Key Reports to Evaluate Usage of Parallel Access Volumes

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

- In this webinar Scott Chapman will walk through and explain several reports that will be useful when evaluating the usage of parallel access volumes (PAVs). The concepts and reasons for PAVs will be discussed; then, some key reports will be reviewed to help analyze the effectiveness of the usage of the PAVs in your z/OS I/O subsystem.

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 20-24, 2023
- WLM Performance and Re-evaluating Goals
  - October 2-6, 2023
- Parallel Sysplex and z/OS Performance Tuning
  - May 2-3, 2023
- 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!

History: Why PAV?

In the "old" days…

- Disks were SLEDs: Single, Large, Expensive Disk
  - Have to wait for rotation, head movement, etc.
  - Little or no cache
  - Could only service 1 operation at a time
- Channels were ESCON (or bus & tag) and could only do 1 I/O at a time
- Lots of things single threaded
  - 1 UCB (Unit Control Block) per volser made sense
  - I/O requesters queue to get to UCB
  - IOSQ response time component is wait for UCB

Today

- All disk is RAID (Redundant Array of Inexpensive Disk)
- Most “disk” is actually solid state
  - And may be able to service more than one operation simultaneously
- Cache is very much a thing
  - Definitely can service multiple operations simultaneously!
- Channels are FICON
  - Can do at least several operations simultaneously

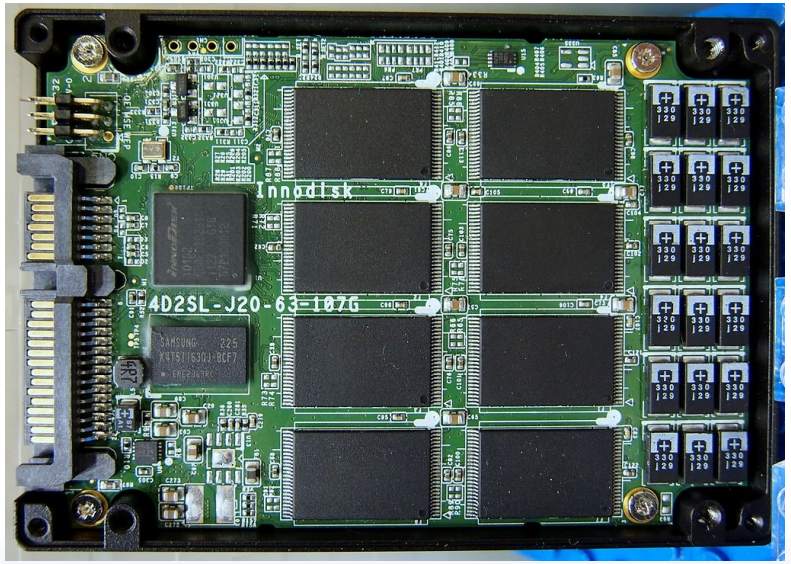

Between then and now…

- As cache started to be introduced, a volume could potentially do more than one simultaneous operation
- But MVS still had limit of 1 UCB per volser
- Hence the idea of “phantom” UCBs was born: multiple UCBs representing a single volser
- Phantom UCBs was renamed **P**arallel **A**ccess **V**olumes
  - Multiple I/Os happening in parallel (simultaneously to the same volume)
- In the early days, ESCON’s single active I/O was still a significant limit for many customers
  - FICON was introduced in 1998, but took some time to become prevalent

DASD Response Time components

- Wait time = IOSQ + PEND
- Service time = CONN + DISC
- Response Time = Wait time + Service Time

- **Response Time = IOSQ + PEND + DISCONNECT + CONNECT**

- Today we're focused on IOSQ:
  - I/O Supervisor Queue time = UCB waiting to be dispatched by IOS

Evolution of PAVs

- Static (1998?)
  - Defined in IO Gen – specify some number of aliases for each base address
- Dynamic (1999)
  - Enabled in WLM Policy
  - Dynamic Alias Tuning = Yes
  - WLM moves PAVs to where they're needed (eventually, maybe)
- HyperPAV (2006)
  - Enabled in IECIOSxx
  - HYPERPAV=YES
  - Assign an alias from the LCU for each I/O
- SuperPAV (2016)
  - Enabled in IECIOSxx
  - HYPERPAV=XPAV
  - Borrow aliases from other LCUs if need be
  - Should eliminate virtually all IOSQ

Reports/metrics to consider

You shouldn’t have a problem today… but do you?

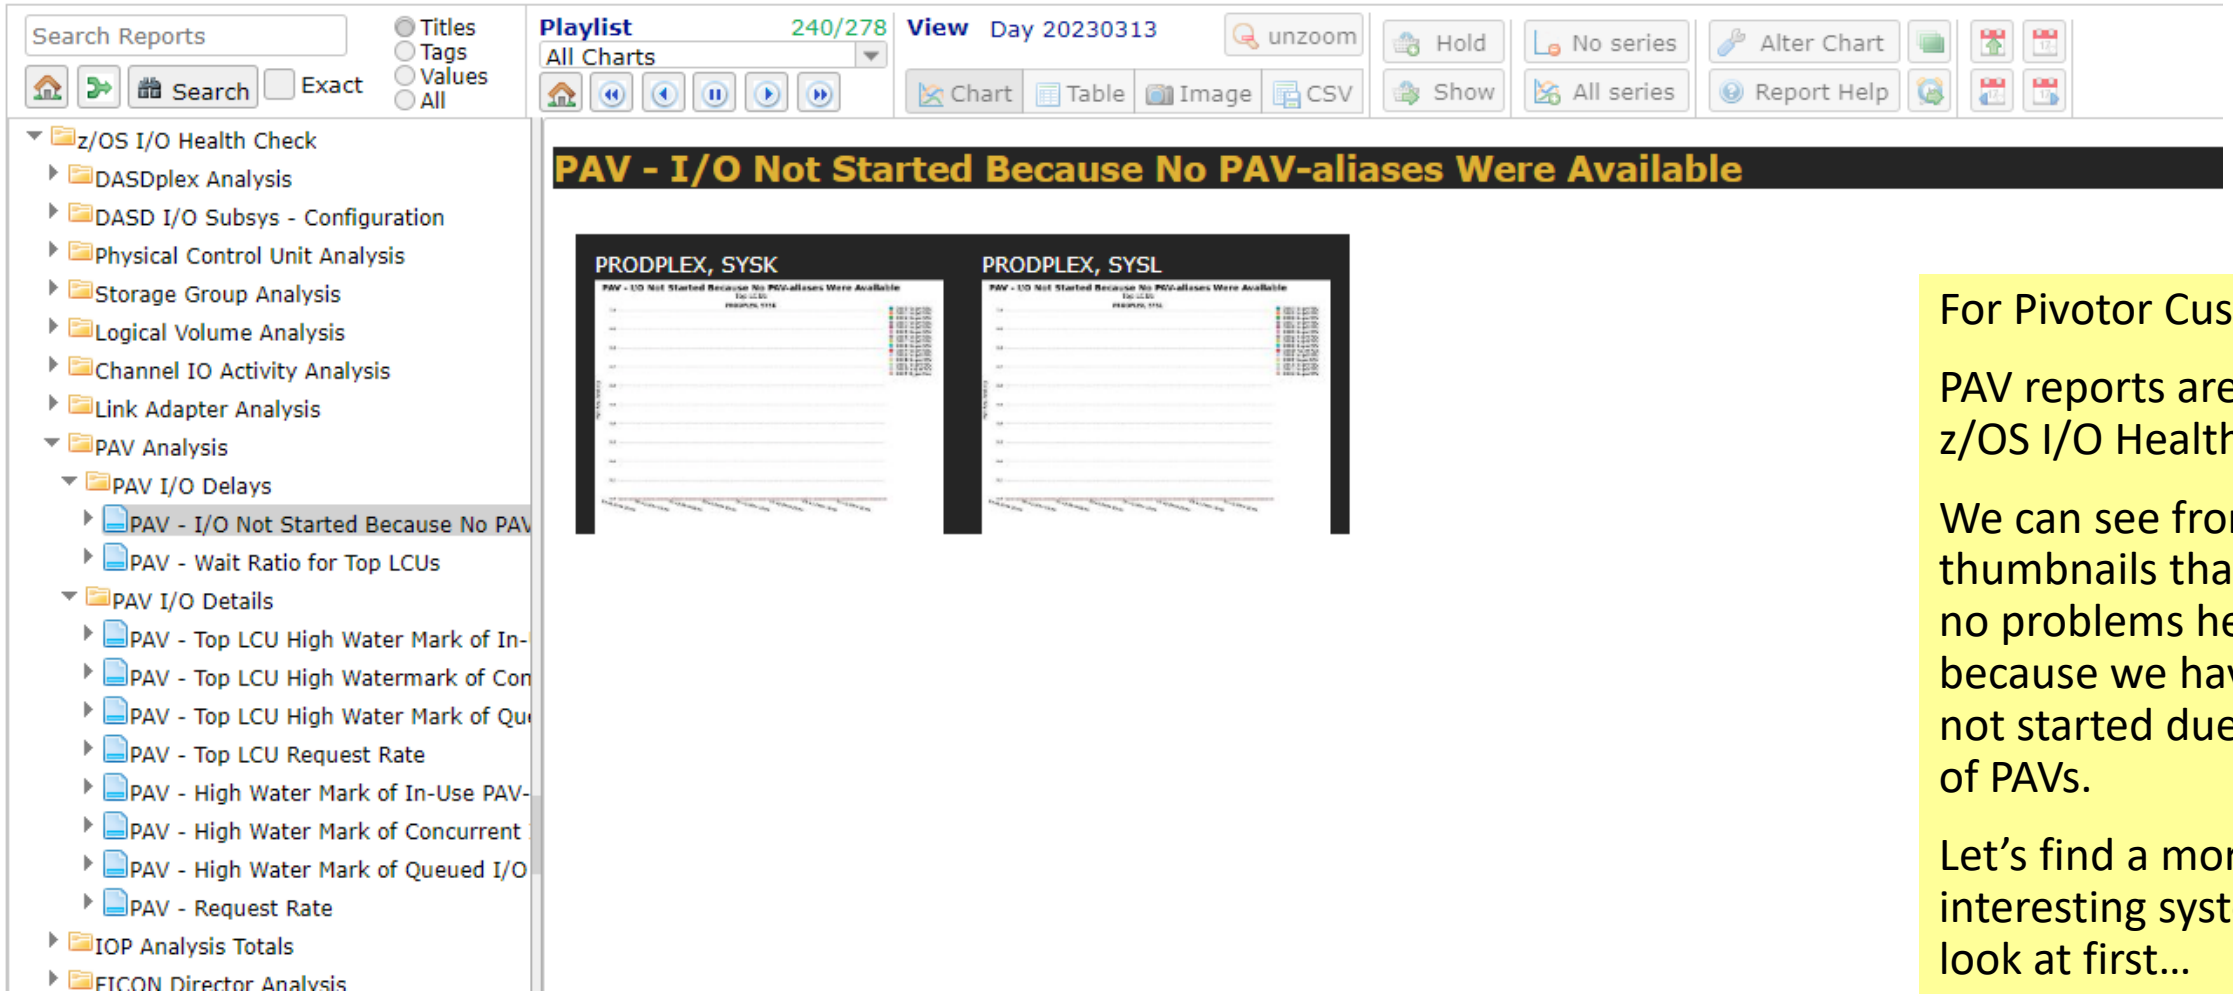

For Pivotor Customers:

PAV reports are in the z/OS I/O Health Check

We can see from the thumbnails that we have no problems here because we have no I/Os not started due to a lack of PAVs.

