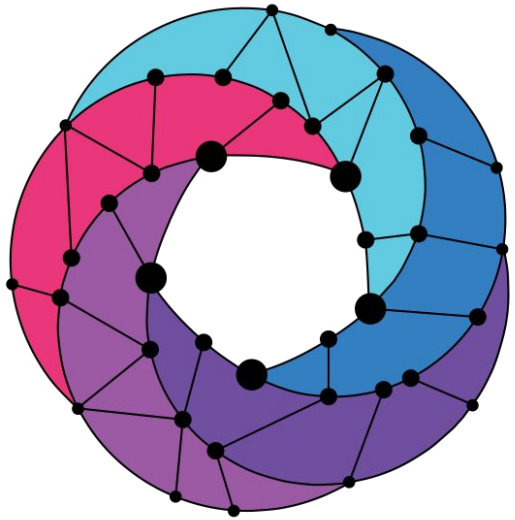

Open Source MANO

Day 2

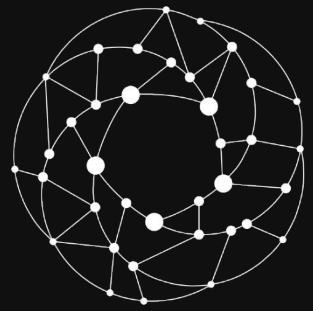

Open Source
MANO

What is a network function, VNF, CNF

Maciej Mazur (Canonical)

ETSI

Telco network is complex

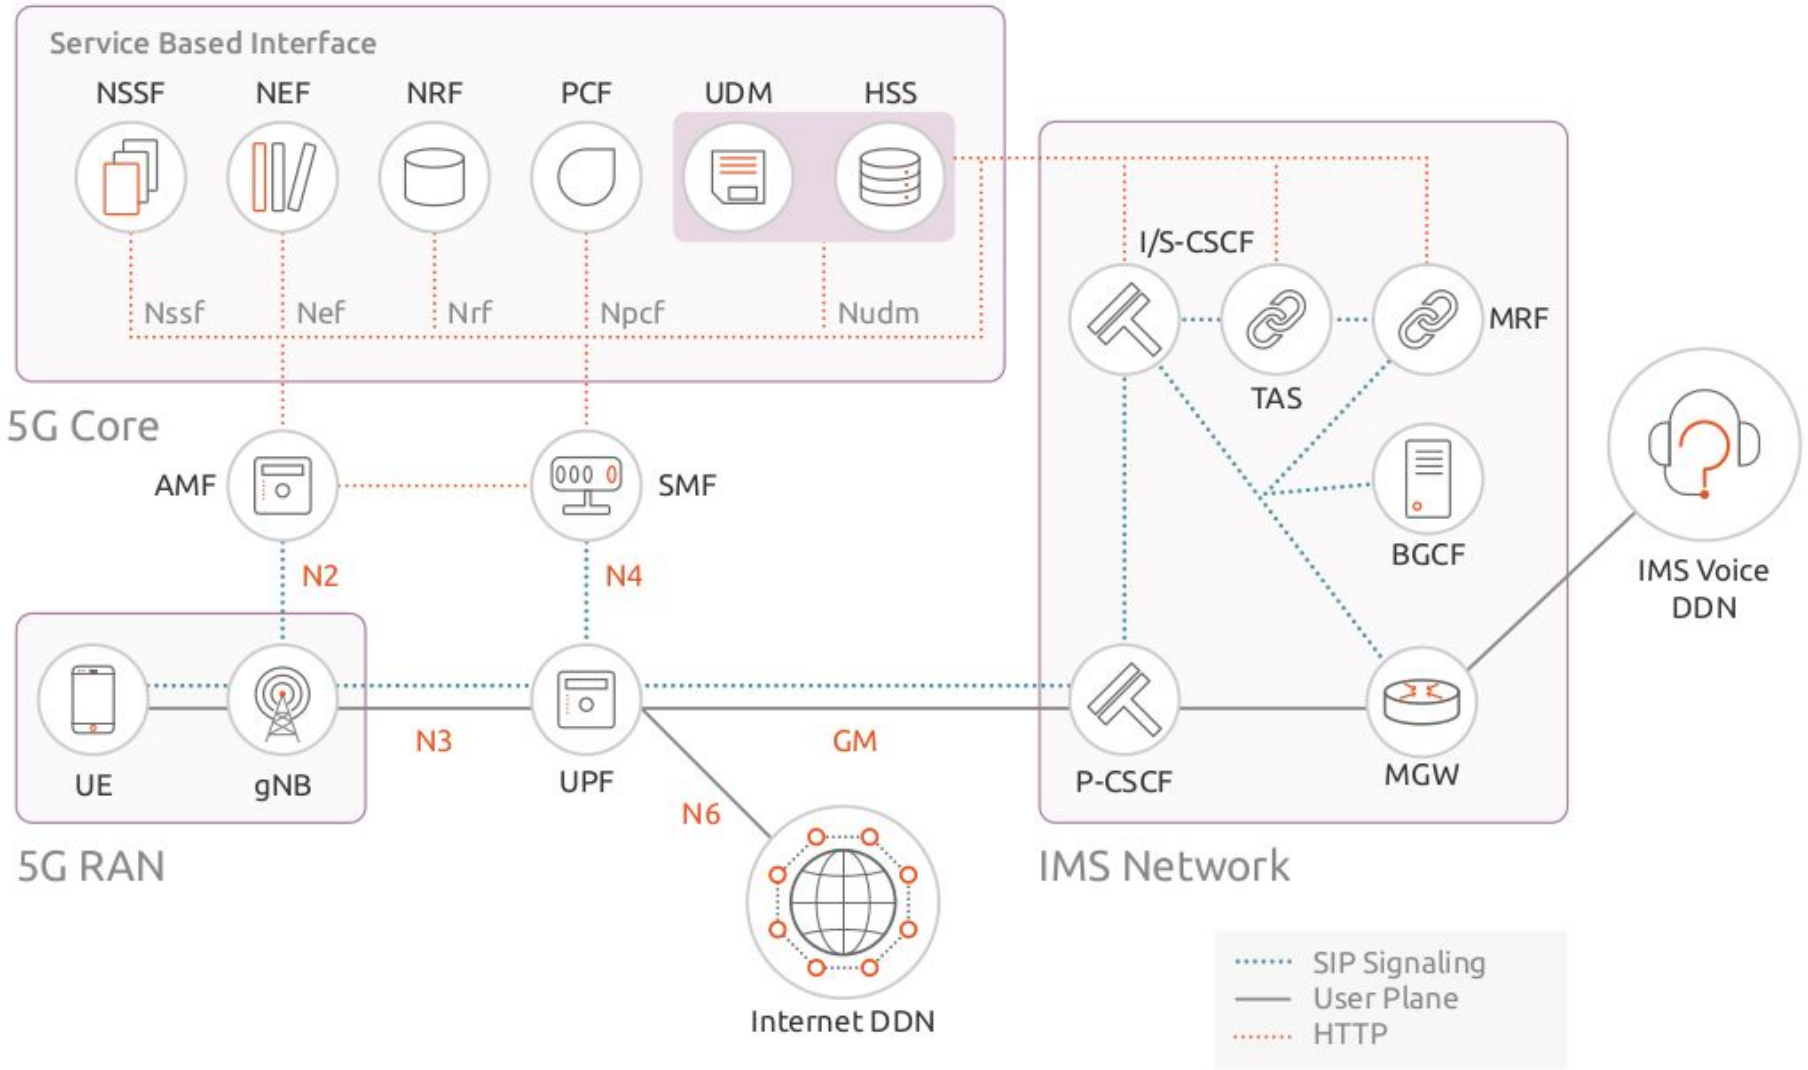

What is NFV trying to address?

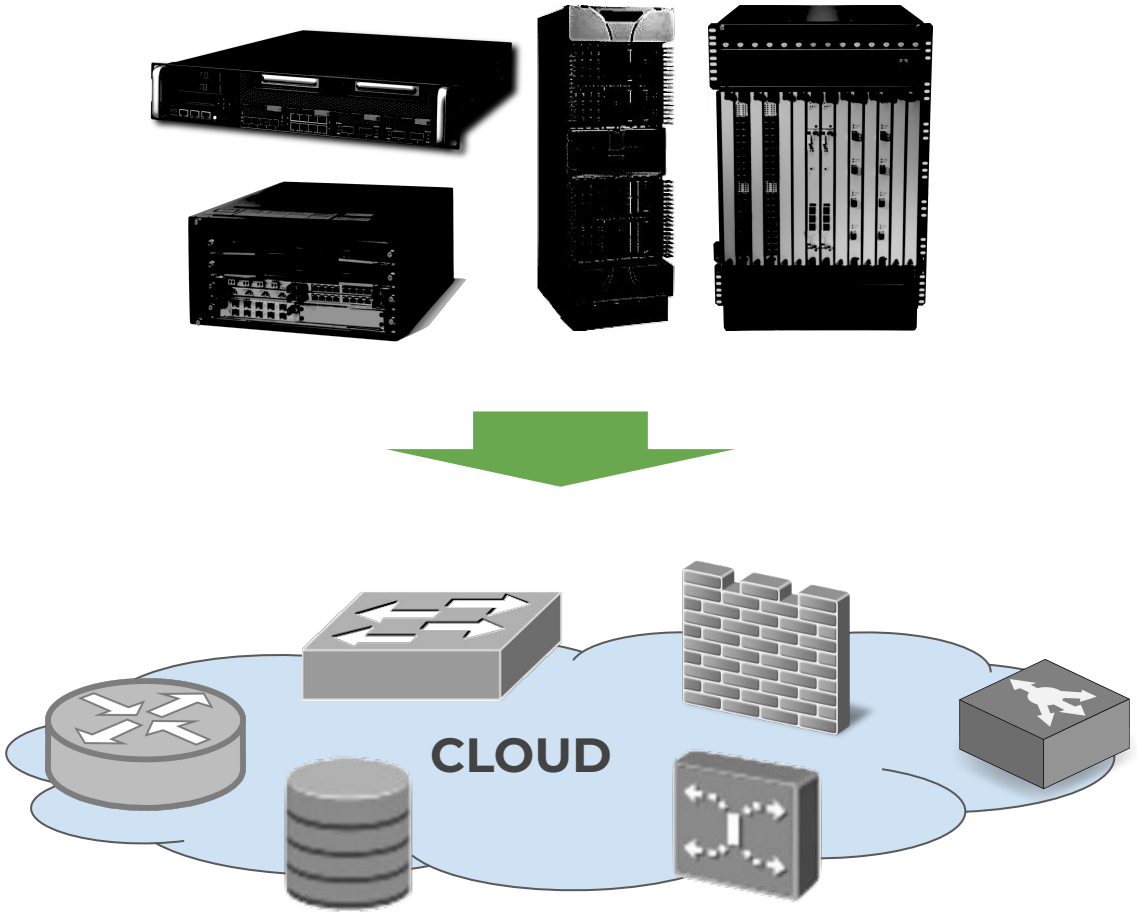

NFV proposes to **virtualize, over commodity hardware, network functions that typically run in dedicated appliances, why?**

- To make network operations more agile and cost-efficient.
- To increase independency of hardware vendors.
- To leverage all the advantages of the Cloud, for network functions.

The original idea triggered an industry movement

Open Source MANO

- Initial white paper was written in 2012 by the world's leading telecom network operators (Europe, América & Asia).
- This group evolved to the ETSI NFV ISG (Industry Specification Group), formed by 300+ companies.
- Their main motivation had to do with reducing TCO of building a network by using open solutions.

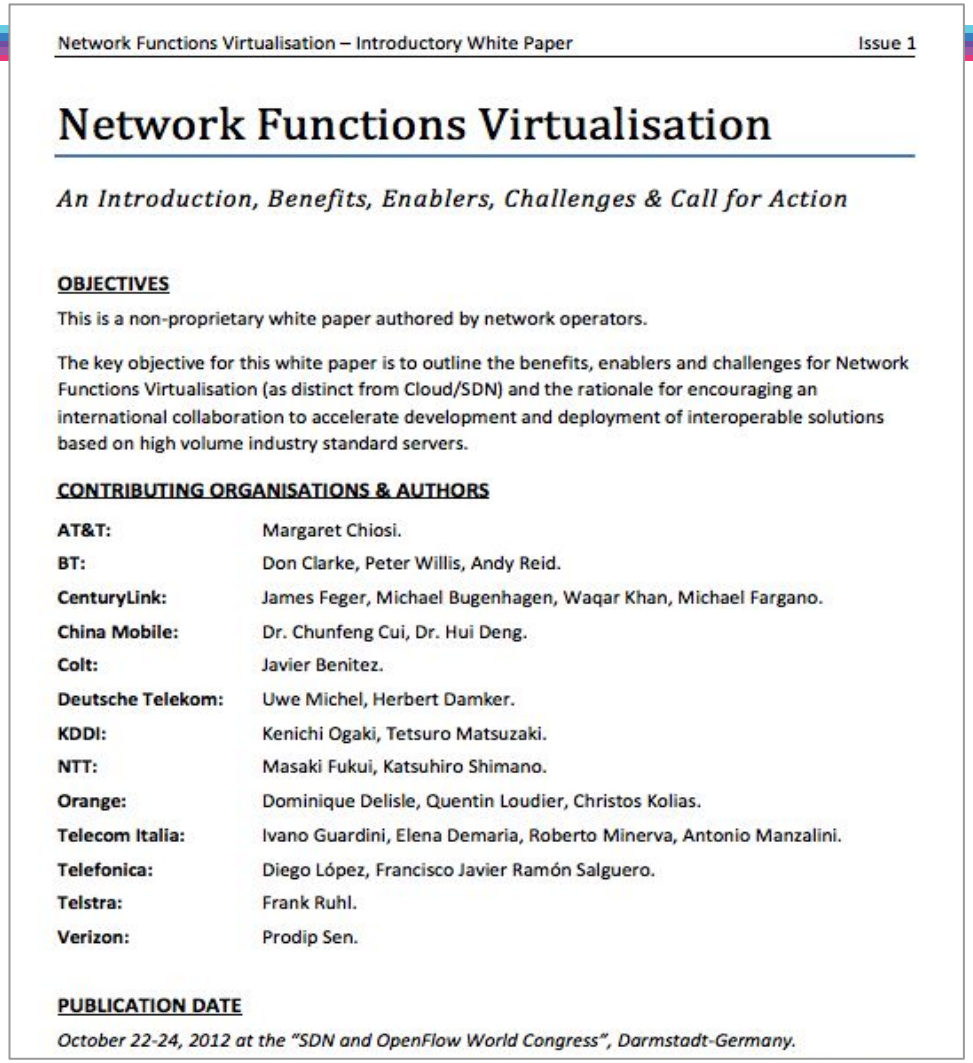

Network Functions Virtualisation – Introductory White Paper Issue 1

Network Functions Virtualisation

An Introduction, Benefits, Enablers, Challenges & Call for Action

OBJECTIVES

This is a non-proprietary white paper authored by network operators.

The key objective for this white paper is to outline the benefits, enablers and challenges for Network Functions Virtualisation (as distinct from Cloud/SDN) and the rationale for encouraging an international collaboration to accelerate development and deployment of interoperable solutions based on high volume industry standard servers.

