THE MATH BEHIND THE INCIDENT AFTERMATH

(A Practical Guide To Measuring Incident Impacts)

SREcon22 Asia/Pacific

- Ashish Patel
- Sriram Srinivasan

PayPal

Overview
Table of Contents

Basics
Incident Impact

1.1 Production Environment
1.2 Incident Lifecycle
1.3 Incident Aftermath
1.4 Manual Assessment Challenges

Design
Of System

2.1 Requirements
2.2 Manual Assessment Process
2.3 Key Abstractions
2.4 Architecture

Math
Models

3.1 6-Week Trimmed Avge. Model
3.2 Machine Learning
3.3 Implementation Considerations
3.4 Models Assessment

Features
& Take-Aways

4.1 Key Features
4.2 Opportunities & Challenges
4.3 Titbits on *How to Setup one*?
4.4 Questions

© 2022 PayPal Inc. Confidential and proprietary.

THE BASICS

On Incidents & Its Aftermath

© 2022 PayPal Inc. Confidential and proprietary.

THE BASICS
Production Environment At a Glance

Behind the technology platform of scale

Business in Q3’22

- 5.6 billion transactions
- $337 billion in total payment volume
- 432 million active accounts
- Connects people and businesses in > 200 markets

Platforms

- On both cloud & on-premises
- Distributed across several availability zones
- Hardware and software components
- 1000s of VMs

Applications

- Over 3200 applications and services
- Database & data warehouse
- Multiple programming languages
- Web, mobile & API offerings

© 2022 PayPal Inc. Confidential and proprietary.

PayPal

THE BASICS
Incident Lifecycle

How does an incident transition?

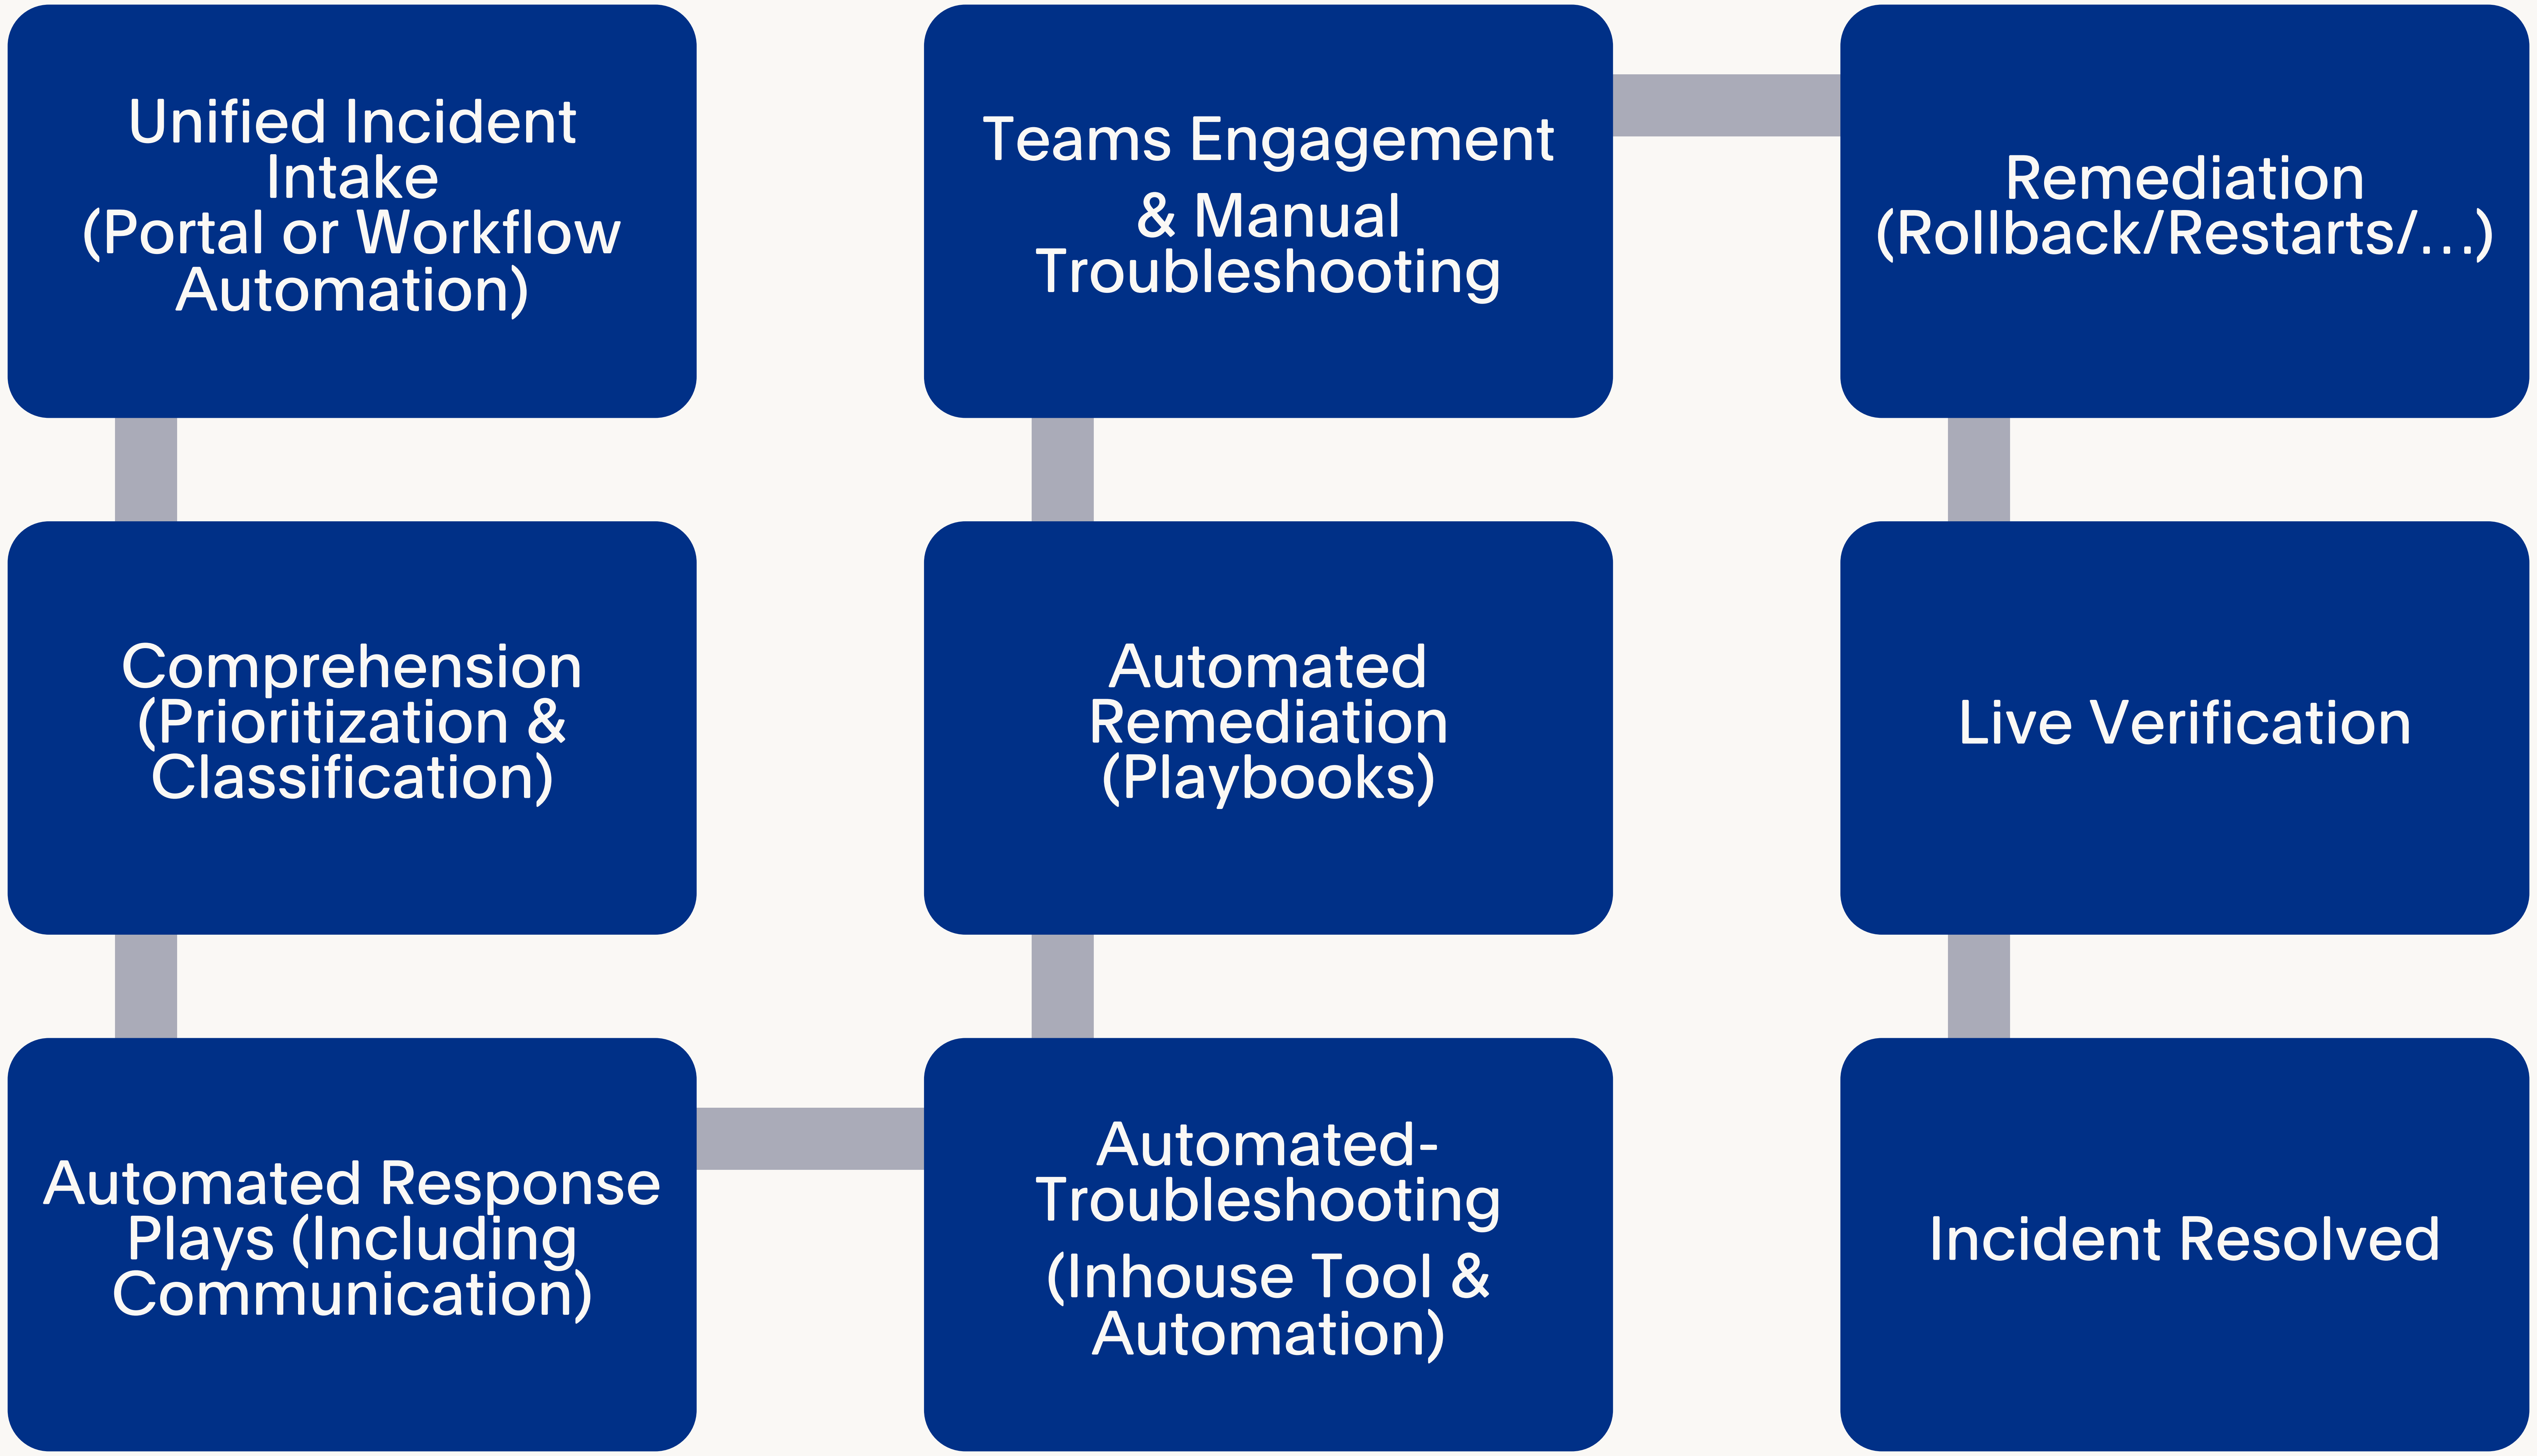

© 2022 PayPal Inc. Confidential and proprietary.

THE BASICS
Incident Aftermath

What is done after an incident?

Incident Timeline	Start Time, Alert Time, Ack Time, Diagnosis Time, Mitigation Time, Impact Duration, Incident End Time.
Impact Assessment	Transaction loss (TPV & Revenue Impact), Availability loss, Segmentation by Customers, Merchants, Products, and Countries.
Notification	Internal and external communication. Regulatory Reporting. Partners, Customers, Merchants, Executives, Status page updates are done.
Root Cause Analysis	Sequence of events. How the process works and what lead to the incident? Caused by change?
Preventive Measures	Any scope for improvements in Detection, Diagnosis, Mitigation, and Recovery? How can we prevent? What lessons are learned (people, processes & technology)?

© 2022 PayPal Inc. Confidential and proprietary.

PayPal

THE BASICS
Manual Assessment Challenges

Challenges

- Time-consuming (several hours)
- Toil to segment the data by products and countries in the trenches
- Knowledge of the area/domain & Challenges in interpreting data
- Error prone
- Inherent urgency in assessing and moving forward
- Not scalable as we grow

Urgency

On the other hand, we need instant assessment for

- Regulatory reporting (varies from 4 to 72 hours)
- Instant communication for customers/merchants & in PayPal Status site
- To know the loss incurred and segment it for next steps

Some Regulators

- CSSF (Luxembourg/EU)
- APRA (Australia)
- KLFB (Japan)
- NYDFS (U.S.)
- SEC (U.S.)
- BACEN (Brazil)
- HKMA (Hong Kong)
- MAS (Singapore)
- BOT (Thailand)
- CBR (Russia)
- CERT-IN (India)

© 2022 PayPal Inc. Confidential and proprietary.

SYSTEM DESIGN

On Requirements & Architecture

© 2022 PayPal Inc. Confidential and proprietary.

SYSTEM DESIGN
Requirements

What are the requirements?

Requirements

- Automate the impact assessment process, i.e., to measure the loss during an incident.

Business Health

- Payment volume is number of payments over a time window (say per minute).
- Total Payment volume is a reliable metric that represents the business health.

Payment Volume

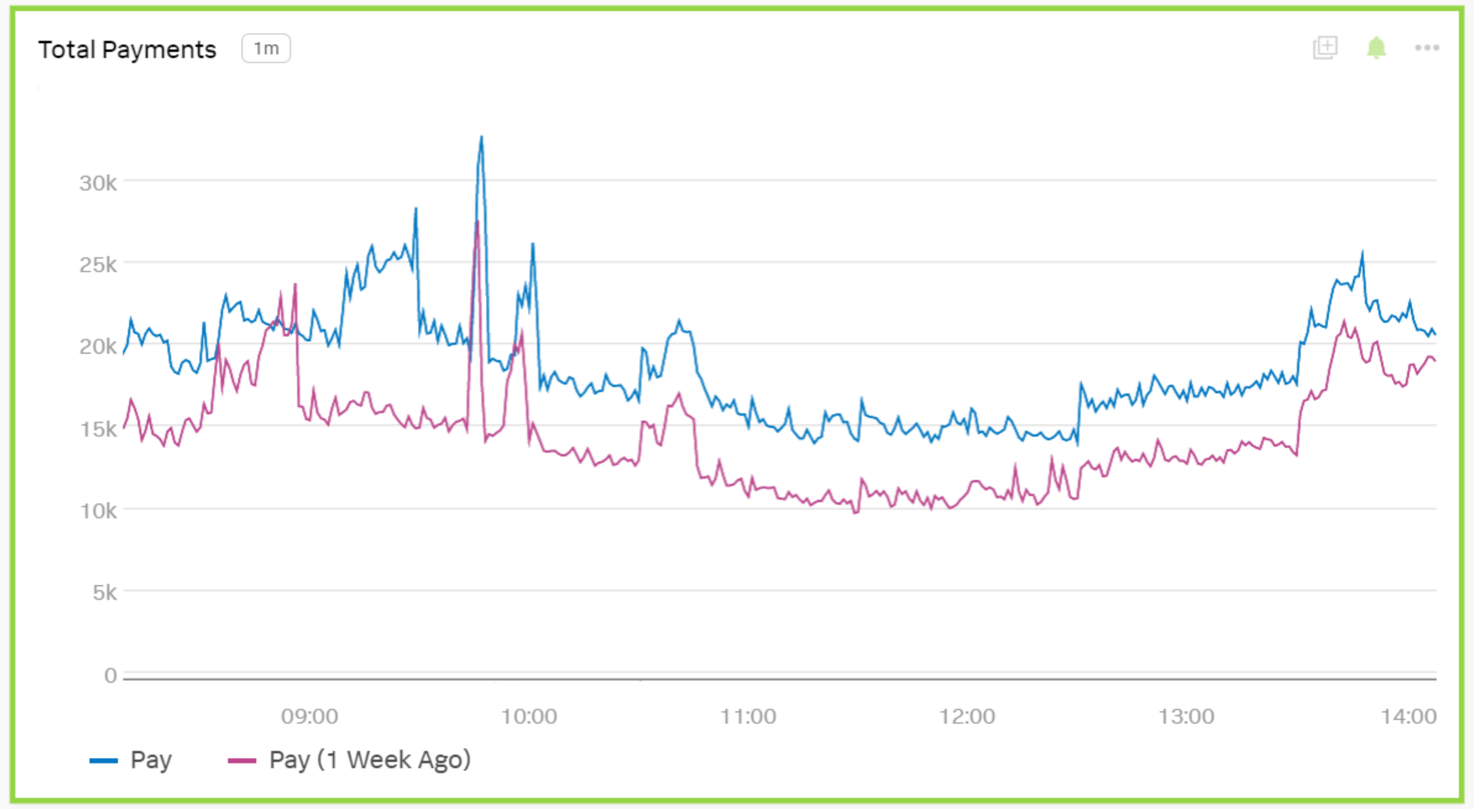

Representational Image

© 2022 PayPal Inc. Confidential and proprietary.

