The IT Performance Management & Capacity Planning Company

Performance Audit Results

IBIS Server Future Project

System "umep1004"

Author: Walter Verhoeven

Doc. Ref.: capacity_audit_example

Version: 1.0

Date: 2 February 2018

I. Document History

Version	Date	Name	Description
1.0	2 Feb/18	W. Verhoeven	Hand over to Customer

II. Stakeholders

Company	Responsible Stakeholder
Customer	Mr. X
Customer	Mr. Y
Customer	Mr. Z
CREATIVE Associates	Walter Verhoeven

CREATIVE Associates

Dondervelstraat 52
1651 Lot/Beersel
Belgium

III. Contents

Table of Contents

1. MANAGEMENT SUMMARY....5

1.1 Introduction....5

2. PRODUCTION SERVER DESCRIPTION....7

3. RECOMMENDATIONS & REMARKS....8

3.1 CPU....8

3.2 I/O Subsystem....8

3.3 Memory....8

3.4 Network....9

3.5 Modeling Results....9

4. CDR ACTIVITY....10

4.1 Full Period....10

5. CPU LOAD....12

5.1 Overall Activity....12

6. MEMORY USAGE....23

7. I/O SUBSYSTEM ACTIVITY....25

7.1 Volume Group “vg00”....25

7.2 Volume Group “vgibip_ibip”....26

7.3 Volume Group “vgibip_data”....29

7.4 Volume Group “vgibip_dbs”....31

7.5 Volume Group “vgibip_index”....32

7.6 Volume Group “vgibip_arch”....34

7.7 Volume Group “vgibip_sysopt”....35

8. NETWORK ACTIVITY.....37

9. CAPACITY PLANNING.....38

9.1 Increase Load With 5% More CDR's By Year.....38

1. Management Summary

1.1 Introduction

Currently at Customer the "umep1004" Superdome HP-UX server, running the production IBIS environment called IBIP, processes between 102,2 – 127,8 Mio Call Data Records per day during the working week. A remark from Mr. X about these figures is *"for VOICE in Rating we are doing a SKIP rating for call without duration = Rating is so doing almost nothing for half of the voice cdrs) and for LTE, MMS, SMS, Sip, STP this a a full rating"*.

The "raw" CDR data, for the period Thursday 28th Dec/17 – Tuesday 30th Jan/18, has been loaded into the TeamQuest environment and has the following lay-out:

```
2018/01/17;00;202454;1146312;1326;608424;294528;946683;;3199727
2018/01/17;01;108283;731082;1297;465018;217800;702944;;2226424
2018/01/17;02;101107;713446;1036;472389;211646;590819;;2090443
2018/01/17;03;68480;789623;1110;530102;232378;727879;;2349572
2018/01/17;04;127284;1372201;1185;523070;269325;686474;;2979539
2018/01/17;05;97969;1282738;855;621640;351999;652076;;3007277
2018/01/17;06;98502;723557;1052;735657;496529;817500;;2872797
2018/01/17;07;105278;746181;1622;996809;737774;821148;;3408812
2018/01/17;08;88795;766957;2180;1317437;1140641;1203613;;4519623
2018/01/17;09;128482;1092356;2555;1970459;1666931;1295780;;6156563
2018/01/17;10;143161;823726;3158;2413944;2107749;1415795;;6907533
2018/01/17;11;150974;838410;3975;2569275;2373583;1530304;3;7466524
```

Figure 1-1: Rating Raw CDR Data Lay-Out

The definition of the records is shown below:

b. Format of the file provided
RNDATE : is date
RNTM : is Hr
DRA : CDR DRA/LTE
_STP : CDR signalling
MMS00 = CDR MMS
NGNUXPROD = CDR Veraz/dialogic
NGNUSON = CDR Sonus
SMS30 = CDR SMS
TITAN : CDR TITAN/Sip

RNDATE	RNTM	_DRA	_STP	MMS00	NGNUXPRO	NGNUXSON	SMS30	TITAN	Grand Total
2018/01/16	11	148674	839159	3248	2508885	2255749	1349080	8	7104803
2018/01/16	12	76701	793295	3581	2444468	2294364	1576788		7189197

Figure 1-2: Rating CDR Data Clarifications

Remark that there are two types of activity, a **CDR Rating** related one that spans the month and a **CDR Bill Run** activity that occurs at specific moments, mostly by the end of the month & begin of a new month.

The Oracle database and batch processes are currently running on the same server. A project named "IBIS Server Future" has been launched in order to evaluate what configuration/platform would be the best fit to handle current and future increased business loads.

As such I have been asked at this stage of the project to provide by the 2th of February the following information:

1. Performance assessment of the current situation. Are there any issues ?

2. Capacity assessment during a simulation of a yearly business/CDR load increase of 5%. Simulate the next 5 years, so 2018 → 2022, or a total load increase of 25%.

The server performance data collection is done by TeamQuest Manager (CMIS), normally we have data at a 1 minute granularity for 2 weeks. But the period we had to analyse was earlier so we work with the 10 minute granularity data.

The capacity part is handled by TeamQuest Predictor and here we took a relevant snapshot at Monday 8th Jan/18 15h as this is a moment where the most CDR's (9 Mio) where processed by hour.

Data analysis is done via TeamQuest Vityl Monitor.

The results of this audit have been documented in version 1 of the report and passed to Customer on the 2th Jan/18.

2. Production Server Description

The hardware configuration (partial) of this server is shown below:

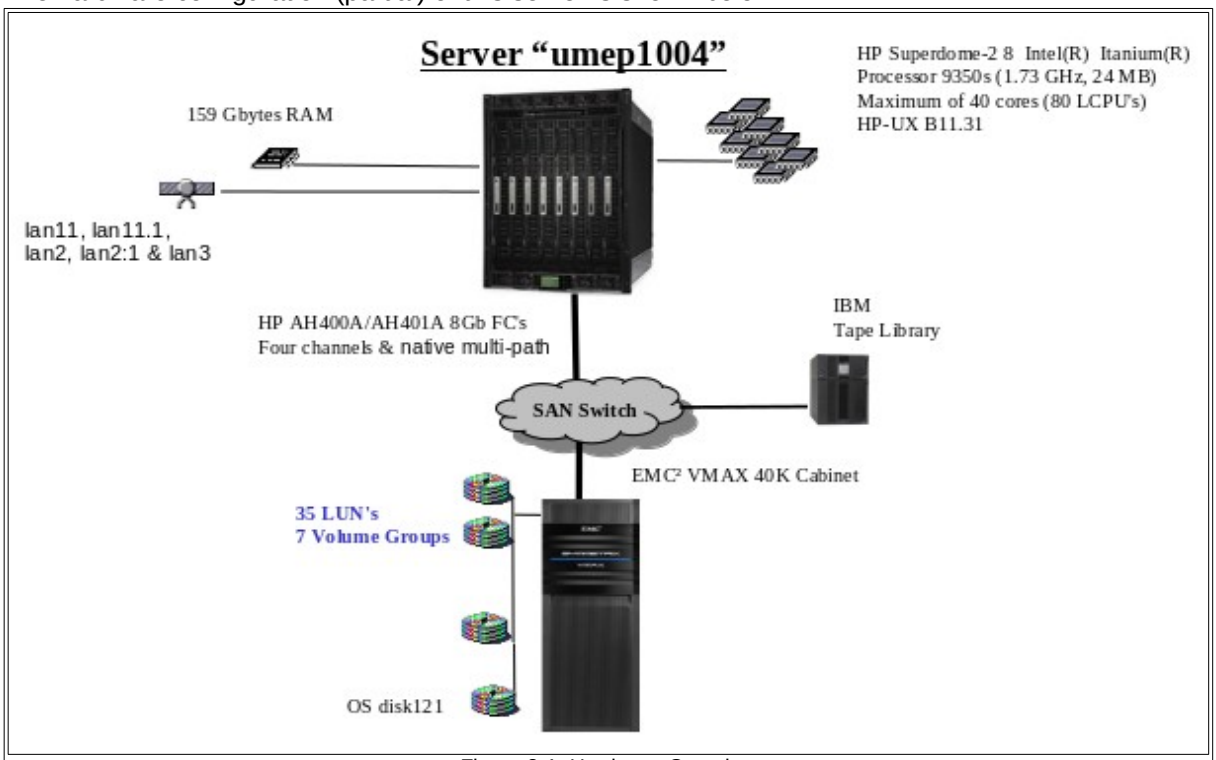

Figure 2-1: Hardware Overview

Globally we have the following hardware resources:

CPU: Server of the type HP Superdome-2 Intel Itanium processors 9350 of 1.73 GHz (24 MB caching, 4 cores, 8 logical processors per socket). Active processor count, during our measurement period 29th Dec/17 → 29th Jan/18 is 8 sockets, 40 cores (2 to 4 per socket) so total of 80 potential logical processors but we saw a maximum of 72 LCPU's active at any time. This is done via the "ticap" feature.

Memory: Total of 159,6 Gbytes installed. The SWAP size is a combination of device & pseudo SWAP. We measured a size of 175,7 Gbytes.

LAN: Several LAN cards where found but only on three we found activity "lan11, lan11.1, lan2, lan2:1 & lan3". All of the type 1 Gbit.

DISK: We found 35 active disks (LUN's) all connected on a SAN EMC² VMAX 40K solution with four cards type "AH400A & AH401A 8Gb PCIe Fibre Channel Adapter" no PowerPath, as hp-ux 11.31 is using its native multi-pathing, to use the channels for high availability & performance reasons. Only ssd & fc technology used. Allocated by default "meta stripped volume" (8 members) and configuration RAID-5 (3+1)

There are 4 connections to each disk, the OS disk is also on SAN configuration.

3. Recommendations & Remarks

3.1 CPU

The CPU load is maximally around 72% and on average about 60%. During the weekend the load is 10% smaller. So sufficient CPU resources available with a good potential to process more CDR's.

Generally there is a daily LCPU's change that occurs at the following moments, we took Tuesday the 2th Jan/18 as an example: (10h 60→64), (11h 64→68), (14h 68→70), (18h 70 → 72), (18h40 72 → 70), (21h 70 → 66), (22h 66 → 64) and (23h 64 → 60) again.

The amount of active LCPU's is linked with the number of "behldnc" processes so for the same day we have for example: (10h 25→28), (11h 28→35), (14h 35→38), (18h 38 → 34), (18h10 34 → 36), (21h 36 → 34) and (22h 34 → 23). Also every day around 23h55 these processes are all restarted.

Largest consumers are the "IBIS_behldnc" processes, using about 10%, and "IBIS_ibip", using about 25% CPU. These workloads represent the Rating & Bill Run process application & DB parts. The "OTHER", occupying about 20%, is due to a lot of small/script commands also called short lived processes.

The remaining processes only consume few CPU resources.

Most of the work is done on weekdays during the "prime shift" between 8h – 18h as confirmed by the CDR statistics we received from Customer. Between 115 – 127 Mio CDR's where processed daily during our measurement period.

When we look at the type of CDR's processed we have the following, from highest to lowest activity "NGNUXPRO", "NGNUXSON", "SMS30", "CUSTOMER_STP", "CUSTOMER_DRA", "MMS00" and "TITAN".

Seen the potential to process more CDR's there are no processes queuing up due to stress on this resource.

3.2 I/O Subsystem

We have about 35 Disks (LUN's) and 7 Volume Groups. There are 4 Fibre Channels connected to SAN switch(es) at 8 Gbits/Sec.

Overall the I/O subsystem is in a very good state but we found quite high activity, 90% to 100% utilization, on the disks, "disk105", "disk108", "disk109", "disk110", "disk118" and "disk334" part of the "vgibip_index" volume group.

This is quite high and need attention since these there is no more room for handling more I/O. Currently the queuing time and service times are still fine. But during the modeling exercise we found that "disk110" is at the limit of it's capacity.

3.3 Memory

The memory resource does not encounter any stress during the observed period since the available physical memory (not allocated) is about 53 Gbytes.

Seen the size of the unallocated memory there are no worries. As such the paging activity, due to memory pressure, on scan & page-out level is almost non-existing.

The swap space has sufficient free space, about 76 Gbytes, to handle more applications/processes allocating memory. For each data page some swap space will be reserved.

We do not physically use the SWAP space as there is no paging-out activity

3.4 Network

Not relevant in this case.

3.5 Modeling Results

Mr. X mentioned that an interesting period to evaluate further is Sunday 31th Dec/17 – Wednesday 10th Jan/18 as then both CDR activities being Rating & Bill Run are ongoing.

Later there is only Rating activity. Also during that period the VOICE rating does a SKIP rating without duration. This means that almost nothing happens for half of these CDR's. For LTE, MMS, SMS, Sip and STP this is a full rating.

The busiest day for the period Sunday 31th Dec/17 – Wednesday 10th Jan/18 was Monday the 8th Jan/ 18 where about 127,5 Mio CDR's where processed.

On the 8th we processed 127.573.254 CDR's per day so a 5% increase represents 6.378.663 CDR's. That's a total of 133.951.917 CDR's for 2018. As agreed by Mr. X we will go for a compound increase for the next years.

We only increase the business workloads with 5%, the supporting tools load with 1%.

current	5.00%	2018	2019	2020	2021	2022
CDR Increase	**Compound (Add 5% each year of the activity of the past year)**					
127.573.254	6.378.663	133.951.917	140.649.513	147.681.988	155.066.088	162.819.392
		6.697.596	7.032.476	7.384.099	7.753.304	

Figure 3-1: CDR Compound Increase

The global CPU load by application (workloads) shows also the highest CPU consumption at that moment which is 54%. There is a difference with 1 hour, compared to the CDR time due to the way the data is logged with a user agent (end of time/begin of time) versus CMIS agent.

So we extracted the 1h data of the 8th Jan/18 of 16h to do the calculations, the number of LCPU's active at that time was 70.

The model revealed that, apart from the "disk110" that needed to be offloaded, the server will handle the load increase up to 2021 with 70 LCPU's. For 2022 we need 2 more cores resulting in 74 LCPU's.

4. CDR Activity

As mentioned in the introduction we received a raw CDR data file, 1 record/hour, that spans the period Thursday 28th Dec/17 – Tuesday 30th Jan/18.

4.1 Full Period

The following amount of CDR's have been processed during this period, clearly less activity during the weekends and up to about 9 Mio CDR's processed by hour on the busiest hour being Monday the 8th January/18.

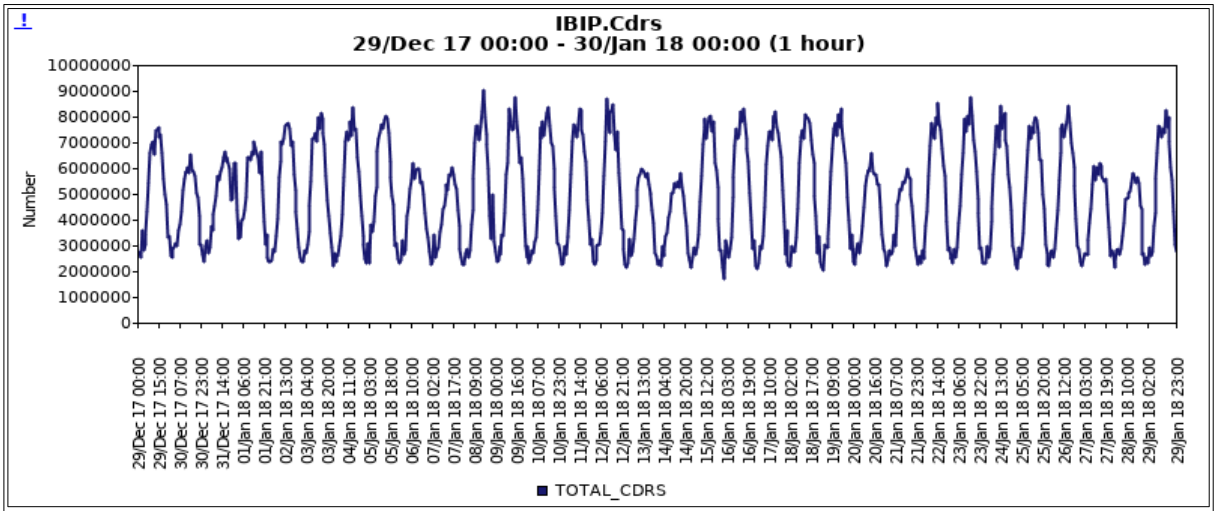

Figure 4-1: Total Number Of Processed CDR's By Hour

When we look at the type of CDR's processed we have the following, from highest to lowest activity "NGNUXPRO", "NGNUXSON", "SMS30", "CUSTOMER_STP", "CUSTOMER_DRA", "MMS00" and "TITAN".

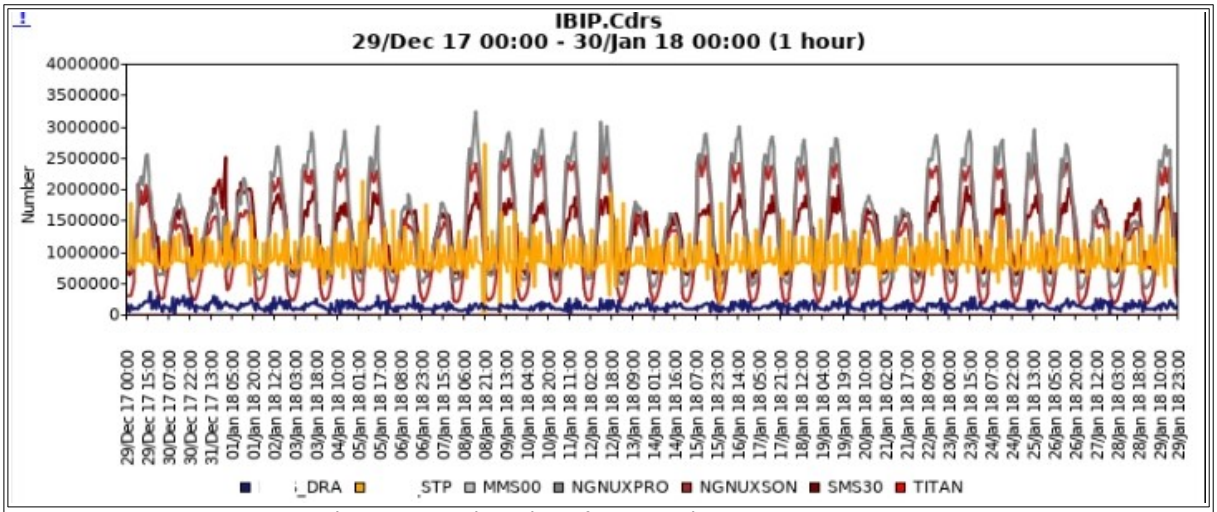

Figure 4-2: Total Number Of Processed CDR Types By Hour

Mr. X mentioned that an interesting period to evaluate further is Sunday 31 Dec/17 – Wednesday 10 Jan/18 as then both CDR activities being Rating & Bill Run are ongoing. Later there is only Rating activity.

As expected is the load the highest during the 8h – 18h business hours period.

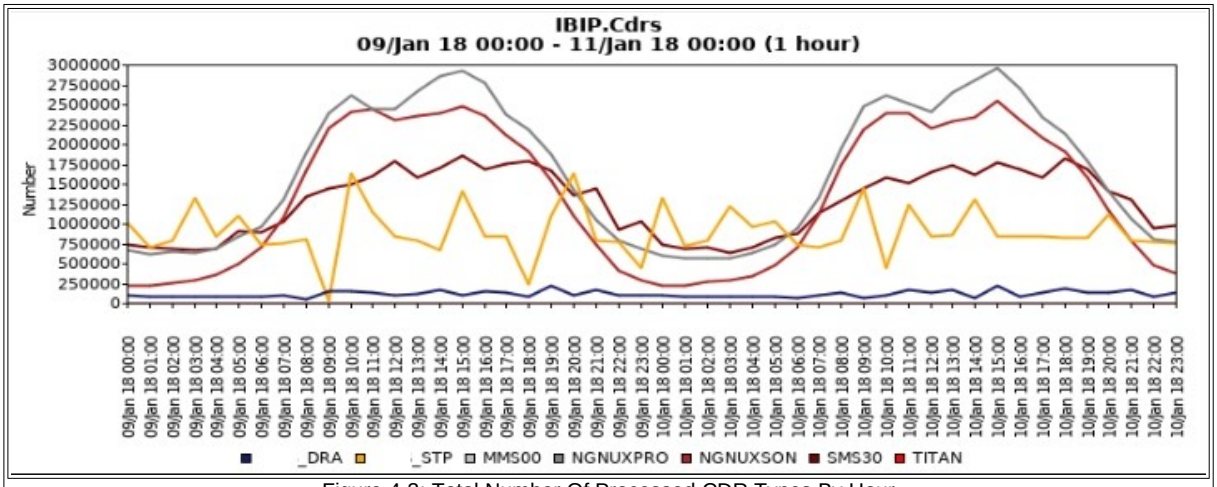

Figure 4-3: Total Number Of Processed CDR Types By Hour

The same activity but now on a 24h bases for the full measured period looks as following.

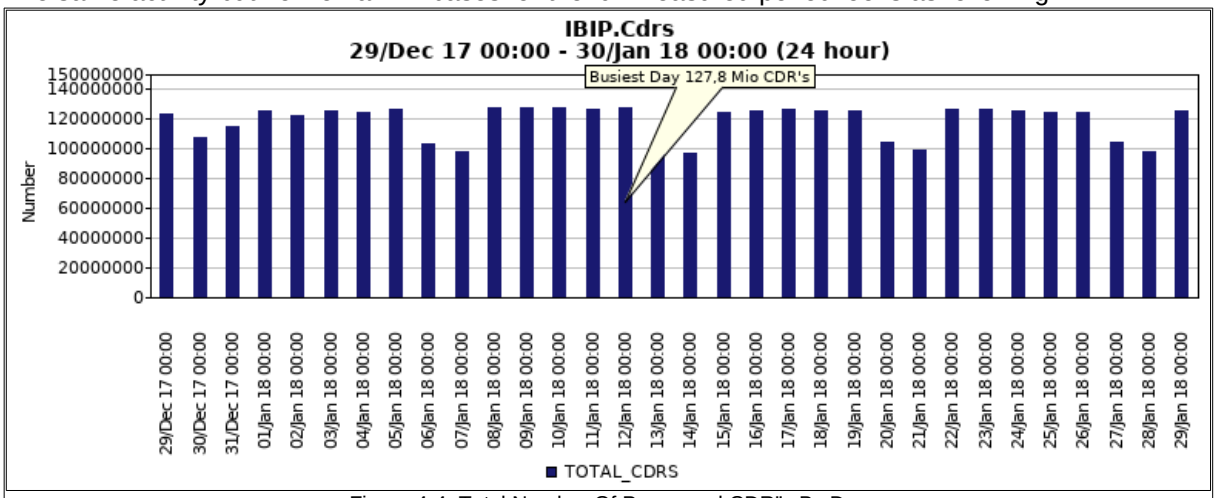

Figure 4-4: Total Number Of Processed CDR"s By Day

The different weeks show a similar activity with about 120 to 127 Mio CDR's processed per day. Busiest day was Friday 12th Jan/18 where 127,8 Mio CDR's where handled.

The busiest day for the period Sunday 31th Dec/17 – Wednesday 10th Jan/18 was Monday the 8th Jan/ 18 where about 127,5 Mio CDR's where processed.

Time	TOTAL_CDRS
31/Dec 17 00:00	115142256.000
01/Jan 18 00:00	125989903.000
02/Jan 18 00:00	122599353.000
03/Jan 18 00:00	125452553.000
04/Jan 18 00:00	124384353.000
05/Jan 18 00:00	127184179.000
06/Jan 18 00:00	103884102.000
07/Jan 18 00:00	98325610.000
08/Jan 18 00:00	127573254.000
09/Jan 18 00:00	127409725.000
10/Jan 18 00:00	127391533.000

Table 4.1: CDR's By Day

5. CPU Load

5.1 Overall Activity

We first look at the overall CPU load and the number of active logical CPU's for the period where we have a combined Rating & Bull Run activity. The number of active LCPU's change during the day.

The CPU load is maximally around 72% and on average about 60%. During the weekend the load is about 10% smaller. So sufficient CPU resources available with a good potential to process more CDR's.

