

September , 2023

Operational Excellence at Zendesk



John Viner
Senior Director of Engineering

A Framework for running reliable software at scale



Agenda

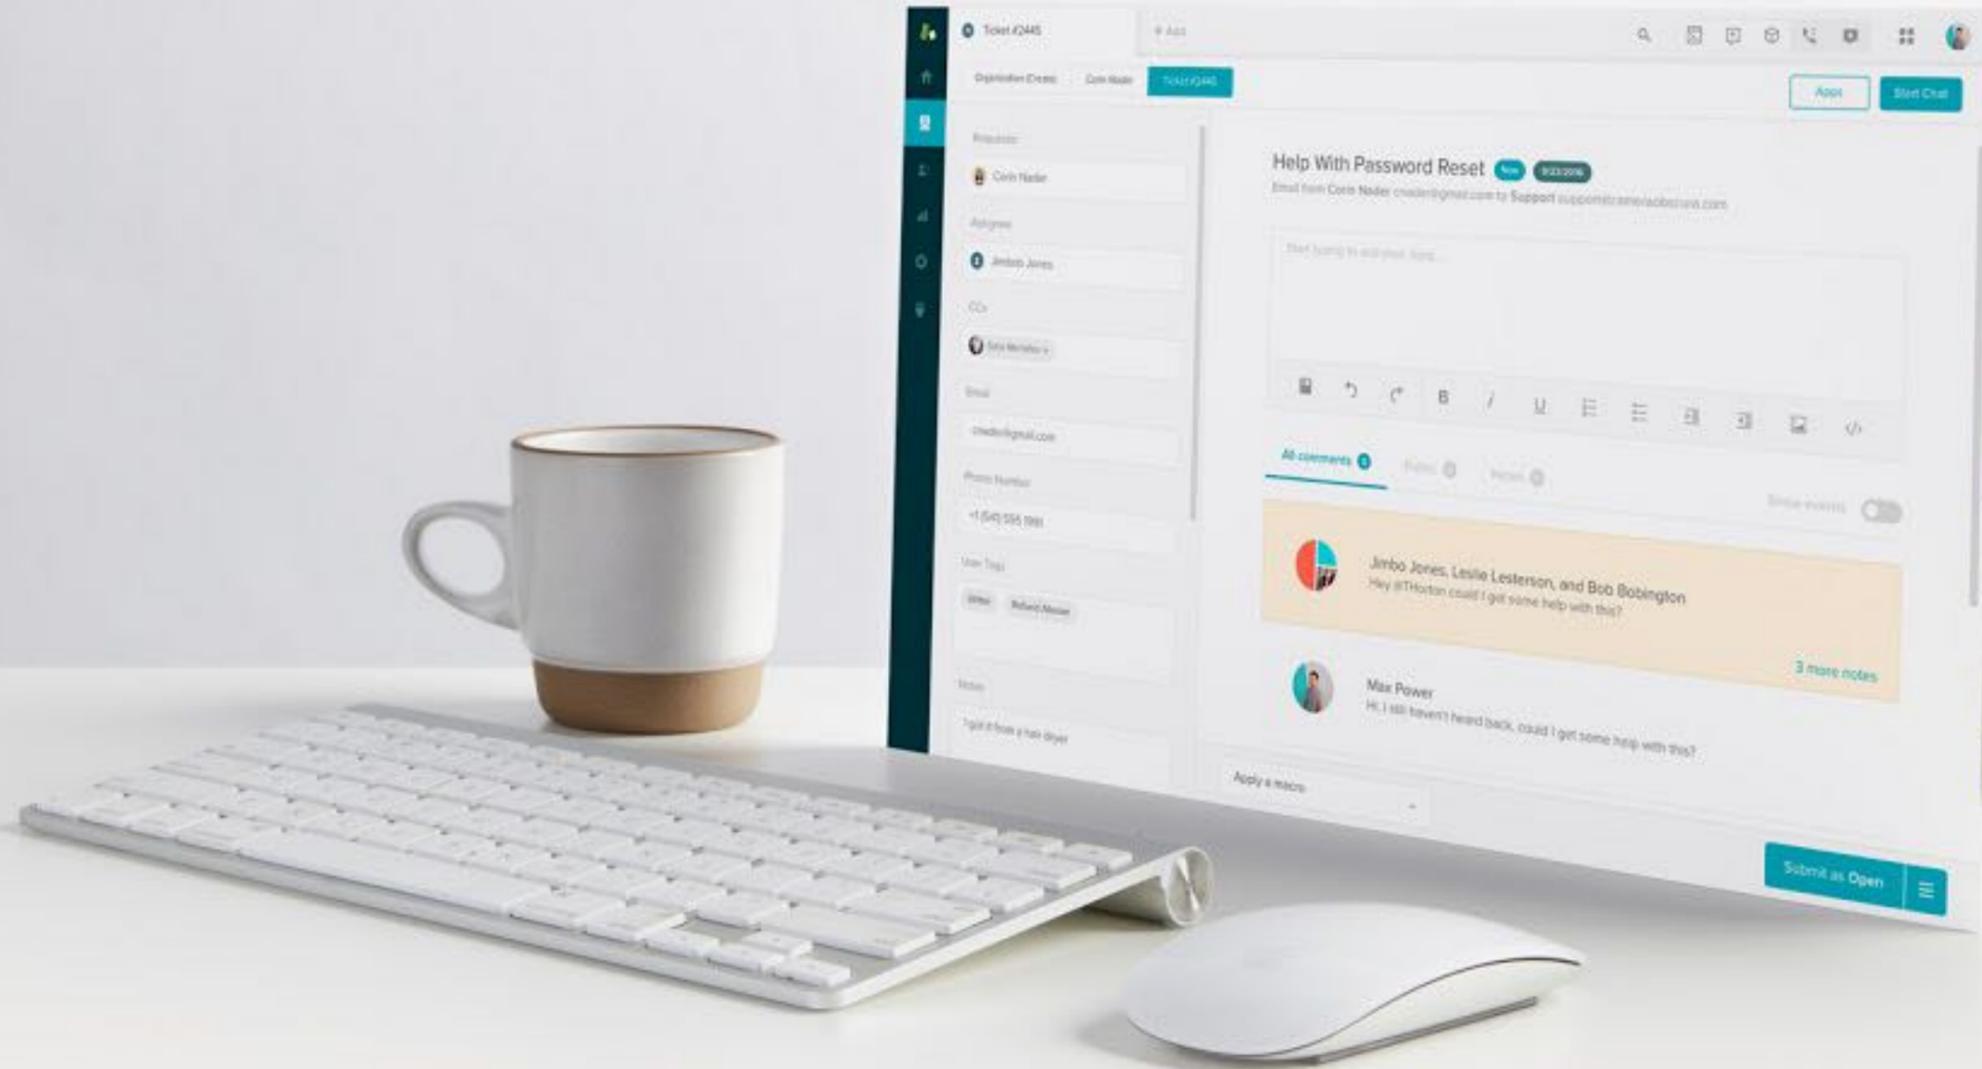
Recap - Error Budgets

Operational Metrics Review

Operational Excellence Dashboard

The Results

Recommendations



Ticket #245

Operator: Cain Nader

Assignees: Jimbo Jones, Eric Marder

Help With Password Reset

Email from Cain Nader: cna@amazon.com to Support: support@amazon.com

Dear typing in with email. Hello,

All comments

Jimbo Jones, Leslie Lesterson, and Boo Bobington
Hey @Horton could I get some help with this?

