Object Storage for Storage Preservation

Designing Storage Architectures 2017

Data Center

Network

Cloud

Deepak Jain
Founder, AiNET
deepak.jain@ai.net

Challenge: Storage Preservation

- Growing Amounts of Data
- Unlimited Preservation Lengths
- Constrained Budgets
- Rapid Evolution of Technology
- Silent Corruption (Bit Rot)

Silent Corruption

Traditional Storage (RAID, clustering, etc) does not address silent corruption

Silent Corruption is data decay over time due to:
hardware, software, radiation and other errors (transient or persistent) that result in very small errors (bit flips) being spread across a data store.

As data preservation times increase, silent corruption increases

As data sizes increase, silent corruption increases

As copies and moves to verify data / manage hardware / media cycles increase, silent corruption increases

Over time, silent corruption increases

Silent Corruption

How do you address silent corruption?

Cryptographically Strong Fixity!

With all the other challenges, that sounds expensive and difficult….

Opportunity: Use Object Storage's advantages to address Preservation

ZSTOR

Use ZSTOR Object Storage to address

- Multi-Tenancy (1000s of accounts)
- Scalability (Exabytes and beyond)
- Cost and Energy Efficiency
- Massive Throughput
- Resilience, Availability, Recovery
- DARE – Data at Rest Encryption
- DIME – Data in Motion Encryption

Problem: Traditional Object Storage is not enough

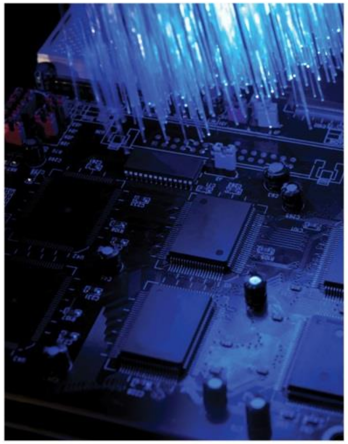

Does not automatically address:

- Verification
- Fixity
- Malware/Virus
- Deletion Protections
- At Rest Data Integrity (cryptographic checksum of objects)
- 3rd party Separated Witness (SEPWIT)

