

Recovery from Intrusions in PaaS Clouds



Miguel P. Correia
(joint work with Dário Nascimento)

IFIP WG Meeting
Jan. 2015

Number of **critical applications** in the Cloud is
increasing



Number of **Intrusions** in these applications is
increasing



Motivation

Intrusions compromise:

- Integrity
- Availability
- Confidentiality

Intrusion/fault causes:

- Software flaws
- Configuration and usage mistakes
- Corrupted legitimate requests (e.g. SQL injection)

Motivation

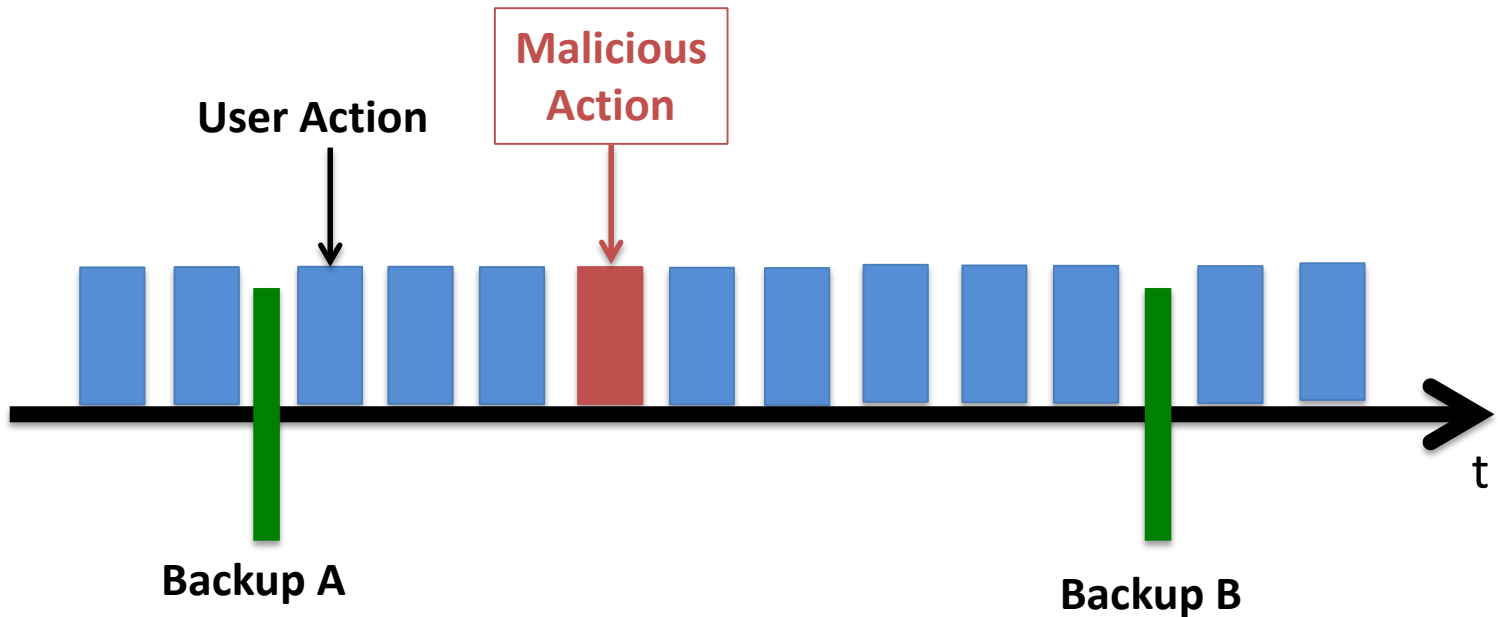
- Personal motivation:
- I've been working on masking faults and intrusions for 15 years
- Industry seems not to care
- Industry does care about recovering from intrusions/faults when they happen

Recover the application's integrity
when intrusions happen

Related Work

Backups

works but removes both bad and good actions



Related Work

Intrusion recovery: remove bad, not good actions

- Operating systems: Taser, Retro
- Databases: ITDB, Phoenix
- Web applications: Goel et. al, Warp, Aire
- Others: Undo for Operators

Limitations:

- All require setup and configuration
- Max. complexity: 1 app server, 1 database instance
- Cause application downtime during recovery

Platform as a Service (PaaS)

- Cloud service = to run applications
- Consumer develops application to run in that environment, using
 - Supported languages, e.g., Java, Python, Go, PHP
 - Supported components, e.g., SQL/noSQL databases, load balancers

