



# ENGINE OPTIMIZATION POWERED BY BIG DATA

**GDC**

March 20–24, 2023  
San Francisco, CA

# HELLO GDC!



**VINCENT  
LECLERCQ**

Lead Programmer  
Ubisoft Paris



**VLADIMIR  
KATCHADOURIAN**

Senior Data Scientist  
Ubisoft Paris




March 20–24, 2023 | San Francisco, CA



SHIPPING A GAME


DOWNGRADES CUTS  
PERFORMANCE  
ISSUES STRESS BUGS



March 20–24, 2023

San Francisco, CA

GDC





A close-up, side-profile shot of a man with a dark beard and glasses, looking intently at a computer screen. The screen is the primary source of light, casting a soft glow on his face and the text overlaid on it. The background is blurred, showing what appears to be a dimly lit room with other people in the distance.

# ANALYSIS ARE (MOSTLY) HANDMADE BY A FEW (VERY BUSY) EXPERTS

IT SHOULD BE A SHARED RESPONSABILITY  
BETWEEN THE TEAM MEMBERS



March 20-24, 2023 | San Francisco, CA

GDC


# DATA ANALYTICS AS A SOLUTION ?



**DATA ANALYSIS**




**MACHINE  
LEARNING MODEL**



**INVESTIGATION  
PROCESSES**



# SUMMARY



March 20–24, 2023 | San Francisco, CA







THE DIVISION 2

# THE BEGINNING OF PERFORMANCE DATA ANALYSIS



GDC



# THE DIVISION 2 : CONTEXT



Multiplayer  
Action RPG



Open  
World



PVP, PVE,  
PVPVE



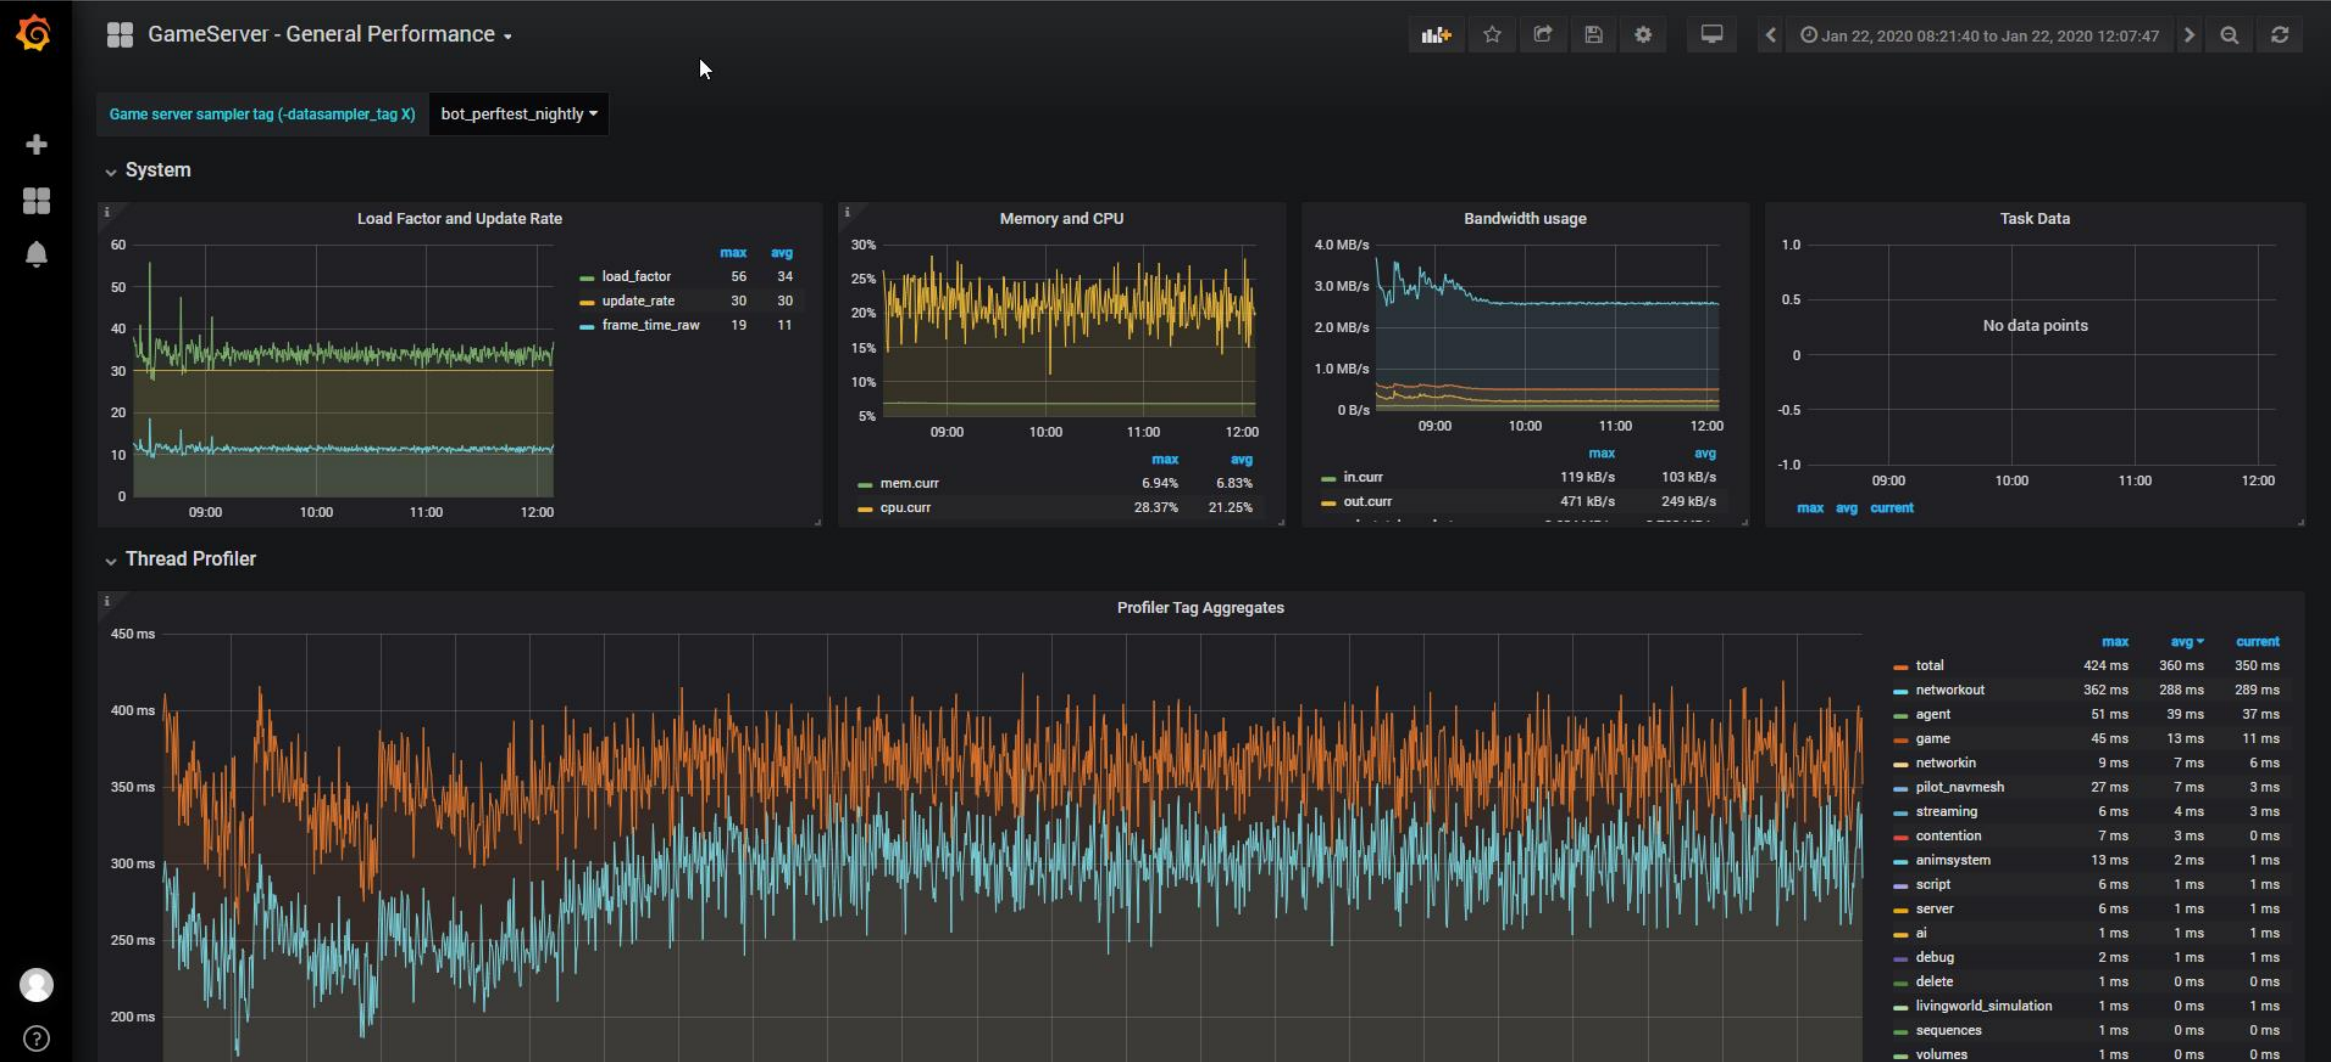
March 20–24, 2023 San Francisco, CA





# TOO HIGH-LEVEL INFORMATION

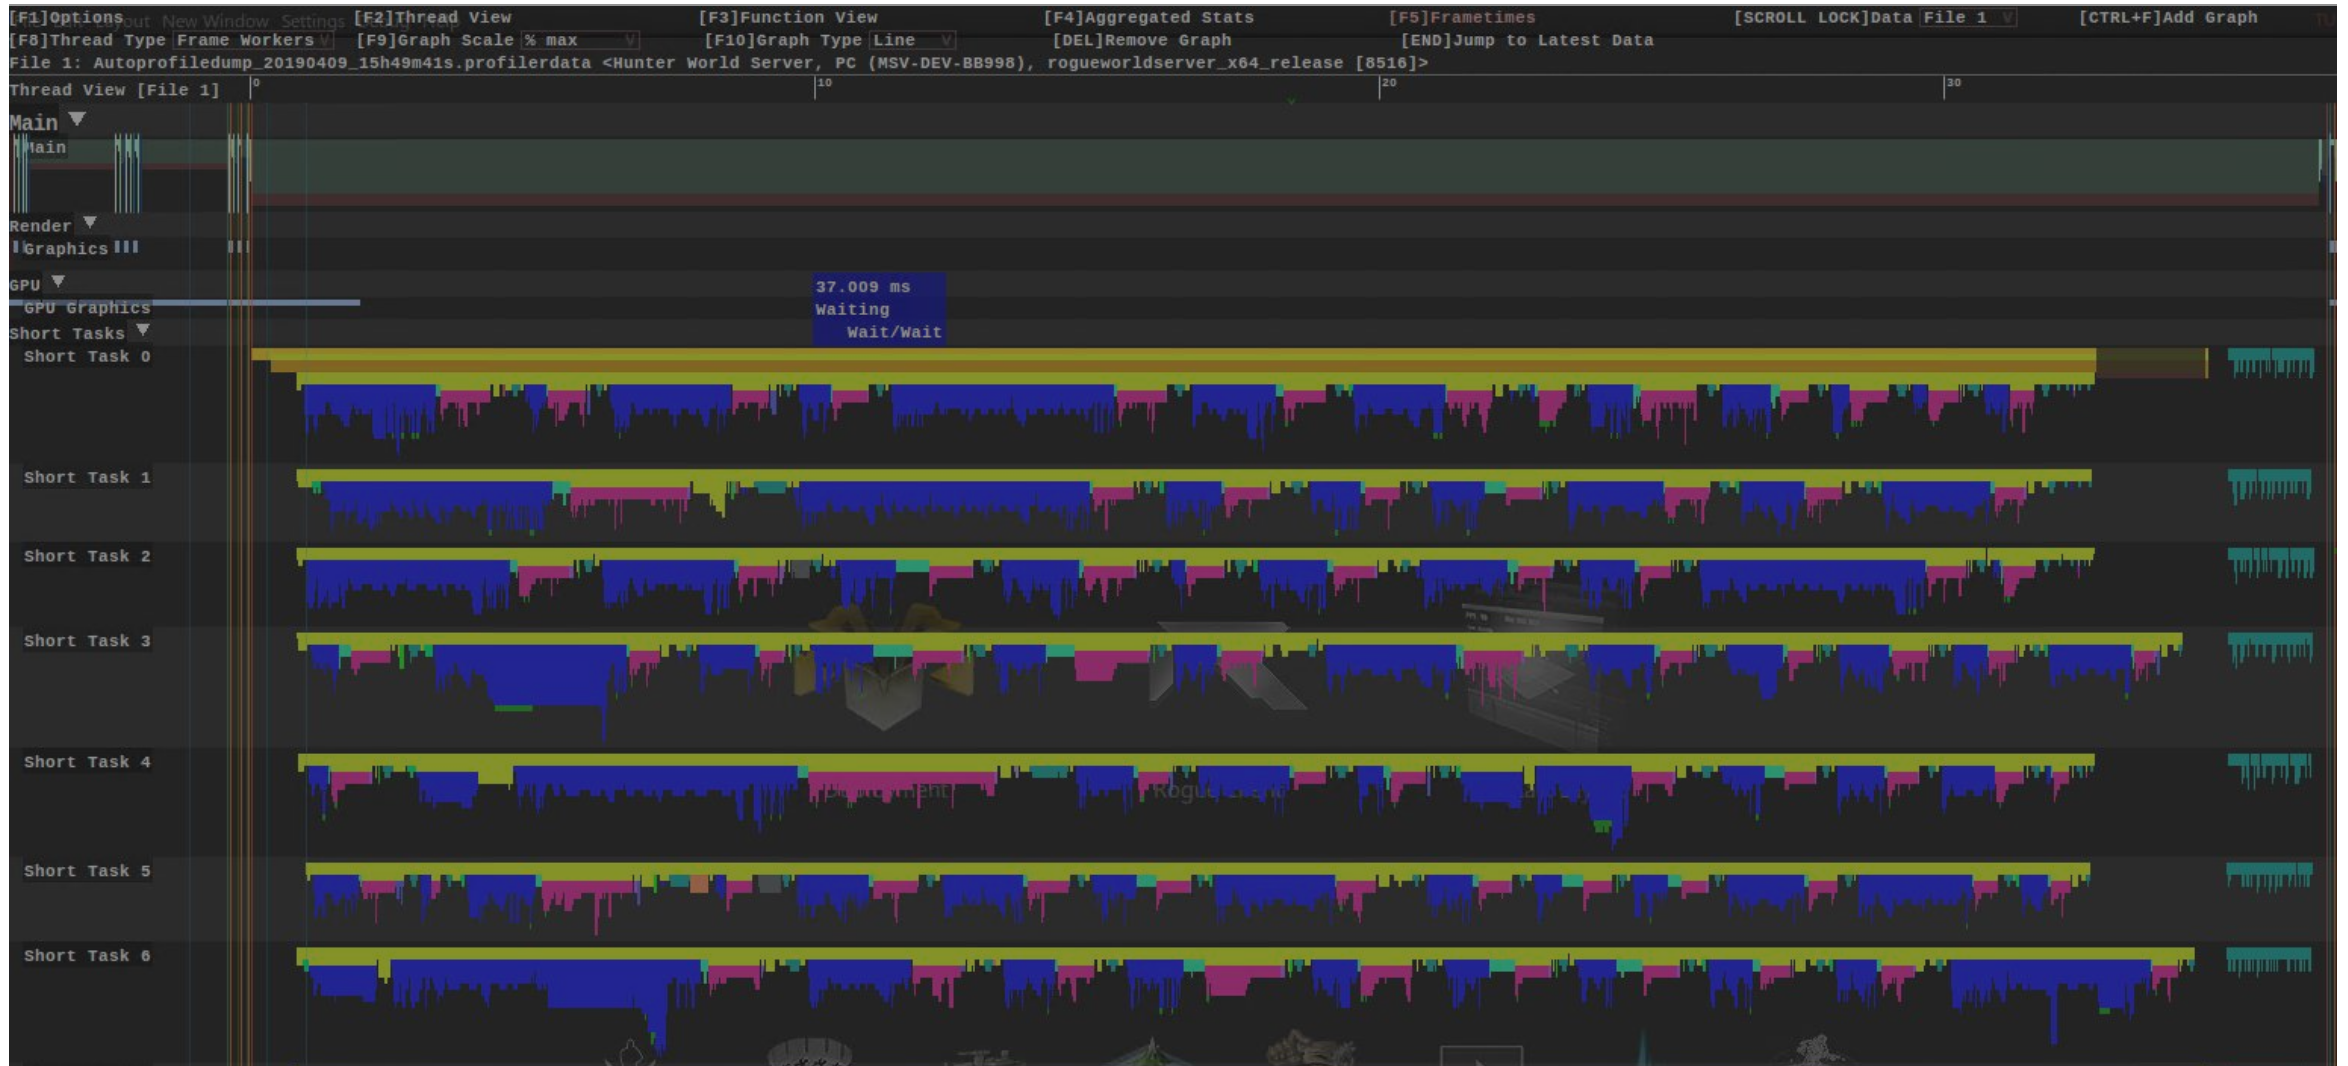