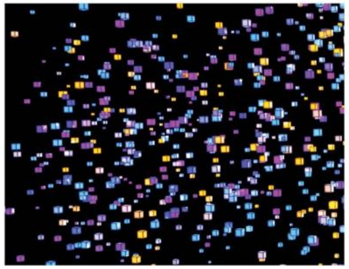

Solution: SANTA

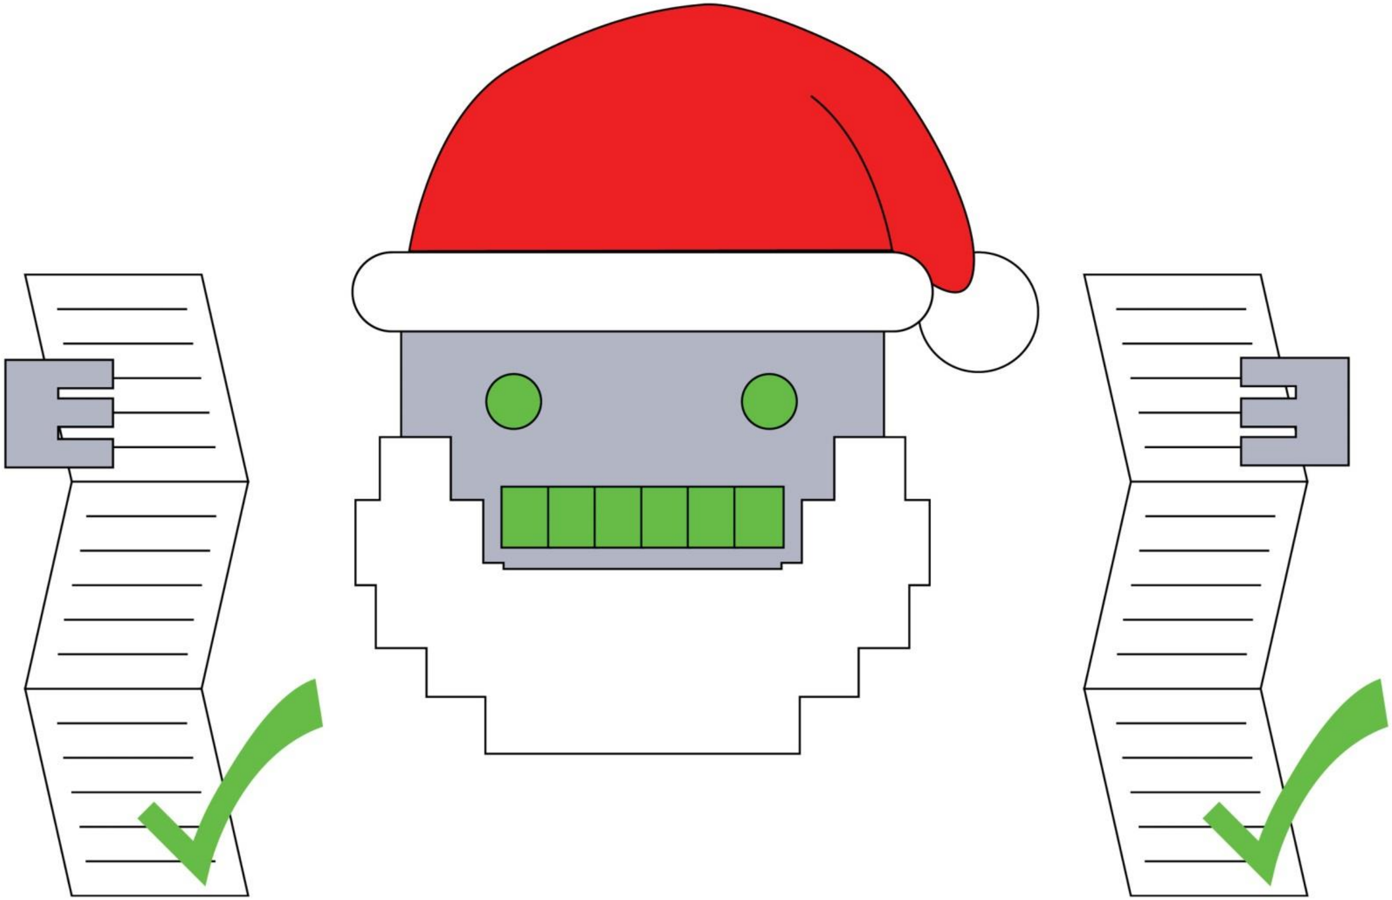

Storage and Network Trust Arbitrator

Why SANTA?

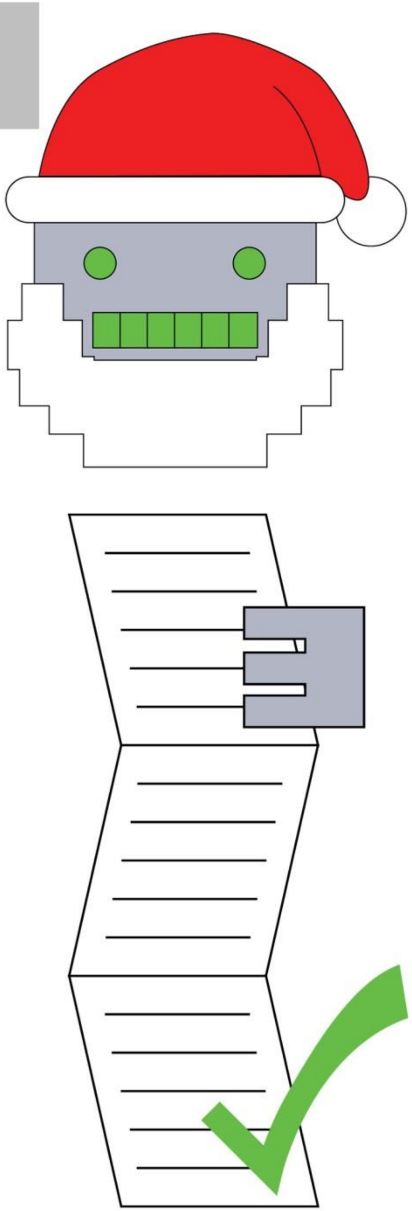

He makes a list, checks it twice – identifies bad hardware, data, processes and quarantines them, isolates them and prevents their propagation.

