Toward Intrusion Tolerant Clouds

Prof. Yair Amir, Prof. Vladimir Braverman
Daniel Obenshain, Tom Tantillo

Department of Computer Science
Johns Hopkins University

Prof. Cristina Nita-Rotaru, Prof. Jennifer Neville
Andrew Newell, Jeff Seibert

Department of Computer Science
Purdue University

Prof. David Evans
Ivan Alagenchev

Department of Computer Science
University of Virginia

http://www.dsn.jhu.edu

Clouds: The Promise

- Cloud computing: A new, cost-effective way to obtain information technology as a service.
  - In contrast to enterprise computing where IT is owned and managed by the user organization.
  - Public clouds: A service provider provides and manages infrastructure shared by many customers.
  - Private clouds: A service provider manages "separate" infrastructure for (large) organizations.
- Large potential advantages:
  - Economic benefits of scale.
  - Efficiency through specialization – amortization of expertise and experience.
  - "Pay per Use" instead of "Build for Peak".

Clouds: The Caveats

- Trend toward fewer, much larger systems.
  - Increased stakes: A cloud service problem affects a large number of users / organizations.
  - Increased complexity: A cloud is a very large distributed system, commonly residing in several data centers distributed globally over many different networks.
- The need for resiliency at scale:
  - Availability.
  - Reliability.
  - Security.

Lessons from LiveTimeNet

- LTN - a cloud networking service provider.
  - Broadcast-quality TV transport and delivery with global reach.
- Adequate scale and availability require two distributed system capabilities:
  - Consistent global state replicated across the network.
  - Distributed messaging system that connects cloud components.
- With these capabilities:
  - From cloud components to a cohesive system that manages itself autonomously.

Toward Intrusion Tolerant Clouds

- Our main premise:
  - The large gap in constructing resilient clouds is the vulnerability to intrusions.
  - No known algorithms to construct consistent global state and distributed messaging system at cloud scale, guaranteeing their integrity and performance even under intrusion attacks.
- Our main goal:
  - Invent, develop and transition the replication and overlay messaging tools necessary to make public and private clouds resilient to sophisticated intrusion attacks.

Introducing the Team

PURDUE UNIVERSITY


Cristina
Secure Distributed Systems


Jennifer
Machine Learning


Andrew


Jeff

UNIVERSITY of VIRGINIA


David
Trusted Systems


Ivan

JOHNS HOPKINS UNIVERSITY


Yair
Distributed Systems


Vladimir
Theory Algorithms


Daniel


Tom

Outline

- Clouds: The promise and the caveats.
- Project goal.
- Intrusion-tolerant replication.
- Diverse attack surface.
- Intrusion-tolerant messaging.
- Real-time attack detection, diagnosis and prediction.
- Validation and integration.
- Summary.

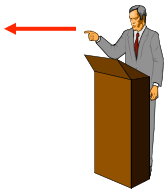

Intrusion-Tolerant Replication
Charting the State of the Art

- Byzantine Fault Tolerance (BFT) [CL99, YMVAD03, CMLRS06, MA06, KADCW07, more]
  - Good performance in fault-free scenarios, small -scale systems, eventual progress guarantees.
- STEWARD [ADDK+06, ACKL07]
  - Good performance in fault-free scenarios, Cloud -scale systems, eventual progress guarantees.
- Prime [ACKL08], Aardvark [CWADM09]
  - Good performance in fault-free scenarios, small -scale systems, performance guarantees under attack.

Byzantine Fault Tolerance

[CL99]

