



JetStor Document - White Paper

SSD Caching Performance Monitor Case Studies

**JetStor 724iF/FD, 724HS 10G/HSD 10G, 712iS/iSD, 712iS 10G/iSD 10G,
716iS/iSD, 716iS 10G/iSD 10G**

Version 1.0
November 2014

Copyright

Copyright@ 1994~2014, Advanced Computer & Network Corporation. All rights reserved. No part of this document may be reproduced or transmitted without written permission from Advanced Computer & Network Corporation.

Trademarks

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

Advanced Computer & Network Corporation

5001 Baum Blvd., Ste 680
Pittsburgh, PA 15213

Tel: 412-683-9010

Fax: 412-683-9070

Email: sales@acnc.com

Website: www.acnc.com

Introduction

In this document we are glad to help users who may encounter some performance issues when they are using JetStor products. Customers may feel slow performance while they operate some routine cases after more loading increased into the RAID group on the SAN storage. Now it is easy to figure out where the bottleneck is. We will take some cases in this document and guide you through how to determine which part could be the main cause of the slow performance.

Usage

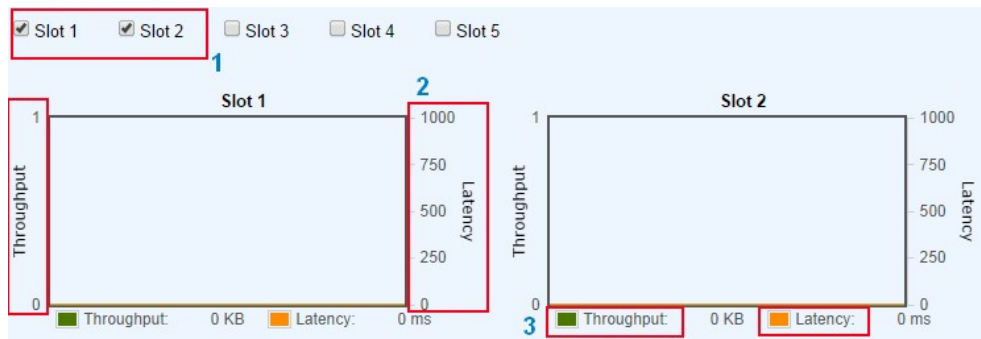
The new feature - performance monitor - is located on the left menu on the web UI of JetStor SAN storage system.



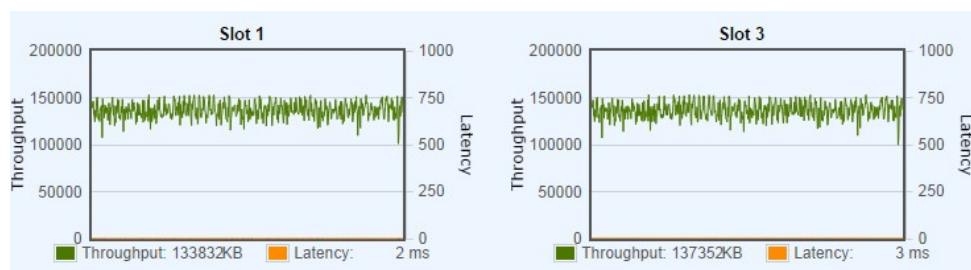
It supports the monitoring of the detected hard drives and the iSCSI flow as well as the flow of Fiber channel.



1. Select **Disk**, **iSCSI**, or **Fiber Channel** (if using FC series models) tab for what you would like to monitor.
2. You may choose to monitor the hard drives which are installed on the local unit or on the JBOD unit.
3. Select the detected hard drives that you want to monitor.

