



Intensity!

Understanding the Concepts and Usage of Intensity Measurements

Peter Enrico



z/OS Performance
Education, Software, and
Managed Service Providers



Creators of Pivotor®

Email: Peter.Enrico@EPStrategies.com

Enterprise Performance Strategies, Inc.

3457-53rd Avenue West, #145

Bradenton, FL 34210

<http://www.epstrategies.com>

<http://www.pivotor.com>

Voice: 813-435-2297

Mobile: 941-685-6789



Contact, Copyright, and Trademark Notices



Questions?

Send email to Peter at Peter.Enrico@EPStrategies.com, or visit our website at <http://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**[®]

Other trademarks and registered trademarks may exist in this presentation

Abstract



- **Intensity!**

- **Understanding the Concepts and Usage of Intensity Measurements**

- In this webinar Peter Enrico will discuss the concept and usage of calculated intensity measurements. When there is many of a particular resource or workload to be evaluated, intensity is a calculated measurement to help highlight which of the many are impacting performance the most. For example, say you have a DASD farm with 3000 logical volumes - certain intensity measurements can help prioritize which of these many logical volumes are impacting performance the most.

Like what you see?



- Free z/OS Performance Educational webinars!
 - The titles for our Fall 2022-2023 webinars are as follows:
 - ✓ *Key Reports to Evaluate z16 Processor Caches*
 - ✓ *Understanding System Recovery Boost's Impact on Performance and Performance Reporting*
 - ✓ *WLM Management of DDF Work: What can you do and what has changed?*
 - ✓ *Intensity! Understanding the Concepts and Usage of Intensity Measurements*
 - *High, Medium, Low: Understanding how HiperDispatch influences performance in z/OS*
 - *Putting a lid on XCF*
 - *How and why Pivotor is different than other performance management reporters*
 - *Key Reports to Evaluate Usage of Parallel Access Volumes*
 - *Key Reports to Evaluate Coupling Facility CPU Utilization*
 - *Understanding how memory management has evolved in z/OS*
 - Let me know if you want to be on our mailing list for these webinars
- 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>

EPS: We do z/OS performance...



- Performance reporting and analysis of your SMF data
 - Pivotor: Not just reporting, but analysis-based reporting based on our expertise
- Performance Educational Workshops (while analyzing your own data)
(All presented via the web)
 - Essential z/OS Performance Tuning
 - March 20 – 24, 2023
 - Parallel Sysplex and z/OS Performance Tuning (2 days)
 - April or May of 2023
 - WLM Performance and Re-evaluating Goals
 - Dates TBD
- Performance Consulting
 - Performance war rooms: concentrated, highly productive group discussions and analysis
 - MSU reduction services

Like what you see?



- The z/OS Performance Graphs you see here come from Pivotor™
- If you don't see them in your performance reporting tool, or you just want a free cursory performance 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>
- We also have a **free** Pivotor offering available as well
 - 1 System, SMF 70-72 only, 7 Day retention
 - That still encompasses over 100 reports!

All Charts (132 reports, 258 charts)

All charts in this reportset.

Charts Warranting Investigation Due to Exception Counts (2 reports, 6 charts, [more details](#))

Charts containing more than the threshold number of exceptions

All Charts with Exceptions (2 reports, 8 charts, [more details](#))

Charts containing any number of exceptions

Evaluating WLM Velocity Goals (4 reports, 35 charts, [more details](#))

This playlist walks through several reports that will be useful in while conducting a WLM velocity goal an.

Basic metrics typically used during an analysis



- Metrics Data

- What is typically thought of as the measurement to gauge the relative value of during an analysis or study
- One should think of metrics as any measured value that could be used to help gain an understanding needed to answer a question
- Many times, when reviewing measurements, metrics can be further defined
 - Cumulative / Interval
 - Snapshot
 - Sampled
 - Etc.
- Examples of metrics (in a computer measurement environment)
 - Counts
 - Averages
 - Rates
 - Utilizations
 - Percentages
 - Ratios
 - Intensities
 - Custom
 - Etc.

