



# Pivotor: Finding and Analyzing Outliers

Peter Enrico  
Email: [Peter.Enrico@EPStrategies.com](mailto:Peter.Enrico@EPStrategies.com)

Scott Chapman  
Email: [Scott.Chapman@EPStrategies.com](mailto:Scott.Chapman@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



z/OS Performance  
Education, Software, and  
Managed Service Providers



Creators of Pivotor®

# 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

# Performance Analysts' Dreams



- Computer performance analysts have long sought out methods to detect when values are unusual and so should be flagged for investigation
- Historically these have relied on various forms of statistical analysis
  - Variance (often  $n$  standard deviations) from the mean
  - Statistical Process Control
  - Multivariate Adaptive Statistical Filtering
  - Analysis of Variance
  - Etc.
- All have some value, but also significant limitations
  - Most require significant (manual) up-front analysis
  - Many are prone to false positives
  - Multimodal data can be especially difficult for some of these

Several good CMG papers about these methods over the years especially c. 1995-2005.

# Yet still...



- It would be great if we knew when something starts using more (or less?) of some resource
  - But we have lots of resource measurements
  - And we have lots of “something”s (batch jobs, system, service classes, etc.)
  - And false positives should be minimized
  - And we should ignore “trivial” changes
- Additionally, from a product / service provider perspective...
  - We can't spend time manually doing analysis
  - It has to work across a variety of very different customer environments
  - It can't be overly costly to run, despite needing to analyze a lot of data
  - And should self-correct as usage changes over time



# Introducing... Pivotor Outlier Reporting!

# Naming things is hard



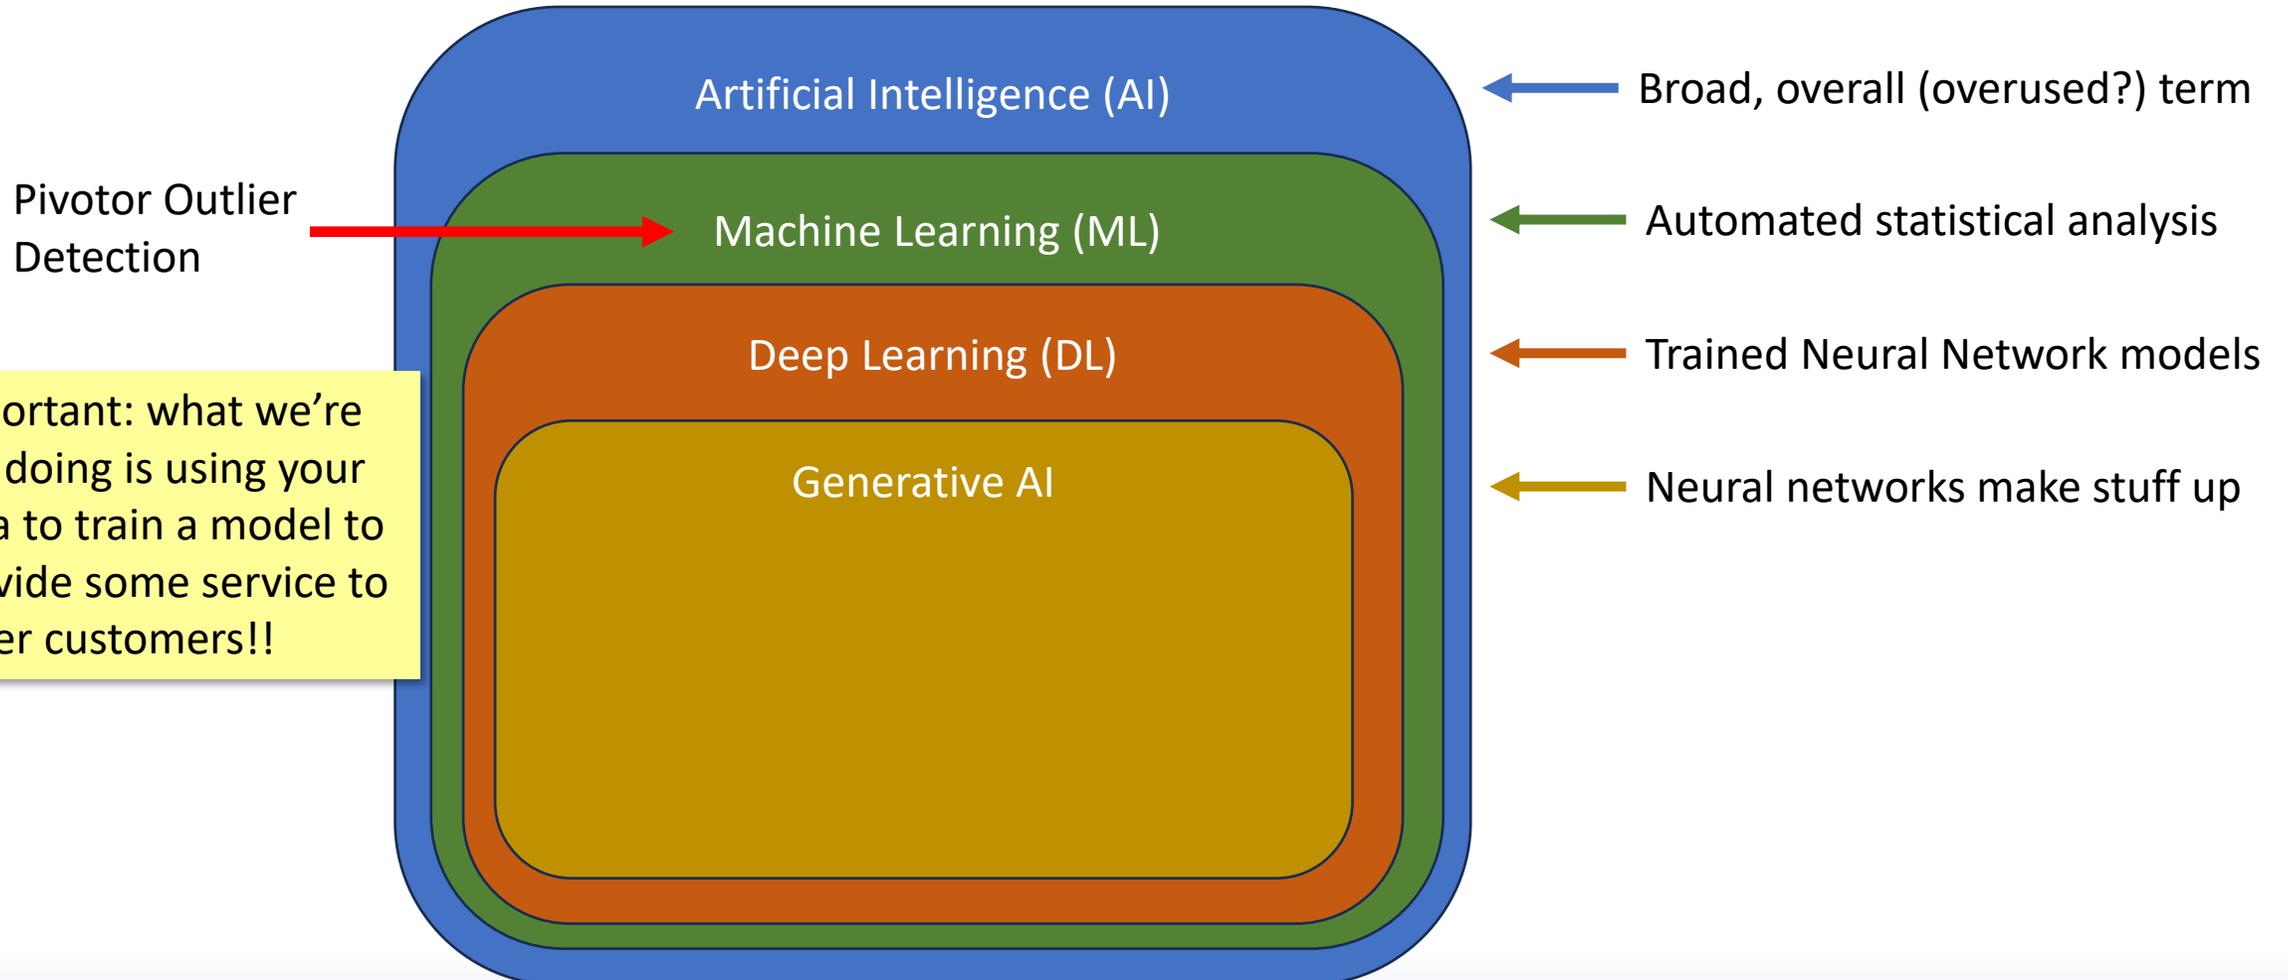
- We decided to name these “outliers” not “anomalies”
- “Anomalies” felt more alarming to us
- An outlier may or may not be cause for concern
  - It’s just different than what the recent past has been
  - May be worth investigating
  - Not necessarily reason to panic
- A focus on just anomalies that are known to be problematic would likely ignore values that are simply unusual but may be an important clue

# What is AI?



- This seems to be a consensus taxonomy:

My simplistic view



# What we *have* done...



- We've combined a couple of Machine Learning techniques in a way that we think works pretty well for determining outliers
  - At their core, these ML techniques are just fancier statistical methods
  - This is *not* generative AI!
  - This is a specific analysis for each customer environment
- Basically, we're clustering historical observations together and then looking for observations that don't fall in those clusters
  - Combination of *k*-means and density-based clustering
- Analysis based on dimensions of:
  - Metric value
  - Day of week
  - Time of day
- Analysis done by specific entities (systems, batch jobs, service classes, etc.)

# Two processes



- Once a week: find the norms
  - Based on last several (generally ~8) weeks of data
  - Sets bounding boxes for each specific entity and metric that encompasses the clusters of values that appear to be normal
  - Can take quite some time (hours), depending on the size of the environment
    - Generally no need to run more than once a week
    - Even once a week might be overkill for many situations
    - Expectation is weekly runs will allow for norms to evolve over time
- Daily: find the outliers
  - Take current (“today”) data and compare to bounding boxes
  - Extract current observations that are outside the bounding boxes
  - Report on those outliers, along with some contextual data
  - Relatively fast (minutes)

# Limitations



- This is for metrics that are not necessarily “right” or “wrong”
- Inherent assumption is that the past is “normal”
  - Holidays will likely become outliers on their own
- False positives are reduced, but not necessarily eliminated
  - Did you expect the machines to learn perfectly? 😊
- At least for now, not available for 1 minute interval data
  - Plexes with 1 minute RMF/SMF intervals excluded
  - Transactional data (CICS, DB2, IMS, etc.) summarized to 1 minute data also excluded
    - Will likely try with the hourly data to evaluate usefulness there
  - In some cases, certain high-frequency entities excluded

# Potential use cases



- We had an incident yesterday around noon.  
What measurements were unusual around that time?
- The batch window ran long last night.  
What batch jobs ran longer than usual?
- CPU consumption was up yesterday.  
What service classes consumed more CPU than usual?
- I'm interested in our common storage utilization.  
Has the utilization of any of those areas changed recently?
- I'm interested in...

Outlier Analysis is inherently a more advanced analysis because you likely need a holistic understanding of the data and your systems.

