


Automating rolling upgrades for RKE2 Kubernetes clusters

Stefano Bovina

Giuseppe Misurelli

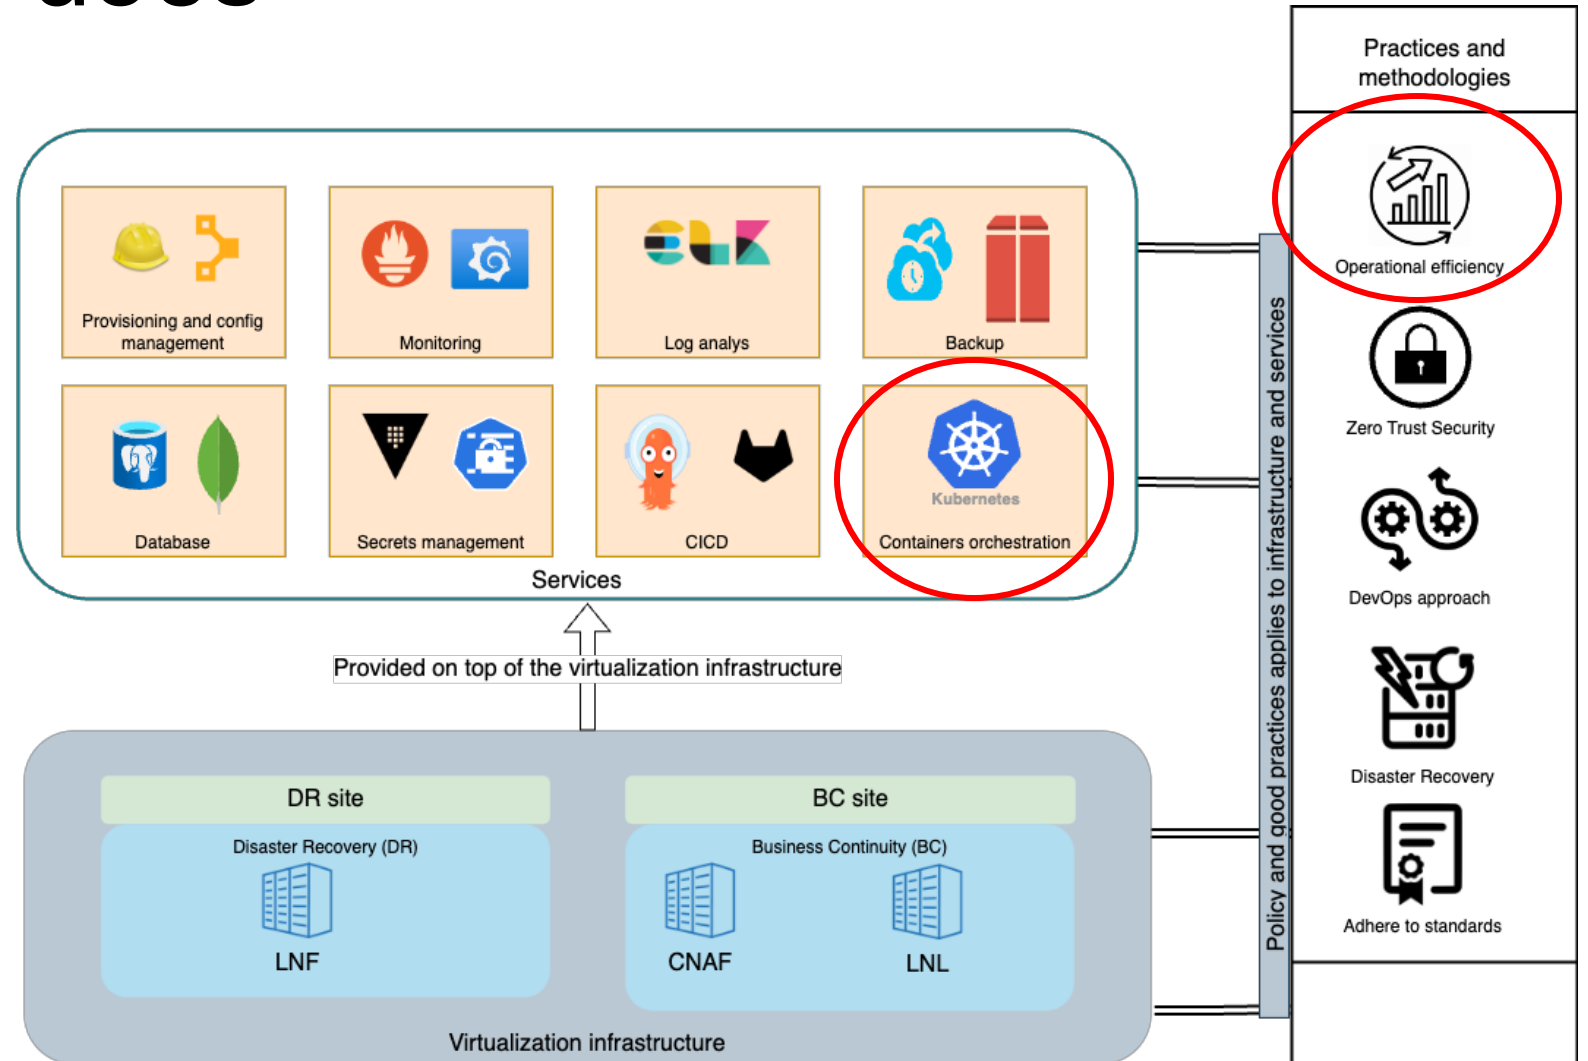
The INFN DSI apps

Ecosystem of apps serving INFN people (business trips, buying computing facilities, managing recruitment process, accounting, payrolls).



What our team does

Platform team providing and operating infrastructure and services to enable the software development lifecycle



The concept of Operational Excellence