SYSTEM DESIGN
Manual Assessment

How do we measure impacts manually?

STEPS

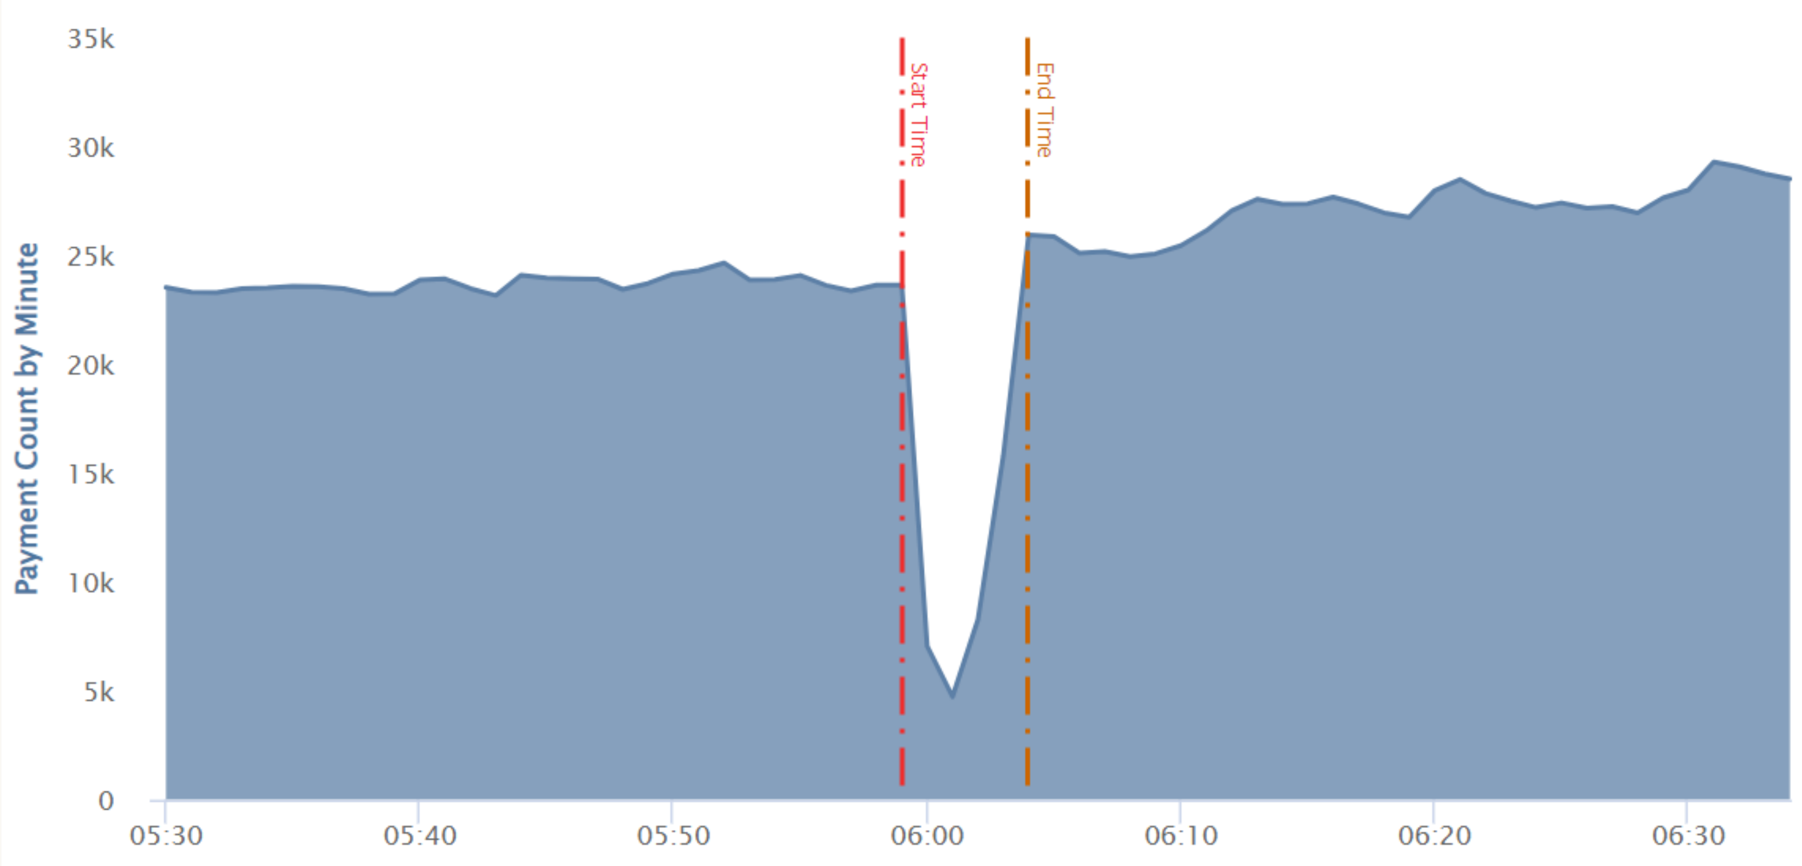

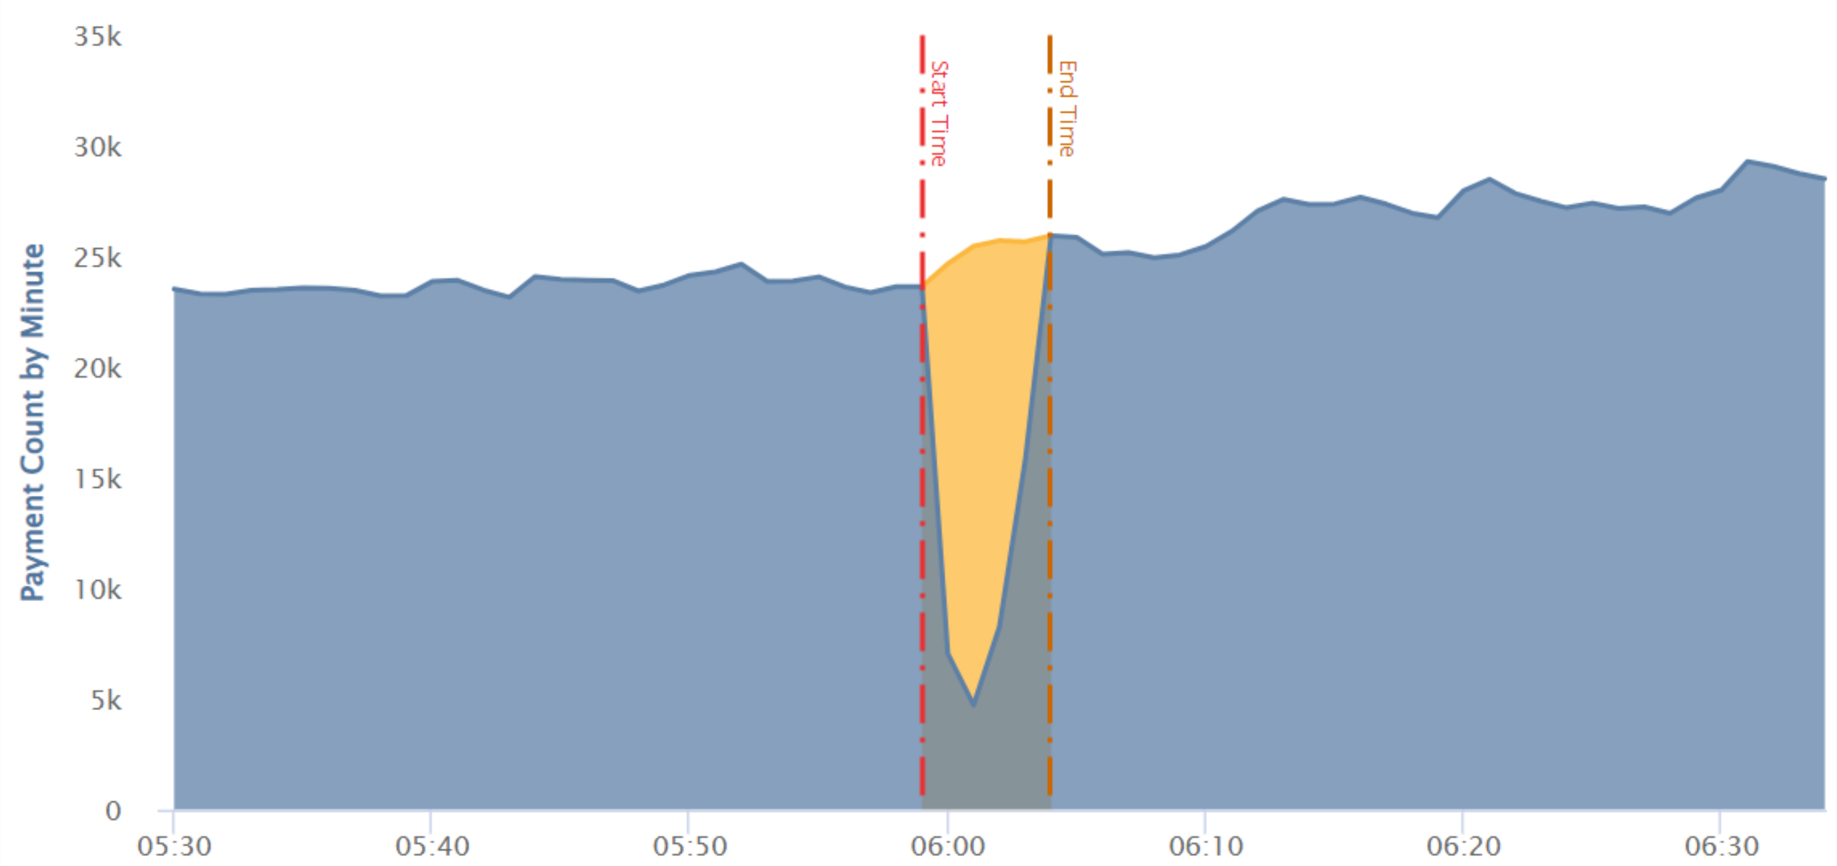

Representational Images

Step 1 – Determine the impact start & end time.

Step 2 – Choose a statistical model. A simple model is ‘x’ week average.

Step 3 – Calculate the modeled count for each minute within the window. For e.g., take average of the same minute for the previous 6 weeks.

Step 4 – Impact for a given minute is *max (0, model payments - actual payments)*.

Step 5 – Total number of lost transactions is the sum of the impact to each minute in the incident timeframe.

Step 6 – Revenue Impact = Total no. of lost transactions x average revenue per payment (in $)

© 2022 PayPal Inc. Confidential and proprietary.

SYSTEM DESIGN
Key Abstractions

Evolution of the Architecture

Ad. Requirements

- Key-in start and end times.
- Provision to talk to disparate real-time data sources.
- Visualize the results.
- Scale up and down the modeled counts as needed.
- Record impact (with screenshots) in ticketing system.

Key Abstractions

- A front-end web-based user interface application.
- Authentication for the tool using SSO.
- Ability to provide multiple models for consumption.
- A good charting library for data visualizations.
- Visual features to play with the modeled data.

Patterns

- n-tier /multitier architecture.
- Database may not be needed.
- Facade design pattern for multiple data sources.
- Existing popular Languages would work.

© 2022 PayPal Inc. Confidential and proprietary.

