



# Challenges in starting an SRE team from scratch in an Enterprise

SRECON20 Americas

# Your Speakers



Graeme Bye



Wayne Bridgman



Anand Bobade



Amreen Firdouse



Shiv Patil



Pauline Narvas

# Purpose, goal and strategy



# SRE at BT

Let there be...



# Where to start?

Automation

TOIL

Security



The Pillars

Observability

Performance

Tools

Reliability

Cost Control

# We've migrated to the cloud... Now what?



# Key challenges...





# OUR KEY CHALLENGES & SOLUTIONS


**The challenge: no SREs!**

# Forming the team



# What do we need to succeed??

SUCCESS



What people think it looks like

SUCCESS



What it really looks like




# The challenge: starting from scratch



# Our goals




Faster rollout with HA  
& reliability



Best practice &  
automation



Reduce cloud  
sprawl



Shared goals  
across teams

# Our principles



1.

## Automation

- If an existing process cannot be automated, it must be redesigned and replaced.
- If a proposed process cannot be automated, it will be rejected.



2.

## Disposable Infrastructure

- Servers are disposable and treated like cattle.
- Servers live in auto-scaling groups that self-heal.
- Servers are provisioned from images that are fully equipped and operational; no post provisioning tinkering allowed.
- Application servers are stateless.



3.

## Continuous Integration

- All code changes are made in Git via pull requests, verified, and approved.
- Functional tests run on every deploy, preventing (or rolling back) deploys when the build fails.



4.

## Continuous Deployment

- Deploys do not require any human interaction other than to instigate them.
- Deploy time matters and engineers should strive to make it faster.
- Rollbacks happen automatically when a failed deploy is detected.
- It is easy to tell which commit is deployed.



5.

## Monitoring & Alerting

- All systems are monitored for the critical "four golden signals" metrics - latency, traffic, error rate and saturation.
- Metrics are easily available and consumable in a single interface.



6.

## Security

- Security is automated and baked into everything.
- Security checks are run as part of CI/CD.
- Intrusion detection systems are in place.
- As few infrastructure components as possible are publicly accessible, ideally zero.

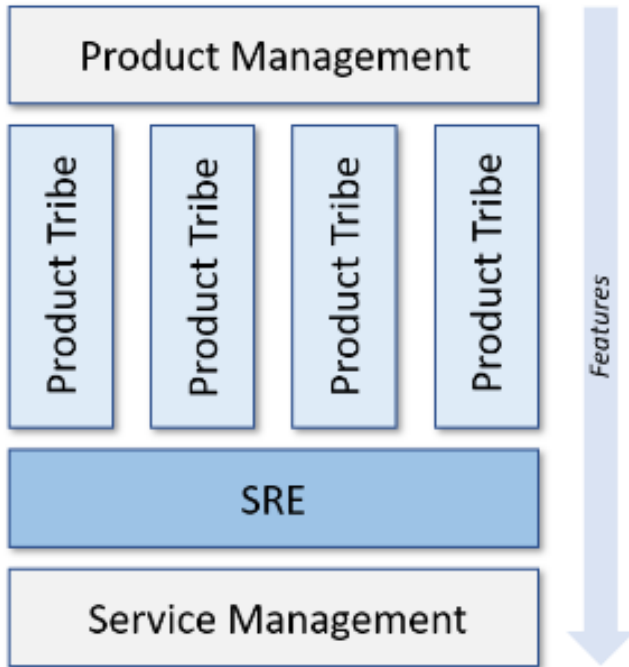


7.

