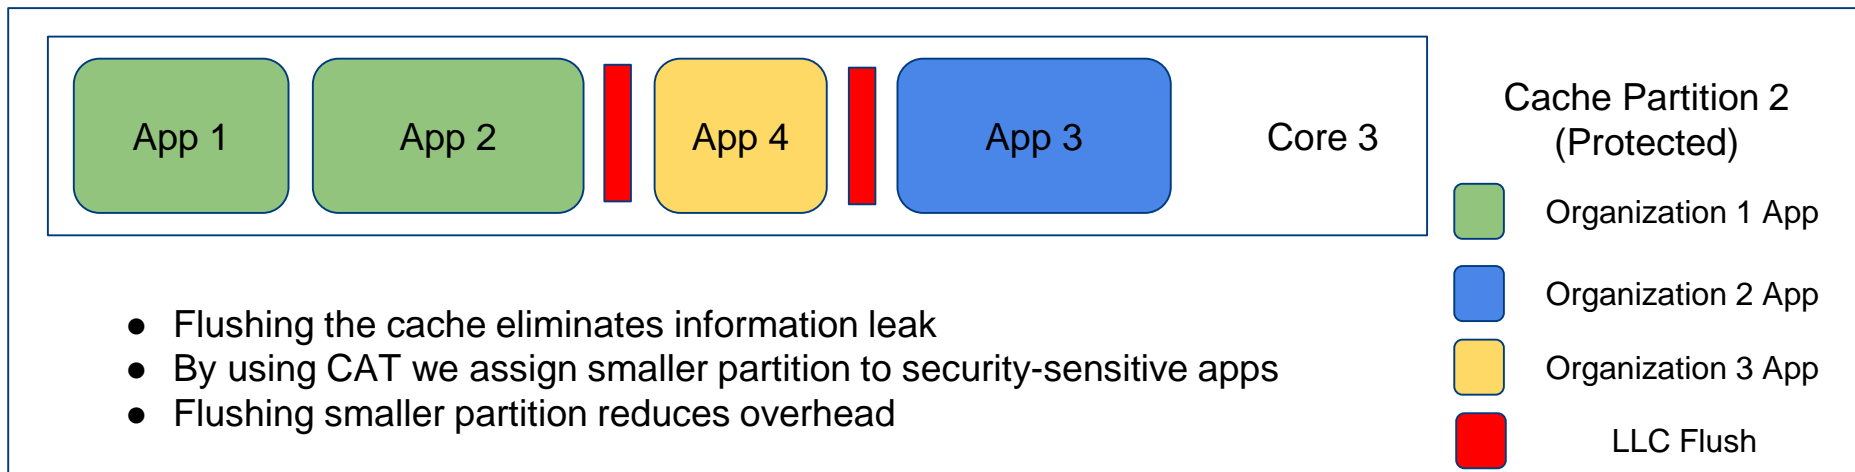
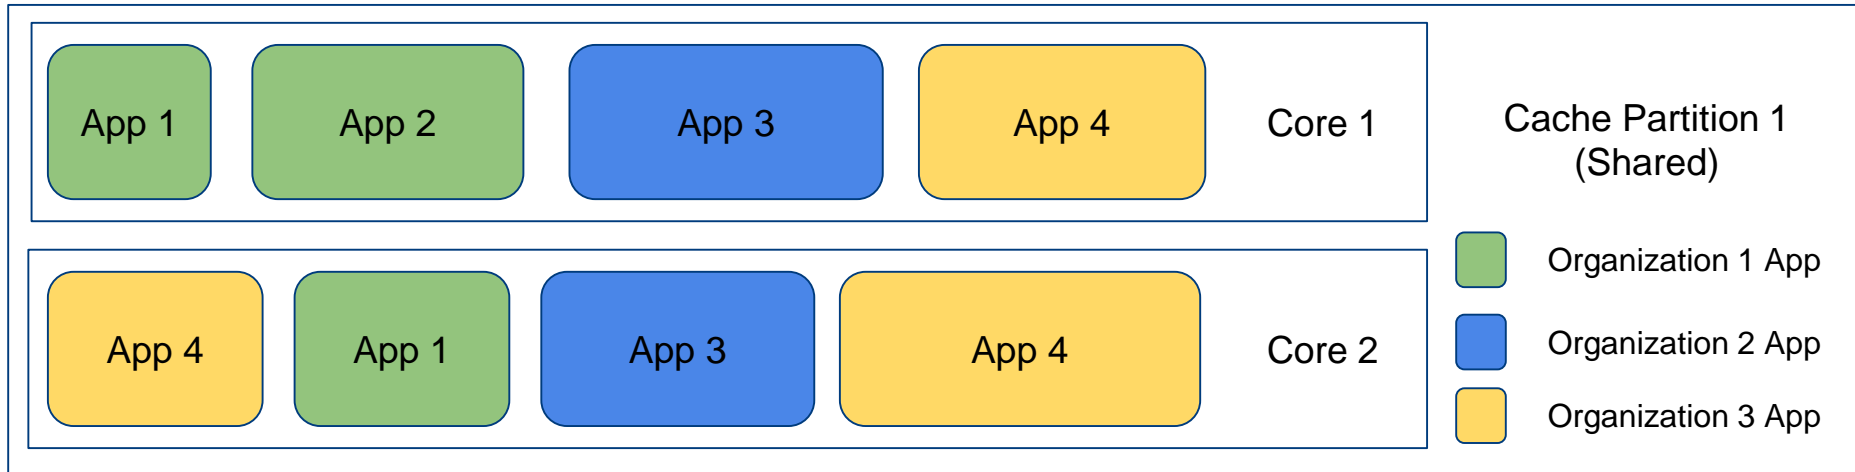


