

# Red Hat Open Tour 2022



---

intel®

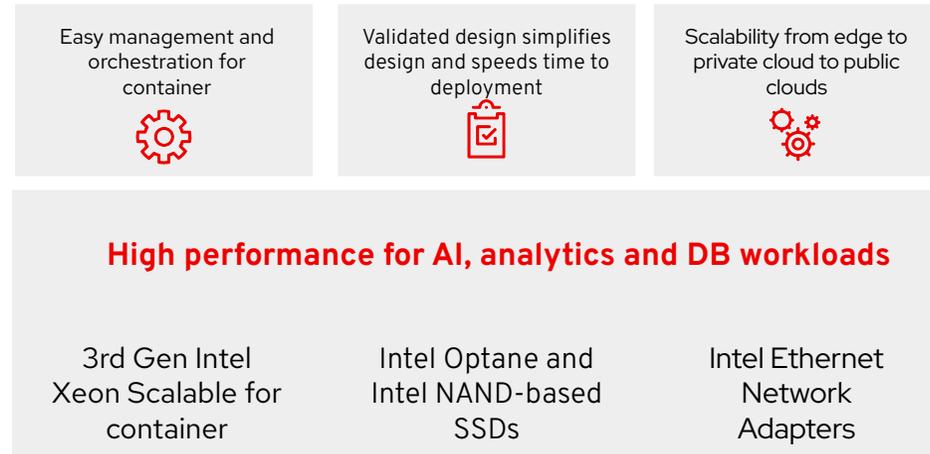
# Red Hat OpenShift Reference Architecture

## Joint Red Hat and Intel OpenShift Reference Architecture

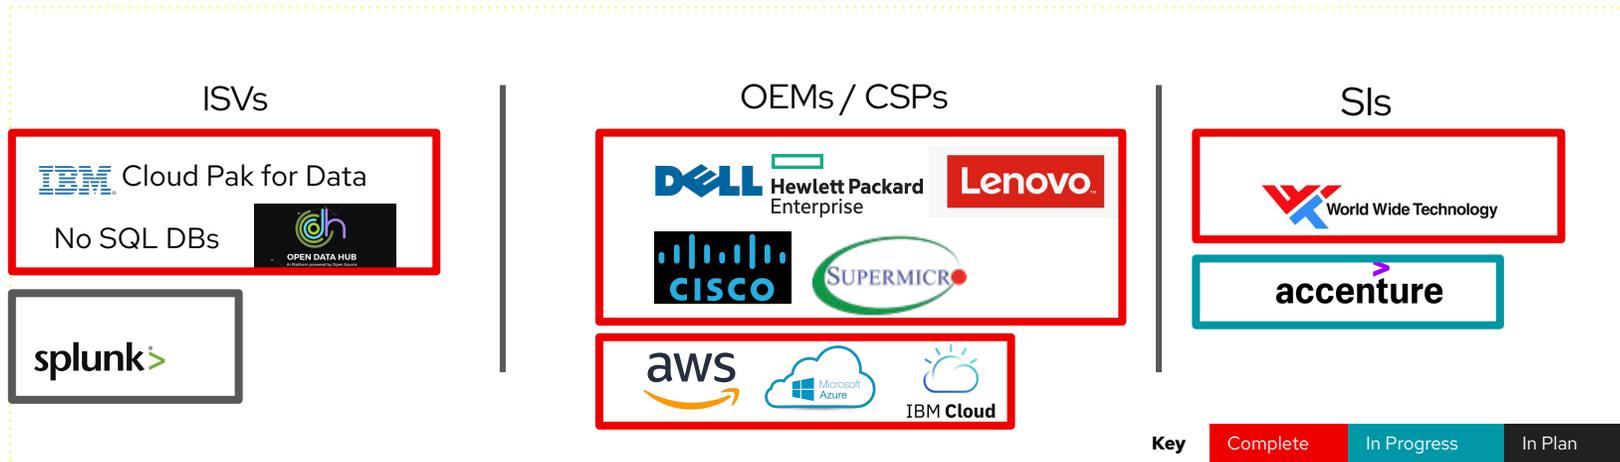
### Solution overview

**Summary:** The RA enables deployment of performant and low-latency container-based workloads onto different footprints, such as bare metal, virtual, private cloud, public cloud, or a combination of these, in either a centralized data center or at the edge

**Purpose:** A general purpose OpenShift reference architecture to showcase the best of Intel and Red Hat products with key workloads



### Solution ecosystem



### Intel enabling status

- Intel® Xeon (2<sup>nd</sup> Gen – Cascade Lake, 3<sup>rd</sup> Gen – Ice Lake)
- Intel Optane (PMEM, SSD); Columbiaville

### Collateral

- [Intel OpenShift RA for 4.6](#)
- [Intel OpenShift Solution Brief for 4.6](#)
- [Red Hat: OpenShift Ref Arch – Multiple OEMs](#)
- [Dell: OpenShift Offering](#)
- [HPE: OpenShift Offering](#)
- [Cisco: OpenShift Offering](#)
- [Lenovo: OpenShift Offering](#)
- [Supermicro: OpenShift Offering](#)
- [Penguin Computing: OpenShift Offering](#)

# Gain robust repeatability as self service, by automating the automation

Speaker:

Senior Solution Architect

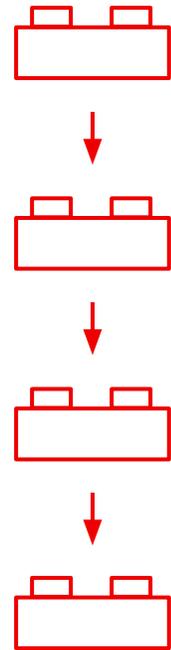
Michael Bang



In this presentation i am going to talk about

- Standardisation
- Automation
- Collaboration

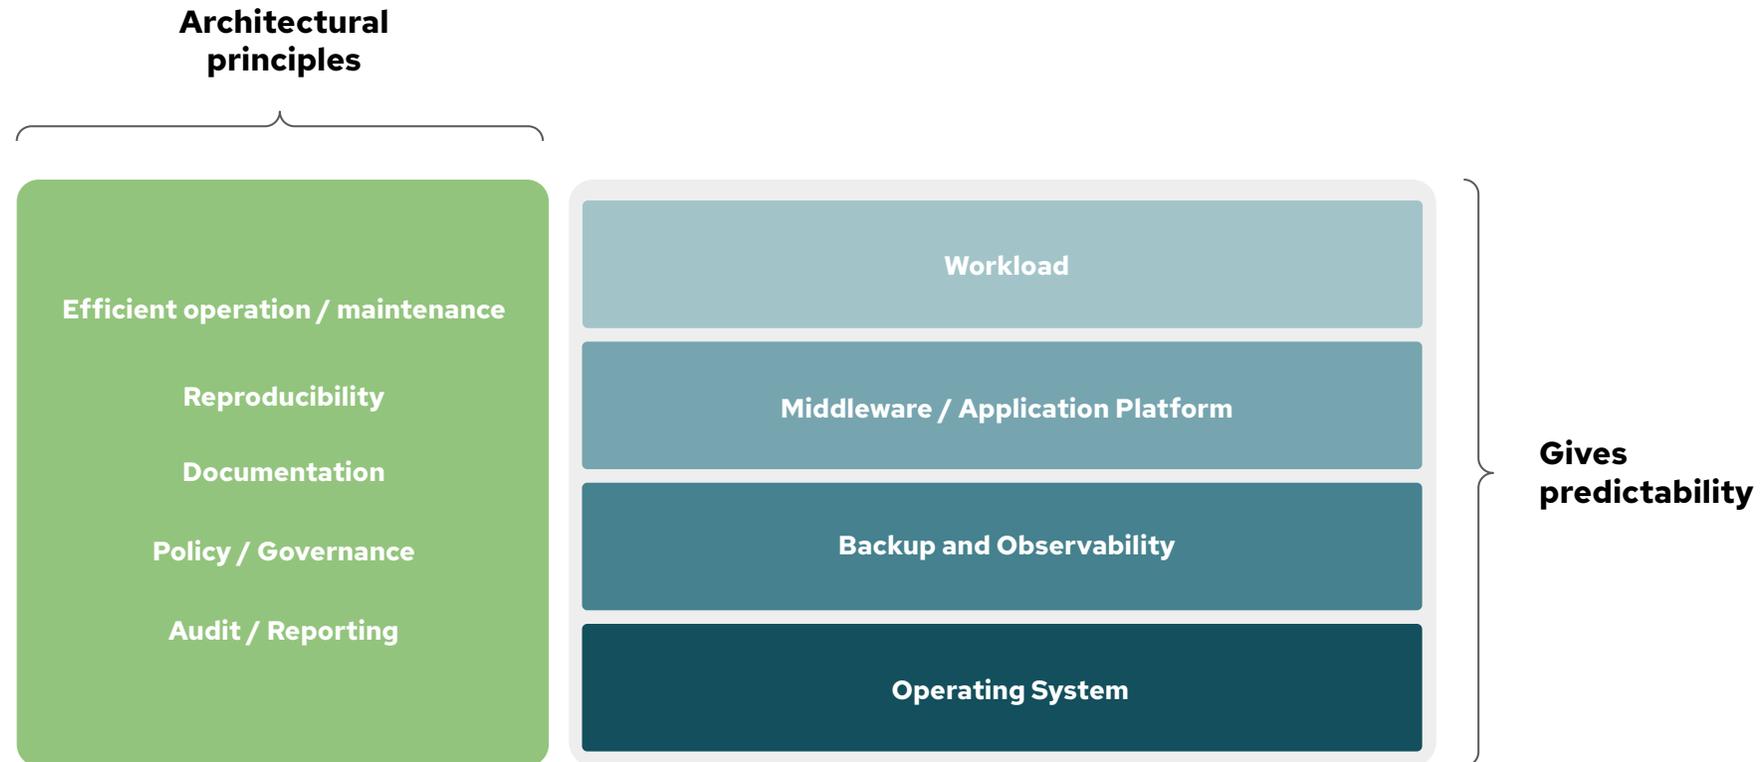
It all starts with  
standardisation





# Standardisation in IT

## Strategic Viewpoint



# Standard Operating Environment (SOE)

## Definition:

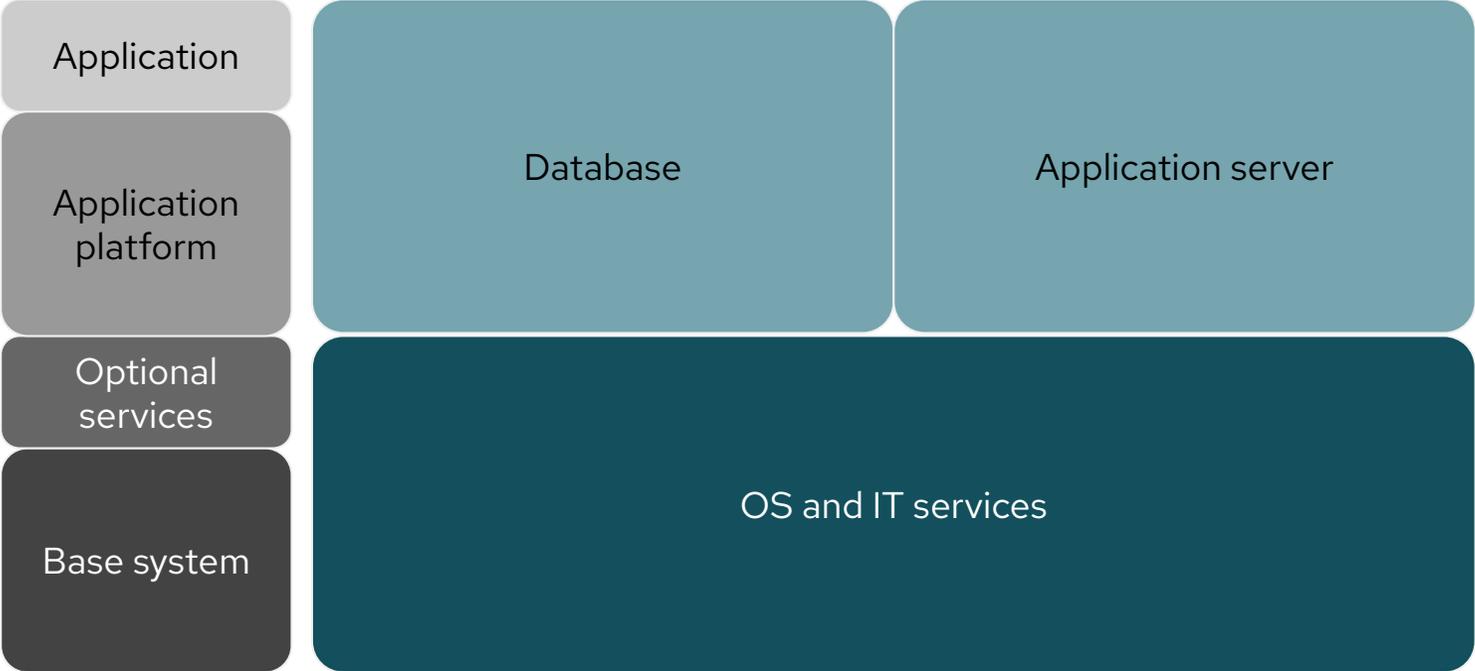
“Provides tools, standards and best practices to manage the lifecycle of an entire, deployed stack – from operating system and infrastructure services through to middleware and applications.”

## What areas does it focus on?

- ▶ Automation
- ▶ Standardisation
- ▶ Lifecycle management
- ▶ Reporting

# Standard Operating Environment (SOE)

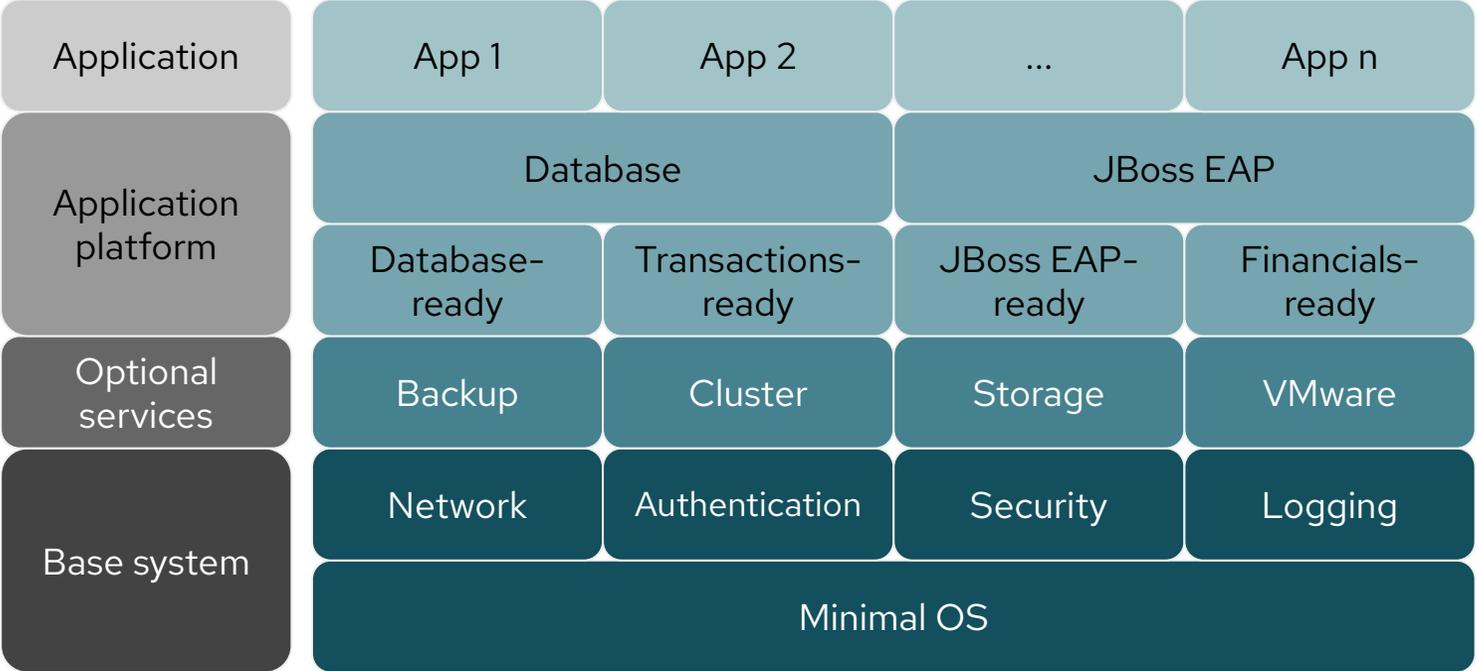
One-size-fits-most, generic servers with functional application blocks



Basic approach

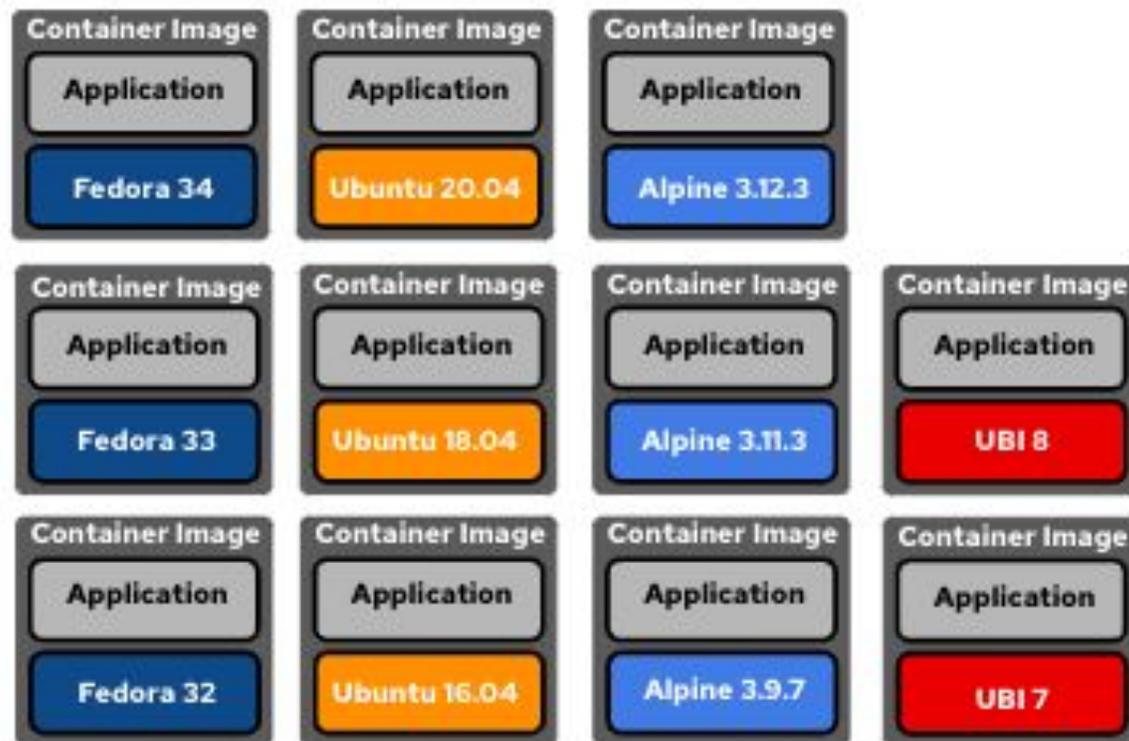
# Standard Operating Environment (SOE)