Common Analysis Problem



- When you have many of something, and are unsure which of the many to focus an analysis on
- What to focus on when...
 - There are thousands of logical I/O volumes
 - There are thousands of datasets / files
 - There are hundreds, if not thousands, of batch jobs
 - There are dozens of coupling facility structures
 - There are many transaction ids
 - There are dozens of WLM report classes
 - There are many channels in the channel subsystem
 - Etc...
- Would be nice to understand which ones are potentially impacting performance the most

Common Analysis Problem



- Classic Example:
 - Say we want to do a logical volume I/O analysis and have 3000 logical volumes
 - Which of the 3000 logical volumes are impacting performance the most?
 - The ones with the highest I/O rates?
 - The ones with the worst response times?
 - The ones with the worst delays (such as IOSQ, disconnect time, pend time)?
- Let's assume an acceptable I/O response time is less than 1 millisecond.

Which volume is probably impacting workload performance the most?

- Volume: ABC001 has an average RT=3 ms, and an I/O rate of 20 I/Os per second
- Volume: XYZ001 has an average RT=0.8ms, and an I/O rate of 200 I/Os per second



Introduction to Intensity

Definition of Intensity



- Intensity: A weighted value that is a measure to gain insight into what components / workloads to concentrate an analysis on
 - The greater the intensity, the potentially greater the impact and interest
- Basic Formula : $Intensity = (Rate) * (Average Time)$
- Common when there are many components being measured, but want to focus on the most interesting ones that are impacting performance the most
- Intensity is really just an average total per (typically) second.
 - E.G. average rt in MS * average I/Os/sec = total ms RT per second (on average)
- Examples of use:
 - I/O intensities and Queuing intensities
 - Helps determine which logical volumes or LCUs or files are impacting I/O performance the most
 - Coupling facility structure intensities
 - Helps determine the structures impacting performance the most
 - Batch job intensities
 - Finding which batch jobs are impacting performance of the stream
 - DB2 query intensities
 - Etc.



Examples of Intensity

Looking at DASD I/O performance

DASD I/O Intensities




- When it comes to I/O many analysts make the mistake on concentrating solely on the following:
 - Poor I/O response times
 - Logical volumes with the most activity
- It is useful to also look at the calculated *Intensity* values
 - The greater the intensity, the greater the potential impact
- Useful formulas include:
 - I/O Intensity (Int) = (I/O Rate) * (Avg CONN + Avg DISC + Avg PEND + Avg IOSQ)
 - Queuing Intensity (QInt) = (I/O Rate) * (Avg DISC + Avg PEND + Avg IOSQ)
 - *This value is particularly useful because it focuses on delays*
 - Connect Intensity (CONNInt) = (I/O Rate) * (Avg CONN time)
 - Disconnect Intensity (DISCInt) = (I/O Rate) * (Avg DISC time)
 - Pend Intensity (PENDInt) = (I/O Rate) * (Avg PEND time)
 - IOSQ Intensity (IOSQInt) = (I/O Rate) * (Avg IOSQ time)

Earlier Example of Concept of Intensity



- Intensity = (I/O Rate) * (Average Response Time)
- Going back to our previous example:
 - Volume: ABC001 has an average RT=3 ms, and an I/O rate of 20 I/Os per second
 - $I/O\ intensity = 3ms\ RT * 20\ I/O/sec = 60$
 - Volume: XYZ001 has an average RT=0.8ms, and an I/O rate of 200 I/Os per second
 - $I/O\ intensity = 0.8ms\ RT * 200\ I/O/sec = 160$
- Chances are that even though volume XYZ001 has a decent response time, it is probably impacting workload performance more than volume ABC001.
 - And to be fair, I/O analysis is always difficult because ideally one would consider I/O differently depending on whether the application is actually waiting for the I/O to complete or not.