- Container-based products become ubiquitous in cloud infrastructure
- Several parties run their containerized applications in a shared environment
- Enables cache-based side-channel attacks (e.g. Prime+Probe and Flush+Reload)
- These attacks can be used to retrieve fine-grained sensitive information (e.g. cryptographic keys)
- Both attacks have been effectively carried in PaaS and IaaS infrastructures, both in a lab and real world environments

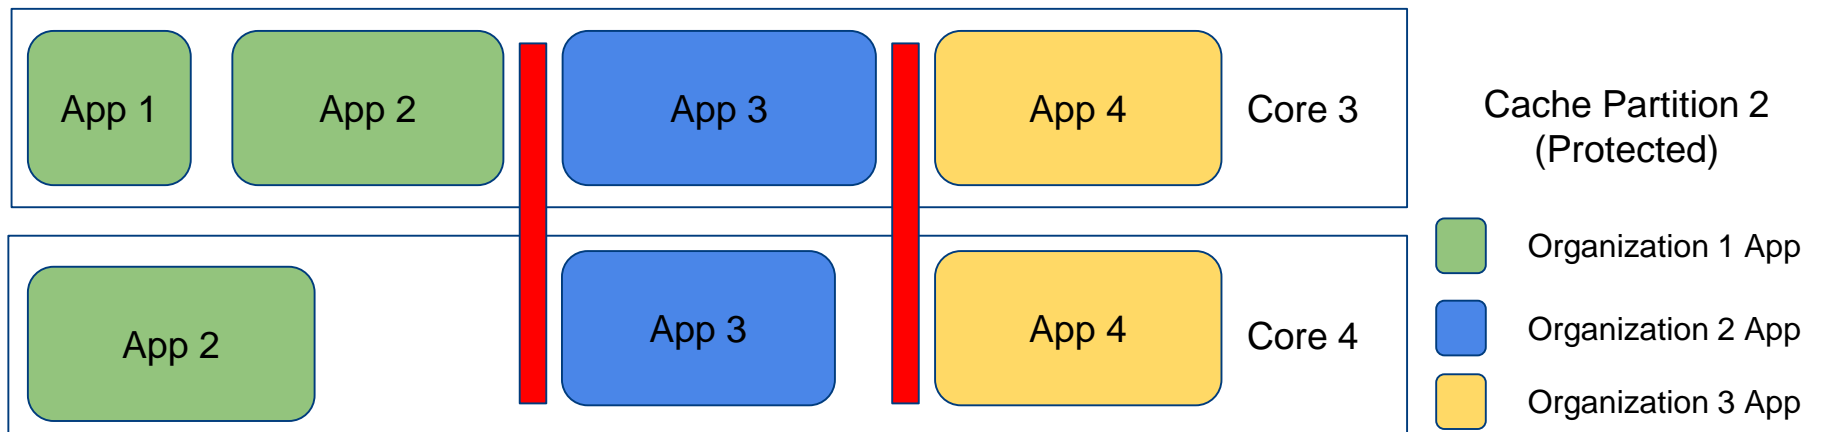
Cauldron Framework Design



Workflow example




Improving performance with Gang Scheduling



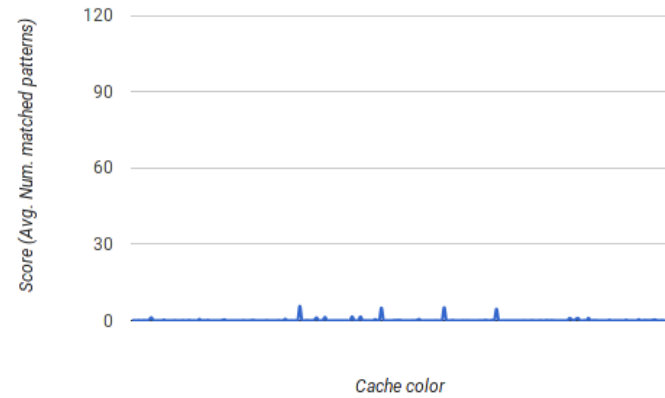
- Gang-schedule apps from the same organization
- Reduces the number of flushes
- Potentially increases idling
- Possible solution: soft gang scheduling
 - If no apps from the same org are available, schedule from other orgs
 - No flushing
 - Might leak some information, but not enough to enable the attack

Initial results

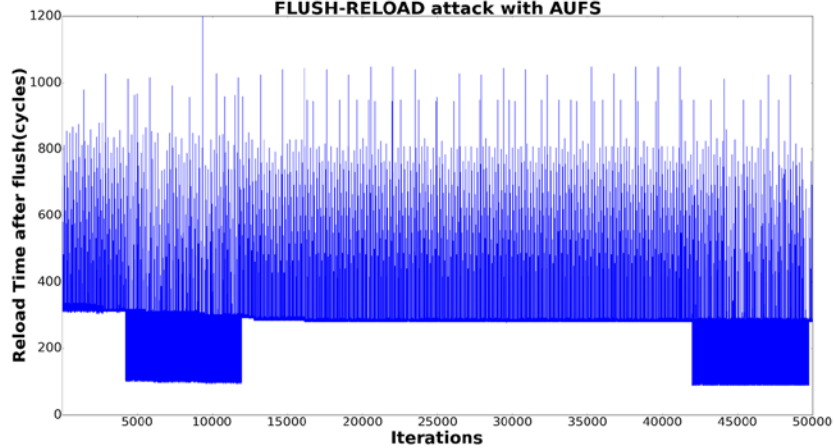
Attacker and victim share cache partition




Attacker and victim don't share cache partition



FLUSH-RELOAD attack with AUFS



FLUSH-RELOAD attack with Devicemapper



- Design a Secure Containers framework with support from multiple layers of the stack including hardware, hypervisor, kernel, compiler and application layer.
 - Hardware supported isolation and sandboxes
 - Novel scheduling techniques for increased isolation and performance
 - Monitoring techniques to detect compromises and protect containers from both co-tenants and host

Getafix: Workload-aware Data Management in Lookback Processing Systems

Presenter: Mainak Ghosh

Collaborators: Le Xu, Thomas Kao, Xiaoyao Qian, Indranil Gupta

Problem

- Lookback Processing Systems -- Warehouse for time series data
- Current systems like Druid, Pinot make workload assumptions in design replication, caching and load balancing strategies
 - Recent segments assigned to “hot tier” -- larger replication
 - LRU used for cache eviction
- Under a different workload, this causes, poor memory utilization, large network overhead
- Our solution, Getafix, proposes a general solution which looks at segment popularity to define replication, caching and load balancing strategies.