Concept: Independent yet compatible and interchangeable components



## Adaptive SOE approach

What about containers? they need SOE's too



- 8 different versions of glibc
- 3 different versions of muslc
- 11 different versions of openssl

# Efficiency Through Automation

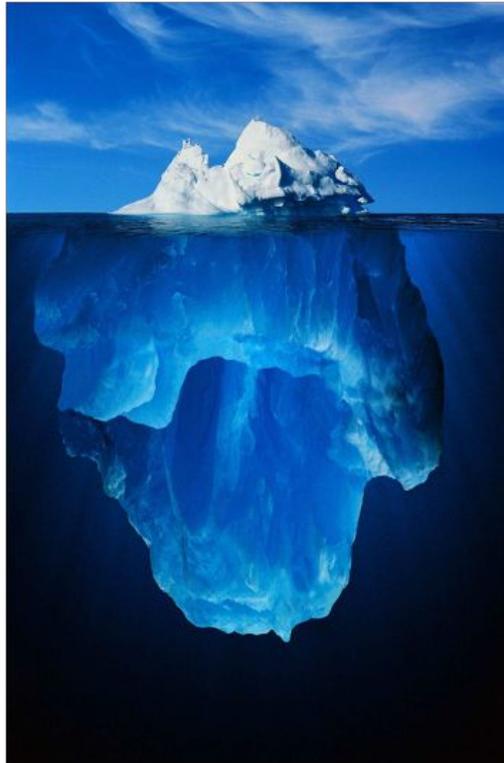
Ok, standards are great, but:

- ▶ only define point-in-time snapshots of the environment
- ▶ take time to maintain in a complex environment
- ▶ By automating the process of implementing standards we achieve:
- ▶ higher flexibility to accommodate change  $\Rightarrow$  higher agility
- ▶ higher operational efficiency
- ▶ eases lifecycle management

What kinds of automation?

# IT Automation

your stack



Business value

Lots of tech

use cases

FULLY AUTOMATED  
PROVISIONING

SECURITY/  
COMPLIANCE

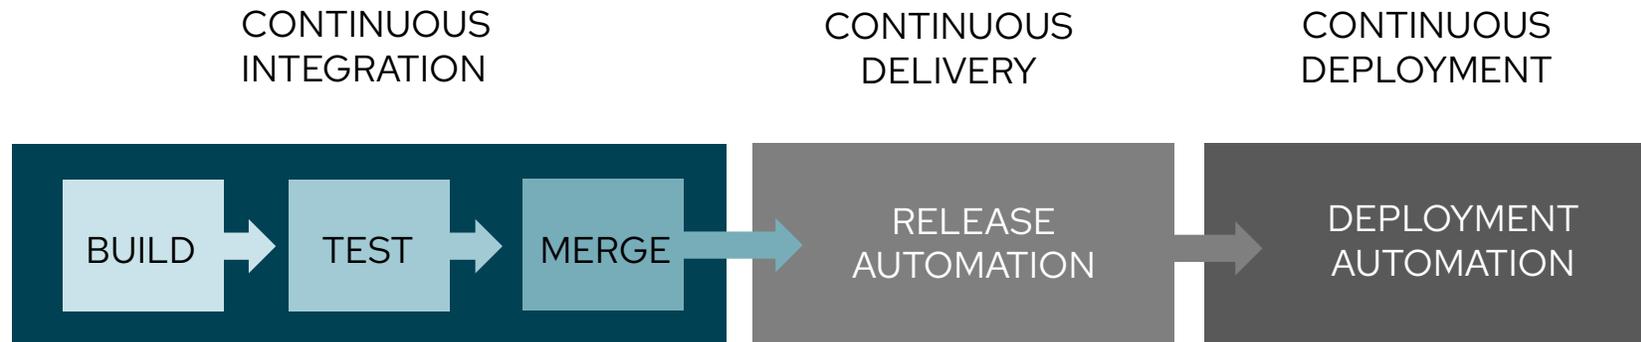
CONFIG  
MANAGEMENT

CONTINUOUS  
DELIVERY

ORCHESTRATION

APP  
DEPLOYMENT

# Application Development and Deployment



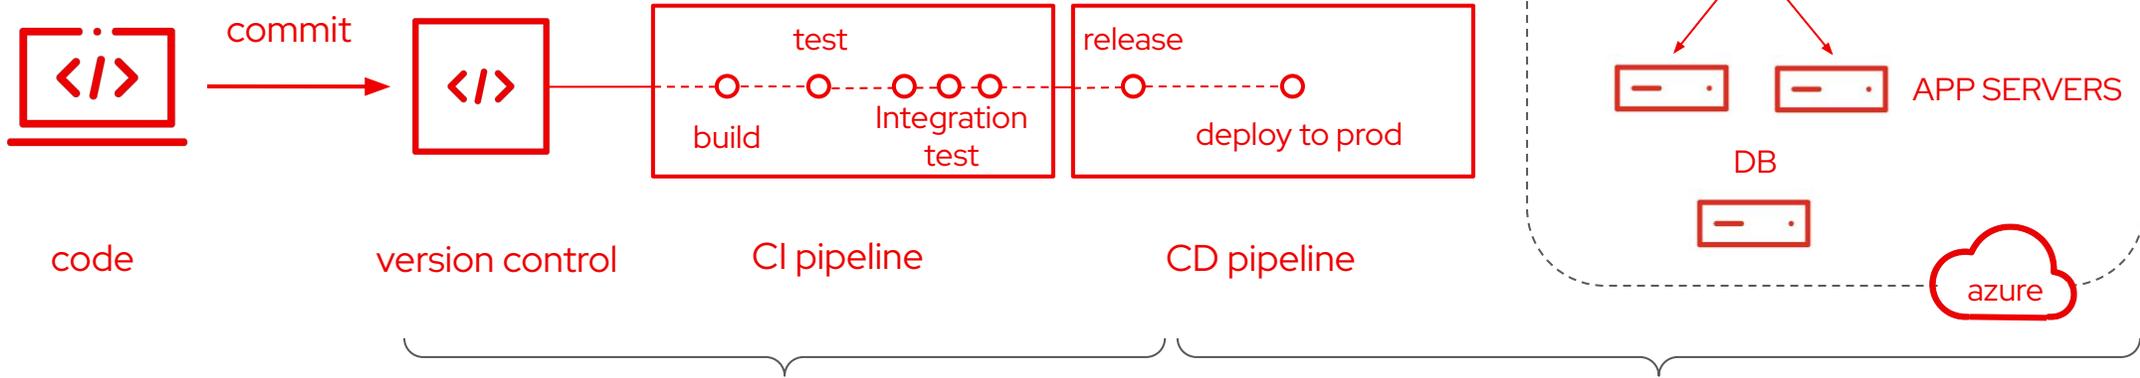
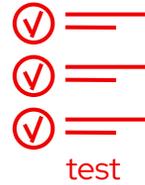
# GitLab & Ansible Automation Platform Application upgrade via CI/CD

In this demo you will see

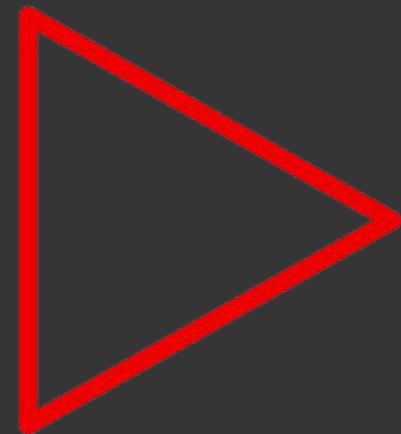
- Commit application change to git
- Triggers pipeline run with tests
- Deploy change to production
- Remove 1st appserver from load balancer
- Update 1st appserver and enable in load balancer
- Remove 2nd appserver from load balancer
- Update 2nd appserver and enable in load balancer.