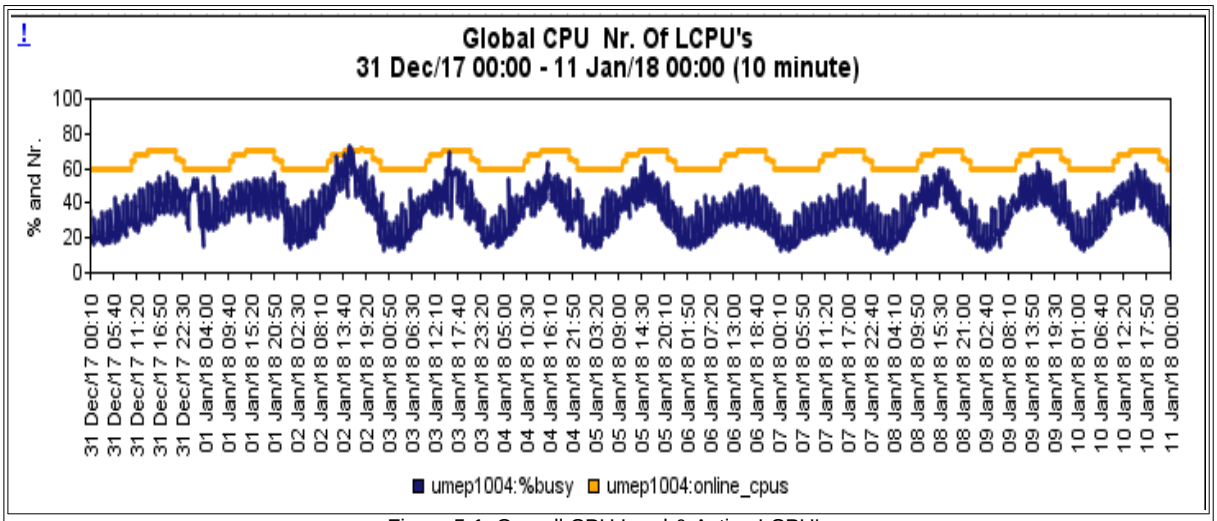

Figure 5-1: Overall CPU Load & Active LCPU's

The daily LCPU's change occurs at the following moments, we took Tuesday the 2th Jan/18 as an example: (10h 60→64), (11h 64→68), (14h 68→70), (18h 70 → 72), (18h40 72 → 70), (21h 70 → 66), (22h 66 → 64) and (23h 64 → 60) again.

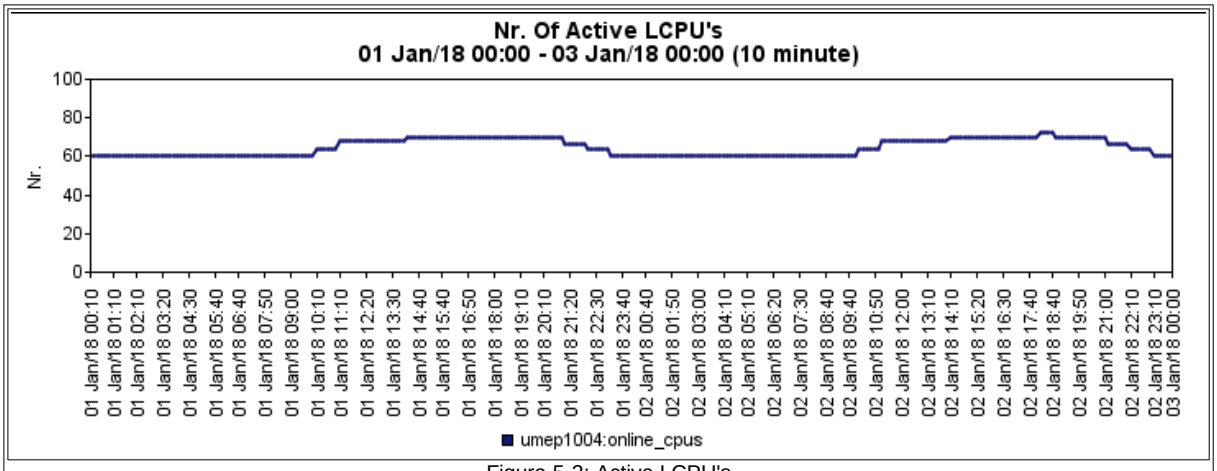

Figure 5-2: Active LCPU's

As the "behldnc" processes are important for processing CDR's and variable in number, we look at the amount of these processes during these two days.

Again we look we took Tuesday the 2th Jan/18 as an example: (10h 25→28), (11h 28→35), (14h 35→38), (18h 38 → 34), (18h10 34 → 36), (21h 36 → 34) and (22h 34 → 23). So not totally in sync with the LCPU's changes !

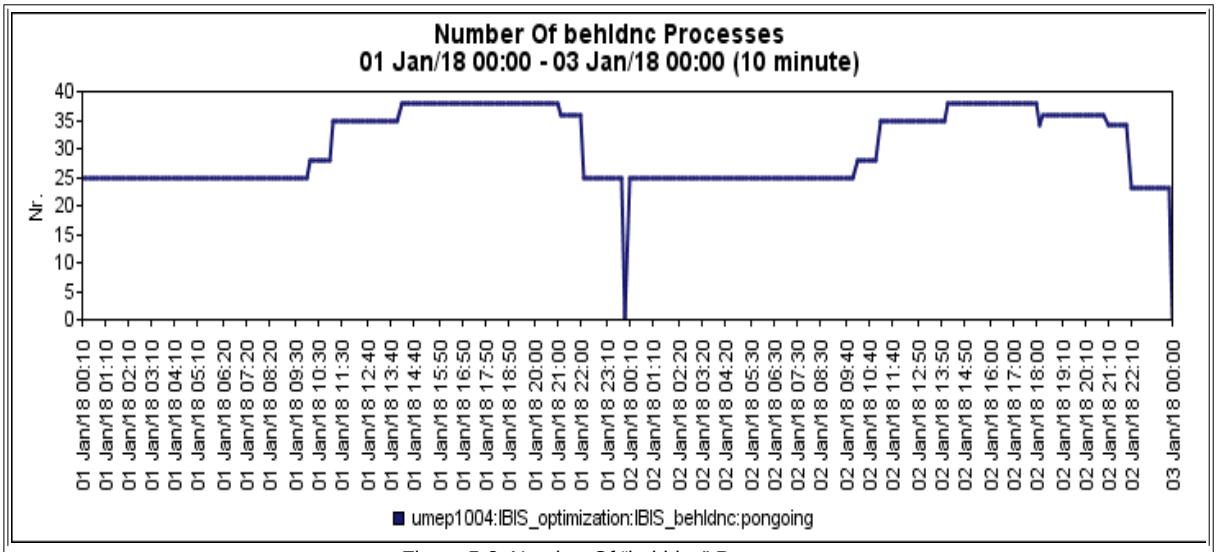

Figure 5-3: Number Of "behldnc" Processes

Also every day around 23h55 these processes are all restarted.

When looking at the CPU by workload (application groups) we have the following. The 100% peaks are measured in the "OTHER" workload but not reflected in the Global CPU consumption and occurs always on 11h. So it is a "workload" measurement issue and not a real consumption, a drill-down does not reveal a specific process neither.

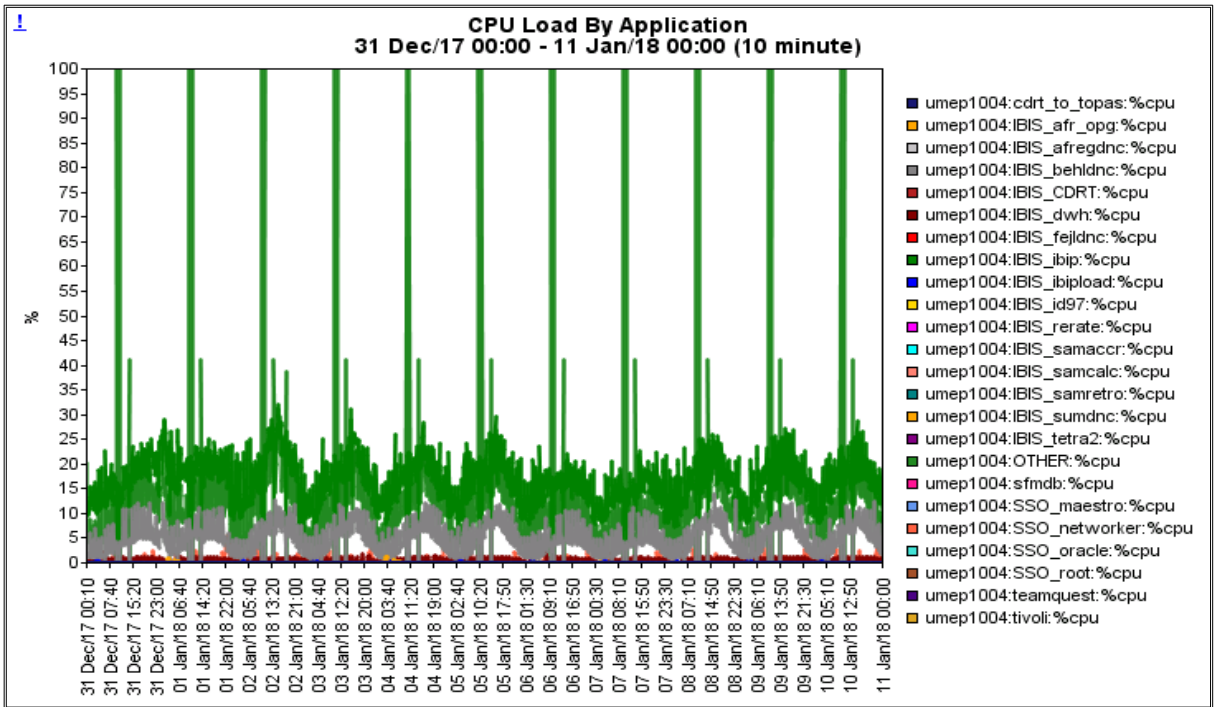

Figure 5-4: CPU Load By Application

Largest consumers are the “IBIS_behldnc”, occupying about 10%, and “IBIS_ibip”, occupying about 25%, workloads representing the Rating & Bill Run process application & DB parts. Filter looks as following (fullcmd = /.*ibip.*/). We have mainly all IBIP related Oracle processes in there.

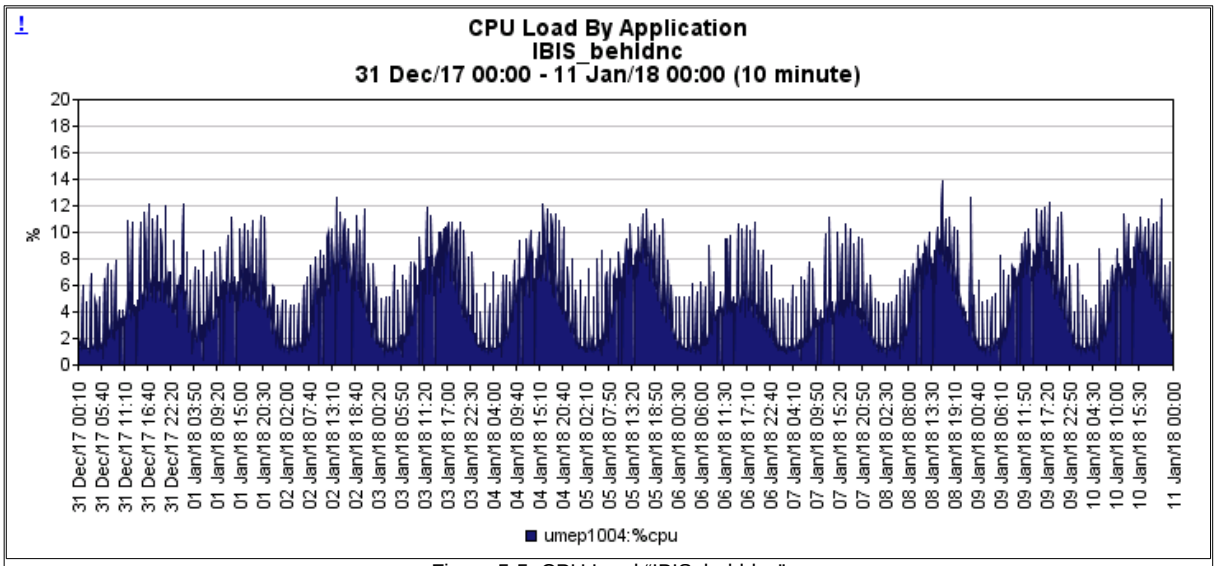

Figure 5-5: CPU Load “IBIS_behldnc”

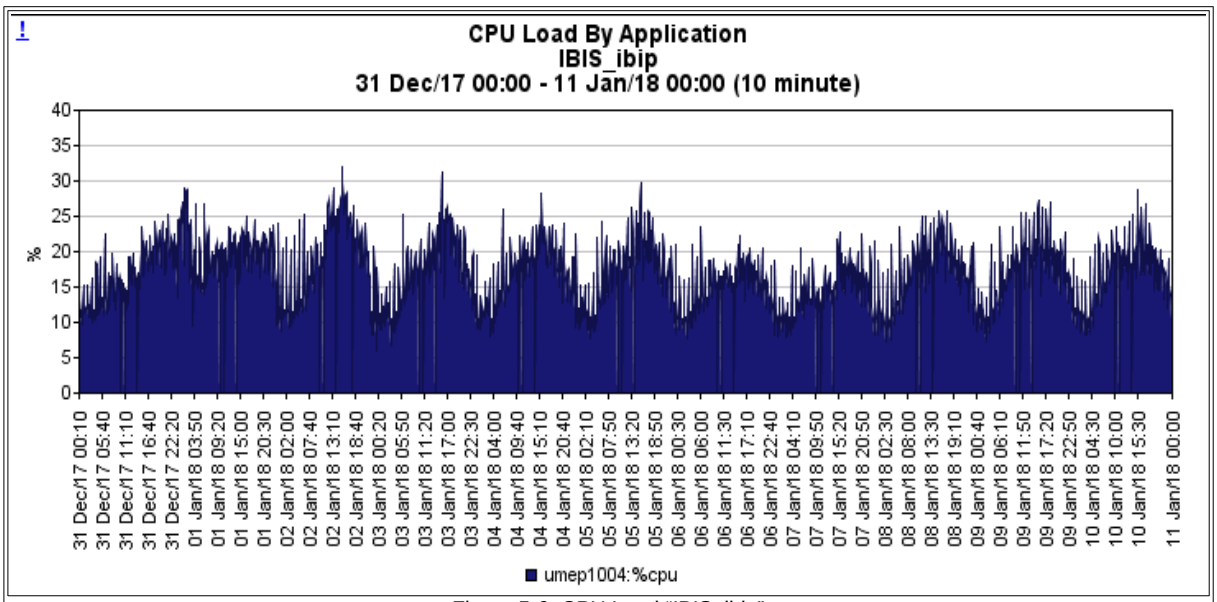

Figure 5-6: CPU Load “IBIS_ibip”

The “OTHER”, occupying about 20%, is due to a lot of small/script commands also called short lived processes.

We already tried, for earlier audits, to decrease this “OTHER” activity by activating the accounting for “tqwarp” but that did not gave us the expected result. So we work further with the data we have.

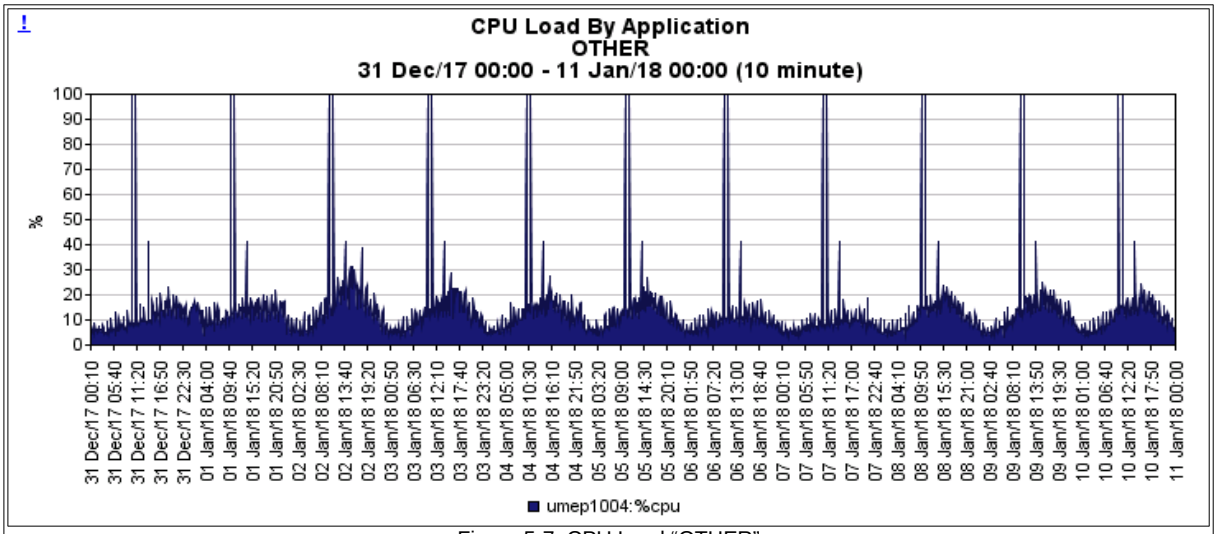

Figure 5-7: CPU Load “OTHER”

The remaining workloads, representing the application site of the rating/billing process, are very small in activity, see the following set of charts below:

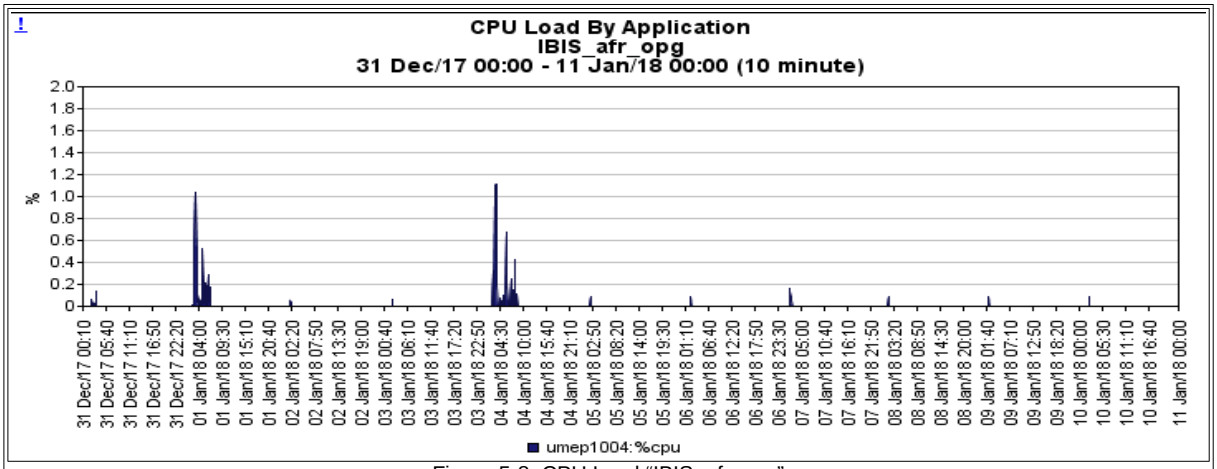

Figure 5-8: CPU Load “IBIS_afr_opg”

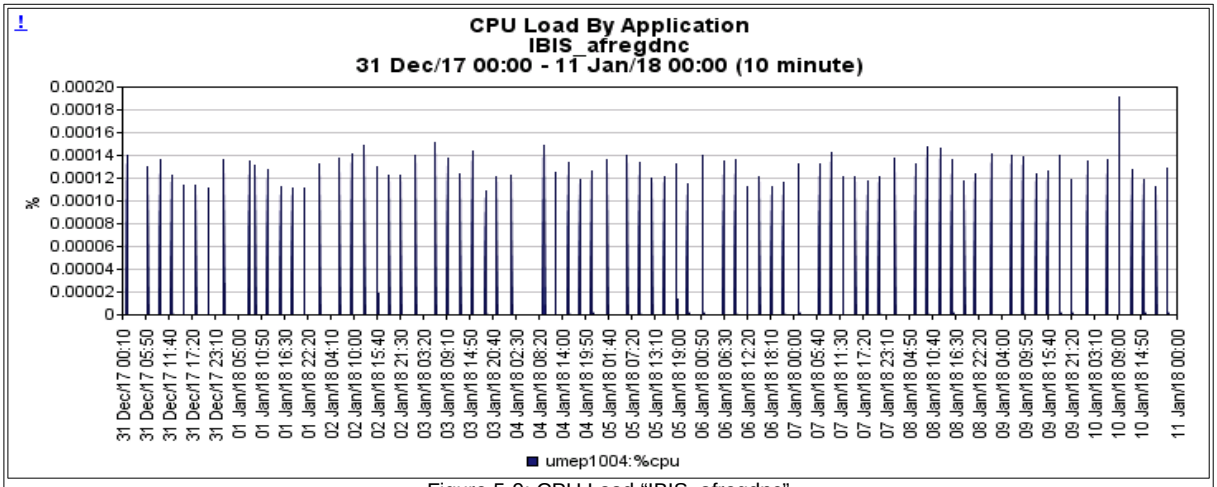

Figure 5-9: CPU Load “IBIS_afregdnc”

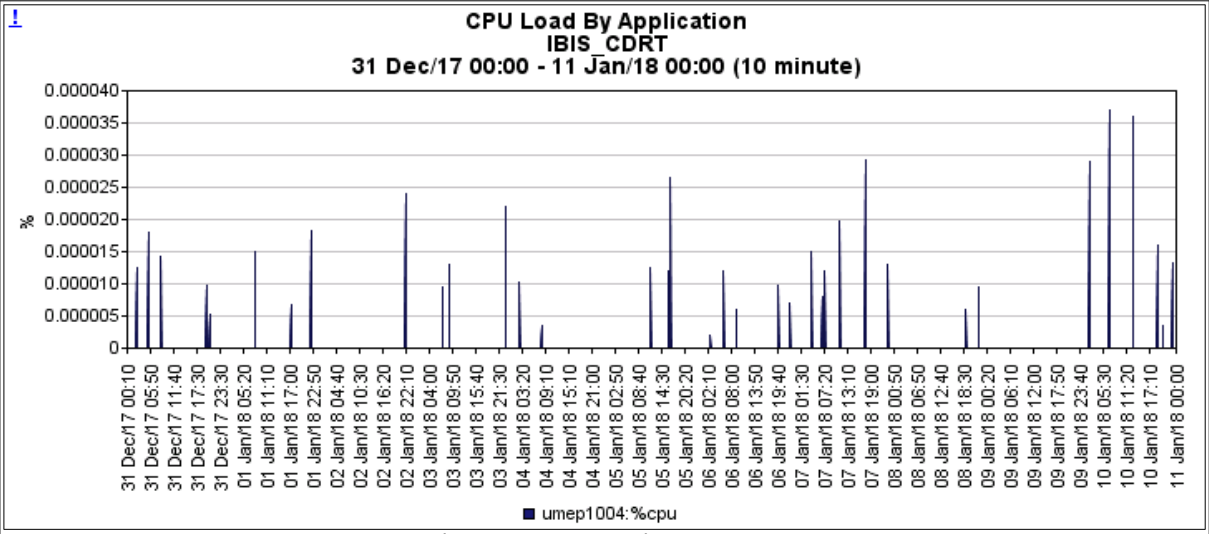

Figure 5-10: CPU Load "IBIS_CDRT"