**We expect you'll come up with more!**  
(Let us know when you do)

# A note about metrics



- We're capturing norms for dozens of metrics
  - Some of the metrics will only rarely have outliers
  - Some metrics (e.g. batch job execution time) may vary quite frequently
  - Some metrics may be stable for a long time, then take a step up/down
    - At which point it may be flagged as an outlier for a while until that becomes the new norm
- But if a metric doesn't have outliers, it won't appear in the report
  - Same for entities for metrics
- Don't expect zero outliers!
  - Even for small environments we're looking at millions of observations
  - Things do change, whether expected or not



# Outlier Reporting Details

# Outlier Reportset on the Calendar



- Runs daily, during long-term reporting (after daily reporting)
- Is a week-to-date report
  - So shows up under the weekly reports section of the calendar
  - And actually contains data for up to 5 weeks
- One reportset for all outliers for a reporting plex

ExamOwl Data Co.

February 2024

Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Weekly Reports
				01 <a href="#">HOURLPLEX</a> <a href="#">PRODPLEX</a> <a href="#">TESTPLEX</a>	02 <a href="#">HOURLPLEX</a> <a href="#">PRODPLEX</a> <a href="#">TESTPLEX</a>	03 <a href="#">HOURLPLEX</a> <a href="#">PRODPLEX</a> <a href="#">TESTPLEX</a>	<a href="#">HOURLPLEX</a> <a href="#">PRODPLEX</a> <a href="#">TESTPLEX</a>
04 <a href="#">HOURLPLEX</a> <a href="#">PRODPLEX</a> <a href="#">TESTPLEX</a>	05	06	07	08	09		<a href="#">PRODPLEX</a> <a href="#">TESTPLEX</a>
11	12	13	14	15	16	17	

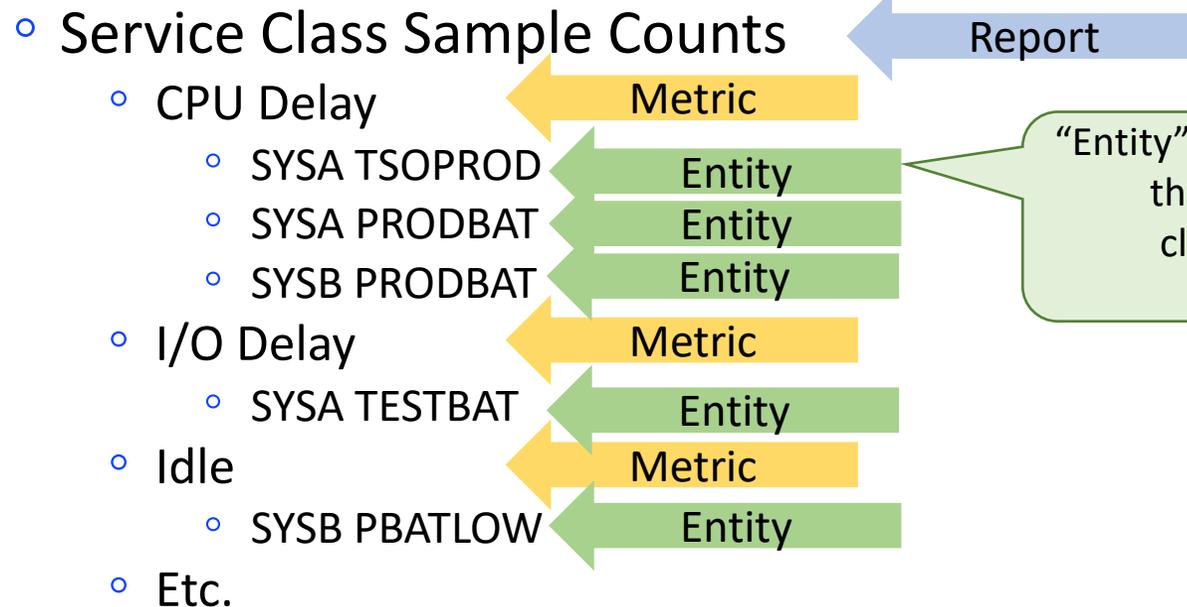
PRODPLEX Week ending 2024-02-10 x

- [Outliers](#)

# Outlier Report Taxonomy

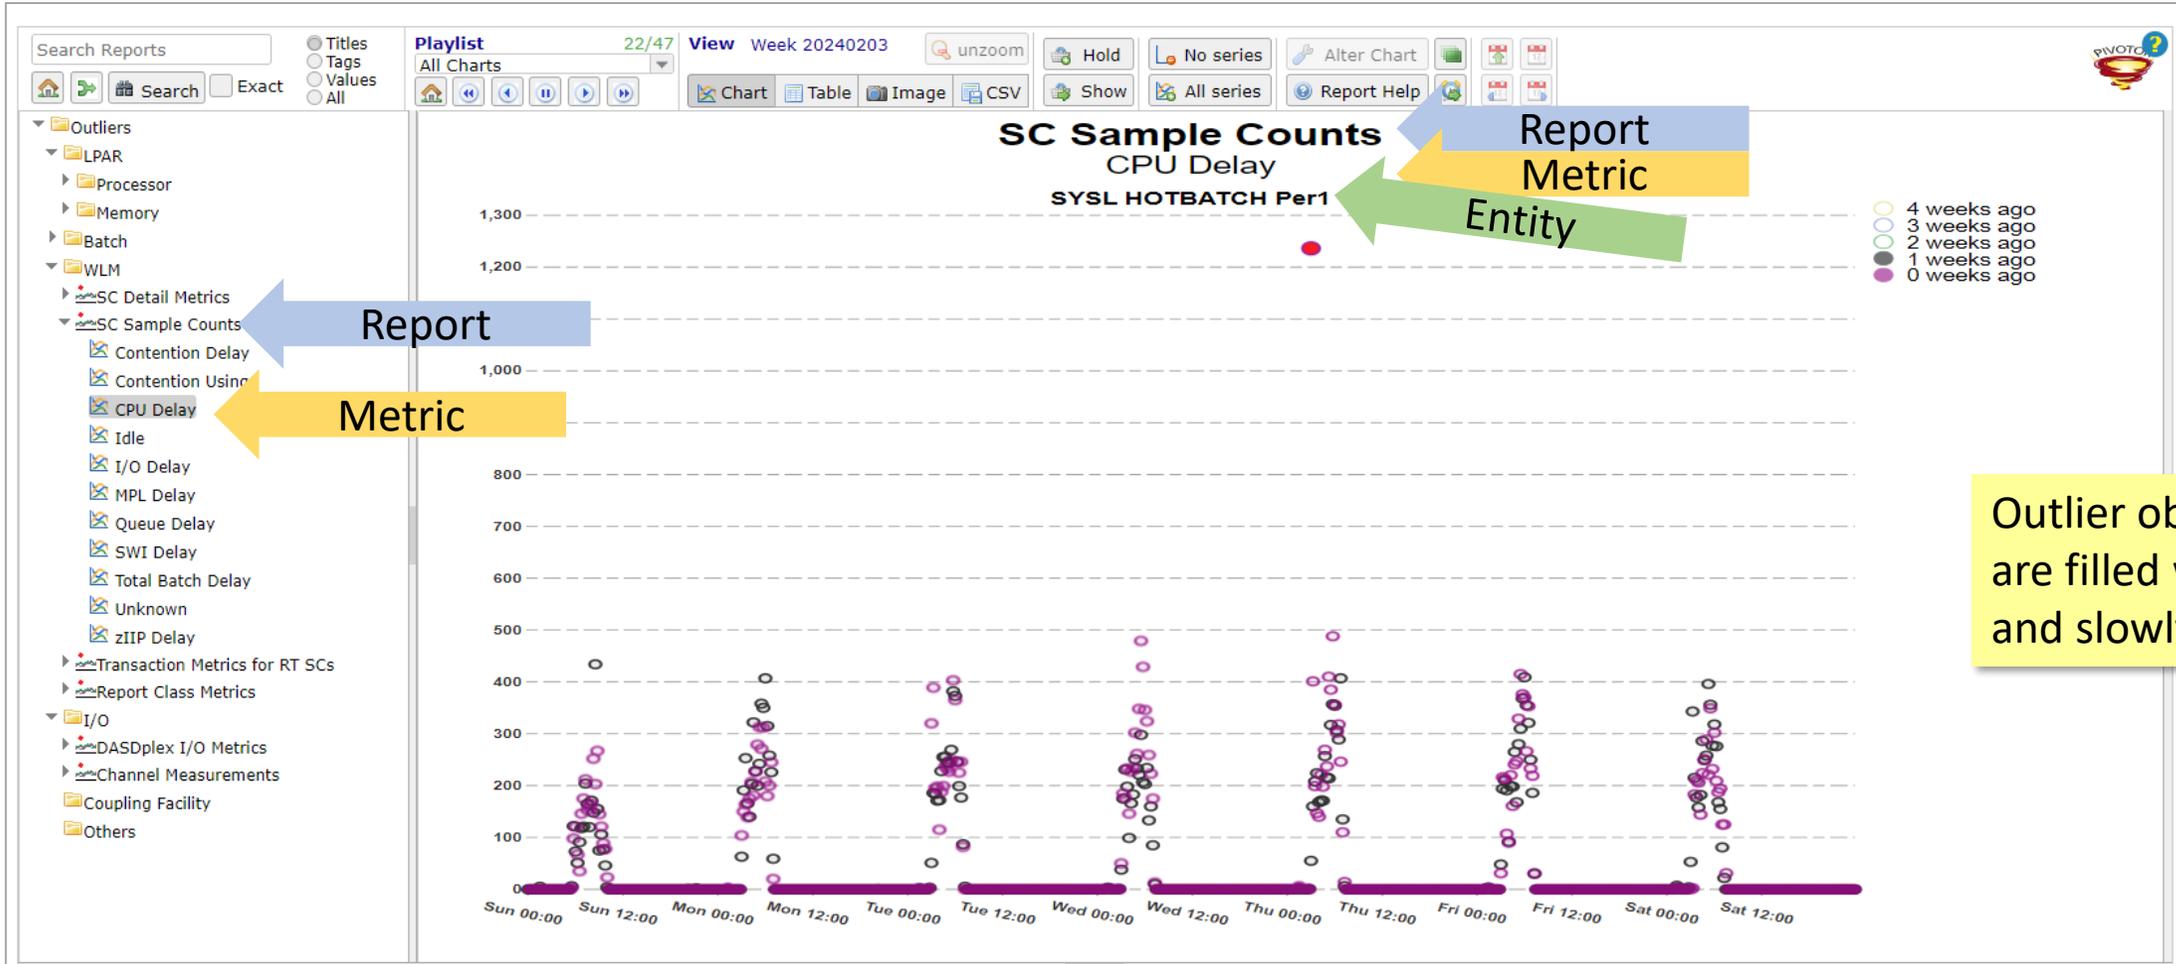


- A single outlier report is for a group of related metrics
- Each metric then may have outliers for multiple entities
  - Entities are specific combinations or groupings such as service class PRODBAT Period 2 on SYSA
- For example:



“Entity” may not be the best name for the specific combination of classification values, but... naming things is hard!

Note that it’s quite possible (even probable) that there may be multiple individual outlier observations for each entity.



Outlier observations are filled with red and slowly pulse.

