



Improving Performance with Multiple Period Service Classes

Peter Enrico

Email: Peter.Enrico@EPStrategies.com



Enterprise Performance Strategies, Inc.
3457-53rd Avenue North, #145
Bradenton, FL 34210
<http://www.epstrategies.com>
<http://www.pivotor.com>

Voice: 813-435-2297
Mobile: 941-685-6789



Instructor: Peter Enrico

Enterprise Performance Strategies, Inc. ©

1

Enterprise Performance Strategies, Inc. ©

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**®

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

2



Abstract



• Improving Performance with Multiple Period Service Classes

- Service classes can have multiple periods. Why is this? Where does it make sense to use multiple periods? How should you set the period duration?
- In this webinar, you will learn when it makes sense to use multiple periods for a service class and how to set intelligent durations for the periods. Intelligent use of multiple period service classes can help overall system performance and throughput.

3

Enterprise Performance Strategies, Inc. ©


EPS: We do z/OS performance...






- We love to bring you these complimentary webinars and to teach you how to do what we do!
- But remember, we would love to work with you!
 - **Pivotor - Performance reporting and analysis of your z/OS measurements (i.e. SMF, etc.)**
 - Not just reporting, but coset effective analysis-based reporting based on our expertise
 - **Performance Educational Workshops (while analyzing your own data)**
 - Essential z/OS Performance Tuning
 - Parallel Sysplex and z/OS Performance Tuning
 - WLM Performance and Re-evaluating Goals
 - **Performance War Rooms**
 - Concentrated, highly productive group discussions and analysis

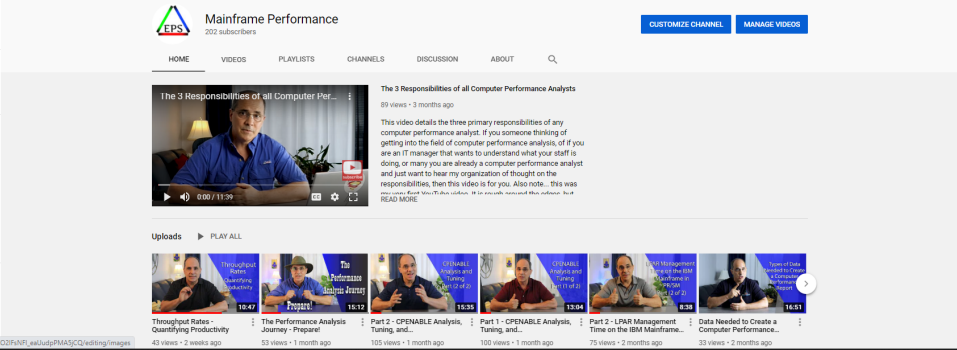
4




YouTube Mainframe Performance Channel







5

5

Enterprise Performance Strategies, Inc. ©

Like what you see?

- Free z/OS Performance Educational webinars!
 - The titles for our Fall 2020 webinars are as follows:
 - ✓ *WLM Response Time Goals Primer*
 - ✓ *Setting Response Time Goals for Modern Applications*
 - ✓ *Evaluating Latent Demand in the Mainframe Environment*
 - ✓ *Adjusting WLM Settings for Latent Demand*
 - ✓ *Improving Performance with Multiple Period Service Classes*
 - *Preparing for Any z/OS Performance Analysis*
 - *Evaluating Coupling Facility Lock Structures*
 - *Exploring Coupling Facility Exploitation by VSAM RLS*
 - *Data in Memory (DIM) Primer*
 - *Counting Instructions: Valuable Insights or More Noise?*
 - 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>

6





Quick Reminder: Convert your WLM Service Definition to HTML for Easy Reading

7

Enterprise Performance Strategies, Inc. ©

Format your WLM Service Definition in HTML



- Today, it is recommended to save the WLM service definition in XML format
 - The ISPF tables are automatically updated when they are touched by new APARs or z/OS releases. This then makes them ineligible to be updated if the APARs or z/OS releases are rolled back, or if an older release needs to edit or access.

