Toil

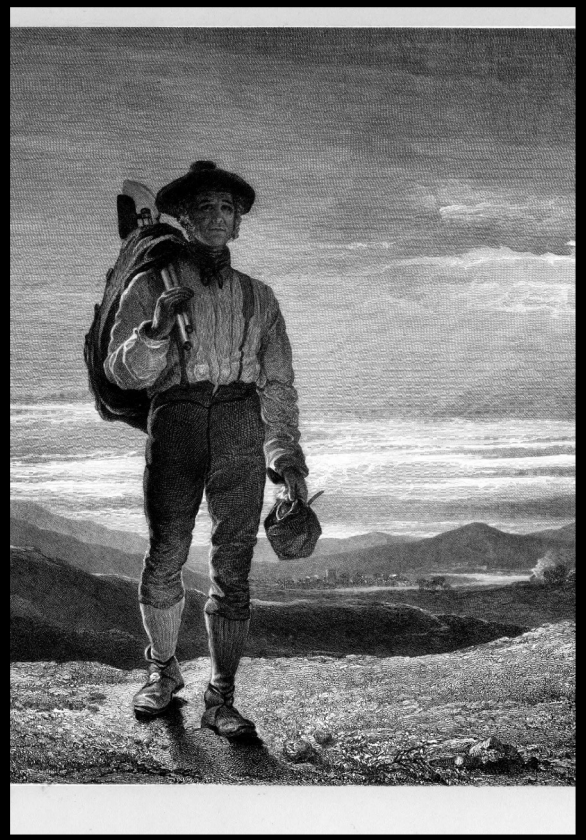

Toil is the kind of work that tends to be manual, repetitive, automatable, tactical, devoid of enduring value, and that scales linearly as a service grows.

SRE teams usually aim for under

50% toil

If an employee is told that 50% of their work has no enduring value, how does this affect their productivity and job satisfaction?

– Byron Miller

The eternal sunshine of the toil-less prod

Sasha Rosenbaum

@DivineOps

Red Hat

Dev

B.Sc. in C.S.

Ops

MBA

Dev + Ops

Cloud Consulting

DevRel

Technical Sales

Sasha Rosenbaum @DivineOps

And you?

