Performance Troubleshooting Information Paper

Quantum

Revision 1.03 02/29/2016

Revision History

With every change to the document the title page will be modified to reflect current version and date. Page numbers listed on a specific revision may not be accurate on later versions of the document.

26-Feb-2016 - Version 1.00
Reformatted original document with TOC, Headings, updated header links.
27-Feb-2016 - Version 1.02
Added Revision History page
Update 1.11.3 important notes and workarounds
29-Feb-2016 - Version 1.03
Update 1.2.4 is there high % file fragmentation

Table of Contents

Performance Troubleshooting Information

Introduction

Purpose: The purpose and intent of this paper is to document some basic and common performance troubleshooting steps SPS support member need to take once the customer has identified the issue to be a performance degradation issue in their StorNext environment. This is not intended to address all performance degrading scenarios a customer can encounter in their environment but only share some of the most common and also establish and share Ideas and steps to troubleshoot these problems as they are typically complex in their nature until you can systematically narrow down to where the problem may be.
Some or most of this information might be second nature to some and not others so excuse the redundancy, so if it helps even one person, this is my intent.

Affected Products: StorNext G300, StorNext Metadata M Series Appliances, StorNext Software, StorNext QXS-1200/2400/5600, StorNext QX-1200/2400, StorNext t Q-Series and StorNext Pro Solutions.

1.1 Infomation Gathering.

On initial contact and conversation, I have been using standard disclaimer that we don't do performance tuning and that this is a professional services offering that is billable, but can check and verify health of the MDC pair and check for some typical best practices an once hardware is verified and issue is more complex then PS engagement will be recommended.

First we need to gather details about the issue and find out some basics by providing him with the Performance Issue Questions standard questionnaire found in Munich, under StorNext, Performance Issue Questions Link and request he answer all those questions to the best of his ability. The Answers he provides will help us determine then next course of action as to where we need to concentrate our efforts to get to the bottom of the issue and resolve it. The document will prompt them to do some homework on their own environment and really think about what and where the issue is and help him and us resolve the issue. I have had customers resolve their issue after they have tried to answer some of the questions after seeing the result they troubleshoot their own environment and resolve. It doesn't happen all the time, only when the customer is intimately familiar with their environment and when they have a real system administrator.

This is a good initial document we need to use to quantify and qualify his degradation, it will come handy if there is a need to do an escalation to engineering because they will request for these questions to be answered

http://10.20.9.228/products/stornext/data/Performance_Issue_Questions.doc

copied content of the document below;

General questions to ask a customer at the beginning of a performance related SR.
Add the customer's answers under the questions, then cut/paste everything into an SR note.

General Questions to qualify the issue:
What is your performance expectation?
What was the original performance and what is the performance now?
When was degradation detected?
Was the degradation gradual?
Is the degradation seen on both reads and writes?
Does the problem exist on only one client machine?
Does the problem exist when only one client is accessing the file system or only when it is under full load?
Does the problem disappear when other load is removed from the system?
Does the problem exist on all file systems or just one?
Do all clients exhibit the same behavior?
Which application(s) exhibit the performance degradation?
Do all applications exhibit the performance degradation?
Do standard tools exhibit the issue? (I.e. dd, lmdd, cvcp)
Has the application performed satisfactorily in the past?
What is the IO profile of the application? (Application docs or cvlog on the client may provide help here)

Environmental questions:
Were any firmware changes made to SAN/LAN components? (FC or network switches, disk arrays, HBA's, etc)
Were any OS changes implemented? (New service packs, major OS upgrades, etc)
Were any new clients added to the SN environment?

Work flow change questions:
Were there any changes made to workflows?
Were any new loads added to the system?
Has there been any new project data added to the system?

StorNext specific questions:
Are client to/from MDC communications slow? (Latency test from cvadmin)
Are there latencies reading/writing from/to meta data and/or journal disk? (cvlog)
Are there latencies reading/writing from/to data disks? (cvlog)
Is there a high percentage of file fragmentation?
Is there a high percentage of free space fragmentation?
Does data go buffer cache or DMA? (application docs and/or cvlog in conjunction with mount options to determine buffer cache or DMA.)
How full is the file system?

Bottom line we need to identify and eliminate common components as much as we can.
For instance, if the response we get is that all his **LAN client** are suffering degradation on a specific file system but his SAN clients are performing as expected, then from this we know the issue is with one specific file system and more specifically just LAN clients, this narrow down for us to troubleshoot his network from the **LAN client** to the **Gateway**(s) and from the **Gateway,** check for fibre to FC Data and Metadata Network to the MDC then from the MDC, check connectivity to the Metadata luns or the health of the Netapp 2600 SAS array and metadata network interface on the MDC. Of course if this is one of the newer pro solutions it won't be a Netapp but QXS array on fibre, so just keep those things in mind and keep the overall picture how stornext

data flows and metadata flows from the client I copied a good diagram to reference for those of us that need to visualize the data and metadata I/O flow.

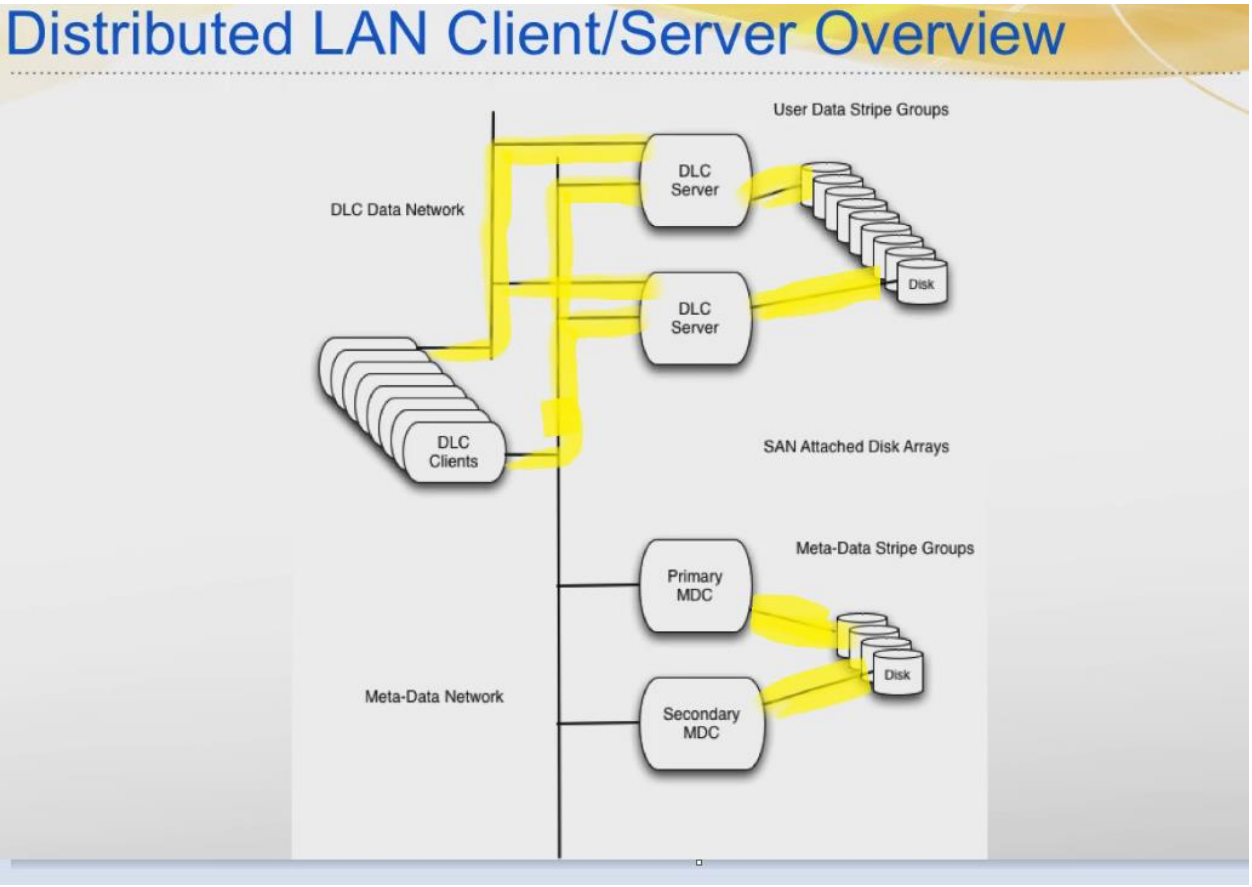

1.2 StorNext Specific Questions, Checks and Best Practices

Check affected file system bufferCacheSize and inodeCacheSize and change to best practice, run ddreadtest to check for a slow lun in the file system and run latency test to check the metadata network for network issues.

1.2.1 Check file system configuration file bufferCacheSize and inodeCacheSize

check if file system still has defaults for bufferCache Size and inodeCacheSize settings and if so we need to increase to at least 2G bufferCacheSize and 256M inodeCacheSize per file system

bufferCacheSize 2GB
inodeCacheSize 262144
or 4G and 512k
bufferCacheSize 4GB

inodeCacheSize 524288

sample check where it shows bufferCacheSize 8G and inodeCacheSize 512k;
/usr/cvfs/config/
cat ppogws.cfgx | grep -i cache
<snfs:bufferCacheSize msgId="bufferCacheSize">**8589934592**</snfs:bufferCacheSize>
<snfs:inodeCacheSize msgId="inodeCacheSize">**524288**</snfs:inodeCacheSize>
<snfs:useL2BufferCache>true</snfs:useL2BufferCache>

M440 has 48GB of memory L2 buffer caching (new in SN 5.X) will not be operational, if bufferCacheSize < 1 GB.

For M440, the suggestion would be per file system:
bufferCacheSize 2GB
inodeCacheSize 262144

M660 has total of 148GB of memory, so you can increase bufferCachesize on your file systems to something higher like 4 or 8GB. Guidance is to allocate memory to all file system, so long as we don't exceed 80% of total memory.

If you have file system that is used more you can set it up to a max of 40GB on the file system. Inode Cache size you can set to 512 on all.

1.2.2 Are there high reading /writing from/to metadata and journal and data disk? Run ddreadtest script to find slow disk

Using this ddreadtest script has help me find a lun with slow read throughput which guides us to investigate the array logs to check what is causing this LUN to have such slow throughput. of 3.6MB/s in this instance.

