


Site Reliability **Evangelism**

Practice start-up within an established web-presence



A bit about us

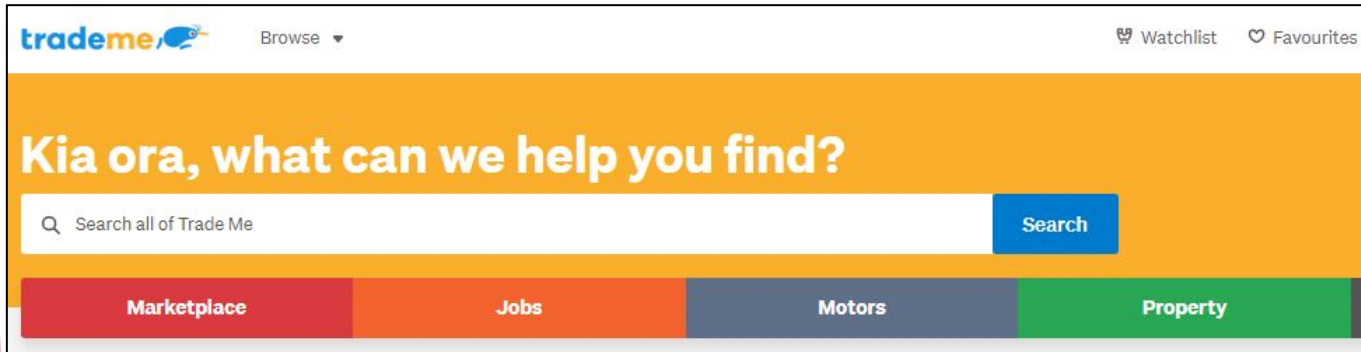
Trade Me is NZ's largest online auction and classifieds website

Piers (SRE Practice Lead) and **Catherine** (Technical Product Owner)



Structure & History

- 300 engineering staff
- Cloud Migration(s) ~18 months ago
- 3 Business Units

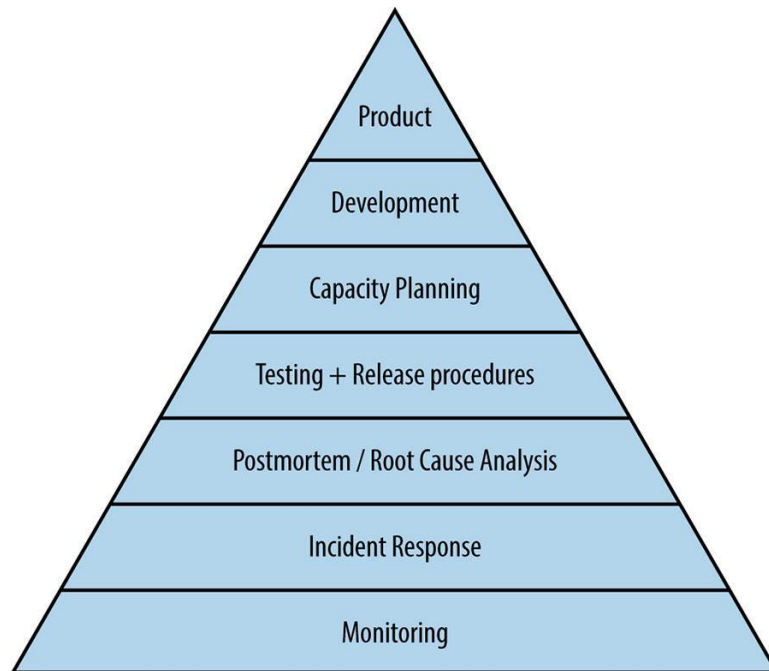


"Marketplace"

"Classifieds" Business Unit

Common Applications and Technology

Dickerson's Hierarchy of Service Reliability



Monitoring

Quality is a shared responsibility


- The squad shares responsibility for releasing their code
- Engineers are invested in tools that can give them more confidence that their deploy was successful
- Main code contributor is involved from beginning to end of the cycle of the case so quality is considered throughout
- We don't have a BAU team so engineers are responsible maintaining their systems