Red Hat OpenShift - Kubernetes Experience

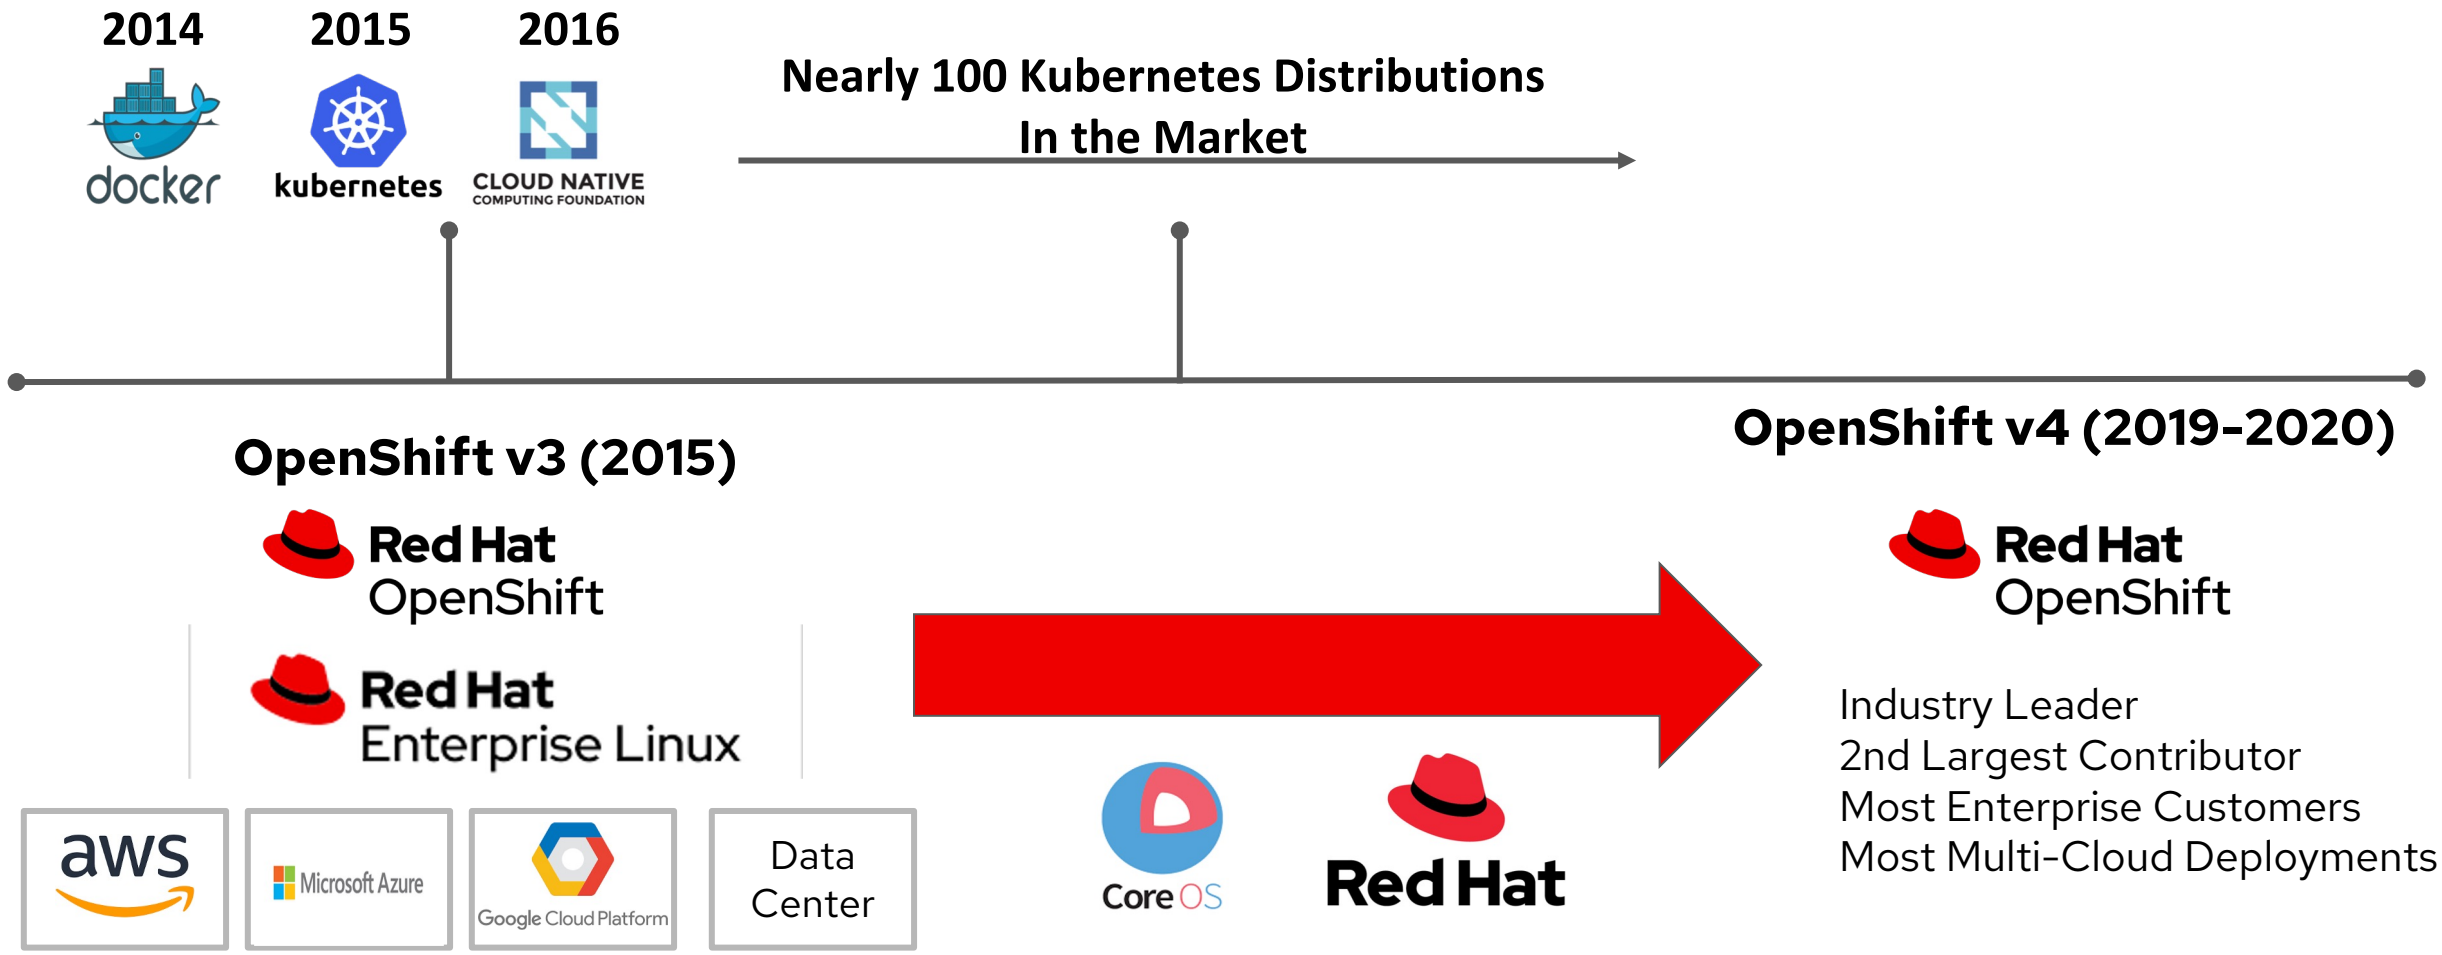

Red Hat

A consistent experience no matter where you run it

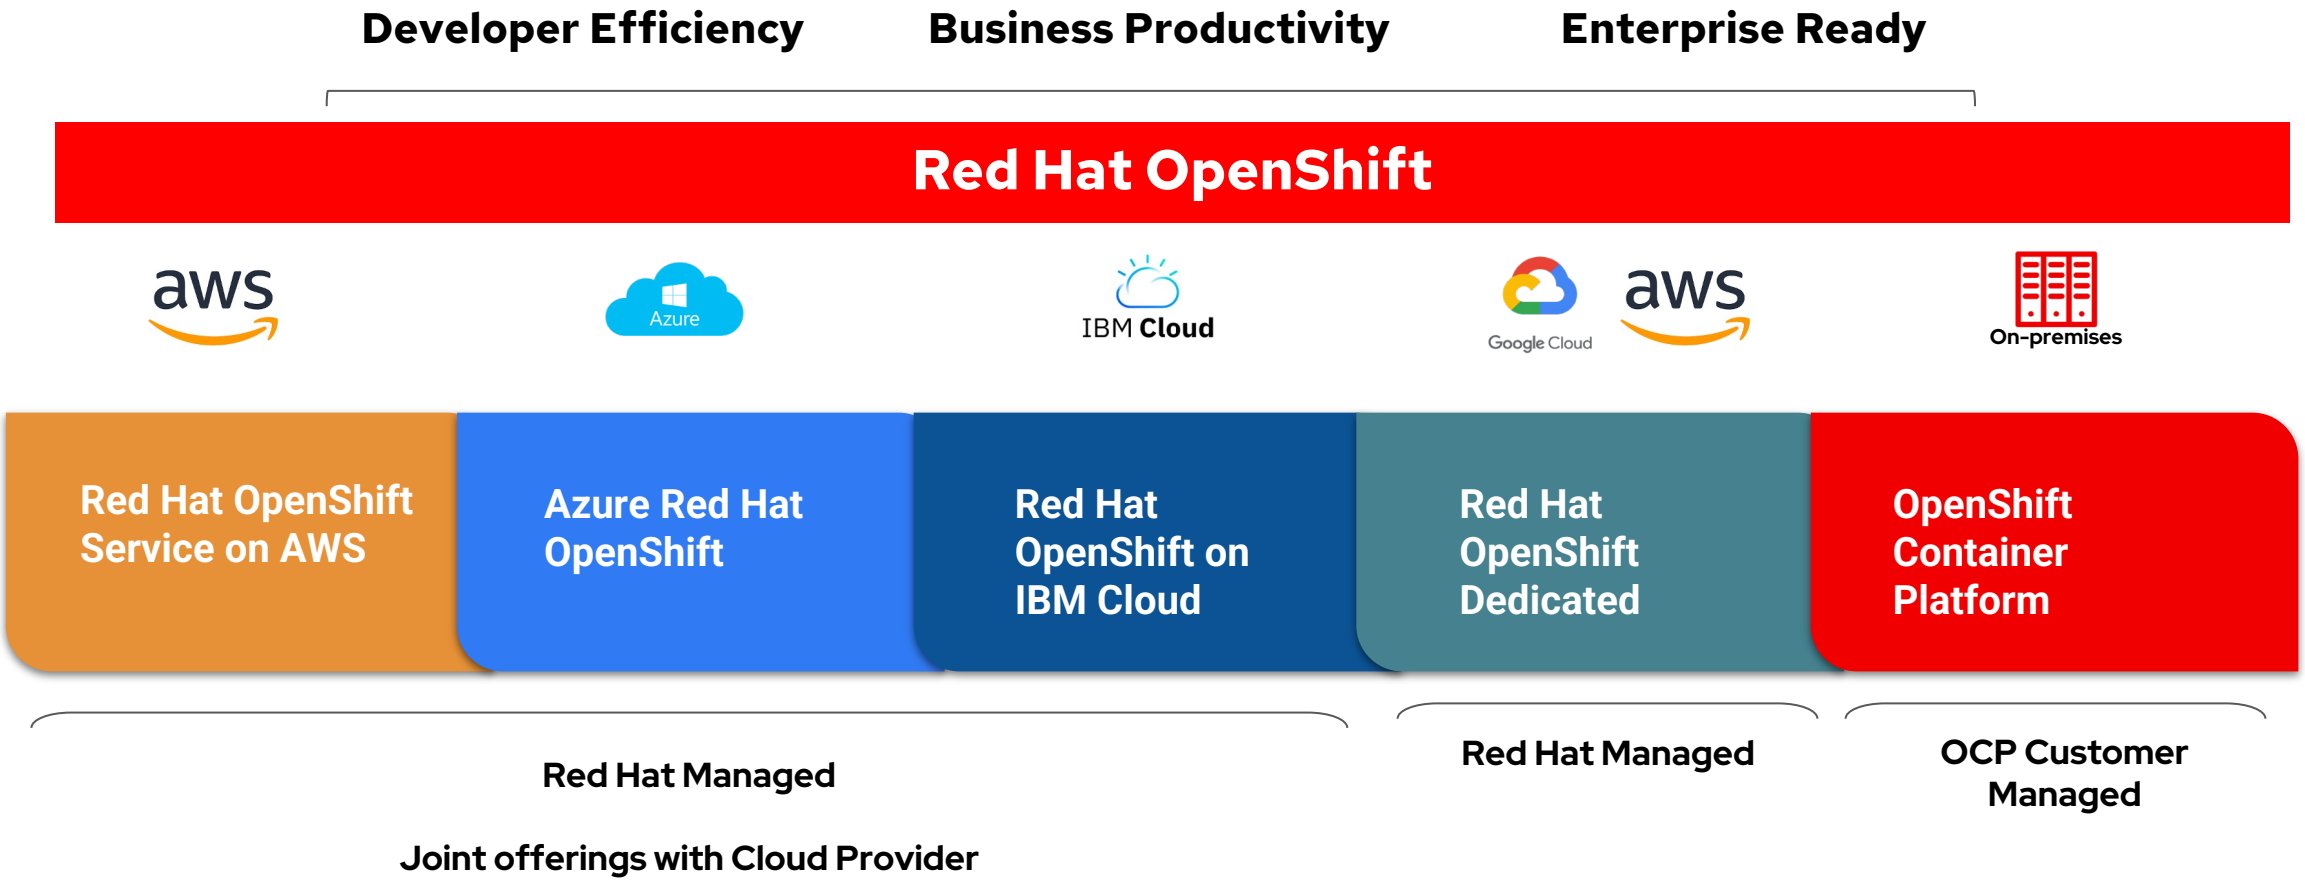

Red Hat

Our Experience

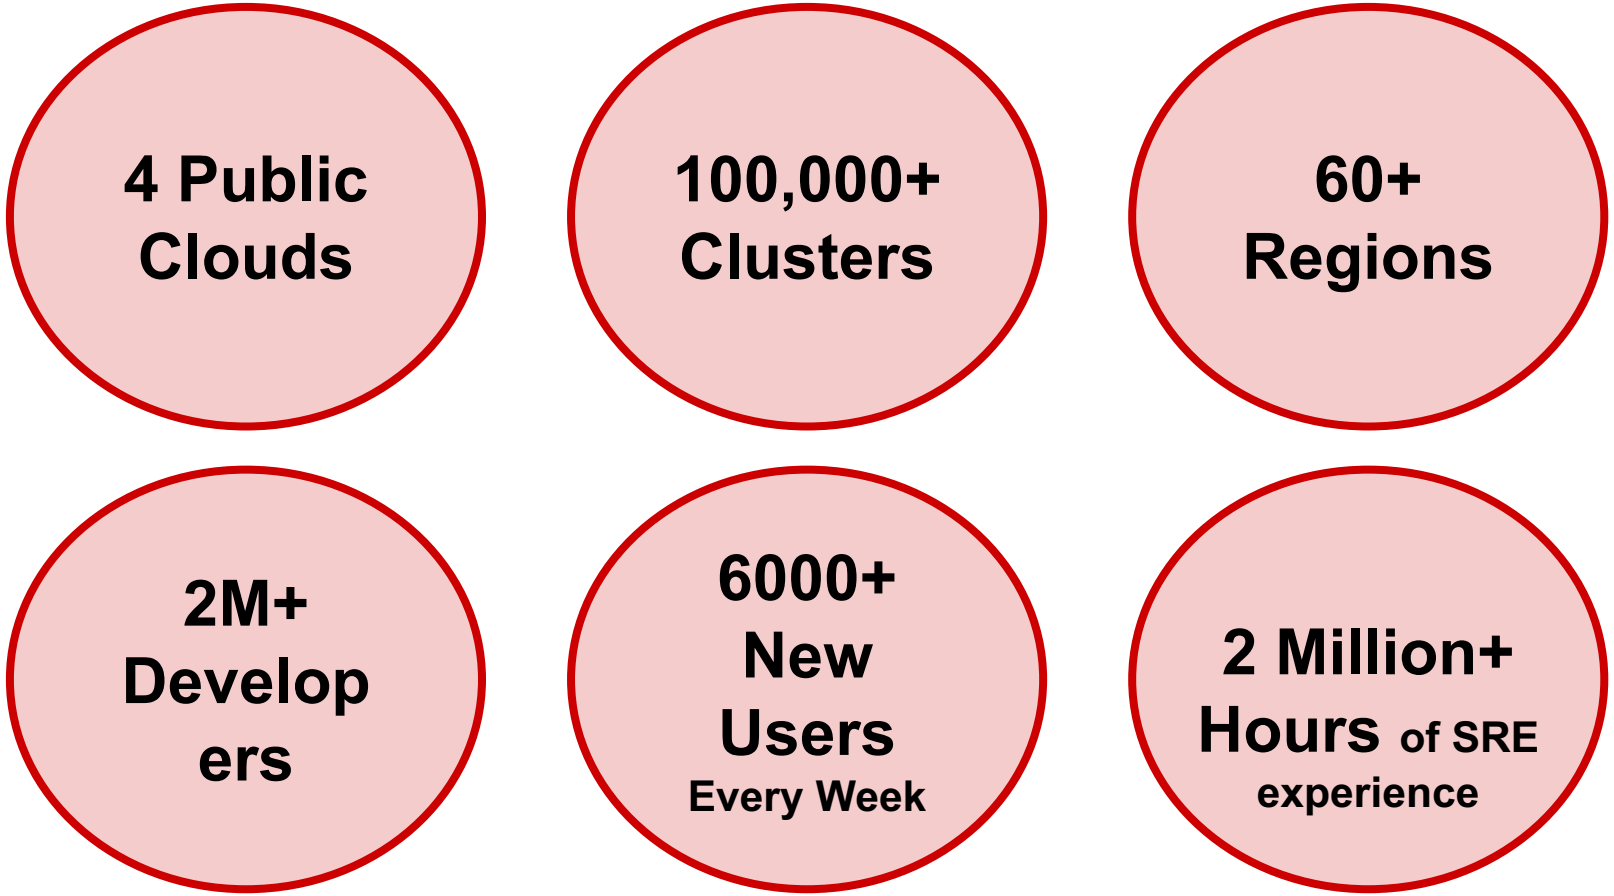

Products to Services

Products and Services

@DivineOps

SRE

What is the most important and innovative thing about SRE discipline?