Let's find a more interesting system to look at first…

PAV - I/O Not Started Because No PAV-aliases Were Available

Top LCUs

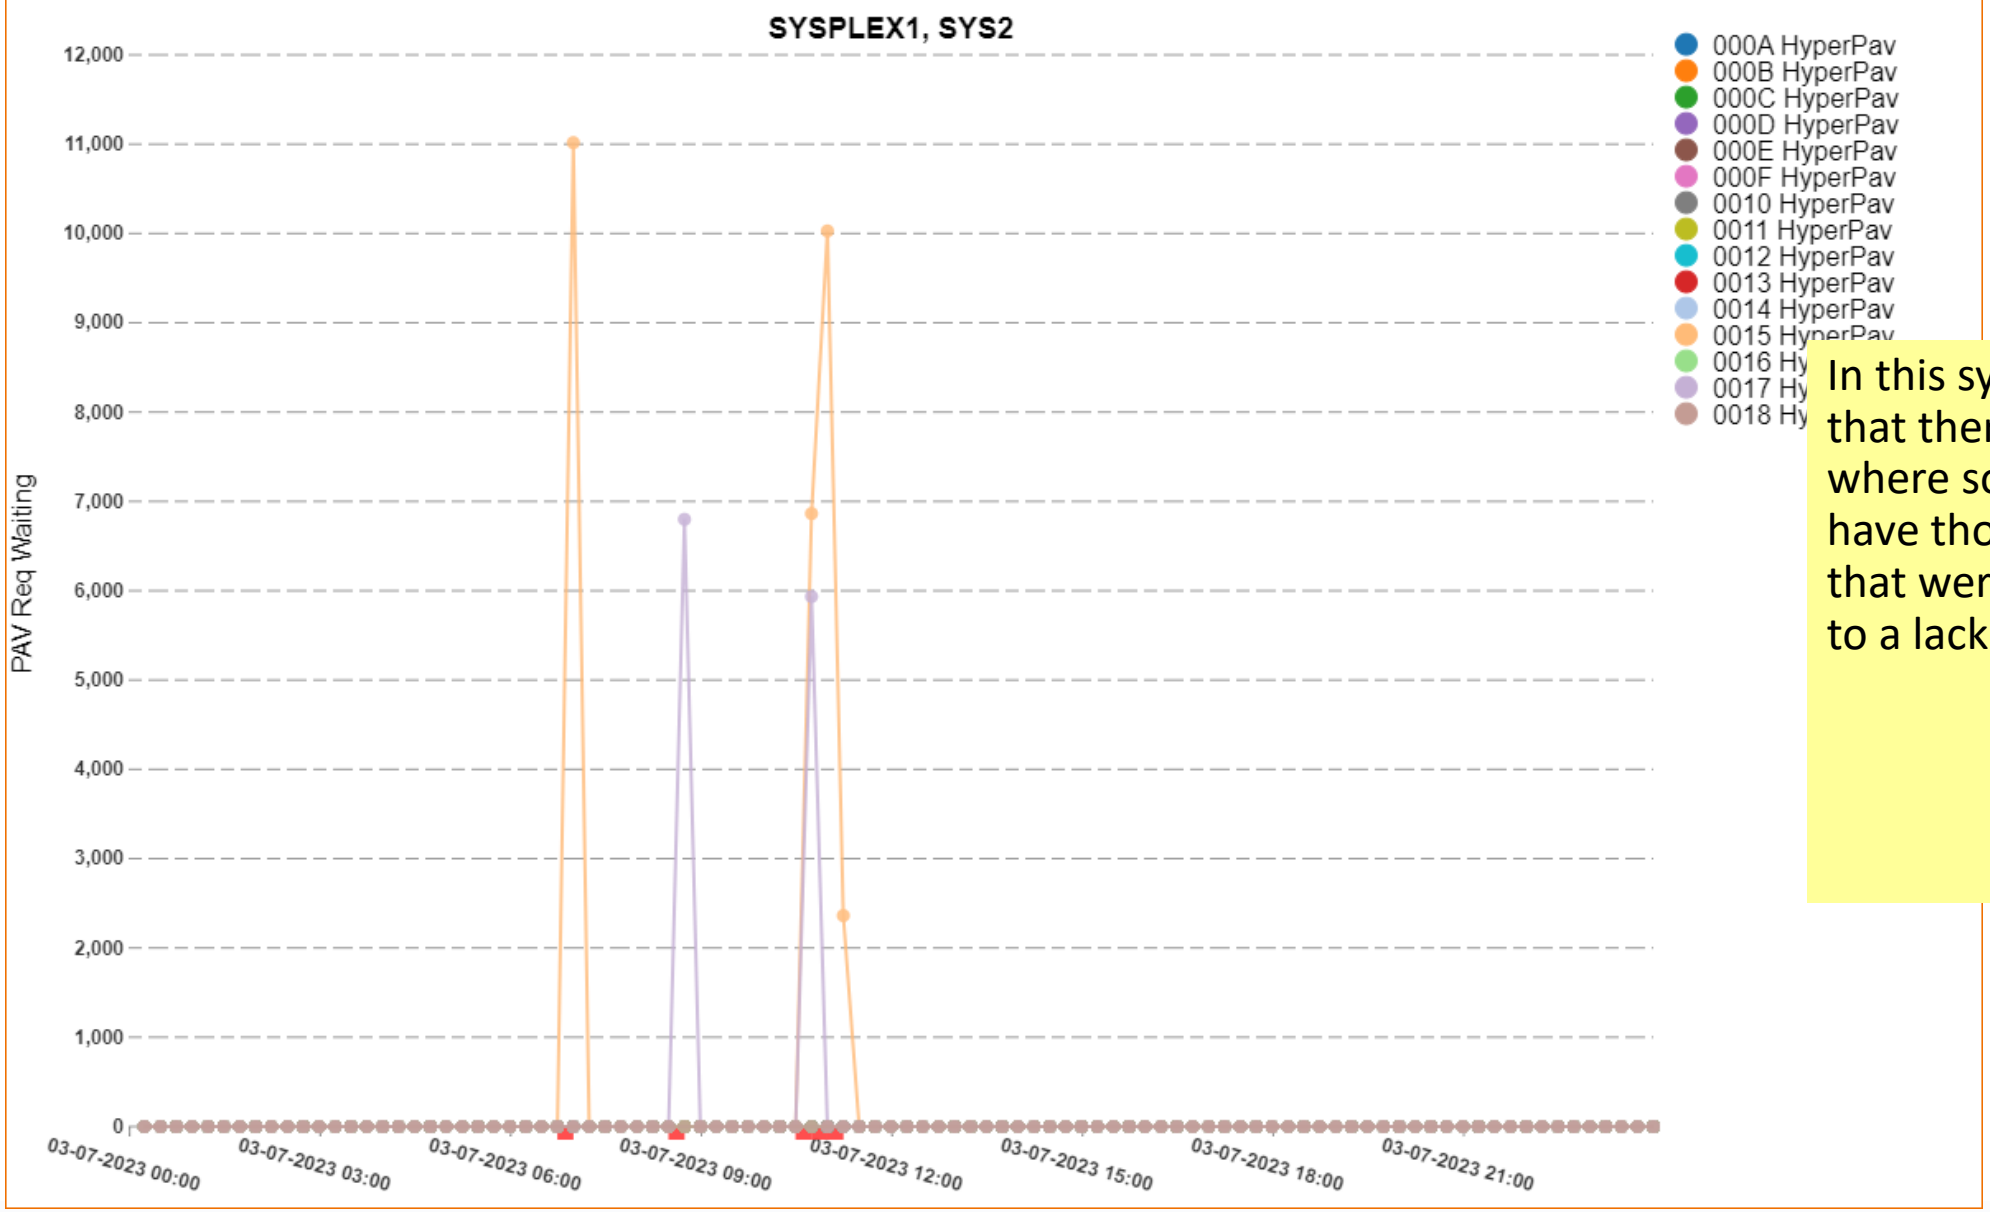

In this system we do see that there are intervals where some LCUs did have thousands of I/Os that were delayed due to a lack of PAVs.

PAV - Wait Ratio for Top LCUs

(Delayed I/Os Due to No PAV-alias to Number of PAV Requests)

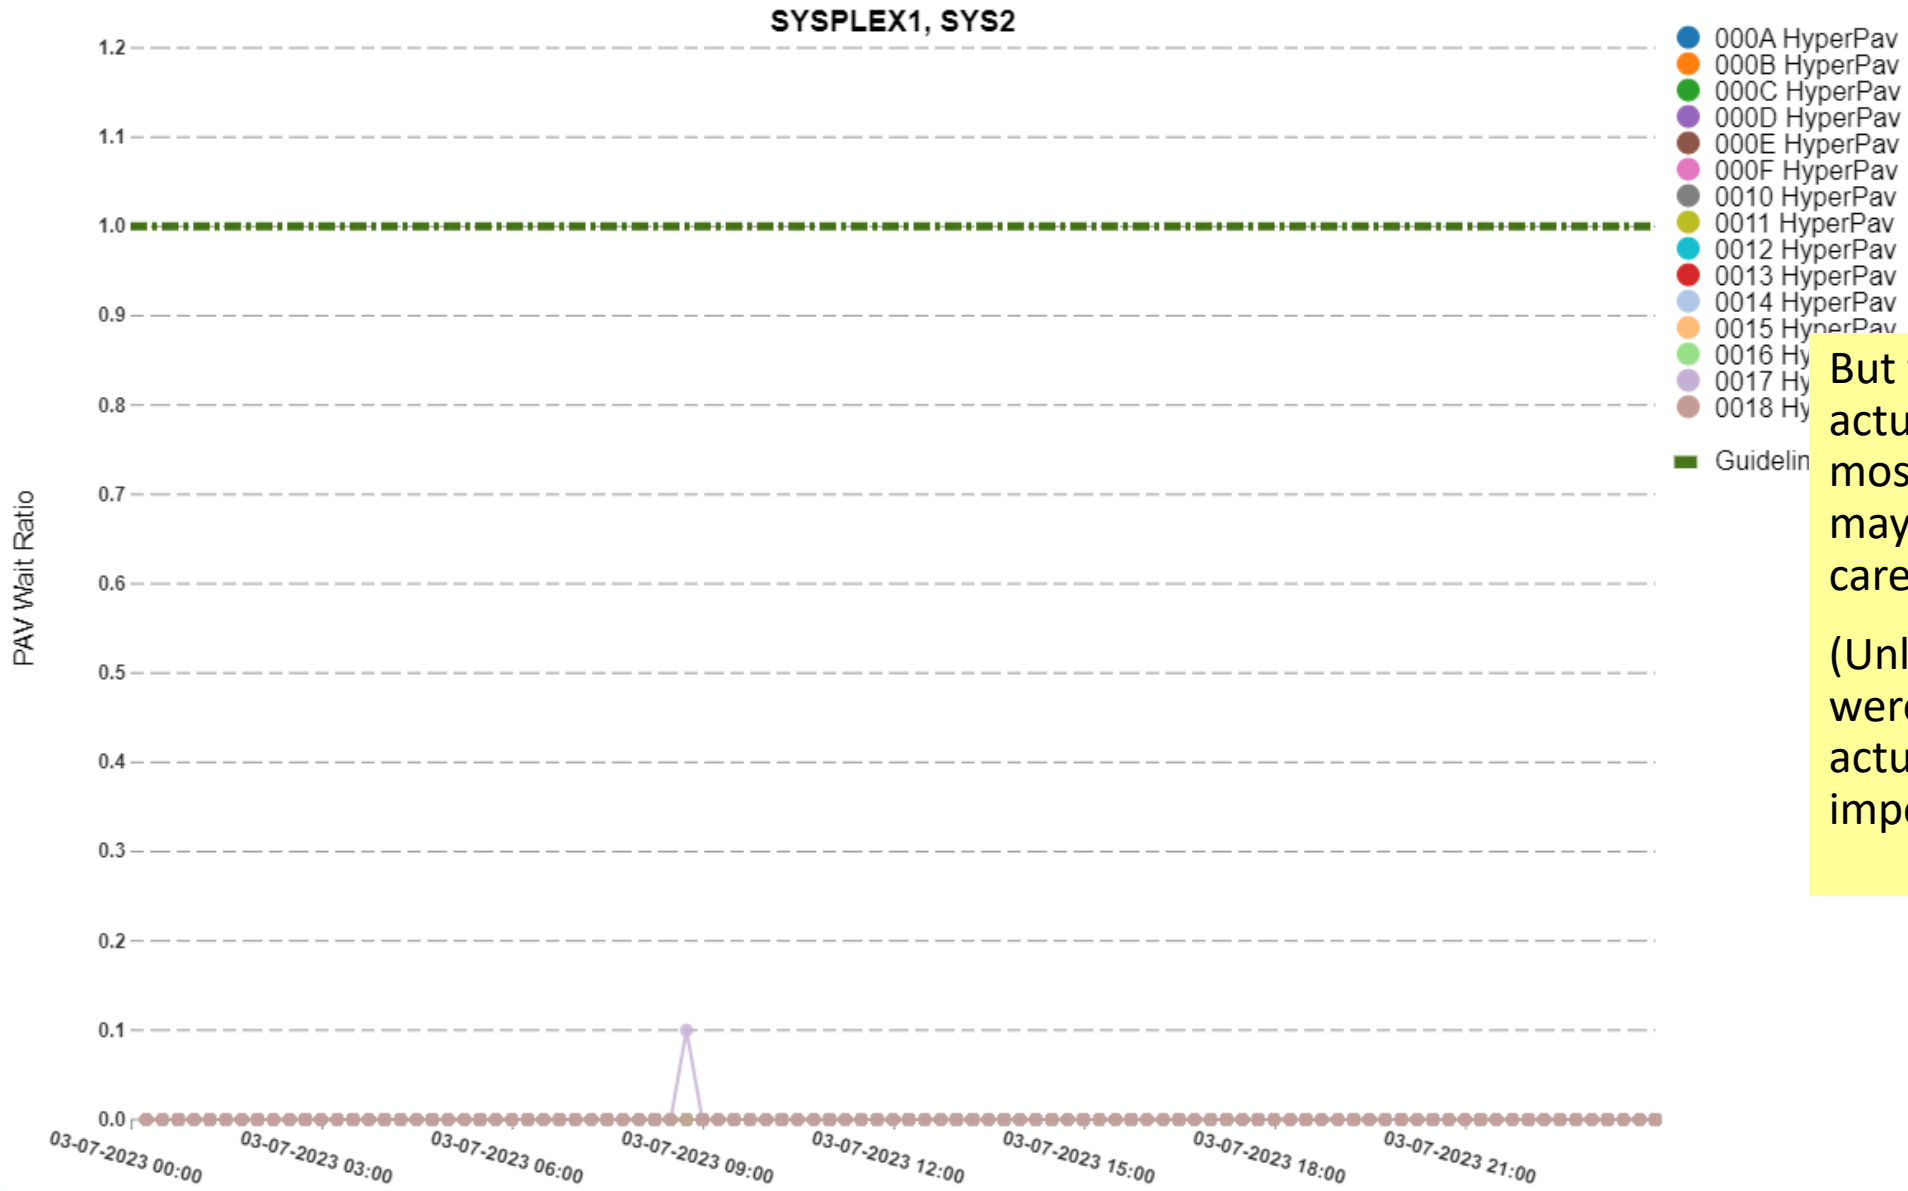

But the wait ratio is actually near zero for most of the intervals, so maybe we don't actually care so much.

