ATRATO™

ACCESS. THE REVOLUTION.

Efficient Data Storage

MSST 2008

Atrato, Inc.

10955 Westmoor Drive, Ste. 300

Westminster, CO 800021

www.Atrato.com

Presenter:

Thomas M. Ruwart

Chief Scientist

612.850.2918

tmruwart@atrato.com

Overview

- Atrato
- Efficiency: Physical Space, Energy Usage, and Performance in general
- Physical Space
- Storage Energy usage
- Secondary energy usage
- Performance
- Storage design considerations
- Why move to smaller disk drives
- Design Considerations
- Conclusions

Atrato

- Atrato was founded on the idea of addressing *Data Access* as opposed to just capacity and/or bandwidth
- Startup formed in January 2004
- Formerly Sherwood Information Partners, Inc., name changed Feb 2008
- Based in Westminster, CO
- Focused on
  - Self-maintaining Array of Identical Disks (SAID)
  - High-density packaging of small-form-factor commodity disk drives
  - Highly scalable storage controller – *Avenger*
- Atrato, Inc. is named after the Rio Atrato in Colombia. The Rio Atrato discharges at least 175,000 cubic feet (5,000 cubic m) of water per second making it the fastest river in the world.

Generic Presentation

- This is NOT a commercial for Atrato
- It is intended to demonstrate some non-intuitive results of the application of small form-factor laptop-class disk drives in a massively parallel array
- Atrato and a former division of Seagate are the only two companies working on this at the time of this writing

Physical Space

- Disk drive Form Factors are 3.5-inch low profile and 2.5-inch laptop
- 3.5-inch = 147mm x 102mm x 26mm
- 2.5-inch = 70mm x 40mm x 9.5mm
- Approximately 5.85:1 2.5-inch disks to 3.5-inch disks in terms of physical volume
- Practical packaging of 2.5-inch drives easily supports a volumetric ratio of 4:1
- Fail-in-place packaging model can boost the volumetric ratio to 10:1
  - 160 disk drives in a single 3RU enclosure versus 16

Energy Usage in General

- Data Center Power Consumers
  - Processors The principle consumer in a data center – anywhere from 50% to 90%
  - ***Storage Devices – about 20%***
  - Networks
  - Cooling units
  - Power Distribution units
  - Displays and Misc
- Focus here is on Storage Energy Usage, specifically disk drives (no tapes)

Storage Energy Usage

- Disk Drive is a primary energy consumer in a storage system
- Typical Storage System Components
  - Disk Drive
    - Motor
    - Electronics
    - Actuator
  - Infrastructure (enclosure, controllers, fans, …etc)
- Disk Drive Energy Usage Relative to Each Other
  - In a 3.5-inch disk it is about 33/33/33
  - In a 2.5-inch disk is it about 20/50/30

Secondary Energy Usage

- Cooling
  - For every watt used it takes 1 watt to remove the heat
  - Air cooling is currently the preferred method
  - Water is 4000 times more effective than air for cooling components
- Keep the heat out of the box
- Get the heat out of the rack
- Move the heat out of the data center

Performance

- In General
  - Instead of making a single disk run ever faster, just use a larger number of smaller disks
  - Common practice in CPU industry
  - Virtually unknown in Disk Storage industry
- IOPS
  - Smaller laptop-class disk drives are individually slower than an Enterprise-class drive – about 2:1 in favor of 3.5-inch disks
  - Can package an order of magnitude more laptop-class drives in an array
  - Aggregate IOP performance for an array of SFF drives is 5:1 in favor of the 2.5-inch disks
- Bandwidth
  - Same argument as above - about 2:1 in MB/sec in favor of 3.5-inch disks
  - Aggregate bandwidth for an array of SFF drives is 5:1 in favor of 2.5-inch disks