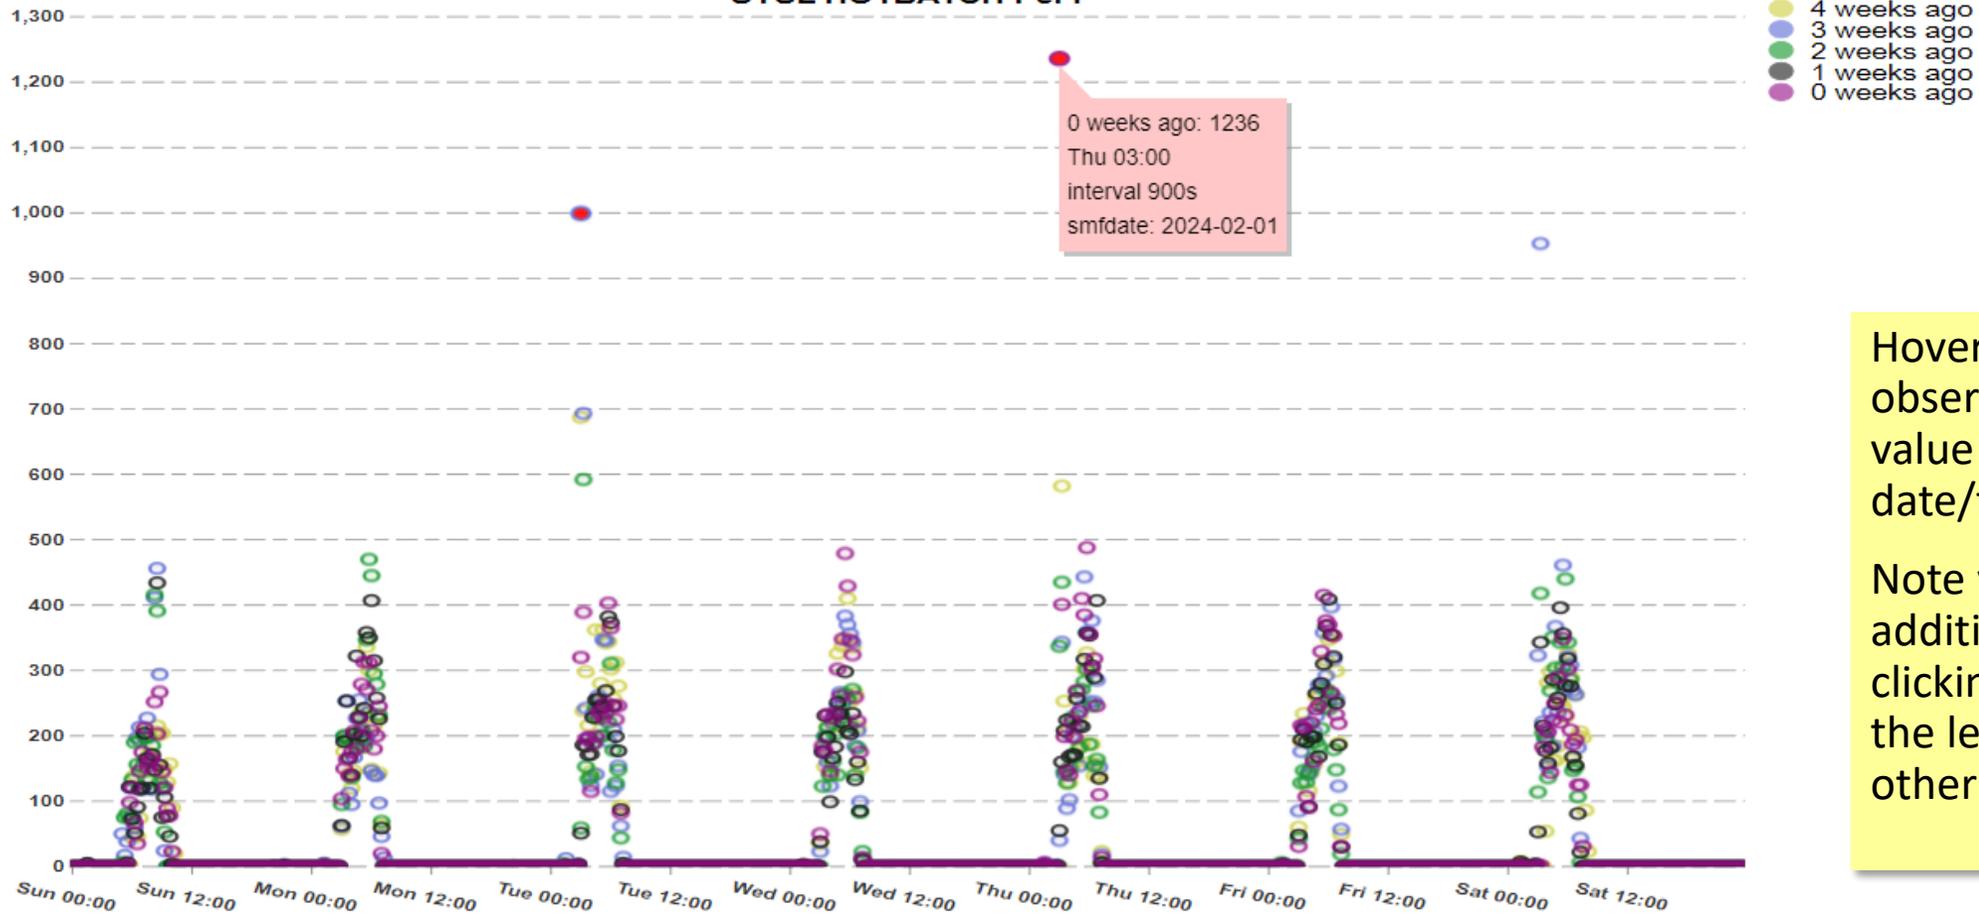
SYSTEM ▲	SC_NAME ▲	SC_PERIOD ▲	Week 0 Count ▲	Score ▼
SYSL	HOTBATCH	Per1	1	39.87
SYSK	BATCHLO	Per2	1	36.4
SYSL	SAPBW	Per1	1	28.42
SYSL	CTCUT	Per1	1	28.41

Pro tip: Once you've clicked in the entity table you can navigate through the entities with the up/down arrow keys instead of clicking.

# SC Sample Counts

## CPU Delay

### SYSL HOTBATCH Per1



Hover over the observation to get the value as well as the date/time.

Note you can turn on additional history by clicking on the circles in the legend. (Just like any other Pivotor report.)

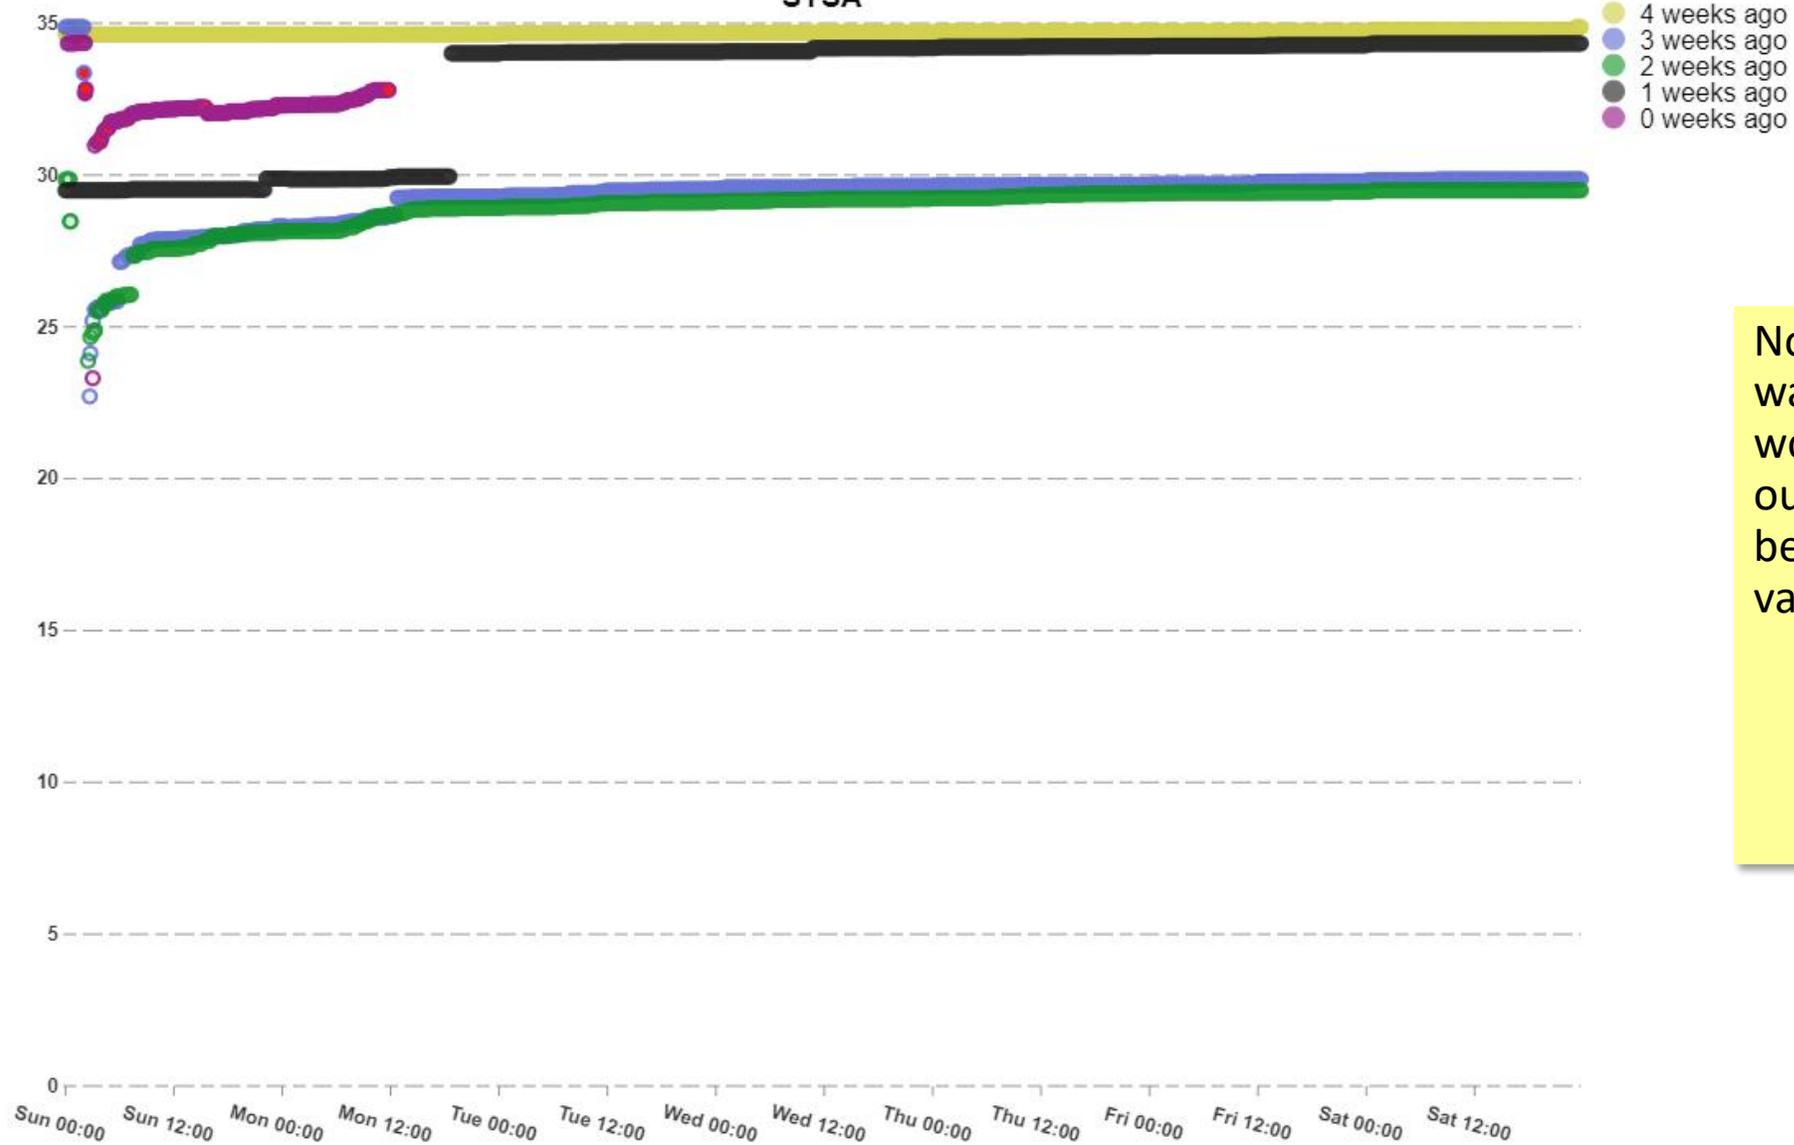
Data for all weeks are shown overlapping on the x-axis that represents Sunday through Saturday.