(Unless those I/Os that were delayed were actually your very important I/Os!)

DASDplex RT Components

Including I/O rate

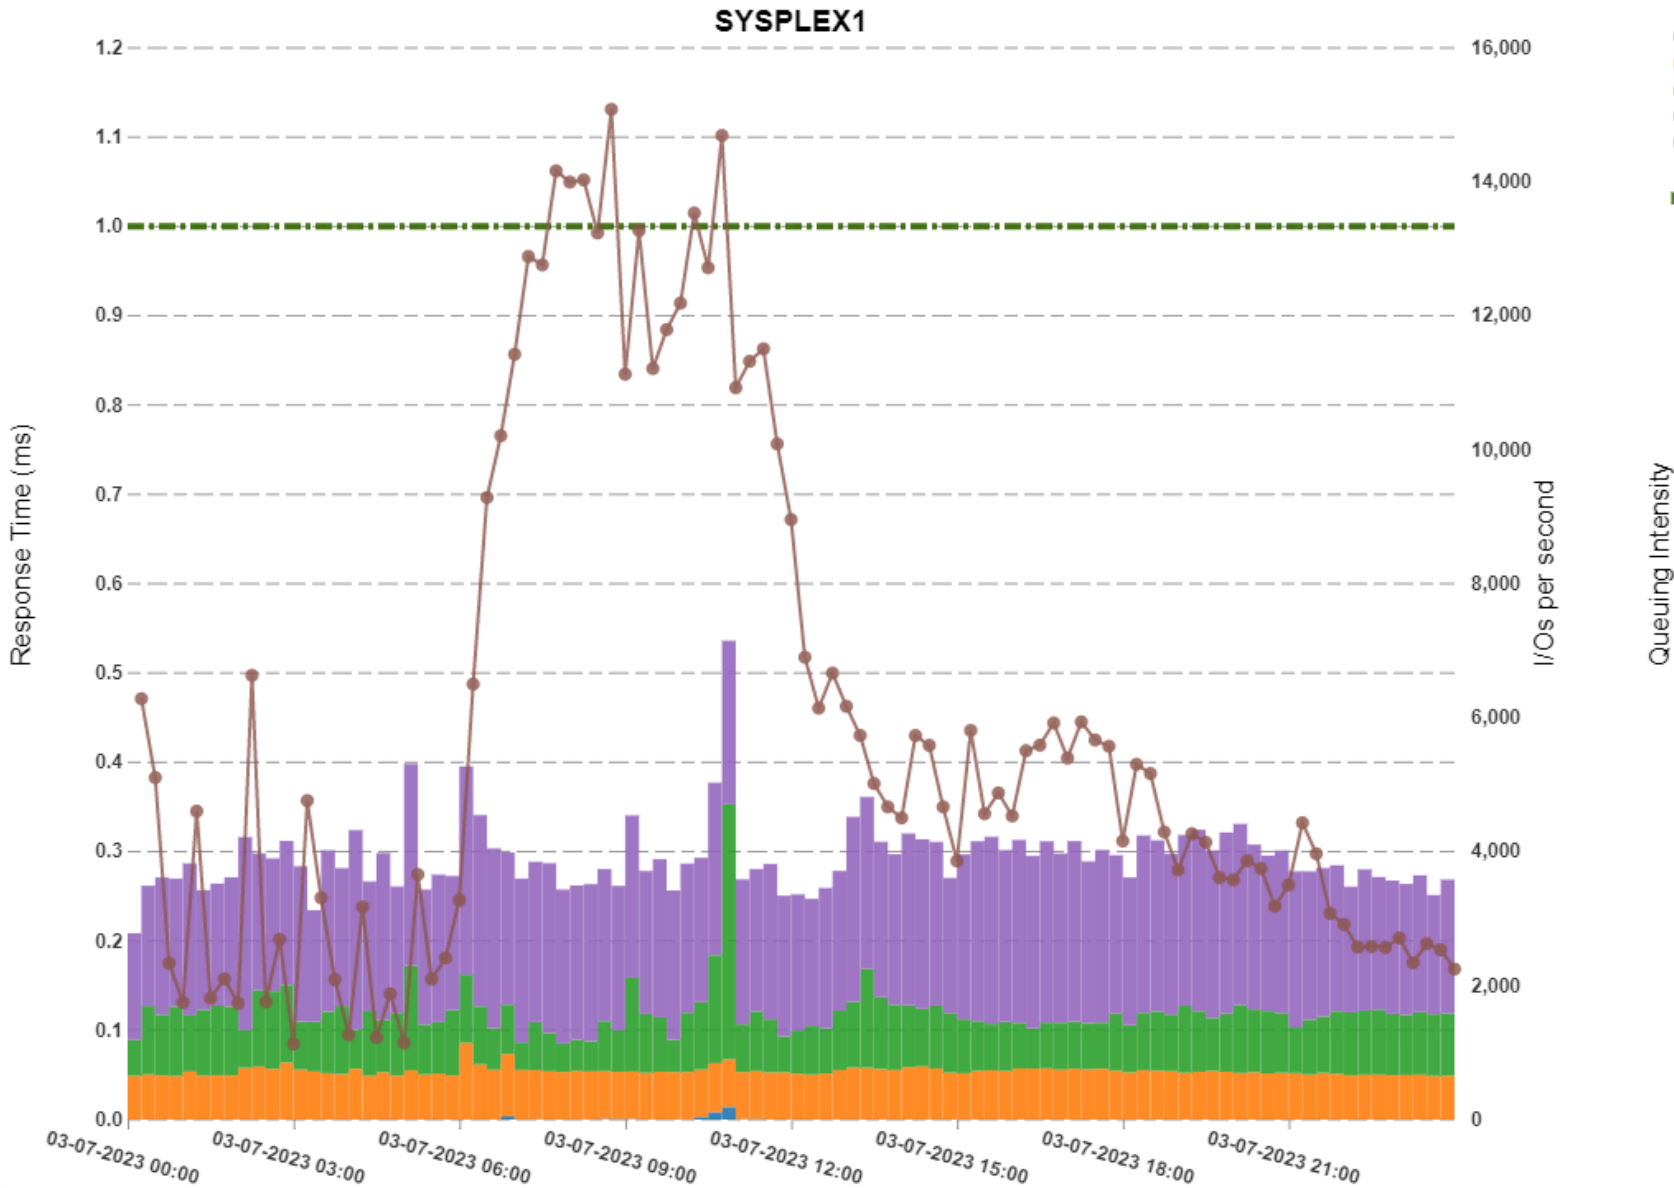

Note that we do see some IOSQ time in the overall DASD response time.

PAV - Top LCU High Water Mark of In-Use PAV-alias Devices

Top LCUs

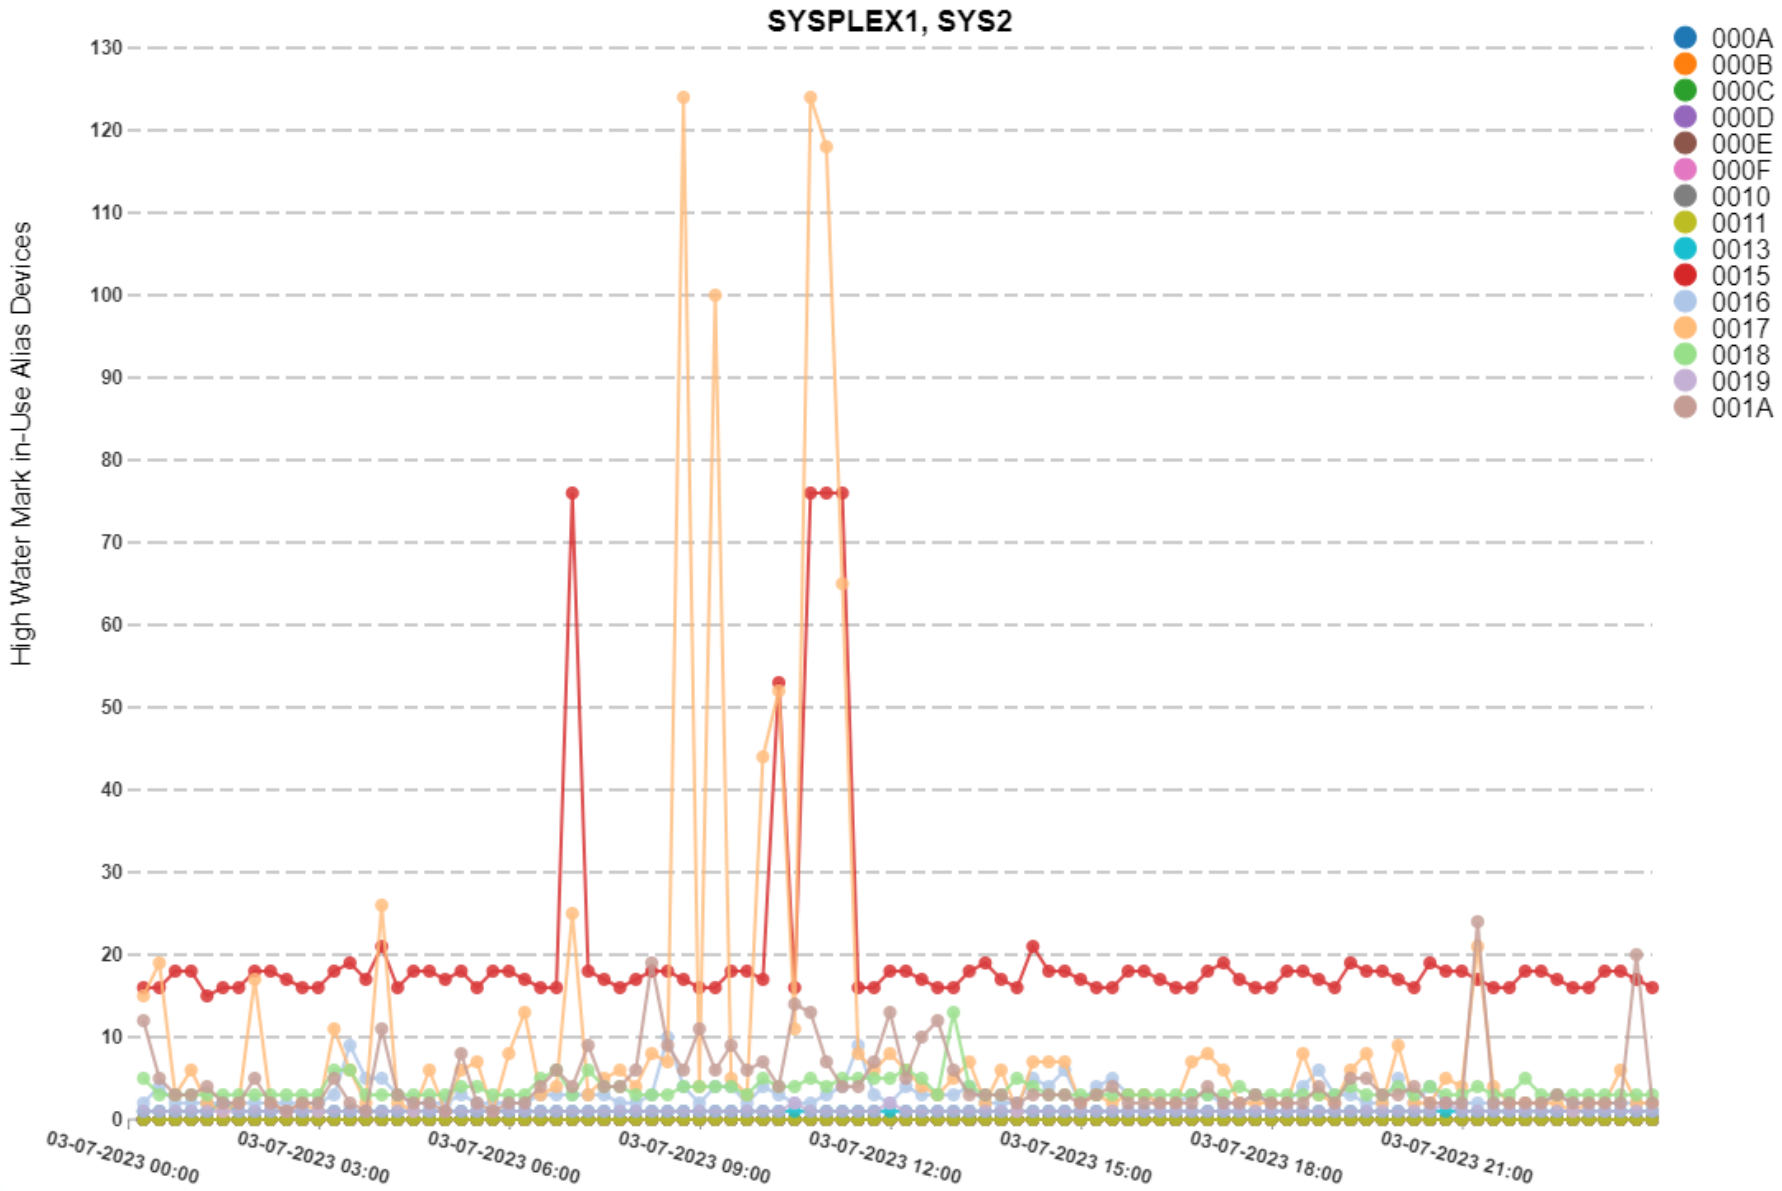

This shows the high water mark of the in-use PAV aliases by LCU. What's interesting here is are you hitting the max number of PAVs defined to one or more of the LCUs? (The plateau there implies it very well may be.)

