Wait…
Do we need to re-evaluate our WLM Goals?

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:

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

Re-evaluating your WLM service definition is generally said to be something that should be done periodically and after changes to the environment. But really, why? We regularly see service definitions that haven't changed in years. Presumably, their systems haven't devolved into a chaotic mess that has put the company out of business, so are we sure we care? Actually, there are some good reasons why you need to periodically re-evaluate your goals. In this webinar, **Scott Chapman** will explain his reasoning and high-level process for re-evaluating a WLM service definition and why yes, many systems could stand to benefit from a refreshed and optimized WLM policy.

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

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)

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!

Agenda

- Do people ignore their WLM policies?
- Why do we have to have WLM?
- What changes matter?
- What if happens goals aren't well-tuned?
- Wrap-up

Do People Ignore Their WLM Policy?

When was your WLM policy* last updated?

- We get dozens of WLM policies submitted every month through our WLM to HTML tool https://www.epstrategies.com/wlm2html.html
  - Presumably this is because people are making changes
  - Or maybe they're just trying to understand an old policy
- But… I've seen policies that haven't been updated in years
  - There is in fact a change log embedded in the service definition!
  - I added a summary on the tool output for how long it's been since it was updated

    This service definition is at functionality level 032

    The ProdId string is: WLM AA zOS V2R4 HBB77C0 LEVEL038

    It has been approximately 4 months since the last recorded change.
  - I've seen the "months" large enough that we should probably automatically switch to expressing it in years at some point.
  - If you're curious about the levels: https://www.ibm.com/docs/en/zos/3.2.0?topic=details-mapping-functionality-levels-zos-releases

* - To be technically accurate: your WLM Service Definition

Years between policy updates?

- Common advice is to re-evaluate and adjust your policy for lots of reasons
- How would a site go years between policy updates??
  - Maybe they were getting off the mainframe?
  - Maybe the WLM person retired?
  - Maybe really nothing changed?
- Is this possibly ok?
- Are we just making excess work for ourselves?
- Linux and Windows don't have a sophisticated dispatching priority management system, so why do we need to maintain one?

Why do we have WLM?

z/OS ≠ Linux / Windows

- Linux and Windows are fundamentally managed differently
- Linux/Windows generally have less competition for resources
  - Generally run at lower utilization levels, in part because CPs are “cheap”
  - Less competition = dispatching priorities matter less
- Linux/Windows servers often run “single” workload
  - Database server vs. web server vs. application server vs. whatever
  - If only one workload, no real need to distinguish between workloads

- z/OS almost always have many workloads per system
- z/OS often sees significant CPU contention
  - Often have few CPs because CPs are “expensive”

z/OS needs WLM

WLM allows z/OS to run diverse workloads on one system, while balancing the performance of those workloads in the face of significant resource contentions

- Today, the “resources” of primary concern is CPU
  - Either GPs or zIIPs
- Primary control for access to CPU is dispatching priority
  - Higher dispatching priority = closer to front of queue to get access to CPU

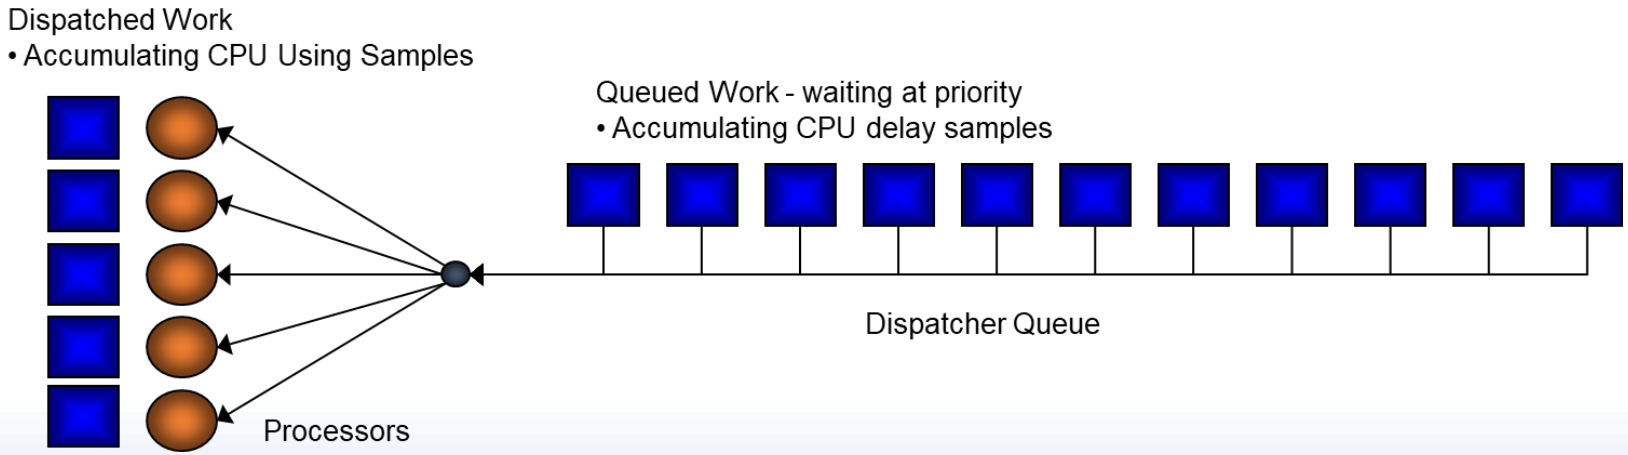

Do I care about dispatching priority?

- If you never have a queue for a CP, dispatching priority doesn't much matter
  - No line = no waiting
  - But there are always times when there are lines!
- "WLM doesn't really matter because our machine is only 70% busy"
  - Is that a 15 minute average? What about at 9:03 – 9:07?
- "We have lots of CPs we shouldn't have a problem"
  - But how many LPARs do you have?
  - How many logical CPs does each LPAR have?

CEC Physical Machine CP Busy% by CEC Serial Number

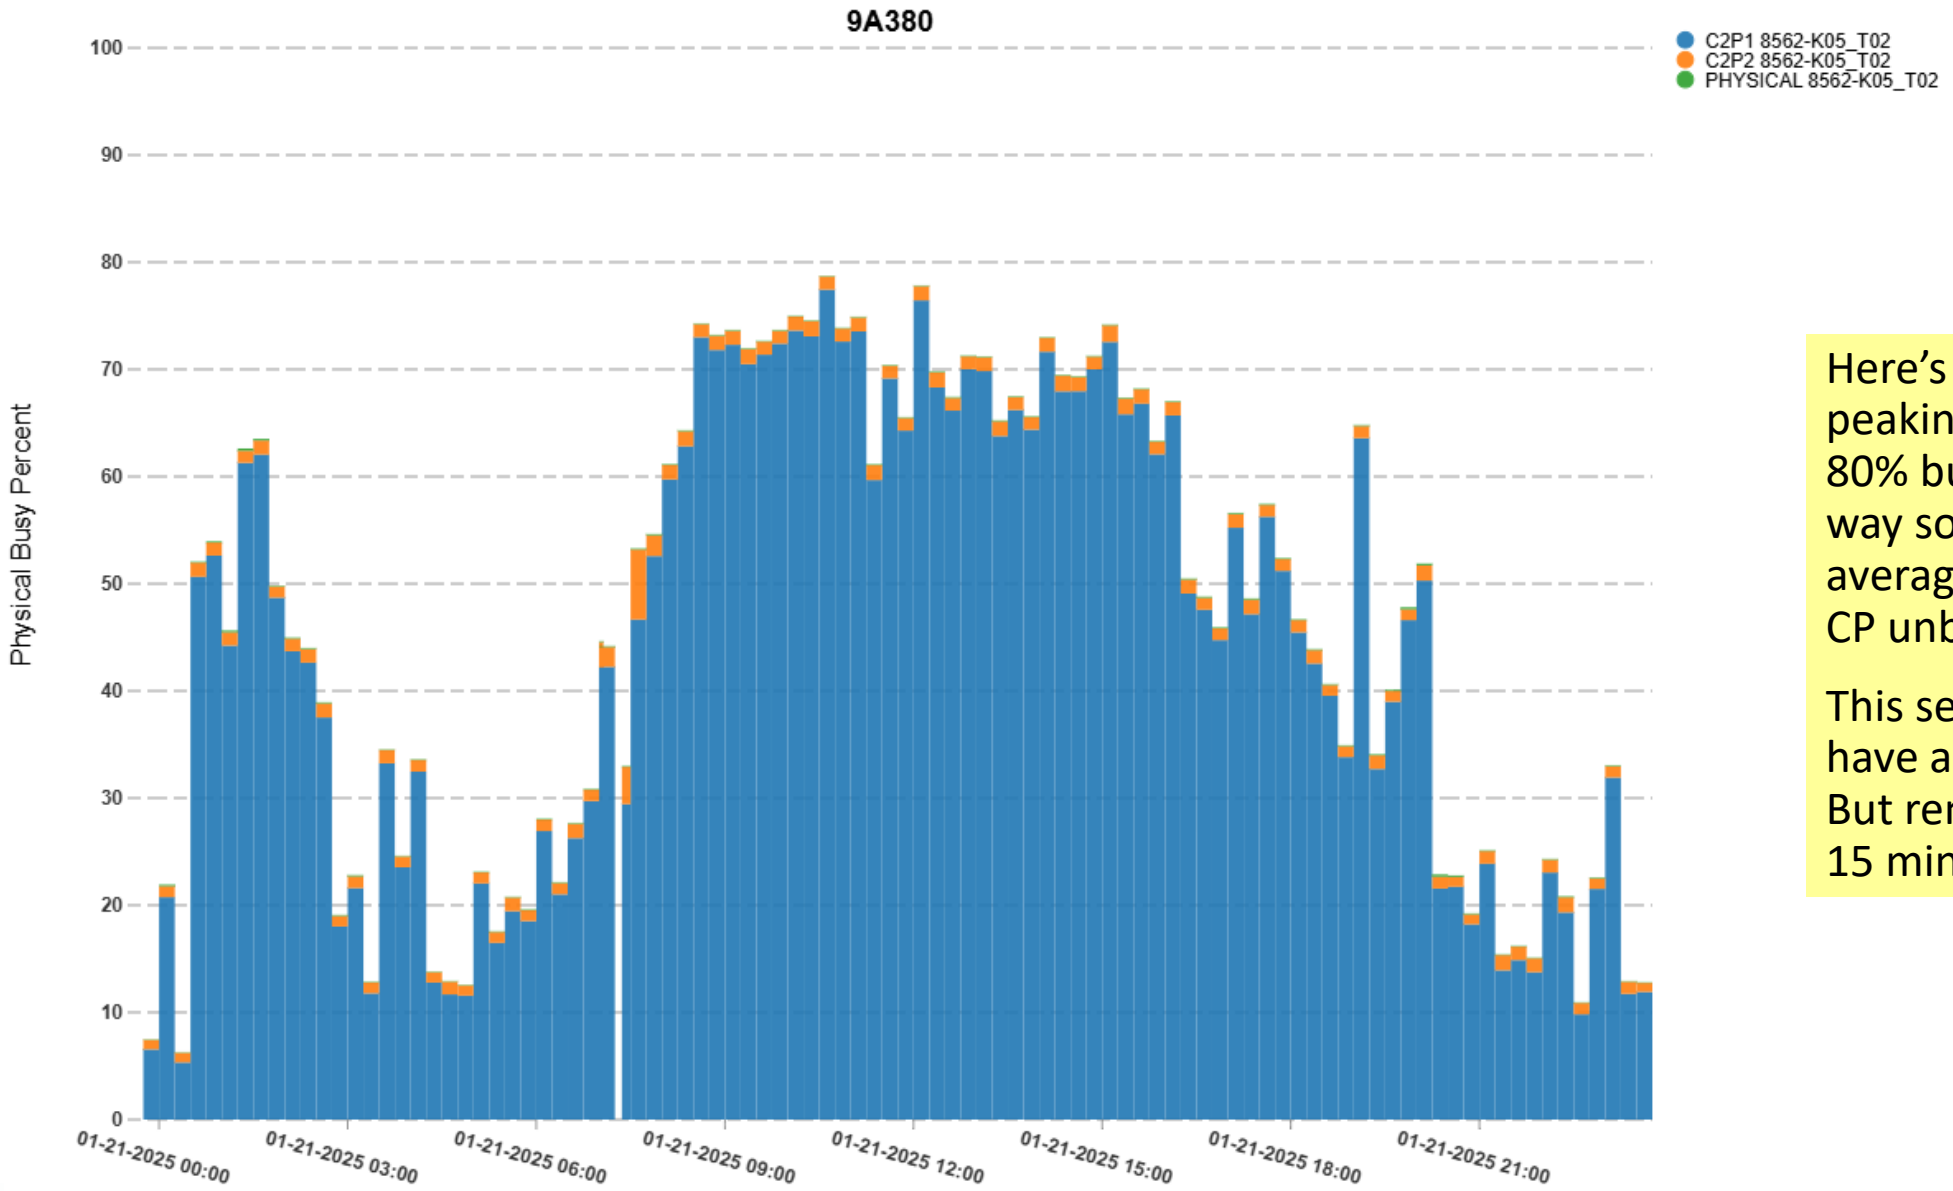

Here's a machine that's peaking out at less than 80% busy. And it's a 5-way so that means on average there's always 1 CP unbusy.

