



z/OS Measurements Basic to Advanced Series: I/O: Response Time and Beyond

Scott Chapman

Enterprise Performance Strategies, Inc.

Scott.Chapman@EPStrategies.com



Contact, Copyright, and Trademarks



Questions?

Send email to performance.questions@EPStrategies.com, or visit our website at <https://www.epstrategies.com> or <http://www.pivotor.com>.

Copyright Notice:

© Enterprise Performance Strategies, Inc. All rights reserved. No part of this material may be reproduced, distributed, stored in a retrieval system, transmitted, displayed, published or broadcast in any form or by any means, electronic, mechanical, photocopy, recording, or otherwise, without the prior written permission of Enterprise Performance Strategies. To obtain written permission please contact Enterprise Performance Strategies, Inc. Contact information can be obtained by visiting <http://www.epstrategies.com>.

Trademarks:

Enterprise Performance Strategies, Inc. presentation materials contain trademarks and registered trademarks of several companies.

The following are trademarks of Enterprise Performance Strategies, Inc.: **Health Check[®], Reductions[®], Pivotor[®]**

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries: IBM[®], z/OS[®], zSeries[®], WebSphere[®], CICS[®], DB2[®], S390[®], WebSphere Application Server[®], and many others.

Other trademarks and registered trademarks may exist in this presentation

Abstract (why you're here!)



We don't have I/O problems anymore because we all have fast solid-state storage, right? Well... maybe not. In this webinar, Scott Chapman will discuss the traditional I/O measurements commonly surfaced in the RMF data, discuss the “hidden” measurement that impacts the actual response time, and finally look at the I/O measurements from other SMF records to get a deeper understanding of why I/O may be impacting some workloads despite an overall good I/O response time measurement.


AI Disclosure




AI Level 0 – AI was not involved in the creation of this content, other than possibly for spelling or grammar checking.



AI Level 1 – AI contributed to non-materially parts of this content, such as transcription, formatting, or clip art. All analysis and conclusions were human-authored.



AI Level 2 – AI assisted in the drafting and analysis. All findings and figures reviewed and confirmed by a human.



AI Level 3 – This was substantially produced by AI and reviewed by a human.

EPS: We do z/OS performance...



- Pivotor - Reporting and analysis software and services
 - Not just reporting, but analysis-based reporting based on our expertise
- Education and instruction
 - We have taught our z/OS performance workshops all over the world
- Consulting
 - Performance war rooms: concentrated, highly productive group discussions and analysis
- Information
 - We present around the world and participate in online forums
<https://www.pivotor.com/content.html>

Like what you hear today?



- Free z/OS Performance Educational webinars!
 - Let us know if you want to be on our mailing list for these webinars
 - Email: contact@epstrategies.com
- If you want a free cursory review of your environment, let us know!
 - We're always happy to process a day's worth of data and show you the results
 - See also: <http://pivotor.com/cursoryReview.html>

z/OS Performance workshops available



During these workshops you will be analyzing your own data!

- Essential z/OS Performance Tuning
 - March 30 – April 3, 2026 (4 days, excl Wednesday the 1st)
- WLM Performance and Re-evaluating Goals
 - July 27 – 31, 2026 (4 days, excl Wednesday the 29th)
- Parallel Sysplex and z/OS Performance Tuning
 - May 12-13 2026
- Also... please make sure you are signed up for our free monthly z/OS educational webinars! (email contact@epstrategies.com)