CONTRIBUTING ORGANISATIONS & AUTHORS

AT&T:	Margaret Chiosi.
BT:	Don Clarke, Peter Willis, Andy Reid.
CenturyLink:	James Feger, Michael Bugenhagen, Waqar Khan, Michael Fargano.
China Mobile:	Dr. Chunfeng Cui, Dr. Hui Deng.
Colt:	Javier Benitez.
Deutsche Telekom:	Uwe Michel, Herbert Damker.
KDDI:	Kenichi Ogaki, Tetsuro Matsuzaki.
NTT:	Masaki Fukui, Katsuhiro Shimano.
Orange:	Dominique Delisle, Quentin Loudier, Christos Kolias.
Telecom Italia:	Ivano Guardini, Elena Demaria, Roberto Minerva, Antonio Manzalini.
Telefonica:	Diego López, Francisco Javier Ramón Salguero.
Telstra:	Frank Ruhl.
Verizon:	Prodip Sen.

PUBLICATION DATE

October 22-24, 2012 at the "SDN and OpenFlow World Congress", Darmstadt-Germany.

https://portal.etsi.org/nfv/nfv_white_paper.pdf

ETSI Publications

- Based on member's feedback, field experiences and proof of concepts, standard documents have evolved.
- 60+ publications exist today, including the following three main documents:
  - NFV Architectural Framework
    http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.02.01_60/gs_NFV002v010201p.pdf
  - NFV Infrastructure Overview
    http://www.etsi.org/deliver/etsi_gs/NFV-INF/001_099/001/01.01.01_60/gs_NFV-INF001v010101p.pdf
  - NFV Management and Orchestration
    http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.02.01_60/gs_NFV002v010201p.pdf

http://www.etsi.org/standards-search

Benefits of a standard NFV architecture

Open Source MANO

The ultimate goal is for operators to have a **unified and generic virtualization infrastructure**, compatible with any vendor's Virtual Networking Function (VNF), **this makes standardization a must**.

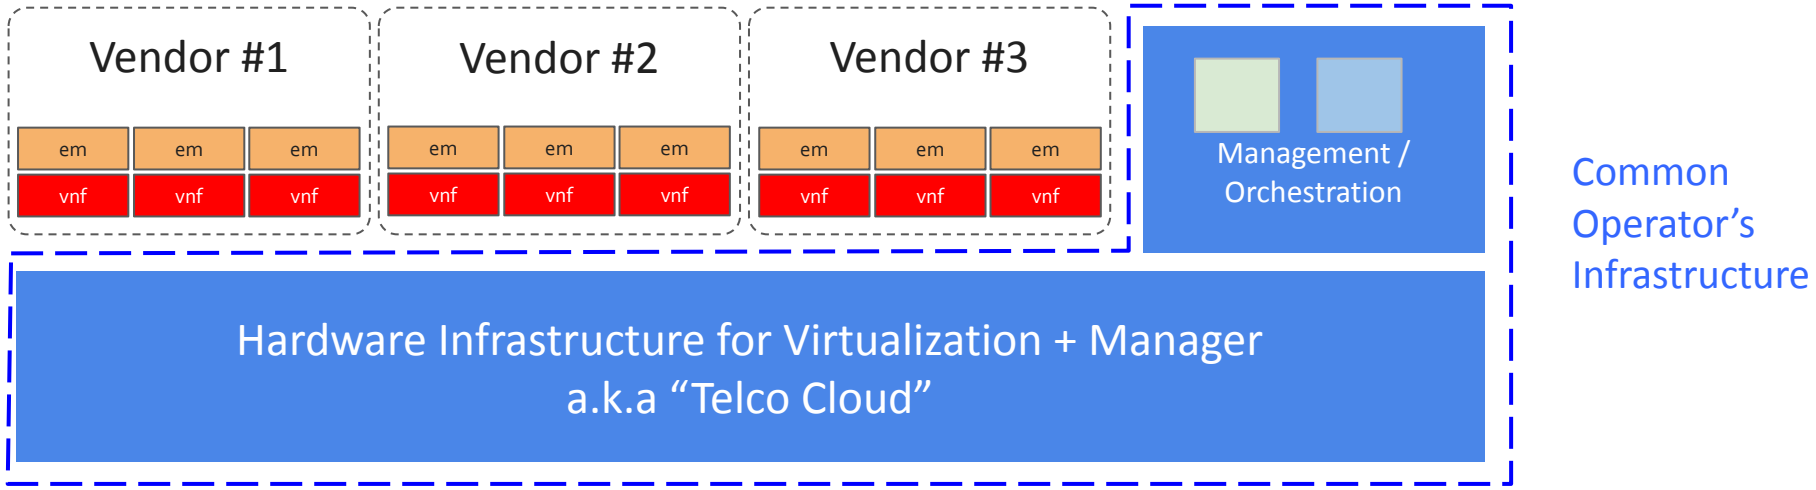

The ETSI NFV Architectural Framework

The standard architecture can be better understood in three blocks:

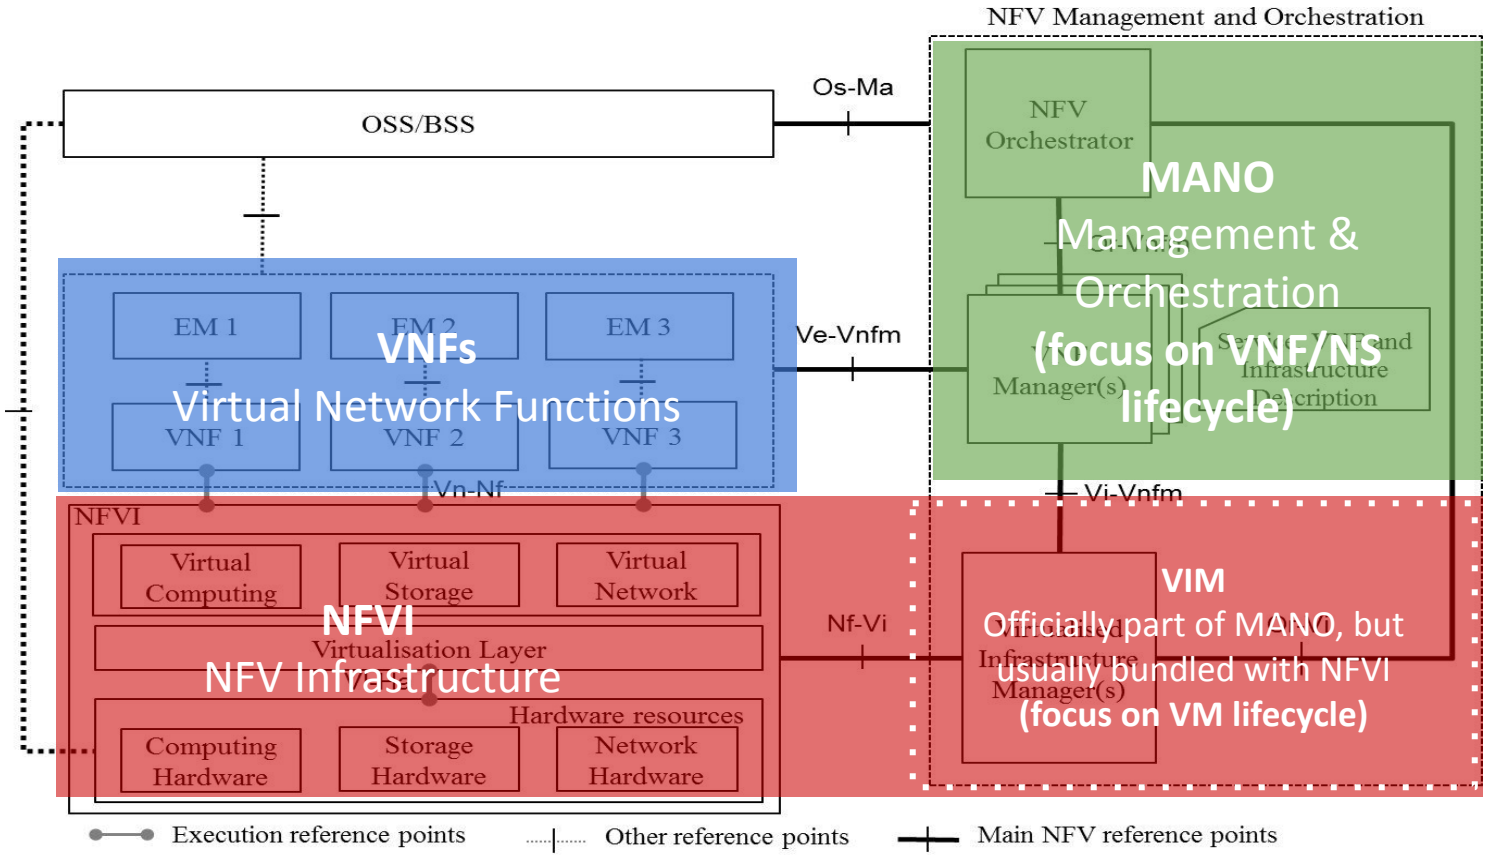

NFVI: NFV Infrastructure

NFVI goal is to provide a virtualization environment for VNFs, including virtual compute, storage and networking resources.

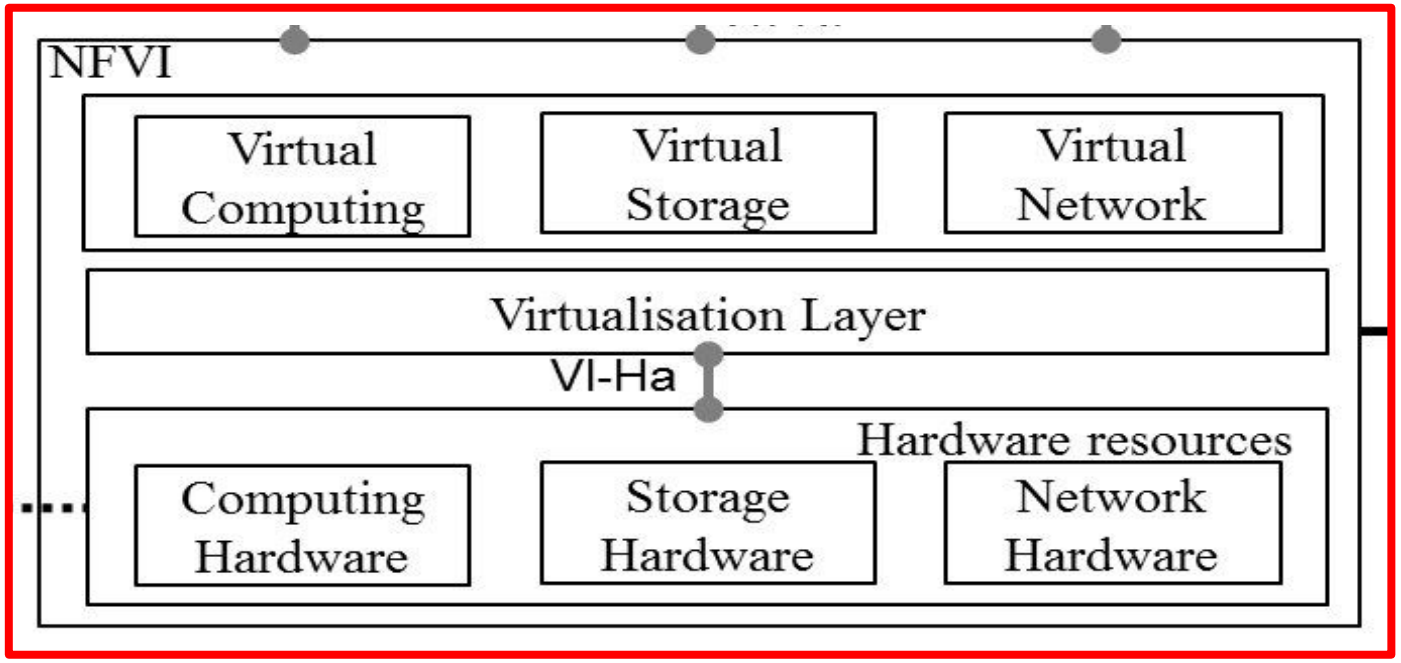

VNF Special Requirements

VNFs, especially data-plane ones, usually have additional requirements than common cloud applications, including:

- **Minor latency** (disk I/O & network)
  → Faster hardware (More cores, SSD disks, faster buses)
  → Dataplane acceleration
- **Higher throughput or PPS**
  → Dataplane acceleration
  → EPA: Enhanced Platform Awareness
- **Geographical distribution**
  → multi-site cloud
- **Horizontal auto-scaling**
  → automated operations (orchestration)

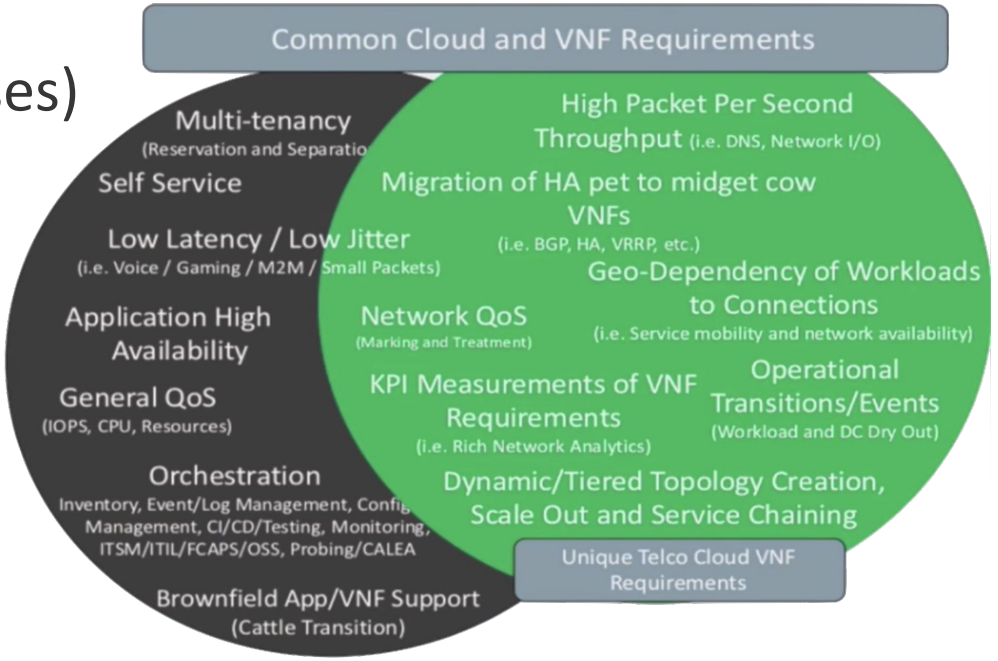

OpenStack Austin 2016: Telco Cloud Requirements: What VNF's Are Asking For

NFVI: NFV Infrastructure

VNF Special Requirements

EPA covers the different approaches that can be taken at the NFVI layer to increase performance while maintaining a generic (COTS) infrastructure. VIM and MANO should be able to request them.

Huge Pages

NUMA Topology Awareness

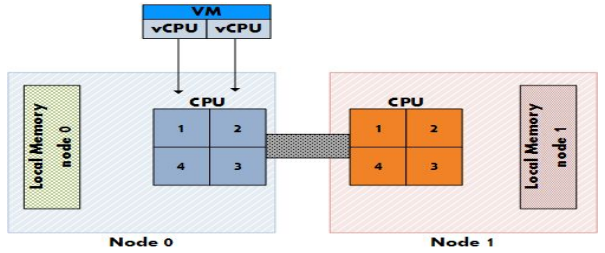

CPU Pinning

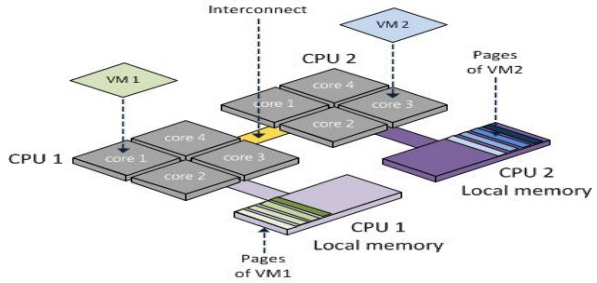

Data Plane assignment

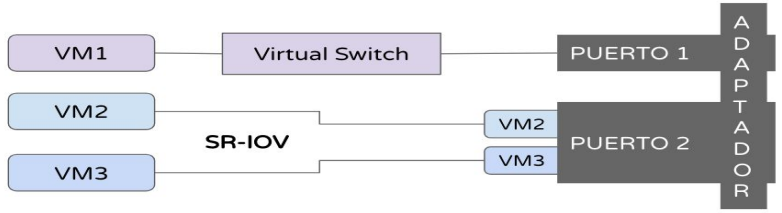

VIM: Virtualized Infrastructure Manager

The Virtualised Infrastructure Manager conceptually part of the ‘MANO Stack’, provides lifecycle management for virtualized resources (VMs, volumes, networking paths and connectivity, etc.)