This seems "we don't have any CPU issues". But remember these are 15 minute averages.

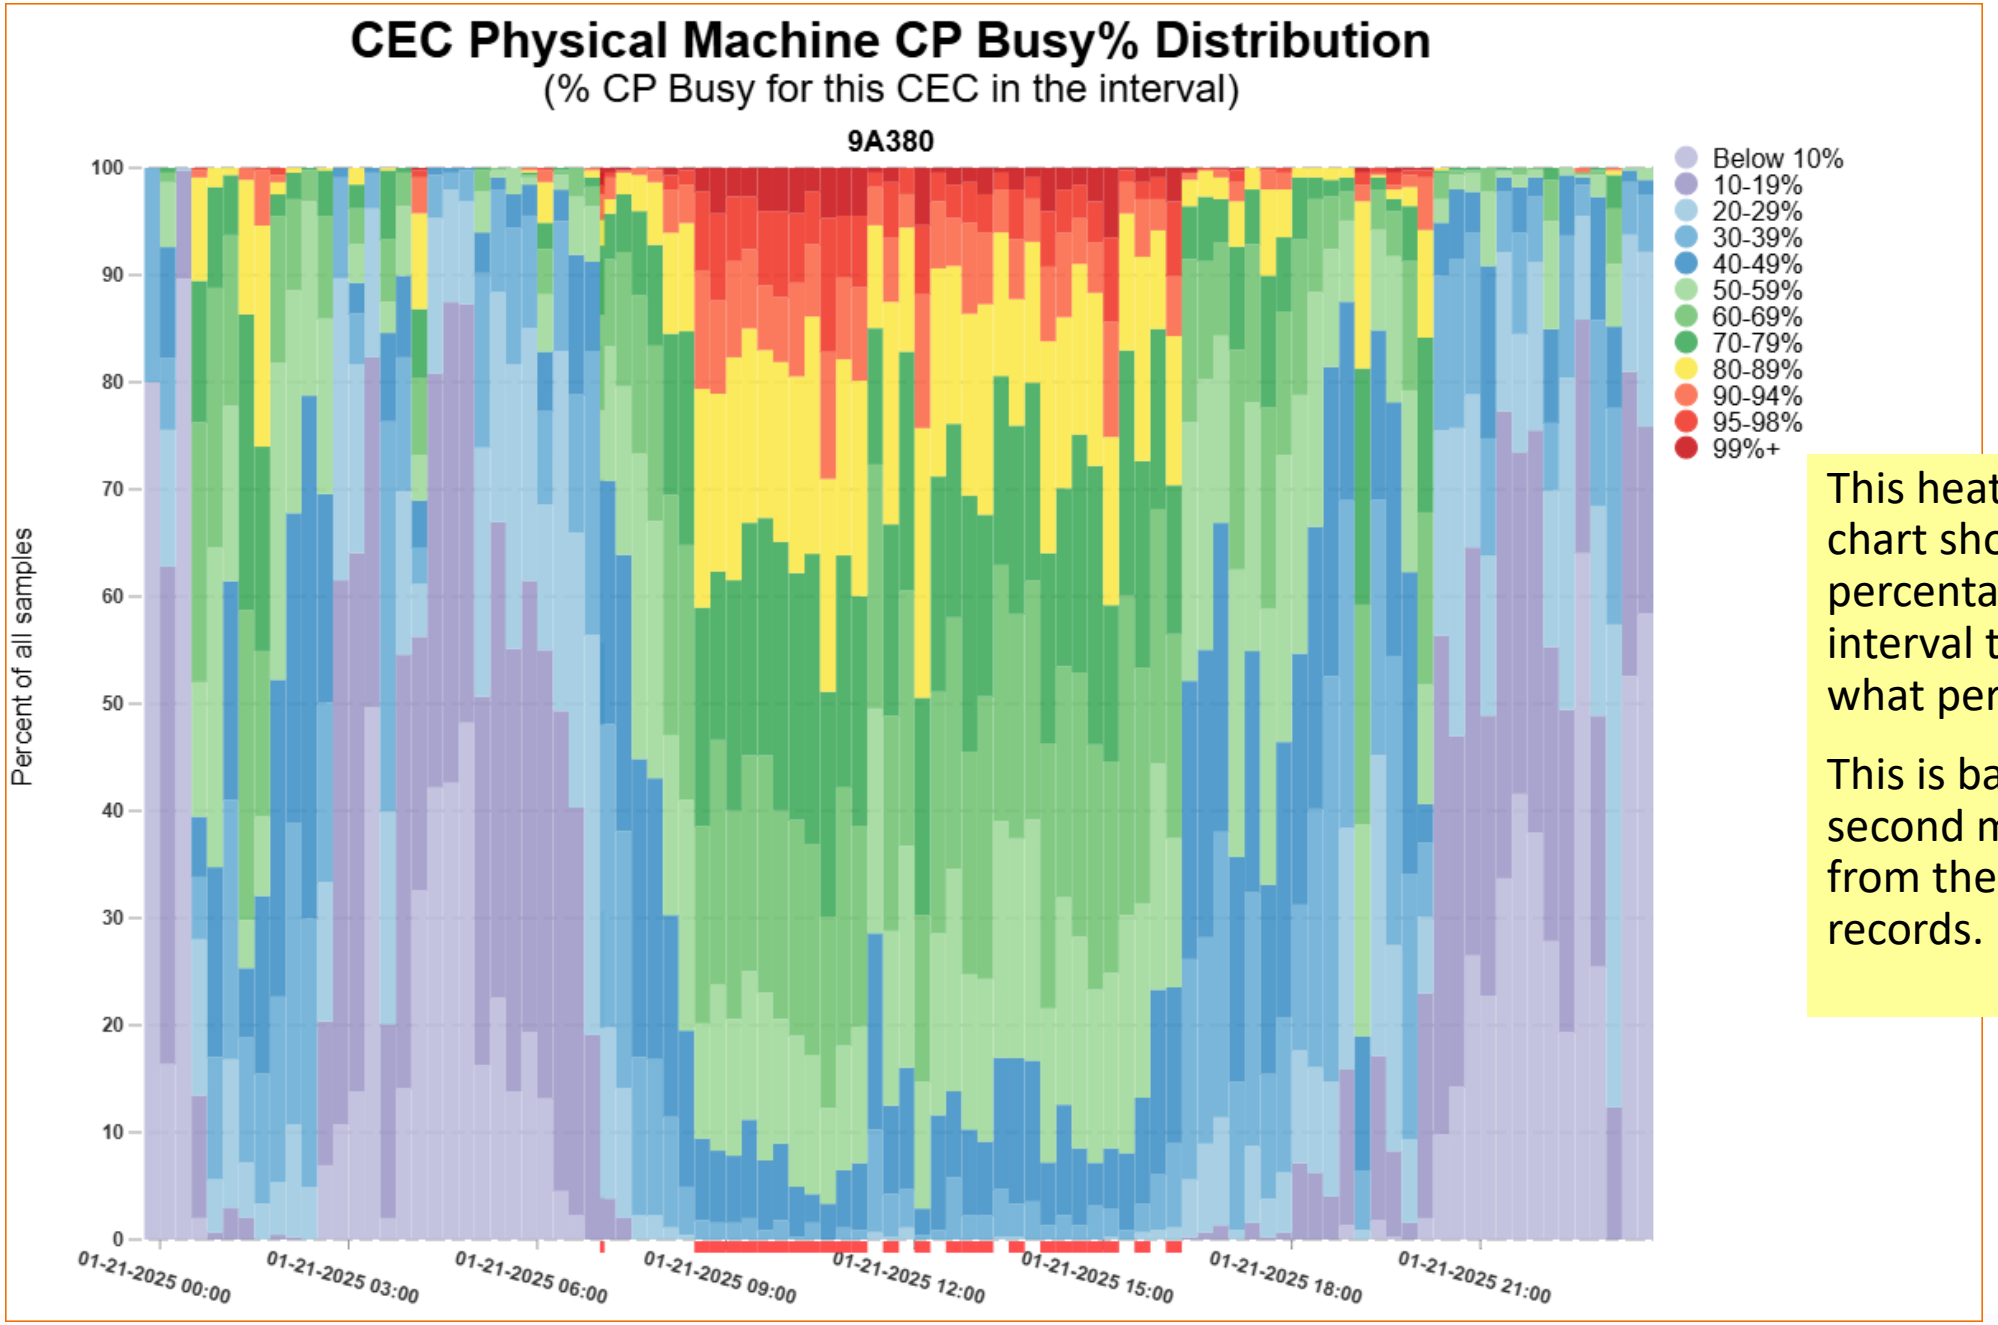

This heat distribution chart shows what percentage of the interval the machine was what percent busy.

This is based on 2-second measurements from the SMF 99.12 records.

