



# Keeping Pinterest Running

Joe Gordon  
2 February 2016

# What is Pinterest?



The visual bookmarking tool that helps you discover and save creative ideas.

©2016 Pinterest. All rights reserved.



# Software v. Service

# Software v. Service

- **Stable branches**
- **Drivers and configurations**
- **Support matrix**
- **Dependency versions**
- **Developers support their own service**
  - On call rotation
  - **Aligns incentives**
  - Monitoring & alerting built in from day one
- **Testing against **production** traffic**



# SRE at Pinterest



*the worlds most intense pit crew*





*fire fighting @scale*



*changing tires while driving at 100 mph!*



**What do SREs **focus** on?**

# Operational Maturity

**Initial (chaotic, ad hoc, individual heroics)** - the starting point for use of a new or undocumented repeat process.



```
graph TD; A[Initial (chaotic, ad hoc, individual heroics) - the starting point for use of a new or undocumented repeat process.] --> B[Repeatable - the process is at least documented sufficiently such that repeating the same steps may be attempted.]; B --> C[Defined - the process is defined/confirmed as a standard business processes.]; C --> D[Managed - the process is quantitatively managed in accordance with agreed-upon metrics.]; D --> E[Optimizing - process management includes deliberate process optimization/improvement.];
```

**Repeatable** - the process is at least documented sufficiently such that repeating the same steps may be attempted.

**Defined** - the process is defined/confirmed as a standard business processes.

**Managed** - the process is quantitatively managed in accordance with agreed-upon metrics.


**Optimizing** - process management includes deliberate process optimization/improvement.

# Operational Excellence

**Availability**

**Efficiency**


**Performance**



**Developer Productivity**

# Operational Excellence

## Software Engineering



## Infrastructure Engineering

## Process Engineering

- ❖ Resiliency features
- ❖ Monitoring & SLAs
- ❖ Incident Management
- ❖ Runbooks

- ❖ Effective Utilization
- ❖ Efficiency improvements
- ❖ Capacity Management

- ❖ Site speed & latency
- ❖ Load / Redline testing

- ❖ Repeatable deploys
- ❖ Automation
- ❖ Change management



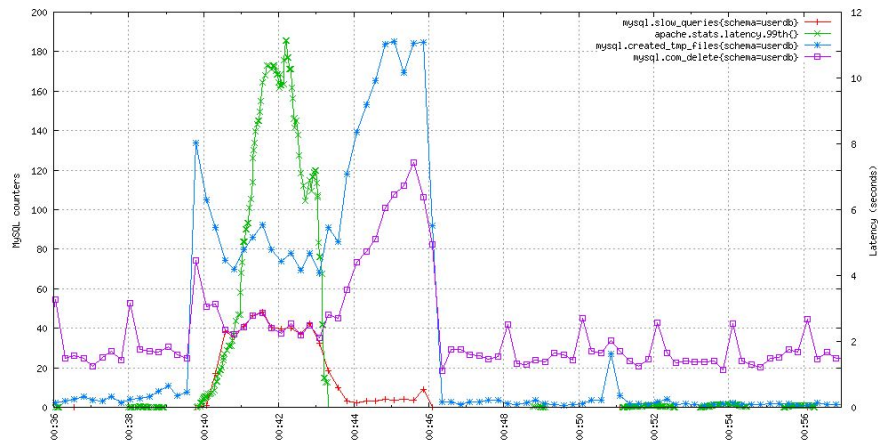
# Visibility

*Insight into the system*



# Visibility

- **Data Driven**
- **Cornerstone for many things we do**
  - Measure and enforce SLA (Service Level Agreement)
  - Debug issues
  - Capacity planning
- **Time series data - TSDB**
- **Metrics**
  - System
  - Service
  - Dependencies
  - Latencies
- **Alerting**
- **ELK stack for real time log collection**





# Deployments

# Deployment Requirements


- No impact to end user
- Change history
- **Easy**

# Staging and Canary

*Canary in a Coal mine.  
Rabbit in a Sarin gas plant*




# Canary vs Staging





# Teletraan

*deploy system*



# Teletraan

## *Features*

- **Rollback**
- **Hotfix**
- **Rolling deploy**
- **Staging and testing**
- **Visibility & Usability**

