Education

Notes & Lessons Learned from a Field Engineer

Robert M. Smith, Microsoft

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Abstract

- Notes & Lessons Learned From a Field Engineer
  - This session is a collection of storage related topics based on firsthand experiences with various customer interactions over the past 5 years.

    The customer base covered in this presentation ranges from manufacturing to government contractors where constraints such as budget and down-level customer requirements drive storage realities.

Notes & Lessons Learned: Overview

- Partition alignment still a big problem
- Storage management issues
- Storage performance issues
- Storage realities
- "Growing pains"
- Tools
- Wish list
- Futures

Many Partitions Still Not Aligned

- Storage partitions on RAID arrays (other than 2-disk RAID-1) are still misaligned in many cases
  - Often from older systems that have been upgraded
  - Are not "re-aligned" during an upgrade to a newer OS
  - Still have to be deleted and recreated
  - 2048 (512-byte) sector starting offset is a compatible value with a tiny loss of disk space (1 MB).
  - Performance penalty for misaligned partitions up to 30% (not anecdotal)
    Source: Jimmy May
    http://blogs.msdn.com/jimmymay/archive/tags/Disk+Partition+Alignment/default.aspx

Storage Management Is A Mixed Bag

- Proprietary vendor tools
- Expensive vendor tools
  - "Yes we have that, but you didn't buy that, you bought this"
- Mixed OS support
  - Native infrastructure is basic, and vendors still implement their own methods

Storage Management Is A Mixed Bag

- SAN group doesn't always interact well with Server group
  - "I need 250 GB, 2500 IOPS, 80/20 read/write, 9 KB Average Request Size,10 ms"….
    "Here's your LUN" (250 GB, unknown performance capability), sometimes shared, sometimes virtualized at more than 1 layer
- The tool or suite either isn't properly used, or managed by a different group

Storage Performance Challenges

- Performance decisions made by purchasing dept.
- Performance decisions dictated by GB
  - Example: 300 GB disk-drives, RAID 1+0, 6 disks, == 900 GB, but only 540 IOPS*
- Performance issues difficult to diagnose
  - Multiple groups involved
  - Multiple vendors involved
  - "The Tool" not purchased
    (Sales team might more strongly assert value of even rudimentary tools)
  - Storage performance more of an art than a science, and not taught widely

Storage Performance Challenges (2)

- Performance issues difficult to diagnose
  - "The SAN is performing fine"
    - What is "The SAN"…the array, the fabric, or both?
    - Performance results not shared
      "The "SAN" is performing fine".
    - Performance results difficult to disseminate
    - Analyzers nearly out of the question, and even then, takes a very experienced analyst to disseminate
    - Not many (ANY???) vendor neutral tools

Storage Realities

- Hardware refresh pushed out
  - Controller batteries failing!
  - No new firmware update
    - Persistent Reserve support
- Still using basics…nothing exotic
- Customers NOT using performance tools at all, or to their potential
- Critical Situations often invoke a circular blame-game

Storage "Growing Pains"

- Increased capability means increased need for space
  - Backup schemes
  - Online user-restore capabilities
  - "Invisible" space used by OS, sometimes not accounted for
  - New capabilities add overhead to servers
  - Poor utilization of fixed volumes
    - Grow and shrink becoming common, which helps
    - Still requires intervention in many cases
      Development…A Little Help Here! ☺

Volume Recovery Tools

- Master Boot Record
  - Who backs that up, and who restores it?
  - Recover from deleted partition
    - Sector editors
    - 3rd-party data recovery tools
- As partition layouts increase in complexity, so does volume recovery (Ex: GUID Partition Table)

Wish List

- Storage container with assignable attributes
  - Response time (MS)
  - Input-Output Operations / Sec (IOPS)
  - Size (MB, GB, TB, etc)
  - Action when over-utilized or under-utilized
    - Auto-grow
    - Auto-shrink, reclaim for other uses
- Vendor-Neutral performance tools (That don't cost a fortune, or better, free)
- SSD appears to be nirvana, other than cost

What's Taking SSD So Long???

- Still perceived as too expensive in many cases
  - Sales…a little help here please. ☺
- Storage Management not taken into account
  - Ease of growing and shrinking LUNs, thereby increasing utilization of purchased GB
- Support Cost not taken into account
  - How many man-hours are spent diagnosing poor performance?

What's Taking SSD So Long??? (2)

- Some previous problems alleviated:
  - Data locality (Short-Stroking)
  - "Spindle-Sharing"
  - Rack Space
  - Power Consumption
  - Cooling
- We know SSD may not be nirvana, but it's close
  - Firmware data protection schemes
  - Security overhead mitigated by increased performance
    - Native hardware encryption
    - Full-Volume Encryption

Related SNIA Tutorials

Check out this SNIA Tutorial:

1. **Solid State Storage Architectures**
2. **The Benefits of Solid State in Enterprise Storage Systems: Today and Tomorrow**
3. **How to Eliminate Configuration Drift Risk**

Q&A / Feedback

- Please send any questions or comments on this presentation to SNIA: trackstoragemgmt@snia.org

Many thanks to the following individuals for their contributions to this tutorial.

- SNIA Education Committee

Jimmy May