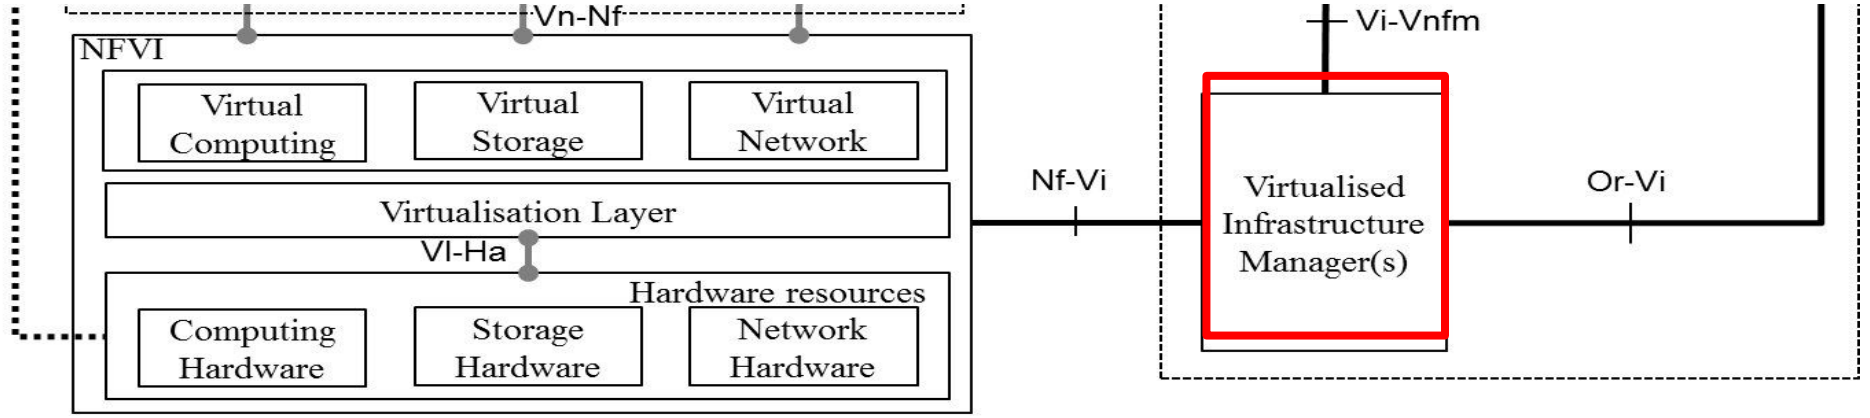

MANO: VNF Manager (VNFM)

- The VNF Manager, also part of the 'MANO Stack', covers **lifecycle management for Virtual Network Functions** (VNFs), either directly or through their own Element Management System (EMS).

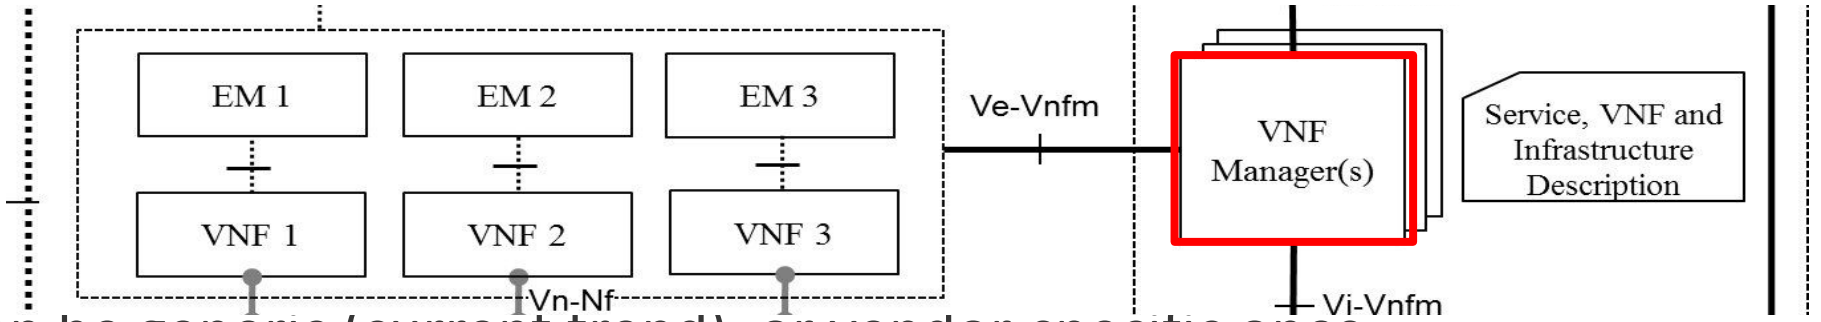

VNF Managers can be generic (current trend), or vendor-specific ones.

MANO: NFV Orchestrator (NFVO)

- The NFV Orchestrator, the higher entity in the ‘MANO Stack’, covers general resource orchestration and services lifecycle, which comprise multiple VNFs and define their roles (traffic paths, scaling decisions, and other service-related requirements)
- It can interact with a generic VNF Manager, or vendor-specific ones.

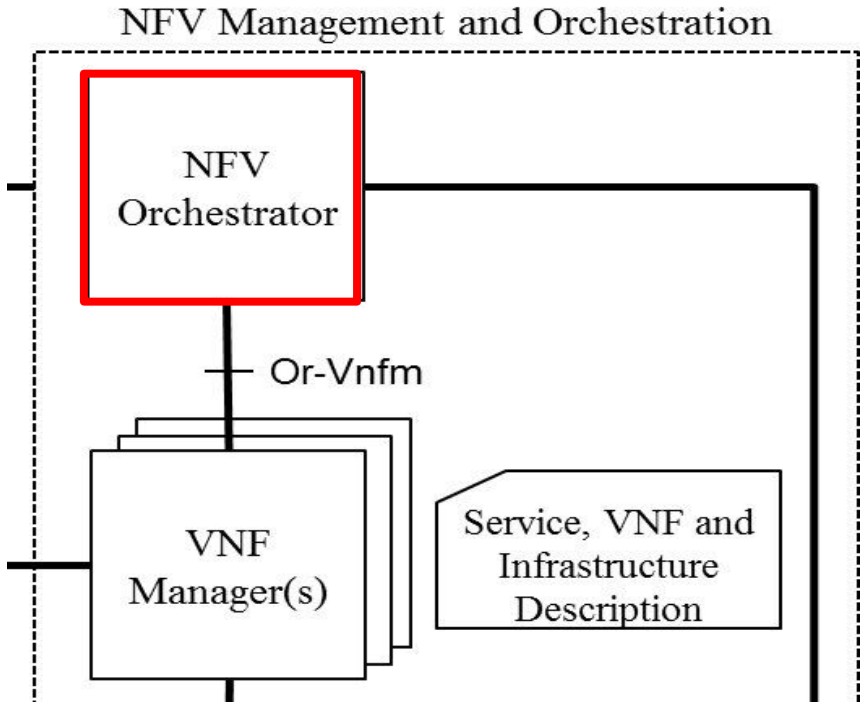

Virtual Network Functions (VNF)

- Finally, the VNFs, which are supported by the underlying NFVI, and managed by their own EM (internal, element manager) and the VNF Manager (external, 'context-aware' manager)
- They should be able to provide any networking function and interact with other VNFs.

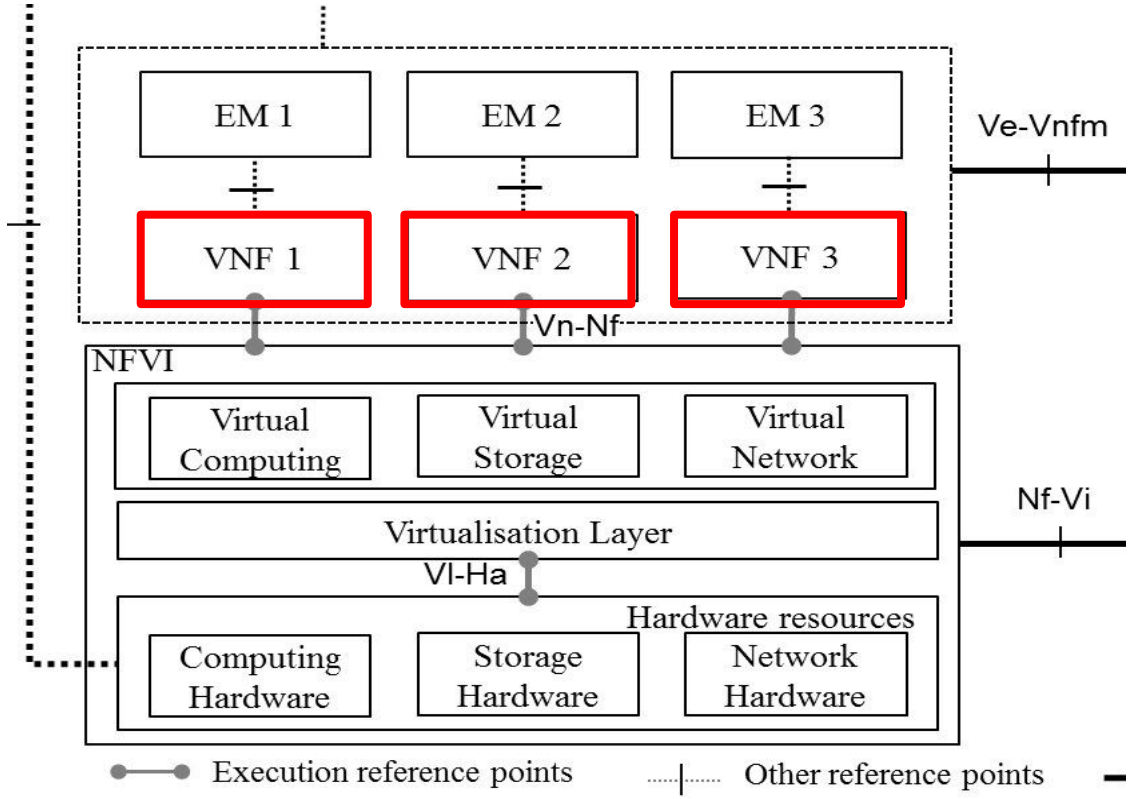

The Open Source MANO Project

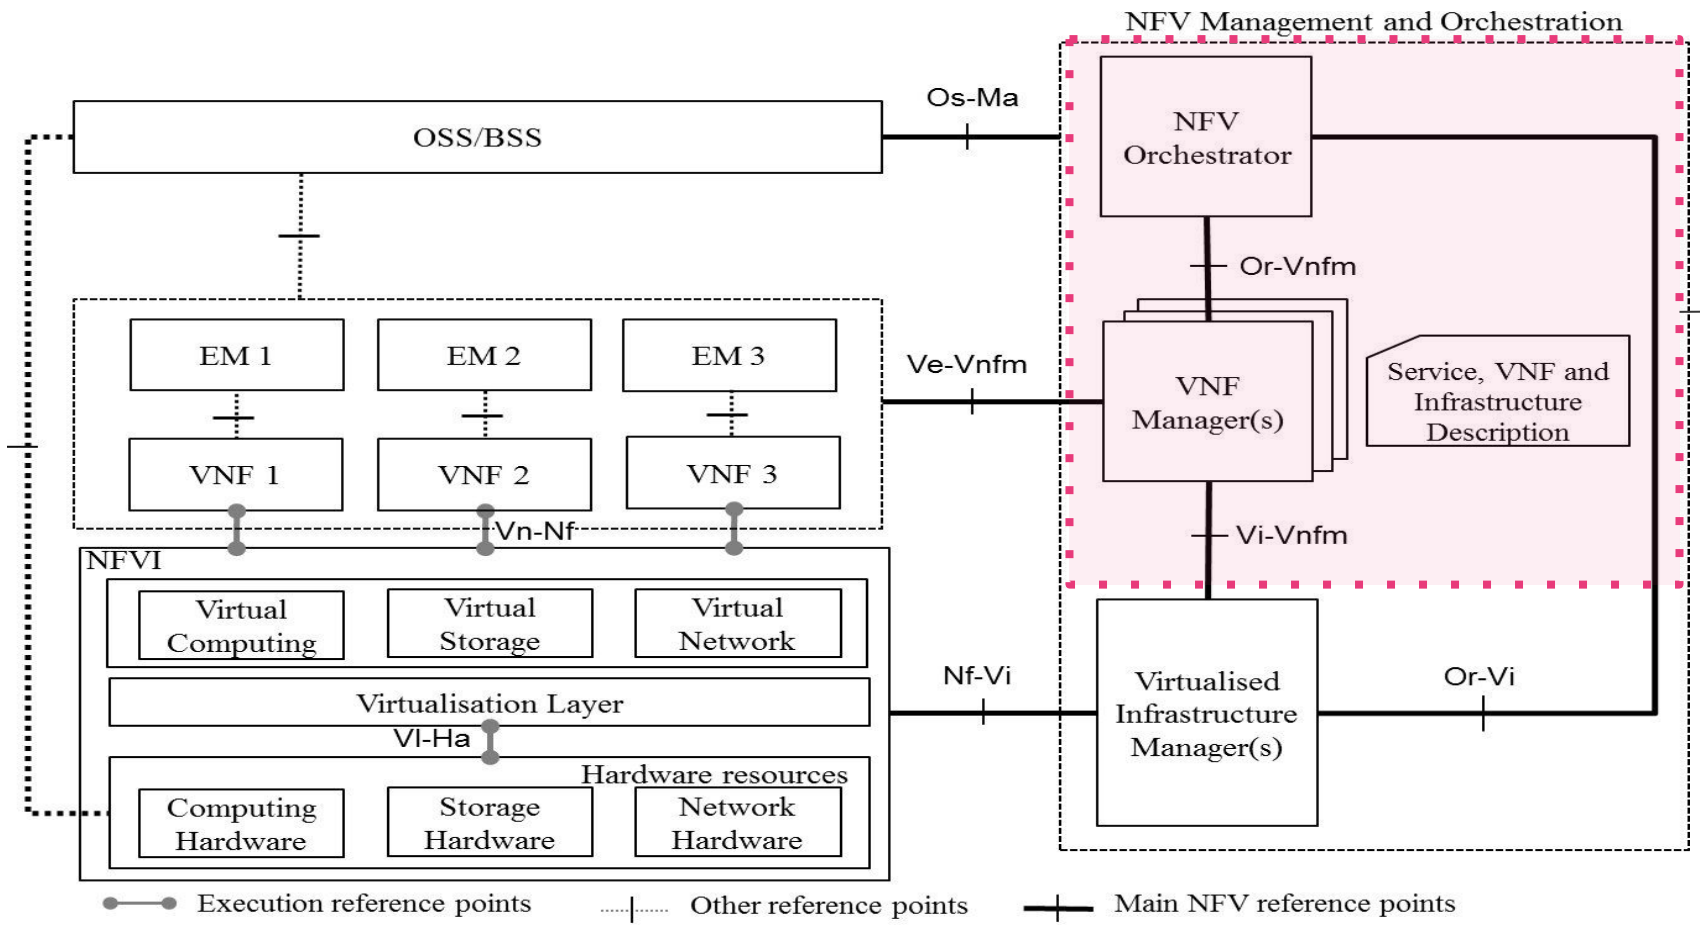

We are here!
Open Source MANO is an ETSI-hosted project developing an Open Source NFV Management and Orchestration (MANO) software stack aligned with ETSI NFV.

