



# Understanding System Recovery Boost's Impact on Performance and Performance Reporting

Scott Chapman

Enterprise Performance Strategies, Inc.

[Scott.chapman@EPStrategies.com](mailto:Scott.chapman@EPStrategies.com)



# Contact, Copyright, and Trademarks



## Questions?

Send email to [performance.questions@EPStrategies.com](mailto: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<sup>®</sup>, Reductions<sup>®</sup>, Pivotor<sup>®</sup>**

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries: IBM<sup>®</sup>, z/OS<sup>®</sup>, zSeries<sup>®</sup>, WebSphere<sup>®</sup>, CICS<sup>®</sup>, DB2<sup>®</sup>, S390<sup>®</sup>, WebSphere Application Server<sup>®</sup>, and many others.

Other trademarks and registered trademarks may exist in this presentation

# Abstract



This webinar will provide a brief overview of System Recovery Boost (SRB) and then discuss the broader implications on performance and performance reporting. As a reminder, during a speed boost, sub-capacity CPs will run faster, and a zIIP boost will provide additional capacity by allowing CP work to run on zIIP processors. SRB can improve the return-to-service time for various planned or unplanned “recovery” activities, but there are some caveats to be aware of, especially if you’re doing performance analysis. Turns out some of the data will lie to us during a boost period.

# 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

# 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](#))

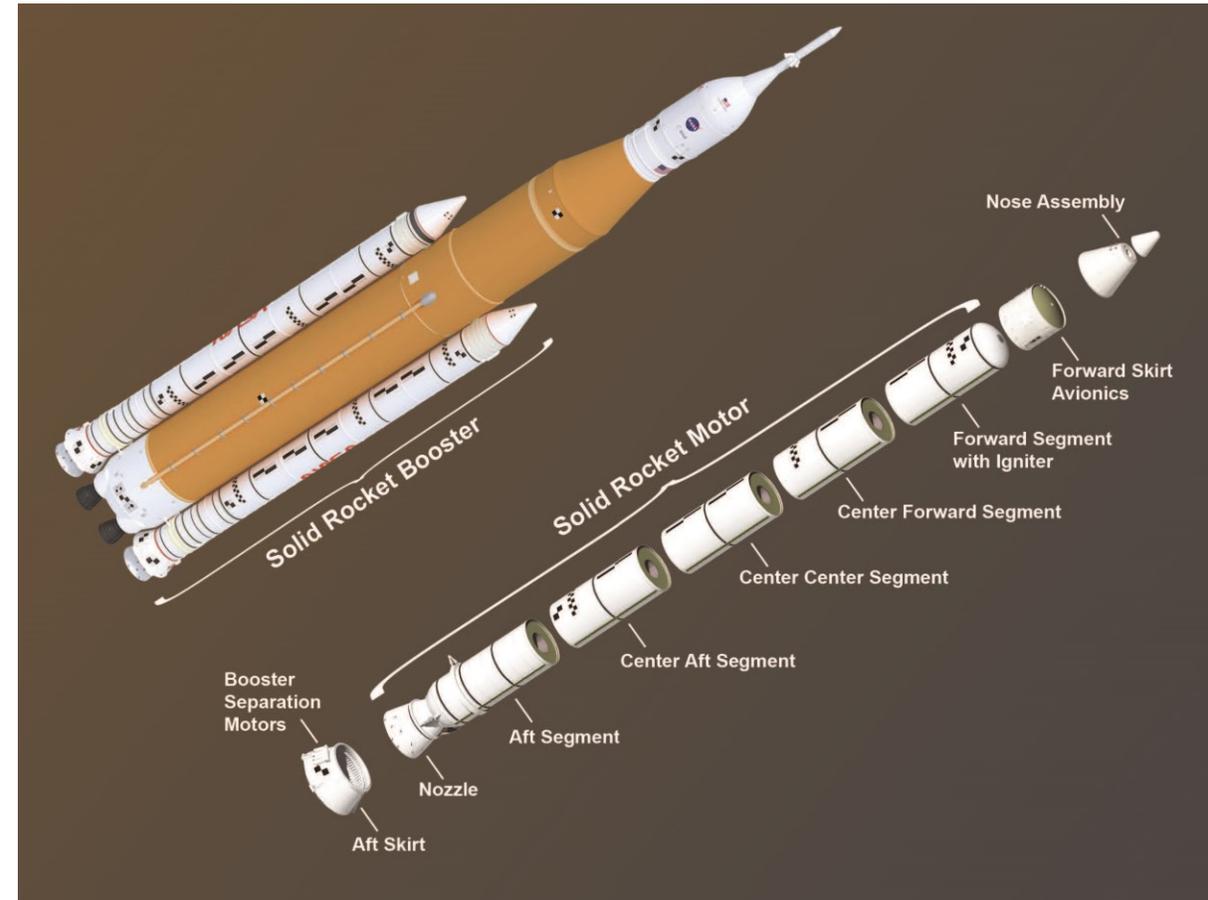
This playlist walks through several reports that will be useful in while conducting a WLM velocity goal an.

# Agenda = Questions to be answered



- What is SRB and what types of boosts are available?
- When will boosts be triggered?
- How is performance impacted?
- How is the SMF performance data impacted?
  
- Summary/Recommendations

What is SRB?  
What are the types of boosts?



# System Recover Boost



- Overall goal: provide additional capacity and performance to better recover from certain planned or unplanned events
  - Shutdown faster
  - Startup faster
  - Faster DR site switch
  - Faster sysplex recovery
  - Faster subsystem restart
- Requires z15 or z16 machine
  - + z/OS 2.3 and up on z15, z/OS 2.4 and up on z16
- No additional charge for basic System Recovery Boost
  - Optional extra cost item discussed in just a moment

# Types of boost



- Speed Boost: *GP engines run as full speed engines*
  - Applicable only to sub-capacity systems
  - I.E. this is of no use to 8561-7xx and 8562-Z0x systems
  - Only boosted LPARs run as full speed!
- zIIP Boost: *zIIPs will run work not eligible to run on zIIPs*
  - I.E. **any** workload may be dispatched to a zIIP
  - Reserved (but physically available) zIIPs brought online to the boosted LPAR
  - Note that zIIPs always run full speed
  - What if you want more zIIPs for zIIP boost?

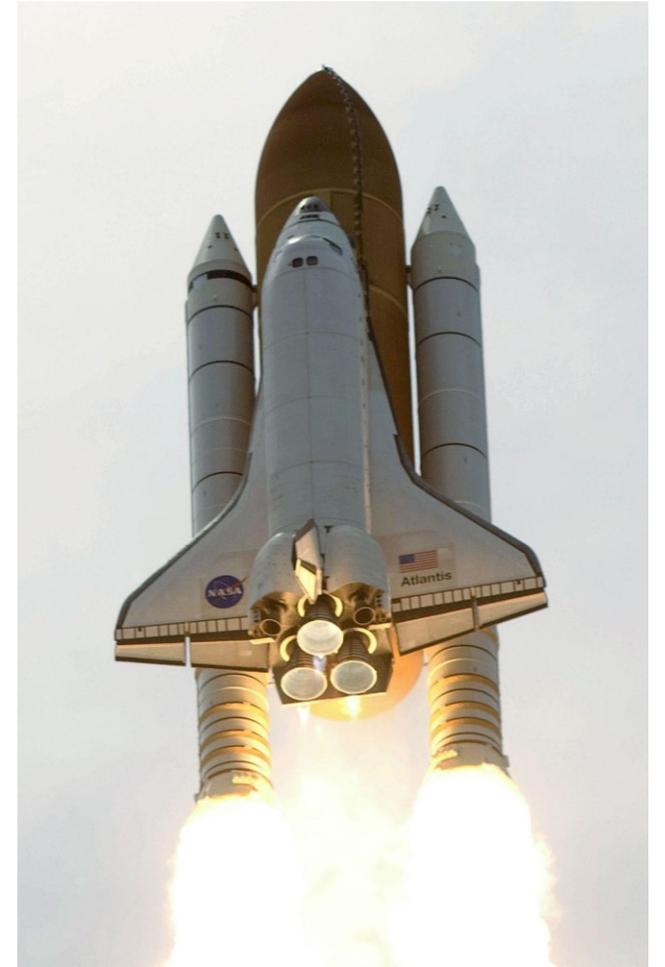


# System Recovery Boost Upgrade Record



- Pay IBM some money and you can have more zIIPs for boost periods!
- Subscription feature that allows for activation of uncharacterized PUs as zIIPs during boost periods
  - Can add up to 20 zIIPs
  - Does allow for violating the 2 to 1 zIIP to GCP ratio
- Handled much like a capacity upgrade record
  - Boost Upgrade record is activated in 6 hour increments for the machine
  - But boosts are still limited to their normal time frames
- Primarily expected to be useful for larger 7xx systems with multiple LPARs that may need to be IPLed in short succession
  - Only available on z15 T01 and z16 A01 machines

When will boosts be triggered?



# Boost classes = triggering events



Class	Where	Duration	WLM Work Routing
Shutdown	Single system	30 minutes	Routes work away from boosted system
IPL	Single system	60 minutes	Routes work to boosted system
Recovery Process	Multiple systems	<=5 minutes	No change

## ● Recovery Processes Boost

- CF data sharing member recovery
  - Triggered by disconnection from lock structure while locks were held
- Sysplex partitioning
  - System removed from a sysplex
- CF structure recovery
  - Structure rebuild, duplex failover, or reduplexing
- HyperSwap Recovery
  - Recovery from storage controller failover

### **New with z16:**

- SVC Dump
- HyperSwap Configuration Load
- **Middleware Region Startup**

# IPL Boost



- Enabled by default
  - Can be controlled in IEASYSxx with BOOST=SYSTEM | ZIIP | SPEED | NONE
    - Selecting ZIIP or SPEED means to only use that type of boost
    - For most customers the default (both) is probably appropriate
- Lasts for 60 minutes
  - Very early IPL only benefits from speed boost because it's single-threaded
  - Hour apparently starts when WLM comes up and zIIP boost is (possibly) enabled

```
13:32:22.40 *MASTER* IEE252I MEMBER IEASYM00 FOUND IN ...
13:32:29.18 *MASTER* IEA681I IPL speed boost is active
13:33:08.33 *MASTER* IEA675I IPL zIIP boost is active with 1 transient zIIP cores
13:33:08.59 WLM IWM064I BOOST ACTIVATED.
14:33:08.35 *MASTER* IEA678I All IPL boosts have ended
14:33:08.62 WLM IWM064I BOOST ENDED.
```

- WLM Routing Services will report boosted capacity
  - May result in more than normal work being routed to boosted LPAR

# Shutdown Boost



- Initiated by running proc IEASDBS
  - Obviously need to change your procedures to do this
  - Shutdown boost can not be invoked again until the LPAR is IPLed
- WLM starts avoiding the LPAR for new work
- Lasts for 30 minutes, until LPAR comes down, or proc IEABE is run
  - At end of 30 minutes, lose boost advantage but WLM will continue to avoid directing work to the LPAR
  - Must figure out when it makes most sense to run IEASDBS to improve shutdown time: probably just before the timeframe of the most CPU consumption even if it means final shutdown steps run unboosted
  - But don't want to run it too far in advance because of the first point
- Respects BOOST= in IEASYSxx

# Recovery Process Boost (RPB)



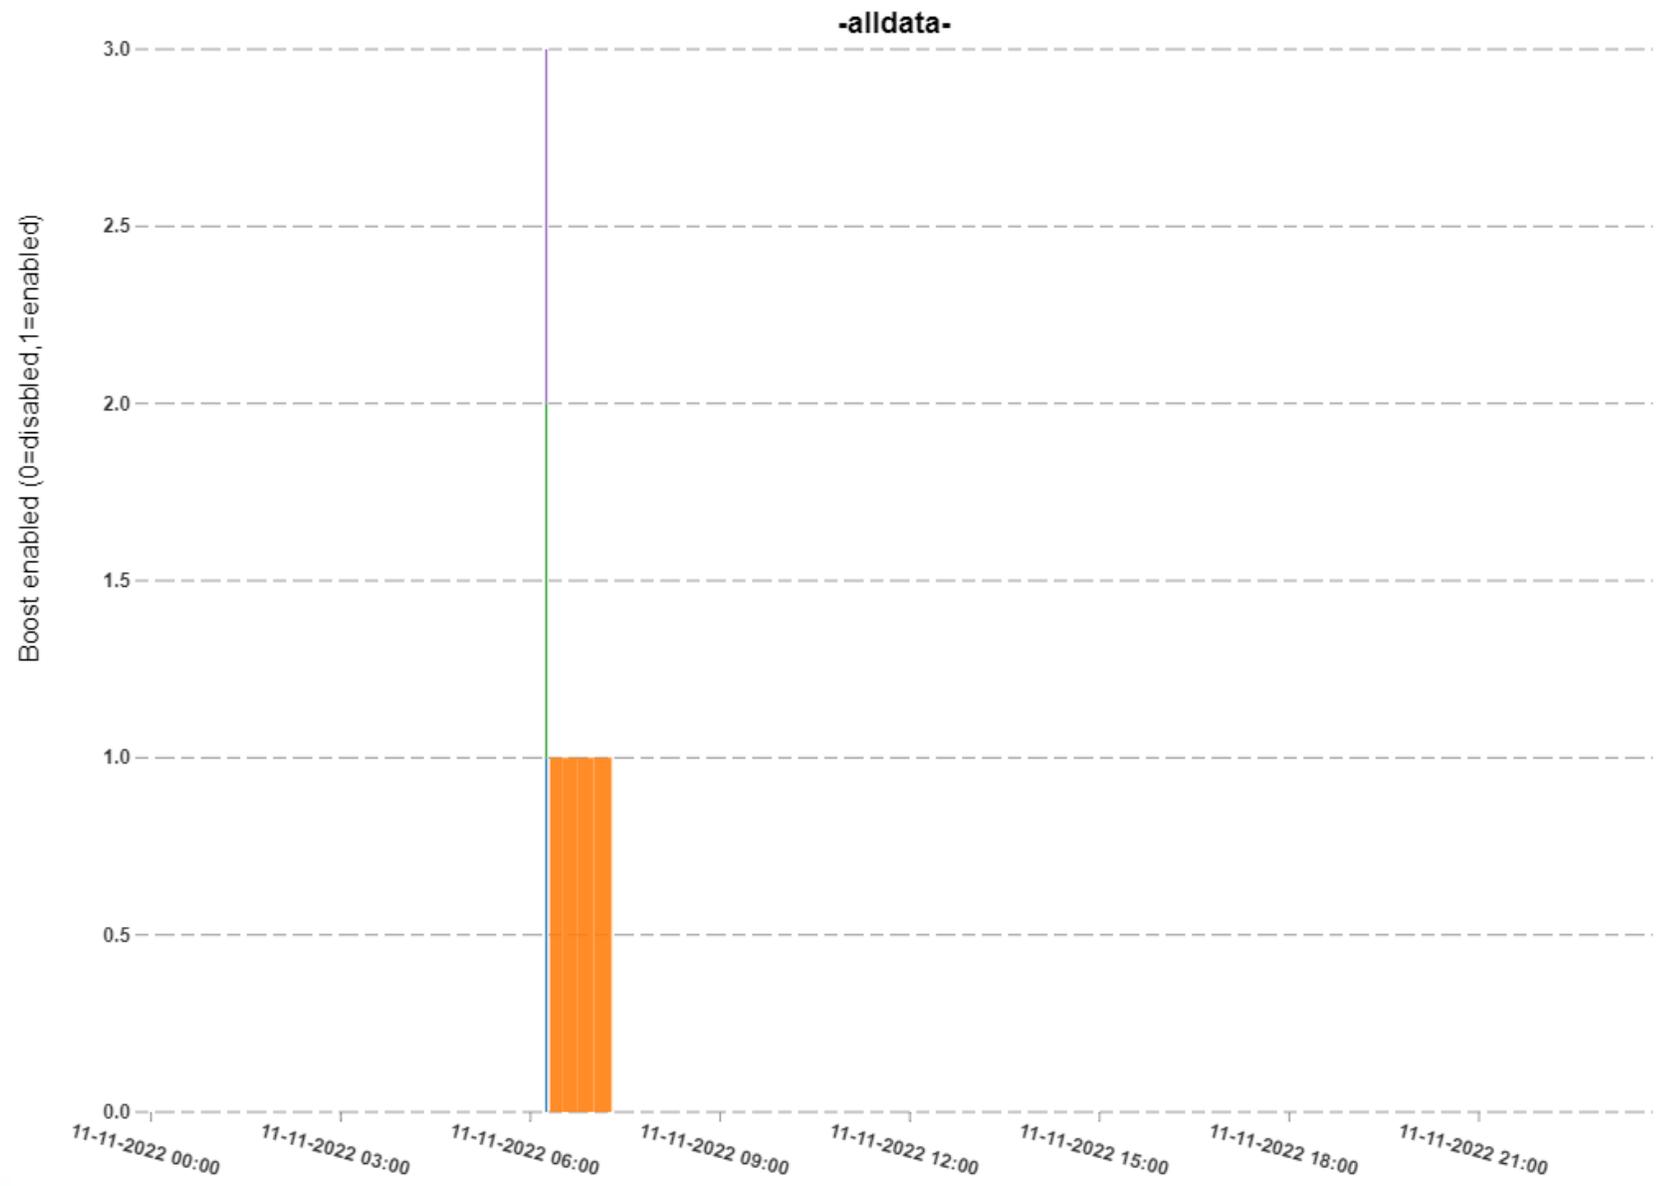
- Automatically initiated by the system
  - Again, respects BOOST= in IEASYSxx
  - Won't be initiated if the LPAR in IPL or Shutdown boost
  - Note: “automatically” might mean in response to an operator starting an address space (Middleware Recovery Boost on z16)
- Short duration ( $\leq 5$  minutes)
  - Max of 30 minutes Recovery Process Boost time / LPAR within any 24 hour period
  - Sysplex partitioning boost seems to be 2 minutes
  - Multiple reasons can overlap, extending the boost
- zIIP Boost limited to bringing online 2 reserved zIIPs
  - Probably doesn't make much sense for more given the short duration

# Which LPARs will be impacted, and how?

Be ready for the turbo to kick in!



