

DASD Reports and Analysis

DASD I/O Overview

Modeling DASD, Historical Perspective

Cache Controllers

Tuning MDC

Tuning DASD Cache

Seek Analysis

OSA Analysis

## DASD

- ESADSD1 – Configuration
- ESADSD2 – Disk Performance
- ESADSD6 – Disk Performance extended

## DASD Cache

- ESADSDC – DASD Cache configuration
- ESADSD5 – DASD Cache performance

## Channels

- ESACHNC – Channel configuration
- ESACHAN – Channel Performance
- ESACHNH – Hypersockets
- ESAQDIO,ESAQDIO2 – Queued I/O analysis
- ESAIOP – I/O Processor (SAP)

## Seek Analysis

- ESASEEK – Seeks by minidisk
- ESAUSEK – Seeks by user

# DASD Configuration Report

## ESADSD1 show configuration, known information

Report: ESADSD1      DASD Configuration      Linux Test  
Monitor initialized: 05/06/08 at 12:00:00 on 2094 serial AEA7D      First record analyze  
Monitor period:      3600 seconds (    1:00:00)      Last record:

```
-----  
Dev Sys                    Device       <CHPIDS OnLn><-Cntrl Unit-> UserID    MDisk <----Extent---->  
No. ID    Serial Type    SHR 01 02 03 04 OBR/CU Model    (if ded) Links Type    Start Size  
-----  
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
```

Report: ESADSD1      DASD Configuration                    ESAMAP 3.7.4 06/06/08    Pg    2815  
Monitor initialized: 05/06/08 at 12:00:00ed: 05/06/08 12:00:00  
Monitor period:      3600 seconds (      05/06/08 13:00:00

```
-----  
Dev Sys                    Device       <CHPIDS OnLn> . <--MDC Status--> <3990 Cache Status> PAV  
No. ID    Serial Type    SHR 01 02 03 04 . Elig Def Now Shr Actv DFW NVS STGID    Base  
-----  
C557 6177 LNXP5F 3390-9 YES 48 49 76 87 . Yes Off    . Yes    Yes    On    On    9C50  
                                                  95 96 D6 5A .  
C558 6178 LNXP55 3390-9 YES 48 49 76 87 . Yes Off    . Yes    Yes    On    On    9C50 .  
                                                  95 96 D6 5A .
```

## ESADSD2 provides

- Interval top dasd
- Interval control unit
- Summary - everything

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 provides cache statistics

- Read/Hit depends on cache size and locality of reference
- Write/Hit depends on NVS (Persistent Memory)
- NVS Full, Cache Inhibit, Cache Bypass are 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***  
0C51 VME090 100 172 172 100 100.0 64.6 64.6 100 108 108 100 0 0 0 0 0 0 0 0  
0B6E VME106 100 131 131 100 100.0 100 100 100 31.3 31.3 100 0 0 0 0 0 0 0 0  
0E78 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  
0B00 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  
0C00 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  
0D00 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  
0E00 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  
-----
```

# Disk Performance history

Assumptions and rules of thumb based on technology:

DASD Assumptions on what is “good”:

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

DASD Cached Assumptions

- 3880-23 (1982), 3380: 15ms
- 3390-1 (1988), 3390, gray cable: 10ms
- 3390-3 (1995), ESCON: 5ms
- 3390-9 (2001), FICON: 2ms
- “3390” (2009) sub 1ms

FCP only “Recently measurable”

## Note change in response time, explain?

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
F3A7 VS2PAC 3390-3   600   0   9.7  9.7  10.2  10.2   9.5   9.5  0.3  0.0   9.2
***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
F5A8 VS2PB6 3390-3   322   0  38.5 38.5   5.5   5.5  70.5  70.5  0.3  57.9 12.4
F795 VS2P56 3390-3   325   0  37.5 37.5   5.5   5.5  68.0  68.1  0.3  52.7 15.0
***End Top DASD by Device busy***

```

## PEND Time

- CPU Channel Connection
- Protocol

## Disconnect Time

- Rotational Delay
- Seek time
- (Internal controller contention)

## Connect Time

- Transmission time

## Queue Time

- Control Unit
- Device
- Throttle



## DASD Response Time

= Service time + Queue Time

## DASD Service Time

= (Pend + DISC + CONN)

Device Busy = rate \* service time

## DASD Response Time (queuing theory: MM1)

= (Service time) / (1 – Device Busy)

## Service Time

= (Pend + DISC + CONN)

## DISCONNECT Time

= (Rotational Delay + Seek + RPS Miss + Internal contention)

## Connect Time

= DataSize / DataRate (On ESCON and prior)

Ficon is different.