## SERVER MONITORING





# TOO LOW-LEVEL INFORMATION

## PROFILER DATA





# LIFE OF AN EXPERT





A cinematic scene showing two soldiers in tactical gear and backpacks standing in a doorway. The doorway is framed by thick, dark green vines and leaves. Bright sunlight streams in from the doorway, creating a strong lens flare and illuminating the soldiers and the foliage. The soldiers are holding rifles and looking out into the bright light.

# IS THERE A BETTER WAY?

A DATA ANALYTICS SOLUTION  
BETWEEN HIGH & LOW LEVEL

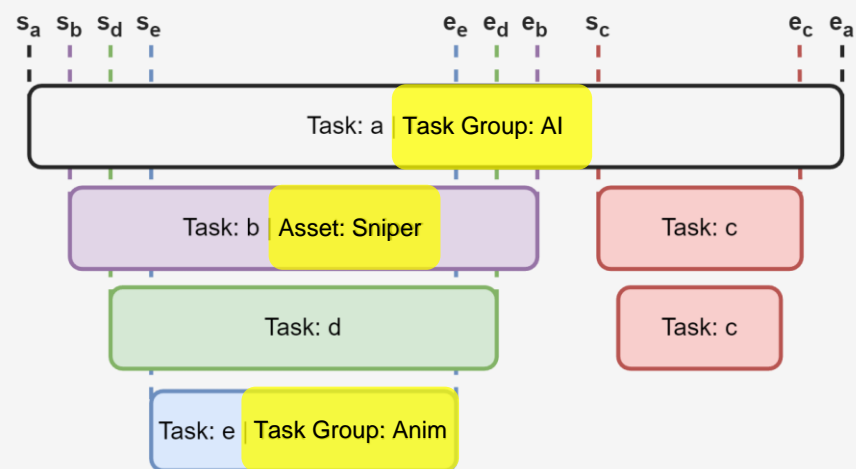


March 20-24, 2023 San Francisco, CA

GDC



# FROM PERFORMANCE PROFILING FILE...

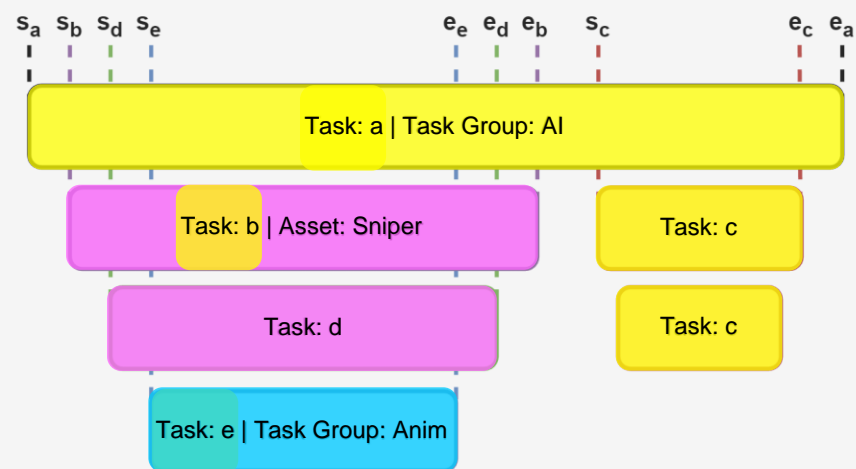


# ...TO ANALYTICAL TABULAR DATA

Task ID	Task Name	Task Group	Asset Name	First task name	Path	Is_first_task group	Is_first_asset	n_recursions	Root Task
1	a	AI		a	a	True	False	0	a
2	b	AI	Sniper	b	b  d	False	True	0	a
3	c	AI		a	a  c	False	False	1	a
4	d	AI	Sniper	b	b	False	False	0	a
5	e	Anim	Sniper	e	e	True	False	0	a



# FROM PERFORMANCE PROFILING FILE...



# ...TO ANALYTICAL TABULAR DATA

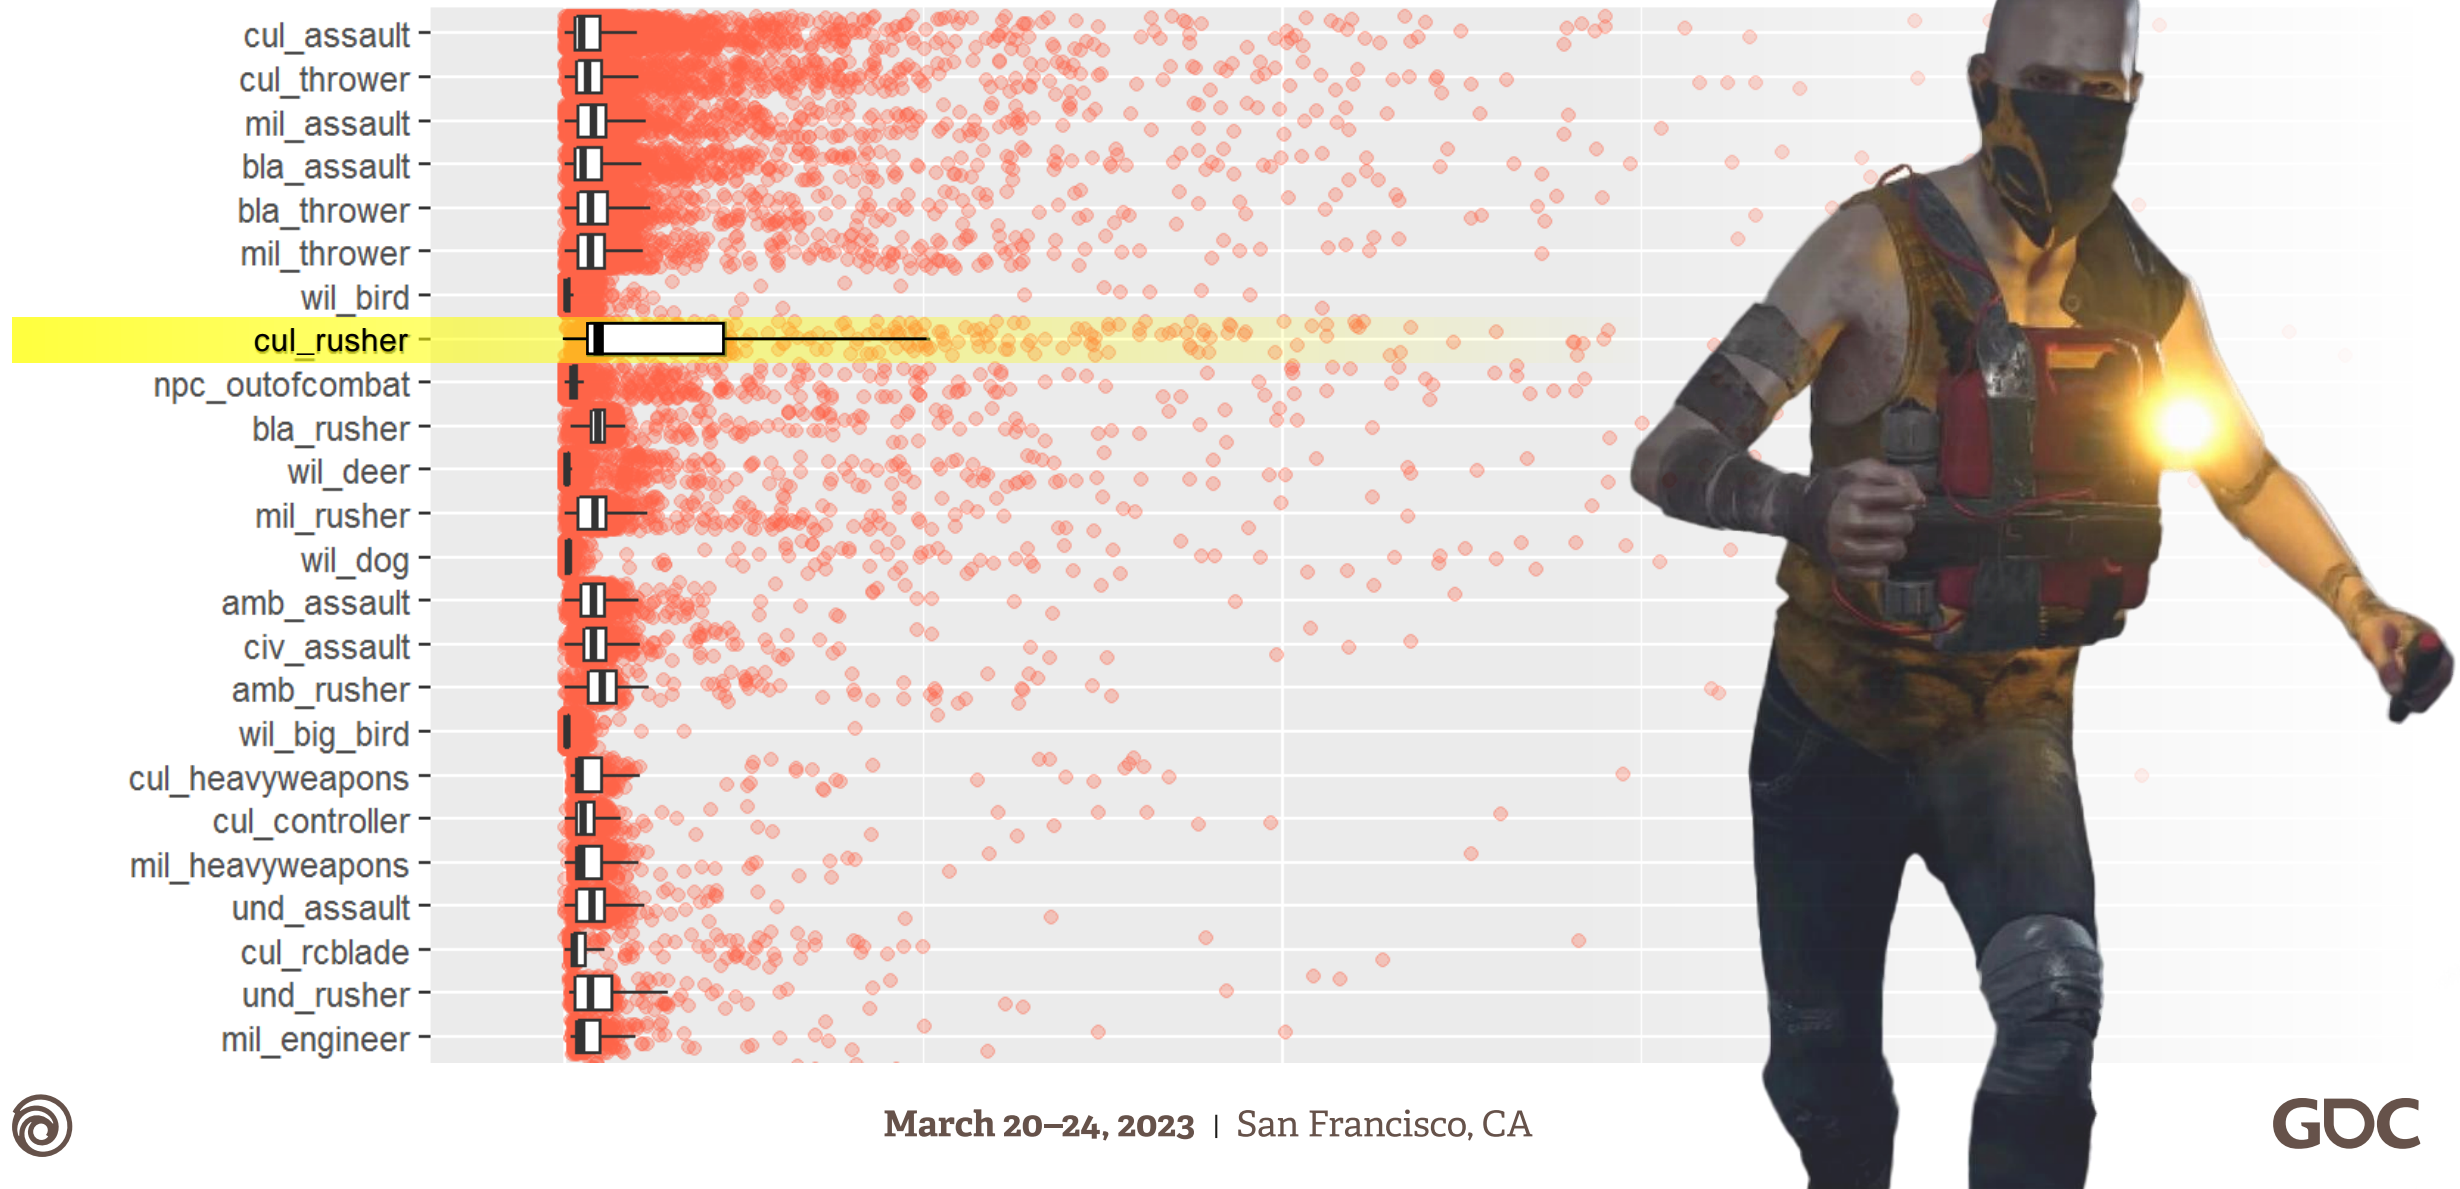
Task ID	Task Name	Task Group	Asset Name	First task name	Path	Is_first_task_group	Is_first_asset	n_recursions	Root Task
1	a	AI		a	a	True	False	0	a
2	b	AI	Sniper	b	b d	False	True	0	a
3	c	AI		a	a c	False	False	1	a
4	d	AI	Sniper	b	b	False	False	0	a
5	e	Anim	Sniper	e	e	True	False	0	a