Basic

RMF Device Activity



- RMF Device Activity report seems like a good place to start

DIRECT ACCESS DEVICE ACTIVITY

PAGE 1

z/OS V2R4

SYSTEM ID SYSP
RPT VERSION V2R4 RMF

DATE 03/14/2023
TIME 22.15.00

INTERVAL 15.00.028
CYCLE 1.000 SECONDS

TOTAL SAMPLES = 900			IODF = 11		CR-DATE: 03/17/2022		CR-TIME: 15.02.14		ACT: ACTIVATE											
STORAGE GROUP	DEV NUM	DEVICE TYPE	NUMBER OF CYL	VOLUME SERIAL	PAV	LCU	DEVICE ACTIVITY RATE	AVG RESP TIME	AVG IOSQ	AVG CMR DLY	AVG DB DLY	AVG INT DLY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	% DEV CONN	% DEV UTIL	% DEV RESV	AVG NMBR ALLOC	% ANY ALLOC
DB2PROD3	04001	33903	3339	DBACA0	1.0H	0024	0.271	.118	.000	.000	.000	.004	.050	.002	.066	0.00	0.00	0.0	6.0	100.0
DB2PROD3	04002	33903	3339	DBACA1	1.0H	0024	0.049	3.76	.000	.000	.000	.003	.064	3.63	.067	0.00	0.02	0.0	6.0	100.0
DB2PROD3	04003	33903	3339	DBACB2	1.0H	0024	0.001	.128	.000	.000	.000	.000	.128	.000	.000	0.00	0.00	0.0	8.0	100.0
DB2PROD3	04004	33903	3339	DBACCC	1.0H	0024	0.078	.666	.000	.000	.000	.002	.053	.521	.091	0.00	0.00	0.0	11.0	100.0
DB2PROD3	04005	33903	3339	DBACD8	1.0H	0024	0.010	.242	.000	.000	.000	.000	.057	.114	.071	0.00	0.00	0.0	7.0	100.0
DB2PROD3	04006	33903	3339	DBACE5	1.0H	0024	0.142	.163	.000	.002	.000	.008	.053	.050	.060	0.00	0.00	0.0	11.0	100.0
DBABIGLD	04029	33903	3339	DBADAA	1.0H	0024	0.001	.128	.000	.000	.000	.000	.128	.000	.000	0.00	0.00	0.0	4.0	100.0
DBABIGLD	0402A	33903	3339	DBADAB	1.0H	0024	15.504	.201	.000	.001	.000	.004	.055	.083	.064	0.10	0.23	0.0	5.7	100.0
DBABIGLD	0402B	33903	3339	DBAG2E	1.0H	0024	0.001	.128	.000	.000	.000	.000	.000	.000	.128	0.00	0.00	0.0	0.0	100.0
DBABIGLD	0402C	33903	3339	DBAG4A	1.0H	0024	0.001	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	1.0	100.0
DBABIGLD	0402D	33903	3339	DBAG77	1.0H	0024	2.415	.639	.000	.002	.000	.001	.073	.311	.255	0.06	0.14	0.0	1.0	100.0
DBABIGLD	0402E	33903	3339	DBAH8E	1.0H	0024	0.001	.256	.000	.000	.000	.000	.128	.000	.128	0.00	0.00	0.0	2.0	100.0
DBABIGLD	0402F	33903	3339	DBAH9A	1.0H	0024	0.001	.256	.000	.000	.000	.000	.128	.000	.128	0.00	0.00	0.0	1.0	100.0
DBABIGLD	04030	33903	3339	DBAJA0	1.0H	0024	0.001	.128	.000	.000	.000	.000	.000	.000	.128	0.00	0.00	0.0	0.0	100.0
DBABIGLD	04031	33903	3339	DBAJA2	1.0H	0024	0.001	.128	.000	.000	.000	.000	.128	.000	.000	0.00	0.00	0.0	0.0	100.0
DBABIGLD	04032	33903	3339	DBAJCC	1.0H	0024	0.001	.256	.000	.000	.000	.000	.128	.000	.128	0.00	0.00	0.0	1.0	100.0

Device activity: what can we learn?




- Activity rate - how many I/Os
- Response time, broken down by the 4 historical components:
 - IOSQ – I/O Supervisor Queue (waiting on UCB, expect 0)
 - PEND – Pending for (e.g.) channel, CU, LV, director
 - CMR (command response) and DEVB (device busy) components of this
 - CONN – Connect time (data transfer time + protocol time)
 - Can be impacted by I/O to other volumes sharing the channel
 - DISC – Disconnect (generally waiting in the controller)
- Average Interrupt Delay
 - Time between I/O operation complete and z/OS begins processing it
 - Impacted by LPAR dispatchability (e.g. expect more on smaller, low-weight LPARs)
 - Notably not included in the response time total

How much and how fast are the two things people are usually looking for.