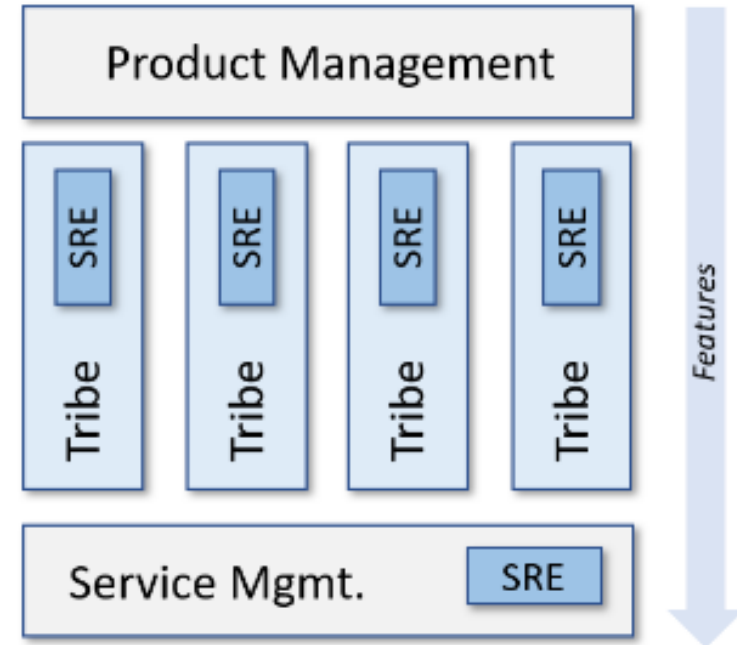
## Cloud Computing Standards



## Phase 1



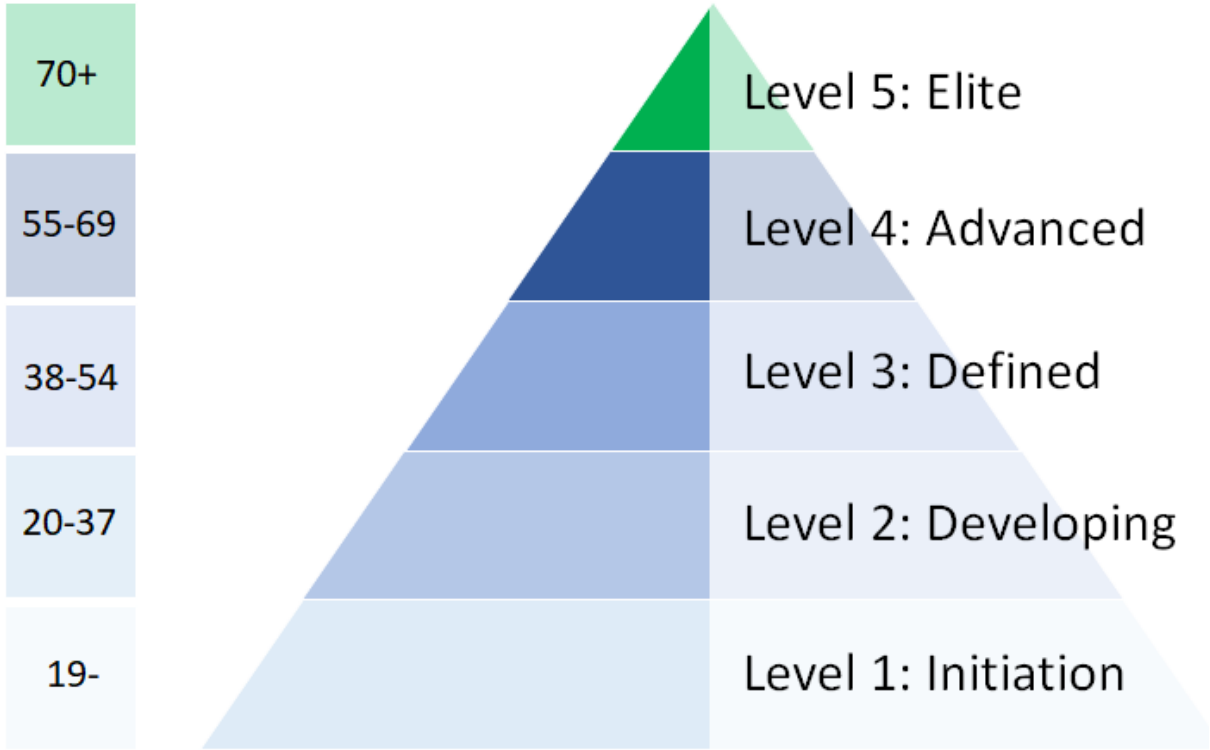
## Phase 2



## Doing some serious self-reflection...



# SRE Maturity Report




Ref	Activity	Rating*				
		1	2	3	4	5
<b>BT Specific SRE Activities</b>						
1	Disaster recovery is periodically (at least annually) tested against non-production environments.	x				
2	SRE standards are documented.					x
3	SRE standards are continuously audited.			x		
4	Manual operations tasks are regularly reviewed and automated.		x			
5	Game days are run regularly to check the efficiency of change processes.	x				
6	Development teams are provided with observability of production performance in real time.	x				
7	Development teams are provided with observability of production incidents.	x				
8	There is a clear SRE roadmap for the next 12 months.		x			
9	The SRE team supports the Digital Security Manager with Security Incidents.		x			
<b>BT SRE Score</b>		<b>16</b>				
<b>Traditional SRE Activities</b>						
10	SRE team members participate in incident response procedures on a regular (i.e. weekly) basis.	x				
11	There is a culture of authoring blameless post-mortems for production incidents.		x			
12	There is a process to manage production incidents.	x				
13	There is an SLI (Service Level Indicators) policy in place.	x				
14	There is an SLO (Service Level Objectives) policy in place.	x				
15	There are periodic reviews of SRE project work and impact with business leaders.		x			
16	There are periodic reviews of SLIs and SLOs with business leaders.	x				
17	TOIL is monitored and work patterns adjusted accordingly.	x				
18	There is a resilience testing program in place	x				
<b>Traditional SRE Score</b>		<b>11</b>				
<b>Overall Score</b>		<b>27</b>				