SYSTEM DESIGN
Architecture

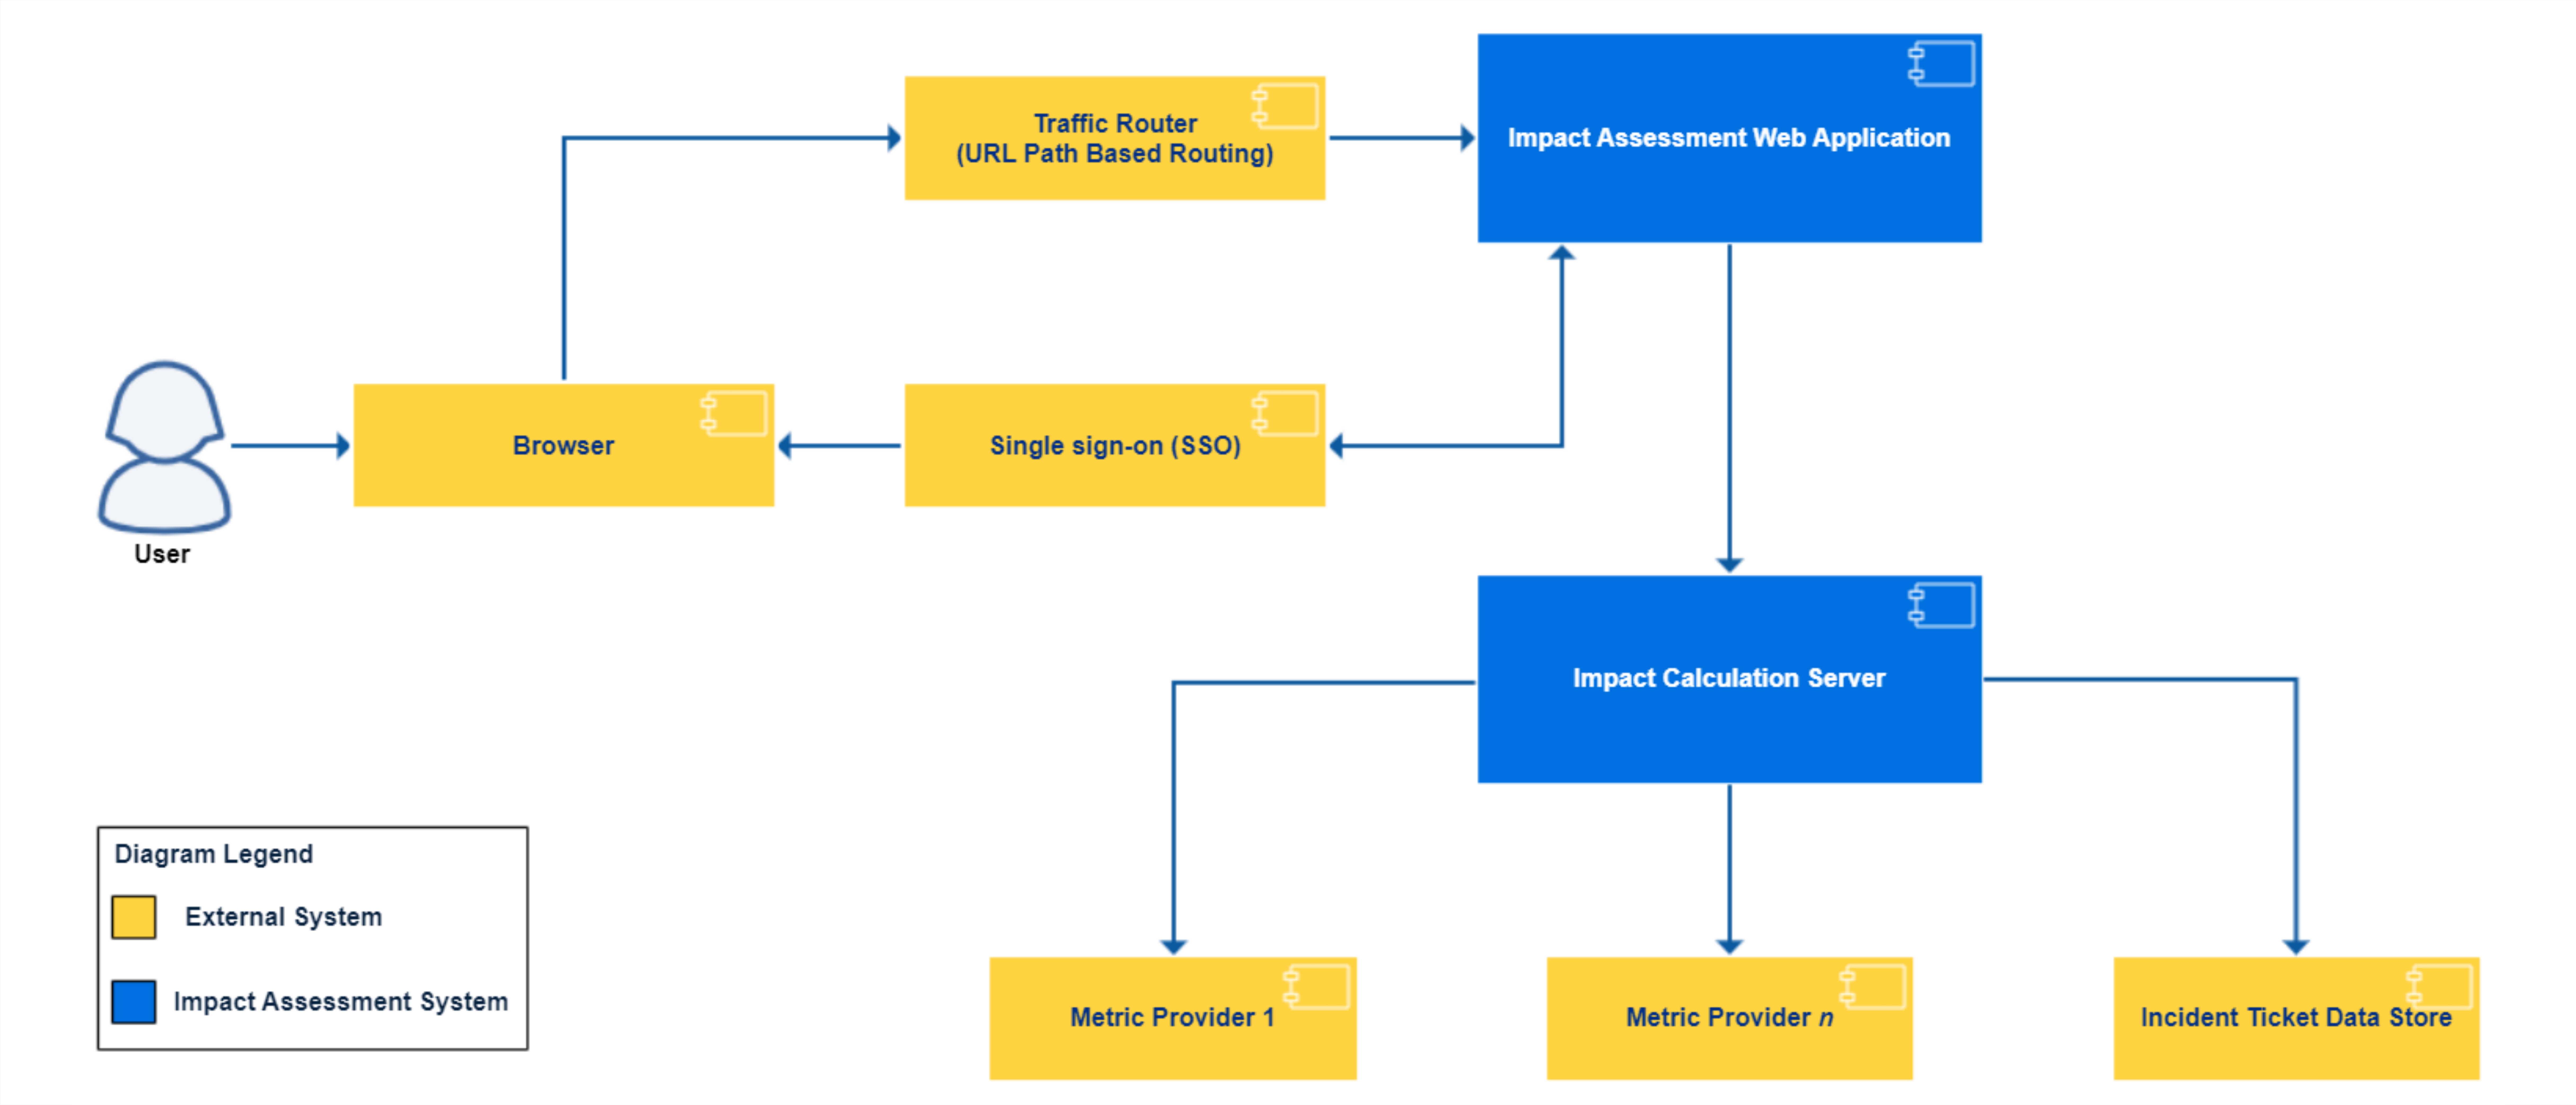

© 2022 PayPal Inc. Confidential and proprietary.

MATHEMATICAL MODELING

On Statistical & Machine Learning

© 2022 PayPal Inc. Confidential and proprietary.

PayPal

MATHEMATICAL MODELING
Scaled 6 Week Trimmed Average

How does the 6-week average model work?

Process

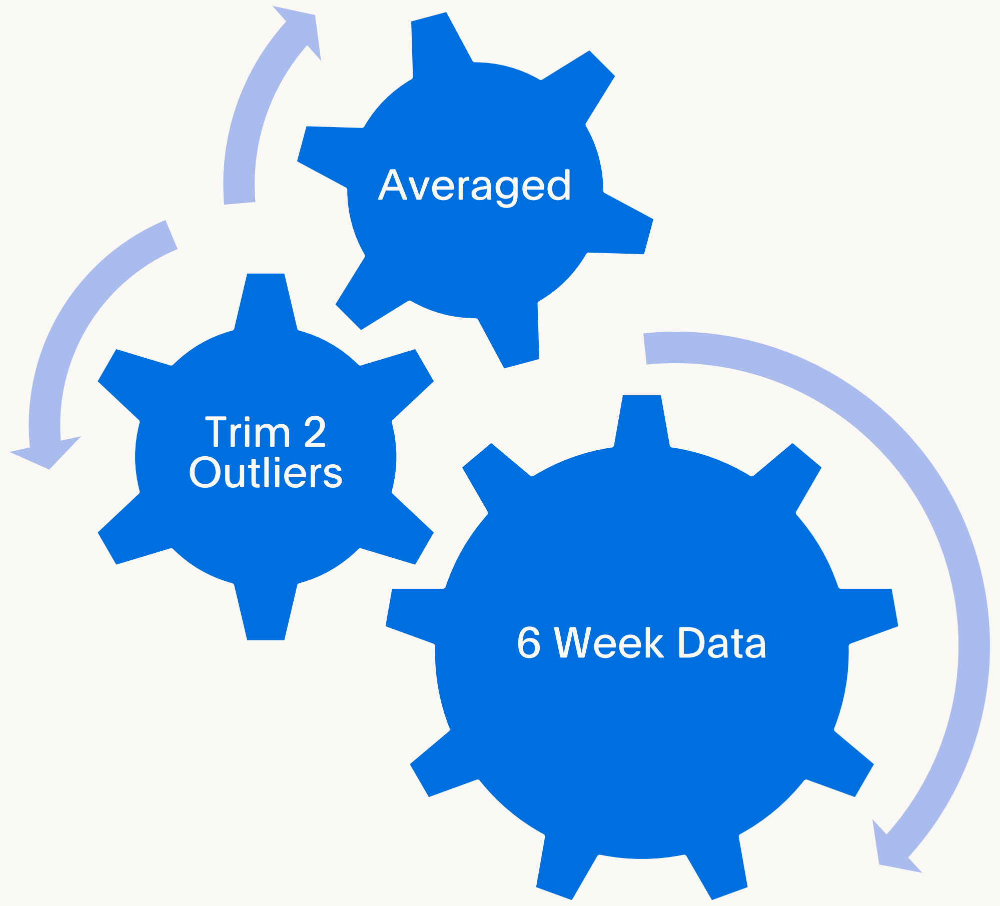

Model

1. Get count of the same minute for each of the previous 6 weeks.
2. Remove the highest and lowest value (2 outliers).
3. Calculate the mean of the remaining 4 values.
4. Repeat the process for all the minutes in the window.
5. Scale by a multiplier to align it with the data preceding the incident.

Example

Description	Incident Time	# Payment	Comment
Incident Time	2022/12/07-10:30:00	4	Actual
Week-1	2022/11/30-10:30:00	7	
Week-2	2022/11/23-10:30:00	8	Outlier
Week-3	2022/11/16-10:30:00	7	
Week-4	2022/11/09-10:30:00	4	
Week-5	2022/11/02-10:30:00	3	Outlier
Week-6	2022/10/26-10:30:00	6	

Mean		6	
Lost		2	

© 2022 PayPal Inc. Confidential and proprietary.

MATHEMATICAL MODELING

Scaling of Trimmed Average

How does scaling work?

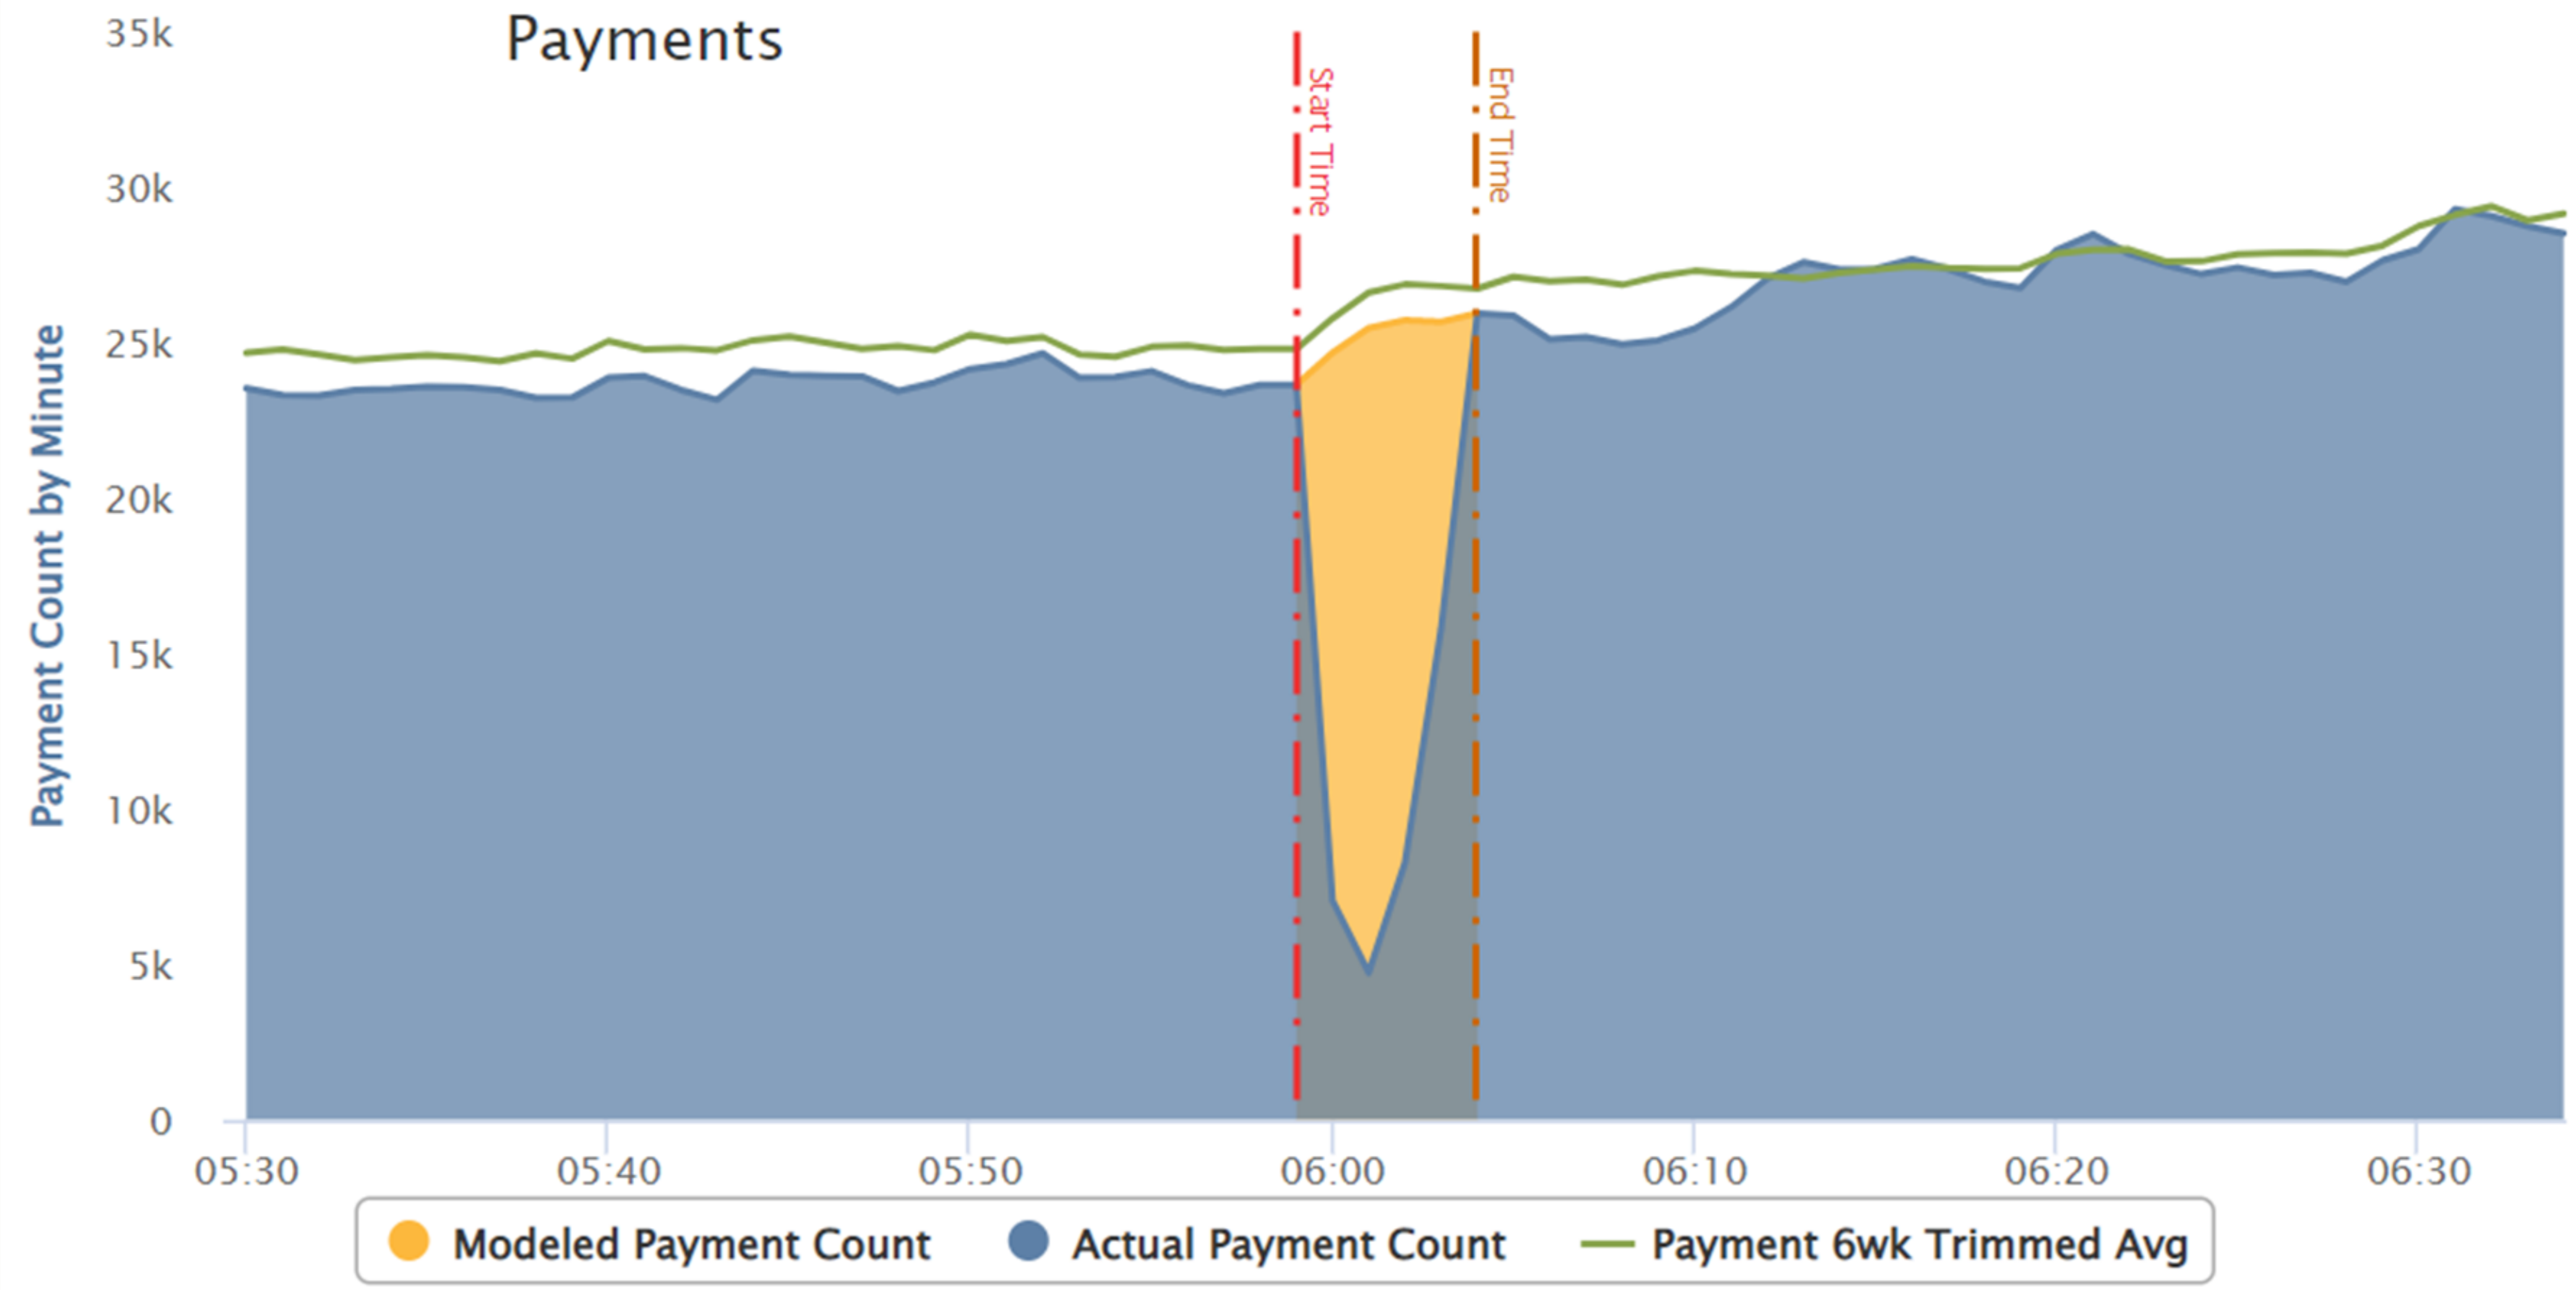

Payment Loss %	64.7
Actual Payments	35,885
Model Payments	101,541
Payment Loss	65,656
Recaptured 0%	0
Net Payment Loss	65,656

Representational Image

© 2022 PayPal Inc. Confidential and proprietary.

MATHEMATICAL MODELING

Recapture & Manual Scaling

How do we account for recapture & can we adjust scaling?