## Pend

= Protocol / (1 - Channel Utilization)

## DISCONNECT Time

= (Rotational Delay + Seek + RPS Miss)

Rotational Delay =  $\frac{1}{2}$  Rotation on average

Seek (3380) is 7ms on Average, Tuneable

## RPS Miss

= Rotation \* misses

= Rotation / (1 - Channel Utilization) - Rotation

# DASD I/O Enhancements (History Lesson)

## Cache (Reduce disconnect time)

- Read only (3880-23, 1980)
- Write with non-volatile cache (8mb, 1985)
- Large Write cache (GB, 1990)
- On board cache (device level)
- MDC
- RAID

## Channels (Reduce connect time)

- VM/XA Multiple Channels
- Grey Cables, 3390 (4.5mb/sec)
- XA Channels
- Escon (17mb/sec, 1995)
- Ficon (110mb/sec, 2003)
- Ficon Express (200/400 mb/sec) 2005/2007

## PAV (z/VM 5.2) – Reduce queue time

VM/XA allowed multiple channels to access same device, reconnect on non busy channel

If 1<sup>st</sup> channel busy, try 2<sup>nd</sup>, and so on

Architecture supports up to 16 channels

Four channels should be considered minimum

Some controllers (IBM 6800) do not support dynamic reconnect

## Why do you care about channel utilization????

### RPS Miss (3380, XA)

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

### RPS Miss Times – 20% Channel Util

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

Utilization can be higher when more channels

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

## Cache reduces “Hardware” DISConnect Time

### DISConnect (non-cache)

$$= (\text{Rotational Delay} + \text{Seek} + \text{RPS Miss})$$

### DISConnect (cache)

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

### Rotational Delay + Seek with cache

$$= (17\text{ms} / 2 + 7\text{ms}) * (\text{CacheMiss} / \text{CacheTotal})$$

## Block sizes vary and impact assumptions

- 4K I/O
- MDC does full track I/O
- Linux does 4k I/O
- Linux can chain 1024 I/O



## Storage “Control Units”

- Raid
- Storage Processors
- Multiple Internal Paths
- Logical Volumes vs Physical Volumes

## PAV (2000)

- Z/OS Guest
- Linux Guest, dedicated devices

## Linux

- LVM

## Subchannel Measurement Block

- Pend
- Connect
- Disc

## Seek Analysis

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

## Channels

- VM Sample
- FICON

## Old Style

- Sampled by CP
- Measured by Channel

Utilization Sampled at  
High frequency rate,  
example is 6 second  
sample

LPAR is total for all  
LPARs, measured.

Report: ESACHAN      Chann  
Monitor initialized: 01/11

```
-----  
                        <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
```

## New Style (z/VM 3.1)

- Sampled by CP (TOTAL)

- Measured by Channel (LPAR)

- What is Channel Utilization?

- (EMIF)

```

Report: ESACHAN          Channel Performance Analysis          Velocit
Monitor initialized: 06/30/03 at 13:21:34 on 2064 serial 5146B  First r
-----
Time/    <Pct Channel>
CHPID    Utilization
-----
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
A8        4.5  17.8  Yes ESCON
B4        16.4 19.4  Yes ESCON
-----

```

## Channel Utilization not obvious

- One fibre for read
- One fibre for write
- Data unit processing
- Channel Bus
  
- Channel balancing takes planning

# Ficon Channel Measurements

## What is channel utilization?

- Sampled by CP
- Measured by Channel, data / second

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>			Chanl Type	<-----Data Units ----->							
	Utilization LPAR Total	Shrd			<---Reads/Second-->				<---Writes/Second-->			
					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
F6	.	3.9	Yes	FICON	0	0	0	82K	0	0	0	82K
F7	.	3.9	Yes	FICON	0	0	0.0	82K	0	0	0	82K
System:	20.8	142.5										

# Ficon Channel Measurements

## New Style

- Sampled by CP
- Measured by Channel, work per second
- What is Channel Utilization?

