



ALTEERVE

Uptime Made *Easy*

SIMPLE

STABLE

TRUSTED

AVAILABLE

About Us

Hi! ^_^

- Focused on resilience for 20 years.
- Specialized in availability clusters for 9.
- Virtual server availability specialists.
 - It is all we do.
 - Users in research, defence, marine shipping, medical imaging and industry.



ALTEEVE

Debunking The Myths of High-Availability

1. The cloud is highly available
2. Availability requires quorum
3. Nodes and a SAN is HA



1. The cloud is not HA

What goes up must come down

- AWS is the gold standard
 - SLA is 99.9% uptime monthly. That is 8h 45m downtime/year
 - 10% credit between 99.0~99.9%
 - 25% credit below 99.0%
- The cloud is focused on Resource Utilization Efficiency (RUE)



ALTEEVE

The Cloud Is Scalable, Not HA

What goes up must come down

- Interdependent services
 - Actual uptime drops as more services used.
 - As of 2017-03-07, Gartner's CloudHarmony reported YTD outages of 205min for AWS, 740min for Azure and 11min for Google^[1].

1. <https://www.theinformation.com/how-aws-stacks-up-against-rivals-on-downtime>



ALTEEVE

2. 3+ Nodes or Bust?

No

- Hello my peer, I miss you.
- Avoiding “split-brains”
 - Quorum doesn't help!
- Fencing; Shoot first, ask questions later
 - It's the only way to be sure

https://www.alteeve.com/w/The_2-Node_Myth



ALTEEVE

Quorum Confusion

The real world is messy

- Case study; Node hung during write
 - Peer has quorum; locks reaped, journals replayed, operation resumes.
- Peer recovers, no concept of time, no reason to check locks or quorate state.
 - Boom goes your data.



ALTEEVE

Apples and Oranges

Scratching different itches

- Quorum is optional
 - A tool for when things are working.
- Fencing is not!
 - A tool for when things go sideways.



ALTEEVE

But I don't have shared storage!

>_>

- If you can run a service in different places at the same time without coordination, then do so, you don't need a cluster.
- If you can't do that, you need fencing.

ALWAYS USE FENCING!



ALTEEVE

3. I have a SAN, I have HA!

It's a start, but no

- My vendor said it's fully redundant, so I am good.
 - If it's one box, it's a single point of failure.
 - Bad firmware update, human error, etc
 - Repairs require a maintenance window
- Your data must be in two places at the same time
 - Those places must be mechanically and electrically isolated.



ALTEEVE

The Foundation Under The Nodes

Yes, you need full stack redundancy

- DC-sized UPSes are not perfect
 - Human error, firmware bugs, we've seen it all
- Network redundancy
 - Can you replace a switch in the stack without interruption?

“But That Won’t Happen”

Yes, yes it will

- The moment you think you’ve thought of all ways things can fail, Murphy will find you. He will have pie, he will make you eat it and it will be bitter.
- “That’s never going to happen”.
 - Yes, yes it will. Just to prove you wrong.



ALTEEVE

Questions?

mkelly@alteeve.ca

<https://alteeve.ca/w/Support>

freenode: #clusterlabs



ALTEEVE