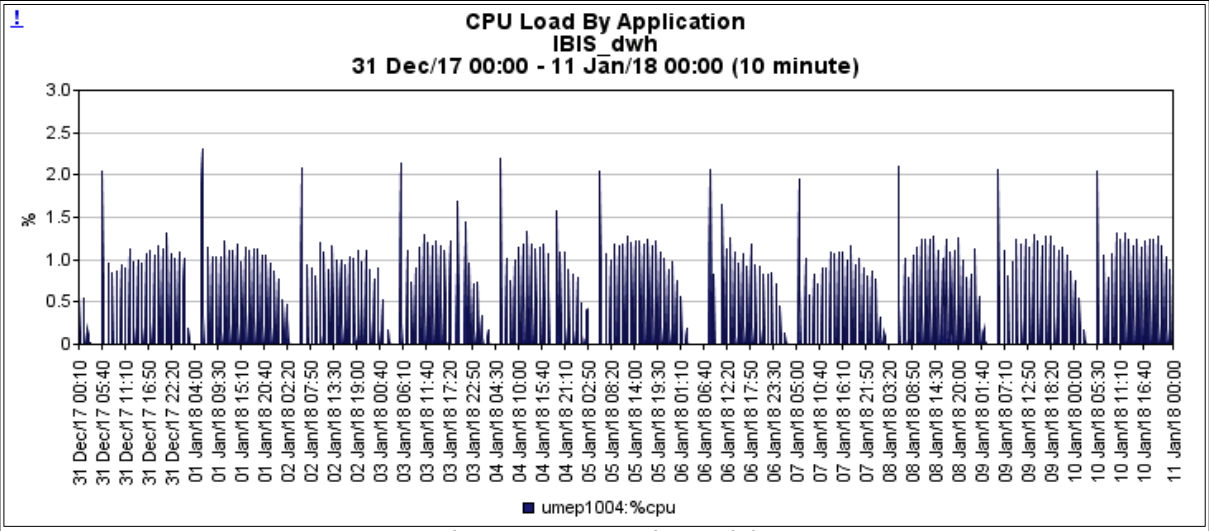

Figure 5-11: CPU Load "IBIS_dwh"

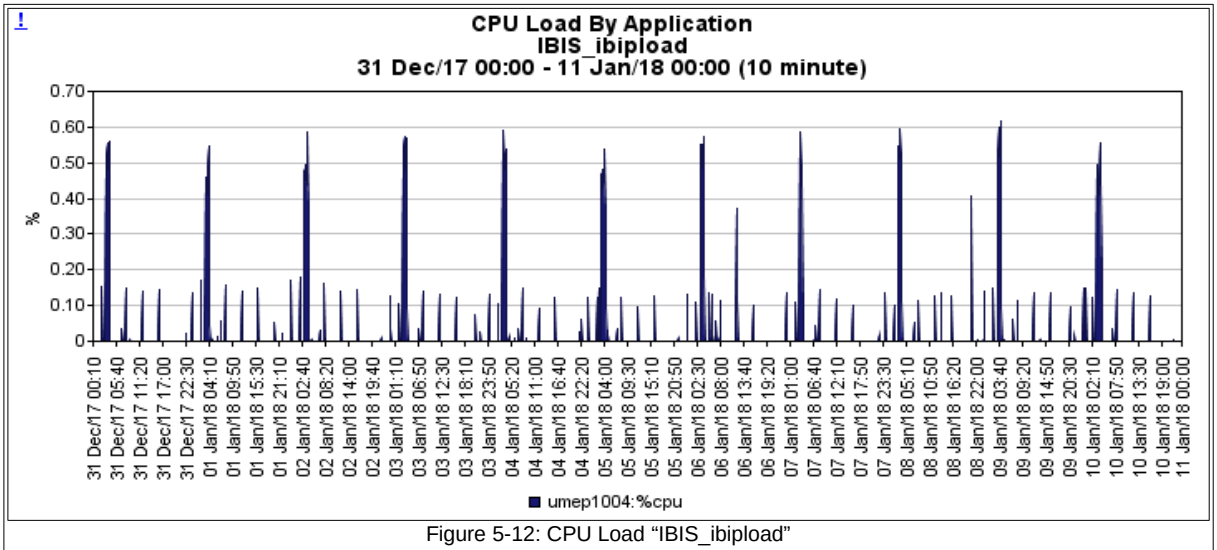

Figure 5-12: CPU Load "IBIS_ibipload"

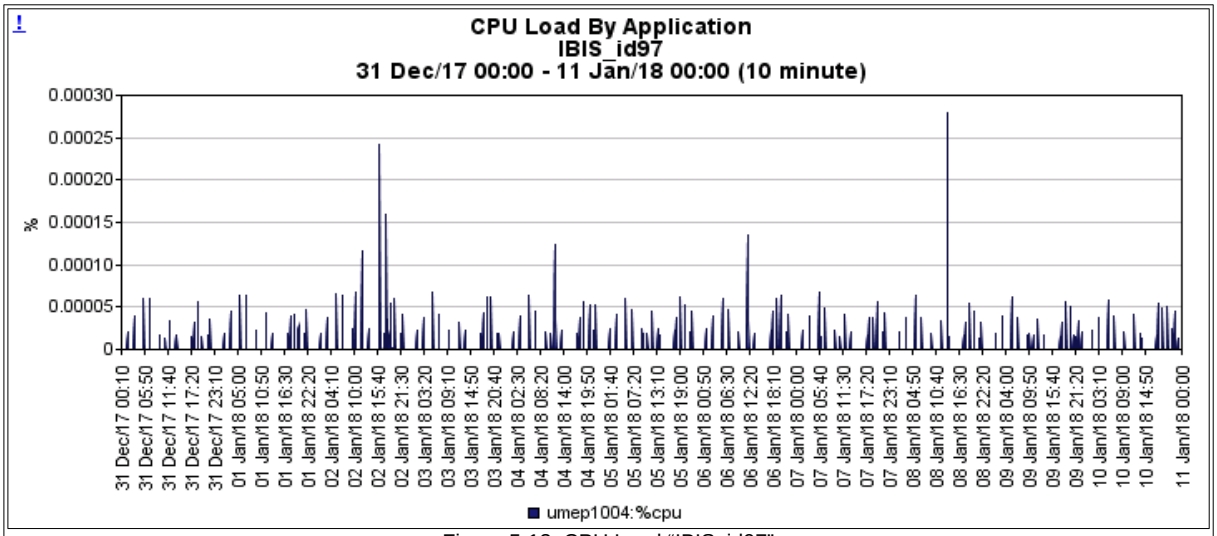

Figure 5-13: CPU Load “IBIS_id97”

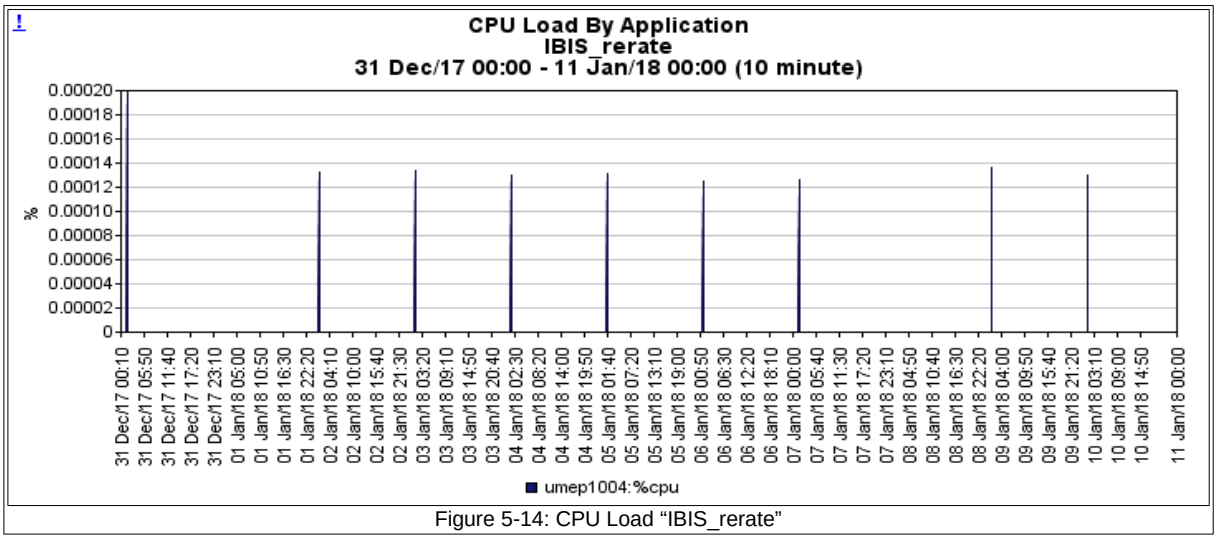

Figure 5-14: CPU Load “IBIS_rerate”

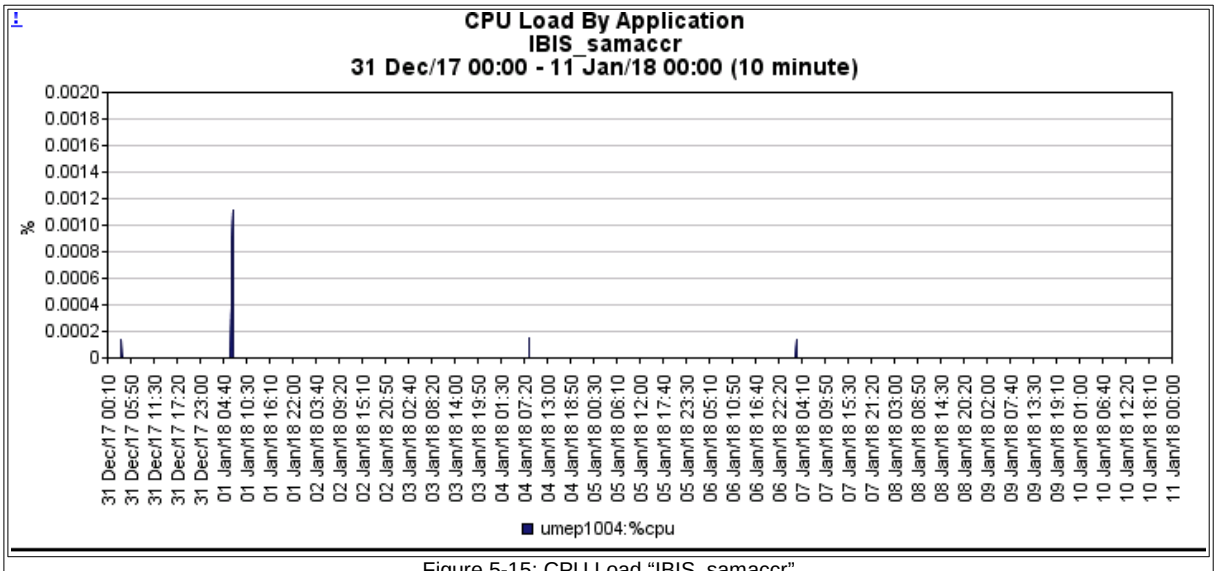

Figure 5-15: CPU Load “IBIS_samaccr”

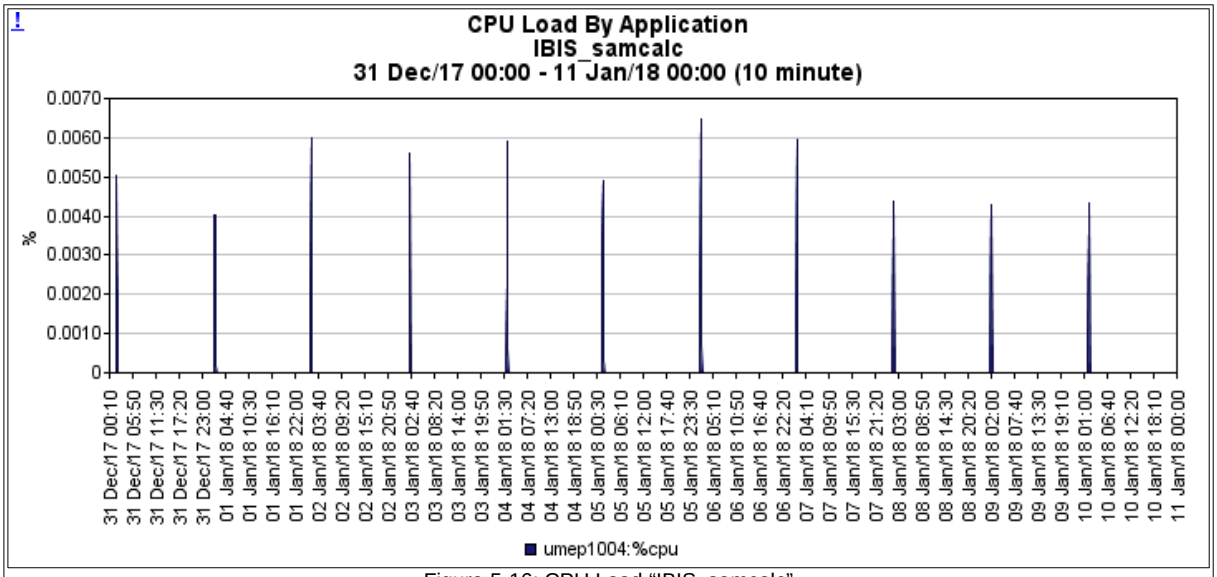

Figure 5-16: CPU Load “IBIS_samcalc”

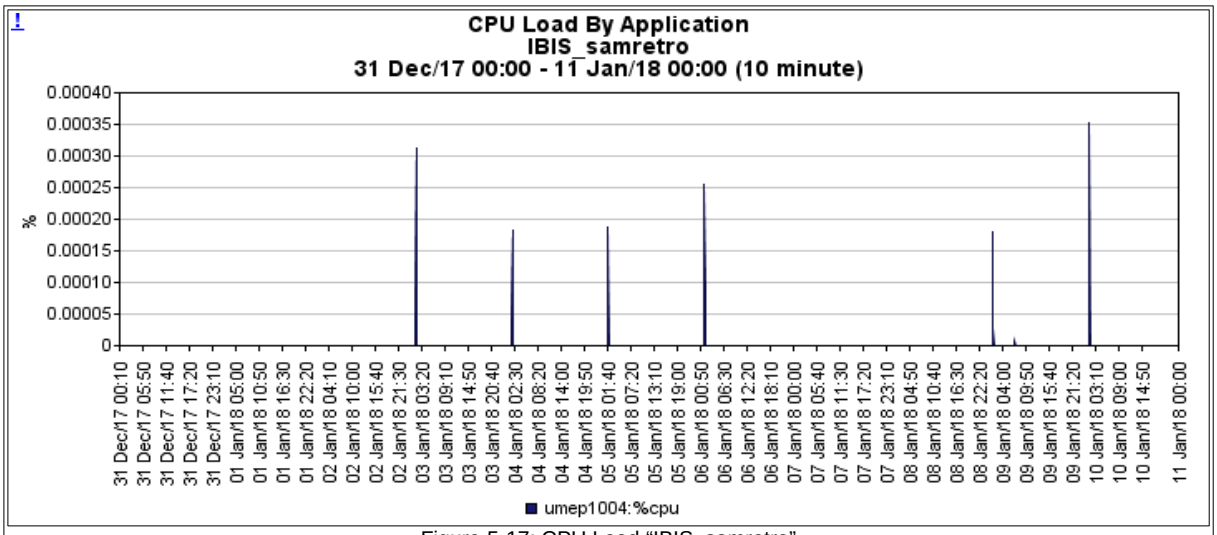

Figure 5-17: CPU Load “IBIS_samretro”

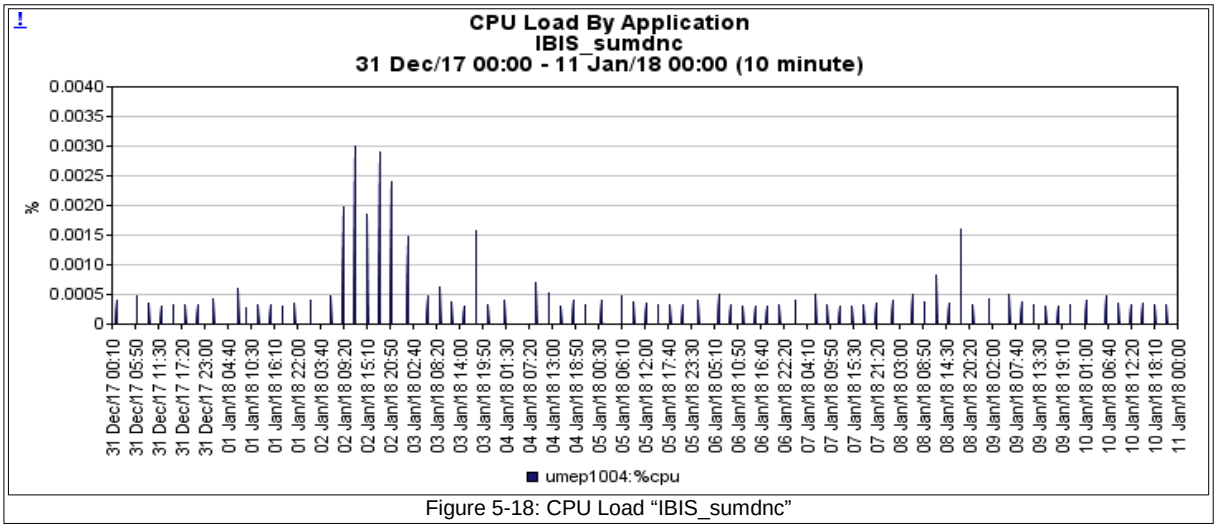

Figure 5-18: CPU Load “IBIS_sumdnc”

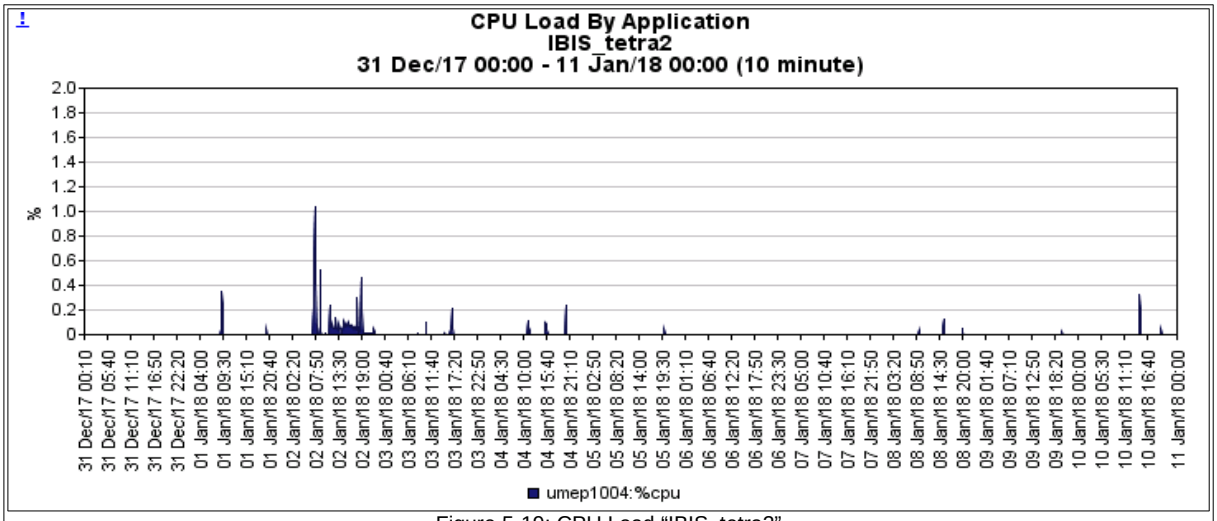

Figure 5-19: CPU Load “IBIS_tetra2”

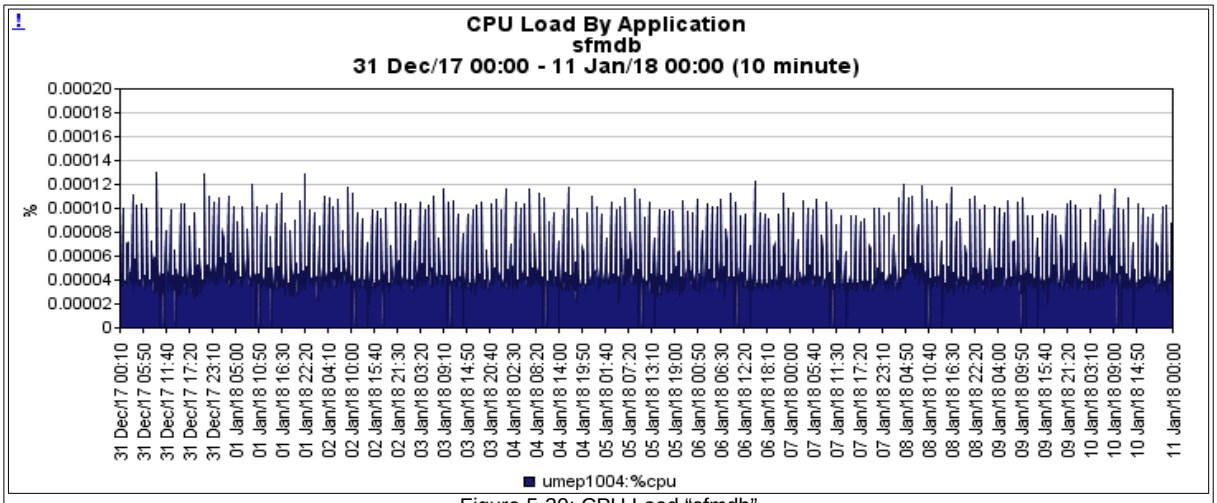

Figure 5-20: CPU Load “sfmdb”

The SSO related load is shown below.