Referenced in many Cloud Operating Models.

Formalized in the AWS Well Architected Framework as

the ability to support development and run workloads effectively, gain insight into their operations, and to continuously improve supporting processes and procedures to deliver business value

Of course we can adopt the concept also in non cloud contexts

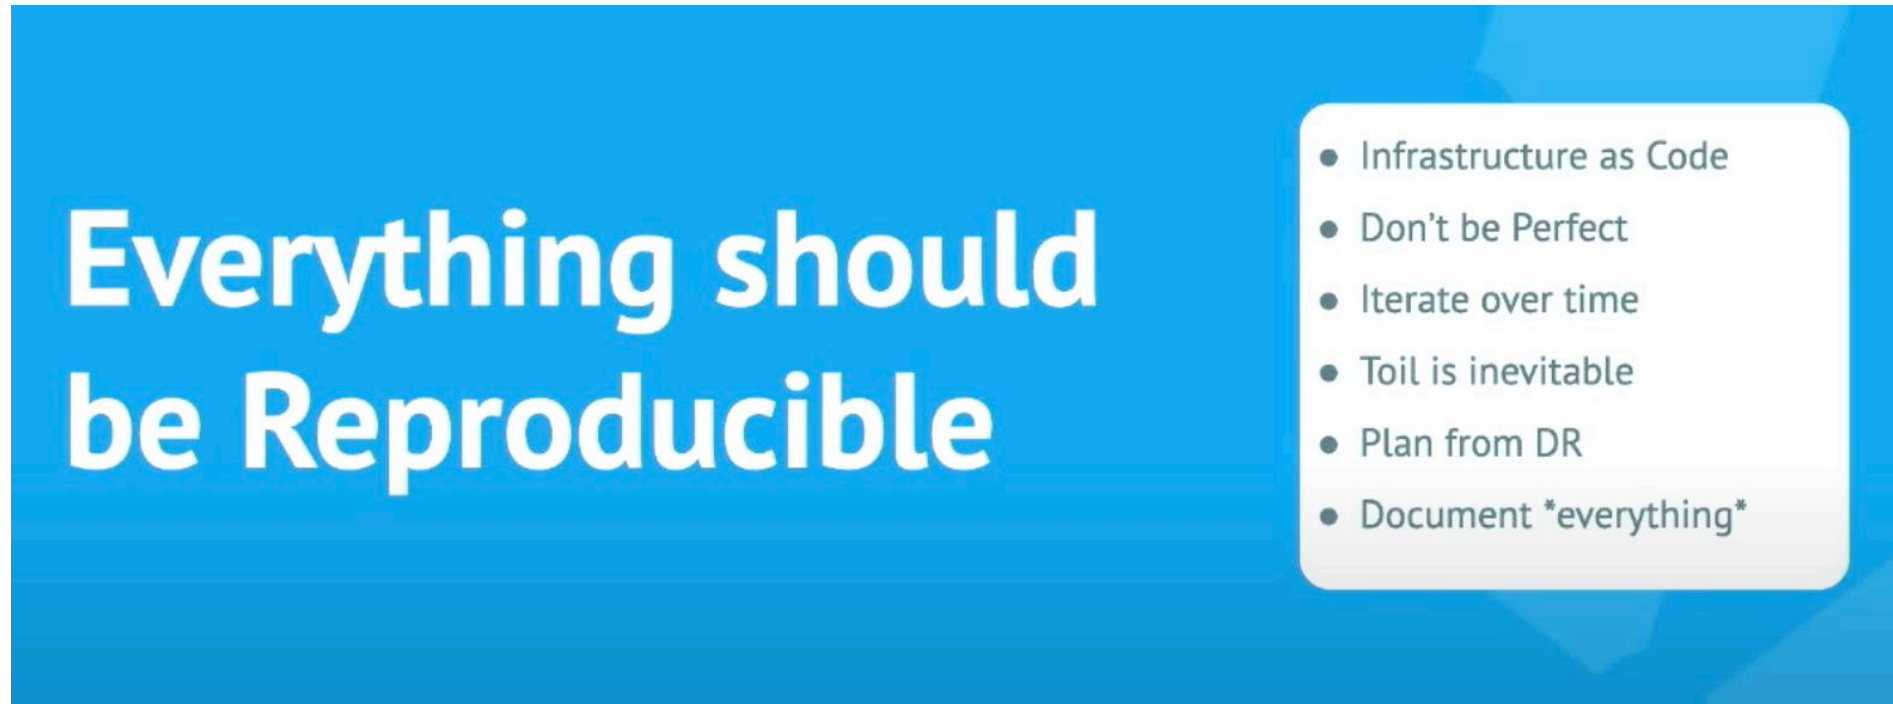
Operational Excellence for the INFN DSI apps