# Boost Period Intervals



What was boosted when?

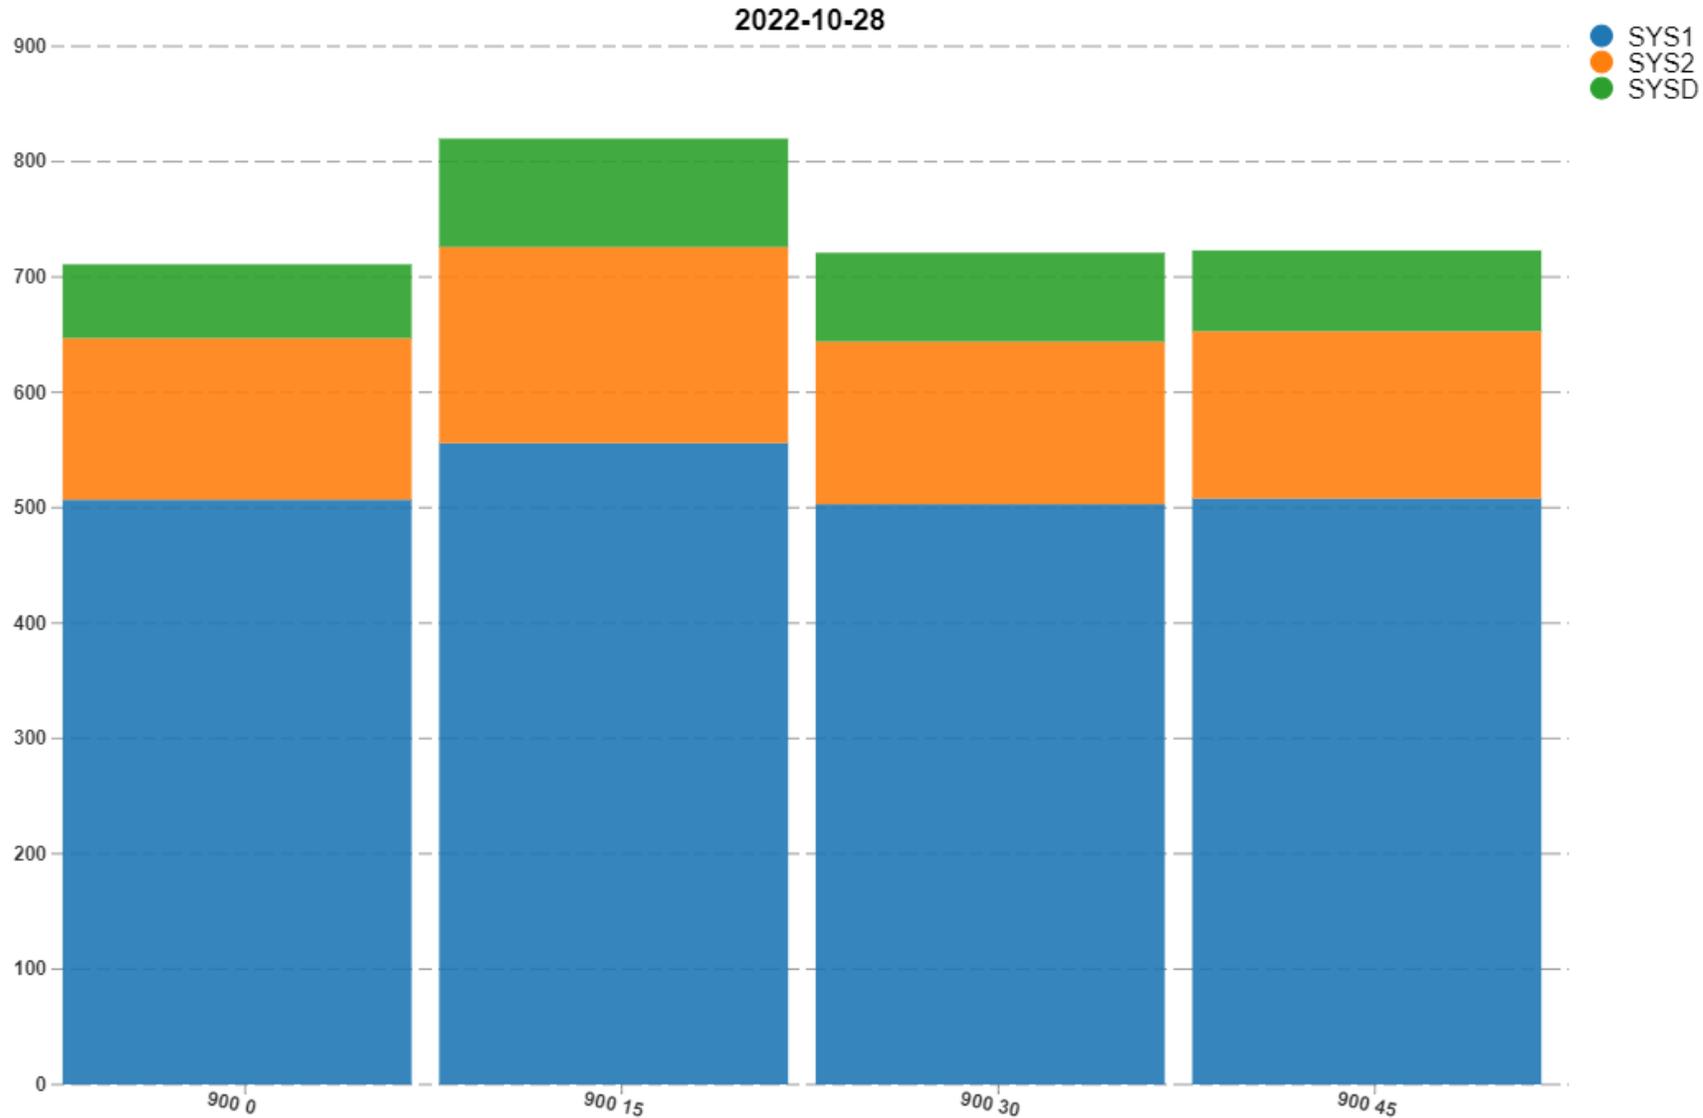
New report for Pivotor customers shows which LPARs were boosted when. Note here the other LPARs in the sysplex were boosted for 2 minutes when SYD1 came out of the sysplex at shutdown. SYD1 then enjoyed IPL boost for an hour. Shutdown boost was not used.

# LPARs being boosted



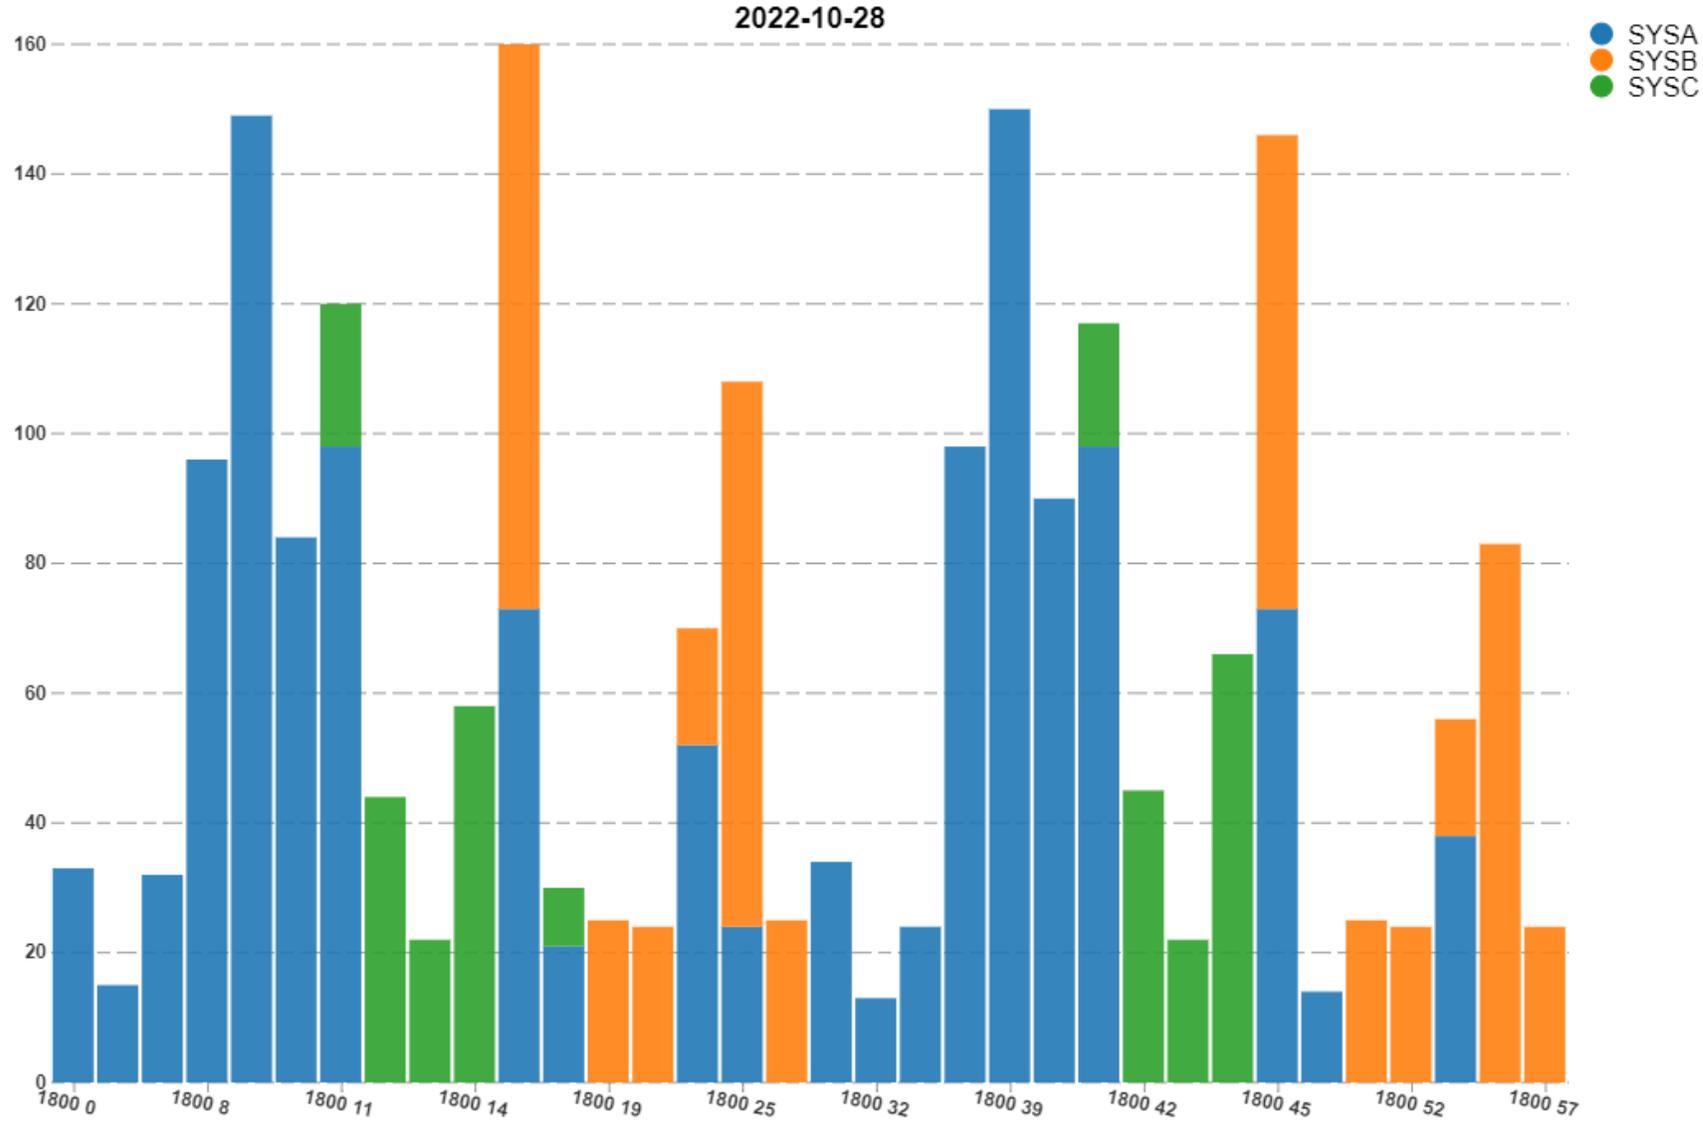
- Seems obvious: they get more work done!
- But may need changes to prepare them to get more work done
  - Maybe: define reserved (offline) zIIPs to be brought online during boost period
  - **May need to add zIIP weight to be able to access newly online zIIPs**
- Increased zIIP usage could cause some zIIP work to cross over to the GPs
  - Work with HONORPRIORITY=NO prevented from crossing over
  - So *may* want to consider changing *from* HONORPRIORITY=NO during IPL boost
    - Probably unnecessary in most cases
    - Recovery process boosts likely short enough to not be a major concern
    - Work should be routed away from the LPAR during shutdown boost anyways
- Boosted LPARs over-achieving work won't be capped to help discretionary
- New SMF interval started when boost starts/ends
  - **Trap: must have SMF intervals SYNCed to get clean SMF intervals!**
  - **Use SUBSYS(xxx,INTERVAL(SMF,SYNC)) in your SMFPRMxx!** ←