PAV - Top LCU High Watermark of Concurrent In-Use aliases for One Base

Top LCUs - includes loaned alias devices

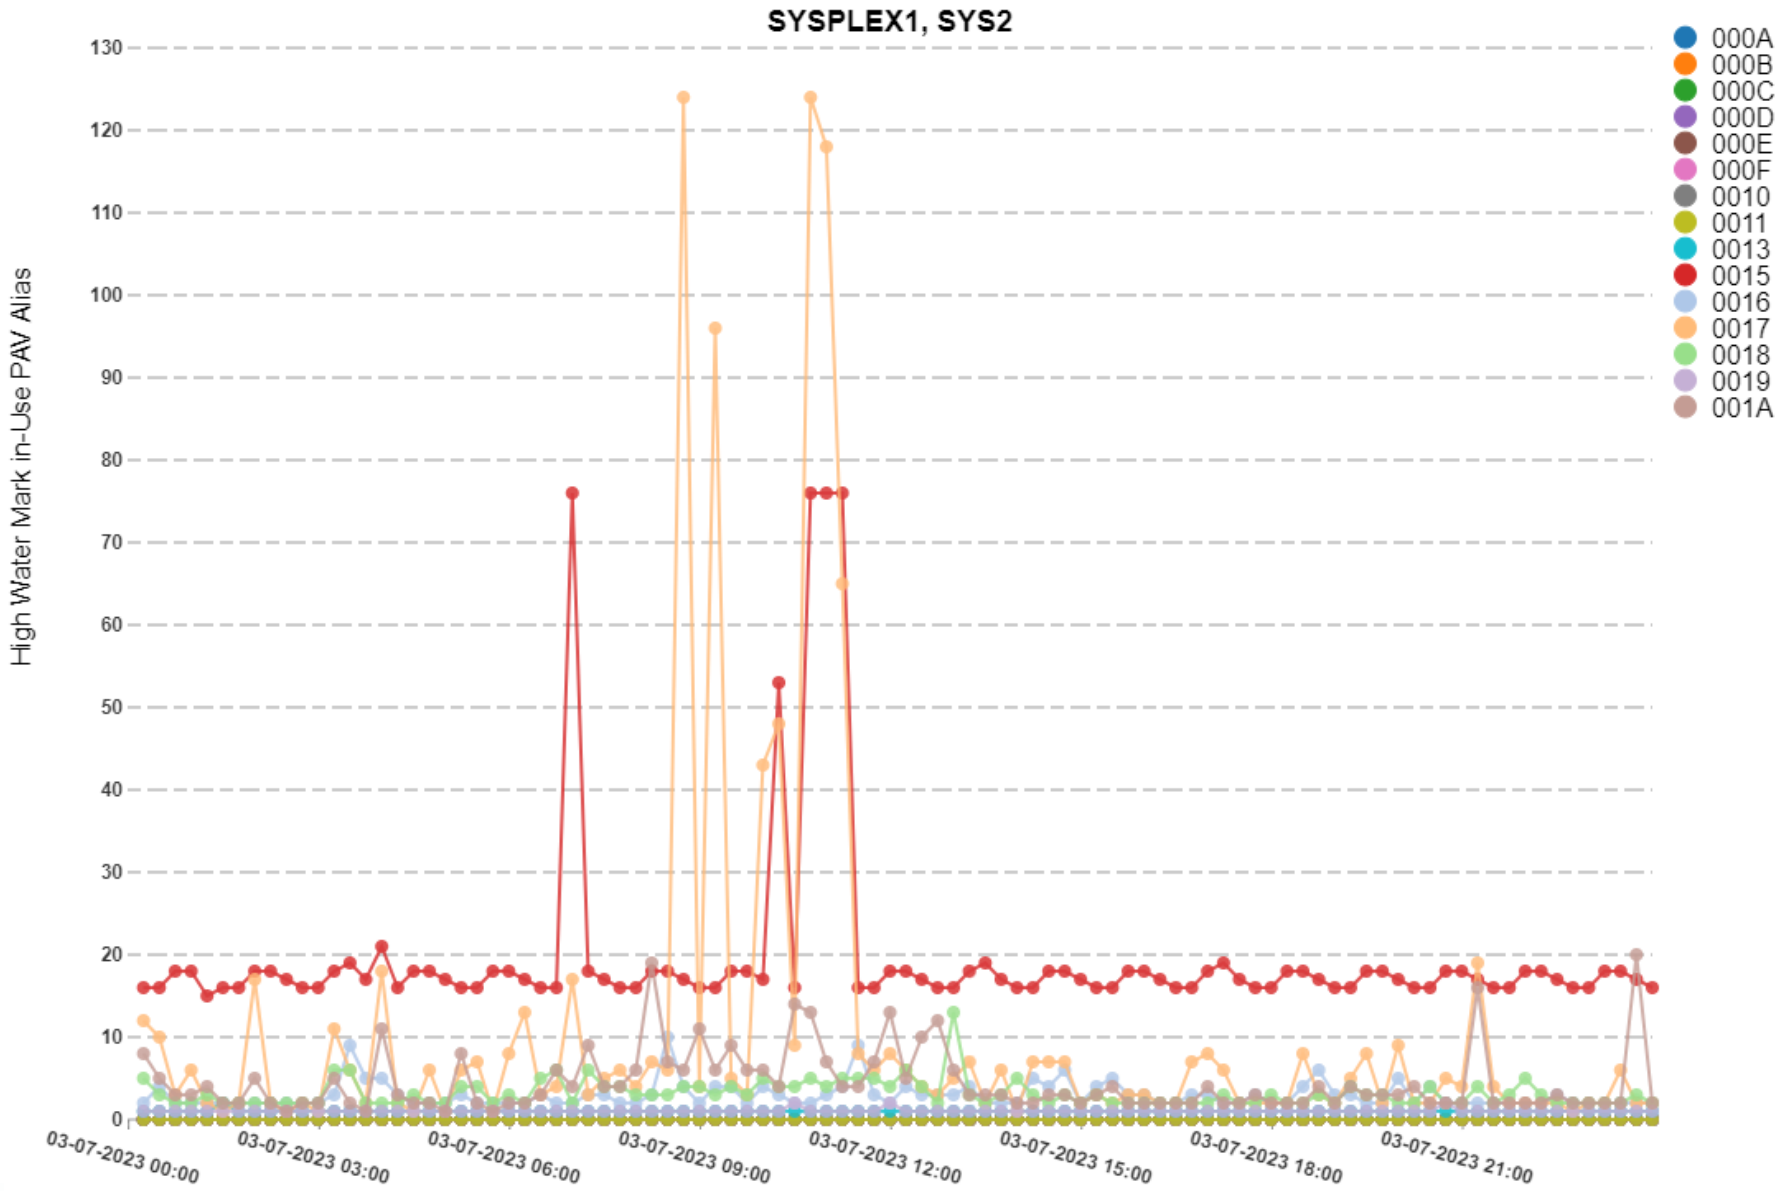

This view includes “loaned” aliases from other LCUs. If you have SuperPAV enabled, these numbers will quite possibly/probably be larger than the numbers on the previous report. That would be goodness.

PAV - Top LCU High Water Mark of Queued I/O Requests

Top LCUs

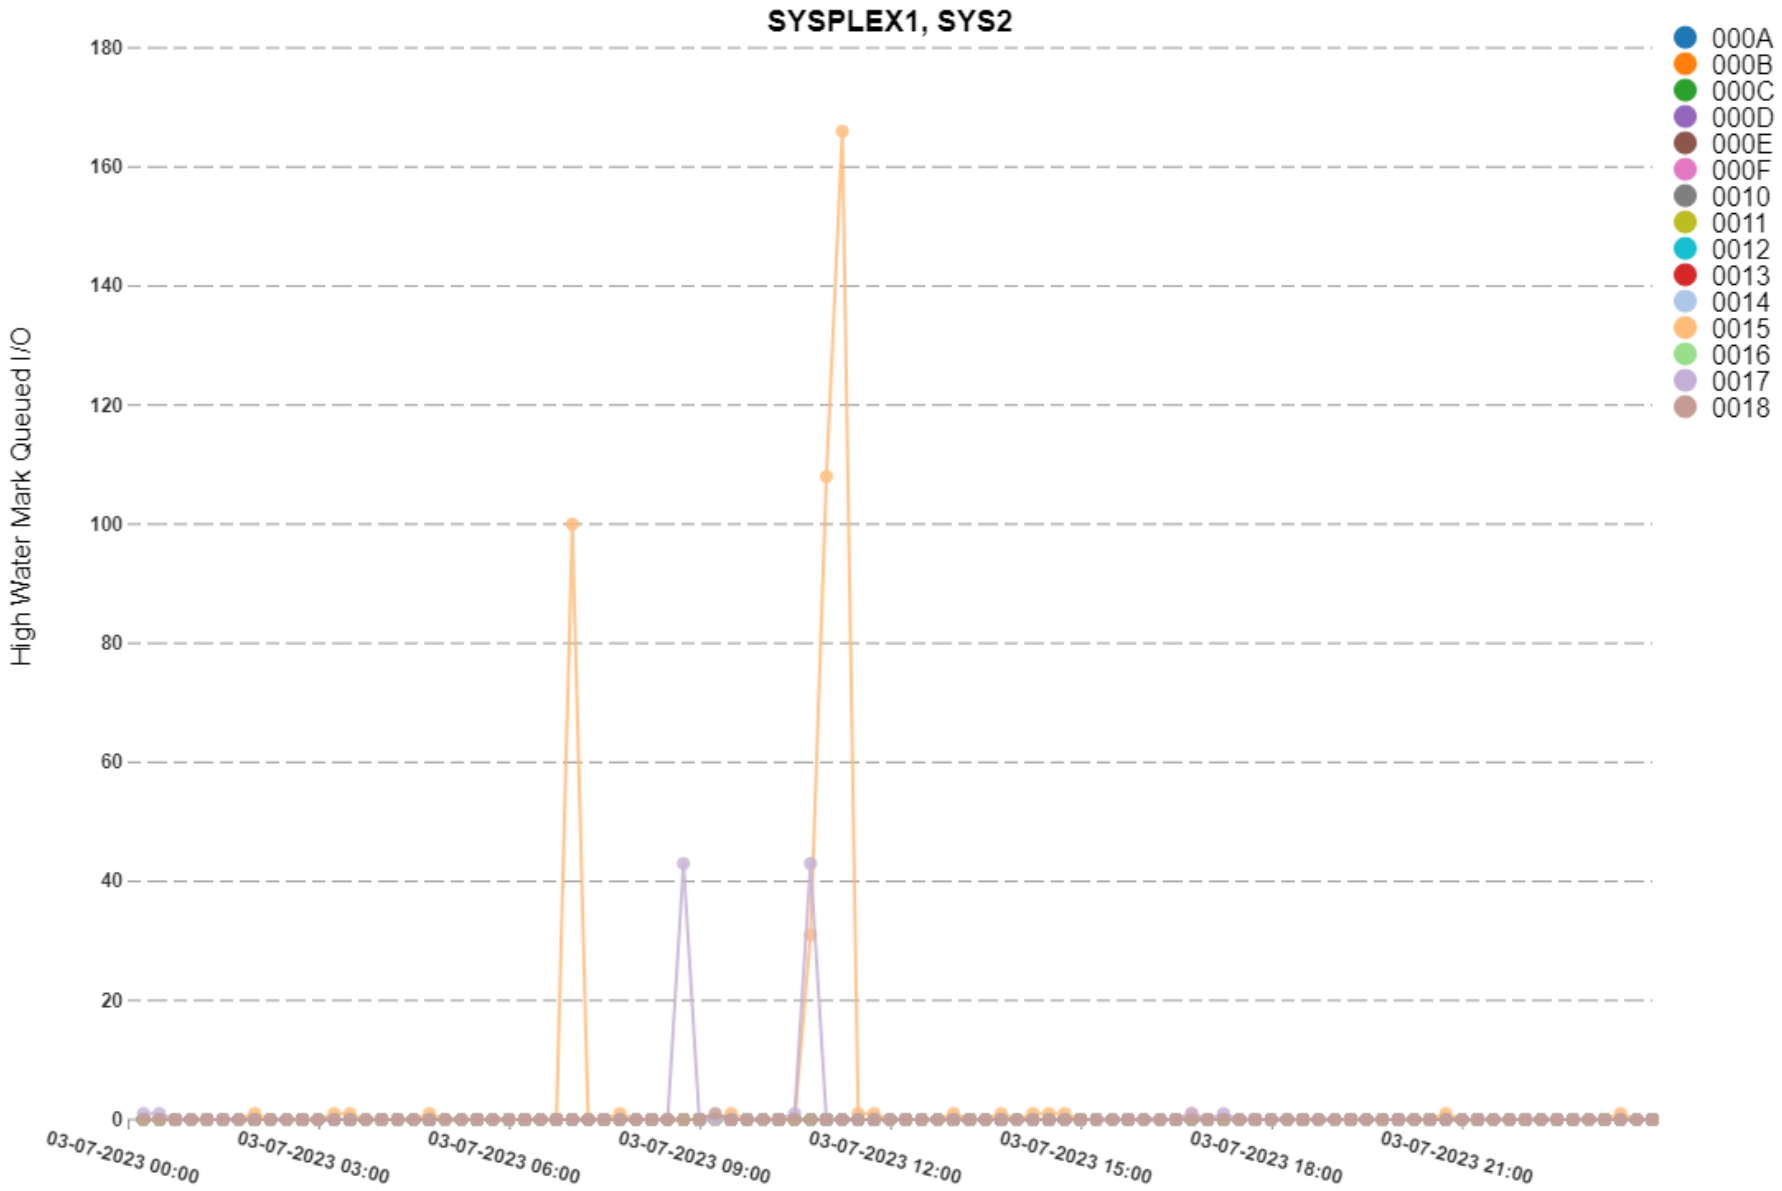

The peak number of queued I/O requests is helpful for getting a sense for your peak wait times and how many more PAVs you may need.