The Software-Defined DataCenter

Software defined data centers are composed of completely virtualized infrastructure, that can be easily managed using software.
When talking about SDDC, we are mainly referencing **virtualization of compute, storage and networking**, being that all the infrastructure is totally programmable.

This new data center building paradigm focuses on:

- Using the full potential of hyper scalable architectures
- More agility and faster application provisioning
- Cost reduction
- Automated management (elastic and programmable)

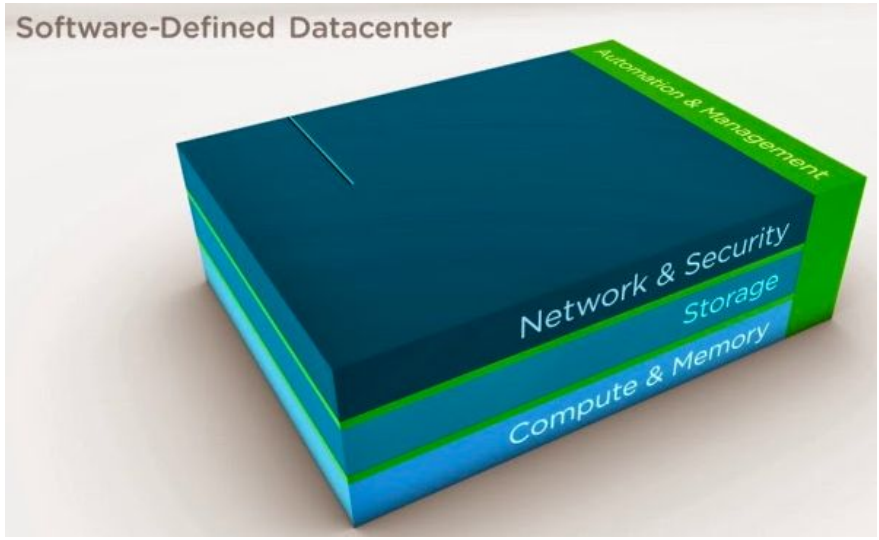

Virtual Machines and Containers

With hypervisors, like KVM, we are capable of partitioning a physical compute node into multiple **“virtual machines”** that use their own Operating System to share the physical resources, providing efficiency on the host resources consumption.

“Containers”, made popular by Docker, produce instances that share a single Operating System while only adding the libraries they need to run a lightweight application. This concept increases application mobility and takes the efficiency on host resource consumption a step further, simplifying modularity to the “micro-services” level.

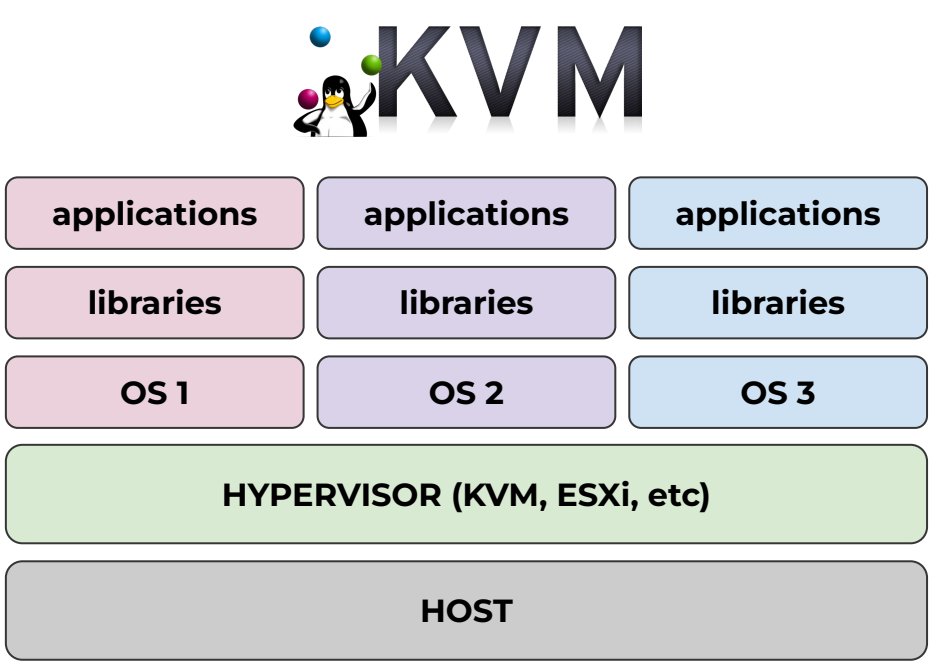

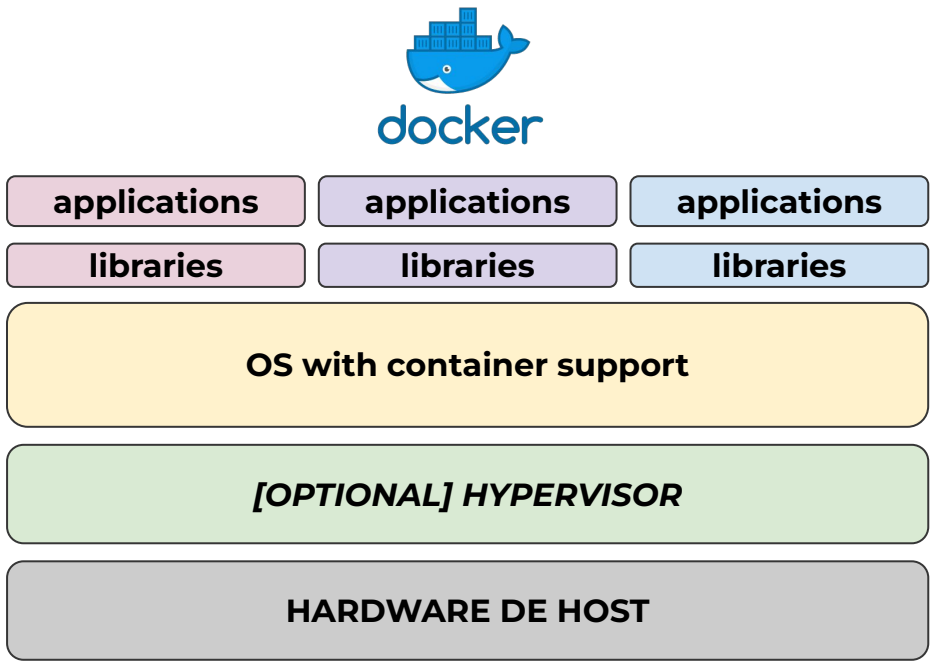

Beginnings were hard

"I think there's a world market for maybe five computers"

Thomas Watson, IBM, 1943

OpenStack Timeline

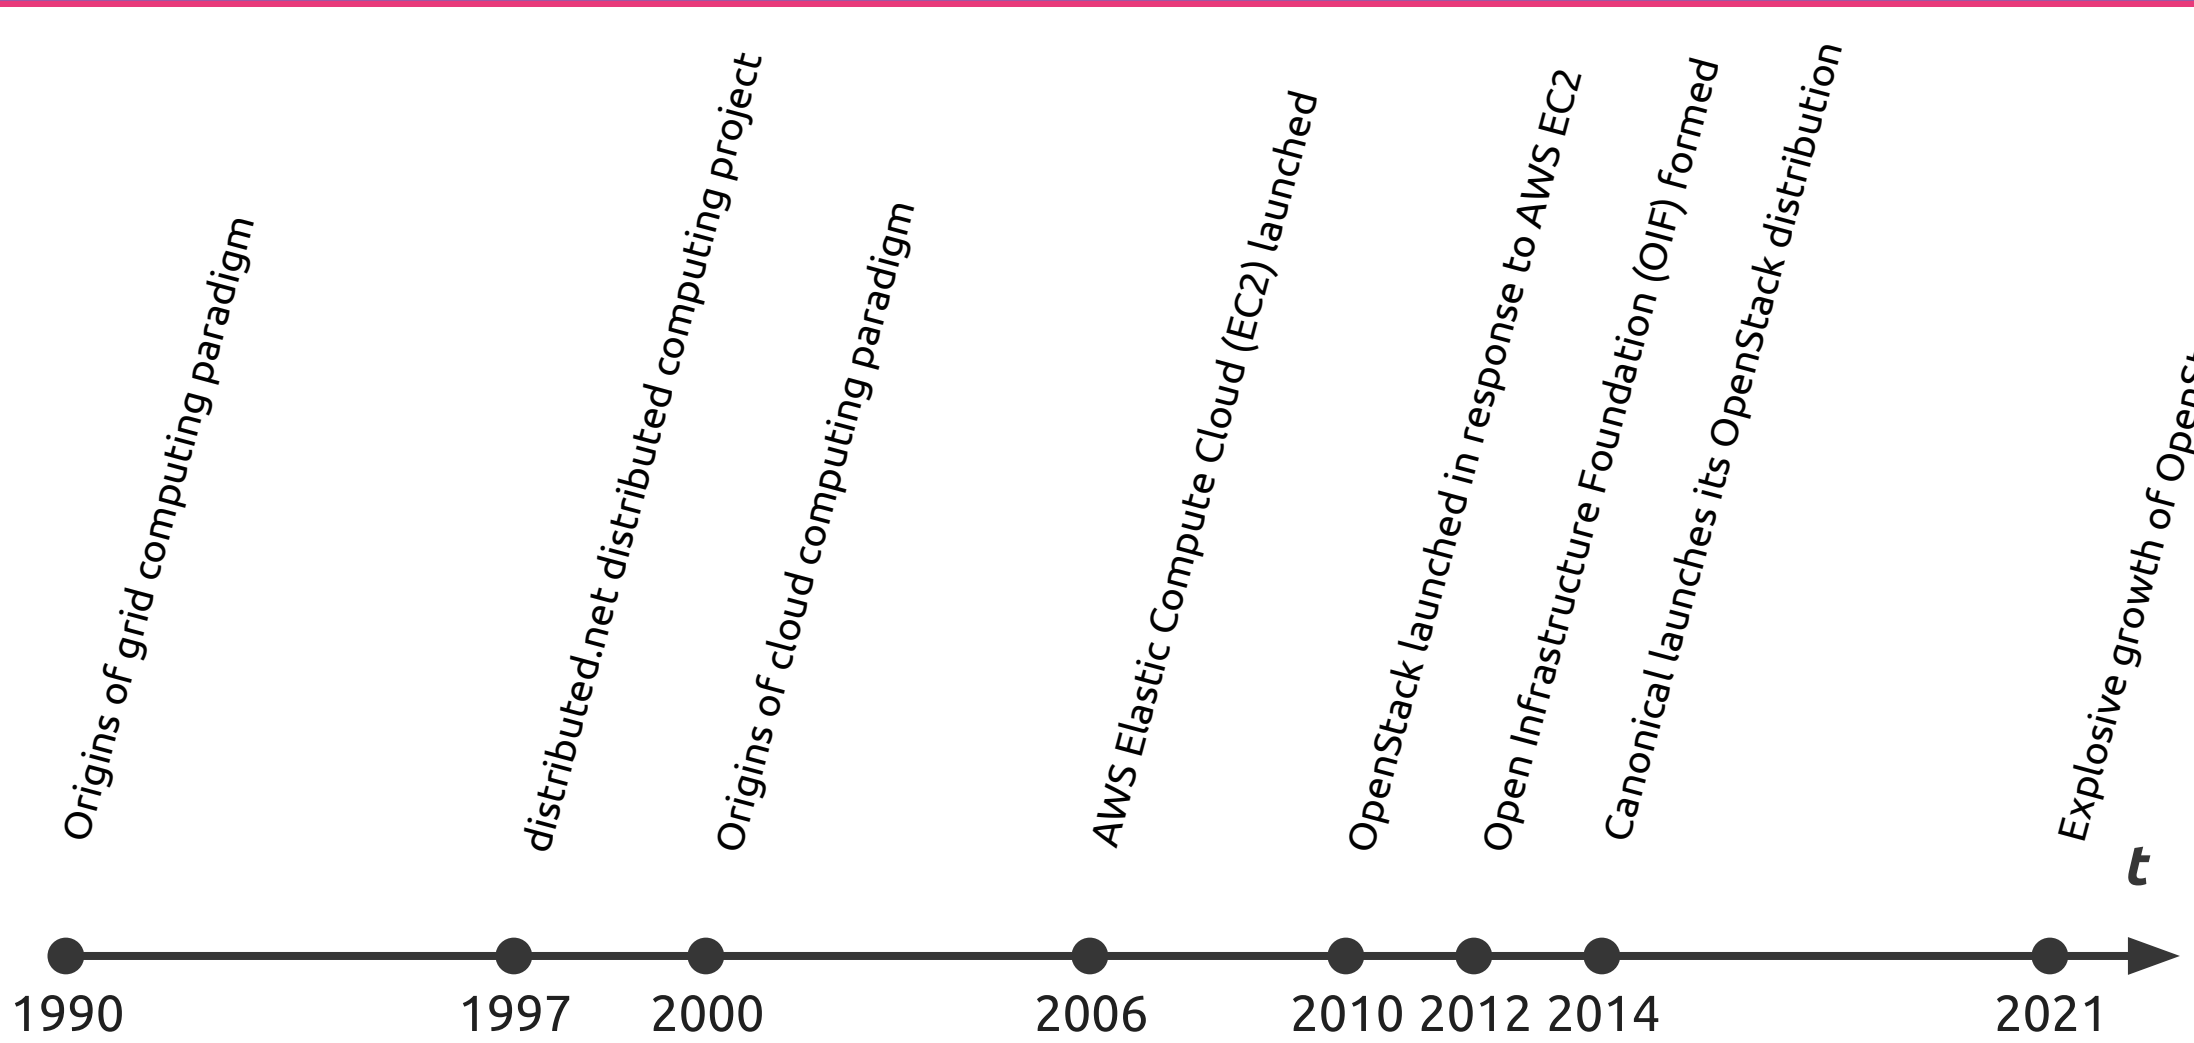

OpenStack vs typical virtualization platforms

	Virtualisation Management	OpenStack
TCO	High	Low
Resource management	“Pets”	“Cattle”
Resource provisioning	Image-based	Template-based
Hardware	Specialised	Commodity
Scalability	Scale up	Scale out
Workload	Stateful	Stateless
Open source	Mostly no	Yes

What is OpenStack?