Storage Controller Queuing Intensity

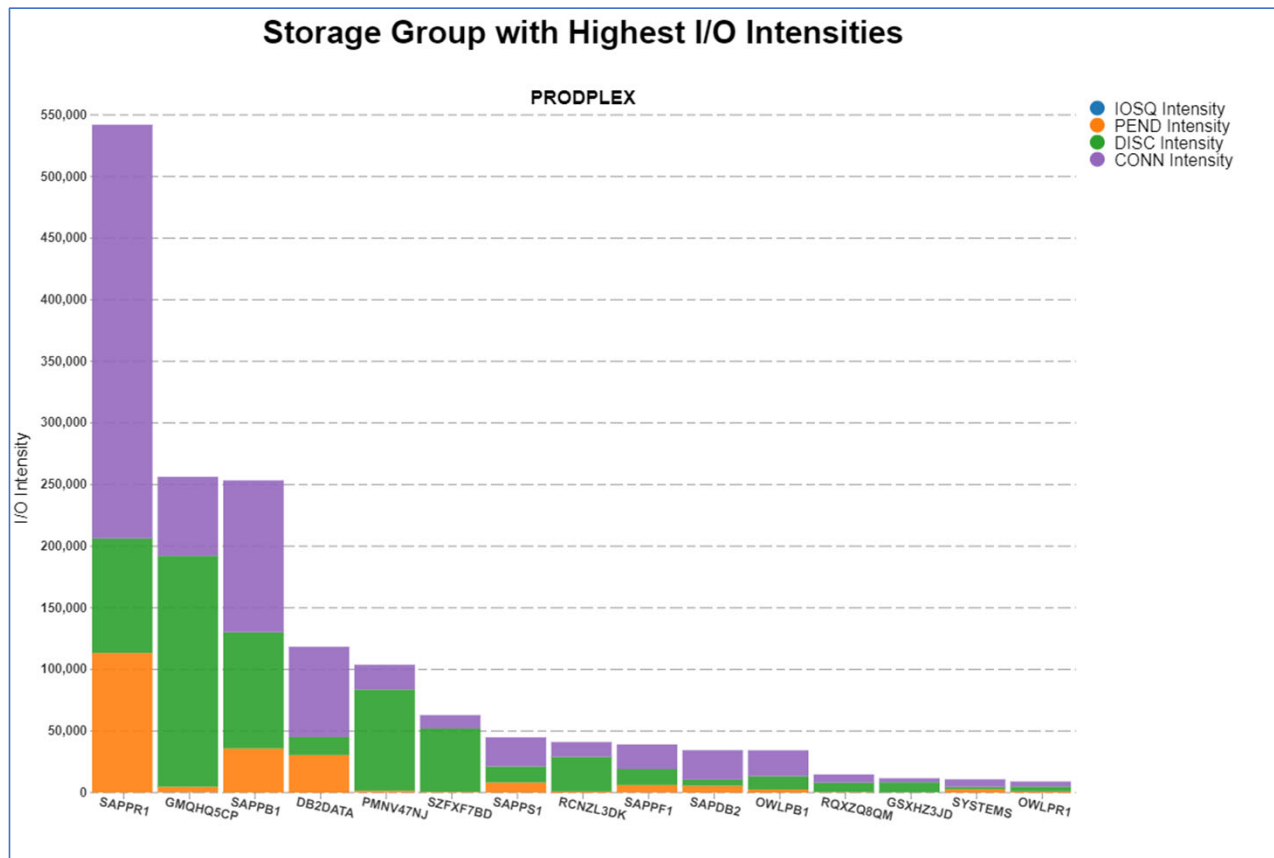


Queuing Intensity is useful to help determine periods of time to concentrate an I/O analysis on.

In this example, we see high I/O rates in the early morning hours.

But queuing intensity is greatest when I/O rates are lower, but response times are higher.