**Seamless  
upgrade of  
application**

Application	App 1	App 2	...	App n
Application platform	Database		JBoss EAP	
	Database-ready	Transactions-ready	JBoss EAP-ready	Financials-ready
Optional services	Backup	Cluster	Storage	VMware
Base system	Network	Authentication	Security	Logging
	Minimal OS			



**DEMO**



- ✓ No manual steps
- ✓ No human errors
- ✓ Predictable outcomes
- ✓ Higher efficiency
- ✓ Faster time to market
- ✓ Less stress

**Seamless  
upgrade of  
application**

# Continuous Integration

# OpenShift Pipelines

Open source, standardised  
cloud-native style



**based on TEKTON**

- Self service application platform
- Build and test your application automatically
- Standardised native tools
- Everything as code
  - Application
  - Deployment
  - Build and test
- Collaborative workflow

**Simplified  
collaborative  
application  
development**

# Automate the automation

Standardise your CI - cloud-native style



**Pipelines as a (cluster) service**



**CI resources - cloud-native**



**Git-centric workflow**



# Why OpenShift Pipelines?



Built for  
Cloud-native



Scale  
on-demand



Secure pipeline  
execution



Flexible and  
powerful

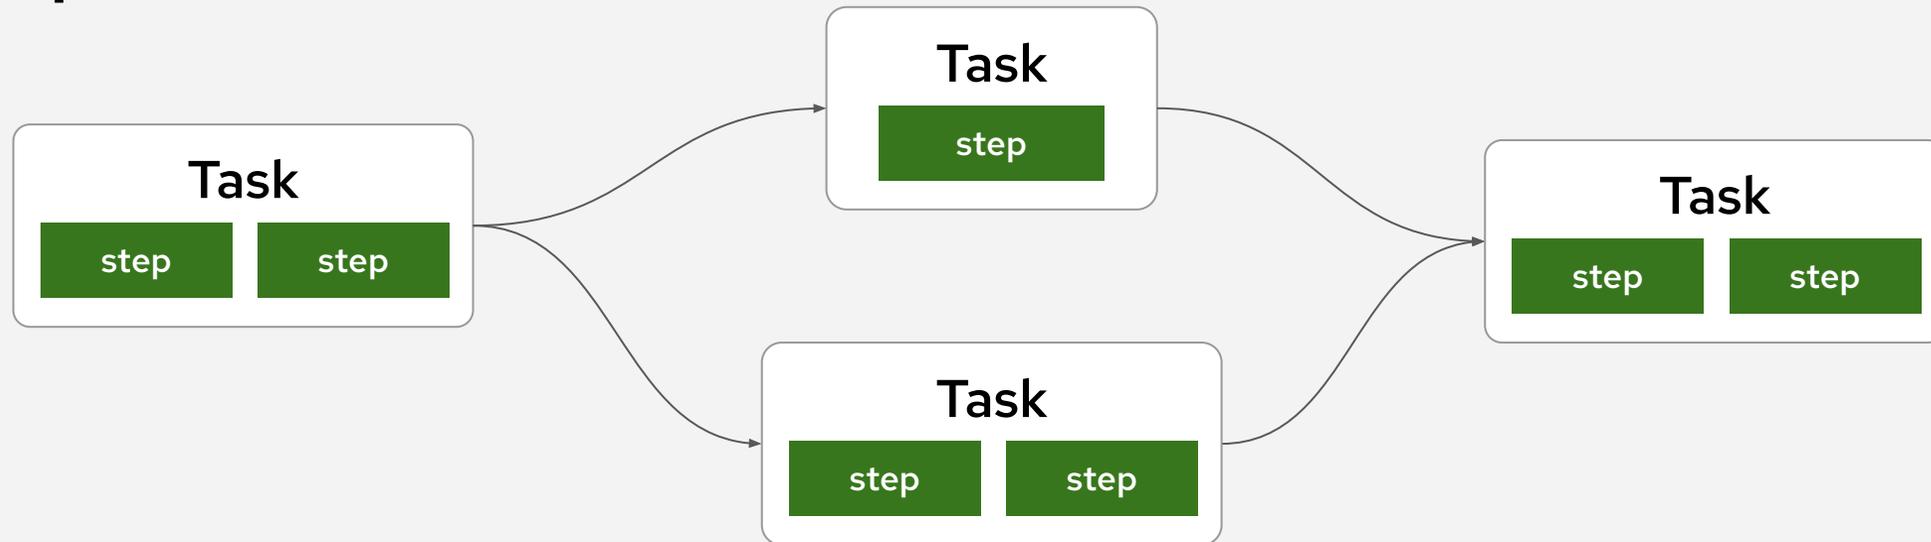
# Pipelines as a service

The screenshot shows the Red Hat OpenShift Container Platform interface. The left sidebar is expanded to 'Installed Operators'. The main content area shows the details for the 'Red Hat OpenShift Pipelines' operator (version 1.7.0). The 'Description' section states: 'Red Hat OpenShift Pipelines is a cloud-native continuous integration and delivery open-source CI/CD framework, which enables automating deployments across multiple environments.' The 'Features' section lists: 'Standard CI/CD pipelines definition' and 'Build images with Kubernetes tools such as S2I, Buildah, Buildpacks, Kaniko'.

The screenshot shows the Red Hat OpenShift Container Platform interface for a specific pipeline run. The left sidebar is expanded to 'Developer'. The main content area shows the details for a pipeline run named 'petclinic-deploy-dev-run-qwqx4', which has a status of 'Succeeded'. Below this, there is a 'Pipeline Run Overview' diagram showing a sequence of steps: 'unit-tests', 'release-app', 'code-analysis', 'generate-r...', 'build-image', 'deploy', 'int-test', and 'perf-test'. All steps are marked with a green checkmark, indicating they completed successfully.

# OpenShift Pipelines - Tekton concepts

## Pipeline

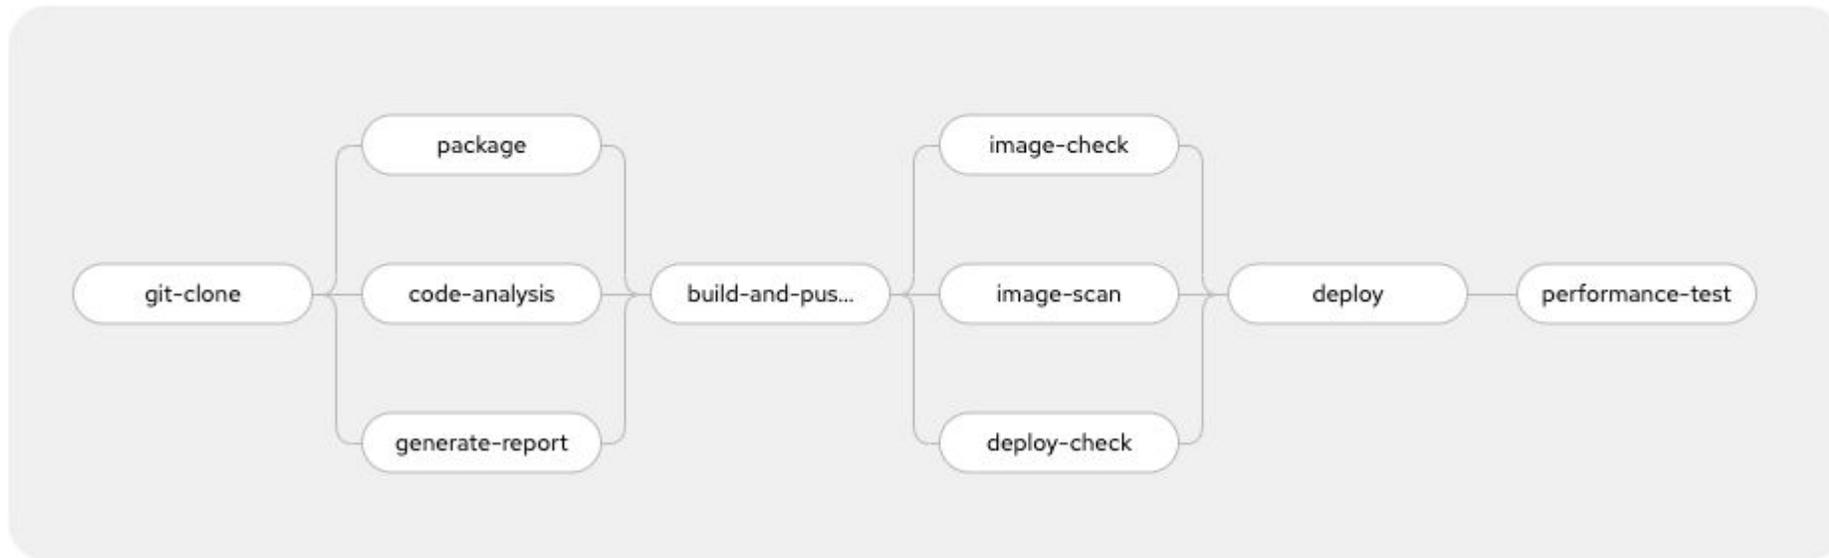


Pipelines > Pipeline details

## PL build-and-push-image

[Details](#) [Metrics](#) [YAML](#) [PipelineRuns](#) [Parameters](#) [Resources](#)