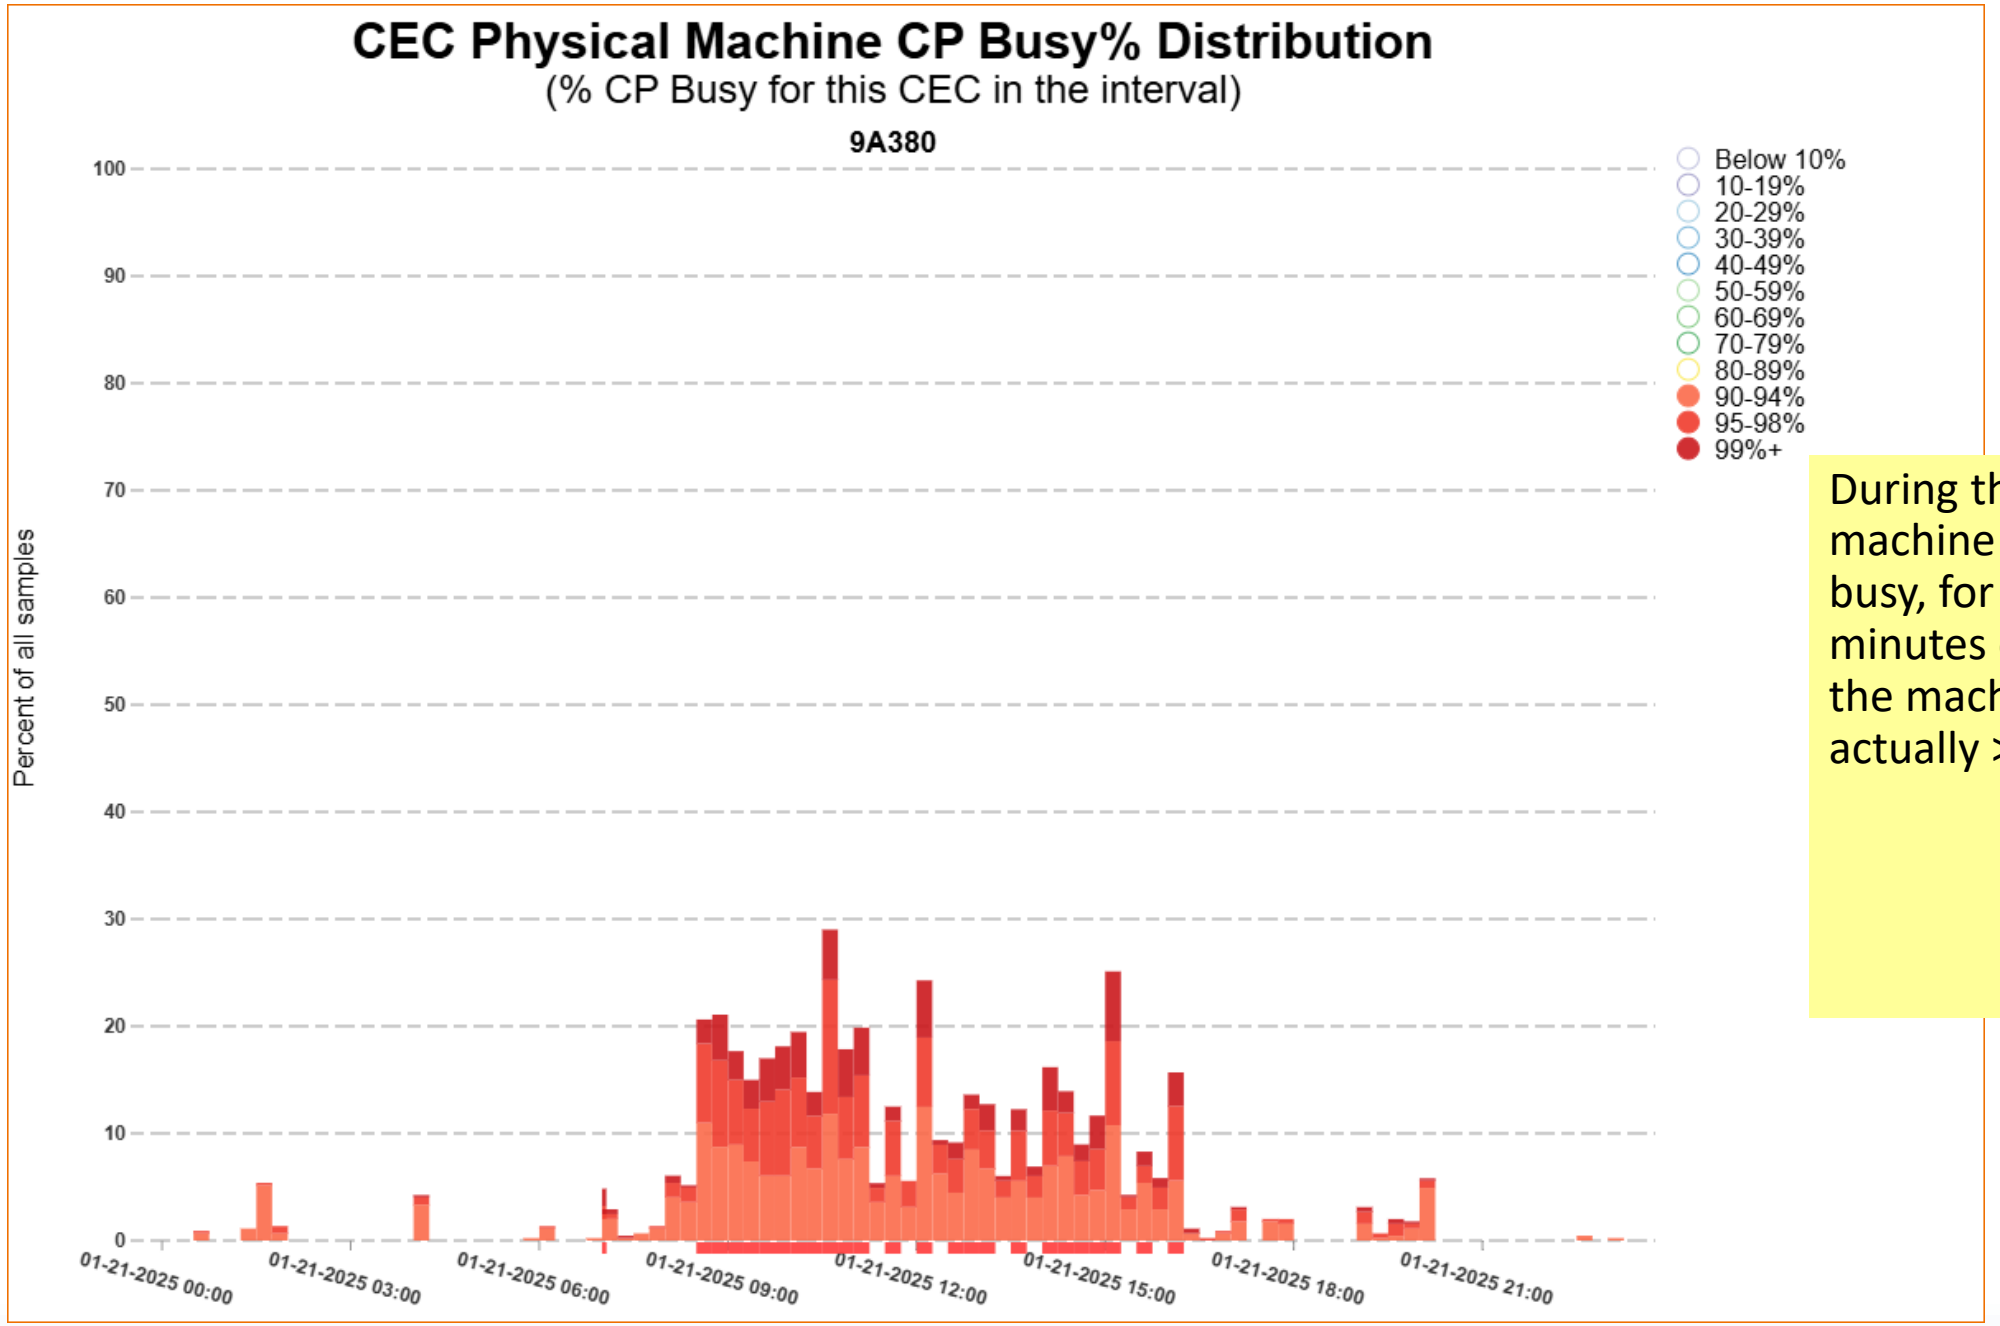

During the day when the machine was under 80% busy, for about 1-3 minutes of each interval the machine was actually >= 90% busy.

WLM Delay Samples by Period

(Top Periods with Delay Samples)

