
 The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.

# The Impact of SSDs on Software Defined Storage

Barbara Murphy  
VP of Marketing, Weka.IO  
barbara@weka.io



# Company Overview

 The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.

- Founded in 2013, R&D in Israel
  - DNA in Storage - XIV lead architects
- US HQ in San Jose, CA
  - DNA in Flash - Michael Raam, SandForce,
- \$32.25M raised to date
- Hyper-converged file storage
- Targeting Web 2.0, Rendering, Life Sciences, EDA and HPC






# All Storage is Software Defined

✘ The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.

## Software Defined

Vs.

## Software *Designed*



✘ The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.


✘ The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.

✘ The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.

✘ The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.



# Built to Last or Built to Fail?

 The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.

## Software Defined

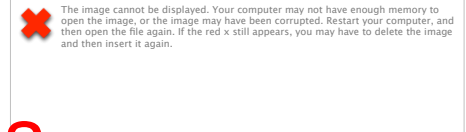
- Homogenous
- High durability components
- Dedicated networks
- Failures are exceptional
- Tight bill-of-materials
- Design to slowest component

## Software Designed

- Heterogeneous
- Cheapest components
- Poor networks
- Failures are normal
- Latest off the truck
- Design to the average of all




# What Modern *Design* Needs



- Designed for tomorrow's IT environments
  - High speed networks
  - Flash first
  - Hot and cold tiers
  - Commodity "unreliable" servers
  - Virtualized agile deployment
  - Cost containment




# *Designed* for Disk or Flash?

 The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.

- Everything changed with SSDs
  - Data Structures
  - Power failure
  - Endurance
  - Reliability
  - Latency





# ZFS – Designed for Disk


 The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.


“ZFS...is designed with a focus on data integrity by protecting the user’s data **on Disk**” Wikipedia


- B Tree
- Write Amplification
- Disk Caching
- Disk Scrubbing
- RAM
- Copy-on-write


 The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your

 The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and

 The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x


 The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.

 The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then

 The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the



## SDS *Designed* for Scale

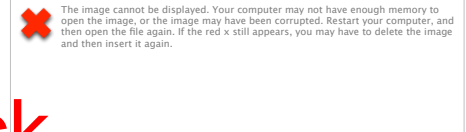
 The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.

- Triple Replication Costs
  - 3x cost, 1/3 life
  - Impacts power, network traffic, footprint, failures,
- Dedupe and Compression
  - A compensation - not a solution
- Erasure coding reduces overhead by over 80%
  - Properly designed to minimize write amplification





# Latency – the New Bottleneck

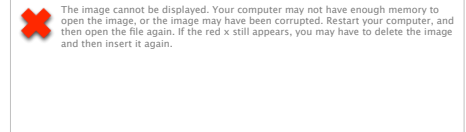


- IOPS - Solved
  - Bottleneck moved to software, controllers and network
- Legacy stacks designed for disk latencies
  - Requires I/O path efficiencies
- Is NVMe the answer?
  - Requires non-standard network settings
  - May be relegated to rack level deployment



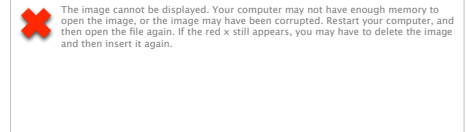
# SSD for the Cloud

- Virtualized storage
  - Expand or contract your infrastructure, add performance on demand
- The larger the scale, the more you will fail
  - Triple replication or N+2 is not sufficient
- File intelligence, block speed
  - POSIX, REST, HDFS
- Tight coupling to leverage cloud
  - Upstream into orchestration environments (Openstack)
  - Downstream to data lake (Disk/archive)





# Summary

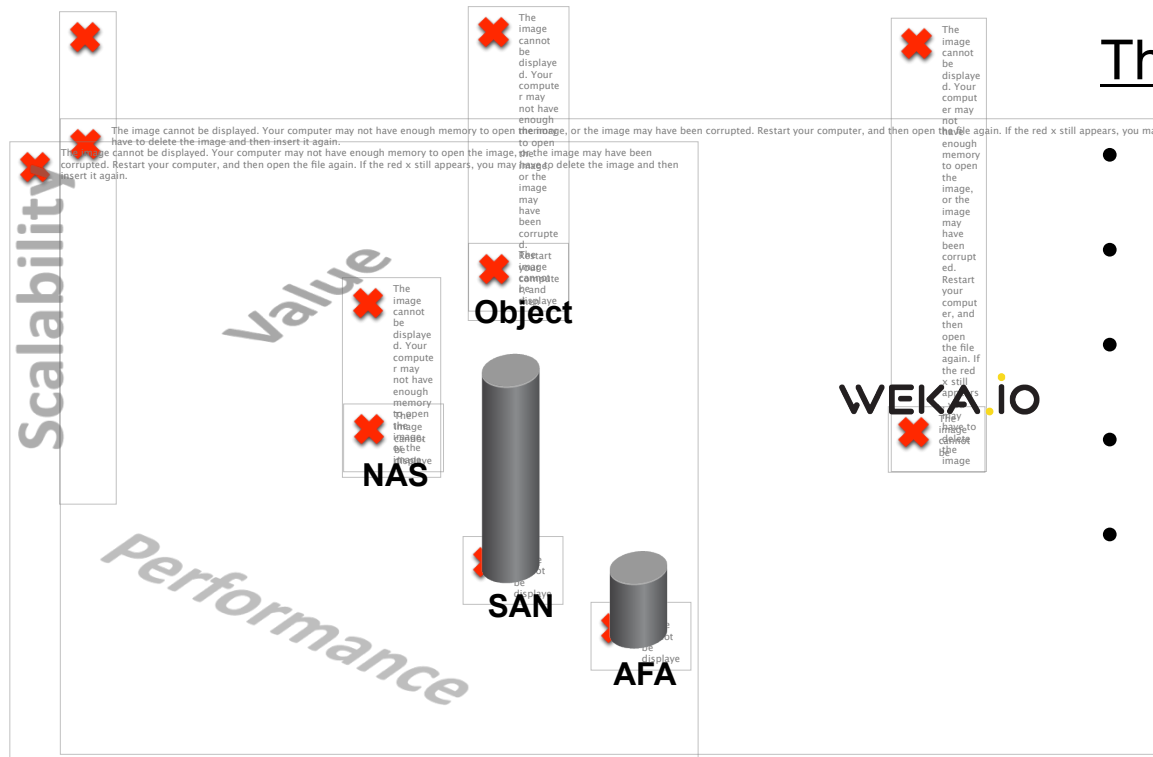


- SSD has a foundational impact on SDS
  - Data structures, I/O path, resiliency, endurance, cost
- Legacy file systems will not solve the SSD challenge
- Next generation file systems:-
  - Flash centric
  - Virtualized, tiered, scalable, modern interfaces



# Performance, Scalability and Value!

✘ The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.



## The Weka.IO File System

- 100% SSD *Designed*
- Performance of Flash
- Intelligence of File
- Latency of SAN
- Cloud scale and economics



Peta

Exa

Zetta

Yotta

Xona

***Weka*** ( $10^{30}$ )

**Thank you**



The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.