Report: ESACHAN Channel Perest ESAMAP 3.3.0 Monitor  
 initialized: 05/09/03 atecord analyzed: 05/09/03 10:54:03

Time/ CHPID	<Pct Channel>			Chanl Type	<----Work Unit---->				<-Bus Cycles>			Bytes /Data Unit	
	Utilization	LPAR	Total		Shrd	<-Rates / Second-->	LPAR	TOTAL	Pct	MAX	<-per Second>		Used
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
F6	.	3.9		Yes	FICON	0	636	3.9	16K	2463	15	16K	1024
F7	.	3.9		Yes	FICON	0	635	3.9	16K	2465	15	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 Utilization		Shrd	Channel Class/Type	<-----Data Units----->				<--Writes/Sec		
	LPAR	Total			<---Reads/Second--->		LPAR	TOTAL			Pct
12:01:00											
48	0.3	8.5	Yes	FICON/FCS	36	7872	6.7	117K	294	2228	1
49	0.3	8.4	Yes	FICON/FCS	35	7827	6.7	117K	315	2127	1
5A	0.4	11.6	Yes	FICON/FCS	53	10570	5.4	195K	441	3312	1
76	0.3	8.5	Yes	FICON/FCS	30	7984	6.8	117K	267	2284	1
87	0.3	8.4	Yes	FICON/FCS	32	7568	6.5	117K	285	2078	1
95	0.5	12.6	Yes	FICON/FCS	57	11393	5.8	195K	455	3583	1
96	0.5	12.4	Yes	FICON/FCS	58	11495	5.9	195K	471	3466	1
D6	0.4	11.9	Yes	FICON/FCS	52	11192	5.7	195K	471	3303	1

Report: ESADSD1		DASD Configuration								Linux test				
Dev No.	Sys ID	Serial	Device Type	SHR	<CHPIDS OnLn>				<-Cntrl Unit->	UserID (if ded)	MDisk Links	<----Extent---->		
					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						





## New Channel reporting

- Specific channel type now reported:

<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

Determine potential problems: Top dasd

What is caching characteristic? (no disc time?)

What could be done to enhance?

Report: ESADSD2      DASD Performance Analysis      Velocity Software,  
Monitor initialized at 16:24:57 on 9021 serial 42105      First record analy

```
-----  
                                     <-----DASD Response times (ms  
Dev          Device  <%DevBusy> <SSCH/sec->          <---Service times---> <---Qu  
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 cache controller configuration

Each controller has Device Address and Identifier

Size of cache, non-volatile cache 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 store often 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    <Channell 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 Performance

- Evaluate value of cache and why
- Read percent/Read cache
- Write percent/Write cache

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***
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***
```

## Dedicate volumes for:

- Spool
- Paging
- Tdisk
- SFS File Pools
- Linux shared disk
- Linux LVM

# Parallel Access Volumes: PAV

Allows multiple real addresses to point to one volume

- One device defined as Base address
- Other addresses defined as Alternate

ESADSD1 reports PAV Addressing

ESADSDx reports combine data and report as base

- Option to report all address separately

Added to z/VM in release 5.1

Linux support available, requires z/VM

Report: ESADSD1 DASD Configuration

```

-----
  Dev Sys      Device      <CHPIDS OnLn><-Cntrl Unit->  PAV
  No.  ID      Serial  Type      SHR  01  02  03  04  OBR/CU  Model      Base
-----
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

```

Base address and parallel volume configuration shown on ESADSD1



Function has moved to Storage Processor

**Internal architecture is important**

Must often understand 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

## How to improve performance of top DASD?

### Connect time high

- Channels

### Pend time

- IOP

### DISC

- Cache

```

Report: ESADSD2          DASD Performance Analysis          Velocity Sof
-----
Dev          Device <--SSCH--> <%DevBusy> <SSCH/sec-> <-----DASD Response tim
No. Serial  Type  Total  ERP  Avg  Peak  avg  peak  Resp  <--Service times-->
-----
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
    
```

## Check configuration

- Channel paths
- Caching status

Report: ESADSD1      DASD Configuration Inc.      ESAMAP 3.4.0 01/25/04    Page 331  
03/25/03 21:17:00  
Monitor period:      1260 seconds (      03/25/03 21:38:00

