

Qualité et vélocité: grâce au
Lean vous n'avez pas à choisir

SRE



Victor Boissière

Lead SRE
Platform & Storage



Mathieu Corbin

Senior staff SRE
Platform

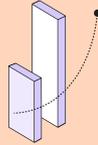
mcorbin.fr

Qonto

Créer la solution de finance qui énergise les PME et les freelances



400k clients



€622M Raised



5 offices
(FR, IT, DE, SP, SE)



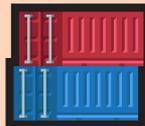
1300+ Qontoers



Quelques chiffres



4 équipes SRE



9 clusters Kubernetes



100 déploiements en prod
par jour



323 microservices



8 comptes AWS



Go, Python & Ruby

Ambition - Eliminer le gaspillage

problèmes courants



Objectifs

1 SRE pour traiter le RUN chaque semaine

0% de RUN

Astreinte qui sonne quasiment toutes les nuits

0 incident

Pas de déploiement le vendredi

Réduire le lead time d'implémentation des fonctionnalités

Trier et modifier régulièrement le backlog

Pas de backlog

MAJ Kubernetes = 1 mois d'un SRE par an

Mettre à jour Kubernetes en moins de 2H

*"The Toyota style is not to create results by working hard. It is a system that says there is **no limit to people's creativity**.*

*People don't go to Toyota to 'work' **they go there to 'think'.**"*



Taichii Ohno

1912 - 1990

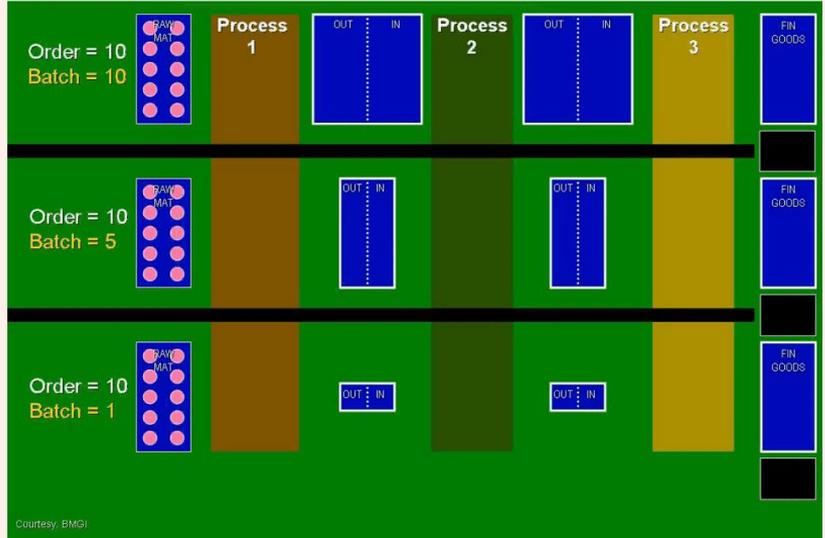
Comment atteindre notre ambition?



Préconception n°1

Travailler sur plusieurs projets à la fois augmente la productivité

Flow – One piece flow versus Batch Production



[link](#)

L'origine du backlog

Hello @sreduty 🙌,
I would like to create a PostgreSQL read replica to ingest data to Elasticsearch. **Can you create it?**

Hello @sreduty 🙌,
I cannot pull docker images locally. **Can you help me?**

Hello @sreduty 🙌,
My CI job is failing because out of disk space. I've relaunched it but same issue. **Can you fix?**

Hello @sreduty 🙌,
We would need an SFTP for a new feature, **can you set it up?**



Surcharge = backlog = priorisation

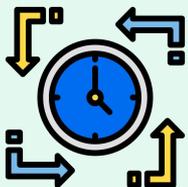
Stopper les demandes avec du self-service

Hello @sreduty 🙌,
I would like to create a PostgreSQL read replica to ingest data to Elasticsearch. **Can you create it?**

```
qontoctl rds create-read-replica  
-i db.r5.large  
--source-instance postgresql-prod
```

Hello @sreduty 🙌,
I cannot pull docker images locally. **Can you help me?**

```
qontoctl config troubleshoot  
Start analyzing your system to find potential known issues:  
✅ Shell environment, no issue found  
✅ AWS Local Configuration, no issue found  
✅ Kubectl Local Configuration, no issue found  
✅ AWS CLI Version (2.7.18)  
✅ Kubectl Version (1.24.3)  
✅ Docker Local Configuration, no issue found
```