**“If you fail to plan,  
you are planning  
to fail.”**

# The challenge: cloud sprawl

# SRE Cloud Standards




Security




Cost  
Optimization



Reliability

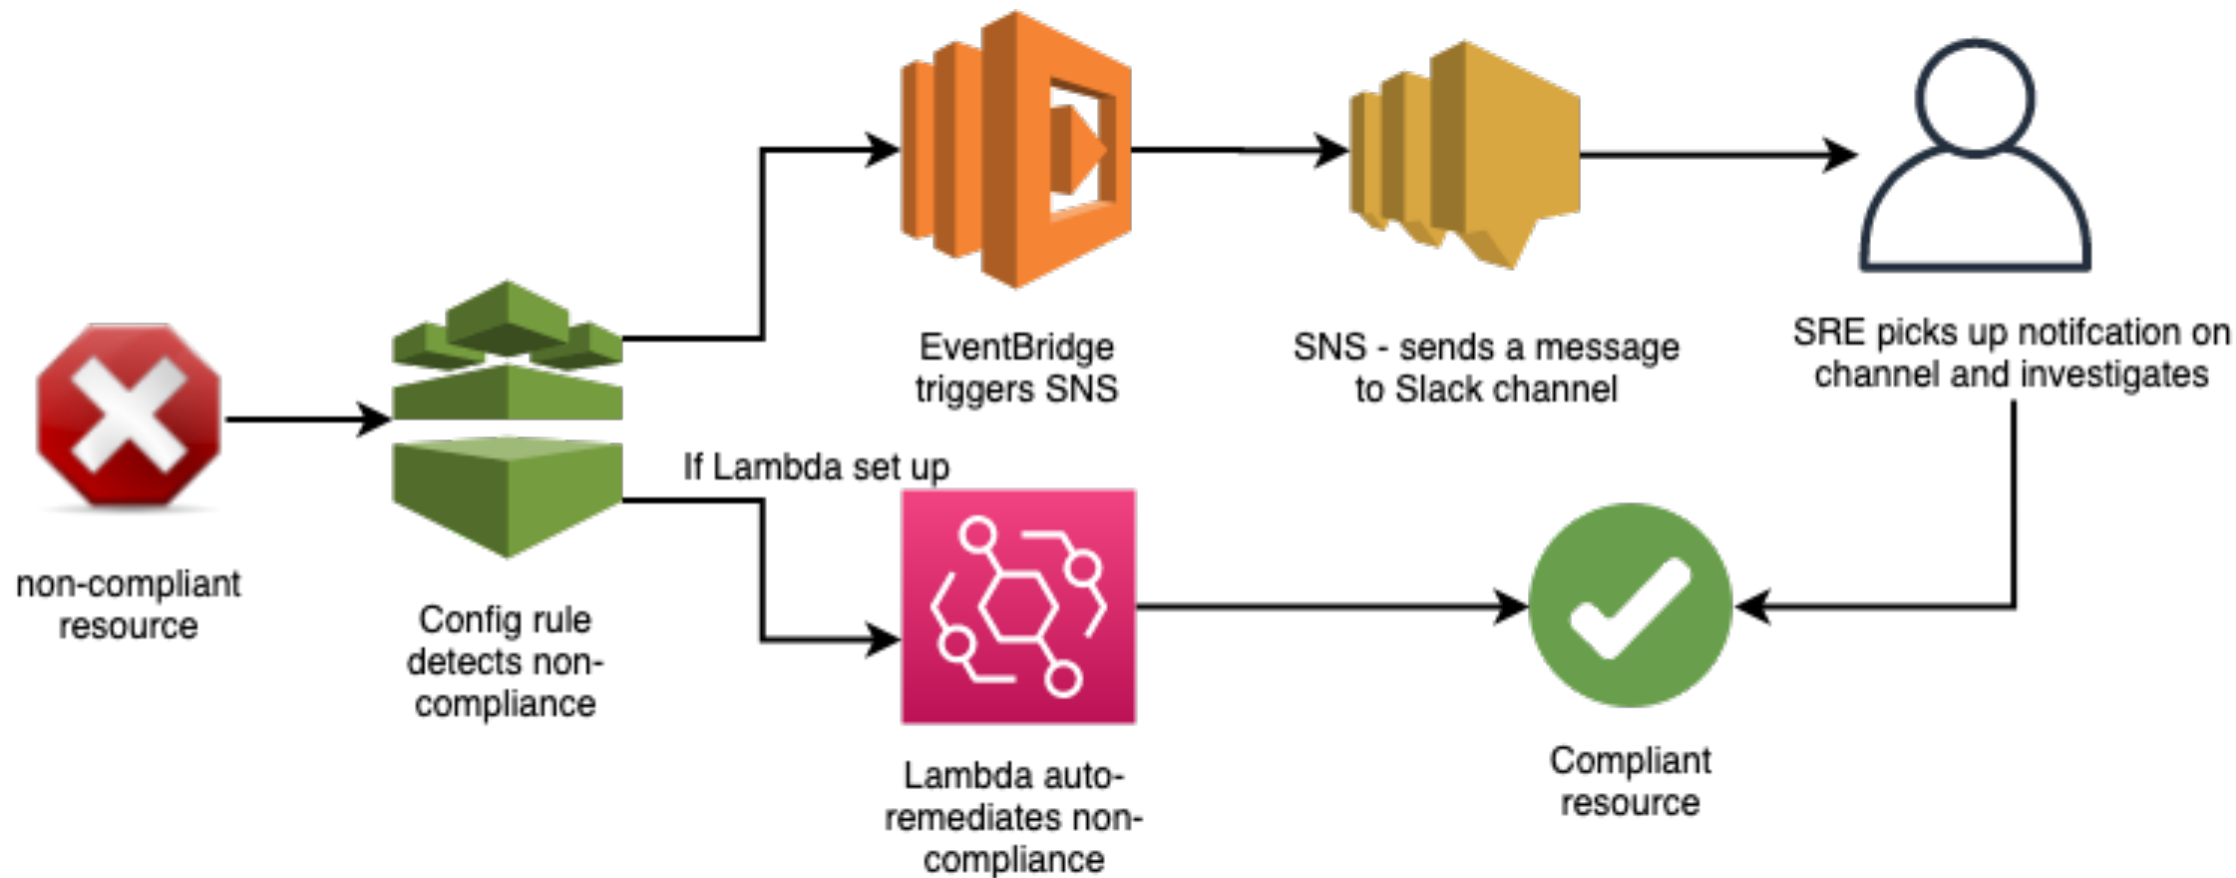


Performance



Operational  
Excellence

- 1) Discovery method
- 2) Resolution method



**The challenge: lack of the "traditional" SRE model**



# The "traditional" SRE