

DASD Performance

- Barton@VelocitySoftware.com
- [HTTP://VelocitySoftware.com](http://VelocitySoftware.com)

“If you can’t Measure it,
I am Just Not Interested™”

DASD Performance Topics

- **DASD Reports and Analysis**
- **Disk Performance – is it a problem?**
- **DASD I/O Overview, I/O Concepts**
- **Cache Controllers**
- **Tuning MDC**
- **Tuning DASD Cache**
- **Seek Analysis**
- **Other Functions**
- **Storage Processors**

When is DASD Performance an Issue?

- When the workload is impacted
- When there is paging
- When there is contention for disk
- When there are high I/O workloads

Disk Performance History

Assumptions and “Rules of Thumb” based on technology:

DASD Assumptions on what is “good”:

- 3350 (1978): 30ms
- 3380 (1983): 20ms
- 3380J/K (1986): 18ms
- 3390 (1990): 15ms
- IBM Drive (2012): 73gb/5ms
- 600gb – 15,000rpm – 3.5ms seek – 2ms rotation

DASD Cached Assumptions:

- 3380-23 (1982) 3380: 15ms
- 3390-1 (1988) 3390 gray cable: 10ms
- 3390-3 (1995) ESCON: 5ms
- 3390-9 (2001) FICON: 2ms
- “3390” (2009) **sub 1ms**

Assumptions keep changing – current SSD is sub-ms

When there is poor DASD performance:

- Check high device utilization (queueing on a device)
- Check poor cache
- Check high write activity
- Check queueing on Non-Volatile Storage (NVS)
- Check channel contention (queueing on channels)
- Check SAP contention

Hardware Subchannel Measurement Block

- Pend
- Connect
- Disc

Cache statistics (storage controller)

Seek Analysis (CP I/O Subsystem)

- Seek by cylinder
- Seek by minidisk
- Seek by user

Channels

- VM sample, hardware measurements
- FICON

DASD Performance Report: ESADSD2

Hardware instrumentation provides metrics:

- Service components
- Rates
- Calculate busy

Report: **ESADSD2** DASD Performance Analysis Velocity Software ESAMAP
Monitor initialized: 05/22/08 at 14:00:00 on 2084 serial 15BAF First record analyzed: 05/22/08 14:0

<-----DASD Response times (ms)----->

Dev	Device	<--SSCH-->	<%DevBusy>	<SSCH/sec->				<--Service times-->	<--Queueing->	QLengths								
No.	Serial	Type	Total	ERP	Avg	Peak	avg	peak	Resp	Serv	Pend	Disc	Conn	DASD	Cntl	THR	avg	max

14:01:00																		
Top DASD by Device busy																		
0C51	VME090	3390-3	8795	0	8.5	8.5	146.6	146.6	0.6	0.6	0.2	0.0	0.3	0	0.0	0	0	0
0B6E	VME106	3390-3	7858	0	8.4	8.4	131.0	131.0	0.6	0.6	0.2	0.0	0.4	0	0	0	0	0
0E78	VME089	3390-3	4186	0	8.1	8.1	69.8	69.8	1.2	1.2	0.3	0.4	0.5	0	0	0	0	0
0E71	VME067	3390-3	989	0	7.5	7.5	16.5	16.5	4.6	4.6	0.4	3.3	0.9	0	0	0	0	0
0***End Top DASD by Device busy***																		

ECKD Service Time Components

Pend Time (Time to start the I/O)

- CPU channel connection
- Protocol (I/O / SAP processor)

Disconnect Time (Hardware delays)

- Rotational delay (not in the cache)
- Seek time (not in the cache)
- Internal controller contention

Connect Time (Channel time, disk rotation)

- Transmission time
- Data size/data rate

Device Queue Time

DASD Service Time	Pend + Disc + Conn
Device Busy	Rate * Service Time
<u>DASD Response Time (MM1)</u>	(Service Time) / (1 – Device Busy)
Queue Time	Response Time – Service Time

Disconnect Time (non-cache)

- Seek + Rotational Delay + RPS Misses

Disconnect Time (cache)

- $(\text{Seek} + \text{Rotational Delay} + \text{RPS Misses}) * (\text{CacheMiss} / \text{CacheTotal})$

Cache – Reduce Disconnect time

- SSD was the original (non-IBM) cached paging device
- Read only - 3880-23 (1980)
- Write with non-volatile cache - 8mb (1985)
- Large write cache - 1gb (1990 – EMC)
- On board cache (device level – Storage Tech)
- RAID (meant for reliability, but slow)
- SSD is current technology

Connect Time – Data transmission

- Transmission time (Connect time)

If Channel Busy is high:

- Cause: High data transmission, large blocks
- Cause: Slow channels
- Delay: Queueing on the channel (I/O Processor/SAP)

Channels – Reduce Connect time

- VM/XA Multiple channels - can reconnect on any channel
- Grey Cables/3390 - 4.5mb/sec (up from 3mb/sec)
- ESCON – 17mb/sec (1995)
- FICON – 110mb/sec (2003)
- FICON Express – 200/400 mb/sec (2005/2007)

Queue Time Types:

- Device Busy Contention (queue time is calculated)
- Channel Path Contention – Control Unit/Channel/SAP
- Throttle – operator command

Throttle: Only allow x number of I/Os per second

- Meant for shared DASD between LPAR situations
- Reduce activity to control contention for other systems
- Invented for z/OS shared disks to allow slower z/OS to IPL

PAV (Parallel Access Volumes) – Reduce

- Multiple definitions of device on multiple channels
- Data available from cache even when device is busy
- HyperPav (dynamic multi-channels)
- FCX...

Why Have Multiple Channels to a Device?

VM/XA allowed multiple channels to access the same device to reduce RPS MISS:

- If the 1st channel is busy, try the 2nd, and so on...
- Architecture supports up to 16 channels
- Four channels was considered minimum

Some controllers no longer support dynamic reconnect

Multiple channels now for PAV - Availability

The Reason for Dynamic Reconnect

RPS MISS (3380, XA)

- One channel: $17\text{ms} * (1 / (1 - \text{ChanUtil} * \text{ChanUtil}))$

RPS MISS Times – 20% Channel Utilization

- Two channels: $17\text{ms} * (1 / (1 - .04)) - 17 = .4\text{ms}$

Channel Utilization can be higher when more channels (used to matter)

Channel	RPS MISS TIME (3380)		
	1 Chan	2 Chan	4 Chan
BUSy			
20%	4ms	.7ms	0ms
40%	11ms	3ms	.4ms
60%	25ms	9ms	2ms
80%	68ms	30ms	12ms

Block sizes vary and impact assumptions

- 4k I/O
- MDC does full track I/O – higher connect times
- Linux does 4k I/O
- Linux can chain 1024 I/O
- Paging may be up to 40 * 4k pages

Channel Measurements

Report: **ESACHAN**

Monitor initialized:

```
-----  
                <Pct Channel>  
Time/      Utilization  
CHPID      LPAR Total  Shrd  
-----  
16:26:00  
26         28.3  10.0   No  
27         27.7  20.0   No  
2E         27.8  50.0   No  
42          6.4  20.0   No  
60         12.3  10.0   No  
64         11.9  10.0   No  
65         10.1  10.0   No  
C2         11.7  20.0   No  
D3          2.1  20.0   No  
D6         28.2  40.0   No  
D7         28.3  30.0   No  
DD         28.1  20.0   No  
DE         28.1  40.0   No  
E8         10.1  10.0   No  
E9         11.2  10.0   No  
EC         10.2    0   No  
ED         10.8  20.0   No  
-----
```

System: 337 370.0

Two measurements:

- LPAR (measured)
- TOTAL (physical channels)

Old Style:

- Sampled by CP (TestSubChannel)
- Measured by channel

Utilization sampled at Hi-Frequency rate

- Example is 6 second example
- The LPAR is “us”

Channel Measurements

Report: **ESACHAN** Channel Performance Analysis Velocit
Monitor initialized: 06/30/03 at 13:21:34 on 2064 serial 5146B First r

Time/ CHPID	<Pct Channel>		Shrd	Chanl Type	<-----Data Units ----->							
	Utilization				<---Reads/Second-->		<---Writes/Second-->					
	LPAR	Total			LPAR	TOTAL	Pct	Max	LPAR	TOTAL	pct	MAX
13:31:53												
16	5.1	14.4	Yes	ESCON								
18	9.9	18.3	Yes	ESCON								
20	3.9	19.4	Yes	ESCON								
25	5.3	12.2	Yes	ESCON								
27	10.3	10.6	Yes	ESCON								
30	3.7	16.1	Yes	ESCON								
34	5.3	12.8	Yes	ESCON								
36	10.0	14.4	Yes	ESCON								
43	5.1	15.6	Yes	ESCON								
45	10.4	15.6	Yes	ESCON								
52	5.6	20.0	Yes	ESCON								
53	1.9	31.1	Yes	ESCON								
54	10.0	14.4	Yes	ESCON								
63	10.4	14.4	Yes	ESCON								
70	5.3	10.6	Yes	ESCON								
71	.	11.7	Yes	ESCON								
72	9.8	15.0	Yes	ESCON								
77	18.1	21.7	Yes	ESCON								
7F	5.3	18.9	Yes	ESCON								
90	10.2	15.0	Yes	ESCON								
98	4.1	15.0	Yes	ESCON								
9D	53.6	56.7	Yes	ESCON								

Two measurements:

- LPAR (measured)
- “Total” (physical channels)

Newer Style:

- Sampled by CP (“Total”)
- Measured by Channel (LPAR)
- What is Channel Utilization?
- (EMIF)