Demandes régulières

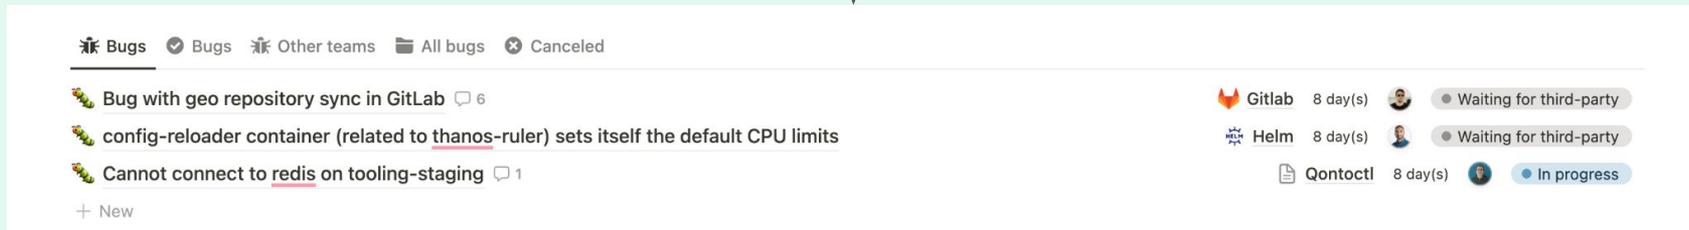
Toil

Traiter un bug à la fois en dehors du RUN

Hello @sreduty 🙌,
My CI job is failing because out of disk space. I've relaunched it but same issue. **Can you fix?**

- Correctif + monitoring
- Plusieurs stratégies nécessitant des tests
- Bug qui se répète

Si investigation > 15min



The screenshot shows a Jira bug list with the following items:

- Bugs** (selected) | Bugs | Other teams | All bugs | Canceled
- Bug with geo repository sync in GitLab** (6 comments) | Assigned to **Gitlab** (8 day(s)) | Status: **Waiting for third-party**
- config-reloader container (related to thanos-ruler) sets itself the default CPU limits** | Assigned to **Helm** (8 day(s)) | Status: **Waiting for third-party**
- Cannot connect to redis on tooling-staging** (1 comment) | Assigned to **Qontoctl** (8 day(s)) | Status: **In progress**

+ New

Traiter un projet à la fois

Hello @sreduty 🙌,
We would need an SFTP for a new feature, **can you set it up?**

Staging?

Accès?

Prix?

Compliance?

Alertes?

Sauvegarde?

Cycle de vie?

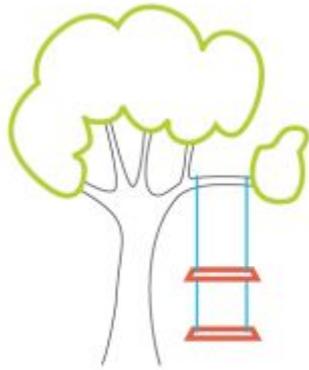
Object storage?

Logs?

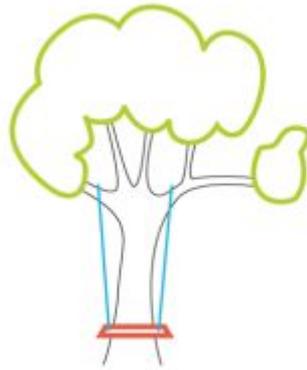
Maintenance?

Préconception n°2

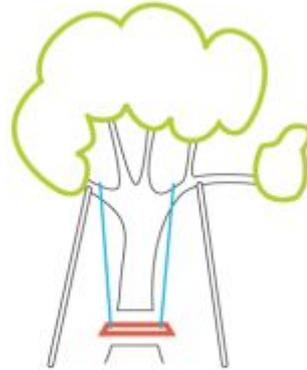
Il vaut mieux se lancer directement sur le projet que de passer une semaine à réfléchir