Looking top storage groups with highest intensities



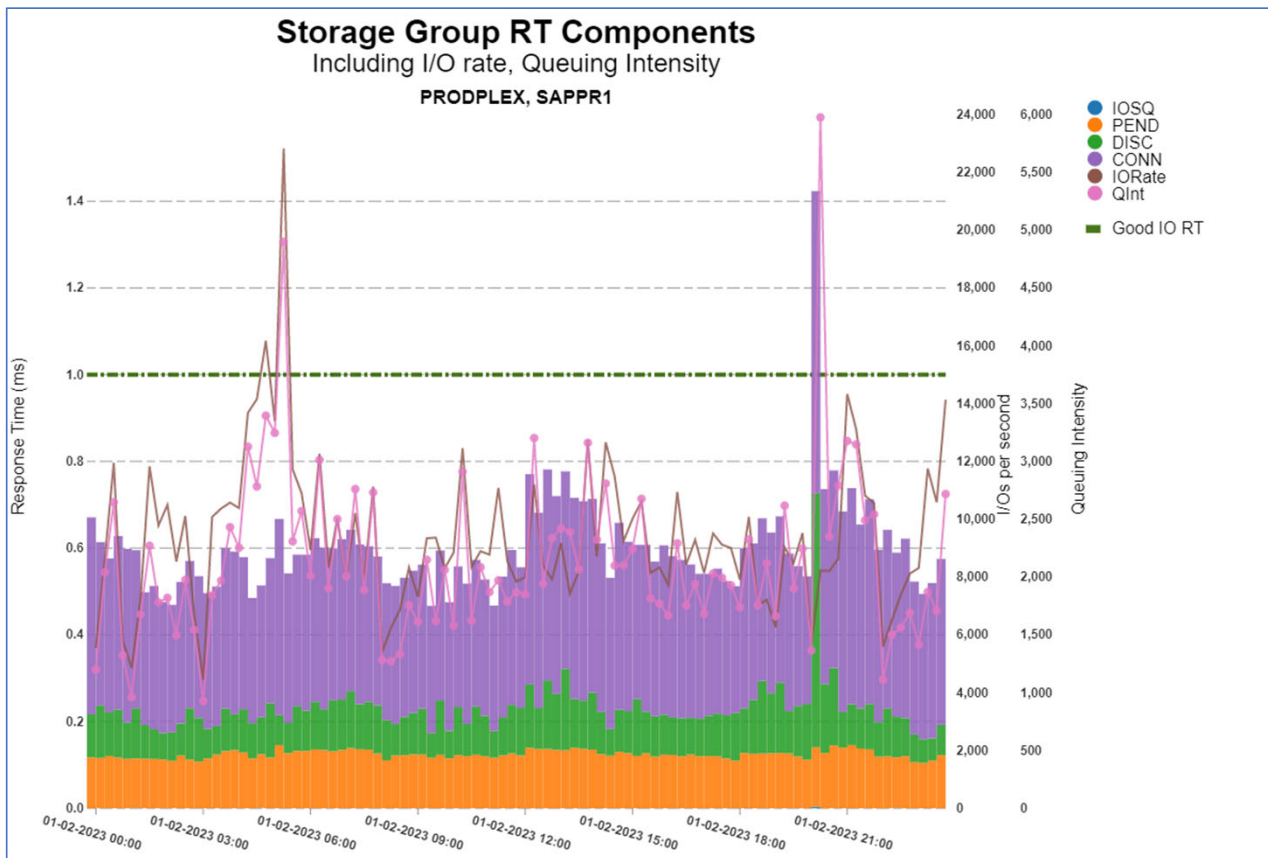
We can use intensities to determine which storage groups we may want to focus our analysis on.

In this example, the 4 primary I/O intensity values are reported.

Connect Intensity is probably less interesting since connect is a productive time.

But we now know which storage groups are impacting workload performance the most.

Looking at Qint for the top storage group



Queuing Intensity is useful to help determine periods of time to concentrate an I/O analysis on.

When calculated over time, we can see which periods of time when I/O is impacting the workloads the most.

Notice the two peak QInt points:

- One has good RT, but high I/O rate
- The other has poor RT, but low I/O rate

Both periods of time are of interest.