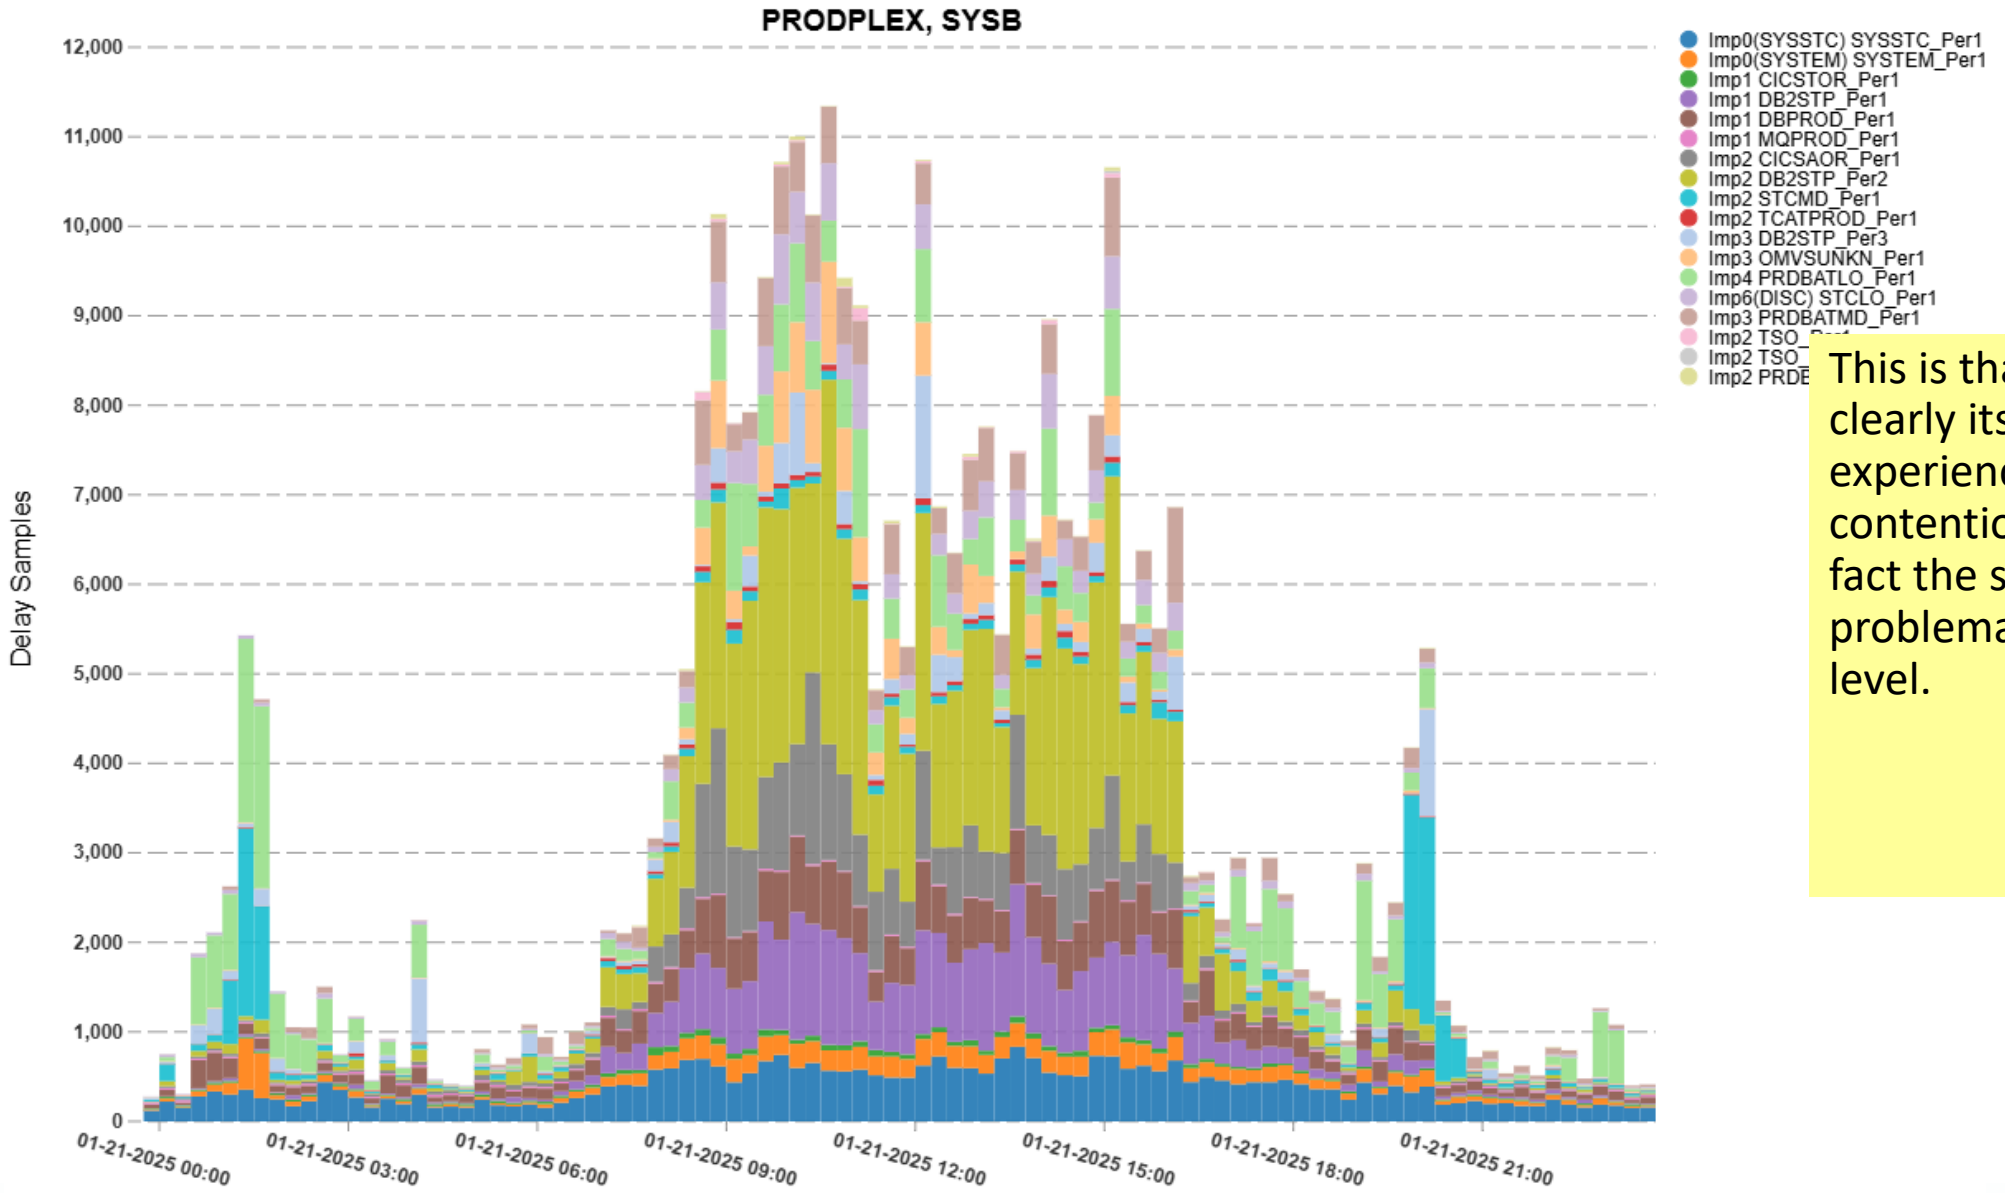

This is that blue LPAR: clearly its workloads are experiencing some CPU contention despite the fact the seemingly non-problematic utilization level.

So…

- Even “unbusy” systems likely experience CPU contention at times
- Dispatching priority impacts how much CPU delay work will experience during times of contention
- WLM manages dispatching priorities
- How well WLM manages dispatching priorities depends on how well your policy matches your current:
  - System configuration
  - Workload characteristics
  - Business needs

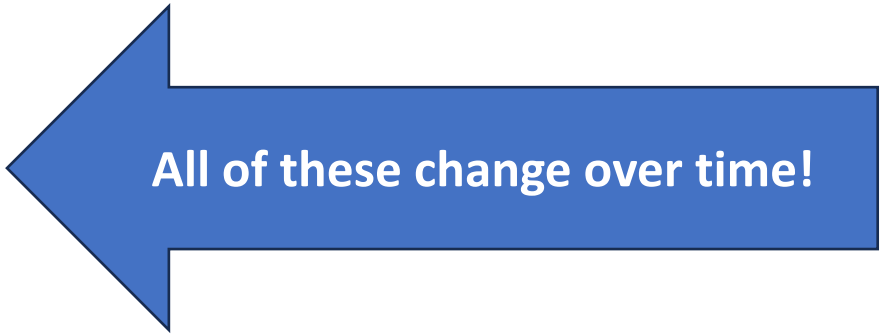

What changes matter?

System Changes

- Achievable velocities can easily change due to system configuration changes
  - Number/speed of CPs
  - Number of logical CPs
  - LPAR weights
  - I/O avoidance due to buffer changes
  - Additional I/O due to… many things
  - I/O RT changes due to… many things
- Similarly, achievable response times can change with all of the above
  - Of course that doesn't change what the business needs for a RT

Workload changes

- An increase or decrease in resource consumption from a workload can impact all workloads
  - Increased CPU consumption means less CPU available and increased contention
  - Decreased CPU means more CPU available and decreased contention
- Increase/decrease can come for many reasons:
  - More transactions (either for technical or business reasons)
  - Application changes
  - Database configuration changes
  - Data growth (or shrinkage)
- All of these can lead to changes in velocities and achievable RTs

Business changes

- Changing business needs can drive application changes
- Faster response times may lead to an expectation of faster response
  - Regardless of actual business need
- Business environment may require change in response time
  - Faster = more competitive?
  - Or… if AI is making the request, maybe you don't need to be as fast?
  - Shipping schedules change may change batch turn around?
- New/changed business may change relative importance of workloads
  - Optimize to best serve your most profitable customers?

Bonus: Recommendation changes

- Less frequently, but sometimes best practices/recommendations change
  - Usually in response to changes in technology
- Examples include:
  - Service definition coefficient recommendations over the decades
    - Now enforced at 1 for CPU and SRB and 0 for I/O and memory
    - Effects period durations
  - I/O Priority management
    - Initially recommended, now not due to I/O technology changes
    - Definitely affects velocities!
  - Average vs Percentile Response Time goals
    - Whether we should exclude or include long transactions may have changed over the decades
      - Trivial transactions without I/O have gotten much faster, I/O bound ones not so much
      - The longer transactions may be the ones we want to manage more closely
  - Usefulness of sysplex PI
    - Current area of interest for me

What happens if goals aren’t well-tuned?

How do "bad" goals impact WLM?

- WLM will only select one service class period (SCP) to help every 10 seconds
- Work doing much better or worse than its goal will tend to "drift" without much WLM management
- Work doing much better than its goal:
  - Will allow WLM to degrade the work to goal during times of stress
  - Will not be helped by WLM, may donate to other work
  - *If* the goal is correct with respect to the business needs these are good outcomes
- Work whose goal is unachievable:
  - May cause WLM to make futile changes (preventing other work from being helped)
    - And consuming a (meaninglessly) tiny bit of CPU at the same time
  - Can be "skipped" by WLM for some time when it realizes this (SCP that cried wolf?)
    - Probably doesn't matter unless it really could be helped at some point

How bad is that?

- Practically, speaking… not so bad
- But there are risks:
  - If over-performing *and* the goal is not reflective of the business needs, then during times of stress the work may not perform as the business needs
  - If under-performing, *possibly* WLM won't be as responsive as it could be to other workloads (and possibly you waste an unmeasurably tiny amount of CPU)
  - If you don't know things have changed: that's an unfortunate knowledge gap
- You don't need to review all your goals every day, but do so:
  - When you know something has changed
  - At least once a year, because probably something will change!

