RAID Performance Comparison

Version1.0
January2013

Copyright

Copyright@2004~2013, Qsan Technology, IncAll rights reserved. No part of this document may be reproduced or transmitted without written permission from Qsan Technology, Inc.

Trademarks

All products and trade names used in this manual are trademarks or registered trademarks of their respective companies.

QsanTechnology, Inc.

4F., No.103, Ruihu St.,
Neihu Dist., Taipei City 114,
Taiwan (R.O.C.)

Tel: +886-2-7720-2118
Fax: +886-2-7720-0295

Email: Sales@QsanTechnology.com
Website: www.QsanTechnology.com

Introduction

It often confuses the end users that, how the performance would be impacted with different RAID group and Volume amount settings applied to a system with the same amount of physical disks. In order to further clarify this technical hazard we tried to perform a test using different combination of virtual disks and RAID group amounts.

With all other configurations fixed, major factors impacting test results are listed and explained as following:

1. Virtual Disk Divergence Level

   Because of the nature of magnetic platters based hard disk drives, the drive head has to move between disk tracks. The divergence level is important because the more aggregated ŭKŜ ØλNŭdzI·f ŘλăʅăI·NŜ ǿKŜƴ ŭKŜNŜQNŜ ϒ2NŜ ŭKI·ƴ 2yŜ ØλNŭdzI·f ŘλăʅăλƴI·w!L5 ƎN2dzLJƶ ǿλŭK ŭKŜ virtual disk capacity evenly distributed on all data disks contained in the RAID, the less the read-write has to move around fetching data acquired by access commands and therefore the performance is expected to be better than that with badly diverged target disk sectors.

2. Effective Data Disk Numbers

   This is the nature of the RAID concept, as with more RAID groups (with parity) reside in one system containing unchanged number hard drives, the more drives are assigned to store LJI·NλŭèŘI·ú·φ ¢KŜ ŜŦŦŜŎŭλØŜ ŭKN2dzƎKLJdzú λă ŭKŜNŜŦ2NŜ ŜÈLJŜŎŭŜŘ ú2 ōŜ NŜŘdzŎλyƎI·ă λƴŭKŜ2Nè λŭΩă directly proportional to the amount of data disks offering space for data inputs/outputs.

3. Parity Calculation Cycles

   Extensively needed to be of our concern is how the parity data of RAID 5 and 6 are calculated, as well as how many parity disks are involved in the data writing process, as parities are not checked during read actions. RAID 6 uses different algorithms for calculating the two parities and therefore requires more calculation cycles than RAID 5.

Comparison Set A

We plan to test the following configurations:

- 1R1V (1 RG contains 1 VD)
- 1R2V (1 RG contains 2 VDs)
- 1R3V (1 RG contains 3 VDs)

Copyright@2004~2013 Qsan Technology, Inc. All Rights Reserved.

- 1R4V (1 RG contains 4 VDs)

Scenario

In this comparison set we start with the simplest setting, one RAID 5 RG and one VD, then adding VD amounts to 4 in order to simulate the scenario that the user adds VD to unused RG space with growing needs.

Results

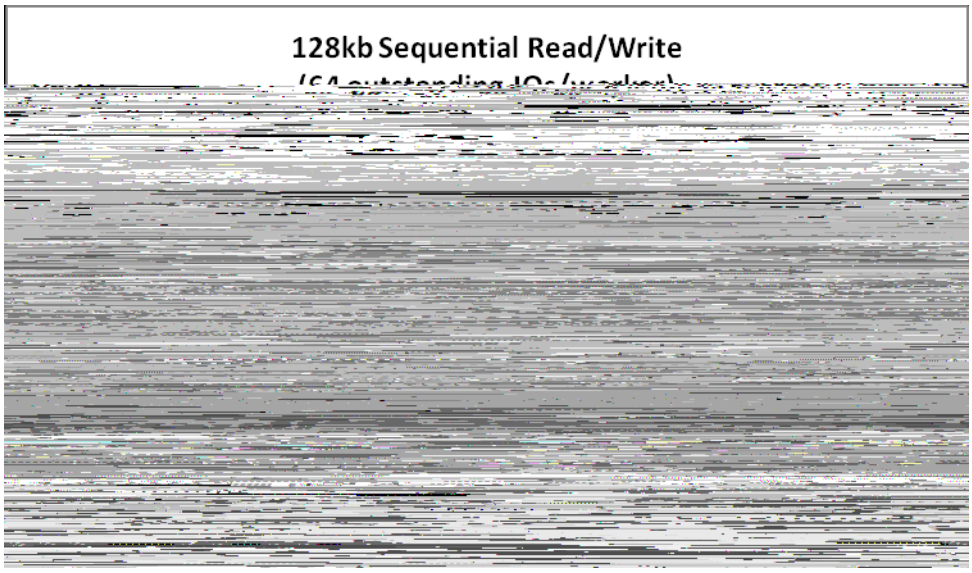

Analysis

A very obvious trend of descending performance can be observed from the comparison of different RG/VD amount settings. Each time with one more VD added to the RG, the read/write performance drops lower. With 4 VDs reside in the RG, the performance compared to only one VD resides in the RG, drops by 24~45% in reading, 30~45% in writing, depending on the worker loadings assigned in IOmeter test configurations, as shown in the following diagram. It is therefore showing the diverged read/writes to different VD segments in the same RG would lead to obvious reduction.

Comparison Set B

We plan to test the following configurations:

- 1R4V (1 RG contains 4 VDs)

Copyright@2004~2013 Qsan Technology, Inc. All Rights Reserved.