Analyzing Performance Data at the LV Level




- Summarize performance at the Logical Volume level
 - Summarize the measurements to obtain a *single* set of condensed numbers for each logical volume (LV).
 - Measurements exist separately for each LV in each LPAR in each interval
 - Consolidation (across LPARS or independent CPUs) means *one* set of numbers for *each* LV in *each* measurement interval
- There are many (typically thousands) of logical volumes to examine
 - Identify problem LVs
 - List of LVs must be narrowed down based on performance criteria
 - Our primary filter mechanism is Queuing Intensity, QI
 - LVs with the highest QI have bad impact because of the highest delays
- Select the worst volumes (by QI)
 - Examine cause for high QI (usually combination of high QDT and high S)
 - Analyze reasons: high DISC, high PEND, high IOSQ
 - Possibly high CONN (usually CONN is productive)

I/O Rates: Top Logical Volumes



LVs with Highest I/O Rates
(Averaged Over Period of Study)

PRODPLEX



This chart shows the logical volumes with the highest average I/O rates for a single day.

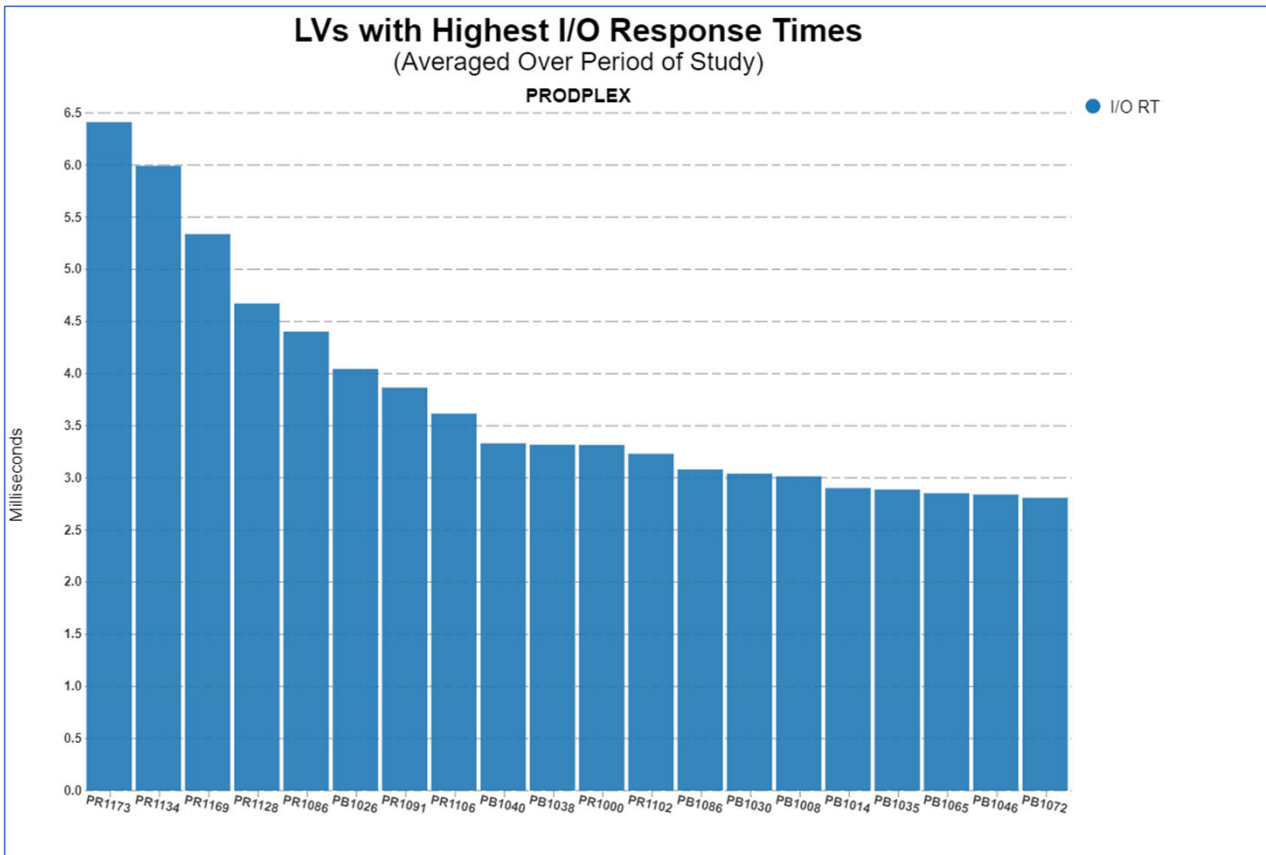
Yes, it is a useful exercise to examine the logical volumes with the highest I/O rates to understand the activity.