# Teletraan

## *Design*

- client-server model
- PRE/POST-DOWNLOAD
- PRE/POST-RESTART
- RESTART
- RBAC

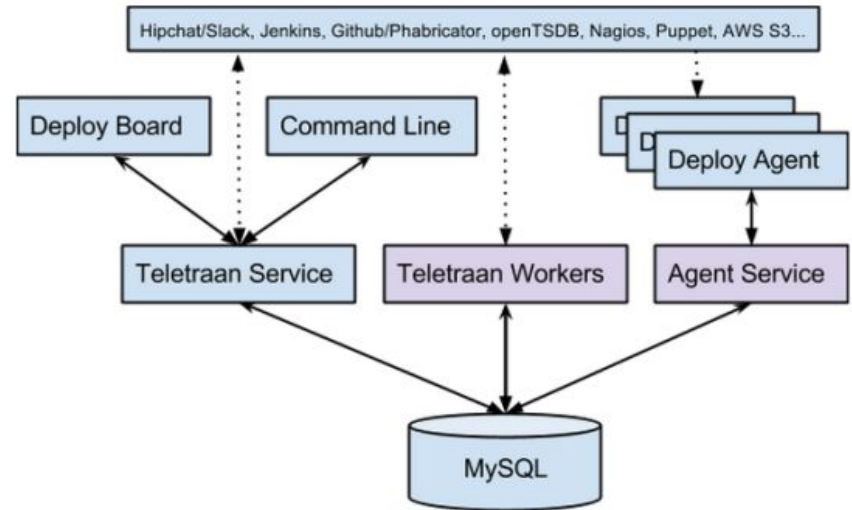
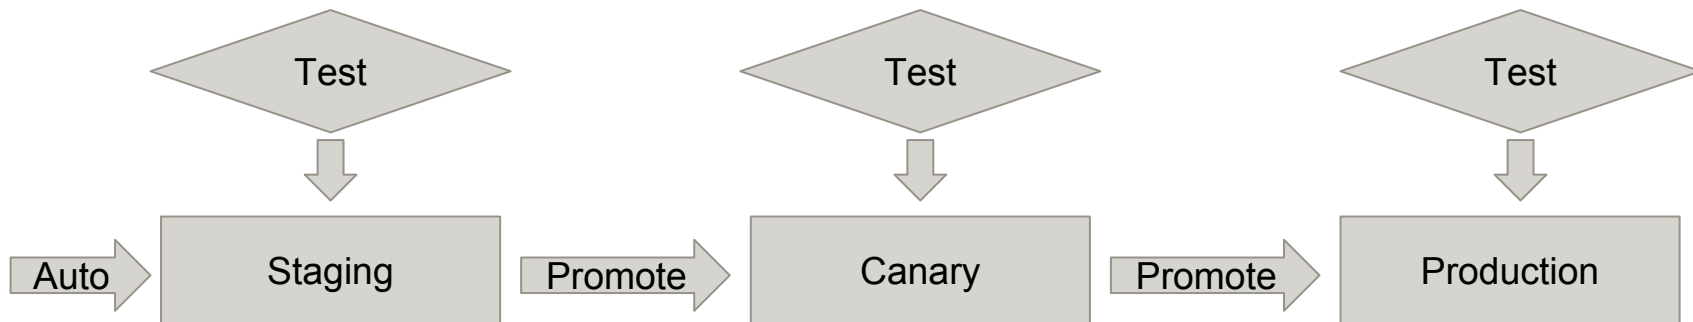


Figure 1. Teletraan design.

# Teletraan

*Advanced Features*

- **Pause/Resume**
- **Acceptance Testing**
- **Auto Deploy**
- **Autoscaling**





# Postmortems

*Learn from our mistakes*



# Postmortems

- **Blameless**
- Incident Manager
- Impact
- Outage Type
- Method of Detection
- Timeline
- Root Cause
- Restoration Details
- Actionable Items



# Production Readiness Review

*Pre-mortem?*

# Production Readiness Review

- Dependencies
- Define an SLA
- Alerting
- Capacity planning
- Testing
- On call rotation
- Decider to turn feature off if needed
- Incremental launch plan
- **Rate limiting**



# Public Cloud

## *Issues*

# Public Cloud

## *Issues*

- “If you get an **InsufficientInstanceCapacity** error when you try to launch an instance, AWS does not currently have enough available capacity to service your request.”
  - Cloud is not infinite
  - Reserved instances
  - Capacity planning
- **RequestLimitExceeded**: “The maximum request rate permitted by the Amazon EC2 APIs has been exceeded for your account.”
  - Includes **DescribeInstances**
  - Use internal mirror (powered by elasticsearch)
- Noise Neighbors
- Rightsizing
- Ownership





# Open Sourced Tools

# Open Sourced Tools

- **mysql\_utils**
  - MySQL Management Tools for the Cloud
- **thrift-tools**
  - thrift-tools is a library and a set of tools to introspect Apache Thrift traffic
- **secor**
  - Secor is a service implementing Kafka log persistence
- **pymemcache**
  - A comprehensive, fast, pure-Python memcached client
- **pinrepo**
  - Artifact Repo
- **Teletraan**

More at: <https://github.com/pinterest>

