

Problem Analysis Report

**Customer : National Capital Commission
AMS2100 83042xxx & 83044xxx**

HITACHI DATA SYSTEMS
Global Support Operations

PROBLEM SUMMARY:

Slow read on the LUNs, end users are feeling performance issues
LUN IDs: AMS2100-PROD-L0501-R6, AMS2100-DEV-L0301-R6

Analysis:

AMS2100 83042xxx

- Processor Busy% is GOOD at 52% peak, with 70% max recommended for optimal performance
- CNTL1 Processor Cache Write Pending is HIGH at 20%, with 12% recommended to avoid more aggressive destaging of cache to disk
- HDP1 RAID5 SAS15K RAID Groups Busy% is GOOD to 12% peak, with 50% recommended for optimal performance for SAS
- HDP3 RAID6 SAS7.2K RAID Groups Busy% is GOOD to 18% peak, with 50% recommended for optimal performance for SAS
- HDP4 RAID6 SAS7.2K RAID Groups Busy% is HIGH to 92% peak, with 50% recommended for optimal performance for SAS
- Port IOPS to 1200 IOPS peak
- Port Transfer Rate to 91 MB/s stacked peak
- LUN Response Time to 32 millisecond peak on HDP4 LUN 401, yet HDP4 92% busy at the time with Cache Write Pending to 20%. If LUN 301 is of interest it's response Time is very good at 4.2 milliseconds
- HDP4 LUN 401 IOPS to 730 peak
- HDP4 LUN 401 Transfer Rate to 60 MB/s peak
- LUN Queue Depth to 95 Outstanding I/Os peak, with 32 Outstanding I/Os max recommended rule of thumb

Summary:

Processor Busy%, HDP1 & HDP3 RAID Groups Busy%, and Port Transfer Rates are good, but CNTL1 Processor Cache Write Pending is very high at 20%. This caused by HDP4 RAID6 SAS7.2K RAID Groups Busy% high to 92% peak, with 50% recommended for optimal performance for SAS. Contributing to subsystem issues is LUN Queue Depth to 95 Outstanding I/Os peak, with 32 Outstanding I/Os max recommended rule of thumb. Recommend first to adjust Host Queue Depth lower so that not seen over 32 Outstanding I/Os per LUN and if this doesn't reduce HDP4 RAID6 SAS7.2K RAID Groups Busy% then see Sales for additional spindle count to get back to good health.

For SAS drives:

Host Queue Depth settings need to be changed to only allow a Maximum QD of 32 to each. (SAS)

This means that if 1 x Server has 2 x HBAs with access to the LUN, then the Max QD = 16 per HBA to that LUN.


This means that if 2 x Servers have 2 x HBAs each with access to the LUN, then the Max QD = 8 per HBA to that LUN.

This means that if 4 x Servers have 2 x HBAs each with access to the LUN, then the Max QD = 4 per HBA to that LUN