Mike Power
Hi, I still haven't heard back, could I get some help with this?

Submit as Open

Recap: 2019 Error Budgets



ERROR BUDGETS RECAP - 2019



SLI - Service Level Indicator

HOME PAGE REQUEST SERVED IN
< 100MS

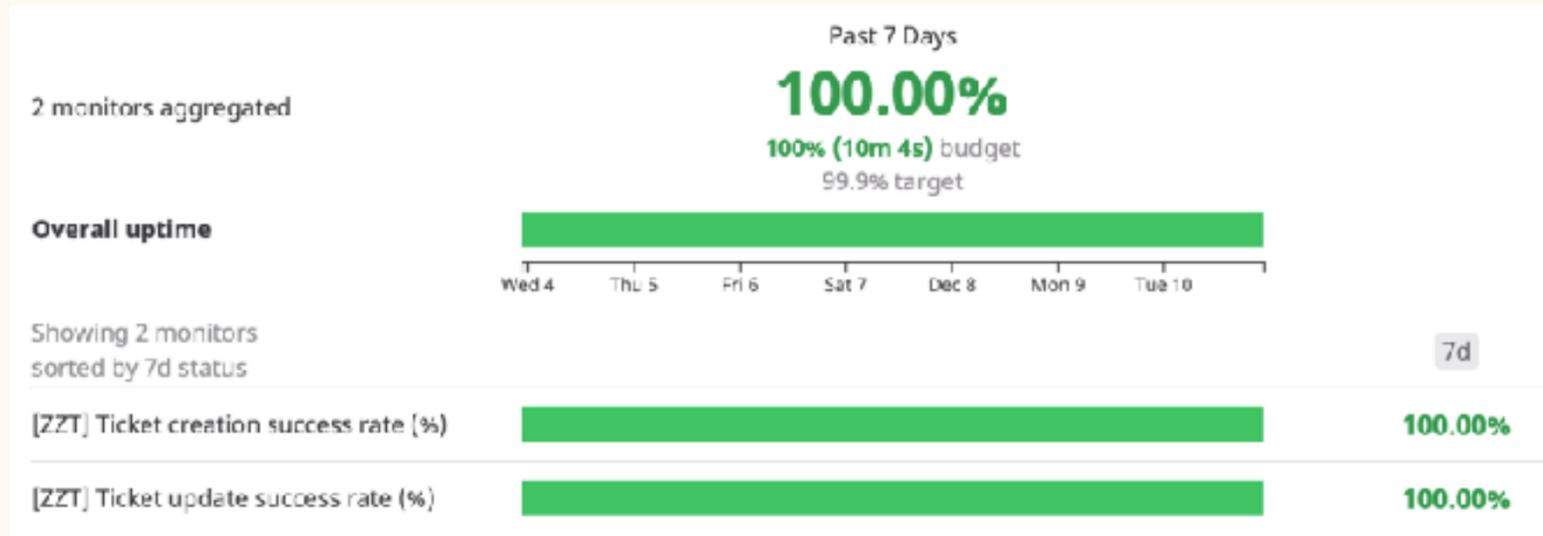
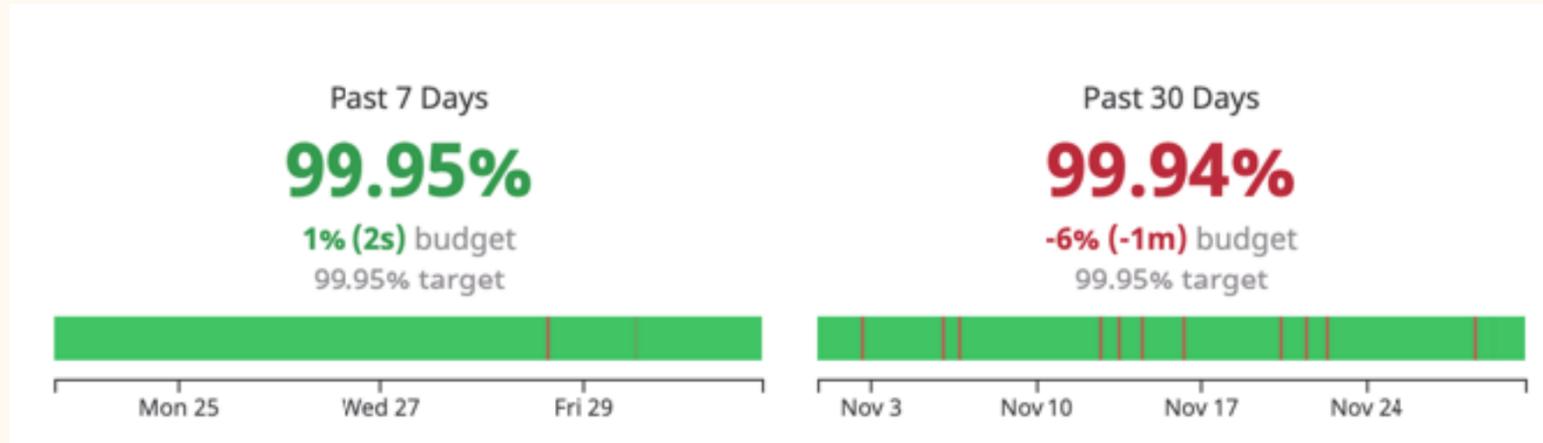
SLO - Service Level Objective