Progress

- Finished so far ...
 - Proposed an optimal algorithm which can minimize the replication while improving throughput
 - Compare different replication strategies using a simulator
 - Implemented Getafix inside Druid
- Moving on ...
 - Evaluate the improvements in memory usage, effect on query throughput while using our adaptive replication scheme
 - Implement popularity aware caching and load balancing strategies and measure their effect
 - Publish this work in a top conference.

Energy-Aware Dynamic Code Offloading in Mobile Cloud Applications

Kirill Mechitov, Atul Sandur, and Gul Agha

IMCM: Illinois Mobile Cloud Manager

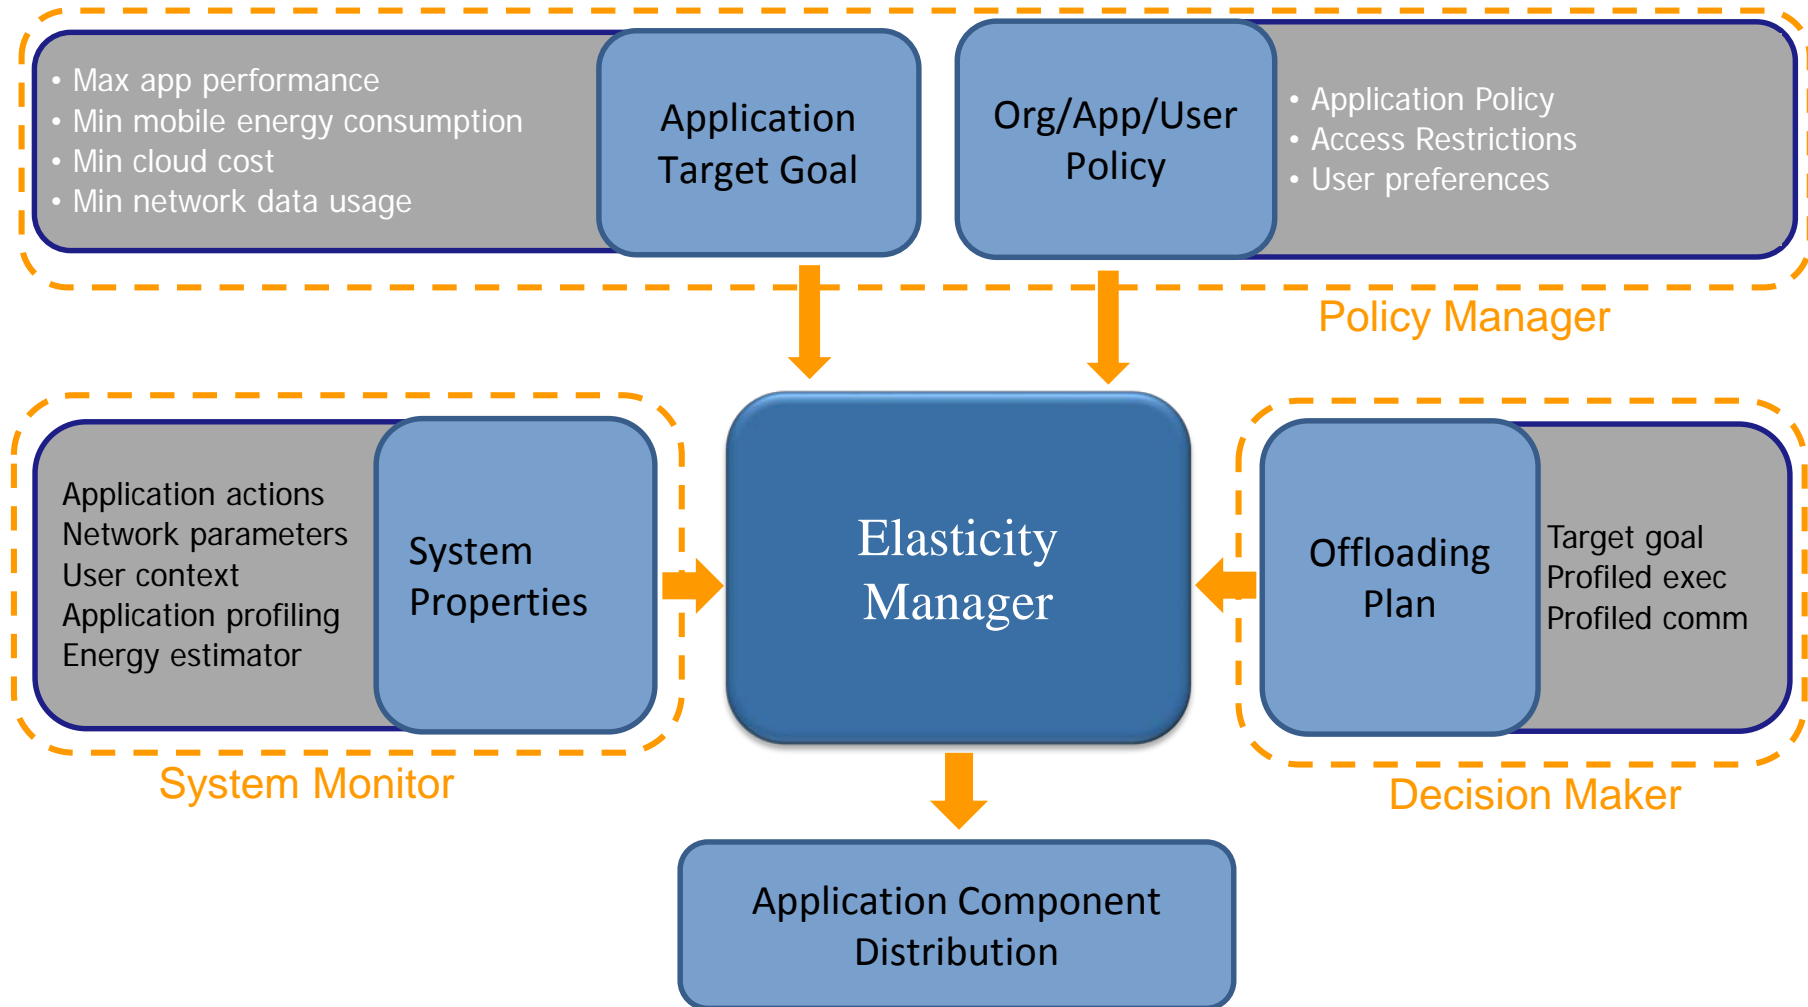
- Code offloading:
 - Automatic
 - Dynamic
 - Fine-grained
 - Parallel
- Supports:
 - Hybrid cloud with multiple cloud spaces
- Provides:
 - Policy-based control by cloud provider, app developer, user