- Select

- File
- -> Save as

Use Save as to save the currently displayed service definition in a PDS as ISPF tables or in a PS as XML

```
File Utilities Notes Options Help
-----
Functionality LEVEL025 Definition Menu WLM Appl LEVEL025
Command ==> _____

Definition data set . . . : none

Definition name . . . . . (Required)
Description . . . . . _____

Select one of the following options.
____ 1. Policies                               12. Tenant Resource Groups
      2. Workloads                             13. Tenant Report Classes
      3. Resource Groups
      4. Service Classes
      5. Classification Groups
      6. Classification Rules
      7. Report Classes
      8. Service Coefficients/Options
      9. Application Environments
     10. Scheduling Environments
     11. Guest Platform Management Provider
```

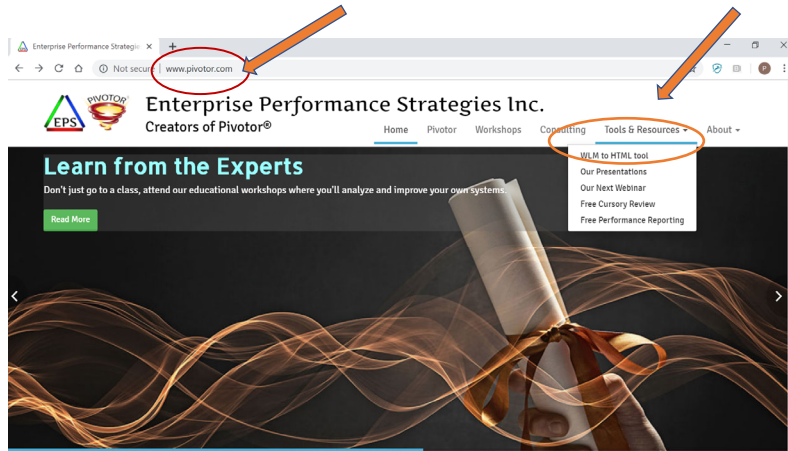
8



Convert it to HTML via our website!



- Once saved as XML, a tool is available to nicely format the XML file into a easy to read format to assist during your WLM analysis
- Visit www.pivotor.com or www.epstrategies.com
- Select 'Tools & Resources' option
- Select WLM TO HTML
- Provide your XML file and email address
 - HTML formatted WLM service definition emailed to you in seconds!



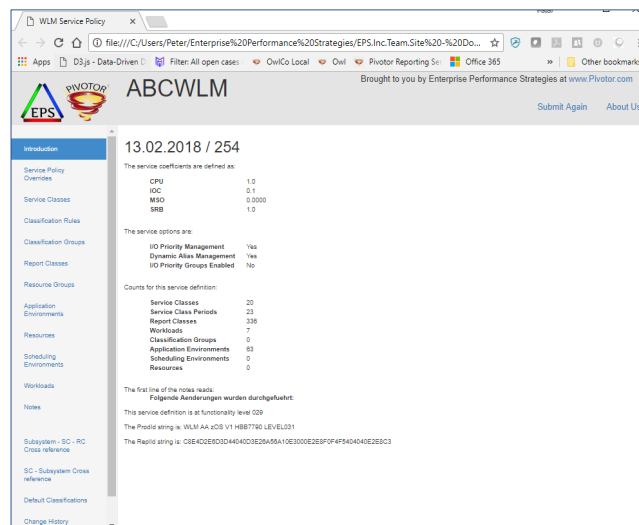
9

Enterprise Performance Strategies, Inc. ©

Example of HTML Formatted Service Definition



- Not only is the XML file nicely formatted
- But there is some analysis built into the file to help you with your service definition cleanup.



10





Refresh of Multiple Period Service Classes

11

Enterprise Performance Strategies, Inc. ©



Purpose of WLM Service Classes

- Why Service Classes?
 - To group work with similar performance characteristics for WLM management
 - ... but this also means that service classes are used to separate unlike work away from each other
- Examples of similar characteristics includes:
 - Work types
 - Resource requirements
 - Workload objectives
 - Business Requirements

STCHIGH Service Class Period 1 Goal = Velocity 50 Importance 1 RGRP = none	DB2PRD Service Class Period 1 Goal = Velocity 60 Importance 1 RGRP = none
SAPHIGH Service Class Period 1 Goal = Velocity 50 Importance 1 RGRP = none	SAPMED Service Class Period 1 Goal = Velocity 60 Importance 2 RGRP = none
CICSTORS Service Class Period 1 Goal = Velocity 60 Importance 1 RGRP = none	CICSAORS Service Class Period 1 Goal = Velocity 60 Importance 2 RGRP = none

12



But separation of work is not always possible



- As mentioned, work with similar performance characteristics is grouped into a Service class
- But at WLM classification, sometimes not enough is known about the work to separate the work into different service classes
 - Yet these different types of work should be managed separately
- Example: when the transaction starts it may not be known:
 - ... if the transaction will have a long or short response time
 - ... if transaction will be a large resource consumer or not
 - ... if the transaction will absorb service very quickly or if uses service more intermittently

PRODBAT Service Class

- Long or short?
- Big CPU consumer?
- CPU intensive or I/O intensive?