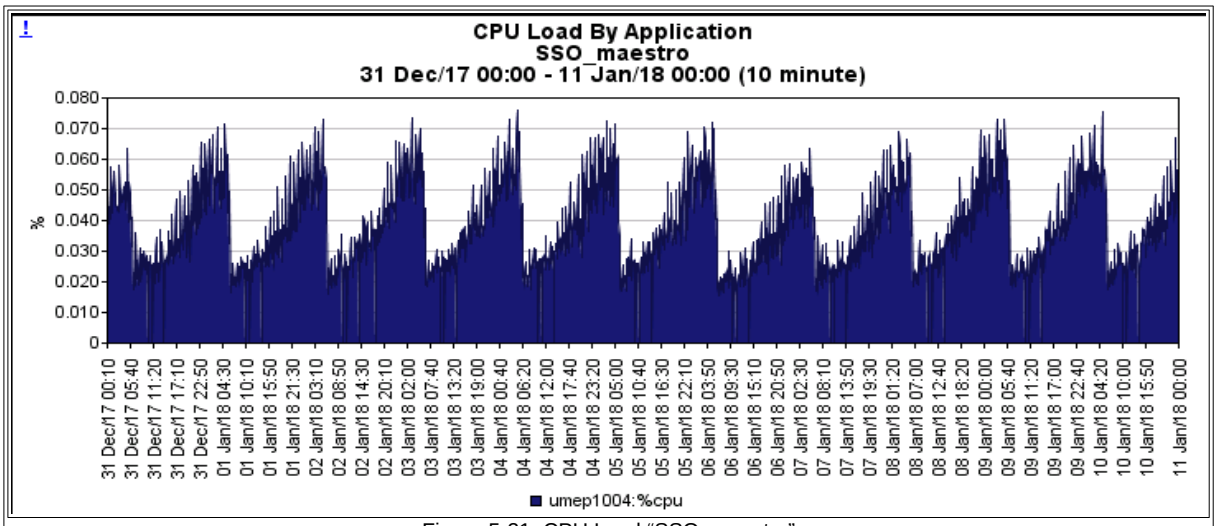

Figure 5-21: CPU Load “SSO_maestro”

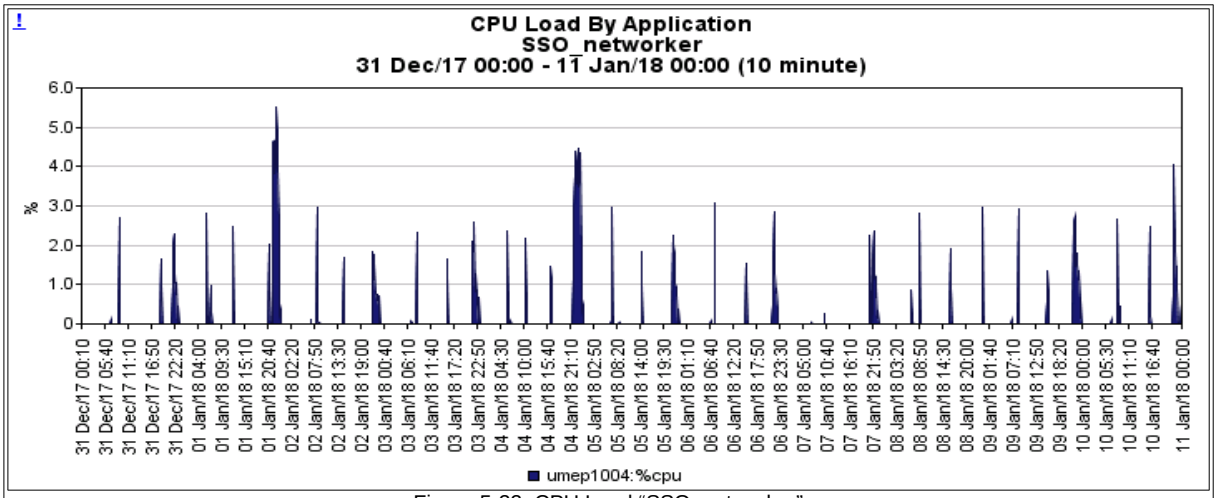

Figure 5-22: CPU Load “SSO_networker”

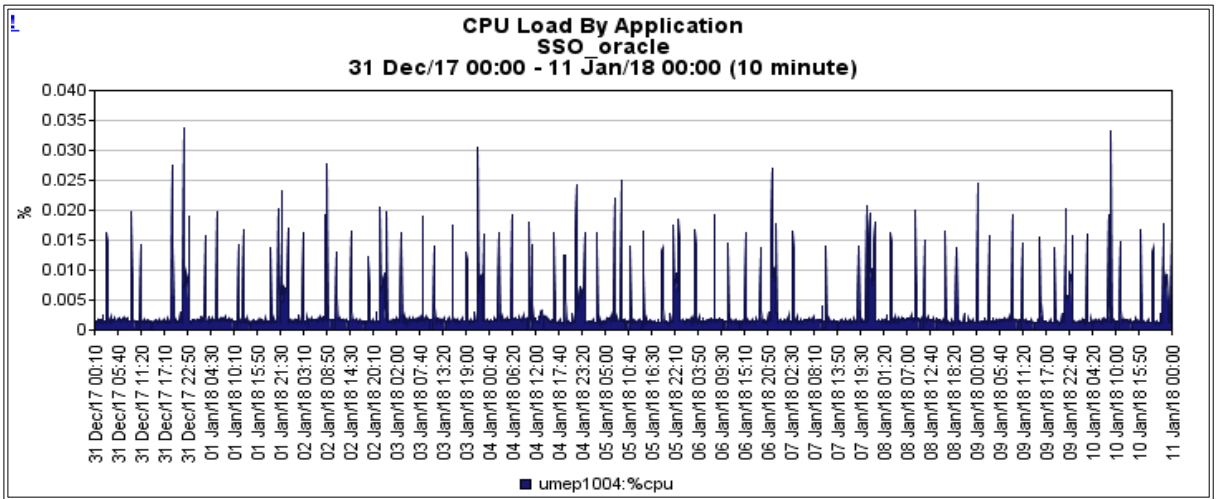

Figure 5-23: CPU Load “SSO_oracle”

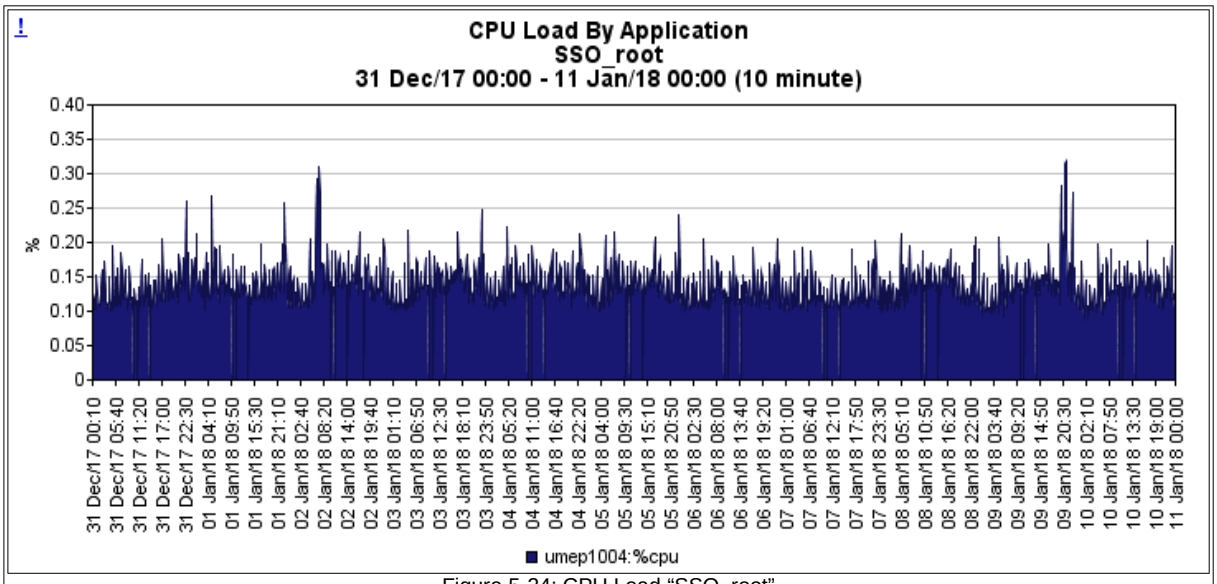

Figure 5-24: CPU Load “SSO_root”

Finally we have the load generated by the management tooling, TeamQuest & Tivoli.

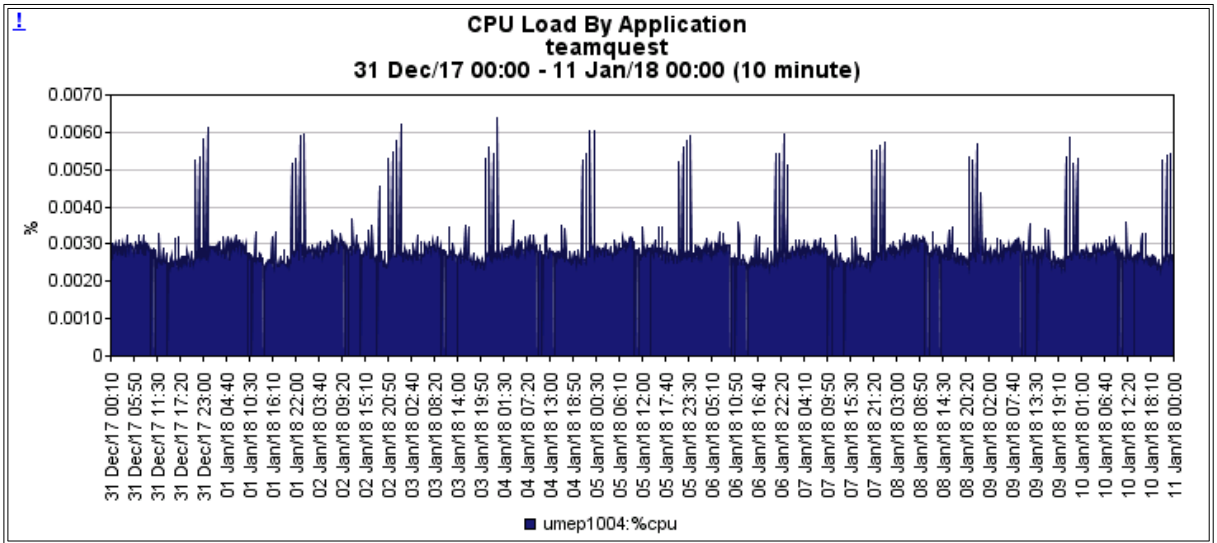

Figure 5-25: CPU Load “teamquest”

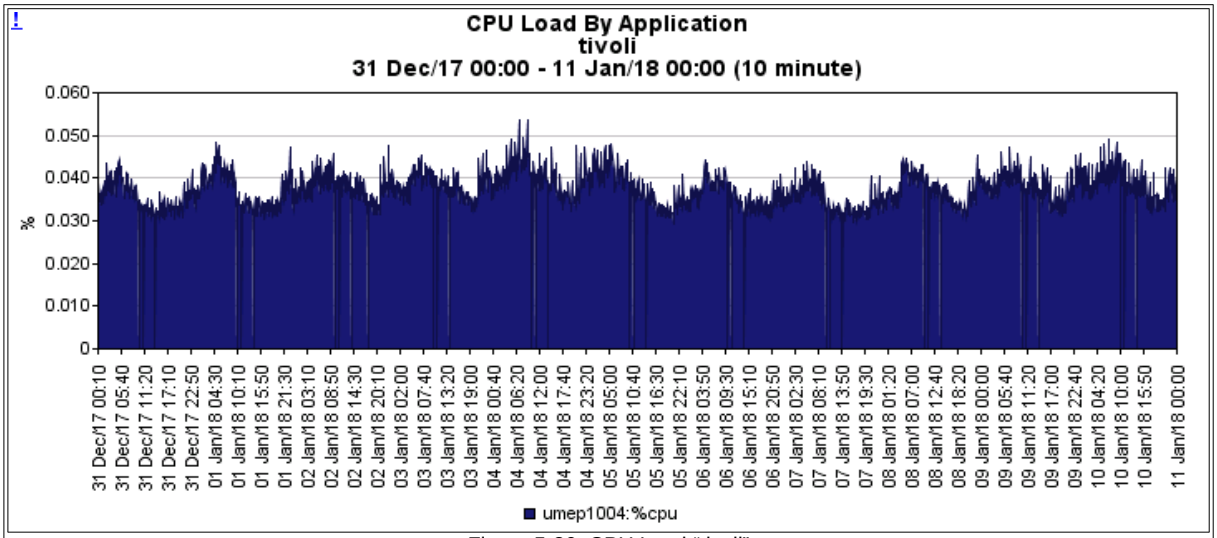

Figure 5-26: CPU Load “tivoli”

Seen the potential to process more CDR's we expect that there are no processes queuing up due to a NO stress on this resource. Figure 5-27 Shows that this is the case.

Also the approach of increasing the number of LCPU’s at periods of high CDR activity helps to keep the server out of the saturation zone.

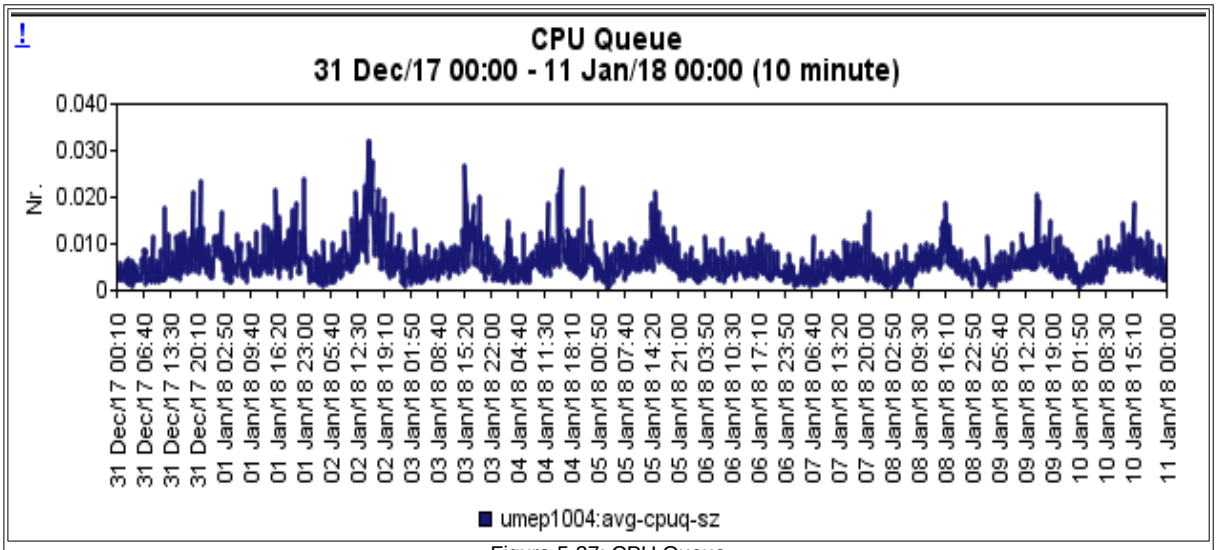

Figure 5-27: CPU Queue

To cross check of one or more LCPU's are 100% used up due to a single thread process we check the %idle activity for each LCPU. So no CPU hogging processes as we did not find any moment of 0%.

Found 8 non active LCPU's, so seems they have not been activated yet. We measured LCPU numbers from 00 → 79 so a total of 80, the maximum number of LCPU's we saw active was 72 so left 8.

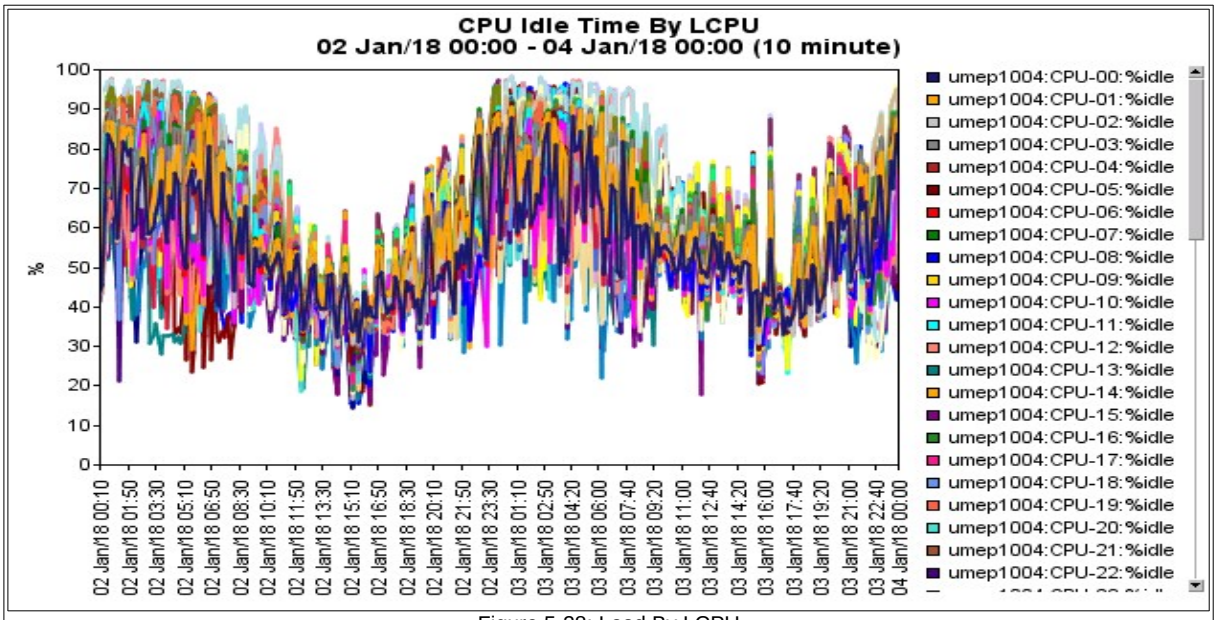

Figure 5-28: Load By LCPU

So basically we did not see any stress on this resource with potential to process more CDR's.

6. Memory Usage

The total physical memory installed is 167.354.016 Kbytes or about 159,6 Gbytes.

This reveals that the memory resource does not encounter any stress during this period since the

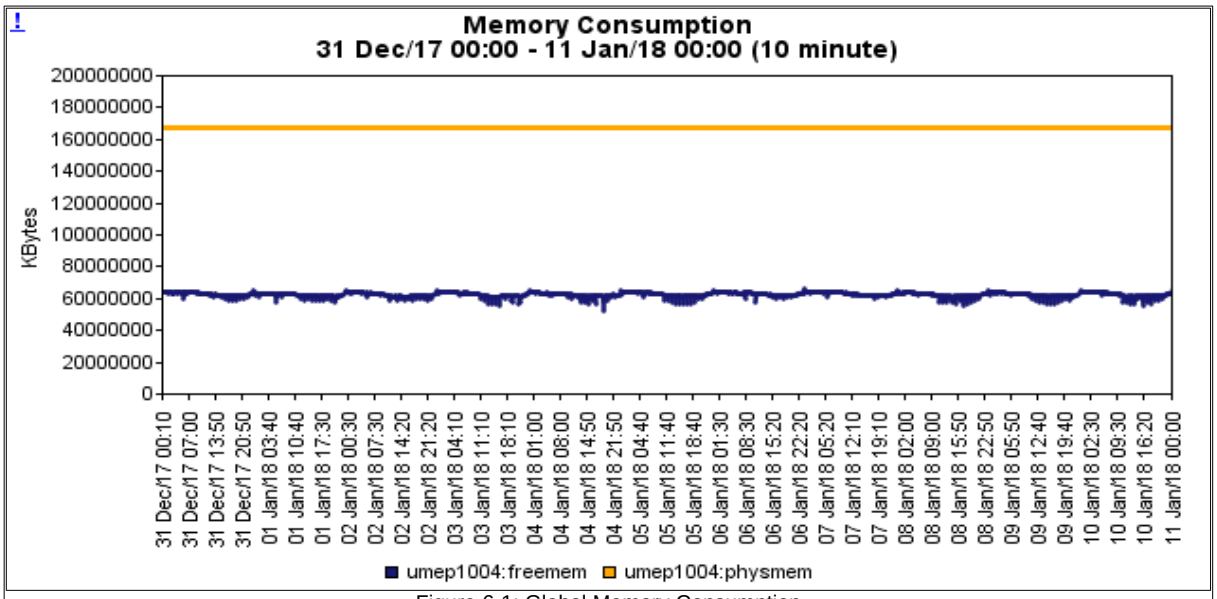

Figure 6-1: Global Memory Consumption

available physical memory (not allocated “freemem”) is never lower than about 53 GBytes.

Seen the size of the unallocated memory there are no eminent issues and largely sufficient to increase memory on the level of Oracle and the number of active “behldnc” processes if that would be needed.

As such the paging activity, due to memory pressure, on scan & page-out level is almost non-existing.

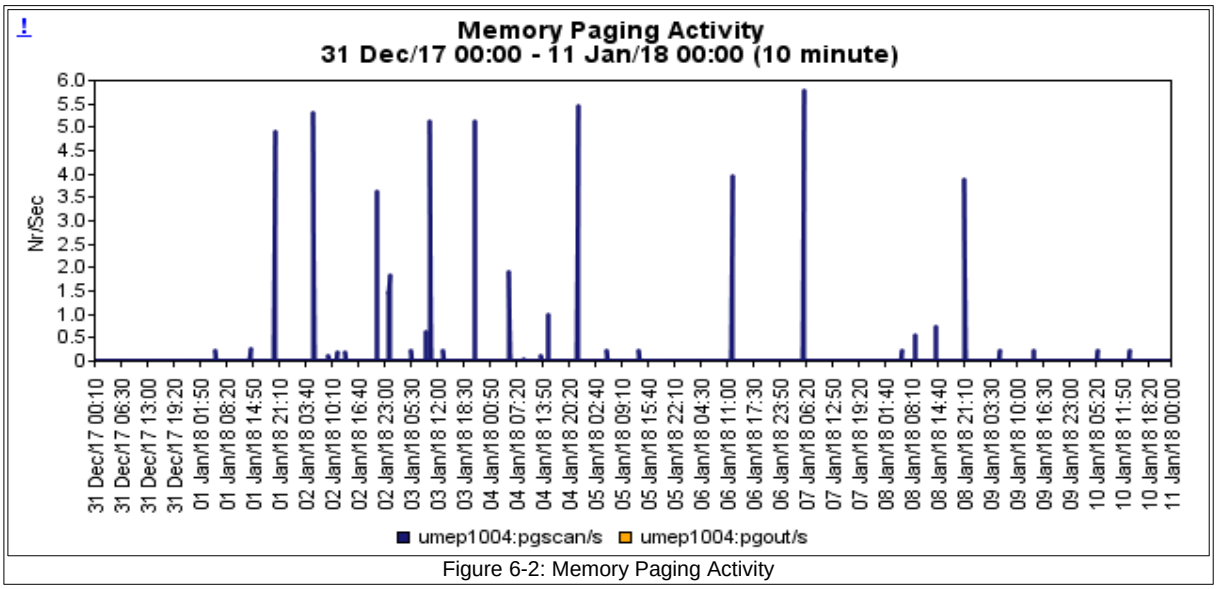

Figure 6-2: Memory Paging Activity

The swap space consumption is shown in the figure below.

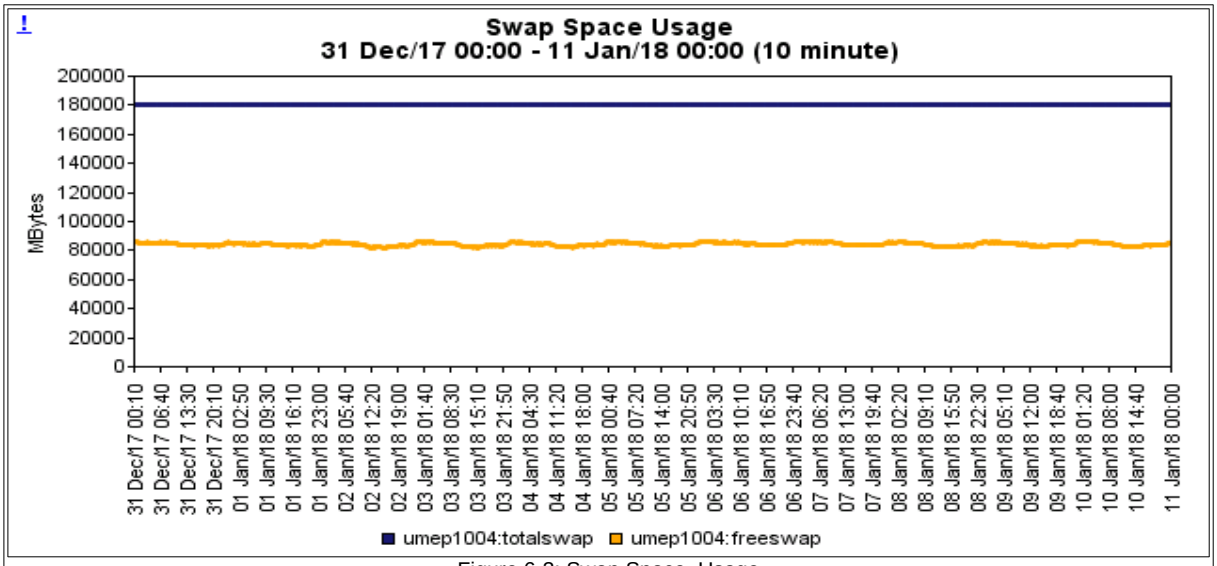

Figure 6-3: Swap Space Usage

The swap space has sufficient free space, about 76 Gbytes, to handle more applications/processes allocating memory. For each data page some swap space will be reserved.

We do not physically use the SWAP space as there is no paging-out activity, see Figure 6-2.

7. I/O Subsystem Activity

We will look at the I/O subsystem for each Volume Group and see if issues occurred. We show in the tables the mapping between mount point and the logical volume. Also we list the activity of the different disks. We exclude the weekends as the load is significant lower at that moment and as such not really relevant.

Each time we have 4 paths to the disks that are part of the Volume Group, only one is active and shows up in the charts.