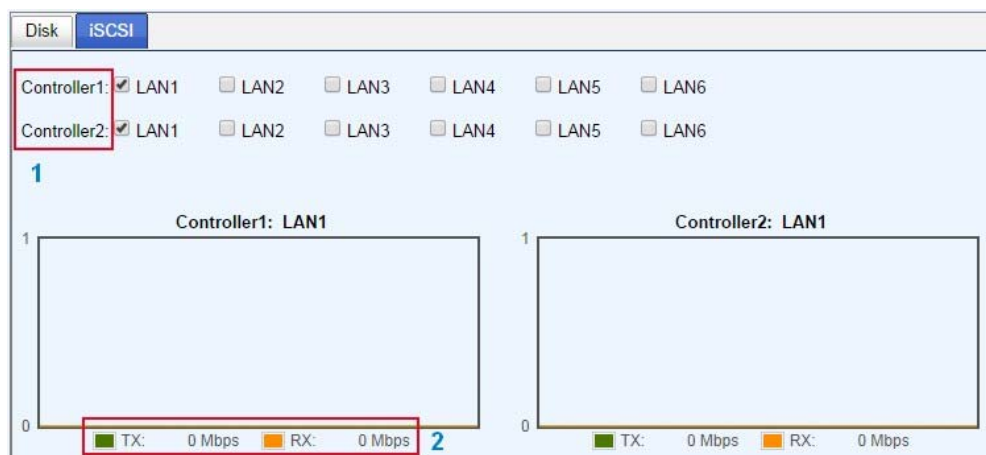


4. For example the **Slot 1** and **Slot 2** have been selected, there will be two charts displayed.
5. You will see the **Throughput** of the selected hard drive at the left side whereas the **Latency** at the right side listed on the showed charts.
6. Green color represents the **Throughput** of the hard drive while the amber is the **Latency**.



7. And there will be numbers displayed below the chart indicating the exact throughput the hard drive delivered.

The iSCSI and Fibre Channel monitoring pages work at the same method.




1. You can monitor two controllers in the same tab.
2. The **TX** with green line displayed represents the Transferred iSCSI traffic which means the read commands generated from the connected host(s)/server(s). For example, the server connected to the iSCSI volume, and copies some files out of the volume onto the local drive of the server.
3. The **RX** with amber line displayed represents the Received iSCSI flow which means the “Write” commands established from the connected host(s)/server(s). For example, the server connected to the iSCSI volume, and copies some files into the volume from the local drive of the server.

Environment

As we mentioned in the beginning of this document, we are going to take some examples as the real scenario that users may encounter. Here we introduce the testing environment in our laboratory.

Host	Intel Xeon® CPU E5620, 2.40GHz, 16GB RAM
Host OS	Windows 2008 R2, 64-bits
Storage	JetStor SAS 716iSD (Dual controller)
Controller Firmware	V1.2.0 (20140731_1100)
HDD	WD2003FYYS-02W0B0, 2TB, 7200 RPM, 64MB cache (x4)
RAID Level	RAID 5
Virtual Disk	1000GB


Diagram




Preparation

Before we start, it is better to know how much throughputs one single HDD (we used in this document) can deliver. We first created a virtual disk of RAID 0 with only one HDD, and attached a LUN to be connected from the server, and tested the maximum throughput via IOMeter.

- **Read** commands triggered from the server, it's around 150MB/s.



- **Write** commands triggered from the server, it's around 100MB/s.



- **100% random read commands with 4K block size** triggered from the server, it's around 150 IOPS.

Display	All Managers	152.62	1000
Total I/Os per Second			


- **100% random write commands with 4K block size** triggered from the server, it's around 100 IOPS.

Display	All Managers	103.46	1000
Total I/Os per Second			

Case Studies

Case 1 - The maximum throughput of one single iSCSI port has reached

Now we start to run the first testing case. Triggering the **read** commands via IOMeter from the server with one single iSCSI connection. In theory, based on the result we got in the pre-test above, 4 hard drives with RAID 5 RAID group created shall deliver up to 450MB/s for read performance. However in the chart below, we can observe that the maximum throughput within one single iSCSI connection (1GbE) has reached – 903.97Mbps, it is around 112MB/s.



In the Disk throughput monitoring, we can also watch the current delivering values from each hard drive in the RAID group, they are around 28MB/s for each one, total is around 112MB/s.




So the issue in this case that users may meet shall be on the following.

I got limited/slow performance after the I/O loadings increased to the iSCSI volume from the server. I have only one iSCSI connection logged in, but I am not sure if it is normal or not. Where may I try to sort out the problem?