Or a series of best practices to adopt in order to:

1. Use as much as possible automation to create and manage infrastructure and services.
 1. Code repo (IaC) and scripts acting as automated and controlled procedures.
2. Be ready to react to events (planned and unplanned) through playbook and recepies.
3. Review continuously procedures and playbook to validate them against the evolution of infrastructure and services (es. changes, upgrades).

Use as much as possible automation

Operations as code is our inspiring principle



**Everything should
be Reproducible**

- Infrastructure as Code
- Don't be Perfect
- Iterate over time
- Toil is inevitable
- Plan from DR
- Document *everything*

Screenshot from SRE for Smaller Organization presentation - <https://youtu.be/ZmvuLLMxhug?si=rglWiJjtEbJWHLsQ>

Kubernetes release cadence

3 releases a year

N-2 support policy. 3 most recent minor versions receive security and bug fixes

Each release being supported for 14 months (12 months of support and 2 months of upgrade period)

 Last updated on 12 September 2024 🌐




[Kubernetes](#) is an open-source container-orchestration system for automating computer application deployment, scaling, and management.

Release	Released	Active Support	Maintenance Support	Latest
1.31	1 month and 3 weeks ago (13 Aug 2024)	Ends in 10 months (28 Aug 2025)	Ends in 1 year (28 Oct 2025)	1.31.1 (11 Sep 2024)
1.30	5 months and 3 weeks ago (17 Apr 2024)	Ends in 6 months and 3 weeks (28 Apr 2025)	Ends in 8 months (28 Jun 2025)	1.30.5 (12 Sep 2024)
1.29	9 months and 4 weeks ago (13 Dec 2023)	Ends in 2 months and 3 weeks (28 Dec 2024)	Ends in 4 months and 3 weeks (28 Feb 2025)	1.29.9 (11 Sep 2024)
1.28	1 year and 1 month ago (15 Aug 2023)	Ended 1 month and 1 week ago (28 Aug 2024)	Ends in 3 weeks (28 Oct 2024)	1.28.14 (11 Sep 2024)
1.27	1 year and 6 months ago (11 Apr 2023)	Ended 5 months ago (28 Apr 2024)	Ended 3 months and 1 week ago (28 Jun 2024)	1.27.16 (17 Jul 2024)

Source <https://endoflife.date/kubernetes>

DSI Kubernetes clusters upgrade



Use as much as possible automation

Start from a procedure

Be ready to react to events (planned and unplanned) through playbook and recipes.