Whether a historical observation was an outlier depends on the state of the norms at that point in time when the data was processed.

# LPAR Storage Outliers

Max LPA MB

SYSA

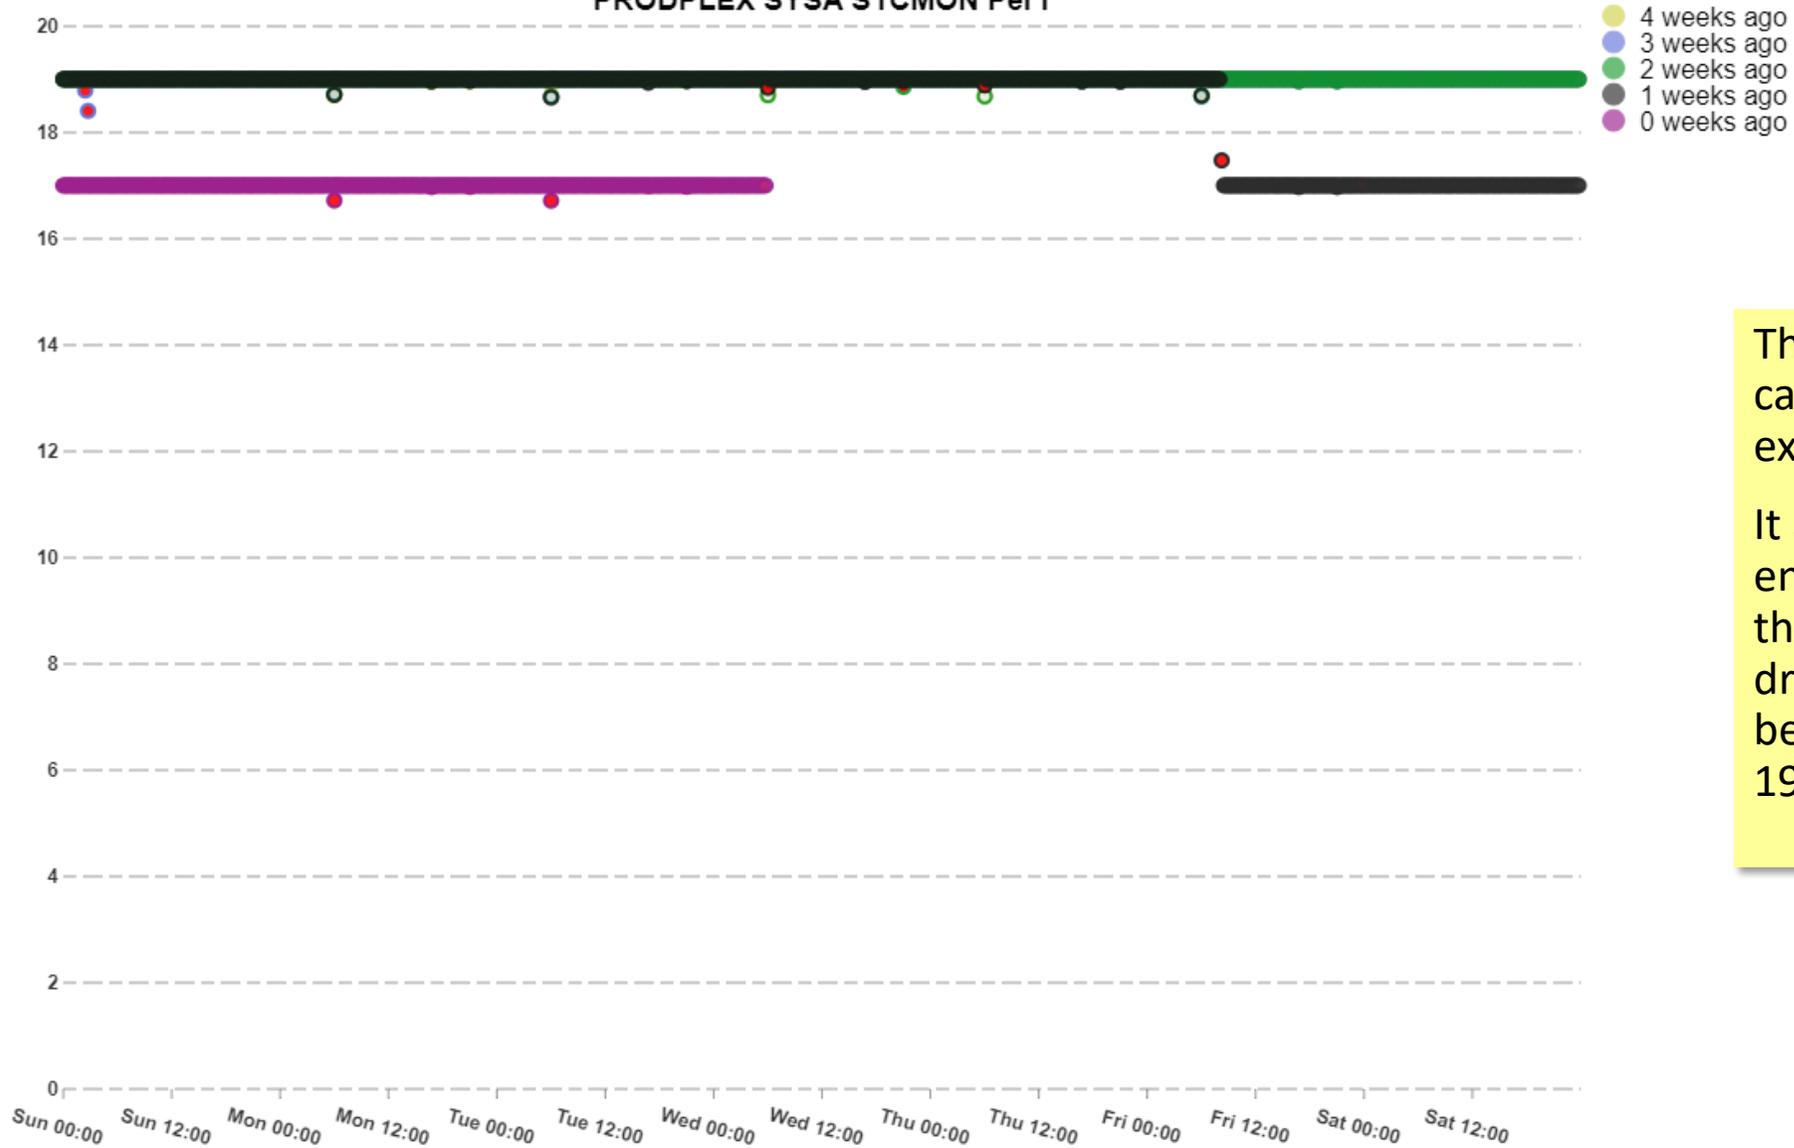


Note that because of the way the ML algorithms work, we can find outliers that might be between common values.