WLM Velocity Goal - Achieved Velocity Across Sysplex

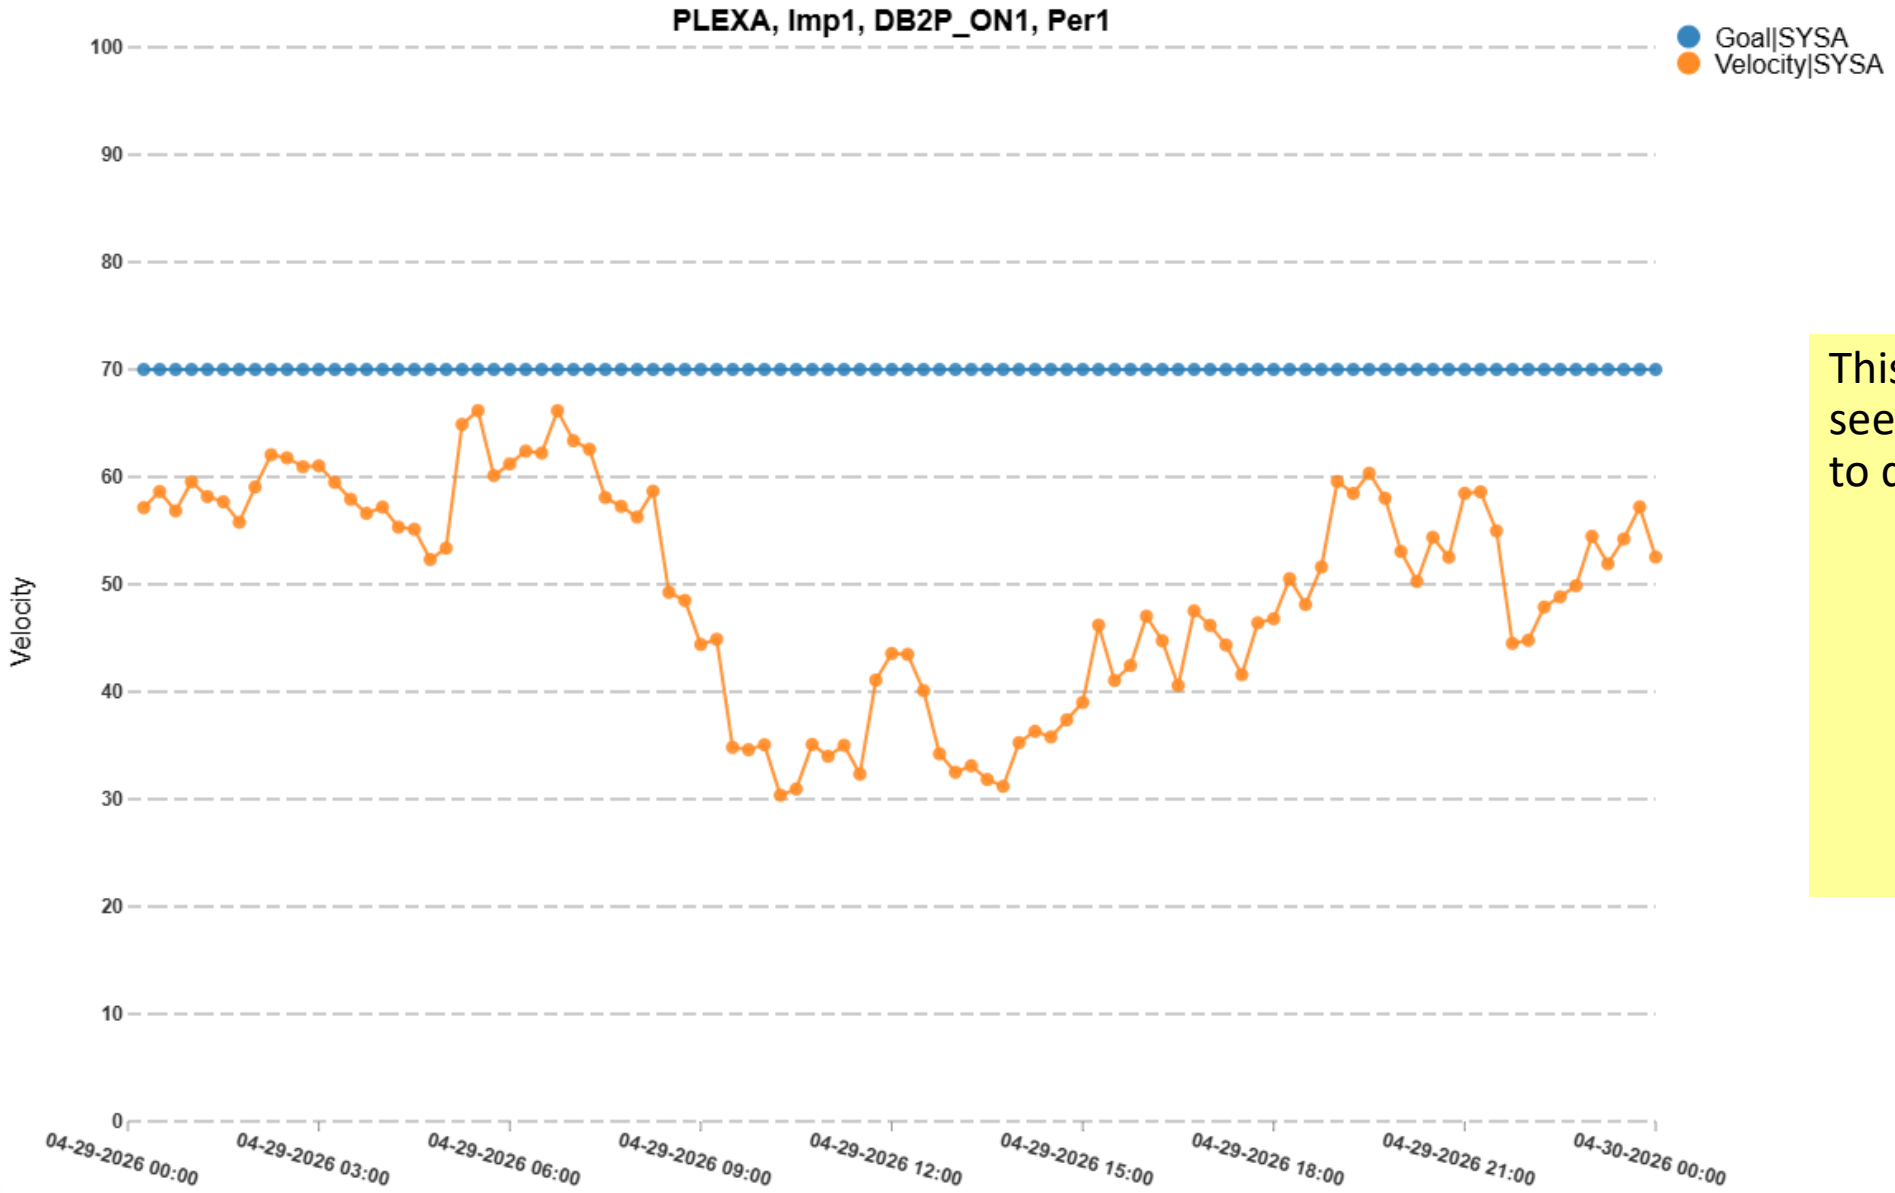

This Service Class Period seems to never be able to quite reach its goal.

WLM PA - Receiver Selections by Hour

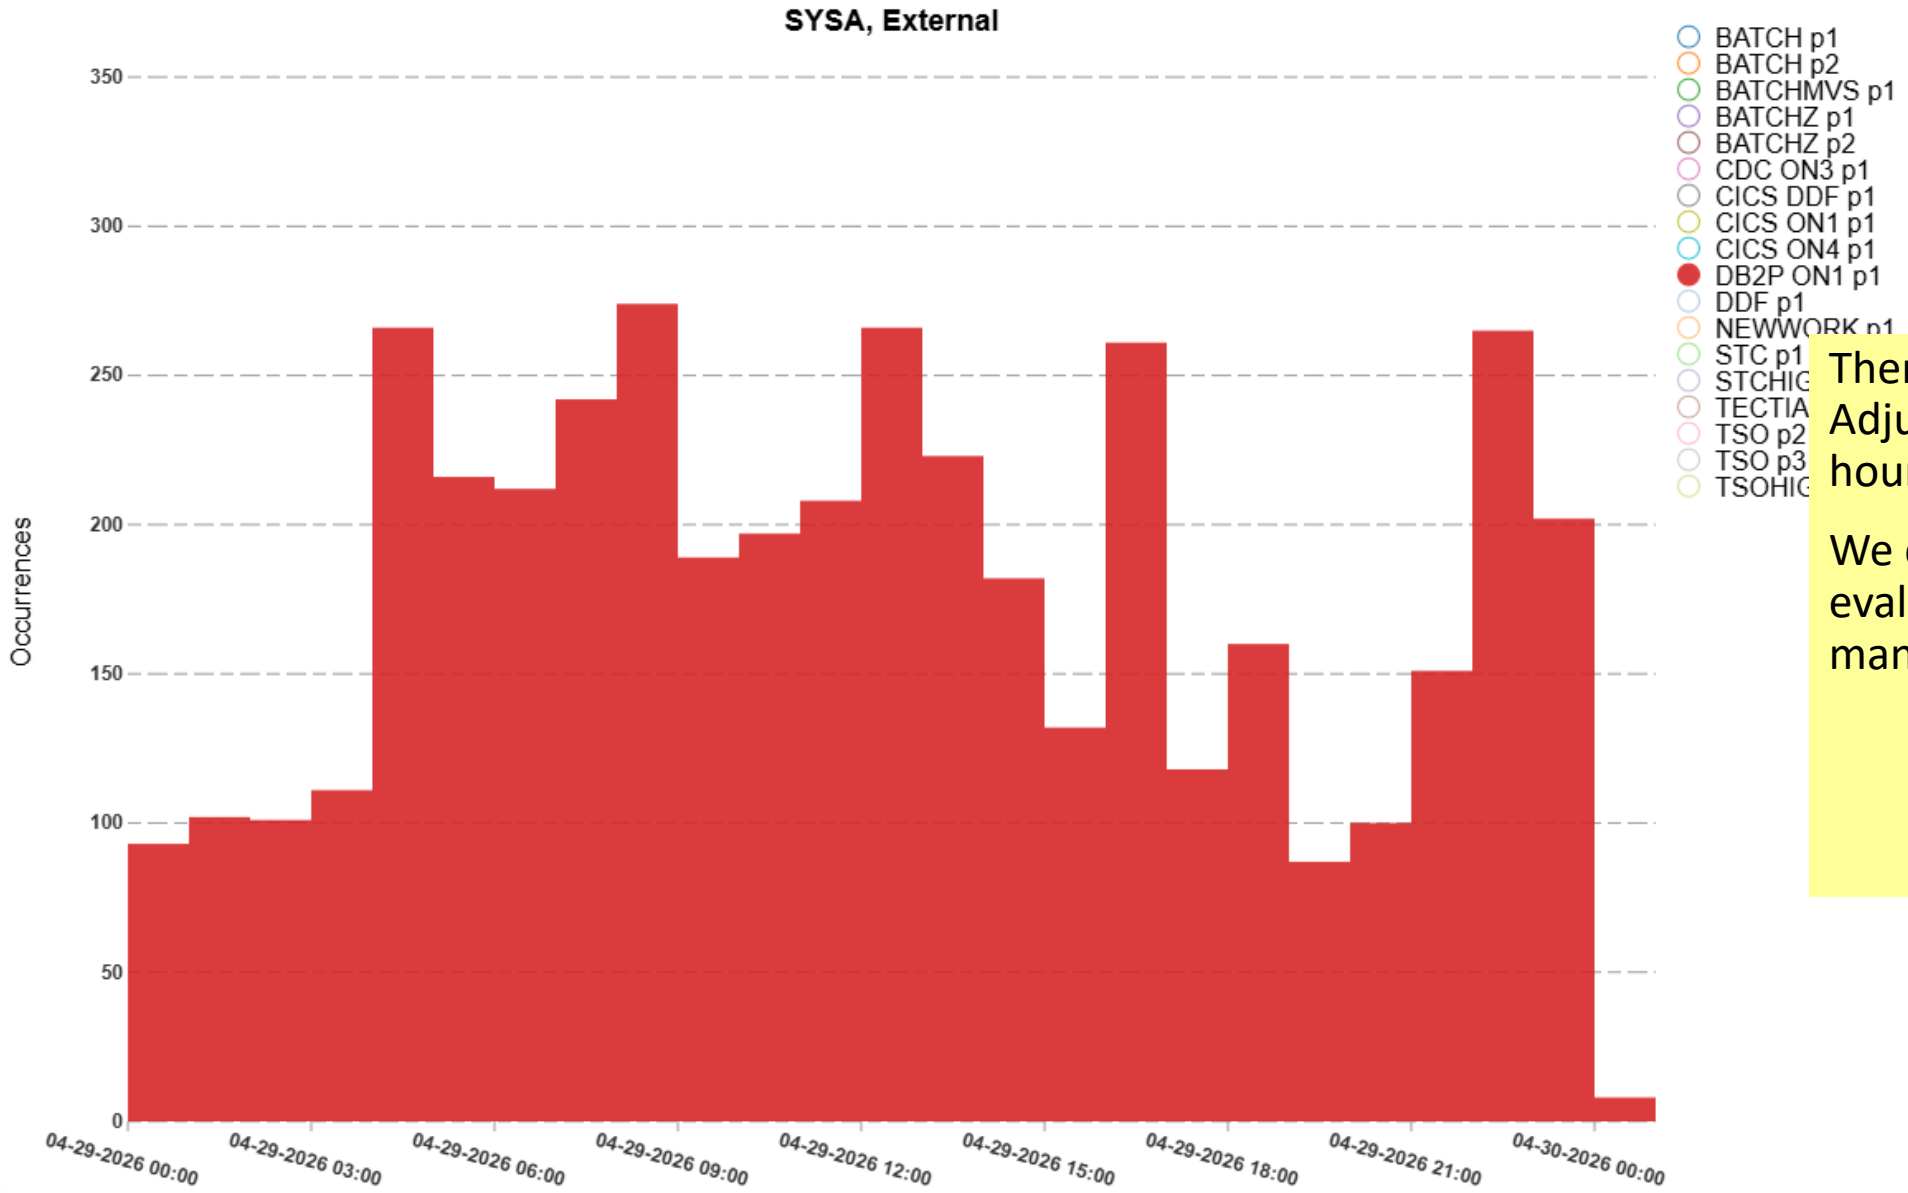

There are 360 Policy Adjustment cycles in an hour.

We can see that WLM is evaluating this SCP in many of those periods.