Review continuously procedures and playbook to validate them against the evolution of infrastructure and services (es. changes, upgrades).

sysinfo_oss / ops / k8s / docs

The screenshot shows a GitHub file view for the file 'docs / upgrade_and_maintenance.md'. At the top, there is a breadcrumb navigation 'docs / upgrade_and_maintenance.md' and a search bar. To the right are buttons for 'Find file', 'Blame', 'History', and 'Permalink'. Below this is a commit summary: 'Post internal review changes' by 'Misu' authored '2 months ago', with a commit hash '6986718d' and a copy icon. The file name 'upgrade_and_maintenance.md' is shown with a size of '8.60 KiB'. A toolbar contains icons for code view, file view, 'Blame', 'Edit' (highlighted in blue), 'Replace', 'Delete', and download icons. The main content area has a 'Table of Contents' section with two items: '1. K8s cluster upgrade preparation' and '2. K8s rolling upgrade'. Below the table of contents, there is a paragraph: 'This is a collections of guides and playbooks used by the [INFN DSI SysInfo-Ops team](#) whenever a k8s cluster upgrade version should be carried out manually. Since the SysInfo-Ops team utilizes [RKE2](#) as k8s distribution, most of the described steps, commands and checks rely on RKE2 utilities.'

Sharing the INFN DSI procedure

Released for the INFN community in the [sysinfo_oss baltig project](#).

Used and evolved during the latest 7 Kubernetes upgrades.