# Service Class Detail Metrics

Multi-Programming Level

PRODPLEX SYSA STCMON Per1



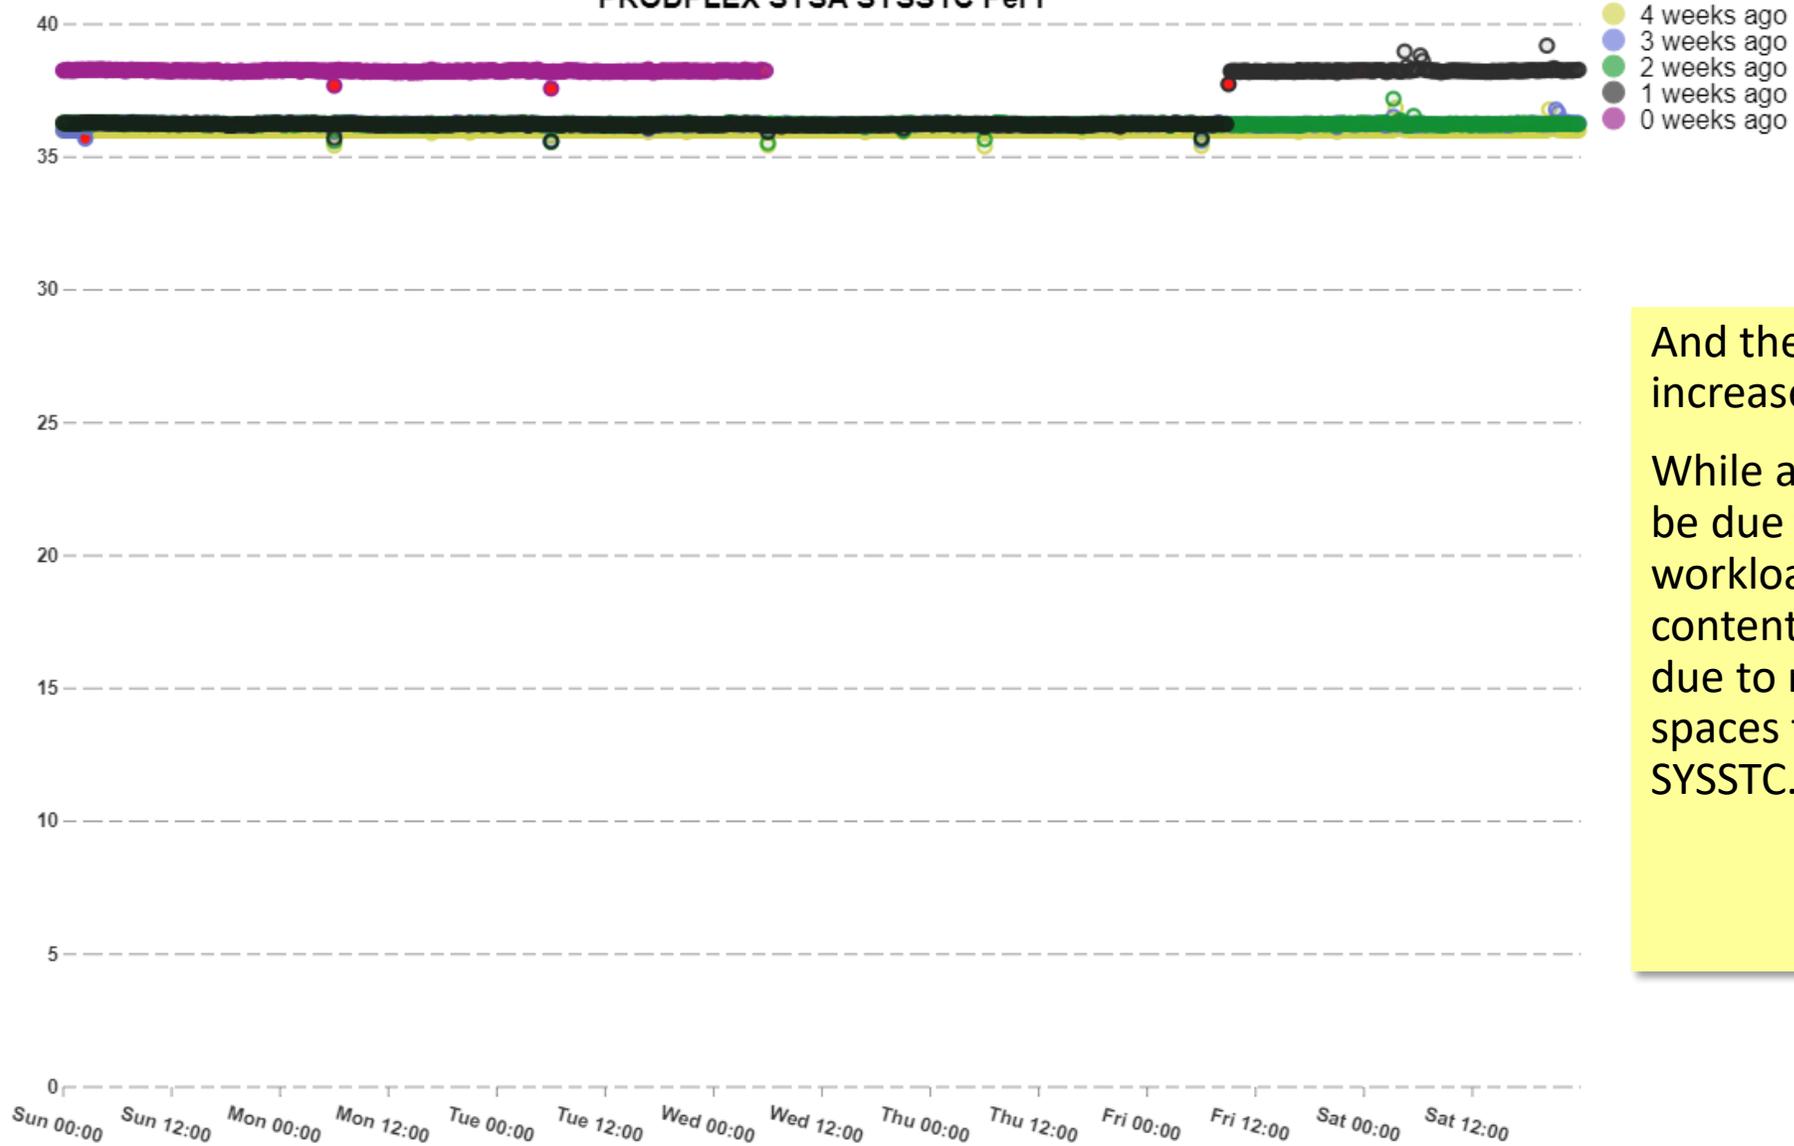
This is an interesting case of noticing an expected change.

It appears that at the end of the prior week, the MPL for STCMON dropped by 2, after being very consistently 19 for weeks.

# Service Class Detail Metrics

Multi-Programming Level

PRODPLEX SYSA SYSSTC Per1



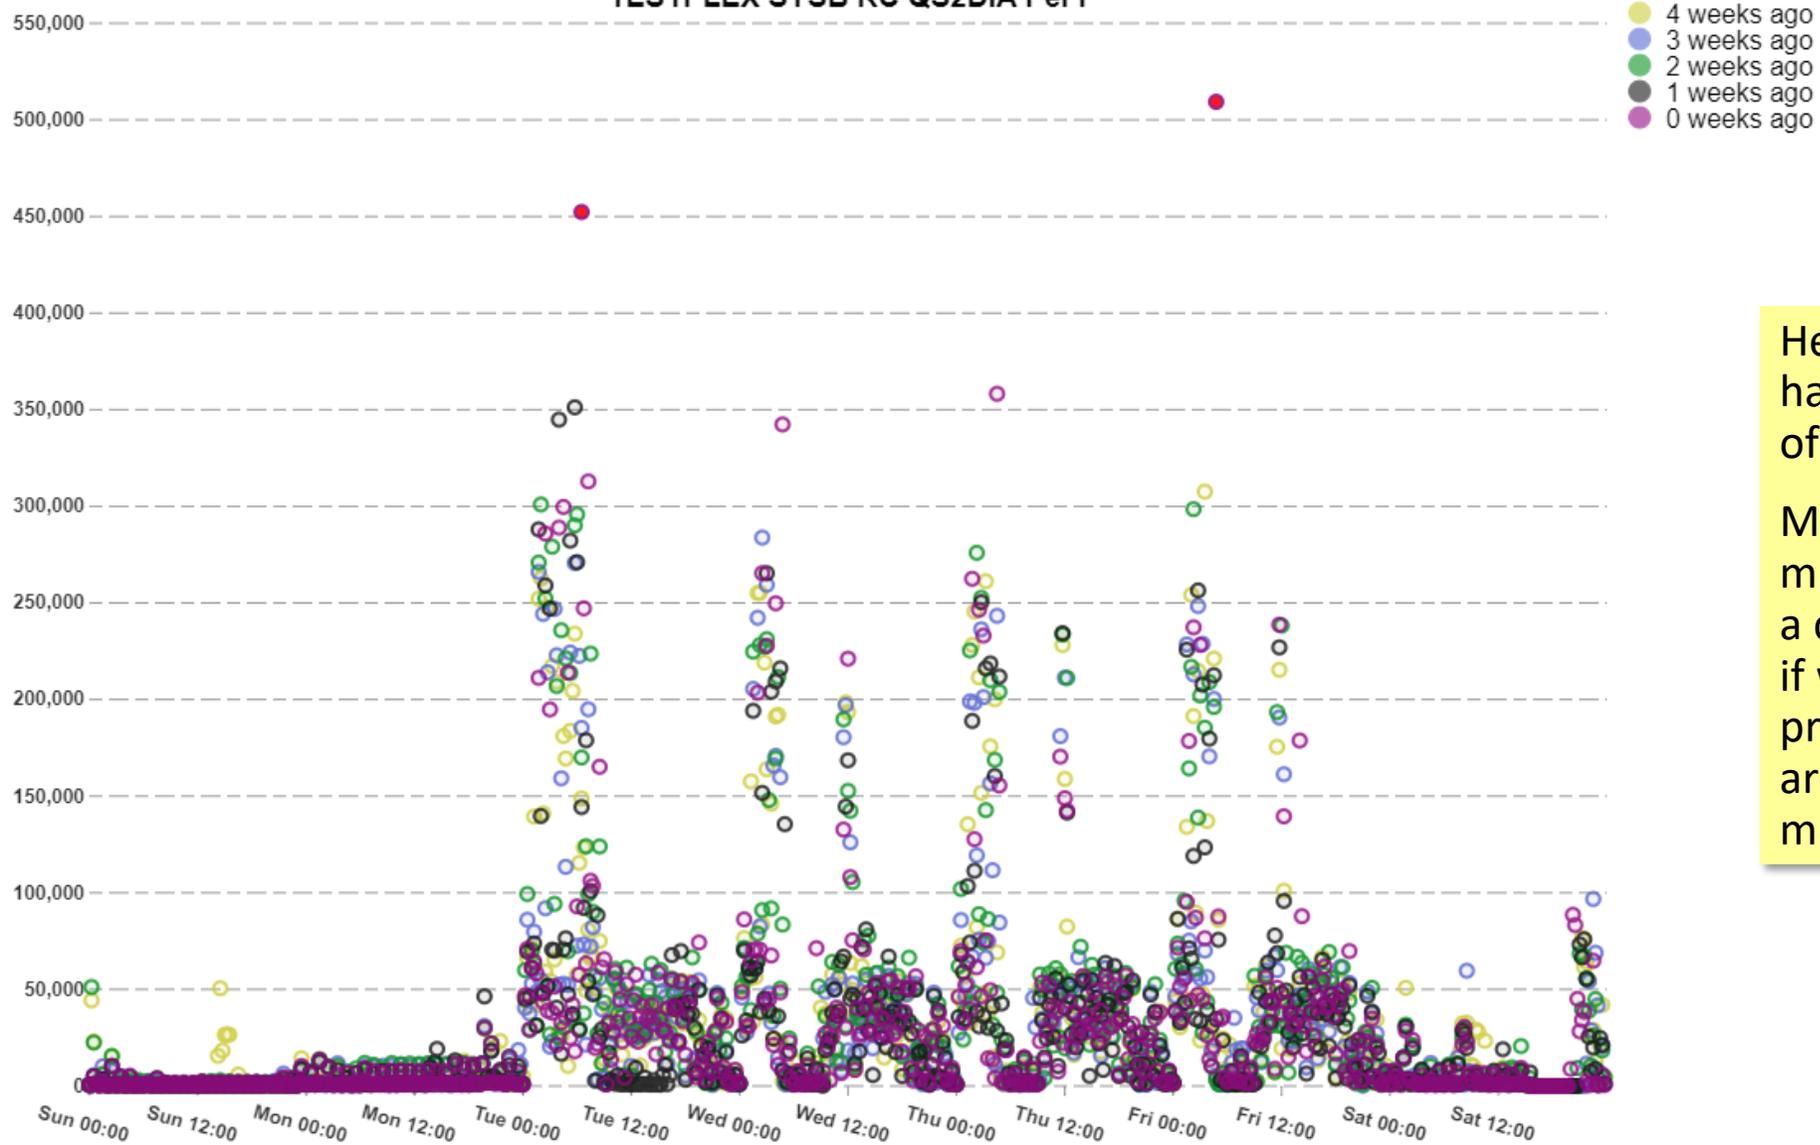
And the MPL of SYSSTC increased by 2.

While a change in MPL could be due to a change in the workload or some storage contention, in this case it was due to reclassifying 2 address spaces from STCMON to SYSSTC.

# Report Class Overview Metrics

SSCH Count

TESTPLEX SYSB RC QS2DIA Per1



Here's a report class that had a couple of intervals of unusually high I/O.