However, these may not be impacting the workload performance the most.

Top 4 logical volumes:

- PR1171
- PR1130
- PR1196
- DB2S12

I/O Response Times: Top Logical Volumes



This chart shows the logical volumes with the highest average I/O response times over a single day.


Yes, it is a useful exercise to examine the logical volumes with the highest I/O response to understand performance.

However, these may not be impacting the workload performance the most.

Top 4 logical volumes:

- PR1173
- PR1134
- PR1169
- PR1128

Queuing Intensity: Top Logical Volumes




This chart shows the logical volumes with the highest average Queuing Intensity over a single day.

These are the logical volumes probably impacting performance the most since both response time and activity rates are both factored in.

Top 4 logical volumes:



- PR1108
- PR1171
- PR1184
- PR1130

PR1108 : I/O Response Time Components



This chart shows the response time components of PR1108



PR1108 : I/O Response Time Components and Rates





PR1108 : I/O Response Time Components and Rates





The left-hand chart shows the average queuing components of the logical volume's response time.
The right-hand chart shows both the I/O rate and the average queuing intensity.
Can now investigate period of time of interest.



IOSQ Int and DISC Int: Top Logical Volumes



PEND Int and CONN Int: Top Logical Volumes





Examples of Intensity


Looking at Coupling Facility Structures

CF Structure Intensities



- When it comes to Coupling Facility analysis, many analysts make the mistake on concentrating solely on the following:
 - Poor structure CF response times
 - Structures with most activity
- It is useful to also look at the calculated *Intensity* values
 - The greater the intensity, the greater the potential impact
- Useful formulas include:
 - Structure Intensity = (Request Rate) * (Structure Response Time)
 - Sync (Spin) Intensity (CF_SyncInt) = (Sync Request Rate) * (Avg Sync Time)
 - Async Intensity (CF_AsyncInt) = (Async Request Rate) * (Avg Async Time)

Host effect CPU seconds




This chart shows the host effect CPU seconds over a 24-hour period of time.

Host effect represents the number of CPU seconds on the z/OS system to do coupling facility operations.

These CPU seconds translate to capacity being consumed, and real MSU towards billing.