SRE is about explicit agreements

that align incentives

@DivineOps

SLA, SLI, SLO

SLA

=

Financially-backed availability

Monthly Uptime Percentage	Service Credit Percentage
Less than 99.95% but equal to or greater than 99.0%	10%
Less than 99.0% but equal to or greater than 95.0%	25%
Less than 95.0%	100%

Monthly downtime > 1.5 days means 100% refund

SLAs are about

aligning incentives

between Vendor & Customer

99% => 99.99%

- SLA usually includes a single metric
- For financial and reputational reasons, companies prefer to under promise and overdeliver

SLO

=

Targeted reliability

Service Level Objectives: What we care about?

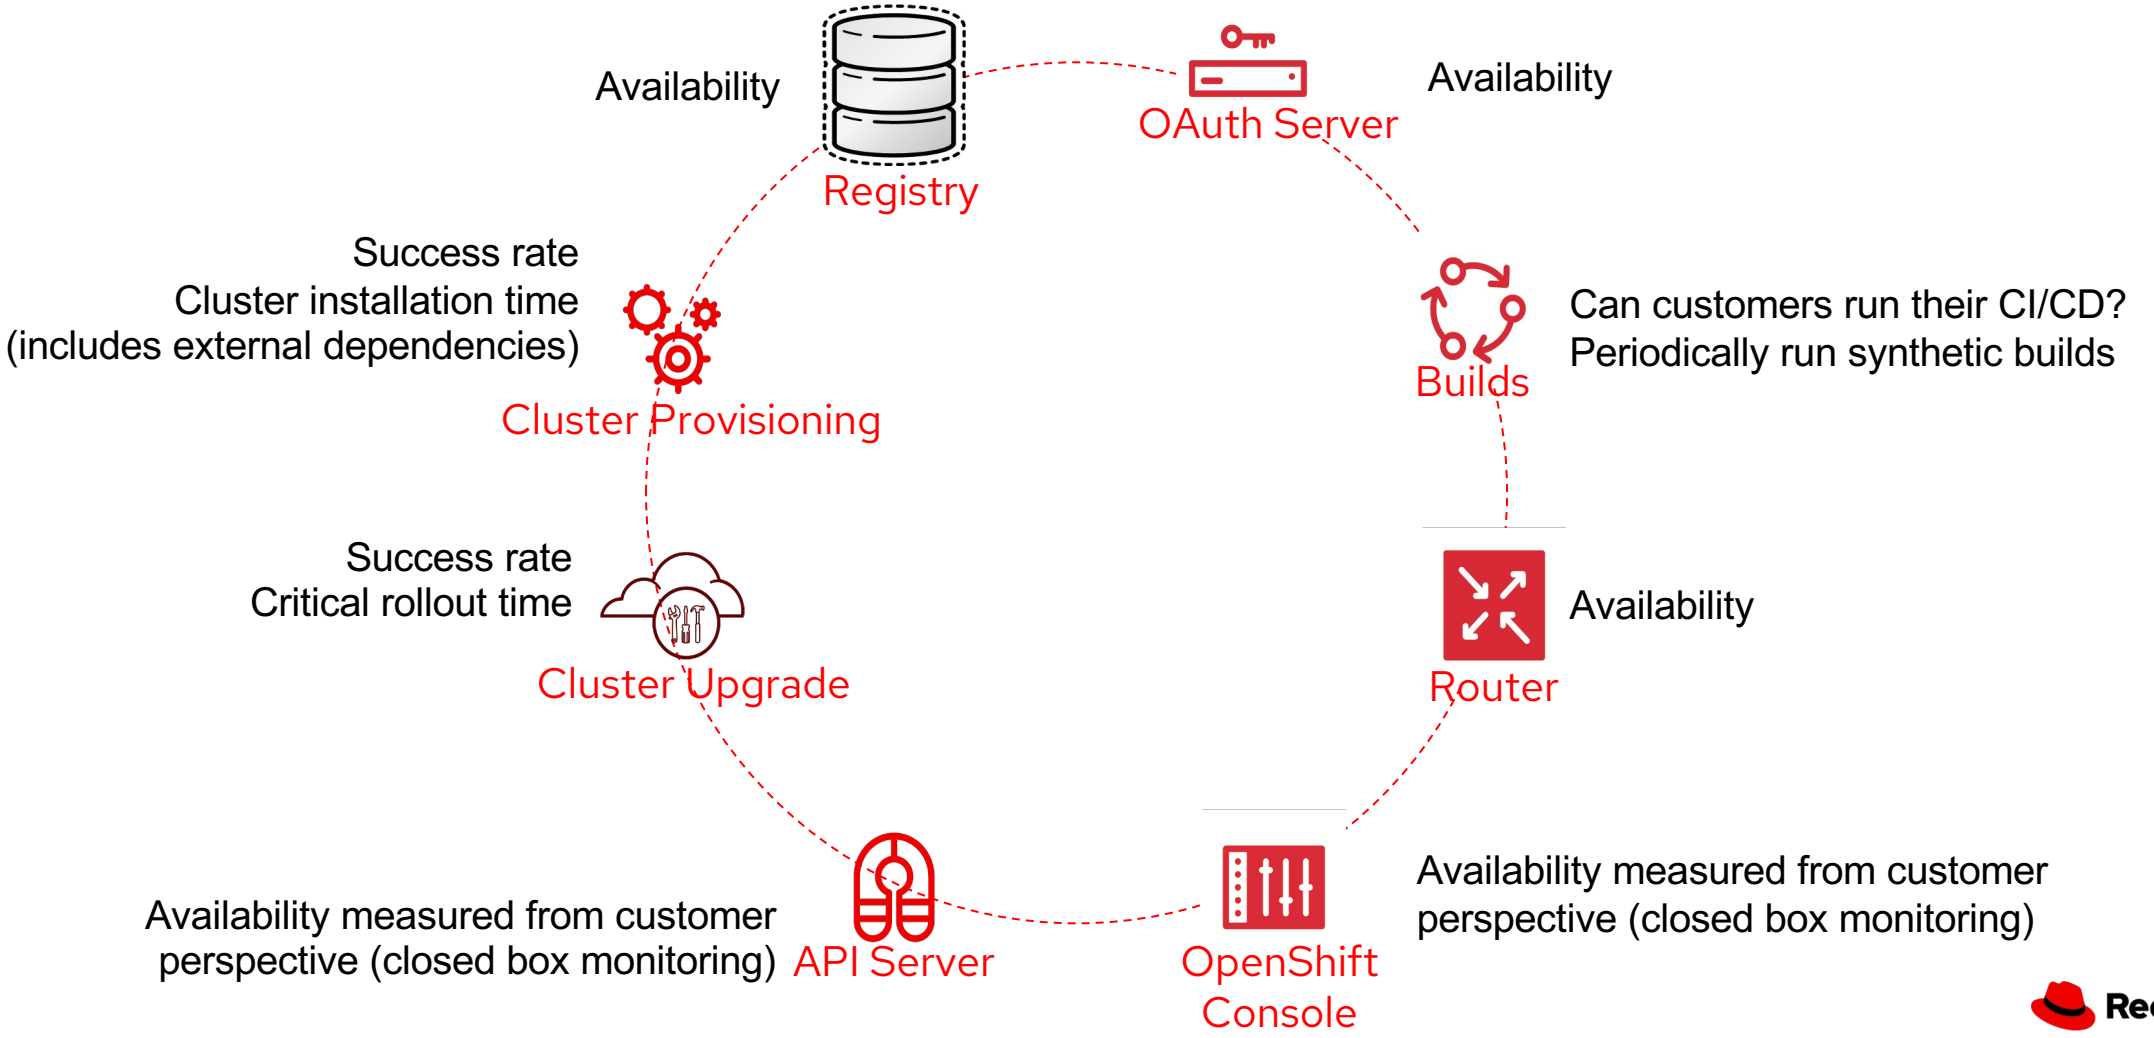

Red Hat

SLI

=

Actual reliability

@DivineOps

Monitoring

Without monitoring, you have no way to tell whether your service even works!

@DivineOps

Good Monitoring

Without good monitoring, you don’t know that the service does what users expect it to do!

@DivineOps

Signal to noise ratio

Early on, one of the major monitoring problems we had is alerts on customer clusters that were intentionally taken offline

Without good monitoring, your SRE is potentially overloaded with unwarranted emergencies and blindsided by real incidents

Periodically incidents may be caught by internal users, rather than the monitoring system

We aim to implement monitoring improvements that will catch future problems of the same kind

@DivineOps

SLO

@DivineOps

SLO

=

Business-approved reliability

100% reliability is...

- Unattainable
- Unnecessary
- Extremely expensive

The five nines

99.999%

5.26 mins / year

Will your users even notice?

The ISP background error rate is

0.01% – 1%

SLOs are about

explicitly aligning incentives

between Business & Engineering

Error Budgets

Acceptable level of unreliability

Error budget = 1 - SLO

EB = 1 – 99.99% = 0.01% ≃ 13 mins / quarter

Error budgets are about aligning incentives between Dev & Ops

If developers are measured on the same SLO, then when the error budget is drained developers shift focus from delivering new features to improving reliability

So, we've written things down

Are we there yet?

The future is already here.

