Qsan Document - White Paper

Performance Monitor Case Studies

Version 1.0
November 2014

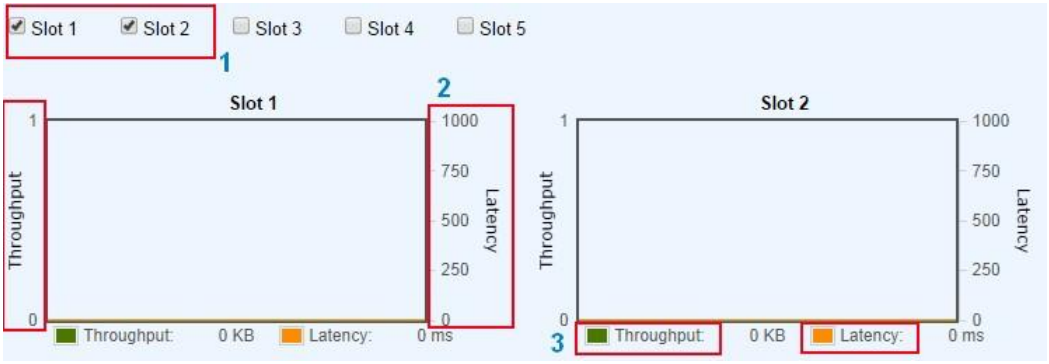

4. For example the Slot 1 and Slot 2 have been selected, there will be two charts displayed.
5. You will see the Throughput of the selected hard drive at the left side whereas the Latency at the right side listed on the showed charts.
6. Green color represents the Throughput of the hard drive while the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

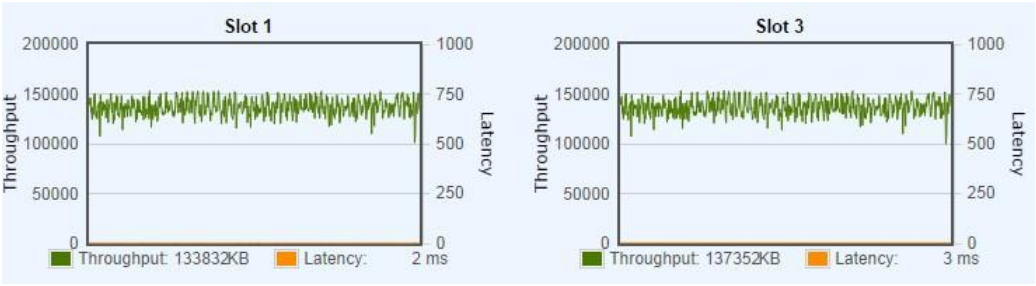

7. And there will be numbers displayed below the chart indicating the exact throughput the hard drive delivered.

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

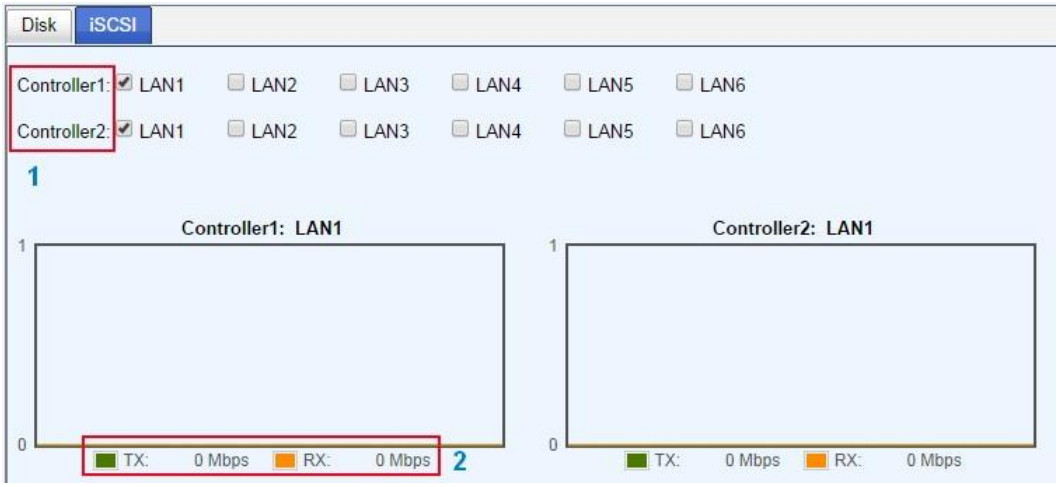

Copyright@2004~2011 Qsan Technology, Inc. All Rights Reserved.