# CULTIST RUSHER ASSET ISSUE DETECTION


Behavior Update Duration per Asset



March 20–24, 2023 | San Francisco, CA

GDC

# THE DIVISION 2: CONCLUSION



We Still need  
an expert but  
**it's promising**

Objective:

**900**  
**PLAYERS**  
PER SERVER



March 20–24, 2023 | San Francisco, CA

GDC



GHOST RECON BREAKPOINT

# INDUSTRIALIZATION & PERFORMANCE ANALYTICS PIPELINE

From prototype to new performance analytics tool



GDC



# GHOST RECON : CONTEXT



Multiplayer



Openworld



PVE, PVP



March 20-24, 2023 San Francisco, CA

GDC



## NEXT STEPS



**1**  
AUTOMATIZE  
REPORTING



**2**  
DEMOCRATIZE  
INVESTIGATION




March 20–24, 2023 | San Francisco, CA

**GDC**

# THE PIPELINE

## GENERATE DATA

Daily Auto-tests  
Daily Manual tests  
Playtest sessions





## Continuous profiling

Record 1 sec every 20 seconds  
Add files when slow frames are detected





# DATA PROCESSING



## File level transform

- Deserialization
- Inheritance
- Frame-wide aggregation

## Session level transform

- Session-wide aggregation





## Analytics Database

Exploration,  
Iteration, R&D



# ANALYZE & VISUALIZE

- R&D, iteration, exploration
- Production agnostic applicative usage



Output is saved in 2 formats



## Applicative Database

Data services  
& exposition







# AD-HOC ANALYSIS

TO ANSWER SPECIFIC QUESTIONS



March 20–24, 2023 | San Francisco, CA



# REGRESSION INVESTIGATION

## Ad-Hoc analysis example 1

### REGRESSION SPOTTED AFTER VERSIONS MERGE



- Auto test detection
- Historical data available

### INVESTIGATION ON DATA SYSTEMS

	asset_name	q50	q50_var
	TGT_SearchLight_BaseEntity	96090	+710%
	TGT_SearchLight_BaseEntity	11861	+710%
	TGT_SearchLight_BaseEntity	99432	+709%
	TGT_SearchLight_BaseEntity	12281	+709%
	TGT_SearchLight_BaseEntity_NoPole_Vertical	97643	+705%
	TGT_SearchLight_BaseEntity_NoPole_Vertical	12117	+705%
	TGT_SearchLight_BaseEntity	13633	+637%
	TGT_SearchLight_BaseEntity	100589	+637%
	TGT_SearchLight_BaseEntity_NoPole_Vertical	13449	+635%
	TGT_SearchLight_BaseEntity_NoPole_Vertical	98952	+635%
	TGT_SearchLight_BaseEntity_NoPole_Vertical	11793	+625%
	TGT_SearchLight_BaseEntity_NoPole_Vertical	85562	+625%
	TGT_SearchLight_BaseEntity	56009	+289%
	TGT_SearchLight_BaseEntity	14387	+289%

### OFFENDING ASSET IDENTIFIED




**+700%**  
**EXECUTION TIME**  
F(Search\_Light)



# REGRESSION INVESTIGATION

Ad-Hoc analysis example 1




# ANTICIPATE NEW FEATURE COST

Ad-Hoc analysis example 2



## AI TEAMMATES

Estimated asset's CPU cost  
& FPS impact



March 20–24, 2023 | San Francisco, CA







# STANDARD DASHBOARDS

TO MUTUALIZE THE BENEFITS.




March 20–24, 2023 | San Francisco, CA



## SELECT A SESSION

### Task Group Time per Profiling Session



**March 20–24, 2023** | San Francisco, CA

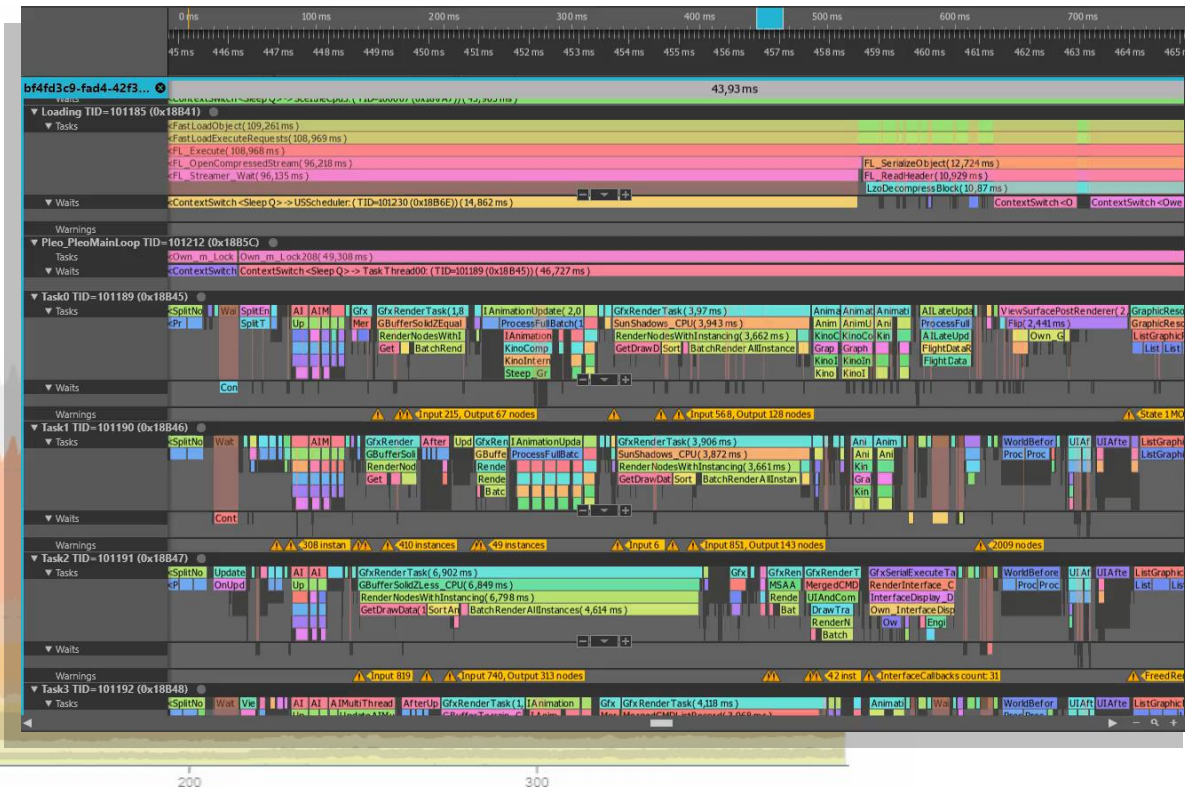


GDC



# ANALYZE A SESSION

## Task Group Time per Frame

Average distribution



## Selected Info



No Frame Selected

Select a frame above to view detailed information and open it in Gear Studio

Copy Link

Download File

Open in Gear Studio



March 20–24, 2023 | San Francisco, CA



# STANDARD DASHBOARD

## ANALYSE SESSION: TASKS



Sorted by cost



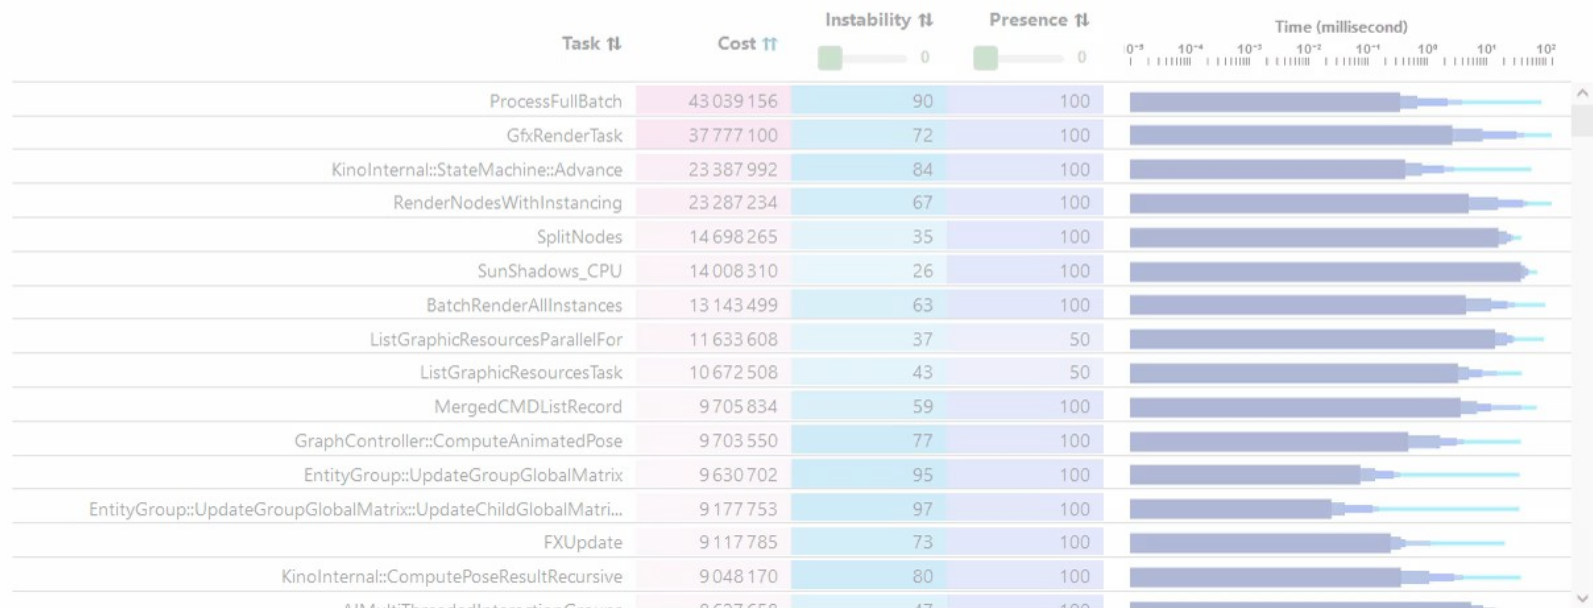
March 20–24, 2023 | San Francisco, CA





# STANDARD DASHBOARD

## ANALYSE SESSION: TASKS




March 20–24, 2023 | San Francisco, CA

GDC

# STANDARD DASHBOARD

ANALYSE SESSION:  
DATA SYSTEMS



Spread

q95 Spread

15.42%

## Stack Distribution

Data

Graph Scales


Average Time ☒ Max Value

Logarithmic ☒ Linear

tag\_name

average time (ms)

WAIT::WaitForG3_DIMalloc_Lock	0.548804
KinoInternal::CrossFadeTransitionInsta	0.358499
KinoInternal::BlendTreeInstance::Relea	0.355197
KinoInternal::StateMachine::Advance	0.142057
IAnimationUpdate::Item	0.131148