Why move toward smaller disks

- Power reduction is non-linear in favor of smaller form factor
- Cooling is simpler because of low power consumption
- Self induced (rotational) vibration modes are significantly reduced
- Pricing takes advantage of the commodity lap-top drives
- Reliability and data integrity is a different talk
- Aggregate performance is significant
- Aggregate head-count per unit space or volume is significantly higher than 3.5-inch packaging methods

Why not Move Toward Smaller Disks

	3.5	2.5	Difference
Capacity Per drive	1TB	320GB	~1/3th
Density TB/unitvol	1TB	1.87TB	87% higher
Space in units of 2.5" drives	5.85	1	~1/6th
IOPS	77	59 / 236-354	Individually slower Aggregate is much higher
BW (MB/s)	105	58 / 232-348	Individually slower Aggregate is much higher

In terms of Power…

	3.5	2.5	Difference
Seek/R/W Power	12W	2W	1/6th
Idle Power	8W	0.6W	1/13th
Capacity	83.3 GB/W	160 GB/W	~ 2 X better
Density	83.3 GB/U/W	312 GB/U/W	3.75 X better
Power Density	2 W/U	2 W/U	Same
IOPS	6 IOPS/W	30 IOPS/W	5X better
BW (MB/s)	9 MB/s/W	29 MB/s/W	3X better

Storage Design Considerations

- Performance
  - Signal Aggregation
  - "We need more disks, not bigger ones" *Gary Grider, NNSA*
- Tight packaging but you must get the heat out
  - Heat is the #1 threat to disk drive life - maybe
  - Bigger disks produce more heat than smaller ones
  - Tight packaging can require sophisticated cooling
- Vibration management
  - 2.5-inch drives have virtually no rotational vibration
- Maintenance – Fail-in-place