Linux of infra

OpenStack is an open source cloud platform

just like Linux is an open source operating system

Big community

Open Source MANO

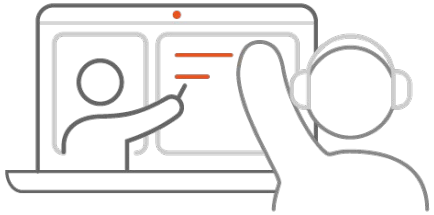

Industry events

~450 organisations

openstack®

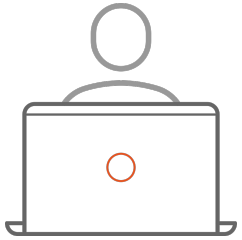

~8,500 developers

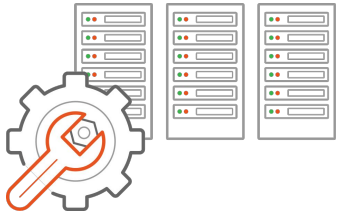

Admin’s guide

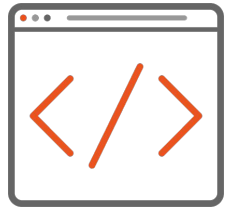

Developer’s guide

Modular

OpenStack has a modular architecture and consists of the following components:

- OpenStack services - expose API endpoints and handle basic cloud functions
- OpenStack dashboard - provides a web-based user interface
- OpenStack client - provides a command-line user interface
- SQL databases - store various records created by OpenStack services
- Message queues - facilitate inter-process communication
- Additional components - NoSQL databases, memcached, etc.

Terms

When in OpenStack, speak as the OpenStackers speak :)

- Tenant - a user or a group of users with isolated cloud resources
- Image - a template containing an OS used for instance provisioning
- Flavor - a template defining cloud resources for instance provisioning
- Instance - a VM provisioned from the image and the flavor
- Ephemeral storage - volatile storage attached to the instance
- Block storage - non-volatile storage attached to the instance
- Object storage - non-volatile, cloud-native storage accessible through the API
- Security group - virtual firewall attached to the instance

OpenStack HLD

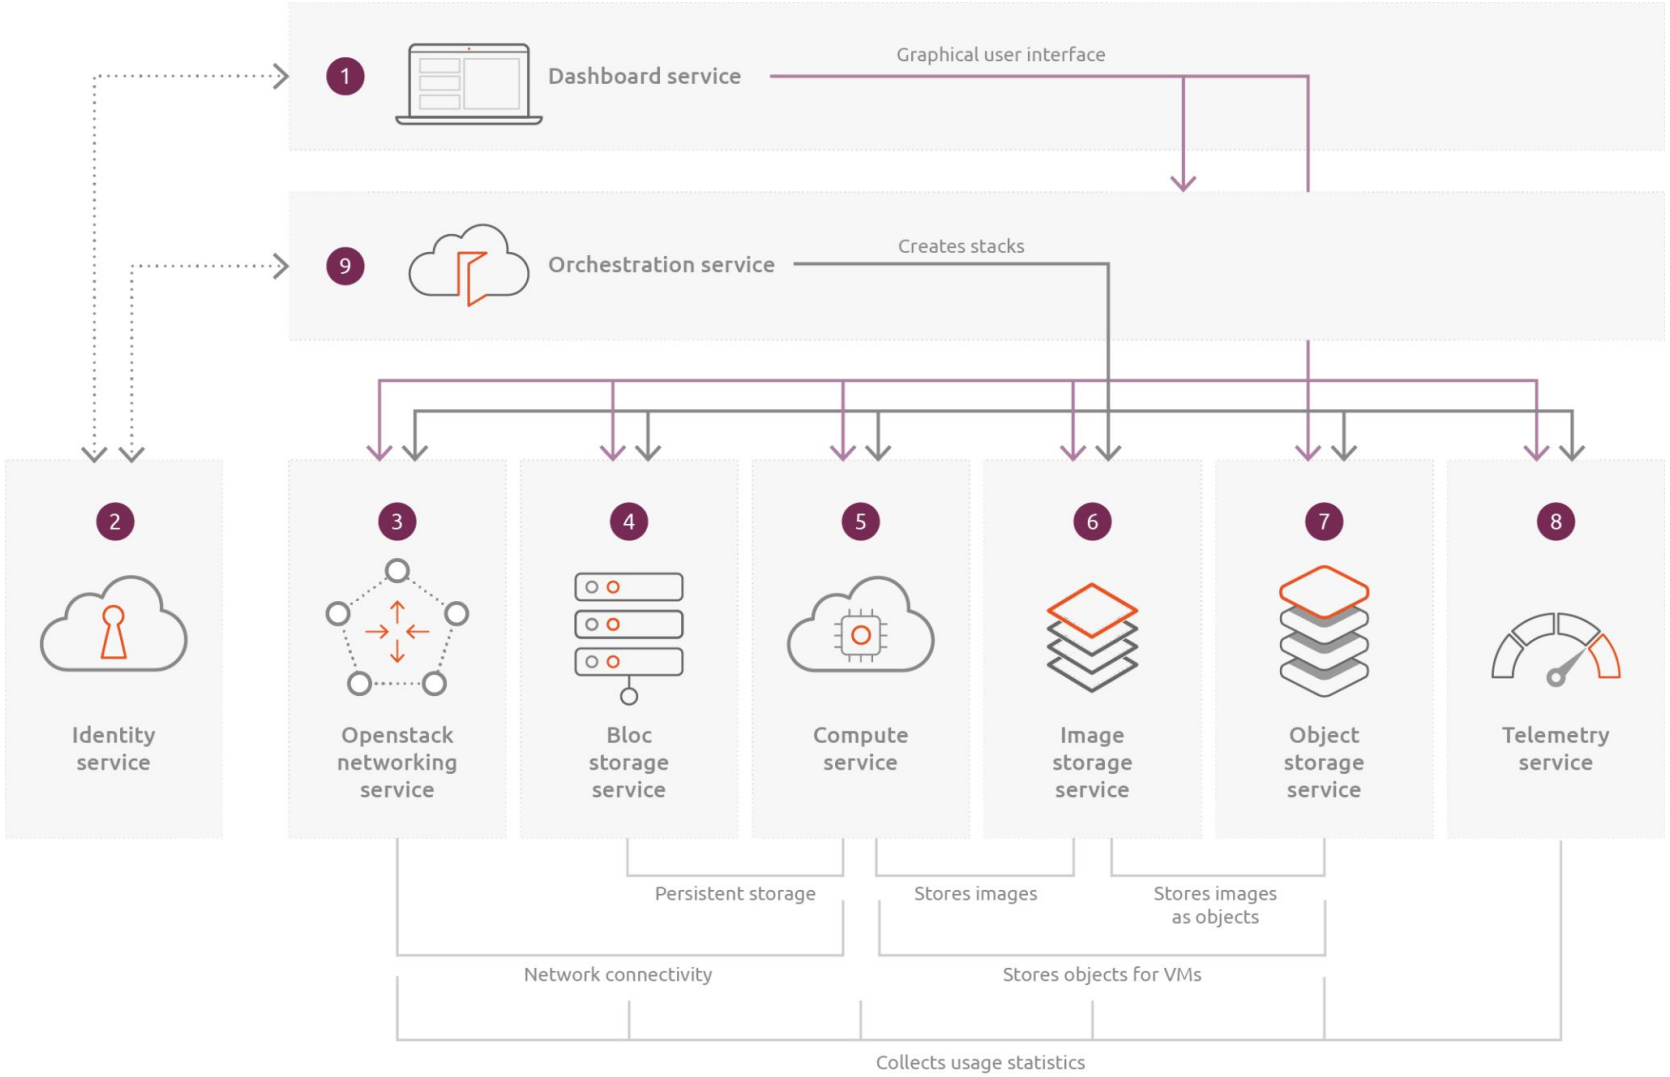

Keystone

- Identity service
- Manages domains, projects, roles, groups and user accounts
- Provides authentication and authorisation functions
- Integrates with LDAP, Active Directory, SAML, etc.

Glance

- Image service
- Manages the catalogue of cloud images
- Can also build and validate images

Nova

- Compute service
- Responsible for instance provisioning, scheduling and termination
- Also manages ephemeral storage devices
- Uses KVM or other hypervisor underneath

Neutron

- Network service
- Manages virtual networks, subnets, routers and security groups
- Uses OVN/OVS or other SDN underneath

Cinder

- Block storage service
- Manages block storage volumes, their snapshots and backups
- Integrates with a variety of storage platforms, including Ceph, LVM, NetApp, PureStorage, etc.