Research topics

- **Monitoring**
 - Real-time fine-grained monitoring of energy use, performance, security of mobile actor-based applications
 - Real-time monitoring of actor energy use and access patterns to identify malicious code
- **Optimization**
 - Refined model for actor deployment and dynamic reconfiguration in hybrid mobile-cloud spaces with security-aware priority management
- **Model-checking**
 - Use model-checking tool for creating valid and sensible initial deployment configurations
- **Policy-based control**
 - Composable per-site, per-actor, per-user policies

IMCM framework



Future work

- Full IMCM Framework proof-of-concept implementation
 - Based on Salsa for Android mobile actor platform
- Model-checking tool for actor deployment
 - Timed Rebeca model of mobile-cloud hybrid applications
- Optimization algorithm for actor deployment and reconfiguration
 - Performance & energy goals
 - Policy-based access restrictions
 - Assurance of performance guarantees/SLAs

A Digital Forensic Analysis Framework

Imani Palmer

Department of Computer Science
University of Illinois at Urbana-Champaign

Roy Campbell

Department of Computer Science
University of Illinois at Urbana-Champaign

Motivation

Cloud is composed of a large number of components vulnerable to attacks

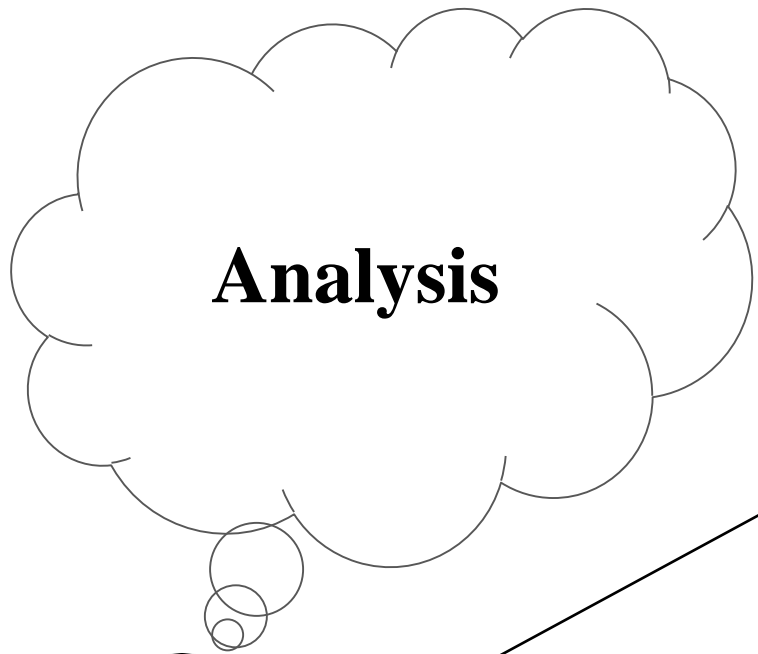
Systems generate an enormous amount of digital evidence

Incident responders/examiners determine the cause of the intrusion


Analysis of digital evidence remains highly subjective to the forensic practitioner

The Problem

The digital forensic investigative process is marred by its lack of knowledge, accreditation, and human bias.




Analysis



Analysis Toolkit

Google Search History
Chat logs
Email
Photos
Internet Activity Logs
Executable Programs
Internet Protocols Address
Financial Asset Records
Address Books
Telephone Records
Maps
Movie Files
Images
Configuration Files