# SMF Interval Check



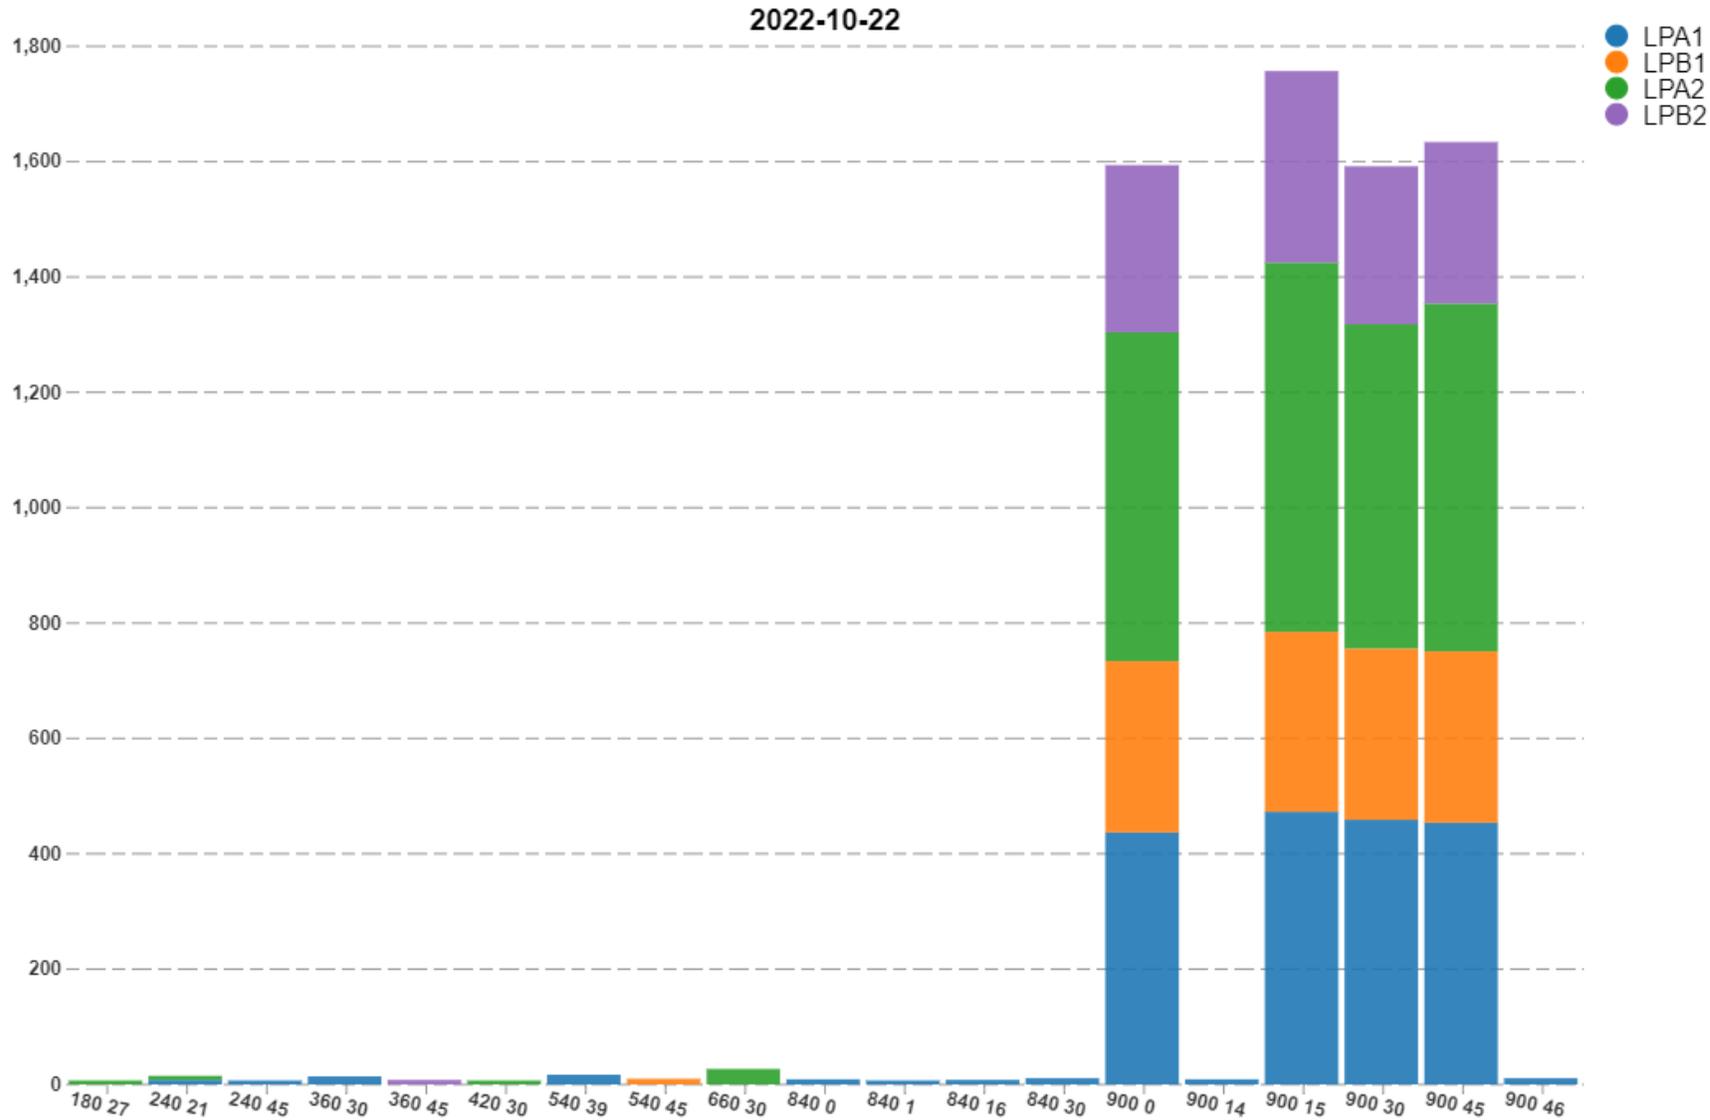
For Pivotor customers: this reports should look something much like this. 15 minute intervals, all synced, all systems using the same settings.

# SMF Interval Check



You don't want it to look like this: 30 minute intervals that aren't synced.

# SMF Interval Check



Note that you should expect that boosts will cause a few short intervals: this is to be expected.

# LPARs *not* being boosted



- These LPARs *could* be negatively impacted
- CPU cache effectiveness may be impacted by boosted LPARs (at least theoretically)
  - Speed boosted LPARs do more work per unit of time on the GPs
  - zIIP boosted LPARs may drive more work to zIIPs
  - zIIP caps ignored during boost periods
- Higher physical zIIP utilization may impact non-boosted LPAR's ability to get work dispatched on zIIPs
  - Could potentially lead to more crossover
  - zIIP caps ignored but weight enforcement still applies
  - **May need to change relative weights during boost periods, if trying to protect unboosted LPARs**
    - Or (in theory) if trying to give more capacity to boosted LPAR
- Resource Group caps with sysplex scope don't count work running on boosted LPARs
  - May allow more than expected work to run on unboosted systems
- Conversely:
  - If boosted systems consume less of GP capacity, might help non-boosted systems

# Weighty issues



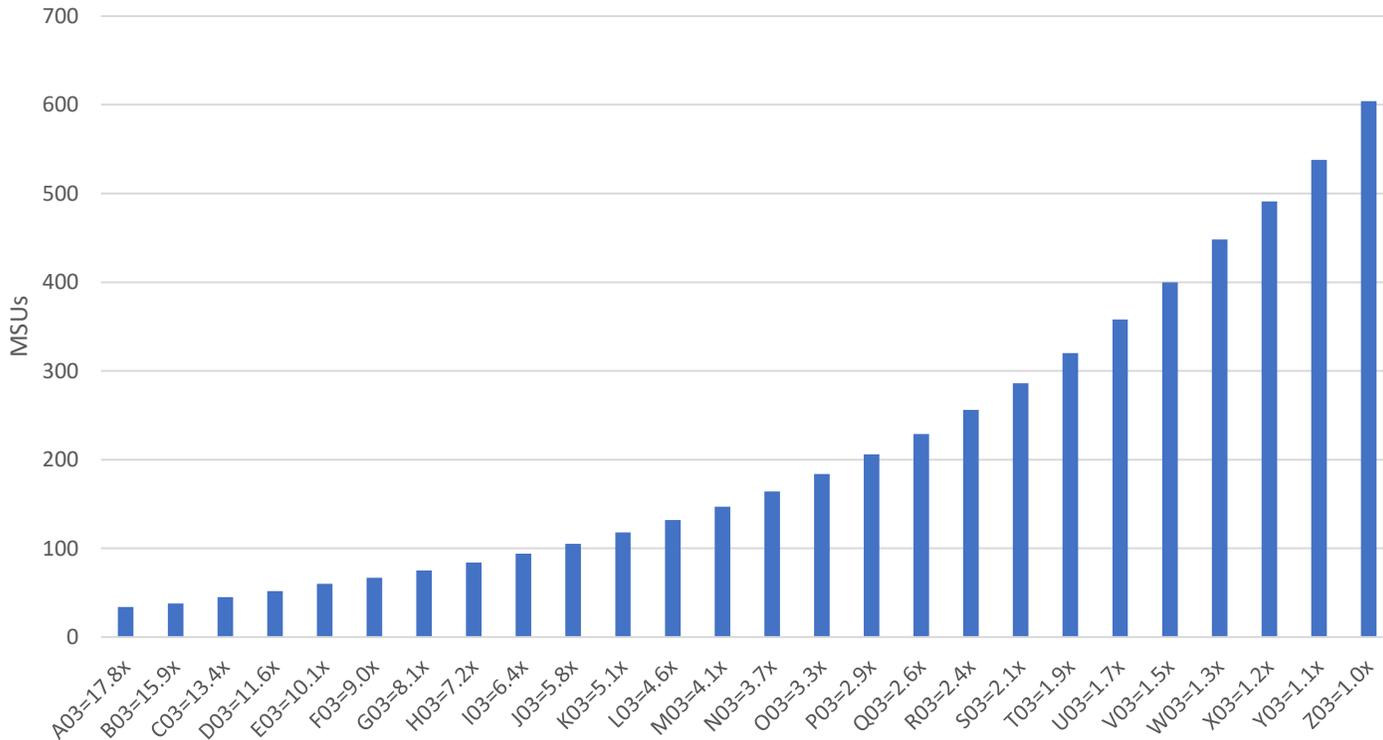
- Boosts happen on a per-LPAR basis
- PR/SM dispatches physical CPs to LPARs as normal
  - Speed boosted LPARs simply get more useful work done in their dispatch interval
- Absent weight changes:
  - zIIP boosted LPARs with low weights may end up using low pool zIIPs
    - May not be able to access that capacity if other LPARs busy
  - Other LPARs using more than their zIIP weight may be limited if they are borrowing weight from the boosted LPAR

# How big of a speed boost?



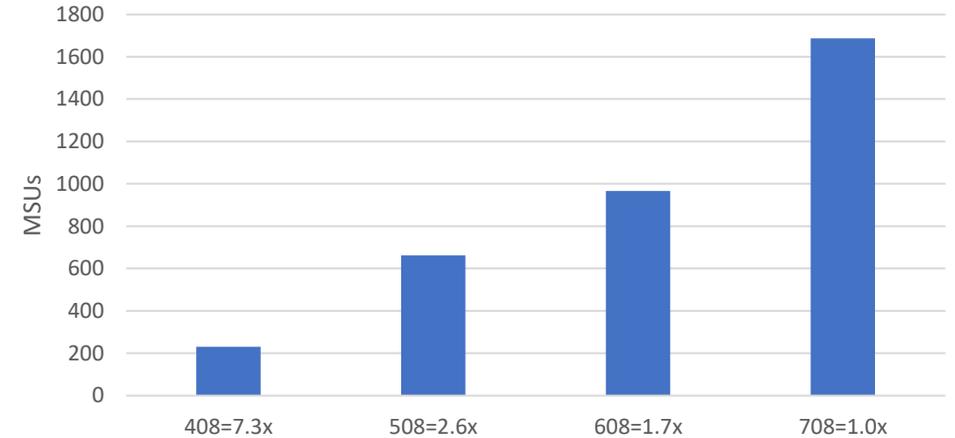
- Customers on slower sub-capacity machines could see a huge speed boost!

8562 3-way machine comparisons

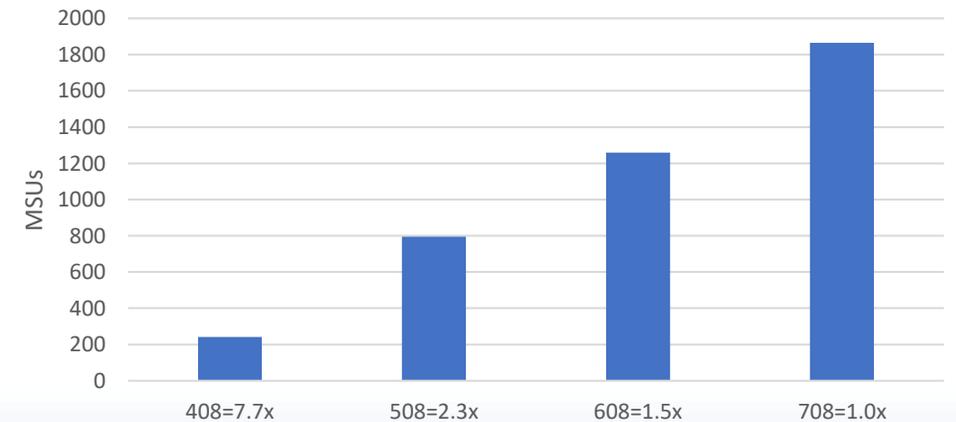


Plus zIIP boost could add even more capacity!