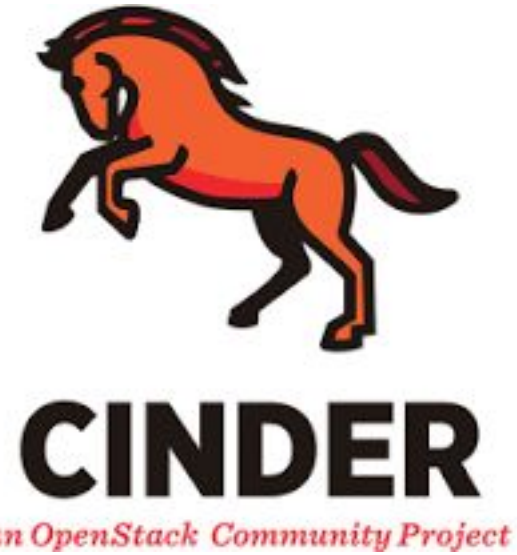

Many other services

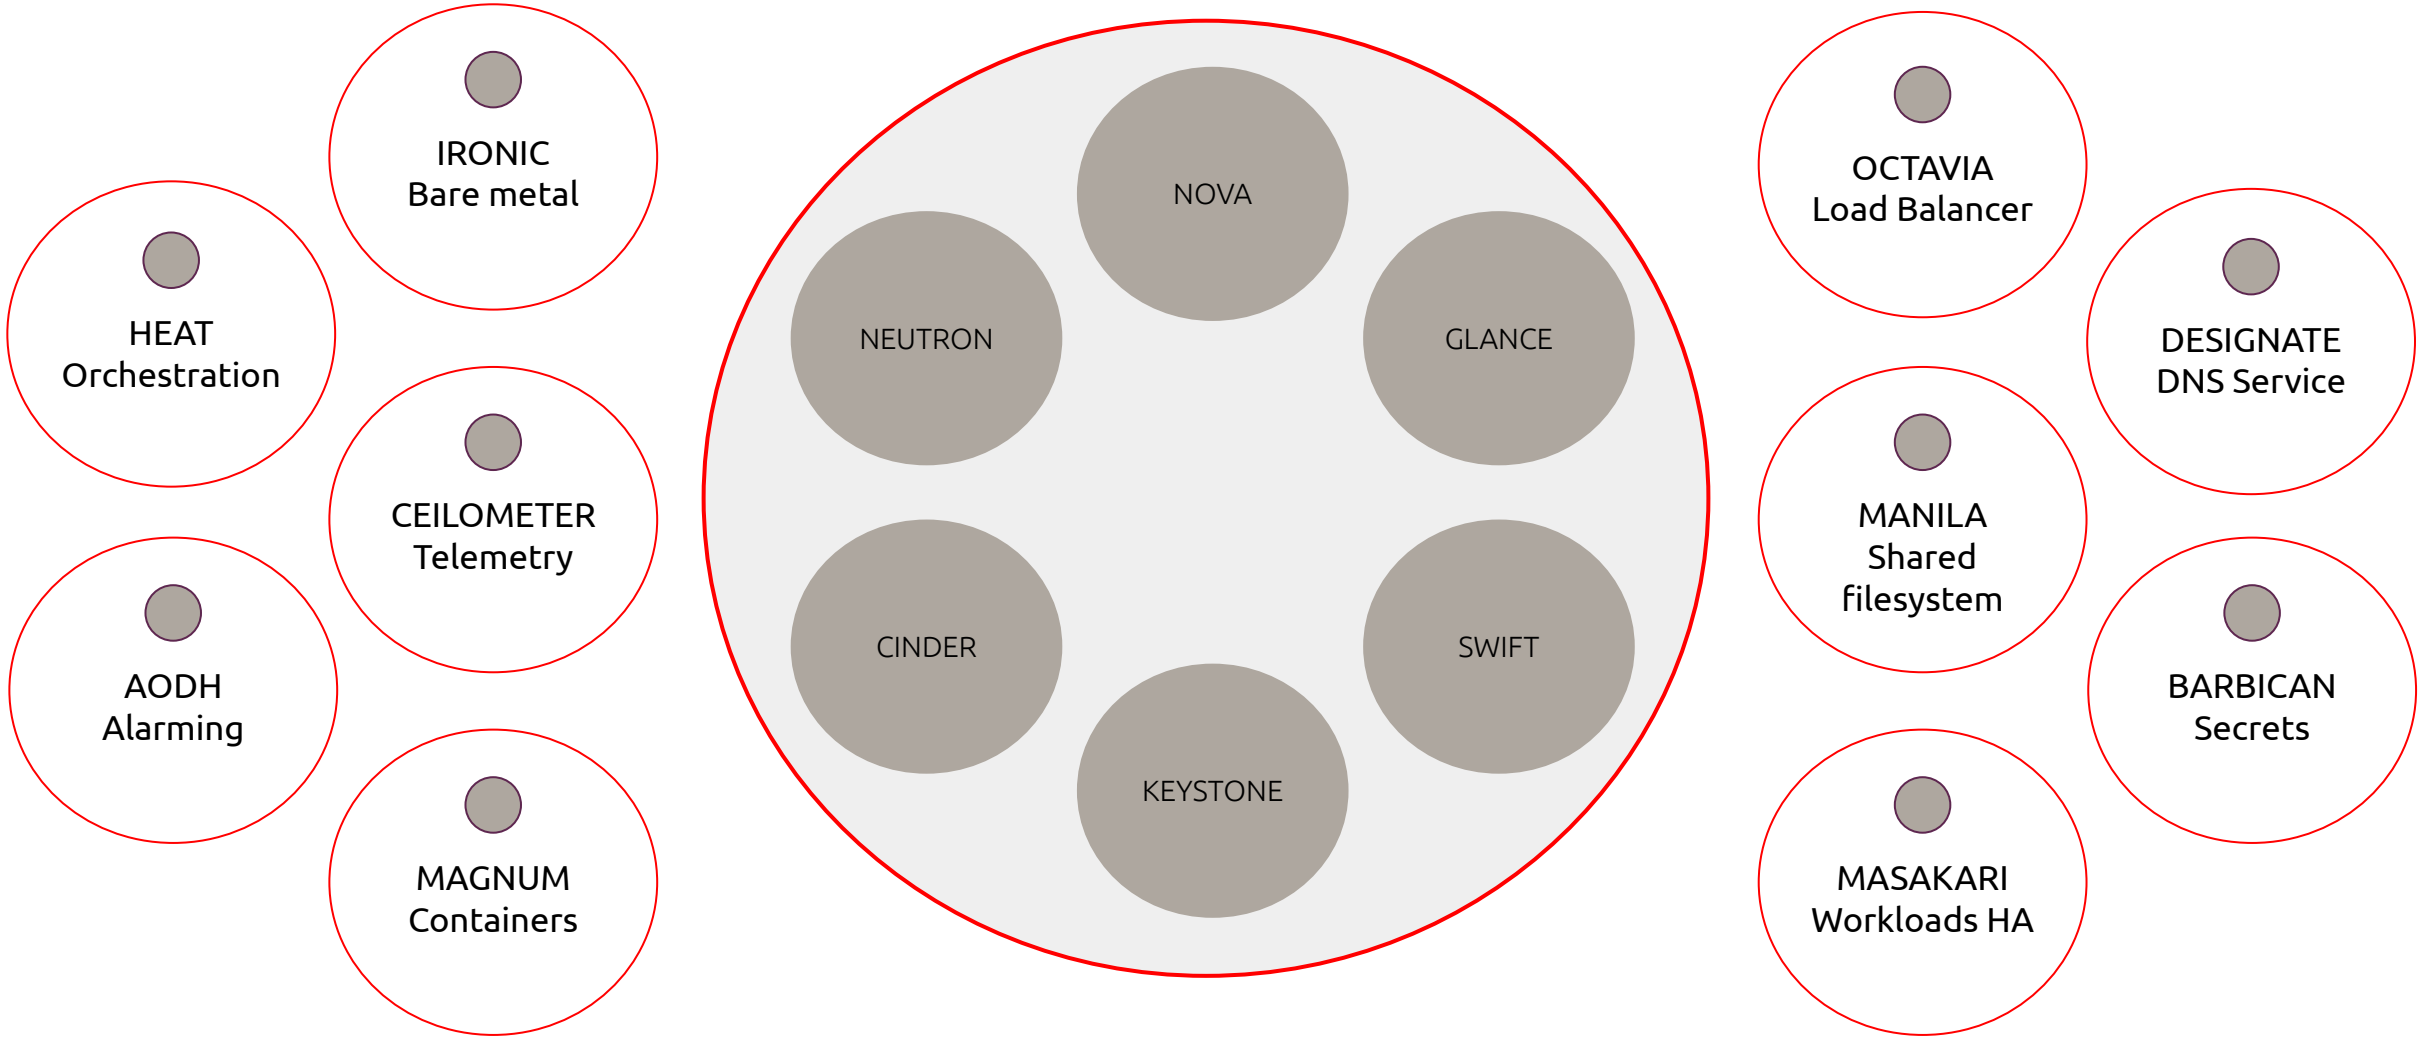

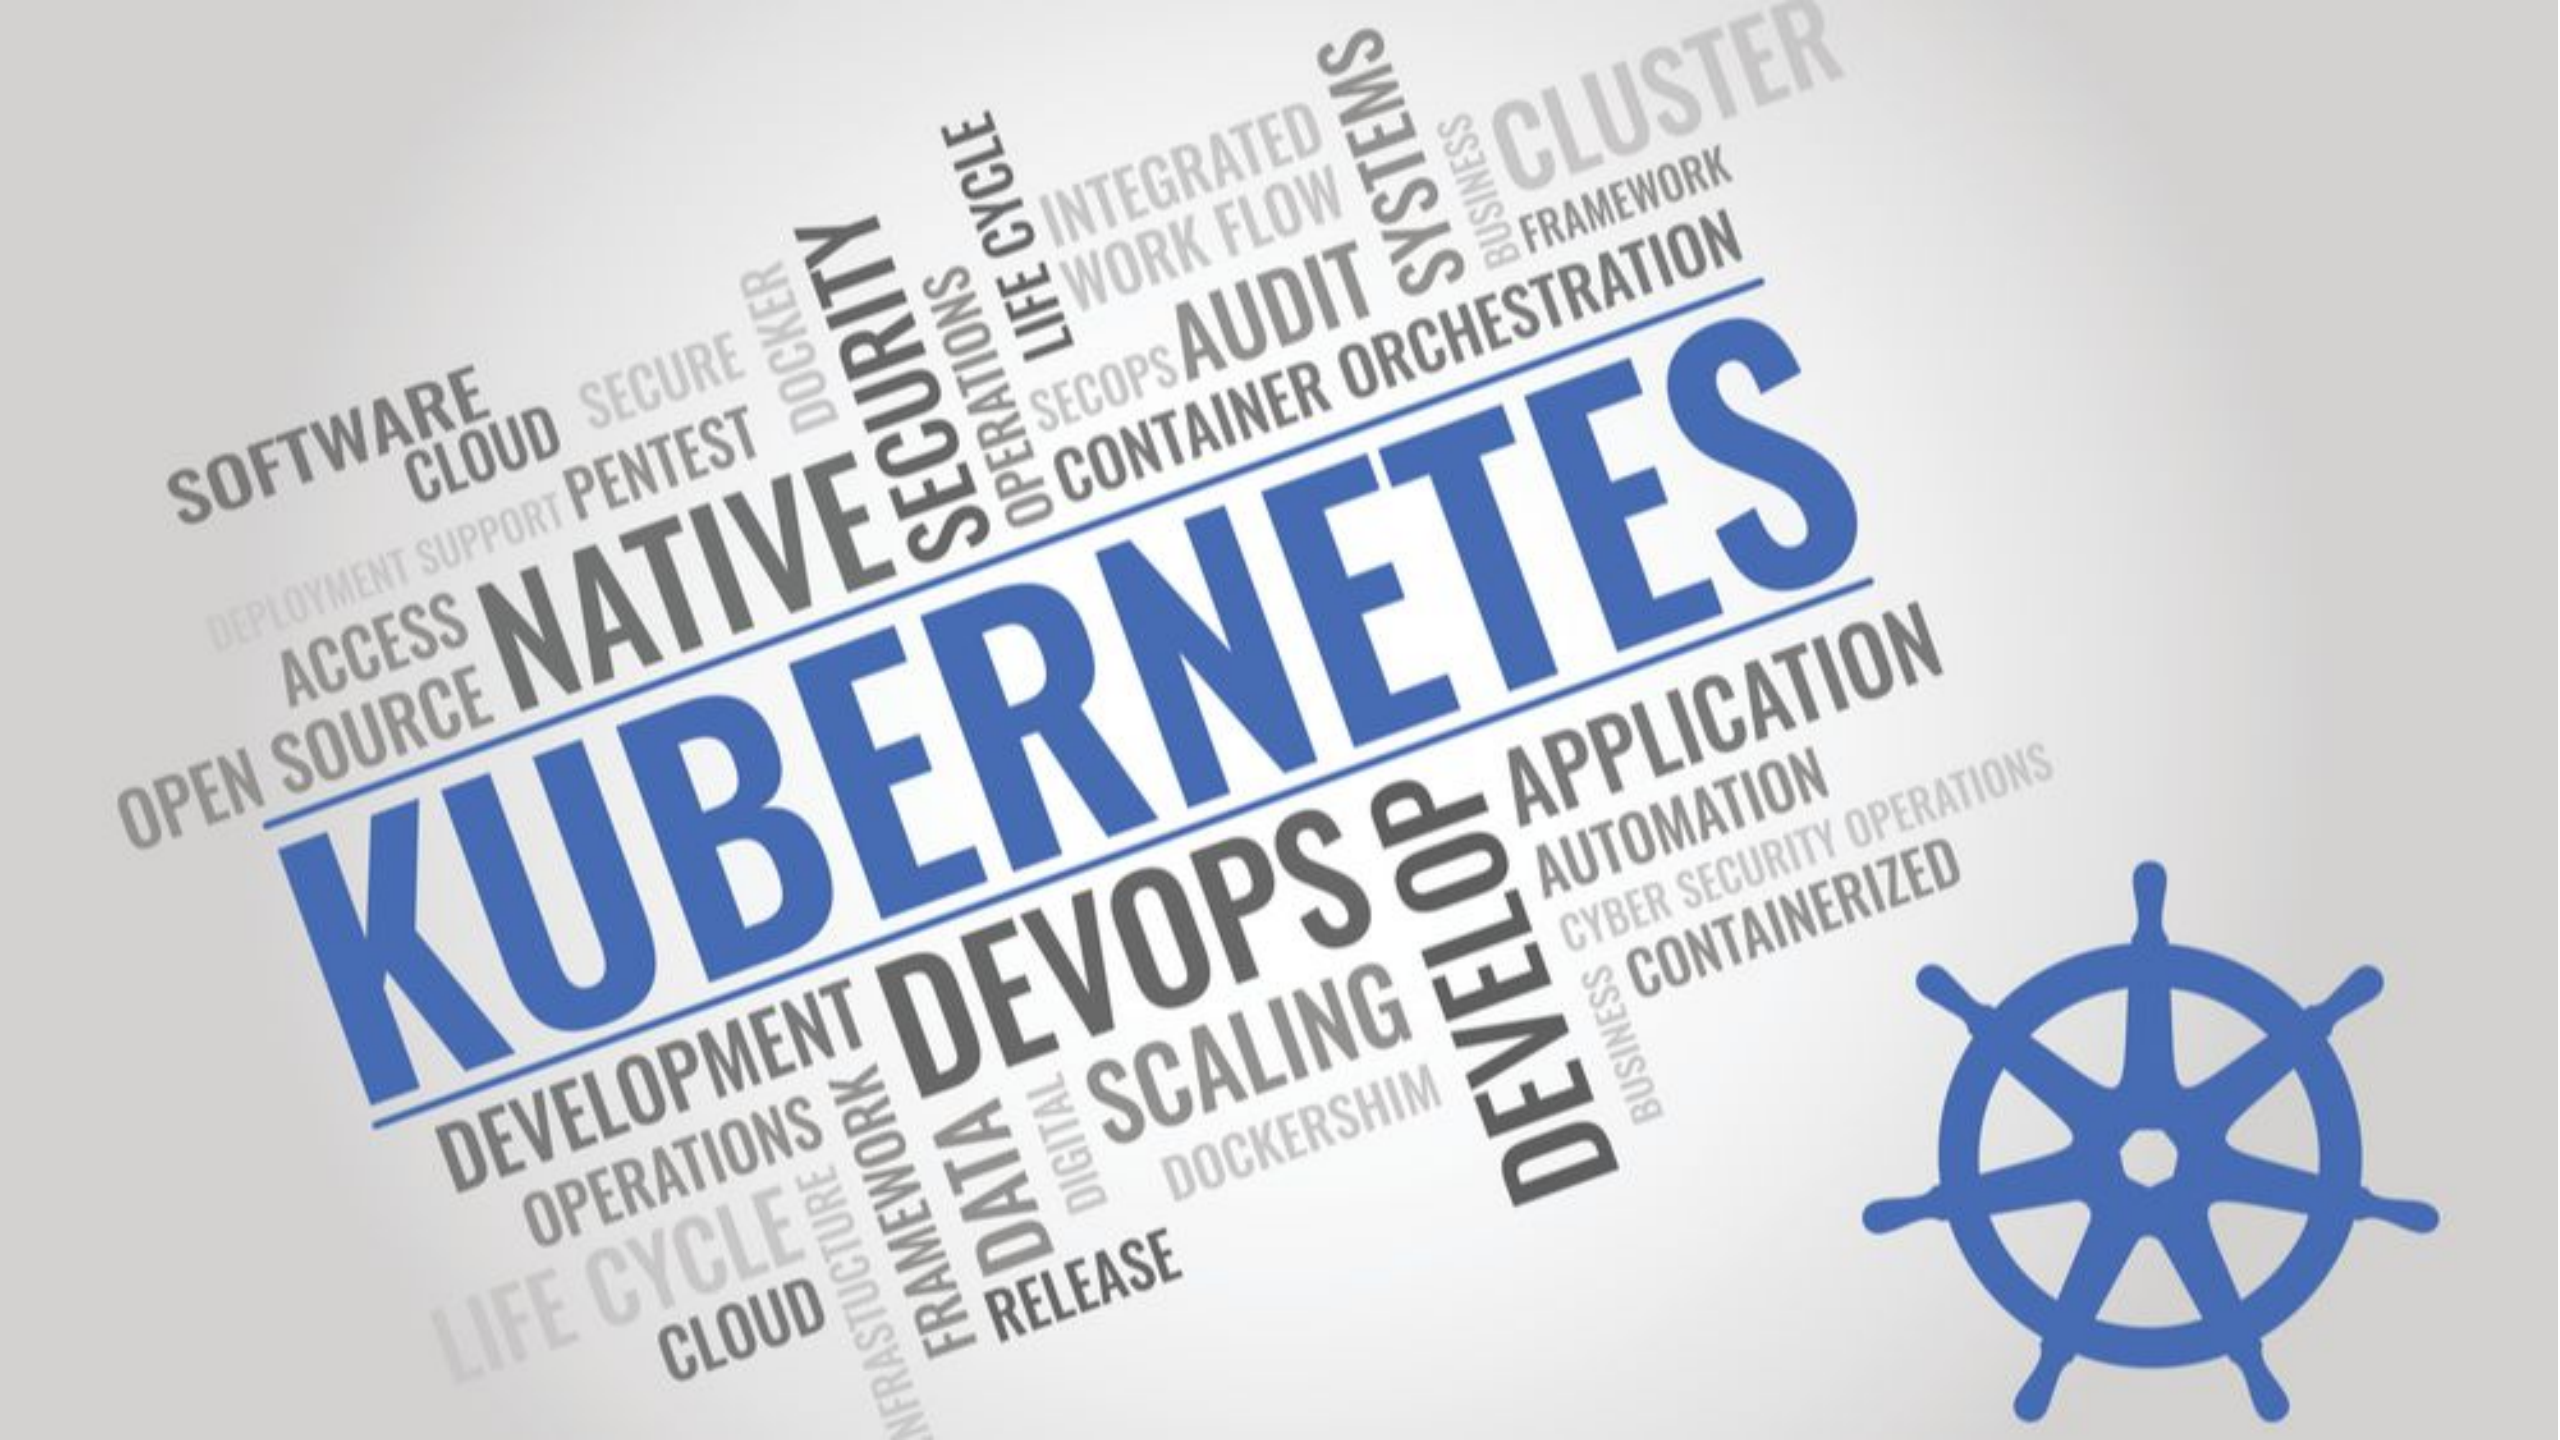
KUBERNETES
NATIVE
SECURITY
DEVOPS
DEVELOP
APPLICATION
AUTOMATION
CONTAINERIZED
CONTAINER ORCHESTRATION
AUDIT
SYSTEMS
CLUSTER
FRAMEWORK
WORK FLOW
INTEGRATED
LIFE CYCLE
OPERATIONS
SECOPS
BUSINESS
SOFTWARE
CLOUD
SECURE
DOCKER
PENTEST
DEPLOYMENT SUPPORT
ACCESS
OPEN SOURCE
DEVELOPMENT
OPERATIONS
LIFE CYCLE
CLOUD
INFRASTUCTURE
FRAMEWORK
DATA
DIGITAL
RELEASE
SCALING
DOCKERSHIM
CYBER SECURITY OPERATIONS

