



Aeron 

Hot deploying low-latency services for 24/7 operation

Martin Thompson - @mjpt777

***Who would like to release
software more often?***

***Do you need scheduled
downtime to release?***

Downtime Cost

\$5600 per minute average

\$2000 - \$9000 per minute range

- Gartner & Avaya (2014)

Downtime Cost
Average \$9000 per minute

- Ponemon Institute (2016)

***How often do you
security patch?***

Are you renting servers?

Quality Matters

NFRs ***(Non-Functional Requirements)***

NFRs



Wat!

24/7

24/7 Operations

1. Pervasive Devices

24/7 Operations

- 1. Pervasive Devices**
- 2. Globalisation**

24/7 Operations

- 1. Pervasive Devices**
- 2. Globalisation**
- 3. Competition**

High-Availability

Uptime

***Does it include
scheduled maintenance?***

High-Availability

1. **MTBF** - (Mean Time Between Failures)

High-Availability

1. **MTBF** - (Mean Time Between Failures)
2. **MTTR** - (Mean Time To Recovery)

High-Availability

1. **MTBF** - (Mean Time Between Failures)
2. **MTTR** - (Mean Time To Recovery)
3. **RTO** - (Recovery Time Objective)

High-Availability

1. **MTBF** - (Mean Time Between Failures)
2. **MTTR** - (Mean Time To Recovery)
3. **RTO** - (Recovery Time Objective)
4. **RPO** - (Recovery Point Objective)

**Why we cannot have
nice things?**

State!

It's always the state

Replicated State Machines

1

2

3

4

5

6

...

X

State

1

2

3

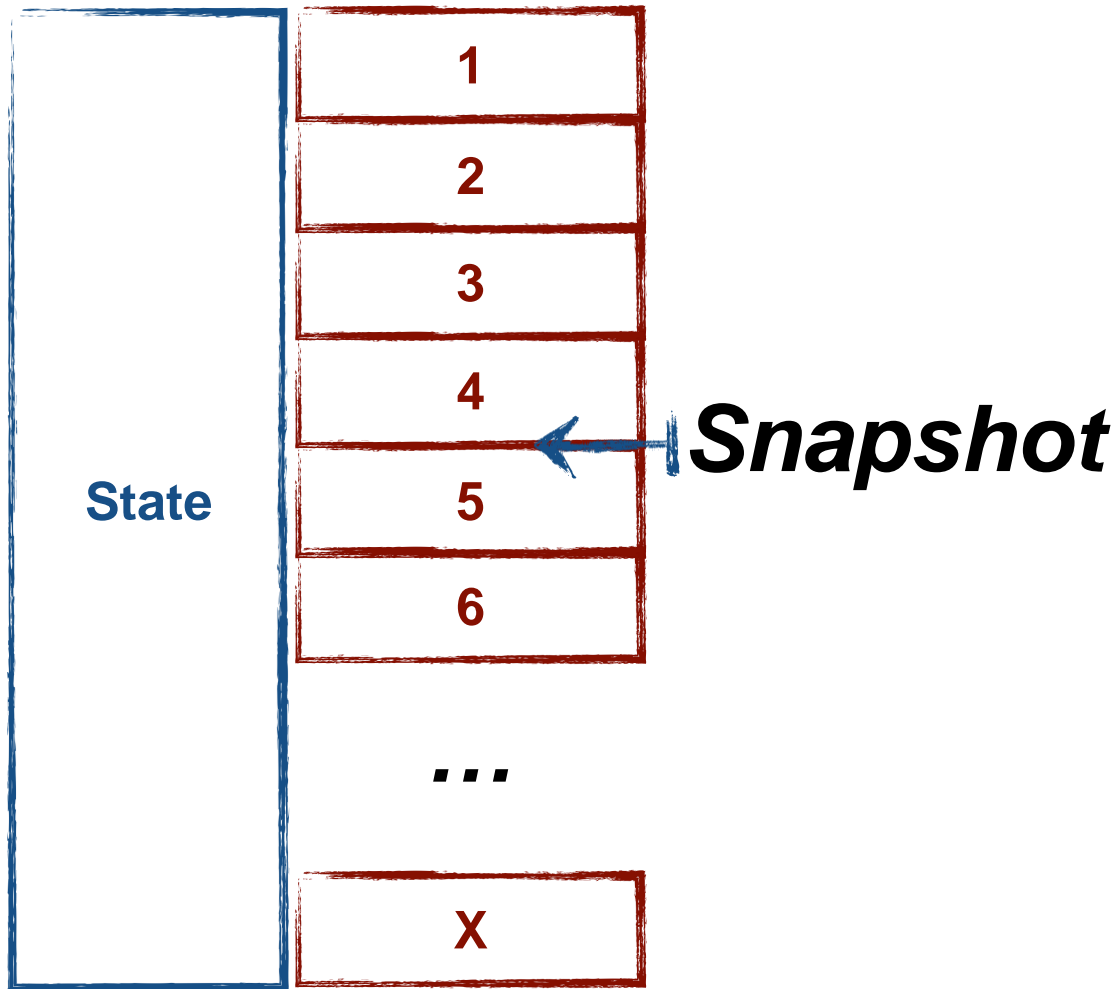
4


5

6

...

X





Replicate log of inputs

Consensus & Fault Tolerance

**In low-latency systems you need
low-latency fault tolerance**

Rolling Upgrades

**The “easy” bit –
upgrading/patching a platform**

**Node Recovery =
Load Snapshot + Log Replay**

How to maintain low-latency?

Are 3 nodes enough?

**What platform support is needed
for hot upgrade?**

What about my services?

We want to be releasing often

Only react to what you know

Version your messages

Version your session protocols

Version your feature flags

**Once something is out there it
can never be taken back**

Be strict with encapsulation

Handover/resumption protocols

Disaster Recovery

Thermal Events

RPO

RTO/MTTR

Warm Standby Nodes

What helps at extremes?

Multicast

Async all the things!

Background long running tasks

Version version version!

Aeron

<https://aeron.io/>

<https://github.com/real-logic/aeron>