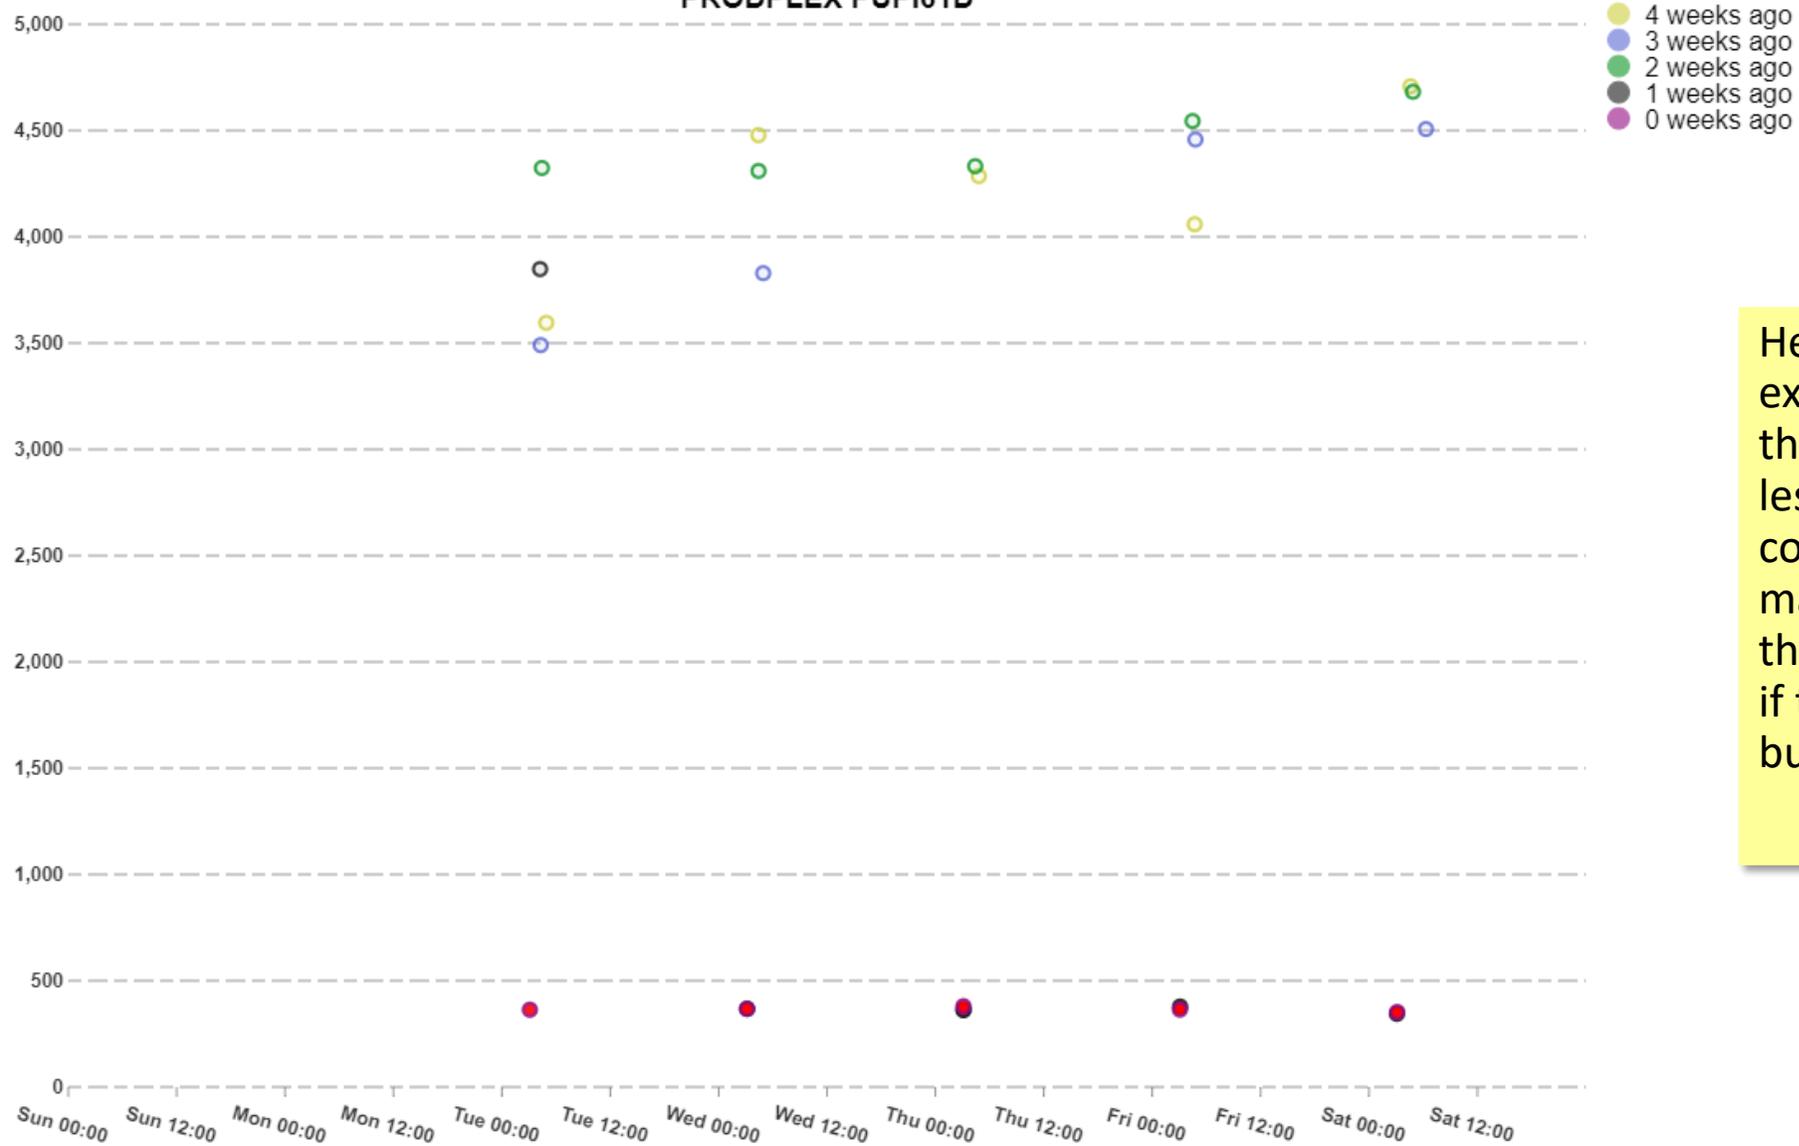
Maybe we don't care too much because it was just a couple of intervals, but if we were looking into a problem that occurred around those times we might!



# Batch Job Outliers

Elapsed time

PRODPLEX PUPI61D



Here's an interesting example of a batch job that is consuming much less CPU in the last couple of weeks. So maybe that's a good thing. (Or maybe it's not if that means there's less business!)

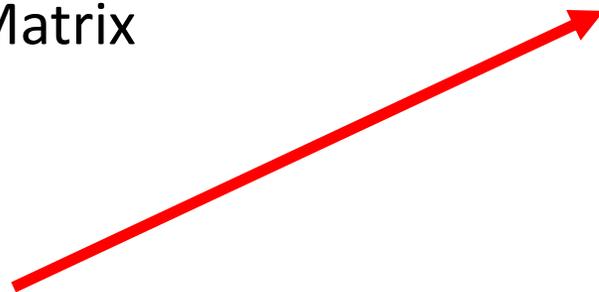


# Outliers Overview Panel

# Other ways to navigate



- Navigating via the report hierarchy makes sense if you are interested in specific metrics
- The new “Outliers Overview” panel provides alternative navigation
  - Idea is to provide useful summarization and links for different analyses
  - Centered around looking at a specific day
- Three sections:
  - Outlier Metrics / Entities Matrix
  - Outliers by Time Interval
  - Outliers Score by Metric
- Start by selecting the day



# Outlier Metrics / Entities Matrix



- This section provides an overview of what entities had outliers for what metrics
- Change the entity type with the drop-down
- Note that the entity and metrics might be shown on either the row or column
  - Which has the most will end up as the rows
  - Scroll bars may appear if you have a lot of rows (common e.g. when showing batch jobs)
- Click on the asterisk to be taken to the report for that metric and also search then for that specific entity

## Outlier Metrics / Entities Matrix

SYSTEM			
	SYSTEM		
	SYSK	SYSL	
▼ Report Class Overview Metrics (4 items)			
CPU Seconds	*	*	
DASD IO RT	*	*	
SSCH Count	*	*	
zIIP on CP Seconds		*	
▼ Service Class Detail Metrics (1 item)			
Avg. Memory MB		*	
▼ Service Class Sample Counts (3 items)			
CPU Delay		*	
Unknown		*	
zIIP Delay		*	
▼ Transaction Metrics for RT SCs (1 item)			
Response Time (secs)	*		



Search Reports  
Exact

- Titles
- Tags
- Values
- All

Playlist 28/35 View Week 20240210

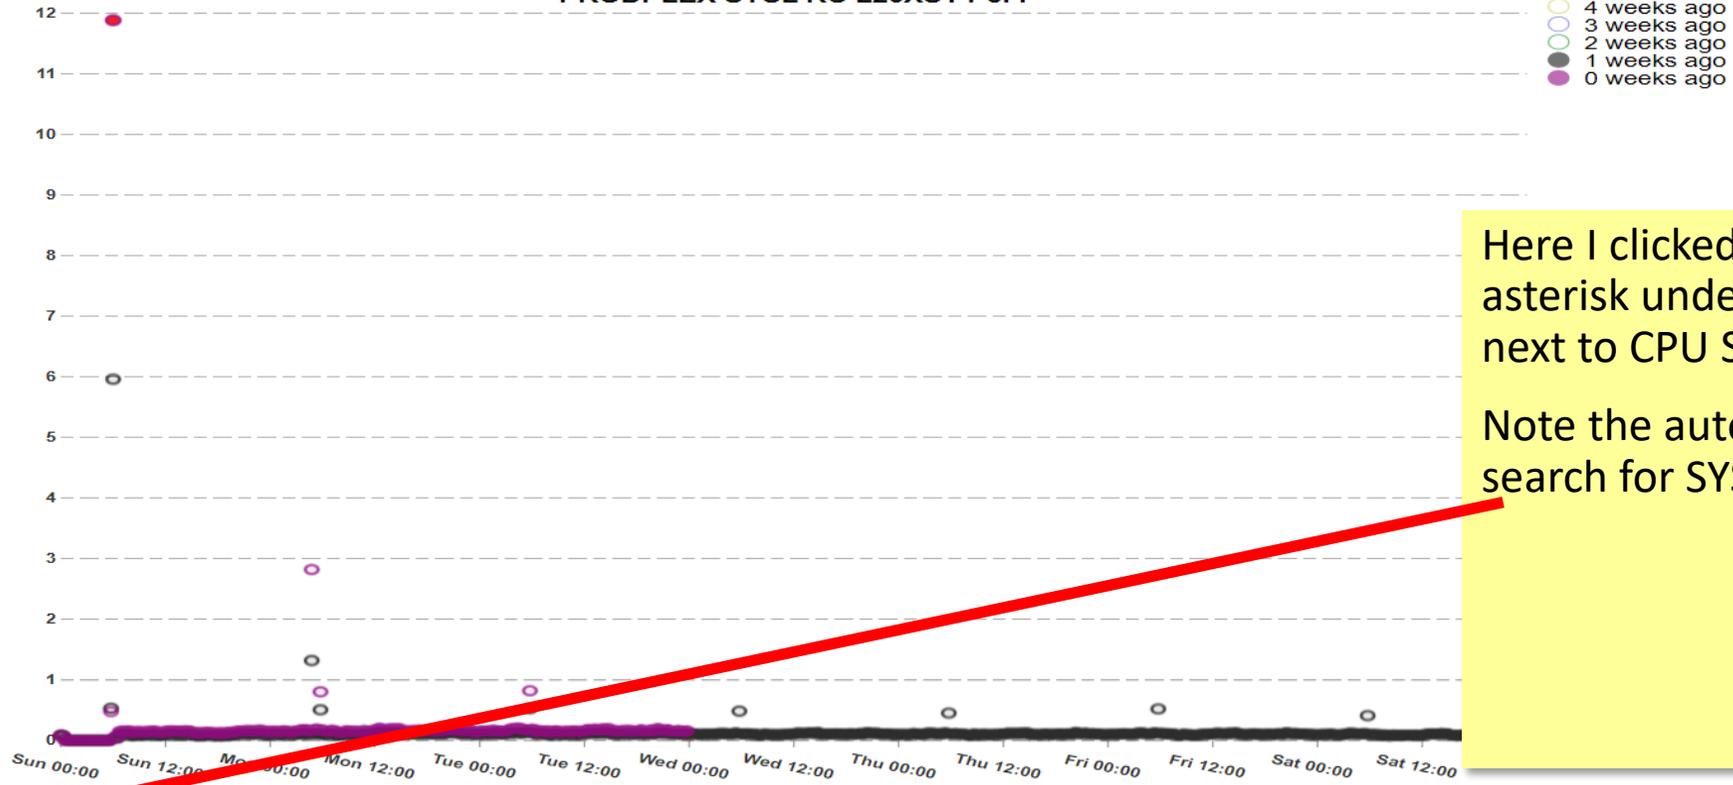
unzoom Hold No series Alter Chart  
Chart Table Image CSV Show All series Report Help

