


IASO: A Fail-Slow Detection and Mitigation Framework for Distributed Storage Services

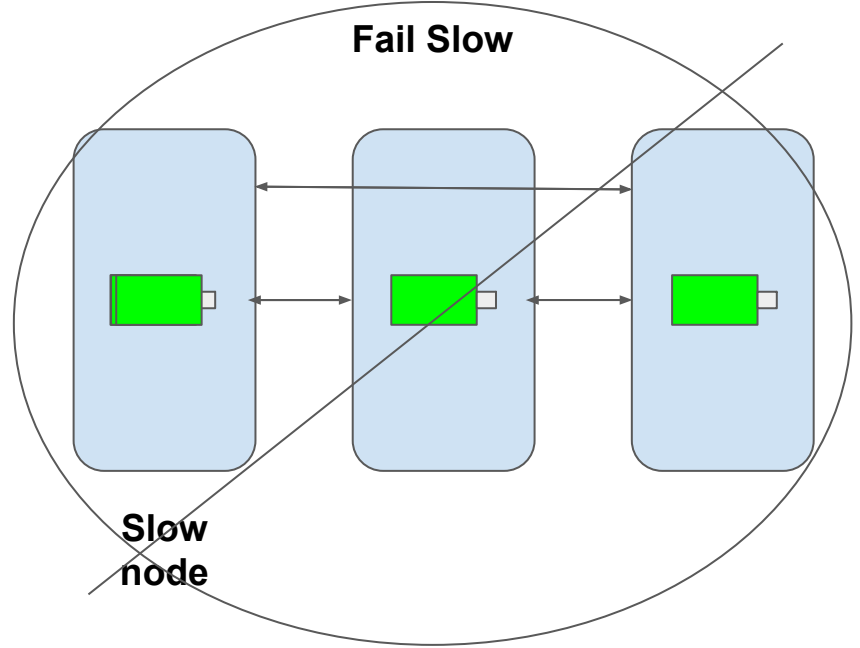
Biswaranjan Panda (Nutanix)
Deepthi Srinivasan (Nutanix)
Huan Ke (University of Chicago)
Karan Gupta (Nutanix)
Vinayak Khot (Nutanix)
Haryadi S. Gunawi (University of Chicago)

Fail Stop



Stopped
node

Fail Slow



Slow
node

*“We present a study of **101 reports** of fail-slow hardware incidents, collected from large-scale cluster deployments in **12 institutions**.”*

Fail-Slow at Scale: Evidence of Hardware Performance Faults in Large Production Systems. In Proceedings of the 16th USENIX Symposium on File and Storage Technologies (FAST), **2018**.

*“The dataset contains **232 validated cases** collected from the deployment of **39,000 nodes** throughout a period of **7 months**”*

*“We found that the fail-slow **annual failure rate** in our field is **1.02%**”*

IASO: A Fail-Slow Detection and Mitigation Framework for Distributed Storage Services

Categories	Root causes
Hardware	Faulty dimm (15), ECC error (10), low memory (9), SATADOM (5), CRC error (1), RAID controller (1), LSI controller (1), unknown (5)
Software	Software upgrade (8), VM issue (6), GC (3), BIOS (1), scheduler (1), unknown (6)
Network	Faulty device (13), network outage (9), device replace(7), unreachability (6), packet drop (5), network contention (2), device reboot (1), unknown (18)
Environment	Incorrect setting (11), high load (1), energy issue (1)
Human error	Misconf (10), network migration (4), install /deploy (3), unplugged cable (2), unknown (4)

*“We benchmark **five cloud systems** (Hadoop, HDFS, ZooKeeper, Cassandra, and HBase) and find that limpware **can severely impact** distributed operations, nodes, and an entire cluster.”*

Limplock: Understanding the Impact of Limpware on Scale-Out Cloud Systems. In Proceedings of the 4th ACM Symposium on Cloud Computing (SoCC), 2013.

IASO

A **peer based non intrusive**

fail-slow detection & mitigation **framework**

<p>Deployment</p> <p>1.5 years</p>	<p>Scale</p> <p>39000 nodes</p>	<p>Effectiveness</p> <p>10 minute bound 3.7% False Positives</p>
------------------------------------	---------------------------------	--

ATC 2019, 12:20 PM, Track 1 on
July 10, 2019