A SAID – Self-maintained Array of Identical Disks

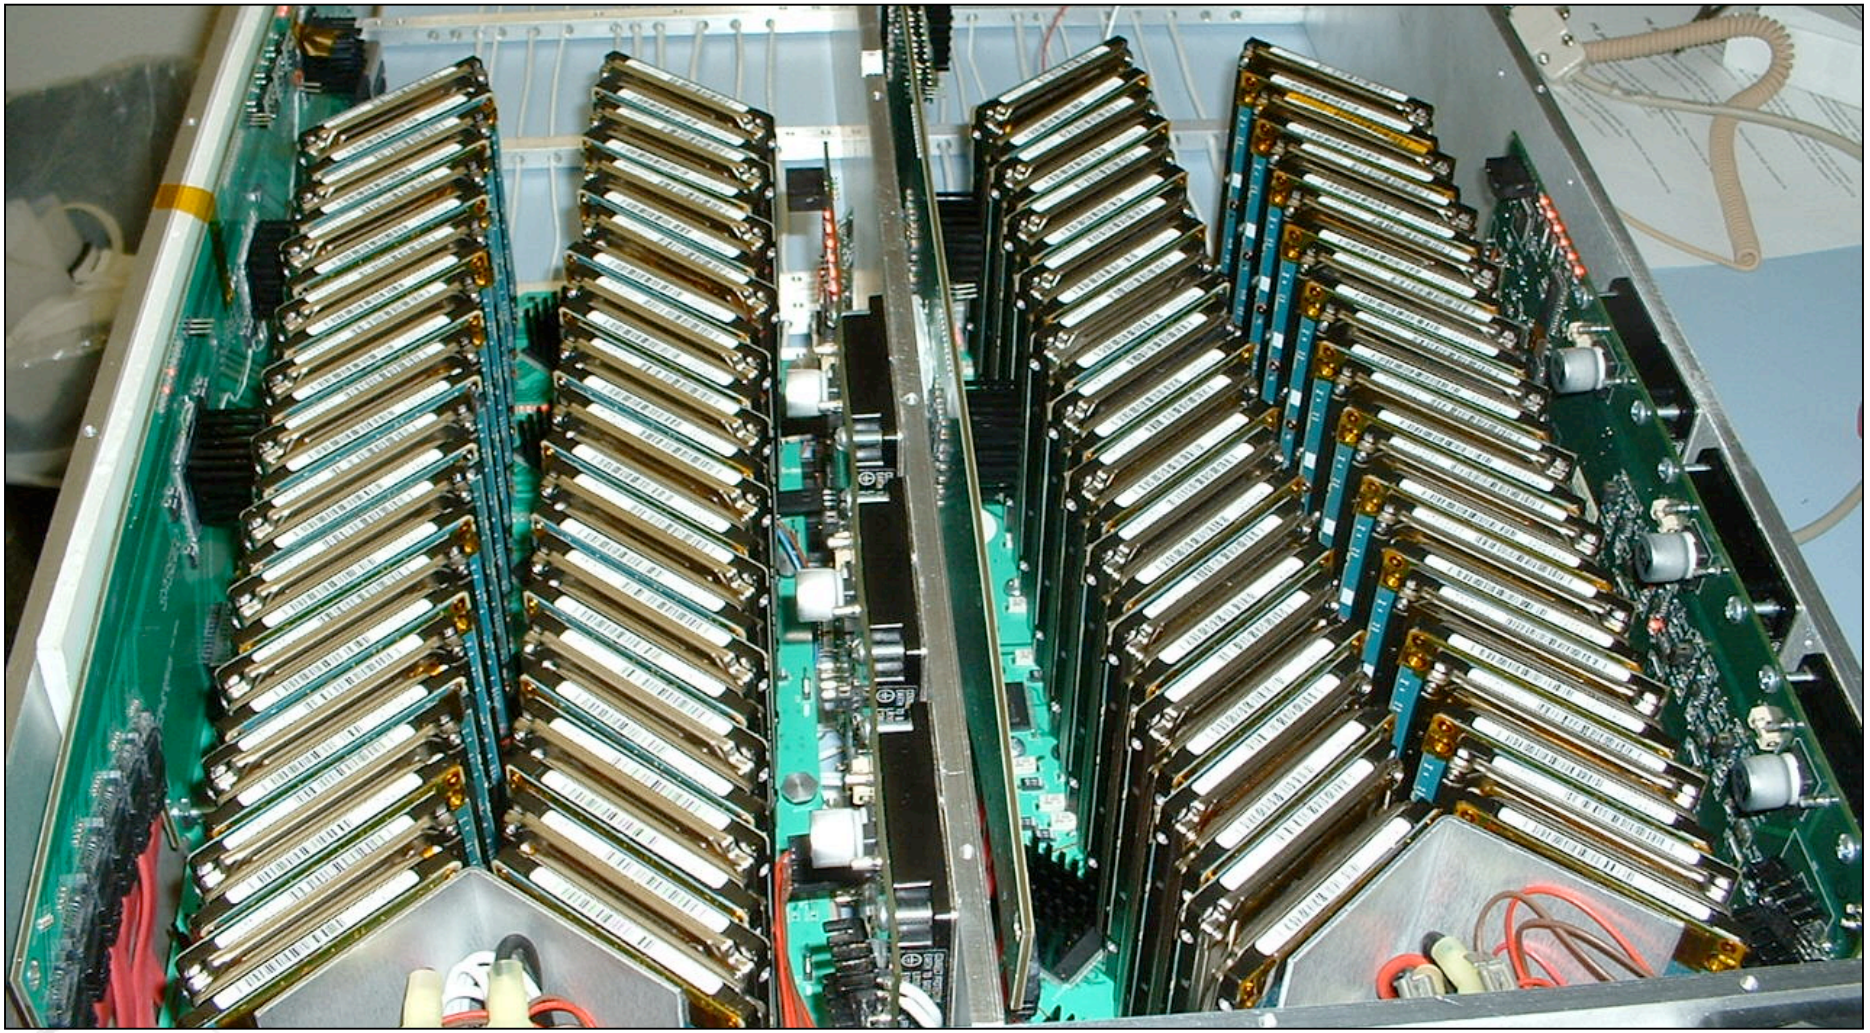

What it all comes down to

Atrato SAID
- 3RU
- 9 GB/sec
- 10,000 IOPS
- 50TB raw capacity
- 700Watts

3.5-inch standard packaging (16 drives per box)
- 30RU
- 16 GB/sec
- 12,320 IOPS
- 160 TB raw capacity
- 2500 Watts