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 file 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	AegisSAN Q500-P10-D316 (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

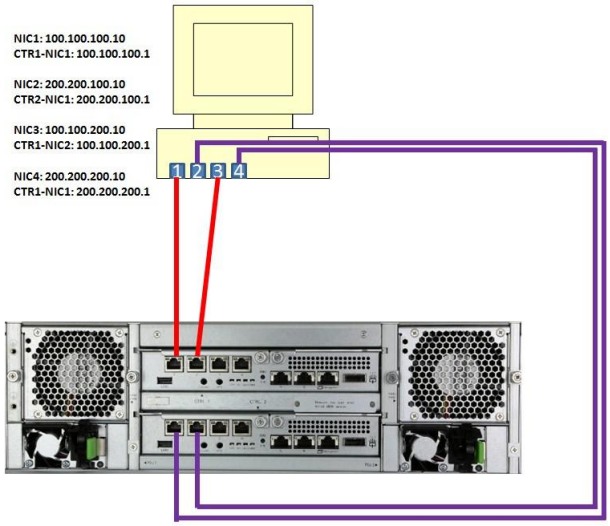

Copyright@2004~2014 Qsan Technology, Inc. All Rights Reserved.

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 is around 150MB/s.

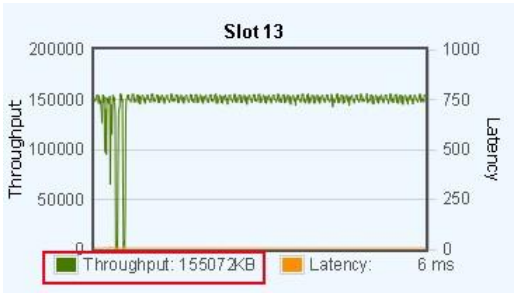

- Write commands triggered from the server, it is around 100MB/s.

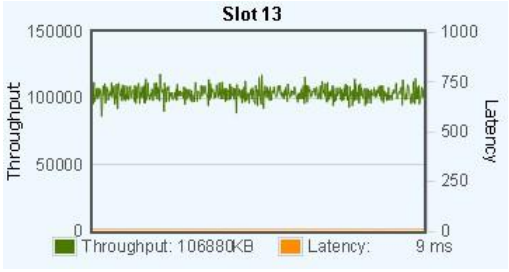

- 100% random read commands with 4K block size triggered from the server, it is around 150 IOPS.

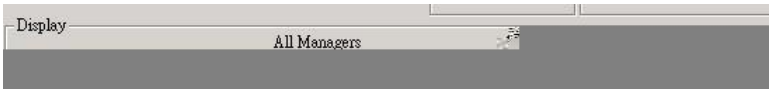

- 100% random write commands with 4K block size triggered from the server, it is around 100 IOPS.

Copyright@2004~2011 Qsan Technology, Inc. All Rights Reserved.

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 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 to 903.97Mbps, it is around 112MB/s.

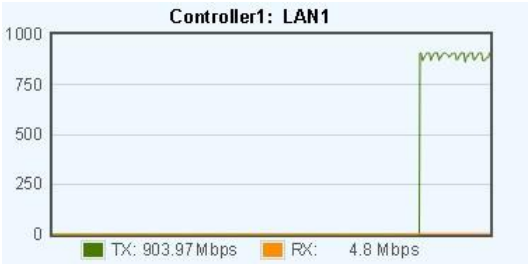

In the Disk throughput monitoring, we can also watch the current delivering values from each hard drive in the RAID group, they are around 28MB/s for each one, total is around 112MB/s.