DASDplex RT Components

Including I/O rate

ZKXXPLEX



100s to 1000s of volumes per interval make the RMF device activity report very difficult to use for understanding overall I/O activity and health.


Pivotor starts at the total I/O activity across all the systems and control units to give an overview of the total I/O activity and responsiveness.

Details that may be surprising



- RMF stores I/O response times in 128-microsecond increments
 - 0.128ms when average response times are frequently under 0.5ms.
 - Can impact accuracy of results for volumes with very few I/Os
- Interrupt delay not included in the calculated response times
 - To be fair, this time is not the “fault” of the I/O subsystem
 - But from an application perspective, it is part of how long an I/O takes
 - But “hidden” in traditional I/O RT reporting
 - For some LPARs it can be a significant portion of the overall time
 - Can be impacted by LPAR weight, number of CPs, CPEENABLE setting (mostly prior to z/OS 3.1)
- IOSQ time should generally be very near zero today with SuperPAV
 - Some IOSQ can still show up either due to peaking or due to replication things


Average Response Time versus Interrupt Delay Time




At times on this system interrupt delay is significant relative to the response time from the I/O subsystem.

In this case, the interrupt delay time corresponded to times when the LPAR was capped.

CEC Capping Actually Limited Percentage vs. Percent Considered Capped



Capping for that LPAR



If the LPAR is capped, it won't be dispatched at times and so won't be able to handle those I/O interrupts.


It takes CPU to process the I/O, and if you have less CPU, that will impact the I/O processing.

That doesn't show up in the traditional I/O response time but does impact the application.

WLM DASD - I/O SSCH Rate by Service Class Period Over Time



ZKXXPLEX, SYSJ (1 of 2)




The SMF 72s let us see I/O rates and response times by service class period.

Handy for a high-level view of what is responsible for doing I/O.




WLM DASD - I/O Response Time by Service Class Period Over Time



Probably those extreme response times are for SCPs with very little activity.



WLM DASD - I/O Response Time by Service Class Period Over Time




Zoomed in, most of the SCPs are showing average RT below 0.2ms



Intermediate

PCU - zHPF Read/Write I/O per Sec




This is fairly typical: most of the I/O being done is read, not writes.

Also: you should be using zHPF! (Most are at this point, but we still sometimes come across somebody who has ZHPF=NO in IECIOSxx)

PCU - Average Read Miss Ratios (Total Reads)

(All Storage Subsystems)

ZKXXPLEX




And most of those reads come from cache.

These cache details are in the 74.5 which records information from the DASD subsystem itself.

zHyperlink Read Summary

For Each System

SYSJ



This system is doing millions of I/Os per hour that could be eligible for zHyperlink, which may take only 0.012-0.016ms.

Compared to the 0.1-0.2ms response times we saw earlier, that sounds amazing: at least 10x faster!

zHyperlink details available in multiple places, here we're looking at SMF 42.6 records.

What is zHyperLink?



- zHyperLink introduced back in 2017 to improve response time for I/O satisfied out of the DS8K cache
 - Targets ~20 microsecond (0.02 milliseconds) response time
 - Processor “spins” waiting for response, much like CF sync request
 - Requires zHyperLink adapter cards
 - Various technical requirements, initially not write, only recently for >4K I/Os
- Potential benefits:
 - Even faster I/O
 - CPU optimization because don’t have to redispatch after I/O
- Potential cost:
 - CPU overhead due to “spinning” waiting for the I/O to complete
- We’re starting to see more customers experiment with zHL
 - Long time coming for multiple reasons including hardware lifecycle timelines

So... should you zHyperLink?