### Pipeline details



## PipelineRun details



PLR build-and-push-image-62yddj ✔ Succeeded

Actions ▾

Details [YAML](#) [TaskRuns](#) [Logs](#) [Events](#)[Download](#) | [Download all task logs](#) | [Expand](#)✔ git-clone✔ package✔ generate-report✔ code-analysis✔ build-and-push-image✔ image-scan✔ image-check✔ deploy-check✔ deploy✔ performance-test

build-and-push-image

```

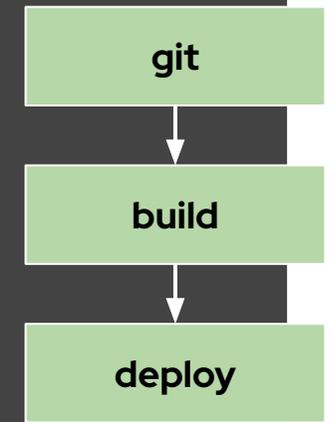
[INFO] Scanning for projects...
[INFO]
[INFO] -----< org.wanja.demo:person-service >-----
[INFO] Building person-service 1.6.0
[INFO] -----[ jar ]-----
[INFO]
[INFO] --- maven-resources-plugin:2.6:resources (default-resources) @ person-service ---
[INFO] Using 'UTF-8' encoding to copy filtered resources.
[INFO] Copying 3 resources
[INFO]
[INFO] --- quarkus-maven-plugin:2.7.5.Final:generate-code (default) @ person-service ---
[INFO]
[INFO] --- maven-compiler-plugin:3.8.1:compile (default-compile) @ person-service ---
[INFO] Nothing to compile - all classes are up to date
[INFO]
[INFO] --- quarkus-maven-plugin:2.7.5.Final:generate-code-tests (default) @ person-service ---
[INFO]
[INFO] --- maven-resources-plugin:2.6:testResources (default-testResources) @ person-service ---
[INFO] Using 'UTF-8' encoding to copy filtered resources.
[INFO] skip non existing resourceDirectory /workspace/source/the-source/person-service/src/test/resources
[INFO]
[INFO] --- maven-compiler-plugin:3.8.1:testCompile (default-testCompile) @ person-service ---
[INFO] Nothing to compile - all classes are up to date
[INFO]
[INFO] --- maven-surefire-plugin:3.0.0-M5:test (default-test) @ person-service ---
[INFO] Tests are skipped.
[INFO]
[INFO] --- maven-jar-plugin:2.4:jar (default-jar) @ person-service ---
[INFO]
[INFO] --- quarkus-maven-plugin:2.7.5.Final:build (default) @ person-service ---
[INFO] [io.quarkus.kubernetes.deployment.KubernetesDeployer] Selecting target 'openshift' since it has the highest priority among the implicitly enabled deploy
[WARNING] [io.quarkus.kubernetes.deployment.KubernetesDeployer] An openshift deployment was requested, but the container image group:mbang1 is not aligned with

```

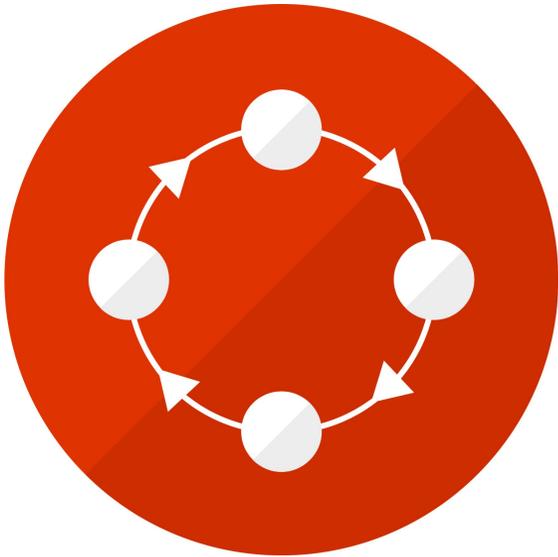
# Everything as code

```
.
├── my-service
│   └── <application-code>
├── manifests
│   ├── deployments
│   │   ├── deployment.yaml
│   │   └── service.yaml
│   └── pipelines
│       ├── tekton-pipeline.yaml
│       └── tasks
│           ├── kustomize-task.yaml
│           └── maven-task.yaml
```

```
kind: Pipeline
metadata:
  name: deploy-dev
spec:
  params:
    - name: IMAGE_TAG
  tasks:
    - name: git
      taskRef:
        name: git-clone
        params: [...]
    - name: build
      taskRef:
        name: maven
        params: [...]
        runAfter: ["git"]
    - name: deploy
      taskRef:
        name: knative-deploy
        params: [...]
        runAfter: ["build"]
```

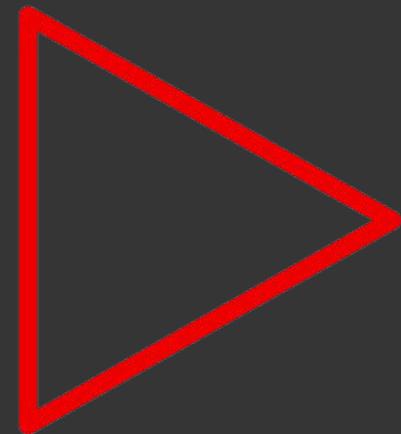


# Pipelines as code



- GitOps enabled - git-centric workflow
- Integrated with Git provider
  - Events, actions
- Pipelines run in cluster
  - No pre-configured infrastructure

**DEMO**



What did we see in the demo

- ✓ No manual steps
- ✓ No human errors
- ✓ Predictable outcomes
- ✓ Higher efficiency
- ✓ Faster time to market
- ✓ Less stress

**Simplified  
collaborative  
application  
development**

# Continuous Deployment

Hybrid cloud pattern: Multicloud GitOps

- Keep delivering no matter the location
- Automate introduction of new features
- Manage risks by replication and scaling out environments
- Everything automated

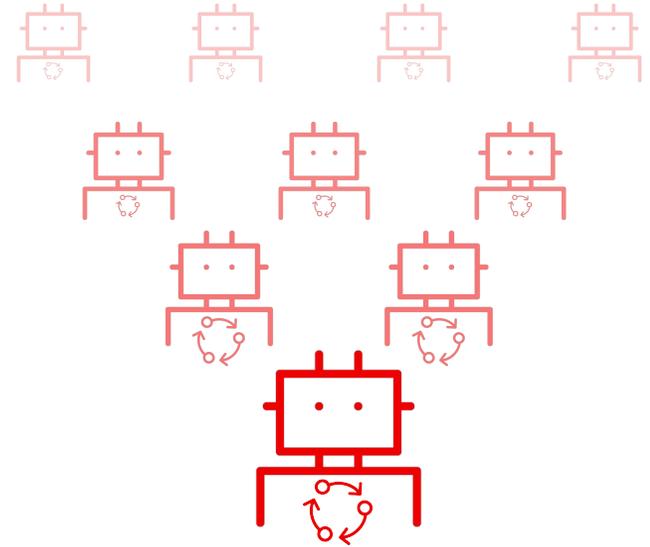
**Automated  
business  
continuity**

We want everything as code.

Applications, configurations  
and secrets delivered to  
autonomous environments.

Visible change history.

Comes with self healing.





## Provided APIs

### **A** Application

An Application is a group of Kubernetes resources as defined by a manifest.

[+ Create instance](#)

### **AS** ApplicationSet

ApplicationSet is the representation of an ApplicationSet controller deployment.

[+ Create instance](#)

### **AP** AppProject

An AppProject is a logical grouping of Argo CD Applications.

[+ Create instance](#)

### **ACD** Argo CD

Argo CD is the representation of an Argo CD deployment.

[+ Create instance](#)

Operator based



# OpenShift GitOps



## Multi-cluster config management

Declaratively manage cluster and application configurations across multi-cluster OpenShift and Kubernetes infrastructure with Argo CD



## Automated Argo CD install and upgrade

Automated install, configurations and upgrade of Argo CD through OperatorHub



## Opinionated GitOps bootstrapping

Bootstrap end-to-end GitOps workflows for application delivery using Argo CD and Tekton with GitOps Application Manager CLI

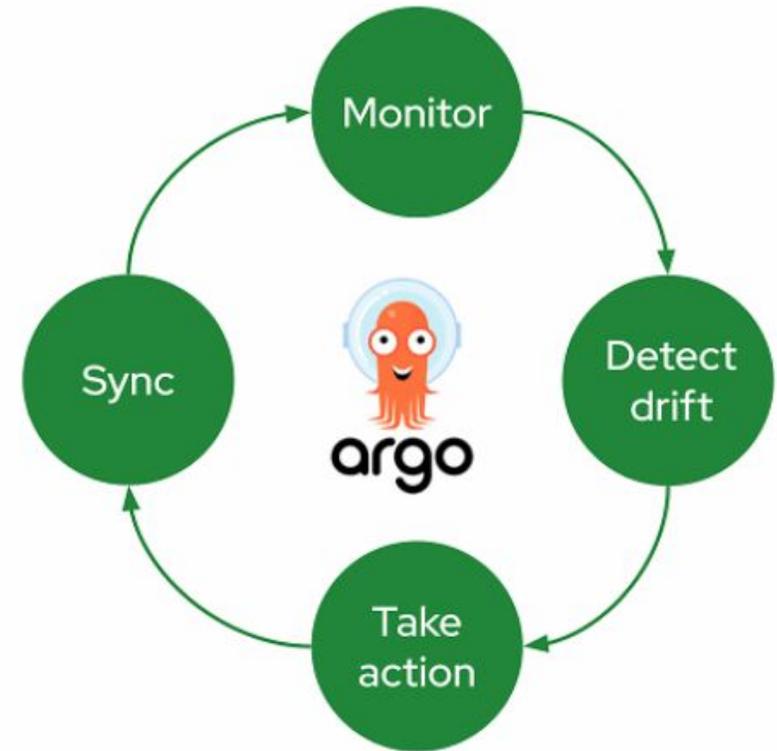


## Deployments and environments insights

Visibility into application deployments across environments and the history of deployments in the OpenShift Console