It’s just not evenly distributed

~ William Gibson

Awkward Segue

What we all got wrong

@DivineOps

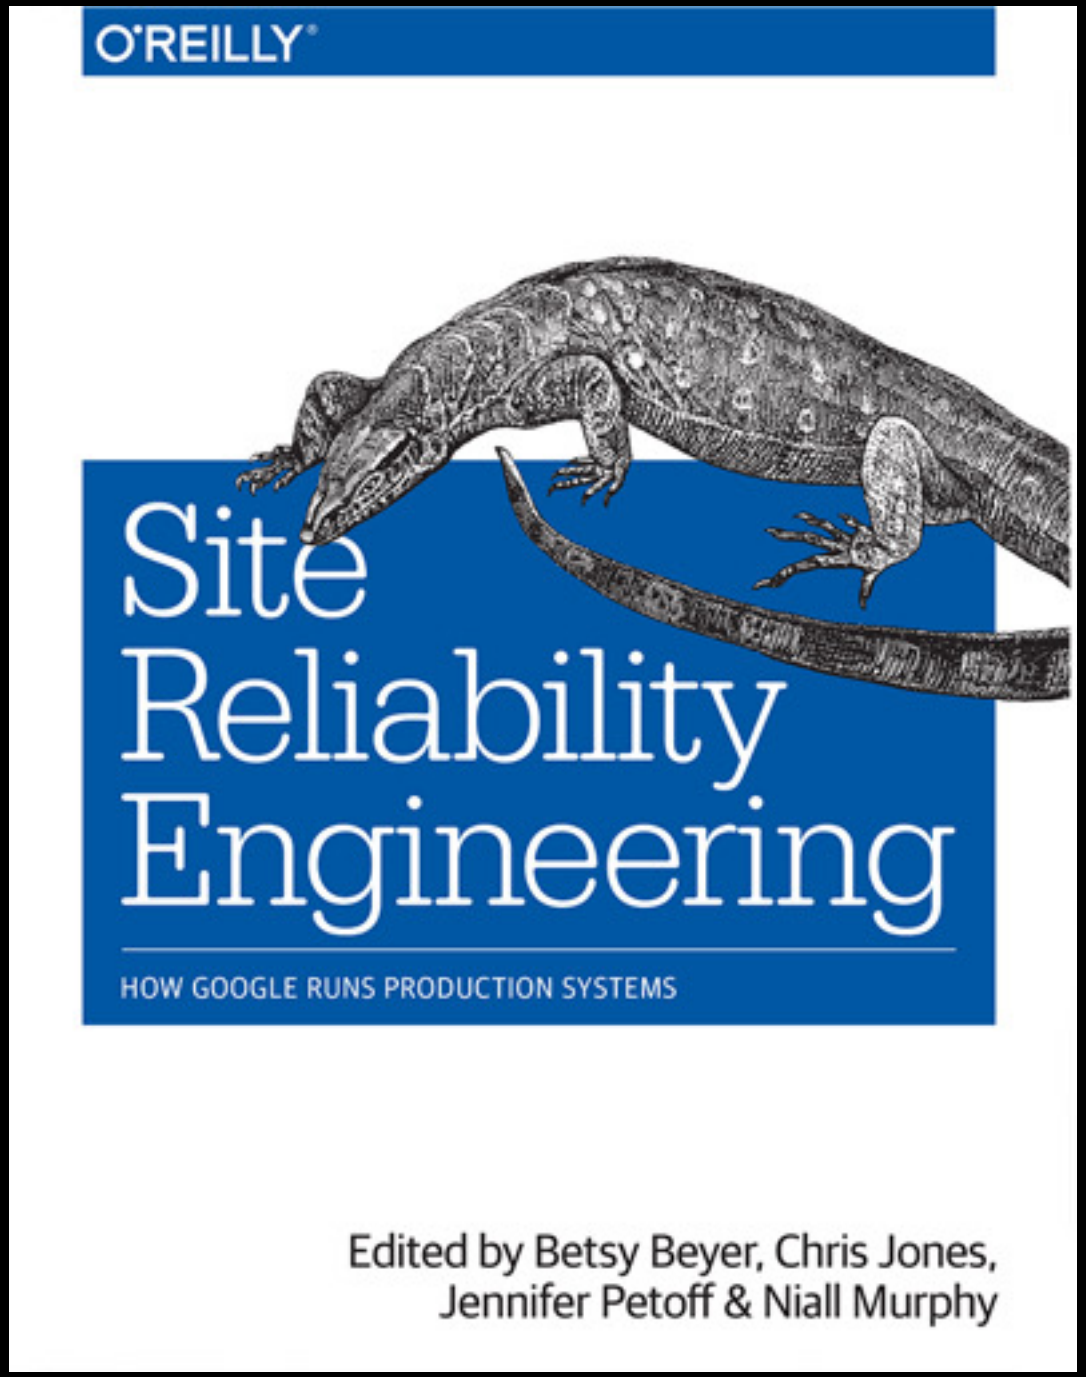
O'REILLY
Site Reliability Engineering
HOW GOOGLE RUNS PRODUCTION SYSTEMS
Edited by Betsy Beyer, Chris Jones, Jennifer Petoff & Niall Murphy

"SRE is what happens when you ask a software engineer to design an operations team."

- Google SRE book, 2017

Is it though?

@DivineOps

DevOps

DevOpsDays Ghent 2009

@DivineOps

Automate ourselves out of a job!

So why didn’t we do it?

Effective automation requires consistent APIs

OS-level APIs

2000

27% of server market

- File-based OS
- Maintains configuration in files
- Every device is a file

41% of server market

- Executable-based OS
- Maintains configuration in registry
- Every device has a different driver mechanism

PowerShell

(Windows) configuration management framework, CLI, and scripting language

GA: 2006

Jeffrey Snover

"Which part of F@!$ng Windows is confusing you?"

"Admins don't want CLIs"

DEVOPSDAYS SEATTLE

Every wave of automation

Enables the next wave of automation

Infrastructure-level APIs

Borg cluster manager

"Central to its success - and its conception - was the notion of turning cluster management into an entity for which API calls could be issued"

- Google SRE book, 2017

Amazon Web Services: 2002

Amazon Cloud Computing: 2006

Azure Cloud Services: 2008

Infrastructure as code

2005

2009

2012

We did **NOT** suddenly get the idea of infrastructure & platform automation

We did **NOT** suddenly get the idea of infrastructure & platform automation

We gradually built the tools required to make it happen

Why does this matter?

If we get the origin story wrong, we end up working to solve the wrong problem!

Corollary 1:

Hiring developers to do operations work
≠ effective SRE

We have seen success from hiring the skillsets across the entire landscape, hiring well-rounded folks with understanding of Ops and Dev concepts, as opposed to just Dev experience

Corollary 2:

The desire to automate the infrastructure & platform operations is insufficient

Corollary 2:

The desire to automate the infrastructure & platform operations is insufficient

We need consistent APIs and reliable monitoring to unblock the automation

Early on, we had to move the Cloud Services Build system from on-prem to the Cloud, because it was not meeting our reliability and agility targets

What we all got wrong

Toil

Toil

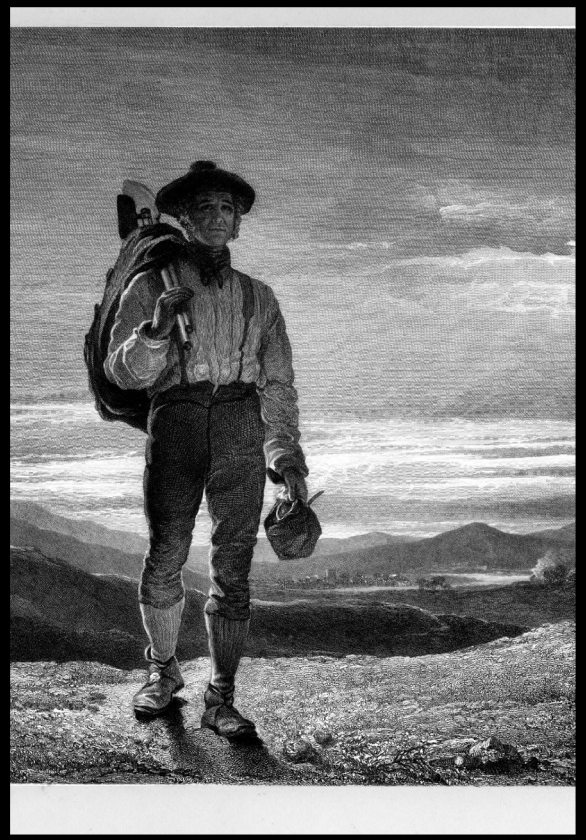

Toil is the kind of work that tends to be manual, repetitive, automatable, tactical, devoid of enduring value, and that scales linearly as a service grows.

SRE teams usually aim for under

50% toil

So, are we striving for a human-less system?

Second law of thermodynamics

With time, the net entropy (degree of disorder) of any isolated system will increase