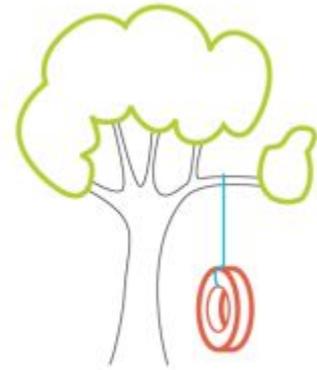
Le besoin
exprimé
par le client



Le besoin
compris par
le chef de produit

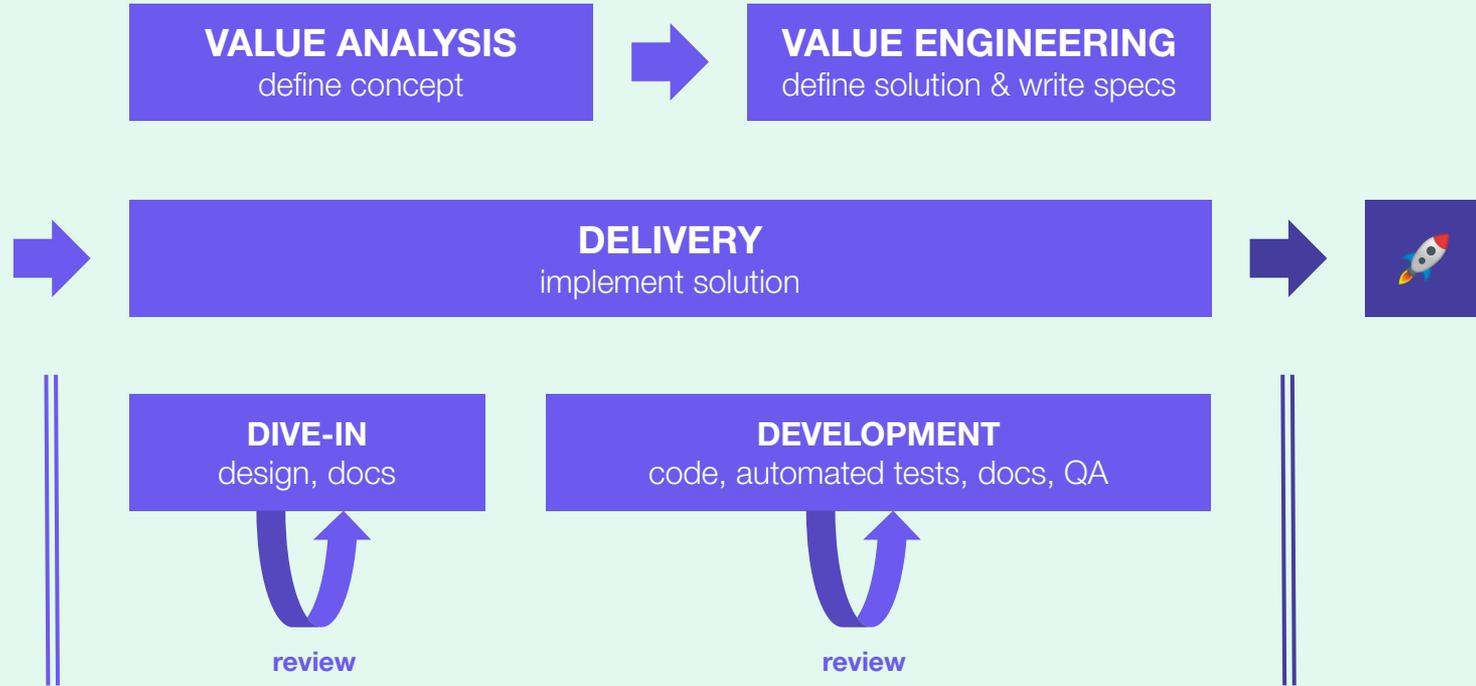


Le produit
conçu par
l'ingénieur



Ce dont
le client avait
réellement besoin

Design avant l'implémentation



Visual management avec le Kanban

The image displays a Kanban board interface with the following sections:

- Value engineering (1):** Card titled "Tooling revamp" with assignees Ramy Nassif and Allan Jacquet. Status: "4 days remaining for the VE", dated December 15, 2023.
- Dive-in (0):** Empty column.
- Work In Progress (2):** Two cards:
 - "Logs to Loki": Assignee Morgan Fourny. Status: "There is no sub-tasks identified", "Delivery is due Today", dated November 17, 2023.
 - "Decommission old microservice": Assignee Allan Jacquet. Status: "There is no sub-tasks identified", "Delivery is due Today", dated November 17, 2023.
- Done (3):** Three cards:
 - "Improve secrets management": Assignee Ramy Nassif. Dated October 13, 2023.
 - "Higher-Level abstraction Helm Chart for generic apps": Assignees Christophe Collot and Mathieu Corbin, with Clovis Delarue. Dated November 15, 2023.
 - "MLOps Setup for the Data Science Team": Assignee Allan Jacquet. Dated November 17, 2023.
- Hidden groups:** A sidebar on the right showing counts for "Waiting maintenance" (0), "Not started" (0), "Archived" (358), "Canceled" (0), and "No Status" (1).
- What's next? (Pull):** A section at the bottom with a "CFT project" card titled "DNS routing" by Mathieu Corbin, with a status of "SRE Platform Reliability" and a wait time of "2 day(s)".