- Maybe... do you have a I/O problems?
 - If you think so, can you articulate the business cost of that problem?
 - If you don't or can't: why do you want to complicate your life?
- Can you solve your problems with memory?
 - Memory access on the order of 15-100x faster than zHL
 - You may have memory on your LPARs/CEC that you're not using
 - You can put more memory on a CEC than on a DS8K
 - 512GB = Max DS8A10 cache = Min z17 ME1 orderable memory
- If you can't... are your problematic I/Os zHL eligible?
 - Must be 4K prior to latest DS8K (gen 10)
 - Various details around replication
- Uniquely addressable with zHL: Db2 log write bottlenecks

Max DS8K is ~4.5 TB
Max z17 is 64 TB



Advanced

I/O Measurement Difficulties



- 15-minute averages can be misleading
 - Average response time of 0.2ms clearly doesn't mean all I/O took 0.2ms
 - How bad could it really be?
- I/O by Service Class is helpful but is still a rather coarse measurement of what work is driving the I/O
 - Would be nice to see I/O at a more detailed level
- Even for I/O coming from production Service Classes, do we care about all the I/Os equally?
 - Probably not: the application isn't necessarily waiting on every I/O
 - Async read-ahead and async writes are common

For at least the first two, the SMF 42s can provide significant detail

Dataset High I/O Response Time Observations

ZKXXPLEX

Freeze: 4 Filter: sysplex = Value Apply Clear All Show: Averages 3035 Initial rows



sysplex	Usage	DS Name	sc name	int end date	int end time	Max RT ms	Avg RT ms	Total RT	I/O Count	Read MiB	Write MiB
3035						159147.9	11889.9	36857803.4	12637554	43233...	100878
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-06	11:23:32	1,092.352	1.147	127,756.301	111,383....	198.254	60%
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-06	13:15:42	1,046.656	0.702	15,491.736	22,068.000	39.242	11%
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-06	11:15:00	1,040.896	0.892	73,242.120	82,110.000	166.688	50%
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-06	18:12:09	1,033.344	0.431	83,488.579	19		
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-07	07:17:27	1,033.344	0.325	62,307.375	19		
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-06	20:31:45	1,026.944	0.416	47,629.920	11		
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-06	19:25:58	1,026.048	0.852	164,533.980	19		
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-06	21:07:46	1,025.280	0.468	90,610.416	19		
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-06	14:38:59	1,022.336	0.674	130,148.052	19		
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-06	13:15:00	1,021.440	0.713	122,249.554	17		
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-06	15:50:58	1,014.784	0.664	115,384.608	17		
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-07	08:52:06	1,014.656	0.354	68,674.938	19		
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-06	20:30:00	1,007.616	0.525	41,311.200	78		
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-07	01:45:00	1,006.976	0.546	101,616.060	18		
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-07	07:36:15	1,004.544	0.251	48,360.170	19		
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-07	06:39:22	1,004.160	0.410	78,692.530	19		
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-06	16:54:24	1,003.904	0.354	68,615.112	19		
ZKXXPLEX	CATALOG	SYS1.VOLCAT.VGENERAL	SYSTEM	2026-10-06	21:00:00	1,003.648	11.268	1,036.656			
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-06	15:45:00	1,003.264	0.734	14,315.202	19		
ZKXXPLEX	CATALOG	SYS1.VOLCAT.VGENERAL.CATINDEX	SYSTEM	2026-10-07	05:45:00	1,002.624	0.316	1,768.020	5		
ZKXXPLEX	OTHER-L...	SYS3.SS2A.DSSP.REMOTE.ACT.V11.DATA	SYSSTC	2026-10-07	01:45:34	1,001.216	0.979	7,493.266	7,654.000	13.512	41
ZKXXPLEX	OTHER-P...	SYS3.TECH.ISPPLIB	TSO	2026-10-06	15:15:00	984.192	127.869	1,918.035	15.000	0.125	(
ZKXXPLEX	TEMP-PS	SYS26279.T223940.RA000.NF050610.DLYSORT....	BATMD	2026-10-07	03:39:41	559.872	559.951	559.951	1.000	0.000	(
ZKXXPLEX	OTHER-P	SYS3.TECH.ISPPLIB	TSO	2026-10-06	15:30:00	549.632	19.807	554.596	28.000	0.499	(


The SMF 42.6 records include maximum response times which can be interesting.