Entropy always wins

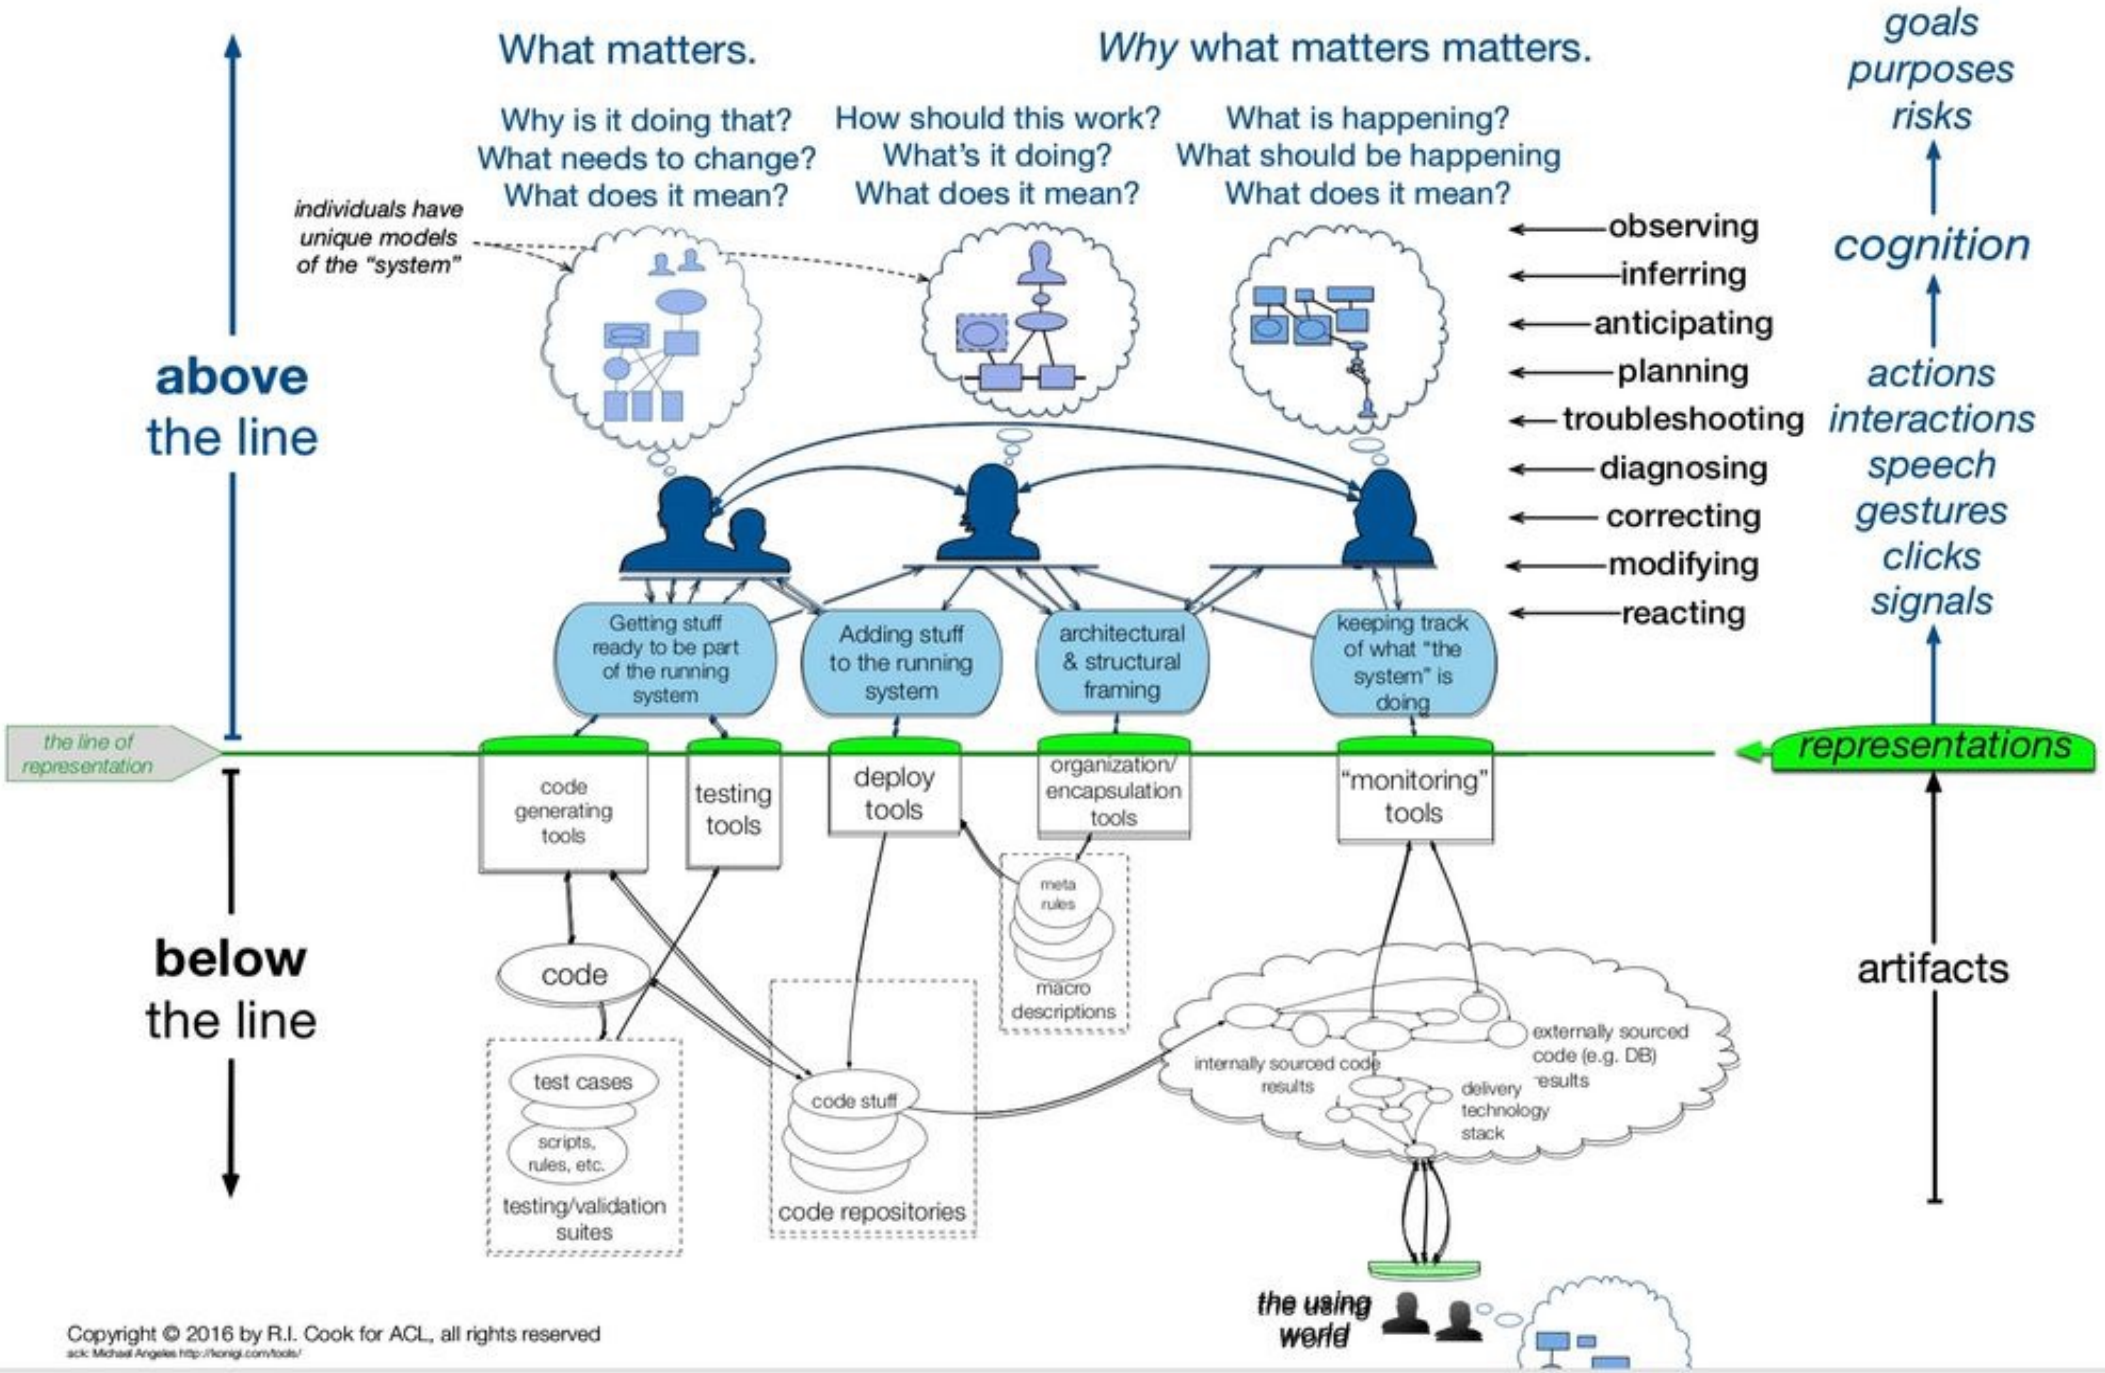

People working above the line of representation continuously build and refresh their models of what lies below the line. That activity is critical to the resilience of Internet-facing systems and the principal source of adaptive capacity.