Host effect CPU seconds: Top Structures



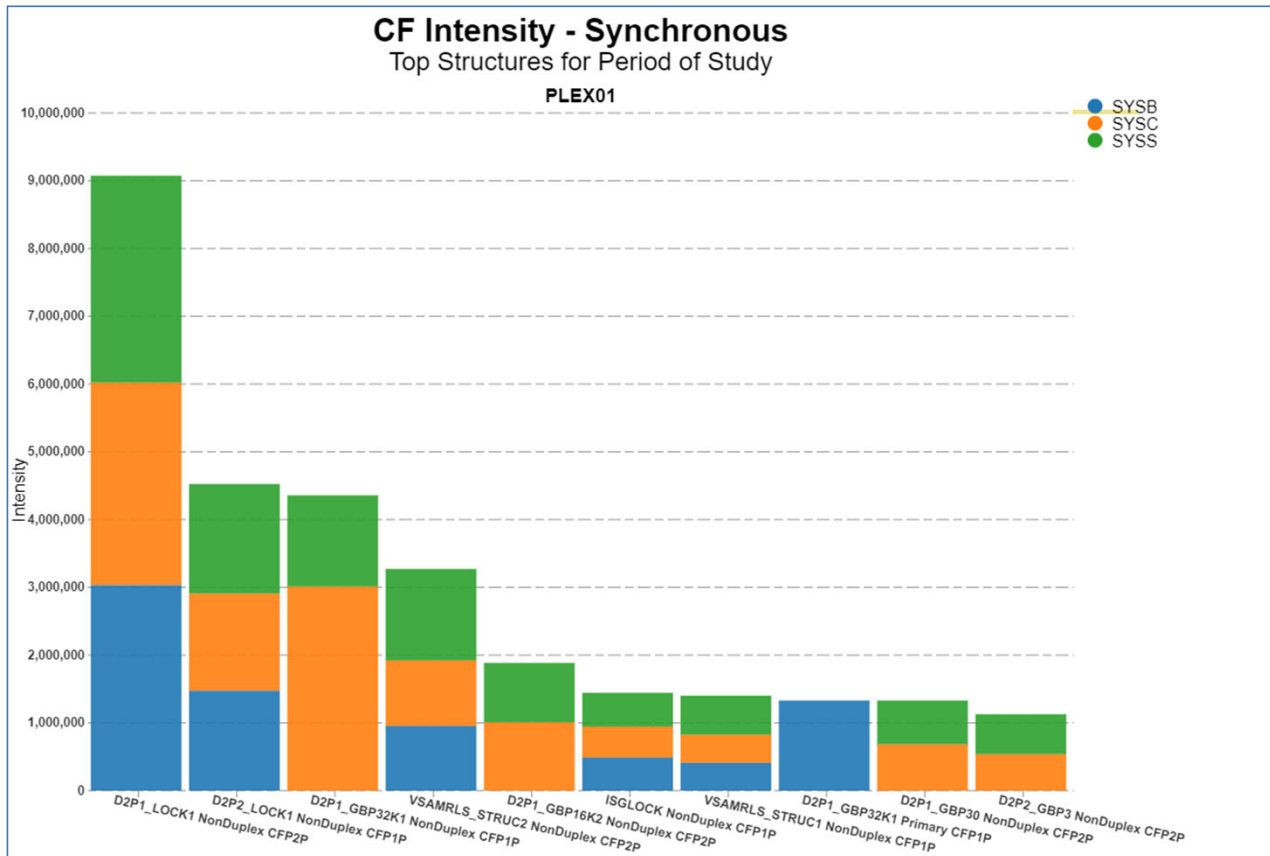
This chart shows the summed host effect CPU seconds for the top 10 structures.

Reported are the CPU seconds consumed on z/OS to do coupling facility operations to these structures.

These CPU consumption measurements and reports are incredibly valuable when selecting which structures to tune first.

Recommended for first analysis

CF Synchronous Intensity: Top Structures





This chart shows the structures with the highest synchronous intensities for 24-hours of study

Selecting structures to analyze based on sync intensity values is a great secondary analysis since it highlights structures that are impacting performance of the workload the most (rather than focusing on CPU consumption).


Typically, the production lock structures show up both here and on the top host effect reports.

However, most times other interesting structures are identified when using sync intensity.

D2P1_LOCK1 Lock Structure: Input into Sync Intensity



CF Asynchronous Intensity: Top Structures





This chart shows the structures with the highest asynchronous intensities for 24-hours of study

Selecting structures to analyze based on async intensity values is a great secondary analysis since it highlights structures that are impacting performance of the workload the most (rather than focusing on CPU consumption).

Typically, the XCF structures show up both here.

XCF List Structure IXCSML3: Input into Async Intensity



Summary



- Intensity is a useful measurement when you have many of something, and are unsure which of the many to focus an analysis on
- Intensity is just one more measurement in your arsenal when doing an analysis.
- Useful when analyzing when...
 - There are thousands of logical I/O volumes
 - There are thousands of datasets / files
 - There are hundreds, if not thousands, of batch jobs
 - There are dozens of coupling facility structures
 - There are many transaction ids
 - There are dozens of WLM report classes
 - There are many channels in the channel subsystem
 - Etc...



Comments from Jamie... and then Q & A

Questions about content of webinar?

Of maybe general performance questions?