

E4

COMPUTER
ENGINEERING

WHEN PERFORMANCE MATTERS

REFERENCE ARCHITECTURES

CONTAINER
PLATFORM





WHAT'S KAPTAIN ?

KUBERNETES-BASED PLATFORM TO EASILY MANAGE COMPLEX MULTI-CONTAINER APPLICATIONS

The complexity of deployment, scheduling and load balancing operations of an infrastructure grows exponentially year on year. When it comes to running **multi-container applications** in a production environment, things get complicated: IT departments must manage hundreds or thousands of containers and different types of applications.

Kaptain is a **ready-to-use, high-performance Kubernetes cluster**, designed to host the most demanding workloads in terms of compute and storage resources and configured to be easy to use.



KAPTAIN FUNDAMENTALS

HIGH PERFORMANCE CONTAINER PLATFORM

KAPTAIN KEY POINTS

HIGH PERFORMANCE PLATFORM

It includes the typical components of HPC infrastructures (all-flash disks, GPUs and RDMA networking), to run the most data intensive workloads.


EASY AND POWERFUL UI

Rancher Server deployment enable easy to use & powerful user interfaces.

PRECONFIGURED CLOUD NATIVE STORAGE

Is it possible to use both an external storage resource (NFS, Gluster, Ceph..), or a distributed block storage solution, for the deployment, in self-provisioning, of persistent volumes to be “combined” with containerized workloads.

KAPTAIN COARSE GRAIN DESIGN



KAPTAIN MAIN ROLES



MASTER NODE


MASTER NODE is the “director” of a Kubernetes cluster and contains:

API Server: is the centralized management entity of the cluster and implements the Kubernetes API, the interface (internal and external) for all the functions;

ETCD: stores all data of the system state (applications, services, workloads, ..);

Control manager: guarantees, at all times, that the system status corresponds to what is defined in the configuration settings registered in etcd;

Scheduler: distributes workloads on available computing resources;



WORKER NODE

WORKER NODE hosts all containerized workloads, running so-called Pods.

These sets of containers are treated by the system as a single entity.

The worker node operates under the close coordination of the Master.



STORAGE NODE

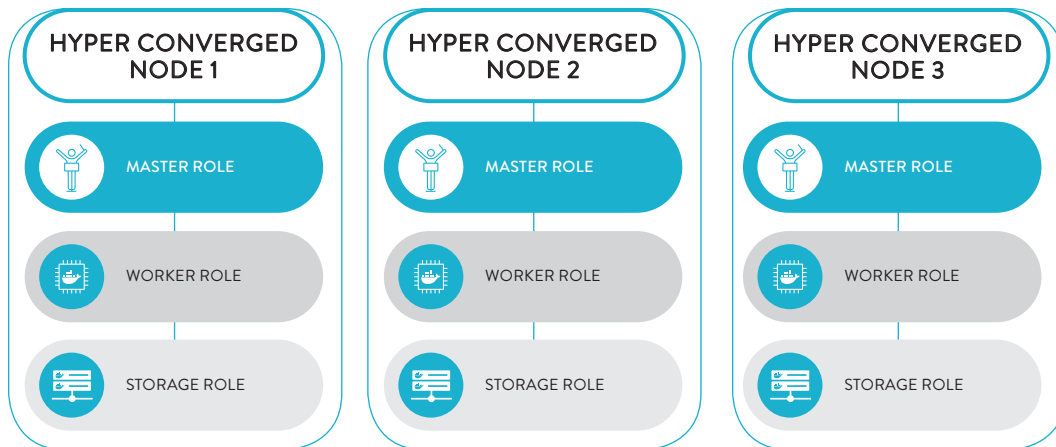
STORAGE NODE is a Worker Node dedicated (totally or partially) to workloads that implement container-based storage services, therefore “cloud native” storage services.

KAPTAIN CONFIGURATION

KAPTAIN HYPERCONVERGED

SMALLER DISTRIBUTED HA CONFIGURATION

The Hyper Convergent configuration includes a total of **3 servers**, each plays both the role of Master Node and Worker Node of the Cluster. In this configuration, each of these 3 servers also hosts the **distributed block storage services**, that is, it also plays the role of Storage Node. It is the minimum configuration able to guarantee high availability to the infrastructure.