Recapture

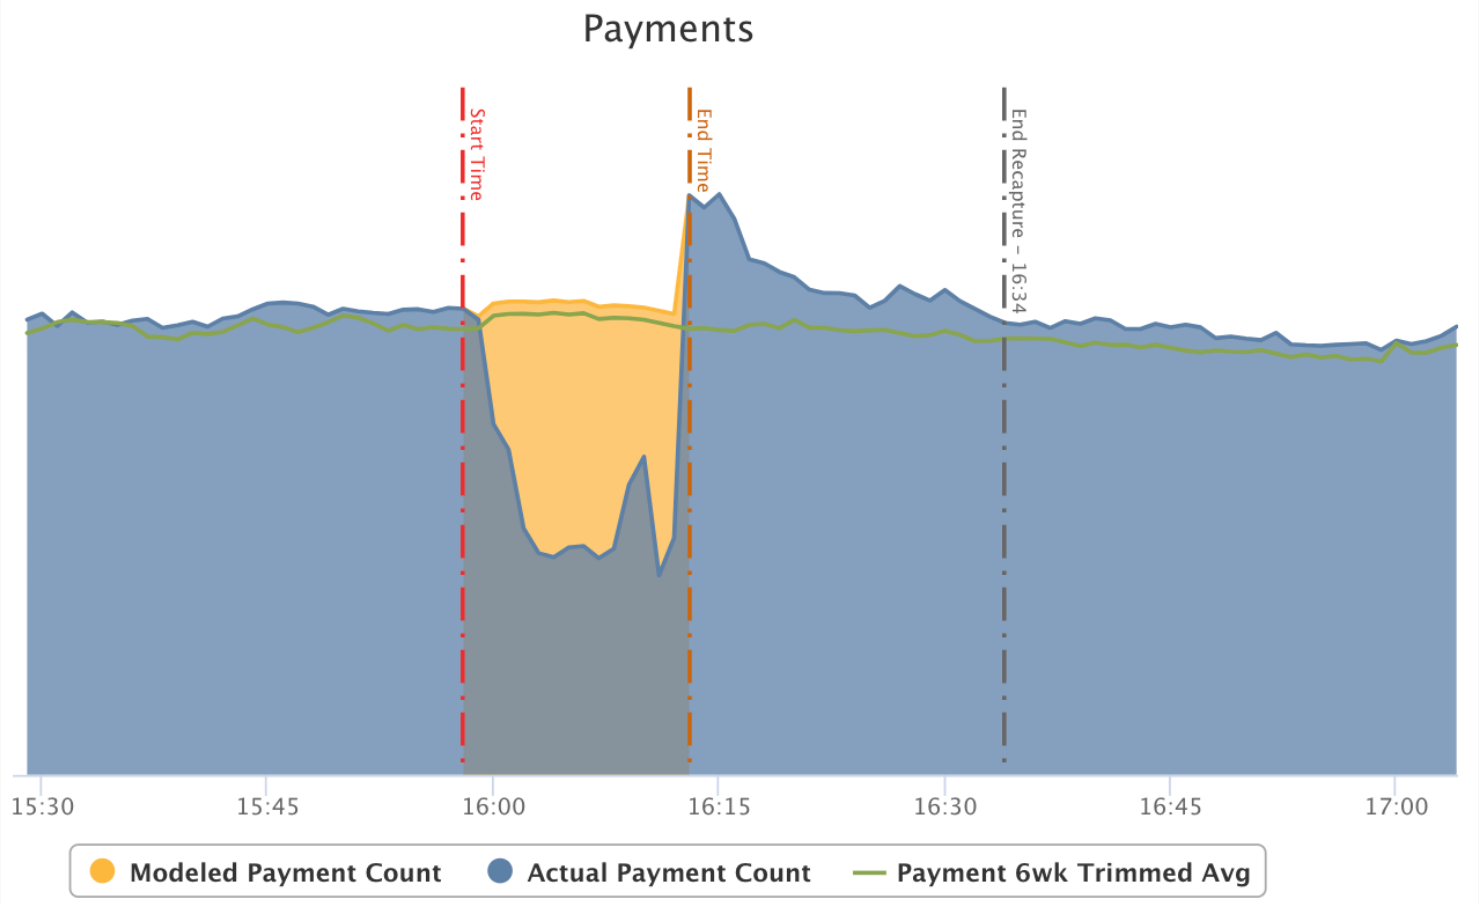

After the incident ends, some failed transactions would be retried. This can be seen as a spike compared to the norm. We discount them from impact.

Default Model @ 1.0

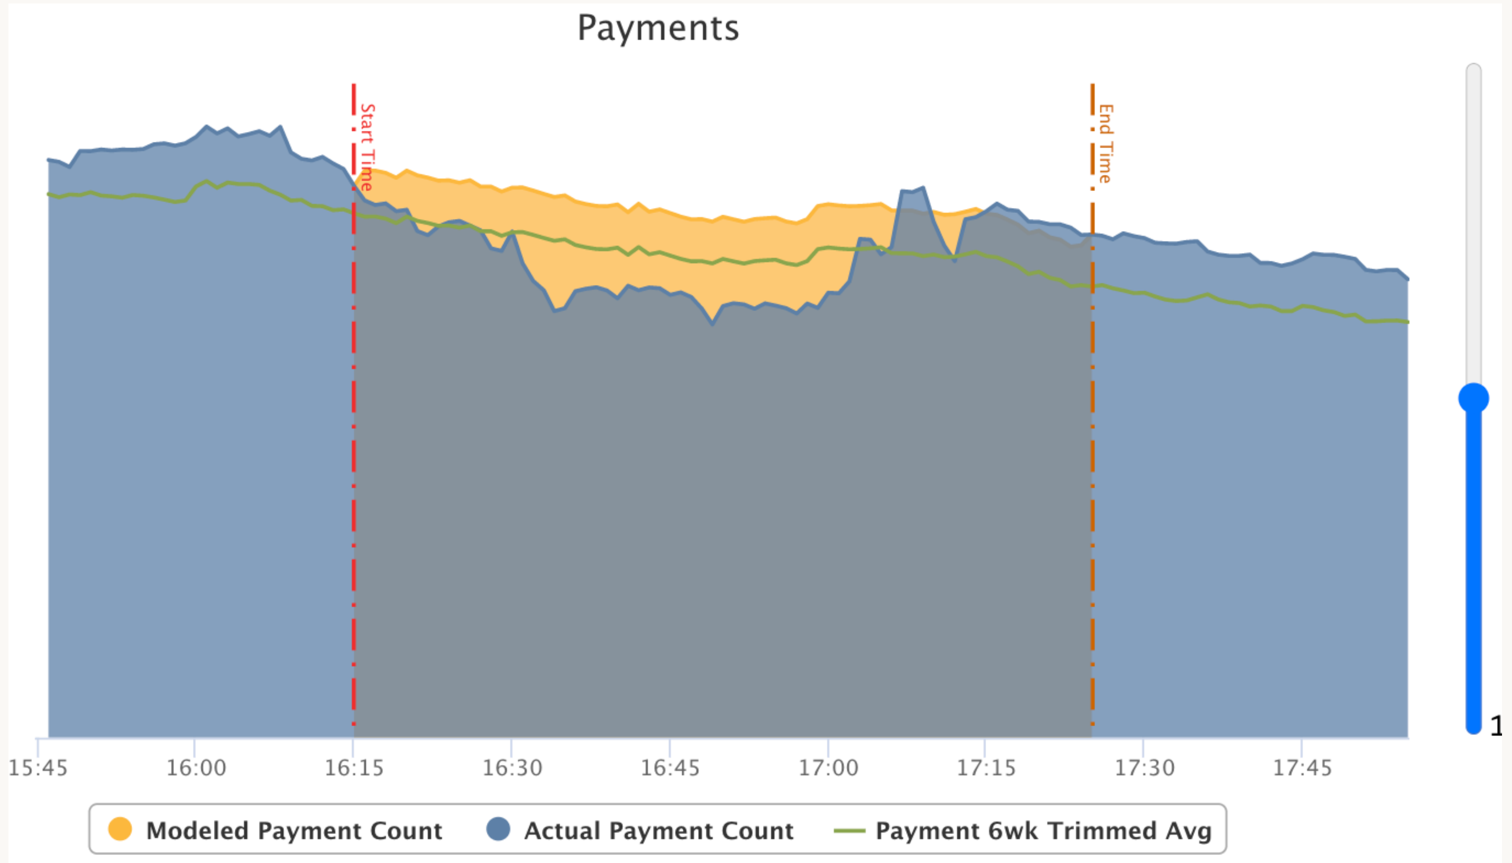

Rarely a model like scaled `6` week trimmed average may not align properly with current trend. This can be observed visually.

Scaled Model @ 1.05

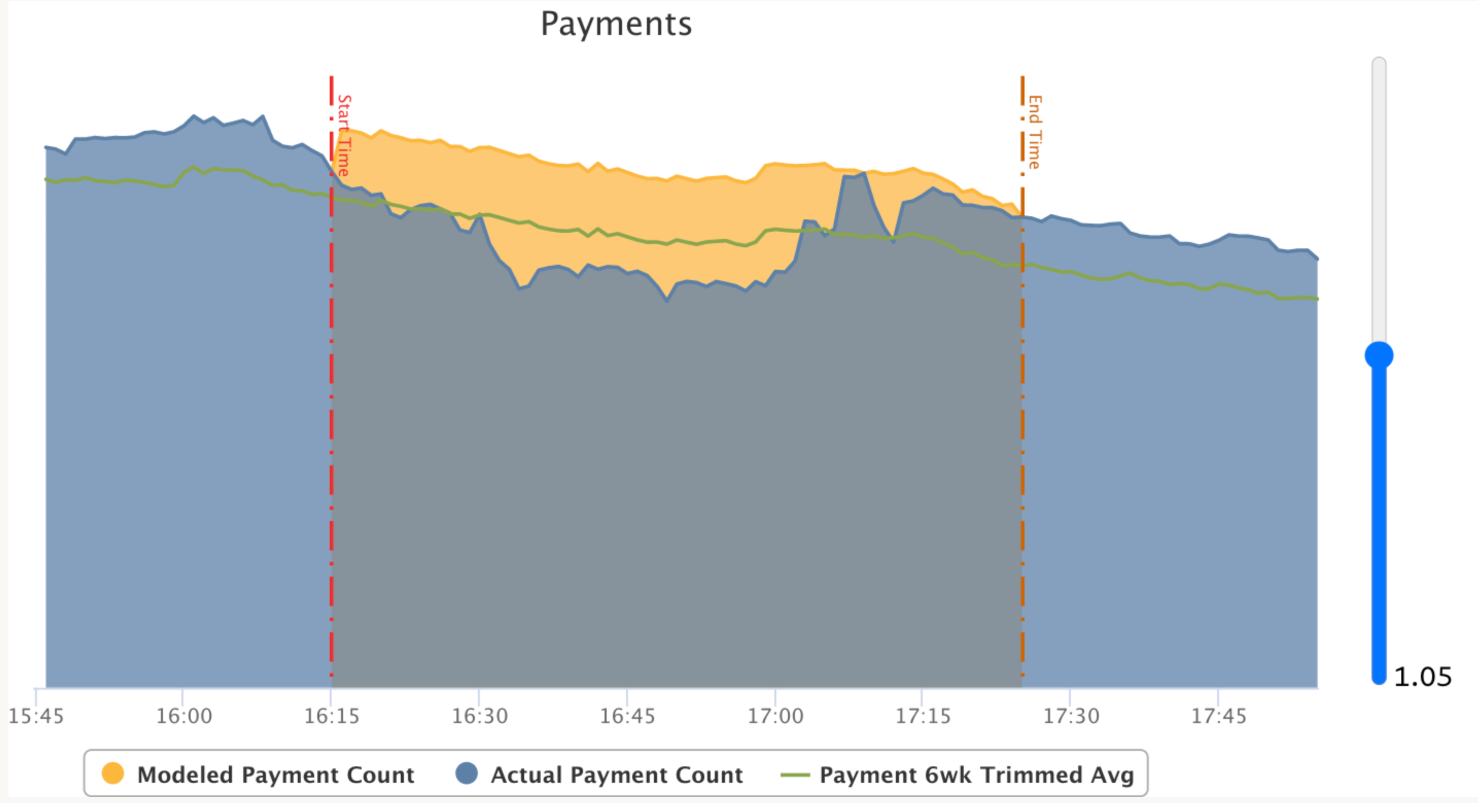

Users can do minor scaling (multiplier) to the model to align it to the current trend. Here the model is scaled by 5% upwards.

© 2022 PayPal Inc. Confidential and proprietary.

MATHEMATICAL MODELING

Formula for ‘x’ Week Trimmed Average

‘X’ Week Trimmed Average

$$impact(start\_time, end\_time) = \sum_{t=start\_time}^{t=end\_time} max(expected\_value[t] - observed\_value[t], 0)$$

$$recaptured(end\_time, recapture\_end\_time) = \sum_{t=end\_time}^{t=recapture\_end\_time} max(observed\_value[t] - expected\_value, 0)$$

$$net\_impact(start\_time, end\_time, recapture\_end\_time) = impact(start\_time, end\_time) - recaptured(end\_time, recapture\_end\_time)$$

$$max\_value(t) = max\{observed\_value[t - week\_to\_minutes(w)]: w = 1..x\}$$

$$min\_value(t) = min\{observed\_value[t - week\_to\_minutes(w)]: w = 1..x\}$$

$$sum\_of\_past\_observations(t) = \sum_{w=1}^{w=x} observed\_value[t - week\_to\_minutes(w)]$$

$$expected\_value(t) = sum\_of\_past\_observation(t) - max\_value(t) - min\_value(t) \,/\, (x-2) : x > 2$$
$$expected\_series = \{expected\_value(t) : t = start..end\}$$

© 2022 PayPal Inc. Confidential and proprietary.

MATHEMATICAL MODELING

Alignment and Scaling

Scaling to align with recent trend

$$padded\_start = start - 30$$

$$padded\_end = end + 30$$

$$ratio(t) = observed\_value(t)/expected\_series(t)$$

$$factor = median(\{ratio(t) : t = padded\_start..start\}$$

$$expected\_series(padded\_start, padded\_end) = \{v * factor : v \in expected\_series\}$$

$$scale(series, factor) = \{v * factor : v \in series\}$$

$$scaled\_expected\_series = scale(expected\_series, factor)$$

$$\boldsymbol{revenu\_loss = net\_impact * avg\_revenue\_per\_txn}$$

© 2022 PayPal Inc. Confidential and proprietary.

MATHEMATICAL MODELING
Employing Machine Learning

Forecasting

Distilled from ~180 Models on Time Series Forecasting:

- ARIMA & Variants
- LightGBM
- Linear/Polynomial Regression
- LSTM
- NBEATS
- Exponential Smoothing (ETS) Algorithm
- The Theta Model
- Transformer Architectures