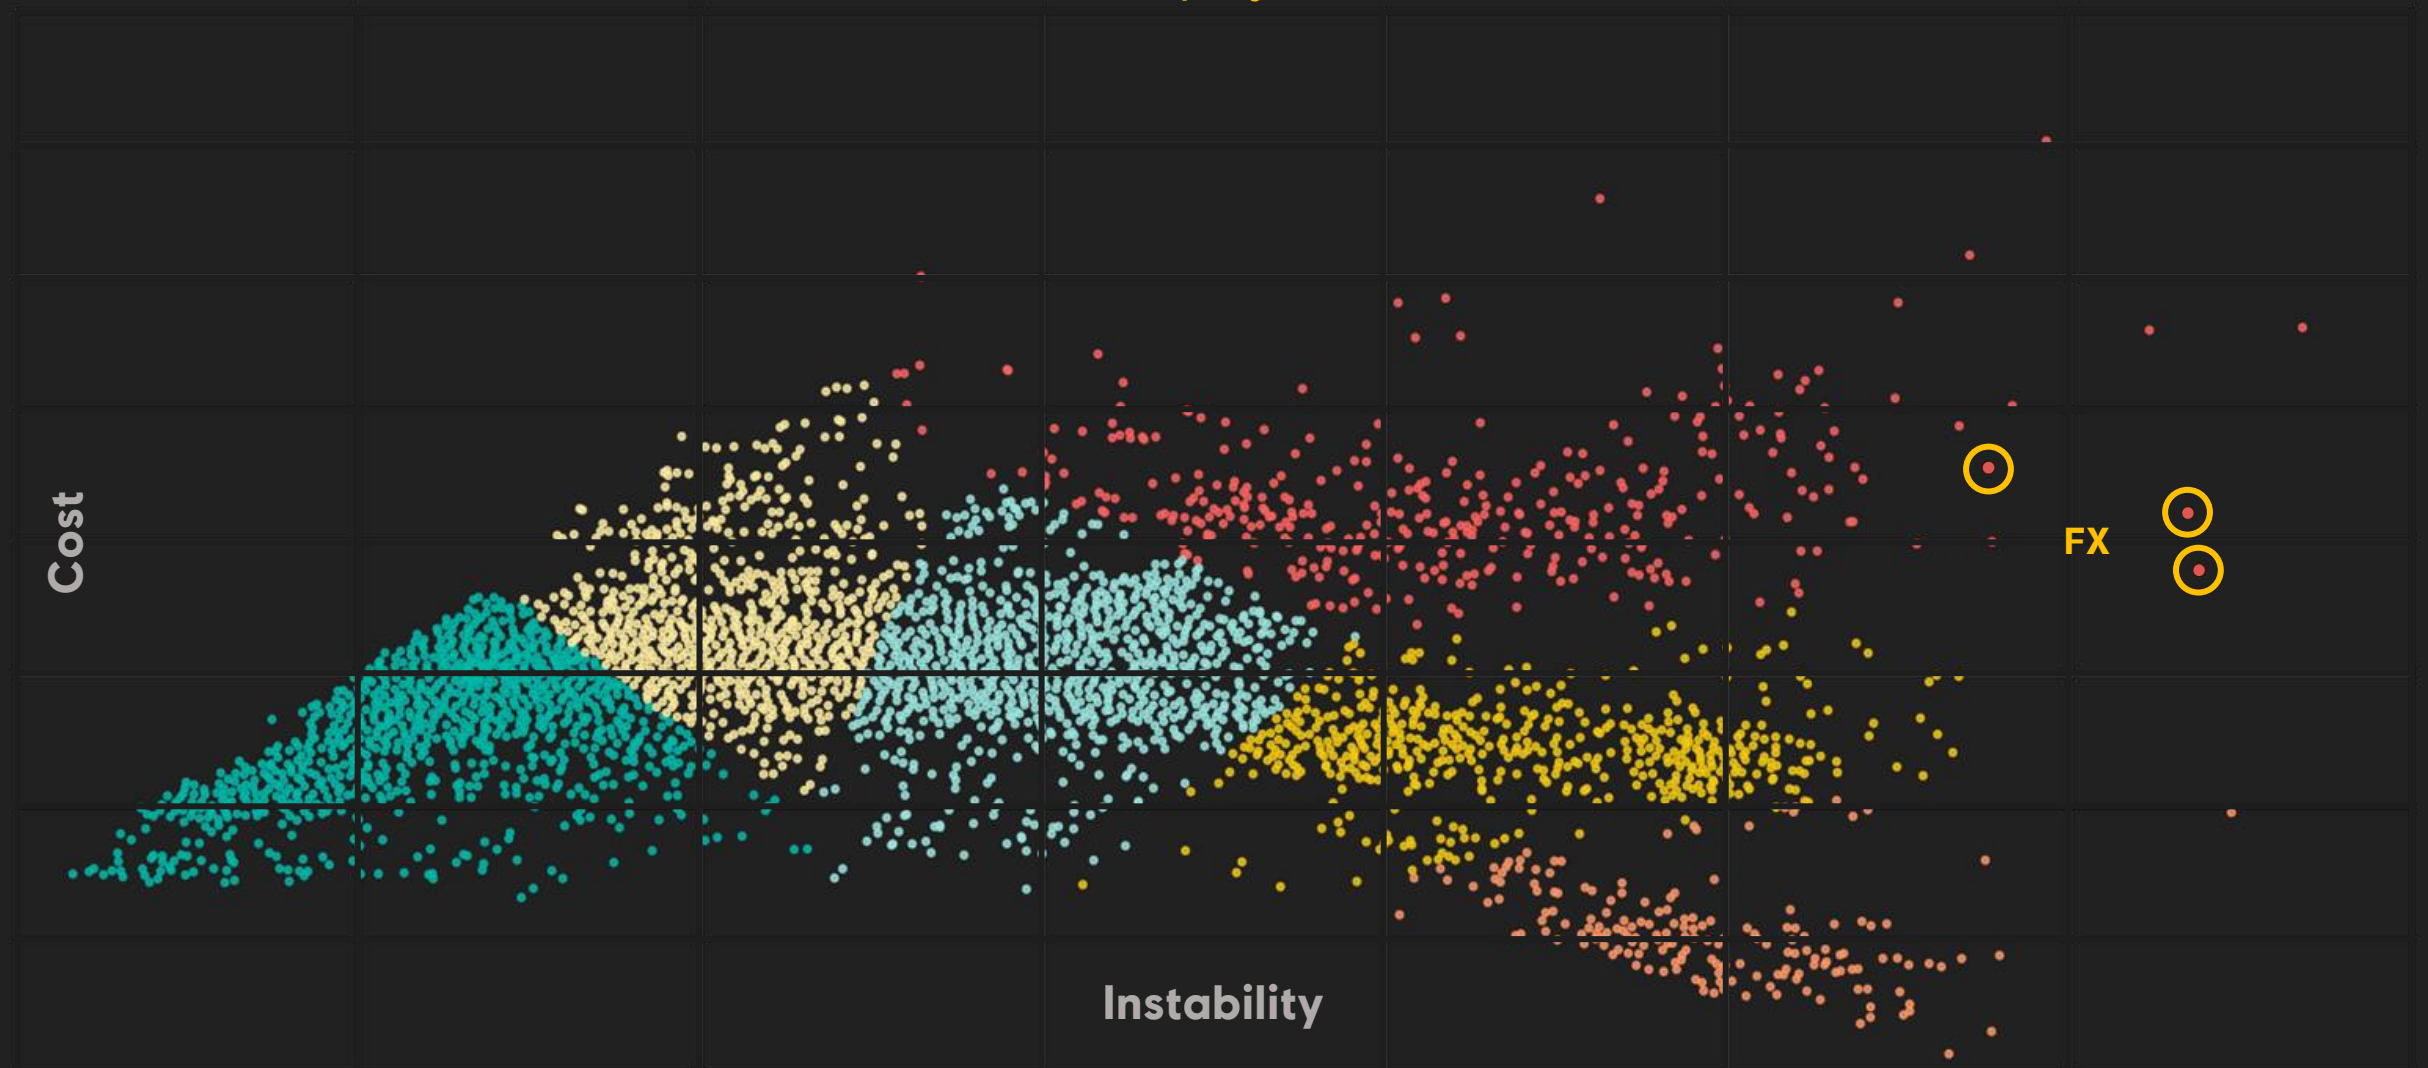
March 20–24, 2023 | San Francisco, CA

GDC



# FX ISSUES


Assets projection



March 20–24, 2023 San Francisco, CA


GDC

# FX ISSUES SOLVED



**100+**

**USELESS  
UPDATE  
REMOVED**



**25%**

**EXECUTION  
TIME SAVED**



**75%**

**BETTER STABILITY  
ACHIEVED**







RIDERS REPUBLIC

# MACHINE LEARNING USING CPU PROFILING DATA

How to save 50% of experts time using ML





# RIDERS REPUBLIC : CONTEXT



Anvil engine  
/Scimitar



Open  
World



Easy  
transition

March 20-24, 2023 San Francisco, CA

GDC



# DASHBOARDS ARE GOOD FOR MONITORING:




Builds health



Engine behaviors



Tasks



Assets



**BUT ARE LIMITED TO PRIORITIZE  
OPTIMISATION EFFORTS**

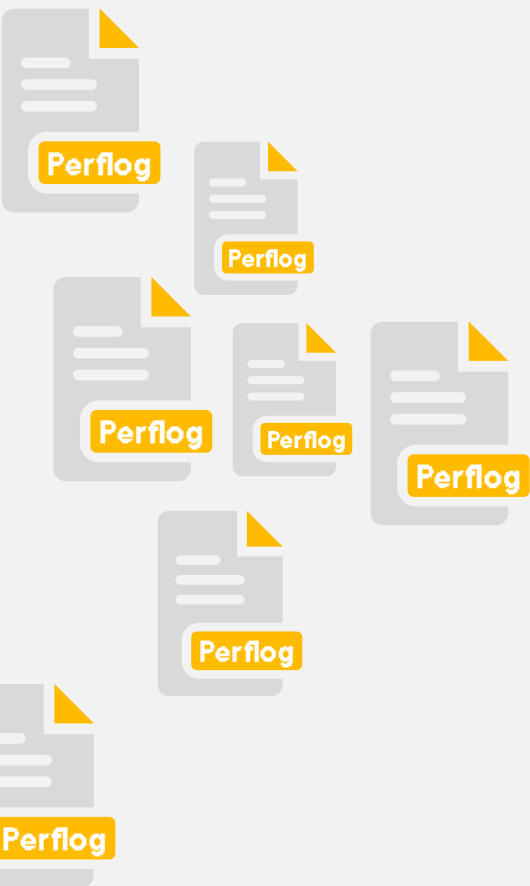


March 20–24, 2023 | San Francisco, CA

**GDC**

# EXPERTS ARE BOTTLENECK

and overwhelmed



We need to ease  
& automate part of their job



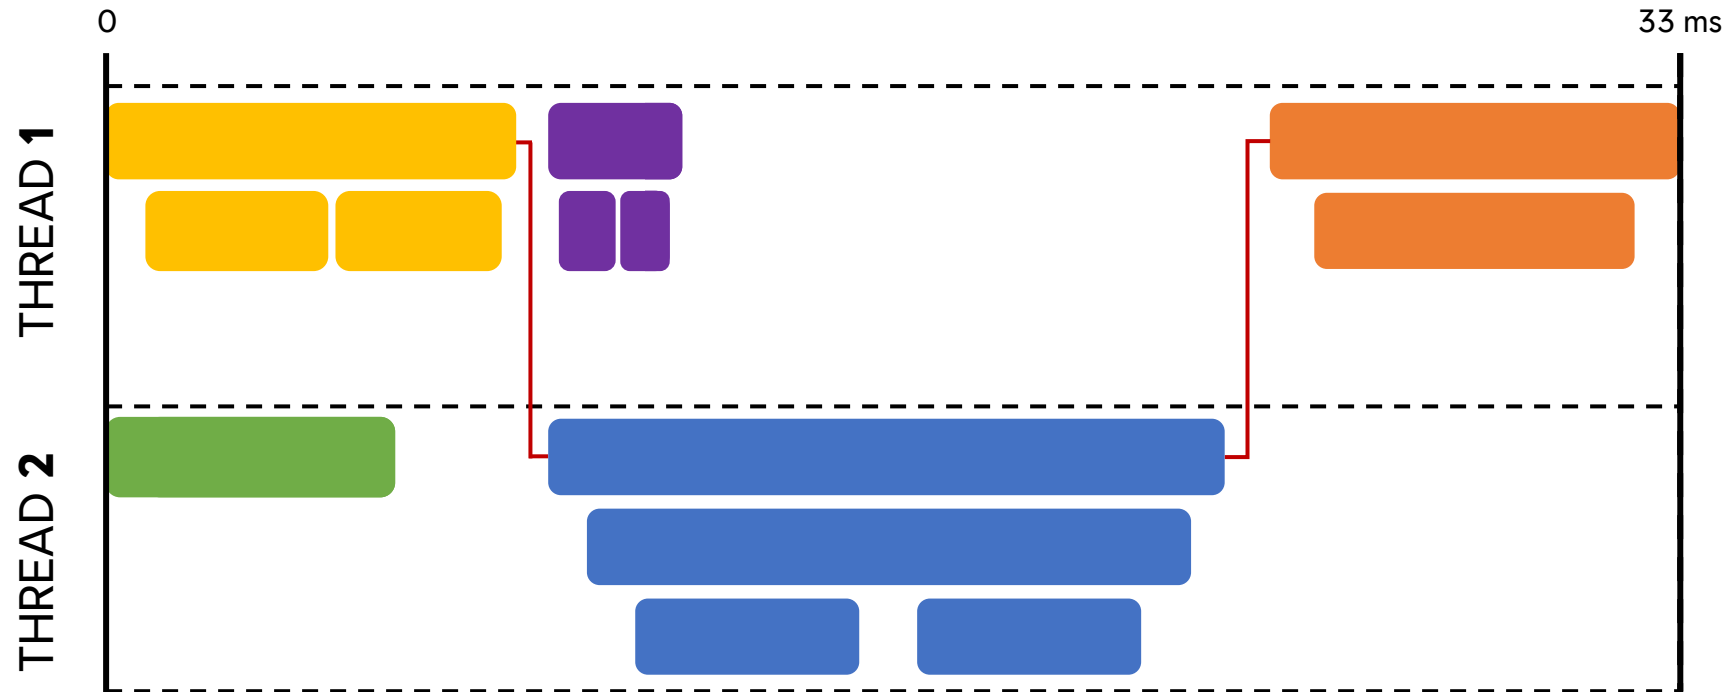
March 20–24, 2023 | San Francisco, CA

GDC



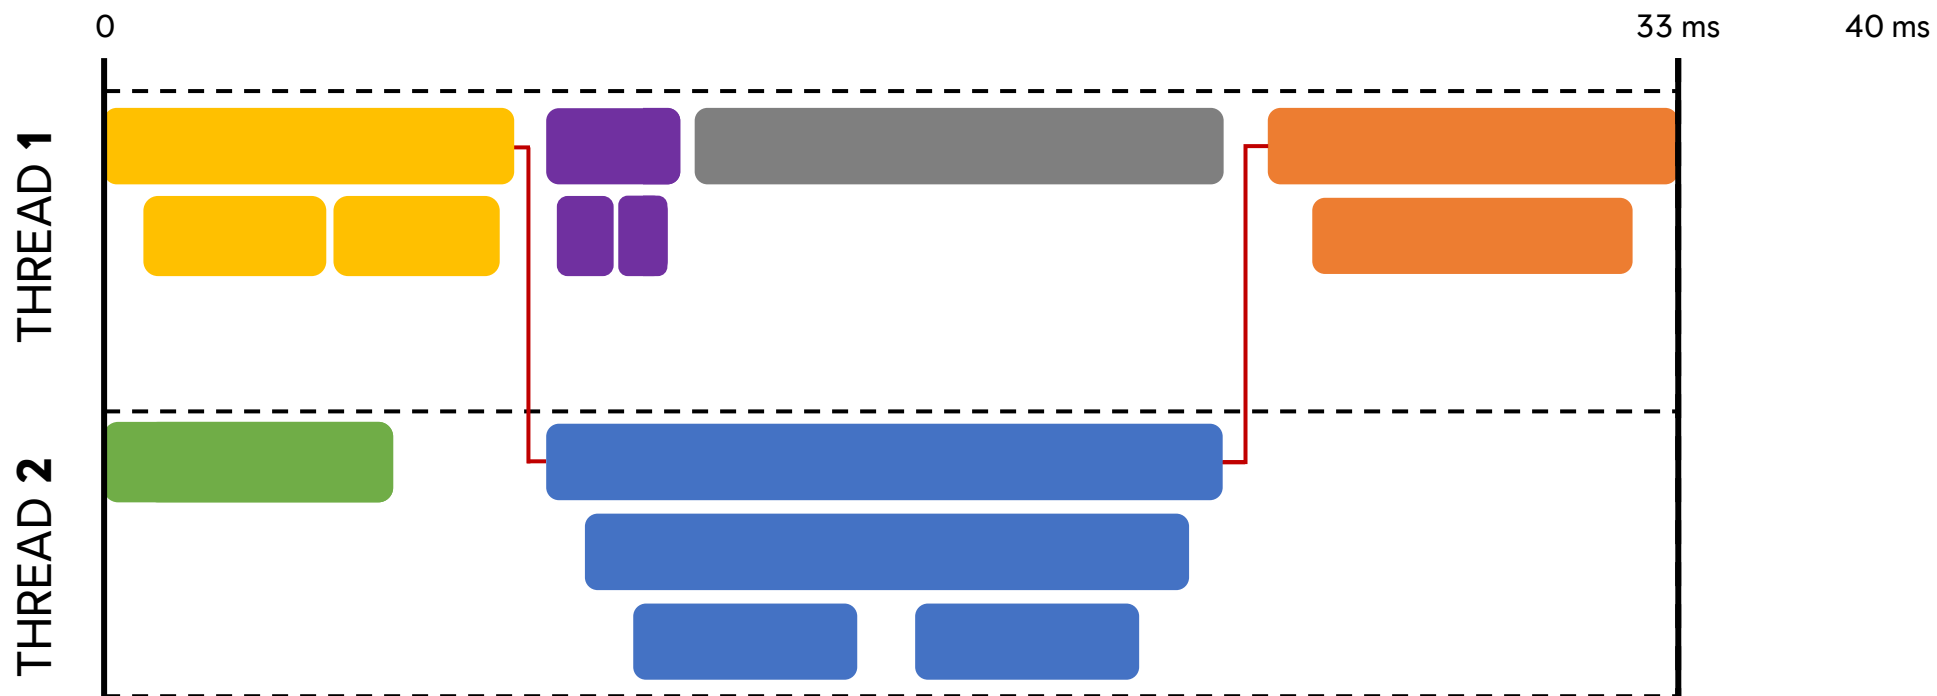
# INSIDE A CPU FRAME

The problem



# INSIDE A CPU FRAME

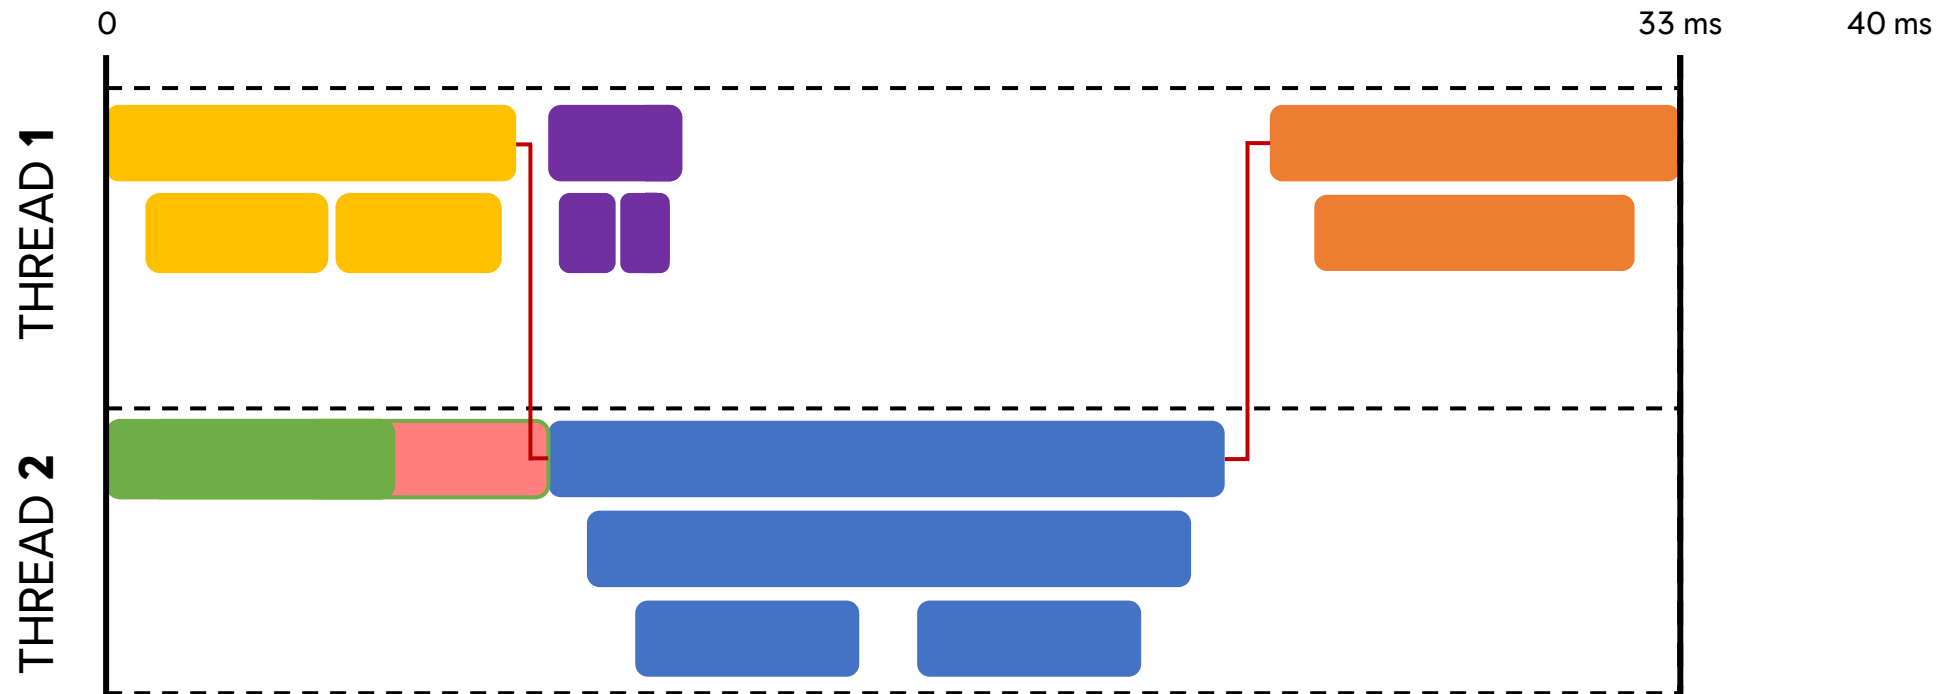
The problem






# INSIDE A CPU FRAME

The problem



# INSIDE A CPU FRAME

The objectives



# FLATTEN VIEW OF THE PROBLEM



Estimate the **time impact** of guilty tasks for each frame