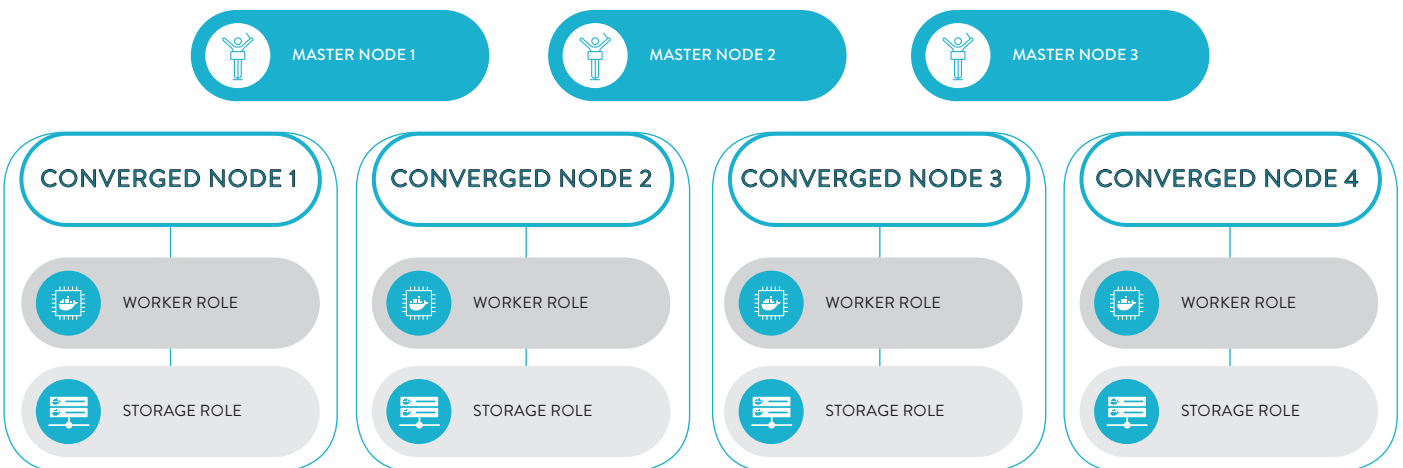


- Access and orchestration services are in High Reliability (3 master roles)
- Applications can scale horizontally (3 worker roles)
- Storage resources are high available (in replica 2 or 3) and scalable in access (3 storage roles)

KAPTAIN CONVERGED

COMPUTING PERFORMANCE AND STORAGE CAPACITY SCALE TOGETHER

The Convergent Configuration provides **3 or 5 servers** dedicated to the role of Master Node and the rest that simultaneously perform the role of Worker Node and Storage Node. It is the ideal configuration for organizations that expect a **simultaneous growth in the computing and storage** capacities they need over time.



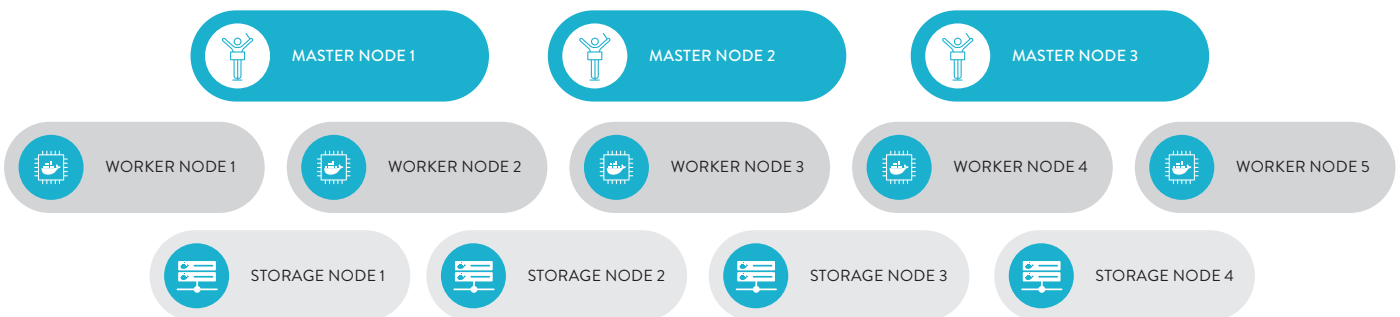
- Master role implemented on 3 dedicated nodes
- Perfect configuration for customers who expect a simultaneous growth in storage capacity and computing resources they need
- Converged Nodes can be configured with Worker and Storage dedicated resources (CPU & RAM) to ensure performance stability

