



Top100io BOF

MSST2005
April 13, 2005

Ben Kobler, NASA GSFC
Jack Cole, ARL

April 13, 2005

IEEE MSST2005
Monterey, California USA



Goals

- ✓ Develop a Method for Tracking Trends in High-Performance Storage
- ✓ Release Annual or Biannual List of 100 Most Powerful Storage Installations
- ✓ Get Started, Don't Wait for Perfection



Inspiration

- ✓ Based on User Production Sites
Not Vendor Test Installations
- ✓ Inspired by Top500 Supercomputing
Sites Project
- ✓ Top 10 Database Sites



Organization

- ✓ Coordinating Committee
 - Jack Cole
 - Ben Kobler
 - Tom Ruwart
 - Gary Sobol
- ✓ Industry Advisors
 - Don Capps
 - Garth Gibson
 - Hoot Thompson
- ✓ General Members
- ✓ Subcommittee of MSSTC



Proposed Metrics

- ✓ Use I/O Bandwidth as Measure of Storage System's Ability to Do Work
- ✓ Similar to Top500
2 Tests: Max & Peak

Proposed Tests: Max & Peak

- ✓ Max: Tests Against Single Mount Point
 - Different Parts of a Storage System Considered as One, If Tied Together At Single Mount Point
 - Single Mount Point Can Encompass a Single File System, or Multiple File Systems
 - Could Use Distributed File Systems Over WAN

Proposed Tests: Max & Peak

- ✓ Peak: Tests Against Multiple Mount Points
 - Tests Cumulative Performance of Multiple Independent Storage Systems
 - Theoretically 'Peak' Could Be Higher Than 'Max' Because There Would Be Less Overhead Managing a Common Mount Point

Proposed Tests: Max & Peak

- ✓ The List Will Be Sorted Based on 'Max', But Also Report 'Peak'.
 - Will minimize incentive to arbitrarily add a lot of unrelated components together in parallel to get a high 'peak' number, since only the 'max' number would determine your place on the list.

Goal of BOF

- ✓ Solicit User, Vendor Feedback
- ✓ Review MSST2005 Vendor (And Other) Test Results
- ✓ Discuss Other Issues:
 - What benchmark to use
 - Recommending iozone and xdd
 - www.iozone.org
 - www.hut.fi/~mweissen/iozone-macosx.html
 - www.ioperformance.com
 - 209.98.16.6

Additional Issues

- ✓ Read vs Write
- ✓ Multi- vs Single-stream
- ✓ Block Sizes to Use
- ✓ Size of Test: Great Enough to Overcome Cache Effects
- ✓ Solution?: Allow Sites to Pick What Works Best for Them and Document Results



Future meetings

- ✓ Top100io Subcommittee
 - May 24, 2005, College Park, Maryland
- ✓ Who of you will come?
- ✓ Supercomputing BoF
- ✓ FAST05?