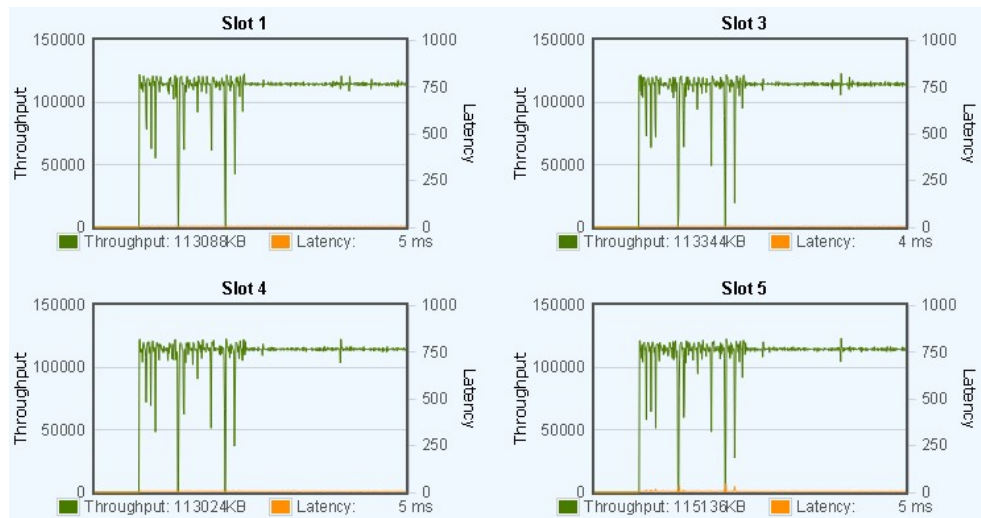
Or some similar questions from users, you may suggest them to check:

1. If the maximum performance on that single iSCSI port has been reached or not.
2. If there is any specific hard drive in the same RAID group generating the high latency value larger than others, it might be the one which is slow down the overall performance. Please suggest the user to replace the problematic hard drive and check again.
3. Add the iSCSI connection(s) one by one between the SAN storage and server, and configure MPIO for improving the performance accordingly.

The result in the following charts is the maximum throughput that 4 hard drives can theoretically provide. Check the values next to TX, 900Mbps for each. So it is 3600Mbps, around 450MB/s total.



The throughput for each hard drive is 113MB/s, and 452MB/s total.

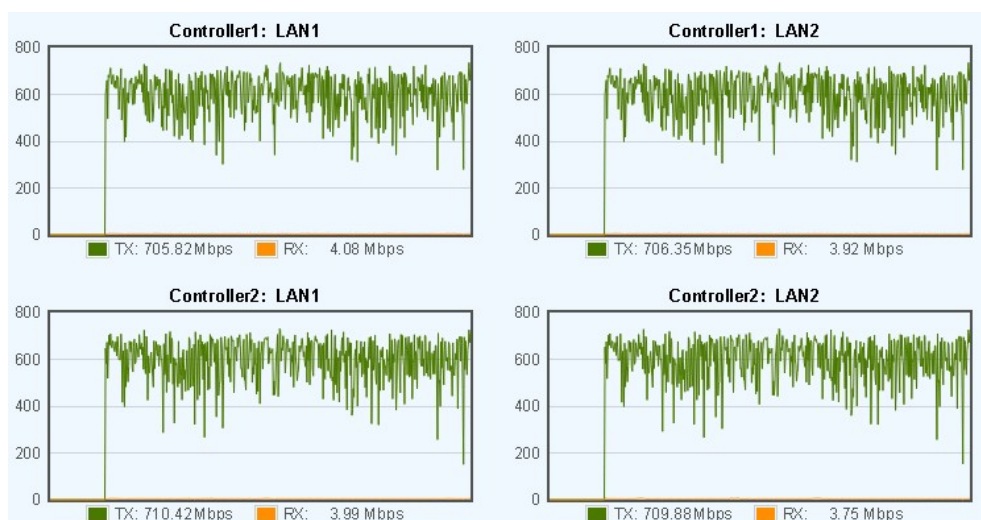


Note that we have tried to add one more (5 connections at the moment) iSCSI connection in this case. The MB/s can be increased up to 500MB/s, and the throughput of the hard drives are close to 150MB/s. Theoretically the performance is not able to reach over 450MB/s due to the parity in the RAID 5 RAID group, but our system has read ahead feature which can help to enhance the sequential read performance, so you might see greater read performance.


Case 2 - The maximum throughput of the hard drives in the RAID group has reached

In this case we removed one of the hard drives and re-created the RAID 5 with 3 ones. According to the basic test above, we shall get around 300MB/s for read and 200MB/s for write in such scenario. Let's see how the results will be.

It's about 700Mbps * 4 = around 350MB/s.