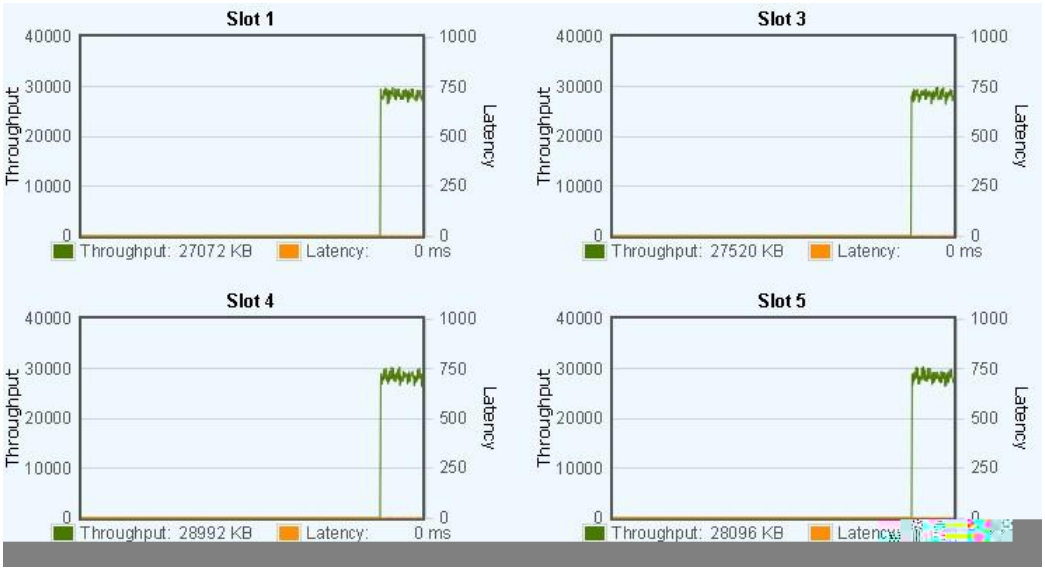

So the issue in this case that users may meet should be the following.

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?

Copyright@2004~2011 Qsan Technology, Inc. All Rights Reserved.

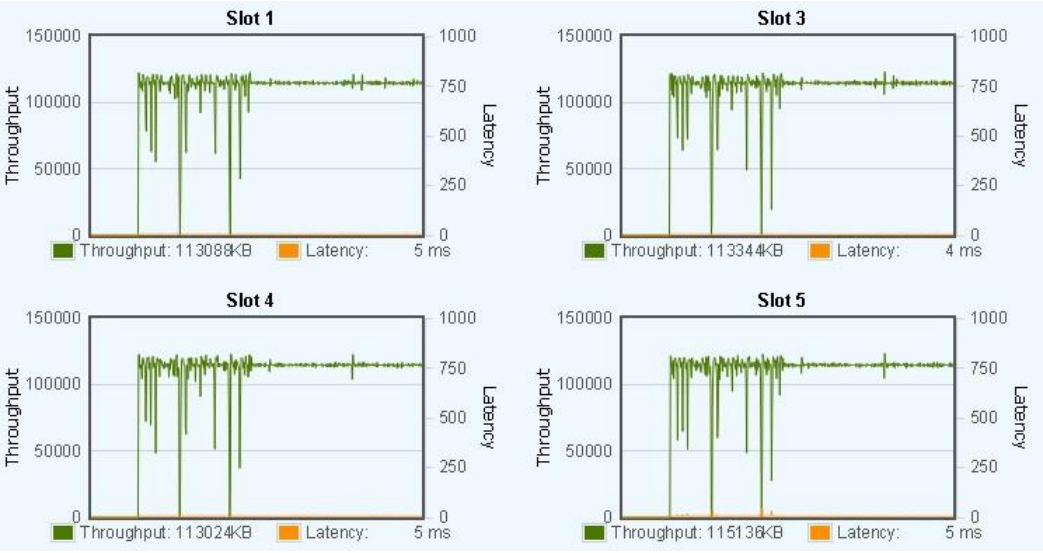

Note that we have tried to add one more (5 connections at the moment) iSCSI connections in this case. The MB/s can be increased up to 500MB/s, and the throughput of the hard drives are close to 150MB/s. Theoretically the performance is not able to reach over 450MB/s due to the parity in the RAID 5 RAID 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 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 is about 700Mbps * 4 = around 350MB/s.