PAV - Top LCU Request Rate

Top LCUs

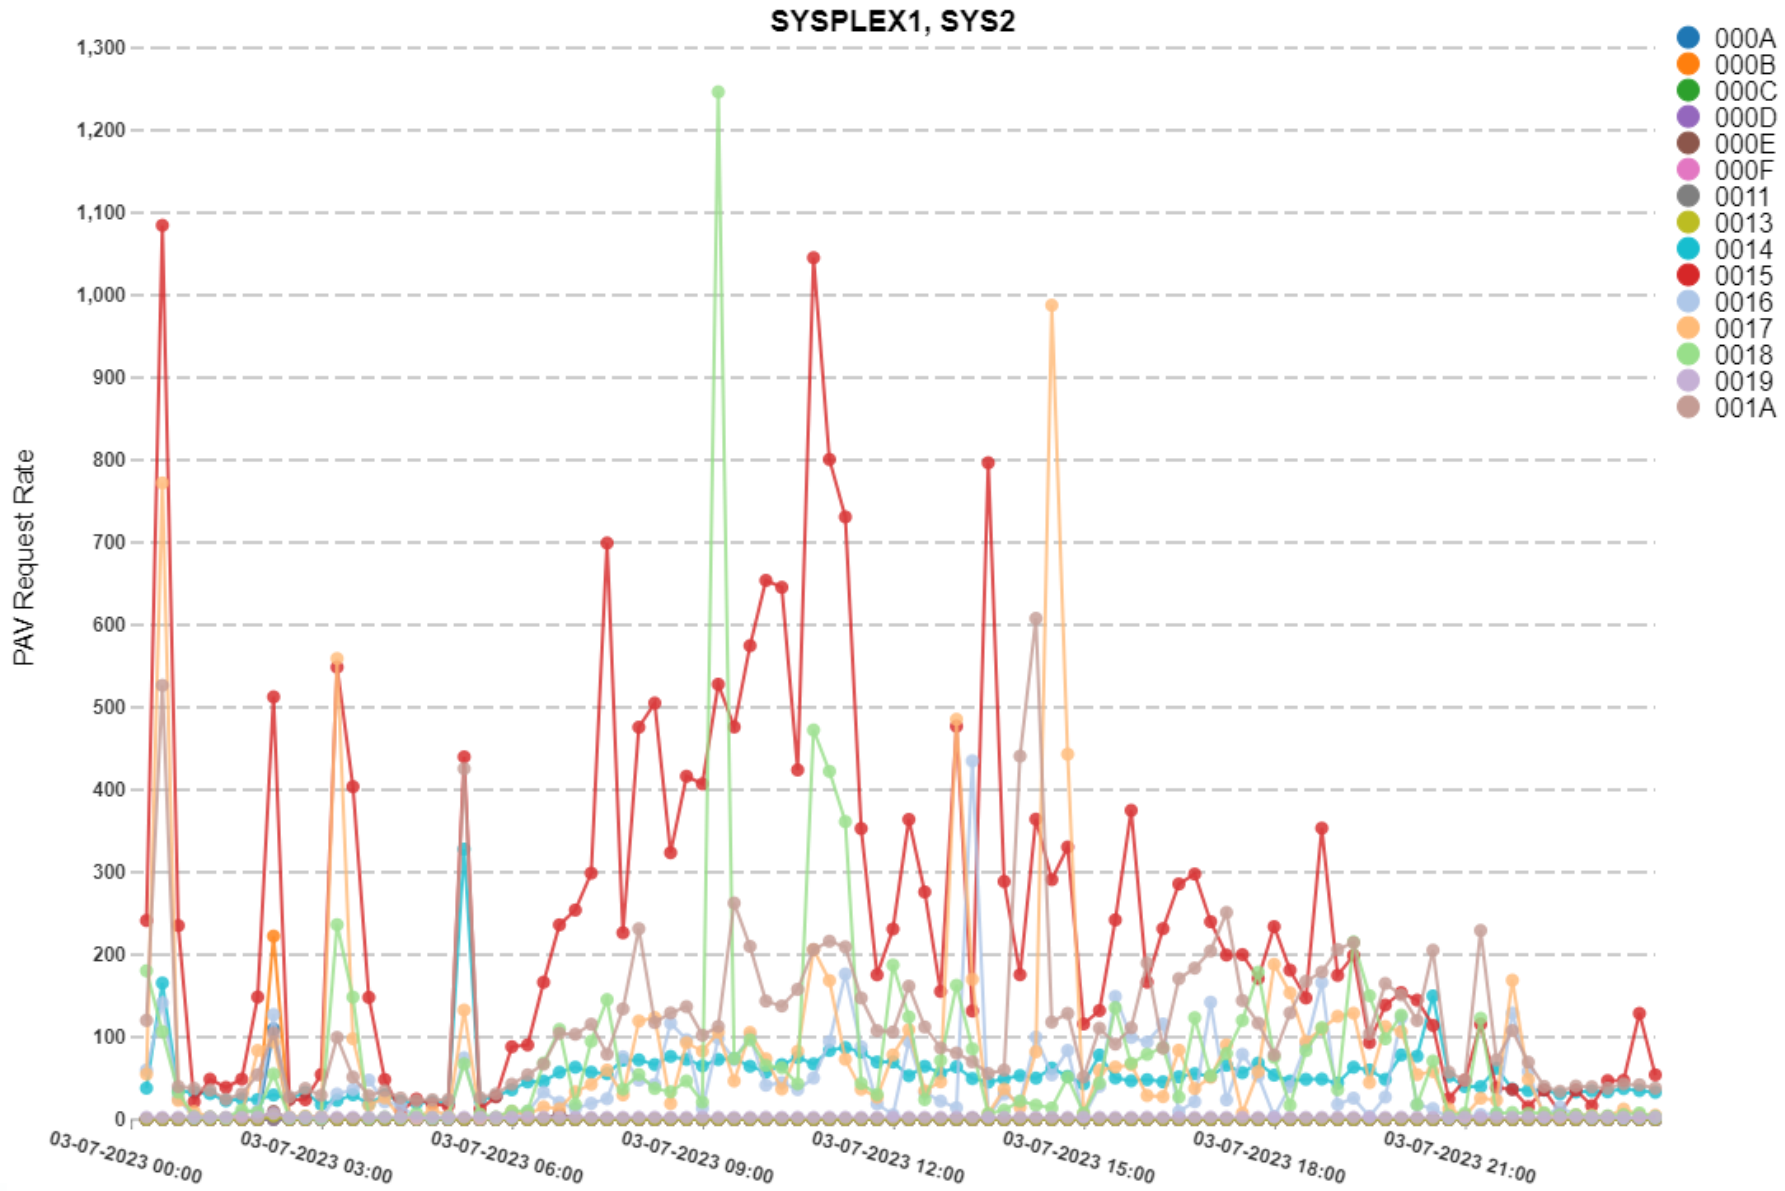

This is the rate of I/Os being done by the top LCUs. Note that there's not necessarily a correlation between I/O rate and PAV waits. (Although in this case LCU 0015 is one of the busiest LCUs and sees the most PAV waits too.)

What does a more optimal situation look like?