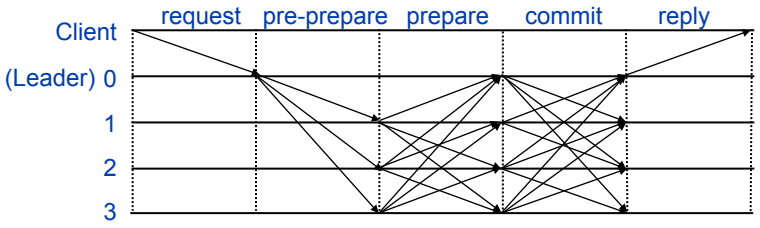

Performance Attack

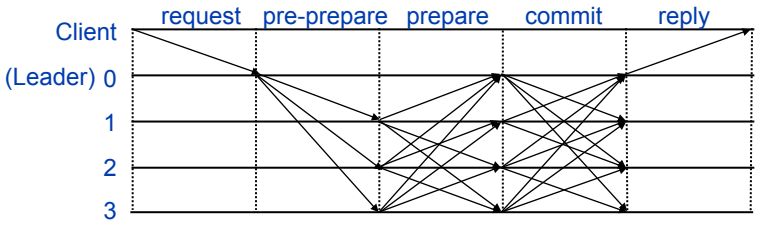

- Attack example: Pre-Prepare Delay
  - **Malicious leader can add delay into the ordering path by withholding its Pre-Prepare.**
  - Non-leaders maintain a FIFO queue of pending updates.
    - Use timeouts to monitor the leader.
    - Timeout placed on execution of first update in queue.
  - Malicious leader can stay in power by ordering one update per queue per timeout period!

Performance Guarantees Under Attack Prime

[ACKL08]

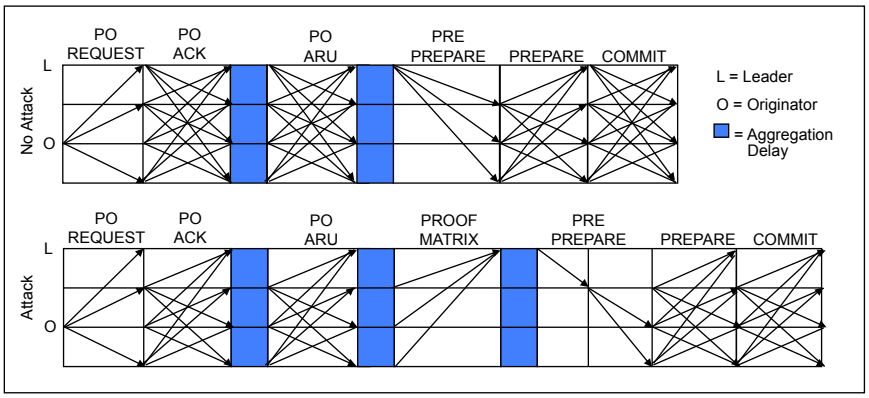

- Limiting the power of a malicious leader.
  - Bounded-delay performance guarantee.

Scalable Prime

- First scalable replication architecture with performance guarantees under attack.
  - Local replication protocol that executes within a cluster in a single data center.
  - Global replication protocol that executes between clusters in different data centers.
  - Threshold signature protocol that ensures validity of messages between the local and global levels.
  - Exchange protocol that disseminates global replication messages between the data centers.

Scalable Prime Replication Engine

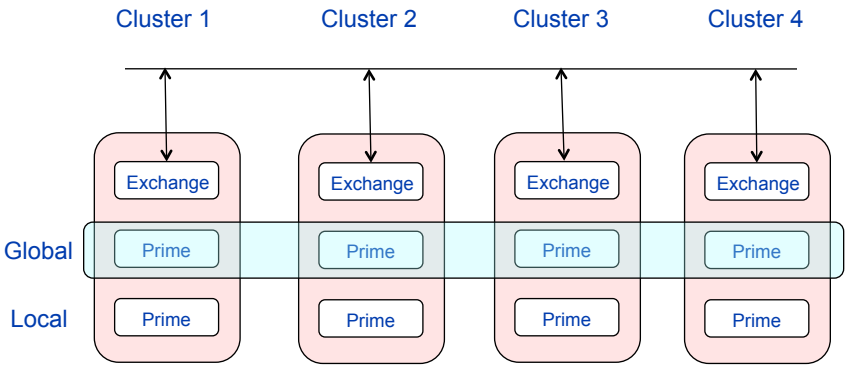

Intrusion-Tolerant Exchange

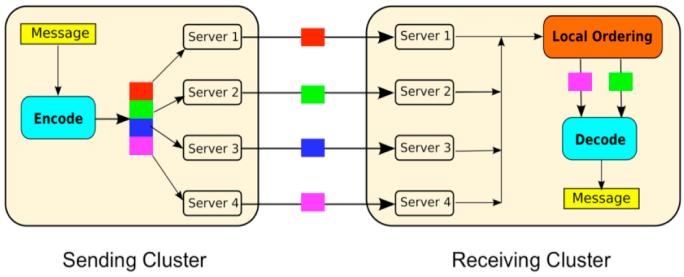

- Erasure encoding-based exchange.
  - Overcomes malicious servers while making efficient use of wide area network bandwidth.