Using & Delay Samples Breakdown

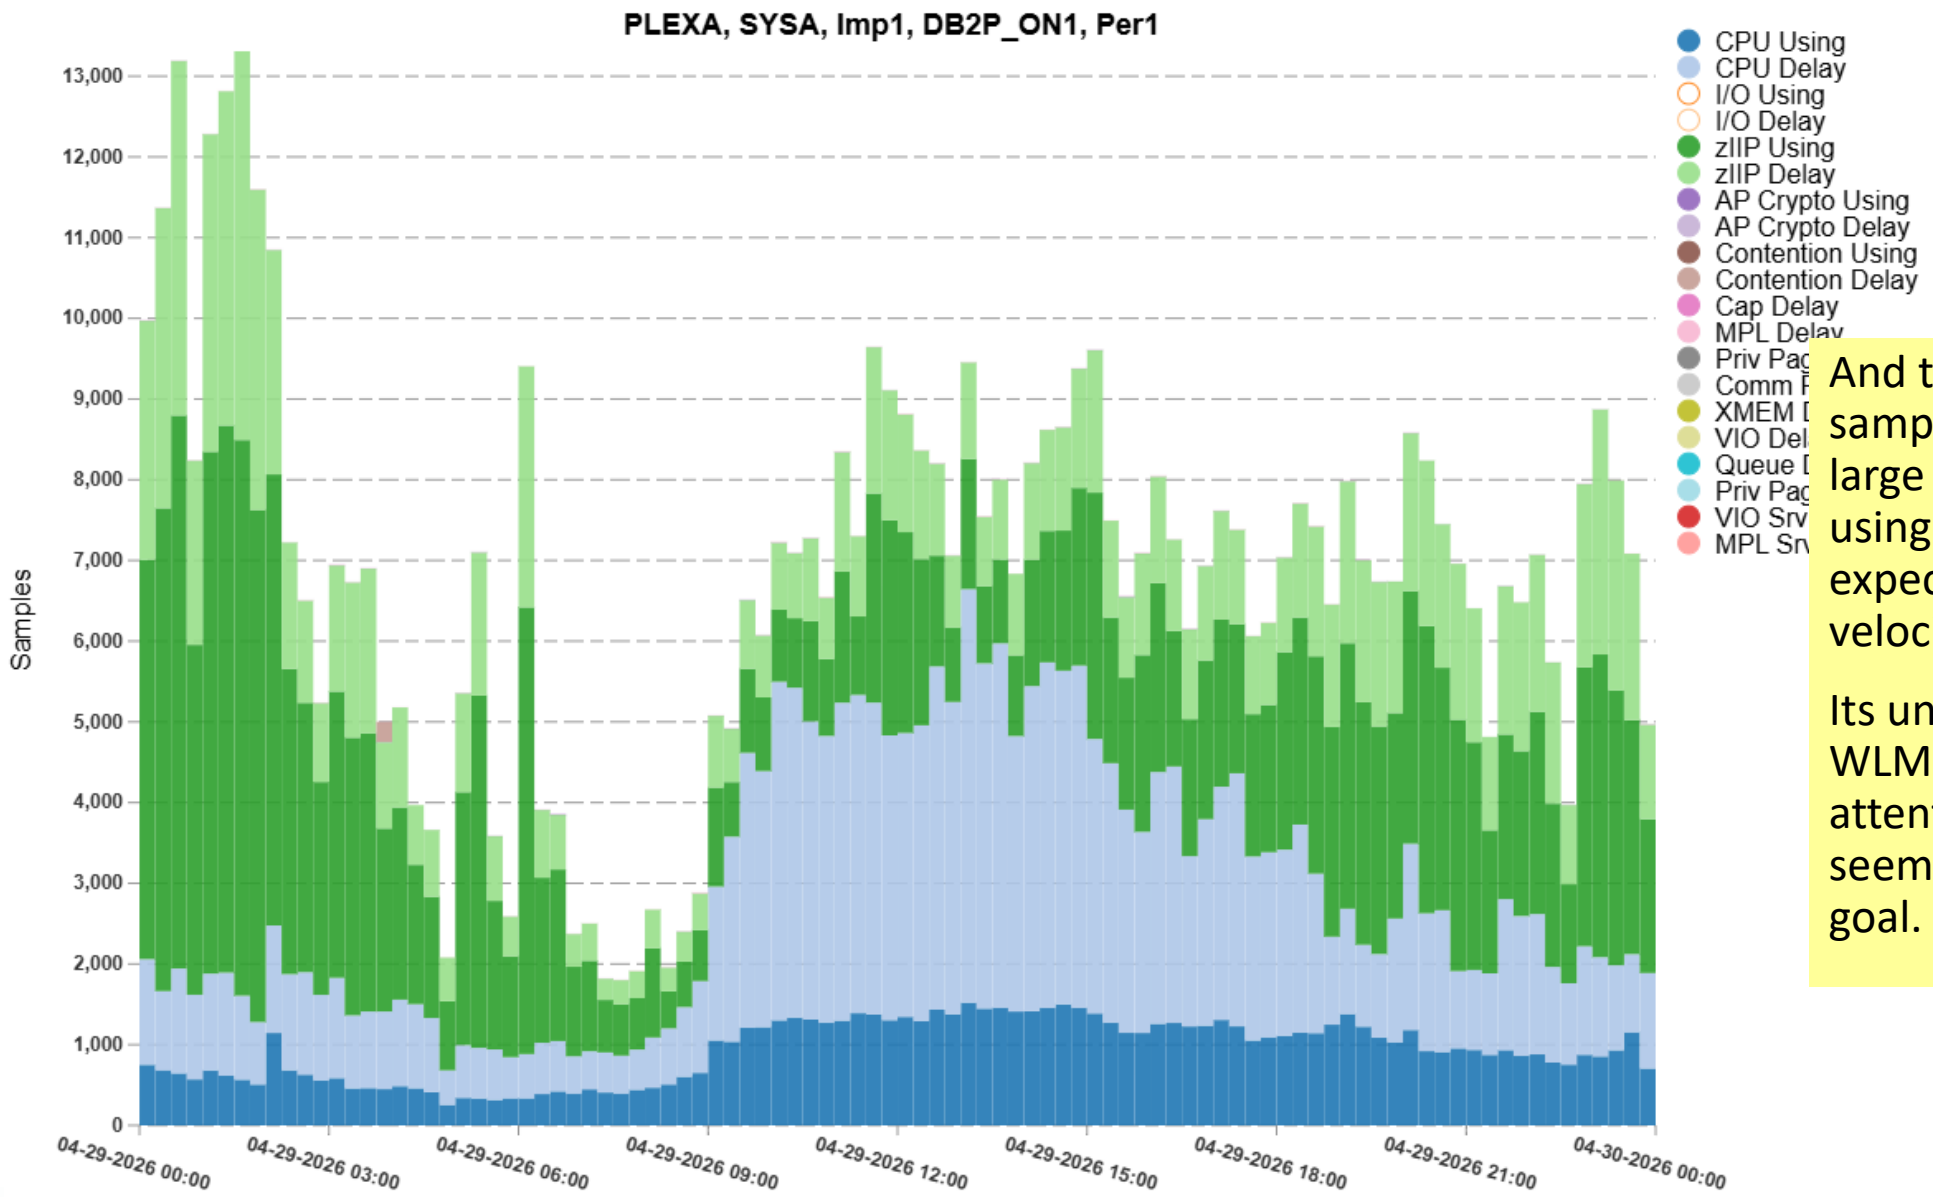

And the CPU delay samples are relatively large compared to the using samples. (As we'd expect from the velocity.)

Its underperforming, WLM is paying close attention it, but it can't seem to get to the given goal. Maybe briefly?

WLM SMF 99.6 Data Explorer

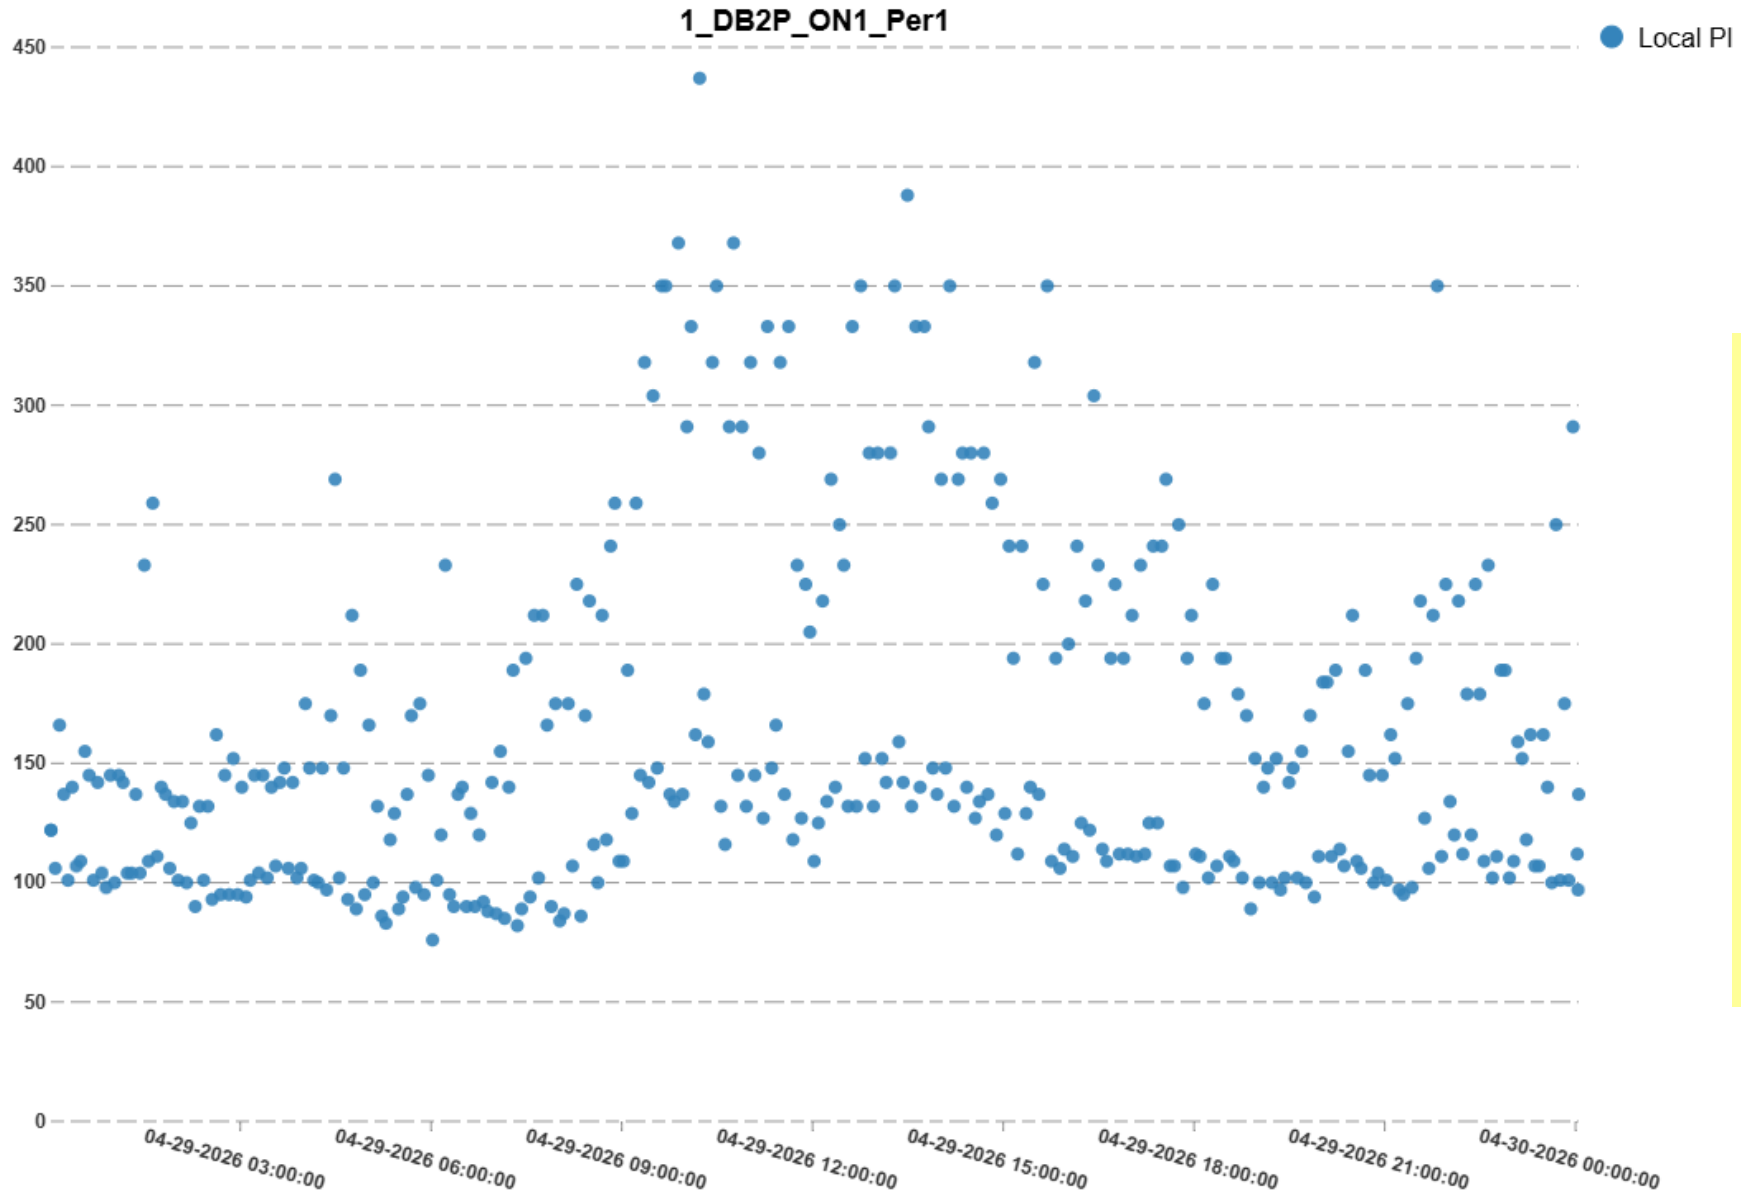

This from the 10 second interval data instead of 15 minute intervals and yes at less busy times it briefly meets its goal, but not during the busier part of the day.