DASD Performance – Is there a problem?

z/VM Virtual Machine SIO and Asynch I/O Wait

- Rule of thumb? Greater than CPU Wait?

```
Report: ESAXACT Transaction Delay Analysis
Monitor initialized: 05/31/16 at 00:00:00ial 2F5A7
-----
<-----Percent non-dormant (Wait states)
UserID <-Samples-> . Tst <Asynch>
/Class Total In Q Run Sim CPU SIO Pag . CF Idl I/O Pag Ldg
-----
00:15:00 1101 1987 11 0 35 0 1.2 . 0 52 0 . .
Hi-Freq: 143K 120K 11 0.1 29 0.1 1.2 . 0 57 0.0 0.6 1.2
***Key User Analysis ***
TCPIP 900 198 1.0 2.5 5.6 0 0 . 0 91 0 0 0
***User Class Analysis***
Servers 8100 1 0 0 0 0 0 0 . 0 100 0 0 0
Velocity 9358 1479 3.0 0.5 3.2 5.2 0.6 . 0 87 0 0 0
LDXmno 9000 9000 13 0.2 40 0 1.3 . 0 45 0.0 0.2 0.1
LDXecs 14400 14400 24 0.2 51 0 0.2 . 0 24 0 0.1 0.1
LDXcii 1800 1800 7.7 0 23 0 33 . 0 27 0.1 0.1 8.6
LDXfrm 7200 7200 3.4 0.0 33 0 0.6 . 0 62 0 0.3 0.6
LDXlfe 8100 8100 13 0.1 31 0 1.2 . 0 53 0 1.2 1.2
TheUsrs 20700 13994 4.4 0.1 21 0 0.7 . 0 71 0.0 0.5 2.1
```

DASD Performance – Is it a problem?

Linux also measures I/O Wait time - ESALNXS

- Rule of thumb? Comparable to steal? Total CPU?

```

Report: ESALNXS          LINUX VSI System Analysis Report          Velocity Sof
Monitor initialized: 05/30/17 at 02:00:00 on 2827 serial 0168A7      First record
-----
Node/      <---Load Numbers--> CPU <Processor Pct Util> NICE <CPU Overhead%> IO
Time      Users Procs MaxProc NBR Total Syst User Idle Time Krnl IRQ  Steal Wait
-----
02:01:00
L21EP      0    448      0 Tot  27.2 13.2  9.0 38.1    0  0.7  4.3 127.3  7.4
           1    16.5  7.8  5.1  4.2    0  0.4  3.2  76.7  2.6
           2    10.7  5.4  3.9 33.9    0  0.3  1.1  50.6  4.7
L212P      0    252      0 Tot  10.9  3.6  3.6 75.0    0  0.5  3.3   9.1  5.0
L218P      0    552      0 Tot  18.3  4.1 13.3 20.8    0  0.4  0.4  58.1  2.8
L22AP      0    440      0 Tot  18.3  4.6 12.6  114    0  0.6  0.5  63.1  5.0
           1    11.4  2.6  8.0 45.0    0  0.3  0.4  40.4  3.2
           2     7.0  2.0  4.6 68.7    0  0.3  0.1  22.6  1.7
L220P      0    248      0 Tot  31.5 10.8 14.2 23.5    0  0.4  6.2 140.9  4.1
           1    16.4  6.1  5.2  6.4    0  0.3  4.8  76.4  0.8
           2    15.1  4.7  9.0 17.1    0  0.1  1.4  64.5  3.3
L21EP      0    446      0 Tot  24.6 14.1  3.4 55.0    0  0.8  6.2 114.7  5.7
           1    13.0  6.3  1.7 24.3    0  0.4  4.5  60.5  2.3
           2    11.6  7.9  1.7 30.7    0  0.3  1.7  54.3  3.4
L258P      0    244      0 Tot  21.0  5.0 13.4  0.2    0  0.4  2.1  77.9  0.9
    
```

z/VM DASD Subsystem Reports

DASD

- ESADSD1 – Configuration
- **ESADSD2** – Disk Performance
- ESADSD6 – Disk Performance Extended

DASD Cache

- ESADSDC – DASD Cache Configuration
- ESADSD5 – DASD Cache Performance

Seek Analysis

- ESASEEK – Seeks by Minidisk
- ESAUSEK – Seeks by User

Channels

- ESACHNC – Channel Configuration
- **ESACHAN** – Channel Performance
- ESACHNH – HiperSockets
- ESAQDIO – QDIO 1
- ESAQDIO2 – QDIO 2
- ESAIOP – I/O Processor (SAP)

DASD Configuration Report – ESADSD1

Monitor configuration data needed for device number

- SysID is an internal value assigned by the system
- DevNo is what is externalized
- If DevNo = SysID then the configuration data is not captured
- Check channel paths online

Report: **ESADSD1** DASD Configuration
Monitor initialized: 05/06/08 at 12:00:00 on 2094 serial AEA7D
Monitor period: 3600 seconds (1:00:00)

```
-----  
Dev Sys      Device      <CHPIDS OnLn><-Cntrl Unit-> UserID   MDisk  
No. ID      Serial Type      SHR 01 02 03 04 OBR/CU Model   (if ded) Links  
-----  
C557 6177 LNXP5F 3390-9 YES 48 49 76 87 35/00 2105           0  
          95 96 D6 5A  
C558 6178 LNXP55 3390-9 YES 48 49 76 87 35/00 2105           0  
          95 96 D6 5A
```

DASD Configuration Report – ESADSD1

DASD Page/Spool Extents, PAV, Cache Status/STGID

Report: **ESADSD1** DASD Configuration
Monitor period: 3600 seconds (1:00:00)

```
-----  
Dev Sys      Device . <----Extent---->  
No. ID      Serial Type . Type Start Size  
-----  
C557 6177 LNXP5F 3390-9 . Page      0 10K  
.  
C558 6178 LNXP55 3390-9 . Page      0 10K
```

Report: **ESADSD1** DASD Configuration
Monitor period: 3600 seconds (05/06/08 13:00:00)

```
-----  
Dev Sys      Device      <. <--MDC Status--> <3990 Cache Status> PAV  
No. ID      Serial Type  SHR . Elig Def Now Shr Actv DFW NVS STGID Base  
-----  
C557 6177 LNXP5F 3390-9 YES . Yes Off . Yes Yes On On 9C50  
.  
C558 6178 LNXP55 3390-9 YES . Yes Off . Yes Yes On On 9C50 .
```

DASD Performance Report – ESADSD2

Performance data provides:

- Interval Top DASD – Look for high device utilization
- Interval Control Unit – Look for channel/cache issues
- Summary – everything. Is Response Time good?

Report: **ESADSD2** DASD Performance Analysis Velocity Software ESAMAP
 Monitor initialized: 05/22/08 at 14:00:00 on 2084 serial 15BAF First record analyzed: 05/22/08 14:0

```

-----<-----DASD Response times (ms)----->
Dev      Device <--SSCH--> <%DevBusy> <SSCH/sec->      <--Service times--> <--Queueing-> QLengths
No. Serial Type  Total  ERP  Avg  Peak  avg  peak  Resp  Serv Pend Disc Conn DASD Cntl THR  avg max
-----
14:01:00
***Top DASD by Device busy***
0C51 VME090 3390-3  8795  0  8.5  8.5 146.6 146.6  0.6  0.6  0.2  0.0  0.3  0  0.0  0  0  0
0B6E VME106 3390-3  7858  0  8.4  8.4 131.0 131.0  0.6  0.6  0.2  0.0  0.4  0  0  0  0  0
0E78 VME089 3390-3  4186  0  8.1  8.1  69.8  69.8  1.2  1.2  0.3  0.4  0.5  0  0  0  0  0
0E71 VME067 3390-3   989  0  7.5  7.5  16.5  16.5  4.6  4.6  0.4  3.3  0.9  0  0  0  0  0
***End Top DASD by Device busy***

0509 Control Unit    4166  0  0.1  0.1  69.4  69.4  1.2  1.2  0.3  0.2  0.6  0  0  0  0  0
0B00 Control Unit    8716  0  0.1  0.1 145.3 145.3  0.7  0.7  0.3  0.1  0.4  0  0  0  0  0
-----
System:              32137  0  0.1  0.1 535.6 535.6  1.0  1.0  0.3  0.3  0.5  0  0.0  0  0  0
  
```


DASD Cache Performance Report – ESADSD5

Cache statistics provides:

- Read/Hit percentage – depends on cache size and locality of reference
- Write/Hit percentage – depends on NVS (Non-Volatile Storage – persistent memory)
- NVS Full, Cache Inhibit and Cache Bypass – not good

Report: **ESADSD5** 3990-3 Cache Analysis Velocity Software ESAMAP 3.7.4 05
 Monitor initialized: 05/22/08 at 14:00:00 on 2084 serial 15BAF First record analyzed: 05/22/08 14:00:00

```

-----
Pct. <-----per second-----> <-----Write activity per second----->
Dev   Actv <-----Total-----> <----Read----> <--Seq Read--> Total DFW DFW SEQ NVS <--Cache-->
No.  Serial Samp I/O Hits Hit% Read% I/O Hits Hit% I/O Hits Hit% I/O I/O Hits I/O Hit% Full Inhib Bypass
-----

```

14:01:00

Top DASD by Device busy

Dev	Serial	Samp	I/O	Hits	Hit%	Read%	I/O	Hits	Hit%	I/O	Hits	Hit%	Total	DFW	DFW	SEQ	NVS	Cache	Full	Inhib	Bypass
OC51	VME090	100	172	172	100	100.0	64.6	64.6	100	108	108	100	0	0	0	0	0	0	0	0	0
OB6E	VME106	100	131	131	100	100.0	100	100	100	31.3	31.3	100	0	0	0	0	0	0	0	0	0
OE78	VME089	100	74.0	72.2	97.6	97.9	58.4	56.7	97.1	14.0	13.9	99.3	1.5	1.5	1.5	1.5	100	0	0	0	

