



Batch Initiators WLM or JES ? Does it matter?

Nancy DiFilippo

Enterprise Performance Strategies, Inc.

Nancy.DiFilippo@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!)



WLM initiators for batch jobs were introduced in the late 1990's but there are still many installations that make use of JES initiators instead. During this webinar, ***Nancy DiFilippo*** will discuss the pros and cons of both JES and WLM initiator batch management to help you understand your choices.

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>



z/OS Performance workshops available



During these workshops you will be analyzing your own data!

- WLM Performance and Re-evaluating Goals
 - February 19-23, 2024
- Parallel Sysplex and z/OS Performance Tuning
 - August 20-21, 2024
- Essential z/OS Performance Tuning
 - September 16-20, 2024
- Also... please make sure you are signed up for our free monthly z/OS educational webinars! (email contact@epstrategies.com)

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.

Like what you see?



- Free z/OS Performance Educational webinars!
 - The titles for our Fall 2023-2024 webinars are as follows:
 - ✓ *LPAR Configurations to Avoid*
 - ✓ *How Different are High, Medium, and Low Pool Processors?*
 - ✓ *CPU Critical: A Modern Revisit of a Classic WLM Option*
 - ✓ *Mainframe Efficiency at High Utilizations (Bob Rogers)*
 - ✓ *I/O, I/O It's Home to Memory We (Should) Go*
 - ✓ *30th Anniversary of WLM : A Retrospective and Lessons Learned*
 - ✓ *Mainframe Efficiency at High Utilizations (presented by Bob Rogers)*
 - ✓ *Understanding and Measuring Warning Track on z/OS*
 - ✓ *30th Anniversary of Parallel Sysplex - A Retrospective and Lessons Learned*
 - ✓ *Introducing Pivotor Outlier Detection and Analysis*
 - *Batch Initiators – WLM Managed or JES Managed?*
 - *Analyzing 'Per CPU' Utilizations*
 - *AI on Z: Exploring new SMF Measurements -TBD*
 - 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>

About Me



- Early 1980's Shared Services Environment
 - Heavily modified MVS/JES2
 - Retrofitting mods to JES2 updates and creating new functionality
- Mid 1980's through 2022 - Commercial product ThruPut Manager
 - Automated Batch Management
 - Part of initial development team – focus JES2
 - Many hats – development / support / training
- Spent many years trolling around JES2 internals

Agenda



- JES initiators
 - Problems
 - Focus on JES2 initiators
 - Majority of z/OS shops
 - JES3 has the concept of dynamic initiators
- WLM initiators
 - Problems they were designed to address
 - New issues introduced
- Futures
 - What lays ahead
- Recommendations
- Questions

JES2 Initiators



Setup is manual

- Define number of initiators and where (what system) they are to be active
- Initiators are defined with job classes they service

```
INIT(1) CLASS=A
```

```
INIT(2) CLASS=B
```

```
INIT(3) CLASS=C
```

- Can be multiple classes per initiator

```
INIT(1) CLASS=AFJKE      /*INITIATOR 1*/
```

```
INIT(2) CLASS=BCDEF     /*INITIATOR 2*/
```

```
INIT(3) CLASS=DEFGH    /*INITIATOR 3*/
```

```
INIT(4) CLASS=XKH      /*INITIATOR 4*/
```

```
INIT(5) CLASS=JKEBF    /*INITIATOR 5*/
```

```
INIT(6) DRAIN          /*SPARE INITIATOR*/
```

```
INIT(7) DRAIN          /*SPARE INITIATOR*/
```