Dev No.	Sys ID	Serial	Device Type	Device	<CHPIDS	OnLn>	<--MDC Status-->				<3990 Cache Status>				PAV Base		
				SHR	01	02	03	04	Elig	Def	Now	Shr	Actv	DFW	NVS	STGID	Base
2214	2214	VLS049	3990	NO	B5	B6	B7	C5	Yes	On	On	No	Yes	On	On	2F00	.
					C6	C7											
2215	2215	VLS050	3990	NO	B5	B6	B7	C5	Yes	On	On	No	.	.	.	.	.
					C6	C7											
2216	2216	VLS051	3990	NO	B5	B6	B7	C5	Yes	On	On	No	.	.	.	.	.
					C6	C7											
2217	2217	VLS052	3990	NO	B5	B6	B7	C5	Yes	On	On	No	.	.	.	.	.
					C6	C7											
2218	2218	VLS053	3990	NO	B5	B6	B7	C5	Yes	On	On	No	Yes	On	On	2F00	.
					C6	C7											
2219	2219	VLS054	3990	NO	B5	B6	B7	C5	Yes	On	On	No	.	.	.	.	.
					C6	C7											

## Check Channel path busy

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

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 Processor (SAP)

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

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  Nmbr 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 Caching status

- High read percent and low read hit: probably backup?
- 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....

```

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***

```

## Two reasons for using LVM

- I/O Performance
- Large files

## Performance:

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

## Large files

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

FCP eliminates 390 I/O CKD translation

Less processing to be done

- (but done in CPU, not the IOP/SAP)

Reportedly faster than FICON

Reportedly lower bandwidth than FICON

No way to effectively measure response times



## Seeks are not optimized in Linux

## Most I/O is write I/O

Report: ESASEEK		DASD Seeks Analysis								Veloc	
Monitor period:		1320 seconds ( 22:00)								Last	
Dev	Device	Ownerid	Mdsk	<Cylinder>	Total	<---Non-zero---	Read				
No.	Serial	Type	/userid	Addr	Start	Stop	Seeks	Seeks	Pct.	Dist.	Pct.
0491	VMR408	3990	Volume:	.	0	.	2644	1587	60.0	342	41.8
			LNK0036	0292	51	2963	708	549	77.5	640	0
			LNK0036				708	549	77.5	640	0
			LNK0036	0293	3140	3173	1936	1038	53.6	184	57.0
			LNK0036				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
							6719	4277	63.7	606	17.9

System:



# 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

## 0 Percent reads at DASD cache level, Linux caches everything?

```
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 MDC reads at cache level