Machine Learning

Emergence of AI Cloud Platform which could automate a lot of the steps

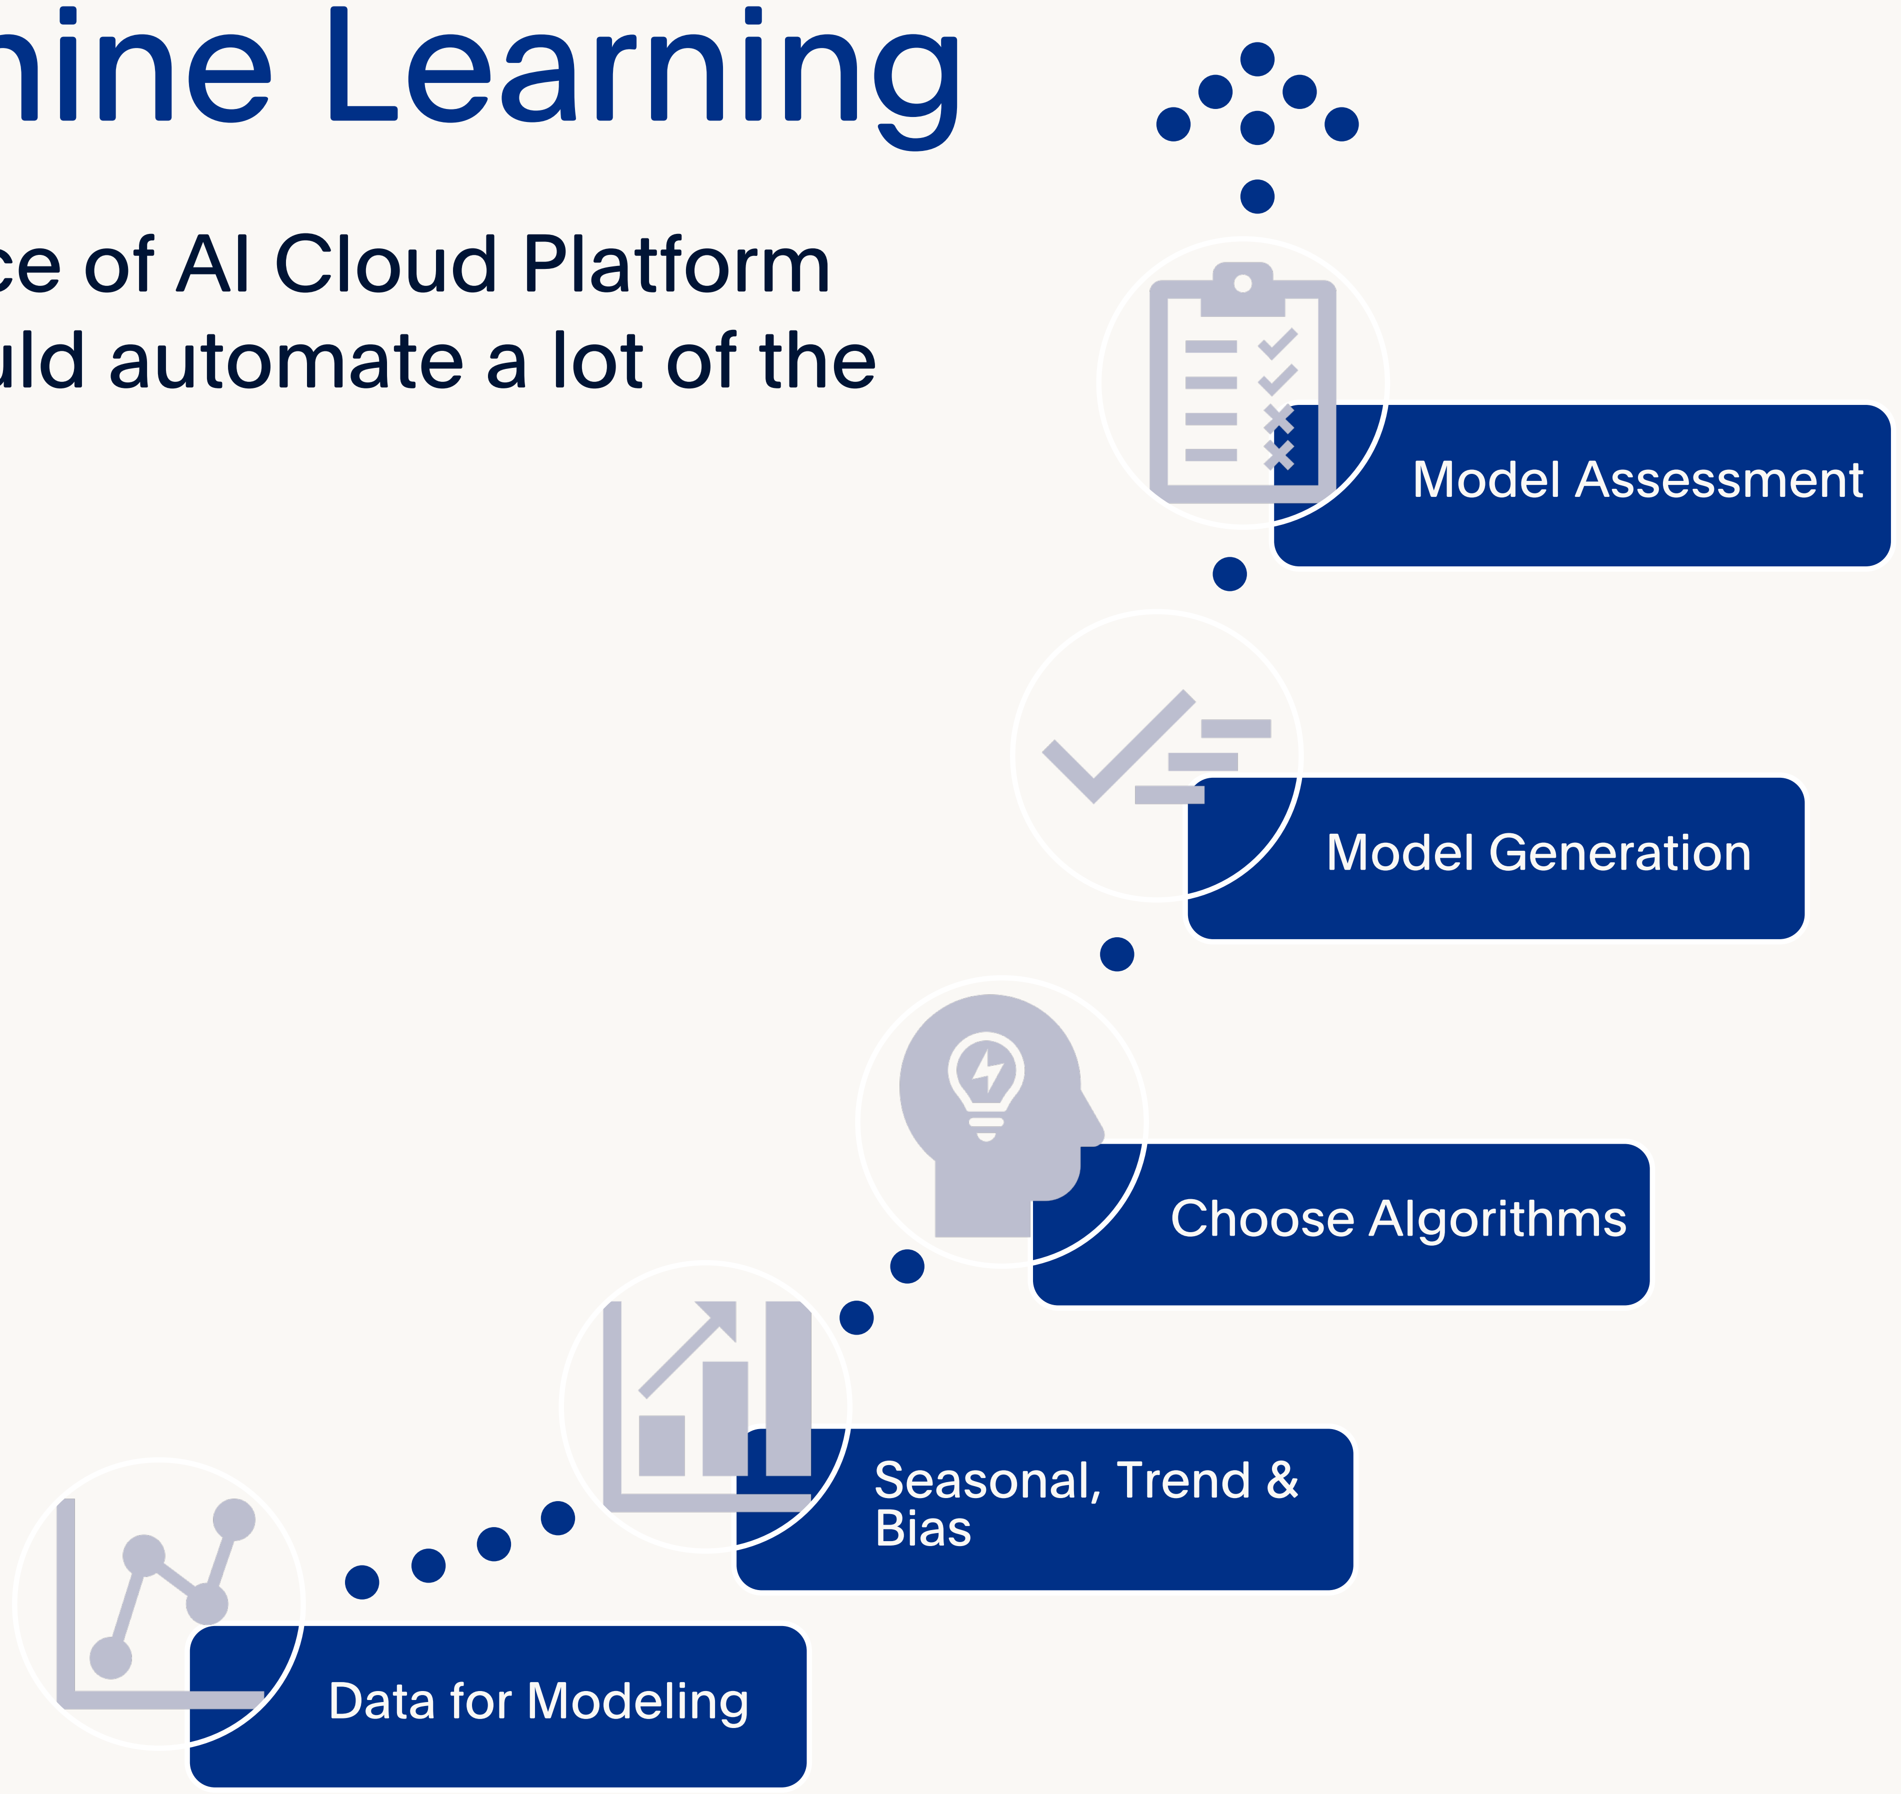

© 2022 PayPal Inc. Confidential and proprietary.

MATHEMATICAL MODELING
Implementation Considerations

How do we develop a good model?

Near Real-time Data	As we need to provide the impact data in near real-time, we may have to ensure that the source data is available.
Understanding Datasets	Seasonal, Trend & Bias. Univariate/multivariate data. Stationary and non-stationary data. Highly correlated or not.
Dataset for Training	Test and train the models on nonimpact data.
Many Models	Tried with many models iteratively and with many possible use cases (test data).
Integration	How are we going to integrate the model?

© 2022 PayPal Inc. Confidential and proprietary.

PayPal

MATHEMATICAL MODELING ARIMA IN ACTION

ARIMA

How did ARIMA model work?

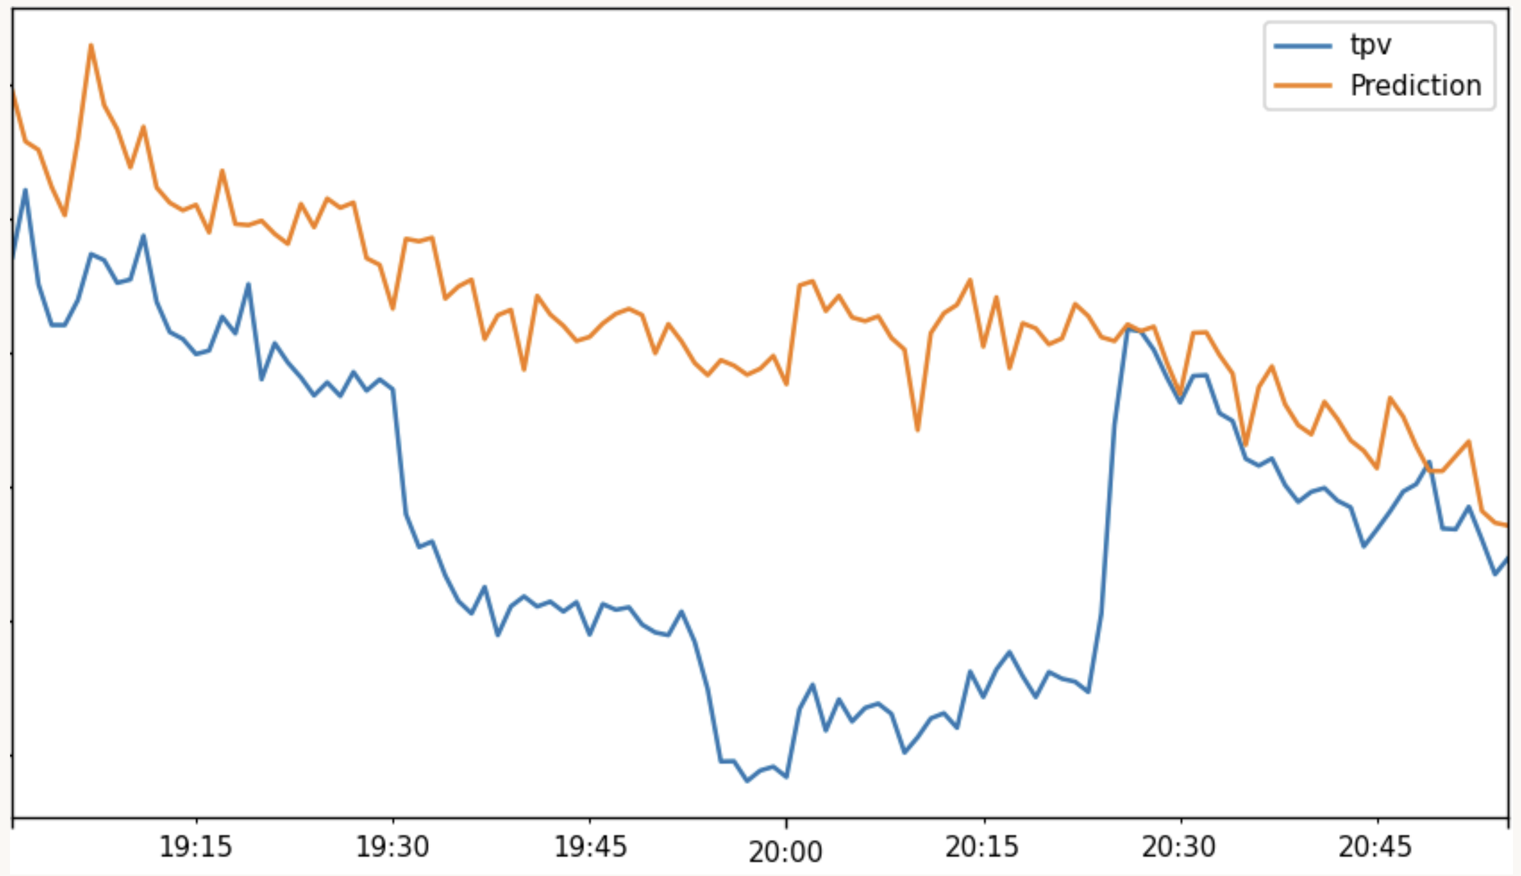

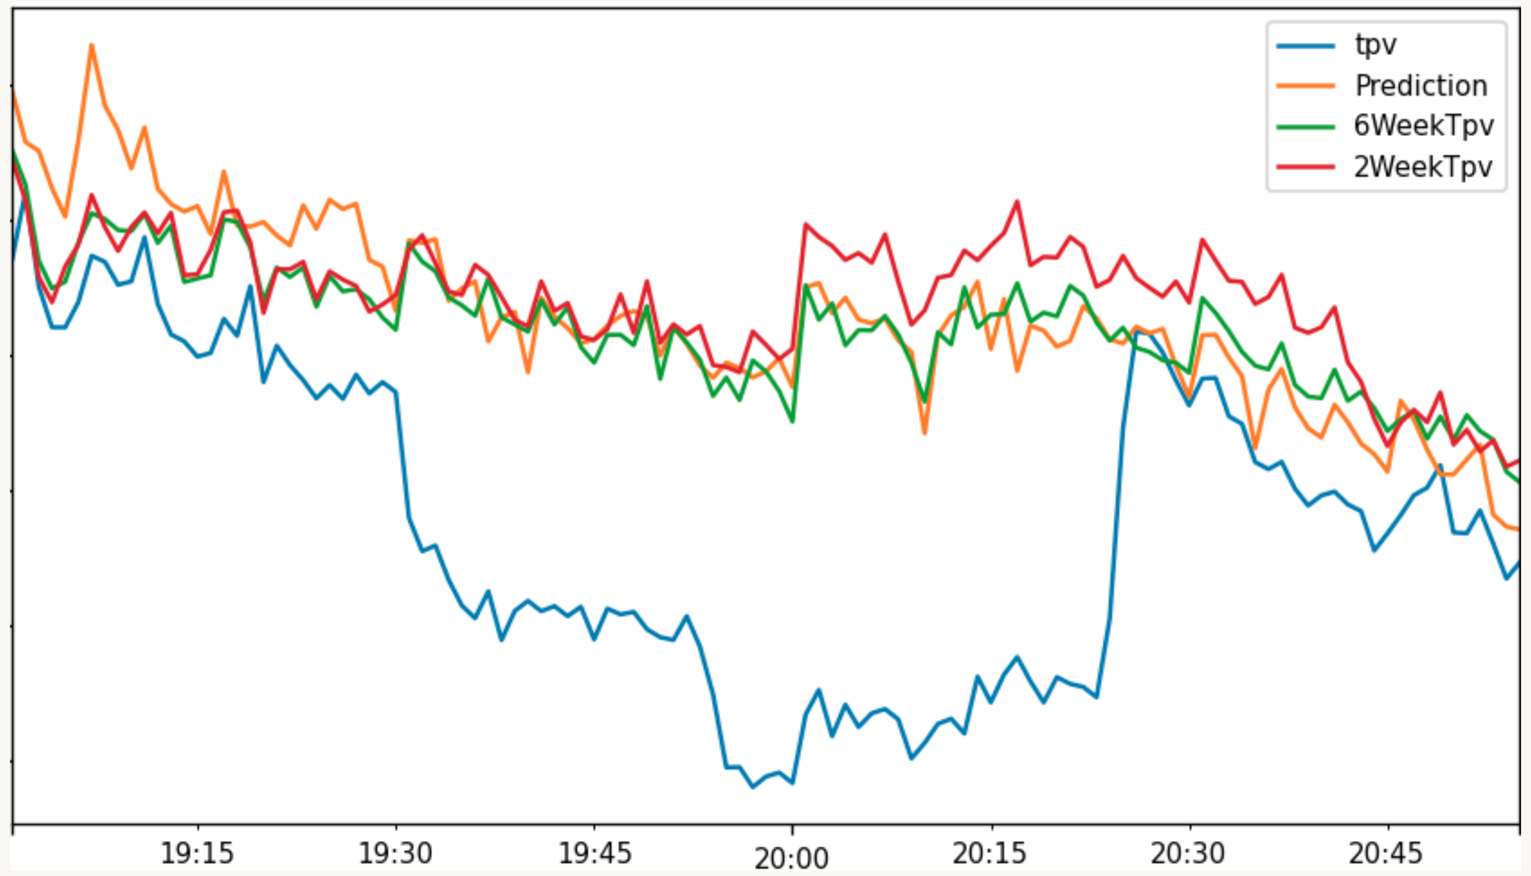

- ARIMA and variants (ARMA, Seasonal ARIMA).
- ARIMA (3,0,2) fared well compared to others.
- Autoregressive integrated moving average.
- The model performed well but the challenge is to train and utilize it in real-time.
- Also, a lot of preprocessing is required.

© 2022 PayPal Inc. Confidential and proprietary.

MATHEMATICAL MODELING

H2O Driverless AI

How do we build models using H2O Driverless AI?