activities

## SLI/SLOs

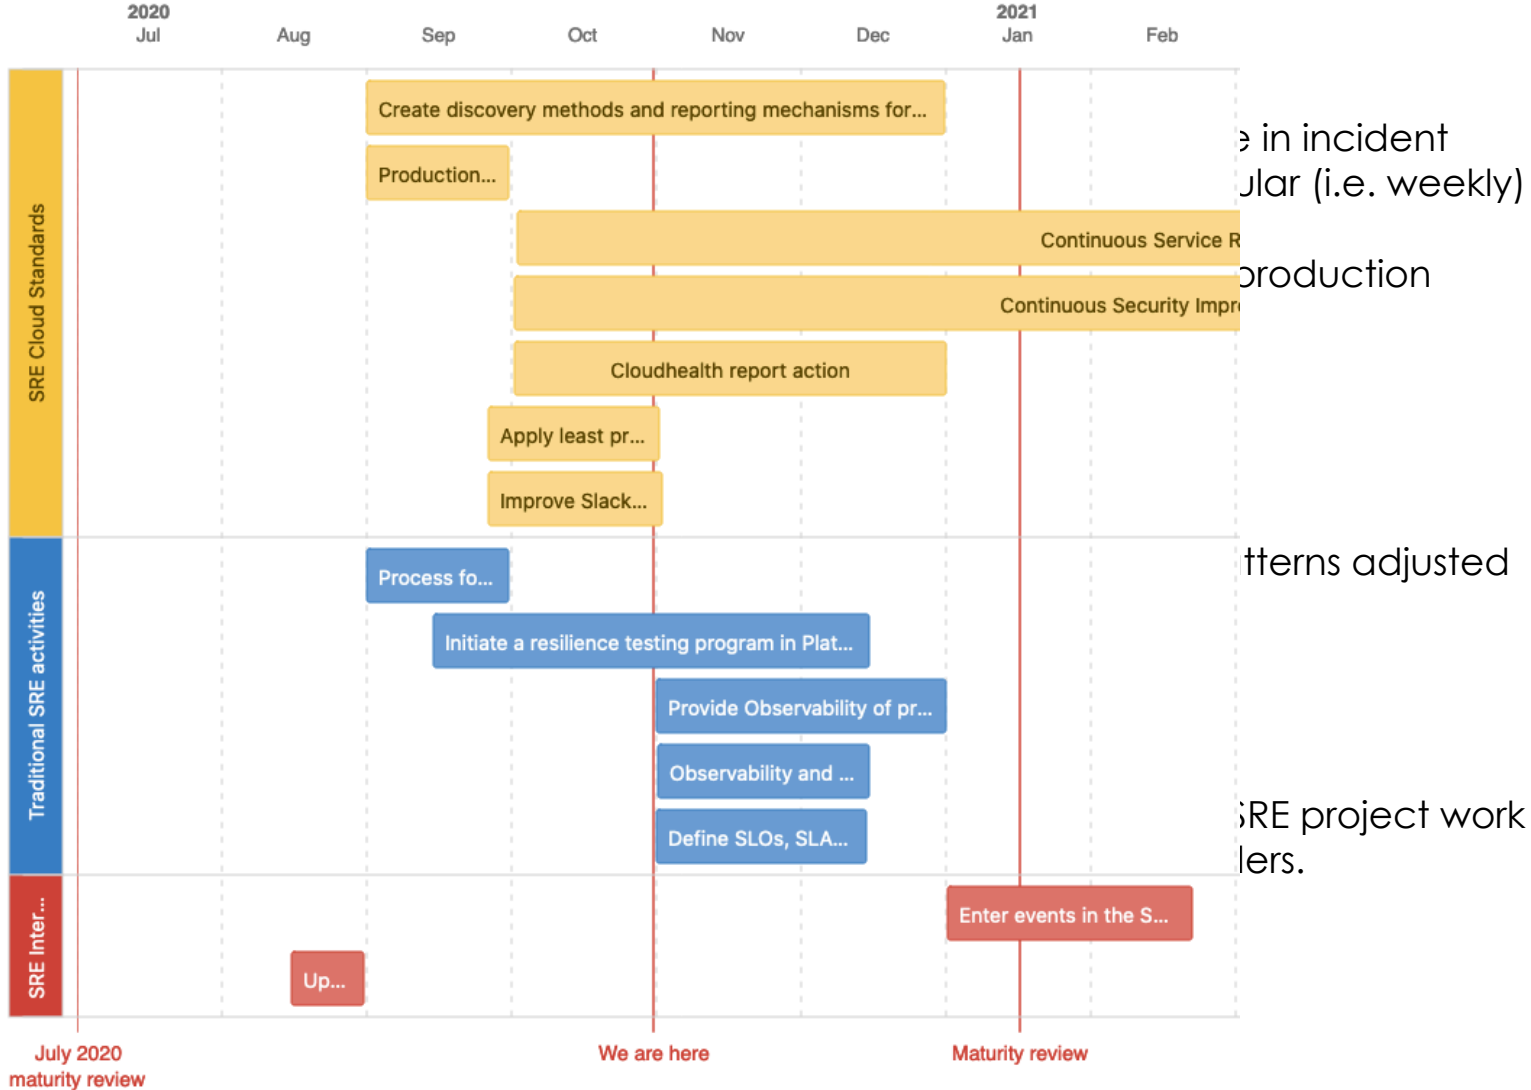
- There is...
- Indicat...
- There is...
- Object...
- There c...
- and SL...

## Blameless

- There is...
- blame...
- produc...

## Disaster R...

- There is...
- progro...



# Chaos Engineering: the art of breaking things purposefully




# KEY TAKEAWAYS


## SRE team wins 🏆

- ★ Saved ££ in our AWS accounts
- ★ Focused on security in our AWS accounts
- ★ Improved collaboration with developers and operations
- ★ Visibility of production to developers
- ★ Resiliency at the forefront
- ★ Positive change!







**Leadership support**



**Plan, plan, plan!**



**What is the problem we're trying to solve?**



**Use SRE principles**

EVERYONE  
IS JUST TRYING  
TO MAKE IT