112640KB+114688KB+119304KB = 346MB/s, close to 350MB/s.



In such case we can obviously see that we have 4 iSCSI connections, (maximum value shall be around 440MB/s) but only 350MB/s delivered. Users might ask like this.

Why did I get performance lower than 400MB/s while there were 4 iSCSI connections logged in?

You may advise users to check:

1. On each iSCSI portal, if they have the similar performance value or not.
2. Check if all hard drives (in identical RAID group) have reached the maximum throughput or not (in this case, the maximum shall be around 112MB ~ 120MB per each).
3. If the maximum value of the hard drives has been reached, try to re-create the RAID group with more drives, otherwise this shall be the limitation that the current installed hard drives can provide in performance.

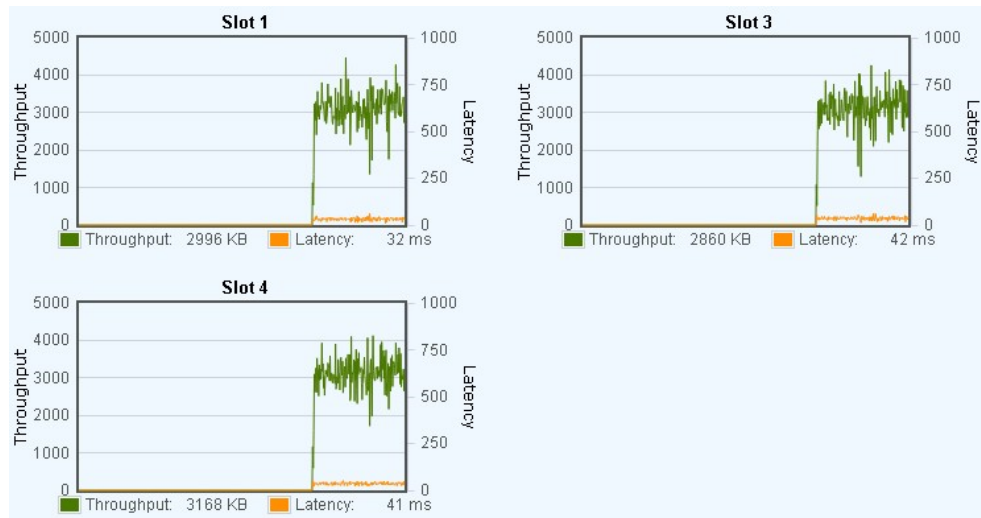
Case 3 - Got less performance from hard drives and iSCSI portals, check the latency

This case inherits the environment from the prior one; we use the same RAID configuration which is a virtual disk of RAID 5 created with only 3 hard drives. You will know that even when you get less performance from hard drives and iSCSI portals, in some situations it will be a normal and reasonable result.


We try to simulate a scenario via Iometer with random write commands inputting to the logged on iSCSI volume. Imagine that there are lots of virtual machines (implemented by such as VMware vSphere or Hyper-V) stored on the volume and generating intensive, small I/O. The I/O might

cause slow performance and lots of latency due to the hard drives are handling or responding to the incoming commands too slowly. Here I use a 100% random write I/O pattern via IOMeter with 512B block size to simulate the behavior.

The latency on each hard drives grows to over 30 ms. In reality, users may see higher values.



The delivered throughput on the iSCSI portals is relatively lower.



In such scenario, users shall be capable to understand that the incoming I/O to this virtual disk is very intensive and heavy. You may consider to:


1. Separate the I/O to a different time frame for decreasing the loading during peak hours.
2. Move some jobs to a different RAID group.