Automated Machine Learning – Experiment Result

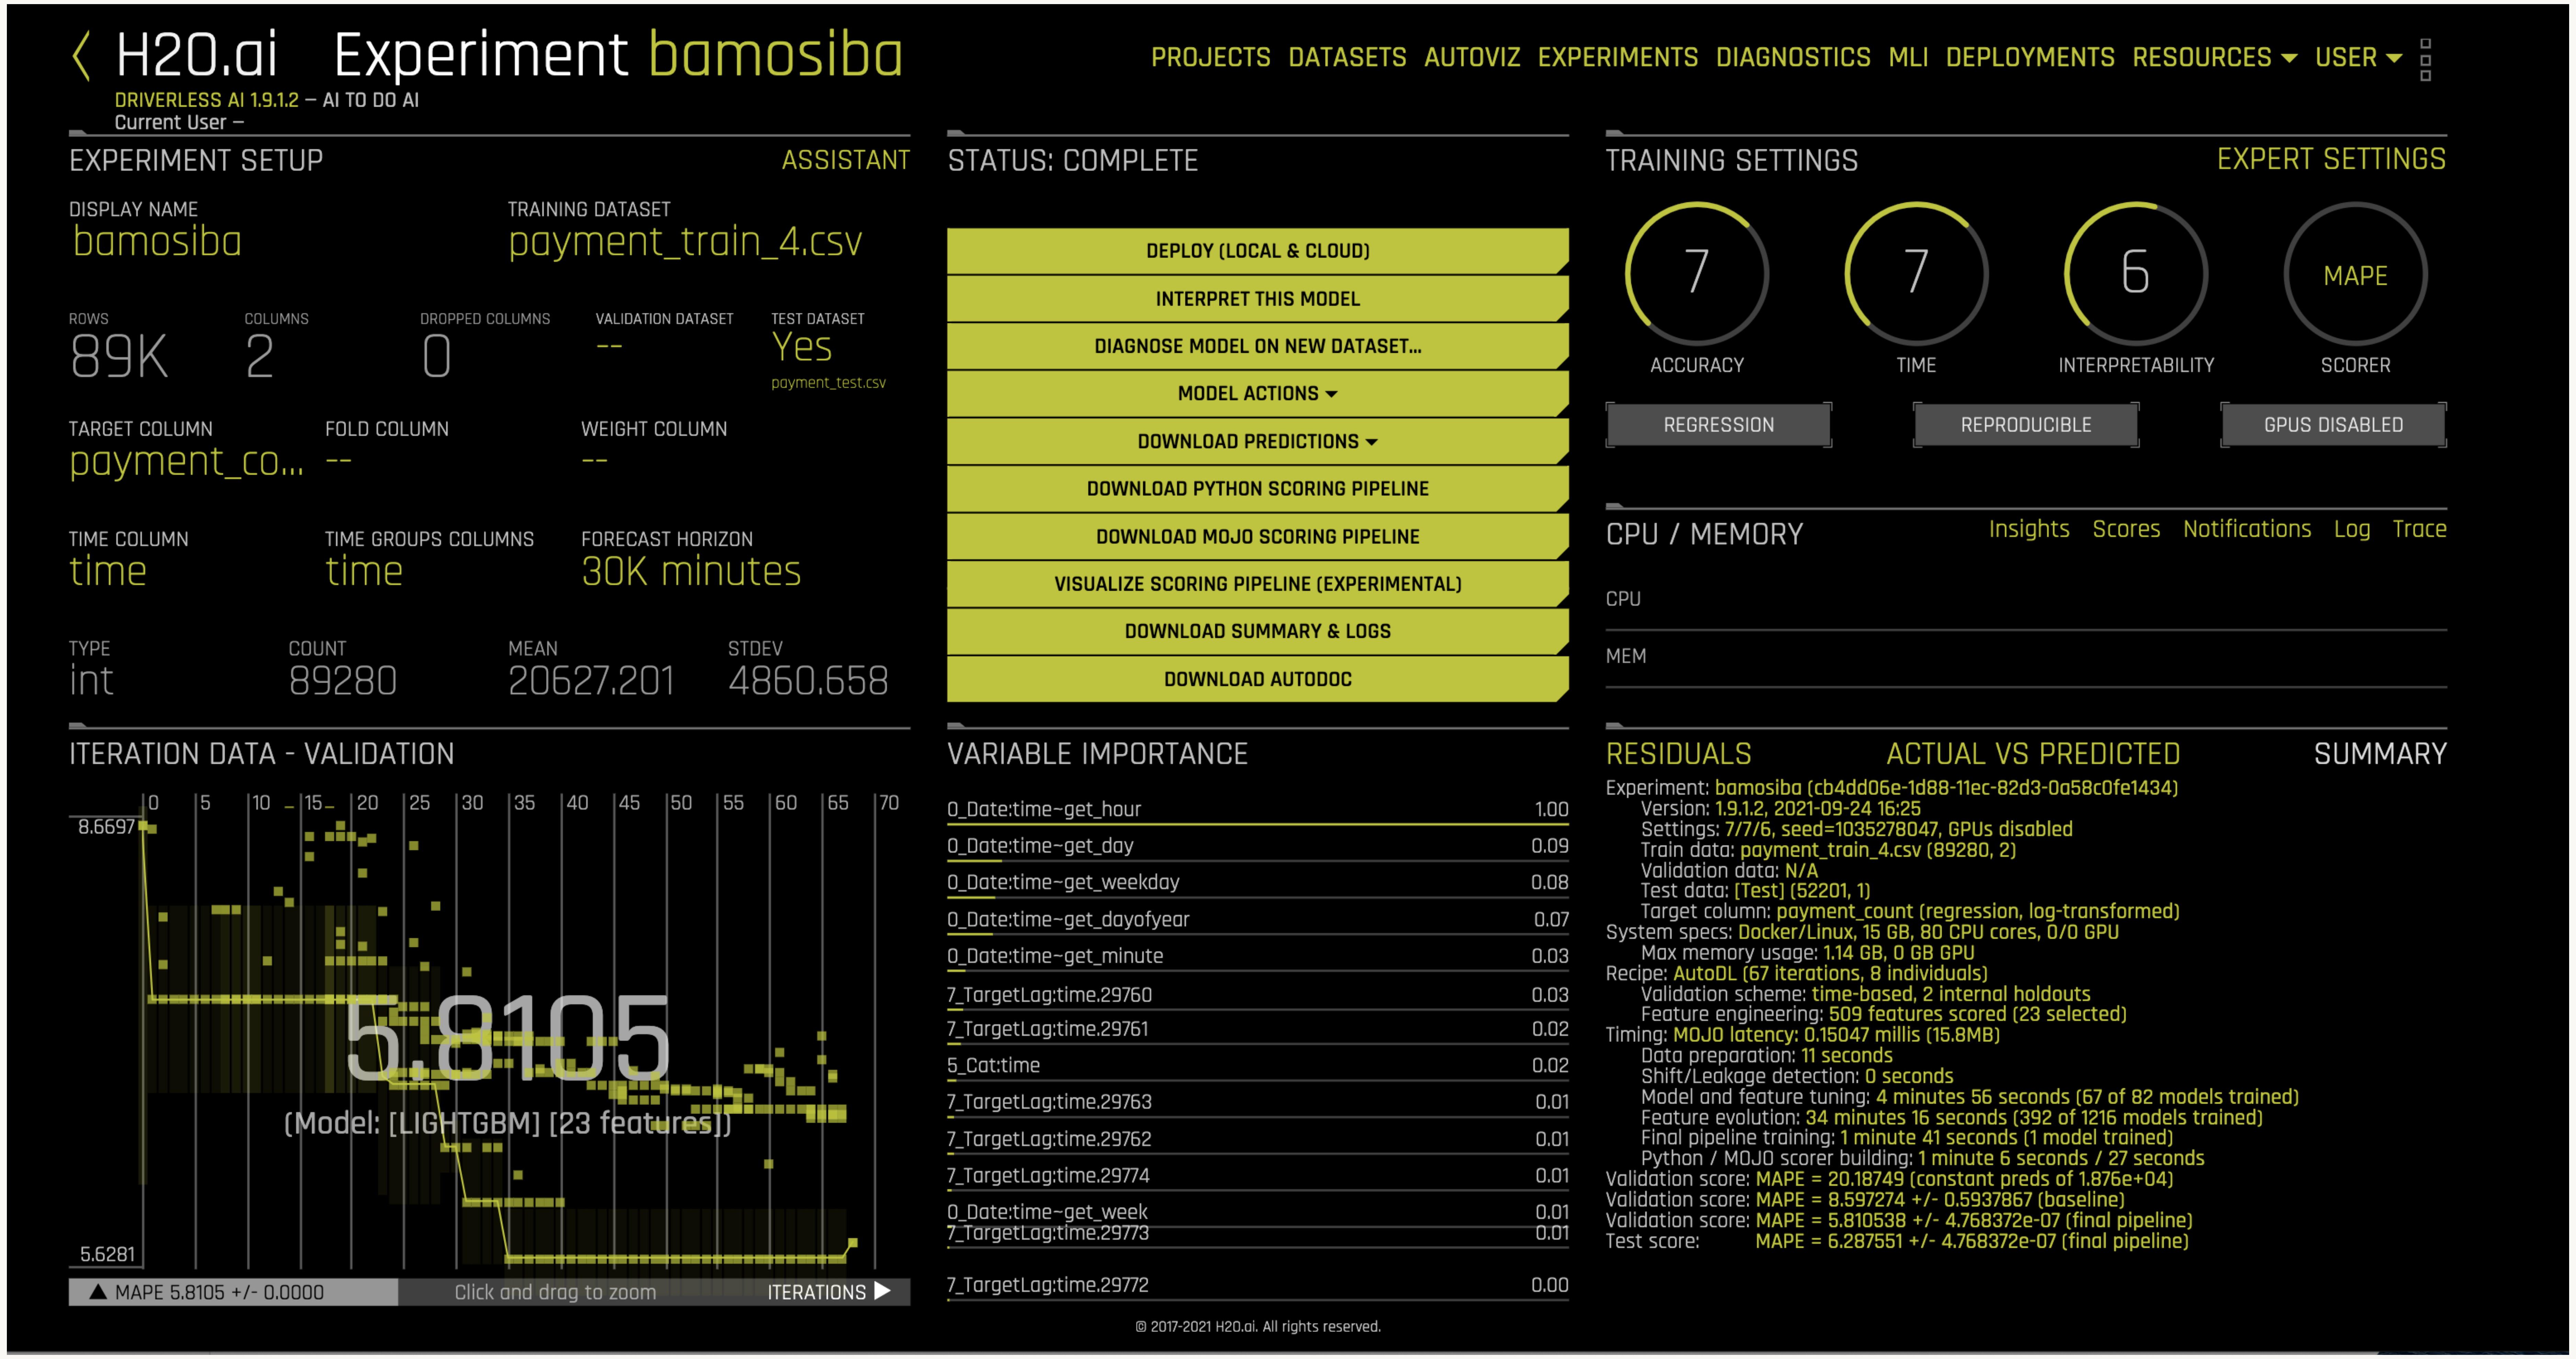

© 2022 PayPal Inc. Confidential and proprietary.

MATHEMATICAL MODELING

LightGBM Performance

How did the LightGBM model perform?

H2O (LightGBM) model still lags six-week trimmed average model. In forecast horizon even for first two days we see a significant difference.

For long intervals six-week trimmed average model performs better but only by a slight margin of 1-2 %.

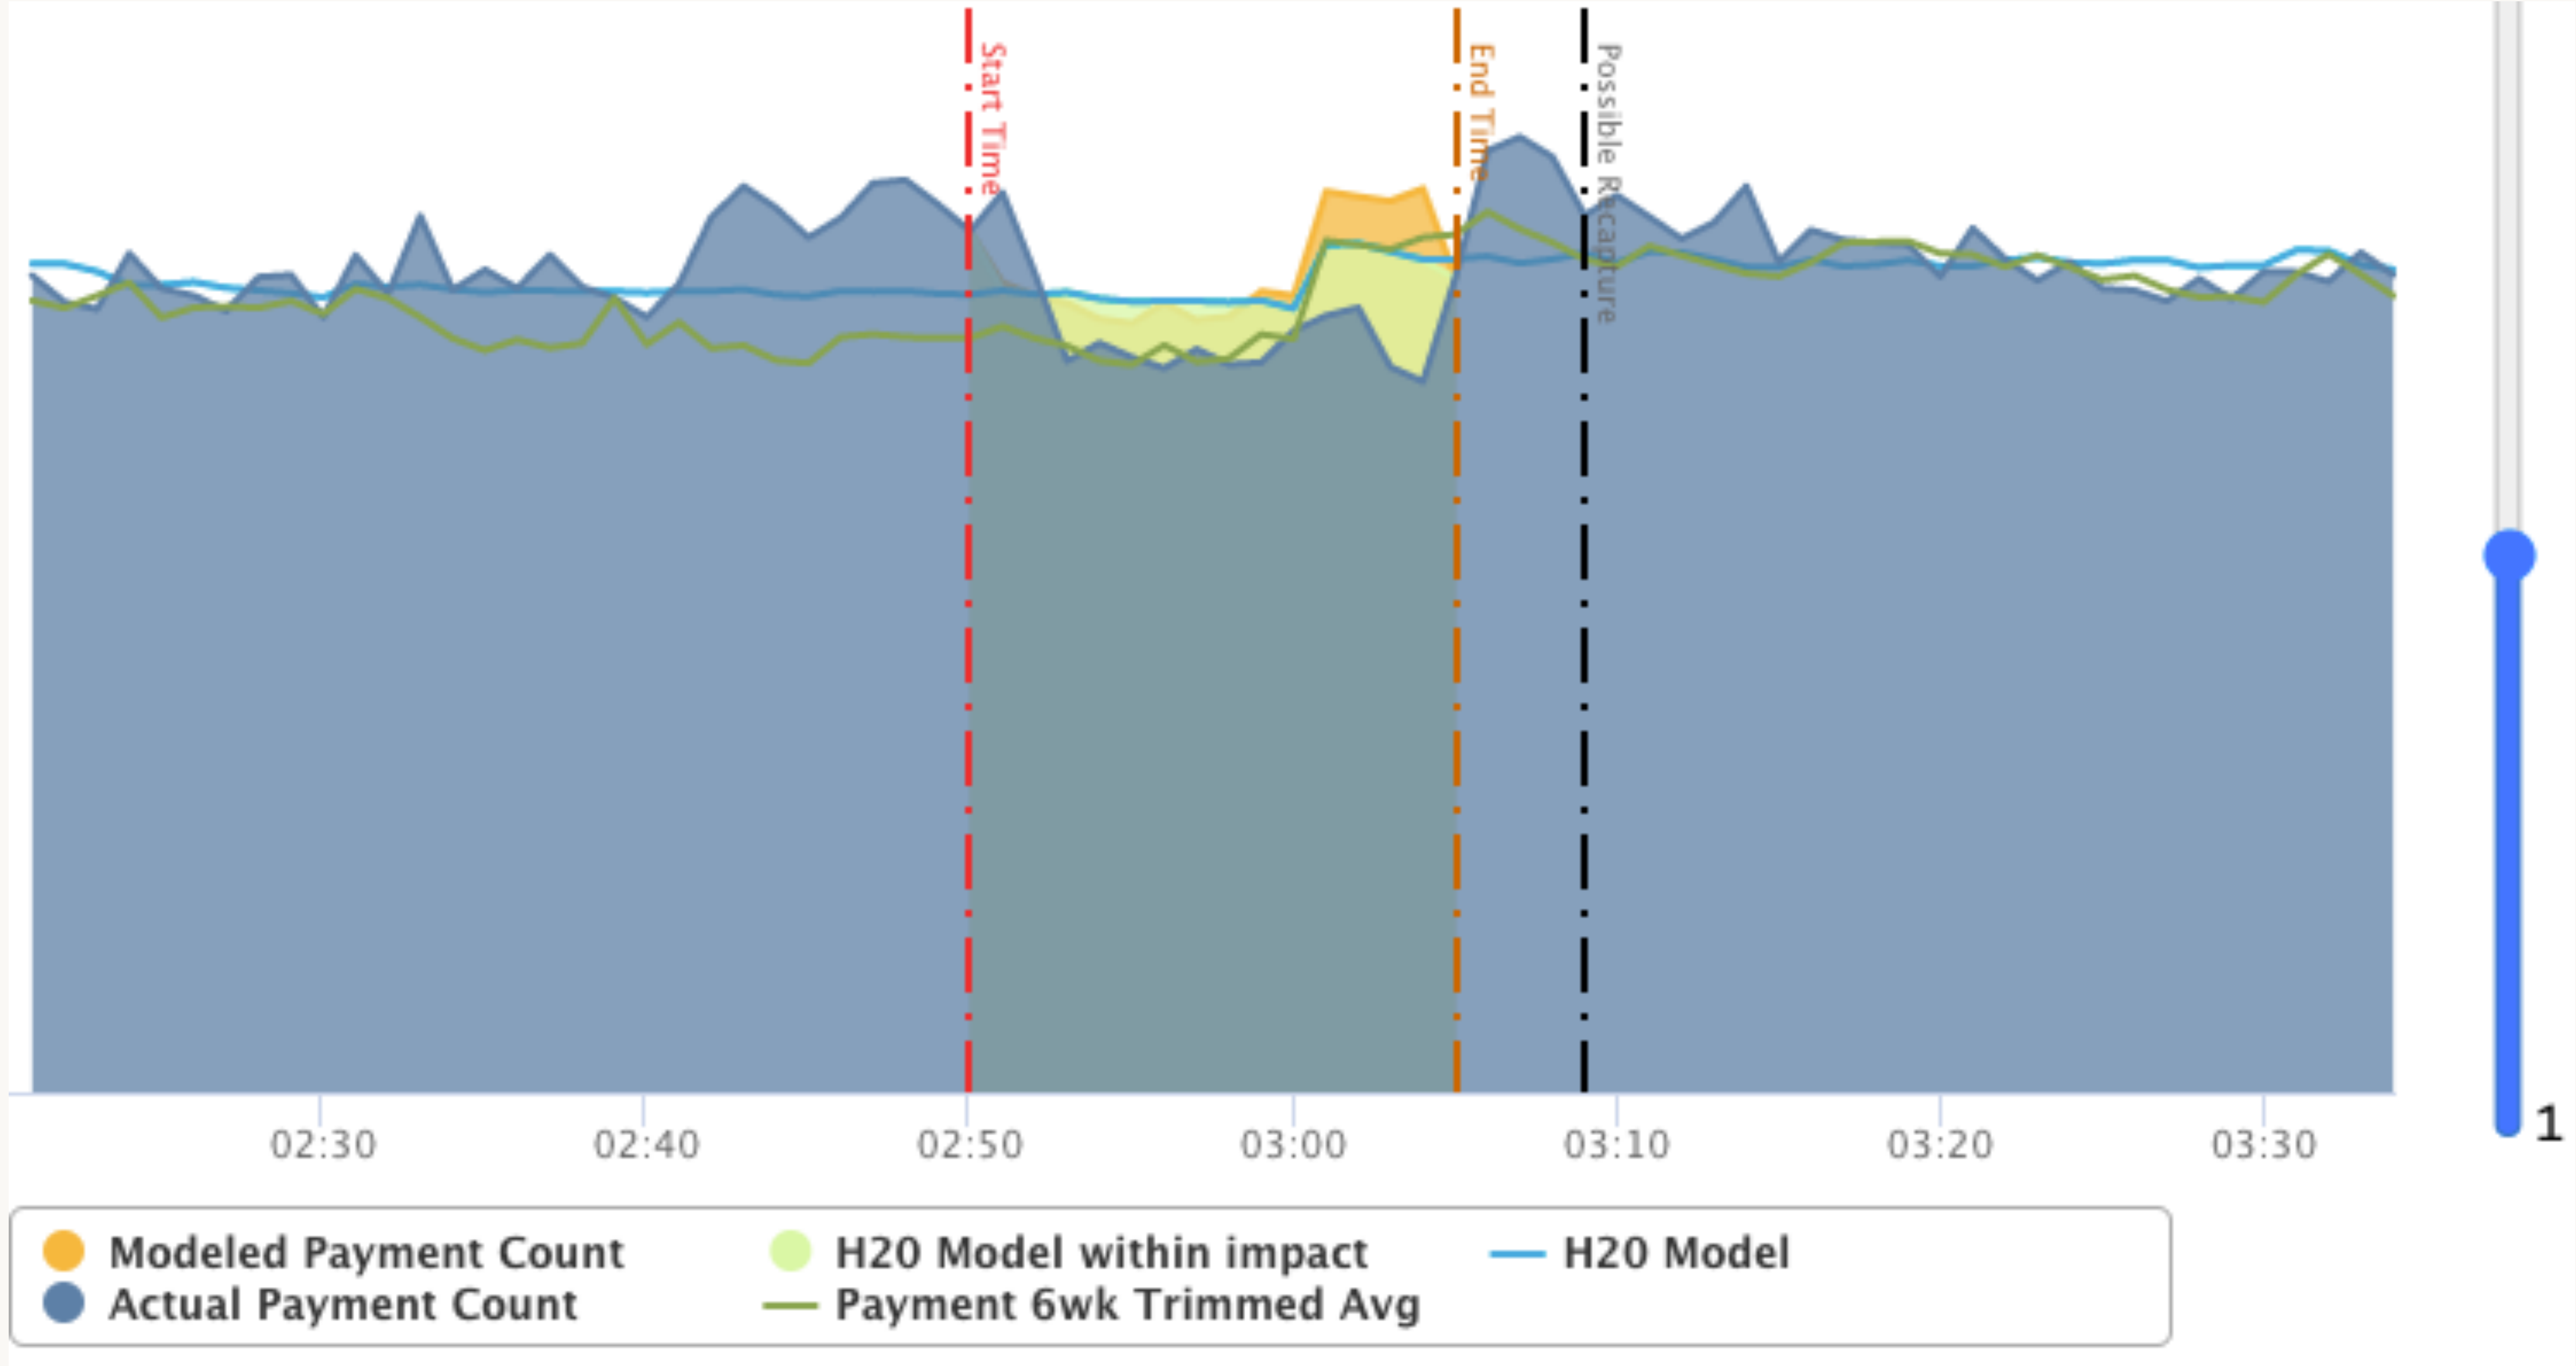

© 2022 PayPal Inc. Confidential and proprietary.