# CONSTRAINTS



**FRAME-WISE  
EXPLICATIVE**



**MEANINGFUL  
OUTPUTS**



**ROBUST &  
TRUSTWORTHY**

# CANDIDATE SOLUTIONS

**Hand-made  
correlation metric**



March 20–24, 2023 | San Francisco, CA

GDC



# CANDIDATE SOLUTIONS



**Hand-made  
correlation metric**

It was a dead end

**Causal profiling,  
Frame simulation**



March 20–24, 2023 | San Francisco, CA



**GDC**



# CANDIDATE SOLUTIONS



**Hand-made  
correlation metric**

It was a dead end



**Causal profiling,  
Frame simulation**

Lack resources & data



**Machine learning  
XAI: eXplainable AI**



March 20–24, 2023 | San Francisco, CA



**GDC**

# METHODOLOGY

1

## DEVELOP

different approaches



2

## EVALUATE


which works better



3

## DEPLOY

the best one



4

## REPORTS

provided to team



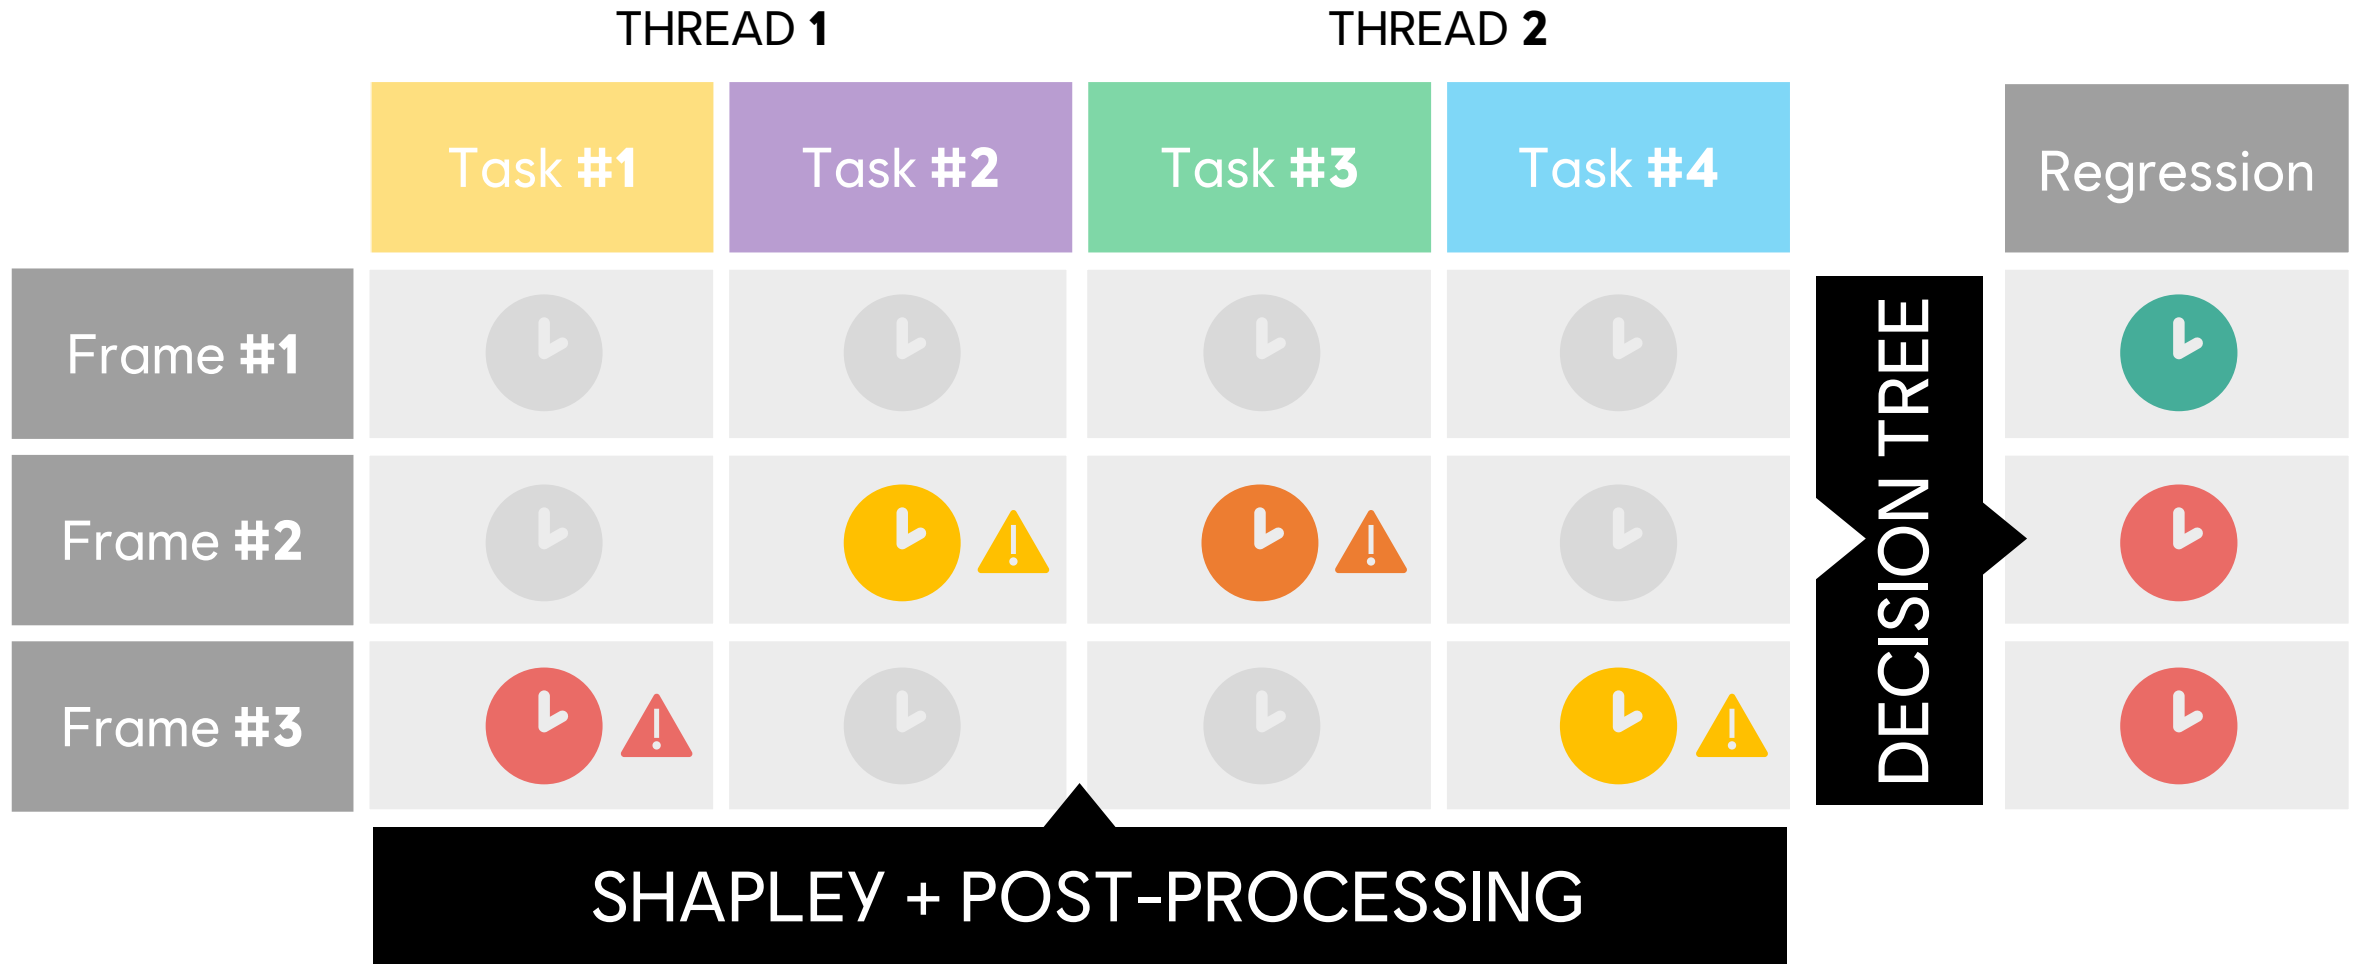
March 20–24, 2023 | San Francisco, CA

GDC



# DECISION TREES & SHAPLEY VALUES

Method 1







# DECISION TREES & SHAPLEY VALUES

Method 1



**+ CUMULATIVE**



**FRAME-WISE  
EXPLICATIVE**



**MEANINGFUL  
OUTPUTS**

Heavy post-processes  
needed



**ROBUST AND  
TRUSTWORTHY**



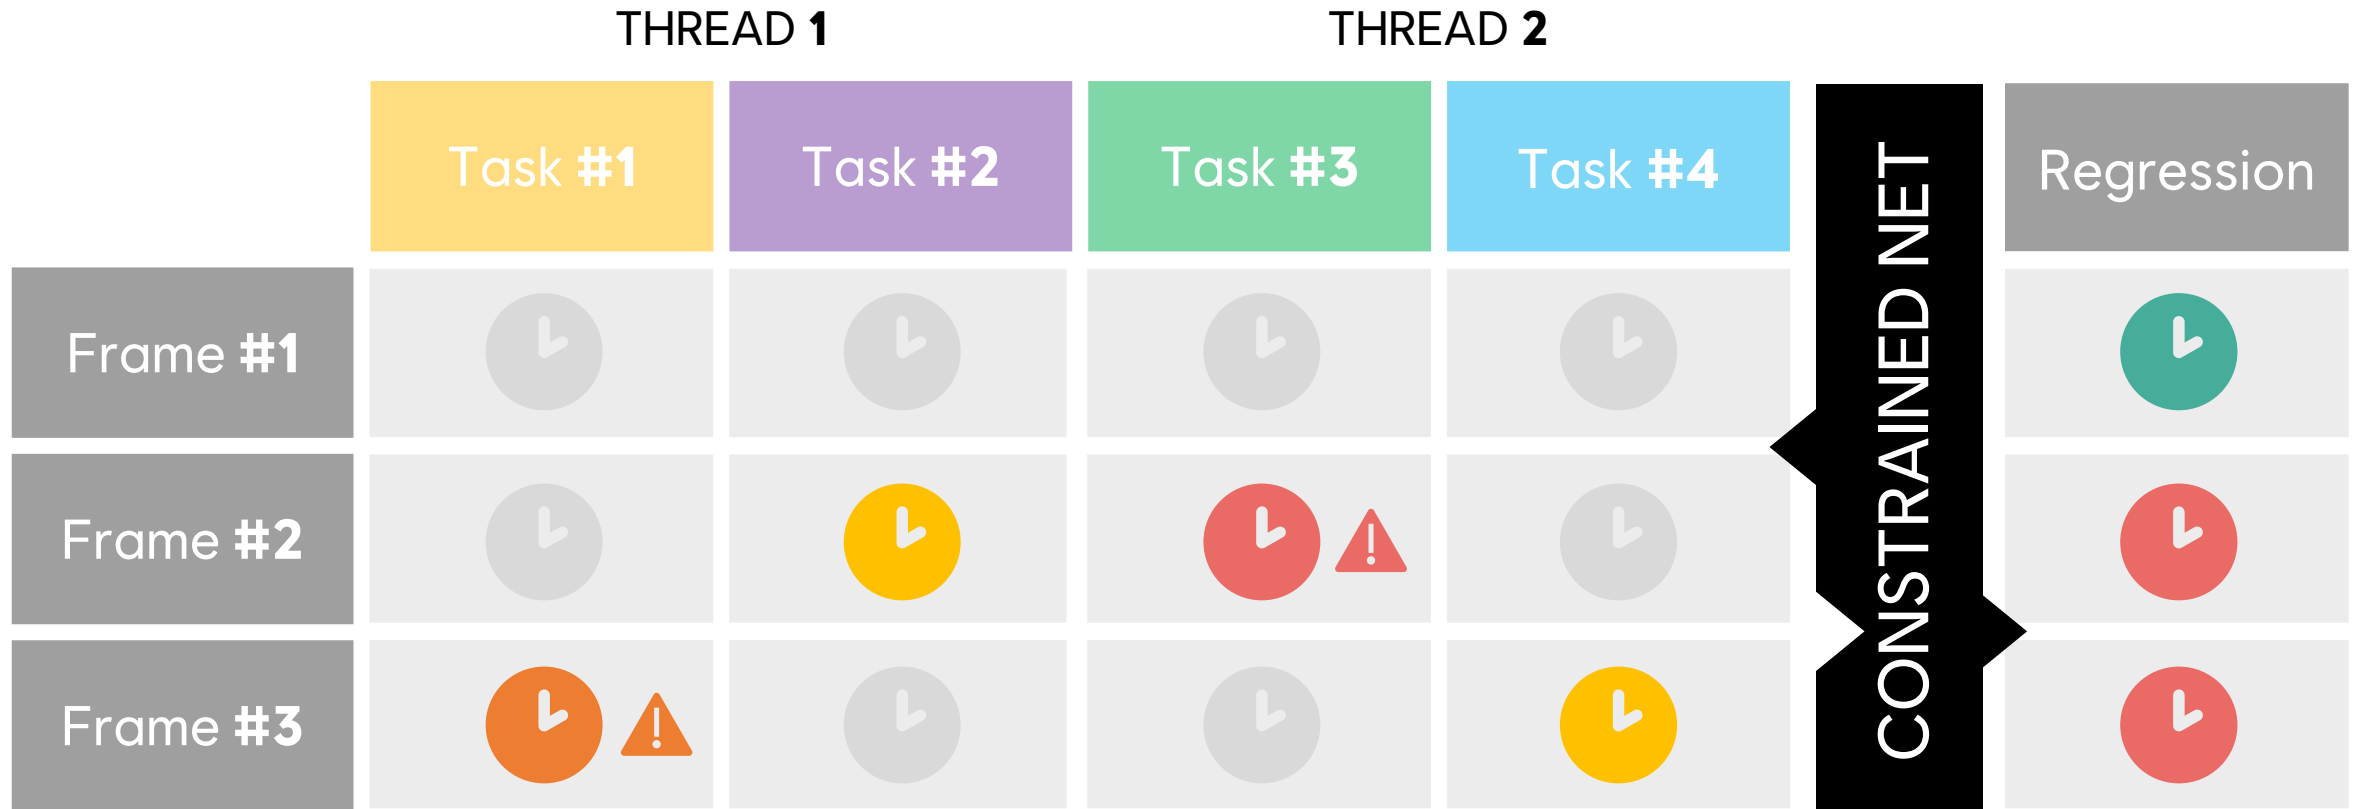
March 20–24, 2023 | San Francisco, CA

**GDC**



# CONSTRAINED NEURAL NETWORK

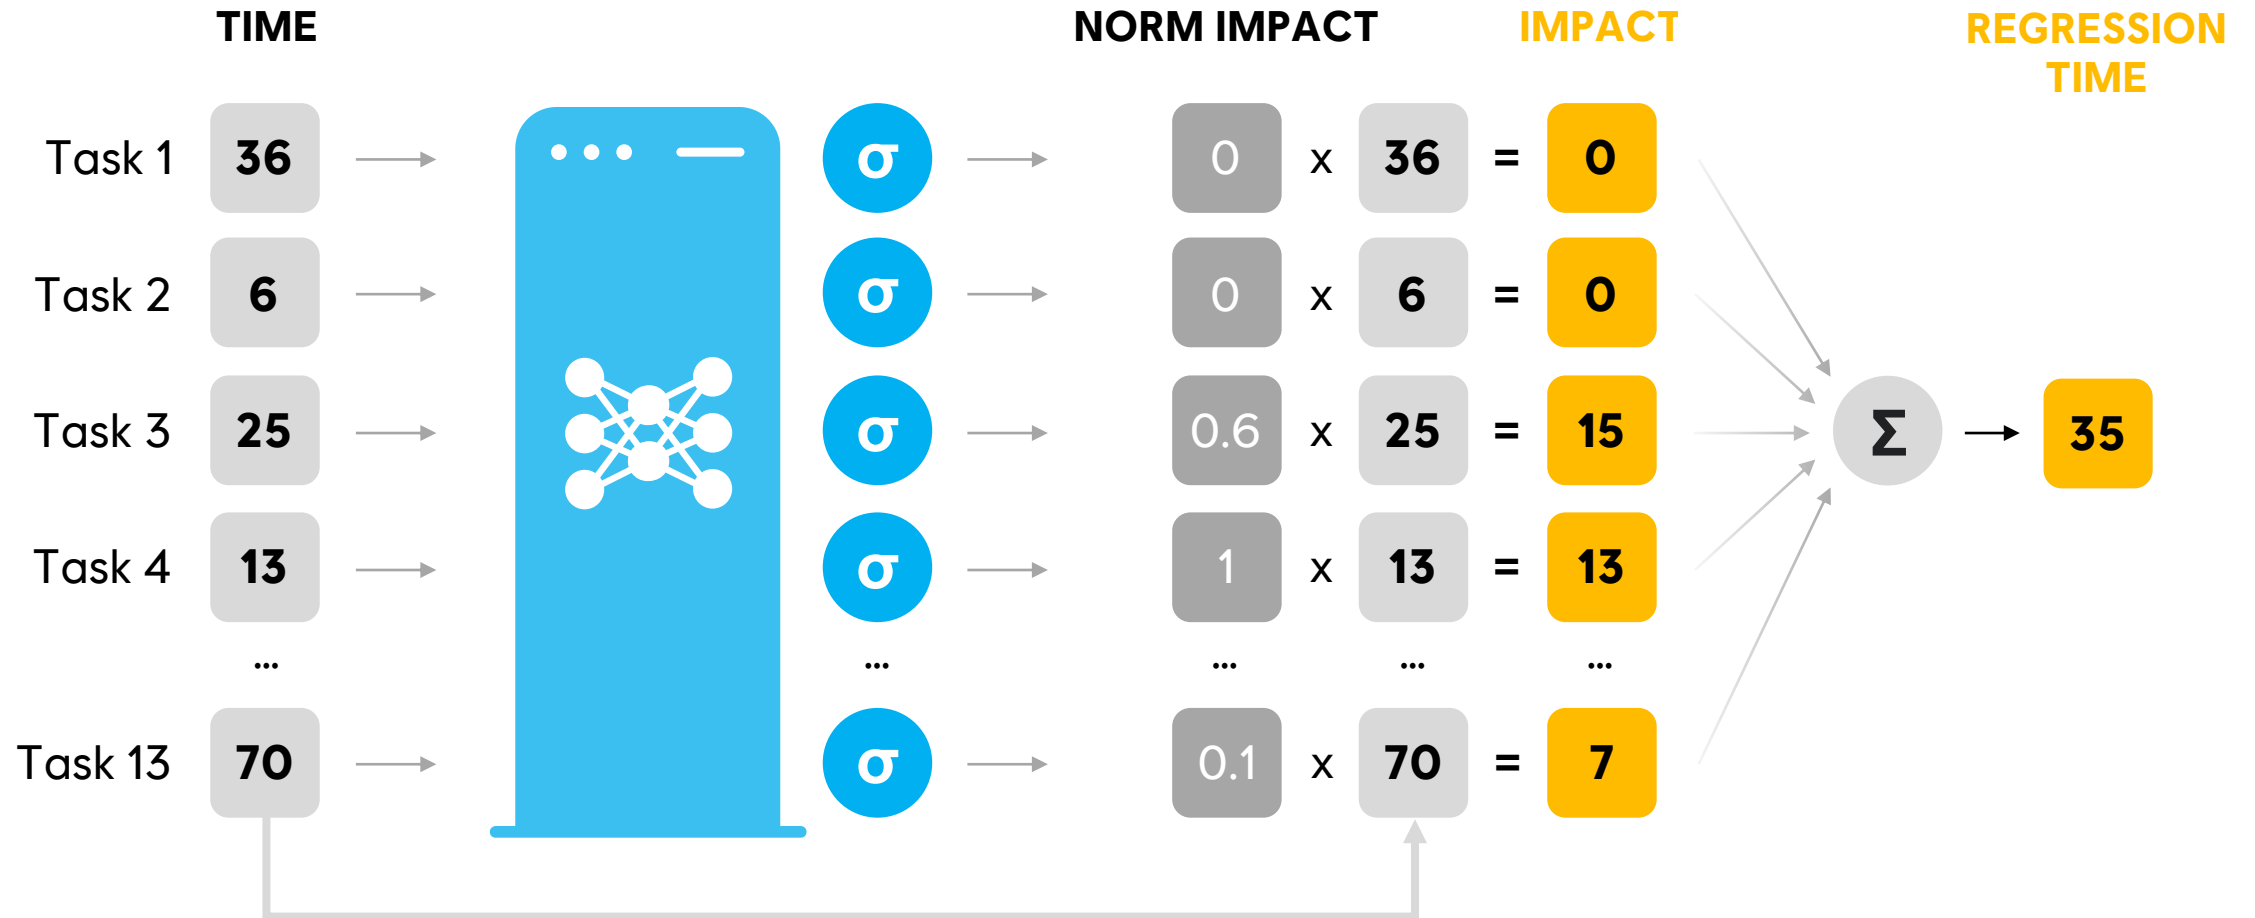