- Outliers
  - LPAR
  - Batch
  - WLM
    - Service Class Detail Metrics
    - Service Class Sample Counts
    - Transaction Metrics for RT SCs
    - Report Class Overview Metrics
      - DASD IO RT
      - SSCH Count
      - CPU Seconds
      - zIIP on CP Seconds
      - zIIP Seconds
  - I/O
    - Coupling Facility
    - Others

# Report Class Overview Metrics

## CPU Seconds

### PRODPLEX SYSL RC L26XST Per1



Here I clicked on the asterisk under SYSL and next to CPU Seconds.

Note the automatic search for SYSL outliers.

SYSplex	System	Class Type	SC Name	SC Period	Week 0 Count	Score
PRODPLEX	SYSL	RC	L26XST	Per1	1	79.21
PRODPLEX	SYSL	RC	SAPRPT	Per1	1	44.34
PRODPLEX	SYSL	RC	PR1DIA	Per1	1	33.62
PRODPLEX	SYSL	RC	PR1DB2	Per1	1	30.06

# Outliers by Time Interval



- The idea for this section is to get an overview of when outliers might be clustered in a particular timeframe
- Or to find what outliers might have happened around a time

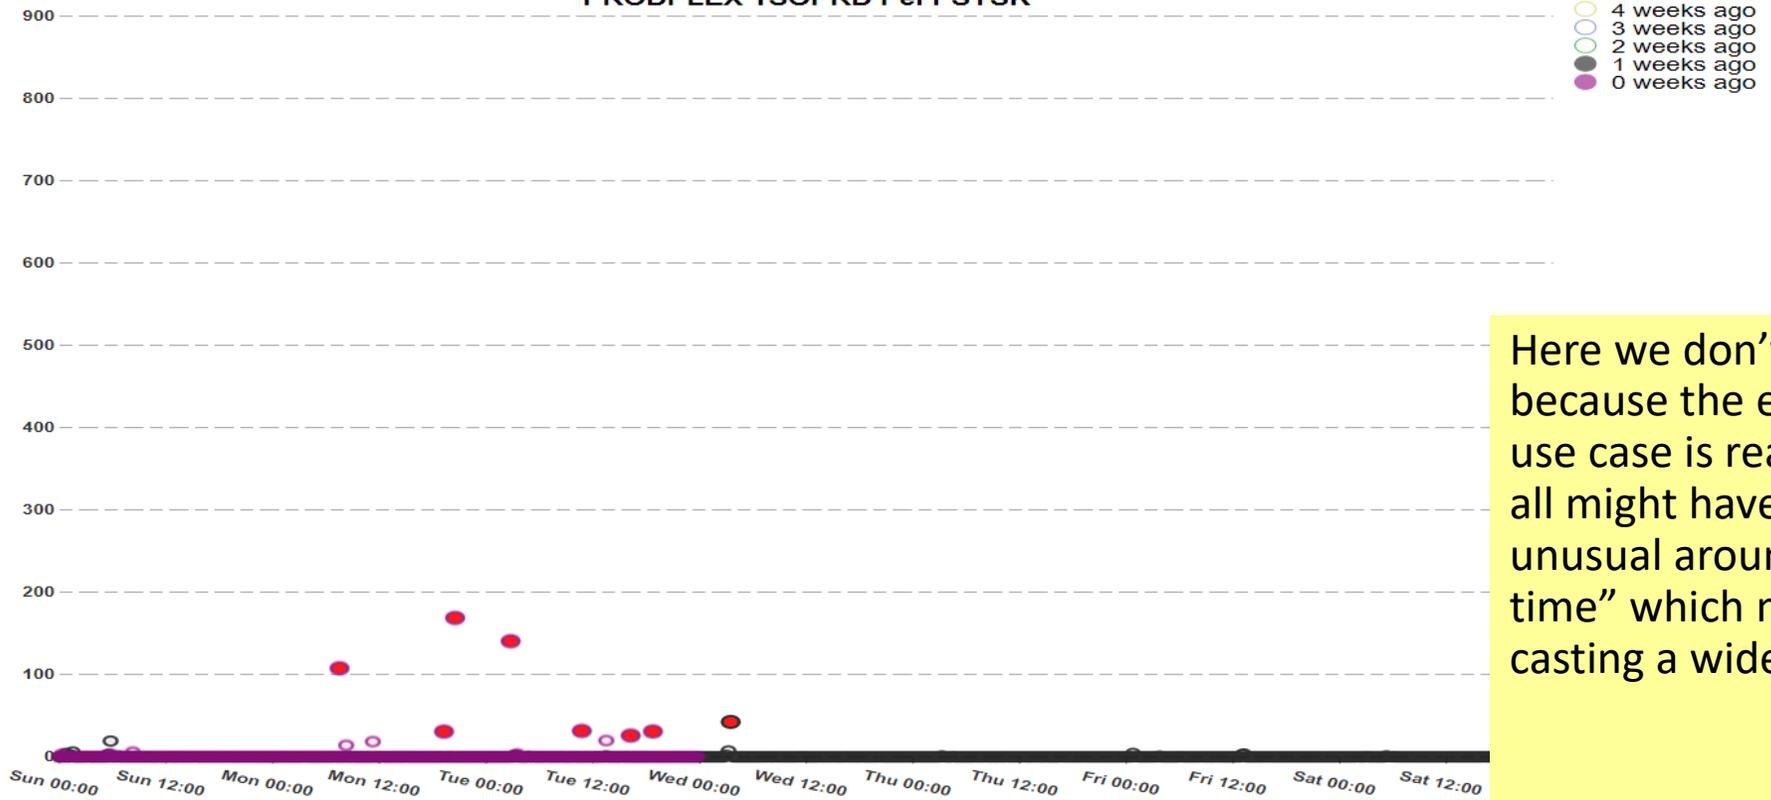
## Outliers by Time Interval

Chart	2:00	12:15	12:30	12:45	13:00	13:15	13:30	13:45	14:00	14:15	14:30	14:45	15:00	15:15	15:30	15:45	16:00	16:15	16:30	16:45	17:00	17:15	17:30	1	
▼ All Reports																									
-									2				1		1			1	1						
▼ Batch Job Outliers																									
<a href="#">CPU time</a>																									
<a href="#">Elapsed time</a>																									
▼ Report Class Overview Metrics																									
<a href="#">DASD IO RT</a>													1		1				1						
<a href="#">SSCH Count</a>									1																
<a href="#">CPU Seconds</a>									1																
<a href="#">zIIP on CP Seconds</a>																									
▼ Service Class Detail Metrics																									
<a href="#">Avg. Memory MB</a>																									
▼ Service Class Sample Counts																									
<a href="#">CPU Delay</a>																									
<a href="#">Unknown</a>																									
<a href="#">zIIP Delay</a>																									
▼ Transaction Metrics for RT SCs																									
<a href="#">Response Time (secs)</a>																									1



## Transaction Metrics for RT SCs

Response Time (secs)  
PRODPLEX TSOPRD Per1 SYSK



Here we don't search because the expected use case is really "what all might have been unusual around this time" which means casting a wider net.

- Outliers
  - LPAR
  - Batch
    - Batch Job Outliers
      - CPU time
      - Elapsed time
  - WLM
    - Service Class Detail Metrics
    - Service Class Sample Counts
    - Transaction Metrics for RT SCs
      - Response Time (secs)
    - Report Class Overview Metrics
      - DASD IO RT
      - SSCH Count
      - CPU Seconds
      - zIIP on CP Seconds
      - zIIP Seconds
  - I/O
    - Coupling Facility
    - Others

SYSPLEX ▲	SC_NAME ▲	SC_PERIOD ▲	SYSTEM ▲	Week 0 Count ▲	Score ▼
PRODPLEX	TSOPRD	Per1	SYSK	7	374.16
PRODPLEX	NE	Per1	SYSL	2	199.79
PRODPLEX	NE	Per1	SYSK	2	141.27
PRODPLEX	ONLINELO	Per1	SYSK	2	76.82