This script will only check reads so if there is an array with cache disabled and set to writethrough this test will not detect this and will have to check the array to verify all arrays and luns have cache enabled

End result was a slow spindle identified in the netapp array logs, see Metadata and Data Array section for specifics on how to tell;

Line 9: 10737418240 bytes (11 GB) copied, 78.543 s, 137 MB/s
Line 16: 10737418240 bytes (11 GB) copied, 90.6032 s, 119 MB/s
Line 23: 10737418240 bytes (11 GB) copied, 146.282 s, 73.4 MB/s
Line 30: 10737418240 bytes (11 GB) copied, 95.1716 s, 113 MB/s
Line 37: 10737418240 bytes (11 GB) copied, 124.066 s, 86.5 MB/s
Line 44: 10737418240 bytes (11 GB) copied, 86.8579 s, 124 MB/s
Line 51: 10737418240 bytes (11 GB) copied, 124.798 s, 86.0 MB/s
Line 58: 10737418240 bytes (11 GB) copied, 91.5078 s, 117 MB/s
Line 65: 10737418240 bytes (11 GB) copied, 162.919 s, 65.9 MB/s
Line 72: 10737418240 bytes (11 GB) copied, 92.801 s, 116 MB/s
Line 79: 10737418240 bytes (11 GB) copied, 181.318 s, 59.2 MB/s
Line 86: 10737418240 bytes (11 GB) copied, 79.1669 s, 136 MB/s
Line 93: 10737418240 bytes (11 GB) copied, 211.743 s, 50.7 MB/s
Line 100: 10737418240 bytes (11 GB) copied, 80.6575 s, 133 MB/s
Line 107: 10737418240 bytes (11 GB) copied, 114.31 s, 93.9 MB/s
Line 114: 10737418240 bytes (11 GB) copied, 75.3306 s, 143 MB/s

Line 121: 10737418240 bytes (11 GB) copied, 141.991 s, 75.6 MB/s
Line 128: 10737418240 bytes (11 GB) copied, 78.0669 s, 138 MB/s
Line 135: 10737418240 bytes (11 GB) copied, 124.58 s, 86.2 MB/s
Line 142: 10737418240 bytes (11 GB) copied, 2946.24 s, 3.6 MB/s
Line 149: 10737418240 bytes (11 GB) copied, 109.833 s, 97.8 MB/s
Line 156: 10737418240 bytes (11 GB) copied, 117.481 s, 91.4 MB/s
Line 163: 10737418240 bytes (11 GB) copied, 100.53 s, 107 MB/s
Line 170: 10737418240 bytes (11 GB) copied, 138.517 s, 77.5 MB/s
Line 177: 10737418240 bytes (11 GB) copied, 105.766 s, 102 MB/s
Line 184: 10737418240 bytes (11 GB) copied, 61.6095 s, 174 MB/s
Line 191: 10737418240 bytes (11 GB) copied, 62.114 s, 173 MB/s
Line 198: 10737418240 bytes (11 GB) copied, 60.9544 s, 176 MB/s
Line 205: 10737418240 bytes (11 GB) copied, 62.7922 s, 171 MB/s
Line 212: 10737418240 bytes (11 GB) copied, 62.1141 s, 173 MB/s
Line 219: 10737418240 bytes (11 GB) copied, 66.47 s, 162 MB/s

Find script in Munich under StorNext > scripts link;

http://munich.quantum.com/products/stornext/scripts/index.shtml

http://munich.quantum.com/products/stornext/scripts/index.shtml#ddreadtest

Usage:
`ddreadtest <file system name>`
Example run on managedfs file system:
`#./ddreadtest managedfs`

ddreadtest uses dd to read the first 10Gb from each disk in a file system and reports the read speed.

Address the issue by checking array log and check for hardware errors and slow disk(MRT)

1.2.3 Check metadata network with latency-test

It is best to run this latency-test during the time of the degradation but recommend running during normal healthy time to establish a good starting benchmark numbers to compare with results during the degradation.

1.2.4 Is there a high percentage of file fragmentation?

To check if file fragmentation is causing the degradation we need to run cvfsck –f and –e

detailed free space fragmentation report :

cvfsck -f <fsname> > **space_frag_report**

detailed file fragmentation report :

cvfsck -e <fsname> > **file_frag_report**

grep the file_frag_report file for none and sort for number on column 5 to get report in ascending order of the files with largest extents

grep none file_frag_report | sort -g -k 5 | less

-g makes it sort for numbers, -k 5 makes it sort on the 5th column.

Sample list of files with largest extents after running the sort command above, sorting from smallest to largest number o extents;
Here is the list of Files with extents larger than 10k, we need to run snfsdefrag on these directories one at a time then at the the parent directory afterwards Dalet/Storage/LSU/

4e6f26 100777 10670346184 81409 10274 1 none
Dalet/Storage/LSU/MXF_File_Ingest/GFP FOOT TEXAS CHARLIE STRONG ON BRIGHAM YOUNG PRESS CONFERENCE 090114_79036.m
xf
791948 100777 10832317420 82645 10371 1 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT SAMFORD TCU 083014_81434.mxf
e4a799 100777 11913840548 90896 11434 12 none
Dalet/Storage/LSU/MXF_File_Ingest/GFP FOOT TEXAS CHARLIE STRONG PRESSER ON UCLA 090814_81131.mxf
f0f212 100777 12027106687 91760 11471 3 none MAM/MAM INGESTS/Jermaine/watch folder/Source/9-19-2014 2-06-55 PM/CBS_Upshaw_Dierdorf_Allen_Matthews_Katie_Keane_XD50.mov
ef1fbc 100777 13169955216 100479 12636 12 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT NFL WEEK 2 IN REVIEW 091914_84015.mxf
eecf85 100777 13619671012 103910 13076 8 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT NFL HOU TEXANS WAS REDSKINS 2 OF 3 090714_84010.mxf
5832c3 100777 15123029536 115380 14423 1 none MAM/MAM INGESTS/Jermaine/watch folder/Source/9-11-2014 1-21-13 PM/08302014 Central Arkansas vs Texas Tech - TT Pageantry.mxf
2769b2d 100777 15965558676 121808 15279 6 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT NFL SF 49ERS KC CHIEFS 100514_90374.mxf
46e4ca 100555 18236448556 139134 17392 3 none MAM/MAM INGESTS/DD DUBS/ENCODES/MISC/DD CHECK 090414/MAS ICF HEISMAN PREVIEW 081414.mxf
2f5029 100555 18256260408 139285 17411 2 none MAM/MAM INGESTS/DD DUBS/ENCODES/MISC/DD CHECK 090414/MAS ICF WHOS BETTER WHOS BEST 081614.mxf
2e2bfda 100777 19921590884 151990 19063 12 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT NFL TEN TITANS CLE BROWNS GROUNDS 100514_90382.mxf
64491c 100555 19952262756 152224 19174 13 none
Dalet/Storage/LSU/MXF_File_Ingest/CBSSN EVS FOOT D2 FOOTBALL ELEMENT REEL 090214_80001.mxf
2b6b5fd 100777 22230348808 169605 21267 12 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT NFL GB PACKERS MIN VIKINGS GROUNDS 100214_90377.mxf
f5bc19 100777 22901183616 174723 22060 11 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT OKLAHOMA TENNESSEE 091314_84728.mxf

2cc72ab 100777 23686932752 180717 22658 6 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT NFL JAX JAGUARS PIT PANTHERS GROUNDS 100514_90380.mxf
c2cf71 100555 23846577448 181935 22890 11 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT AUBURN ARKANSAS 083014_80522.mxf
670635 100777 25164167704 191988 24107 11 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT BAYLOR SMU 3 OF 3 083014_81435.mxf
8054b8 100555 25744096816 196412 24735 11 none
Dalet/Storage/LSU/MXF_File_Ingest/AR MOMENTS OF NCAA MARCH MADNESS FINAL 4 042014_86084.mxf
c4f622 100777 26113792204 199233 25117 13 none
Dalet/Storage/LSU/MXF_File_Ingest/GFP FOOT MEMPHIS UCLA 090914_81440.mxf
790874 100777 28110085276 214463 27040 12 none Dalet/Drop Folders/Drop To Production/failed/GML FOOT CFB KANSAS TEXAS 20140927.mxf
801211 100555 30212522420 230504 29341 11 none
Dalet/Storage/LSU/MXF_File_Ingest/GML 9809 FOOT UTAH TCU 110610_80511.mxf
f9e669 100777 32559462938 248409 31052 6 none MAM/MAM INGESTS/Lahey/CFB 20140927 RICE at USM Melt FSN.mov
fa0e3d 100777 32409269588 247264 31128 1 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT TEXAS AM ARKANSAS 092714_87587.mxf
eded4c 100777 36089574368 275342 34550 9 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT OHIO STATE VIRGINIA TECH 090614_81556.mxf
c32ff4 100777 36521355524 278636 34965 7 none
Dalet/Storage/LSU/MXF_File_Ingest/FLD TOPS PHOTOS CARD 1 v3 090514_80293.mxf
f77692 100777 36891032688 281457 35668 6 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT NFL CAR PANTHERS DET LIONS FOX 092114_85847.mxf
eecf34 100777 37671181516 287409 36100 3 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT ALABAMA MISSISSIPPI 100414_88845.mxf
5ce72b 100777 40188201172 306612 38663 5 none
Dalet/Storage/LSU/MXF_File_Ingest/FLD FOOT NAVY KEENAN REYNOLDS INTERVIEWS 092514_87424.mxf
28c9949 100777 42959514316 327756 41089 3 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT NFL DET LIONS BUF BILLS GROUNDS 100514_90376.mxf
933f41 100777 48396682220 369238 46552 8 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT CFB IOWA STATE BAYLOR 20140927_87431.mxf
c2cdfd 100777 57250763500 436789 54717 11 none
Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT NFL DEN BRONCOS ARI CARDINALS GROUNDS 100514_90187.mxf
8686e1 100555 62930831104 480125 60586 13 none
Dalet/Storage/LSU/MXF_File_Ingest/FLD FOOT NFL CIN BENGALS GREEN SCREEN 090214_79637.mxf
a787d0 100555 65590410464 500416 63096 6 none
Dalet/Storage/LSU/MXF_File_Ingest/FLD FOOT NFL MIA DOLPHINS GREEN SCREEN 090214_79664.mxf
3bcb58 100555 75299902404 574493 72386 7 none
Dalet/Storage/LSU/MXF_File_Ingest/FLD FOOT NFL OAK RAIDERS GREEN SCREEN 090214_79689.mxf

30c9b5 100555 78895935897 601929 75242 9 none MAM/MAM INGESTS/DD DUBS/ENCODES/MISC/DD CHECK 090414/MAS COURT OF CHANGE SPLIT.mov
2f528a 100555 82915488564 632595 79761 13 none Dalet/Storage/LSU/MXF_File_Ingest/FLD FOOT NFL NY GIANTS GREEN SCREEN 090214_79672.mxf
10a9403 100777 90173178972 687967 86215 1 none Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT NFL CAR PANTHERS CHI BEARS GROUNDS 100514_90229.mxf
f16c5d 100555 97280012652 742188 93363 2 none Dalet/Storage/LSU/MXF_File_Ingest/FLD FOOT NFL PACKERS RAIDERS TNF TECH EMMY SHOOT 082814_1_84075.mxf
e4a5be 100777 102726085968 783738 98226 5 none Dalet/Storage/LSU/MXF_File_Ingest/GML FOOT NFL IND COLTS BAL RAVENS 100514_90190.mxf

To defrag a heavily fragmented directory;

1. To defrag a **heavily fragmented**, run **#snfsdefrag –rv** <target directory>

Fyi, this customer's initial largest extents was 90K extents and had to run snfsdefrag for several weekends to defrag to this point and was able to get noticeable difference in performance once he defragged the largest extents.

1.3 Run Latency_test to test metadata network client latencies

It is best to run this latency-test during the time of the degradation but recommend running during normal and healthy time to establish a good starting benchmark numbers to compare with results during the degradation.

Getting the customer to see the results is key, so they can agree if this is the issue or we can move on to something else after eliminating this as the source of the degradation.

How can you tell if the latencies are ok or to high? Depends on what normal looks like when they run latency-test from the mdc to all clients when they said all is normal now customer are not complaining.

Usually the problem client will stickout like a sore thumb and investigation should be done by the customer's network team. As you can see from the sample output half second is to high, and have seen customer with even greater latencies which prompts investigation for that client. Usually are networks issue on the client, DNS resolution or jumbo frames were enabled on that client but not on the rest of the environment.