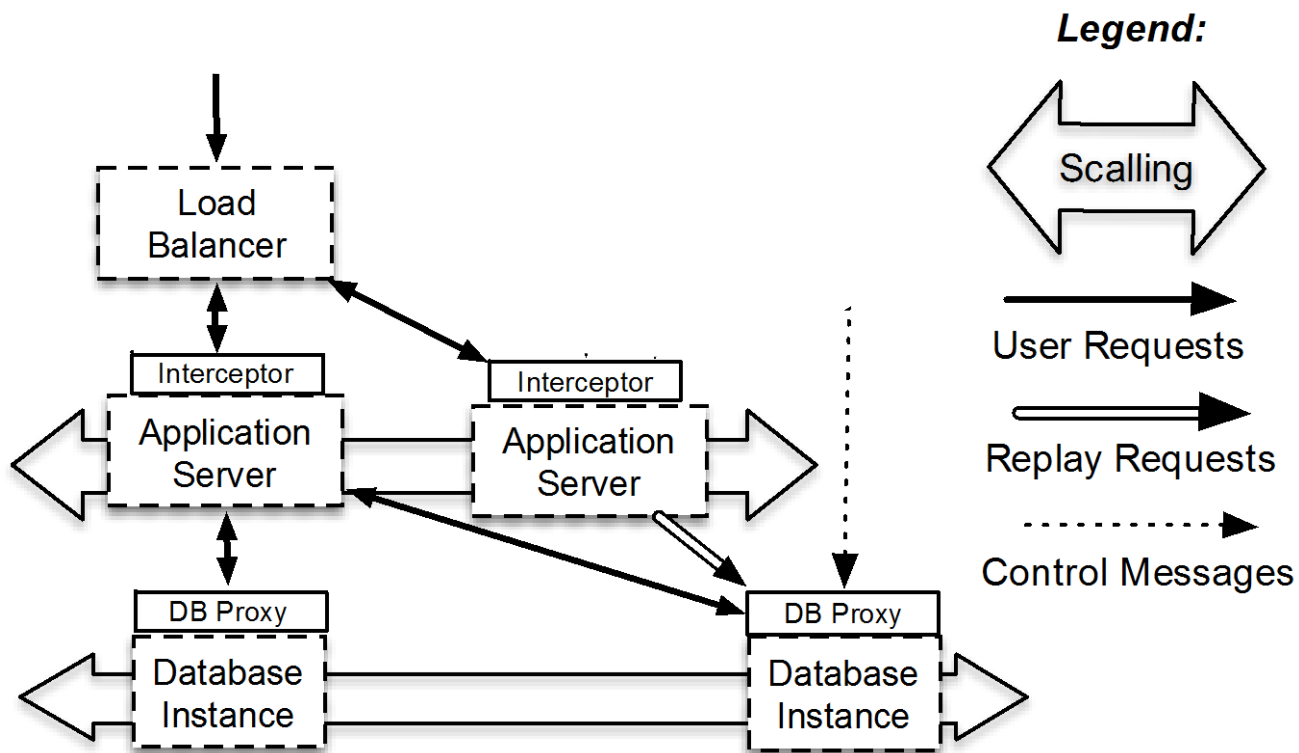
Objective

Intrusion recovery system for PaaS

- Supported by the PaaS: available without setup
- Remove the intrusion effects
- Support applications deployed in various instances
- Avoid application downtime
- Cost effective
- Recover fast

Shuttle

User requests



Replay Process

1. Identify the malicious actions
2. Start new application and database instances
3. Load a snapshot previous to intrusion instant
Create a new branch
4. Replay requests
Database operations shall replay in same order as original
5. Block incoming requests; replay last requests
6. Change branch

Replay Modes

	Full-Replay	Selective-Replay
1 Cluster (Serial)	✓	✓
Clustered	✓	X

Full-Replay: Replay every operation after snapshot

Selective-Replay: Replay only affected (tainted) operations

Serial: Replay all dependency graph sequentially

Clustered: Independent clusters can be replayed concurrently



Environment

- Amazon EC2, c3.xlarge instances, Gb Ethernet
- WildFly (formely JBoss) application servers
- Voldemort database
- Ask Q&A application; data from Stack Exchange

Accuracy with intrusion scenarios:

1. Malicious requests
2. Software vulnerabilities
3. External channels (e.g. SSH)

	#tampered intrusion	#tainted	#replayed (selective rep.)	#replayed (full replay)
1a	110	0	[0, 605]	> 38 620
1b	58	14	[0, 379]	> 38 620
1c	48	52	[0, 253]	> 38 620
2a	4 338	0	-	> 38 620
2b	18 286	1 278	-	> 38 620
3	> 2 000	-	-	> 38 620