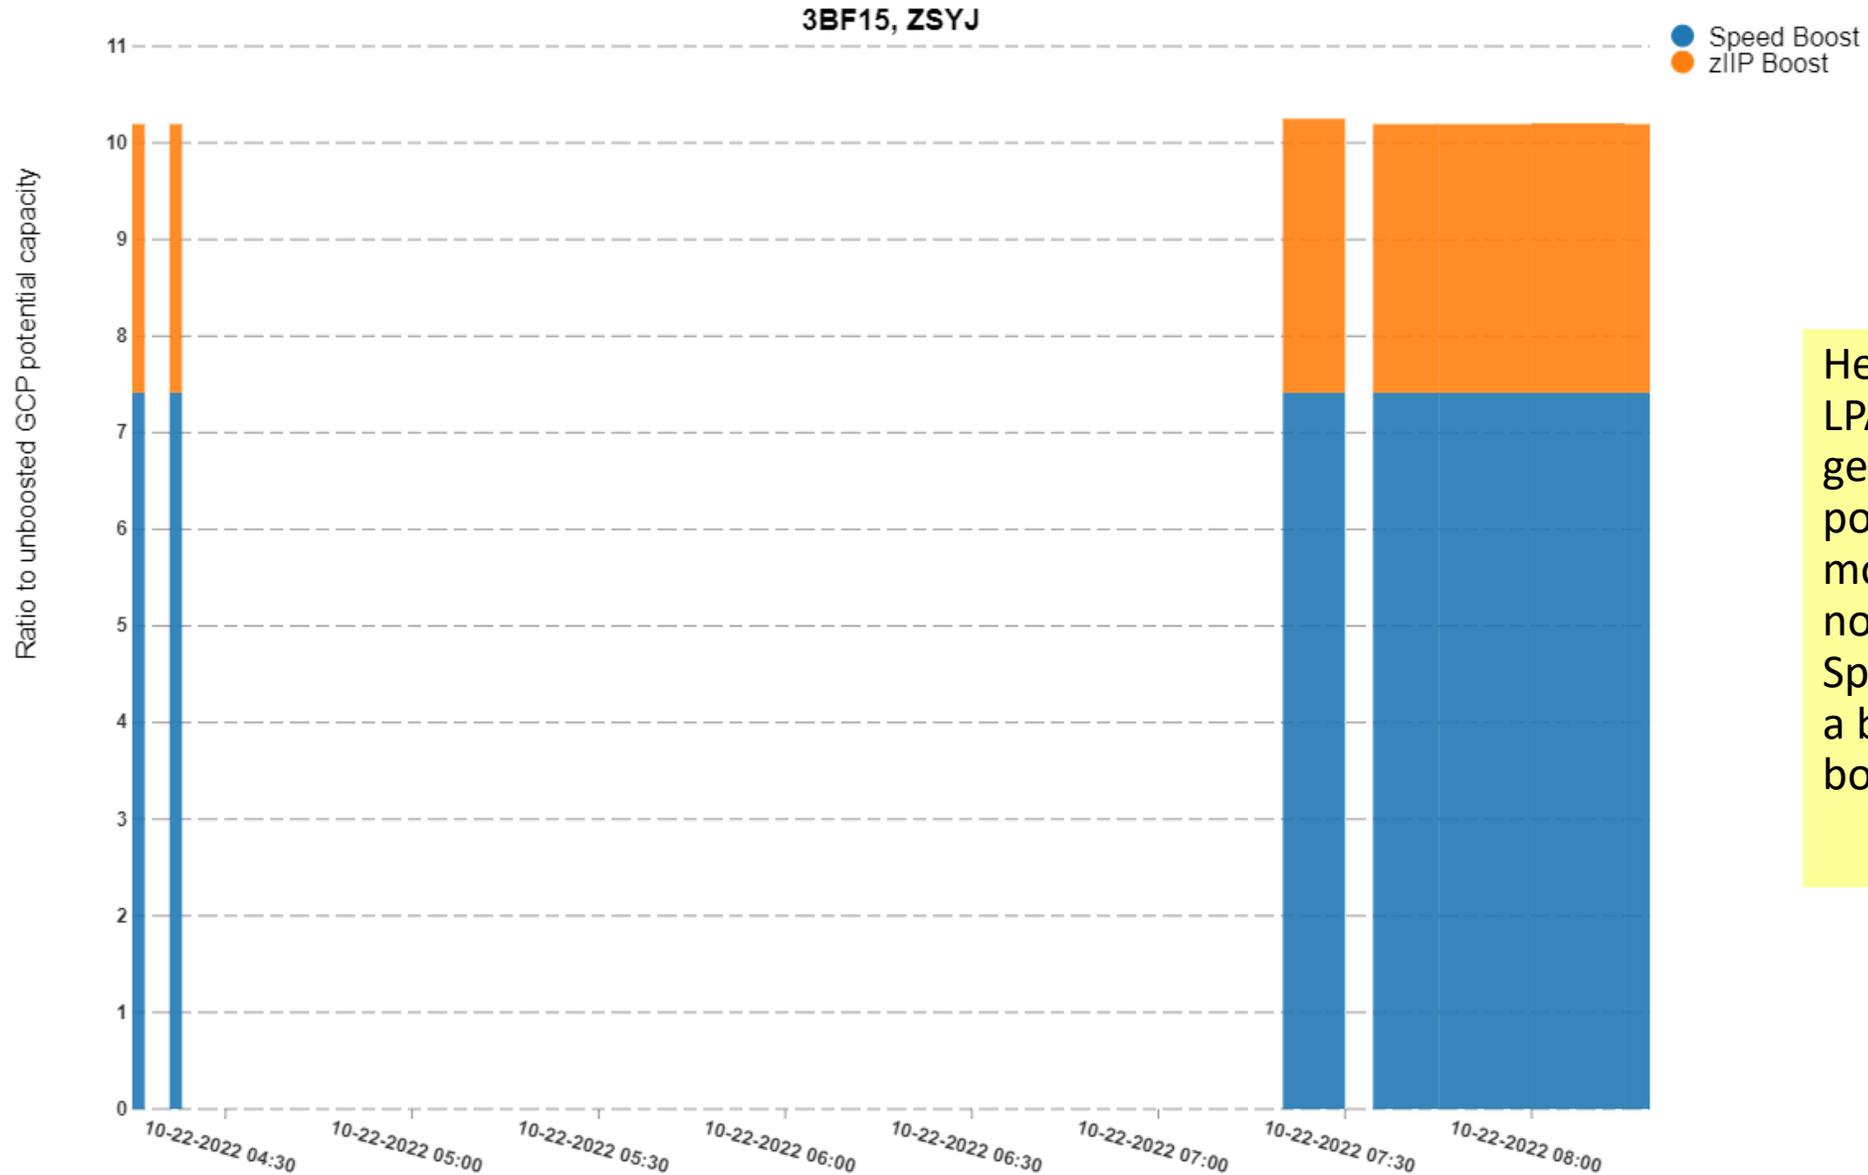
8561 8-way machine comparisons



3961 8-way machine comparisons

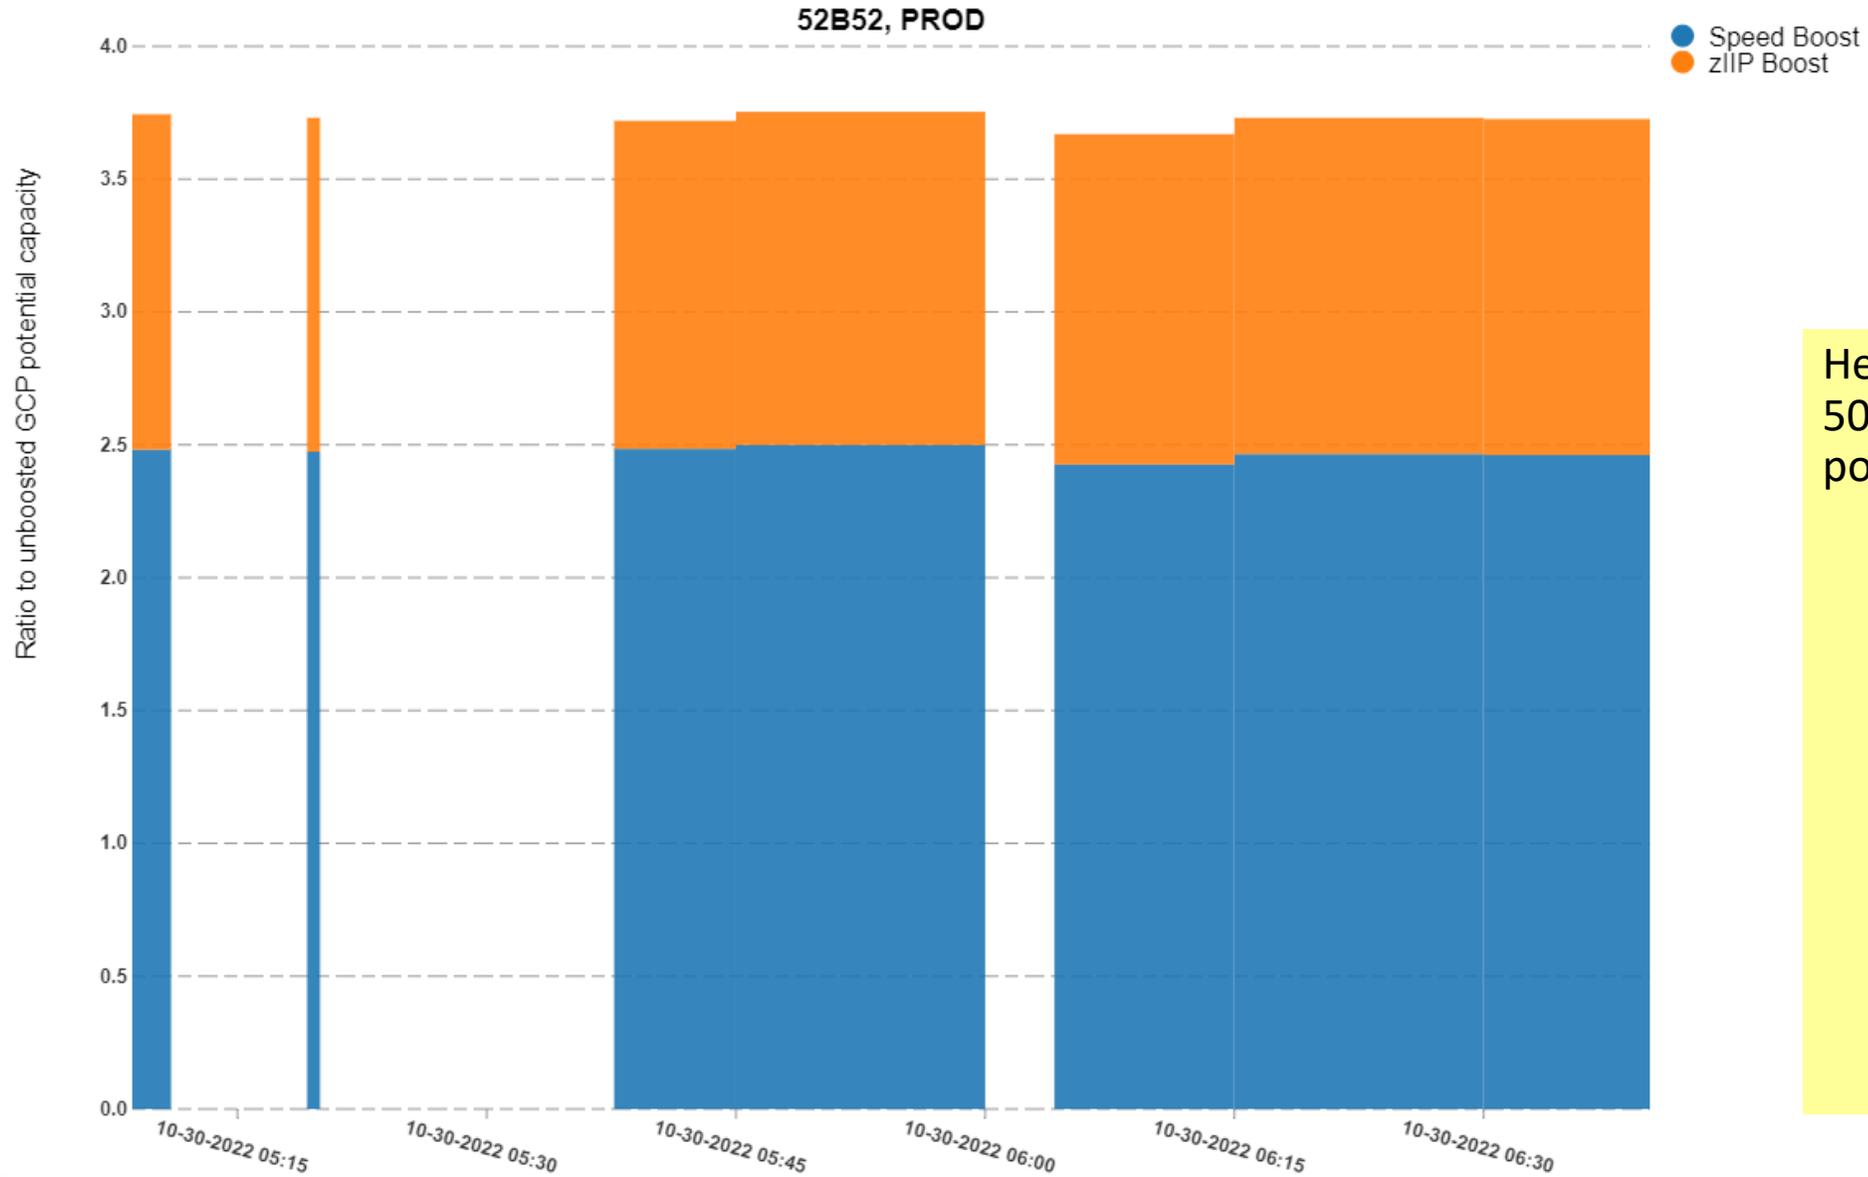


# Boost Potential Capacity Ratio



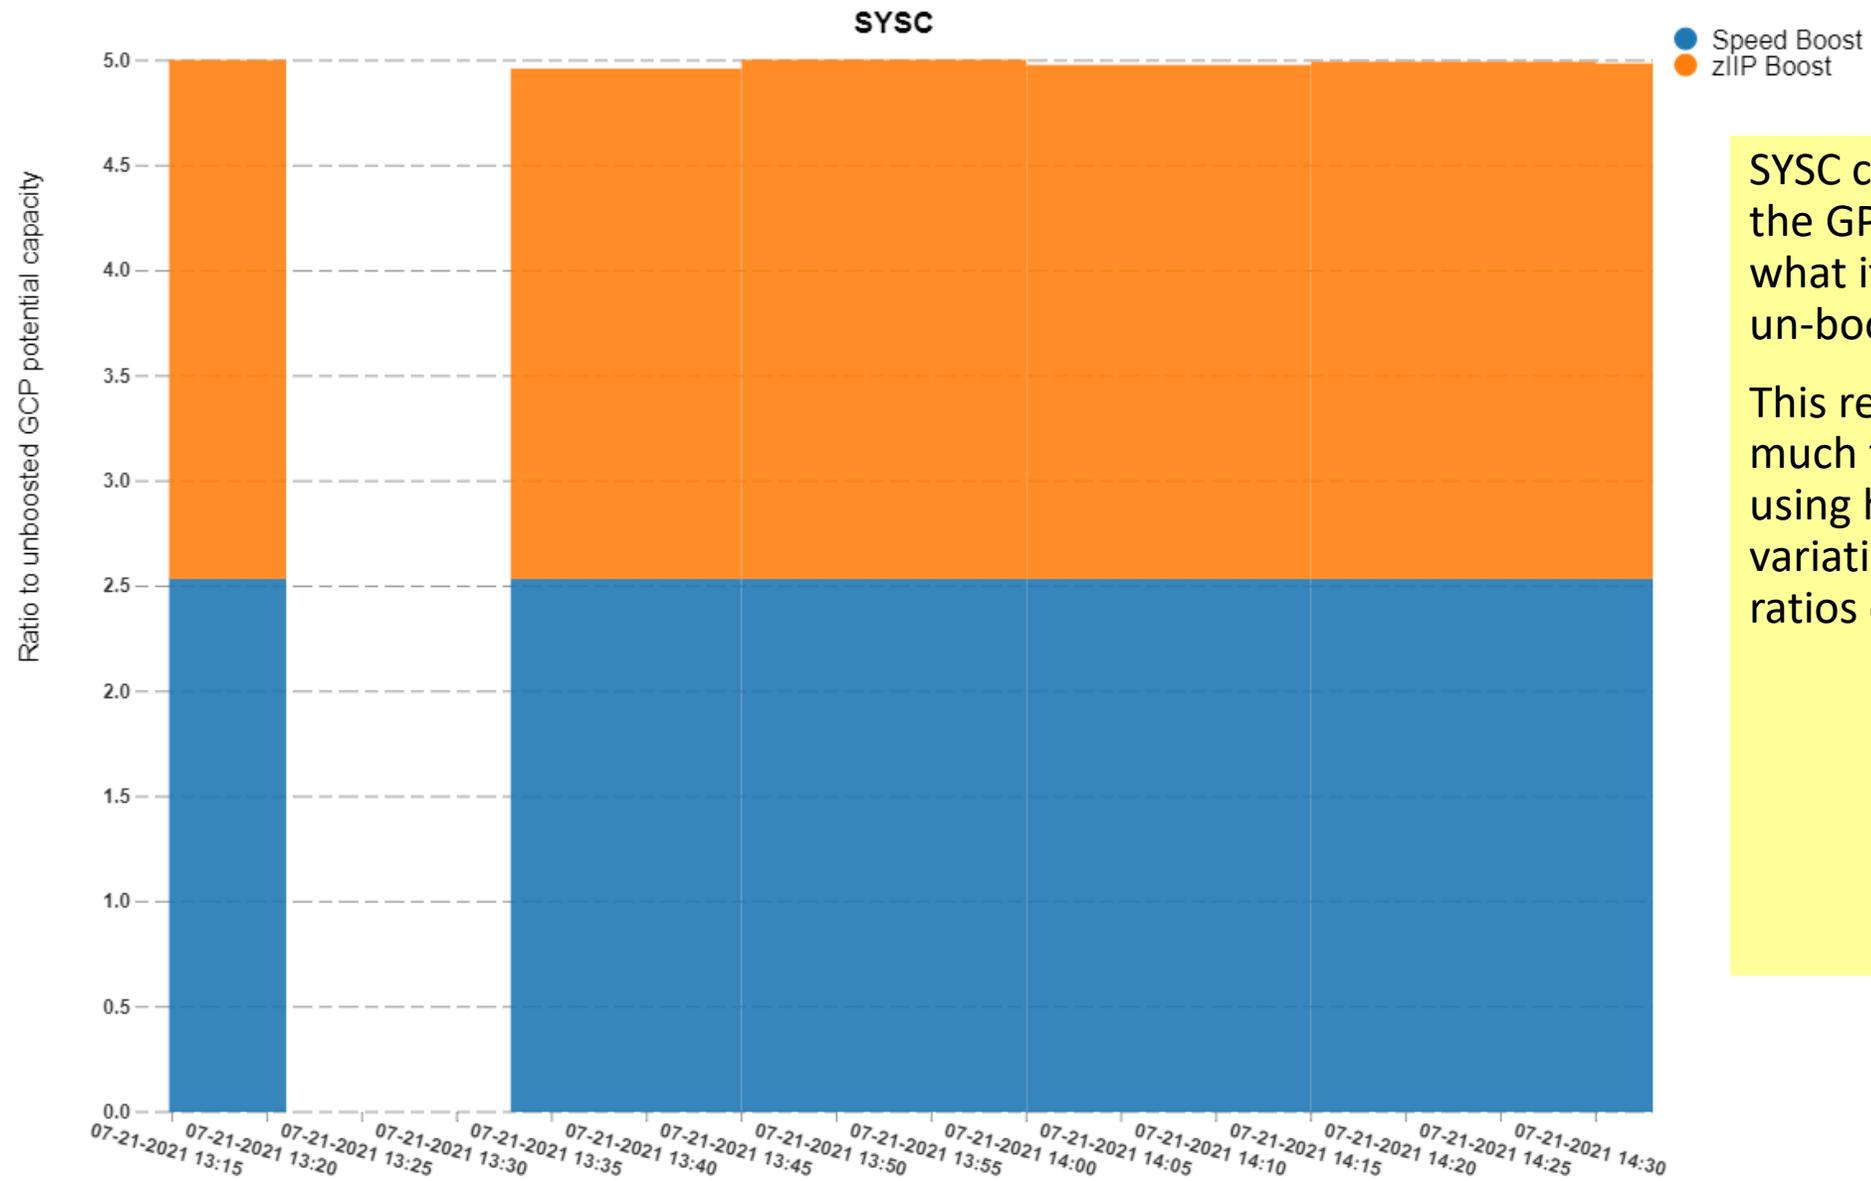
Here's an example of an LPAR on a z15-411 getting access to potentially over 10x more capacity than it normally has access to. Speed boost is obviously a big part of this, but zIIP boost is helping too.

# Boost Potential Capacity Ratio



Here an LPAR on a z15-504 has a boost potential of around 3.7x.

# Boost Potential Capacity Ratio



SYSC could have done ~5x the GP work compared to what it could have done un-boosted.