Sample latency-test syntax and output;
#cvadmin
#>Select snfs1
#snadmin (snfs1) > **latency-test all 10**

Test started on client 1 (MDC-SRV01.)...latency 38us
Test started on client 8 (10.232.86.37)...latency 147us
Test started on client 10 (10.232.86.57)...latency 146us
Test started on client 12 (10.232.86.40)...latency 123us

Test started on client 13 (10.232.86.38)...latency 105us
Test started on client 16 (192.168.150.49)...latency 500000us
Test started on client 43 (192.168.150.47)...latency 38461us
Test started on client 46 (192.168.150.72)...latency 1203us

Sample 2 where highest latency is not that bad but sticks out from the rest identifying client with network issues compared to the rest that may merit investigation;

snadmin> latency-test all 10
No FSS selected.

snadmin> select 2
Select FSM "2"

Created	: Sat Mar 1 12:42:01 2014
Active Connections:	13
Fs Block Size :	64K
Msg Buffer Size :	8K
Disk Devices :	5
Stripe Groups :	3
Fs Blocks :	570972768 (34.03 TB)
Fs Blocks Free :	83661235 (4.99 TB) (14%)

snadmin (fsname) > select 2 [25Glatency-test all 10
Test started on client 1 (172.16.122.19)...latency 168us
Test started on client 2 (172.16.122.30)...latency 167us
Test started on client 3 (172.16.250.188)...latency 83333us
Test started on client 4 (10.11.12.40)...latency 5847us
Test started on client 6 (mdc2) latency 27us
Test started on client 9 (10.11.12.31)...latency 168us
Test started on client 10 (10.11.12.30)...latency 177us
Test started on client 11 (10.11.12.13)...latency 273us
Test started on client 12 (10.11.12.45)...latency 212us
Test started on client 13 (172.16.122.20)...latency 170us
Test started on client 14 (10.11.12.56)...latency 167us
Test started on client 15 (mdc1)...latency 88us
Test started on client 16 (172.16.122.48)...latency 168us

1.3.1 Important notes and workarounds

.
To defrag a heavily fragmented directory;

1. To defrag a **heavily fragmented**, run **#snfsdefrag –rv** <target directory>

We can also take advantage of Sustaining Engineering's StorNext Qustats Repository but this is not my expertise and usually Escalate if I need assistance deciphering what those charts mean and recommendations from those
http://10.20.234.57/cgi-bin/susrepo

SAN Clients

1.4 Linux

1.4.1 Description

This document contains information specific to Linux and performance related issues . Typical logs are needed to check if hardware issue exists on the fibre adapter or metadata network issues. Capture cvgather from the affected client(s).

If customer has Brocade switches purchased from quantum, we need to capture switch logs, as they will provide a better overall picture of the san perturbations if any.

See Fabric Switches section for more details.

1.4.2 Log analysis, indicators of fibre SAN perturbations and indicators of hardware issues on the fabric

Messages log, analysis.

Sample reset events that gets logged are shown below; these are indicators of fabric issues or array issues , these are logged by the host/clients trying to access the device and recover.

.

if you see these DEVICE, TARGET, BUS and ADAPTER RESETs in the client logs and if they show up in the MDC also, then these, merit checking/analyzing fabric switches and or the Array hosting the data luns as these indicate the issue is something they all share in common, whether the fabric or the array or lun specific.

hba resets in the logs, grep for "reset" in messages logs to identify hba, san issue or target issues;

Nov 6 10:58:20 QMDC2 kernel: qla2xxx [0000:0e:00.1]-8009:9: **DEVICE RESET ISSUED**
Nov 6 10:58:41 QMDC2 kernel: qla2xxx [0000:0e:00.1]-800f:9: **DEVICE RESET FAILED**: Task
Nov 6 10:58:41 QMDC2 kernel: qla2xxx [0000:0e:00.1]-8009:9: **TARGET RESET ISSUED**
Nov 6 10:59:01 QMDC2 kernel: qla2xxx [0000:0e:00.1]-800f:9: **TARGET RESET FAILED**: Task
Nov 6 10:59:01 QMDC2 kernel: qla2xxx [0000:0e:00.1]-8012:9: **BUS RESET ISSUED** nexus=9:44:2.
Nov 6 11:01:16 QMDC2 fsm[34427]: StorNext FSS <fsname>[0]': OpHangLimitSecs exceeded VOP-Class-0 Type-9 Subtype-7 362 ticks Conn[23] Thread-0x7f41a8221700 Pqueue-0x29ecea0 Workp-0x7f40e8003028 MsgQ-0x7f40e8003010 Msg-0x7f40e8003088 now 476947 started 476585 limit 360 (180 secs.)
Nov 6 11:01:16 QMDC2 fsm[34427]: StorNext FSS '<fsname>[0]': PANIC: /usr/cvfs/bin/fsm "OpHangLimitSecs exceeded VOP-Class-0 Type-9 Subtype-7 362 ticks Conn[26] Thread-0x7f43f3e79700 Pqueue-0x29ecea0 Workp-0x7f418c0008d8 MsgQ-0x7f418c0008c0 Msg-0x7f418c000938 now 476947 started 476585 limit 360 (180 secs.) " file /scm/nightly/VM-0-RedHat60AS-26x86-64-SP0/sn/snfs/fsm/queues.c, line 372
Nov 6 11:01:23 QMDC2 kernel: qla2xxx [0000:0e:00.1]-802b:9: **BUS RESET SUCCEEDED** nexus=9:44:2.

Nov 6 11:01:33 QMDC2 kernel: qla2xxx [0000:0e:00.1]-8018:9: **ADAPTER RESET ISSUED**
Nov 6 11:01:37 QMDC2 kernel: qla2xxx [0000:0e:00.1]-8017:9: **ADAPTER RESET SUCCEEDED** nexus=9:44:2.
Nov 6 11:05:04 QMDC2 kernel: imklog 5.8.10, log source = /proc/kmsg started.

1.4.3 Fixes

- Reference StorNext Tuning guide, page 47, under Cpuspeed issue on Linux.

  Cpuspeed Service Issue on Linux

  Cpuspeed, an external Linux service on recent Intel processors, is not correctly tuned to allow StorNext to take advantage of processor speed. Suse systems may also be impacted, as may AMD processors with similar capabilities. On processors with a variable clockspeed (turboboost), the cpuspeed service on Redhat controls the actual running speed of the processors based on system load. A workload such as a heavily used FSM and probably Storage Manager does not register as something which needs a faster cpu. **Turning off the cpuspeed service has been shown to double metadata performance on affected hardware**. Looking at the reported CPU clock speed by doing `cat /proc/cpuinfo` while the system is under load shows if a system is impacted by this issue

1.4.4 Important notes and workarounds

Taken from tuning guide;

Disabling CPU Power Saving States

It is strongly recommended that CPU power saving states be disabled if they are supported by your system's
CPUs. Making this change may result in increased power consumption but improves stability and
performance. This applies to all systems running StorNext including clients.

How to Disable CPU Power Saving States on a Redhat Or Suse Linux System

To disable CPU power saving states on a RedHat or SuSE Linux system, perform the following procedure:
1. Add the following text to the "kernel" line in `/boot/grub/menu.lst`:
`intel_idle.max_cstate=0 processor.max_cstate=1`
2. Reboot the system for the change to take effect.

How to Disable CPU Power Saving States on a Debian Linux System

To disable CPU power saving states on a Debian Linux system, perform the following procedure:
1. Modify the `GRUB_CMDLINE_LINUX` string in `/etc/default/grub` so that it contains:
`intel_idle.max_cstate=0 processor.max_cstate=1`
2. Run "**update-grub**".
3. Reboot the system for the change to take effect.
Note: Disabling CPU power saving states in the system BIOS has no effect on Linux.
In some cases, performance can also be improved by adjusting the idle kernel parameter. However, care

should be taken when using certain values. For example, `idle=poll` maximizes performance but is
incompatible with hyperthreading (HT) and will lead to very high power consumption. For additional
information, refer to the documentation for your version of Linux.
On Windows, disable CPU power saving states by adjusting BIOS settings. Refer to system vendor
documentation for additional information.

1.5 Windows

1.5.1 Performance best practice recommendations;

- Updating Individual Buffer Size to 512K and Minimum Total Cache Size to 256M is typical best practice setting for best performance, from defaults

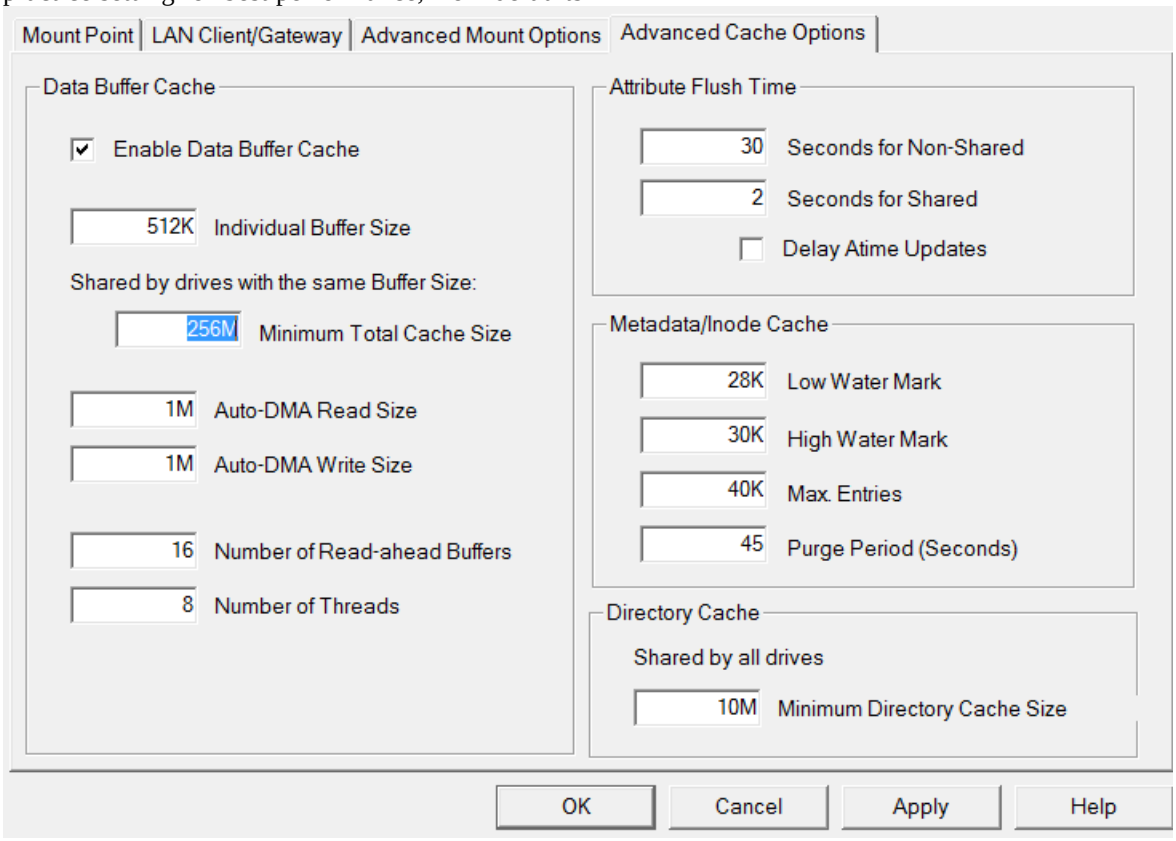

- Enable and use Multipathing and Vendor device specific module DSM when this is made available by the array vendor, as it will increase throughput when doing i/o through all fibre ports to all target ports and typically best to use Shortest Queue Service Time(SQST) load balancing if the array is ALUA aware, but other load balancing options can be used and available. Recommend customer engage with the array vendor for those specifics as they will depend on the hba used, array firmware, OS version etc.