Aligner d'abord

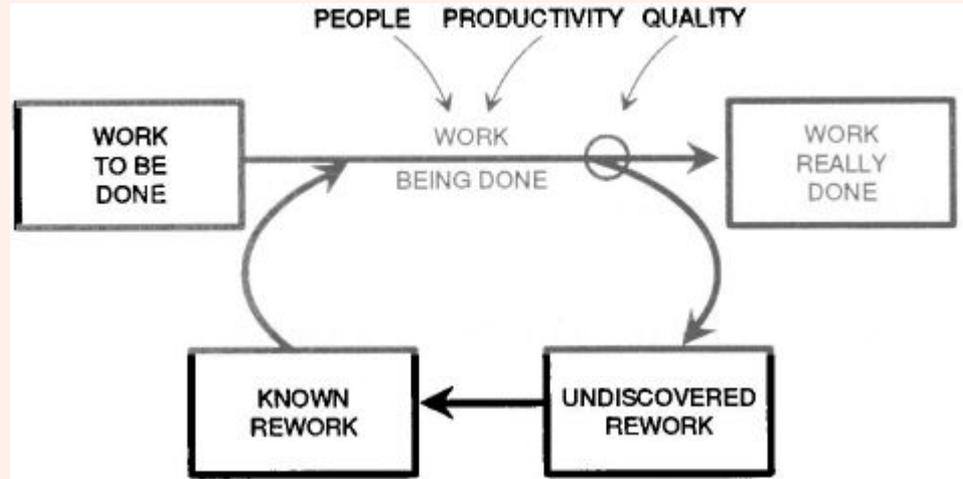
- 180 biais cognitifs
- Top-down contre productif
- Intuitions ou faits



Regarder ensemble la même chose grâce aux faits

Préconception n°3

Si on est pas sûr on tente et on pourra toujours corriger à la fin



On s'arrête en cas de doute



Amine Benseddik 5 days ago

🙋 FYI, we had an **andon** this morning, with [@Victor Boissiere](#) [@Stéphane Tang](#) [@Mathieu Corbin](#) and me, regarding the Gitlab Autoscaling project. We discovered a blocking **issue** in the "Gitlab Runner Autoscaler" that prevented us from proceeding with this solution.

👉 As a result, we made the decision to abandon this solution and develop our own autoscaler. Here is the new **VE**.