95% OF HOME PAGE REQUESTS SERVED
IN < 100MS OVER LAST 24 HOURS

Error Budget

ALLOW 5% FAILURE OF HOME PAGE
REQUESTS SERVED IN < 100MS
OVER LAST 24 HOURS

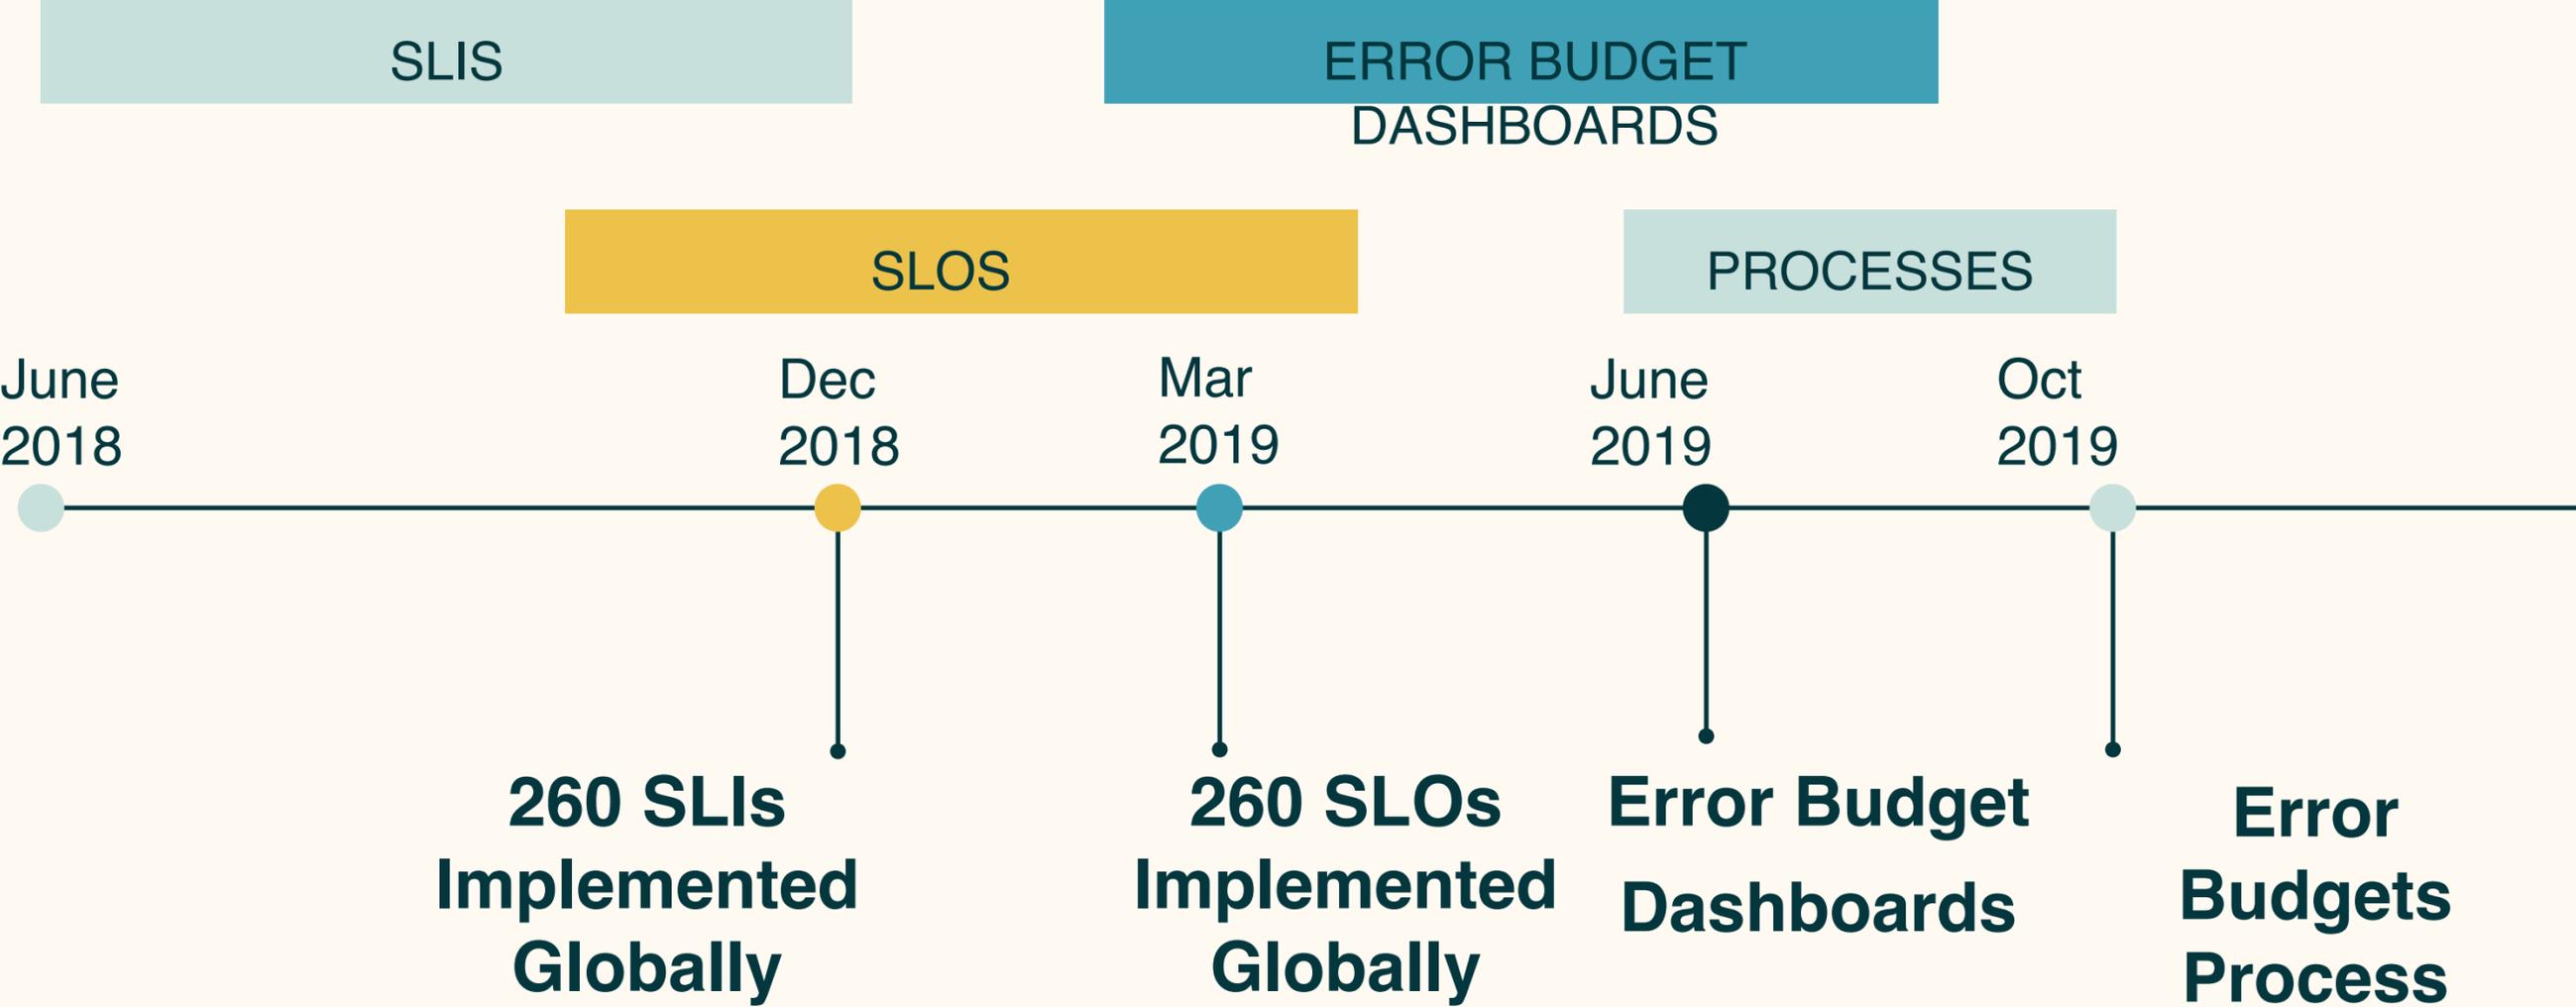
ERROR BUDGETS RECAP



[View Controls](#) | Showing 1221-1240 of 2K results

TYPE	NAME ↑	TIME ↑	STATUS	ERROR BUDGET LEFT
🕒	Api:V2::WorkspacesController#update laten...	30d	🕒	—
🕒	Api:V2beta::Tickets::RelatedController#sho...	30d	100.00%	100% (14h 24m)
🕒	Api:V2beta::Tickets::RelatedController#sho...	30d	100.00%	7% budget left (1m 36s remaining)
🕒	Article Page "Up"	30d	99.95%	7% (1m 36s)
🕒	Articles API "Up"	30d	99.99%	91% (19m 35s)
🕒	AttachmentsController#token Duration (95p...	30d	100.00%	100% (14h 24m)
🕒	AttachmentsController#token Error Rate 🔒	30d	100.00%	100% (14h 24m)
🕒	AttachmentTokenController#show Duration ...	30d	99.08%	54% (7h 46m)
🕒	AttachmentTokenController#show Error Rat...	30d	100.00%	100% (14h 24m)

ERROR BUDGETS RECAP



Freeze or thaw?