1.5.2 Important notes and workarounds

This section identifies important notes and workarounds for

1.6 Mac OS X

1.6.1 Mac OS X Performance best practice recommendations

There are a few best practice setting for Apple Mac OS X client to ensure best performance is achived and higher availability

- Disable sleep on a MAC to prevent this client from going to sleep and keeping a lock on a file that another client will try to access and during this time before StorNext releases the lock, it will cause observable degradation to the rest of the clients in the environment.
  Usual indicator of a client with a lock on a file that may have gone to sleep, is a Timeout while attempting to force data flush for file RAS event in MDC log and log in the syslog.log on the mac client.

  Sample error from an actual client where we disabled sleep to resolve;

  Feb 19 09:52:43 mdc fsm[383]: StorNext FSS 'snfs1[0]': Timeout while attempting to force data flush for file '/VOD/Ingest/PROD/Provider_Delivery/TVNAVAILUS/TVN_Avail_US_SNOW_ANGELS_16x9_5_1_ENGLISH_FEATURE_E0095188_8165355.m2t' (inode 52385225) for file system 'snfs1' on host **172.16.6.56**. Allowing host **172.16.6.21** to open file. **This may cause data coherency issues**. The data path for host 172.16.6.56 should be inspected to confirm that I/O is working correctly. Further flush failure messages will be suppressed until the file is closed on all systems.

  Not all the times does the event above is caused by a mac client going to sleep, It could also be that the client has the file lock because its busy rendering the file or it could be that the client lost complete access to the file.

  To disable Sleep on a MAC run these two commands:

  # sudo pmset -a sleep 0

  # sudo pmset –a disablesleep 1

- Mac does not rely on hosts file, but have seen customers using hosts file only and resolving ok, nevertheless, it relies on DNS to resolve forward and reverse lookup and if any Mac client is not able to resolve either forward or reverse zones, it will flood the network trying to resolve and cause performance degradation when accessing the metadata or if LAN client the data network
  To check if the affected client resolves run host <ip> of the local machine public and metadata network ips and should return a Fully Qualified Domain Name (FQDN) and if you run command host <FQDN> or <alias> it should return the IP.

  # host admins-Mac-mini
  admins-Mac-mini.quantum.com mail is handled by 5 ppointrelay1.quantum.com.

  #host <IP>
  #nslookup <IP> should return <fqdn> if you run from linux or window system.

  If this test fails on any mac client then this is causing the degradation in their environment.

1.6.2 Important notes and workarounds

Mac client going to sleep, indicators to check if the client is going to sleep “Previous Sleep Cause” “Maintenance wake”, “ or “SuspendDevice” and clock jumps forward and backwards in mac client syslog.log logs.

To disable Sleep on a MAC run these two commands:

 # sudo pmset -a sleep 0

 # sudo pmset –a disablesleep 1

LAN clients

Visual DLAN Client/Server diagram highlighting the paths to troubleshoot when only LAN clients are suffering from degradation and if only one LAN client is degraded then that narrows it down to single client and troubleshoot both the DLC Data Network and the Metadata Network from the client perspective.

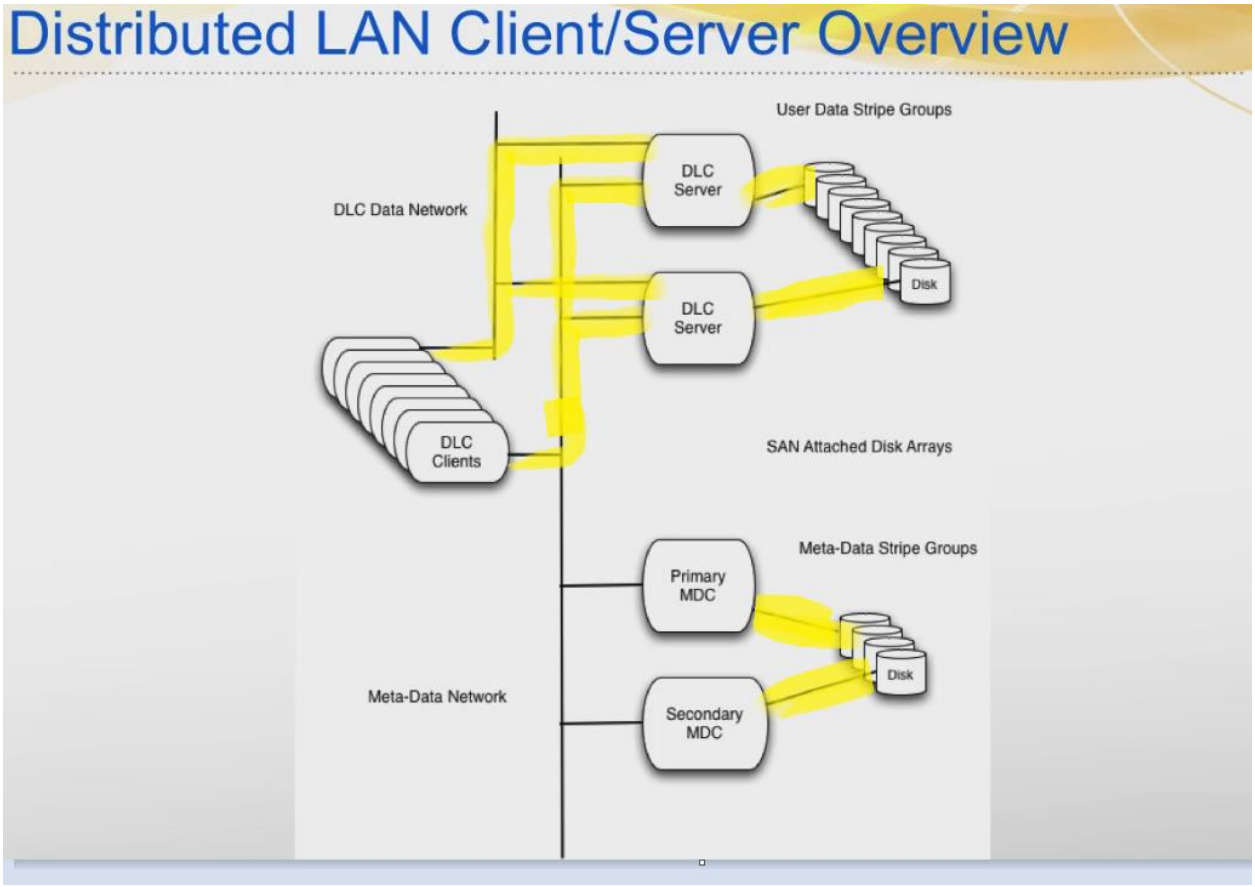

If all clients are suffering from degraded performnac then we need to capture cvgather from DLC client, snapshot, cvgather and DSET from DLC Server and snapshot from MDC.

Note from Performance Tuning Guide:
StorNext Gateway Server Performance
If your configuration includes StorNext LAN Clients, Quantum strongly recommends that the machines you use for your gateway servers should not also be configured as metadata controllers. The exception to this recommendation is the StorNext M660 Metadata Appliance, which is specifically manufactured to handle this workload. Doing so may not only cause performance degradation, but also expose the virtual IPs to additional vulnerability. For best performance, machines used as gateway servers should be dedicated machines.

What is the theoretical maximum throughput of a Gigabit Ethernet interface?

We need to know this in the event a customer is expecting more than the theoretical max. (happened to me)
Theoretical throughput of Gigabit Ethernet with jumbo frames, and using TCP:
997Mbps - .886 - 1.33 - 1.55 - .443 - 2.21 - 2.21 - 1.33 = **987Mbps** or **123MB/s.**

https://kb.netapp.com/support/index?id=3011938&page=content&locale=en_US

check messages log for timeouts to the proxy

sample errors from one of customers;

Grep for "proxy write timed" out to confirm proxy clients disconnecting and timing out;

[0714 18:23:01] 0x14451f0 FSD_NOTICE Proxy client disconnecting from host **'10.146.176.21'**
[0714 18:23:04] 0x14451f0 FSD_NOTICE **Proxy write timed out after 30s** on server **10.146.176.21**
[0714 18:23:06] 0x14451f0 FSD_INFO Reconnecting to Disk Proxy Server **10.78.78.1** via **10.146.176.21** port 43426
[0714 18:23:06] 0x14451f0 FSD_INFO Connected to Disk Proxy Server **10.78.78.1** via **10.146.176.21** port 43426

[0714 18:24:04] 0x14451f0 FSD_NOTICE Proxy client disconnecting from host '10.78.78.1'
[0714 18:24:07] 0x14451f0 FSD_NOTICE **Proxy write timed out after 31s on server 10.78.78.1**
[0714 18:24:09] 0x14451f0 FSD_INFO Reconnecting to Disk Proxy Server **10.78.78.1** via **10.78.78.1** port 59971
[0714 18:24:09] 0x14451f0 FSD_INFO Connected to Disk Proxy Server **10.78.78.1** via **10.78.78.1** port 59971

Jumbo Frames, in all cases we had to disable because there will always be some one that adds new host or changes are made to the enviornement with and without jumbo frames which is bad if you want best performance.

MTU Issues

Jumbo Ethernet frames can increase performance by a factor of 2-4 on 10G paths on older hardware. Even on modern hardware, we see 9.9Gbps vs 5Gbps for UDP, and 9.9Gbps vs 9.3Gbps for TCP. For 40G paths, the performance gains are even greater. Single stream tests of TCP were observed to be 25Gbps vs. 10Gbps. UDP testing showed an improvement of 15Gbps vs. 8Gbps.

Use of Jumbo frames has two advantages: One is that for a given data throughput, the packet rate is less for jumbo frames vs. standard frames, requiring less CPU for packet processing. The other advantage is that recovery in data throughput after a loss event is proportional to the maximum segment size - so with jumbo frames you get a ~6x faster recovery rate from loss events. For more information on the benefits of Jumbo Frames, see: http://staff.psc.edu/mathis/MTU/

ping can be used to verify the MTU size. For example, on Linux you can do:

```
ping -s 8972 -M do -c 4 10.200.200.12
```

Scamper or tracepath are good tools to help verify the path MTU size too.

Note that there are some downsides to jumbo frames as well. **All hosts in a single broadcast domain must to be configured with the same MTU, and this can be difficult and error-prone.** Ethernet has no way of detecting an MTU mismatch - this is a layer 3 function that requires ICMP signaling in order to work correctly. (Unfortunately some sites block ICMP, which breaks path MTU discovery.

Therefore a good approach is often to create a new jumbo frame enabled subnet for your high-speed data transfer hosts.

For Linux hosts using Jumbo Frames, we recommend setting tcp_mtu_probing = 1 to help avoid the problem of MTU black holes. Setting it to 2 sometimes causes performance problems. For more information on this setting, see http://kb.pert.geant.net/PERTKB/PathMTU

The use of jumbo frames may slightly increase latency when performing very small I/O.

1.6.3 Important notes and workarounds

Jumbo Frames Notes and Cveats on DLC clients and Gatway systems;
To be effective, Jumbo frames must be supported and enabled on all networking hardware and software components between the sender and receiver. Not all Jumbo frame implementations are compatible. When specifying a larger MTU, depending on the context, the value may need to include the header. For example, `9216` versus `9000`. The use of jumbo frames may slightly increase latency when performing very small I/O. The benefit of reduced CPU utilization may not be significant if TCP offload is also enabled.

If multiple gateways;
Proxypath Mount Option algorithm on the client had **'File Sticky'** ie mean that all the IO for that file goes via a single Proxy which kept all the I/O on a single Gateway instead of using both Gateways and recommended to switch the algo to only **balance** so the load can be shared among the Proxy servers.. changing also to **Balance** may not make a difference if environment is experiencing network issues.