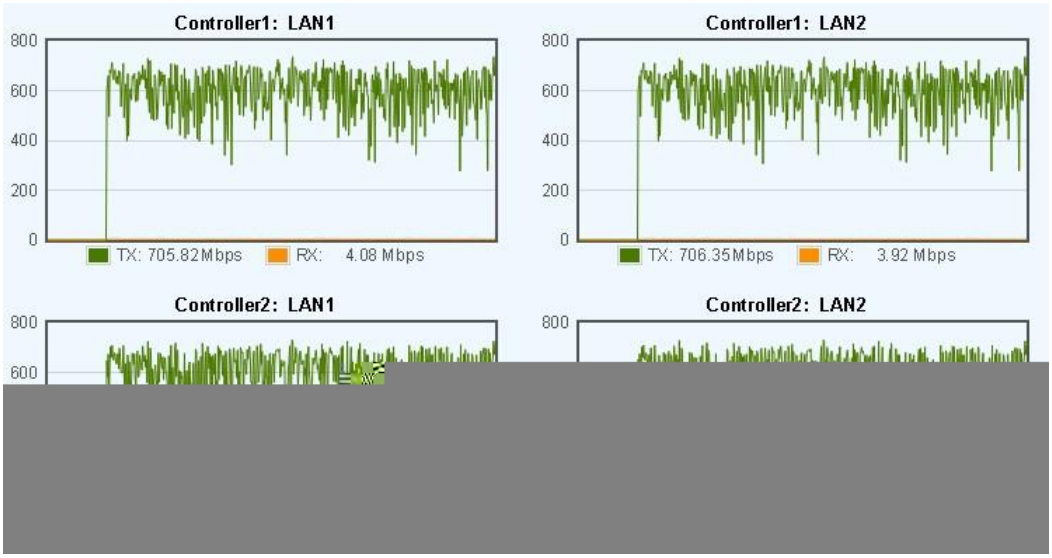

Copyright@2004~2011 Qsan Technology, Inc. All Rights Reserved.

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.

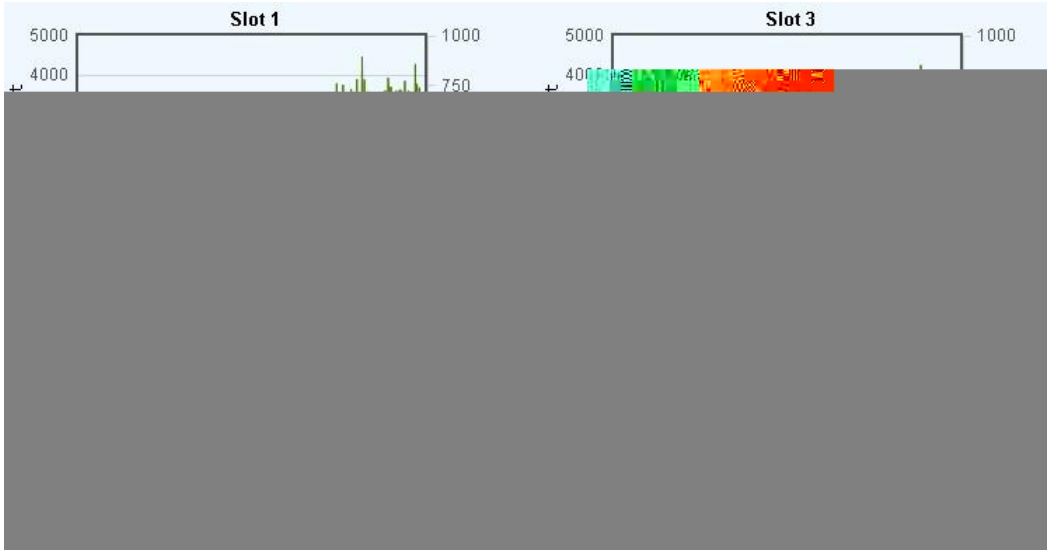

The delivered throughput of the iSCSI portals is relatively lower.

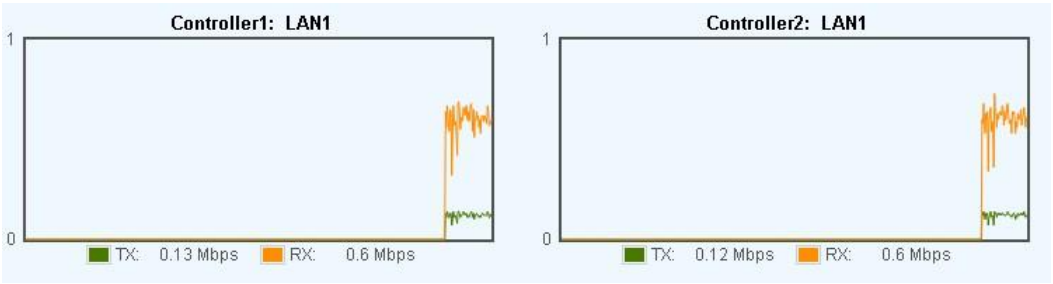

In such scenario, users shall be capable to understand that the incoming I/O to this virtual disk is very intensive and heavy. You may consider to:

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

Case 4 - Got less performance, check the latency ç 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

Copyright@2004~2011 Qsan Technology, Inc. All Rights Reserved.

