

ATRATO™

ACCESS. THE REVOLUTION.



Energy Efficient Storage

MSST 2008

Atrato, Inc.

10955 Westmoor Drive, Ste. 300

Westminster, CO 80021

www.Atrato.com

Presenter:

Thomas M. Ruwart

Chief Scientist

612.850.2918

tmruwart@atrato.com

Overview



Atrato

Energy Usage in general

Storage Energy usage

Secondary energy usage

Storage design considerations

Why move to smaller disk drives

Design Considerations

Conclusions



9/22/2008

MSST September 08

ATRATO™

Atrato

Atrato was founded on the idea that we could help people achieve nearly instantaneous access to the data and content they need.

Our purpose is to reshape and redefine the markets by challenging the traditional thinking about how data is accessed.

Startup formed in January 2004

Formerly Sherwood Information Partners, Inc., name changed Feb 2008

Based in Westminster, CO

Currently about 70 employees

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.

www.AtratoInc.com

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 the 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 watts 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

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 laptop drives

Reliability and data integrity is a different talk



Why not Move Toward Smaller Disks



	3.5	2.5	Difference
Capacity Per drive	1TB	250GB	1/4 th
Density TB/unitvol	1TB	1.45TB	45% higher
Space in units of 2.5" drives	5.85	1	~1/6 th
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	1/6 th
Idle Power	8W	0.6W	1/13 th
Capacity	83.3 GB/W	125 GB/W	50% better
Density	83.3	745	9 times better
Power Density	$\frac{\text{GB}}{2 \text{ W/UV}}$	$\frac{\text{GB}}{2 \text{ W/UV}}$	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
- 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



9/22/2008

MSST September 08

ATRATO™

What it all comes down to



Atrato SAID


- 3RU
- 9.3GB/sec
- 10,000 IOPS
- 40TB raw capacity
- 400Watts

3.5-inch standard packaging (16 drives per box)

- 30RU
- 16GB/sec
- 12,320 IOPS
- 160TB raw capacity
- 2500Watts



Traditional 3.5-inch enclosures



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



Traditional 3.5-inch enclosures