Diverse Attack Surface

- The fundamental assumption:
  - Intrusion-tolerant replication protocols assume some maximum number of compromised replicas.
  - If all replicas are vulnerable to the same exploit, an attacker may be able to simultaneously compromise many of them, breaking that assumption and compromising the system.
- SS-Prime: Survivable Scalable Prime
  - Scalable Prime with diverse attack surface.
  - Engine-level diversity: automatically diversify attack surface of each replica.
  - System-wide diversity: deploy diversified replicas according to system-wide view (e.g., topology) to maximize protection.

Resiliency Through Diversity

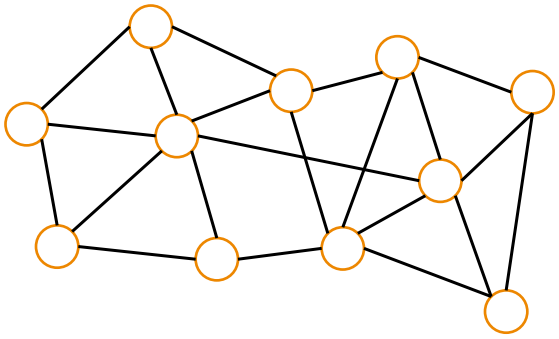

Engine-level diversity: each node has different attack surface

Protocol-layer diversity: different protocols between different endpoints

System-level diversity: deploy diversity with a system-wide view

Structured Diversity

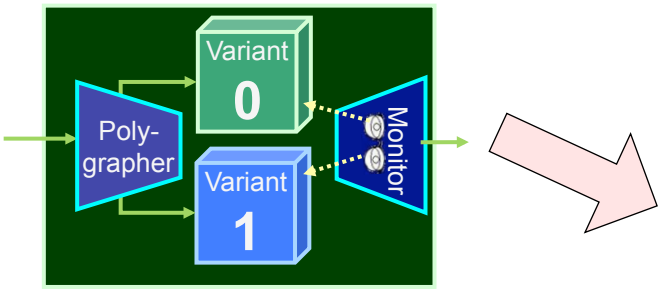

N-Variant Systems
[USENIX Sec 2006, DSN 2008]

Structured Diversity with limited synchronization

Outline

- Clouds: The promise and the caveats.
- Project goal.
- Intrusion-tolerant replication.
- Diverse attack surface.
- Intrusion-tolerant messaging.
- Real-time attack detection, diagnosis and prediction.
- Validation and integration.
- Summary.

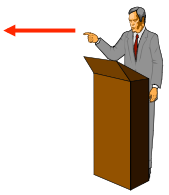

Intrusion-Tolerant Messaging

- No practical intrusion-tolerant messaging that can perform well in cloud environments.
- Underlying approach:
  - Using overlay infrastructure to limit the damage that can be inflicted by compromised nodes.
  - Diverse attack surface (similar to replication).
  - Willingness to pay "a lot" of overhead.
- Two overlay routing protocols
  - Controlled authenticated K-Paths routing.
  - Controlled authenticated flooding.

Controlled Authenticated K-Paths Routing

- Uses K node-disjoint overlay paths to disseminate a message from source to destination.
- Uses overlay topology and authentication to limit the power of compromised nodes.
- Can overcome K-1 compromised nodes.
- Overall cost about K times secure link-state routing.

Controlled Authenticated Flooding

- Uses overlay topology and authentication to limit the power of compromised nodes.
- Floods messages at most twice on each overlay link.
- Optimal intrusion tolerance.
- Optimal latency.
- High cost – up to 15-30 times higher than secure link-state overlay routing on relevant topologies.