End Top DASD by Device busy

0509	CtlUnit	100	1106	1052	95.1	92.9	283	231	81.5	744	743	100	78.5	77.8	77.8	23.0	99.1	0	0	0
OB00	CtlUnit	100	518	489	94.5	94.9	341	313	91.7	150	150	100	26.6	26.2	26.2	1.8	98.5	0	0	0
OC00	CtlUnit	100	233	230	98.7	91.7	88.5	88.5	100	125	124	99.1	19.4	17.4	17.4	4.8	89.9	0	0	0
OD00	CtlUnit	100	127	110	86.6	64.8	80.9	65.0	80.3	1.5	1.4	96.6	44.7	43.7	43.7	3.2	97.7	0	0	0
OE00	CtlUnit	100	316	294	93.1	48.8	139	125	90.3	15.6	15.5	99.2	161.7	153	153	103	94.8	0	0	0

System:		100	2332	2206	94.6	84.8	941	831	88.3	1036	1034	100	355.1	341	341	152	96.0	0	0	0
---------	--	-----	------	------	------	------	-----	-----	------	------	------	-----	-------	-----	-----	-----	------	---	---	---

Channel Utilization is not obvious

- One fibre for read
- One fibre for write
- Data unit processing
- Channel bus

Channel balancing takes planning!

FICON Channel Measurements

Channel Metrics:

- Sampled by CP (“Total”)
- Measured by Channel – Data/second (“LPAR”)
- What is Channel Utilization?

Report: **ESACHAN** Channel Performance Analysis Linux T
 Monitor initialized: 05/09/03 at 10:53:54 on 2064 serial 31539 First r

Time/ CHPID	<Pct Channel> Utilization		Shrd	Chanl Type	<-----Data Units -----> <---Reads/Second---> <---Writes/Second--->								
	LPAR	Total			LPAR	TOTAL	Pct	Max	LPAR	TOTAL	pct	MAX	
10:54:33													
02	0.2	3.0	Yes	FICON	0	0	0	82K	0	0	0	82K	
03	0.2	3.0	Yes	FICON	0	0	0	82K	0	0	0	82K	
04	.	3.0	Yes	FICON	0	0	0	82K	0	0	0	82K	
C0	.	12.2	Yes	FICON	0	0	0	82K	0	0	0	82K	
C1	.	12.3	Yes	FICON	0	0	0	82K	0	0	0	82K	
C2	.	12.2	Yes	FICON	0	0	0	82K	0	0	0	82K	
C3	.	12.3	Yes	FICON	0	0	0	82K	0	0	0	82K	
C9	12.9	12.9	Yes	FICON	28K	27550	24	117K	55	55	0.0	117K	
CC	.	12.2	Yes	FICON	0	0	0	82K	0	0	0	82K	
CD	.	12.2	Yes	FICON	0	0	0	82K	0	0	0	82K	
CE	.	12.2	Yes	FICON	0	0	0	82K	0	0	0	82K	
CF	.	12.2	Yes	FICON	0	0	0	82K	0	0	0	82K	
F4	1.3	3.9	Yes	FICON	0	0	0	82K	0	0	0	82K	
F5	3.7	3.9	Yes	FICON	1	1	0.0	82K	1	1	0.0	82K	

System:	20.8	142.5											

FICON Channel Measurements

New Style:

- Sampled by CP (“Total”)
- Measured by Channel – Work per second
- What is Channel Utilization? The Peak of any component

Report: **ESACHAN** Channel
 Monitor initialized: 05/09/03

Time/ CHPID	<Pct Channel> Utilization		Shrd	Chanl Type	<----Work Unit----> <--Rates / Second-->				<-Bus Cycles> <-per Second>			Bytes /Data
	LPAR	Total			LPAR	TOTAL	Pct	MAX	Used	pct	Max	Unit
10:54:33												
02	0.2	3.0	Yes	FICON	33	495	3.0	16K	1376	8.5	16K	1024
03	0.2	3.0	Yes	FICON	31	495	3.0	16K	1290	7.9	16K	1024
04	.	3.0	Yes	FICON	0	495	3.0	16K	1376	8.5	16K	1024
C0	.	12.2	Yes	FICON	0	1993	12	16K	1440	8.8	16K	1024
C1	.	12.3	Yes	FICON	0	1994	12	16K	1440	8.8	16K	1024
C2	.	12.2	Yes	FICON	0	1992	12	16K	1441	8.9	16K	1024
C3	.	12.3	Yes	FICON	0	1994	12	16K	1439	8.8	16K	1024
C9	12.9	12.9	Yes	FICON	17K	16811	13	130K	10K	16	65K	1024
CC	.	12.2	Yes	FICON	0	1993	12	16K	1345	8.3	16K	1024
CD	.	12.2	Yes	FICON	0	1993	12	16K	1440	8.9	16K	1024
CE	.	12.2	Yes	FICON	0	1993	12	16K	1344	8.3	16K	1024
CF	.	12.2	Yes	FICON	0	1993	12	16K	1440	8.9	16K	1024
F4	1.3	3.9	Yes	FICON	212	636	3.9	16K	2464	15	16K	1024
F5	3.7	3.9	Yes	FICON	596	638	3.9	16K	2304	14	16K	1024
System:	20.8	142.5										

FICON Channel Measurements

New Style:

- Channels are defined

Report: **ESACHAN** Channel Performance Analysis Linux

Time/ CHPID	<Pct Channel>			Channel Class/Type	<---Data Units--->				<---Writes/Sec	
	Utilization	Shrd	Channel		<---Reads/Second---	Max	LPAR	TOTAL	p	
	LPAR	Total			LPAR	TOTAL	Pct		LPAR	TOTAL
12:01:00										
48	0.3	8.5	Yes	FICON/FCS	36	7872	6.7	117K	294	2228
49	0.3	8.4	Yes	FICON/FCS	35	7827	6.7	117K	315	2127
5A	0.4	11.6	Yes	FICON/FCS	53	10570	5.4	195K	441	3312
76	0.3	8.5	Yes	FICON/FCS	30	7984	6.8	117K	267	2284
87	0.3	8.4	Yes	FICON/FCS	32	7568	6.5	117K	285	2078
95	0.5	12.6	Yes	FICON/FCS	57	11393	5.8	195K	455	3583
96	0.5	12.4	Yes	FICON/FCS	58	11495	5.9	195K	471	3466
D6	0.4	11.9	Yes	FICON/FCS	52	11192	5.7	195K	471	3303

Report: **ESADSD1** DASD Configuration Linux test

Dev No.	Sys ID	Serial	Device Type	<CHPIDS SHR	OnLn	<-Cntrl Unit-->	UserID (if ded)	Mdisk Links	<---Extent--->	Start	Size
				01 02 03 04		OBR/CU Model			Type		
C557	6177	LNXP5F	3390-9	YES	48 49 76 87	35/00 2105		0	Page	0	10K
					95 96 D6 5A						
C558	6178	LNXP55	3390-9	YES	48 49 76 87	35/00 2105		0	Page	0	10K
					95 96 D6 5A						
C559	6179	LNXP56	3390-9	YES	48 49 76 87	35/00 2105		0	Page	0	10K
					95 96 D6 5A						
C55A	617A	LXPC22	3390-9	YES	48 49 76 87	35/00 2105		0	.	.	.
					95 96 D6 5A						
C55B	617B	LNXP57	3390-9	YES	48 49 76 87	35/00 2105		0	Page	0	10K
					95 96 D6 5A						
C55C	617C	LNXP58	3390-9	YES	48 49 76 87	35/00 2105		0	Page	0	10K
					95 96 D6 5A						

Channel Types – z/VM 5.2

<u>class</u>	<u>type</u>	<u>description</u>
CTC	CTC	Channel-to-channel
CTC	CTP	Channel-to-channel point-to-point
CTC	CTS	Channel-to-channel swt-point-to-point
ESCON	CBY	Fiber extended (byte pacer)
ESCON	CNC	Serial-channel-path
ESCON	CNP	Serial-point-to-point channel
ESCON	CNS	Serial-switched-point-to-point channel
ESCON	CVC	Fiber extended (block pacer)
ESCON	DSD	Direct-system-device channel
ESCON	EIO	Emulated I/O channel
ESCON	FCV	Fibre-channel converted channel FCV
ESCON	ISD	Internal-system-device-channel
FCP	FCP	Fibre-channel-protocol channel FCP
FICON	FC	Fibre-channel point-to-point channel FC
FICON	FC?	Fibre channel
FICON	FCS	Fibre-channel switched channel
HIPER	IQD	Internal-queued-direct-communication channel IQD
OSA	OSA	Open-Systems-adapter channel
OSA	OSC	OSA 3270-console (OSC) channel
OSA	OSD	Open-Systems-Adapter Direct-Express (OSD) channel
OSA	OSE	Open-Systems-Adapter Express (OSE) channel
OSA	OSN	OSA NCP (OSN) channel
PAR	BL	Reserved, no longer valid, previously Parallel-block
PAR	BY	Reserved, no longer valid, previously Parallel-byte

New Channel Reporting:

- Specific channel type now reported

DASD Measurements

**Determine potential problems – Top DASD
Caching Characteristic? – No disk time?
What could be done to enhance this? (PAV?)**