This report considers how much the other LPARs are using hence the slight variations you see in the ratios on each chart.

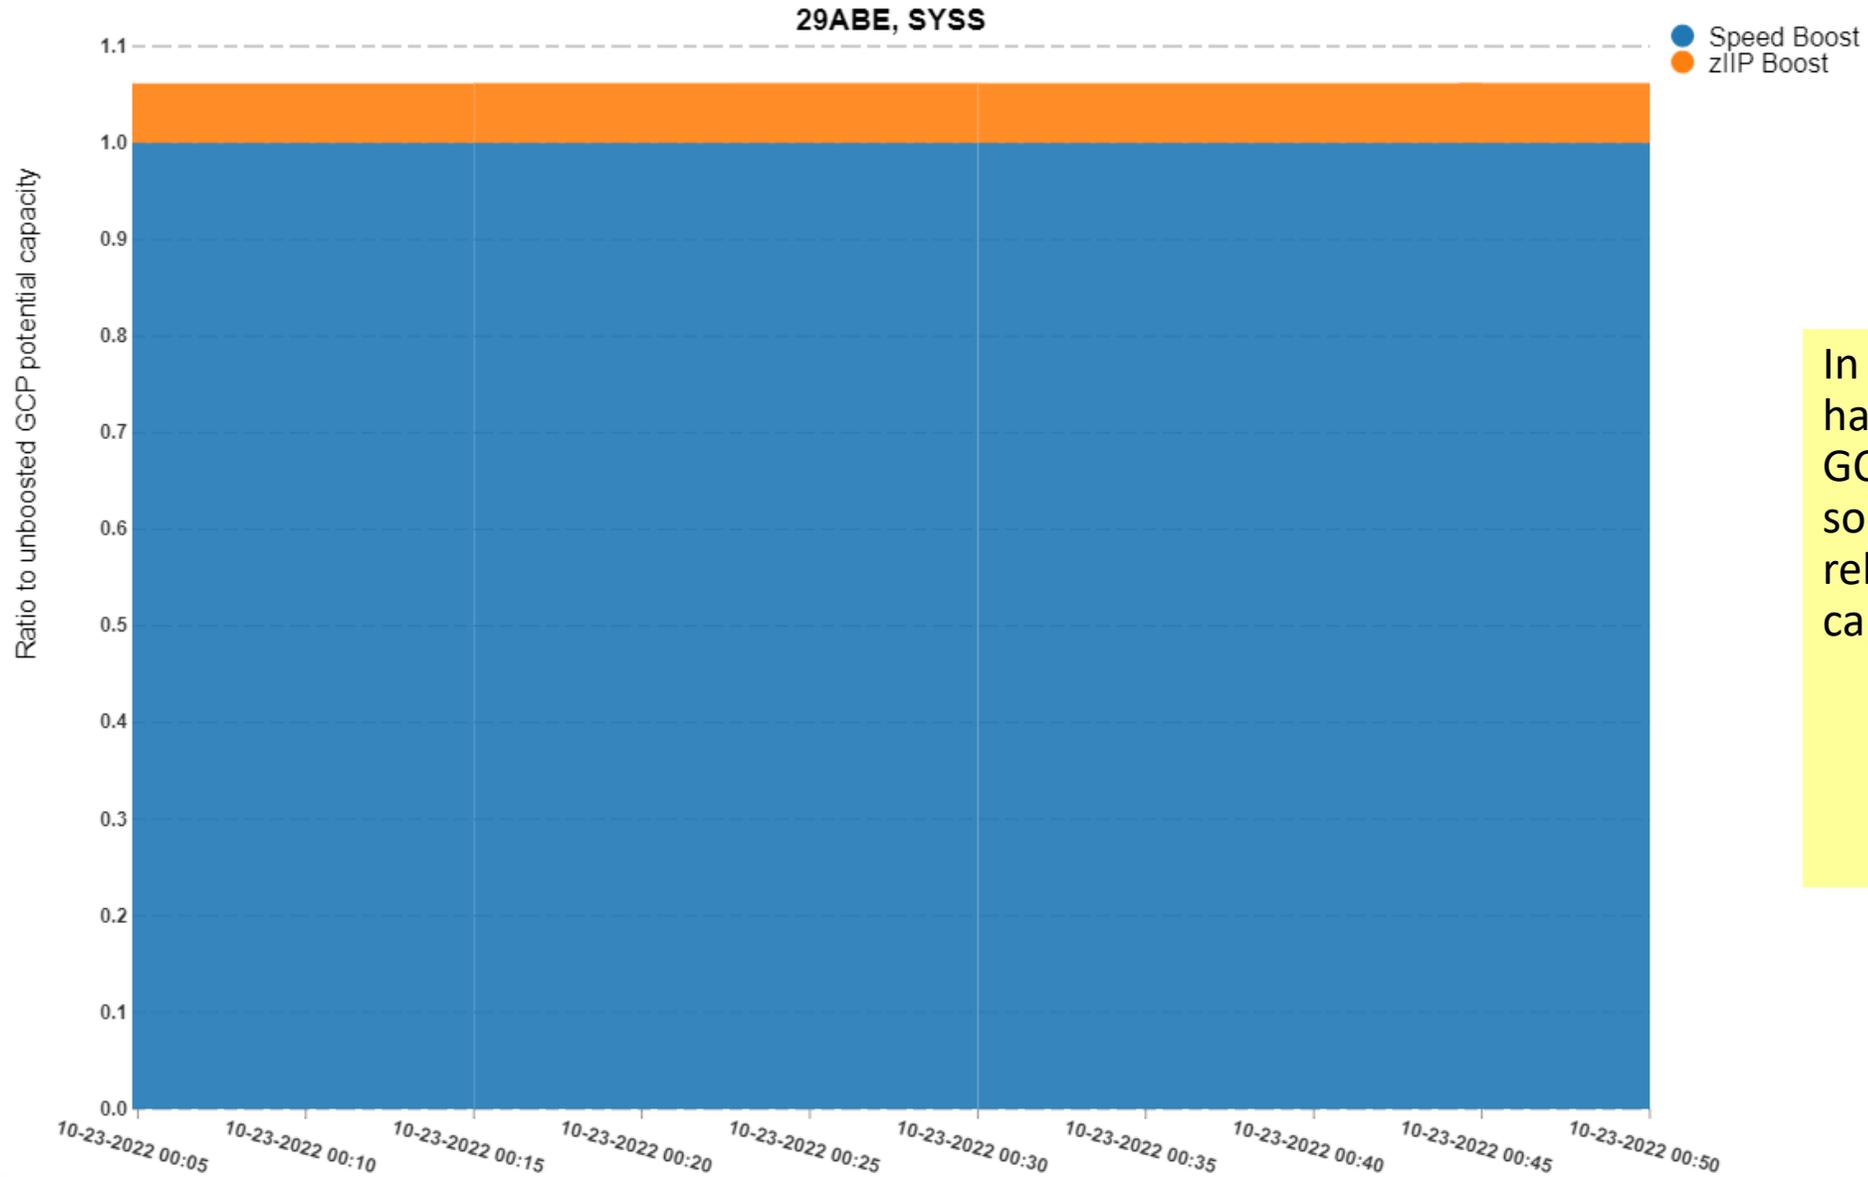
# Things that limit SRB's effectiveness



- Full speed GCPs (no speed boost for full-speed processors)
- Few or no zIIPs available to the boosted LPAR
- Busy zIIPs
  - Less available capacity to the boosted LPAR
- Single physical zIIP shared among several LPARs
  - PR/SM will still move the zIIP between LPARs limiting the time the boosted LPAR might be able to use it
- Weights are still enforced
  - zIIPs and GPs on low-weighted LPARs may not get the processors as often as it would like if the other LPARs are busy and demanding their weight

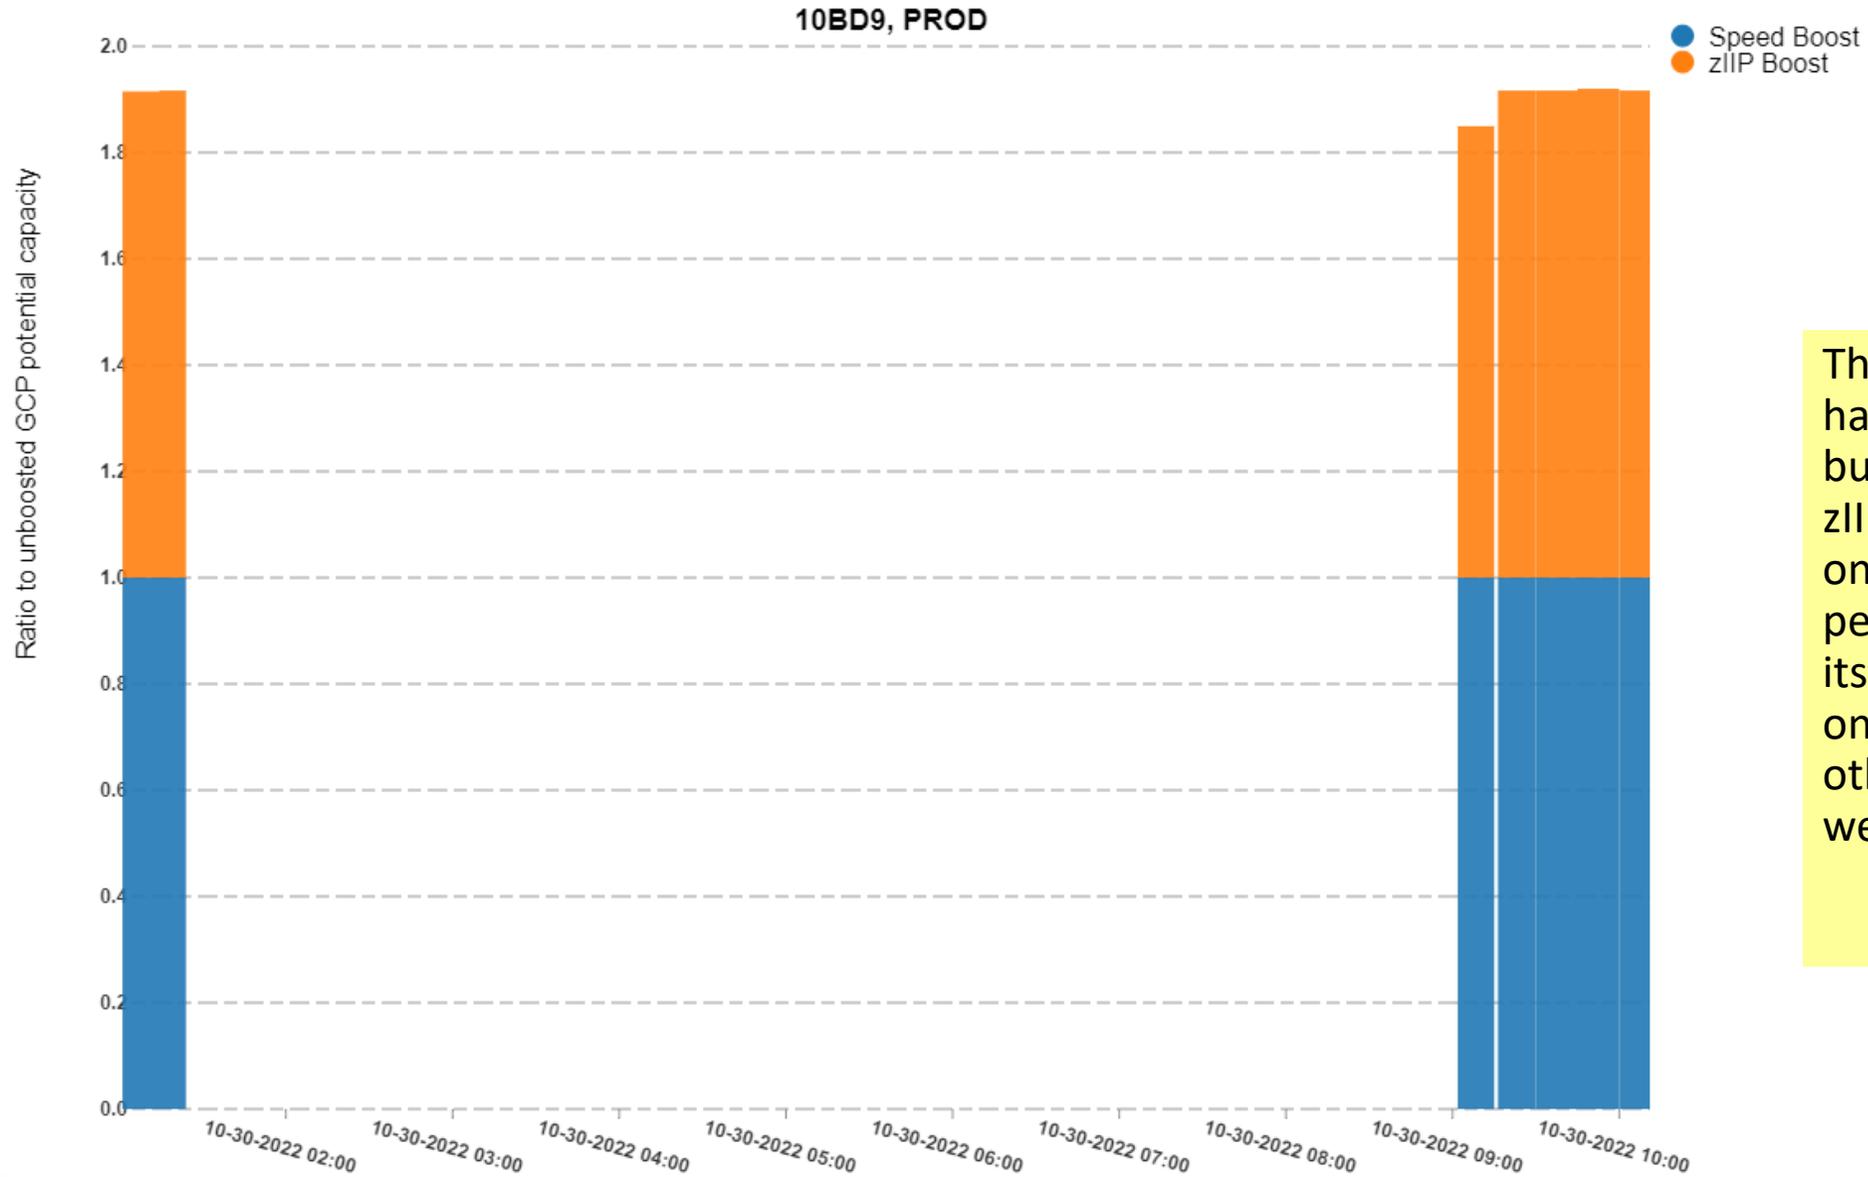
Scott's Opinion: In the 2020s, almost every machine should have at least 2 zIIPs!

# Boost Potential Capacity Ratio



In this example, the LPAR had about 20 full-speed GCPs and only one zIIP, so boost provides relatively little additional capacity.

# Boost Potential Capacity Ratio



This LPAR on a z15-740 had 16 GPs and 10 zIIPs, but another 8 reserved zIIPs were brought online during the boost periods, almost doubling its available capacity once the weights and other LPARs' activity were taken into account.

How is the SMF data affected?