That average RT of a fraction of a ms is immaterial compared to the I/Os that take 1 second to complete!

I/O Interrupt Delay Maximums (UTC Time)

For maximums > 1ms

SYSJ




I/O interrupt delays
sometime extend to
more than a second too,
even when the system
wasn't capped.

(These come from the
42.5 records)

I/O Interrupt Delay Maximums (UTC Time)

For maximums > 1ms


SYSJ



It seems like there may be something to learn from these interrupt delay patterns: probably represents bursts of activity.

But may be worth looking at if you're chasing a transient performance problem.

System I/O Interrupt Delay Distribution



Even better the 42.5 provides us with a distribution of the interrupt delays.

I'd consider all that yellow to be "significant" relative to our expected I/O response times.



Top Datasets by Cache Read Hits

ZKXXPLEX

Freeze: 4 Filter: sysplex = Value Apply Clear All Show: Averages 200 Initial rows

sysplex	Usage	DS Name	M Read Ops	M Cache Read Hits	Allocated MiB	Read MiB	Read-Allocated Ratio	M Write Ops	Cache Hit Pct	Read Hit RT secs	4 R
200			477	209	753626	40883909	46100	8	18493	47758	
ZKXXPLEX	DB2-OBJ	VS62.DSNCDBD.DBZCPA.IX20460A.J0001.A001	68.60	0.99	2,361	8,128,659	3,442.5	0.00	98.8	208.8	
ZKXXPLEX	DB2-OBJ	VS66.DSNCDBD.DBZAAD.TS17613.I0001.A001	62.41	15.05	511	6,599,041	12,922.4	0.00	99.0	3,804.8	
ZKXXPLEX	DB2-OBJ	VS50.DSNCDBD.DBZABW02.TS06435.J0001.A...	36.32	1.42	876	4,276,064	4,880.0	0.00	87.9	614.0	
ZKXXPLEX	DB2-OBJ	VS62.DSNCDBD.DBZCPA.IX20460A.J0001.A002	27.47	0.42	993	3,313,556	3,337.0	0.00	99.6	77.0	
ZKXXPLEX	DB2-OBJ	VS75.DSNCDBD.DBZAAD.TS17613.I0001.A001	22.54	10.59	511	1,921,258	3,762.2	0.00	99.4	2,353.7	
ZKXXPLEX	DB2-OBJ	VS49.DSNCDBD.DBZAEG.TS07292.J0001.A001	22.26	21.85	4,721	152,764	32.4	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS50.DSNCDBD.DBZAM4.TS06874.J0001.A001	19.13	18.52	4,722	586,103	124.1	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS49.DSNCDBD.DBZAEG.TS07292.J0001.A002	18.88	18.57	3,774	145,002	38.4	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS50.DSNCDBD.DBZAEG.TS07288.I0001.A001	10.19	9.79	4,829	322,908	66.9	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS50.DSNCDBD.DBZAEG.TS07292.J0001.A001	10.14	9.75	4,722	88,369	18.7	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS50.DSNCDBD.DBZAEG.TS07288.I0001.A002	7.87	7.69	3,280	249,257	76.0	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS50.DSNCDBD.DBZAEG.TS07292.I0001.A002	7.12	6.88	3,393	73,496	21.7	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS67.DSNCDBD.DBZAIR.IRFSAUD1.I0001.A002	6.69	1.06	4,722	706,411	149.6	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS67.DSNCDBD.DBZAIR.IRFSAUD1.I0001.A003	6.68	1.05	4,722	706,722	149.7	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS67.DSNCDBD.DBZAIR.IRFSAUD1.I0001.A001	6.60	1.06	4,722	655,733	138.9	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS62.DSNCDBD.DBZMEW.TS20158.I0001.A001	5.71	0.34	481	672,017	1,395.7	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS50.DSNCDBD.DBZATI.TS08957.J0001.A004	4.49	4.09	4,175	314,552	75.3	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS50.DSNCDBD.DBZABK.TS01451.J0001.A001	3.98	3.15	353	92,967	263.7	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS50.DSNCDBD.DBZATI.TS08957.J0001.A003	3.93	3.50	4,722	301,042	63.7	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS50.DSNCDBD.DBZAM4.TS06874.I0001.A002	3.67	3.60	570	77,944	136.8	0.04	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS50.DSNCDBD.DBZMBP.TS20684.I0001.A009	3.31	3.06	4,829	196,071	40.6	2.47	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS50.DSNCDBD.DBZAET.TS07315.I0001.A001	3.24	2.46	3,774	255,218	67.6	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS50.DSNCDBD.DBZAHT.TS02809.I0001.A009	3.18	2.53	4,134	96,868	23.4	0.00	98.8	2,603.5	
ZKXXPLEX	DB2-OBJ	VS50.DSNCDBD.DBZATI.TS08957.I0001.A002	3.15	2.79	4,722	223,058	47.2	0.00	98.8	2,603.5	