PAV - I/O Not Started Because No PAV-aliases Were Available

Top LCUs

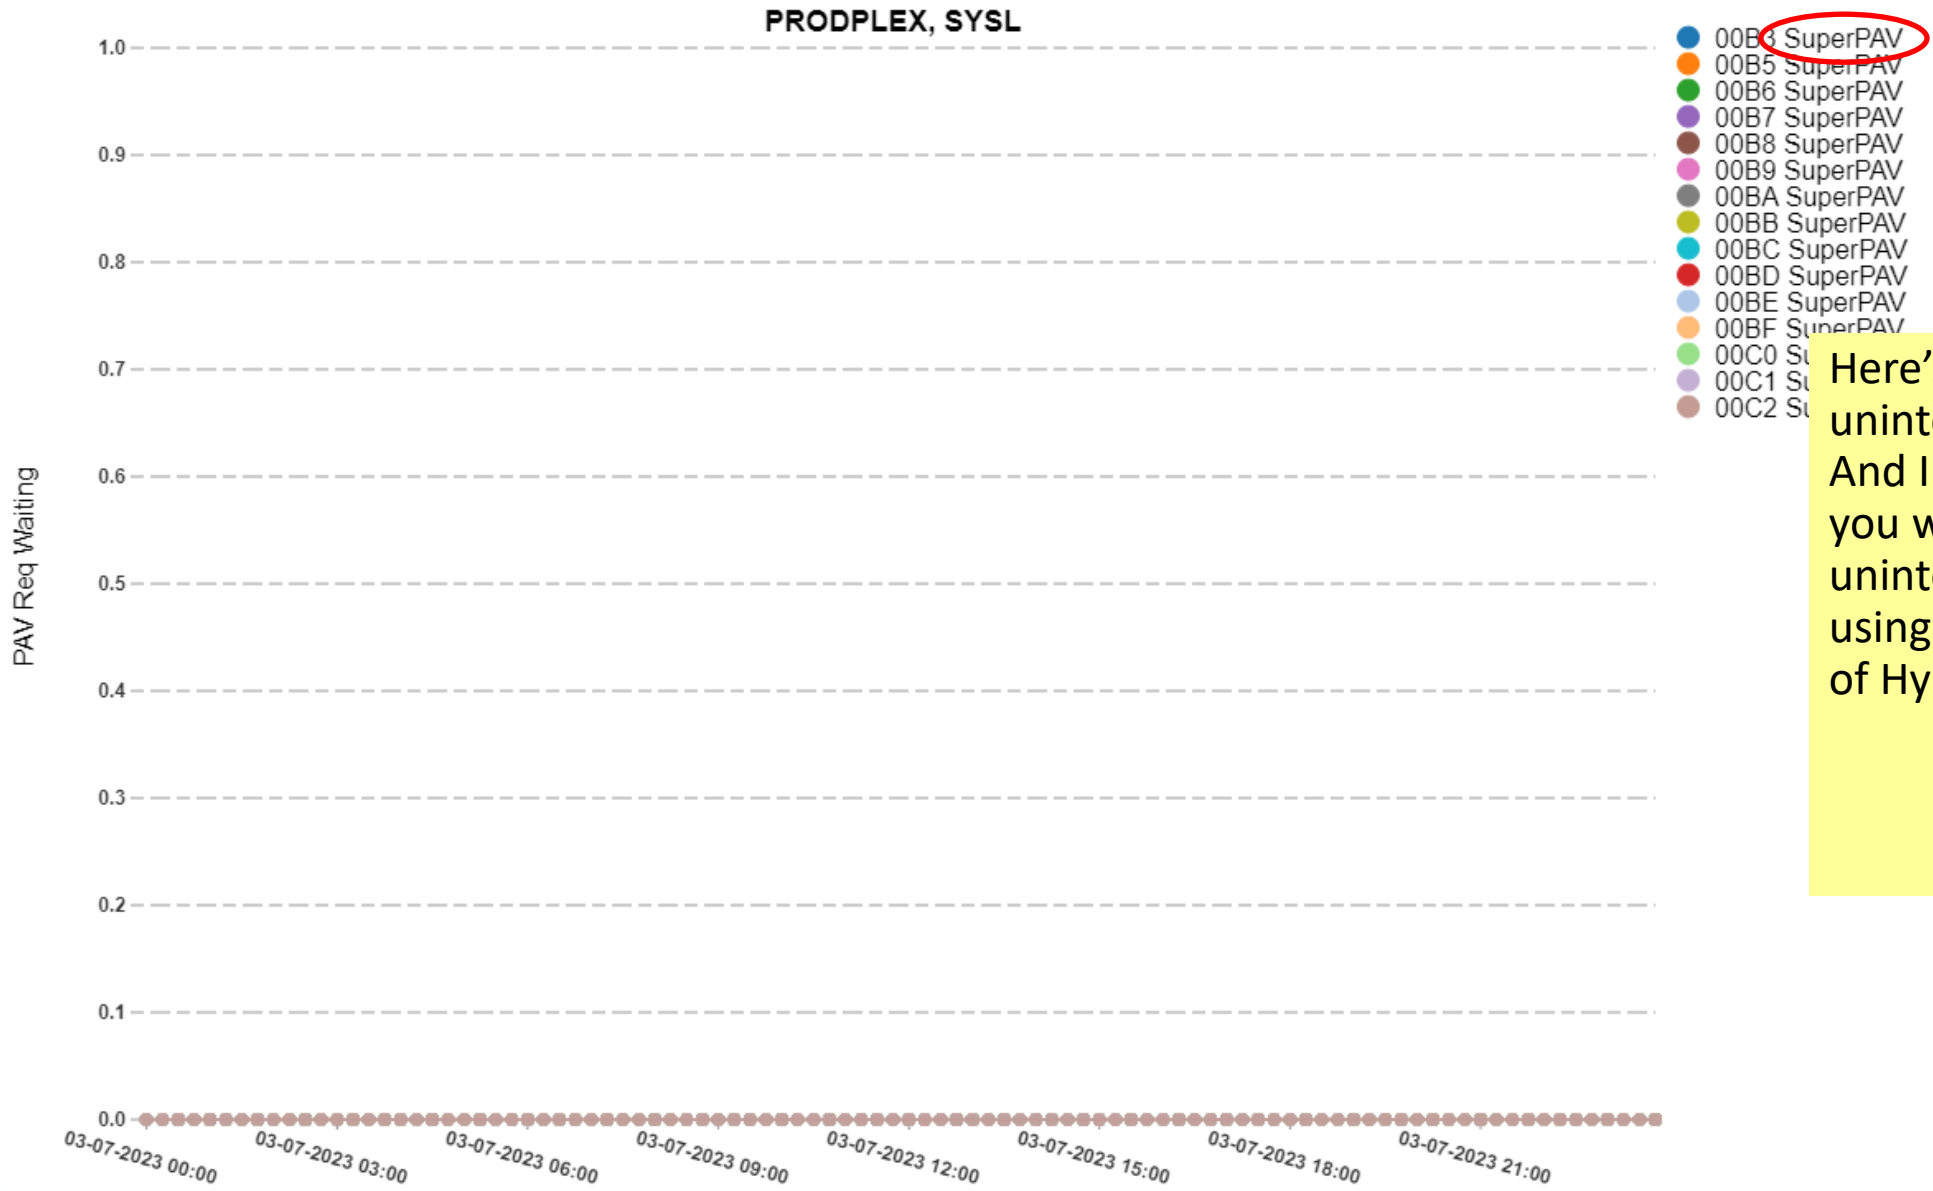

Here's one of those uninteresting systems. And I can already tell you why it's probably uninteresting: they're using SuperPAV instead of HyperPAV.

PAV - Wait Ratio for Top LCUs

(Delayed I/Os Due to No PAV-alias to Number of PAV Requests)

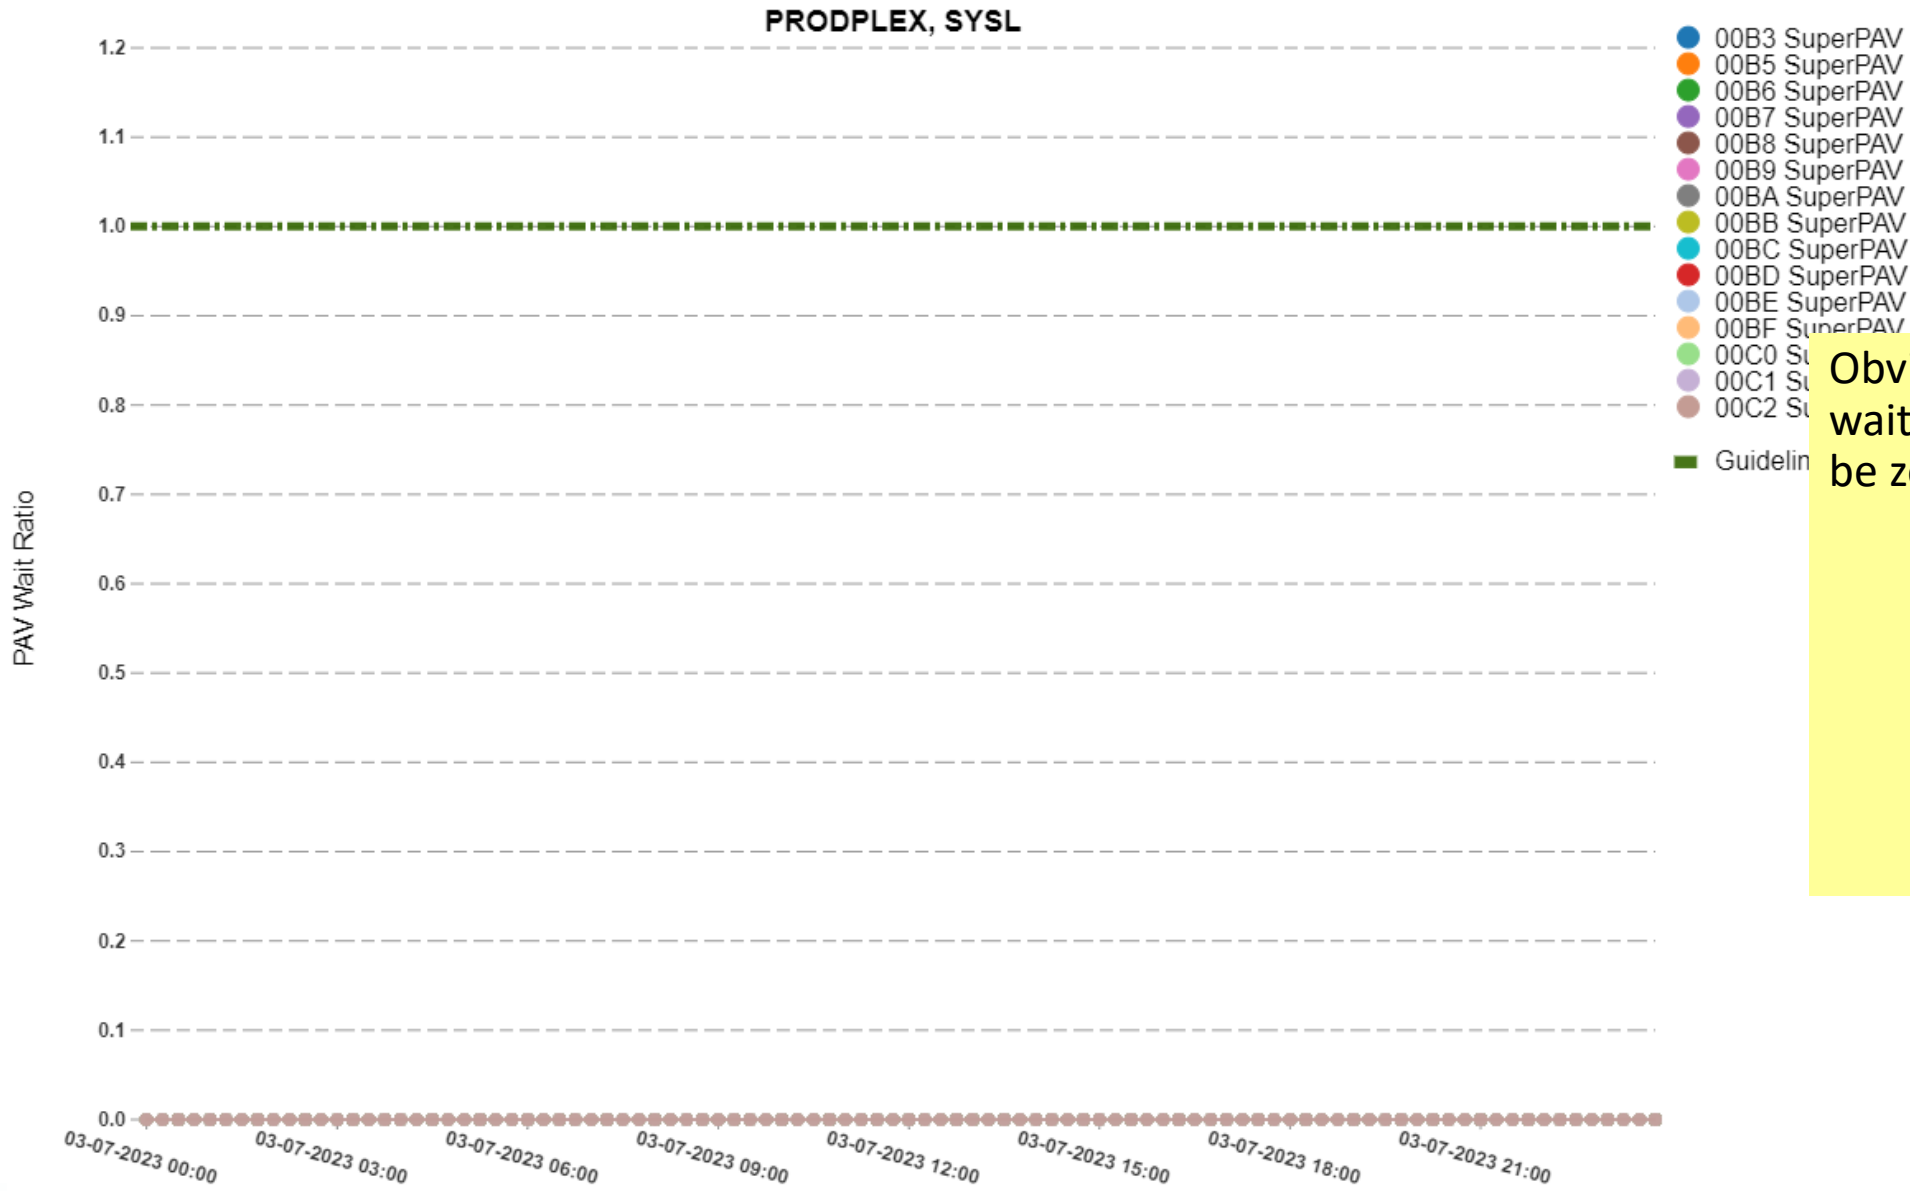

Obviously if there's no waits, the wait ratio will be zero too.