Method 2





# CONSTRAINED NEURAL NETWORK

## Method 2

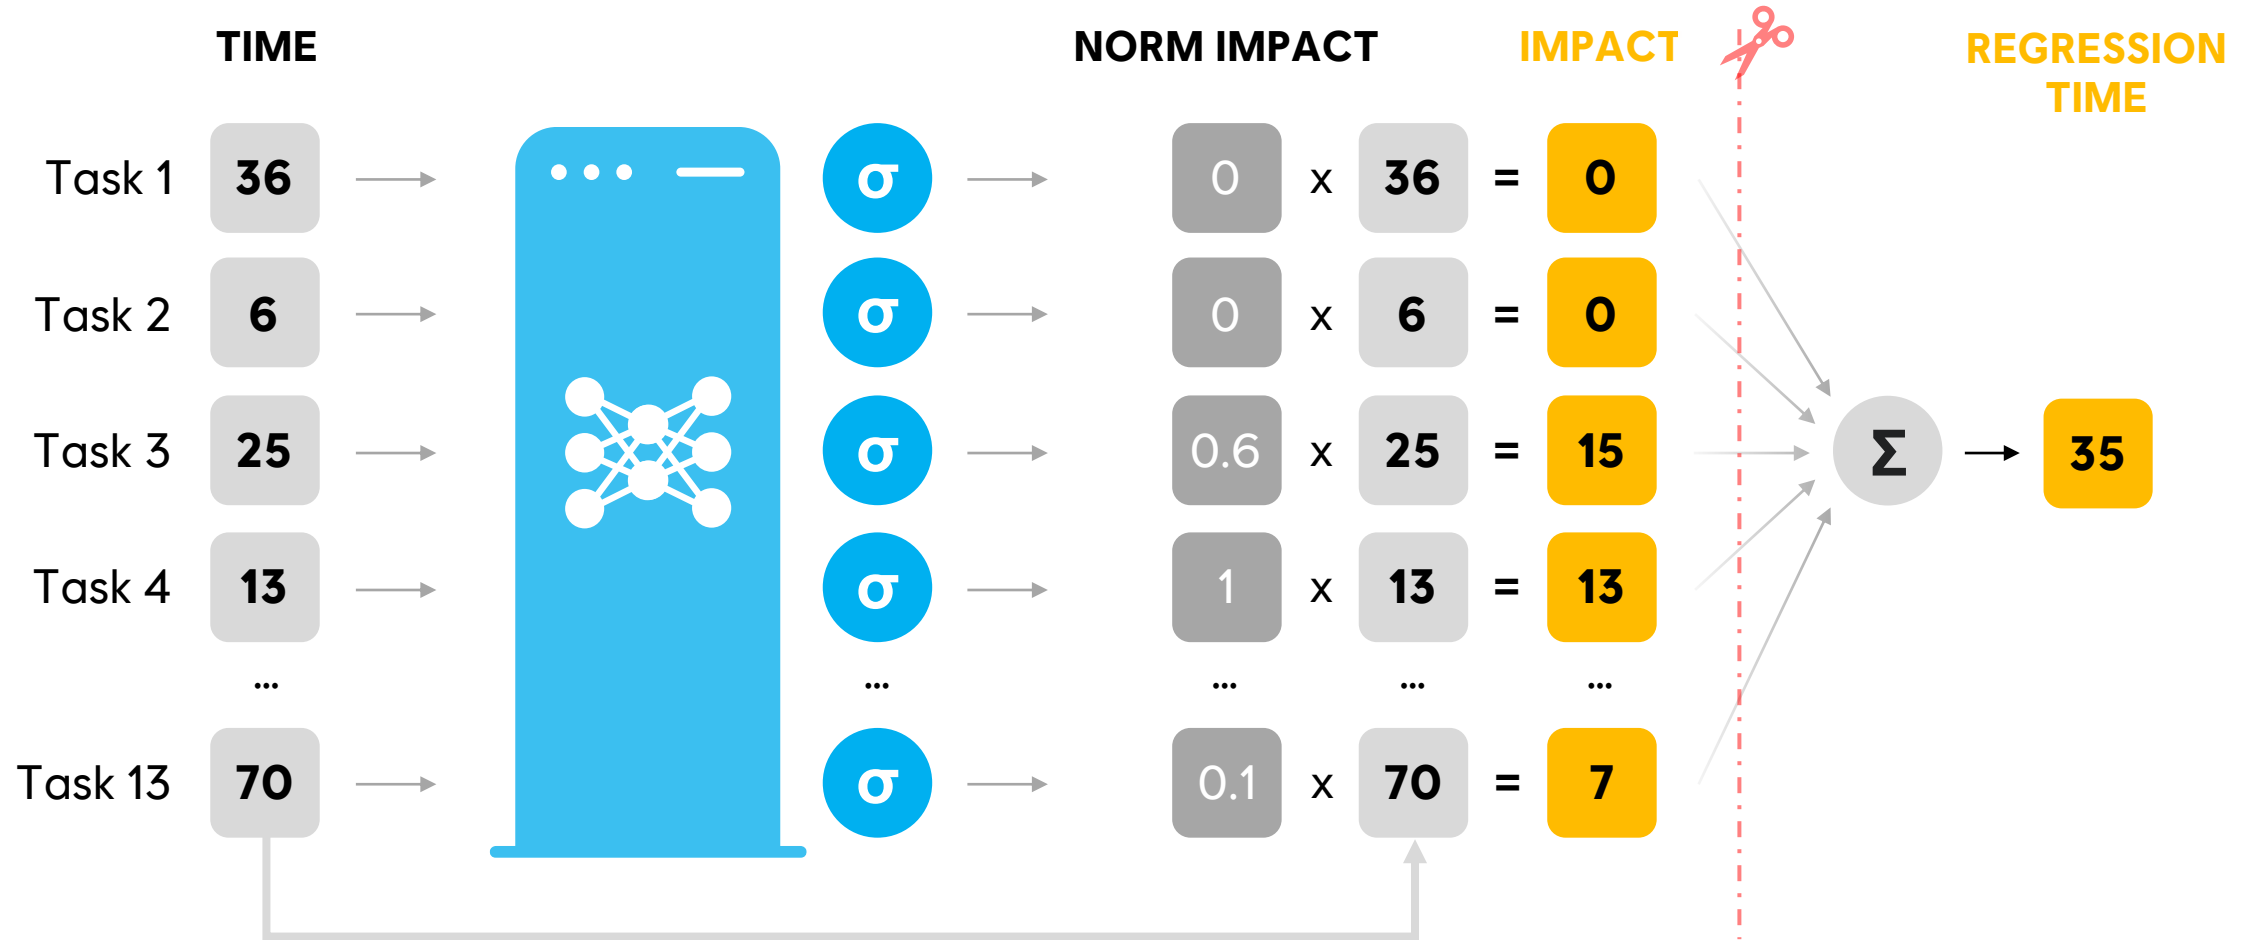






# CONSTRAINED NEURAL NETWORK

## Method 2






# DECISION TREES & SHAPLEY VALUES

Method 1



+ CUMULATIVE



+ DATA FORMAT  
FLEXIBILITY



+ COST  
EFFECTIVE



**FRAME-WISE  
EXPLICATIVE**



**MEANINGFUL  
OUTPUTS**



**ROBUST AND  
TRUSTWORTHY**



# EVALUATION

- **CHOOSE THE BEST METHOD**
- **MEASURE RELEVANCE**
- **BUILD TRUST**




March 20–24, 2023 | San Francisco, CA








# EVALUATION CHALLENGES



**No standard methods  
for XAI evaluation**



**Good prediction  
≠ good explanation**



**Single output evaluation  
= 1-5min of busy expert**

---

**Solution:** Applicative survey to limit evaluation cost & minimize expert time

---



# EVALUATION



## Explainable models validation

### Evaluation How-to

1. **Open** GearStudio
2. **Select** best offenders estimation
3. **Evaluate** relevance of each estimation
4. **Validate** or **inform** all true offenders you are sure about (click on bars in charts or "search and insert" just below)
5. **Comment** (optional)
6. **Confirm** or **skip** evaluation (when skipping, comment section will be stored and example won't be proposed again if filled-in)

### Session description :

Platform : PS4

Version : 6597877