He assumes everyone is naughty – SANTA doesn’t trust checksums, or signatures. He verifies everything each time data is read or written to make sure it started and ended exactly as it should (on top of ZFS and other protections)

Very fast algorithms available for cryptographic acceleration: SKEIN, SHA-3 (NIST), BLAKE, SHA512. Up to 80% faster than SHA-256 and stronger. Salted hash prevents collisions.

At scale, less than 1 defect in 2^600 operations (not bytes – each opp could be multiple MB). This means error free exabytes for every blade of grass on the planet for centuries.

ZSTOR with SANTA overview

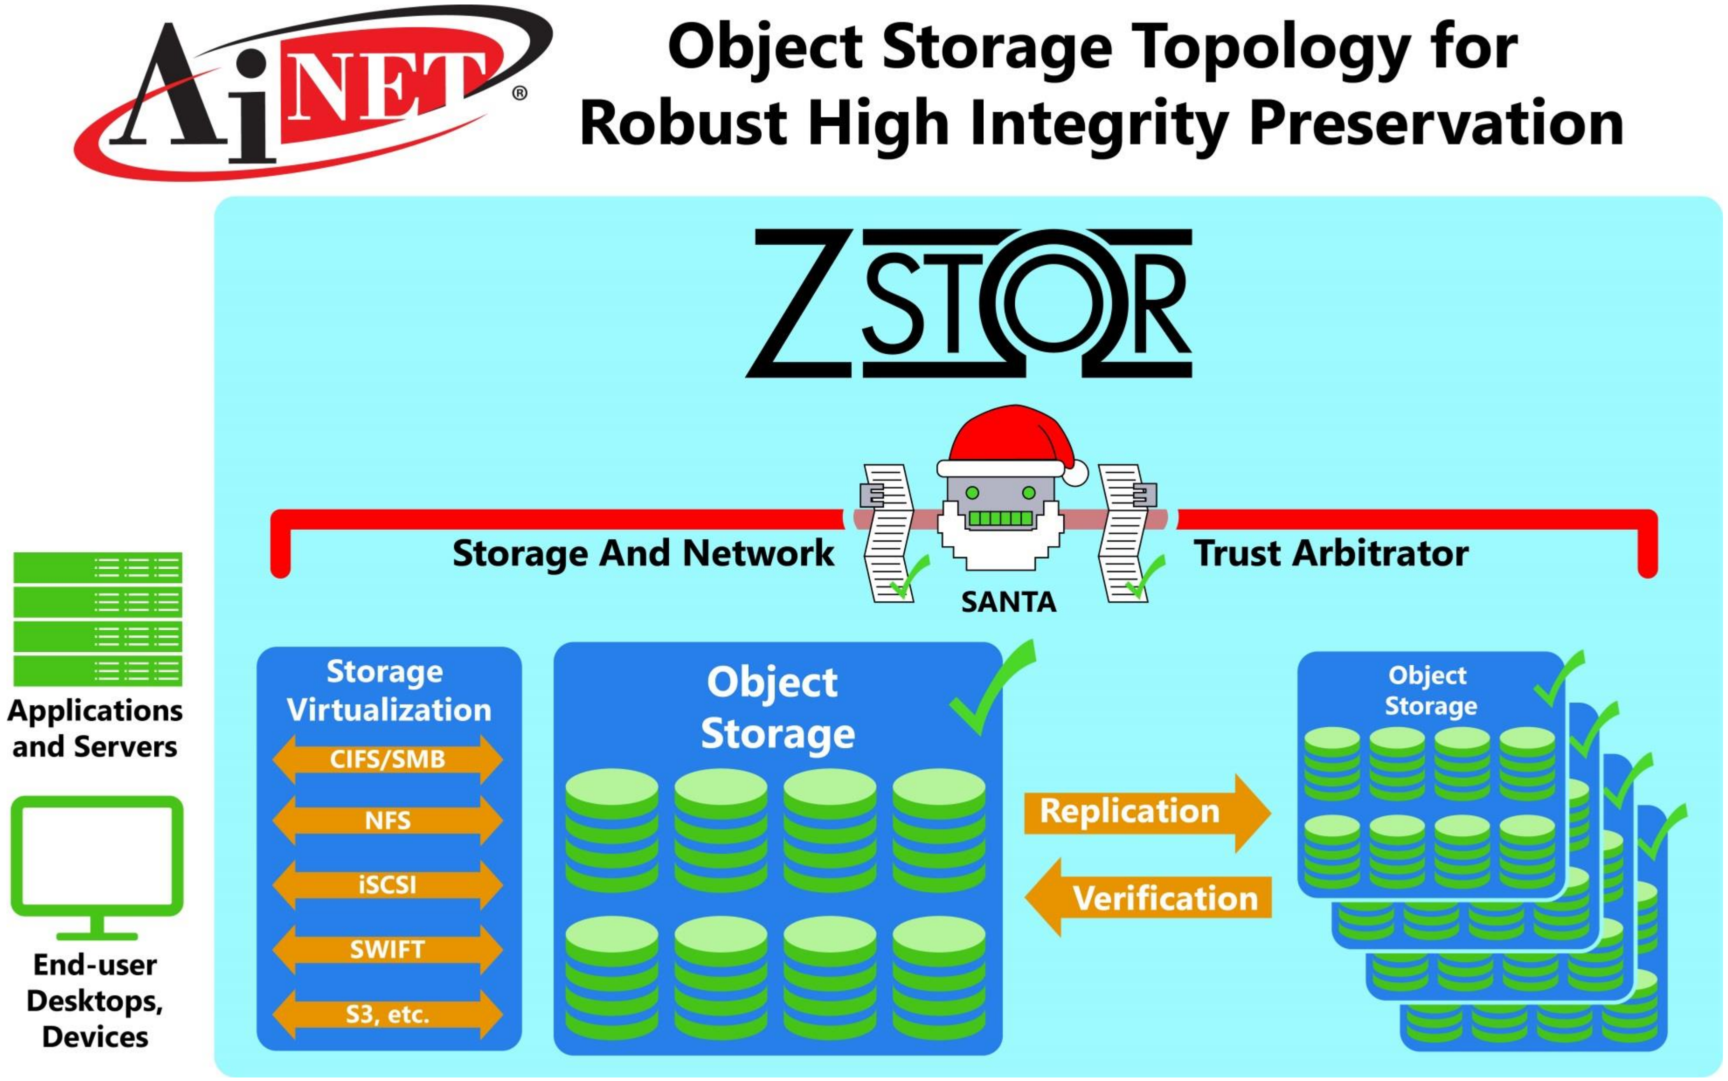

Object Storage Topology for Robust High Integrity Preservation

How does this all work? ZSTOR + SANTA

- All files are objects, all media types are an object