DDFPROD Service Class

- Long or short?
- Big CPU consumer?
- CPU intensive or I/O intensive?

TSOPROD Service Class

- Long or short?
- Big CPU consumer?
- CPU intensive or I/O intensive?

13

Enterprise Performance Strategies, Inc. ©

Reason for Multiple Period Service Classes



- Because some work may have variable resource requirements, service classes can be defined with multiple periods
 - Periods are a way of defining different goals for work depending on the amount of resources the work consumes
- Typically, periods are used to
 - Give shorter transactions more aggressive goals
 - Give longer transactions less aggressive goals
- Each period consists of
 - Goal and importance
 - Duration (except for last period)
- Durations for period 'aging'
 - The amount of resources, in service, that work consumes
 - As work consumes service and consumption exceeds duration, work is transitioned to the next period and managed to goal of next period
 - Way of aging transactions

PRODTSO Service Class

- Period 1 – 500 Service
Goal = RT 0.5 sec, 95%
Importance 2
RGRP =
- Period 2 – 1500 Service
Goal = RT 1.5 sec, 90%
Importance 3
RGRP =
- Period 3
Goal = RT 3.0 sec, 80%
Importance 4
RGRP =

14



Durations Allow for Transaction 'Aging'



- For multi-period Service Classes, work transitions to lower periods as it consumes **Service Units** to control distribution of resources
 - Assumption: If WLM goals and importance are setup correctly, then as a transaction uses more service it gains access to resources only after transactions in earlier periods

PRODTSO Service Class	
Period 1 – 500 Service	Goal = RT 0.5 sec, 95%
Importance 2	RGRP =
Period 2 – 1500 Service	Goal = RT 1.5 sec, 90%
Importance 3	RGRP =
Period 3	Goal = RT 3.0 sec, 80%
Importance 4	RGRP =



15

Enterprise Performance Strategies, Inc. ©

What is measured for a duration?



- As transactions are processed, they consume system resources
 - The measure of resource consumption is service units
 - CPU and SRB service units (i.e. processor)
 - I/O service units
 - MSO service units (i.e. storage)
- Traditionally, durations are in terms of 'service' and not 'service units'
 - Service is service units **weighted** by service definition coefficients (**SDCs**)
 - When duration is set for a period, the service consumed determines period switch
 - Why weight? Historical...

$$\text{Service} = \left(\begin{array}{l} (\text{CPU SDC} * \text{CPU Service Units}) \\ + (\text{SRB SDC} * \text{SRB Service Units}) \\ + (\text{IOC SDC} * \text{IOC Service Units}) \\ + (\text{MSO SDC} * \text{MSO Service Units}) \end{array} \right)$$

16



Reminder: Modernize your SDCs for after z/OS 2.4



- As a reminder...
 - MSO = 0 has been recommended for many years
 - IOC = 0 is not recommended since going away in z/OS after 2.4 (so get ready now)
 - CPU and SRB will be 1 (so processor service and service units are the same)
- The only impact should be to durations of your multiple period service classes

```

Coefficients/Options  Notes  Options  Help
-----
Service Coefficient/Service Definition Options
Command ==>

Enter or change the Service Coefficients:
CPU . . . . . 1.0 (0.0-99.9)
IOC . . . . . 0.0 (0.0-99.9)
MSO . . . . . 0.0 (0.0000-99.9999)
SRB . . . . . 1.0 (0.0-99.9)

Enter or change the service definition options:
I/O priority management . . . . . NO (Yes or No)
Enable I/O priority groups . . . . . NO (Yes or No)
Dynamic alias management . . . . . NO (Yes or No)
    
```