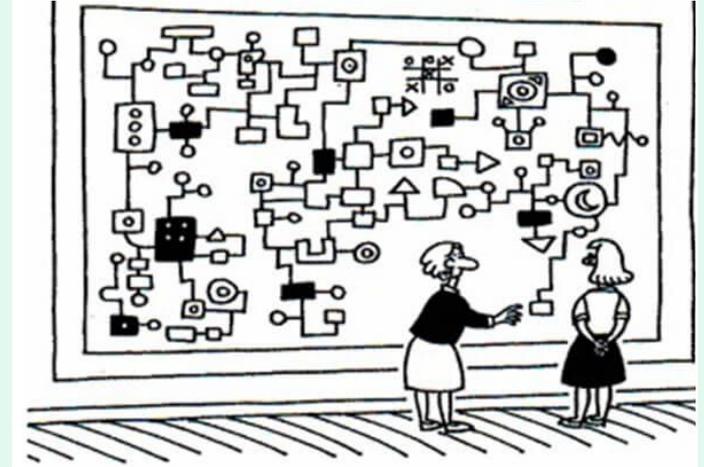


Stopper le développement augmente la vélocité

Les incidents sont des problèmes de qualité

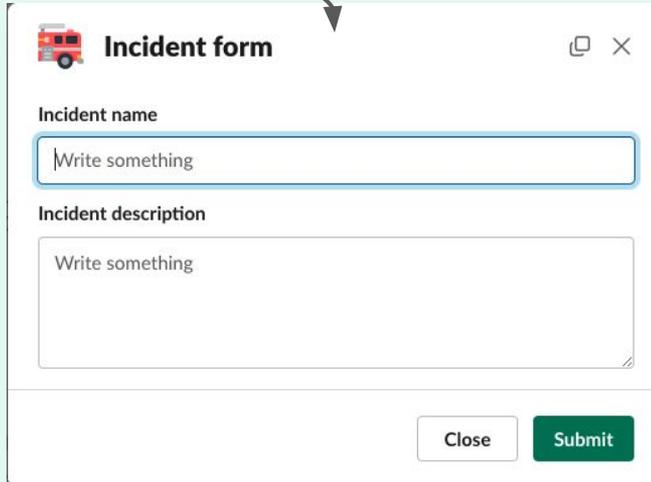
Problèmes courants

- Incidents non déclarés
- Pas de centralisation des incidents
- Processus lourd pour créer un incident



Tout le monde est concerné

#incident-report slack channel
/incident



Incident form

Incident name

Write something

Incident description

Write something

Close Submit



Incident automation APP 12:59 PM

New incident reported by @Mathieu Corbin

Staging: Kubernetes cluster 3 is misbehaving: We have performance issues and pods crashing on the staging 3 kubernetes cluster

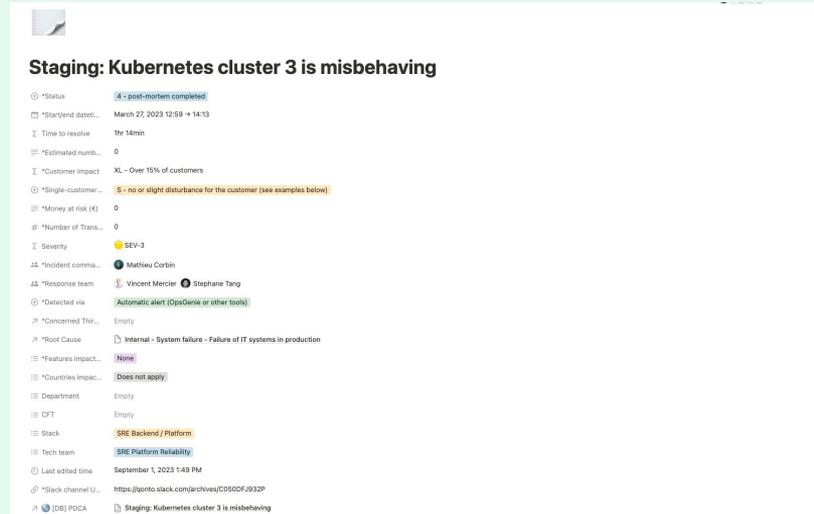
Slack channel

#2023-03-27-staging-kubernetes-cluster-3-is-misbehaving

Report

Incident report

Page Notion créée automatiquement



Staging: Kubernetes cluster 3 is misbehaving

- Status: 4 - post-mortem completed
- Start/end date: March 27, 2023 12:59 - 14:13
- Time to resolve: 1hr 14min
- Estimated number: 0
- Customer impact: XL - Over 10% of customers
- Single customer impact: S - no or slight disturbance for the customer (see examples below)
- Money at risk (€): 0
- Number of Trans: 0
- Severity: SEV-3
- Incident command: Mathieu Corbin
- Response team: Vincent Mercier, Stéphane Tang
- Detected via: Automatic alert (OpsGenie or other tools)
- Concerned Teams: Empty
- Root Cause: Internal - System failure - Failure of IT systems in production
- Features impacted: None
- Countries impacted: Does not apply
- Department: Empty
- CFT: Empty
- Stack: SRE Backend / Platform
- Tech team: SRE Platform Reliability
- Last edited time: September 1, 2023 1:49 PM
- Slack channel URL: https://ponto.slack.com/archives/C05C0DFJ932P
- DB/POCA: Staging: Kubernetes cluster 3 is misbehaving

Les postmortems pour apprendre

Problèmes courants

- Ne pas en faire
- Analyse biaisée du problème
- Le problème existe encore à certains endroits



Correctifs efficaces avec le PDCA (Plan Do Check Act)

- Des services ont eu des problèmes de résolution DNS. Why ?
 - Les résolutions DNS ont échouées. Why ?
 - CoreDNS ne reçoit pas les requêtes. Why ?
 - Nous avons atteint une limite de paquets/sec sur nos instances CoreDNS. Why ?
 - Les requêtes sont centralisées sur seulement 2 instances. Why ?
 - Nous n'avions pas connaissance de ces limites

Conclusion

Est ce que la qualité et la vélocité ne sont finalement pas la même chose ?