7.1 Volume Group "vg00"

This contains the following mapping:

< System	< Name	< Device
umep1004	/	/dev/vg00/lvol3
umep1004	/export/home	/dev/vg00/lvol11
umep1004	/opt	/dev/vg00/lvol5
umep1004	/root	/dev/vg00/lvol12
umep1004	/stand	/dev/vg00/lvol1
umep1004	/tmp	/dev/vg00/lvol10
umep1004	/usr	/dev/vg00/lvol4
umep1004	/usr/ecc	/dev/vg00/lvol13
umep1004	/var	/dev/vg00/lvol9
umep1004	/var/adm/crash	/dev/vg00/lvol6
umep1004	/var/adm/sa	/dev/vg00/lvol7
umep1004	/var/adm/sw	/dev/vg00/lvol8

Table 7.1: vg00

The disk activity, only one disk "disk121", looks as following:

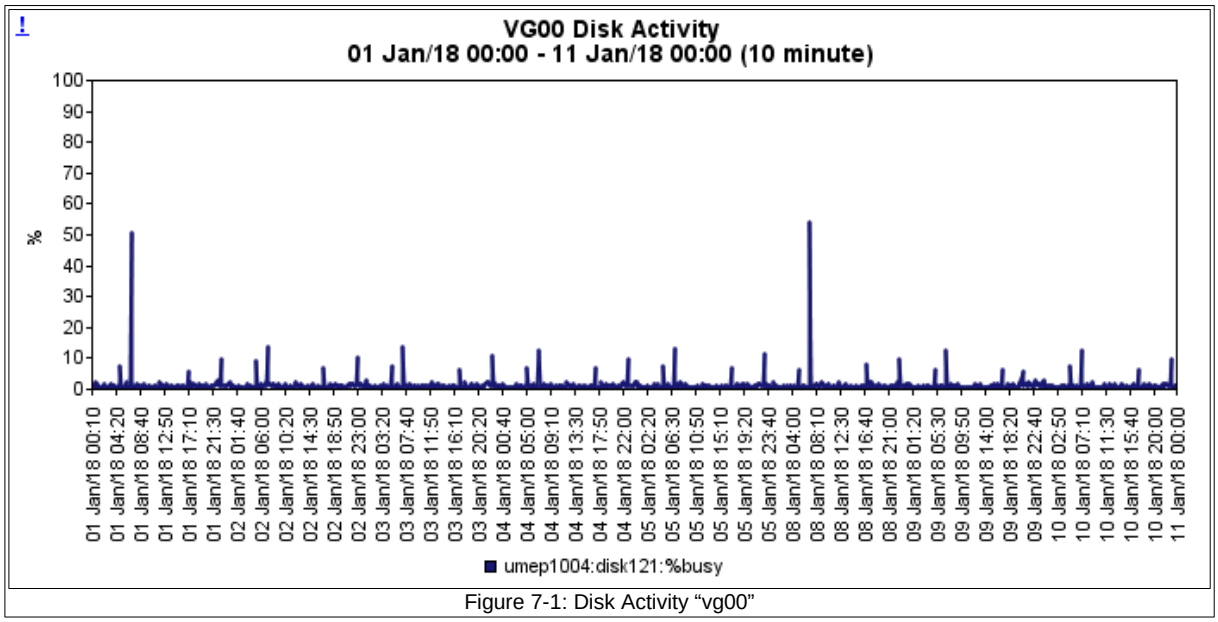

Figure 7-1: Disk Activity "vg00"

Few load is encountered.

7.2 Volume Group “vgibip_ibip”

This contains the following mapping:

System	Name	Device
umep1004	/ibip	/dev/vgibip_ibip/ibip
umep1004	/ibip/CDR_T/error	/dev/vgibip_ibip/CDRerr
umep1004	/ibip/CDR_T/import	/dev/vgibip_ibip/CDRimp
umep1004	/ibip/CDR_T/input_TOPAS	/dev/vgibip_ibip/CDRinTOP
umep1004	/ibip/CDR_T/transfer	/dev/vgibip_ibip/CDRtrns
umep1004	/ibip/dwh/bctdata	/dev/vgibip_ibip/bct
umep1004	/ibip/dwh/data1	/dev/vgibip_ibip/dwhdat1
umep1004	/ibip/dwh/data2	/dev/vgibip_ibip/dwhdat2
umep1004	/ibip/etl	/dev/vgibip_ibip/etl
umep1004	/ibip/greenbox	/dev/vgibip_ibip/GrnBox
umep1004	/ibip/greenbox/inbox	/dev/vgibip_ibip/GB_inbx
umep1004	/ibip/greenbox/transfer	/dev/vgibip_ibip/GB_trns
umep1004	/ibip/ibi	/dev/vgibip_ibip/ibi
umep1004	/ibip/storage_prod	/dev/vgibip_ibip/stor_prd
umep1004	/ibip/topas	/dev/vgibip_ibip/optTOP
umep1004	/ibip/topas/var	/dev/vgibip_ibip/voptTOP
umep1004	/ibip/topas_arch	/dev/vgibip_ibip/toparch

Table 7.2: vgibip_ibip

This volume group contains five disks named “disk113”, “disk116”, “disk342”, “disk120” and “disk164”.

Device “disk113” encounters the highest load.

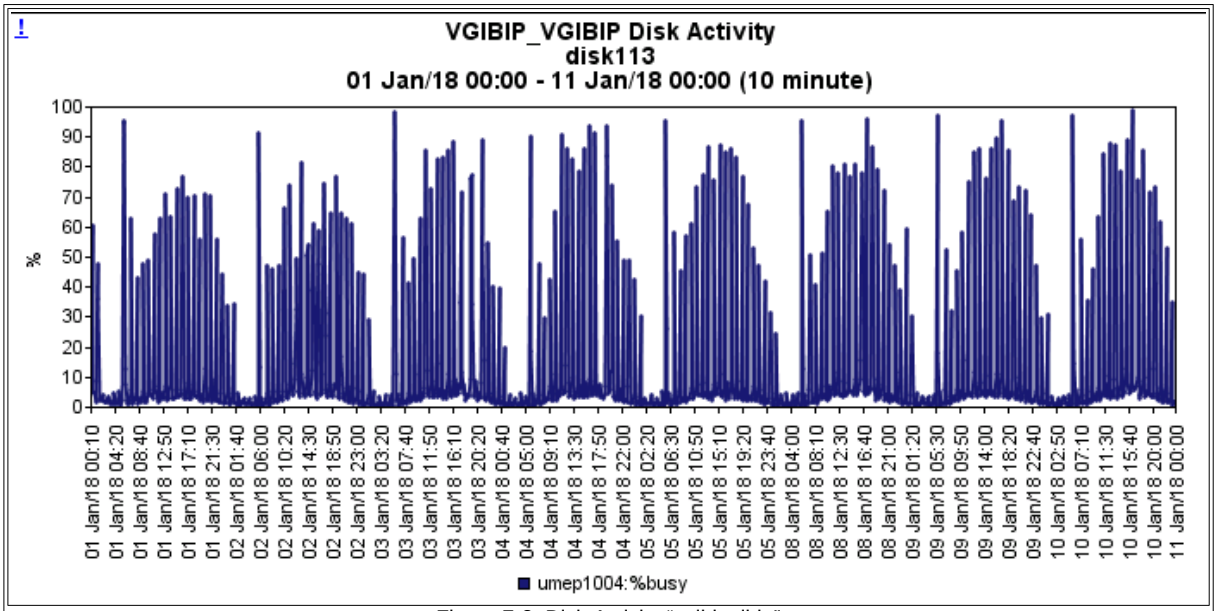

Figure 7-2: Disk Activity “vgibip_ibip”

The remaining disks have a quite low load, see below.

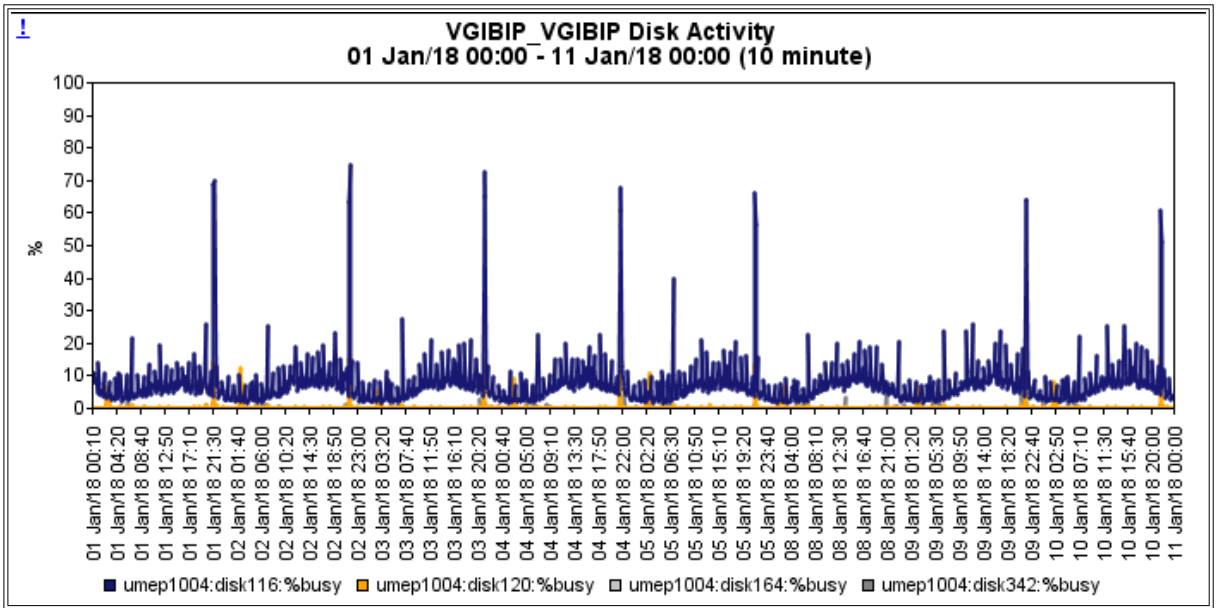

Figure 7-3: Disk Activity “vgibip_ibip”

Let's look at Tuesday the 2th Jan/18 in more detail since that was the busiest day, see Figure 5-1.

This “disk113” is related to the Logical Volumes “CDRerr CDRimp CDRinTOP CDRtrns GB_inbx GB_trns GrnBox bct”.

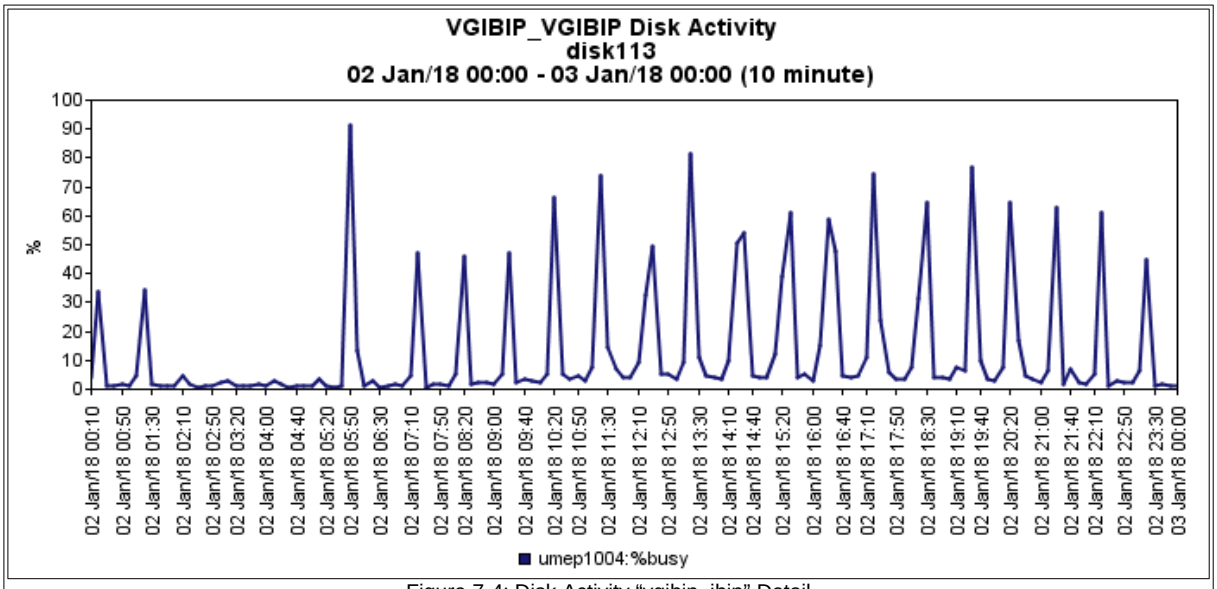

Figure 7-4: Disk Activity “vgibip_ibip” Detail

These peaks correlate with the load of “IBIS_dwh” and can last for 15 minutes.

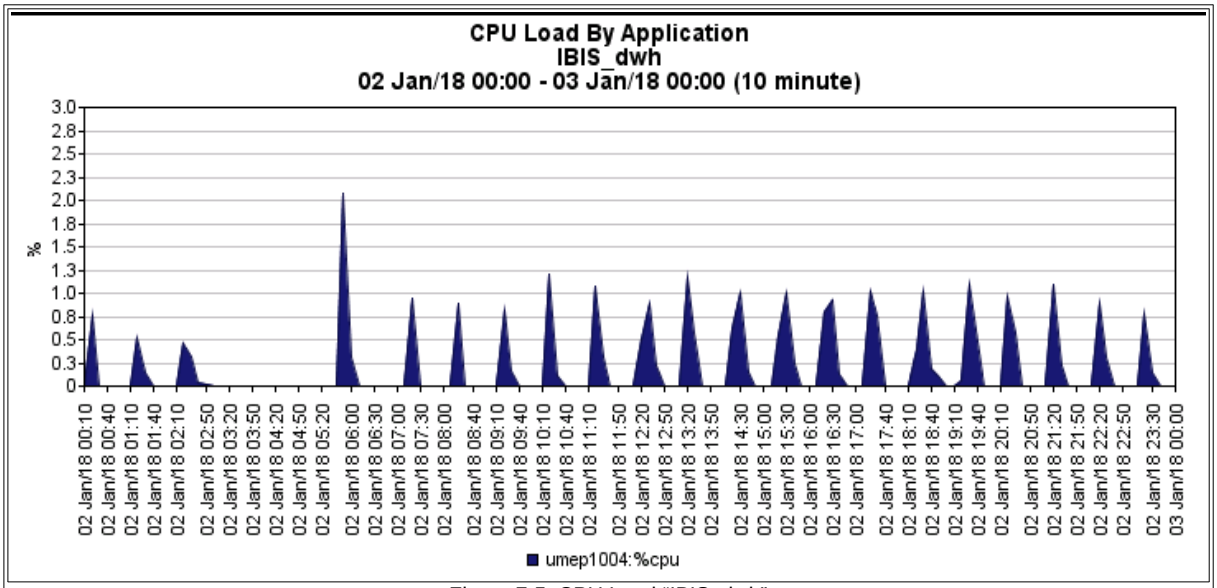

Figure 7-5: CPU Load "IBIS_dwh"

During these peaks some queuing occurs but the wait time during these periods is not that dramatic, up to 40Msec.

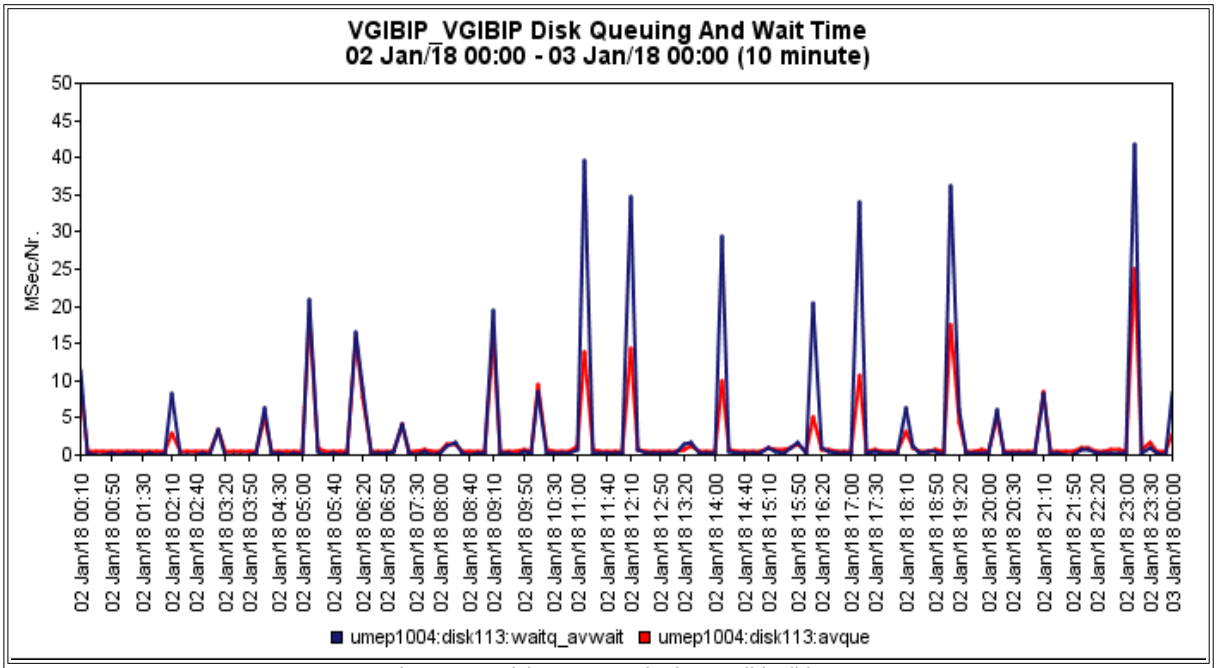

Figure 7-6: Disk Queue Wait Time "vgibip_ibip"

For the remaining disks the queuing activity is very low.

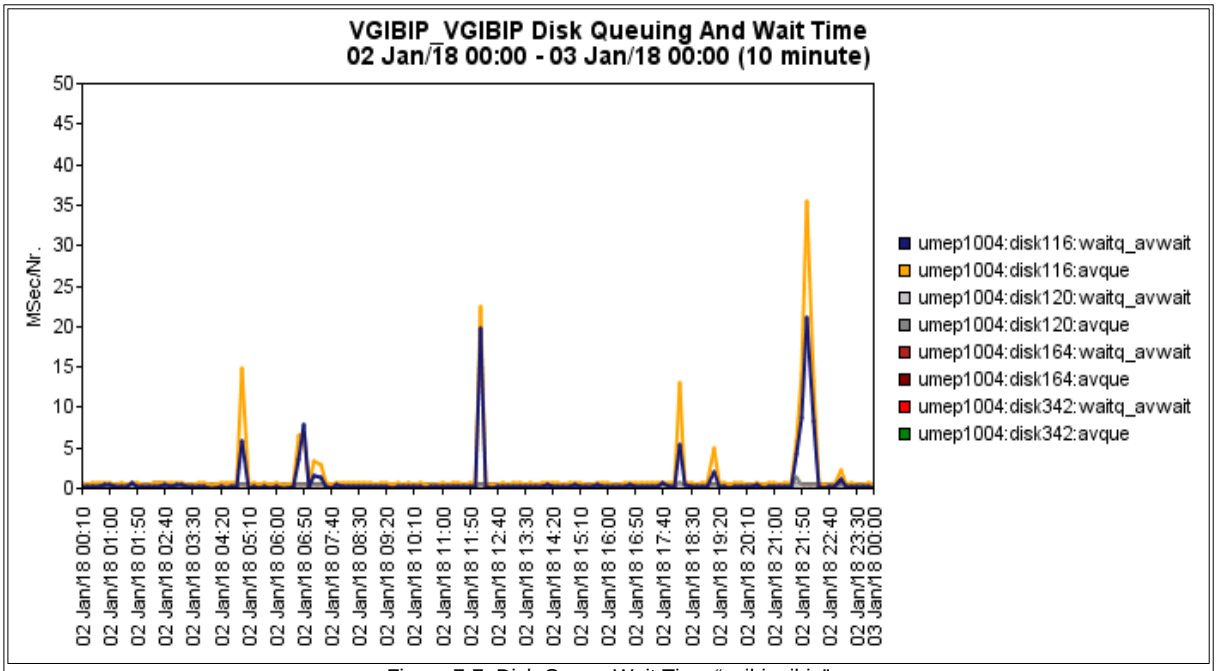

Figure 7-7: Disk Queue Wait Time "vgibip_ibip"

Also the Service Time reveals to be fine as it is around or below 13 Msec for "disk113" and even lower for the other disks.

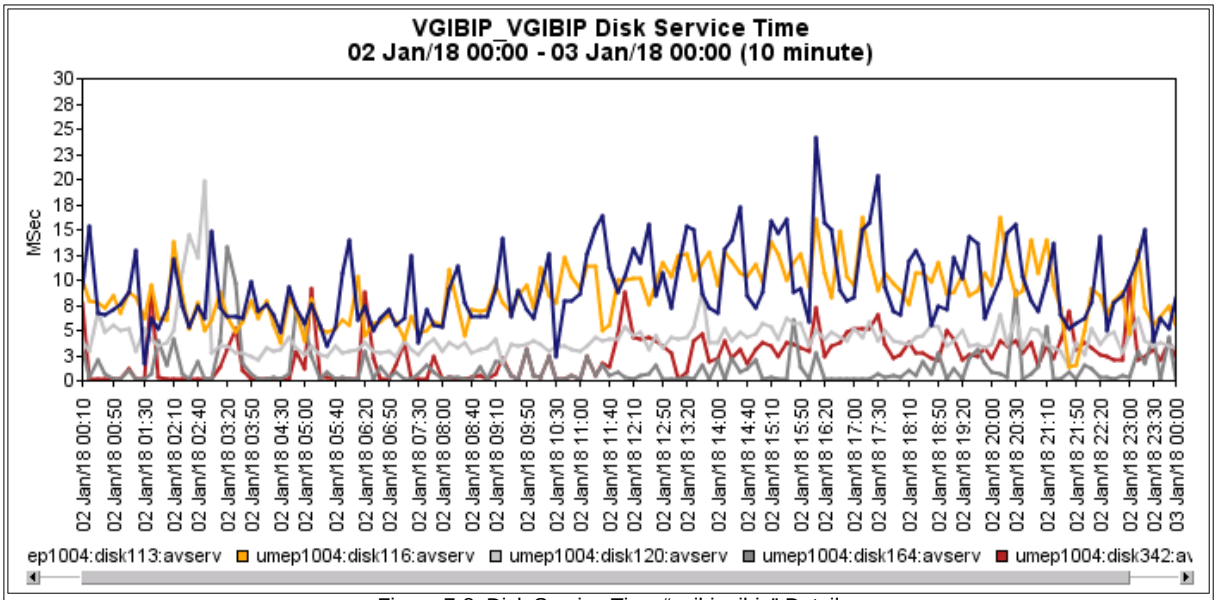

Figure 7-8: Disk Service Time "vgibip_ibip" Detail

7.3 Volume Group "vgibip_data"

This contains the following mapping:

< System	< Name	< Device
umep1004	/ibip/dbs/data	/dev/vgibip_data/data

Table 7.3: vgibip_data

The disk activity, twelve disks in total, “disk318”, “disk316”, “disk315”, “disk111”, “disk107”, “disk106”, “disk317”, “disk117”, “disk112”, “disk76”, “disk104” and “disk73” show a moderate activity.

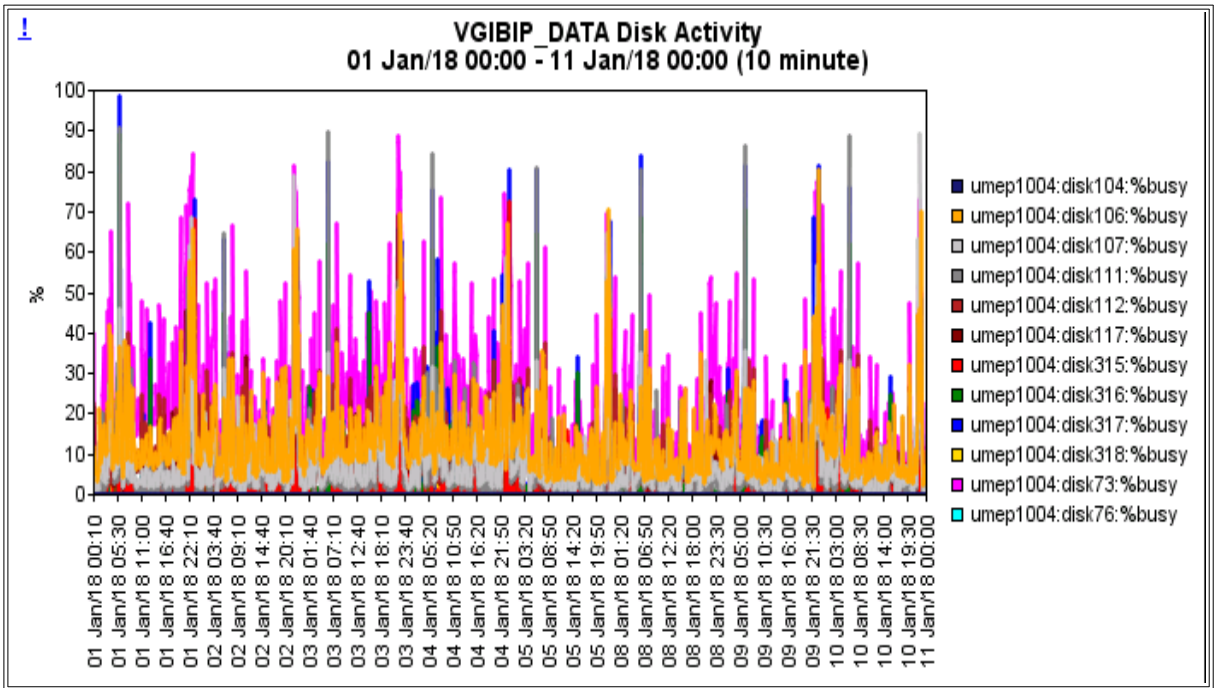

Figure 7-9: Disk Activity “vgibip_data”

The wait time in the Wait Queue is very small.

