

A new world order

Slow ≈ Byzantine

In most modern systems, users perceive:

"slow is the new down."

In most distributed systems:

"slow is indistinguishable from byzantine operations."



We had to be "very sure" in the

Days of Failover

Primary : Replica system **usually** have non-zero operational costs in performance failover.

- dataloss (in asynchronous systems)
- operational downtime
- operational rebuild time (reversing the flows)





For well-designed, available systems,

Constraints Have Changed

Deciding to fail a node is no longer a

"last resort" decision.



What do I mean by

well-designed?

The failure of a node does not cause

- service interruption
- significant performance regressions

The recovery of a node does not cause

- unnecessary work (only minimal replay)
- significant performance regressions







- * Need: zero-downtime
- * Know: Agreement is hard.
- * Know: Consensus is expensive.
- * CAP theorem tradeoffs suck.
- CRDT (Commutative Replicated Data Type)

Snowth design



















A look at adaptive algorithms in

Replication

How do you choose the right unit of work for tasks?



What does it sound like when a system

Backfires

Batch it faster than single ops

- less latency impact
- less transactional overhead

What with QoS enforcement & circuit breakers?

Flogging

TCP (and everything else) can teach us something.



This provides us

Opportunities

What if we had relative homogeny of systems and workloads?







| * | Jul 02 2014, 10:30 (5M) |
|---|-------------------------|
| | 1.1150449523579k |
| | 28.010637338557k |
| | 15.4763937582k |
| | 3.1138964814012002k |
| | 1.0534715994518k |
| | 2.3435219867851003k |



New things become possible

Predicting Future Conditions

With higher volume data,

statistical models offer higher confidence.





It takes good understanding of statistics to ask the right questions.

Misleading yourself

This is a q(0.99) - 99th percentile.

It obviously goes off the rails around 1am.

No.

We have a new tool in the tool chest:

Intentionally Failing Nodes

When nodes are cattle, not pets...

Expect more from you systems.

Thank You

You can observe better,

know more,

don't settle.