Monolith vs microservices

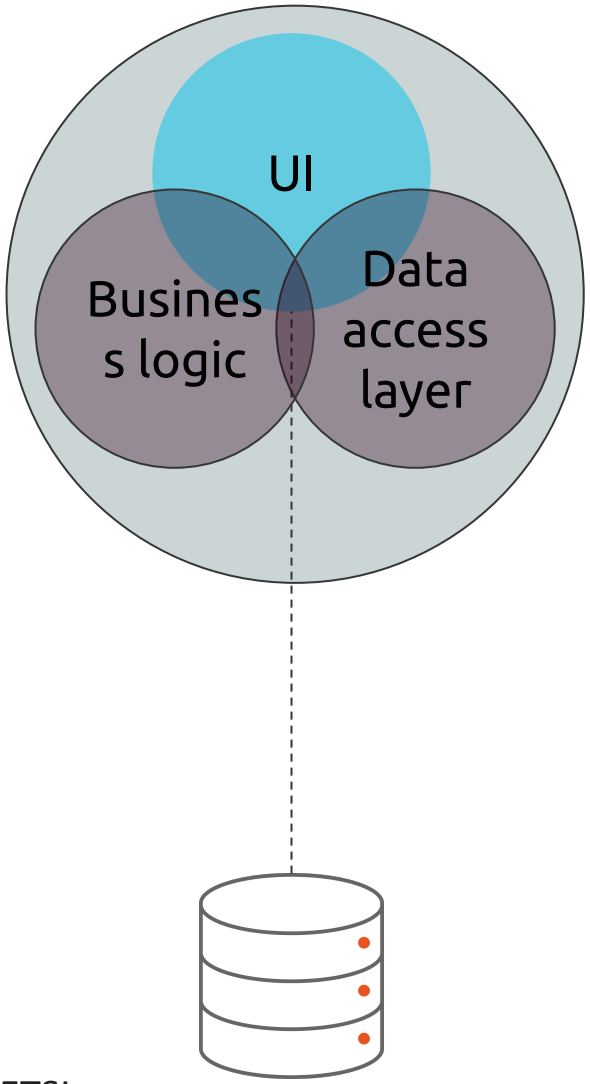

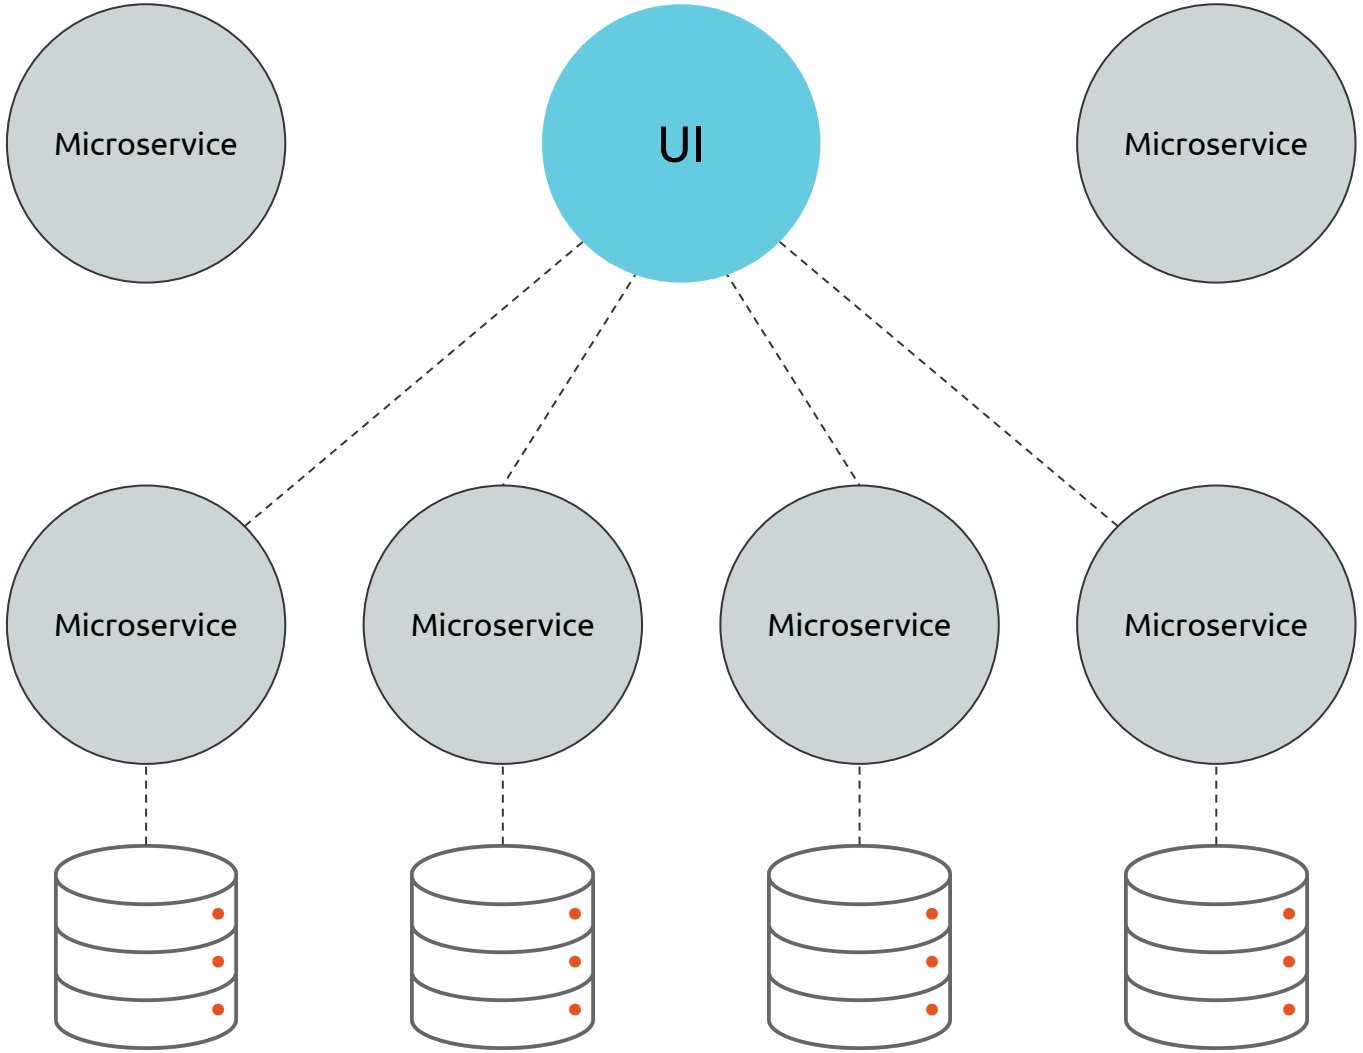

Containers vs machines

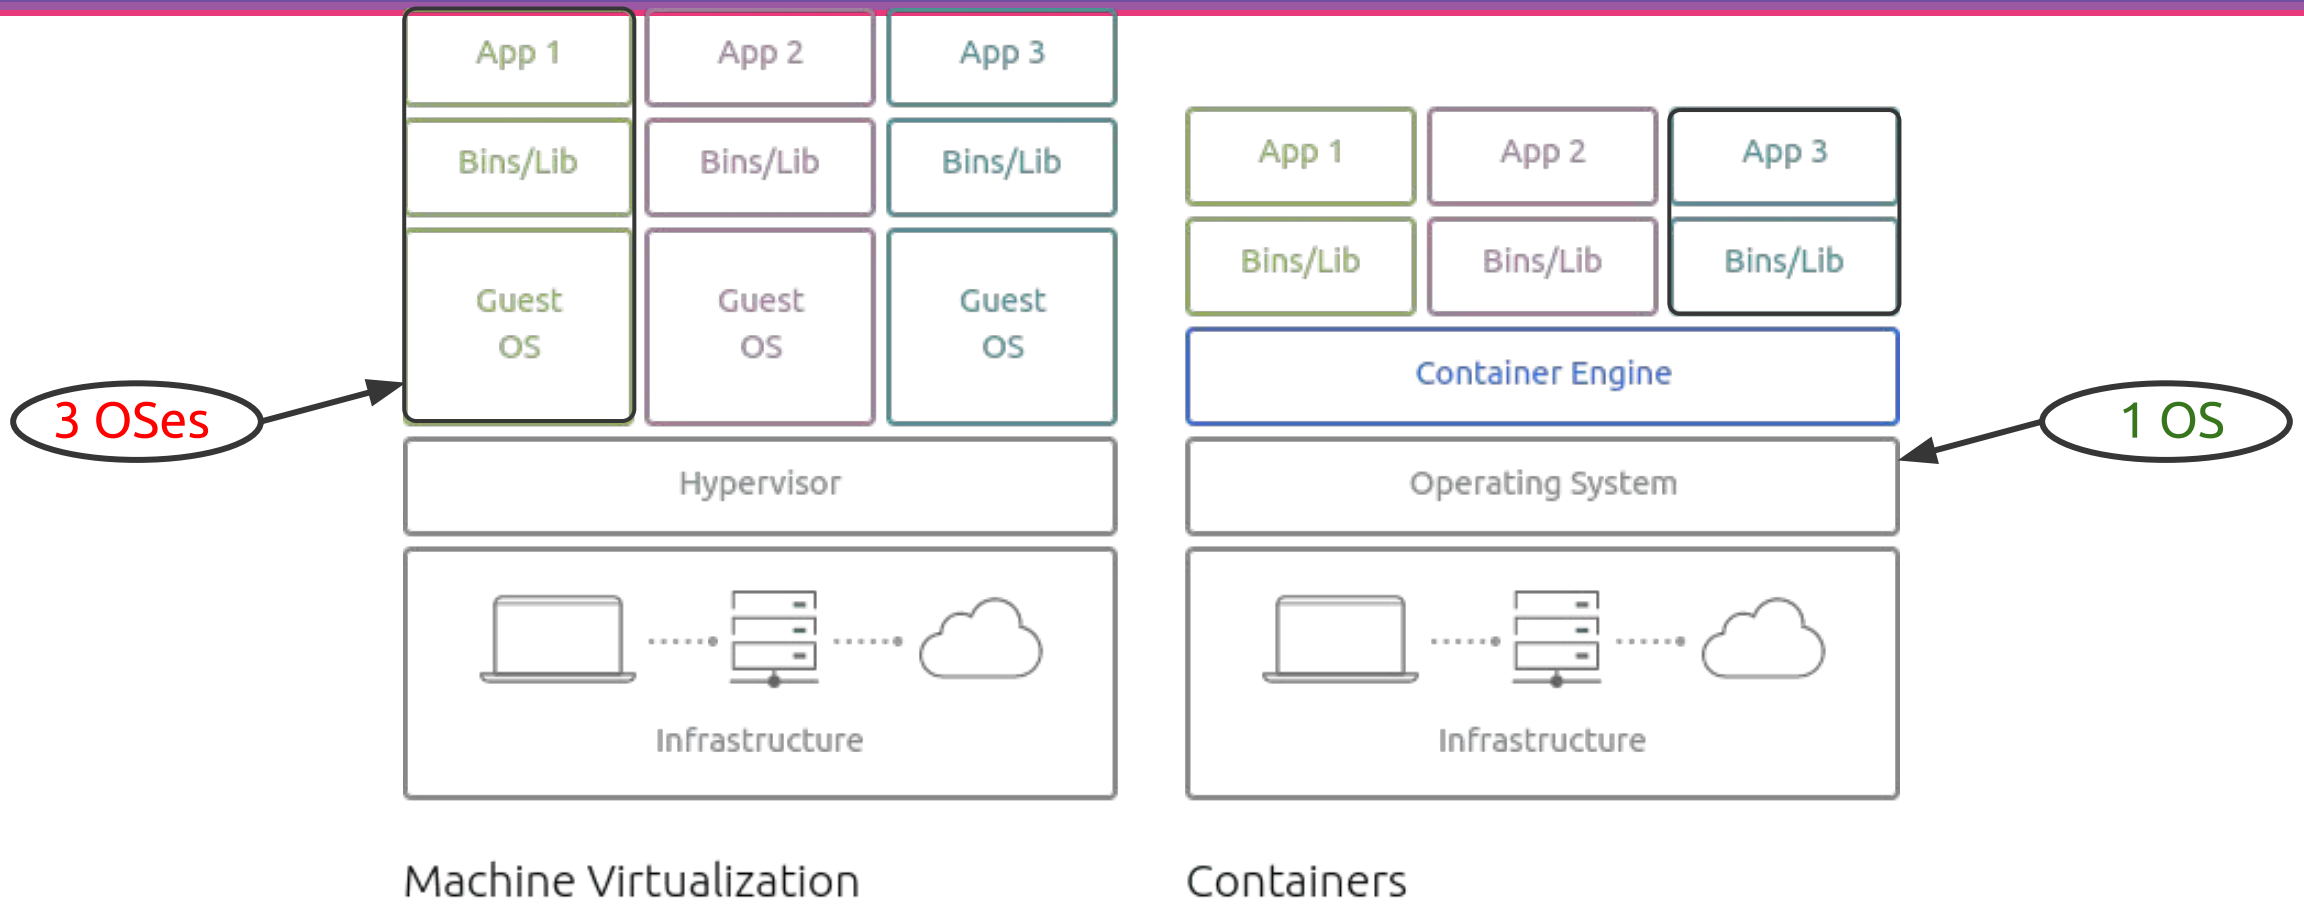

A modern way to virtualise infrastructure, more lightweight than traditional VMs.

A container includes *everything* required to run a single software application.