AMS2100


83042XXX

Processor Utilisation % - 83042405 (AMS) - 2017/01/19 18:42 to 2017/01/19 19:41



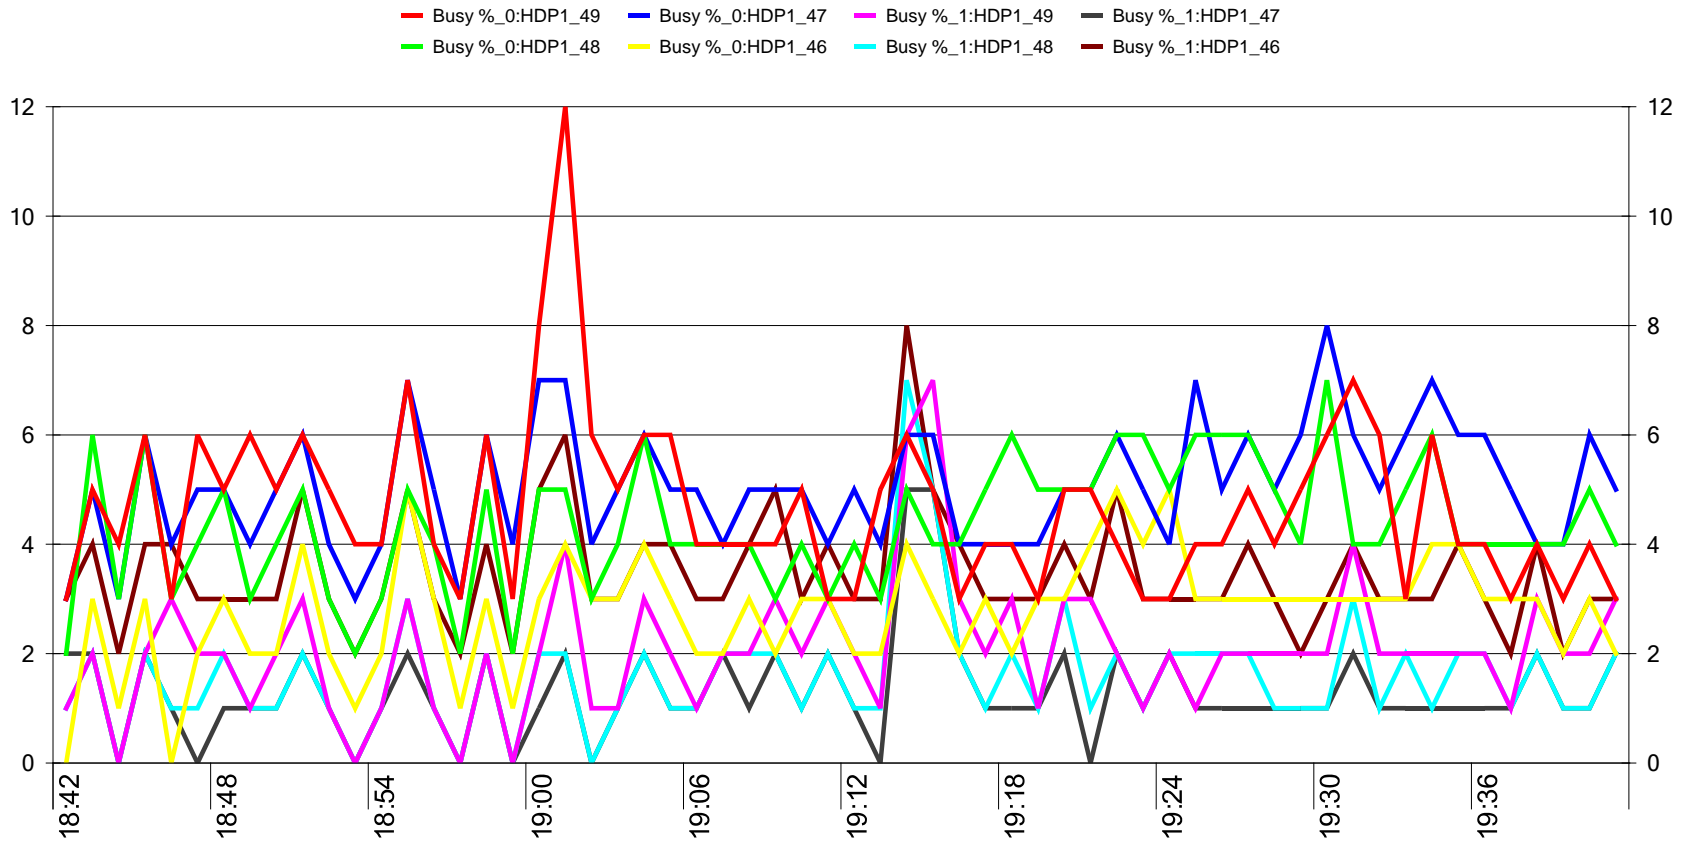
Processor Busy% is GOOD at 52% peak, with 70% max recommended for optimal performance

Processor Write Pending % - 83042405 (AMS) - 2017/01/19 18:42 to 2017/01/19 19:41




CNTL1 Processor Cache Write Pending is HIGH at 20%, with 12% recommended to avoid more aggressive destaging of cache to disk