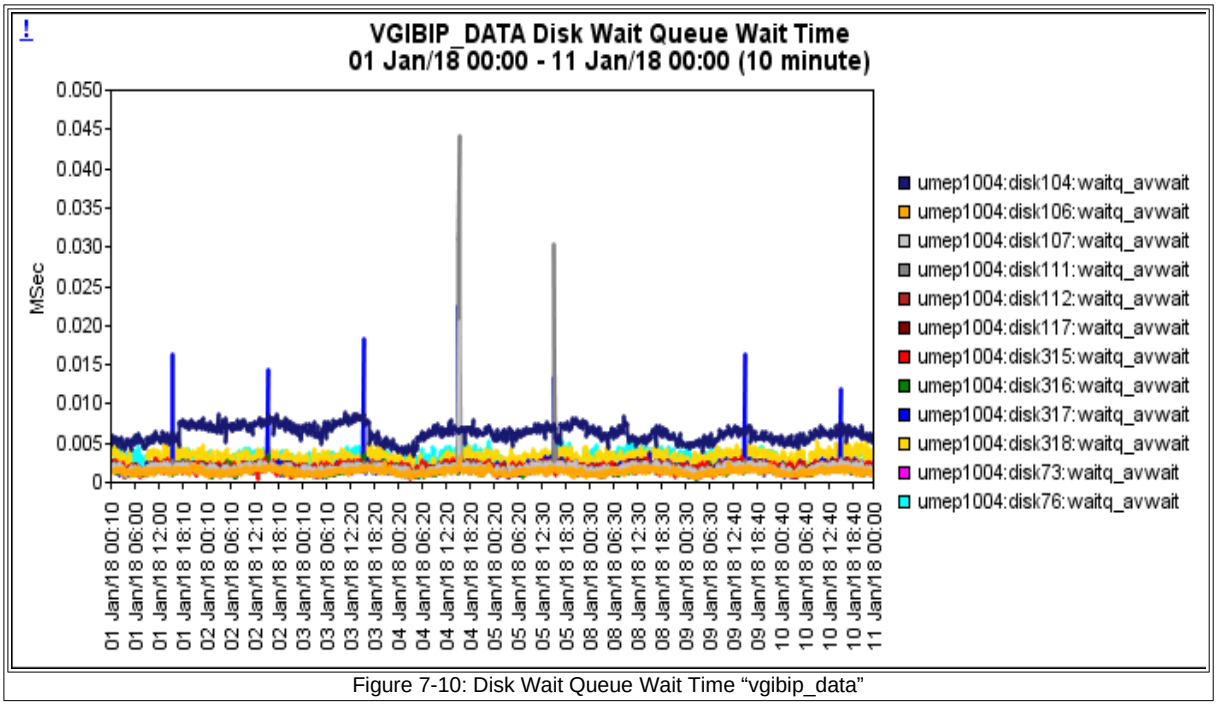

Figure 7-10: Disk Wait Queue Wait Time “vgibip_data”

The service time is excellent since below 6Msec.

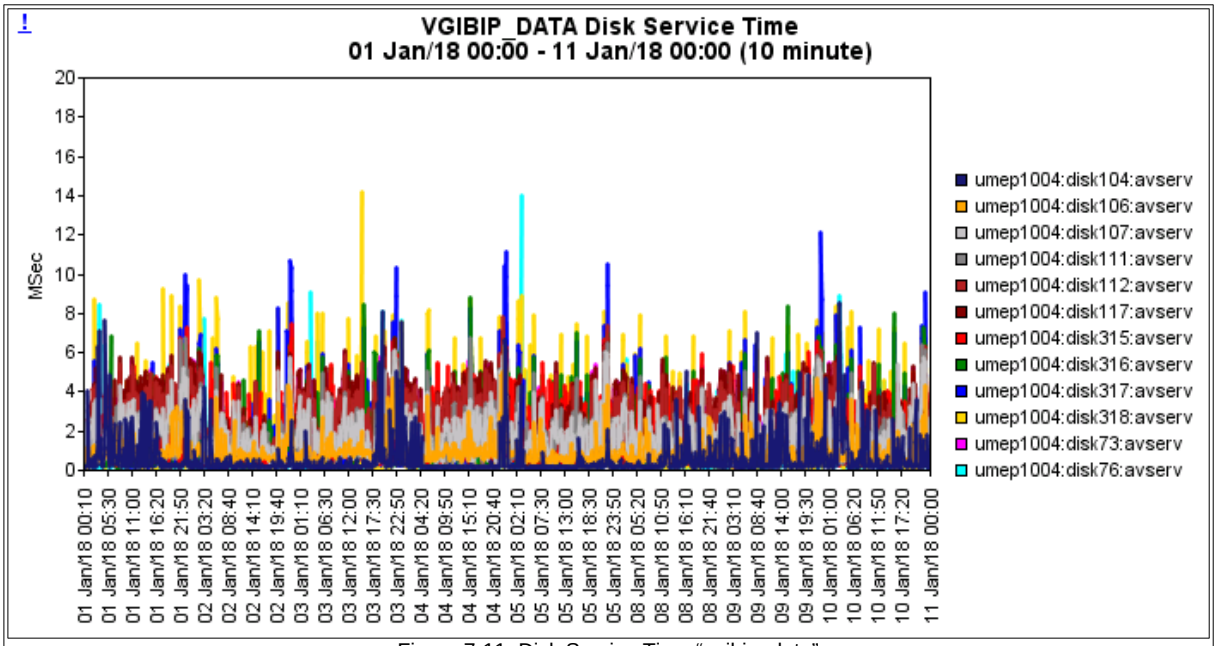

Figure 7-11: Disk Service Time “vgibip_data”

7.4 Volume Group “vgibip_dbs”

This contains the following mapping:

< System	< Name	< Device
umep1004	/ibip/dbs	/dev/vgibip_dbs/dbs
umep1004	/ibip/dbs/redo	/dev/vgibip_dbs/redo
umep1004	/ibip/dbs/sys	/dev/vgibip_dbs/sys
umep1004	/ibip/dbs/tmp	/dev/vgibip_dbs/tmp
umep1004	/ibip/dbs/undo	/dev/vgibip_dbs/undo

Table 7.4: vgibip_dbs

The disk activity, total of seven disks, “disk349”, “disk348”, “disk71”, “disk92”, “disk86”, “disk93” and “disk72” shows also a moderate activity.

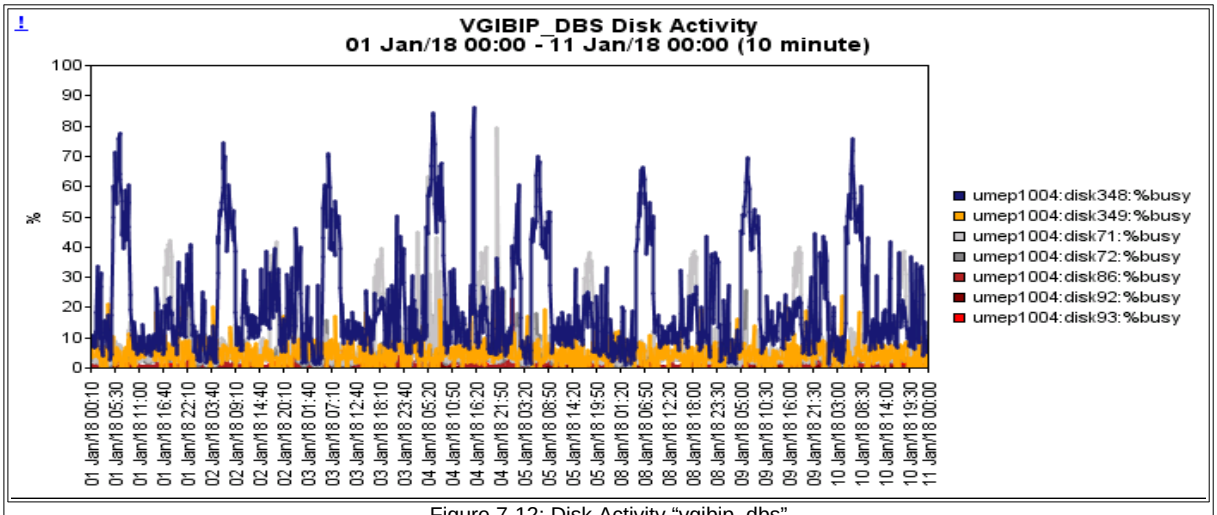

Figure 7-12: Disk Activity “vgibip_dbs”

On “disk348” we have a daily recurring load up to 70% from 4h40 to 9h40.

The wait time in the Wait Queue is very small.

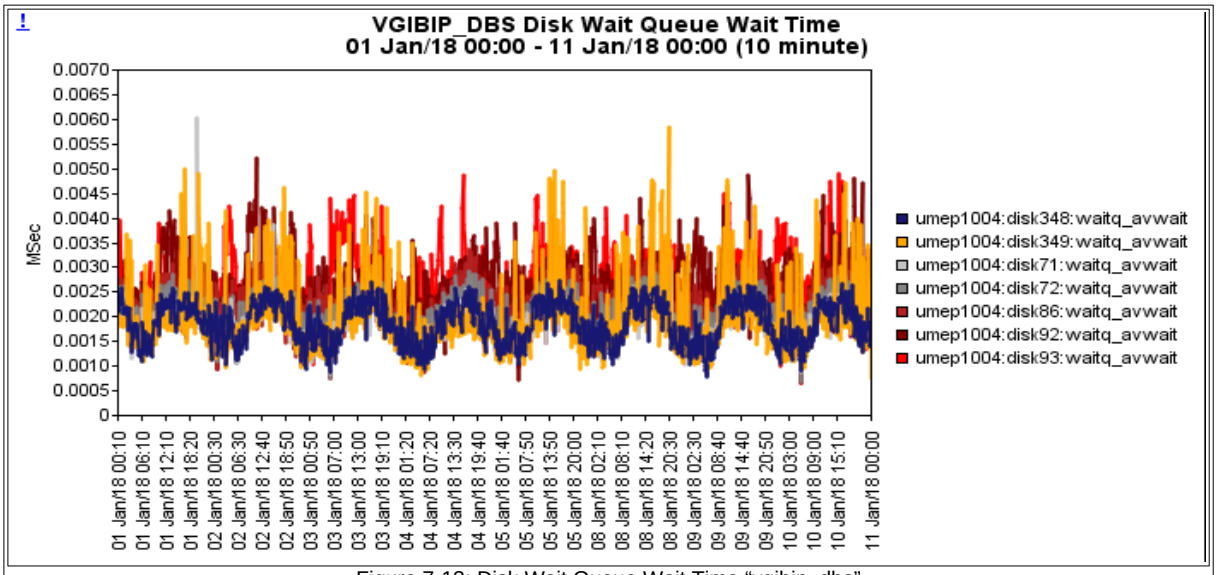

Figure 7-13: Disk Wait Queue Wait Time “vgibip_dbs”

The service time is excellent since below 6Msec.

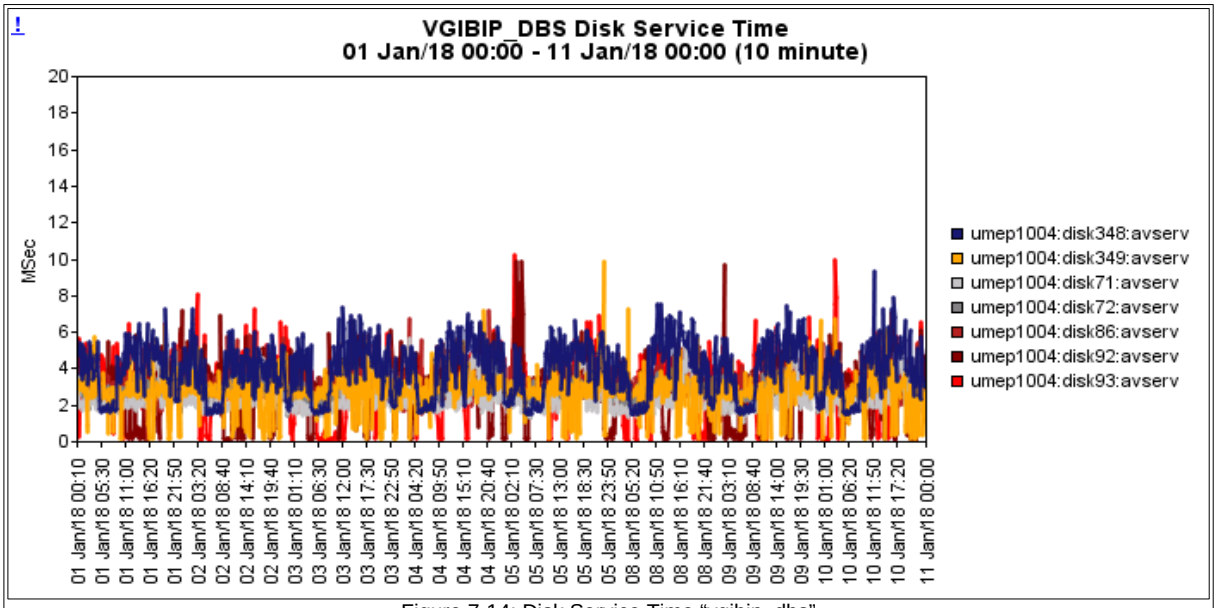

Figure 7-14: Disk Service Time “vgibip_dbs”

7.5 Volume Group “vgibip_index”

This contains the following mapping:

HINV.FileSystem 03 Jan/18 01:00

< System	< Name	< Device
umep1004	/ibip/dbs/index	/dev/vgibip_index/index

Table 7.5: vgibip_index

The disk activity, for ten disks “disk334”, “disk333”, “disk109”, “disk118”, “disk108”, “disk110”, “disk119”, “disk81”, “disk139” and “disk105”, is shown below.

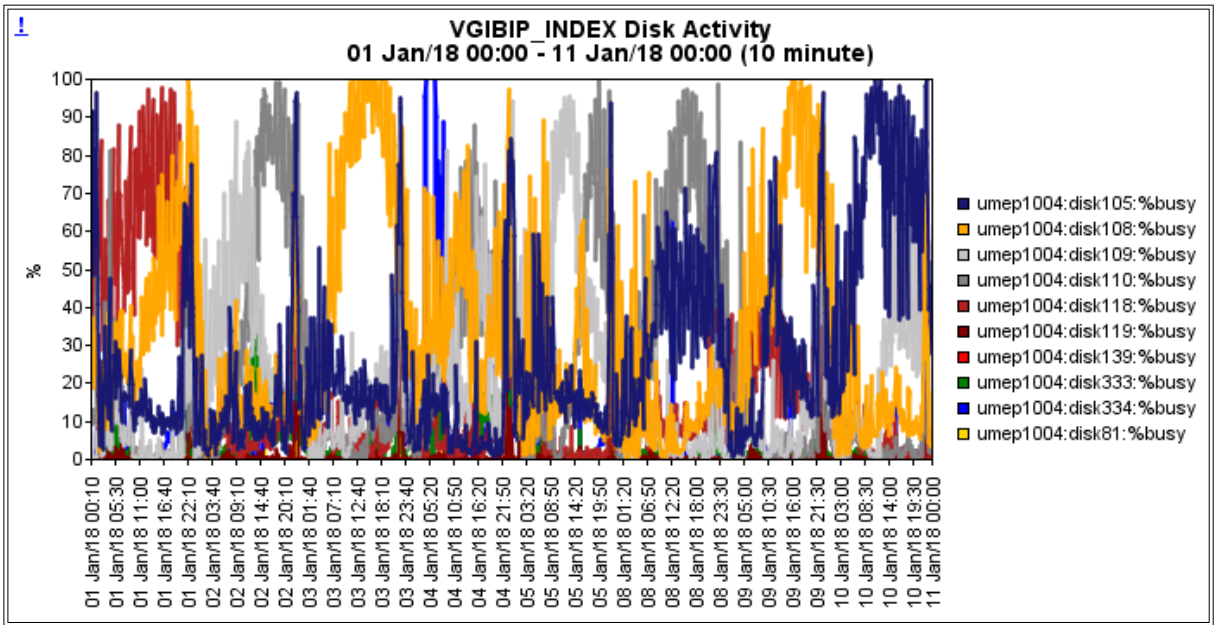

Figure 7-15: Disk Activity “vgibip_index”

The disks, “disk105”,”disk108”, “disk109”, “disk110”, “disk118” and “disk334” show moments of 90% to 100% utilization. This is quite high and need attention since these there is no more room for handling more I/O. Currently the queuing time and service times are still fine. We will see that during the modeling exercise that “disk110” is at the limit.

The wait time in the Wait Queue is very small.

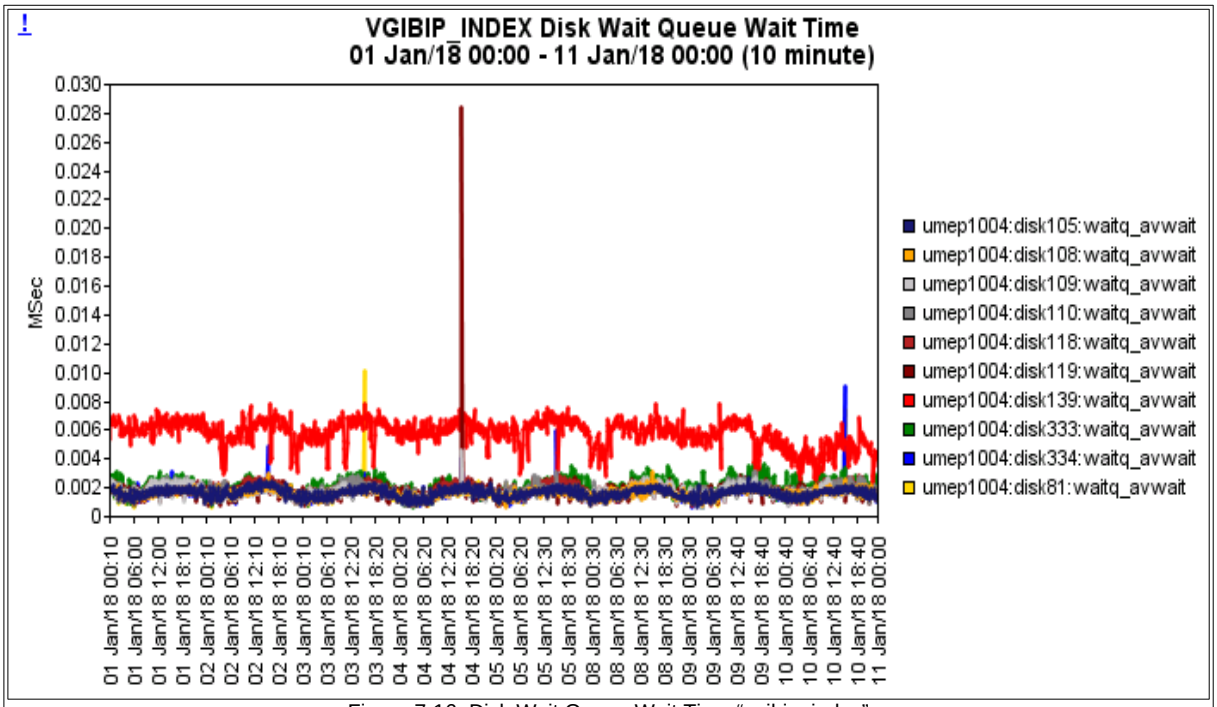

Figure 7-16: Disk Wait Queue Wait Time “vgibip_index”

Also the Service Time reveals to be fine as it is around or below 4 Msec which is excellent.

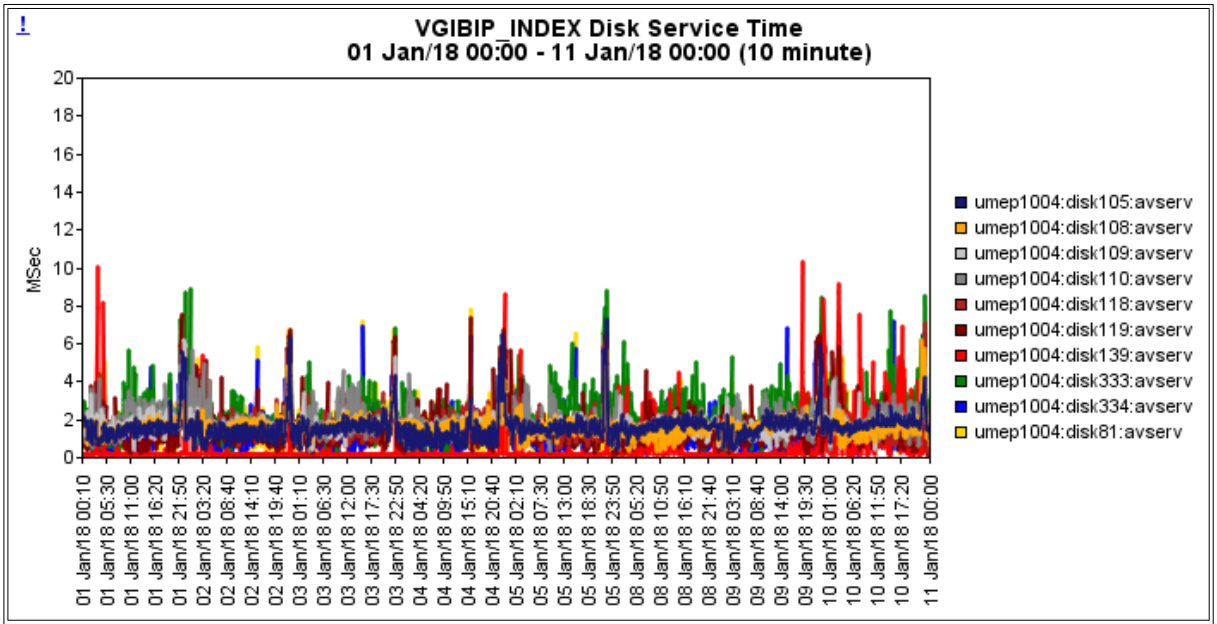

Figure 7-17: Disk Service Time “vgibip_index”

7.6 Volume Group “vgibip_arch”

This contains the following mapping:

System	Name	Device
umep1004	/ibip/dbs/arch	/dev/vgibip_arch/arch

Table 7.6: vgibip_arch

The disk activity, of the four disks “disk115”, “disk12”, “disk114” and “disk22” shows a quite small activity.

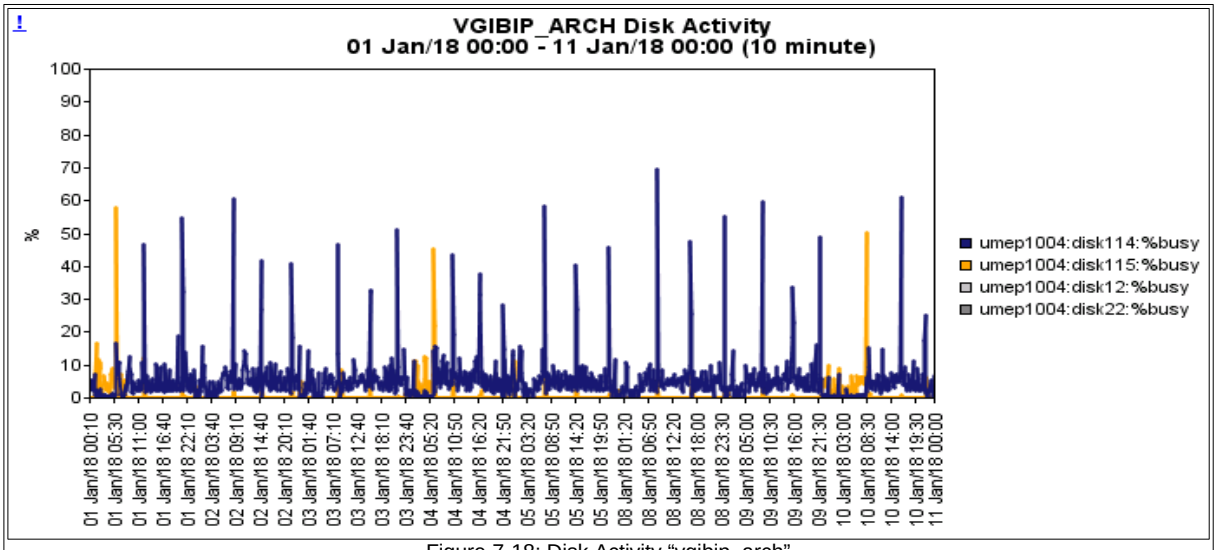

Figure 7-18: Disk Activity “vgibip_arch”

The wait time in the Wait Queue is very small.