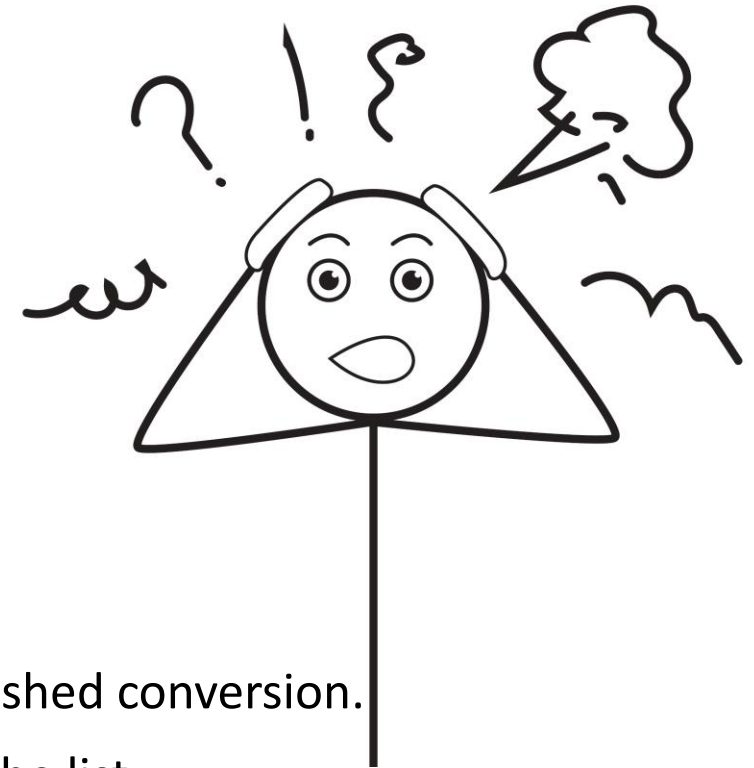
```
INIT(8) DRAIN          /*SPARE INITIATOR*/
```

- **Definitions are system specific, and CAN be, and likely are, different on a given system in a JESplex**

JES2 Alphabet Soup



```
INIT(1) CLASS=AFJKE /*INITIATOR 1*/
INIT(2) CLASS=BCDEF /*INITIATOR 2*/
INIT(3) CLASS=DEFGH /*INITIATOR 3*/
INIT(4) CLASS=XKH /*INITIATOR 4*/
INIT(5) CLASS=JKEBF /*INITIATOR 5*/
INIT(6) DRAIN /*SPARE INITIATOR*/
INIT(7) DRAIN /*SPARE INITIATOR*/
INIT(8) DRAIN /*SPARE INITIATOR*/
```



- Jobs are queued by job class, priority, and the order in which they finished conversion.
- Selection is by **priority** in a given class in the order the classes are in the list
- Unless using Class Groups – round robin !

WHAT JOB will be selected next and where ??

Operator Controls



Operations is manual

- Responsible for “running the system”
 - Start inits
 - Drain inits
 - Load Balancing
- Auto-operators
 - Time specific changes to initiators
 - Often Shift Based – daytime/overnight
 - Month End etc
- Remember : Start of JES initiators is Manual – No performance considerations !
 - Won't prevent starting of more than one high CPU consumption job per system in the MAS.
 - Won't prevent start of any high CPU consuming job classes inits on a system that processes mainly online