```
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
```

# 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	<SSCH/sec> peak	Response times (ms) Resp	Response times (ms) Serv	Response times (ms) Pend	Response times (ms) Disc	Response times (ms) 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	=====> MDC OFF
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	=====> MDC ON
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 backup, then MDC “reads ahead”,  
if database, random 4k I/O, MDC wastes bandwidth

# LINUX I/O Problem? Turn Off MDC???

## Linux chains up to 128 / 1024 I/O

- Overflows the DASD Non-volatile cache

## MDC: Minidisk cache Track Level Caching (Default)

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

## Minidisk caching Block level

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

## Watch MDC hits per device, per server

- Still research to do

## DASD CACHE: choose dasd cache algorithm

- `tunedasd --get_cache` to get caching behaviour
- `tunedasd --cache`
  - sequential | normal | prestage | record | bypass

## DASD Drivers

- Diagnose – designed to assist virtualized environment
  - 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

## QDIO

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

## ESAQDIO

- Provides traffic by user
- QDIO – Network traffic
- FCP – Fiber channel network
- HPER – HiperSockets
- “This record is emitted for only real QDIO devices for which z/VM CP is maintaining shadow queues. (If CP is not shadowing the queues, it doesn't see the I/O and therefore can't emit this record.) Typically Linux guests would use shadow queues but z/VM TCP/IP would not. z/VM TCP/IP drives its real QDIO devices using Diag X'98' and does not use shadow queues. “

## ESAQDIO provides traffic by user

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	owner	Virt DevN	QDIO Fmt	Number		<QDIO SIGA Instructions/Sec->			<-Throughput / sec->			<--PCI Interupts-->			Cause			
					Queues	In	Out	<---Guest-->	<---CP-->	<Buffers>	<--Bytes-->	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

- Device Number - is used to get to get to channel
- QDIO Format – FCP, HiperSockets, QDIO
- QDIO Instruction rate
- QDIO traffic

## ESADEV provides configuration

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	TSMSEV
0801	0003	1732-3	33 . . .	00	00	1731-3	TSMSEV
0802	0004	1732-3	33 . . .	00	00	1731-3	TSMSEV
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

- Device Number – gives Channel Path ID
- Device owner (dedicated to userid) same as ESAQDIO
- Hipersocket devices are on Channel path ID FC-FF

## ESACHAN

- No distinction on channel FICON vs FCP until z/VM 5.2
- What is device busy? Maximum value of (Read, Write, 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 Type	<-----Data Units ----->								<----Work Unit---->				<-Bus Cycles>			Bytes /Data Unit	
	Utilization	LPAR	Total		<---Reads/Second-->				<--Writes/Second-->				<-Rates / Second-->				<-per Second>				
	LPAR	Total	Shrd		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	
<b>33</b>	<b>0.3</b>	<b>0.4</b>	<b>Yes</b>	<b>FICON</b>	<b>10</b>	<b>10</b>	<b>0.0</b>	<b>123K</b>	<b>3948</b>	<b>3948</b>	<b>3.2</b>	<b>123K</b>	<b>21</b>	<b>30</b>	<b>0.4</b>	<b>7291</b>	<b>72</b>	<b>1.0</b>	<b>7291</b>	<b>2048</b>	
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																			

## ESAQDIO – show rates, device number:

Date/Time	Dev. Nmbr	owner	Virt DevN	QDIO Fmt	Queues In	Out	Number <QDIO SIGA Instructions/Sec->		<-Throughput / sec->				<--PCI Interupts-->				Cause		
							Read	Write	"s"	Read	Write	"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	TMSERV	0800	FCP	1	1	0	0	46	0	15	0	73	15	3843K	0	0	0	0
	0802	TMSERV	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-> Code	UserID Model	(if ded)
0800	0002	1732-3	33 . . .	00	00	1731-3	TMSERV
0801	0003	1732-3	33 . . .	00	00	1731-3	TMSERV
0802	0004	1732-3	33 . . .	00	00	1731-3	TMSERV

## ESACHAN – Shows channel path utilizations

Time/	<Pct Channel>			Chanl	<-----Data Units ----->								<----Work Unit----->				<-Bus Cycles>			Bytes	
CHPID	LPAR	Total	Shrd	Type	<---Reads/Second-->				<--Writes/Second-->				<-Rates / Second-->				<-per Second>			/Data	
					LPAR	TOTAL	Pct	Max	LPAR	TOTAL	pct	MAX	LPAR	TOTAL	Pct	MAX	Used	pct	Max	Unit	
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

## Channel utilization on OSA is high, why?

Report: ESACHAN Channel Performance  
Monitor initialized: 06/10/19 at 00:00:0

```
<Pct Channel>
Time/      Utilization      Channel
CHPID     LPAR Total  Shrd FCX Class/Type
-----
00:15:11
78         18.1  18.1   Yes No FCP  /FCP
7C         18.1  18.1   Yes No FCP  /FCP
81          0.4  13.6   Yes No OSA  /OSD
A1          2.3  16.1   Yes No OSA  /OSD
DC         24.1  82.0   Yes No OSA  /OSD
F8          5.8  18.8   Yes No OSA  /OSD
-----
System:    79.4 189.1
-----
00:30:11
78         18.4  18.4   Yes No FCP  /FCP
7C         18.4  18.4   Yes No FCP  /FCP
81          0.2  11.4   Yes No OSA  /OSD
A1          3.0  14.0   Yes No OSA  /OSD
DC         23.4  78.6   Yes No OSA  /OSD
F8          3.5  12.4   Yes No OSA  /OSD
-----
System:    74.8 166.1
-----
00:45:11
78         18.5  18.5   Yes No FCP  /FCP
7C         18.5  18.5   Yes No FCP  /FCP
81          0.1  10.4   Yes No OSA  /OSD
A1          2.7  13.0   Yes No OSA  /OSD
DC         26.9  90.1   Yes No OSA  /OSD
F8          3.0  11.3   Yes No OSA  /OSD
-----
System:    76.3 170.3
```

## What do we know?

- OSA
- Virtual switch
- QDIO
  
- Channel is “90% BUSY”

## Channel utilization on OSA is high, why?

Report: ESACHAN Channel Performance Analysis CENTERS FOR  
 Monitor initialized: 06/10/19 at 00:00:00 on 2964 serial 13C7B7 First record

<Pct Channel>							<-----Data Units ----->					
Time/	Utilization		Channel				<---Reads/Second---				<---Writes/Second	
CHPID	LPAR	Total	Shrd	FCX	Class/Type	LPAR	TOTAL	Pct	Max	LPAR	TOTAL	pct
-----												
00:45:11												
78	18.5	18.5	Yes	No	FCP /FCP	579	2210K	135	2M	670	3416K	208
7C	18.5	18.5	Yes	No	FCP /FCP	579	2209K	135	2M	670	3411K	208
81	0.1	10.4	Yes	No	OSA /OSD	0	2	0.0	128K	0	0	0
A1	2.7	13.0	Yes	No	OSA /OSD	120	433	0.3	128K	68	356	0.3
DC	<b>26.9</b>	<b>90.1</b>	Yes	No	OSA /OSD	164	3313	2.6	128K	3032	252K	<b>197</b>
F8	3.0	11.3	Yes	No	OSA /OSD	54	211	0.2	128K	205	560	0.4
-----												
System:	76.3	170.3										

## Channel wire utilization

- Utilization is rate / max
- As defined in monitor
- Note – this lpar is 27%....

## Channel utilization on OSA is high, why?

Report: ESACHAN

Monitor initialized