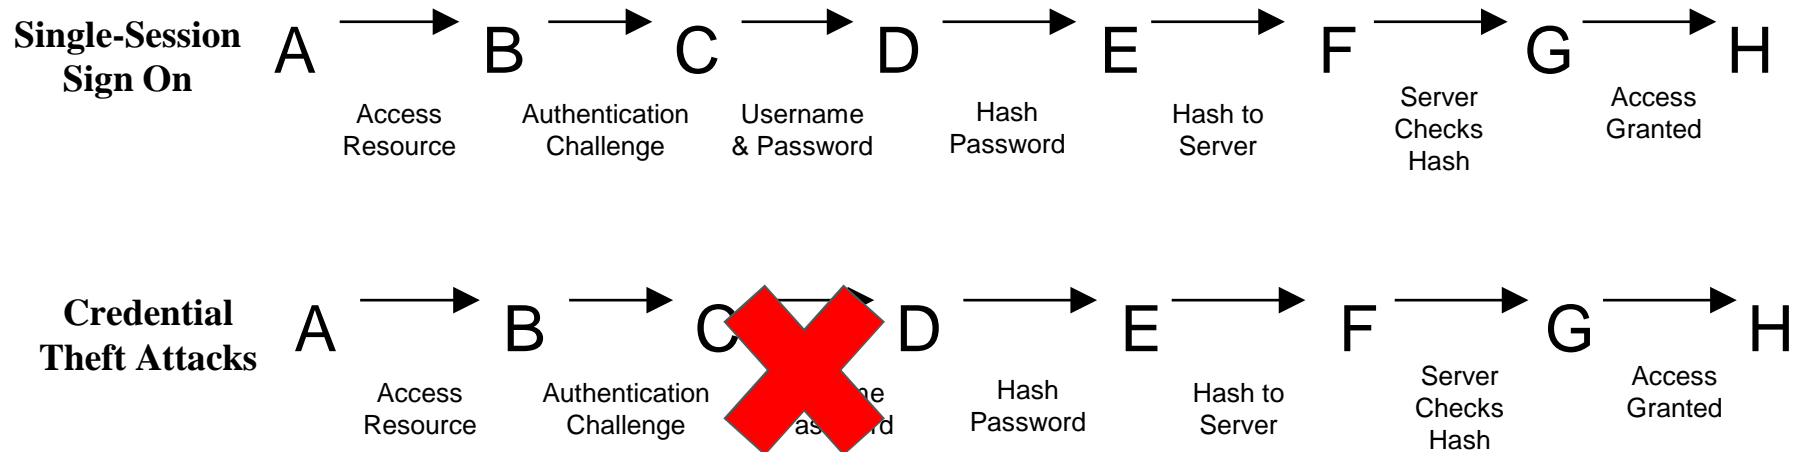


Analysis Toolkit

This actor took action X is supported by facts with strength and quantity




Objective Analysis

Provide quantitative assessments to detect user actions




Framework

Extract Events



Construct Mappings



Identify Actions



Continued Work

Implement Framework

Run case study evaluations

Provide a tool for digital forensic investigators



Verification of Distributed Key-Value Stores Using Reachability Logic

Stephen Skeirik
PI: José Meseguer

Introduction



- Knowing that a distributed system design satisfies certain
 - Consistency
 - Latency
 - Securityrequirements *before* being fully built saves time and money
- Model checking and deductive theorem proving-based techniques can both be used to verify distributed systems meet such requirements

Project Status



- We have successfully used model-checking techniques to explore the behavior of key-value stores, e.g. Cassandra
- To provide even greater assurance, we plan to develop and verify models of key-value stores in reachability logic
- *Reachability logic* naturally models behaviors of *complex, concurrent* systems with *recursive* behavior (as a generalization of both Hoare and Separation logic)
- We have already performed simple experiments modeling and verifying mutual exclusion algorithms

Future Outlook



- Our work will proceed in two directions:
 - Modeling and verifying a selection of distributed key-value stores (e.g. Cassandra and G-DUR are potential targets)
 - Using the ACC case studies to improve the effectiveness of our reachability logic analysis tools for distributed systems (esp. heuristics and techniques for handling undecidable theories)

Project Topic: Dynamic security monitor selection and data analysis for intrusion detection

Student: Uttam Thakore

P.I.: William H. Sanders

Previous work

- A Quantitative Methodology for Security Monitor Deployment
 - A methodology for monitor deployment to **meet intrusion detection goals** and **minimize monitoring cost**
 - 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
 - **Best Paper Award** at DSN 2016
- Intrusion Detection in Enterprise Systems by Combining and Clustering Diverse Monitor Data
 - Applied unsupervised clustering to **fused** network- and host-level security logs to identify potentially malicious behavior **without administrator labeling**
 - Presented at HotSoS 2016

Current work: Data-driven monitor selection in enterprise clouds

- Using **statistical correlation techniques** to identify data that would promote **earlier investigation and detection of incidents**
 - Intuition: Data sources with high correlation to incident-specific alerts with temporal lag are likely useful for detection
- Prioritize monitor deployment/alert investigation based on strength of correlation and administrator security requirements
- Plans for this year:
 - Refine approach and evaluate on NCSA historical security data
 - Will submit paper (potentially to DSN 2017)

Contribution: Can be used to more effectively monitor clouds for security, reliability, and performance incidents

Planned work: Host behavior analysis across heterogeneous logs using unsupervised learning

- Extension of HotSoS 2016 work
- Using unsupervised learning over heterogeneous logs to **classify and track behavior of hosts** over time and identify likely malicious behavior in early stages
- Plans for this year:
 - Identify unsupervised learning techniques and features that strongly separate behavior classes in heterogeneous logs
 - Evaluate on NCSA historical security data

Contribution: Can be used to more effectively detect advanced intrusions in clouds