# Argo CD

- Cluster and application configuration versioned in Git
- Automatically syncs configuration from Git to clusters
- Drift detection, visualization and correction
- Granular control over sync order for complex rollouts
- Rollback and rollforward to any Git commit
- Manifest templating support (Helm, Kustomize, etc)
- Visual insight into sync status and history



**Applications**

+ NEW APP SYNC APPS REFRESH APPS Search applications... /

**FILTERS**

FAVORITES ONLY

**SYNC STATUS**

- Unknown 0
- Synced 1
- OutOfSync 0

**HEALTH STATUS**

- Unknown 0
- Progressing 0

**kustomize-dev-demo** ★

Project: default

Labels:

Status: ♥ Healthy ✔ Synced

Reposito... <https://github.com/RedHatNordicsSA/ad...>

Target R... HEAD

Path: overlay/dev

Destinati... in-cluster

Namesp... dev

REFRESH DELETE

**APPLICATION DETAILS**

[MORE](#)

To 0a58111

(GMT+0200)

g@redhat.com> - sync-hooks

Hybrid Cloud Patterns

Search Hybrid Cloud Patterns

Validated Patterns Blog

## Hybrid Cloud Patterns

Hybrid Cloud Patterns are a natural progression from reference architectures with additional value.

This effort is focused on customer solutions that involve multiple Red Hat products. The patterns include one or more applications that are based on successfully deployed customer examples. Example application code is provided as a demonstration along with the various open source projects and Red Hat products required to for the deployment to work. Users can then modify the pattern for their own specific application.

How do we select and produce a pattern? We look for novel customer use cases, obtain an open source demonstration of the use case, validate the pattern with its components with the relevant product engineering teams, and create GitOps based automation to make them easily repeatable and extendable.

The automation also enables the solution to be added to Continuous Integration (CI), with triggers for new product versions (including betas), so that we can proactively find and fix breakage and avoid bit-rot.

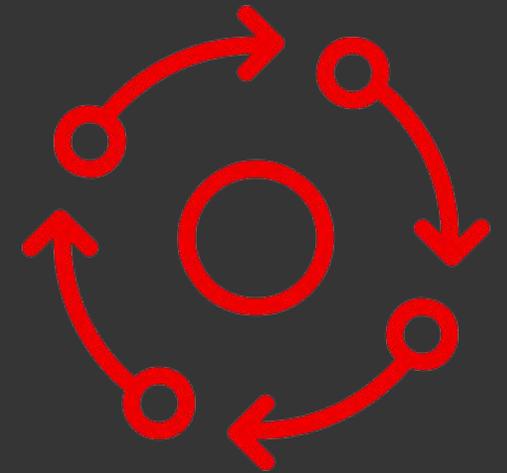
### Who should use these patterns?

It is recommended that architects or advanced developers with knowledge of Kubernetes and Red Hat OpenShift Container Platform use these patterns. There are advanced [Cloud Native](#) concepts and projects deployed as part of the pattern framework. These include, but are not limited to, OpenShift Gitops ([ArgoCD](#)), Advanced Cluster Management ([Open Cluster Management](#)), and OpenShift Pipelines ([Tekton](#))

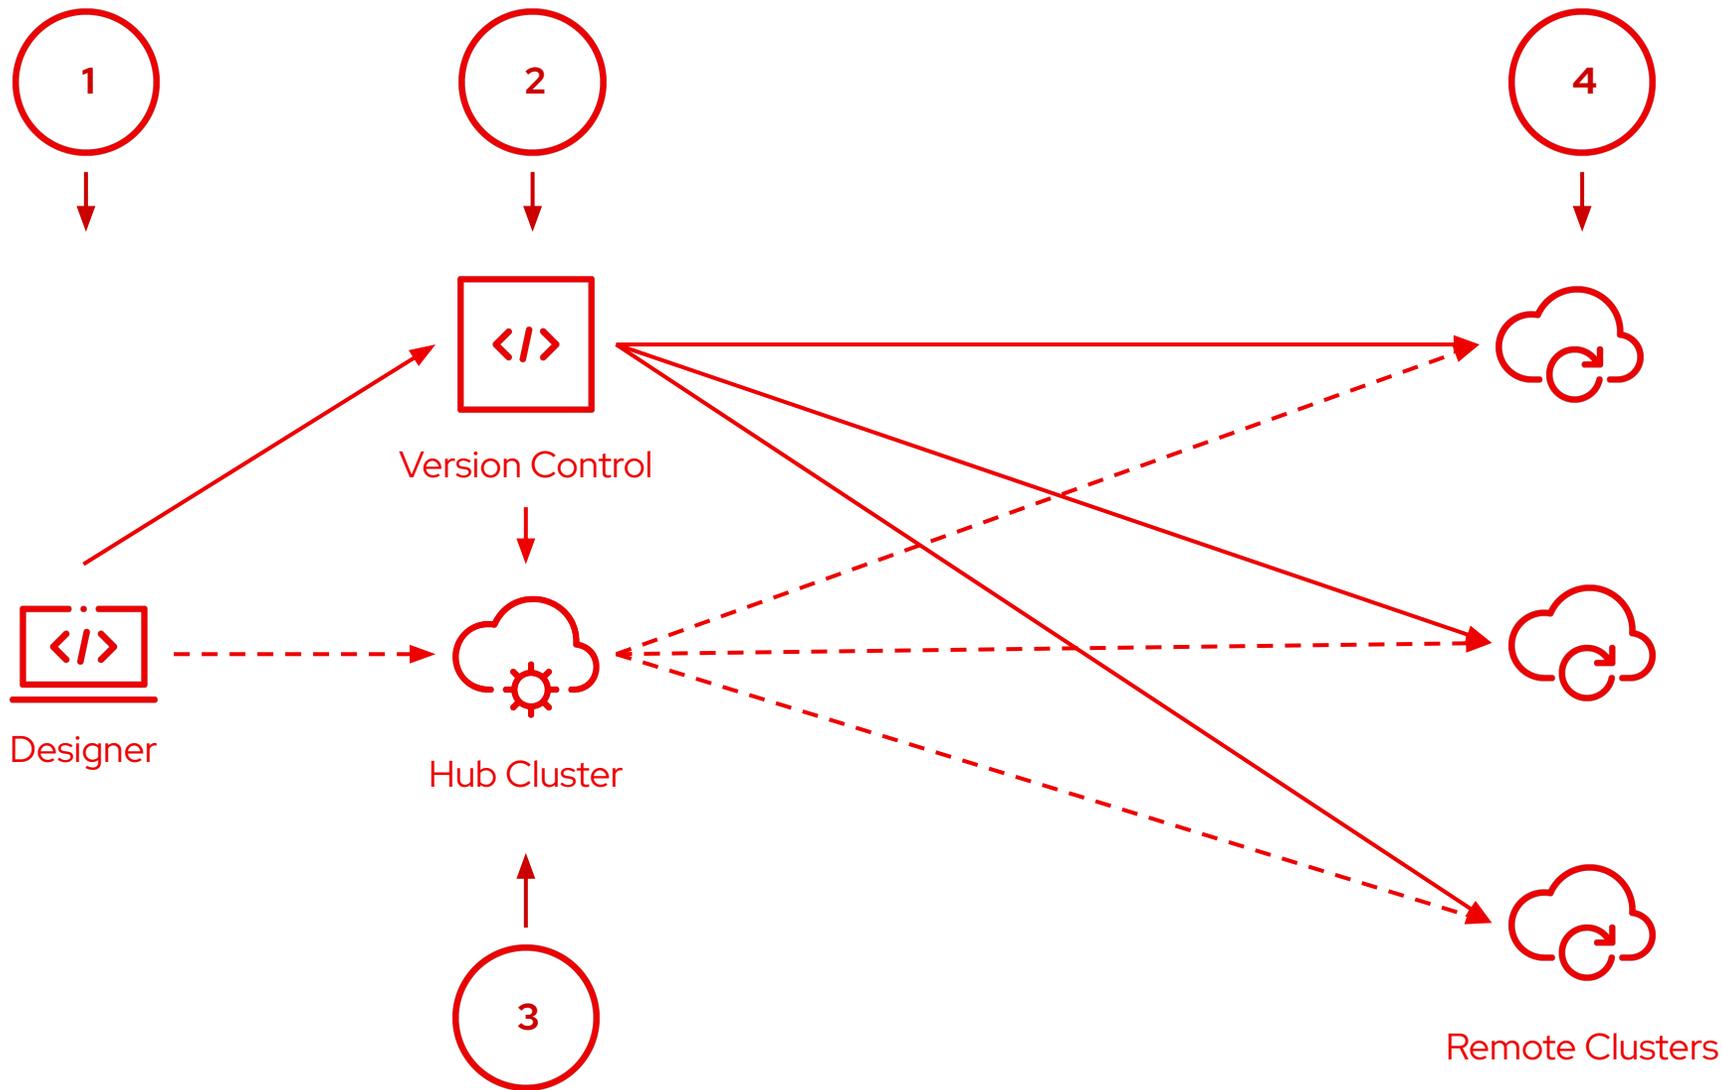
This site uses [Just the Docs](#), a documentation theme for Jekyll.