```
-----  
      <Pct Channel>  <-----Work Unit----->  <-Bus Cycles>  Bytes  
Time/   Utilization  <--Rates / Second-->  <-per Second>  /Data  
CHPID   LPAR Total    LPAR  TOTAL Pct  MAX  Used pct  Max  Unit  
-----  
00:45:11  
78      18.5  18.5      678  1381  19 7453  1381  19 7453  2048  
7C      18.5  18.5      677  1378  18 7453  1378  18 7453  2048  
81      0.1  10.4         0   746  10 7204    9 0.1 7204  2048  
A1      2.7  13.0      195   933  13 7204    8 0.1 7204  2048  
DC      26.9  90.1       89  6492  90 7204  1940  27 7204  2048  
F8      3.0  11.3      216   815  11 7204    10 0.1 7204  2048  
-----  
System:  76.3 170.3
```

## Channel utilization

- OSA
- Virtual switch
- QDIO



## Channel utilization on OSA is high, why?

Report: **ESADEV1** Device Configuration (non-DASD)  
 Monitor initialized: 06/10/19 at 00:00:00 on 2964 serial 13  
 Monitor period: 86400 seconds ( 24:00:00)

```
-----
```

Dev No.	SysID	Device Type	<CHPIDs OnLn>				OBR Code	<-Cntrl Code	Unit Model	UserID (if ded)
			01	02	03	04				
8106	3AFE	1732-1	81	.	.	.	00	00	1731-1	DTCVSW3
8400	3B1A	1732-1	F8	.	.	.	00	00	1731-1	
8403	3B1D	1732-1	F8	.	.	.	00	00	1731-1	DTCVSW2
8406	3B20	1732-1	F8	.	.	.	00	00	1731-1	DTCVSW2
9000	3B5D	1732-1	A1	.	.	.	00	00	1731-1	DTCVSW1
9006	3B63	1732-1	A1	.	.	.	00	00	1731-1	DTCVSW1
9009	3B66	1732-1	A1	.	.	.	00	00	1731-1	DTCVSW3
9400	3B8B	1732-1	DC	.	.	.	00	00	1731-1	
9406	3B91	1732-1	DC	.	.	.	00	00	1731-1	DTCVSW1
9409	3B94	1732-1	DC	.	.	.	00	00	1731-1	DTCVSW2
9800	3BAC	1732-1	9C	.	.	.	00	00	1731-1	DTCVSW3

## Channel utilization on OSA is high, why?

Report: **ESAVSW2** Virtual Switch Report Part Two  
Monitor initialized: 06/10/19 at 00:00:00 on 2964 serial 13C