Report: **ESADSD2** DASD Performance Analysis Velocity Software,

```

-----
Dev          Device  <%DevBusy>  <SSCH/sec->  <-----DASD Response times (ms
No. Serial  Type      Avg  Peak      avg  peak      Resp  Serv Pend Disc Conn DASD
-----
16:26:00
***Top DASD by Device busy
DE0E VM984F 3390-3    97.9 96.3 403.6 396.9    5.4   2.4 0.4 0.0 2.0 3.0
D2F1 VM9C1B 3390-3    38.2 37.6  47.6  46.8    8.0   8.0 0.3 5.3 2.5  0
DE51 VMD11E 3390-3    32.8 32.3  66.6  65.4    4.9   4.9 0.3 0.2 4.4  0
D2EC VM9C16 3390-3    27.9 27.4  59.7  58.8    4.7   4.7 0.3 0.4 4.0  0
D786 VMSPL1  3390-3    25.0 24.6  91.2  89.7    2.7   2.7 1.8 0.1 0.9  0
DE1F VM9860 3390-3    21.4 21.1  30.2  29.7    7.1   7.1 0.3 0.1 6.7  0
D787 VMSPL2 3390-3    18.2 17.9  66.6  65.5    2.7   2.7 1.7 0.1 0.9  0
D2C7 VM9C02 3390-3    17.3 17.0  23.3  22.9    7.4   7.4 0.3 4.6 2.5  0
D2CB VM9C00 3390-3    10.7 10.5   18.6  18.3    5.7   5.7 0.3 0.5 5.0  0
DE4E VMD11B 3390-3    10.6 10.4   18.7  18.4    5.7   5.7 0.3 0.5 4.9  0
***End Top DASD by device busy***

```

Cache Measurements

Evaluate the cache controller configuration
Each controller has a Device Address and Identifier
Size of the cache and non-volatile cache is shown

```
Report: ESADSDC          Cache Control Unit Configuration          Velocity Software, Inc.  ESAMAP 3.4.0
-----
<--Control--> <-Storage--> <-Cache Storage in MB> <Non-volatile> Cache
<---Unit----> <-Director-->          Off Not      Storage (MB) Fast  <Channel Paths Online>
Dev#  Model    <ID/Status->  Size  Avail line Avail  Avail Pinned Write  01 02 03 04 05 06 07 08
-----
02CB  3990-3E  0053/OnLine  3742 3742.0    0    0    4.00      0  Active 33 3B B3  .  .  .  .  .
D2C0  3990-6E  7733/OnLine  3072 3072.0    0    0   1024      0  Active 60 61 64 65 E8 E9 EC ED
D780  3990-3E  0051/OnLine  3750 3750.0    0    0    4.00      0  Active 42 4A C2  .  .  .  .  .
DAA1  3990-3E  0058/OnLine  1702 1701.7    0    0    4.00      0  Active 23 2A 2C D3 D5 DA DC  .
DB40  3990-3E  005B/OnLine  2638 2637.6    0    0    4.00      0  Active 23 2A 2C D3 D5 DA DC  .
DD40  3990-3E  0075/OnLine  1702 1701.7    0    0    4.00      0  Active 26 27 2E D6 D7 DD DE  .
DDA0  3990-3E  0076/OnLine  1702 1701.7    0    0    4.00      0  Active 26 27 2E D6 D7 DD DE  .
DE00  3990-3E  0078/OnLine  2638 2637.6    0    0    4.00      0  Active 26 27 2E D6 D7 DD DE  .
DE40  3990-3E  0079/OnLine  2638 2637.6    0    0    4.00      0  Active 26 27 2E D6 D7 DD DE  .
DE80  3990-3E  007A/OnLine  2638 2637.6    0    0    4.00      0  Active 26 27 2E D6 D7 DD DE  .
DEC0  3990-3E  007B/OnLine  2638 2637.6    0    0    4.00      0  Active 26 27 2E D6 D7 DD DE  .
```


Cache Configuration

- NVS (Non-Volatile Storage can be a limiter with Linux)

```
Report: ESADSDC      Cache Control Unit Configuration
-----
<---Control--> <-Storage--> <-Cache Storage in MB> <Non-volatile> Cache
<---Unit----> <-Director->           Off  Not   Storage (MB)  Fast  <Channel Paths Online >
Dev#  Model  <ID/Status-> Size  Avail line Avail  Avail Pinned  Write  01 02 03 04 05 06 07 08
-----
F000  2105    F000/OnLine  51K  50627  0    0  192.0    0  Active A4 A5 B4 B5 . . . .
F100  2105    F100/OnLine  51K  50627  0    0  192.0    0  Active A4 A5 B4 B5 . . . .
F200  2105    F200/OnLine  51K  50627  0    0  192.0    0  Active A4 A5 B4 B5 . . . .
F300  2105    F300/OnLine  51K  50627  0    0  192.0    0  Active A4 A5 B4 B5 . . . .
F400  2105    F400/OnLine  51K  50627  0    0  192.0    0  Active A4 A5 B4 B5 . . . .
F500  2105    F500/OnLine  51K  50627  0    0  192.0    0  Active A4 A5 B4 B5 . . . .
F600  2105    F600/OnLine  51K  50627  0    0  192.0    0  Active A4 A5 B4 B5 . . . .
F700  2105    F700/OnLine  51K  50627  0    0  192.0    0  Active A4 A5 B4 B5 . . . .
```

Cache Instrumentation

Cache Performance

- Evaluate the value of cache and why
- Read Percent/Read Cache Hit Percent – should be high
- Write Percent/Write Cache Hit Percent – should be 100

Report: **ESADSD5** 3990-3 Cache Analysis Velocity Software, Inc. ESAMAP 3.7.2 09/07/07 Pg 2241
 Monitor initialized: 02/06/07 at 13:00:00 on 2094 serial 2BFBD First record analyzed: 02/06/07 13:00:00

```
-----
                Pct. <-----per second-----> <-----Write activity per second-----> <tracks/second>
Dev            Actv <-----Total-----> <-----Read-----> <---Seq Read--> Total DFW DFW SEQ      NVS <---Cache---> <Staged-> De-
No.  Serial Samp I/O Hits Hit% Read% I/O Hits Hit% I/O Hits Hit% I/O I/O Hits I/O Hit% Full Inhib Bypass Seq Nseq stged
-----
```

13:15:00

Top DASD by Device busy

No.	Serial	Samp	I/O	Hits	Hit%	Read%	I/O	Hits	Hit%	I/O	Hits	Hit%	I/O	DFW	DFW	SEQ	NVS	Full	Inhib	Bypass	Seq	Nseq	stged
F794	VS2P52	100	9.7	5.4	55.4	83.8	8.1	3.8	46.8	0	0	0	1.6	3.2	1.6	0	100	0	0	0	0	0	1
F181	VS2P13	100	10.7	6.2	57.5	82.3	8.8	4.3	48.3	0	0	0	1.9	3.8	1.9	0	100	0	0	0	0	0	2
F4A8	VS2PB5	100	11.0	6.3	57.3	83.3	9.2	4.5	48.7	0	0	0	1.8	3.7	1.8	0	100	0	0	0	0	0	2
F2A6	VS2PA3	100	11.7	6.7	57.5	82.8	9.7	4.7	48.7	0	0	0	2.0	4.0	2.0	0	100	0	0	0	0	0	2
F598	VS2P66	100	11.0	6.3	56.9	82.0	9.0	4.3	47.4	0	0	0	2.0	4.0	2.0	0	100	0	0	0	0	0	2
F39C	VS2P78	100	10.9	6.1	56.6	83.3	9.0	4.3	47.9	0	0	0	1.8	3.6	1.8	0	100	0	0	0	0	0	2
F79F	VS2P95	100	10.1	5.6	55.3	83.4	8.5	3.9	46.3	0	0	0	1.7	3.4	1.7	0	100	0	0	0	0	0	1
F096	VS2P57	100	11.0	6.3	57.3	81.8	9.0	4.3	47.8	0	0	0	2.0	4.0	2.0	0	100	0	0	0	0	0	2
F3A6	VS2PA4	100	11.3	6.4	56.5	83.2	9.4	4.5	47.8	0	0	0	1.9	3.8	1.9	0	100	0	0	0	0	0	2
F0A9	VS2PB9	100	11.0	6.4	58.1	82.7	9.1	4.5	49.4	0	0	0	1.9	3.8	1.9	0	100	0	0	0	0	0	2

End Top DASD by Device busy



PAV – Parallel Access Volume

Base address and Parallel Volume configuration:

Report: **ESADSD1** DASD Configuration

Dev No.	Sys ID	Serial	Device Type	SHR	<CHPIDS OnLn>				<-Cntrl Unit->	PAV Base
					01	02	03	04	OBR/CU Model	
6700	053D	TEST1	3390-3	NO	2C	2E	.	.	35/1B 3990-6E	Base
6701	053E	CMSPAV	3390-3	NO	2C	2E	.	.	35/1B 3990-6E	Base
6702	053F	CMSPV2	3390-3	NO	2C	2E	.	.	35/1B 3990-6E	Base
6703	0540		3390-3	NO	2C	2E	.	.	35/1B 3990-6E	Base
67F0	062D		3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6703
67F1	062E		3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6703
67F2	062F		3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6703
67F3	0630		3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6703
67F4	0631	CMSPV2	3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6702
67F5	0632	CMSPV2	3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6702
67F6	0633	CMSPV2	3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6702
67F7	0634	CMSPV2	3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6702
67F8	0635	CMSPAV	3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6701
67F9	0636	CMSPAV	3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6701
67FA	0637	CMSPAV	3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6701
67FB	0638	CMSPAV	3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6701
67FC	0639	TEST1	3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6700
67FD	063A	TEST1	3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6700
67FE	063B	TEST1	3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6700
67FF	063C	TEST1	3390-3	NO	2C	2E	.	.	35/1B 3990-6E	6700