Physical Drive Busy % - 83042405 (AMS) - 2017/01/19 18:42 to 2017/01/19 19:41




HDP1 RAID5 SAS15K RAID Groups Busy% is GOOD to 12% peak, with 50% recommended for optimal performance for SAS

Physical Drive Busy % - 83042405 (AMS) - 2017/01/19 18:42 to 2017/01/19 19:41



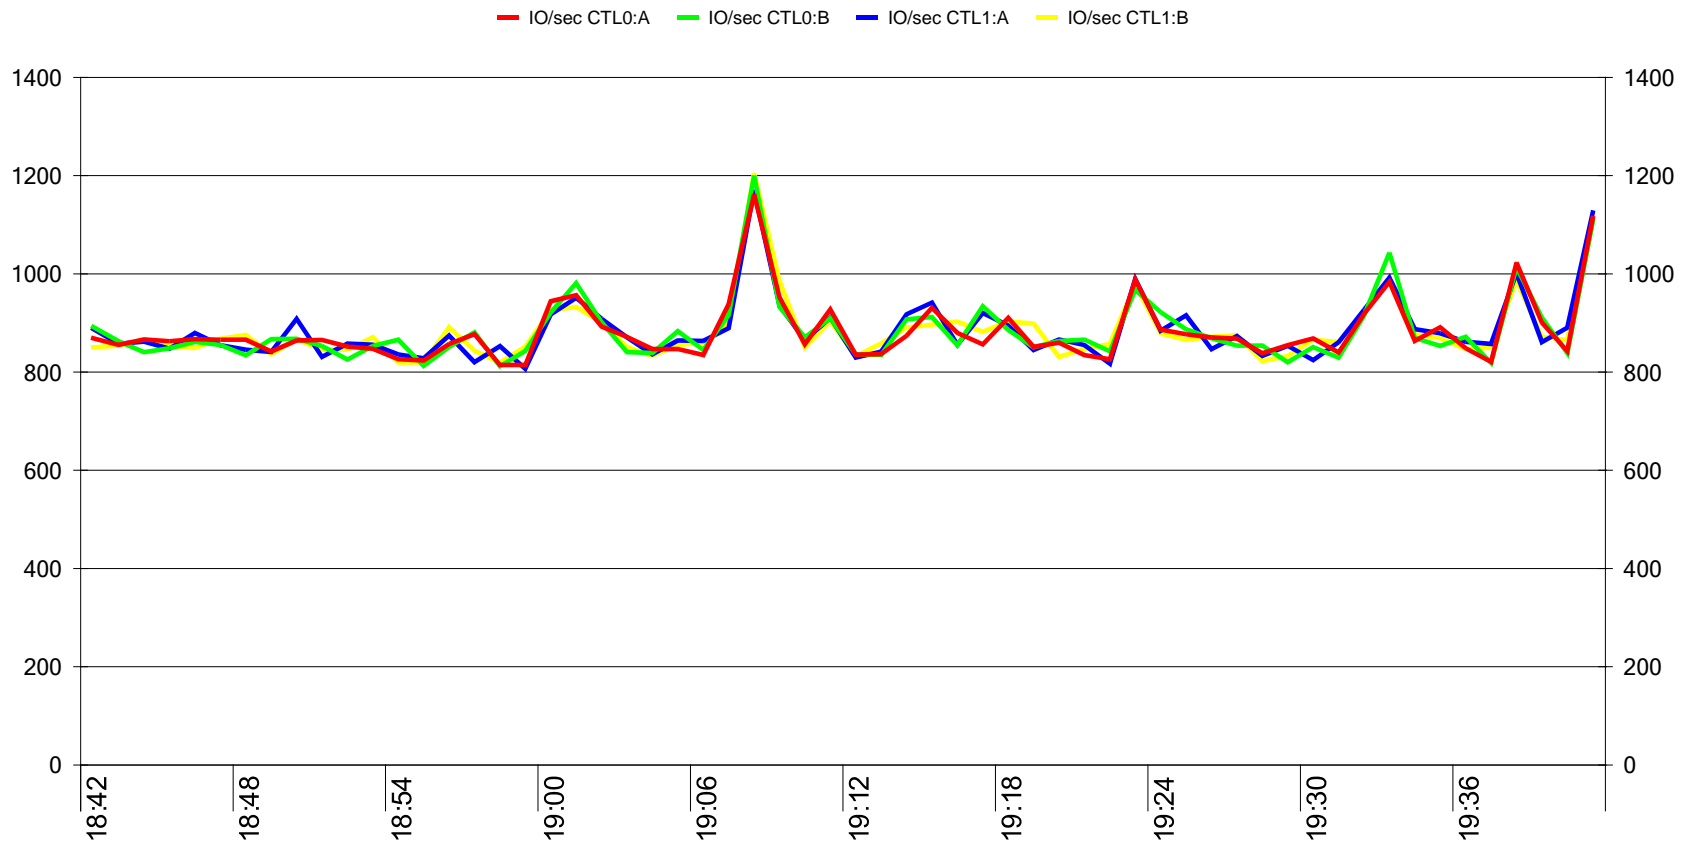
HDP3 RAID6 SAS7.2K RAID Groups Busy% is GOOD to 18% peak, with 50% recommended for optimal performance for SAS