The **balance** algorithm - attempts to keep the same amount of time's worth of I/O outstanding on each connection
The **rotate** algorithm - attempts to keep the same number of bytes of I/O pending on each Gateway connection
The **sticky** algorithm - assigns I/O to specific LUNs to specific Gateway connections.

Proxypath Mount Option - Controls the algorithm used to balance I/O across Gateway connections. The following options may be specified:

Balance - Use the balance algorithm and do not use file sticky behavior
Rotate - Use the rotate algorithm and do not use file sticky behavior
Sticky - Use LUN sticky behavior
File Sticky Balance - Use the balance algorithm and use file sticky behavior
File Sticky Rotate - Use the rotate algorithm and use file sticky behavior

snapshot and cvgather from both gateways and DSET will capture necessary logs and in addition, Capture output of **cvdb –x** from both Gateways and clients and **cvdb –bv** on affected DLC clients during the busy hours while you are experiencing the worse degradation.

1.7 Windows

1.7.1 Performance Best Practice

- Same setting apply to LAN client;
- Updating Individual Buffer Size to 512K and Minimum Total Cache Size to 256M is typical best practice setting for best performance, from defaults

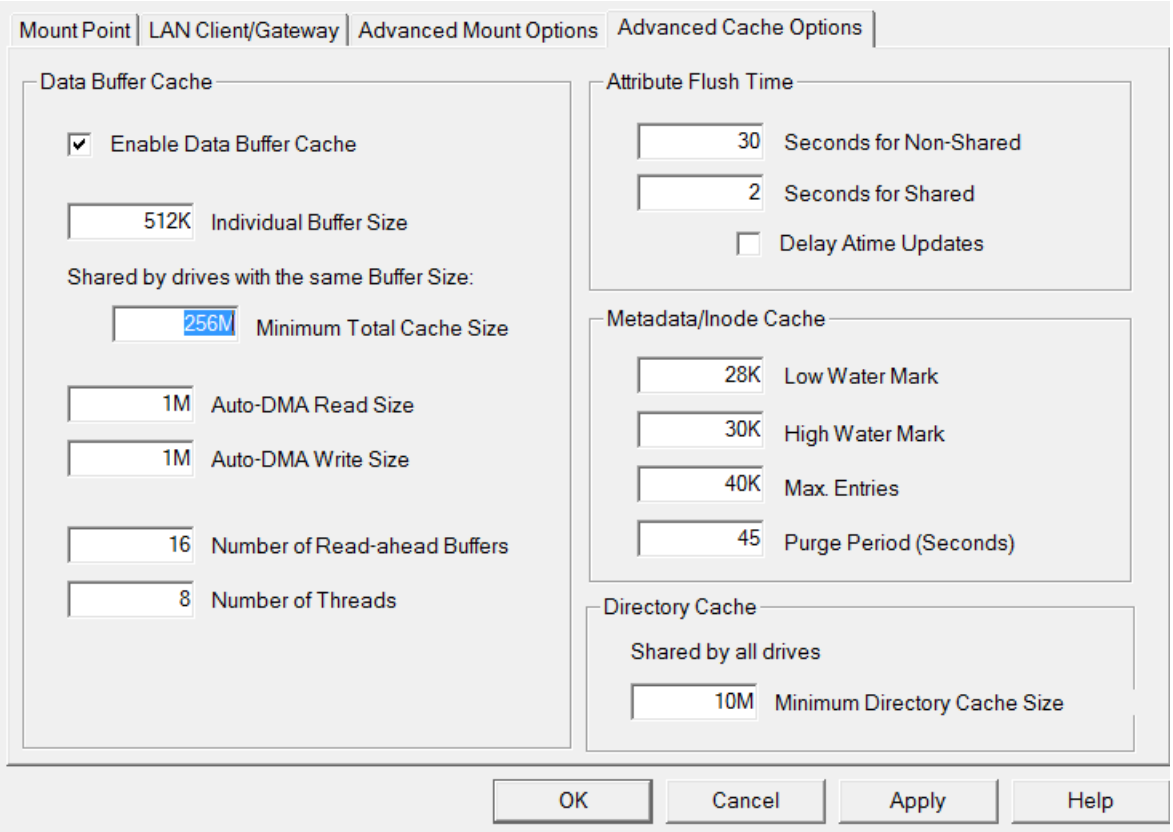

1.7.2 Important notes and workarounds

This section identifies important notes and workarounds for

1.8 Mac OS X

1.8.1 Mac OS X Performance best practice recommendations

Same recommendations as SAN client for best practice setting for Apple Mac OS X client to ensure best performance is achived and higher availability

Disable sleep on a MAC to prevent this client from going to sleep and keeping a lock on a file that another client will try to access and during this time before StorNext releases the lock, it will cause observable degradation to the rest of the clients in the environment and make sure client resolves forward and reverse lookup

1.8.2 Important notes and workarounds

Mac client going to sleep, indicators to check if the client is going to sleep “Previous Sleep Cause” “Maintenance wake”, “ or “SuspendDevice” and clock jumps forward and backwards in mac client syslog.log logs.

To disable Sleep on a MAC run these two commands:

 # sudo pmset -a sleep 0

 # sudo pmset –a disablesleep 1

Metadata and Data Arrays

1.9 Metadata Array (Netapp 2600 SAS Array)

1.9.1 Qarray1 Metadata Array

These arrays we can do some specific checks for but recommend we engage Netapp for assistance and array health confirmation if not sure of health status and confirmation if the array is logging any hardware errors that may cause degradation.

1.9.2 Check for slow disk on metadata array QD6000

Slow disk can identified in the Netapp array logs by checking the "states-capture-data.txt" log under "luall 3" section from both controllers, the counters of interest are Average Response Time **ART(uSec)** and Max Response Time **MRT(uSec).**

Sample counters from an actual customer issue whre replacing disk on slot 53, fixed the performance degradation. QD6000 on t99, s53 and fibre issues on these initiator ports identified below;
Disk on tray 99 slot 53 logged Max Response Time of 1.7 and 1.7 seconds which is to long for a spindle and latencies up to 4.8 seconds on the initiator, way to long.

Read/Write Completions
...IDITN..:.L Unit..:.....R E A D S.....:.....W R I T E S.....:
iditn Ch :Location : #Success BlksXfered ART(uSec) MRT(uSec) : #Success BlksXfered ART(uSec) **MRT(uSec)** :#Errs
------ -- :-------- :---------- ---------- --------- --------- :---------- ---------- --------- --------- :-----

115 0 :t99,s53 : 622718 158791813 100116 **1840505** : 248210 63394923 95797 **1724546** : 0

115 1 :t99,s53 : 640509 163583215 73070 **1659238** : 248661 63447758 77521 **1873901** : 0

Read/Write Completions
.....TDITLN.....:Initiator.:.....R E A D S.....:.....W R I T E S.....:
tditn LUN LU(SSID) Ch : Location : #Success BlksXfered ART(uSec) **MRT(uSec)** : #Success BlksXfered ART(uSec) MRT(uSec) :#Errs
------- --- -------- -- :--------- :---------- ---------- --------- --------- :---------- ---------- --------- --------- :-----

```
   7  5 00000005  2 :  ICE01 :   143043  247220332    6165  2841325 :   82900  115537097
37643   2128149 :   0

   8  5 00000005  2 :  ICE02 :   139324  241792518    3856  2976025 :   11939    215816
1282   1944757 :   0

   9  5 00000005  2 :  ICE03 :   194656  338035107    3163  2902039 :   13562    309168
   942   1154885 :   0
  10  5 00000005  2 :  ICE101 :   143140  248297178    6419  2841460 :      1     128
157     157 :   0

  11  5 00000005  2 :  ICE201 :   192071  334169876    3577  4608265 :      0       0      0
0 :   0
  14  5 00000005  2 :SAN1-FTP2 :   904405  115763840     749  1471414 :  2050434
262455552    8194   5870410 :   0
```

1.10 Data Array QDSeries(QD6000/7000)

1.10.1 Description

These arrays we can do some specific checks for but recommend we engage Netapp for confirmation if the array is logging any hardware errors that may cause degradation.

1.10.2 Check for slow disk

- **Use same counters as the metadata array for Netapp array.**

1.10.3 Check array health Cache Battery status optimal status

If cache battery is failed, then read/write cache will be suspended an degrade performance and if cache battery is failed on only one controller then recommend placing the controller with failed cache battery offline to enable cache on the surviving controller then disable "write cache with mirroring" on single remaining controller. This will set the remaining single controller with cache enabled to improve performance while waiting on the replacement controller cache battery or controller

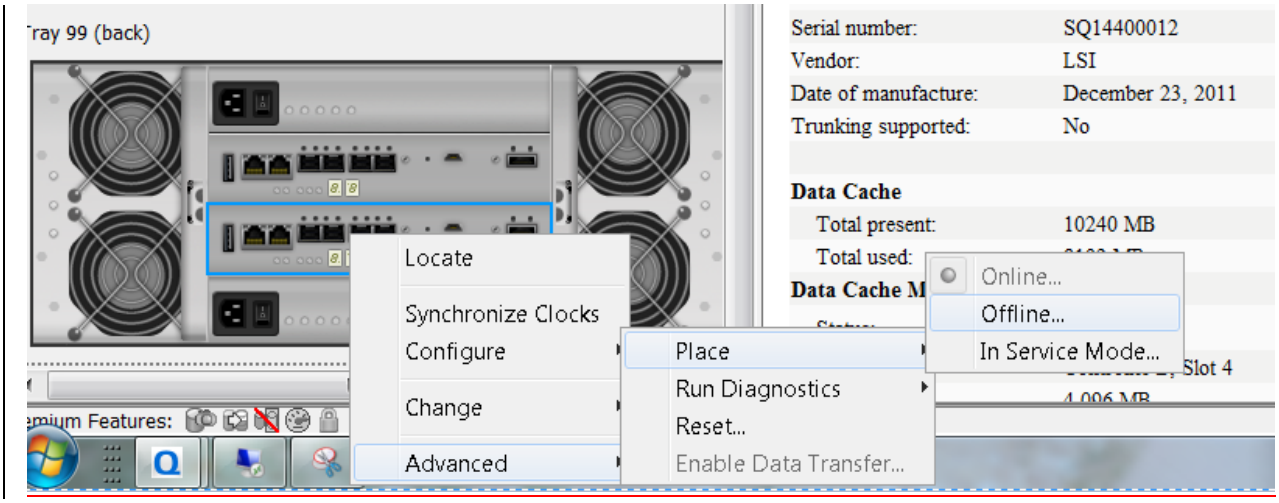

Check SAN Fabric Switches

1.11 Brocade Switches

switch logs have a lot of details about their environment that will help determine if the SAN is logging fibre errors significant enough to cause performance degradation, whether hardware errors or misconfiguration

1.11.1 Switch log (supportshow)

Supportshow or supportsave counters of interest to check for errors are switchshow, porterrshow, portshow <port#> and portstatsshow.

Switchshow will tell us which hba wwpn is connected to which port, so we need to know the wwpn of the affected client or the alias we can look up in the switchlogs.

Sample output and use of switchshow, porterrshow and porterrshow, help determine the connected devices logging hardware errors, error counters of concern are (enc in , crc err, crc g_eof, too short, too long, and bad eof. As it turned out for this customer, these ports with errors were ISL ports going to where most clients were connected, so to address this issue we had to replace cables and sfps and move some clients local to this switch where the QX array was located.