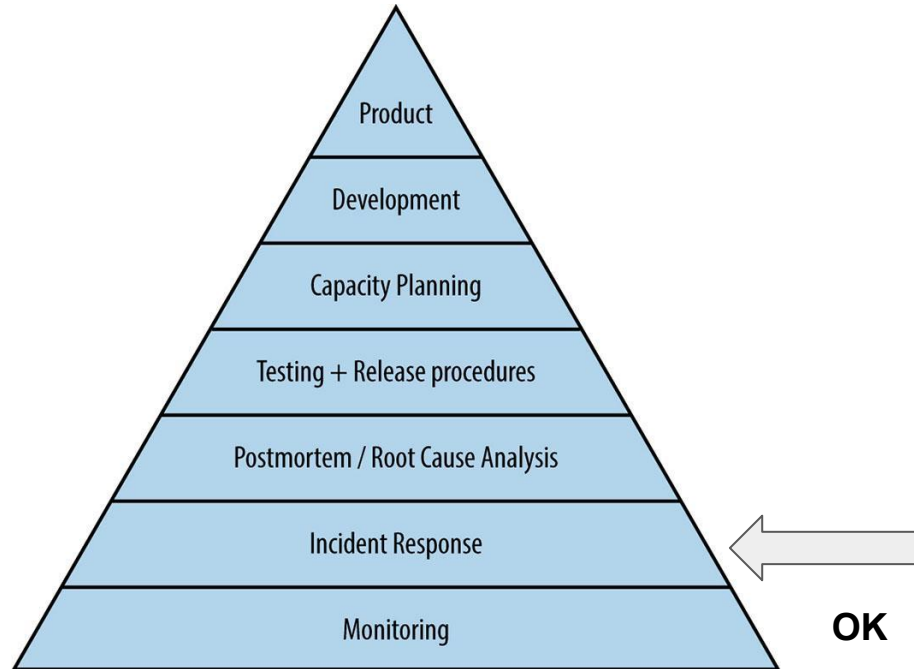
Monitoring

How we were monitoring our systems

- Synthetic production tests
- Deploy dashboards
- Golden Signal dashboards



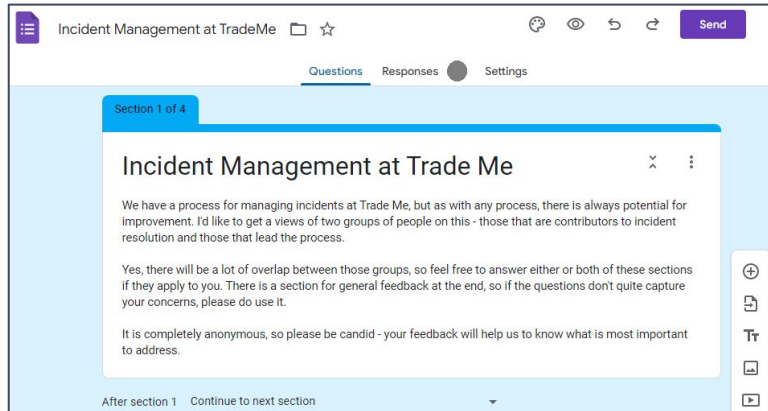
Dickerson's Hierarchy of Service Reliability



Incident Review

We had

- a very structured process
- a Trade Me built bot - “Bergerac”
- a lot of active incidents
- but also some avoidance



Incident Review

The biggest reasons for delays in resolving and/or closing an incident are....

Writing the POM and documenting all the actions

The daisy chain of incidents and the feeling that everyone is too busy/reluctant to perform actions and that POMs go unread so wasted effort.


People not wanting to write a POM

Waiting for POMs to be written.