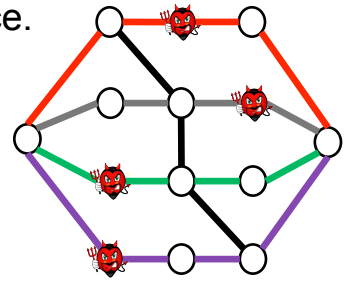

Outline

- Clouds: The promise and the caveats.
- Project goal.
- Intrusion-tolerant replication.
- Diverse attack surface.
- Intrusion-tolerant messaging.
- Real-time attack detection, diagnosis and prediction.
- Validation and integration.
- Summary.

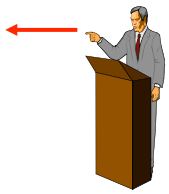

Real-time Attack Detection, Diagnosis and Prediction

Machine-learning approach based on *statistical relational learning*.

- Real-time collection of global cloud entities' log streams.
  - Based on intrusion-tolerant messaging.
- Real-time analysis of log streams to detect faults and attacks and predict their progress.
- *Adversarial statistical relational learning*:
  - Identifies relational inconsistencies stemming from falsified reports by compromised nodes.

Outline

- Clouds: The promise and the caveats.
- Project goal.
- Intrusion-tolerant replication.
- Diverse attack surface.
- Intrusion-tolerant messaging.
- Real-time attack detection, diagnosis and prediction.
- Validation and integration.
- Summary.

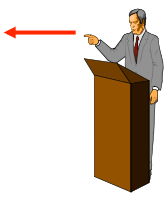

BAA: DARPA IO2 BAA 11-55, MRC

Proposal Title: Toward Intrusion Tolerant Clouds

Offeror: Yair Amir, Johns Hopkins University

Date: 11/4/2011

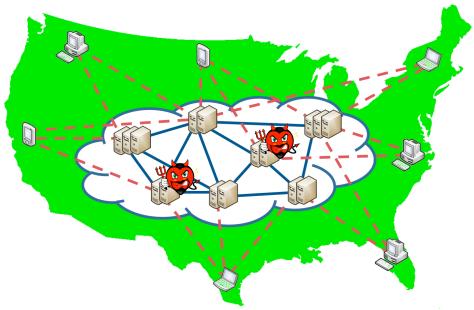

Key Innovations

Year 1+2:

First scalable intrusion-tolerant replication with performance guarantees under attack and a diverse attack surface.

First scalable tunable intrusion-tolerant overlay messaging.

Real-time statistical relational learning-based fault and attack detection, diagnosis and prediction.

Year 3+4:

First prediction-based reactive intrusion-tolerant replication with performance guarantees under attack.

First prediction-based reactive intrusion-tolerant messaging.

Dynamic, protocol-constraints-based and topology-based diversity for higher assurance against sophisticated intrusions.

Adversarial statistical relational learning overcoming falsified reports by compromised entities.

Expected Impact

Improved understanding of intrusion handling in large-scale distributed systems.

Open source intrusion-tolerant replication and overlay messaging engines that enable real-time continuous management, monitoring and control of public or private cloud infrastructure under sophisticated intrusion attacks.

The same engines also enable the construction of cloud applications that withstand sophisticated intrusion attacks.

Open source detection, diagnosis, and prediction engine enabling above intrusion-tolerant replication and overlay messaging engines to react to attacks in real time.

Intrusion-tolerant replication engine, overlay messaging engine and detection, diagnosis and prediction engine validated in an actual cloud environment.

Schedule, Cost, Deliverables & Contact

Diversified Replication & Messaging engines, Prediction Engine	Integration Validation Evaluation Analysis	Reactive Replication & Messaging engines, with holistic diversity	Integration Validation Evaluation Analysis

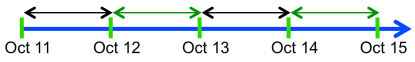

Johns Hopkins POC: Yair Amir, e-mail: yairamir@cs.jhu.edu
Purdue POC: Cristina Nita-Rotaru e-mail: crisn@cs.purdue.edu
UVA POC: David Evans e-mail: evans@cs.virginia.edu
With: Jennifer Neville (Purdue) & Vladimir Braverman (Hopkins)