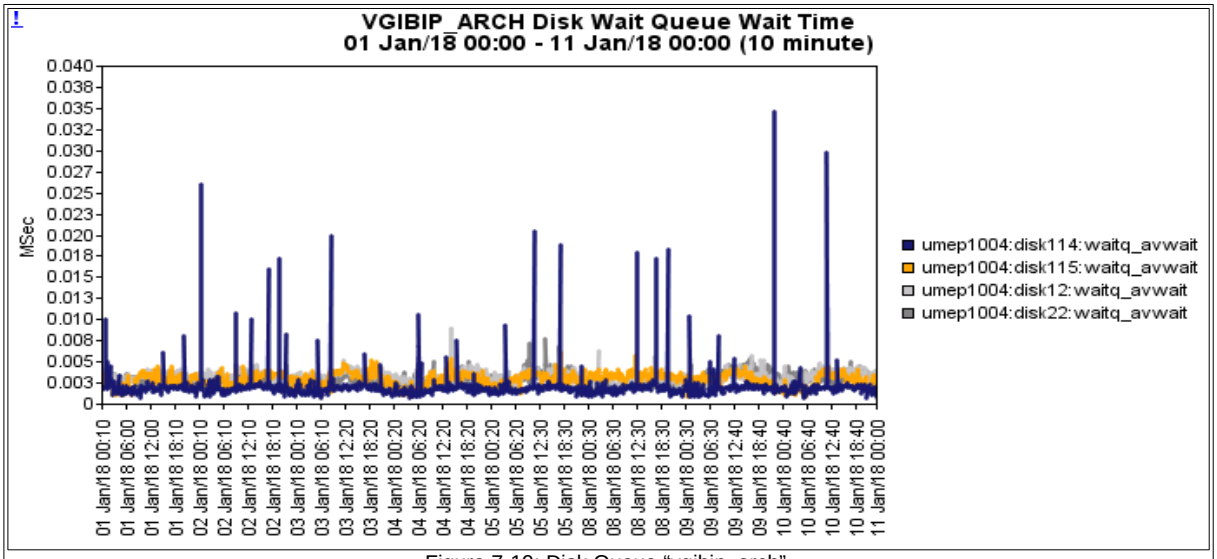

Figure 7-19: Disk Queue “vgibip_arch”

Also the Service Time reveals to be fine as it is around or below 8Msec with peaks to 14Msec.

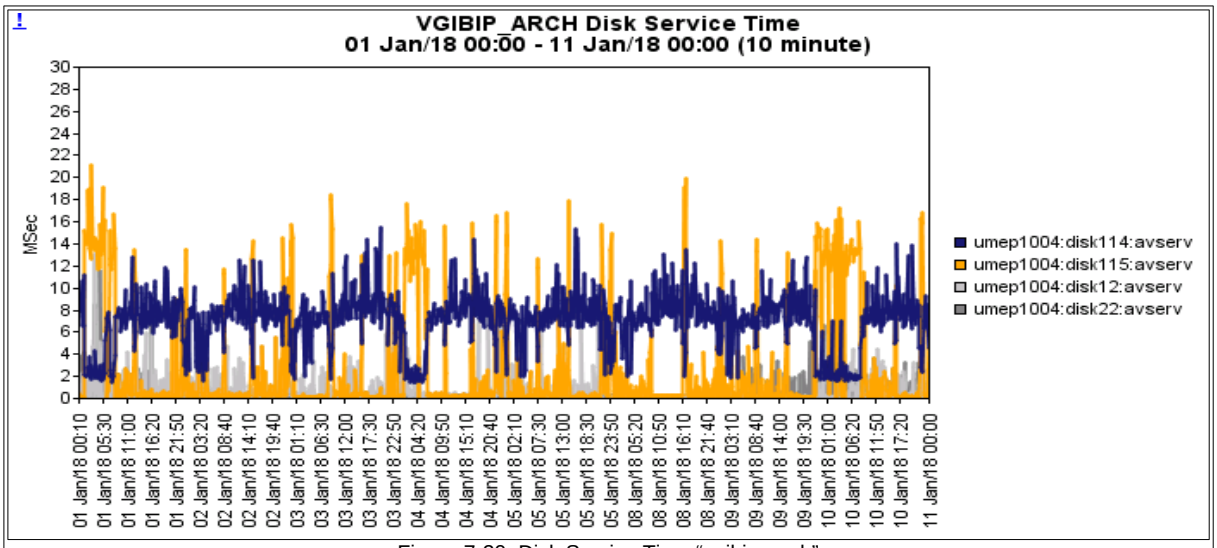

Figure 7-20: Disk Service Time “vgibip_arch”

7.7 Volume Group “vgibip_sysopt”

This contains the following mapping:

< Name	< Device
/etc/lvmconf/other	/dev/vg_sysopt/lvmconf_other
/opt/ipde	/dev/vg_sysopt/ipde
/opt/oracle	/dev/vg_sysopt/oracle
/opt/oracle/dbtools	/dev/vg_sysopt/dbtools
/root/adm/maint	/dev/vg_sysopt/maint
/tivp	/dev/vg_sysopt/tivp
/twsp	/dev/vg_sysopt/twsp
/var/teamquest	/dev/vg_sysopt/teamquest

Table 7.7: vgibip_sysopt

The disk activity, only one disk “disk123” looks as following.

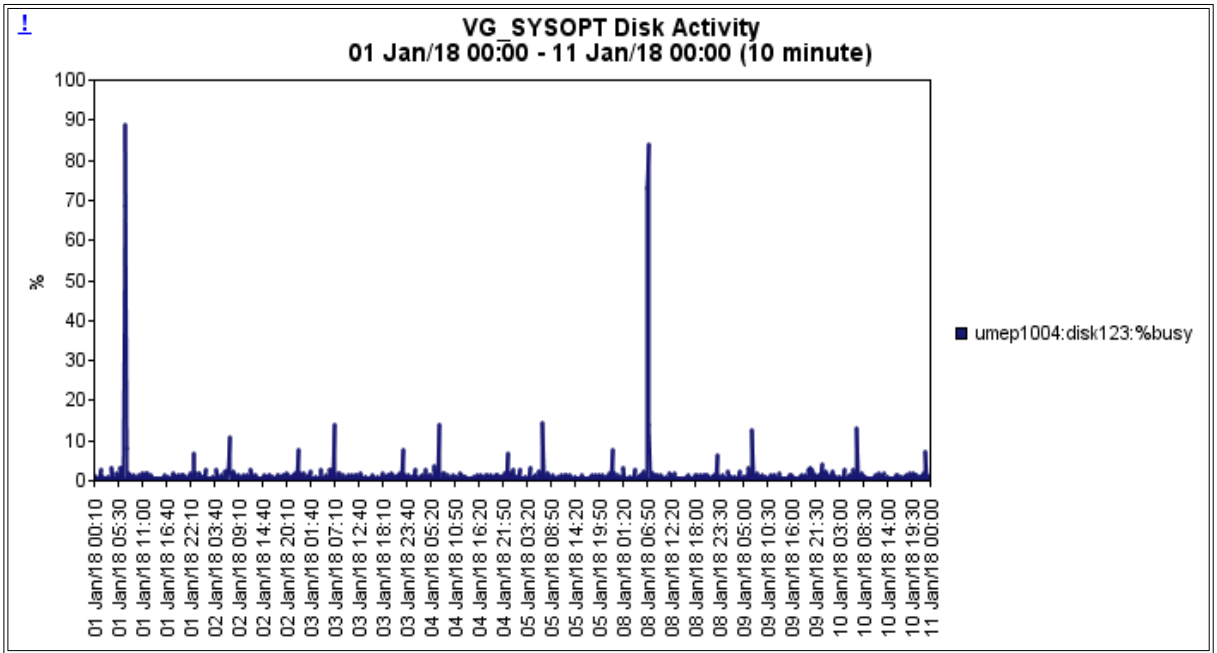

Figure 7-21: Disk Activity “vg_sysopt”

Small activity has been measured and not relevant in the context of this audit.

8. Network Activity

This is currently not relevant as the files are loaded via a local file system load so we did not investigate this resource.

9. Capacity Planning

The first request we got is to evaluate the current HP-UX hardware capacity when the load is increased by steps of 5% per year. We look what happens during the next 5 years.

9.1 Increase Load With 5% More CDR's By Year

The total amount of CDR's processed by day from Friday 29th Dec/17 → Monday 29th Jan/18 is shown below. Remember the best period to use is the first week of January as we have both the Rating and Bill Run activity ongoing. But overall during the working days the nr. of processed CDR's is quite stable and about 120 to 127 Mio CDR's processed per day. Busiest day was Friday 12th Jan/18 where 127,8 Mio CDR's where handled.

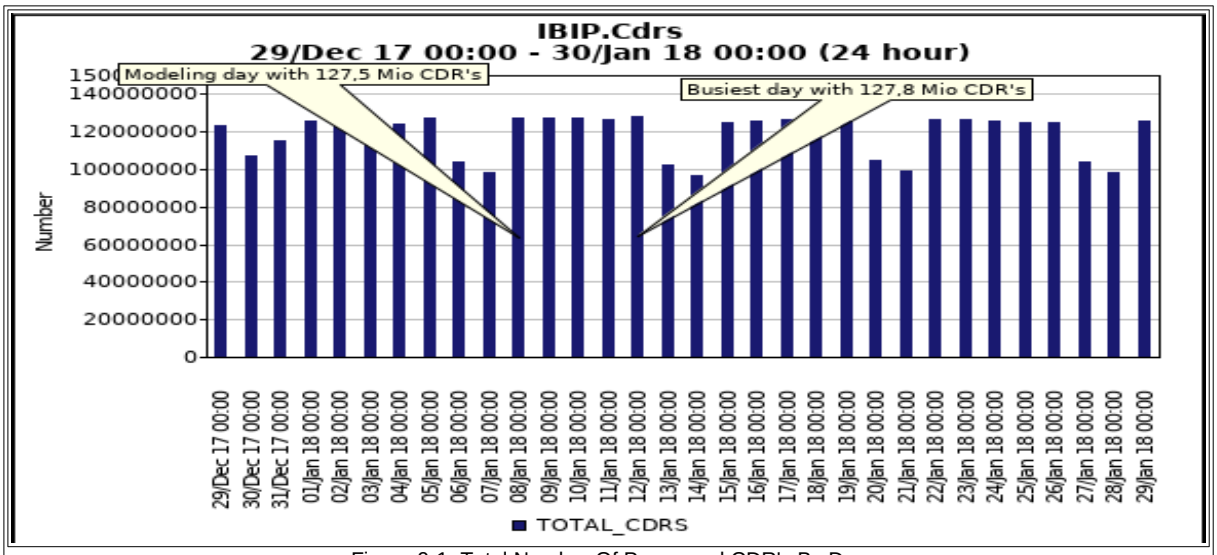

Figure 9-1: Total Number Of Processed CDR's By Day

Mr. X mentioned that an interesting period to evaluate further is Sunday 31th Dec/17 – Wednesday 10th Jan/18 as then both CDR activities being Rating & Bill Run are ongoing. Later there is only Rating activity.

Also during that period the VOICE rating does a SKIP rating without duration. This means that almost nothing happens for half of these CDR's. For LTE, MMS, SMS, Sip and STP this is a full rating.

The busiest day for the period Sunday 31th Dec/17 – Wednesday 10th Jan/18 was Monday the 8th Jan/18 where about 127,5 Mio CDR's where processed.

Time	TOTAL_CDRS
31/Dec 17 00:00	115142256.000
01/Jan 18 00:00	125989903.000
02/Jan 18 00:00	122599353.000
03/Jan 18 00:00	125452553.000
04/Jan 18 00:00	124384353.000
05/Jan 18 00:00	127184179.000
06/Jan 18 00:00	103884102.000
07/Jan 18 00:00	98325610.000
08/Jan 18 00:00	127573254.000
09/Jan 18 00:00	127409725.000
10/Jan 18 00:00	127391533.000

Table 9.1: CDR's By Day

So we will select that day and the busiest hour in order to do the calculations which is 15h.

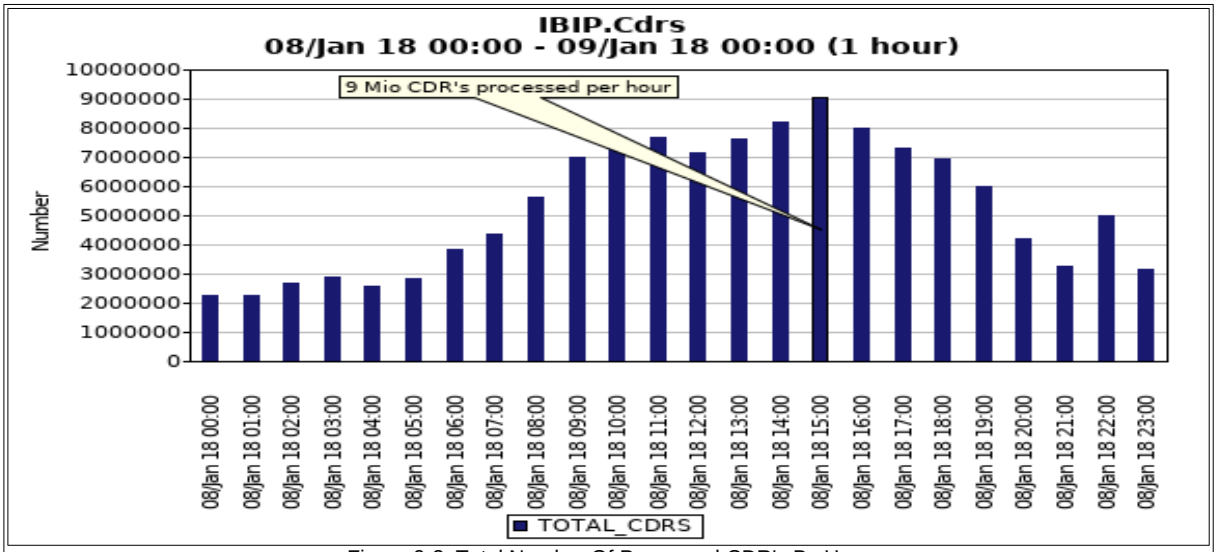

Figure 9-2: Total Number Of Processed CDR's By Hour

The CDR types executed are the following.

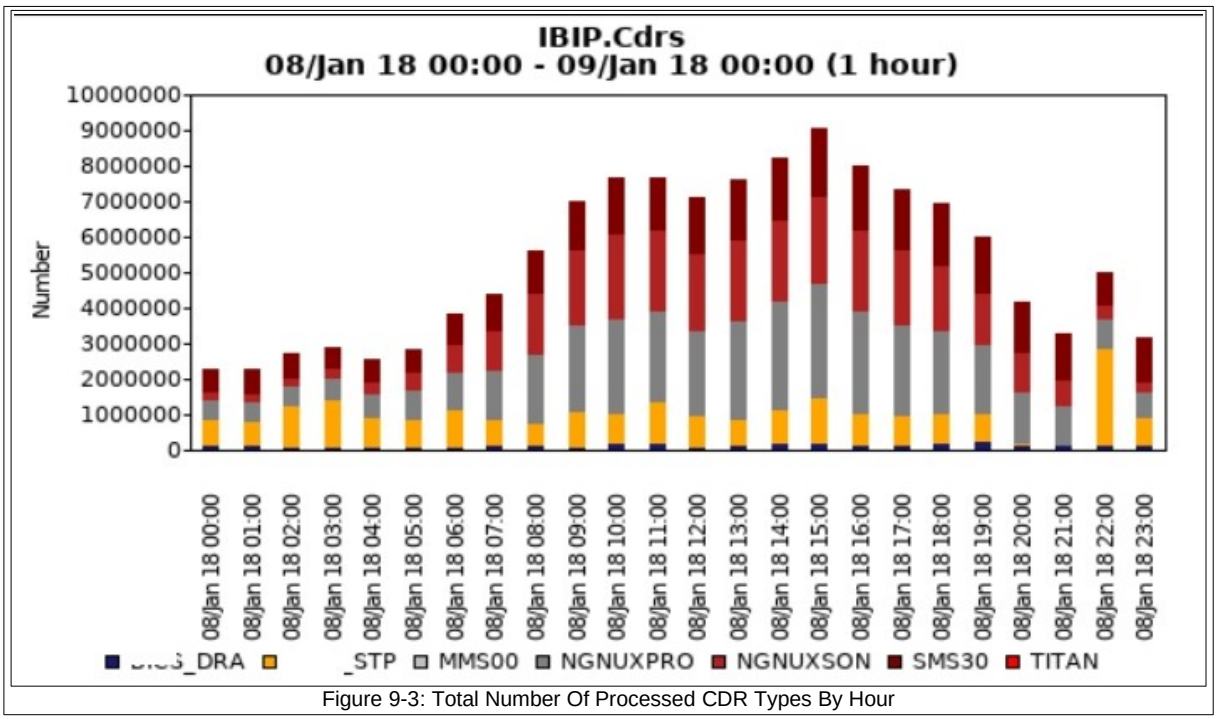

Figure 9-3: Total Number Of Processed CDR Types By Hour

On the 8th we processed 127.573.254 CDR's per day so a 5% increase represents 6.378.663 CDR's. That's a total of 133.951.917 CDR's for 2018. As agreed by Mr. X we will go for a compound increase for the next years.

We will only increase 5% on the business workloads for the supporting tools/software only 1%.

current	5.00%	2018	2019	2020	2021	2022
CDR Increase	Compound (Add 5% each year of the activity of the past year)					
127.573.254	6.378.663	133.951.917	140.649.513	147.681.988	155.066.088	162.819.392
		6.697.596	7.032.476	7.384.099	7.753.304	

Figure 9-4: CDR Compound Increase

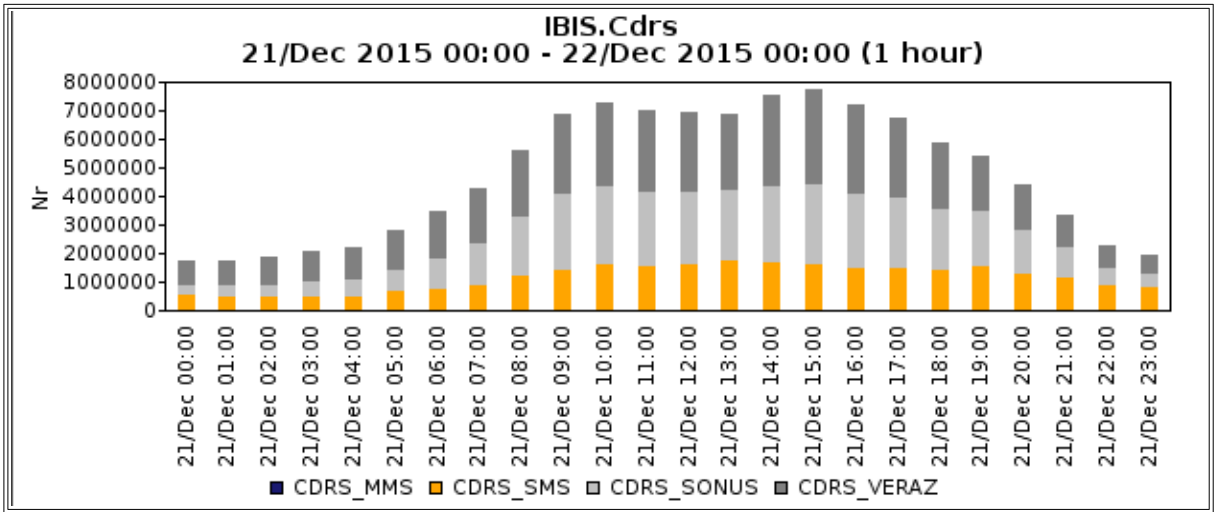

Figure 9-5: Total Number Of Processed CDR's By Hour On 21 Dec/15

The global CPU load by application (workloads) shows also the highest CPU consumption at that moment which is 54%. There is a difference with 1 hour, compared to the CDR time due to the way the data is logged with a user agent (end of time/begin of time) versus CMIS agent.

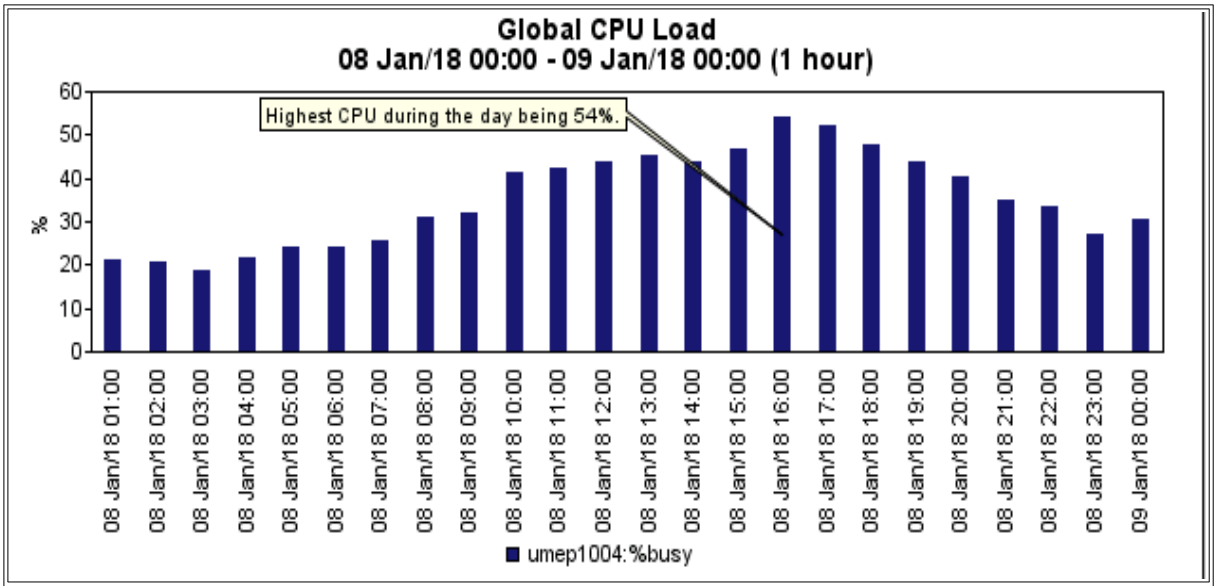

Figure 9-6: CPU Load By Application By Hour On 21 Dec/15

So we extract the 1h data of the 8th Jan/18 of 16h to do the calculations, the number of LCPU's active at that time was 70.

We adapt the model by consolidating Workloads and apply the correct hardware definitions.

System Name	Workload	Consolidate with Workload
umep1004	IBIS_afregdnc	IBIS_behldnc
umep1004	IBIS_behldnc	
umep1004	IBIS_dwh	
umep1004	IBIS_ibip	
umep1004	IBIS_ibipload	IBIS_behldnc
umep1004	IBIS_sumdnc	IBIS_behldnc
umep1004	IBIS_tetra2	IBIS_behldnc
umep1004	OTHER	
umep1004	SSO_maestro	SSO_root
umep1004	SSO_networker	SSO_root
umep1004	SSO_oracle	SSO_root
umep1004	SSO_root	
umep1004	sfmdb	
umep1004	teamquest	
umep1004	tivoli	

Figure 9-7: Model Workload Consolidation

System Name	Active Resource	Equipment Name	Equipment Type	Path	Active Attr
umep1004	CPU	HP ia64 hp Superdome2 16s	CPU		Servers = 70; Hyper-threading = On;
umep1004	fc1	Fibre-Channel 8Gb	Controller		Servers = 1; Discipline = FCFS
umep1004	c60t3d2	EMC SYMMETRIX	Disk Unit	fc1	Servers = 1; Discipline = FCFS
umep1004	fc3	Fibre-Channel 8Gb	Controller		Servers = 1; Discipline = FCFS
umep1004	disk104	EMC SYMMETRIX	Disk Unit	fc3	Servers = 1; Discipline = FCFS
umep1004	disk139	EMC SYMMETRIX	Disk Unit	fc3	Servers = 1; Discipline = FCFS
umep1004	fc0	Fibre-Channel 8Gb	Controller		Servers = 1; Discipline = FCFS
umep1004	disk105	EMC10Kdrive	Disk Unit	fc0	Servers = 1; Discipline = FCFS
umep1004	disk106	EMC10Kdrive	Disk Unit	fc0	Servers = 1; Discipline = FCFS
umep1004	disk107	EMC10Kdrive	Disk Unit	fc0	Servers = 1; Discipline = FCFS

Figure 9-8: Model Hardware Definitions

We have 53,9% measured versus 53,1% modeled. This is excellent so we use that model.