Would lowering the velocity goal help anything? I'm not convinced it would be significant.

WLM Velocity Goal - Achieved Velocity Across Sysplex

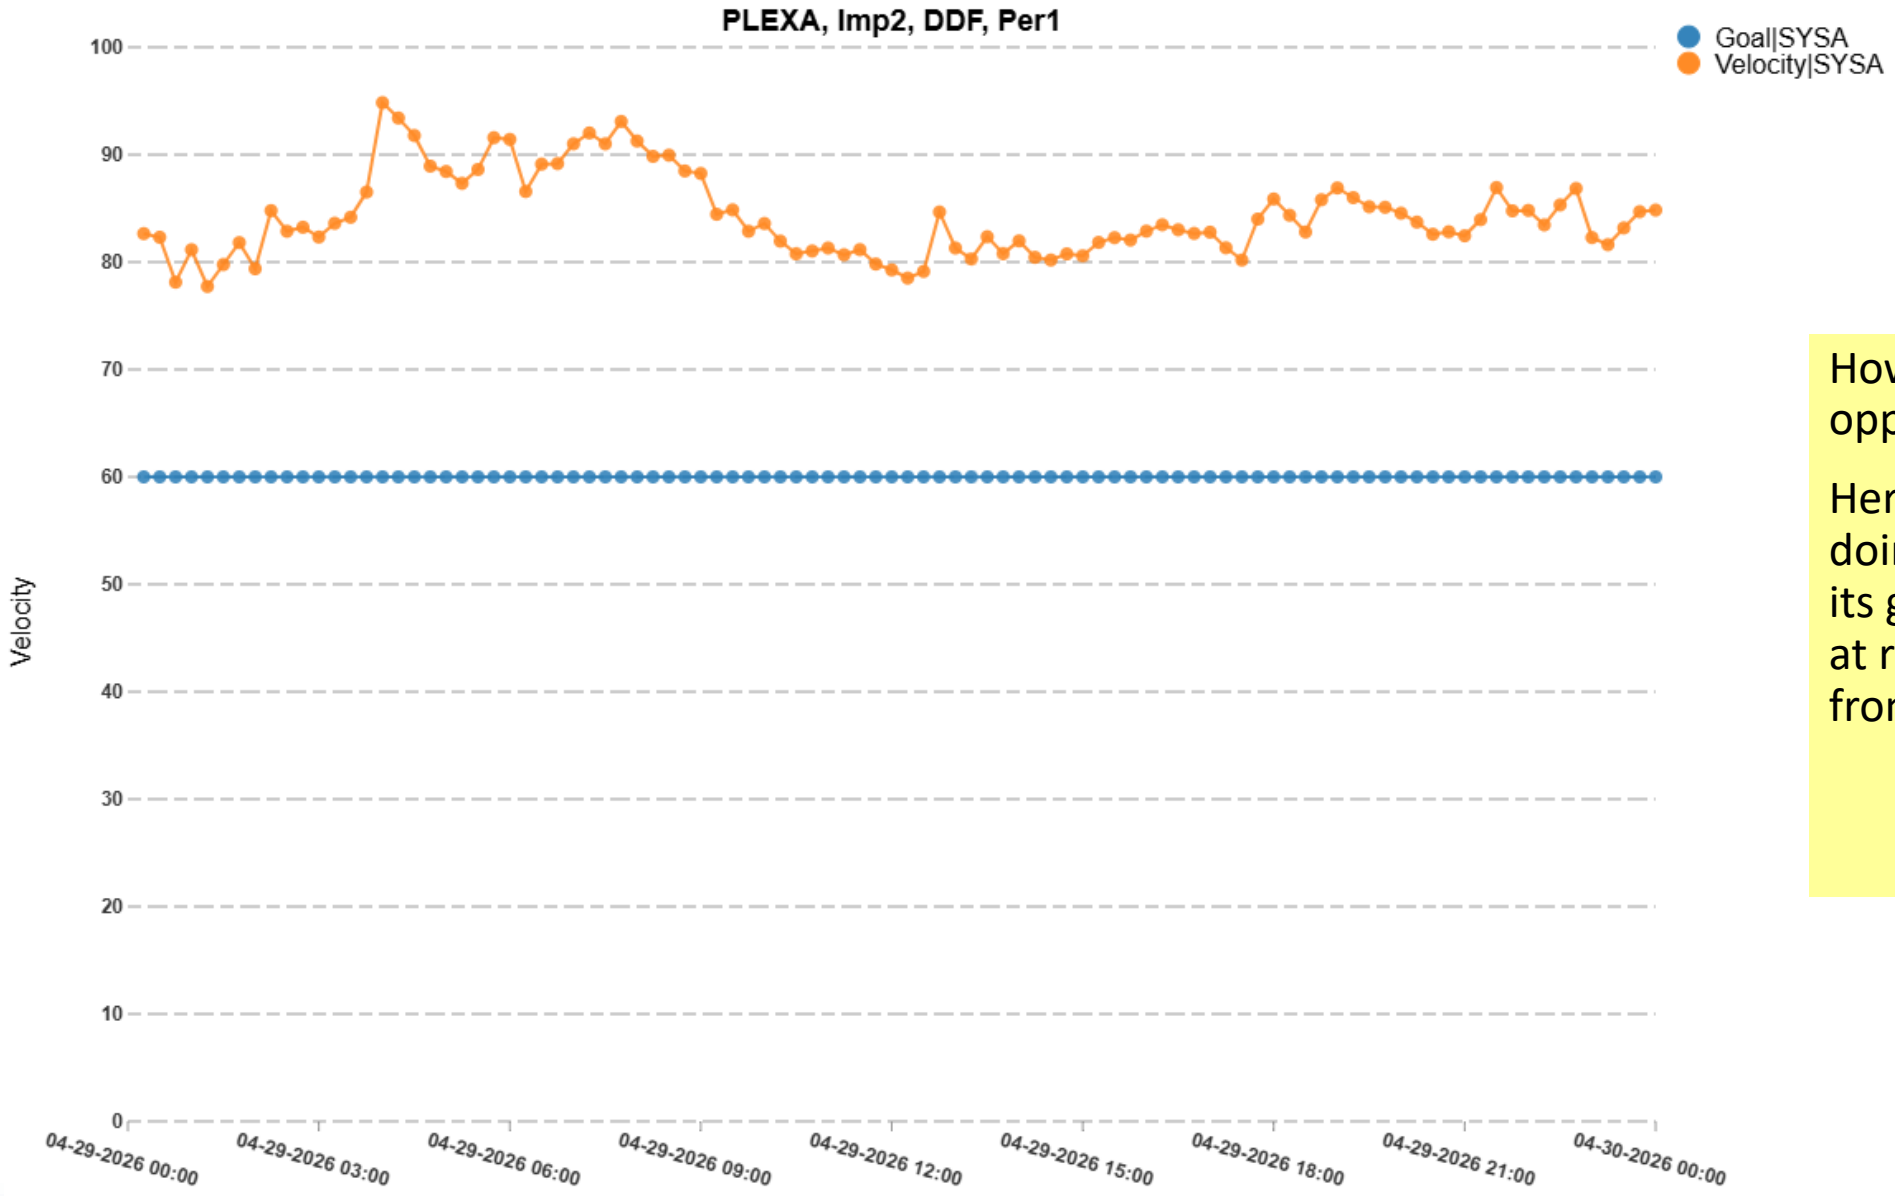

How about the opposite?

Here, this SCP is always doing much better than its goal. But it's certainly at risk of being stolen from.