– Dr. Richard Cook

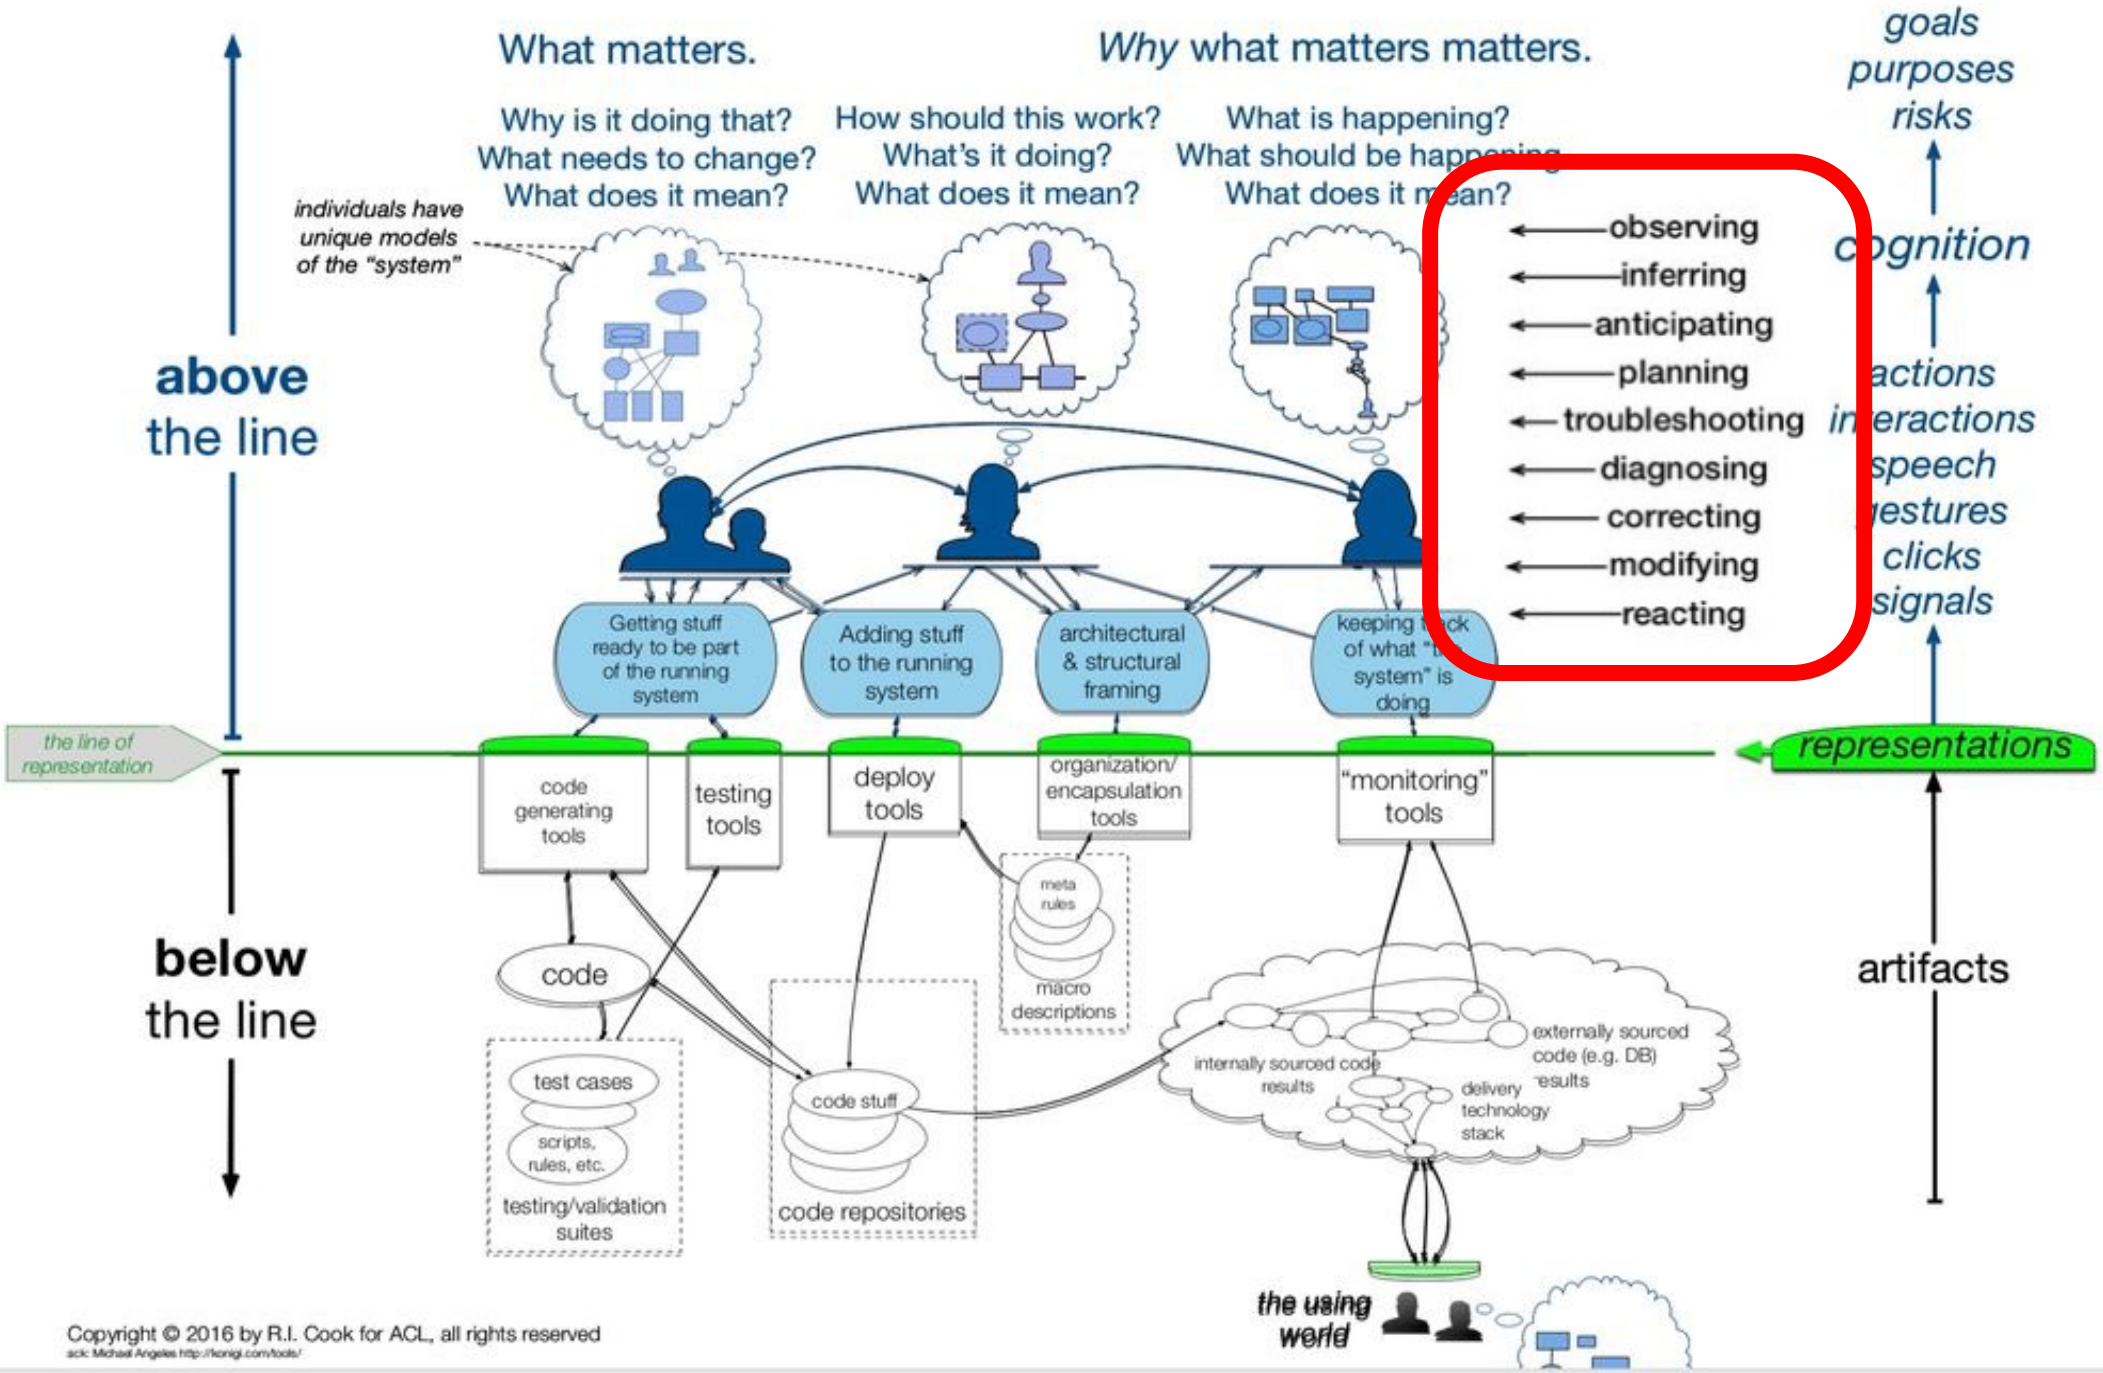

Resilience

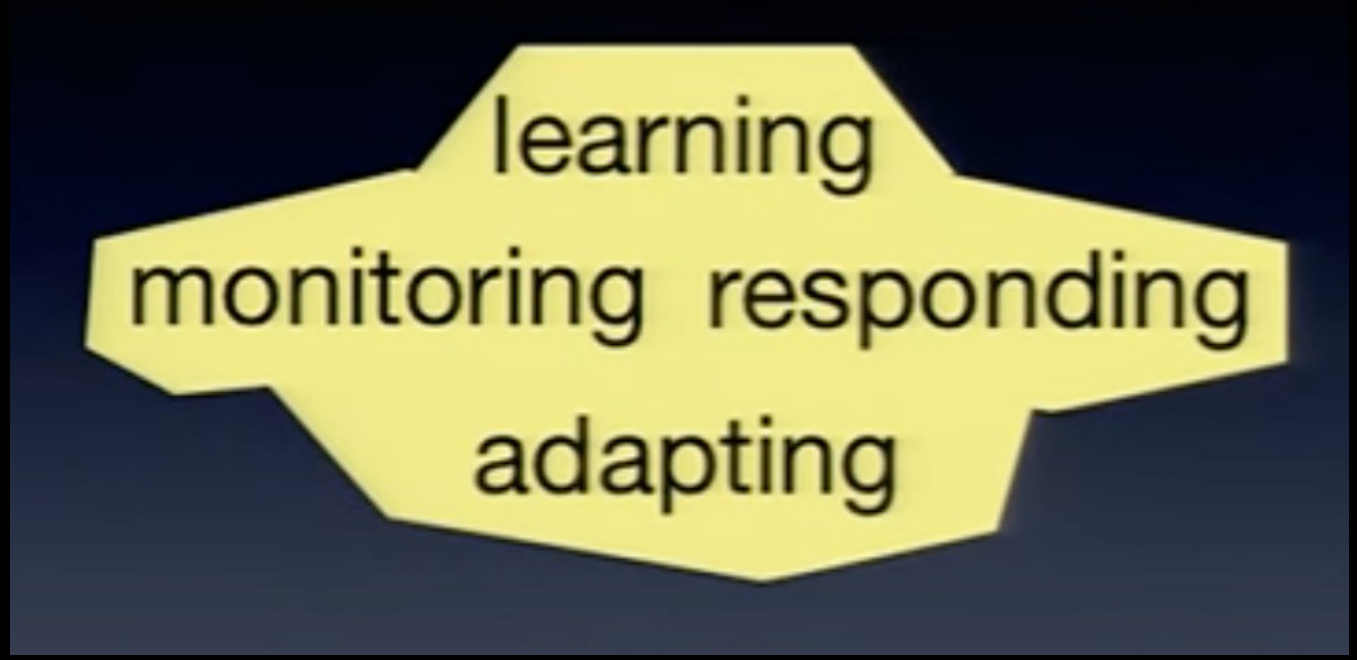

Velocity 2012: Richard Cook,
"How Complex Systems Fail"

What we call toil

is a major part of resilience

and adaptive capacity

Perhaps we need a better way

to look at toil

SRE folks worry that if they spend significant parts of their day focusing on toil, it will negatively affect their bonuses, chances of promotions etc.

If an employee is told that 50% of their work has no enduring value, how does this affect their productivity and job satisfaction?

– Byron Miller

SRE Work Allocation

Work Allocations:

SRE P and SRE O

Traditional IT

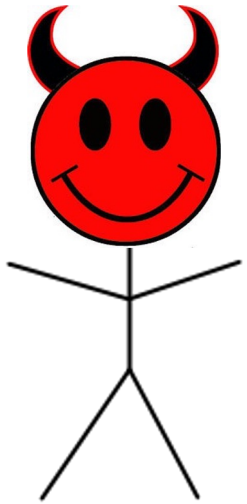

dev

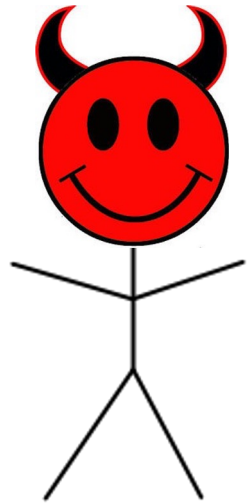

ops

wall of confusion

Work Allocations:

On-call once a month

Red Hat

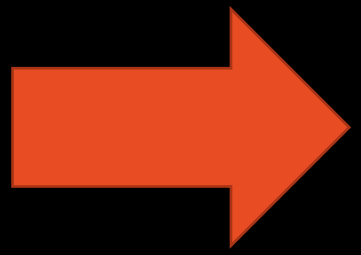

SRE teams asked management for more on-call, as they were losing their "Ops muscle"

Work Allocations:

Rotate engineers working on toil-reduction tasks

Red Hat

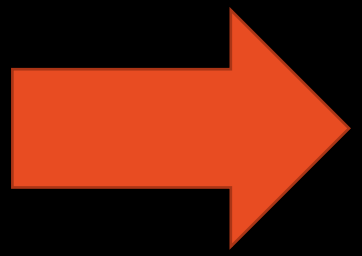

Lack of continuity severely impacted team's ability to deliver

Work Allocations:

The search for the perfect system is still in progress!