Ubi session id : a7b508d6-4664-4c05-86bb-d4ee35836672

Xp type : WalkHub32Pautotests


Frame id : 20161

Frame length : 37975 $\mu$ s

Regression time : 4642 $\mu$ s

[Open in GearStudio](#)


#### Top contributors to regression



Do you find those estimations relevant ?

☐ No ☐ Yes

#### Top contributors to regression



Do you find those estimations relevant ?

☐ No ☐ Yes




# DEPLOYMENT

For each new game version & platform, daily

**1**  
**GATHER**  
sessions data



**2**  
**TRAIN**  
on multiple game  
versions



**3**  
**EXTRACT**  
& serve impact



Hyperparameters are updated **Bi-monthly**



March 20–24, 2023 | San Francisco, CA







# REPORTS

NAVIGATE GAME  
VERSIONS




## Impact


Close Impact History


### No Platform Selected

Select a platform to show impact history graph

Platform 


Version 

Sessions 

XP Types 

FPS Threshold

Aggregation

Task group 

First task 

Task 

Apply

### No Session Selected

Select a system above by filling the filters

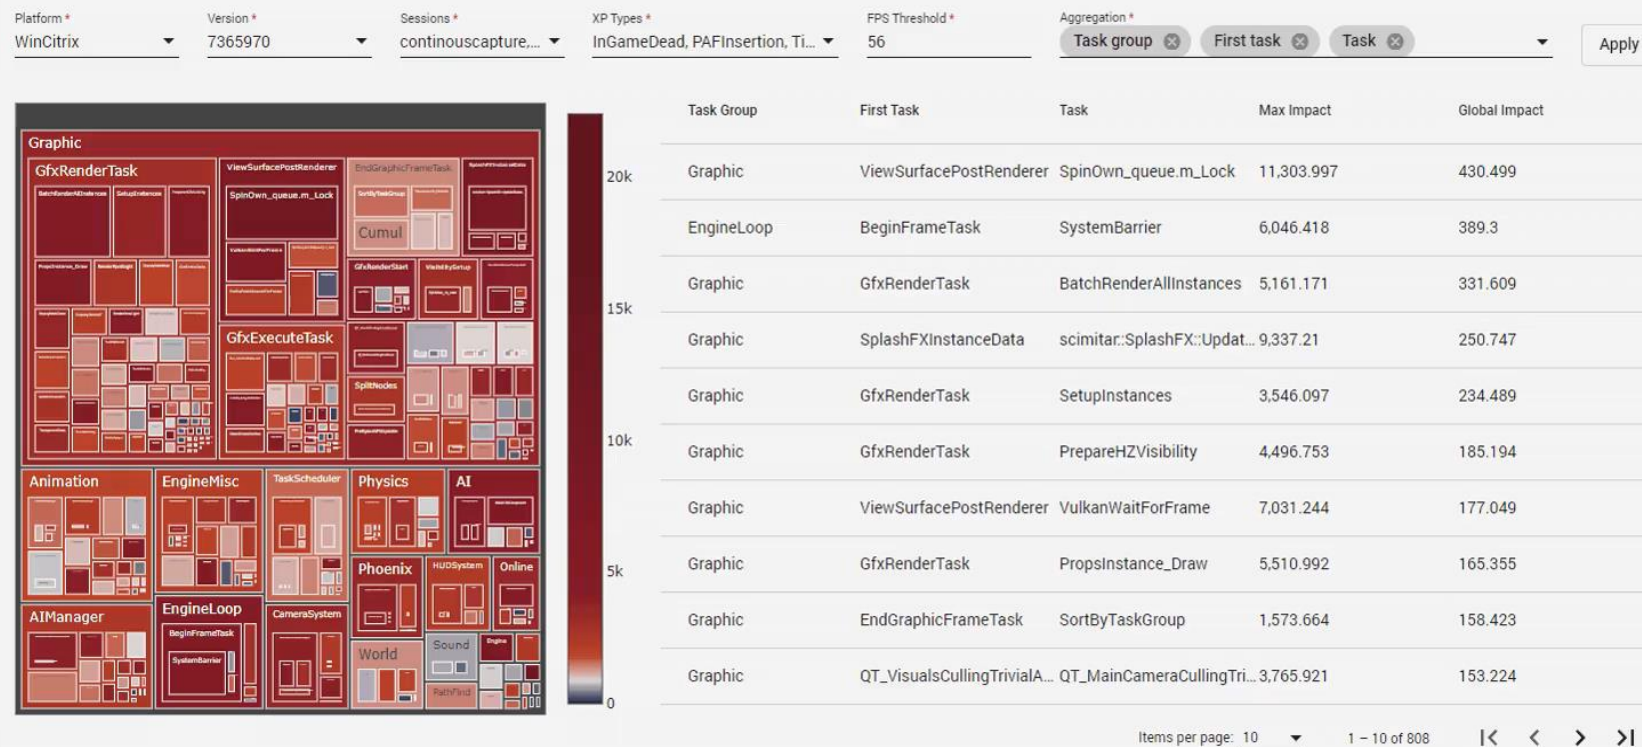
March 20–24, 2023 | San Francisco, CA

GDC



# REPORTS

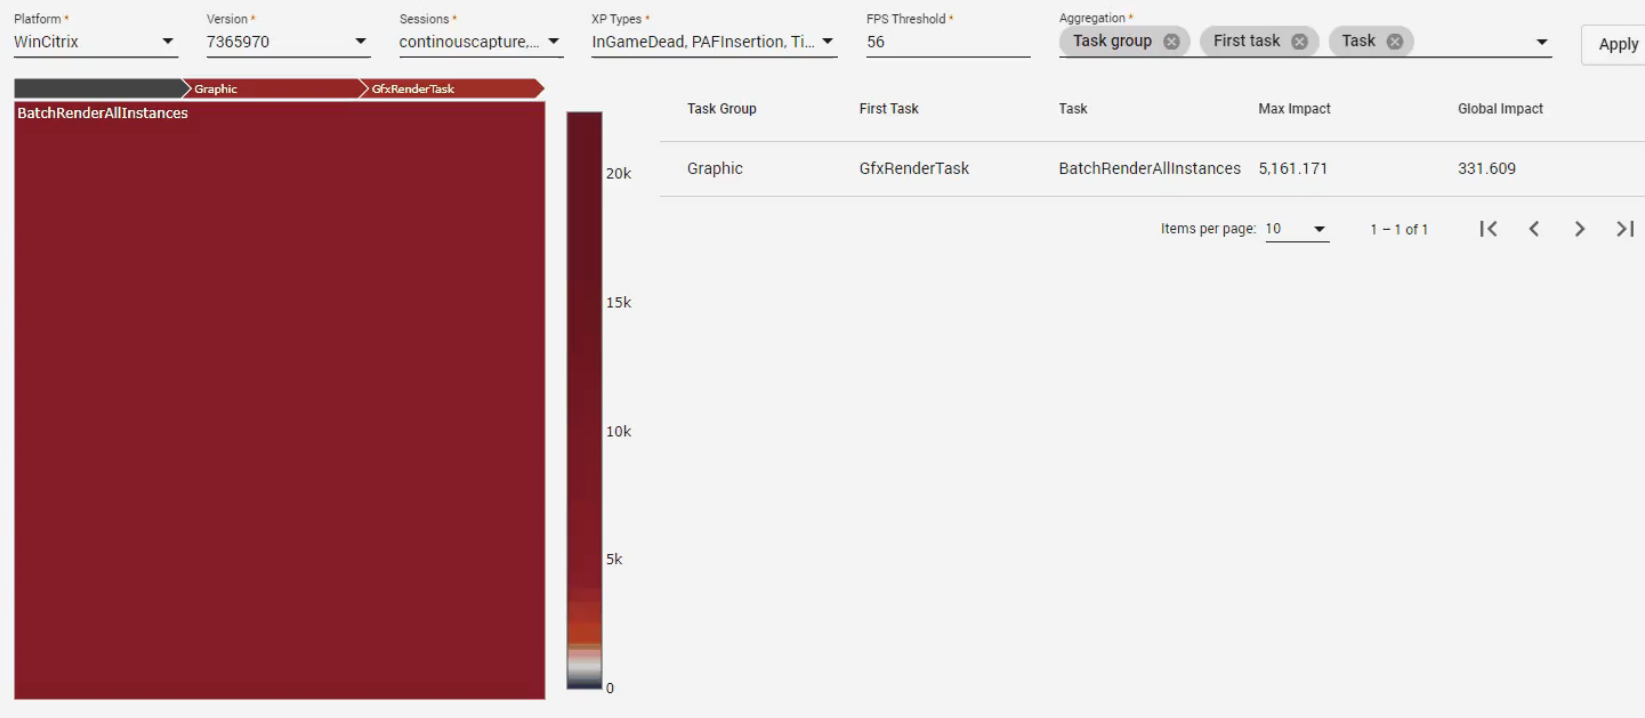
EXPLORE GAME  
VERSIONS





# REPORTS

FIND EXAMPLES



## Related Sessions (5)

Related sessions where selected system were estimated as contributing to regressions.