PAV - Top LCU High Water Mark of In-Use PAV-alias Devices

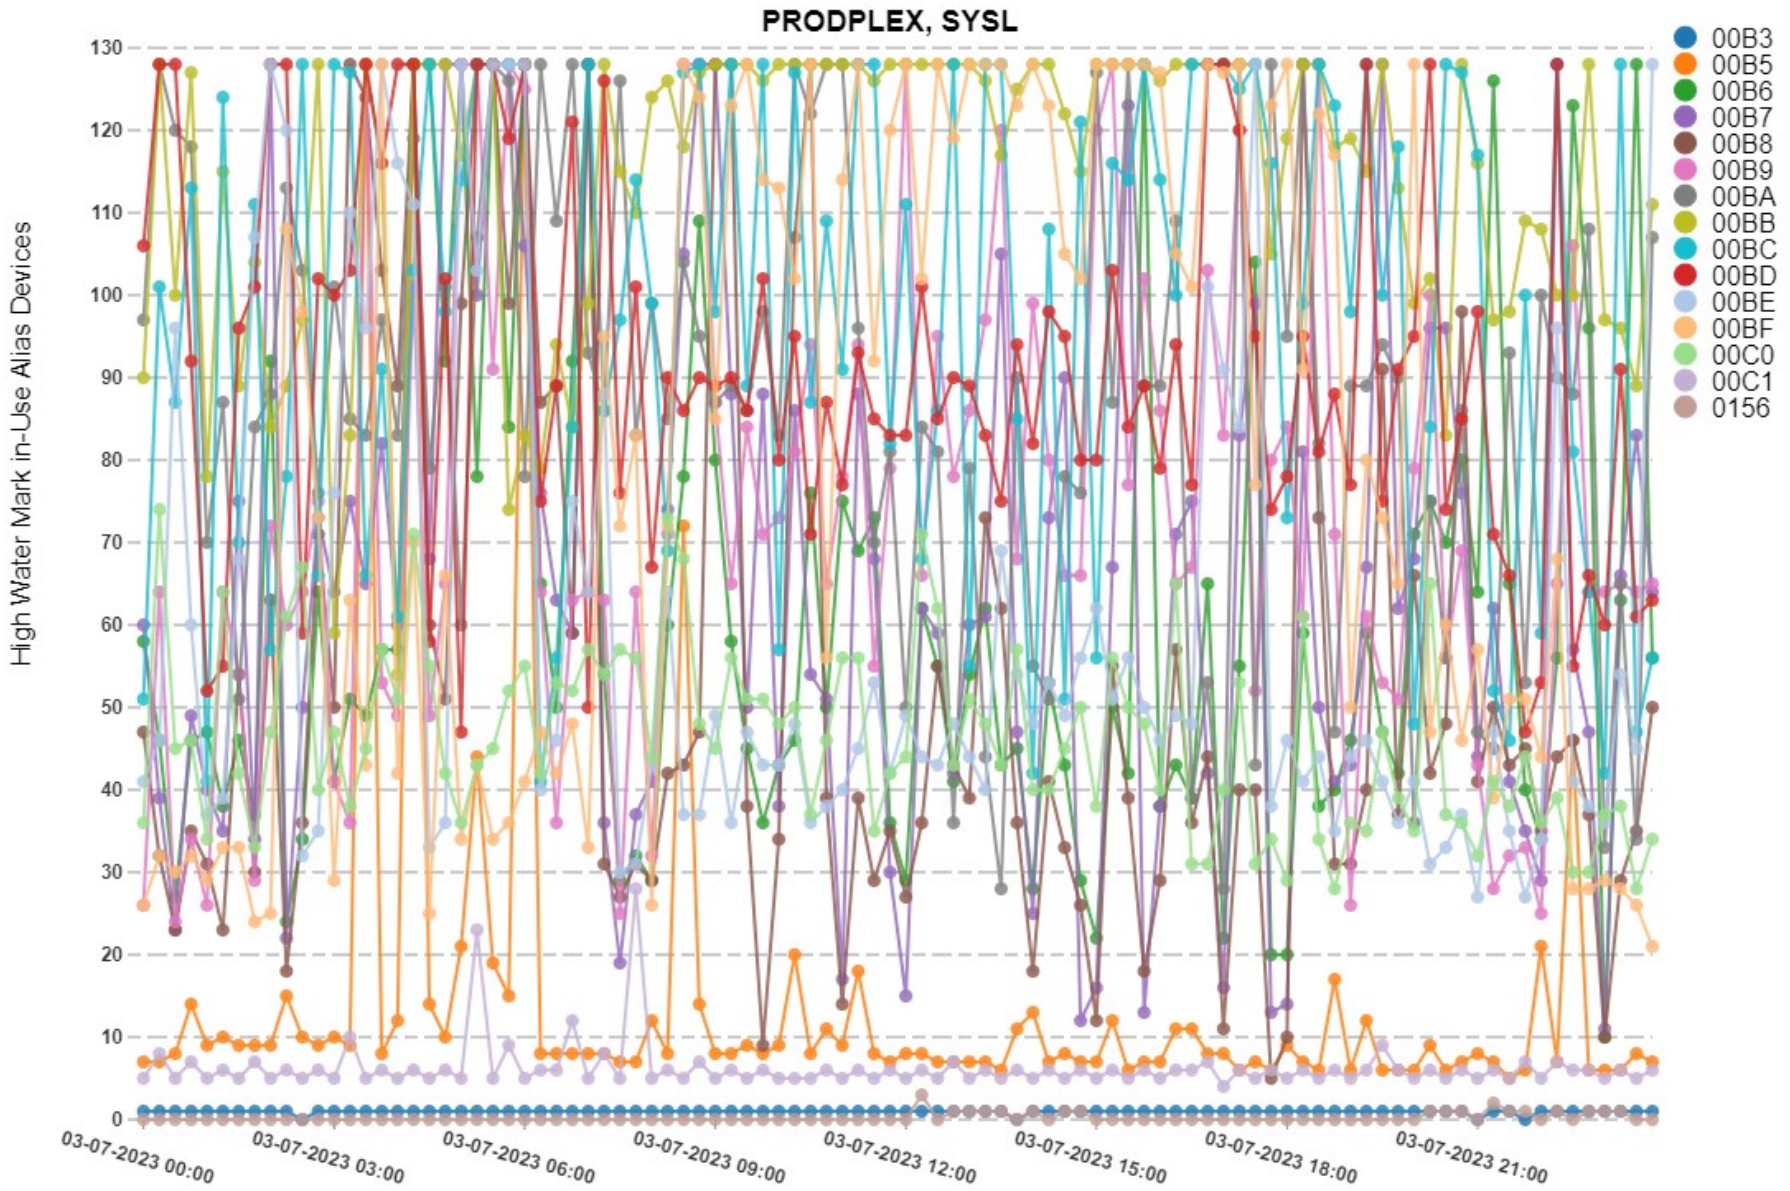

Here we see that it appears that they probably have 128 PAVs per LCU. At least they have that many for some LCUs.

PAV - Top LCU High Watermark of Concurrent In-Use aliases for One Base

Top LCUs - includes loaned alias devices

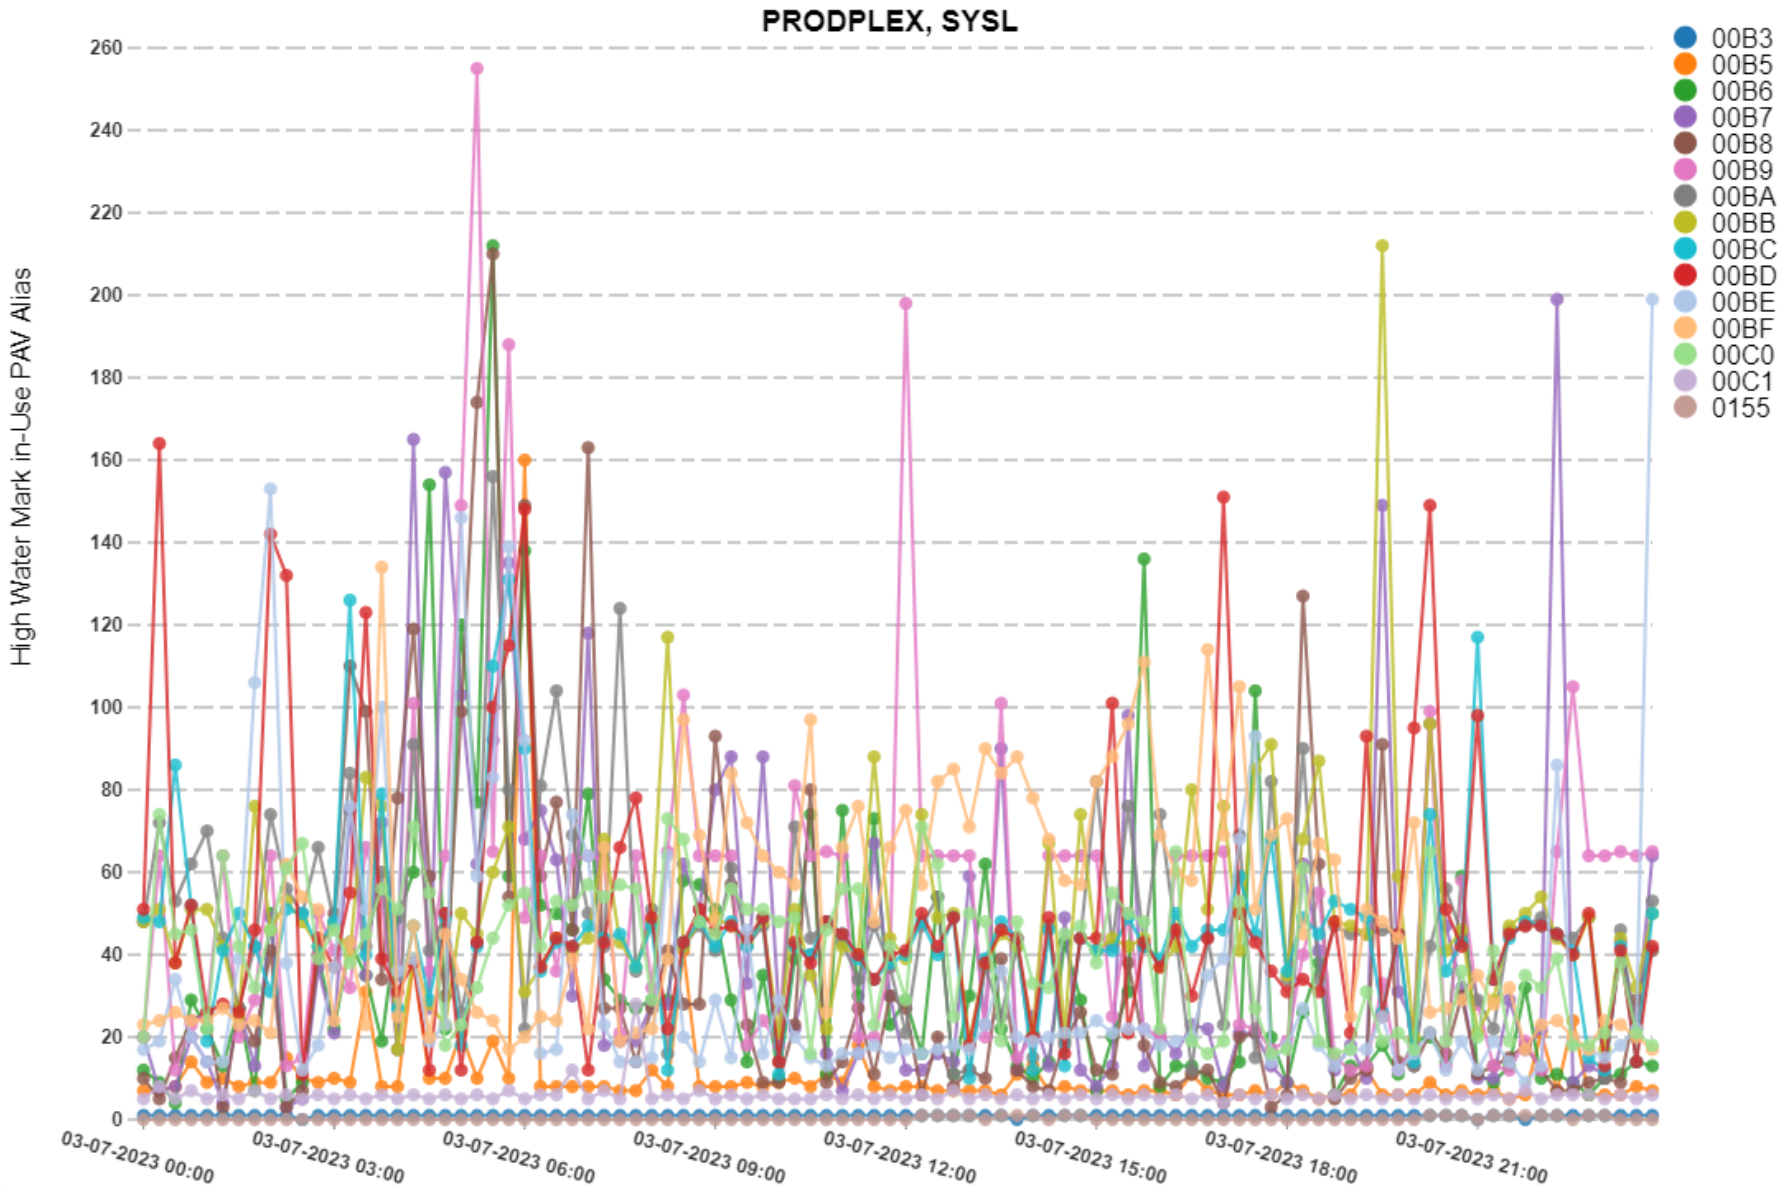

Here now we see the benefit of SuperPAV and we even see that one LCU had a volume that hit the max of 255 aliases for a single base address at one point in time. So despite having 128 aliases/LCU, they probably need them. I'd keep half an eye on this to make sure it doesn't become a problem.