Performance overhead

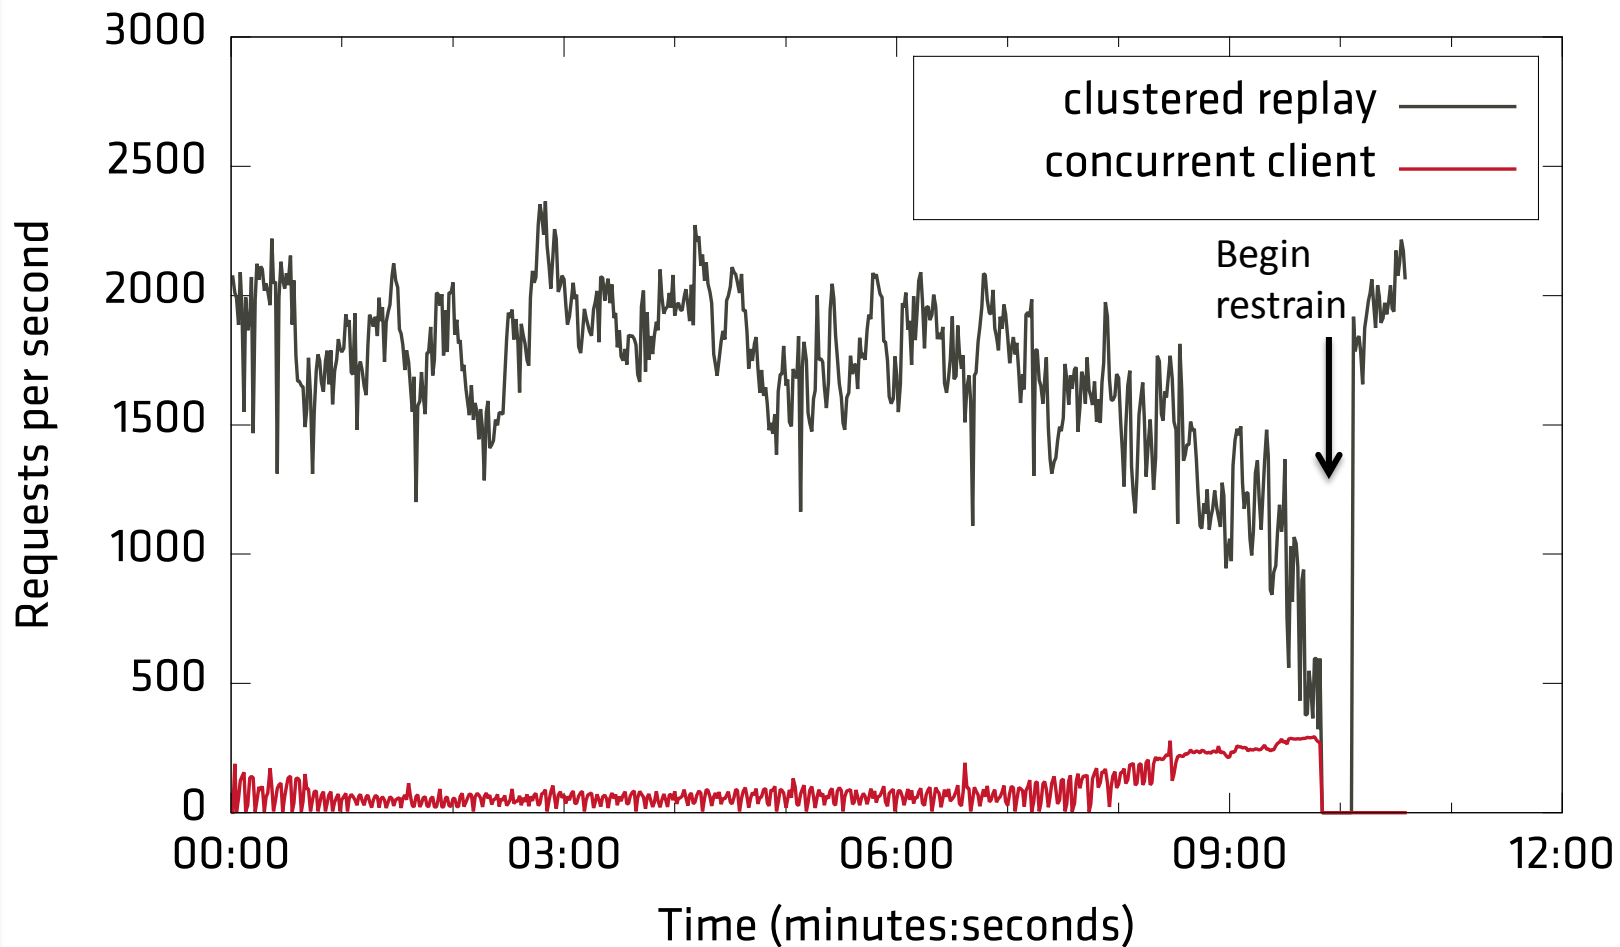
in normal execution

	Workload A	Workload B
Shuttle	6325 ops/sec [5.78 ms]	15346 ops/sec [3.62 ms]
No Shuttle	7148 ops/sec [5.07 ms]	17821 ops/sec [3.01 ms]
overhead	13% [14%]	16% [20%]

Recovery time

1 million requests

Restrain duration



Restrain: 46 seconds

Storage overhead

for 1 million requests

	# objects	size (MB)
Shuttle Storage:		
Request	1 million	212
Response	1 million	8 967
Start/end timestamps	2 million	16
Keys	137 million	488
Total		9 684
Database node:		
Version List	14 593	1.4
Operation list	9 million	277
Total		282
Manager:		
Graph	1 million	718

Conclusion

- New intrusion recovery service to be integrated in PaaS offerings
- Supports applications running in various instances backed by distributed databases
- Leverages the resource elasticity and pay-per-use model to reduce the recovery time and costs