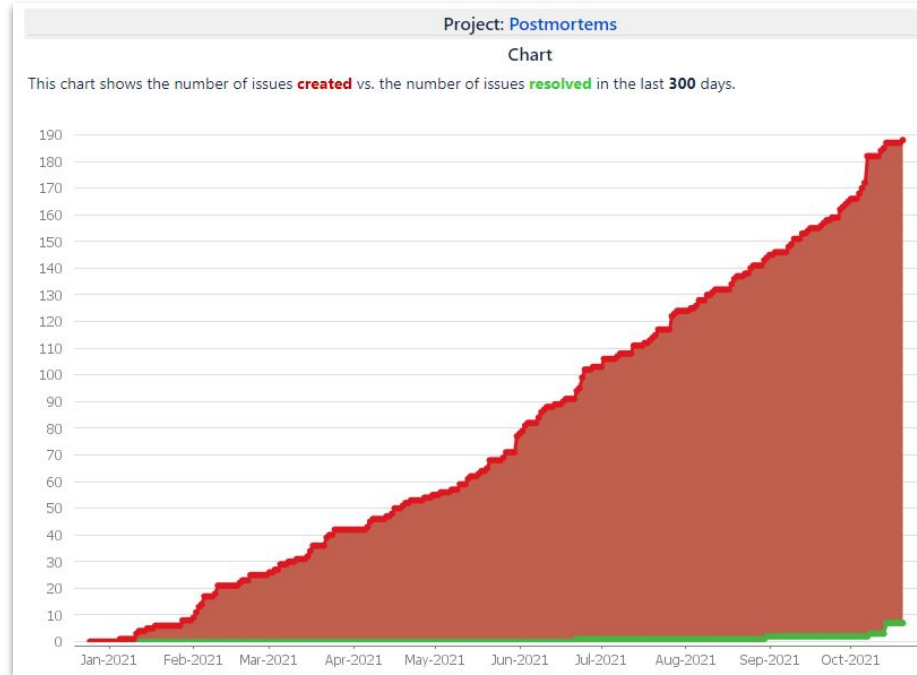
Getting POMs done AND reviewed to be signed off

From my limited experience it seems to be writing up the POM and getting everyone together for the post-mortem meeting.

From what I've seen people dread taking the time to fill out a potentially massive Post Mortem document and having a meeting after the fix, I guess people just want to get on with their work - maybe they're also fearing being blamed, are yet to see the value of it, or have had bad experiences? im just guessing




Incident Review



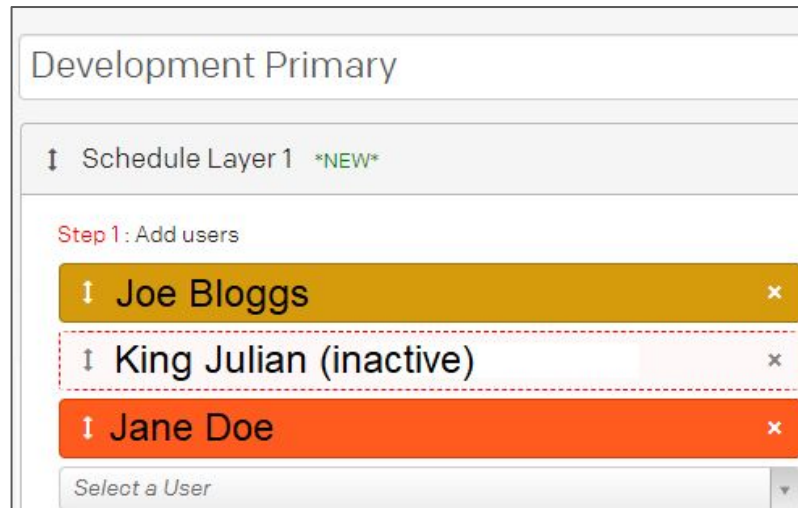
Incident Review - Changes Made

- Reviews expected on Sev1+ Sev2 incidents only
- Dropped Sev4's completely
- Started a community of debrief facilitators
- SaaS incident management tooling with big benefits
- Learning as primary objective
 - Laundry lists of actions discouraged
 - Challenge Dashboard proliferation
 - Prefer transformative change