Measured AR%	Modeled AR%
53.984	53.134

Figure 9-9: Model Calibration

Now we will verify the Stretch Factor to see if any issue occurs already. Afterwards we will apply the growth on the workloads and check when and for what reason the server will crash having the current set-up.

First we will look at the Stretch Factor which is the ((Queue Time + Service Time) / Service Time), this will indicate if we are near the point of exponential degradation of performance. The critical point has the value of 2.

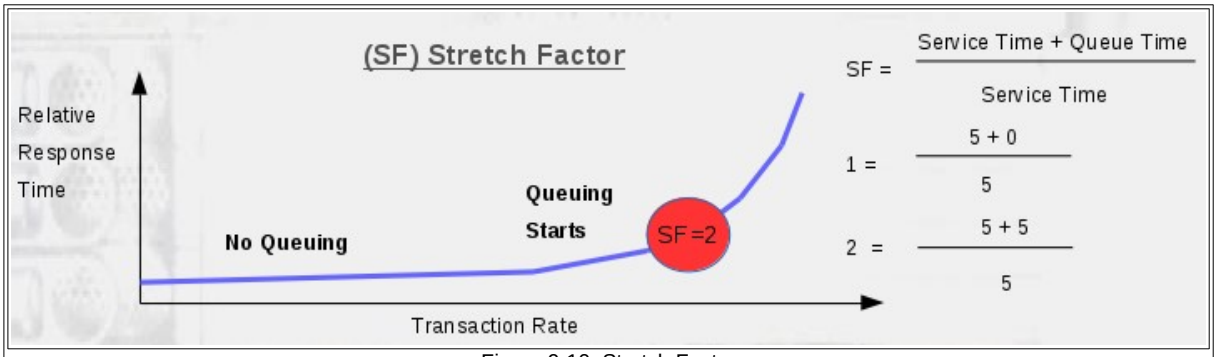

Figure 9-10: Stretch Factor

Seems we already have an issue for the “sfmdb” Workload as the SF is above 2.

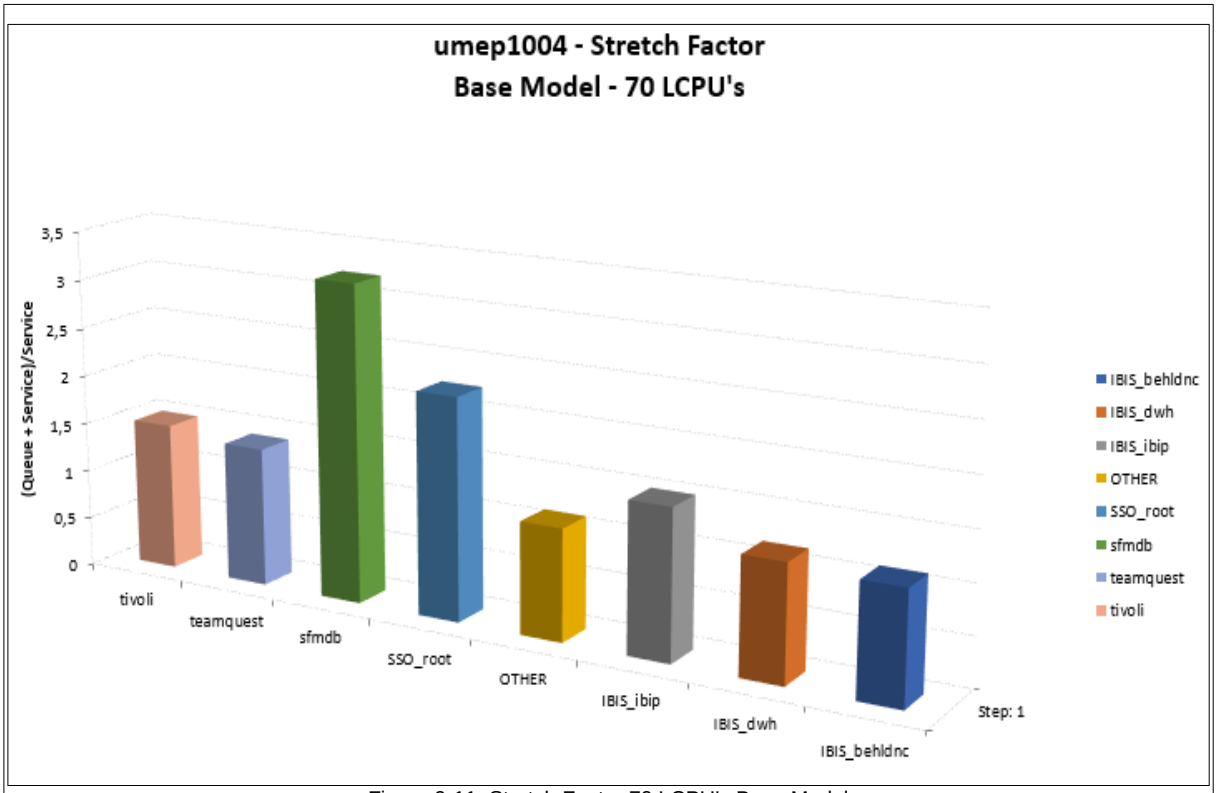

Figure 9-11: Stretch Factor 70 LCPU's Base Model

In fact this is due to the disk named “disk110” part of the volume group “vgibip_index”.

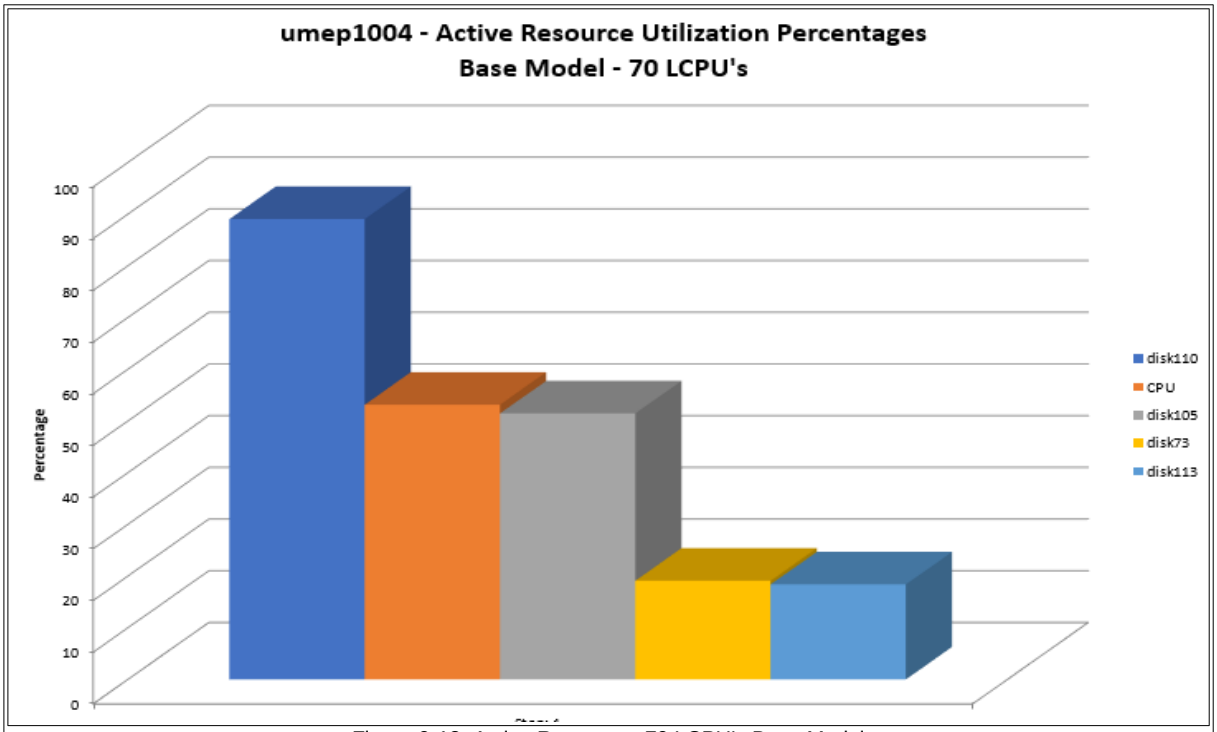

Figure 9-12: Active Resources 70 LCPU's Base Model

So we will offload this device first by adding a disk taking part of the activity.

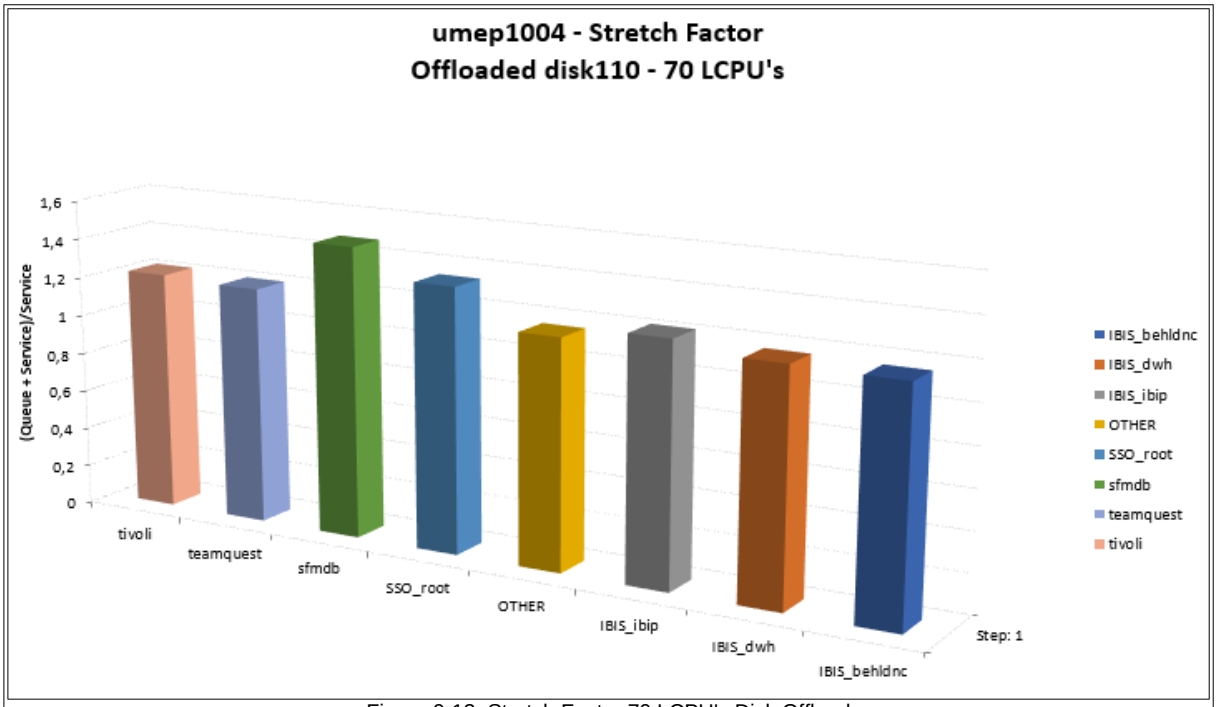

Figure 9-13: Stretch Factor 70 LCPU's Disk Offload

The corresponding impact on the active resources is shown below, as expected the activity of “disk110” dropped significantly.

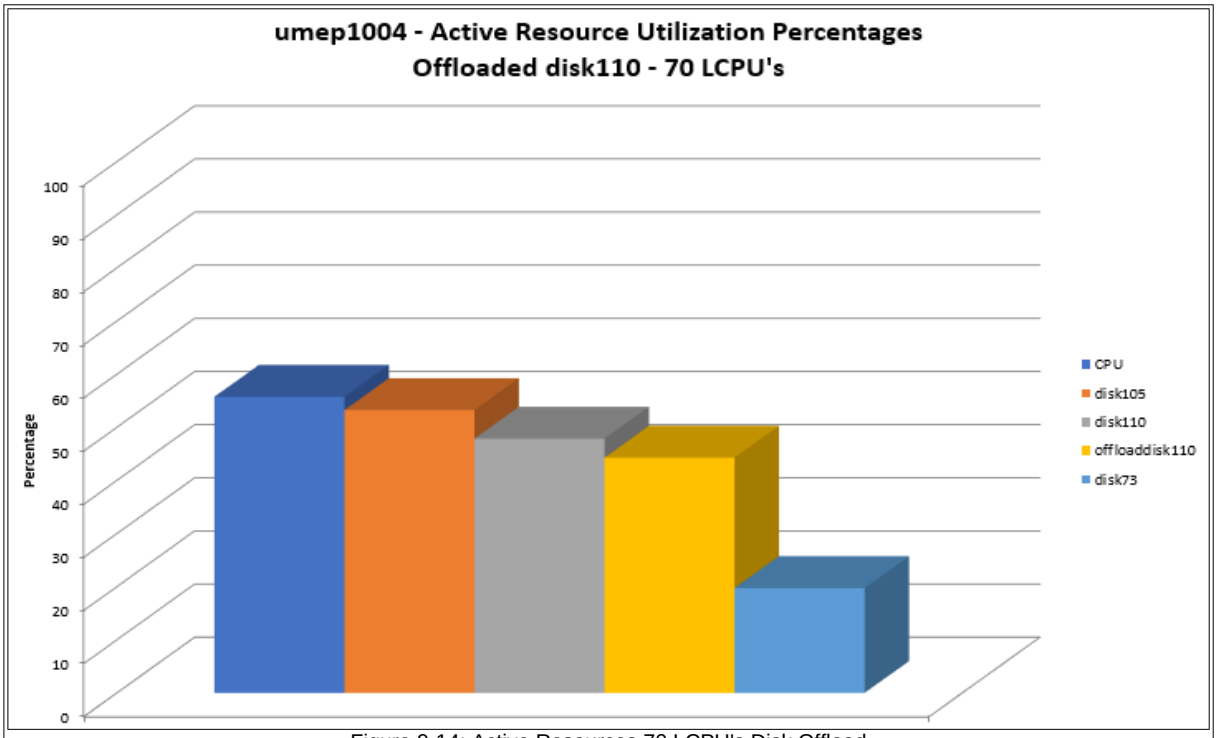

Figure 9-14: Active Resources 70 LCPU's Disk Offload

Finally we will simulate the 5% yearly compound increase on all the “IBIS_<>” , “OTHER” and “sfmdb” Workloads. The remaining Workloads have a 1% yearly compound increase.

The forecasted period is 5 years, starting from 2018 (included). The Stretch Factor tells us that by the end of 2021 the server starts to have issues.

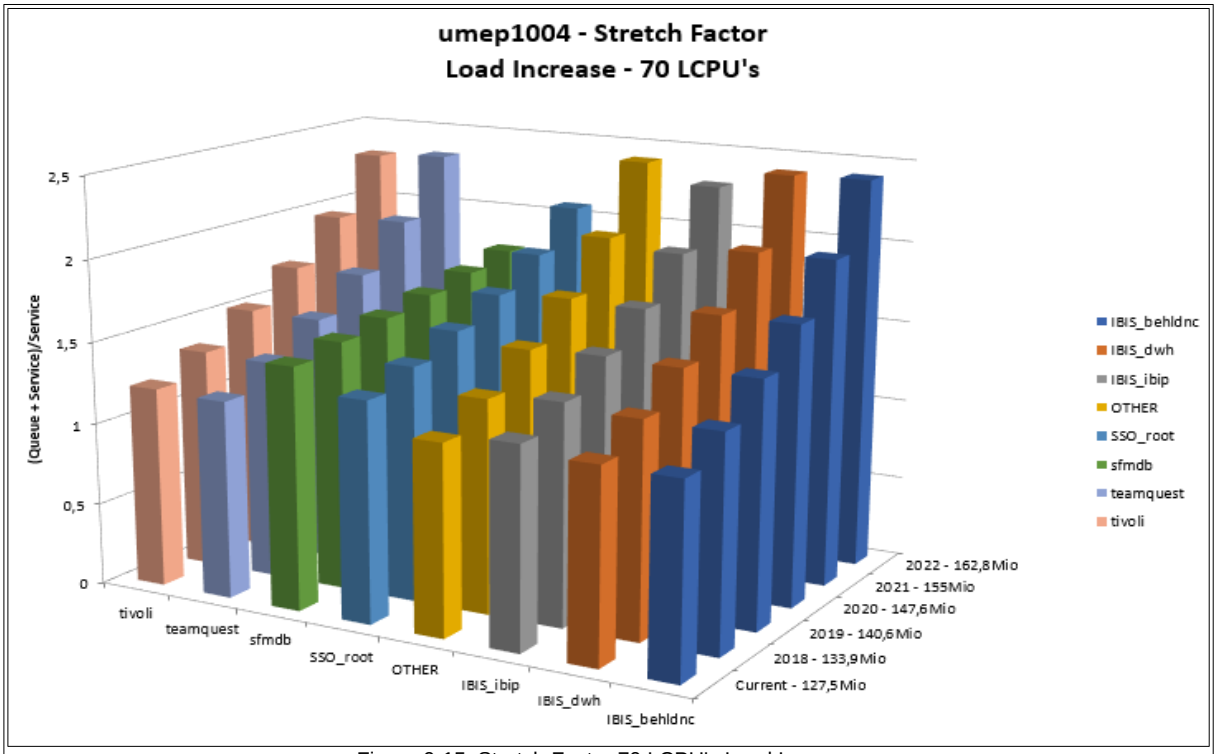

Figure 9-15: Stretch Factor 70 LCPU's Load Increase

The limiting resource is the CPU as we can see below.

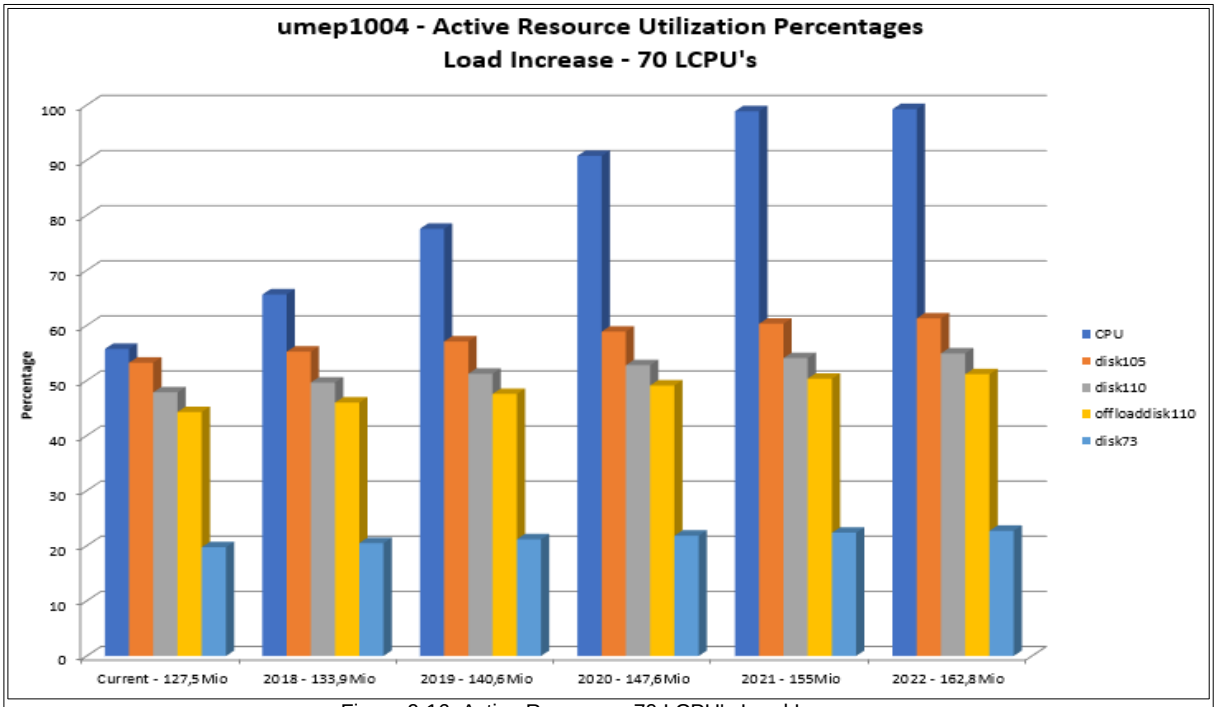

Figure 9-16: Active Resources 70 LCPU's Load Increase

We need two extra cores or 4 LCPU's to handle the increased load.

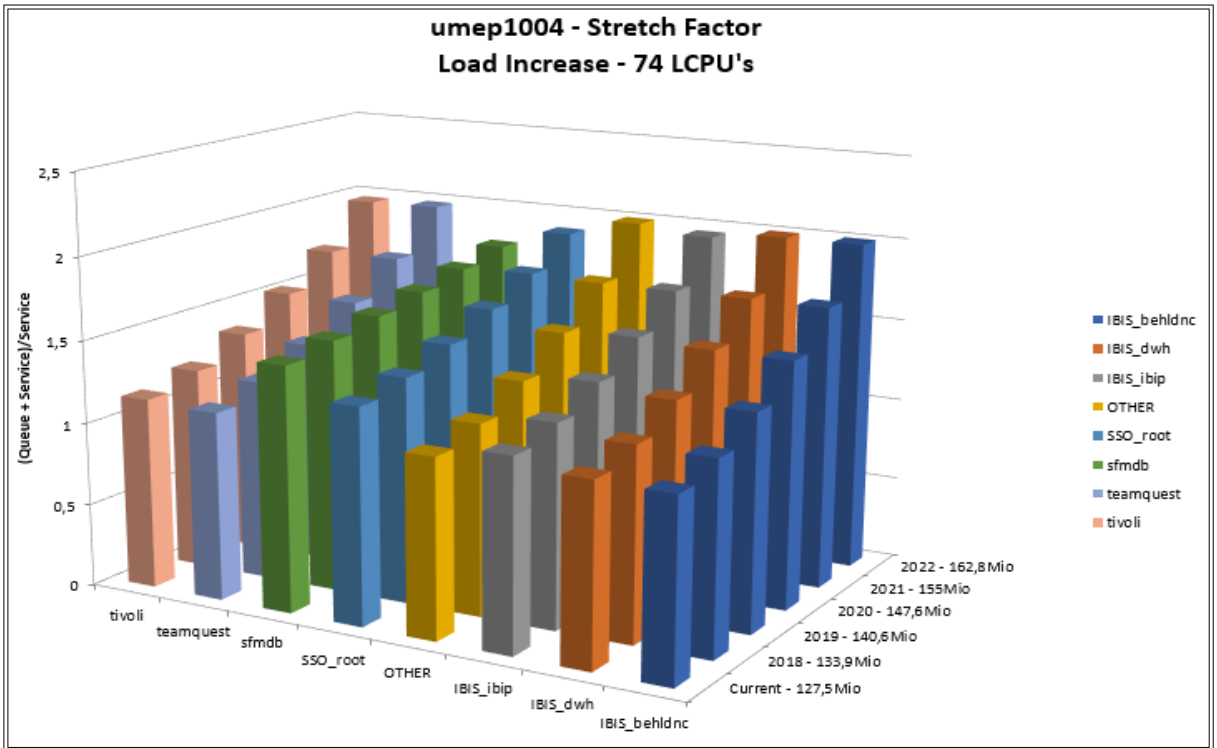

Figure 9-17: Stretch Factor 74 LCPU's Load Increase

-END-