The SMF 42.6 data also can help us find datasets that appear to be read-only, get read a lot, and are relatively small so maybe could be held in memory.

The top 5 in this case are 215 million I/Os but only a little over 5 GB of data.

Why not cache this data and avoid doing those I/Os??

If you want to go even further...



- I/O related information is all over in the SMF records
 - Type 14/15 – “Old” DD-related I/O
 - Type 30 – Summary I/O at job/step
 - Type 42 – Volume and dataset level I/O
 - Type 64 – VSAM Status
 - Type 72 – I/O by service class/report class
 - Type 74 – Volume level I/O details
 - Type 75 – I/O by page dataset
 - Type 100-102 – Various Db2 details, including Db2 I/O (depending on config)
 - Type 110 – CICS details, including CICS I/O
 - Others – Sort, Vendor-specific DASD measurements, etc.

We have more I/O presentations...



https://www.pivotor.com/content.html



Enterprise Performance Strategies Inc.
Creators of Pivotor®

Home Pivotor Workshops Consulting Webinars Tools & Resources About

EPS Papers and Presentations

Peter and Scott present on and write about many mainframe performance oriented topics. Some of their "most hit" papers are found below.

Click on a title to see the abstract for the presentation.

Click on the "Download" button to access the presentation.

BONUS - Click [here](#) to watch Peter Enrico's "Not a Legacy but a Legend" interview with Planet Mainframe.

Year view

Topic view

Performance Pointers

Download

Peter Enrico and Scott Chapman - z/OS Performance Topics: The Ones That Divide Us

Switch to
topic view

Then scroll
down to I/O

I/O

Download

Nancy DiFilippo - What I learned about VSAM RLS SMF Measurements

Download

Scott Chapman - z/OS I/O Performance: Do You Have a Problem?

Download

Scott Chapman and Larry Strickland - Optimizing Performance at the Speed of Light: Why I/O Avoidance is Even More Important Today

Download

Scott Chapman - SHARE - I/O, I/O It's Home to Memory We (Should) Go

Download

Scott Chapman - Key Reports to Evaluate Usage of Parallel Access Volumes

Download

Peter Enrico - Exploring z/OS SMF 14/15 Records for Tape and DASD File Activity

Pivotor



Wrap-up

Summary



- I/O rates and response times are the base everybody looks for
- But response time averages can hide significant outliers
- And response time doesn't traditionally include interrupt delay time
 - But that can be quite significant relative to the I/O time
 - So I/O may be even more expensive than you think for that reason too
- SMF 42s can provide significant details important to finding what specific workloads / datasets may have opportunities for improvement
- The only good I/O is still no I/O
 - Look to see what high read-I/O datasets you might be able to pin in memory



Questions?