# CPU Accounting during boost periods

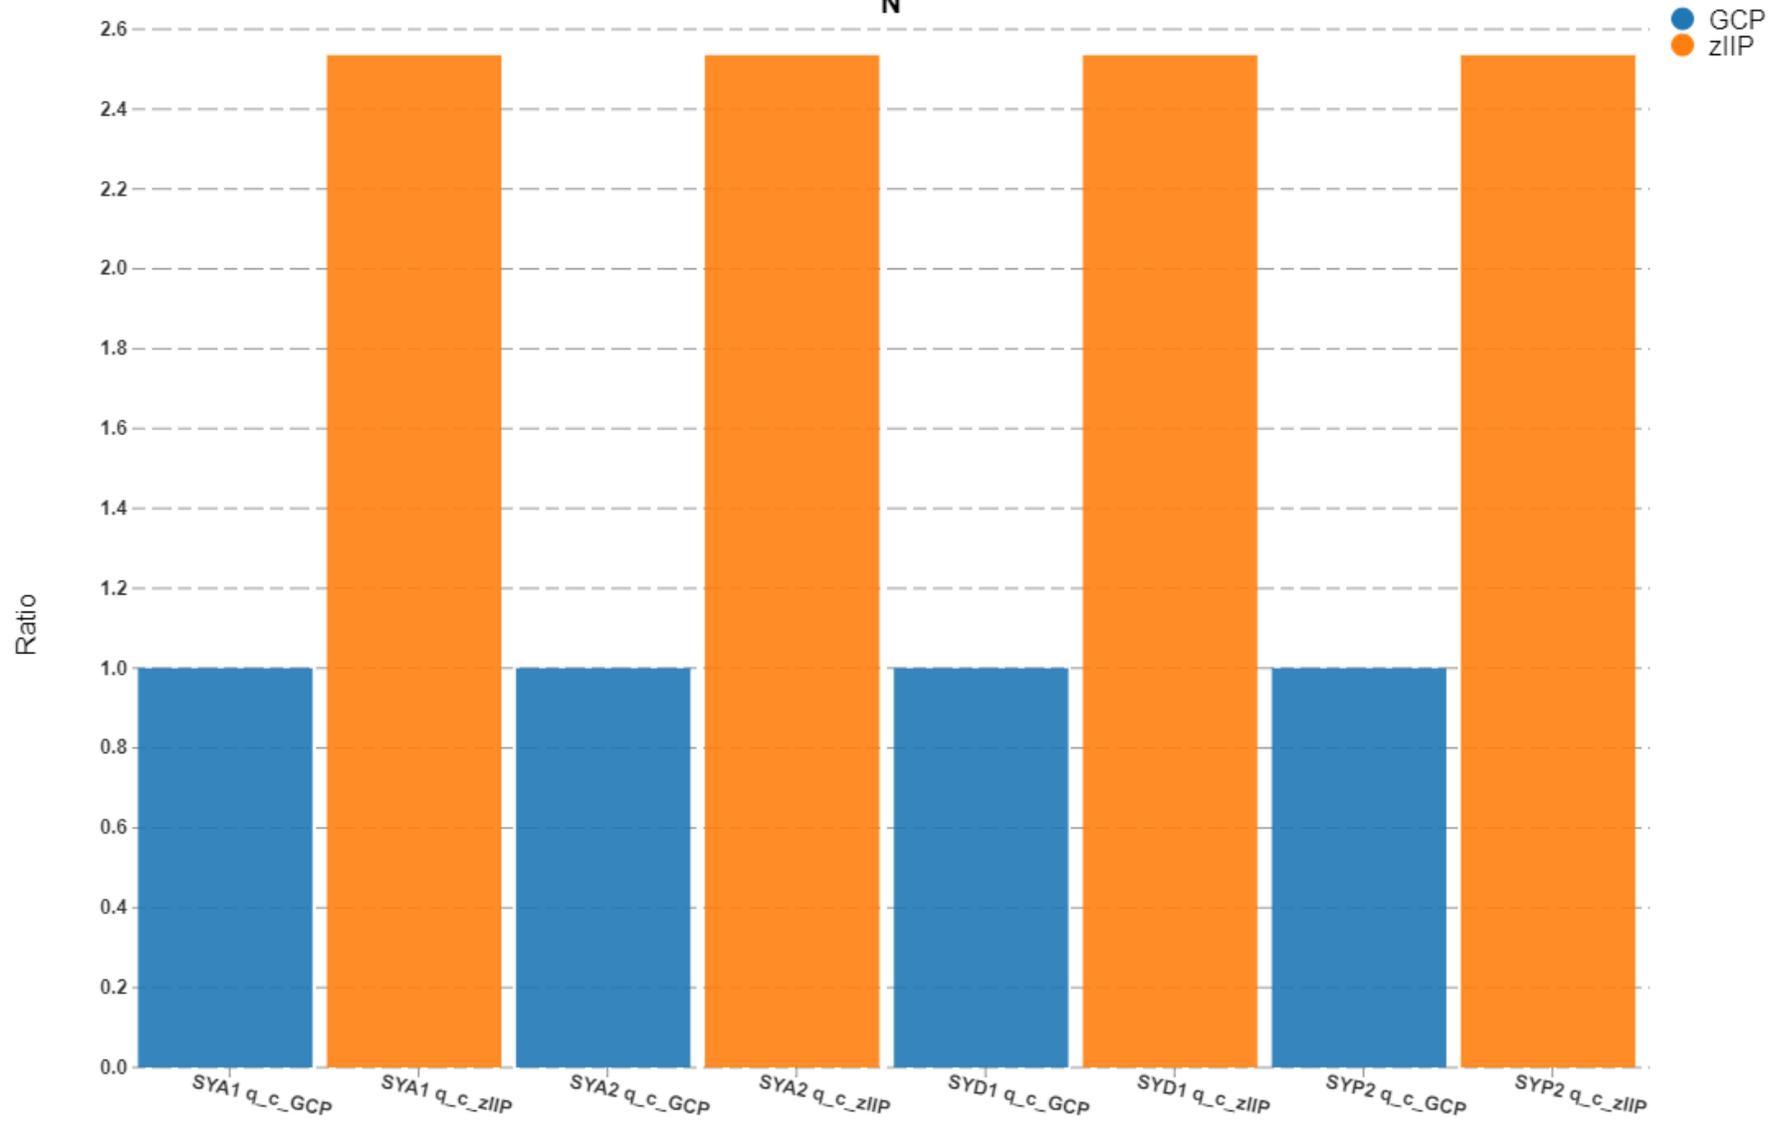


- Generally (across record types), with speed boost:
  - SU/SEC increases
  - Normalization changes to indicate zIIP = CP
  - MSU rating of machine does *not* change
- SMF 70 – Dispatch times unaffected and report LPARs’ “true” usage
  - I.E. work run on the zIIP increases the zIIP’s effective dispatch time
- SMF 30 & 72 – Records GP-only work on GP, even if it ran on zIIP!
  - GP work running on zIIP in SMT is recorded as MT1ET
  - For sub-capacity machines, CPU times will be lower than non-boosted periods
- **In short: measurements during boost period should not be used for capacity planning or performance tuning!**

# zIIP to GCP Ratio

Speedboost:

N



● GCP  
● zIIP

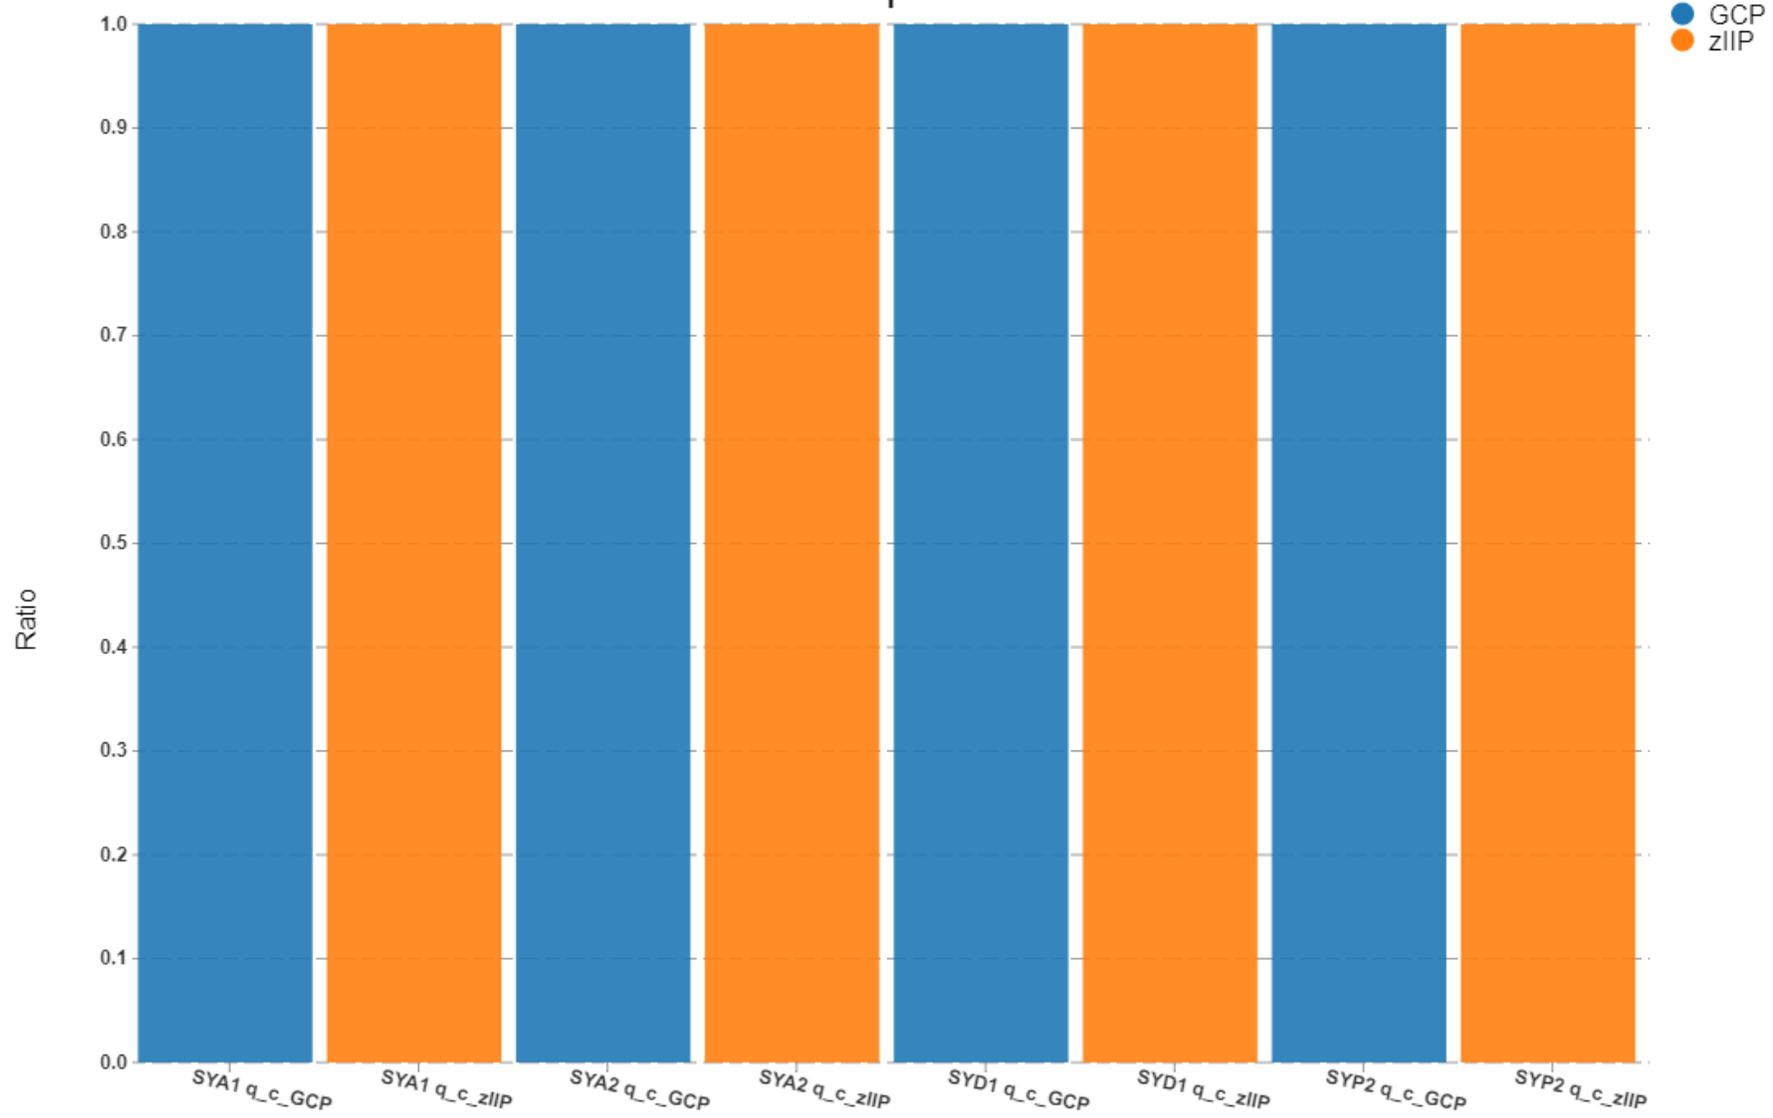


zIIP to GCP ratio on a z15-514: a zIIPs provides a bit over 2.5x the capacity of a GP.

# zIIP to GCP Ratio

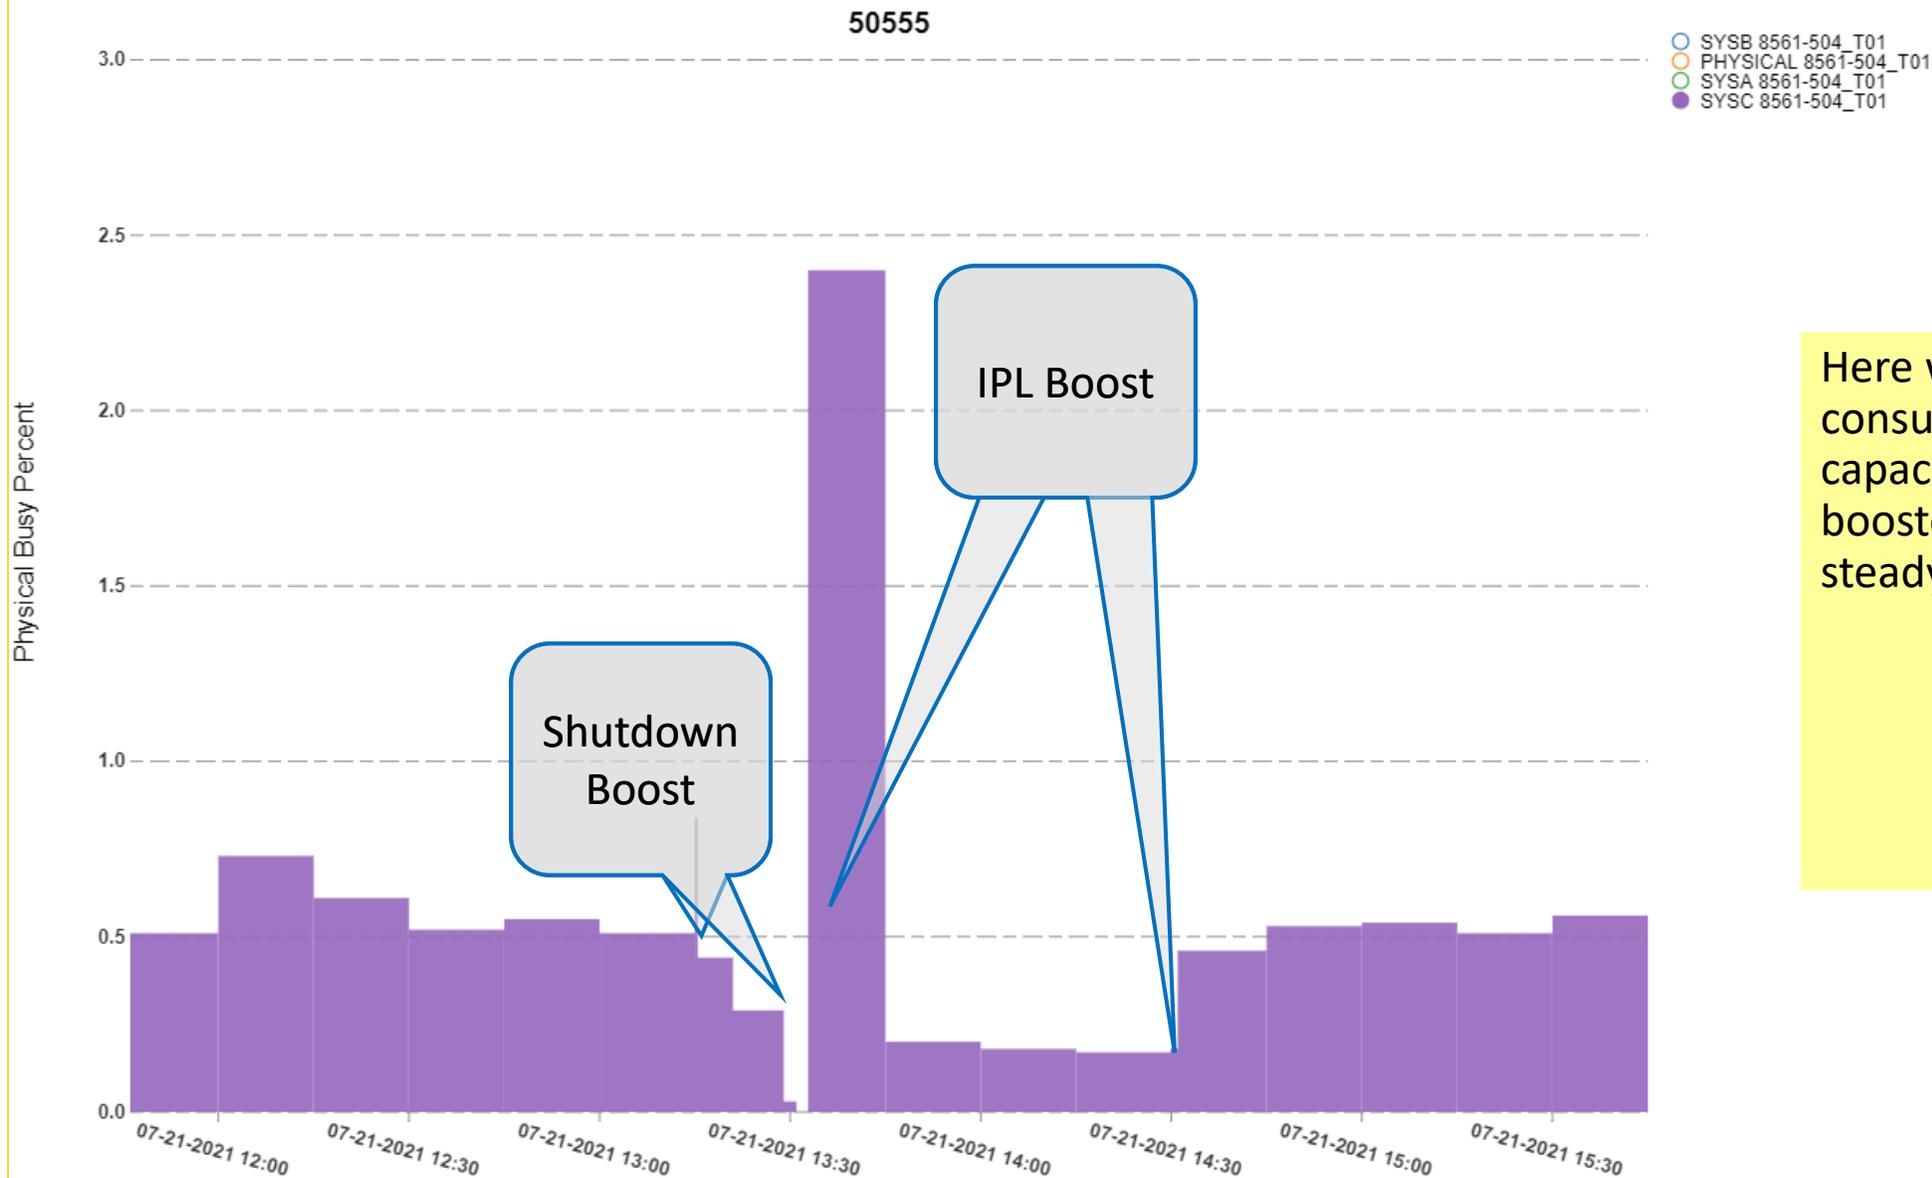
Speedboost:

Y



Same system, when speed boost activated: GPs provide the same capacity as zIIPs.

# CEC Physical Machine CP Busy% by CEC Serial Number



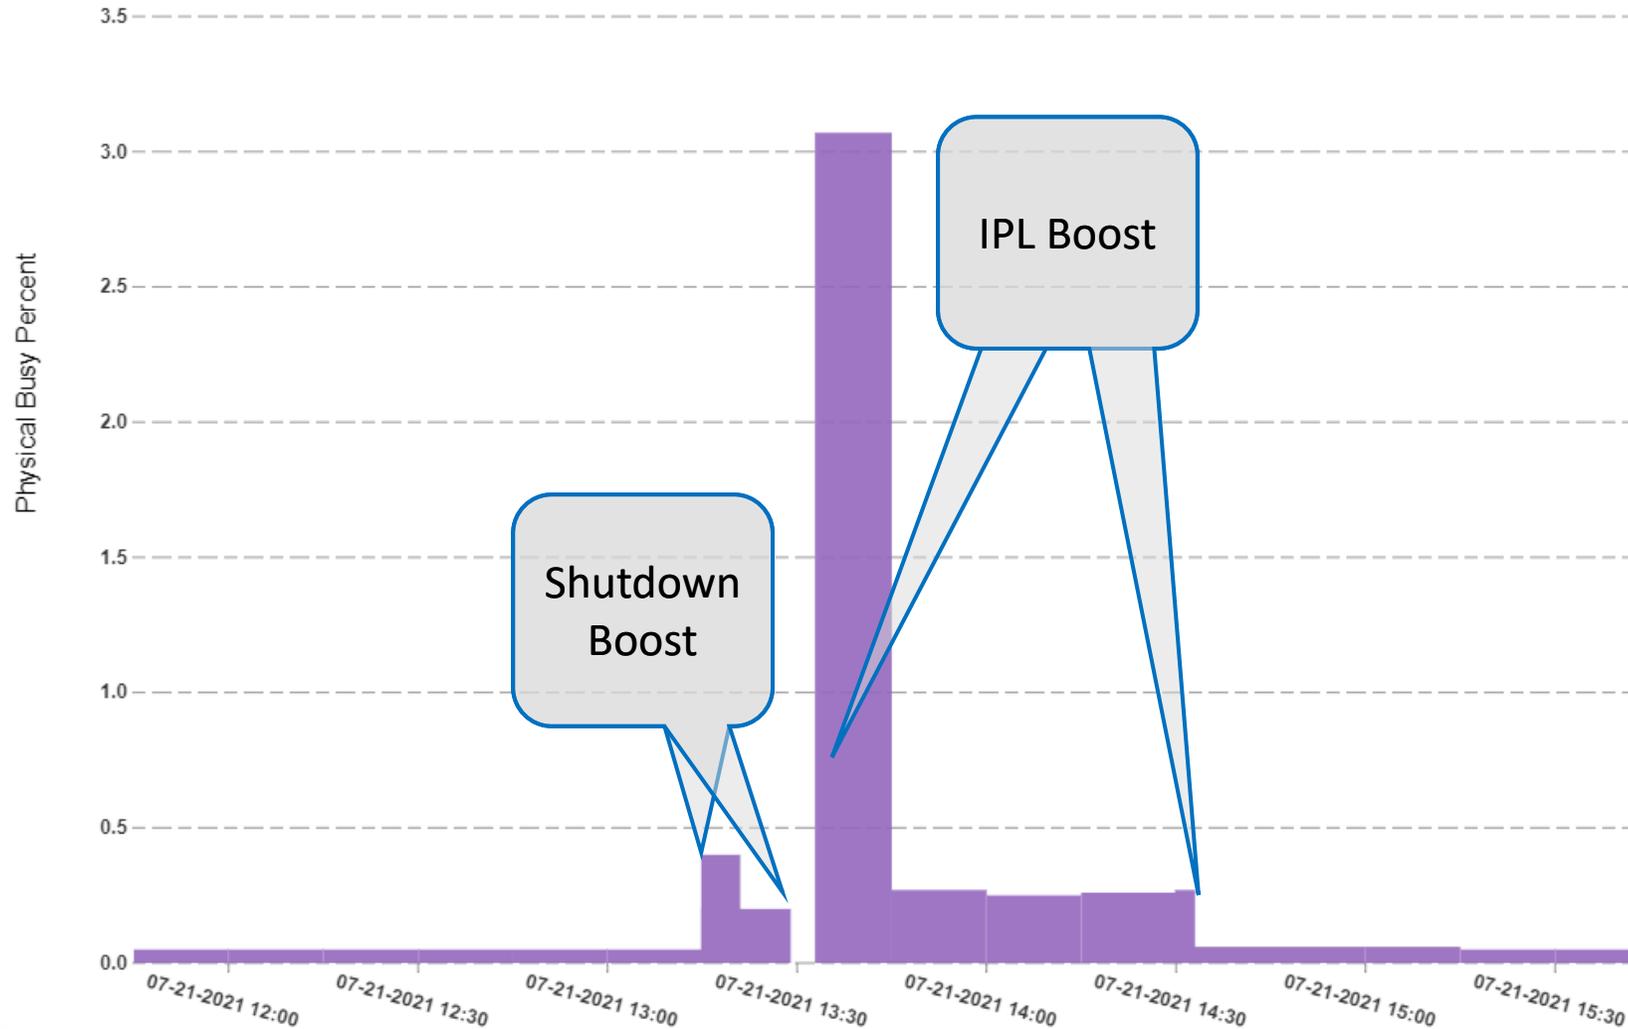
Here we see SYSC consumed less GCP capacity while it was boosted vs. normal steady state.

# CEC Physical Machine zIIP Busy%



50555

- SYSB 8561-504\_T01
- PHYSICAL 8561-504\_T01
- SYSA 8561-504\_T01
- SYSC 8561-504\_T01

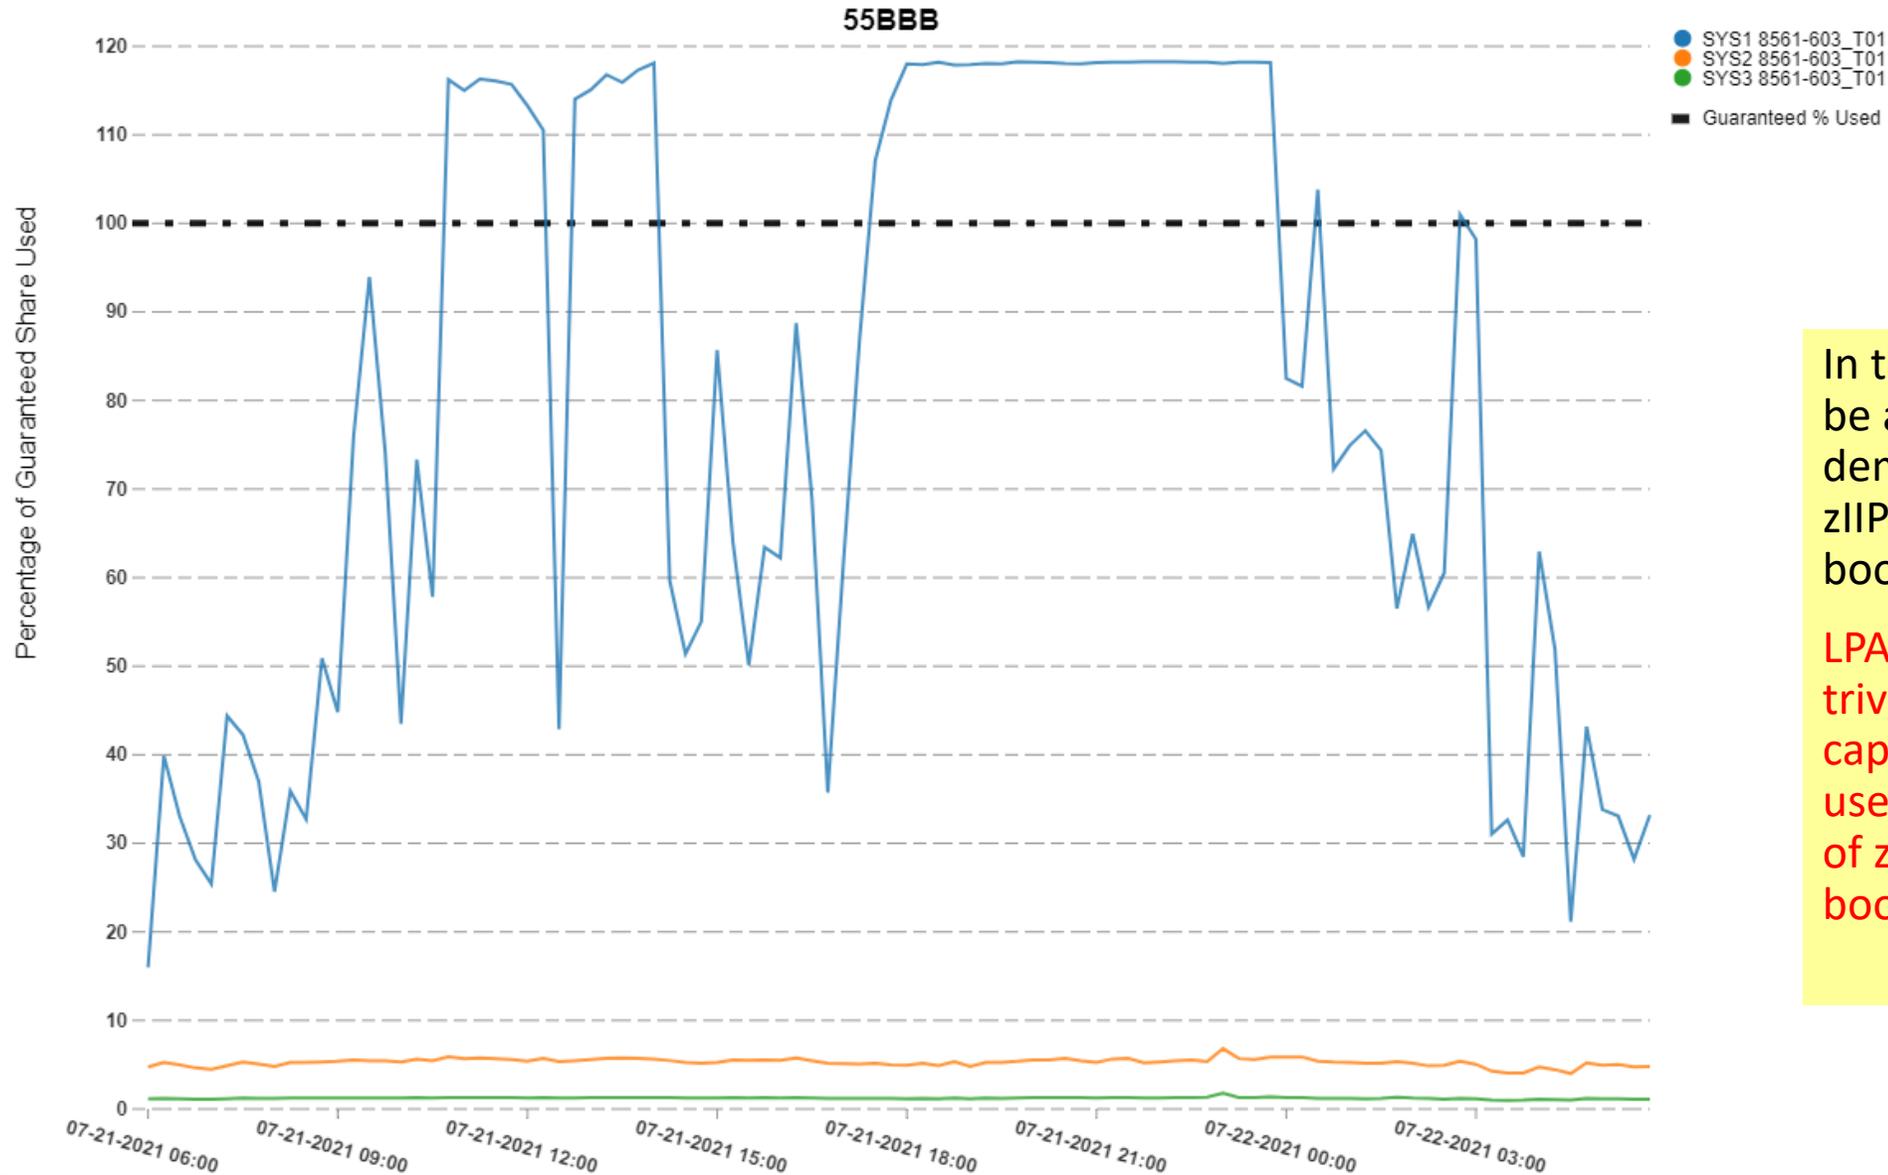


Here we see SYSC using more zIIP because some of that GP work ran on the zIIPs.

Note this is a small system the relative change might not be as dramatic on larger systems. Or it might be larger if more zIIPs were enabled relative to the GPs and there was more work to do.