The main thing to learn here is to identify the ports logging errors and provide customer with a course of action, whether to engage Brocade to help resolve or troubleshoot on their own.

Most time we need to capture a second set of logs couple days later and check logs again to get a delta see which ports are incrementing the errors and verify these are not stale old errors.

switchName: swname
switchType: 118.1

```
switchState:       Online
switchMode:      Native
switchRole:       Subordinate
switchDomain:   1
switchId:          fffc01
switchWwn:       10:00:50:eb:1a:86:68:f7
zoning:             ON (Default)
switchBeacon:   OFF
FC Router:         OFF
FC Router BB Fabric ID:   1
Address Mode:   0

Index Port Address Media Speed      State   Proto
==================================================
  0   0   010000   id   N8     Online     FC  F-Port  21:00:00:24:ff:8b:00:7a
  1   1   010100   id   N8     Online     FC  F-Port  21:00:00:24:ff:8b:00:ee
  2   2   010200   id   N16            No_Light    FC
  3   3   010300   id   N16            No_Light    FC
  4   4   010400   id   N16            Online      FC  F-Port  20:70:00:c0:ff:25:c4:11
  5   5   010500   id   N16            Online      FC  F-Port  21:70:00:c0:ff:25:c4:11
  6   6   010600   id   N16            Online      FC  F-Port  24:70:00:c0:ff:25:c4:11
  7   7   010700   id   N16            Online      FC  F-Port  25:70:00:c0:ff:25:c4:11
  8   8   010800   id   N16            Online      FC  F-Port  20:70:00:c0:ff:25:c9:23
  9   9   010900   id   N16            Online      FC  F-Port  21:70:00:c0:ff:25:c9:23
 10  10   010a00   id   N16            Online      FC  F-Port  24:70:00:c0:ff:25:c9:23
 11  11   010b00   id   N16            Online      FC  F-Port  25:70:00:c0:ff:25:c9:23
 12  12   010c00   id   N8             Online      FC  F-Port  10:00:00:10:86:04:43:c6
 13  13   010d00   id   N8             Online      FC  F-Port  10:00:00:10:86:04:43:c7
 14  14   010e00   id   N8             Online      FC  F-Port  21:00:00:24:ff:8a:31:52
 15  15   010f00   id   N8     Online     FC  F-Port  21:00:00:24:ff:8a:31:53
 16  16   011000   id   N16            Online      FC  F-Port  50:01:43:80:23:1c:75:d4
 17  17   011100   id   N16            Online      FC  F-Port  50:01:43:80:23:1c:75:d6
 18  18   011200   id   N8             Online      FC  G-Port
 19  19   011300   id   N4             Online      FC  F-Port  10:00:00:10:86:04:43:ac
 20  20   011400   id   4G             Online      FC  LE E-Port  10:00:50:eb:1a:86:60:38 "swname"
(upstream)(Trunk master)
 21  21   011500   id   4G             Online      FC  LE E-Port  (Trunk port, master is Port 20 )
 22  22   011600   id   4G             Online      FC  LE E-Port  (Trunk port, master is Port 20 )
 23  23   011700   id   4G             Online      FC  LE E-Port  (Trunk port, master is Port 20

porterrshow       :
        frames      enc    crc    crc    too    too    bad    enc    disc   link   loss   loss   frjt   fbsy  c3timeout    pcs
        tx     rx   in     err    g_eof  shrt   long   eof    out    c3     fail   sync   sig                  tx     rx     err
  0:    2.1m   1.3m 0      0      0      0      0      0      0      82     0      0      0      0      0      0      0      0
  1:    2.3m  48.6m 0      0      0      0      0      0      0      87     0      0      0      0      0      0      1      0
  2:   13.4k   5.3k 0      0      0      0      0      0      2.0k   0      0      0      1      0      0      0      0      0
  3:    0      0    0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0
  4:   62.0m  40.5m 0      0      0      0      0      0      0      32     0      0      0      0      0      0      32     0
  5:   41.7m  21.2m 0      0      0      0      0      0      0      45     0      0      0      0      0      0      45     0
  6:   41.0m  20.8m 0      0      0      0      0      0      0      21     0      0      0      0      0      0      21     0
  7:   62.3m  40.7m 0      0      0      0      0      0      0      25     0      0      0      0      0      0      25     0
  8:   47.8m  34.4m 0      0      0      0      0      0      0      29     0      0      0      0      0      0      29     0
  9:   52.1m  37.3m 0      0      0      0      0      0      0      79     0      0      1      0      0      0      79     0
 10:   24.3m  15.3m 0      0      0      0      0      0      0      49     0      0      0      0      0      0      49     0
 11:   29.2m  18.6m 0      0      0      0      0      0      0      19     0      0      1      0      0      0      19     0
 12:    5.9m   5.5m 0      0      0      0      0      0      0      59     7      0      7      0      0      0      23     0
 13:    5.5m   5.5m 0      0      0      0      0      0      0      57     7      0      7      0      0      0      21     0
 14:    4.6m   1.9m 0      0      0      0      0      0      58     280    2      0      2      0      0      86     6      0
 15:    1.6m 748.4k 0      0      0      0      0      0      0      723    4      0      3      0      0      706    5      0
 16:  120.6m 168.0m 0      0      0      0      0      0      0      16     4      0      4      0      0      0      16     0
 17:   68.0m 111.2m 0      0      0      0      0      0      0      18     4      0      4      0      0      0      18     0
 18:    9.0m   9.0m 108    30     28     0      0      2      3.5g   18     10     0      7      0      0      0      0      0
 19:    9.0m   8.5m 0      0      0      0      0      0      206.4k 33     18     0      7      0      0      0      15     0
 20:  117.7k 107.0k 10.2k  6.4k   6.3k   0      0      93     4.0g   657    1      0      7      0      0      66     454    0
 21:  134.2k 119.6k 191    54     49     0      0      5      2.4g   95     2      0      7      0      0      66     0      0
 22:  102.8k  92.3k 25.8k  10.7k  10.5k  0      0      177    4.2g   302    8      0      8      0      0      66     0      0
 23:  141.7k 127.2k 14     10     10     0      0      0      94.2m  103    1      0      7      0      0      66     0      0

portshow 18
```

```
porterrshow        :
          frames      enc    crc    crc    too    too    bad    enc   disc   link
loss   loss   frjt   fbsy  c3timeout    pcs
       tx     rx      in    err    g_eof  shrt   long   eof     out   c3    fail
sync   sig                  tx    rx    err
  0:    2.1m   1.3m   0      0      0      0      0      0       0     82      0
0      0      0      0      0      0      0
  1:    2.3m  48.6m   0      0      0      0      0      0       0     87      0
0      0      0      0      0      1      0
  2:   13.4k   5.3k   0      0      0      0      0      0      2.0k   0       0
0      1      0      0      0      0      0
  3:    0      0      0      0      0      0      0      0       0      0      0
0      0      0      0      0      0      0
  4:   62.0m  40.5m   0      0      0      0      0      0       0     32      0
0      0      0      0      0     32      0
  5:   41.7m  21.2m   0      0      0      0      0      0       0     45      0
0      0      0      0      0     45      0
  6:   41.0m  20.8m   0      0      0      0      0      0       0     21      0
0      0      0      0      0     21      0
  7:   62.3m  40.7m   0      0      0      0      0      0       0     25      0
0      0      0      0      0     25      0
  8:   47.8m  34.4m   0      0      0      0      0      0       0     29      0
0      0      0      0      0     29      0
  9:   52.1m  37.3m   0      0      0      0      0      0       0     79      0
0      1      0      0      0     79      0
 10:   24.3m  15.3m   0      0      0      0      0      0       0     49      0
0      0      0      0      0     49      0
 11:   29.2m  18.6m   0      0      0      0      0      0       0     19      0
0      1      0      0      0     19      0
 12:    5.9m   5.5m   0      0      0      0      0      0       0     59      7
0      7      0      0      0     23      0
 13:    5.5m   5.5m   0      0      0      0      0      0       0     57      7
0      7      0      0      0     21      0
 14:    4.6m   1.9m   0      0      0      0      0      0      58    280      2
0      2      0      0     86      6      0
 15:    1.6m 748.4k   0      0      0      0      0      0       0    723      4
 0      3      0      0    706      5      0
 16:  120.6m 168.0m   0      0      0      0      0      0       0     16      4
0      4      0      0      0     16      0
 17:   68.0m 111.2m   0      0      0      0      0      0       0     18      4
0      4      0      0      0     18      0
 18:    9.0m   9.0m 108     30     28      0      0      2     3.5g   18     10
0      7      0      0      0      0      0
 19:    9.0m   8.5m   0      0      0      0      0      0    206.4k  33     18
0      7      0      0      0     15      0
 20:  117.7k 107.0k  10.2k   6.4k   6.3k   0      0     93     4.0g  657      1
0      7      0      0     66    454      0
 21:  134.2k 119.6k 191     54     49      0      0      5     2.4g   95      2
0      7      0      0     66      0      0
 22:  102.8k  92.3k  25.8k  10.7k  10.5k   0      0    177     4.2g  302      8
0      8      0      0     66      0      0
 23:  141.7k 127.2k  14     10     10      0      0      0    94.2m  103      1
0      7      0      0     66      0      0

portshow 18
portDisableReason: None
portCFlags: 0x1
portFlags: 0x4203      PRESENT ACTIVE G_PORT U_PORT NOELP LED
LocalSwcFlags: 0x0
portType:  24.0
POD Port: Port is licensed
portState: 1   Online
Protocol: FC
portPhys:  6   In_Sync        portScn:   1   Online
port generation number:    166
state transition count:    27

portId:    011200
portIfId:    43020017
portWwn:   20:12:50:eb:1a:86:68:f7
portWwn of device(s) connected:
```

```
        None
Distance:  normal
portSpeed: N8Gbps

FEC: Inactive
Credit Recovery: Inactive
LE domain: 0
FC Fastwrite: OFF
Interrupts:        0          Link_failure: 10         Frjt:         0
Unknown:           13         Loss_of_sync: 0          Fbsy:         0
Lli:               828147     Loss_of_sig:  7
Proc_rqrd:         319        Protocol_err: 0
Timed_out:         0          Invalid word: 3778430167
Rx_flushed:        0          Invalid_crc:  28
Tx_unavail:        0          Delim_err:    2
Free_buffer:       0          Address_err:  0
Overrun:           0          Lr_in:        17
Suspended:         0          Lr_out:       16
Parity_err:        0          Ols_in:       9
2_parity_err:      0          Ols_out:      15
CMI_bus_err:       0



portshow 20
portDisableReason: None
portCFlags: 0x1
portFlags: 0x10000907  PRESENT ACTIVE E_PORT T_PORT T_MASTER G_PORT LOGICAL_ONLINE
LOGIN
LocalSwcFlags: 0x0
portType:  24.0
POD Port: Port is licensed
portState: 1   Online
Protocol: FC
portPhys:  6   In_Sync        portScn:   16  E_Port    Trunk master port
port generation number:    150
state transition count:    17

portId:    011400
portIfId:    43020011
portWwn:   20:14:50:eb:1a:86:68:f7
portWwn of device(s) connected:
        20:12:50:eb:1a:86:60:38
Distance:  standard <= 10km
portSpeed: 4Gbps

FEC: Inactive
Credit Recovery: Active
LE domain: 0
FC Fastwrite: OFF
Interrupts:        0          Link_failure: 1          Frjt:         0
Unknown:           0          Loss_of_sync: 0          Fbsy:         0
Lli:               32046      Loss_of_sig:  7
Proc_rqrd:         80724      Protocol_err: 0
Timed_out:         0          Invalid word: 4009841749
Rx_flushed:        0          Invalid_crc:  6358
Tx_unavail:        0          Delim_err:    23
Free_buffer:       0          Address_err:  0
Overrun:           0          Lr_in:        6
Suspended:         0          Lr_out:       73
Parity_err:        0          Ols_in:       1
2_parity_err:      0          Ols_out:      4


portshow 21
portDisableReason: None
portCFlags: 0x1
portFlags: 0x903       PRESENT ACTIVE E_PORT T_PORT G_PORT LOGICAL_ONLINE LOGIN
LocalSwcFlags: 0x0
portType:  24.0
POD Port: Port is licensed
portState: 1   Online
Protocol: FC
```

```
portPhys:  6   In_Sync        portScn:   128 T_Port    Trunk port
port generation number:    152
state transition count:    19

portId:    011500
portIfId:    43020011
portWwn:   20:15:50:eb:1a:86:68:f7
portWwn of device(s) connected:
        None
Distance:  standard <= 10km
portSpeed: 4Gbps

FEC: Inactive
Credit Recovery: Active
LE domain: 0
FC Fastwrite: OFF
Interrupts:        0          Link_failure: 2          Frjt:         0
Unknown:           0          Loss_of_sync: 0          Fbsy:         0
Lli:               116975     Loss_of_sig:  7
Proc_rqrd:         45881      Protocol_err: 0
Timed_out:         0          Invalid_word: 2415122012
Rx_flushed:        0          Invalid_crc:  49
Tx_unavail:        0          Delim_err:    5
Free_buffer:       0          Address_err:  0
Overrun:           0          Lr_in:        8
Suspended:         0          Lr_out:       10
Parity_err:        0          Ols_in:       1
2_parity_err:      0          Ols_out:      6
CMI_bus_err:       0


portshow 22
portDisableReason: None
portCFlags: 0x1
portFlags: 0x903       PRESENT ACTIVE E_PORT T_PORT G_PORT LOGICAL_ONLINE LOGIN
LocalSwcFlags: 0x0
portType:  24.0
POD Port: Port is licensed
portState: 1   Online
Protocol: FC
portPhys:  6   In_Sync        portScn:   128 T_Port    Trunk port
port generation number:    154
state transition count:    21

portId:    011600
portIfId:    43020011
portWwn:   20:16:50:eb:1a:86:68:f7
portWwn of device(s) connected:
        None
Distance:  standard <= 10km
portSpeed: 4Gbps

FEC: Inactive
Credit Recovery: Active
LE domain: 0
FC Fastwrite: OFF
Interrupts:        0          Link_failure: 8          Frjt:         0
Unknown:           0          Loss_of_sync: 0          Fbsy:         0
Lli:               414923     Loss_of_sig:  8
Proc_rqrd:         45537      Protocol_err: 0
Timed_out:         0          Invalid_word: 4209057332
Rx_flushed:        0          Invalid_crc:  10581
Tx_unavail:        0          Delim_err:    23
Free_buffer:       0          Address_err:  0
Overrun:           0          Lr_in:        9
Suspended:         0          Lr_out:       32
Parity_err:        0          Ols_in:       3
2_parity_err:      0          Ols_out:      8
CMI_bus_err:       0

portshow 23
portDisableReason: None
portCFlags: 0x1
```

```
portFlags: 0x903       PRESENT ACTIVE E_PORT T_PORT G_PORT LOGICAL_ONLINE LOGIN
LocalSwcFlags: 0x0
portType:  24.0
POD Port: Port is licensed
portState: 1   Online
Protocol: FC
portPhys:  6   In_Sync        portScn:   128 T_Port    Trunk port
port generation number:    156
state transition count:    17

portId:    011700
portIfId:    43020011
portWwn:   20:17:50:eb:1a:86:68:f7
portWwn of device(s) connected:
        None
Distance:  standard <= 10km
portSpeed: 4Gbps

FEC: Inactive
Credit Recovery: Active
LE domain: 0
FC Fastwrite: OFF
Interrupts:         0          Link_failure: 1          Frjt:          0
Unknown:            0          Loss_of_sync: 0          Fbsy:          0
Lli:                130        Loss_of_sig:  7
Proc_rqrd:          45562      Protocol_err: 0
Timed_out:          0          Invalid_word: 94219668
Rx_flushed:         0          Invalid_crc:  10
Tx_unavail:         0          Delim_err:    0
Free_buffer:        0          Address_err:  0
Overrun:            0          Lr_in:        6
Suspended:          0          Lr_out:       8
Parity_err:         0          Ols_in:       1
2_parity_err:       0          Ols_out:      4
CMI_bus_err:        0
```