17

Enterprise Performance Strategies, Inc. ©

Example of a Multiple Period Service Class



- The following is an example of a three-period WLM service class for TSO
 - Note that the last period has no duration
 - Sometimes more than 3 periods, extremely rarely would you need more than 4

```

Service-Class Xref Notes Options Help
-----
Modify a Service Class Row 1 to 4 of 4
Command ==>

Service Class Name . . . . . : TSO
Description . . . . . : All TSO Activity
Workload Name . . . . . : TSO (name or ?)
Base Resource Group . . . . . : (name or ?)
Cpu Critical . . . . . : NO (YES or NO)
I/O Priority Group . . . . . : NORMAL (NORMAL or HIGH)
Honor Priority . . . . . : NO (DEFAULT or NO)

Specify BASE GOAL information. Action Codes: I=Insert new period,
E=Edit period, D=Delete period.

---Period--- -----Goal-----
Action # Duration Imp. Description
---
1 2000 2 90% complete within 00:00:00.250
2 3000 3 90% complete within 00:00:00.500
3 80% complete within 00:00:01.000
***** Bottom of data *****
    
```

18





When does it make sense to use multiple period service classes to improve performance?

In our *WLM Performance and Re-evaluation of Goals* workshop we discuss much more about how to set a duration and how to tune a duration.

Here I will just be discussing when to use to improve performance.

Instructor: Peter Enrico

Enterprise Performance Strategies, Inc. ©

Latent Demand - 19

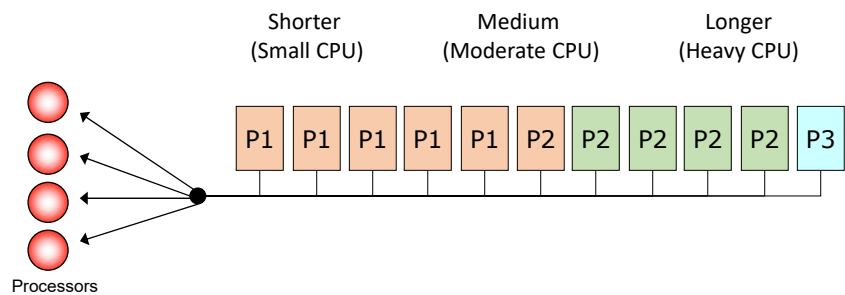
19

Enterprise Performance Strategies, Inc. ©



To simply distribute different work

- As mentioned, sometimes at classification not enough is known to allow the work to be assigned the correct goals and importance level
 - Multiple period service classes helps address this
- Distribute CPU dispatching priority of transactions based on CPU demands
 - Lighter transactions up front
 - Heavy CPU transactions at end



20



To simply distribute different work



- As mentioned, sometimes at classification not enough is known to allow the work to be assigned the correct goals and importance level
 - Transactions using more resources are 'separated' for WLM management based on service consumed
- In this example, say the LPAR has an SU/sec constant of 50,000 SU/sec
- I decided the duration
 - 3,000 represents 0.06 seconds of CPU
 - 2,000 represents 0.04 seconds of CPU
 - So, anything using more than 0.08 seconds of CPU will run in period 3

```

Service-Class Xref Notes Options Help
-----
Modify a Service Class Row 1 to 4 of 4
Command ==>

Service Class Name . . . . . : TSO
Description . . . . . : All TSO Activity
Workload Name . . . . . : TSO (name or ?)
Base Resource Group . . . . . : (name or ?)
Cpu Critical . . . . . : NO (YES or NO)
I/O Priority Group . . . . . : NORMAL (NORMAL or HIGH)
Honor Priority . . . . . : NO (DEFAULT or NO)

Specify BASE GOAL information. Action Codes: I=Insert new period,
E=Edit period, D=Delete period.

---Period--- ---Goal-----
Action # Duration Imp. Description
--- 1 3000 2 90% complete within 00:00:00.200
--- 2 2000 3 90% complete within 00:00:00.400
--- 3 4 4 80% complete within 00:00:01.000
***** Bottom of data *****
    
```