In other words: YMMV

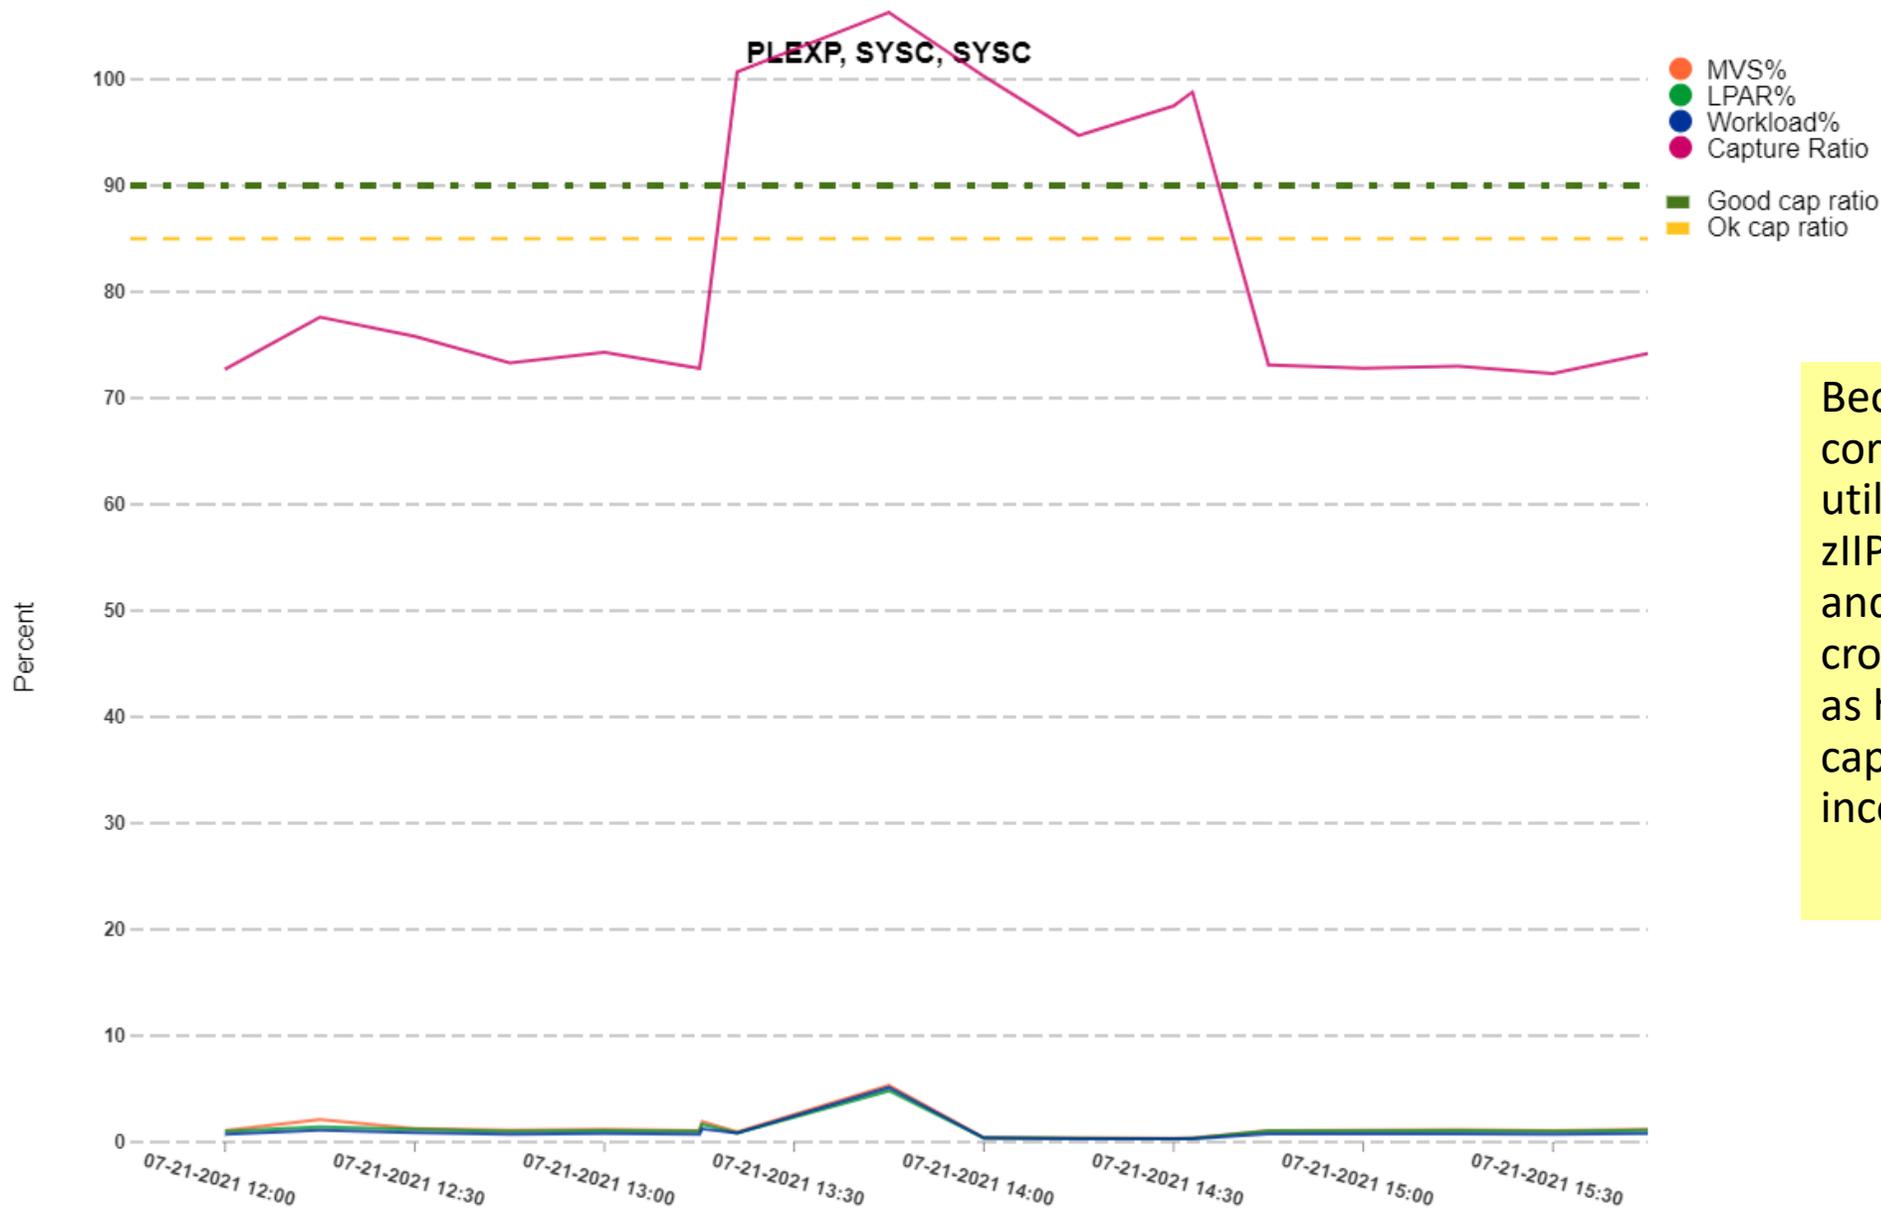
# CEC Percent zIIP Weight Used



In this case SYS1 might be at risk if SYS2 or SYS3 demanded their entire zIIP weight due to a boost.

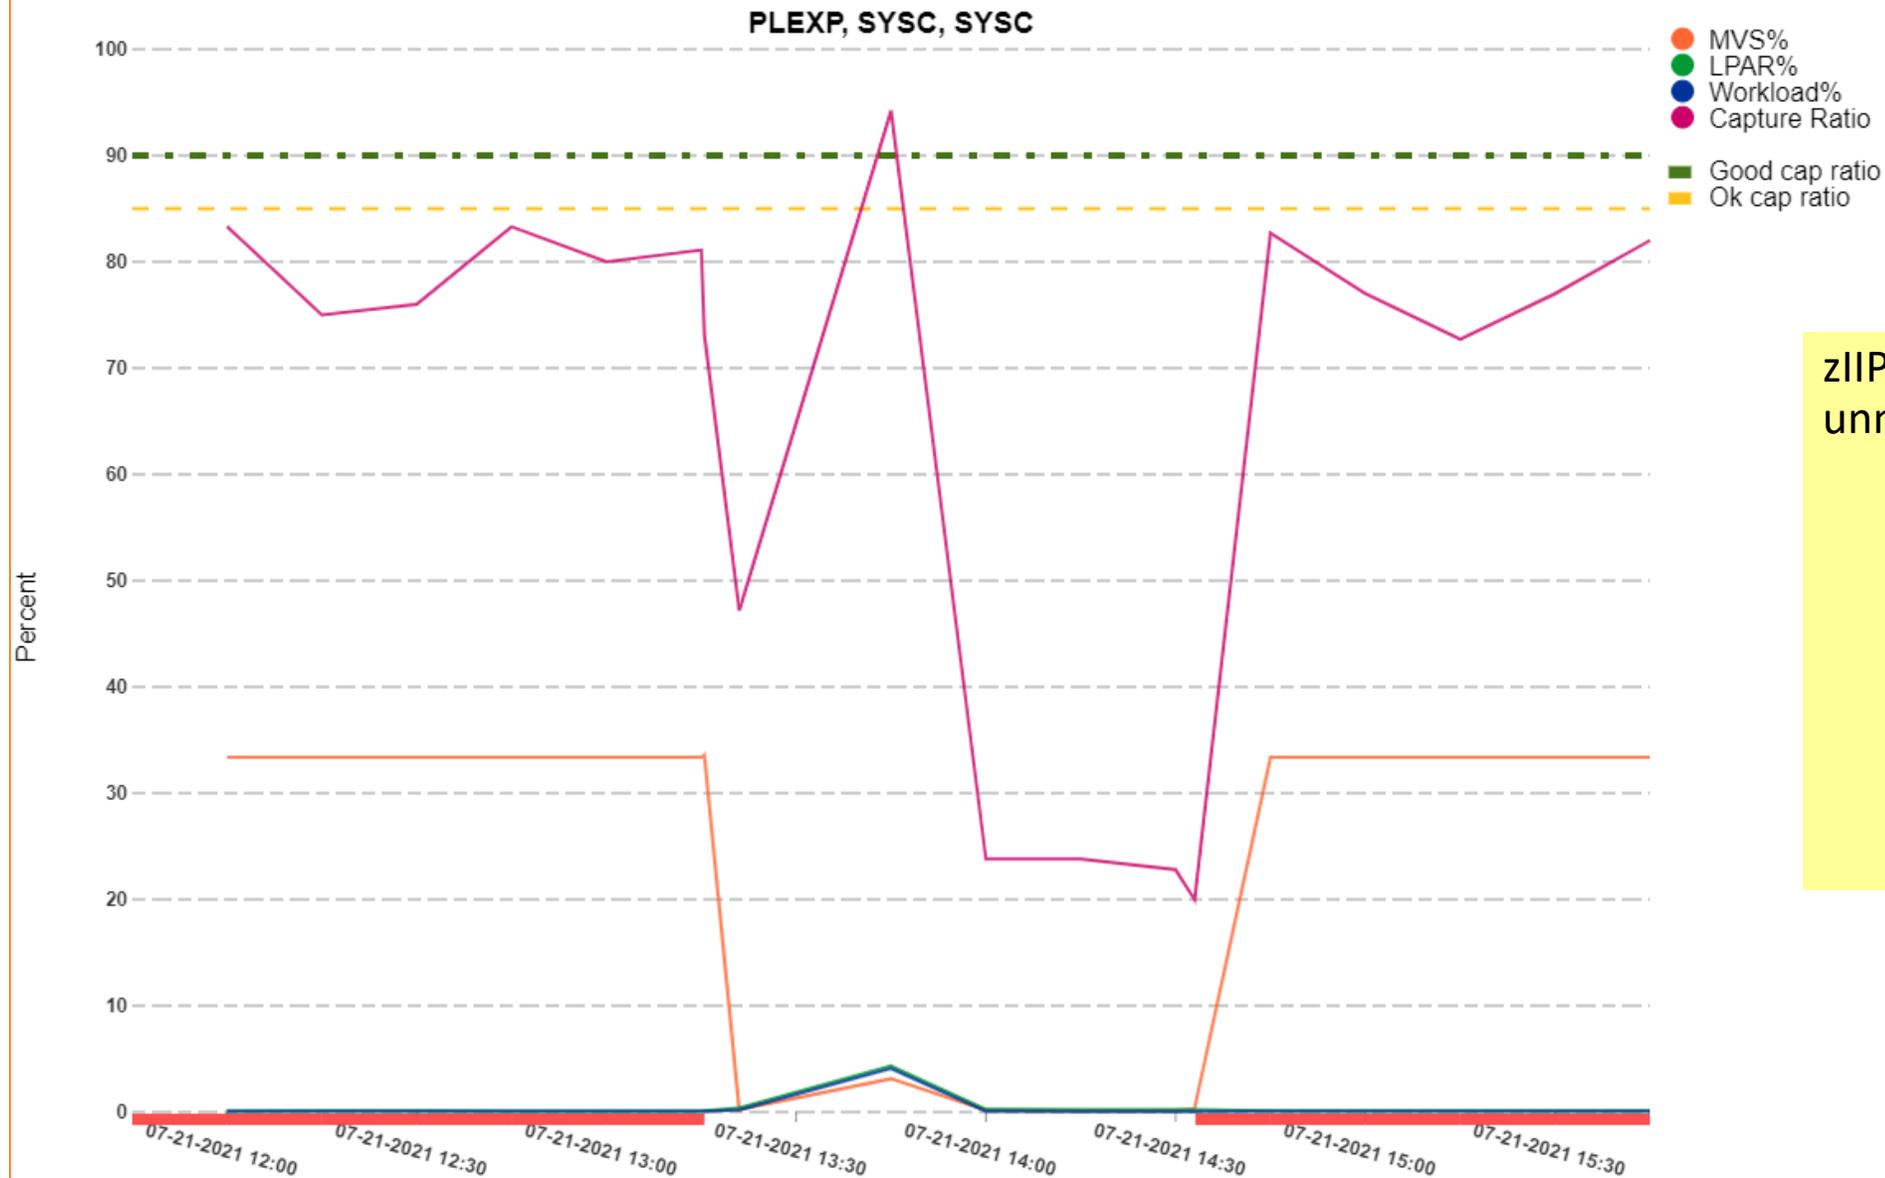
LPARs that normally use trivial amounts of zIIP capacity may suddenly use substantial amounts of zIIP capacity during boost periods!

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



Because SMF 70 correctly records the utilizations of the GPs vs. zIIPs, but the SMF 72 and SMF 70 records crossover from GP to zIIP as having been on GP, capture ratios will be incorrect.

# LPAR, MVS, and Workload zIIP Busy% with Capture Ratio



zIIP capture ratio goes unnaturally low!



# Parting thoughts

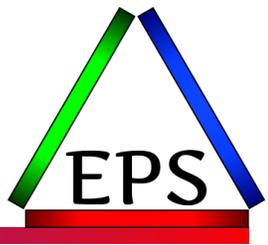
# SRB: The Good



- Overall, this is goodness: faster shutdown, startup, and recovery
- Sub-capacity systems will see more benefit than full-cap machines
  - It's not hard to come up with scenarios where a boosted LPAR might have access to 5-10x more capacity than normal for GP-only work!
- No additional cost
- Simple scenarios where IPLs happen at low-utilization times will see benefit while likely not causing any significant concerns
- Despite design goal of no software cost impact, seems like it could potentially reduce costs slightly

Boost baby, boost!

# SRB: The Bad



- Have to update your shutdown procedures to invoke shutdown boost
- LPAR weights may require attention to make sure they're appropriate during boost periods
  - May need to set up automation to change weights when boosted
- Increased zIIP utilization *might* cause more zIIP delay on non-boosted LPARs
  - Depending on weights, utilization, etc.
- Boosted LPARs might impact cache efficiency of non-boosted LPARs
  - Only in theory; in practice, haven't seen any evidence of that

# SRB: The Ugly



- Recorded CPU time during boost periods is... confusing to misleading
  - SMF 30 intervals that cross boost/non-boost intervals can be especially so
    - So sync your intervals!  
Use `SUBSYS(xxx,INTERVAL(SMF,SYNC))` in your `SMFPRMxx!`
    - Have seen at least 1 case where it appears this may have impacted CMF/RMF intervals
      - Or something else may be at play there
- 2 Minute boosts could be useful in some situations, but may cause confusion when doing performance analysis
  - Performance testing especially impacted!
- Boosted performance might make your non-boosted periods seem lackluster

# Recommended To-Do List



- Add call to proc IEASDBS during shutdown to enable boost
  - Still seeing customers with IPL boost but no shutdown boost
  - If using PRESCPU in IEASYSxx, stop boost before final shutdown
- Add reserved zIIPs as appropriate to LPARs to take advantage of zIIP Boost
- Verify that zIIP weights are appropriate during boost periods
- Understand when Recovery Boost periods happen

# Not necessarily recommended, but...



- What if...
  - You have a small sub-capacity machine that gets a significant CPU boost
  - You have some free zIIP capacity too
  - You have a significant CPU-limited month-end process
- Should you IPL immediately before month-end processing starts?
  - I.E. could boost get you multiple hours of work done in < 1 hour?
- Or if on a z16... could you create a dummy started task that you could start to get a 5 minute boost
  - Are limited to just 30 minutes of this sort of boost per day, but if you're getting access to 5-10x more capacity ... might be worth thinking about
  - Would need to have work queued that could take advantage of that boost
  - My understanding is that current IBM T's and C's don't ban this



# Thanks!

Questions/comments:  
[Scott.Chapman@epstrategies.com](mailto:Scott.Chapman@epstrategies.com)