Effective and Efficient Compromise Recovery for Weakly Consistent Replication

Prince Mahajan (UT Austin), Ramakrishna Kotla, Cathy Marshall, Venugopalan "Rama" Ramasubramanian, Tom Rodeheffer, Doug Terry, Ted Wobber (Microsoft Research Silicon Valley)

#### Scenario



















































#### Our Contributions

Polygraph: A framework that

- Extends weakly consistent replication
- Removes corrupted updates
- Recovers uncorrupted updates
- While being
  - Effective: Retain most uncorrupted updates
  - Efficient: Incur less bandwidth cost

#### Outline

- Motivation
- System Model
- Backup-based approach
- Polygraph: Effective and Efficient Recovery
- Results
- Conclusion

### System Model

### System Model

- Replicas can independently update items
- Each update produces a new version of item
- New versions propagate asynchronously

### System Model

- Replicas can independently update items
- Each update produces a new version of item
- New versions propagate asynchronously
- Replicas retain the most recent version
- Archive replica logs all received versions

### Example System



### Update Timeline



### Update Timeline



#### Threat Model

- Compromises can result from malice or misuse
- Corrupted versions
  - Versions injected by compromised replicas
  - Versions influenced by corrupted versions
- An external agent detects and reports compromises after the fact
- Archive and replication layer is not compromised

### Update Timeline



# Update Timeline (with compromise)



# Update Timeline (after recovery)




















## Drawbacks of Backup-based Approach

- Inefficient: Re-propagation from backup to replicas
- Ineffective: Updates subsequent to checkpoint are lost



- Innocent version identification
  - Effectiveness: innocent versions created postcompromise are recovered

- Innocent version identification
  - Effectiveness: innocent versions created postcompromise are recovered
- Replica-local retention

- Innocent version identification
  - Effectiveness: innocent versions created postcompromise are recovered
- Replica-local retention
  - Replicas retain innocent versions

- Innocent version identification
  - Effectiveness: innocent versions created postcompromise are recovered
- Replica-local retention
  - Replicas retain innocent versions
  - Effectiveness: newer versions recovered

- Innocent version identification
  - Effectiveness: innocent versions created postcompromise are recovered
- Replica-local retention
  - Replicas retain innocent versions
  - Effectiveness: newer versions recovered
  - Efficiency: retained versions save bandwidth

#### Innocent Versions

Version is innocent if it is

- Generated before compromise, or
- Not influenced by any corrupt version from the compromised replica



Archive Log



Archive Log









precompromise cut

#### Precompromise cut summarizes versions archived prior to compromise

 Sufficient to check: is the most recent version from the compromised replica that influenced v is corrupt?

- Sufficient to check: is the most recent version from the compromised replica that influenced v is corrupt?
- Each version has a *taint vector*
- Taint vector of a version v tracks the most recent version from each replica that has influenced v

- Sufficient to check: is the most recent version from the compromised replica that influenced v is corrupt?
- Each version has a taint vector
- Taint vector of a version **v** tracks the most recent version from each replica that has influenced **v**



- Sufficient to check: is the most recent version from the compromised replica that influenced v is corrupt?
- Each version has a taint vector
- Taint vector of a version **v** tracks the most recent version from each replica that has influenced **v**



A version  $\mathbf{v}$  is innocent if the influencing version from the compromised replica in  $\mathbf{v}$ 's taint vector is innocent









- Replicas retain innocent versions
- Key mechanism: innocence predicate

#### Innocent Version Identification at Archive



#### Innocence Predicate: Innocent Version Identification at Replica

















### Implementation

- Implemented on Cimbiosys
- Multiple compromises
  - Can recover simultaneously
- Multiple independent archives
  - Improves effectiveness, efficiency, and fault-tolerance

#### Evaluation

#### Evaluation

#### Setup

- I0 replicas and I000 items
- 3000 updates overall
- I000 updates after compromise
- Random updates and synchronizations
- 5 updates between synchronizations
# Evaluation

#### Setup

- I0 replicas and I000 items
- 3000 updates overall
- I000 updates after compromise
- Random updates and synchronizations
- 5 updates between synchronizations

#### Metric

- (in)effectiveness:
  lost items
- (in)efficiency: network overhead



Network transfers

Lost items





Lost items



Network transfers























## Conclusion

- In a weakly consistent system, Polygraph reverses the effect of corrupt updates
  - Effective: retains most uncorrupted updates
  - Efficient: recovery uses less bandwidth
- Implemented on Cimbiosys replication system