FEATURES & TAKE-AWAYS

On Capabilities & Take-aways

© 2022 PayPal Inc. Confidential and proprietary.

PayPal

FEATURES & TAKE-AWAYS
Availability

Availability

- To measure the availability loss during an incident (in addition to payment volume loss).
- Measured in count of requests or time.
- Formulae - uptime / (uptime + downtime), or successful requests / (successful requests + failed requests).
- Reported in a percentage like 99.9% or 99.999%.

Measured through FCI

- Failed Customer Interactions (FCI) - Number/percentage of intended actions that a Customer* is unable to complete using functionality offered by PayPal and allowed by PayPal policies.
- Customer* - Consumers, Merchants/Partners, anyone consuming the results of an interaction
- Functionalities are broken down into a set of interactions and the failures of these interactions are what we are going to measure.
- PayPal Availability mostly trends above >= 99.99%

© 2022 PayPal Inc. Confidential and proprietary.

FEATURES & TAKE-AWAYS
Availability Measurement

MEASURING FCI LOSS FOR A WINDOW

Sample Availability Loss

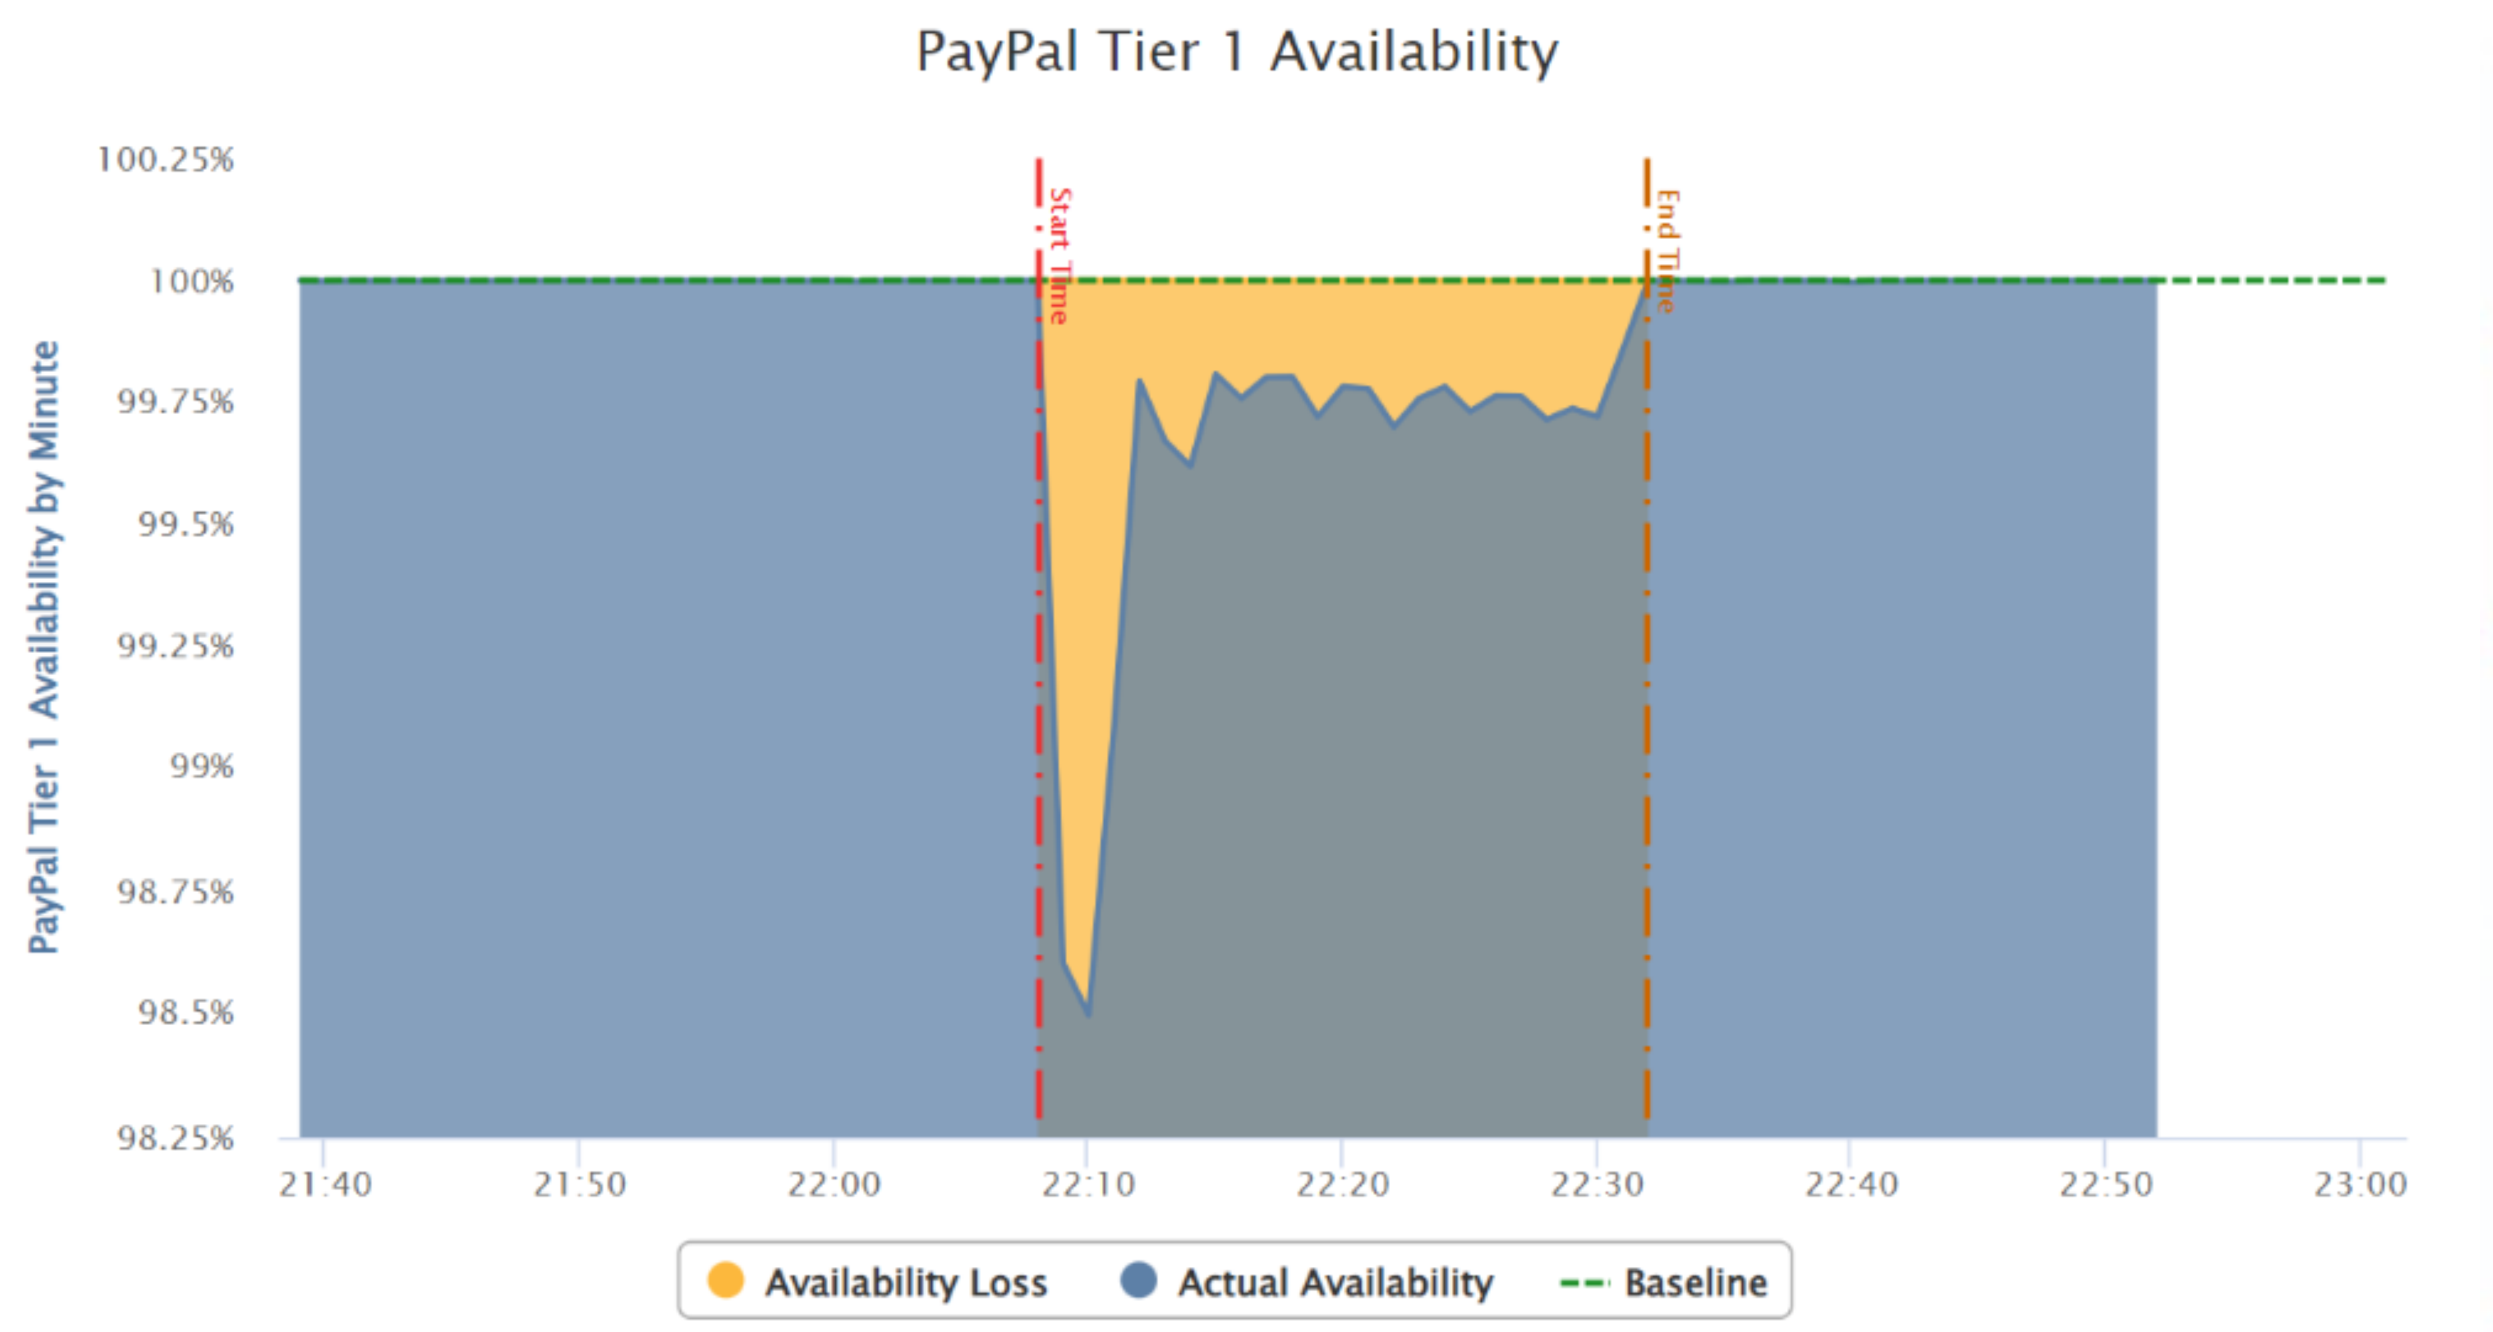

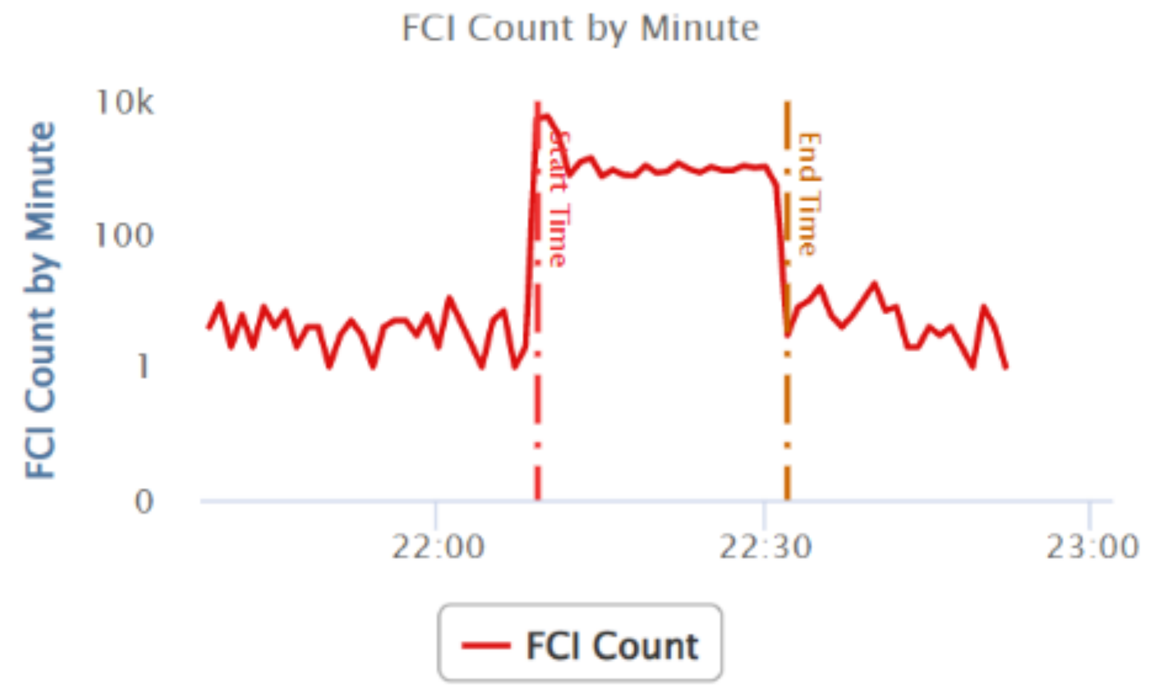

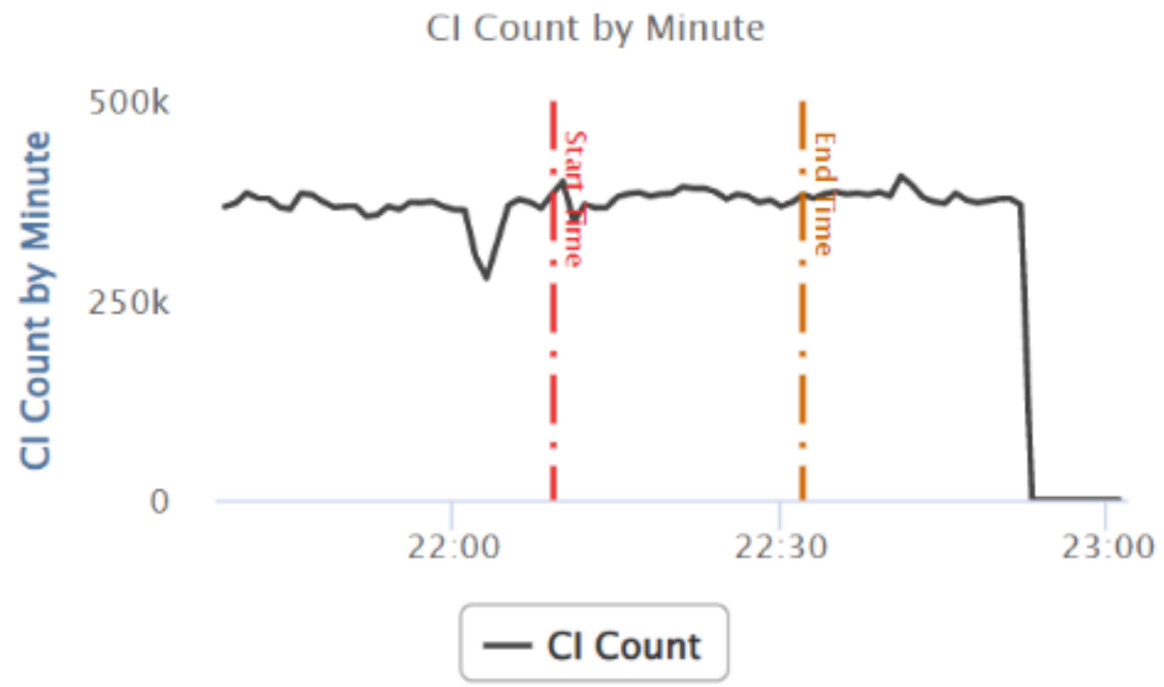

Baseline	99.9990%
Availability	99.6171%
FCIs	33,322
Missed CIs	0
FCI Impact	33,322

Representational Image

PayPal

© 2022 PayPal Inc. Confidential and proprietary.