What is K8s

How it started

An open-source tool to simplify container orchestration

How it’s going

The first, truly universal cloud platform

landscape.cncf.io

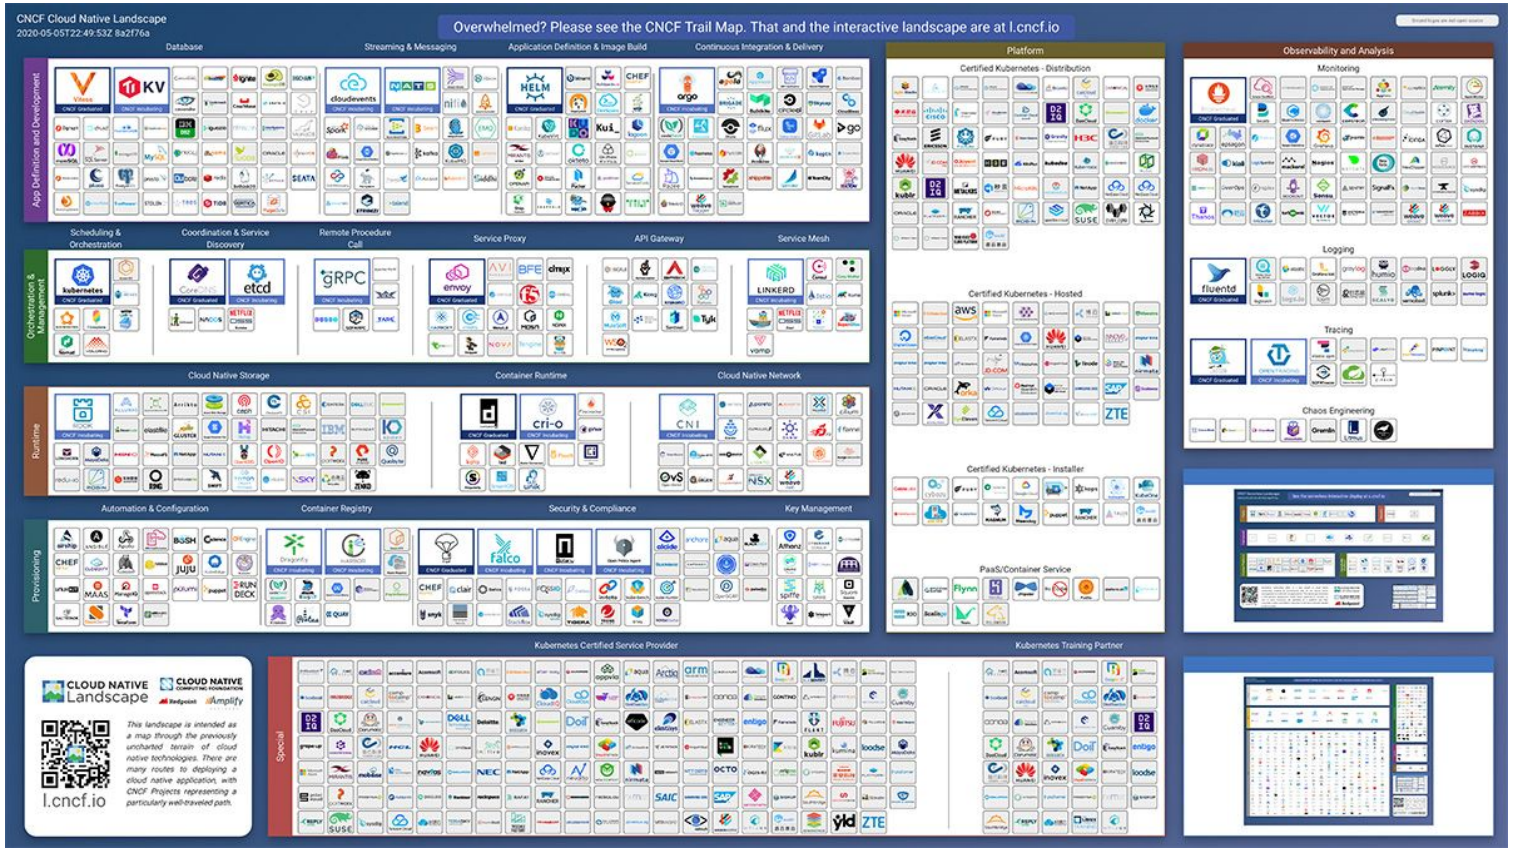

How is it technically possible

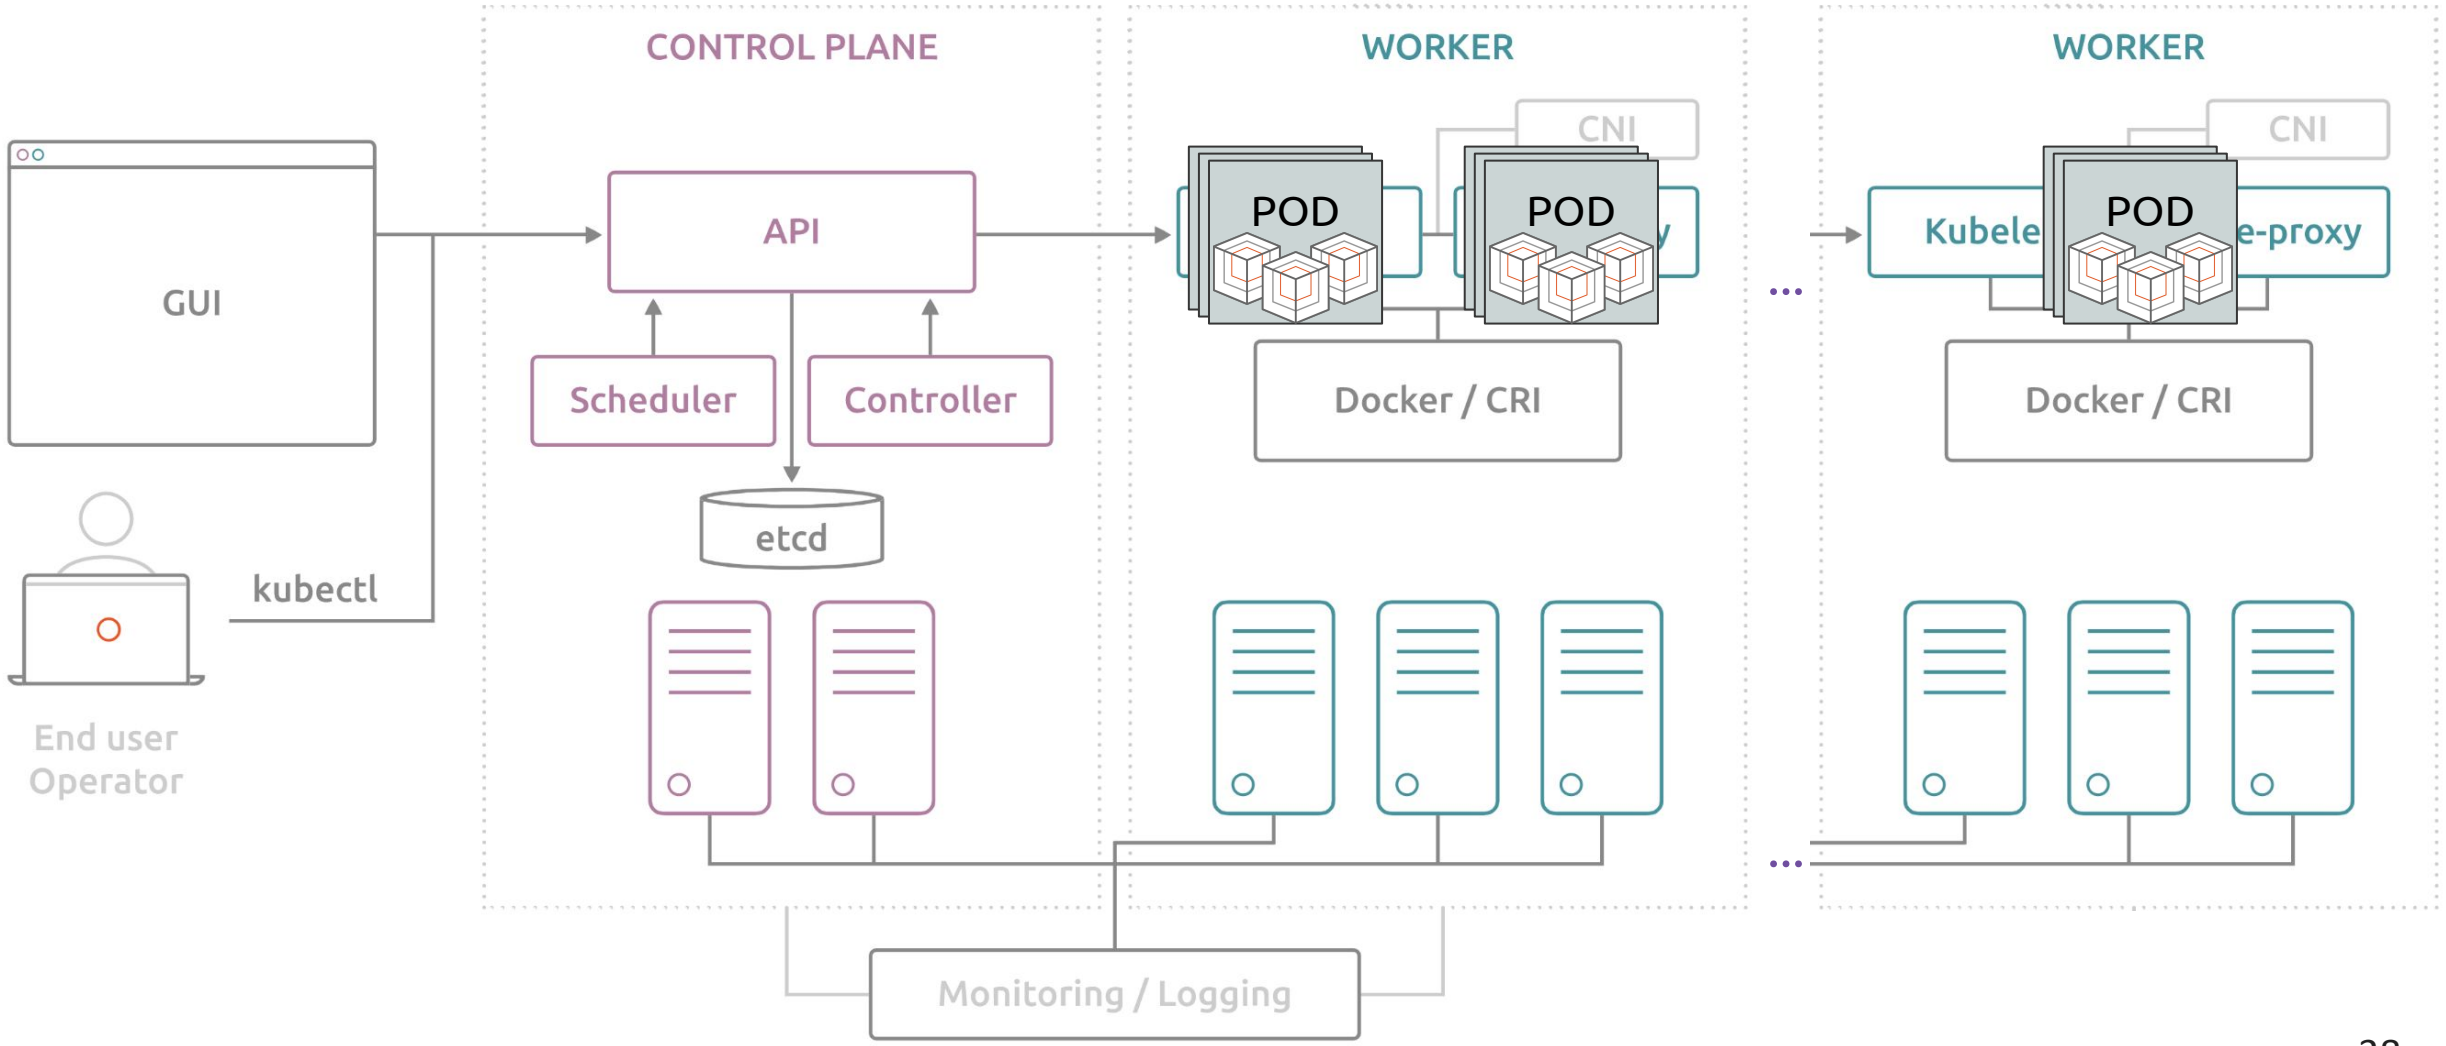

Kubernetes story so far

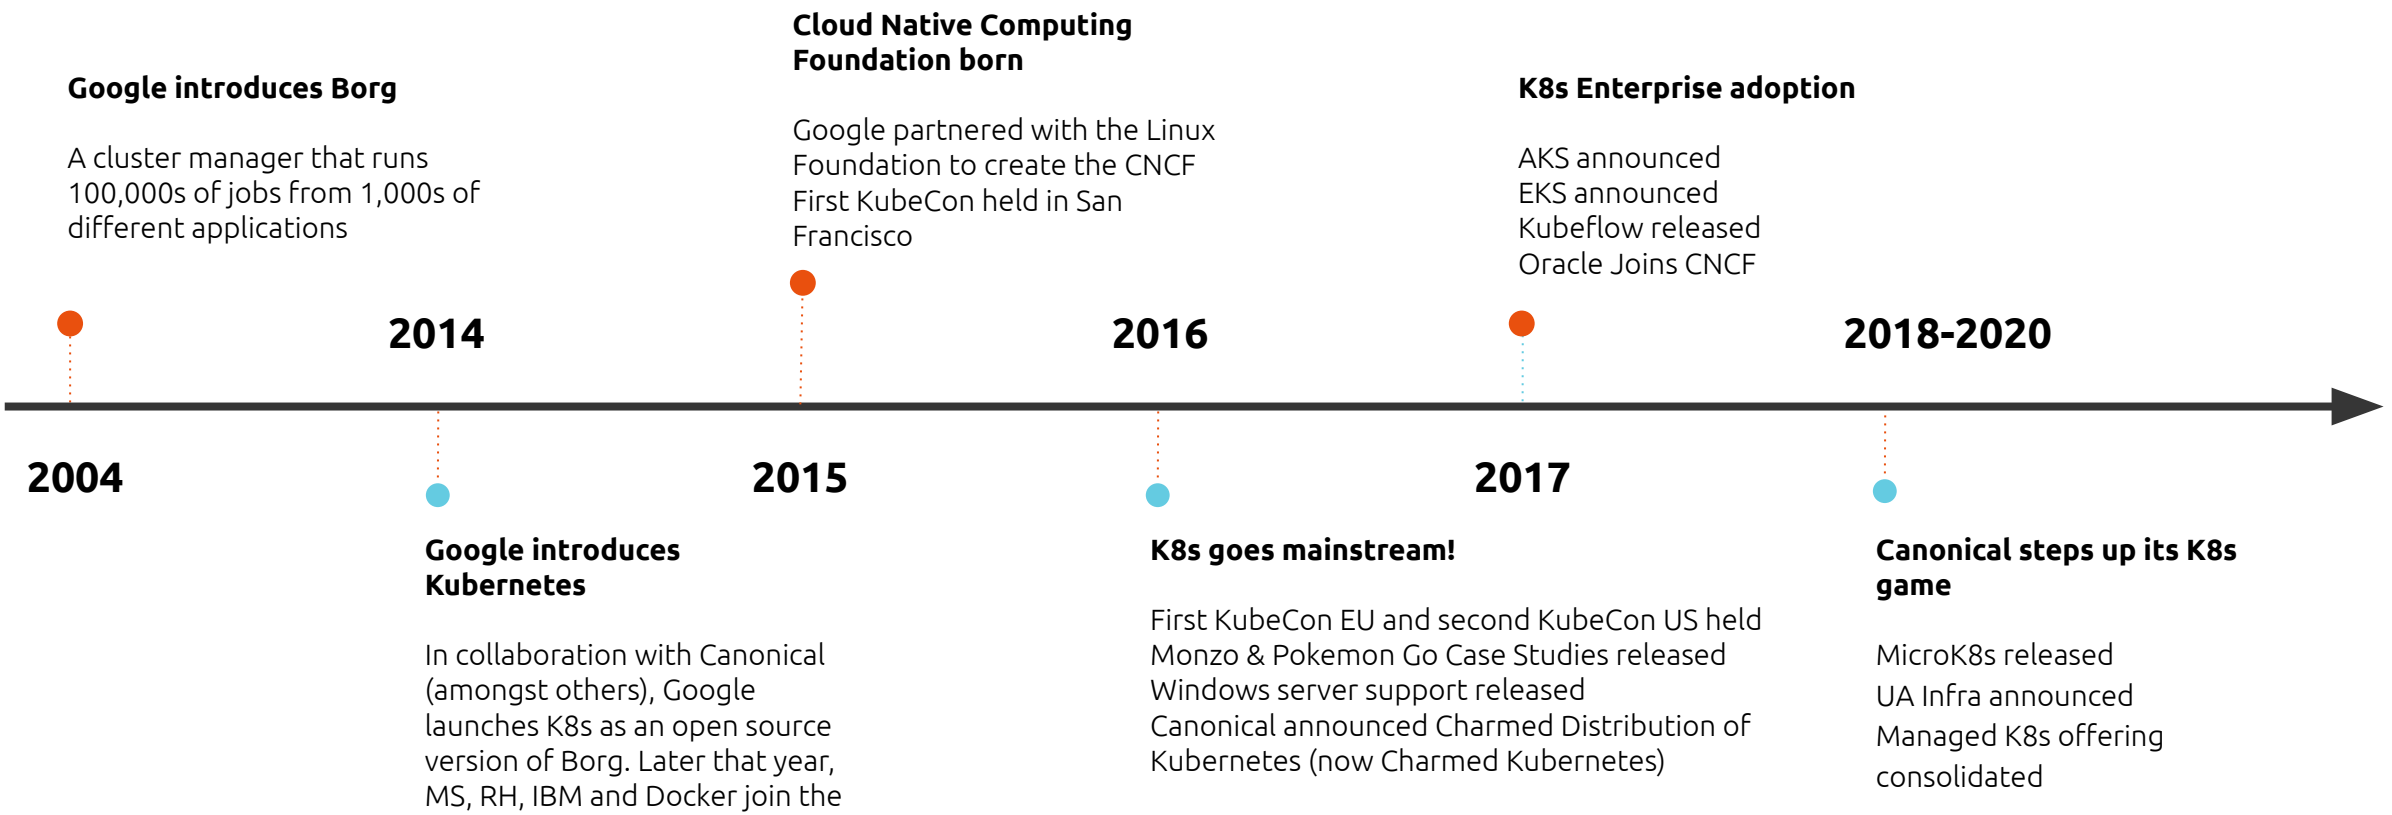

Who uses K8s and why ?

Open Source MANO

Industries

- BFSI
- Telecommunications
- Information Technology
- Retail and e-commerce
- Education
- FedGov
- Media and entertainment
- Healthcare and life sciences
- Hospitality

Use cases

- Simple deployments of stateless applications (e.g. nginx)
- Stateful data services (e.g. MySQL)
- Building modern CI/CD pipelines
- Faster/dynamic application deployments and updates
- Dynamic resource management and scale
- Productivity - less time spent in config mgnt
- Global reach of services
- List of popular real-world K8s use cases (Tinder, NYT, Airbnb, Pinterest, PokemonGo)

K8s benefits

- Elasticity and agility
- Resource optimisation
- Developer productivity
- Faster time-to-market
- Simpler operations
- Portability
- Cutting-edge technology
- Reduced OpEx and CapEx

Evolution from VNFs to CNFs