Operators are asked to “get a job going” so they do

Over/Under Initiation

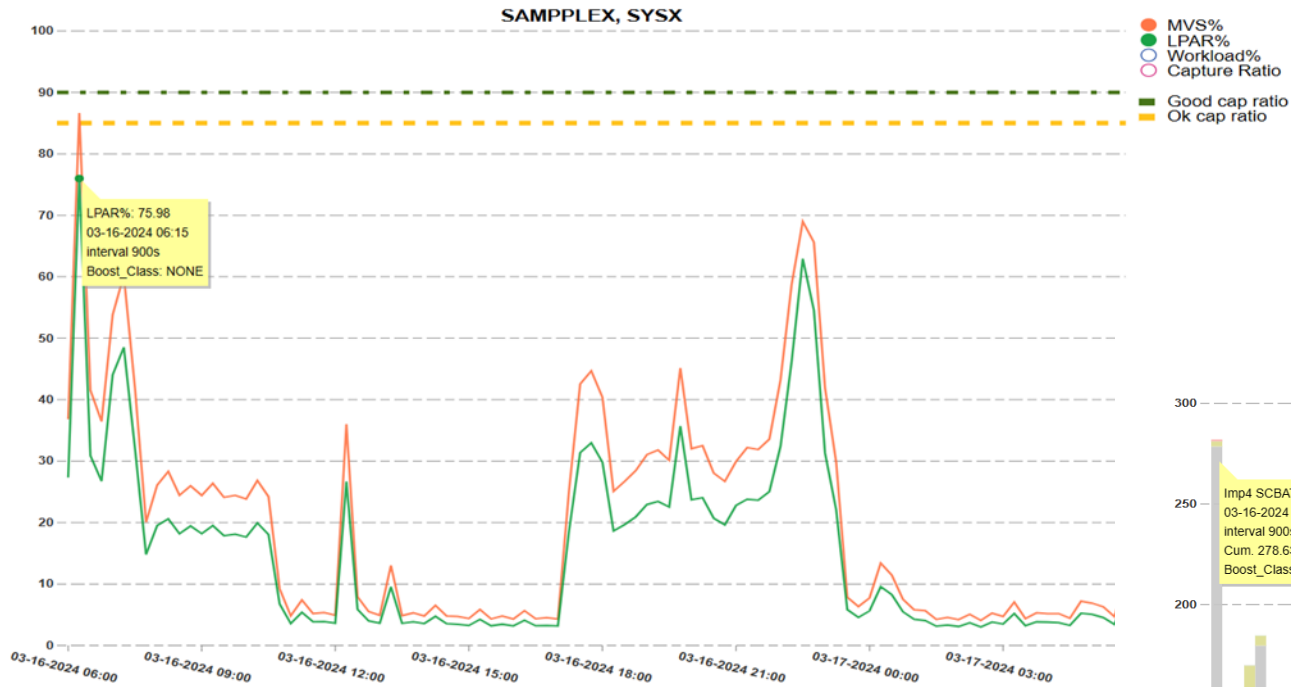


- Finding the right balance is difficult
- Impacts of over Initiation – Real \$\$
 - For every 10% increase in CPC usage there is a 3-5% increase in CPU usage for the same workload
 - (IBM)
 - Contention for system resources increases (memory/processors i/o)
 - Large Systems Performance Reference (LSPR)
 - “Tuning to reduce the number of simultaneously active address spaces to the proper number needed to support a workload can reduce RNI and improve performance”
- Impacts of Under Initiation
 - Batch Window may be longer – how does that impact your business?
 - System is not used efficiently – want to make the most of your \$\$’s

