the green grid®
get connected to efficient IT

Best Available Technologies: External Storage

Overview of Opportunities and Impacts

November 18, 2015

www.thegreengrid.org

Energy Savings: Background

- When examining external storage energy savings opportunities it is important to note
  - Product usage accounts for the vast bulk of energy consumption
    - From 80% - >90%, depending on configuration
  - Since a storage product is shared by multiple applications and servers, it can be difficult to judge whether there is user demand on the device
  - Drive and data hygiene functions execute when user demand is low, to avoid performance losses, significantly reducing potential “idle” time
- Reducing the number of drives (spinning or solid state) directly reduces the energy consumed

The Role of Disk Drives

- The dominant consumers of energy when a storage system is in use are the HDDs and SSDs
- HDD power draw is driven by their rotational speed, not their capacity; thus high performance drives draw far more power than very high capacity drives
- SSDs draw power whenever they read or write, but they draw less power than fast HDDs because less capacity is required to achieve performance than with the fastest HDDs

Best Available Technologies

- The best (largest) energy savings for storage systems is to use less storage.
- This may seem contradictory, since more data is being created daily.
- Several techniques exist that can radically reduce the number of HDDs/SSDs deployed without compromising the needs of the clients
- These technologies are all relatively early in their adoption cycles, and thus are under-deployed
- The next several slides discuss the best ways to minimize the number of physical instances of HDDs and SSDs deployed in any storage system

Best Energy Savings for Disks

- Capacity Optimization Methods (COMs) are software techniques that allow more data to be stored (less media to be used) than would be possible with older techniques.
- **Thin Provisioning:** Virtualizes allocation of storage so applications believe they have all they need, while only a fraction is present.
- **Deduplication:** Elimination of duplicate files or data chunks from one or more files, while maintaining integrity of all records.
- **Compression:** Reduction of data sets (often within files) to reduce the number of bits stored while still presenting the original data
- **RAID5/6, Erasure Encoding:** Methods for improving data integrity without using full data copies; savings range from 70%-90% based on technique
- **Delta Snap-shots:** Method to allow reconstruction of different versions of a file without multiple full copies.

Auto-Tiering

- Data used by applications has patterns of use that reflect its immediacy for business purposes. By seamlessly migrating that data to less performant (energy-intensive) media as the pattern of use changes, the need for increasing amounts of high-performance storage can be greatly reduced.
- Auto-Tiering makes this migration invisible to applications, and is configured to reflect policy decisions appropriate for each application.
- Auto-Tiering maintains business value while reducing the energy burden of storage for each application.

Copyright © 2015, The Green Grid

Why Idle Really Isn't

- Two methods of "idling" storage systems are often proposed
  - Spinning down some or all of the disk drives when not in use
  - "Sleeping" the controllers that speak to the storage network
- There are several technical issues that interfere with the purpose and operation of external storage systems, precluding the use of these approaches

Drive Spin-down

- Storage systems are shared among multiple servers and applications.
- Storage systems cannot determine if each application or server is still operating, nor when the next I/O request will be made.
- The drives in each system are rarely dedicated to a single application; if contact is lost with a specific application or server, it is not possible to determine when the next request will occur.
- The time to return a drive to operating speed is measured in minutes. A new I/O request will consider a drive to be off-line if the data transfer does not occur in a matter of milliseconds, interfering with the application in use.
- To date, there have not been significant adoption of “spin-slow” drives due to the difficulties in identifying which drives to slow, as well as the limited availability of such drives.

“Sleeping Controllers”

- The controllers on a storage system consume far less energy than the drives, so the savings from this would be much smaller.
- As discussed on the last slide, the response times required by applications are very narrow, and the number of applications addressing each storage system can vary.
- For these reasons, storage systems are not good candidates for “network standby” conditions, which would result in application failures

Technology Adoption & Innovation

- Storage technology is continuing to innovate new ways to reduce the impact of both the growth of data and storing it.
  - It is important to avoid regulating in ways that could limit new discoveries
- As innovations are brought to market, data center operators will need time to vet them for reliability and savings prior to full-scale implementation in the majority of installations.
- Other innovations in storage systems may be forthcoming in the areas of improvements in connectivity bandwidth. While these may cause some increase in energy use, they will provide the benefit of increased data transfer per watt where the connection has been the limiting factor.

Copyright © 2015, The Green Grid