We Needed to do More

1

Pre Deployment Testing

2

Deployment Standards

3

Production Smoke Tests

4

Alerting and Monitoring

5

Completing Remediation Items

6

SLI Coverage and Error Budget Status

Across 1500 Engineers, 250+ teams and 7 worldwide locations

V1: The Operational Metrics Review



Learnings

- Screenshot Busy Work for Engineering Managers
- Content Quickly Diverged from Template
- Reporting Dimensions Not Consistent
- Data Munging Effort Duplicated Throughout the Company

V2: Operational Excellence Dashboard



Standardising

1

Single Reporting System

2

Single Data Ingestion Pipeline

3

Standard Org Structure

4

Same Monthly Cadence

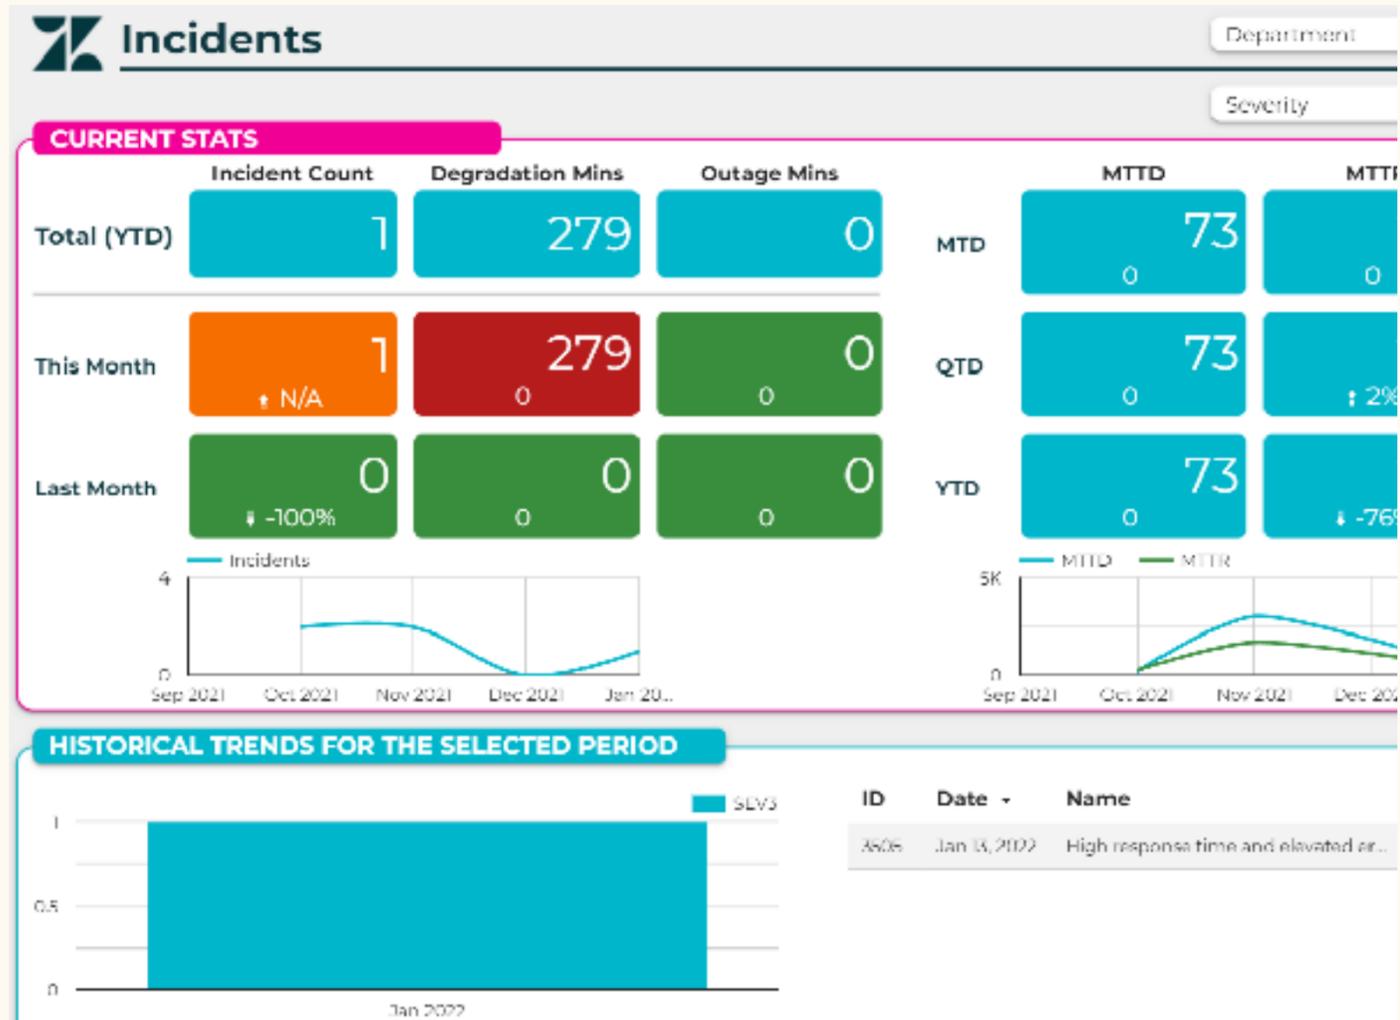
5

Single Notification and Recording