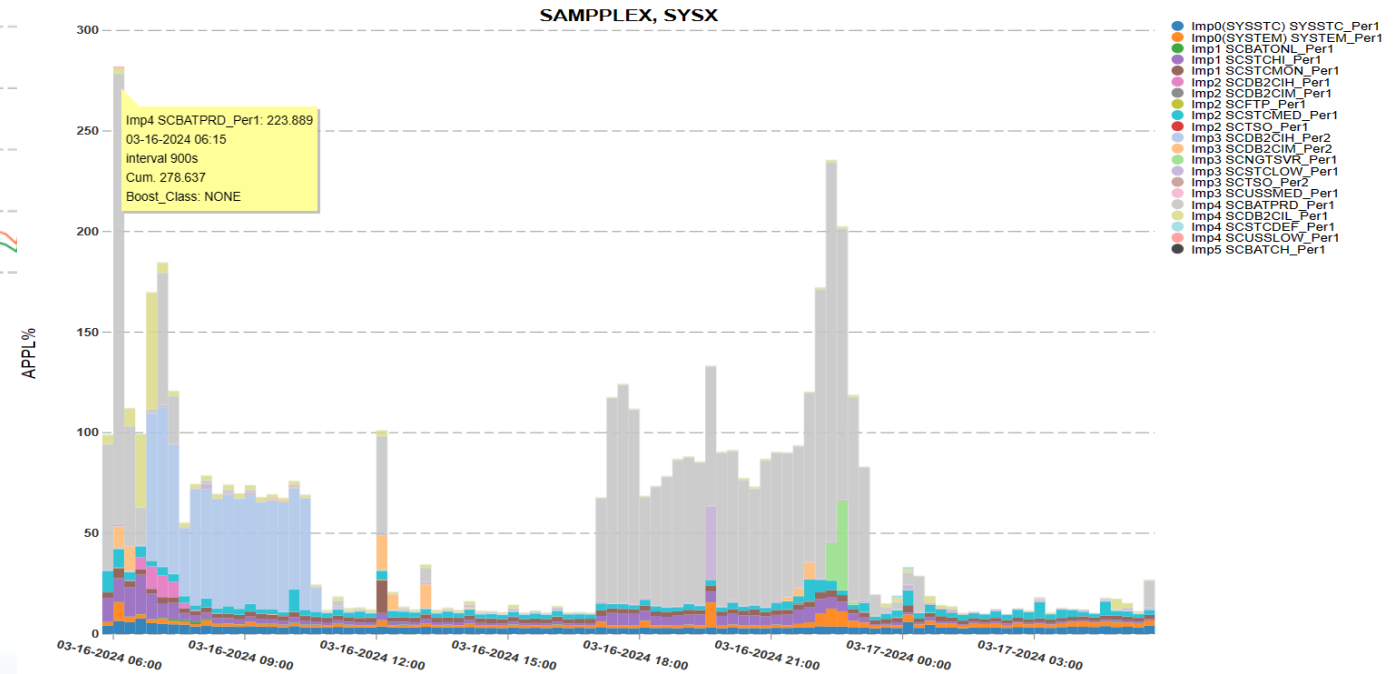
Over initiation/Under initiation



LPAR, MVS, and Workload CP Busy% with Capture Ratio




WLM CPU - CP APPL% by Service Class
(CP + zIIP on CP)




Over initiation/Under initiation



Batch Jobs - Min, Avg, Max



WLM CPU - CP CPU Delay Samples By Period



It's Complicated




- Knowing when and where to start
- Understanding what jobs need help
- Understanding performance of the system

It's Complicated



- Over years expertise developed to fine tune these specifications
 - What if that expertise leaves?
 - What if environment changes ?
 - add more engines ?
 - Jobs become more CPU-intensive due to improvements in I/O (large memory exploitation/reduced I/O time)
 - What if new loads added ?
 - What if systems are merged ?
 - Who has the knowledge to perform the analysis and make the adjustments ?




WLM Inits to the rescue !



WLM initiators are intended to reduce these complications

- Simpler
 - Fewer and simpler externals to control WLM- managed initiators and to perform workload balancing.
 - Automatically change its selection criteria – different service class
- Dynamic: allows the system to adapt to changing environmental conditions
- WLM adjusts the number of initiators per system based on:
 - The **queue** of jobs awaiting execution in WLM managed classes.
 - The performance goals and relative importance of this work.
 - The success of meeting these goals.
 - The capacity of each system to do more work



WLM Inits - what happened?



Sounds pretty good – Why isn't everyone using WLM inits alone?

- WLM inits were introduced in the late 1990's
 - Many shops still using ONLY JES inits or a mixture of WLM and JES
 - EPS data shows almost half customer environments don't use WLM inits
- Two main complaints
 - Tends to over-initiate
 - Analysis done every 10 seconds and only one change per cycle – takes a while to adjust
 - System performance degrades - doesn't halt inits quickly enough
 - Slow to start initiators
 - Large arrival of jobs – not enough initiators – takes time to add them
 - Advice from IBM is if jobs can't tolerate queue delay, then use JES inits
 - Queue Time is important !
 - But what about schedulers that are being used to effectively eliminate queue time ?
- Remember: WLM initiators are best suited for steady stream of work

WLM Inits –what happened?