Case 4 - Got less performance, check the latency – drive failure

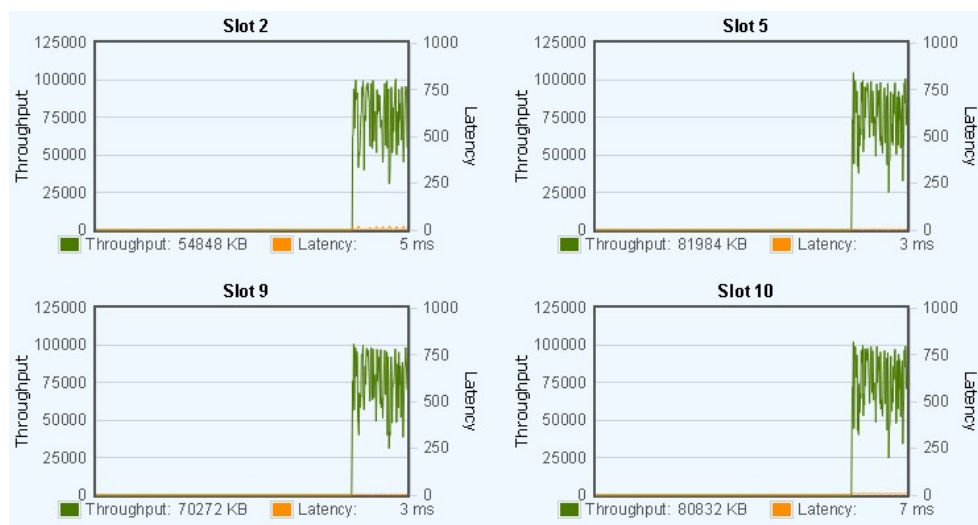
Here comes an abnormal case that users may encounter during daily works. We include one problematic hard drive (with S.M.A.R.T. error) in this test to simulate the case of “one of hard drives gets failure”. We created a virtual disk of RAID 5 with 4 hard drives (problematic one

included). Basically, you will see slow performance from the disk throughput and the iSCSI portals, please refer to the following charts.

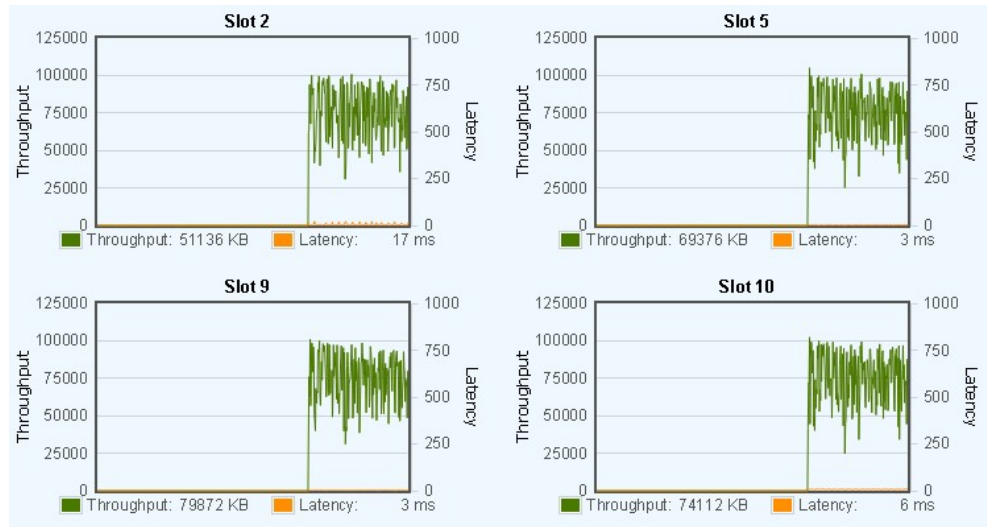
You may see the iSCSI performance is unstable and slow on each portal.



And the problematic drive (in Slot 2) will deliver lower throughputs than others.



Sometimes you may see the latency of the problematic drive is relatively higher than others.



So significantly, you shall suggest users to:

1. Replace the problematic hard drive with a new one, try to run the daily jobs again and check the status after the rebuilding process of the virtual disk is complete.
2. If the same behavior still occurs, you may need to replace the installed MUX board (if any), then try again.
3. Unfortunately if the identical problem still happens, please try to remove the hard drive (with the MUX board) to another disk slot in the chassis, and try again to clarify the issue.

After the operation 2 and 3 above, it is suggested to download the debug info file (**Maintenance\System Information** on the web interface, click **Download System Information** button to download), and report this situation to the JetStor support team (support@acnc.com) for further assistance.

Summary

We took some case studies to show the function of performance monitor. Remember that the performance monitor is used for debugging or verifying purpose only. It is not suggested to keep enabling the monitoring objects all the time. The storage systems spend extra CPU and memory resource for the function, and this might affect the overall reaction of the storage system while operating on the web interface, and the performance/throughput as well.

Applies To

- JetStor 724iF/FD, 724HS 10G/HSD 10G, 712iS/iSD, 712iS 10G/iSD 10G, 716iS/iSD, 716iS 10G/iSD 10G: FW 1.3.0