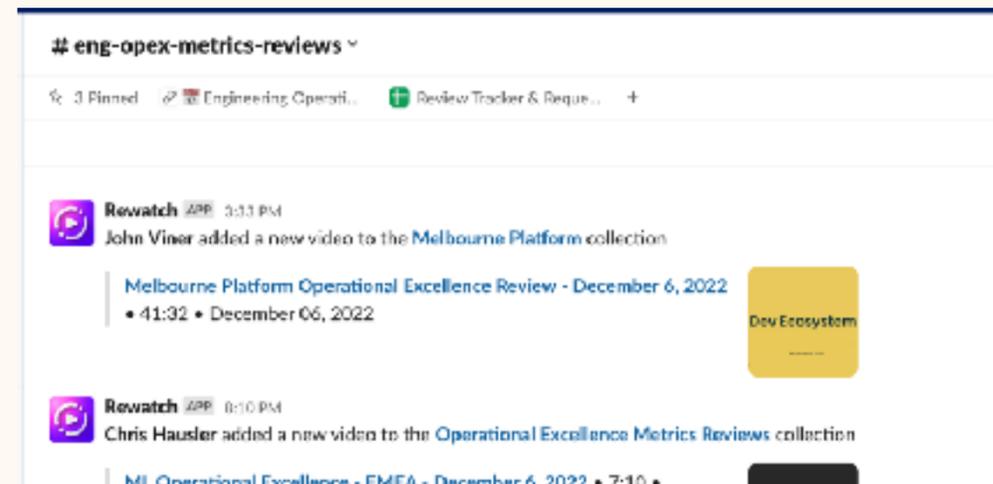
6

One Team to Own It

Bringing it together



And Sharing



Learnings

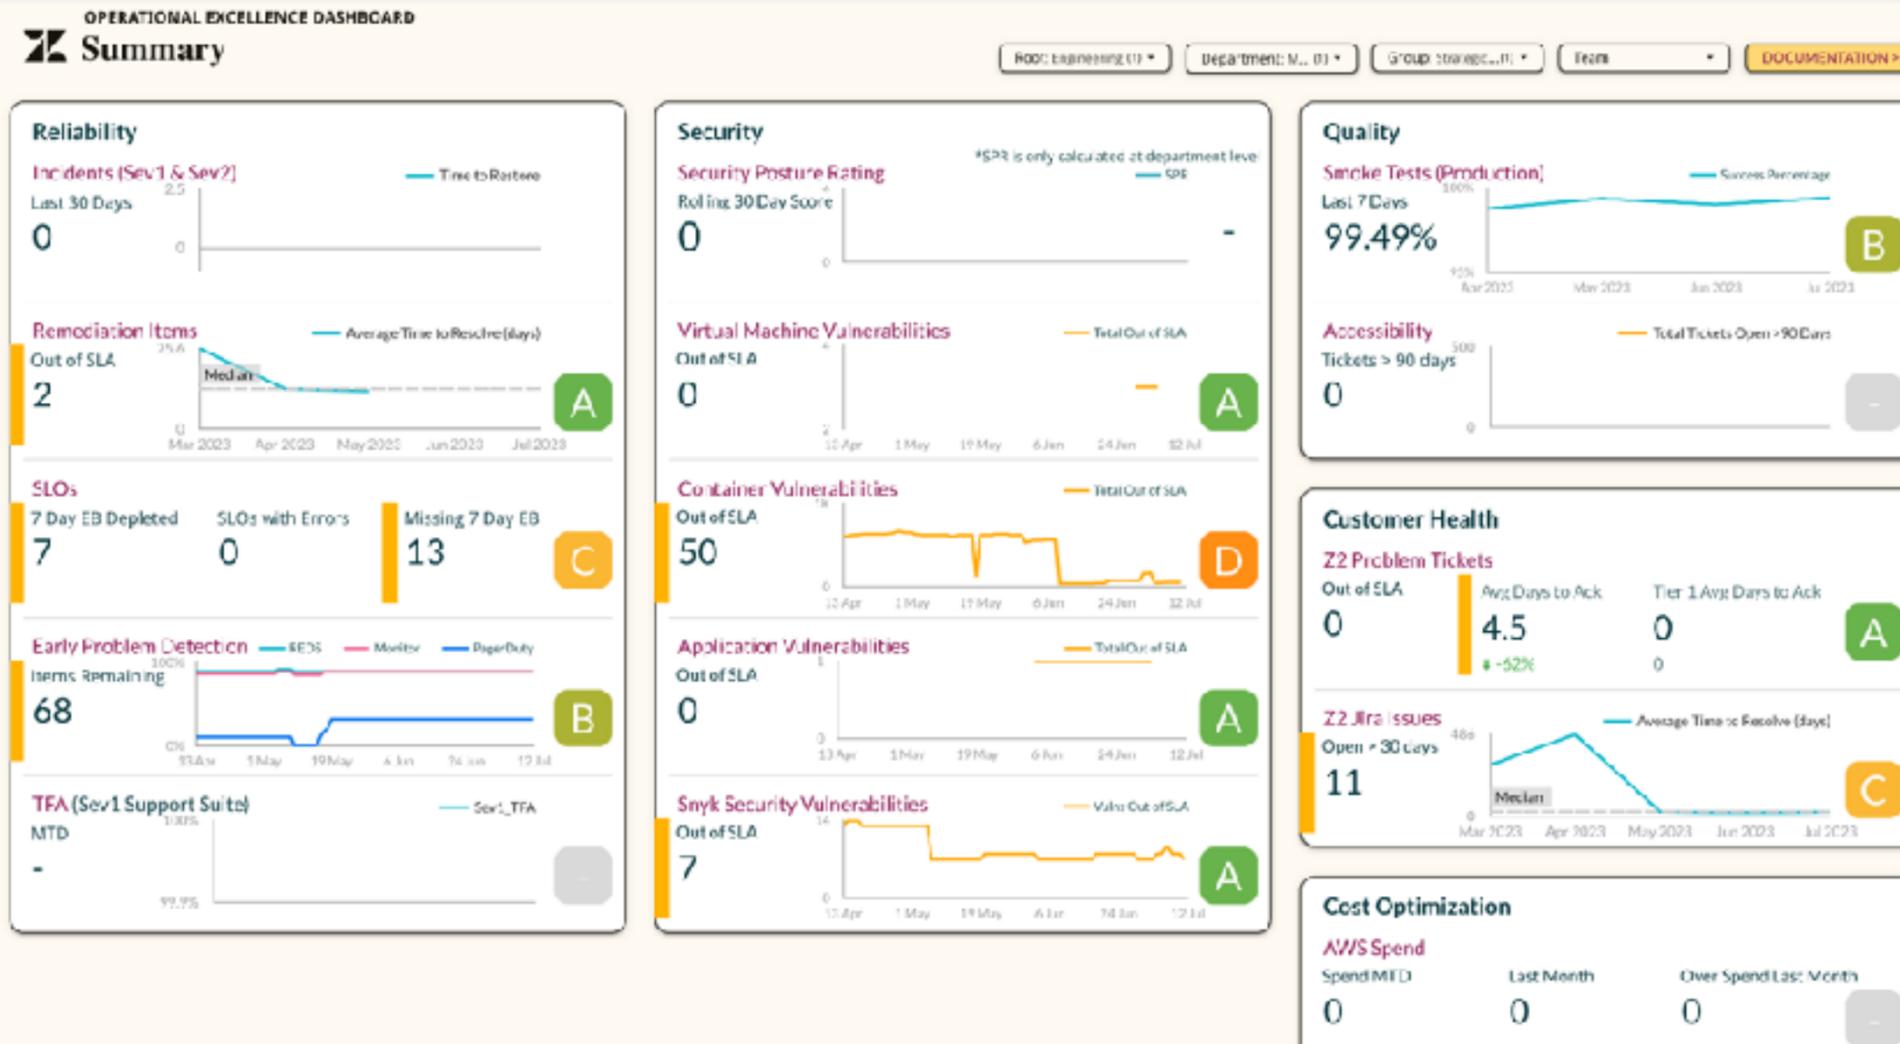
- Much More Efficient for Managers
- Proliferating Amount of Data and Visualisations
- Hard To Pinpoint Problem Areas

V3: Back to School

(Introducing grades)



Grading



OpEx Dashboard

	Jira projects will be combined and added to your report.
Update Frequency	Data is automatically refreshed every 4 hours.
SLAs	https://docs.google.com/document/d/1pB4wGiwS0HnENa0SMmM-0fWn1eQ4wJ5ACVCHmfgXU/edit
Data Owners	Incident Management team #ask-incident-management
Help	If you see some discrepancies in the data on the dashboards you can go to #ask-opex-dashboard .