Physical Drive Busy % - 83042405 (AMS) - 2017/01/19 18:42 to 2017/01/19 19:41



HDP4 RAID6 SAS7.2K RAID Groups Busy% is HIGH to 92% peak, with 50% recommended for optimal performance for SAS


Port IOPS - 83042405 (AMS) - 2017/01/19 18:42 to 2017/01/19 19:41



Port IOPS to 1200 IOPS peak

Port Transfer Rate (MB/Second) - 83042405 (AMS) - 2017/01/19 18:42 to 2017/01/19 19:41


R MB/sec CTL0:A R MB/sec CTL0:B R MB/sec CTL1:A R MB/sec CTL1:B W MB/sec CTL0:A W MB/sec CTL0:B W MB/sec CTL1:A W MB/sec CTL1:B



Port Transfer Rate to 91 MB/s stacked peak


LU Response (ms) - 83042405 (AMS) - 2017/01/19 18:42 to 2017/01/19 19:41

— Resp_HDP1_LU_0:107-0A-ESX Cluster DEV[*]
 — Resp_HDP1_LU_1:109-1A-ESX Cluster DEV[*]
 — Resp_HDP3_LU_1:302-1A-ESX Cluster DEV[*]
— Resp_HDP4_LU_0:402-0A-ESX Cluster DEV[*]
 — Resp_HDP3_LU_1:301-1A-ESX Cluster DEV[*]
 — Resp_HDP4_LU_1:401-1A-ESX Cluster DEV[*]




LUN Response Time to 32 millisecond peak on HDP4 LUN 401, yet HDP4 92% busy at the time with Cache Write Pending to 20%. If LUN 301 is of interest it's response Time is very good at 4.2 milliseconds

LU 1:401-1A-ESX Cluster DEV[*] IOPS - 83042405 (AMS) - 2017/01/19 18:42 to 2017/01/19 19:41



HDP4 LUN 401 IOPS to 730 peak


LU 1:401-1A-ESX Cluster DEV[*] Transfer Rate - 83042405 (AMS) - 2017/01/19 18:42 to 2017/01/19 19:41



HDP4 LUN 401 Transfer Rate to 60 MB/s peak

LU Maximum Queue Depth - 83042405 (AMS) - 2017/01/19 18:42 to 2017/01/19 19:41

- Max Tag_HDP1_LU_0:107-0A-ESX Cluster DEV[*]
- Max Tag_HDP4_LU_0:402-0A-ESX Cluster DEV[*]
- Max Tag_HDP1_LU_1:109-1A-ESX Cluster DEV[*]
- Max Tag_HDP3_LU_1:301-1A-ESX Cluster DEV[*]
- Max Tag_HDP3_LU_1:302-1A-ESX Cluster DEV[*]
- Max Tag_HDP4_LU_1:401-1A-ESX Cluster DEV[*]



LUN Queue Depth to 95 Outstanding I/Os peak, with 32 Outstanding I/Os max recommended rule of thumb