FEATURES & TAKE-AWAYS
Segmentation

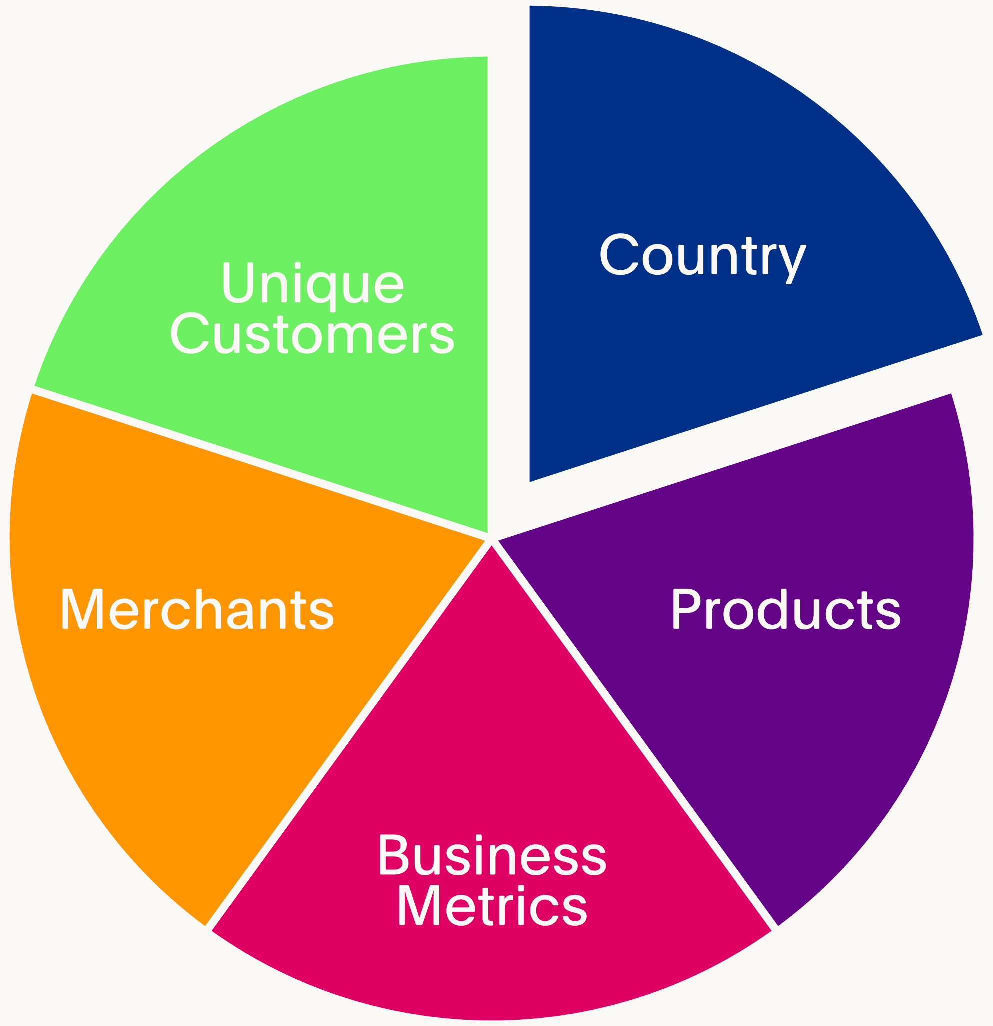

IMPACT BY COUNTRY

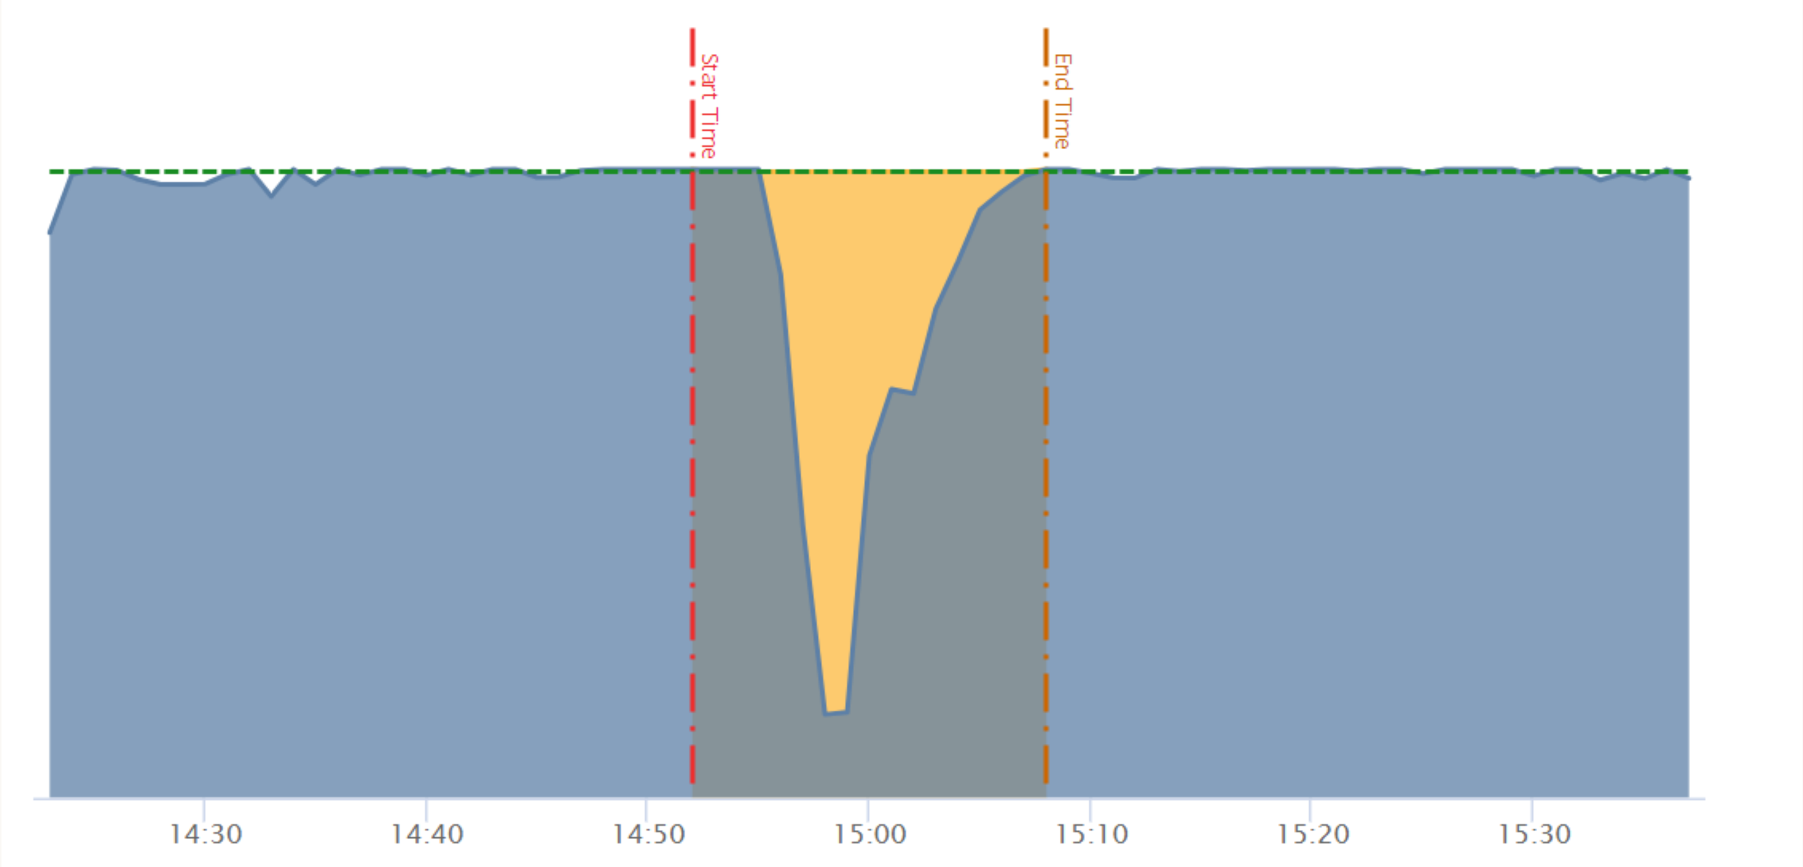

Country	Impact	% Of Impact
United Kingdom	352	30
Germany	313	27
France	143	12
Italy	82	7

Representational Image

- Segmentation by Country and Merchant are available.
- Segmentation computation using Machine Learning has been challenging when limited data is available.

© 2022 PayPal Inc. Confidential and proprietary.

FEATURES & TAKEAWAYS
BU Expansion

IMPACT FOR XOOM & BRAINTREE

Xoom and Braintree

xoom
A PayPal Service

Braintree
A PayPal Service

Framework for BUs

- Key Business Metrics (Transfers, Payments)
- Understanding of Products and Incidents
- BU Metric Data Store Integration
- Generic or Custom Models Deployed

© 2022 PayPal Inc. Confidential and proprietary.

FEATURES & TAKEAWAYS
Opportunities

What's the road ahead?

Superpose multiple models

Track Model use for learning

Calculate & post automatically

Suggest incident window

Expand to Other BUs like venmo

© 2022 PayPal Inc. Confidential and proprietary.

FEATURES & TAKEAWAYS
Auto-suggestions

Can the tool suggest me automatically?

MULTIPLIER USAGE & RECAPTURE END TIME

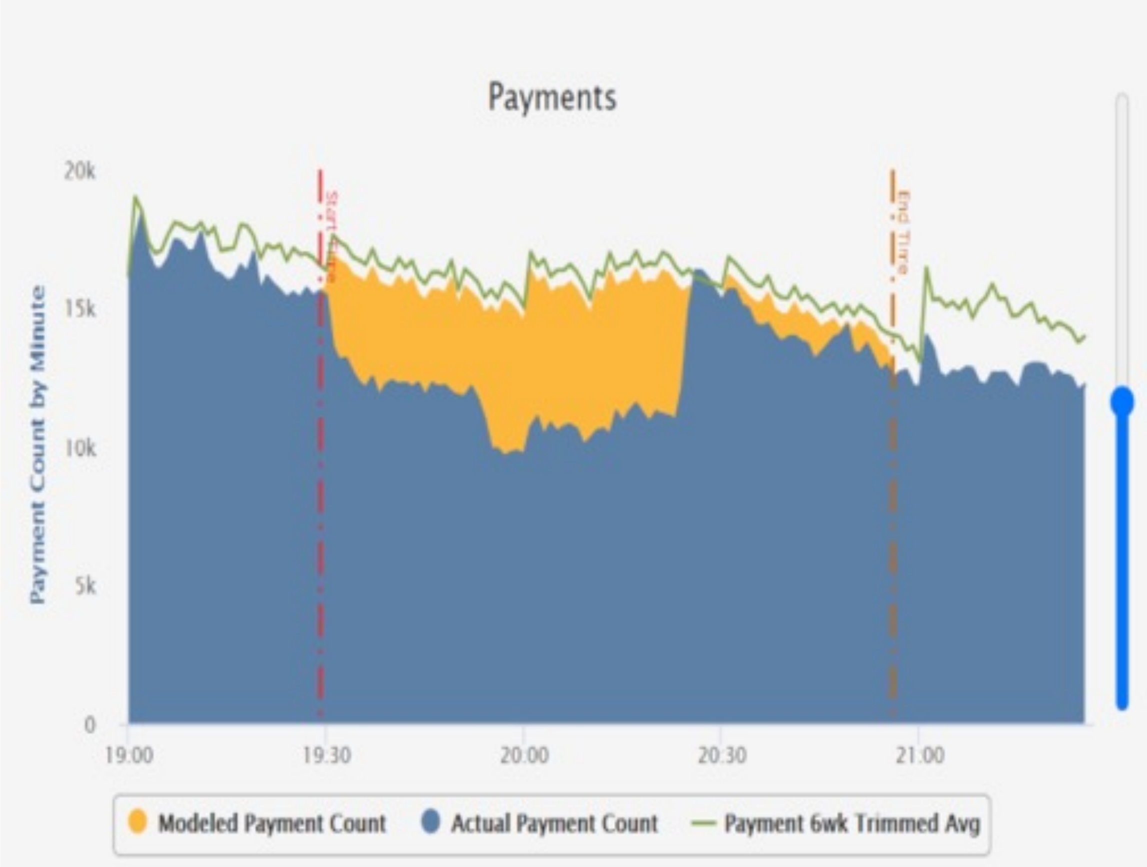

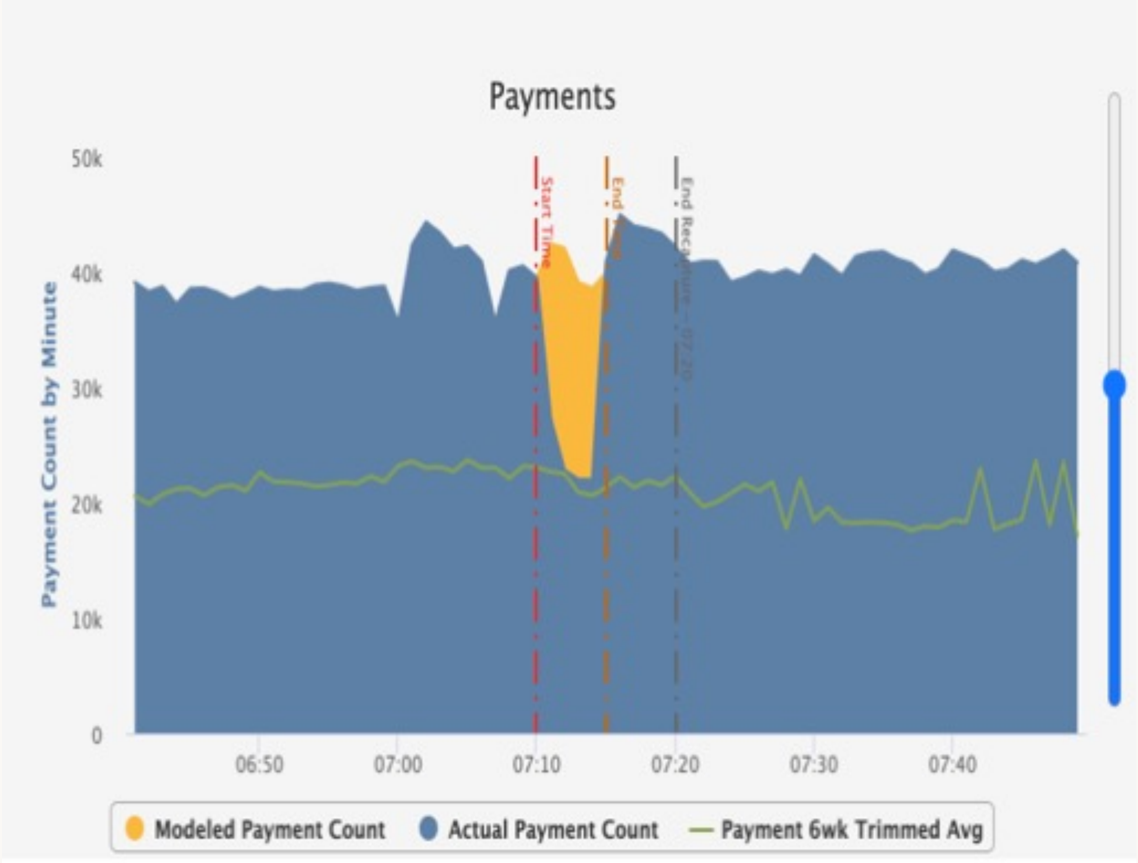

Representational Images

Can the tool auto-suggest the following?

- To use multipliers (as in the left diagram)?
- To re-capture end time (as in the right diagram)?
- Drop in Customer Interactions (CI) count?

© 2022 PayPal Inc. Confidential and proprietary.

FEATURES & TAKEAWAYS
Top Challenges

What challenges do we have?

Heavy Reliance on User	Tool relies heavily on the user for providing impact window, to apply visual reasoning etc.
Testing	Testing in general with a variety of use cases for modeling and segmentation.
Modeling	Machine Learning Modeling for segmentation and other use cases with very little data has been challenging.
Building Intelligence	We found that building auto-suggestions for CI drop, incident window suggestions is challenging.
Approximation	Data is aggregated at minute level. If the current window suffers high volume despite taking an impact, model can't find it correctly. Take rate across products may vary.

© 2022 PayPal Inc. Confidential and proprietary.

FEATURES & TAKEAWAYS

How to Start Building?

What are the considerations?

BUSINESS METRICS

- What are the key business metrics?
- Is it granular enough for segmentation?

DATA SOURCES

- Is your use case real-time?
- How would you access?

MANUAL ASSESSMENT

- How do you assess manually?
- Do we need to change the process?

MODELLING

- Do you understand dataset?
- What model would you use?

TEST & REMODELLING

- How do you test for various use cases?
- Is it consistent? Need remodeling?

INTEGRATION

- How do you integrate and run?
- Is your product working well?

© 2022 PayPal Inc. Confidential and proprietary.

Questions

© 2022 PayPal Inc. Confidential and proprietary.

Thank You

PayPal