1.11.2 Check Brocade switch for hosts/initiators/hba's hogging the bandwidth

Brocade switch counter we can use to identify an initiator using up all the bandwidth from the target ports is portstatsshow <port#> and within that output, counter of interest is " **tim_txcrd_z** - Time TX Credit Zero" which is the number of times the Transmit buffer to buffer credit reach zero.

This means that device which is connected to the SAN ports with a high number of times it reached zero, is not able to handle the i/o traffic because the san switch can not send any frames to the connected device because the device has not given back the BB credits back to the switch. This is typically because the connected device may have queue depth set to high if initiator or if target it has exceeded the number requests it can handle.

In the sample below it was a combination of both, hba queue depth set to high and zoning more target ports to load balance. The array had to many clients zoned to a single array target port and the initiators didn't have multipathing enabled, once enabled, customer added more paths to the array lun and set all hba queue depth to 32 which was the recommended setting by Hitachi VSP array.

```
portstatsshow 5
stat_wtx        2813706502  4-byte words transmitted
stat_wrx        2785666420  4-byte words received
stat_ftx        3732101299  Frames transmitted
stat_frx        3768664615  Frames received
```

stat_c2_frx	0	Class 2 frames received
stat_c3_frx	3767112356	Class 3 frames received
stat_lc_rx	0	Link control frames received
stat_mc_rx	0	Multicast frames received
stat_mc_to	0	Multicast timeouts
stat_mc_tx	0	Multicast frames transmitted
tim_rdy_pri	0	Time R_RDY high priority
tim_txcrd_z	**1281848581**	**Time TX Credit Zero (2.5Us ticks)**
tim_txcrd_z_vc 0- 3:	0 0 0	1281848581
tim_txcrd_z_vc 4- 7:	0 0 0	0
tim_txcrd_z_vc 8-11:	0 0 0	0
tim_txcrd_z_vc 12-15:	0 0 0	0
er_enc_in	0	Encoding errors inside of frames
er_crc	0	Frames with CRC errors
er_trunc	0	Frames shorter than minimum
er_toolong	0	Frames longer than maximum
er_bad_eof	0	Frames with bad end-of-frame
er_enc_out	0	Encoding error outside of frames
er_bad_os	1	Invalid ordered set
er_rx_c3_timeout	0	Class 3 receive frames discarded due to timeout
er_tx_c3_timeout	0	Class 3 transmit frames discarded due to timeout
er_unroutable	0	Frames that are unroutable
er_unreachable	3	Frames with unreachable destination
er_other_discard	0	Other discards
er_type1_miss	0	frames with FTB type 1 miss
er_type2_miss	0	frames with FTB type 2 miss
er_type6_miss	0	frames with FTB type 6 miss
er_zone_miss	0	frames with hard zoning miss
er_lun_zone_miss	0	frames with LUN zoning miss
er_crc_good_eof	0	Crc error with good eof
er_inv_arb	0	Invalid ARB
er_single_credit_loss	0	Single vcrdy/frame loss on link
er_multi_credit_loss	0	Multiple vcrdy/frame loss on link

Switch1

Index	Port Address	Media	Speed	State	Proto	**tim_txcrd_z - Time TX Credit Zero** (2.5Us ticks)	
0	0 010000	id	N8	Online	FC F-Port 50:06:0e:80:16:6f:d1:16	827	
1	1 010100	id	N8	Online	FC F-Port 50:06:0e:80:16:6f:d1:1e	87,070,812	dvc2_PORT_1
2	2 010200	id	N8	Online	FC F-Port 21:00:00:24:ff:6d:9f:0c	1,011	
3	3 010300	id	N8	Online	FC F-Port 50:01:43:80:24:29:3e:94	494	
4	4 010400	id	N8	Online	FC F-Port 50:06:0e:80:16:6f:d1:14	1,772	
5	5 010500	id	N8	Online	FC F-Port 50:06:0e:80:16:6f:d1:1c	1,281,848,581	dvc2_PORT_5
6	6 010600	id	N8	Online	FC F-Port 21:00:00:24:ff:6d:a8:80	13	
7	7 010700	id	N8	Online	FC F-Port 50:01:43:80:24:29:47:4c	41	

```
  8  8  010800  id   N8  Online    FC  F-Port
10:00:8c:7c:ff:65:51:59                                  1,118,132
  9  9  010900  id   N8  Online    FC  F-Port
10:00:8c:7c:ff:65:52:65                                      1,413
 10 10  010a00  id   N8  Online    FC  F-Port
50:01:43:80:24:28:93:94                                      2,519
 11 11  010b00  id   N8  Online    FC  F-Port
50:01:43:80:24:29:3e:90                                    729,619
 12 12  010c00  id   N8  Online    FC  F-Port
10:00:8c:7c:ff:65:52:63                                    562,340
 13 13  010d00  id   N8  Online    FC  F-Port
50:01:43:80:24:28:92:2c                                          3
 14 14  010e00  id   N8  Online    FC  F-Port
50:01:43:80:24:28:82:64                                        602
 15 15  010f00  id   N8  Online    FC  F-Port
50:01:43:80:24:29:3e:98                                    436,933
 16 16  011000  id   N8  Online    FC  F-Port
50:01:43:80:24:29:46:04                                    981,873
 17 17  011100  id   N8  Online    FC  F-Port
50:01:43:80:24:29:44:c4                                    390,699
 18 18  011200  id   N8  Online    FC  F-Port
50:01:43:80:28:cb:84:8c                                          0
 19 19  011300  id   N8  Online    FC  F-Port
50:01:43:80:28:cb:87:d4                                          0
 20 20  011400  id   N8  Online    FC  F-Port
50:01:43:80:24:29:46:2c                                    442,628
 21 21  011500  id   N8  Online    FC  F-Port
50:01:43:80:24:29:3e:88                                    350,169
 22 22  011600  id   N8  Online    FC  F-Port
50:01:43:80:28:cb:8b:b4                                          0
 23 23  011700  id   N8  Online    FC  F-Port
50:01:43:80:28:cb:88:3c                                          0
 24 24  011800  id   N8  Online    FC  F-Port
50:01:43:80:24:28:82:44                                          0
 25 25  011900  id   N8  Online    FC  F-Port
50:01:43:80:24:28:93:60                                          0
 26 26  011a00  id   N8  Online    FC  F-Port
50:01:43:80:28:cb:4b:e8                                      4,738
 27 27  011b00  id   N8  Online    FC  F-Port
50:01:43:80:24:29:67:54                                          0
 28 28  011c00  id   N8  Online    FC  F-Port
50:01:43:80:24:28:87:40                                          0
 29 29  011d00  id   N8  Online    FC  F-Port
50:01:43:80:24:29:65:e8                                          0
 30 30  011e00  id   N8  Online    FC  F-Port
10:00:8c:7c:ff:65:52:5d                                         93
 31 31  011f00  id   N8  Online    FC  F-Port
50:01:43:80:28:cc:6a:18                                          0
 32 32  012000  id   N8  Online    FC  F-Port
50:01:43:80:28:cc:7f:60                                        696
```

```
 33 33 012100 id  N8 Online  FC F-Port
50:01:43:80:28:cb:8a:74                                  8
 34 34 012200 id  N8 Online  FC F-Port
50:01:43:80:24:2a:23:80                                  0
 35 35 012300 id  N8 Online  FC F-Port
50:01:43:80:24:28:96:e4                                  0
 36 36 012400 id  N8 Online  FC F-Port
50:01:43:80:28:cb:84:e0                                  0
 37 37 012500 id  N8 Online  FC F-Port
10:00:8c:7c:ff:65:52:53                      1,374,870   P37_655253
 38 38 012600 id  N8 Online  FC F-Port
50:01:43:80:24:2a:24:5c                                  0
 39 39 012700 id  N8 Online  FC F-Port
50:01:43:80:24:28:90:f4                                  0
 40 40 012800 id  N8 Online  FC F-Port
10:00:8c:7c:ff:65:52:51                                  0
 41 41 012900 id  N4 Online  FC F-Port
50:06:01:68:44:60:50:b2                                  0
 42 42 012a00 id  N4 Online  FC F-Port
50:06:01:68:44:60:50:9e                      4,917,308   dvc3_PORT_42
 43 43 012b00 id  N4 Online  FC F-Port
50:06:01:61:44:60:4b:2f                                  0
 44 44 012c00 id  N8 Online  FC F-Port
10:00:8c:7c:ff:65:52:6b                                  0
 45 45 012d00 id  N4 Online  FC F-Port
50:06:01:60:44:60:50:b2                                  0
 46 46 012e00 id  N4 Online  FC F-Port
50:06:01:60:44:60:50:9e                     95,912,220   dvc3_PORT_46
 47 47 012f00 id  N4 Online  FC F-Port
50:06:01:69:44:60:4b:2f                                  0
```