PAV - Top LCU High Water Mark of Queued I/O Requests

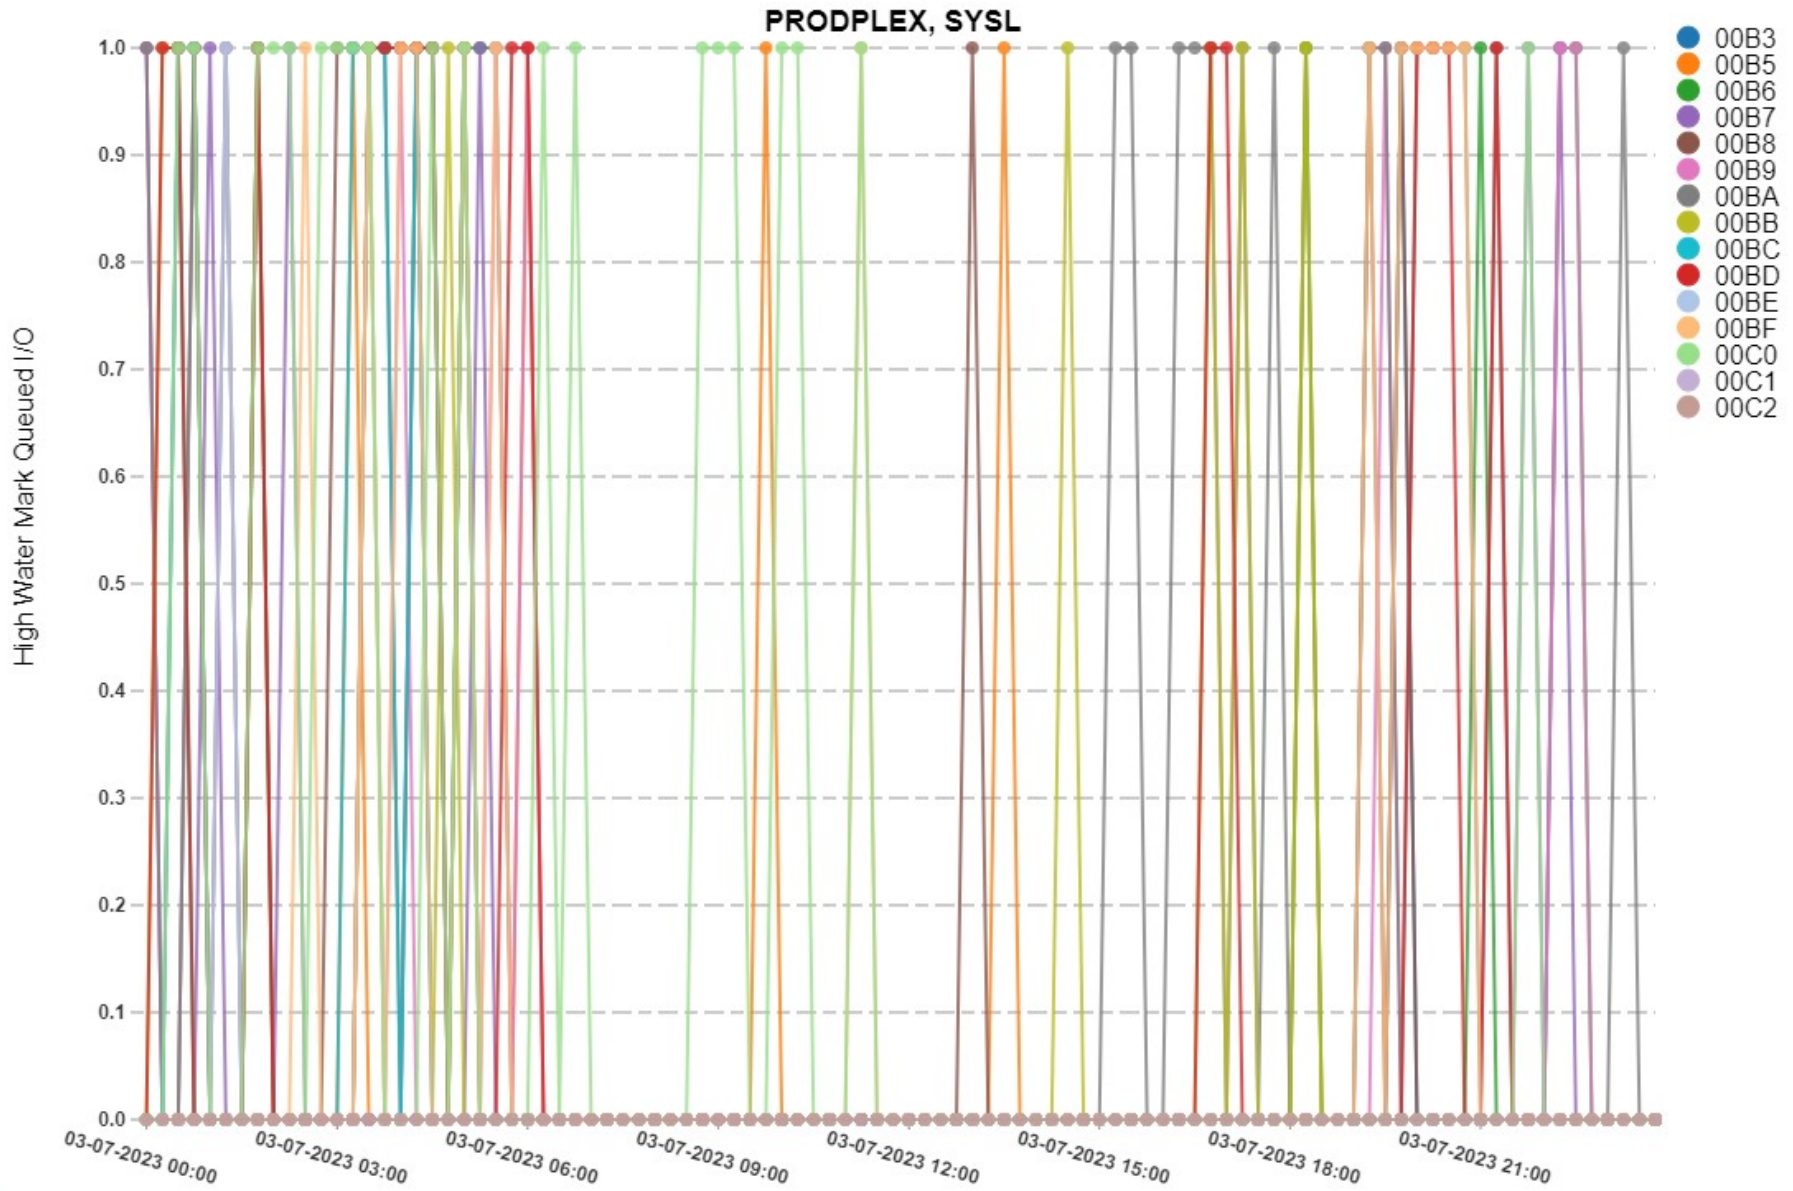

I'm not sure why with SuperPAV we still often see this HWM showing 1, but it's fairly common. I wonder if there may be an OBOB somewhere in the RMF code that captures this.

PAV - Top LCU Request Rate

Top LCUs

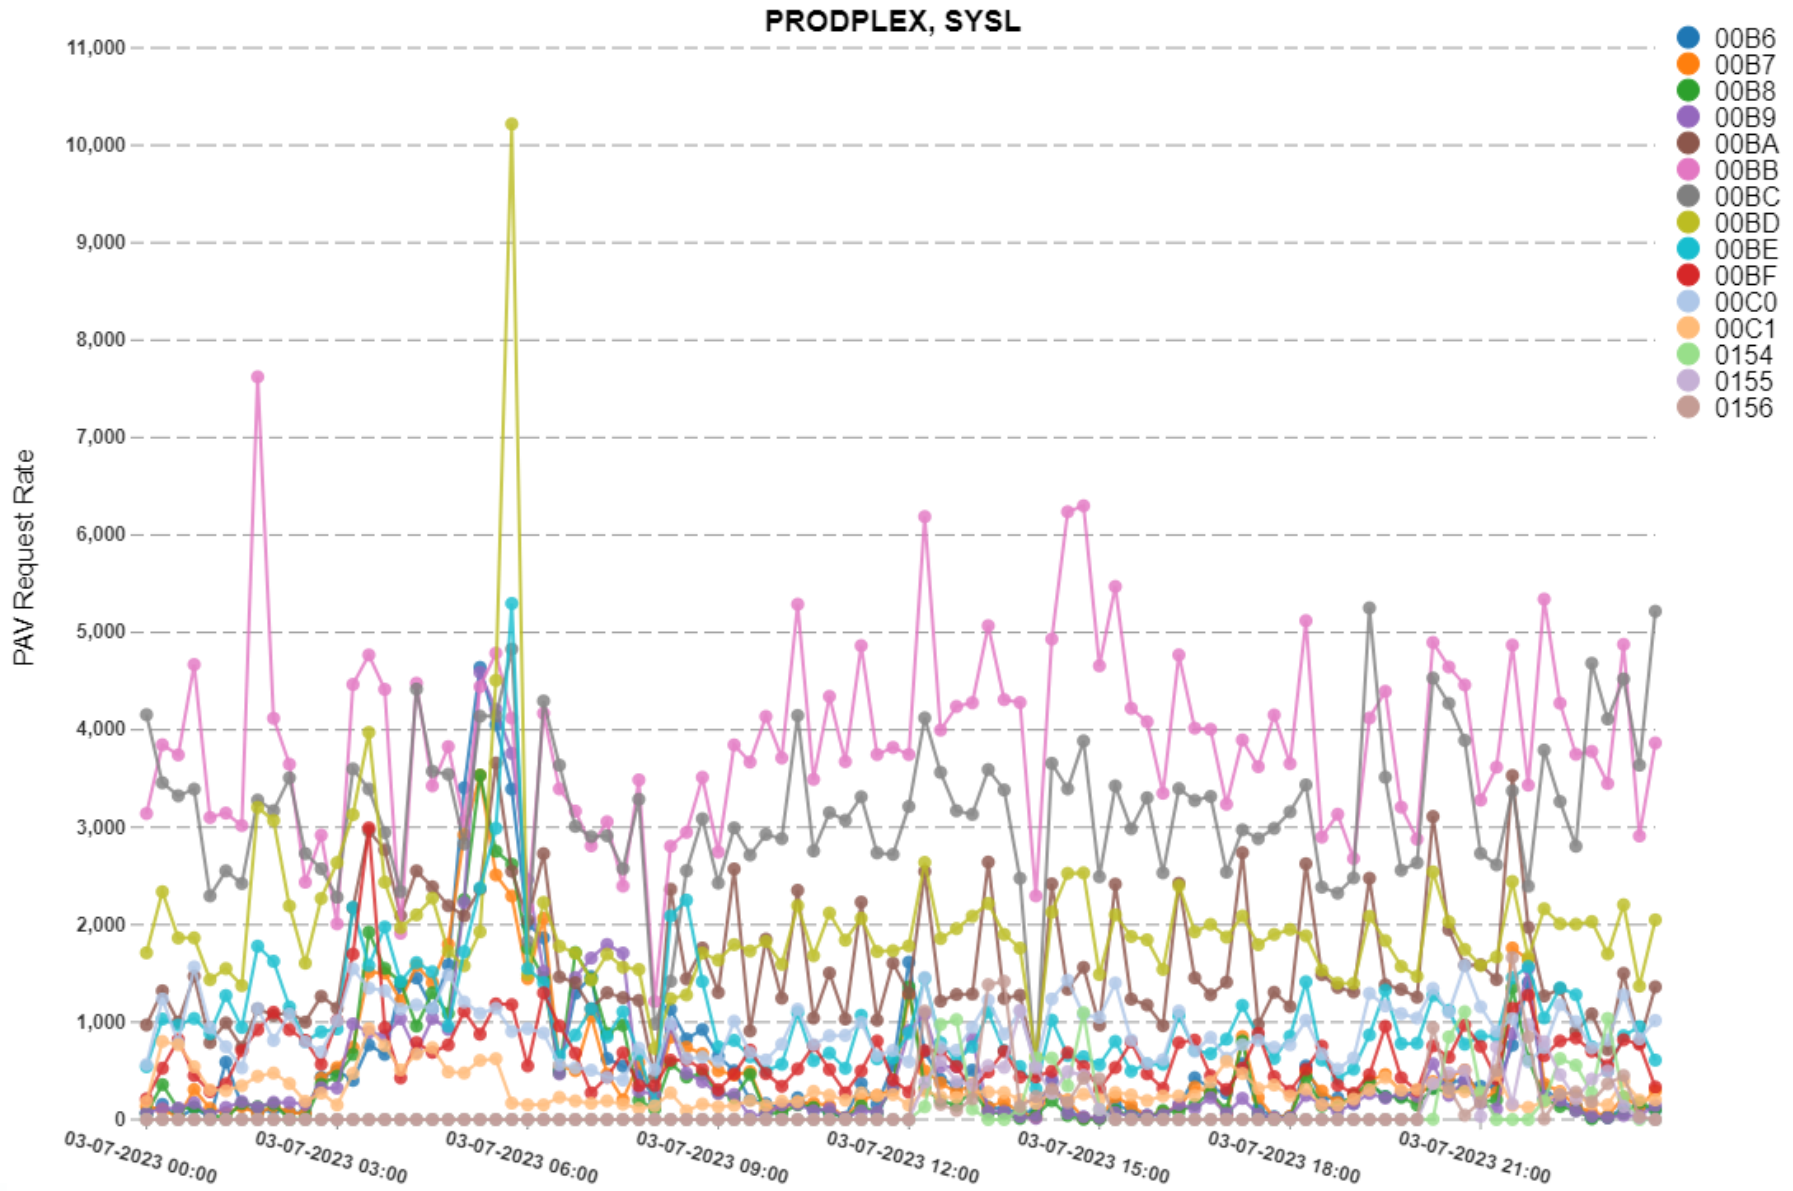

Here we can see that this system is doing a whole lot more I/O than the previous system and still has zero IOSQ time thanks to SuperPAV.

What to do with an “opportunity”

Perfection isn’t necessary, but what if you’re well short of it?

So you still have IOSQ time…

- First, check HYPERPAV=XPAV in IOSQ
  - Very common for people to have missed changing from HYPERPAV=YES
  - If you've replaced your DASD in the last several years it almost certainly supports it
    - But as always, if you're not sure check with your vendor…
- If you still have a problem, you may have an overly busy subset of LCUs
  - Increasing the PAVs per LCU during the next opportunity may be a good idea
  - May want to consider rebalancing I/O across the LCU groups
- Or… avoid doing the I/O altogether!
  - Can you use memory to avoid I/O
  - Avoiding I/O will also improve performance and may even reduce CPU consumption

Summary

- PAVs are good
- SuperPAV is best!
  - Use them if you have them, and you almost certainly do
- If you have I/Os delayed for PAVs look for how much delay you're suffering
  - Some tiny amount of delayed I/Os may be ok
  - But it's hard to know whether the delayed I/Os are important or not!
  - In ***most*** cases, SuperPAVs can eliminate need to for rebalancing I/O across LCUs
    - But does depend a lot on the I/O patterns

Questions??