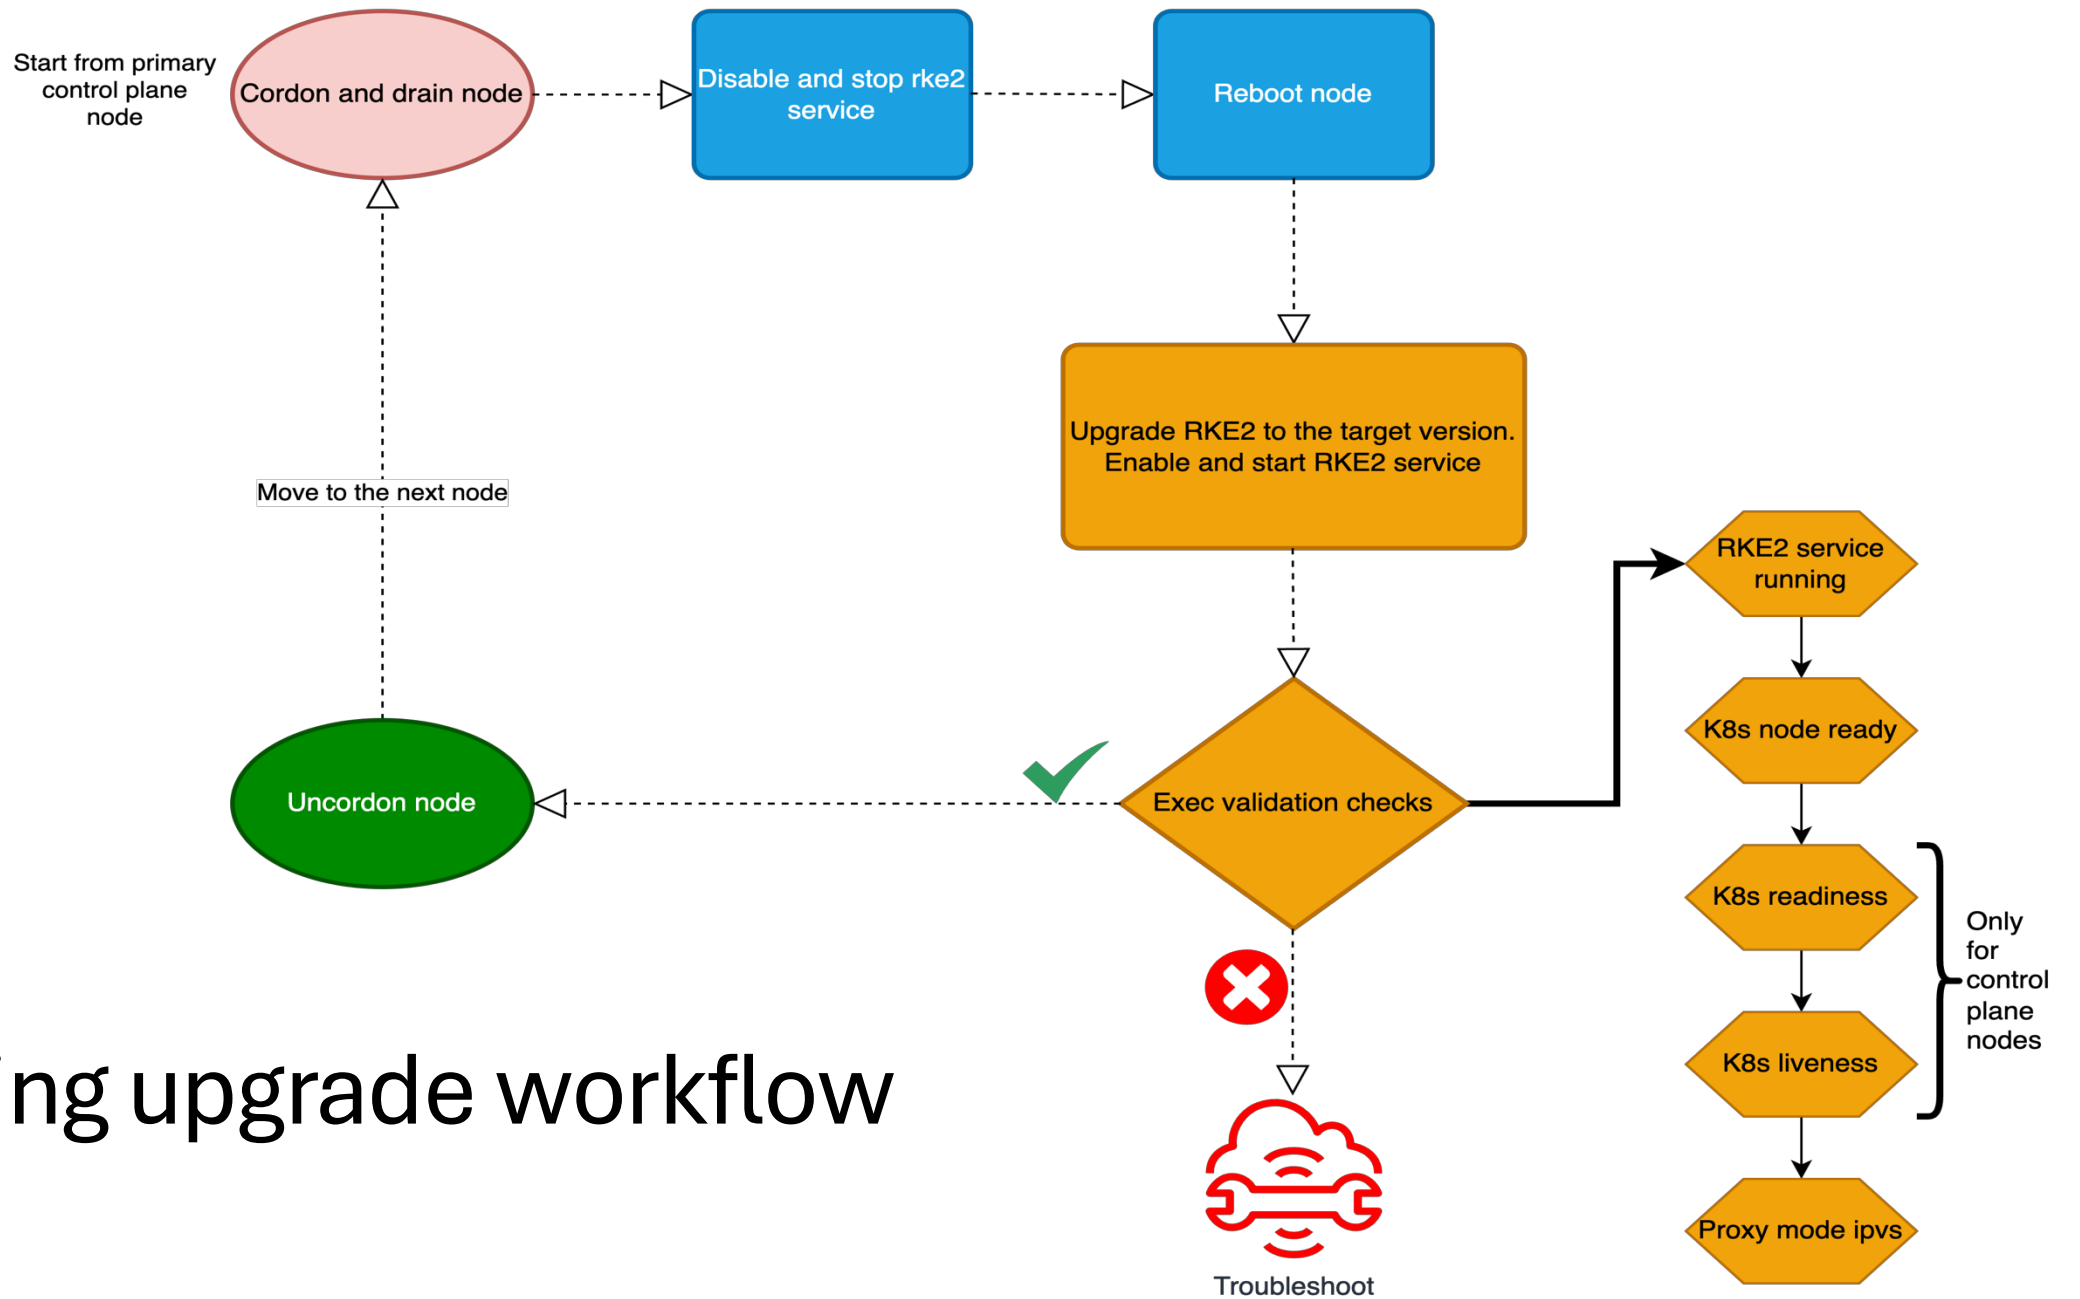
Based on RKE2 Rancher Kubernetes Engine.