Remediation Health Score Criteria

Grade	Criteria
A	<10 Normal out of SLA and 0 Major, Critical & Blocker out of SLA
B	0 Critical out of SLA and under 3 Major out of SLA and under 15 Normal out of SLA
C	under 3 Critical out of SLA and under 7 Major out of SLA and under 20 Normal out of SLA
D	>=3 Critical out of SLA or >=7 Major out of SLA or >=20 Normal out of SLA or
F	>0 Blocker out of SLA

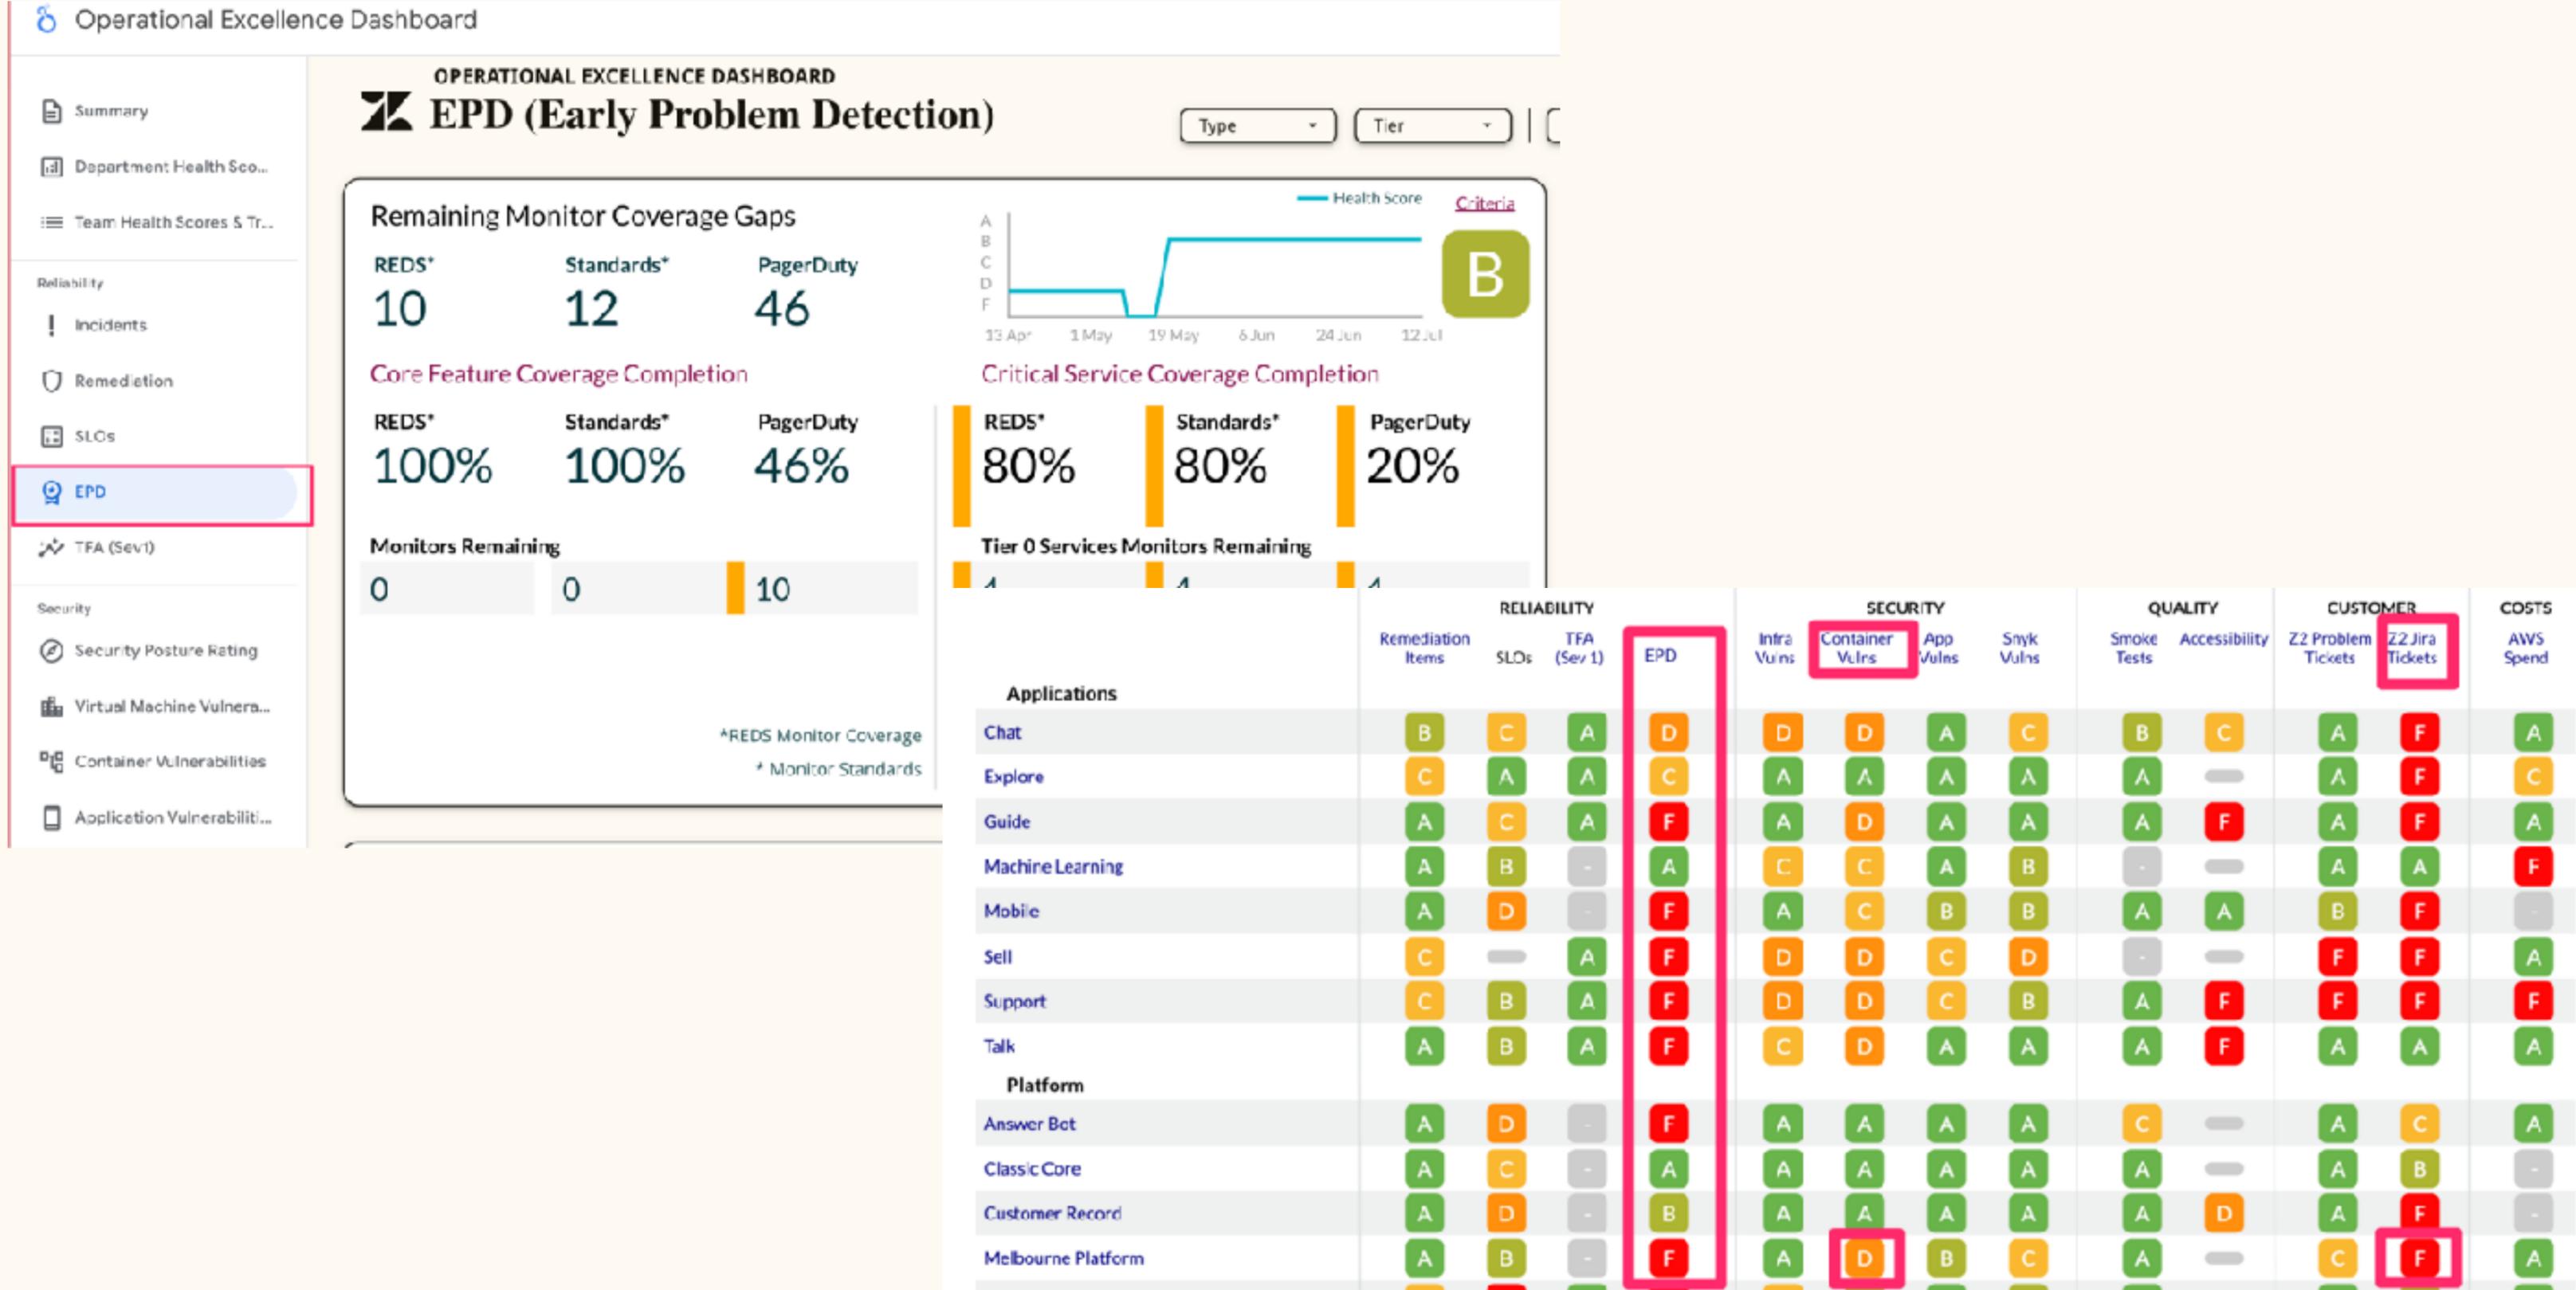
Learnings

- Quickly Surfaced Problem Areas
- Reduced Time Spent Discussing “Good” Areas
- Opportunity to Layer on New Standards

V4: Tip of the Spear



Rolling out Early Problem Detection



Rolling out “First Response Times” to our Customer Tickets

Unreads  All of Zendesk  Sorted by recommended order 

▼ # eng-announcements

 **Kelvin Pham** 10:22 AM
Post 



First Response Time (FRT) macro/metric for Z2 Customer Health

Post

TL;DR: On September 11th, we are launching a new process to measure the First Response Time metric for Product Development Z2 tickets. This will be a required step for all Z2 tickets moving forward and will allow us to measure how quickly we're reviewing and responding to customer tickets.

What you need to do:

- Head on over to the Confluence page and view the [training video](#)
 - In short, we're asking every team to run this new First Response Time (FRT) macro upon receipt of the ticket so we're able to track how quickly we're able to get eyes on a ticket once Advocacy has escalated it to Prod Dev.

When we go live on September 11, The OpEx dashboard will reflect the new FRT so we're able to see out of the TtA bucket, how quickly in that timeframe teams were able to look at tickets escalated by Advocacy.

Learnings

- Powerful Extrinsic Motivation to Adopt a Standard
- Can be Demotivating to start with an “F” Grade ... ratchet grades.

V5: Executive Zoom Out