Incident Response - On call

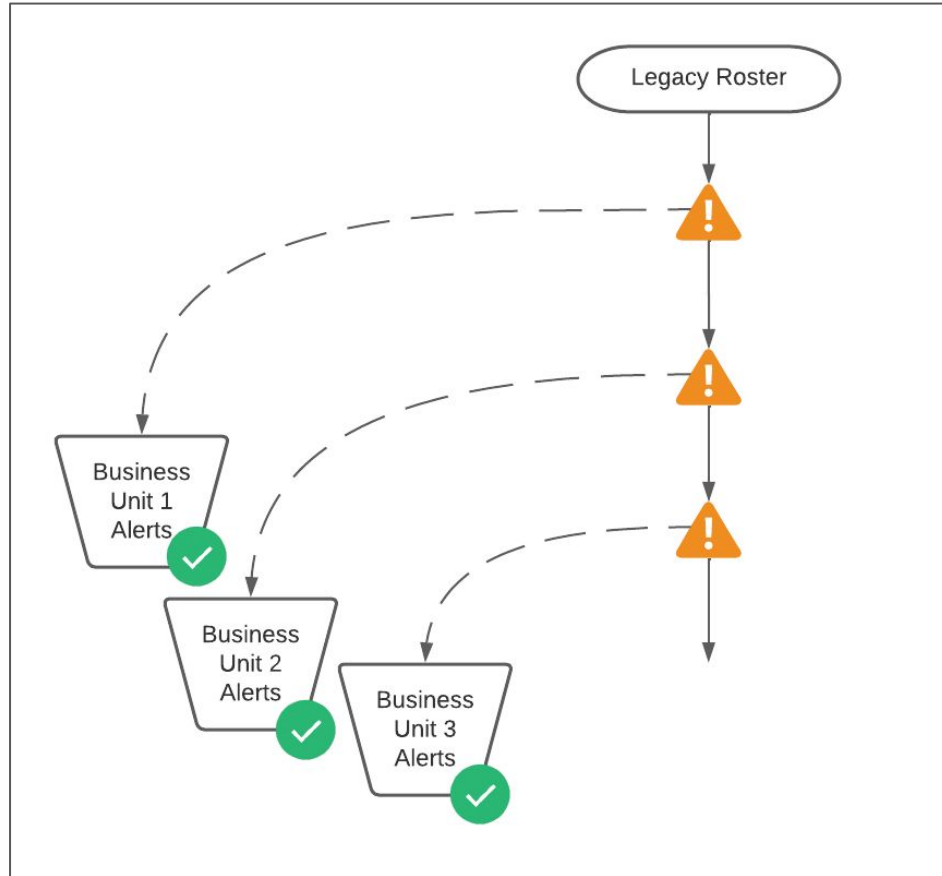
This time last year...



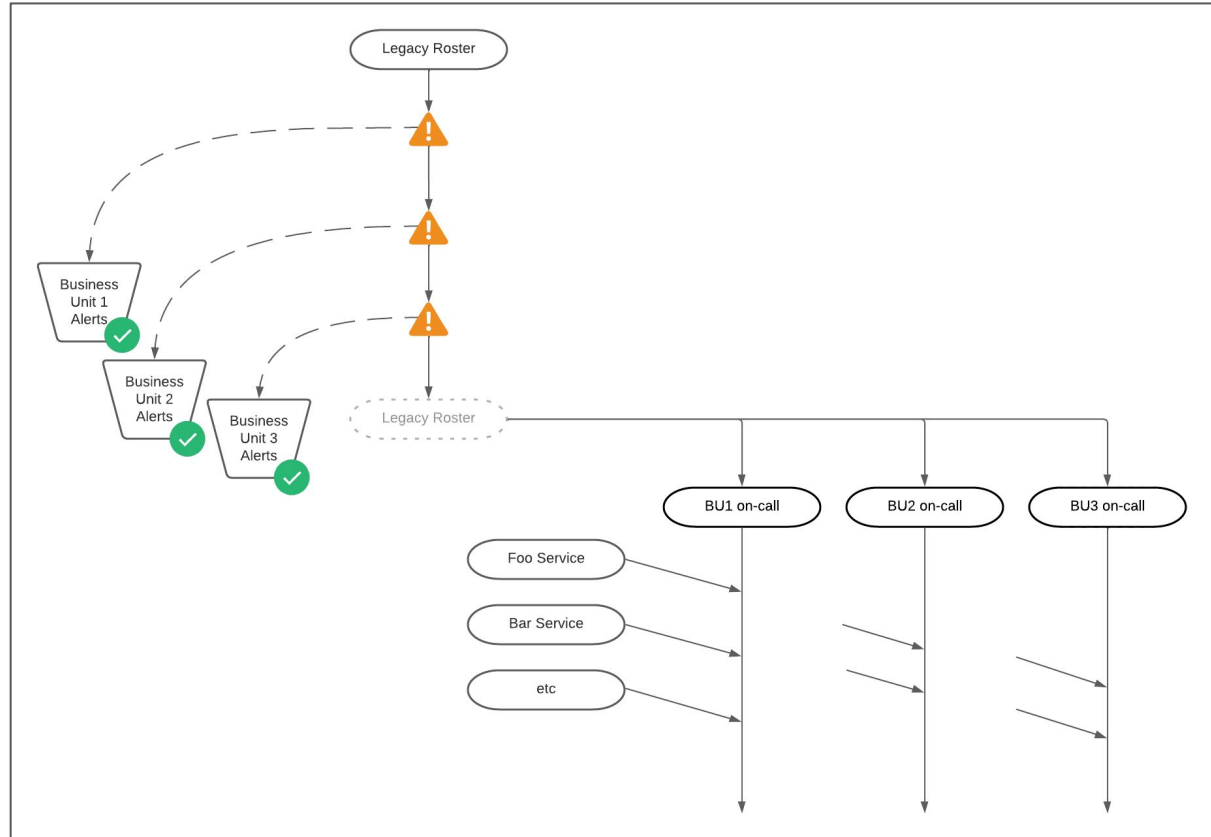
The screenshot shows a web interface for configuring an on-call schedule. At the top, there is a header "Development Primary". Below it is a section titled "Schedule Layer 1" with a green "*NEW*" indicator. Underneath, "Step 1: Add users" is displayed. Three users are listed in a vertical stack, each with a dropdown arrow on the left and a close button (x) on the right. The first user, "Joe Bloggs", is highlighted with a yellow background. The second user, "King Julian (inactive)", is highlighted with a light pink background and is enclosed in a dashed red border. The third user, "Jane Doe", is highlighted with an orange background. At the bottom of the list is a search box labeled "Select a User" with a dropdown arrow.

