

# Proposed Agreement

We would like to define the SRE/Foo Service relationship as follows:

- 1 Software Reliability Engineer (SRE) would be dedicated to work with Foo Service,
- This individual would continue to report into Alex, but work as close as required with the Foo Service team
- This individual would have a desk in the SRE/Observability space and expected to work closely with other SRE's
- The Foo Service focused SRE will initially ensure that:
  - We have understood service level objectives (SLO's) for each service/job defined in the onboarding list.
  - Ensure appropriate monitoring/observability is in place by defining Service Level Indicators (SLI's) for all defined SLO's
- As part of on-boarding, the SRE would initially shadow then actively participate in triaging issues. SRE can shadow after-hours work, or debrief the morning after on any issue, but not be the primary escalation yet
- The Foo Service focused SRE will cross train the SRE team on Foo Service operations
- Once the proper SLI observability is in place and we are familiar with operational processes, the SREs as a whole will begin to be the afterhours escalation. While we expect this transition to occur in a short time frame, the actual cut off date will need to be decided between the SRE, SRE management, and Foo Service management.
- SRE will then be the primary escalation. The embedded SRE will continue towards becoming a Subject Matter Expert (SME) for Foo Service services and jobs, but we will require always having product team engineers available as needed for secondary escalations and severe incidents
- While engaged, An SRE should be self-directing work and focus on tasks that improve the reliability of the service. SRE cannot unilaterally make changes to services, but needs justification. Changes and PR's should be discussed with

and approved by the product team. Any work can reside in existing jira projects but should have the “SRE-Foo Service” component assigned. Work could include:

- Becoming incrementally more involved in pushes
  - Monitoring error aggregation systems and alerts, understanding how to improve signal to noise ratio and automated classification
  - Optimizing services to improve performance, decrease risk, and increase reliability
  - Coordinating Post Mortems
  - Being an advocate to the product team on architecture/design that minimizes risk
  - Configuring/Tuning Productions Systems
  - Exposing additional metrics/logging/tracing information
  - Automating build processes or any manual tasks
  - Documenting and exercising tools, processes, and automation needed to manage the service
- To help with work prioritization, and to track our performance the SRE team will be defining and referencing KPIs based on:
    - The number of escalations. Work to reduce interruptions should be prioritized, and ensuring these are low are a KPI
    - The amount of “toil” (defined as manual work, manual script runs, repetitive tasks, work that does not permanently improve the service). Our initial target for this work is expected to be high, but ideally should take up no more than 50% of an SRE’s time
    - NPS type questionnaires from SRE team members around quality of life/engagement levels
- ***Dedicated SRE resourcing is not permanent, we should be actively trying to automate our need away.*** SRE and Team management will meet periodically to review performance against Key Performance Indicators (KPI’s) and re-evaluate the engagement model. Should performance against these KPI’s not meet expectations, or if the type of work SRE is doing not be focused on improving reliability, or if the service is moving to long-term support or end-of-life, possible outcomes would be:
    - Reduction or suspension of the SRE engagement (Giving back the pager or Business hours only support)
    - Changes to SLOs to reduce alerts/interruptions (effectively agreeing to a larger acceptable range of reliability)