# Outlier Scores



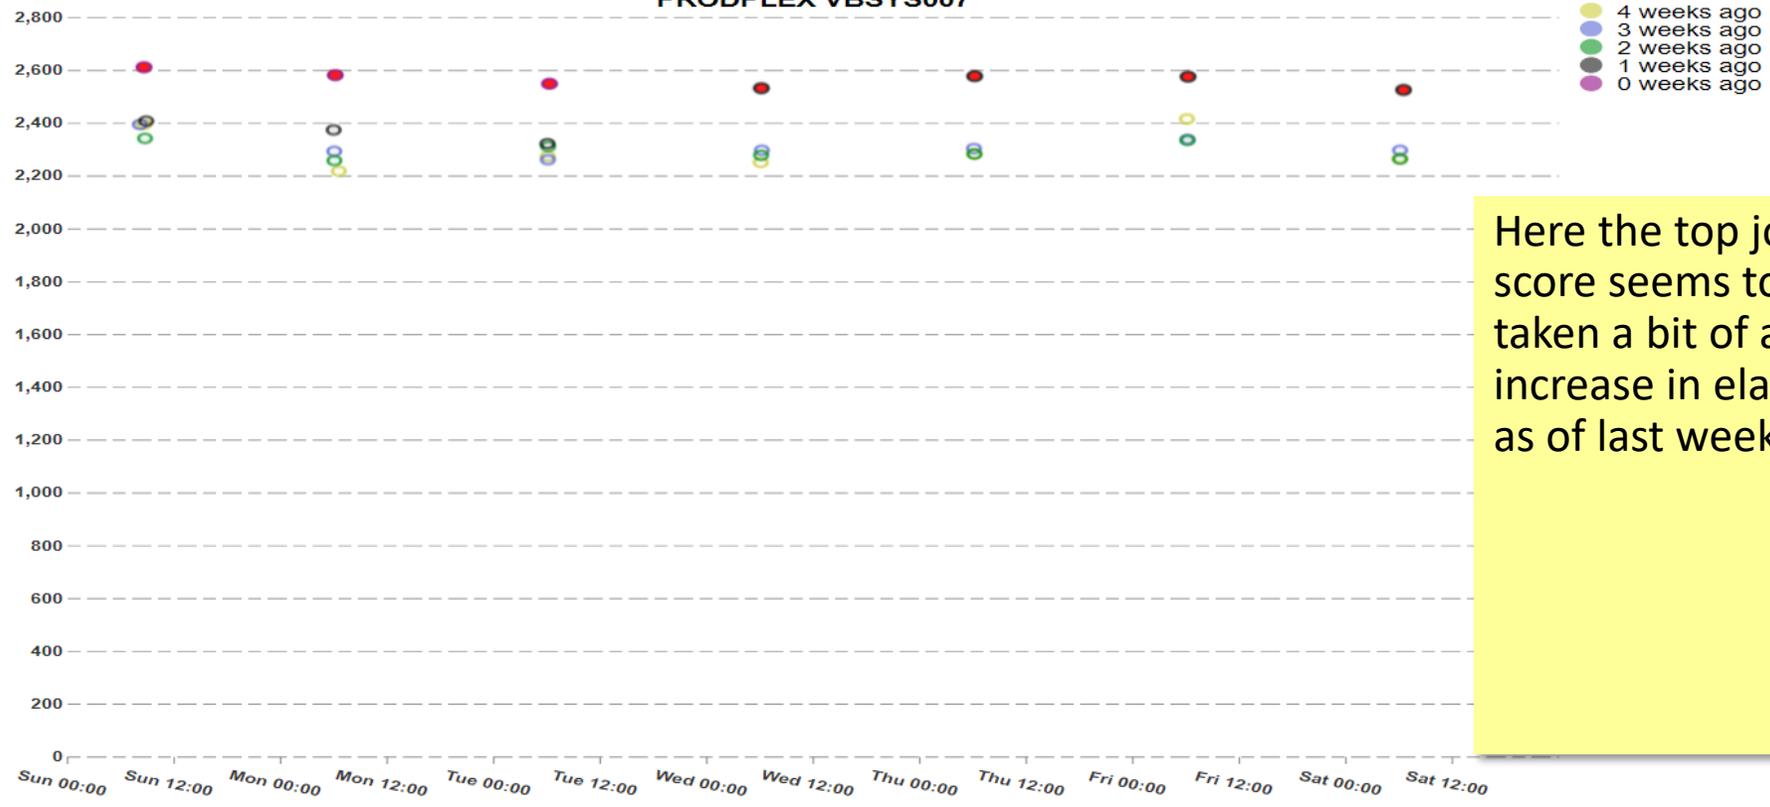
- The purpose of the outlier scores is really to help triage entities to investigate first when you have a large number of entities
  - Primarily: batch jobs
- Generally speaking, the score is computed from the 3-dimensional distance from the normal boxes
  - Dimensions being metric value, day of week, hour of day.
  - Larger values indicate larger deviation from norm and larger metric values
  - Scores are also cumulative by entity so can indicate more outlier observations
- **You cannot compare scores between metrics!**
  - I.E. there's no attempt to normalize CPU time to memory to CPI to whatever
  - There is math involved that makes the numbers also not pure metric values

**Scores are not meaningless, but they're also not *specifically* meaningful!**  
They also may not align with what you think should be investigated first



- Outliers
  - LPAR
    - Processor
      - HIS Metrics Outliers
        - Est Instruction Complexity CPI
        - SIIS Indicator Pct
        - TLB CPU Miss Pct
    - Memory
      - LPAR Storage Outliers
        - cs mb avail less okthresh min
        - Min Avail Storage MB
        - Max CSA MB
        - cs mb hv common max
        - cs mb regionsswa max
        - Max SQA MB
        - Max Shared MB
    - Batch
      - Batch Job Outliers
        - CPU time
        - Elapsed time
    - WLM
      - Service Class Detail Metrics
      - Service Class Sample Counts
        - Contention Delay
        - Contention Using
        - CPU Delay
        - Idle
        - I/O Delay
        - MPL Delay
        - Unknown
        - zIIP Delay
    - Transaction Metrics for RT SCS
      - Response Time (sec)

### Batch Job Outliers Elapsed time PRODPLEX VBSYS007

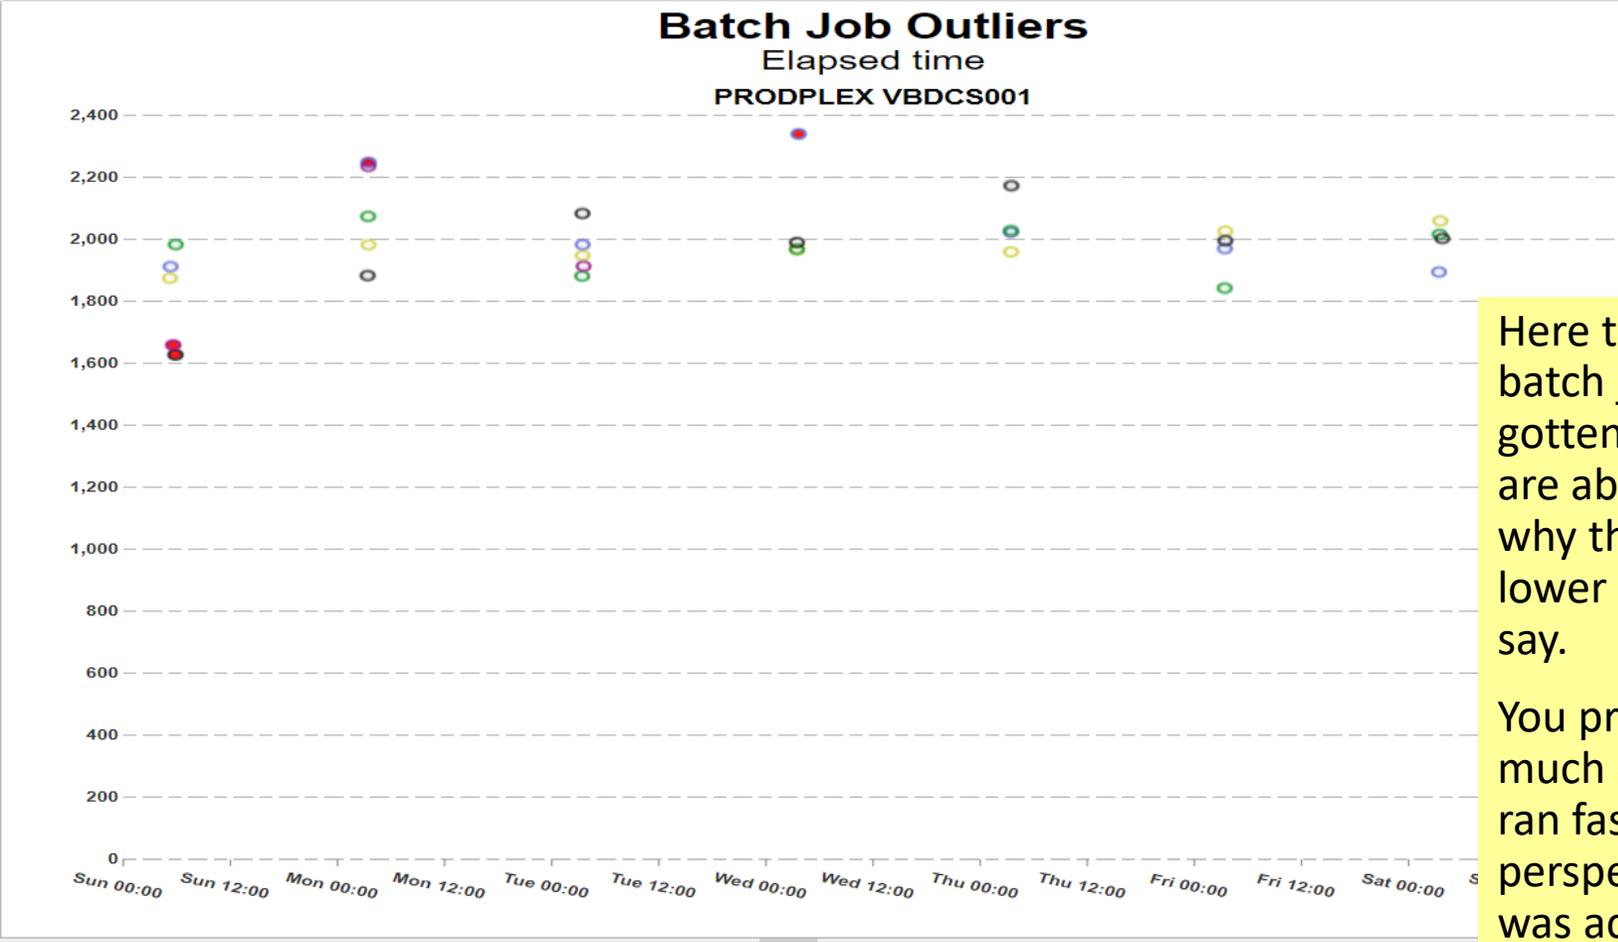


Here the top job by score seems to have taken a bit of a step increase in elapsed time as of last week.

SMF30SYP	JOB_NAME	Week 0 Count	Score
PRODPLEX	VBSYS007	3	104.73
PRODPLEX	PUPI11W	1	99.8
PRODPLEX	PPC1I21D	1	94.24
PRODPLEX	PPR062D	1	93.9



- Outliers
  - LPAR
    - Processor
      - HIS Metrics Outliers
        - Est Instruction Complexity CPI
        - SIIS Indicator Pct
        - TLB CPU Miss Pct
      - Memory
        - LPAR Storage Outliers
          - cs mb avail less okthresh min
          - Min Avail Storage MB
          - Max CSA MB
          - cs mb hv common max
          - cs mb regionsswa max
          - Max SQA MB
          - Max Shared MB
      - Batch
        - Batch Job Outliers
          - CPU time
          - Elapsed time
      - WLM
        - Service Class Detail Metrics
        - Service Class Sample Counts
          - Contention Delay
          - Contention Using
          - CPU Delay
          - Idle
          - I/O Delay
          - MPL Delay
          - Unknown
          - zIIP Delay
      - Transaction Metrics for RT SCs
        - Processor Time (secs)



- 4 weeks ago
- 3 weeks ago
- 2 weeks ago
- 1 weeks ago
- 0 weeks ago

Here the lowest scoring batch job seems to have gotten better. The scores are absolute values, but why the score is so much lower I can't specifically say.

You probably don't care as much about the jobs that ran faster. So from that perspective, the low score was accurate. But tbh, that was probably just lucky.

SMF30SYP	JOB_NAME	Week 0 Count	Score
PRODPLEX	VBP11008	1	22.84
PRODPLEX	PDI370D	1	22.45
PRODPLEX	VBPR1007	1	21.96
PRODPLEX	VBDCS001	1	20.57



# Demo Time!