Allows multiple real addresses to point to one volume

- Allows multiple concurrent I/O
- One device is defined as the Base Address
- **Pool of alternate addresses**
- Alternate is assigned for the duration of the I/O
- Less alternates needed

ESADSD1 reports PAV Addressing

ESADSDx reports combine data and report as a base

- Option to report all addresses separately

HyperPAV Configuration

Report: **ESADSD1** DASD Configuration .3.1 05/18/17 Pg 142
 Monitor initialized: 10/01/14 at 13:08:46 on 2817 seri9:00
 Monitor period: 720 seconds (12:00) 1:00

Dev No.	Sys ID	Serial	Device Type	<CHPIDS SHR	OnLn	<-Cntrl Unit- OBR/	Cache Model	PAV STGID	BASE	<-HiperPav Type	Pool
C801	2A39	VS4W03	3390-3	NO	4B B0 4D 6F	3C/00	2107	C801	.	Base	0005
C901	2B09	VS4W02	3390-3	NO	4B B0 4D 6F	3C/00	2107	C901	.	Base	0007

Report: **ESADSDC** Cache Control Unit ZMAP 4.3.1 05/18/17 Pg 151
 Monitor initialized: 10/01/14 at 13:08:414 13:09:00
 Monitor period: 720 seconds (14 13:21:00)

<--Control--> <--Unit-->	<--Storage--> <--Director-->	<--Cache Stor <ID/Status-->	<-----HiperPav Devices----->						
Dev#	Model	<ID/Status-->	Size	Avail	Nmbr	Base	Alias	min	max
C800	2107	C801/OnLine	406K	405528	5	148	32	31	32
C900	2107	C901/OnLine	406K	405528	7	148	32	31	32
CA00	2107	CA01/OnLine	406K	405528	10	148	32	31	32
CB00	2107	CB01/OnLine	406K	405528	9	148	32	31	32

HyperPAV Performance

Rate going to alternate addresses is VERY low
Ensure the Device Utilization justifies the expense
If Serv = Resp, there is no queueing – PAV eliminated queueing

Report: **ESADSD2** DASD Performance Analysis TEST MAP Pg 1504
 Monitor initialized at 13:08:46 on 2817 serial 0CE0A6 First record

```

-----
                                <-----DASD Response time/HiperPAV>
Dev      Device <%DevBusy> <SSCH/sec->      <--Service times--> rate/Second->
No. Serial Type   Avg  Peak   avg  peak  Resp  Serv Pend Disc Conn SSCH Alias
-----
13:21:00
***Top DASD by Devi
C801 VS4W03 3390-3   6.0 11.8 182.6 361.2   0.3  0.3  0.2   0  0.2  0.26  0
C901 VS4W02 3390-3   5.9 11.7 182.6 360.8   0.3  0.3  0.2  0.0  0.2  0.27  0
C507 VS4S06 3390-3   4.8  9.6  27.5  54.5   1.8  1.8  0.2  0.0  1.6   0  0
C701 VS4W06 3390-3   4.0  7.7  66.5 131.0   0.6  0.6  0.2  0.1  0.3  6.38  0
C601 VS4W07 3390-3   3.9  7.7  67.3 133.0   0.6  0.6  0.2  0.1  0.3  6.81  0
C201 VS4S01 3390-3   3.9  7.7  22.7  45.0   1.7  1.7  0.2  0.0  1.5   0  0
C60A VS4W05 3390-9   3.1  6.1  64.7 127.7   0.5  0.5  0.2  0.0  0.3  1.30  0
C101 VS4S02 3390-3   2.8  5.5  16.8  33.2   1.7  1.7  0.2  0.0  1.5   0  0
  
```

HyperPAV Performance

**Rate going to alternate addresses is VERY low – as a percentage
Ensure the Device Utilization justifies the expense**

Report: **ESADSD2** DASD Performance Analysis 27/21 Pg 1670
Monitor initialized: 11/21/19 at 11:04:57 on 3906 serial 02

```
-----  
Dev          Device <--SSCH--> <%DevBusy> <SSCH/sec-> <----- <-PAV/HiperPAV>  
No. Serial  Type  Total  ERP  Avg  Peak  avg  peak  Resp Q-d  SSCH Alias  
-----
```

11:06:00

Top DASD by Device busy

4F3B	LF4F3B	3390-9	92509	0	18.4	18.1	1568	1542	0.1	0	10.5	0
4F3A	LF4F3A	3390-9	91049	0	17.7	17.4	1543	1517	0.1	0	67.8	0
4F3C	LF4F3C	3390-9	81521	0	16.4	16.2	1382	1359	0.1	0	14.2	0
4F3D	LF4F3D	3390-9	73770	0	15.8	15.5	1250	1230	0.1	0	49.1	0

“Top DASD” Case Study

Report: **ESADSD2** DASD Performance Analysis Linux Test

```

-----
                                <-----DASD Response time
Dev          Device <--SSCH--> <%DevBusy> <SSCH/sec->          <--Service times-->
No. Serial  Type  Total  ERP  Avg  Peak  avg  peak  Resp  Serv  Pend  Disc  Conn
-----
14:24:00
***Top DASD by Device busy***
F68E VS2P12 3390-3    594   0  11.3 11.3  10.1 10.1  11.3  11.3  0.3  0.0  11.0
F490 VS2P11 3390-3    617   0  11.1 11.1  10.5 10.5  10.6  10.6  0.5  0.0  10.1
F49A VS2P71 3390-3    572   0  10.3 10.3   9.7  9.7  10.6  10.6  0.2  0.0  10.3
F38F VS2P39 3390-3    558   0  10.1 10.1   9.5  9.5  10.6  10.6  0.2  0.2  10.2
F1A0 VS2P97 3390-3    542   0   9.9  9.9   9.2  9.2  10.8  10.8  0.2  0.2  10.3
F78E VS2P40 3390-3    530   0   9.8  9.8   9.0  9.0  10.9  10.9  0.3  0.0  10.7
***End Top DASD by Device busy***

```

```

-----
14:25:00
***Top DASD by Device busy***
F68E VS2P12 3390-3    316   0  74.2 74.2   5.4  5.4  138.5  139  0.3  127  10.7
F78E VS2P40 3390-3    353   0  74.2 74.2   6.0  6.0  123.9  124  0.3  112  11.3
F3A7 VS2PAC 3390-3    368   0  59.7 59.7   6.2  6.2  95.7  95.7  0.3  83.2  12.1
F491 VS2P15 3390-3    325   0  51.5 51.5   5.5  5.5  93.6  93.6  0.3  81.1  12.2
F6A6 VS2PA7 3390-3    316   0  39.7 39.7   5.4  5.4  74.1  74.1  0.3  60.3  13.5
***End Top DASD by Device busy***

```

Note the change in response time – why?

- Evaluate the components
- Which one is impacted?

“Top DASD” Case Study – The DASD Cache

Report: **ESADSD5** 3990-3 Cache Analysis
 Monitor ed: 06/02/07 at 13:00:00 on 2094 seriirst record analyzed: 06/02/07

Dev	<-----per second----->							<-----Write activity per se						
	<-----Total----->			<-----Read----->				Total	DFW	DFW	SEQ	NVS		
No.	Se	I/O	Hits	Hit%	Read%	I/O	Hits	Hit%	I/O	I/O	Hits	I/O	Hit%	Full
System:		2325	1527	65.7	65.3	1519	721	47.5	806.0	1724	806	0	100	0
System:		2719	1723	63.4	67.0	1822	827	45.4	896.7	1910	896	0	100	0
System:		2498	1626	65.1	65.6	1639	767	46.8	859.2	1842	859	0	100	0
System:		2238	1504	67.2	65.1	1458	724	49.7	780.6	1667	781	0	100	0
System:		3057	2126	69.6	57.9	1771	840	47.5	1286	2704	1286	0	100	0
System:		2934	2013	68.6	63.3	1857	937	50.4	1076	2264	1076	0	100	0
System:		2383	1677	70.4	65.4	1558	852	54.7	824.8	3964	825	0	100	0
System:		1846	1268	68.7	62.1	1147	569	49.6	699.4	5423	699	0	100	274
System:		1358	934	68.8	58.2	791	368	46.5	567.4	5047	567	0	100	1035
System:		1165	763	65.5	59.1	688	289	42.0	476.3	4753	474	0	99.4	1091
System:		1269	764	60.2	63.1	800	298	37.2	468.7	4255	467	0	100	1520
System:		1230	702	57.1	66.8	822	296	36.0	408.2	3684	406	0	99.4	1807
System:		1160	661	56.9	66.1	766	270	35.2	393.9	3198	391	0	99.2	2391
System:		1218	678	55.7	66.8	813	278	34.2	404.4	3228	400	0	99.0	2397
System:		1107	615	55.6	65.8	729	242	33.2	378.1	3153	373	0	98.6	2377
System:		1134	650	57.3	64.1	726	248	34.2	407.5	3214	402	0	98.6	2384
System:		3192	1751	54.9	73.5	2346	905	38.6	846.0	1864	846	0	100	106
System:		5660	2720	48.1	86.5	4899	1959	40.0	762.1	1605	761	0	100	5.2
System:		4454	2246	50.4	82.0	3651	1444	39.5	802.9	1732	802	0	100	54.8

Flags?

14:26 – DASD Cache?

What really happened?