Let's look at the multicloud gitops pattern today

- 1 **Create or Enhance**
- 2 **Version Control**
- 3 **Automate**
- 4 **Continuously Deployment**



LANDSCAPE

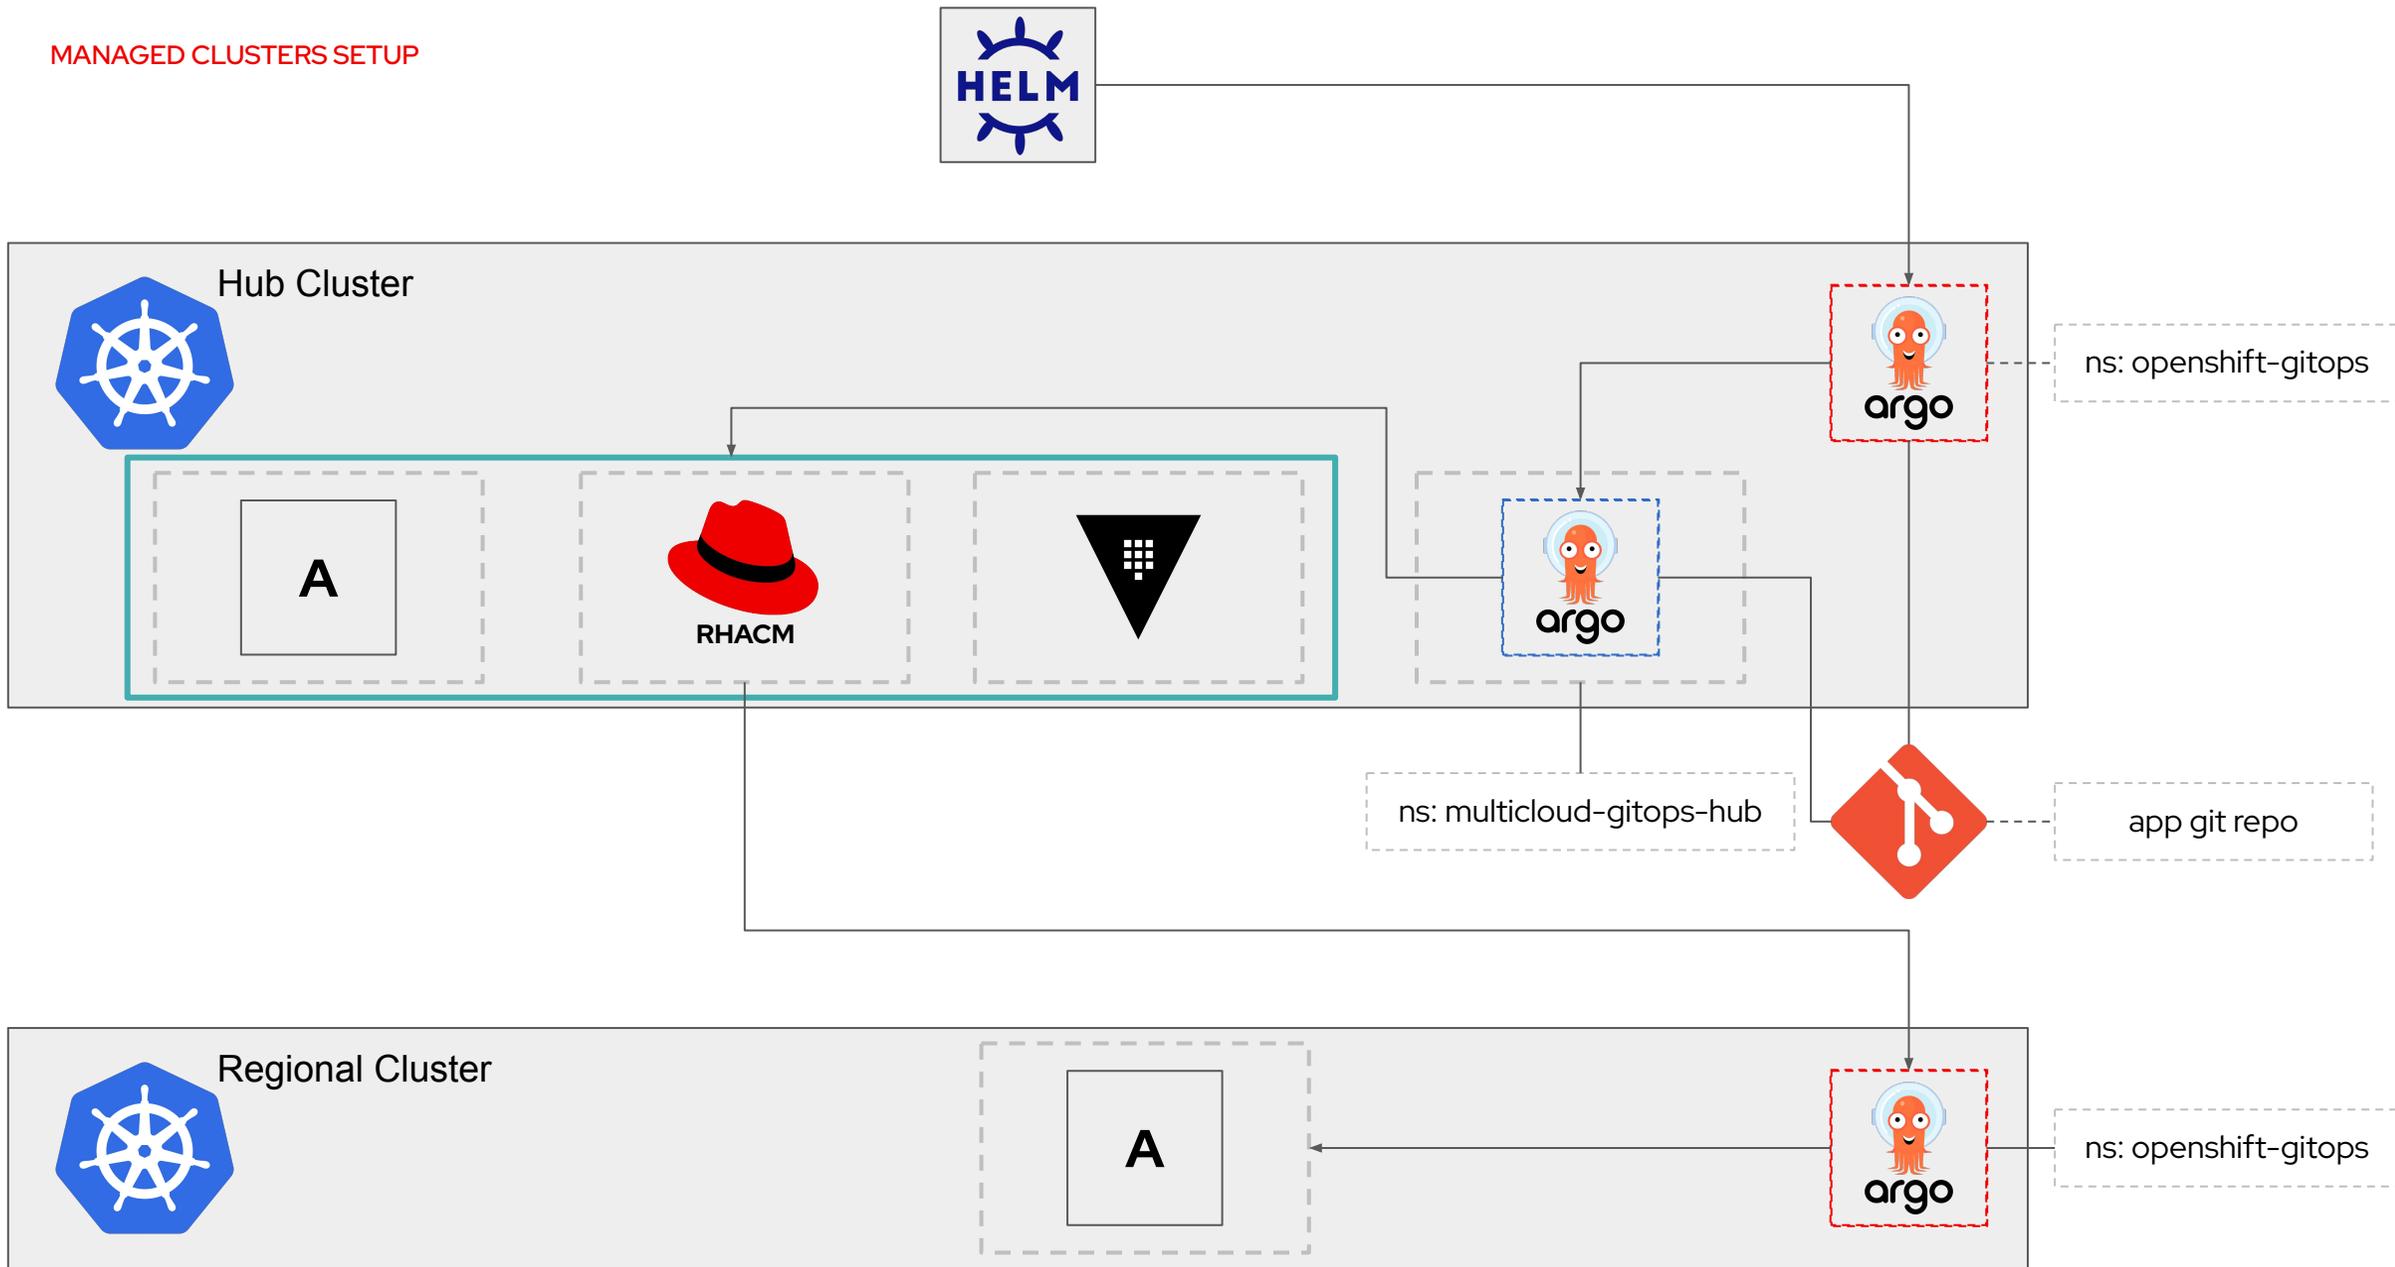


One repository to control delivery versions

Several environments in hybrid clouds to automatically adapt to configuration or application changes.