K8s cluster upgrade preparation

- Get familiar with changes in target version.
- Verify addons compatibilities and api deprecations.
- Use testbed to run initial checks.

K8s rolling upgrade

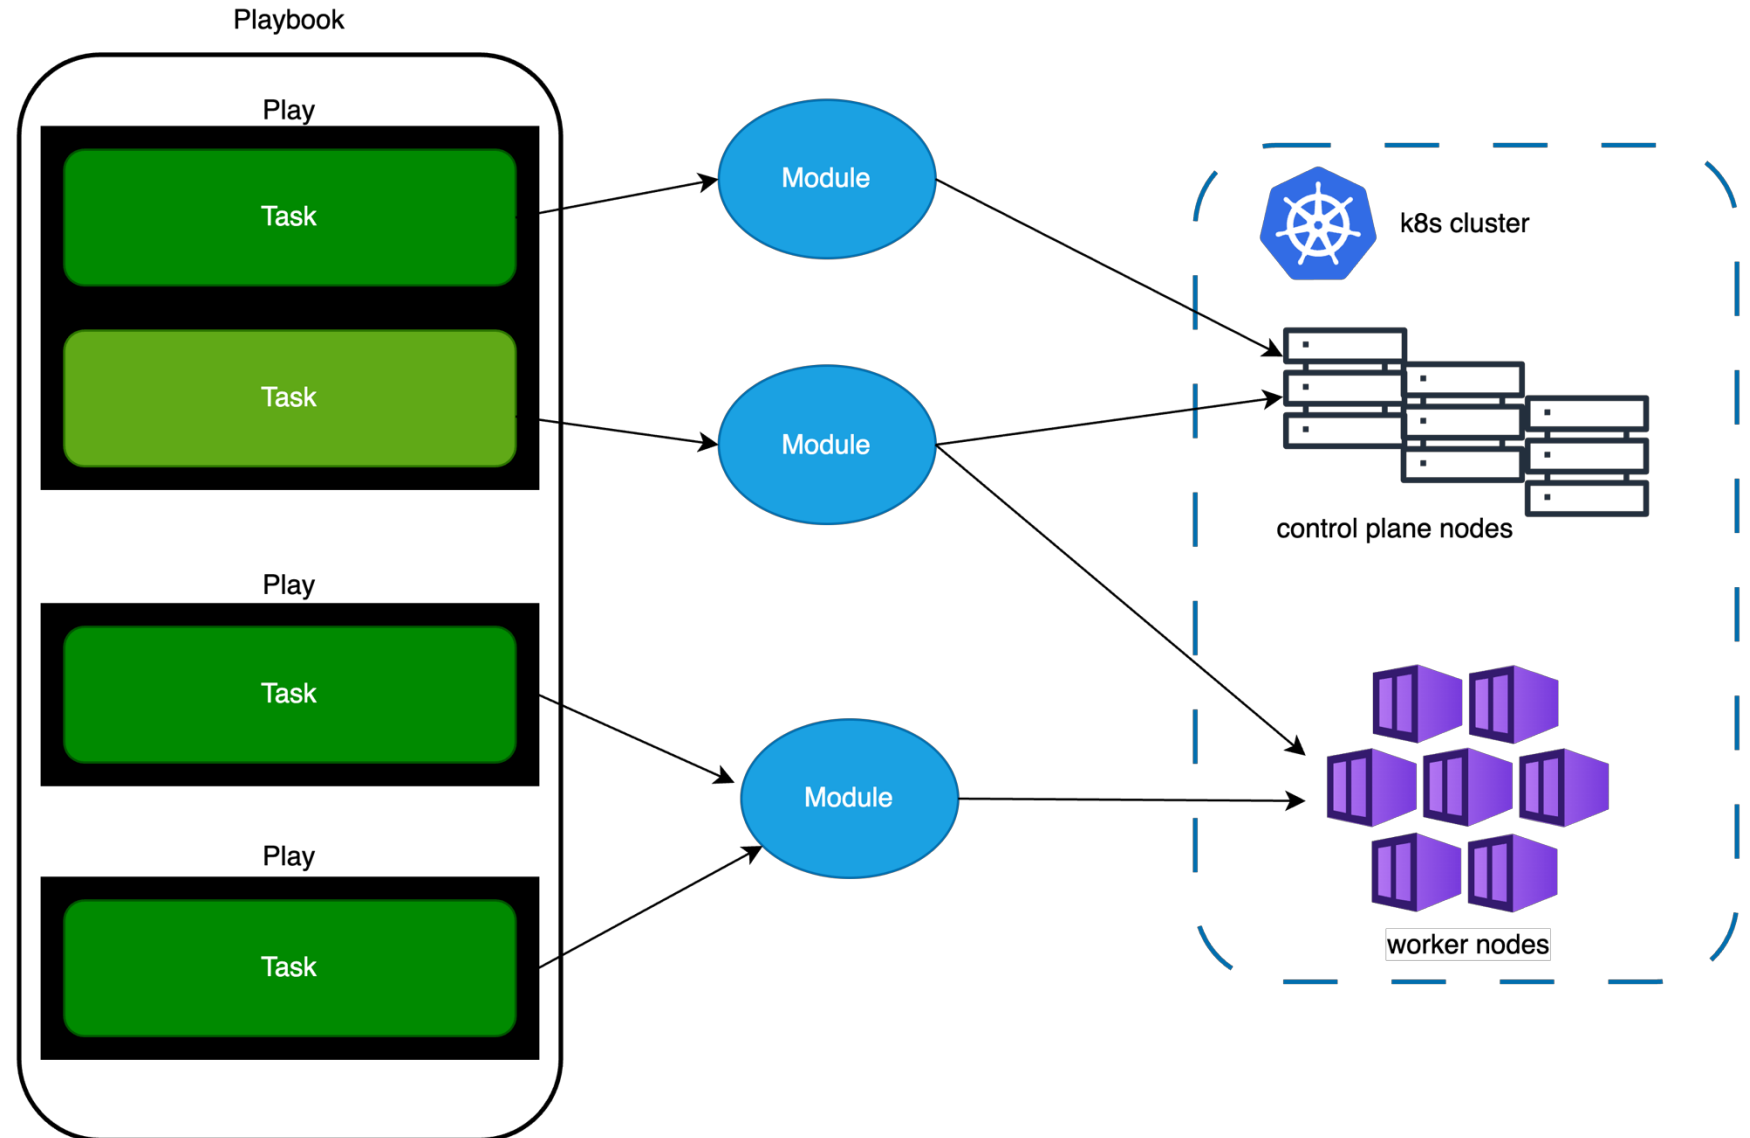
- Upgrade the primary control plane (rke2-server).
- Upgrade additional control plane (rke2-server).
- Upgrade worker nodes (rke2-agent).



K8s rolling upgrade workflow

Automating k8s rolling upgrades

Ansible playbook as a list of tasks to translate the upgrade workflow in desired states



Sharing the DSI Ansible playbook

Released for the INFN community in the [sysinfo_oss baltig project](#).

Engineered for a RKE2 HA cluster setup described by a proper inventory file.

Support to RKE2 rpms upgrade or delegate to puppet agent run.

Sequentially ssh into cluster nodes to exec tasks as sudo user.

Leverage RKE2 utilities (eg. kubectl bin, kubeconfig, rke2 etcd snapshot)

Upgrade all OS packages and reboot node if needs-restarting

DEMO