Head of Engineering getting a 10,000 ft View

OPERATIONAL EXCELLENCE DASHBOARD Department Health Scores & Trends

		RELIABILITY		
		Remediation Items	SLOs	EPD
APPLICATIONS	Chat	C ↘	C ↔	B ↔
	Explore	A ↔	B ↗	A ↔
	Guide	A ↔	C ↘	C ↔
	Machine Learning	A ↔	B ↗	A ↔
	Mobile	D ↔	D ↔	C ↔
	Sell	B ↘		C ↗
	Support	B ↔	B ↔	C ↔
	Talk	A ↔	C ↘	A ↔
GLOBALIZATION	Internationalization	A ↔	A ↗	A ↔
	Localization	A ↔		A ↔
GROWTH AND MONETIZATION	Growth And Monetization Dept	D ↔	D ↘	B ↔
PLATFORM	Answer Bot	A ↔	A ↗	A ↔
	Classic Core	A ↔	B ↔	A ↔
	Customer Record	B ↘	B ↔	A ↔
	Melbourne Platform	D ↘	C ↘	A ↗
	Sunshine Conversations	B ↔	B ↗	A ↔
	Support Platform	B ↔	B ↔	C ↔
	Web Widget	A ↔	A ↔	A ↔
ZENDESK OS	Core Services	A ↔	B ↗	A ↔
	Engineering Productivity	B ↔	B ↔	D ↘
	Foundation	A ↗	C ↘	A ↔
	Resilience and Reliability Engineering	A ↔	A ↔	A ↔

Department Head gets a 1,000 ft View

OPERATIONAL EXCELLENCE DASHBOARD
Team Health Scores & Trends

Root: Engineering (1) | Department: A... (2) | Group | Team

	RELIABILITY				SECURITY					QUALITY		CUSTOMER		COSTS
	Remediation Items	SLOs	EPD	TFA (Sev 1)	SPR	VM Vulns	Container Vulns	App Vulns	Snyk Vulns	Smoke Tests	A11y	Z2 Problem Tickets	Z2 Jira Tickets	AWS Spend
APPS & EXTENSIBILITY	A ↔	B ↔	A ↔	n/a	D ↔	A ↑	A ↔	B ↘	A ↔	B ↔	n/a	B ↔	A ↔	n/a
APPS & EXTENSIBILITY TEAMS														
Customer DevX (Vegemite)	A ↔	A ↔	A ↗	n/a	n/a	A ↔	A ↔	A ↔	A ↔	B ↘	n/a	B ↘	A ↔	n/a
Dugong	A ↔		A ↔	n/a	n/a	A ↔	A ↔	A ↔	A ↔		n/a	A ↔	A ↔	n/a
Tardigrade	A ↔		A ↔	n/a	n/a	A ↔	A ↔	A ↔	A ↔		n/a	A ↔	A ↔	n/a
Wattle	A ↔	B ↔	A ↔	n/a	n/a	A ↔	A ↔	B ↘	A ↔	A ↗	n/a	A ↗	A ↔	n/a
Yowie	A ↔	A ↗	B ↔	n/a	n/a	A ↑	A ↔	A ↔	A ↔	C ↘	n/a	A ↗	A ↔	n/a
STRATEGIC INTEGRATIONS	A ↔	A ↔	A ↗	n/a	A ↔	A ↔	C ↔	A ↔	A ↔	A ↗	n/a	B ↘	B ↗	n/a
STRATEGIC INTEGRATIONS TEAMS														
Ocean	A ↔	A ↔	F ↔	n/a	n/a	A ↔	A ↔	A ↔	A ↔	A ↔	n/a	A ↔	A ↔	n/a
Pegasus	A ↔	A ↔	A ↔	n/a	n/a	A ↔	C ↔	A ↔	A ↔	B ↗	n/a	B ↘	A ↗	n/a