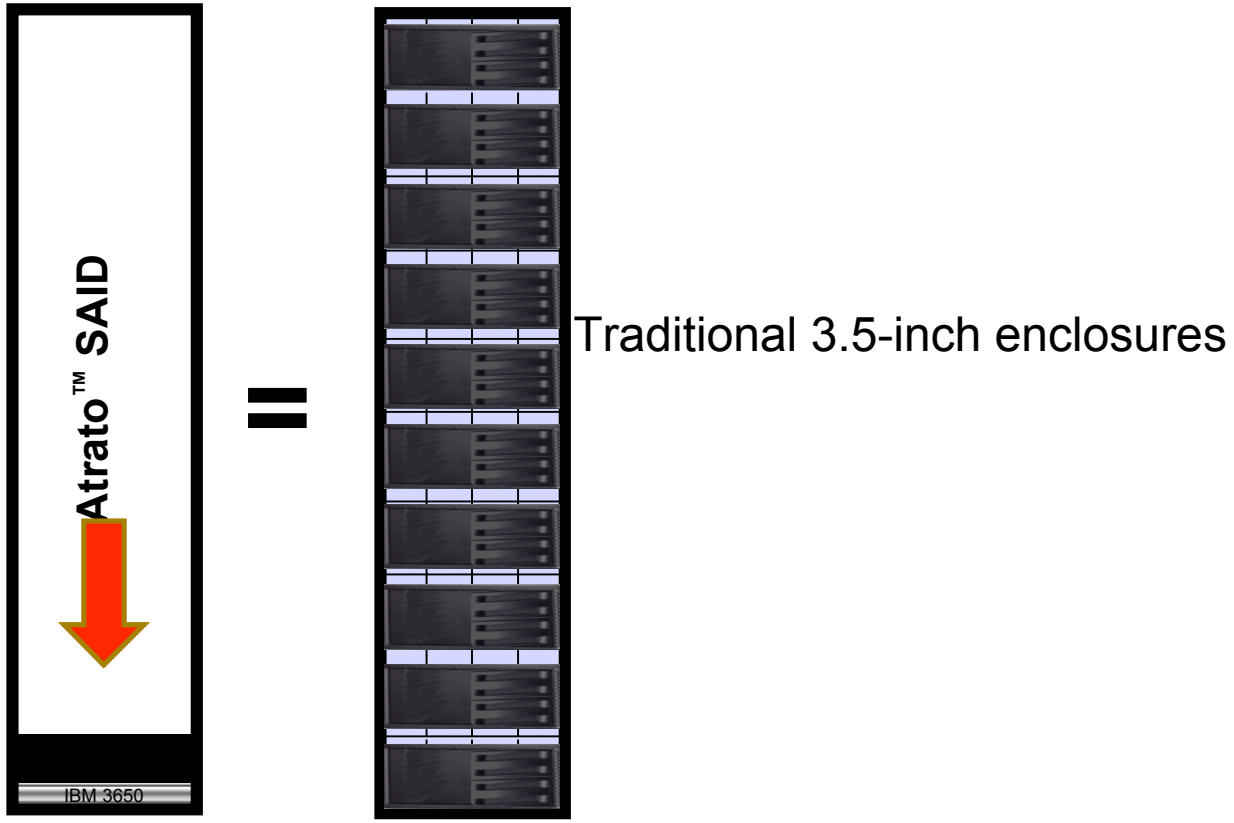

Conclusions

- Small disks are non-intuitively better than 3.5-inch disks when it comes to power
  - Better performance/watt
  - Better capacity/watt
- Requires different engineering practices
- Requires different maintenance philosophy

ATRATO™

ACCESS. THE REVOLUTION.

Thank-you