- Still need some controls
 - Jobs queued FIFO by the assigned WLM service class .
 - No ability to assign priority within the service class.
 - Remember: goals of the SERVICE CLASS are assessed – individual jobs within that service class may suffer.
 - **Operations may still need to help a particular job.**
 - May need to set limits for particular systems (JES2 XEQCOUNT and XEQMEMBER parameters)
- WLM initiators resolve much of the reliance on expert operations but there still are issues

What's in the Future ?



AI INFUSED WLM INITIATORS



What's In the Future ? AI WLM Inits



Introduced in z/OS 3.1 with the z16

Two New Features

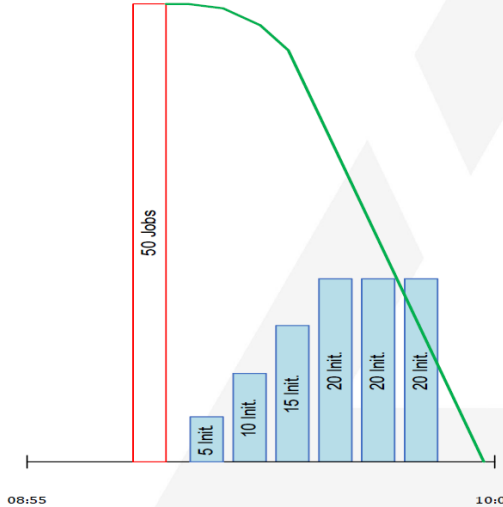
Use of Modelling/Training

- Process period of SMF (99.2) to predict how many jobs to expect
- What time period is appropriate for successful model
 - Current recommendation – minimum 30 days
 - Peaks ?
- WLM will start inits where and when expected.
- If load changes must rerun model.
 - Unexpected changes will impact the model
- Setup is complicated 😊 See <https://www.ibm.com/downloads/cas/A6LG1M38>

JES2 Futures




WLM managed Initiators today



- Reactive mode
- Initiators are started after the workload arrives
- Workload can't be processed right away

AI powered WLM managed Initiators tomorrow



- Workload forecast from AI allows WLM to start/stop initiators proactively
- Workload can start right away
 - Resource assignment of the correct priority is done before the work arrives
- Possibility to switch between non-AI and AI mode

Charts courtesy of Share Session 31286 Managing WLM Batch Initiators with AI
 Steve Partlow / Dieter Wellerdiek Share New Orleans Aug 2023

What's In the Future ? AI WLM Inits



Use of JES2 HOLD UNTIL

- HOLDUNTIL Keyword of SCHEDULE JCL statement introduced in 2015
- Allows specification of a date and time for which the job will be held while it's on the queue
- New Functionality: WLM will not halt initiators where there are jobs on the queue scheduled to start in next 2 min
- Designed to improve WLM initiator availability

How much production batch make use of HOLDUNTIL?

- Schedulers will often avoid submitting jobs to the queue until immediately ahead of their start time
- How successful – Yet to be determined

Recommendations



- Ideally - Today
 - Mix Mode is best
 - Minimize operator intervention
 - Simplify batch initiator structure
 - Easier adaption to change
 - Better overall load balance
 - Better performance for important workloads

JES INITS → WLM INITS ?



Have a batch problem?

NO

Wait / Watch AI Space

Analysis

CPU Delays ?

Large Queues?

Unbalanced Batch Load ?