Where do we go from here?

Let's look at some of the insights from the talk:

Effective automation requires consistent APIs

Cloud

Cloud provides an industry standard for consistent infrastructure-level APIs

@DivineOps

Are you in the datacenter management business?

Kubernetes

85% of global IT leaders agree that Kubernetes is key to cloud-native application strategies

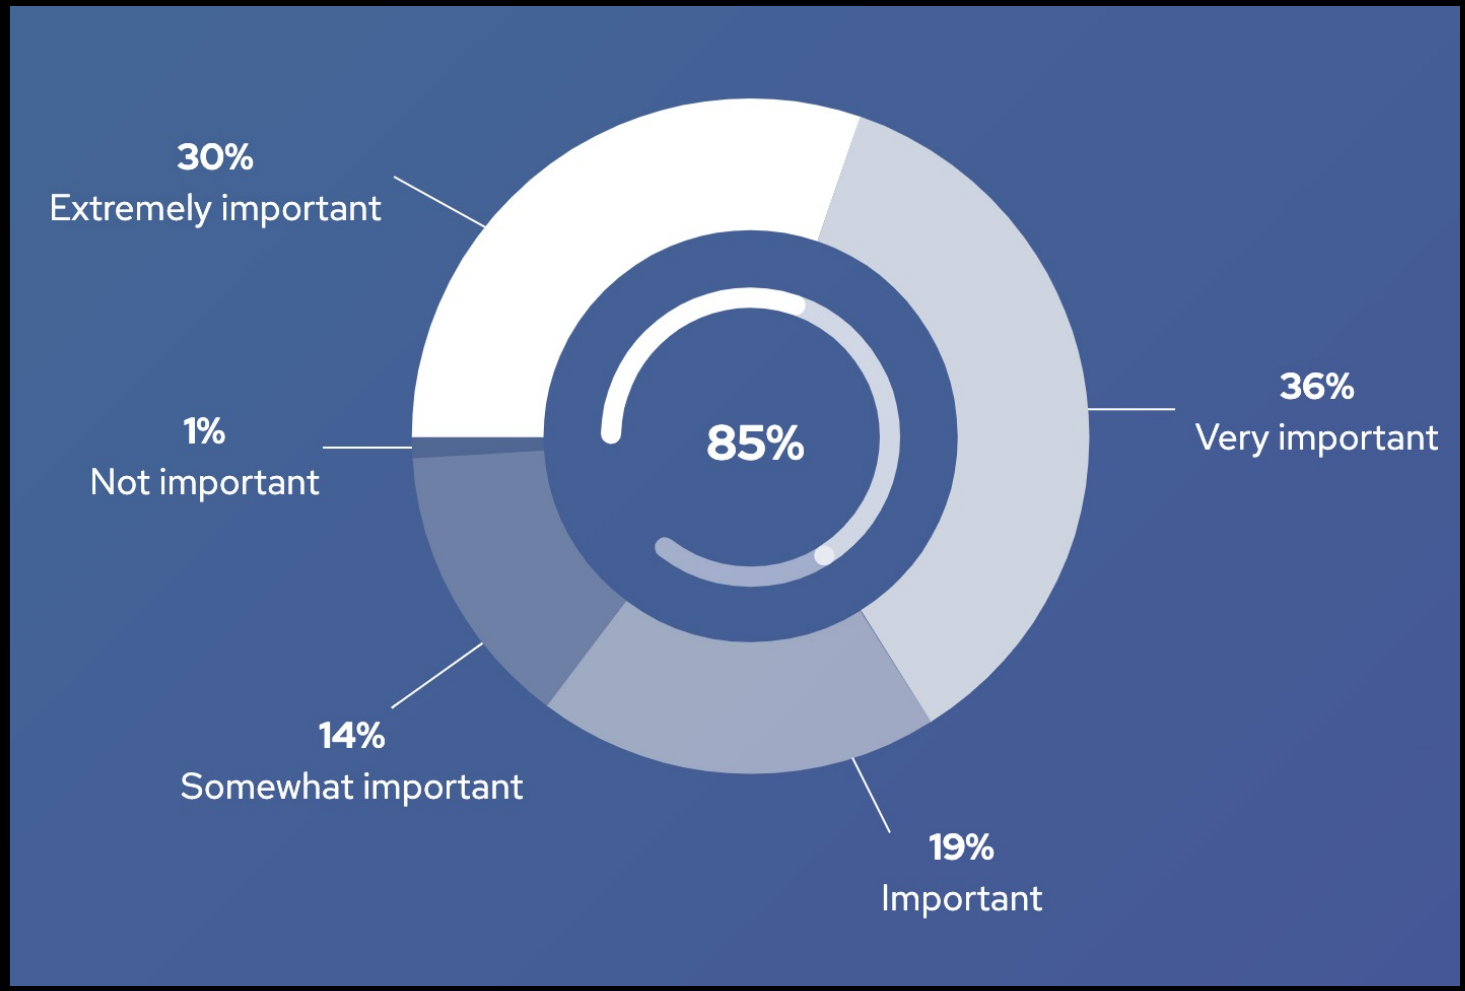

Source: Red Hat State of Open Source Report 2021

Kubernetes could provide the industry standard for consistent platform-level APIs

Red Hat

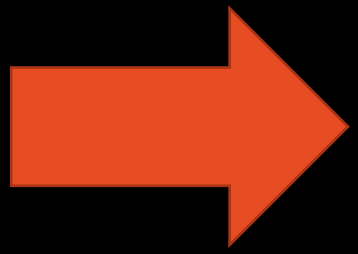

If building PaaS isn't your company's core business

Red Hat

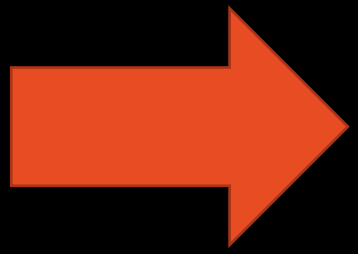

If building PaaS isn't your company's core business

Allow your provider to toil for you

A consistent experience no matter where you run it

Developer Efficiency **Business Productivity** **Enterprise Ready**

Red Hat OpenShift

Red Hat OpenShift Service on AWS	Azure Red Hat OpenShift	Red Hat OpenShift on IBM Cloud	Red Hat OpenShift Dedicated	OpenShift Container Platform
Joint offerings with Cloud Provider			Red Hat Managed	OCP Customer Managed

Offered as a Native Console offering on equal parity with cloud provider Kubernetes service

or

OCP Customer Managed

Red Hat

@DivineOps

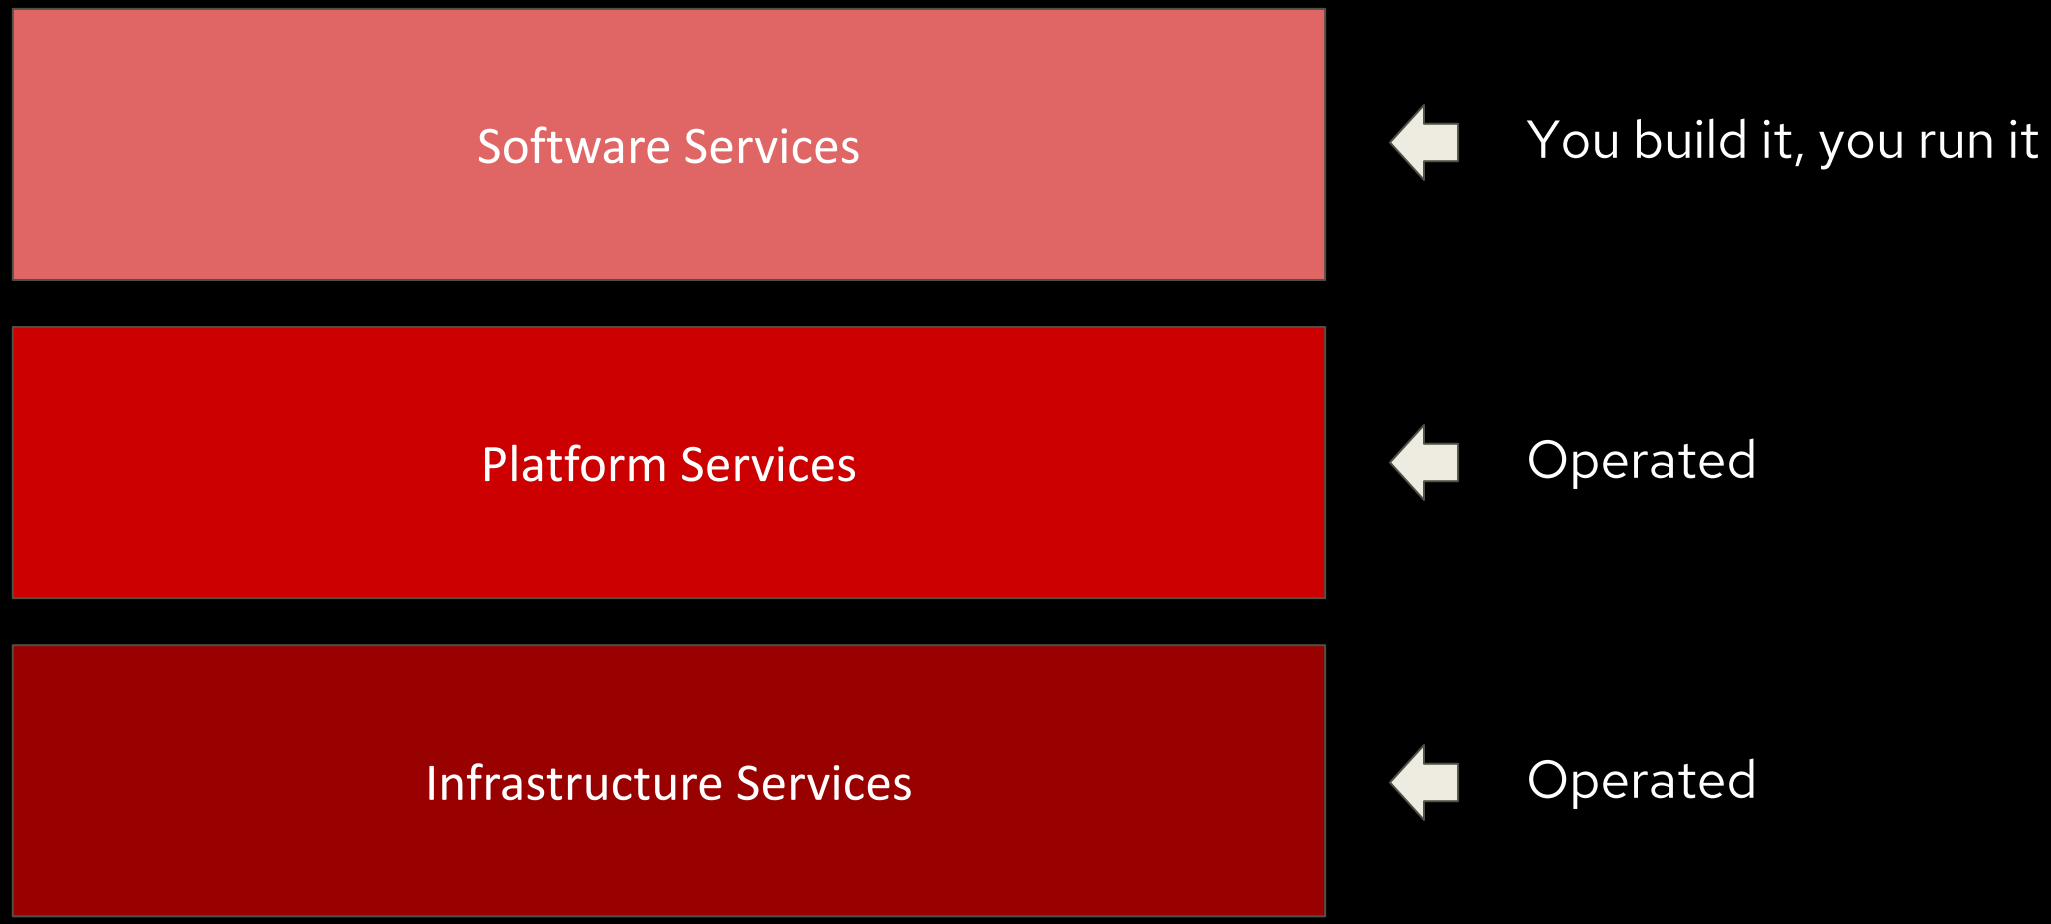
Software Services
You build it, you run it
Platform Services
Operated
Infrastructure Services
Operated