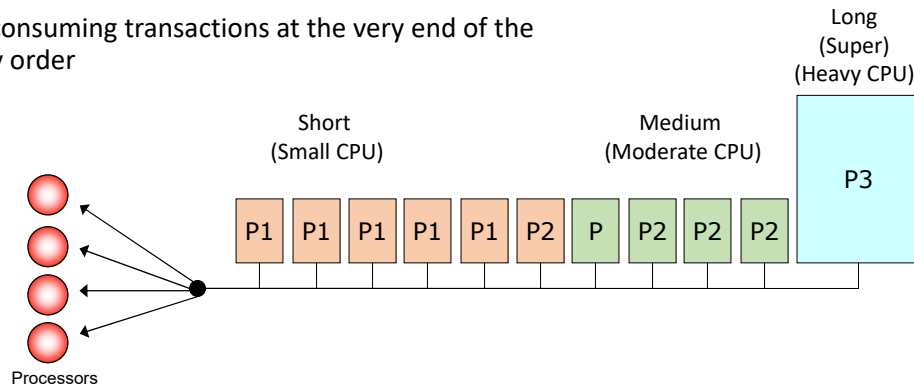
21

Enterprise Performance Strategies, Inc. ©

Improve performance with penalty periods



- A penalty period is the last period of a WLM service class
 - The durations of the service class are setup such that the only work that transitions to the last period is one that has used an unacceptable amount of CPU service
- Put the large CPU consuming transactions at the very end of the dispatching priority order



22



Improve performance with penalty periods



- A penalty period is the last period of a WLM service class
 - The durations of the service class are setup such that the only work that transitions to the last period is one that has used an unacceptable amount of CPU service

- In this example, say the LPAR has an SU/sec constant of 50,000 SU/sec

- Then duration **10,000,000** represents 200.04 seconds of CPU

- I made the decision that any DEVBATCH job using more than 200 seconds needs to go to discretionary

```

Service-Class Xref Notes Options Help
-----
Command ==> _____ Row 1 to 4 of 4

Service Class Name . . . . . : DEVBATCH
Description . . . . . : Development Batch
Workload Name . . . . . : BATCH (name or ?)
Base Resource Group . . . . . : _____ (name or ?)
Cpu Critical . . . . . : NO (YES or NO)
I/O Priority Group . . . . . : NORMAL (NORMAL or HIGH)
Honor Priority . . . . . : NO (DEFAULT or NO)

Specify BASE GOAL information. Action Codes: I=Insert new period,
E=Edit period, D=Delete period.

---Period--- -----Goal-----
Action # Duration Imp. Description
---
1 2000 3 Execution velocity of 50
2 10000000 4 Execution velocity of 60
3 3 Discretionary
***** Bottom of data *****
    
```

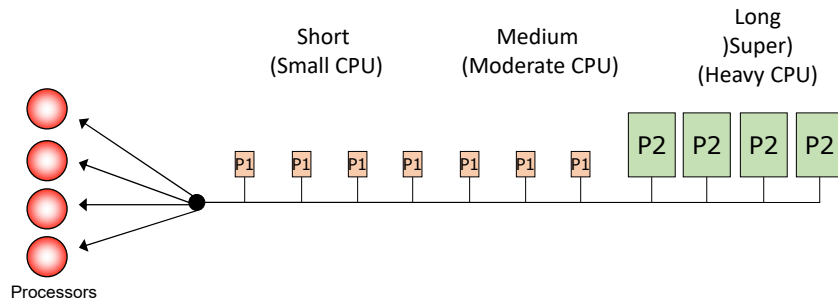
23

Enterprise Performance Strategies, Inc. ©

To get the short transaction in and out ASAP



- Sometimes a workload can improve when we let all transactions run a high CPU dispatching priority when they start
 - But anything that is not super quick goes to a lower period
- So lets get the quick / short small CPU transactions in and out right away



24



To get the short transaction in and out ASAP



- Sometimes a workload can improve when we let all transactions run a high CPU dispatching priority when they start
 - But anything that is not super quick goes to a lower period

- In this example, say the LPAR has an SU/sec constant of 50,000 SU/sec
- Then duration **250** would represent 0.005 seconds of CPU
- I made the decision that DDF trans run at a higher imp until they use up 5 ms of CPU.

```

Service-Class Xref Notes Options Help
-----
Command ==>          Modify a Service Class          Row 1 to 4 of 4

Service Class Name . . . . . : DDFPROD
Description . . . . . : Production DDF
Workload Name . . . . . : DDFDB2 (name or ?)
Base Resource Group . . . . . : (name or ?)
Cpu Critical . . . . . : NO (YES or NO)
I/O Priority Group . . . . . : NORMAL (NORMAL or HIGH)
Honor Priority . . . . . : NO (DEFAULT or NO)

Specify BASE GOAL information. Action Codes: I=Insert new period,
E=Edit period, D=Delete period.

---Period--- -----Goal-----
Action # Duration Imp. Description
-----
1 250 1 95% complete within 00:00:00.015
2 3 3 90% complete within 00:00:00.150
***** Bottom of data *****
    
```