- NVS filled up – writes go to disk (**Takes a LONG time**)
- Disks become highly utilized – everything slows down
- (Hardware design for z/OS track – not Linux records)

DASD Configuration Guidelines

Dedicate volumes for:

- Spool
- Paging
- T-disk
- SFS File Pools
- Linux shared disks
- Linux LVM (Logical Volume Manager)

HyperSockets Performance

Shows rates and errors:

Report: **ESACHNH** HiperSocket Channel Report
Monitor initialized: 03/02/17 at 15:02:00 on 2964 serial 496067

Time/ CHPID	<> Shrd	FCX	Channel Type	<Messages-> <Sent/Sec-> LPAR TOTAL	<DataUnits> <Sent/Sec-> LPAR Total	Failed Sends /Sec	<Failed Receives <Unavail Buffers LPAR Total			

15:03:00										
E4	Yes	No	HIPER/IQD	0.4 3.1	103.2	0	0 0	0.7		
E5	Yes	No	HIPER/IQD	0	0	0	0	0		
E7	Yes	No	HIPER/IQD	0	0	0	0	0		
E8	Yes	No	HIPER/IQD	0	0.6	0	0	0		
EA	Yes	No	HIPER/IQD	0	0	0	0	0		

15:04:00										
E4	Yes	No	HIPER/IQD	0.3 2.9	48.7	0	0 0.3	1.1		
E5	Yes	No	HIPER/IQD	0	0	0	0	0		
E7	Yes	No	HIPER/IQD	0	0	0	0	0		
E8	Yes	No	HIPER/IQD	0	0.6	0	0	0		
EA	Yes	No	HIPER/IQD	0	0	0	0	0		

Function has been moved to the Storage Processor (SAP)*

Internal architecture is important

Must often understand the limiting resource:

- Channel paths (CHPID)
- Host Adapters (should be one CHPID/HA)
- Ranks (one rank is one RAID-5 array of disks)
- Disks

Sequential devices are in one rank

- Concurrent I/O to devices in rank are delayed

*SAP – System Assistance Processor – executes internal code to provide information to the I/O subsystem for better performance.

SAP Case Study

Report: **ESADSD2** DASD Performance Analysis Velocity Sof

Dev No.	Device Serial	Type	---SSCH---		<%DevBusy>		<SSCH/sec->		<-----DASD Response tim				
			Total	ERP	Avg	Peak	avg	peak	Resp	Serv	Pend	Disc	Conn
21:19:00													
Top DASD by Device busy													
2218	VLS053	3990	747	0	12.1	12.1	12.4	12.4	9.8	9.8	0.5	5.1	4.1
E690	VLPPG7	3990	387	0	6.2	6.2	6.4	6.4	9.5	9.5	0.5	9.1	0.0
CE00	VLPPG1	3990	271	0	5.9	5.9	4.5	4.5	13.0	13.0	0.6	12.4	0.0
E691	VLPPG8	3990	306	0	5.1	5.1	5.1	5.1	10.0	10.0	0.5	9.5	0.0
1201	VLPPG5	3990	330	0	3.7	3.7	5.5	5.5	6.7	6.7	0.5	6.2	0.0
1202	VLPPG6	3990	314	0	3.6	3.6	5.2	5.2	6.9	6.9	0.5	6.4	0.0
1203	VLPPG3	3990	311	0	3.5	3.5	5.2	5.2	6.7	6.7	0.5	6.2	0
1334	VLPPG2	3990	332	0	3.3	3.3	5.5	5.5	6.0	6.0	0.5	5.5	0
1E07	VLS005	3990	426	0	2.5	2.5	7.1	7.1	3.5	3.5	0.5	1.4	1.7
1E08	VLS006	3990	116	0	1.6	1.6	1.9	1.9	8.3	8.3	0.6	4.5	3.2
End Top DASD by Device busy													
1CCC	Control Unit		1121	0	0.1	0.1	18.7	18.7	5.9	5.9	0.5	5.4	0.0
1DCC	Control Unit		2183	0	0.2	0.2	36.4	36.4	4.8	4.8	0.5	2.6	1.7
1FCB	Control Unit		923	0	0.1	0.1	15.4	15.4	3.3	3.3	0.5	1.4	1.4
20CB	Control Unit		311	0	0.0	0.0	5.2	5.2	1.6	1.6	0.5	0.1	1.0
21CB	Control Unit		3003	0	0.3	0.3	50.0	50.0	4.7	4.7	0.5	2.0	2.2
CE00	Control Unit		1506	0	0.1	0.1	25.1	25.1	4.9	4.9	0.6	3.1	1.2
E690	Control Unit		954	0	0.4	0.4	15.9	15.9	7.5	7.5	0.5	6.8	0.2
System:			12003	0	0.1	0.1	200.0	200.0	4.3	4.3	0.5	2.5	1.3

How to improve performance of top DASD?

High Connect time

- Check Channels

High Pend time

- Check IOP (I/Os per second)

High DISC time

- Check Cache

SAP Case Study

Report: **ESACHAN** Channel Performance Analysis Velocity
First

Time/ CHPID	<Pct Channel>		Chanl Shrd Type	<-----Data Units ----->				<---Reads/Second--->				<--Writes/Second-->			
	LPAR	Total		LPAR	TOTAL	Pct	Max	LPAR	TOTAL	pct	MAX	LPAR	TOTAL	pct	MAX
21:19:00															
15	0.1	13.8	Yes FICON	24	4610	3.9	117K	39	5172	4.4	117K				
16	0.1	13.9	Yes FICON	32	4759	4.1	117K	38	4889	4.2	117K				
17	0.1	13.8	Yes FICON	29	4476	3.8	117K	36	4961	4.2	117K				
1F	.	96.7	Yes ESCON												
35	0.1	14.1	Yes FICON	29	4842	4.1	117K	39	4845	4.1	117K				
36	0.1	13.9	Yes FICON	23	4839	4.1	117K	26	5147	4.4	117K				
37	0.1	13.9	Yes FICON	23	4861	4.1	117K	47	5234	4.5	117K				
3F	.	85.0	Yes ESCON												
9F	.	90.0	Yes ESCON												
B5	1.0	1.1	Yes FICON	273	274	0.2	117K	73	74	0.1	117K				
B6	1.0	1.1	Yes FICON	266	266	0.2	117K	72	72	0.1	117K				
B7	1.0	1.1	Yes FICON	260	260	0.2	117K	82	82	0.1	117K				
C5	1.1	1.1	Yes FICON	292	292	0.2	117K	81	81	0.1	117K				
C6	1.1	1.1	Yes FICON	278	278	0.2	117K	85	85	0.1	117K				
C7	1.0	1.0	Yes FICON	269	269	0.2	117K	81	81	0.1	117K				
D5	.	19.3	Yes FICON	0	5143	4.4	117K	0	10296	8.8	117K				
DF	.	81.7	Yes ESCON												
E5	.	16.7	Yes ESCON												
E7	.	15.0	Yes ESCON												
E8	.	26.7	Yes ESCON												
System:	14.6	603.2													

Check channel path busy:

- FICON channel paths
- Percent busy is ok
- The busy channels don't belong to us
- Are they impacting the IOP?

SAP Case Study

Report: **ESAIOP** I/O Processor Analysis

Velocity Softwa
First record an

```
-----  
I/O <-----I/O Processor----->  
Proc <Pct Util> <Rate/Second> <-Percent of Strts busy->  
Time Nbr Busy Idle SSCH Intrpts chan switch CtlUnit Device  
-----  
21:18:00 0 81.1 18.9 2387 2038 2152 17.1 0.4 0.9  
1 76.9 23.1 2670 3417 1379 6.7 0.3 0.1  
2 76.7 23.3 1664 3121 2520 0.7 0.4 0.4  
21:19:00 0 60.8 39.2 1990 1763 1888 12.3 0.3 1.7  
1 50.5 49.5 2241 2914 846 7.5 0.2 0.2  
2 50.8 49.2 1547 2754 1962 0.8 0.2 0.3  
21:20:00 0 55.6 44.4 2885 2564 1081 12.3 0.1 1.7  
1 42.6 57.4 3106 3851 527 3.9 0.1 0.1  
2 39.1 60.9 1787 2957 1089 0.6 0.2 0.3  
-----
```

Check channel processor (SAP):

- Processors are very busy
- Lots of I/O Starts because of the channel busy
- Probably not impacting our data
- Other LPARs need more IOP power

SAP Case Study

Report: **ESADSD5** 3990-3 Cache Analysis

```
-----  
          Pct. <-----per second-----  
Dev      Actv <-----Total-----> <-----Read----->  
No.  Serial Samp  I/O Hits Hit% Read%  I/O Hits Hit%  
-----  
21:19:00  
***Top DASD by Device busy***  
2218 VLS053 100 12.2 2.0 16.6 86.1 10.5 0.3 3.2  
E690 VLPPG7 100 6.1 2.9 47.5 75.7 4.6 1.4 30.7  
CE00 VLPPG1 100 5.7 4.6 80.9 65.3 3.7 2.6 70.7  
E691 VLPPG8 100 8.0 4.1 51.6 67.2 5.4 1.5 28.2  
1201 VLPPG5 100 5.9 3.3 56.0 63.6 3.7 1.1 30.8  
1202 VLPPG6 100 5.6 3.4 59.6 59.0 3.3 1.1 32.5  
1203 VLPPG3 100 3.9 2.0 51.9 72.1 2.8 0.9 33.3  
1334 VLPPG2 100 5.6 3.4 60.8 61.4 3.5 1.2 36.1  
1E07 VLS005 100 10.5 8.3 79.0 24.5 2.6 0.4 14.3  
1E08 VLS006 100 2.1 0.6 29.0 90.3 1.9 0.4 21.4  
***End Top DASD by Device busy***
```

Check caching status:

- High read percent and low read hit – probably backups running?
- Page device – read percent about 65%, low hit percent
- Evaluate cache controller technology – “writes” do not enter cache
- Should reduce load on paging devices too...