WLM PA - Donor Occurrences by Hour

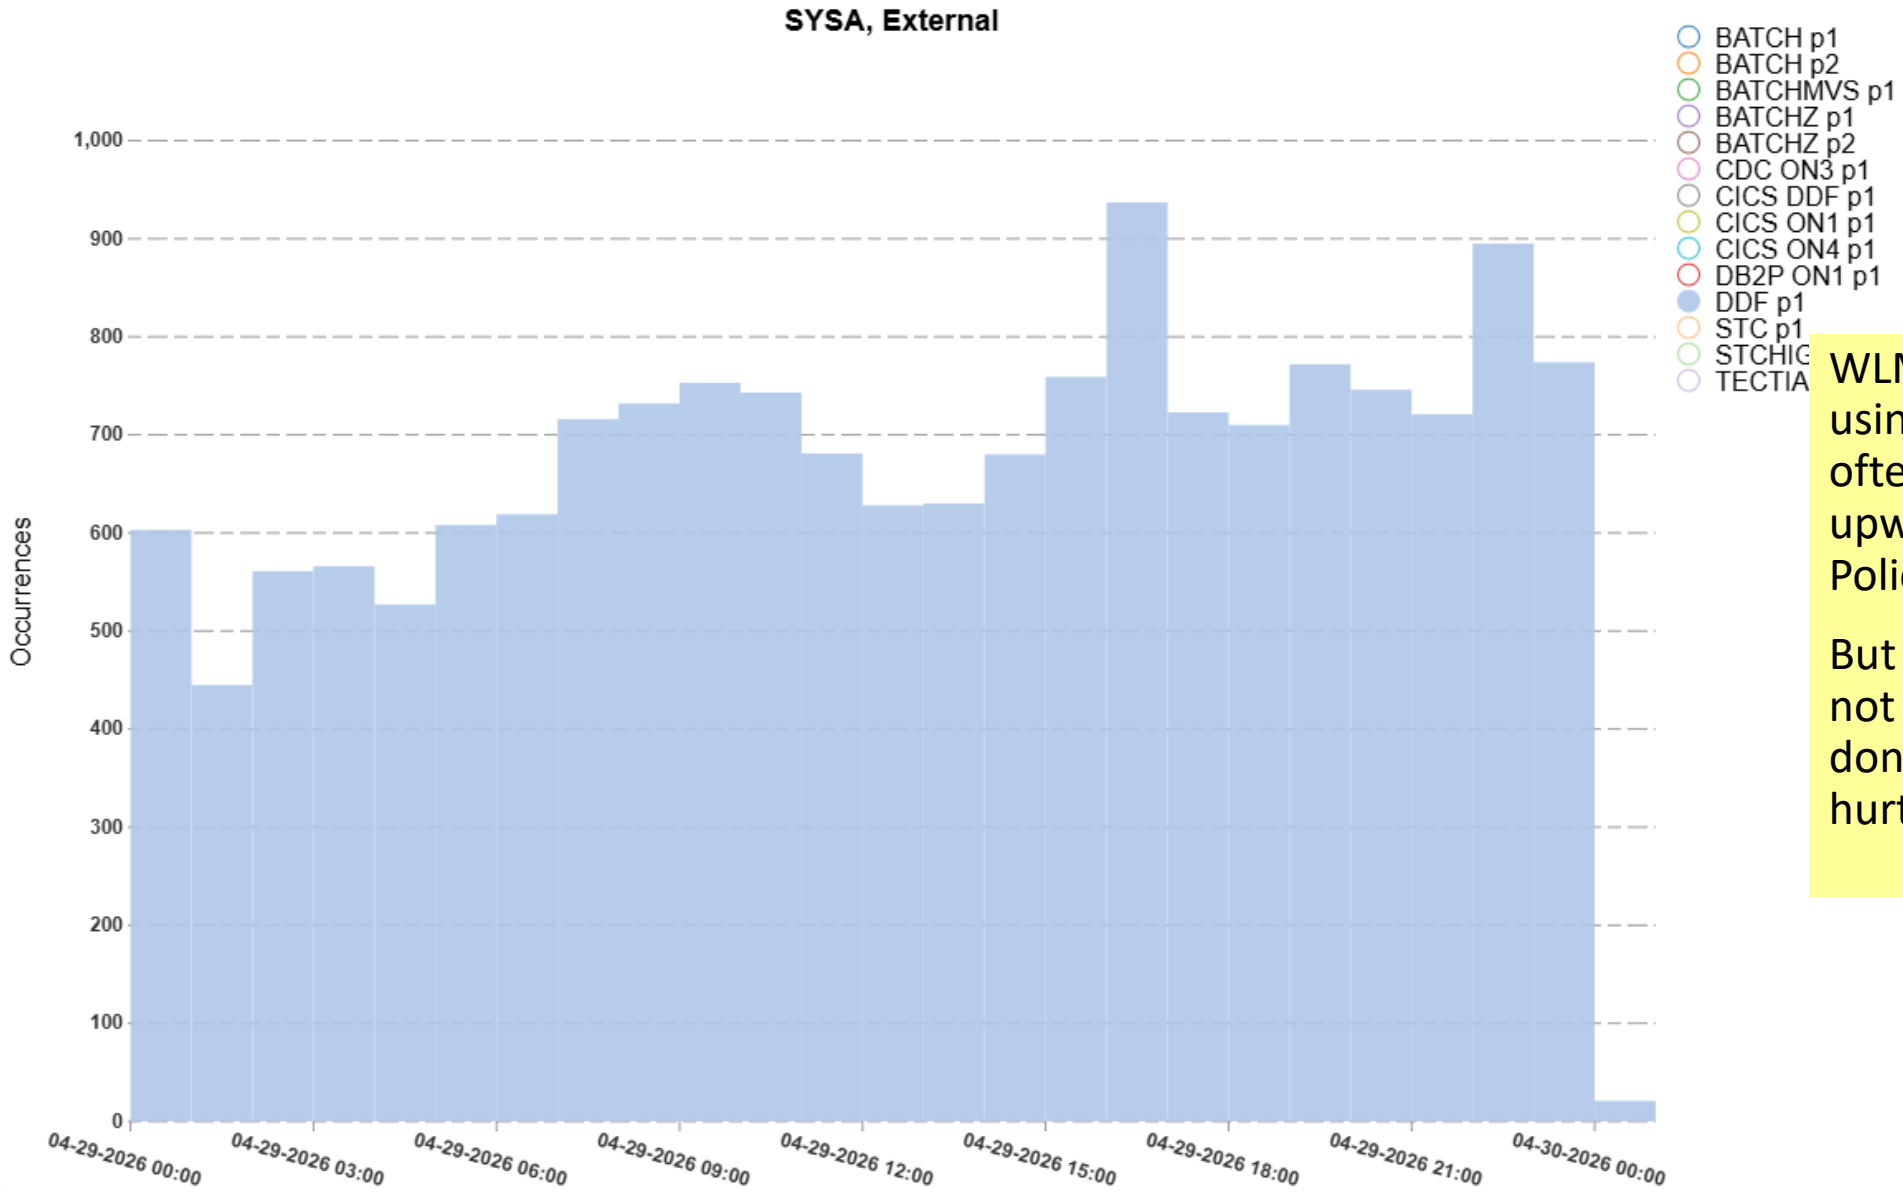

WLM is considering using it as a donor quite often—on average upwards of twice every Policy Adjustment cycle.

But apparently WLM is not selecting it as the donor, or the action isn't hurting the work much.

Using & Delay Samples Breakdown

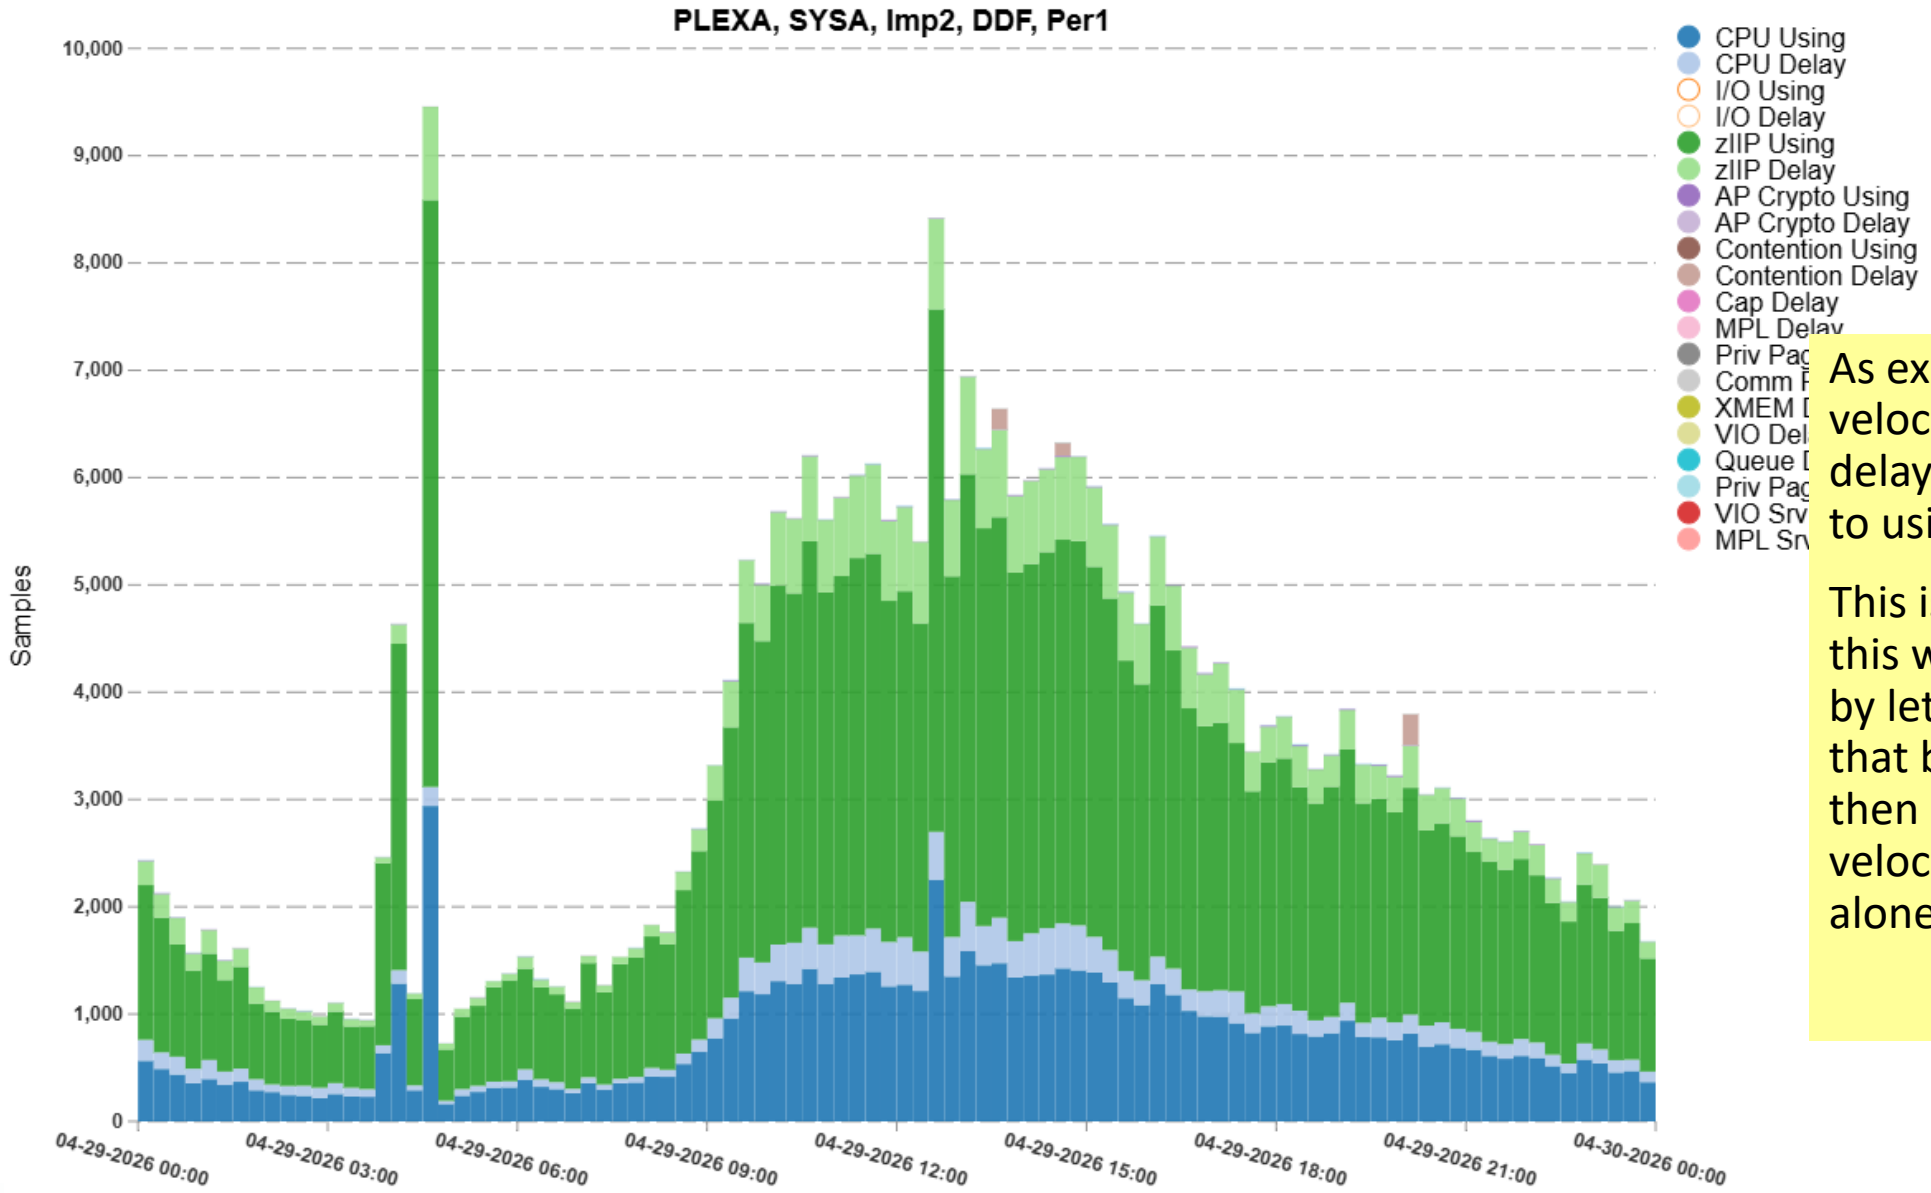

As expected from the velocity, relatively few delay samples compared to using.

This is more difficult: if this workload degraded by let's say 25%, would that be a problem? If so, then I'd raise the velocity. If no, I'd leave it alone.

Wrap-up

Summary

- Not regularly reviewing your WLM policy risks:
  - Work not being actively managed by WLM
  - Work possibly at risk of suddenly degrading during periods of stress
  - Work possibly over-performing at the expense of some other work
  - WLM maybe not helping lower-importance work as much as it could
  - Loss of knowledge about how your system is working
  - Missed opportunity for optimization
- But some sites do get away without changing their policy for years
  - But that doesn't mean you should ignore your WLM policy!
    - But you also don't need to micro-manage it either
- Understanding whether the current performance is meeting your business needs is a very important input to setting your goals!

Thanks!
Questions?

Email: scott.chapman@epstrategies.com