25

Enterprise Performance Strategies, Inc. ©

To assist in latent demand management



- Sometimes in heavy CPU constrained environments when there is just too much work running in certain workload periods, it is best to force the distribution of work
- For example, batch workloads with many jobs, too many initiators relative to the number of processors
- CPU demands result in certain HiperDispatch pooling combinations that result in the thrashing of the processor caches
- Sometimes just better to let the smaller CPU consumers in to get them out of the way
- Let the longer transactions sink in CPU DP order



26



To assist in latent demand management



- Sometimes in heavy CPU constrained environments when there is just too much work running in certain workload periods, it is best to force the distribution of work

- In this example, say the LPAR has an SU/sec constant of 50,000 SU/sec

- I decided the duration
 - 50,000 represents 1 seconds of CPU
 - 200,000 represents 4 seconds of CPU
 - So, anything using more than 5 seconds of CPU will run in period 3

```

Service-Class Xref Notes Options Help
-----
Modify a Service Class Row 1 to 4 of 4
Command ==>

Service Class Name . . . . . : PRDBATCH
Description . . . . . : Production Nighttime Batch
Workload Name . . . . . : BATCH (name or ?)
Base Resource Group . . . . . : (name or ?)
Cpu Critical . . . . . : NO (YES or NO)
I/O Priority Group . . . . . : NORMAL (NORMAL or HIGH)
Honor Priority . . . . . : NO (DEFAULT or NO)

Specify BASE GOAL information. Action Codes: I=Insert new period,
E=Edit period, D=Delete period.

-----Period-----Goal-----
Action # Duration Imp. Description
-----
1 50000 3 Execution velocity of 50
2 200000 4 Execution velocity of 60
3 5 Execution velocity of 60
***** Bottom of data *****
    
```

27

Enterprise Performance Strategies, Inc. ©


Summary






- Because not enough is known about work at classification
 - Let's distribute the work
- As a penalty period for run-away work
 - Let's penalize possible loopers, or just transactions whose CPU consumption is way out of the norm
- To get super short and super small CPU consuming transactions on and off the CPU as quickly as possible
 - Get them out of the way!
- Prioritize work during periods of heavy latent demand
 - Can enable a more effective usage of the processor caches


28




YouTube Mainframe Performance Channel






Mainframe Performance
202 subscribers

[CUSTOMIZE CHANNEL](#)
[MANAGE VIDEOS](#)


HOME
VIDEOS
PLAYLISTS
CHANNELS
DISCUSSION
ABOUT




The 3 Responsibilities of all Computer Performance Analysts
89 views · 3 months ago

This video details the three primary responsibilities of any computer performance analyst. If you someone thinking of getting into the field of computer performance analysis, if you are an IT manager that wants to understand what your staff is doing or many you are already a computer performance analyst and just want to hear my organization of thought on the responsibilities, then this video is for you. Also note... This was recorded at the Fall 2020 Mainframe Summit. It is available on the channel. Read more


Uploads ▶ PLAY ALL




45 views · 2 weeks ago




52 views · 1 month ago




102 views · 1 month ago



100 views · 1 month ago



72 views · 2 months ago





32 views · 2 months ago

29

Enterprise Performance Strategies, Inc. ©

Like what you see?

- Free z/OS Performance Educational webinars!
 - The titles for our Fall 2020 webinars are as follows:
 - ✓ *WLM Response Time Goals Primer*
 - ✓ *Setting Response Time Goals for Modern Applications*
 - ✓ *Evaluating Latent Demand in the Mainframe Environment*
 - ✓ *Adjusting WLM Settings for Latent Demand*
 - ✓ *Improving Performance with Multiple Period Service Classes*
 - *Preparing for Any z/OS Performance Analysis*
 - *Evaluating Coupling Facility Lock Structures*
 - *Exploring Coupling Facility Exploitation by VSAM RLS*
 - *Data in Memory (DIM) Primer*
 - *Counting Instructions: Valuable Insights or More Noise?*
 - 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>

30

15