Linux Logical Volume Manager

Two reasons for using LVM:

- I/O performance
- Large files

I/O Performance:

- Stripe volumes
- May not add volume/space to LVM

Large Files:

- May add volume/space if not striped
- All new writes go to the new volume
- Bad performance

FCP eliminates 390 I/O CKD translation

Less processing to be done

- **But done in the CPU, not in the IOP/SAP**

Reportedly faster than FICON

Reportedly lower bandwidth than FICON

No way to effectively measure response times?

Linux I/O Analysis

Report: **ESASEEK** DASD Seeks Analysis Veloc
 Monitor period: 1320 seconds (22:00) Last

Dev No.	Serial	Device Type	Ownerid /userid	Mdisk Addr	<Cylinder> Start Stop	Total Seeks	<---Non-zero---> Seeks	Pct. Dist.	Read Pct.
0491	VMR408	3990	Volume:	.	0 .	2644	1587	60.0	342 41.8
			LNX0036	0292	51 2963	708	549	77.5	640 0
			LNX0036			708	549	77.5	640 0
			LNX0036	0293	3140 3173	1936	1038	53.6	184 57.0
			LNX0036			1936	1038	53.6	184 57.0
			cylinders	.	0 99	30	24	80.0	2379 0
			cylinders	.	200 299	24	24	100	1882 0
			cylinders	.	400 499	26	14	53.8	812 0
			cylinders	.	700 799	46	36	78.3	866 0
			cylinders	.	900 999	30	29	96.7	402 0
			cylinders	.	1100 1199	107	58	54.2	504 0
			cylinders	.	1300 1399	28	27	96.4	311 0
			cylinders	.	1500 1599	61	50	82.0	782 0
			cylinders	.	1600 1699	102	92	90.2	340 0
			cylinders	.	1700 1799	22	22	100	447 0
			cylinders	.	1800 1899	47	45	95.7	749 0
			cylinders	.	2000 2099	61	32	52.5	515 0
			cylinders	.	2200 2299	57	29	50.9	294 0
			cylinders	.	2400 2499	23	23	100	207 0
			cylinders	.	2700 2799	24	24	100	365 0
			cylinders	.	2900 2999	20	20	100	247 0
			cylinders	.	3100 3199	1936	1038	53.6	184 57.0
System:						6719	4277	63.7	606 17.9

Seeks are not optimized in Linux

Most I/O is write I/O

Analyzing Linux Disks

Report: **ESAUSEK** User DASD Seeks Report

```
-----  
Userid      Dev Volume <--Minidisk-> <Cylinder> Total  
/Time      No. Serial Ownerid  Addr Start Stop Seeks  
-----  
21:02:00  
Linux2      01E7 LX0200 LINUX2    0200      1 3326  107  
Linux4      01BB LX0407 LINUX4    0206      6 2982   82  
            01A9 LX0401 LINUX4    0200      1 2265   61  
            01B7 LX0403 LINUX4    0202    3180 3180    1  
Linux5      01C0 LX0501 LINUX5    0200    1365 1365    2  
ESAWRITE    022D ESALPS ESAWRITE 0191      51  186   21
```

Analyzing Linux Disks

```

Dev          Device Ownerid  Mdisk <Cylinder> Total <---Non-zero---> Read
No.  Serial Type    /userid Addr  Start Stop  Seeks  Seeks Pct. Dist. Pct.
-----
21:02:00
01E7 LX0200 3990  Volume:  .    0    .    107    84 78.5 1514  0
      LINUX2 0200  1 3326  107    84 78.5 1514  0
      LINUX2                107    84 78.5 1514  0
      cylinders .    0    9    12    12 100 2831  0
      cylinders .   10   19    24    12 50.0 2369  0
      cylinders .   40   49    12    12 100   31  0
      cylinders .  540  549    1    1 100  542  0
      cylinders .  550  559    1    1 100  555  0
      cylinders . 1720 1729    2    2 100 1719  0
      cylinders . 1890 1899    2    2 100  948  0
      cylinders . 2360 2369   12   12 100 1873  0
      cylinders . 2490 2499   12   12 100 1870  0
      cylinders . 2800 2809    1    1 100 2762  0
      cylinders . 2910 2919    4    4 100  339  0
      cylinders . 3060 3069    2    2 100 1791  0
      cylinders . 3070 3079   18    7 38.9  610  0
      cylinders . 3220 3229    1    1 100  309  0
      cylinders . 3240 3249    2    2 100  333  0
      cylinders . 3320 3329    1    1 100  104  0

```

Analyzing Linux Disks

```
Report: ESADSD5          3990-3 Cache Analysis
-----
          Pct. <-----per second----->
Dev       Actv <-----Total-----> <----Read---->
No.  Serial Samp  I/O Hits Hit% Read%  I/O Hits Hit%
-----
21:02:00
***Top DASD by Device busy***
01E7  LX0200  100  1.8  1.8  100    0    0    0    0
01BB  LX0407  100  1.4  1.4  100    0    0    0    0
022D  ESALPS  100  0.7  0.7  100   5.0  0.0  0.0  100
01A9  LX0401  100  1.0  1.0  100    0    0    0    0
***End Top DASD by Device busy***

2000 CtlUnit  100  2.4  2.4  100    0    0    0    0
01E6 CtlUnit  100  1.8  1.8  100    0    0    0    0
2200 CtlUnit  100  0.7  0.7  100   5.0  0.0  0.0  100
-----
System:      100  4.9  4.9  100   0.7  0.0  0.0  100
```

0 Percent reads at the DASD cache level

Analyzing Linux Disks

Report: **ESAUSR3** User Resource Uti

```
-----  
                DASD MDisk Virt Cache  
UserID          DASD Block Cache Disk Hit  
/Class          I/O   I/O   Hits  I/O  Pct  
-----  
21:02:00       289     0    13    0   4.5  
***User Class Analysis***  
*Servers        35     0    12    0  34.3  
*LINUX          254     0     1    0   0.4  
***Top User Analysis***  
LINUX2          107     0     0    0    0  
LINUX4          145     0     1    0   0.7  
LINUX5           2     0     0    0    0  
ESATCP           0     0     0    0    0  
ESAWRITE        35     0    12    0  34.3
```

0 Percent MDC reads at the cache level

Linux I/O Problem? Turn off MDC?

```
Screen: ESADSD2  xxxxxxxxxxxxxxxxxxxx          ESAMON V2.2  05/04 11:14-1
```

Time	Dev No.	Serial	Device Type	%Dev Busy	<SSCH/sec-> avg	peak	<-----Response times (ms)----> Resp	Serv	Pend	Disc	Conn
11:16:23	1007	LIN501	3390-3	80.8	31.0	31.0	26.1	26.1	0.2	15.1	10.8
	1008	LIN502	3390-3	98.7	36.1	36.1	27.3	27.3	0.2	16.1	11.0
	1009	LIN503	3390-3	45.5	18.9	18.9	24.1	24.1	0.2	13.4	10.5
11:17:23	1008	LIN502	3390-3	96.8	36.3	36.3	26.7	26.7	0.2	15.6	10.9
11:18:23	1008	LIN502	3390-3	99.2	37.1	37.1	26.8	26.8	0.2	15.7	10.8
11:19:23	1008	LIN502	3390-3	98.5	37.5	37.5	26.2	26.2	0.2	15.3	10.7
11:20:23	1007	LIN501	3390-3	28.8	13.4	13.4	22.7	21.5	0.2	12.1	9.1
	1008	LIN502	3390-3	41.7	17.0	17.0	24.5	24.5	0.2	14.0	10.3
11:30:23	1007	LIN501	3390-3	26.7	10.7	10.7	24.9	24.9	0.3	14.3	10.3
11:31:23	1007	LIN501	3390-3	8.0	5.2	5.2	15.3	15.3	0.2	3.5	11.5
	1008	LIN502	3390-3	10.7	20.6	20.6	5.2	5.2	0.2	1.3	3.6
12:00:23	1007	LIN501	3390-3	35.1	114.8	114.8	3.1	3.1	0.2	0.1	2.7
	1008	LIN502	3390-3	21.6	64.8	64.8	3.3	3.3	0.2	0.3	2.8
	1009	LIN503	3390-3	22.4	37.7	37.7	5.9	5.9	0.2	0.1	5.6
12:01:23	1007	LIN501	3390-3	4.5	13.0	13.0	3.5	3.5	0.2	0.7	2.6
	1008	LIN502	3390-3	35.6	160.0	160.0	2.2	2.2	0.2	0.1	1.9
12:06:23	1007	LIN501	3390-3	9.4	10.9	10.9	8.6	8.6	0.2	0.3	8.2
12:41:23	1007	LIN501	3390-3	24.3	20.3	20.3	12.0	12.0	0.2	1.4	10.4
12:42:23	1007	LIN501	3390-3	61.0	47.1	47.1	13.0	13.0	0.2	1.8	10.9
	1008	LIN502	3390-3	81.8	65.2	65.2	12.5	12.5	0.2	1.3	11.0
	1009	LIN503	3390-3	30.0	27.0	27.0	11.1	11.1	0.2	0.3	10.7
12:43:23	1008	LIN502	3390-3	95.9	74.1	74.1	12.9	12.9	0.2	1.9	10.9
12:44:23	1008	LIN502	3390-3	34.8	27.3	27.3	13.4	12.7	0.2	2.0	10.6

If a backup, then MDC “reads ahead”

If a database, random 4k I/O, MDC wastes bandwidth

Basis for recommending “MDC OFF” by IBM

==> MDC OFF

==> MDC ON

Linux I/O Problem? Turn off MDC?