YES

YES

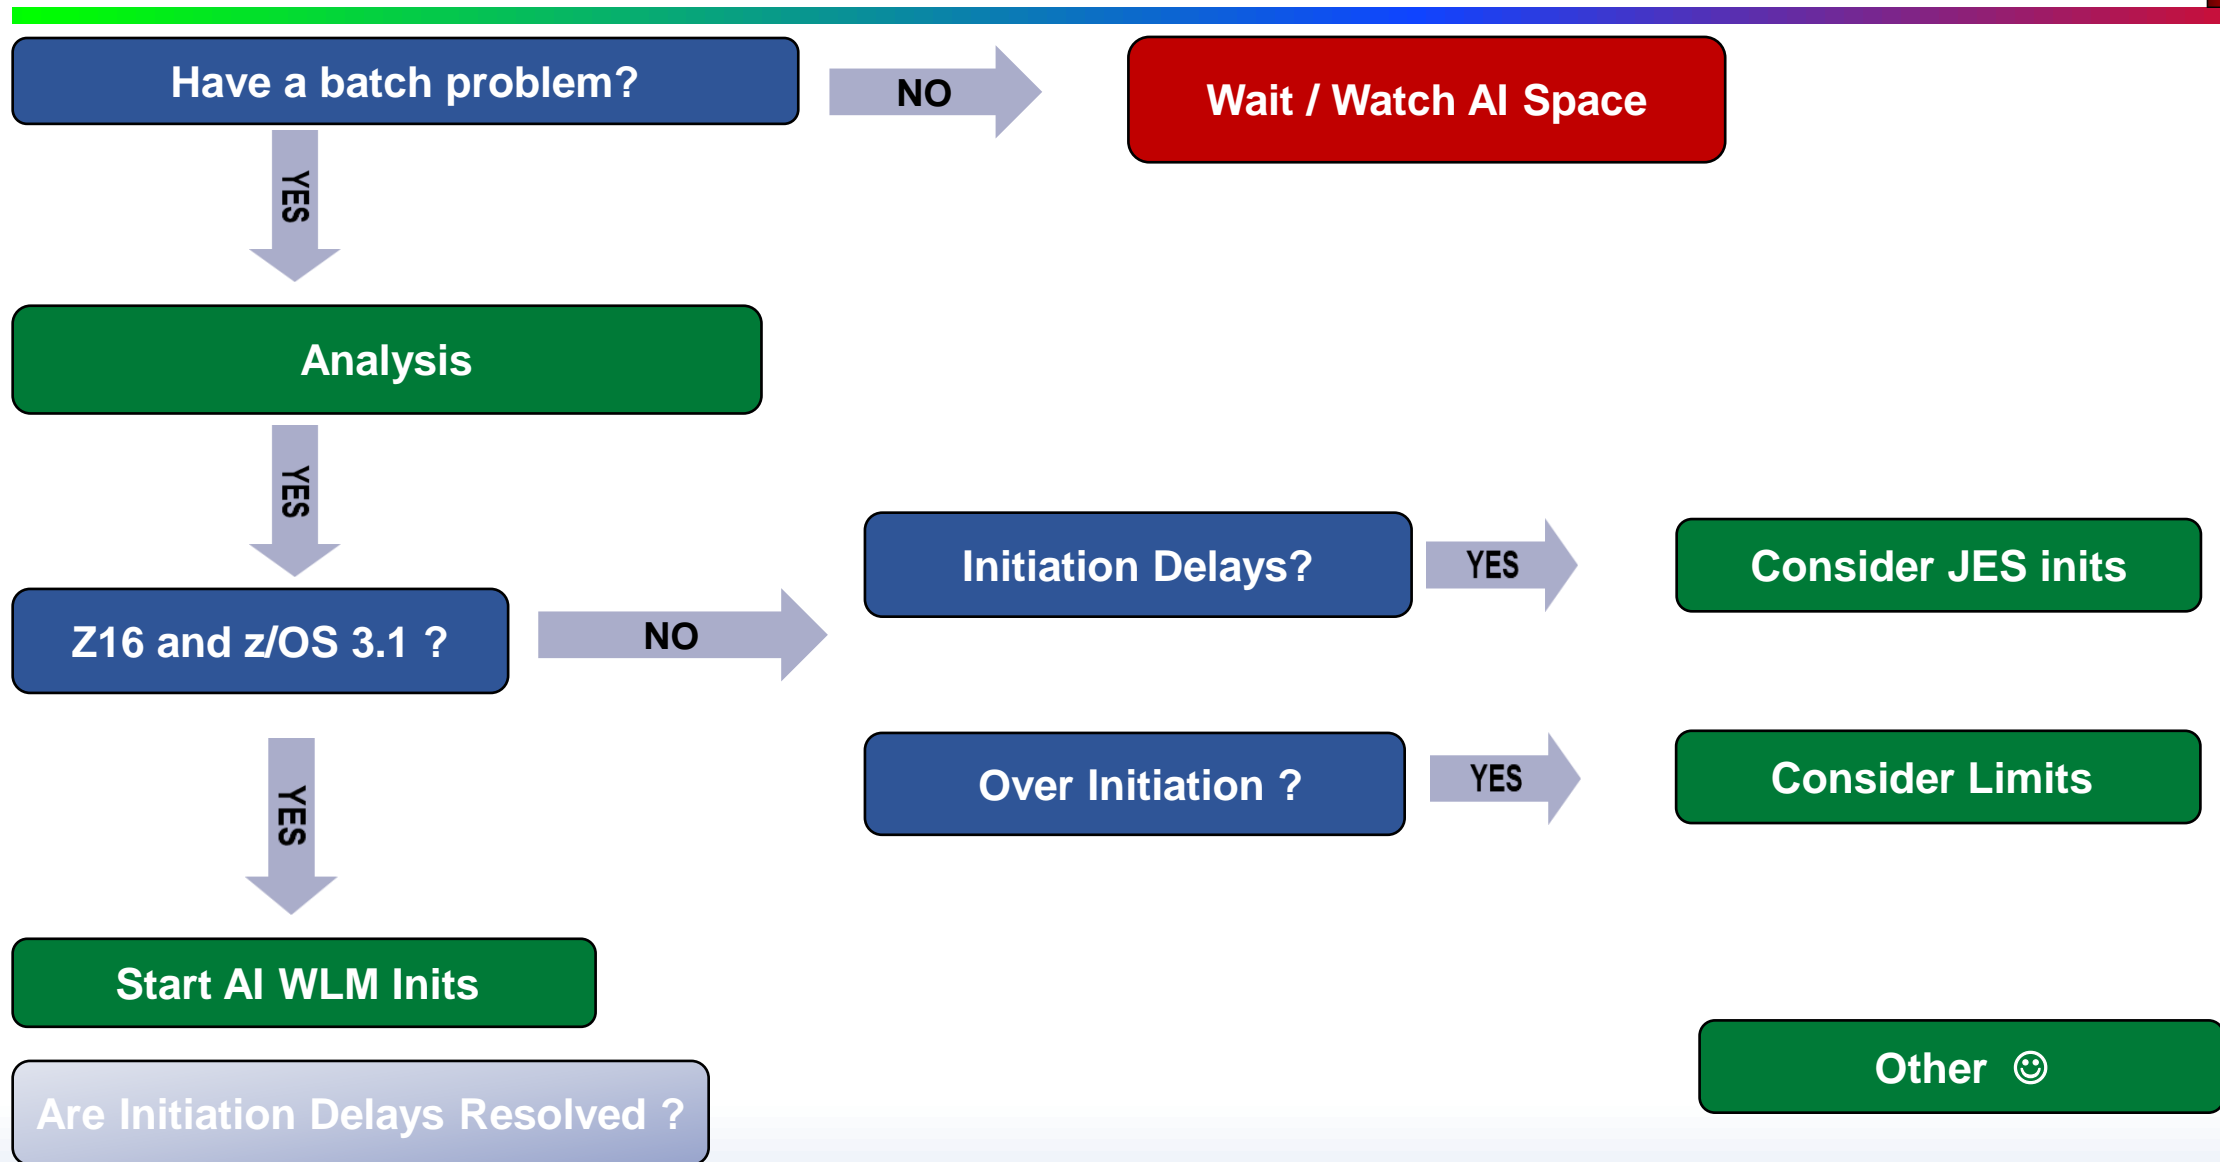
YES

Consider reducing # inits

Consider increasing # inits

Consider WLM Inits

WLM Inits → AI infused Inits ?



Recommendations



- Reality
 - Wait - is better initiator management enough ?
 - Identify those “can’t wait” jobs and run in JES inits
 - Watch for more AI enhancements – More coming



Questions?