- 1. Create
- 2. Commit
- 3. Automate
- 4. Keep Delivering

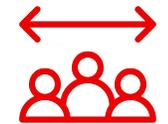
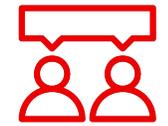
MANAGED CLUSTERS SETUP



- ✓ No manual steps
- ✓ No human errors
- ✓ Predictable outcomes
- ✓ Higher efficiency
- ✓ Faster time to market
- ✓ Less stress

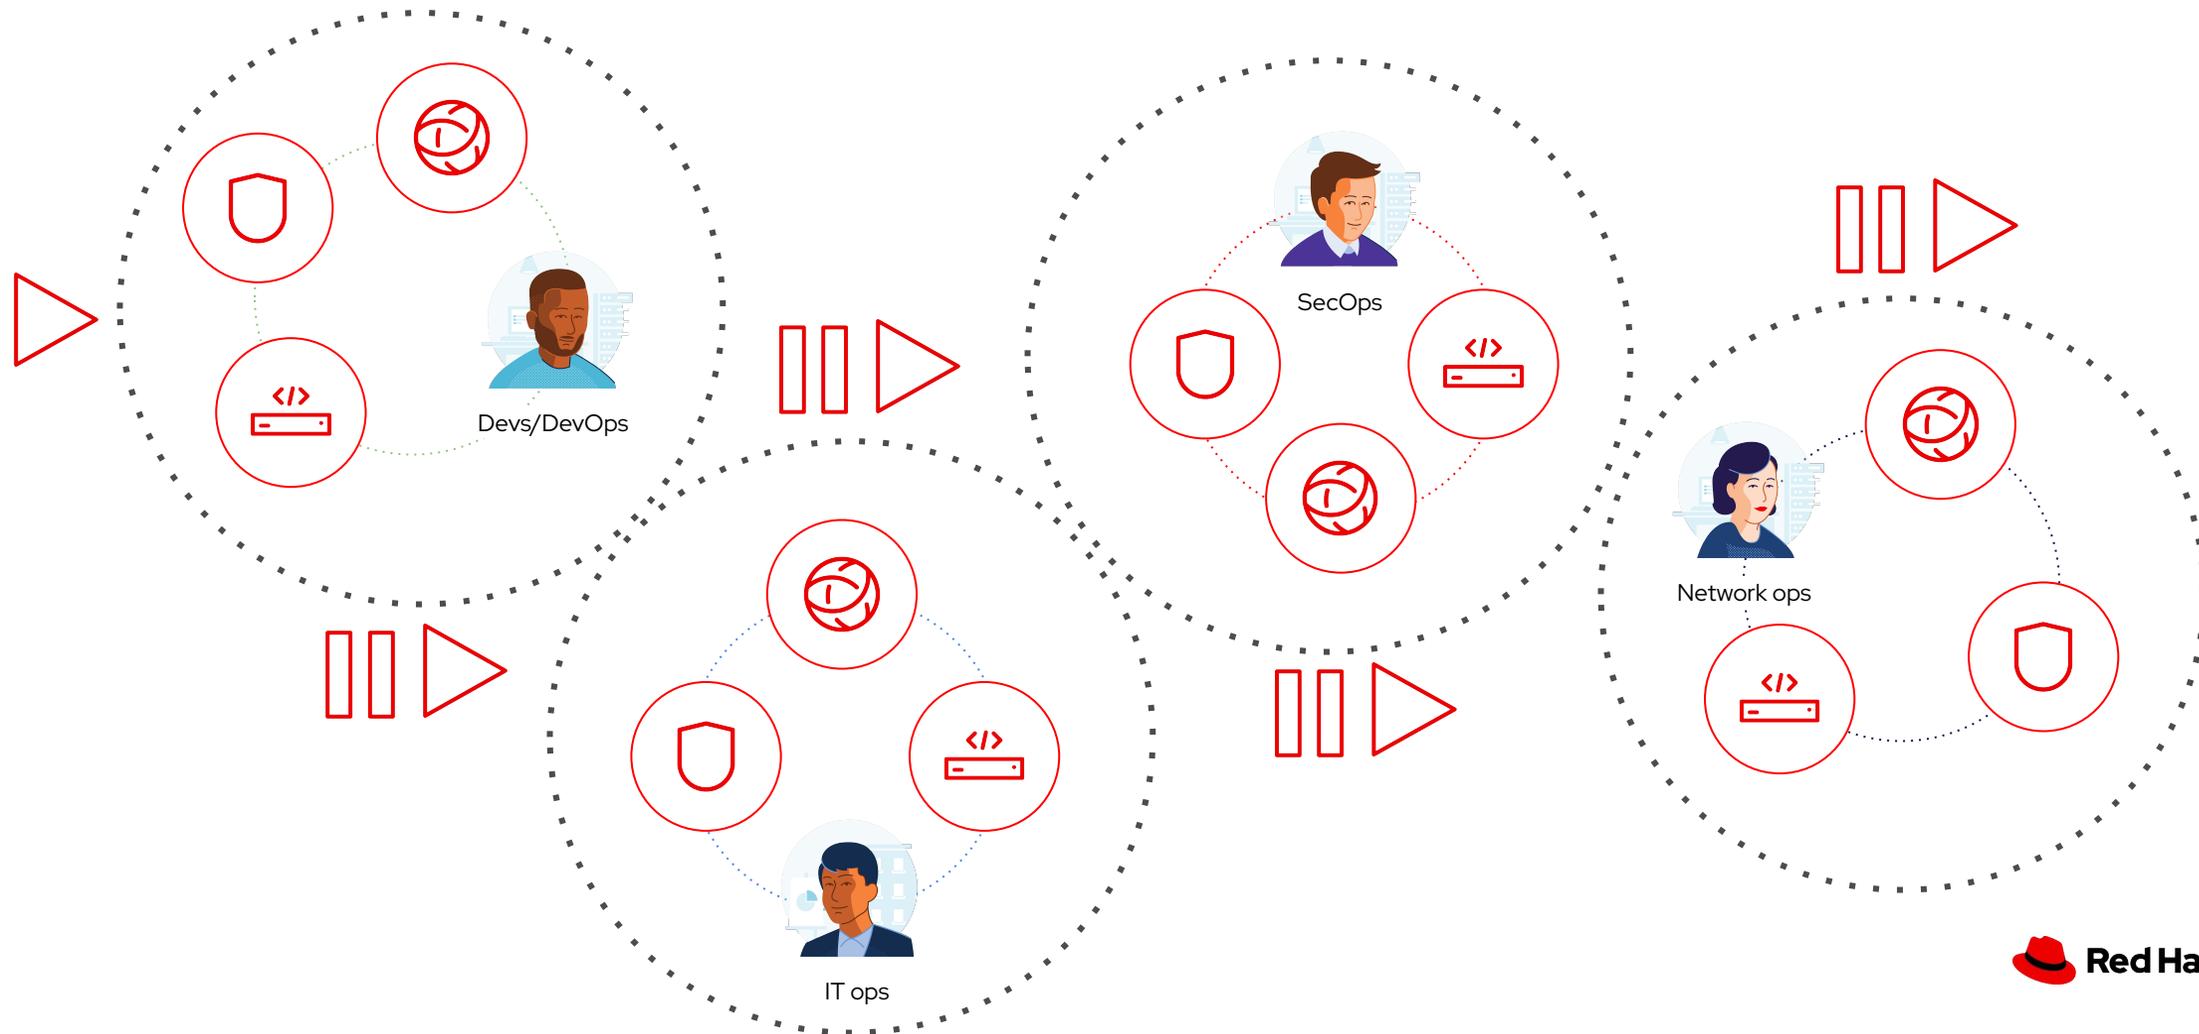
**Automated  
business  
continuity**

# Provide business value through collaboration



# But many organizations have a common problem...

Too many unintegrated, domain-specific tools, limited collaboration and scale



In this presentation you learned about

- Standardisation
- Automation
- Collaboration

To gain robust repeatability as self service, by automating the automation

Red Hat Services get you going!