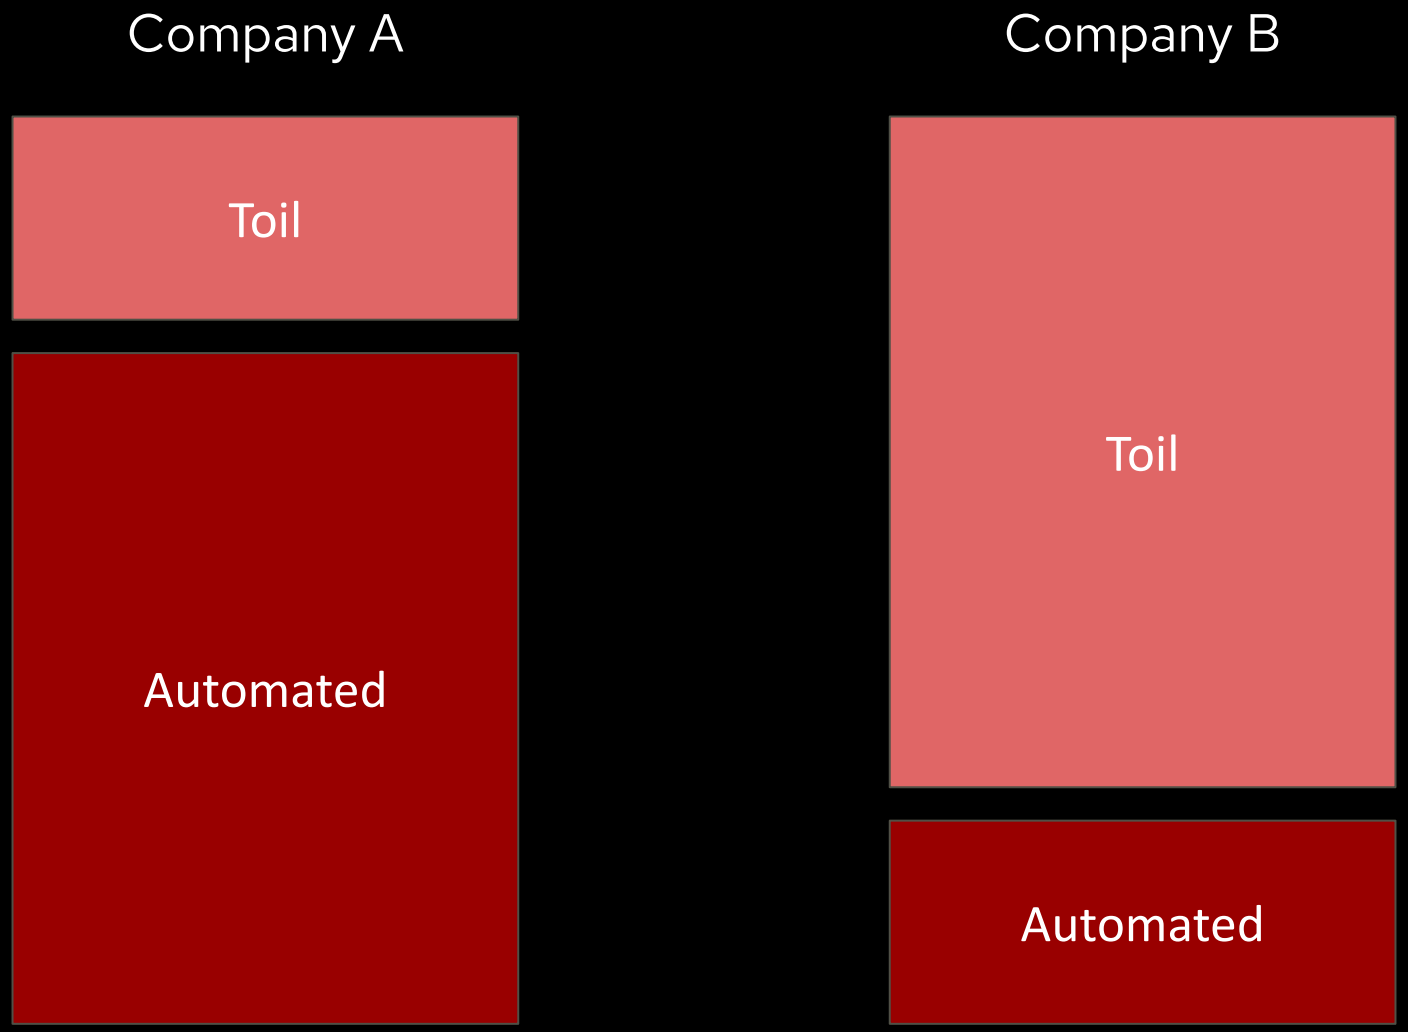
Company A
Toil
Automated
Company B
Toil
Automated

@DivineOps

Get your skills

above the API!

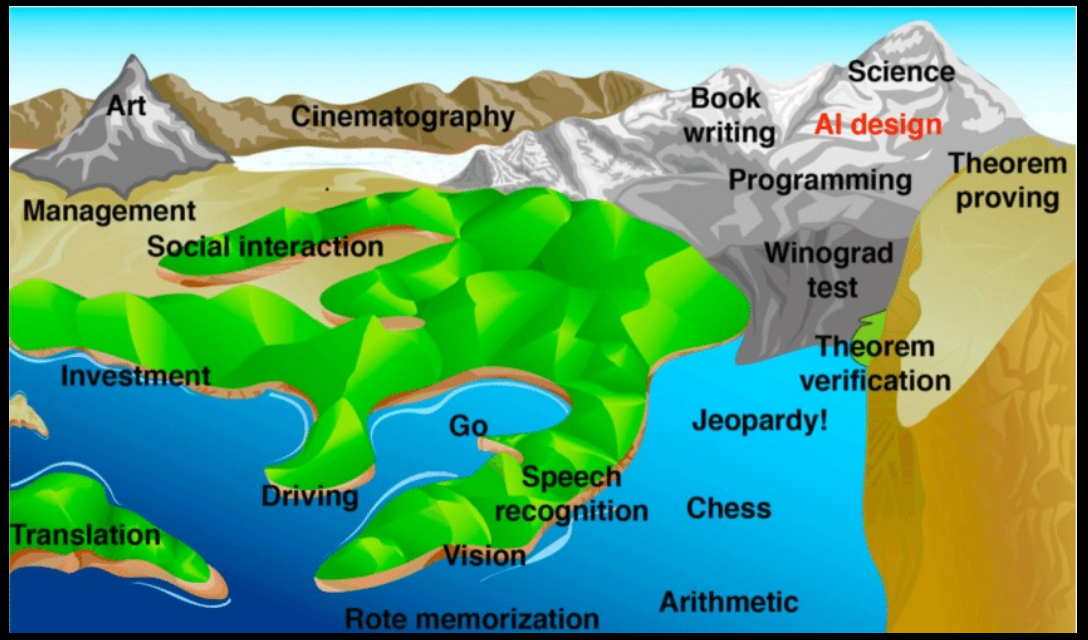

Image Source: Hans Moravec's illustration of the rising tide of the AI capacity. From Max Tegmark (2017)

Red Hat

If building PaaS **IS** your company's core business

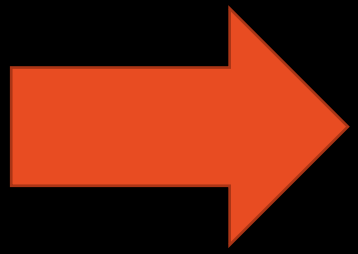

Remember that SRE is about

explicit agreements

that align incentives

Red Hat

Focus your toil

where your business value is

Last, but not least

Ideas are open source

Red Hat

Operate First

A concept of incorporating operational experience into software development

@DivineOps

Red Hat

OPERATE FIRST

https://operate-first.cloud/operations/sre

Thank you!

@DivineOps

Slides:
speaking.sasharosenbaum.com