User	Status
Joe Bloggs	Active
King Julian (inactive)	Inactive
Jane Doe	Active

Incident Response - On call



Incident Response - On call



Cultural Reinforcement

The technical challenge of standing up on-call was just the first step

- Needed to work with our People Leaders to reinforce key ideas
- We have a great DELIVERY culture at Trade Me
- We need to turn “Build Teams” into “Build & **Run** Teams”

November 2022

**Growing Operational
Maturity**



Service Level Objectives


We have an annual OKR to get SLOs wrapped around our core systems

- Struggled to get interest even when the Dev benefits were called out
- Prioritisation was hard (both order of approach, and getting time)
- Leveraged DR plans to source the true core systems
- Proper Work requests into our quarterly planning

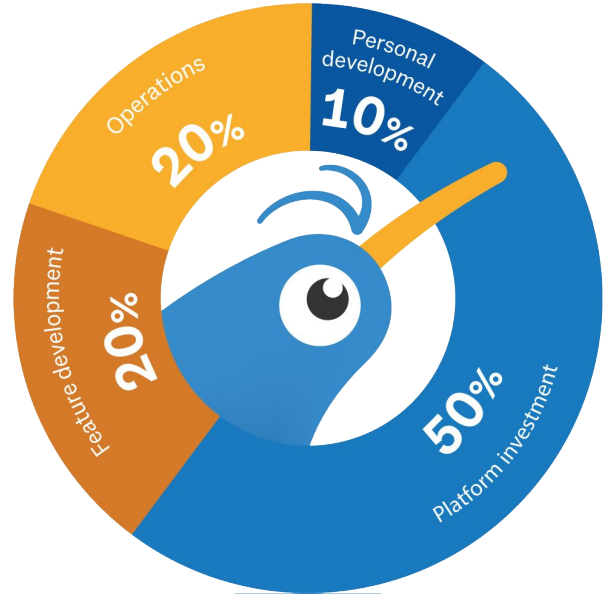


The “Effort Pie”

Stream aligned teams



Platform teams



Operational Review Meetings


- SLOs / Key Performance Stats
- Any incidents, review or actions since the last meeting
- Any emerging risks and / or related actions that need to go into the risk register
- Maintenance plans - what to spend the 20% time on next quarter

SLO 	Last 2 hours... 	Last 24 hours... 	Last 7 days... 	Last 28 days... 
95 % 1 day	94.91 %	93.6 %	93.86 %	94.08 %
99.9 % 1 day	100 %	100 %	100 %	100 %
95 % 1 day	96.68 %	96.27 %	96.12 %	96.19 %



In Summary

- *Choose a safe, foundational project to build early credibility*
- *Simplify the messages / over - communicate*
- *Investing in cultural change is a long game*
- *SRE is not just for Site Reliability Engineers*
- *Get reliability work onto committed roadmaps*



Referenced Resources

- John Tucker's System ERD diagram
<https://john-tucker.medium.com/backstage-by-example-part-1-a18e74849240>
- Ask Ms O11y blog - ...am I a monster?
<https://www.honeycomb.io/blog/devops-on-call>
- Team Topologies by Manuel Pais and Matthew Skelton
- Site Reliability Engineering by Betsy Beyer, Chris Jones, Jennifer Petoff and Niall Richard Murphy