- All of object storage advantages
- Eliminates all object storage disadvantages

- Multiple physical replication locations
- ZFS to each drive
- Each drive is powered down when not in use
- Drives are powered on as objects are accessed or being verified

- Unlimited Metadata per object (author, ownership, copyright data, EXIF, etc).
- Metadata contains cryptographic signatures of each verification pass including data, time, and hardware/software identifiers

- Soft and Hard delete functions (think Recycling Bin + Dumpster)
- Storage Virtualization abstracts object API from traditional file system.
- SANTA arbitrates and supervises every operation

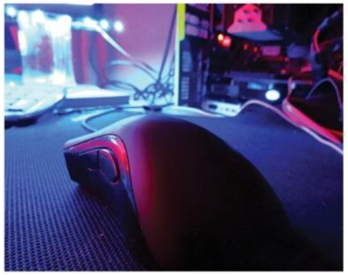

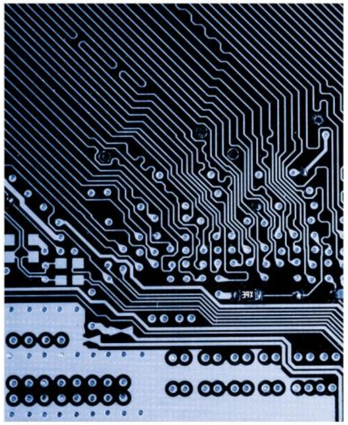

More Info

www.ai.net/santa

Questions?