NOTE:

QX and QXS array best practice and recommended queue depth is 16, similarly the same for Netapp arrays.

1.11.3 Important notes and workarounds

Note1:Similar counter in cisco switches is the TBBZ, but would recommend they engage Cisco or whoever supports their switch so they can do the analysis, if we don't have to do the analysis.

• **TBBZ** refers to the total count of Tx B2B credits transitions to zero and is applicable to Generation 4 and Cisco MDS 9700 Series modules. The count is a representation over a period of time (since the module is up or the counter wraps around) of the number of times that a port has no credits to transmit or receive

Here is a link to cisco that will somewhat explain the Time TX credit zero which is the number of times the port hit zero B2B credits (TBBZ or lack_of_transmit_credit).

http://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9700-series-multilayer-directors/white_paper_c11-729444.html#wp9000190

Note2: If the MDC is being used as a Data Mover, such as Proxy Server, we need to make sure the change the queue depth/Execution Throttle from the default 65535 to something smaller to whatever the array supports which is typically queue depth of 16, or 32, 64, 128 or 256.
Defaults for QX/QXS is 16 or Qseries is the same 16.

These parameters will also apply to any other hba vendor since the configuration is array capability driven, but defined as queue depth on hbas such as Emulex and brocades, instead of excecution throttle. We typically don't have to remember these because these defaults are in the hba BIOS by vendor and customers just have to apply the array vendor defaults but most customers don't know about it, so if there is mixed environment san design rules should be followed, if customer is not aware of these rules they should engage the array vendor for guidance.

Option	EMC	HP	IBM	Qlogic	Sun
Connection Options	Loop Pref.	Loop Pref.	Loop Pref.	Loop Pref.	Loop Pref.
Data Rate	Auto	Auto	Auto	Auto	Auto
Frame Size	2048	2048	2048	2048	2048
Enable Hard Loop	No	No	No	No	No
Hard Loop ID	0	0	0	0	0
Loop Reset Delay	5	5	5	5	5
Enable BIOS	No	No	No	No	No
Enable FC Tape	Yes	Yes	Yes	Yes	Yes
Operation Mode	0	0	0	0	0
Interrupt Delay Timer	0	0	0	0	0
Execution Throttle	256	16	256	16	16
Login Retry Count	8	8	8	8	8
Port Down Retry Count	45	16	30	30	30
Link Down Timeout	45	8	30	30	30
LUNs Per Target	256	128	256	128	128
Enable LIP Full Login	Yes	Yes	Yes	Yes	Yes
Enable Target Reset	Yes	Yes	Yes	Yes	Yes
Enable Receive OoOFrame	No	No	No	No	No

This needs to be done to both hba ports and both MDC nodes if changes are nod done and the mdc is used as proxy server, this mdc will starve the fibre ports of bb credits and will hog the bandwidth;

Need root privs with sudo rootsh and change each fibre port at a time for the qlogic the default is 65535 , it will reset the adapter and can see those resets in var/log messages, so recommend this changes be done on Secondary MDC, then failover from primary to secondary then change original primary node hba settings;

Run san surfer command line with scli command, select 2 2: HBA Information > **3: HBA Parameters > 1: Port > 2: Configure HBA Parameters > 11: Execution Throttle change to 16 > 19: Commit Changes**

```
[root@upm440 ~]# scli
Scanning QLogic FC HBA(s) and device(s), please wait...
-
      SANsurfer FC/CNA HBA CLI

      v1.7.3 Build 14

   Main Menu

   1:  General Information
   2:  HBA Information
   3:  HBA Parameters
   4:  Target/LUN List
   5:  iiDMA Settings
   6:  Boot Device
   7:  Utilities
   8:  Beacon
   9:  Diagnostics
  10:  Help
  11:  Exit


      Enter Selection:2


      SANsurfer FC/CNA HBA CLI

      v1.7.3 Build 14

   Main Menu

   1:  General Information
   2:  HBA Information
   3:  HBA Parameters
   4:  Target/LUN List
   5:  iiDMA Settings
   6:  Boot Device
   7:  Utilities
   8:  Beacon
   9:  Diagnostics
  10:  Help
  11:  Exit


      Enter Selection: 3



      SANsurfer FC/CNA HBA CLI

      v1.7.3 Build 14

   HBA Parameters Menu

   HBA Model QLE2562
     1: Port   1: WWPN: 21-00-00-24-FF-43-DC-24 Online
```

2: Port 2: WWPN: 21-00-00-24-FF-43-DC-25 Online
3: All HBAs
4: Return to Previous Menu

Note: 0 to return to Main Menu
Enter Selection: **1**

SANsurfer FC/CNA HBA CLI

v1.7.3 Build 14

HBA Parameters Menu

HBA Instance 0 (QLE2562 Port 1) : Online
WWPN: 21-00-00-24-FF-43-DC-24
Desc: QLE2562 PCI Express to 8Gb FC Dual Channel

1: Display HBA Parameters
2: Configure HBA Parameters
3: Restore Defaults
4: Return to Previous Menu

Note: 0 to return to Main Menu
Enter Selection: 2

SANsurfer FC/CNA HBA CLI

v1.7.3 Build 14

Configure Parameters Menu

HBA Instance 0 (QLE2562 Port 1) : Online
WWPN: 21-00-00-24-FF-43-DC-24
Desc: QLE2562 PCI Express to 8Gb FC Dual Channel

1: Connection Options
2: Data Rate
3: Frame Size
4: Hard Loop ID
5: Loop Reset Delay (seconds)
6: Enable BIOS
7: Enable HBA Hard Loop ID
8: Enable Fibre Channel Tape Support
9: Operation Mode
10: Interrupt Delay Timer (100ms)
11: Execution Throttle
12: Login Retry Count
13: Port Down Retry Count
14: Enable LIP Full Login
15: Link Down Timeout (seconds)
16: Enable Target Reset
17: LUNs per Target
18: Enable Receive Out Of Order Frame
19: Commit Changes
20: Abort Changes

Note: 0 to return to Main Menu
Enter Selection: **11**
Enter Execution Throttle [1-65535] [65535]: **16**

SANsurfer FC/CNA HBA CLI

v1.7.3 Build 14

Configure Parameters Menu

HBA Instance 0 (QLE2562 Port 1) : Online
WWPN: 21-00-00-24-FF-43-DC-24
Desc: QLE2562 PCI Express to 8Gb FC Dual Channel

1: Connection Options
2: Data Rate
3: Frame Size
4: Hard Loop ID
5: Loop Reset Delay (seconds)
6: Enable BIOS
7: Enable HBA Hard Loop ID
8: Enable Fibre Channel Tape Support
9: Operation Mode
10: Interrupt Delay Timer (100ms)
11: Execution Throttle
12: Login Retry Count
13: Port Down Retry Count
14: Enable LIP Full Login
15: Link Down Timeout (seconds)
16: Enable Target Reset
17: LUNs per Target
18: Enable Receive Out Of Order Frame
19: Commit Changes
20: Abort Changes

Note: 0 to return to Main Menu
Enter Selection:**19**
Error: Invalid selection!
SANsurfer FC/CNA HBA CLI

v1.7.3 Build 14

Configure Parameters Menu

HBA Instance 0 (QLE2562 Port 1) : Online
WWPN: 21-00-00-24-FF-43-DC-24
Desc: QLE2562 PCI Express to 8Gb FC Dual Channel

1: Connection Options
2: Data Rate
3: Frame Size
4: Hard Loop ID
5: Loop Reset Delay (seconds)
6: Enable BIOS
7: Enable HBA Hard Loop ID

8: Enable Fibre Channel Tape Support
9: Operation Mode
10: Interrupt Delay Timer (100ms)
11: Execution Throttle
12: Login Retry Count
13: Port Down Retry Count
14: Enable LIP Full Login
15: Link Down Timeout (seconds)
16: Enable Target Reset
17: LUNs per Target
18: Enable Receive Out Of Order Frame
19: Commit Changes
20: Abort Changes

Note: 0 to return to Main Menu
Enter Selection:19

1.12 Cisco Switches

1.12.1 Log analysis (show tech-support) hardware errors and latency troubleshooting - TBBZ counters bb credits transitions to zero

Similar counter in cisco switches is the TBBZ, but would recommend they engage Cisco or whoever supports their switch so they can do the analysis, if we don't have to do the analysis.

• **TBBZ** refers to the total count of Tx B2B credits transitions to zero and is applicable to Generation 4 and Cisco MDS 9700 Series modules. The count is a representation over a period of time (since the module is up or the counter wraps around) of the number of times that a port has no credits to transmit or receive

Here is a link to cisco that will somewhat explain the Time TX credit zero which is the number of times the port hit zero B2B credits (TBBZ or lack_of_transmit_credit).

http://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9700-series-multilayer-directors/white_paper_c11-729444.html#wp9000190

1.12.2 Important notes and workarounds

http://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9700-series-multilayer-directors/white_paper_c11-729444.html#wp9000190

1.13 Qlogic switches

1.13.1 Description

Performance Troubleshooting Information Paper

#show port <port #> will display similar error counters as brocade but named diferent,for crc, it tracks InalidCRC.

Latency troubleshooting counter to track is txwaits which is displayed counter in show port <portnumber>

TxWait Count is the number of times the port entered a wait state because it was out of buffer-to-buffer credits.

Link to Qlogic switch best practice guide, look on page 21 for the cabling diagram I captured below;

http://filedownloads.qlogic.com/Files/TempDownlods/86218/TechnicalGuide_2500-2600_BestPracticesWin2012-ESXi5x_SN0454502-00A.pdf

.

1.14 Empty Word header structure

1.14.1 New Features and Enhancements

provides the following new features and enhancements.

-

1.14.2 Resolved Issues

The following issues have been resolved in

General Issues

- Eliminated erroneous event log entries for
-
-