Learnings

- Being Compared to other Leaders is Motivating!
- Monthly Review Cadence was too long
- Pulling Insights news to move to Pushing Insights

V6: What's Coming



The Next Few Features

- Scheduled Slack Notifications with a team's score - **SHIPPED**
- Coming Soon:
 - Immediate notifications when scores go below thresholds
 - Are these Unit Tests or CI for your Organisation ?
 - New Standard: Feature Flags
 - New Standard: Error Reporting

pegasus-p Pegasus: :shorturl.at/arwGQ Edit

For more insights visit the [OpEx Dashboard](#). Today ▾

= # Out of SLA
 = # in Warning

RELIABILITY

		Incidents	Last 30 days (Sev 1&2): 0	
		Remediation	4	2
		SLOs	7 Day EB Depleted: 0	
		EPD	Monitor Coverage Gaps: 1	

SECURITY

		VM Vulns	0	0
		Container Vulns	25	1
		Application Vulns	0	0
		Snyk Vulns	7	0

QUALITY

		Smoke Tests	Last 7 Days: 99.3%	
		Accessibility	0	0

CUSTOMER HEALTH

		Z2 Problem Tickets	1	
		Z2 Jira Issues	4	0

COST OPTIMIZATION

AWS MTD Spend: \$0.82K Prev MTD Spend: \$13.53K

Please report any bugs to [#foundation-insights-testing](#). Thank you!

21 replies Last reply today at 2:26 PM

The Results

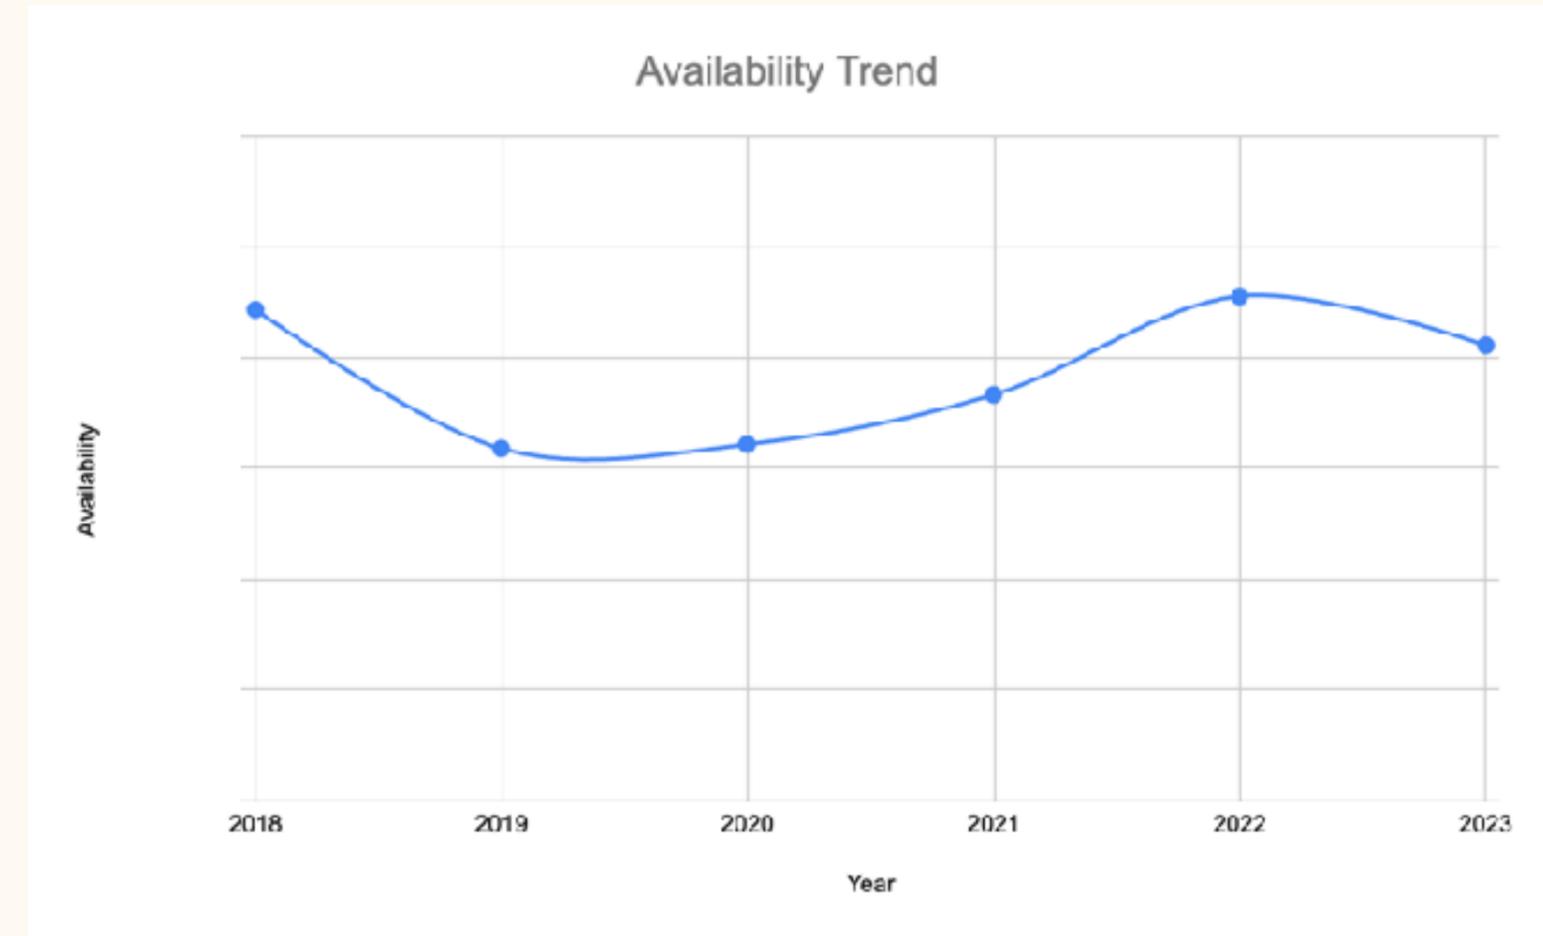


The Results

Has this helped our customers?

Has this helped our engineering teams?

- No more company wide production freezes since early 2021
... local freezes only.
- Empowered teams rolling out new standards, reduced rollout times
- Ensure ALL Systems (including Heritage Listed Systems !)
are in good health
- Improved availability for our customers
.... during 30% YoY traffic and data growth



Recommendations

- Don't "Part Time" the Data Warehouse and Reporting ... Dedicate People and Money
- Data Modelling
 - Consistent Org Structure in all Data Sources
 - Central Service Catalog
- It is a Product ... Seek Feedback from Users

The End

John Viner - www.linkedin.com/in/johnviner