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

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

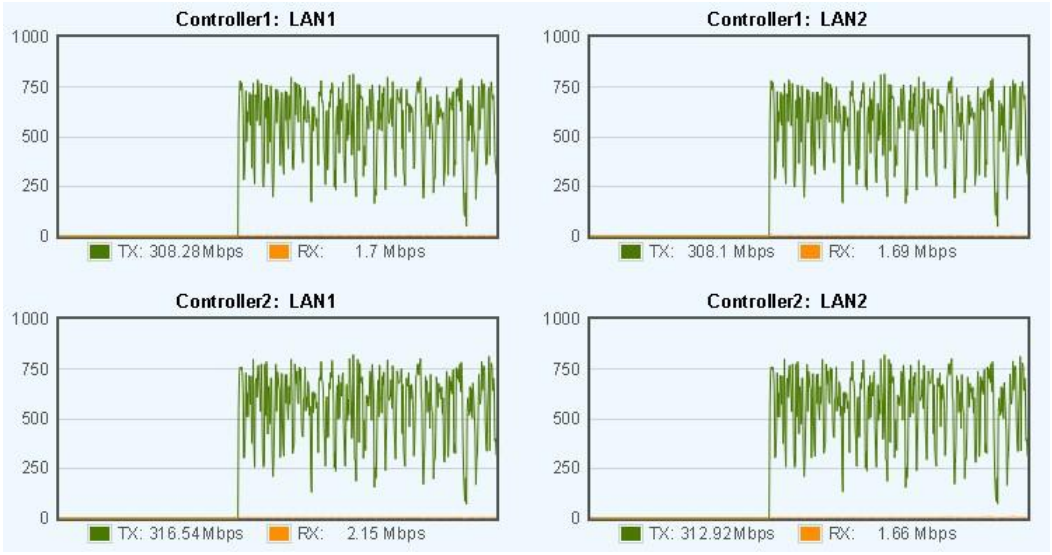

And the problematic drive (in Slot 2) will deliver lower throughputs than others.

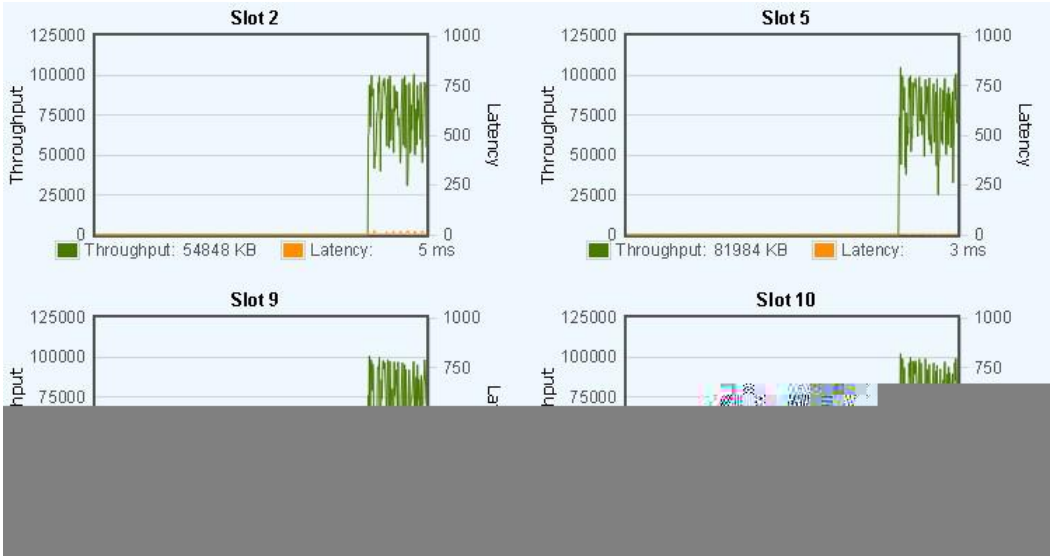

Sometimes you may see the latency of the problematic drive is relatively higher than others.

Copyright@2004~2011 Qsan Technology, Inc. All Rights Reserved.