### Selected System

Task group Graphic First task GfxRenderTask  
Task BatchRenderAllInstances

f55b5690-ffd3-48d5-a3e0-a937b7f4e633  
dd993181-1c90-411a-bc30-7711cf5744f8  
b7efff05-a979-4c31-9c4a-670bcf996f1f  
34fcb50b-d41b-4911-bd5b-fd86884c260d

March 20–24, 2023 | San Francisco, CA


GDC



# A NEW WORK CYCLE



**Look at  
the results**



**Check &  
Prioritize.**



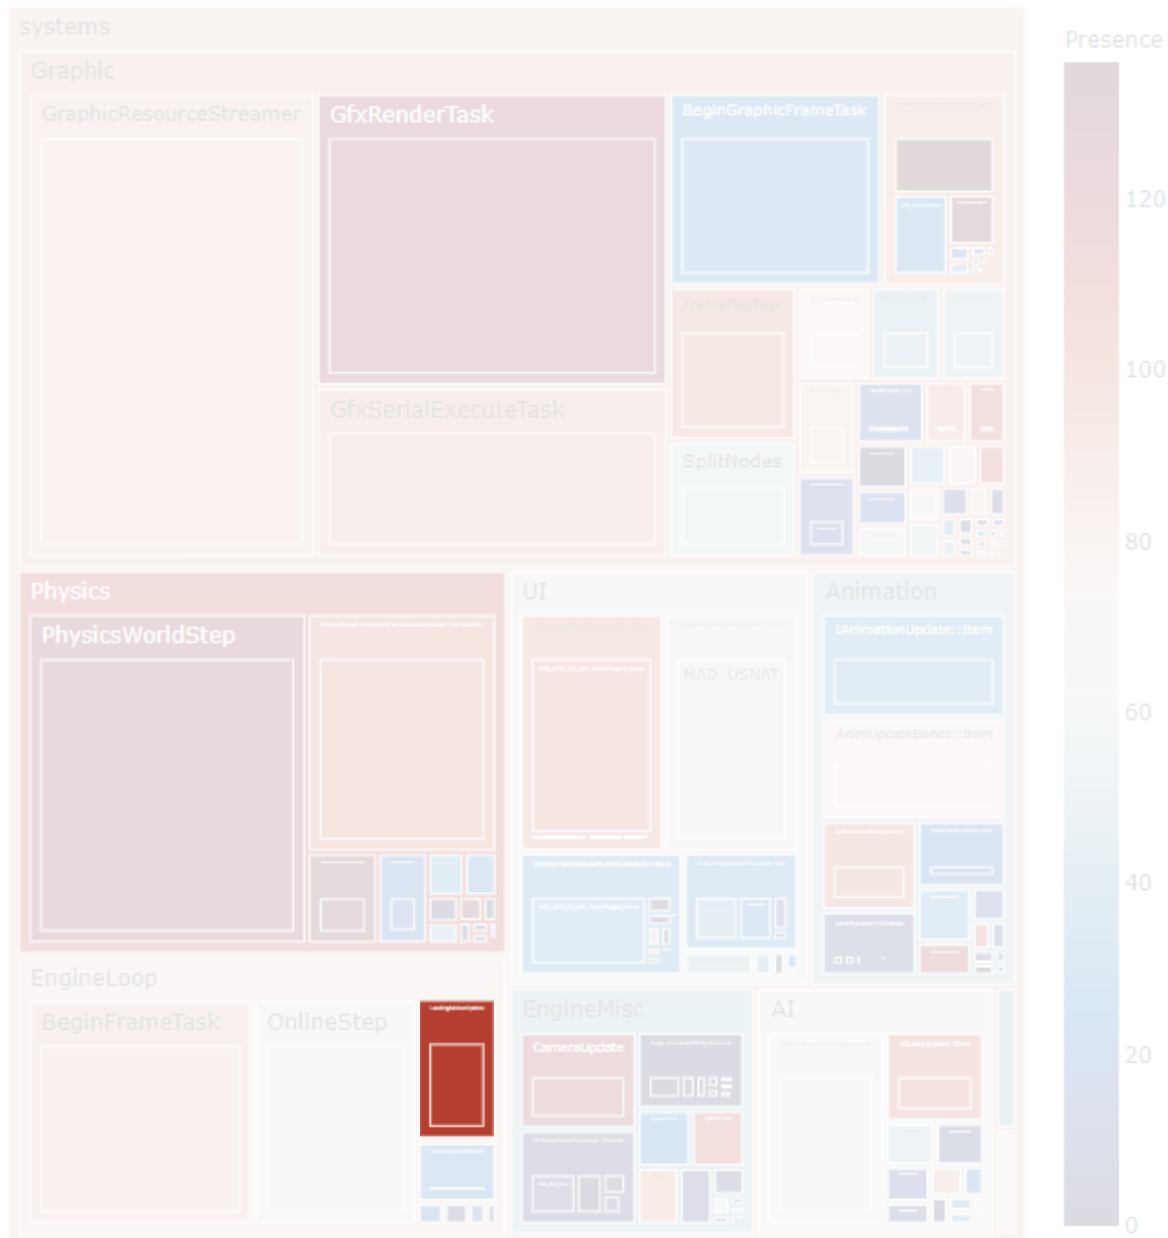
**Distribute  
to the team**

Repeat  
every  
week



March 20–24, 2023 | San Francisco, CA

**GDC**



# « LOADING ADVISOR UPDATE » SYSTEM

Riders Republic, example 1



**BEFORE**

LoadingAdvisorUpdate :

- Impact **30% frames**
- Avg regression of **1ms**


**AFTER**

LoadingAdvisorUpdate :

- Impact **10% frames**
- Avg regression of **0,1ms**



# FRAMEWORK : THE TOOLS



**Dashboard**



**Ad-Hoc analysis**



**Machine learning model**

## ORGANIZATION AROUND PERFORMANCE

March 20–24, 2023 | San Francisco, CA



**GDC**

MARIO RABBIDS : SPARKS OF HOPE

# A NEW ENGINE, A NEW MINDSET





# MARIO + RABBIDS : CONTEXT



Snowdrop  
engine



Multiple  
levels



Switch  
platform



March 20–24, 2023 | San Francisco, CA

GDC

# PERFORMANCE TEAM

Key roles



## PERFORMANCE CHAMPION

Ensure data quality,  
check data availability,  
train the team



## PROJECT CHAMPION

Identify priorities,  
reports to core team  
about game's health



## DATA ANALYST

Create custom  
dashboards  
& ad-hoc analysis



March 20–24, 2023 | San Francisco, CA

GDC

# THE SUCCESSES



## STRUCTURE

Thanks to recommended  
organization



## EARLY DETECTION

Of issues



## SERENITY

About CPU  
performance topic





# TO BE IMPROVED



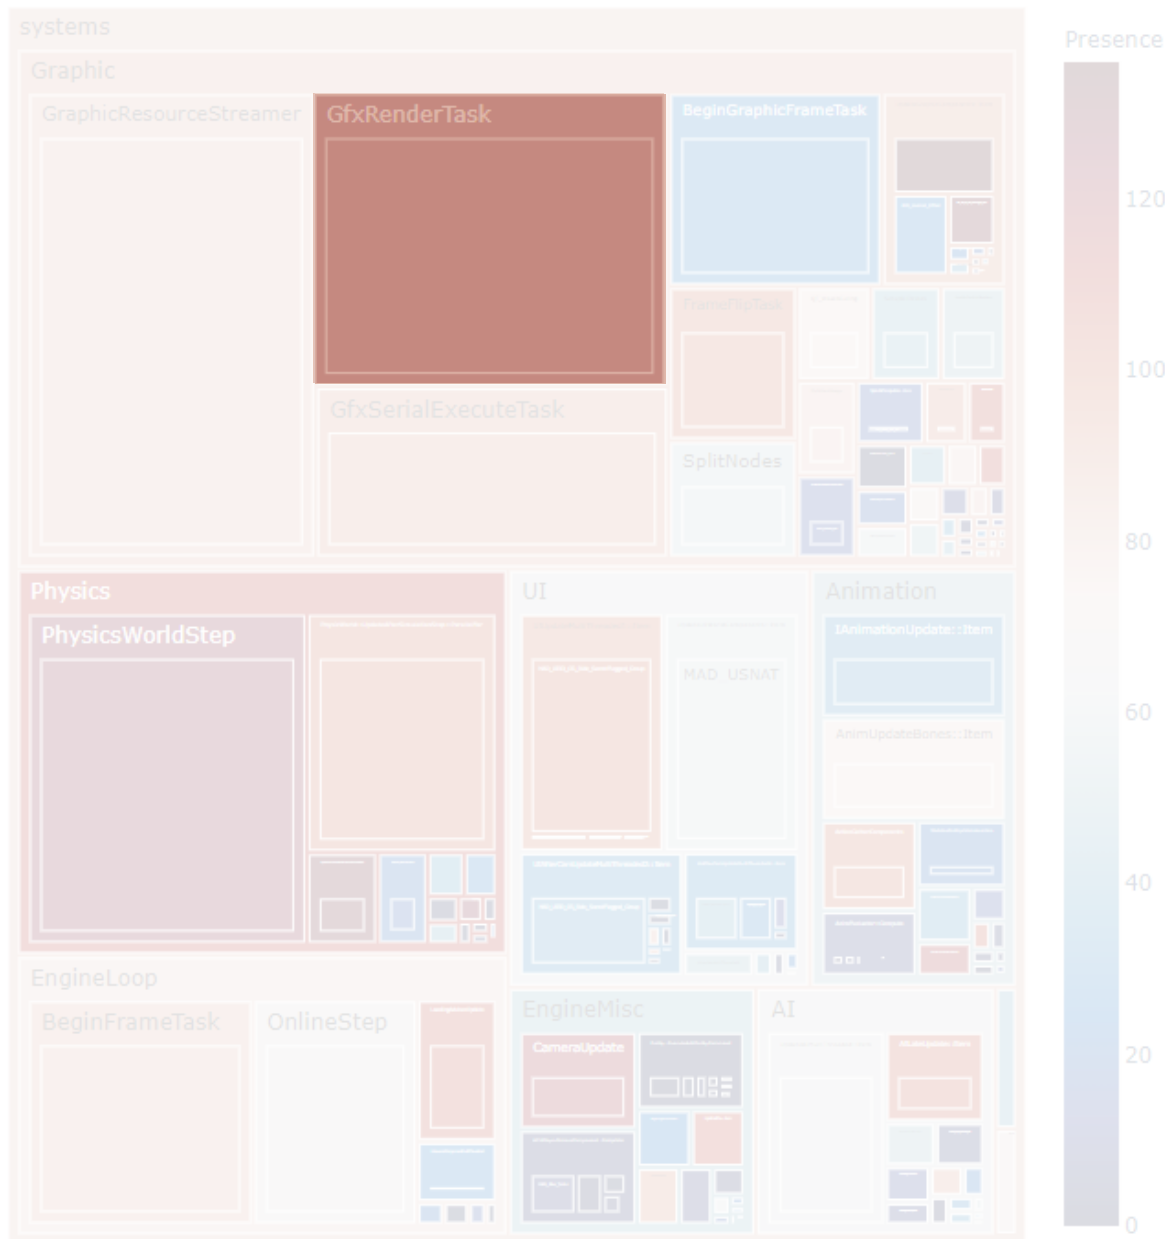
**RELEVANCE**  
of some results



**PRECISION**  
of data







# POSSIBLE SOLUTION

UX Post processing



## DO NOT:

- Remove completely
- Force assumptions into the model



# WHAT CHANGED



## BETTER GLOBAL VISIBILITY

on project health



## BETTER EXPERTS' WORKFLOW

They are less bottleneck



## RELIEF

On CPU performance topic

---

But we can still improve knowledge democratization!


---



March 20–24, 2023 | San Francisco, CA



# WHERE TO GO NEXT?




## CPU

Locks, context  
switches & More  
game context



## GPU



## MEMORY



## BANDWIDTH



March 20–24, 2023 | San Francisco, CA

GDC



# SPECIAL THANKS



**JEAN-PHILIPPE BARETTE**  
Senior Technical Architect



**MATTHIEU FOURNAISON**  
Lead Programming Architect



**FRANÇOIS DETEMMERMAN**  
Gameplay Architect



**THOMAS BIRON**  
Product Manager



**MAXIME MERY**  
Lead Production Analytics



**PIERRE-THOMAS MEISELS**  
Data Engineer



**MATTHEW CHARLTON**  
Senior Web Developer



**CARL LÉTOURNEAU**  
Data Developer



March 20–24, 2023 | San Francisco, CA



# SPECIAL THANKS



**BAPTISTE DECAIN**  
Lead Programmer

**GUILLAUME NEWTON**  
Programmer

**GUILLAUME GRODWOHL**  
Senior Programmer

**XAVIER BILLAULT**  
Senior Tech Lead Programmer



**STEPHANE MARQUIS**  
Technical Lead

**SIMON POLIQUIN**  
IT Project Manager

**MICHEL PERRAULT**  
IT Team Lead

**SIVA VELOUMOUROUGANE**  
Data Scientist Director

**THOMAS GUILLEMBET**  
Data Visualization specialist



**GAUTHIER TANGUY**  
Lead Programmer

**NORBERT REDON**  
Senior Programmer

**ELIA WRZESNIEWSKI**  
Associate Producer

**ALEXANDRE NOVAK**  
QA Analyst

**ARTHUR BLIN**  
Data Analyst



March 20–24, 2023 | San Francisco, CA



# RESSOURCES



**CHARLIE CURTSINGER,  
EMERY D. BERGER**

Coz: Finding Code that  
Counts with Causal Profiling



**BRENDAN GREGG**  
Flamegraphs visualization



# Q&A



March 20–24, 2023 | San Francisco, CA

GDC





# THANK YOU



March 20–24, 2023 | San Francisco, CA

GDC