```
-----  
Date/Time          <-packets><---network lock requests---->  
  Dev Sys  Control <-/second><--Per Second-> waits/sec wait  
  No. ID    Userid  Sent Rcvd Ntwrk send rcv send rcv time  
-----  
00:45:11  
8106 3AFE DTCVSW3      0      0      0      0      0      0      0      0  
8403 3B1D DTCVSW2      0      0      0      0      0      0      0      0  
8406 3B20 DTCVSW2    691    563    367    651    365      0      0      0  
8409 3B23 DTCVSW4      0      0      0      0      0      0      0      0  
8700 3B4C DTCVSW4      0      0      0      0      0      0      0      0  
9000 3B5D DTCVSW1      0      0      0      0      0      0      0      0  
9006 3B63 DTCVSW1      0      0      0      0      0      0      0      0  
9009 3B66 DTCVSW3    451    462    408    445    406      0      0      0  
9406 3B91 DTCVSW1    206    200    190    205    190      0      0      0  
9409 3B94 DTCVSW2   4188    936    372   3653    366      0      0      0  
9800 3BAC DTCVSW3      0      0      0      0      0      0      0      0
```



## Channel utilization on OSA is high, why?

Report: ESATCP4 TCPIP Hardware Layer/Interfaces Report

Date/ Time Node	IFT	<Total Octets> <-Per second-> Input	<Total Octets> <-Per second-> Output	Avg Q Len	<-Subnet packets / Sec-> <-Unicast-> <NonUnicast>			
					Input	Outpt	Input	Output
00:45:11								
***Node Groups***								
DB2	- 0	199K	233878	0	547.7	553.8	13.46	0
TheUsrs	- 0	2667K	<b>9284K</b>	0	4968	8432	178.3	0
*** Nodes *****								
Z144105D	- 1	1609	1608.7	0	6.77	6.8	0	0
	- 2	40.80	0	0	0	0	0.17	0
	- 3	94812	274917	0	587.3	588.1	0	0
	- 4	17762	<b>2409K</b>	0	339.0	1618	0.95	0
	- 5	3657	4010.0	0	22.55	22.6	0.91	0
	- 6	1967	18460	0	10.07	17.0	1.33	0
Z144106D	- 1	1147	1147.2	0	4.98	5.0	0	0
	- 2	445.8	330.43	0	1.16	1.2	0.91	0
	- 3	40.80	0	0	0	0	0.17	0
	- 4	477.9	433.33	0	8.06	4.1	0	0
	- 5	161.8	1236.9	0	0.29	1.0	0.91	0
	- 6	1776	16303	0	9.59	15.4	1.25	0

## Channel utilization on OSA is high, why?

Report: ESAQDIO      Queued I/O R      CENTERS FOR MEDICAID AND

---

Date/ Time	Dev. Nmbr	owner	Virt DevN	QDIO Fmt	structions/Sec->			<-Throughput / sec->				
					>	<-----CP----->	"s"	<Buffers>		<--Bytes-->		
					Read	Writ		Sent	From	Sent	From	
00:45:11	0000	Totals	0000	QDIO	6	5.6	1805	0	12K	6512	158M	2140K
	8700	DTCVSW4	063C	QDIO	0	0	0.2	0	0	0	175	
	9009	DTCVSW3	0630	QDIO	0	0	396	0	445	410	158M	1580K
	9406	DTCVSW1	061E	QDIO	0	0	179	0	205	190	5191K	8055
	9409	DTCVSW2	0618	QDIO	0	0	101	0	3652	376	6306K	78998

Report: ESAQDI2      Queued I/O Report

---

Date/ Time	Dev. Nmbr	owner	Virt DevN	QDIO Fmt	Number		<-Throughput / sec->			
					Queues	In Out	<Buffers>		<--Bytes-->	
					In	Out	Sent	From	Sent	From
00:45:11	0000	Totals	0000	QDIO	0	0	12K	6512	158M	2140K
	8700	DTCVSW4	063C	QDIO	1	1	0	0	175	
	9009	DTCVSW3	0630	QDIO	1	1	445	410	158M	1580K
	<b>9406</b>	<b>DTCVSW1</b>	<b>061E</b>	<b>QDIO</b>	<b>1</b>	<b>1</b>	<b>205</b>	<b>190</b>	<b>5191K</b>	<b>8055</b>
	<b>9409</b>	<b>DTCVSW2</b>	<b>0618</b>	<b>QDIO</b>	<b>1</b>	<b>1</b>	<b>3652</b>	<b>376</b>	<b>6306K</b>	<b>78998</b>

## Spread load

- across multiple channels, RANKS
- across devices
- across internal SAP paths

Define many devices 3390-3, not few 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