Linux chains up to 128/1024 I/O

- Overflows the DASD non-volatile cache

MDC: Minidisk cache defaults to Track Level caching

- Read in one track for every read
- Linux writes 4k blocks, not synchronous
- Long connect times

MDC: Minidisk cache Block Level caching **CAN BE USED!!!**

- Requires diagnose I/O driver
- Reads one block per I/O

Watch MDC hits per device and per server

DASD Cache: Choose a DASD cache algorithm

- `tunedasd --get_cache` - to get caching behavior
- `tunedasd --cache sequential | normal | prestage | record | bypass`
- io scheduler – noop or none for oracle...

DASD Drivers

- Diagnose – designed to assist virtualized environments
 - Proven more efficient
 - Allows MDC Record Level caching
- Non-Diagnose
 - Default

SCSI I/O is not instrumented as is ECKD

- Monitor data for FBA emulated devices
- Performance is measured by each Linux

QDIO

- Eliminates hardware interrupts
- Used for FCP, HiperSockets, FICON

ESAQDIO

- Provides traffic by user
- QDIO – Network traffic (if CP maintains Shadow Queues)
 - TCPIP uses DIAG98, no Shadow Queues and no QDIO monitor record
- FCP – Fiber Channel network
- HPER - HiperSockets

SCSI Device Analysis - ESAQDIO

Provides traffic by user:

- Virtual Device Number – is used to get to the channel
- QDIO Formats – FCP, HiperSockets, QDIO
- QDIO Instruction Rate
- QDIO Traffic

Report: **ESAQDIO** Queued I/O Report Linux Test ESAMAP 3.6.2 10/19/06 Page 233
Monitor initialized: 10/09/06 at 08:00:00 on 2096 serial 29EBD First record analyzed: 10/09/06 08:00:00

Date/Time	Dev. Nmbr	Virt owner	QDIO DevN	QDIO Fmt	Number <QDIO SIGA Instructions/Sec->		<-Throughput / sec->				<--PCI Interupts-->		Cause						
					In	Out	<---Guest---->	<----CP----->	<Buffers>	<--Bytes-->	Sent	From	Sent	From	Count	Input	Complete		
							Read	Writ	"s"	Read	Writ	"s"	Sent	From	Sent	From	Count	Input	Complete
08:15:00	0000	Totals	0000	QDIO	0	0	0	237	928	0	393	0	678	404	20M	297K	7	7	0
	0800	TSMSEV	0800	FCP	1	1	0	0	46	0	15	0	73	15	3843K	0	0	0	0
	0801	TSMSEV	0801	FCP	1	1	0	0	0	0	0	0	0	0	94	0	0	0	0
	0802	TSMSEV	0802	FCP	1	1	0	0	44	0	14	0	71	14	3733K	0	0	0	0
	0C11	SAP000	0C05	HPER	1	4	0	55	109	0	55	0	55	55	17907	35100	0	0	0
	0C14	SAP010	0C05	HPER	1	4	0	95	184	0	95	0	95	97	34264	107K	0	0	0
	3D02	TCPVSWC1	3D02	QDIO	1	1	0	0	0	0	6	0	7	7	9533	1430	7	7	0
	3D02	TCPVSWC1	3D02	QDIO	1	1	0	0	0	0	6	0	7	7	9533	1430	7	7	0

SCSI Device Analysis – ESADEV1

Provides configuration:

- Device Number – gives Channel Path ID (CHPID)
- Device Owner (dedicated to UserID) same as ESAQDIO
- Hipersocket devices are on Channel Path FC-FF

Report: **ESADEV1** Device Configuration (non-DASD)

```
-----
```

Dev	Device	<CHPIDs	OnLn>	OBR	<-Cntrl	Unit->	UserID
No.	SysID	Type	01 02 03 04	Code	Code	Model	(if ded)
0800	0002	1732-3	33 . . .	00	00	1731-3	TSMSERV
0801	0003	1732-3	33 . . .	00	00	1731-3	TSMSERV
0802	0004	1732-3	33 . . .	00	00	1731-3	TSMSERV
0C11	0021	1732-5	FF . . .	00	00	1731-5	SAP000
0C12	0022	1732-5	FF . . .	00	00	1731-5	SAP010
0C13	0023	1732-5	FF . . .	00	00	1731-5	SAP010
0C14	0024	1732-5	FF . . .	00	00	1731-5	SAP010
0C15	0025	1732-5	FF . . .	00	00	1731-5	SAP015

SCSI Device Analysis - ESACHAN

Provides:

- Start with “channel busy”
- Maximum value of Read, Write and BUS

Report: **ESACHAN** Channel Performance Analysis Linux Test ESAMAP 3.6.2 10/19/06 Page 210
 Monitor initialized: 10/09/06 at 08:00:00 on 2096 serial 29EBD First record analyzed: 10/09/06 08:00:00

Time/ CHPID	<Pct Channel>		Chanl Shrd Type	<-----Data Units ----->				<----Work Unit---->				<-Bus Cycles>		Bytes /Data Unit					
	Utilization LPAR Total			<---Reads/Second-->		<--Writes/Second-->		<-Rates / Second-->		<-per Second>									
				LPAR	TOTAL	Pct	Max	LPAR	TOTAL	pct	MAX	LPAR	TOTAL	Pct	MAX	Used	pct	Max	
08:15:00																			
14	0.2	7.6	Yes FICON	37	3617	1.9	195K	198	1800	0.9	195K	285	9444	7.6	125K	1226	1.1	109K	1024
16	0.2	7.5	Yes FICON	35	3622	1.9	195K	188	1780	0.9	195K	284	9420	7.5	125K	1223	1.1	109K	1024
31	.	0.3	Yes FICON	0	0	0	123K	0	0	0	123K	0	19	0.3	7291	0	0	7291	2048
33	0.3	0.4	Yes FICON	10	10	0.0	123K	3948	3948	3.2	123K	21	30	0.4	7291	72	1.0	7291	2048
35	0.3	0.4	Yes FICON	2671	2671	2.2	123K	28	28	0.0	123K	21	29	0.4	7291	38	0.5	7291	2048
37	.	0.3	Yes FICON	0	0	0	123K	0	0	0	123K	0	19	0.3	7291	0	0	7291	2048
39	0.3	0.4	Yes FICON	2670	2670	2.2	123K	30	30	0.0	123K	23	30	0.4	7291	39	0.5	7291	2048
3B	0.1	0.3	Yes FICON	2	2	0.0	123K	785	785	0.6	123K	6	25	0.3	7291	14	0.2	7291	2048

System:	1.9	36.4																	

SCSI Device Data

ESAQDIO - show rates, device number:

Date/Time	Dev. Nbr	owner	Virt DevN	QDIO Fmt	Queues In	Queues Out	Number <QDIO SIGA Instructions/Sec->		<-Throughput / sec->				<--PCI Interrupts-->		Cause				
							<---Guest---->	<----CP----->	<Buffers>		<--Bytes-->		Count	Input	Complete				
							Read	Writ	"s"	Read	Writ	"s"	Sent	From	Sent	From			
08:15:00	0000	Totals	0000	QDIO	0	0	0	237	928	0	393	0	678	404	20M	297K	7	7	0
	0800	TSMSEV	0800	FCP	1	1	0	0	46	0	15	0	73	15	3843K	0	0	0	0
	0802	TSMSEV	0802	FCP	1	1	0	0	44	0	14	0	71	14	3733K	0	0	0	0
	0C14	SAP010	0C05	HPER	1	4	0	95	184	0	95	0	95	97	34264	107K	0	0	0

ESADEV1 - Shows channel path ID

Dev No.	SysID	Device Type	<CHPIDs OnLn>				OBR Code	<-Cntrl Unit-> Model	UserID (if ded)	
			01	02	03	04				
0800	0002	1732-3	33	.	.	.	00	00	1731-3	TSMSEV
0801	0003	1732-3	33	.	.	.	00	00	1731-3	TSMSEV
0802	0004	1732-3	33	.	.	.	00	00	1731-3	TSMSEV

ESACHAN - Shows channel path utilizations

Time/CHPID	<Pct Channel> Utilization			Chanl Shrd	Type	<-----Data Units ----->				<----Work Unit---->				<-Bus Cycles>			Bytes /Data Unit				
	LPAR	TOTAL	Pct			<---Reads/Second-->		<--Writes/Second-->		<-Rates / Second-->		<-per Second>		Used	pct	Max	Unit				
						LPAR	TOTAL	Pct	MAX	LPAR	TOTAL	pct	MAX	LPAR	TOTAL	Pct	MAX				
08:15:00																					
33	0.3	0.4		Yes	FICON	10	10	0.0	123K	3948	3948	3.2	123K	21	30	0.4	7291	72	1.0	7291	2048
35	0.3	0.4		Yes	FICON	2671	2671	2.2	123K	28	28	0.0	123K	21	29	0.4	7291	38	0.5	7291	2048

System: 1.9 36.4

Spread the load

- Across multiple channels and ranks
- Across devices
- Across internal SAP paths

Define many smaller devices (3390-3), not few larger devices (3390-9)

FICON can help

Reduce seek time

Utilize DASD cache

Utilize MDC

Block up

SET IOPRIORITY if queueing

SET SRM DSPSLICE 1 to favor I/O workloads