KAPTAIN DISTRIBUTED

HIGHEST PERFORMANCE AND SCALABILITY

The Distributed Configuration provides **3 or 5 servers** dedicated to the Master Node role and the remaining servers that are dedicated to the Worker Node role or the Storage Node role.

This is the **maximum performance configuration**, because, in addition to providing servers dedicated to the ControlPlane, it allows you to implement the cloud native distributed block storage on the number of Storage Nodes, ideal for optimizing performance according to the net disk space required and to configure the remaining Worker Nodes solely based on the requirements of user workloads and related horizontal scaling capabilities.



Each role is implemented on a set of dedicated nodes, which can be configured independently to optimize the solution on the specific use case

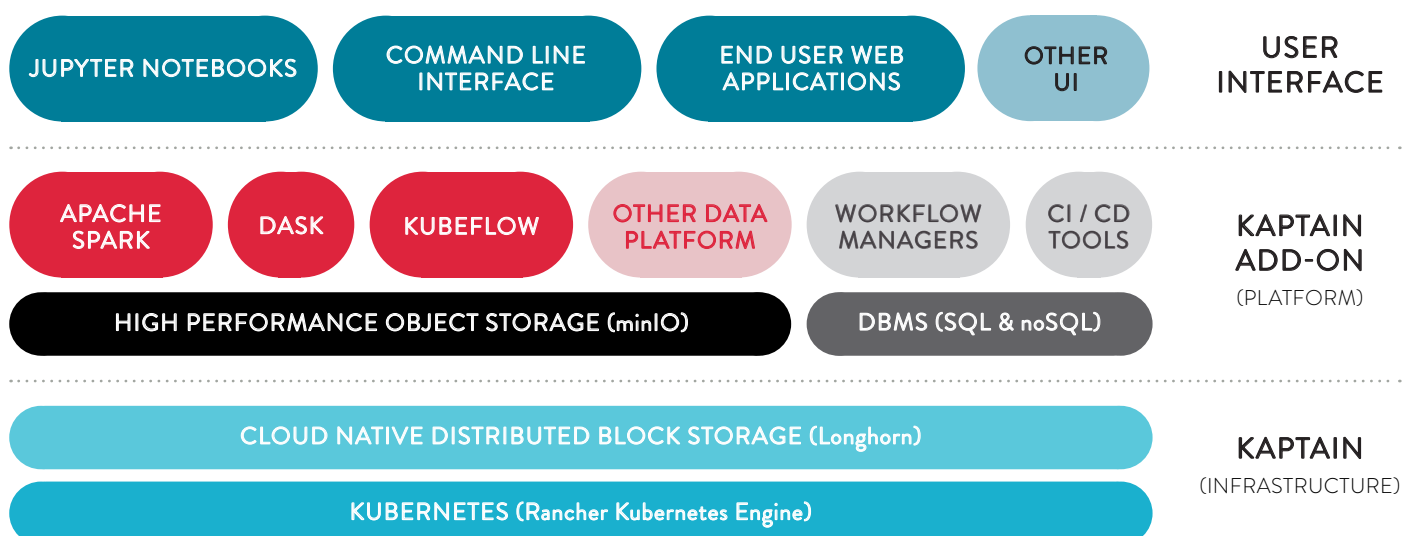
- Master role implemented on 3 or 5 nodes
- Storage role implemented on the number of nodes that optimizes performance and availability according to the overall capacity
- Worker role implemented on the number of nodes ideal for horizontal scaling of applications

CLOUD NATIVE DATA SCIENCE INFRASTRUCTURE

Kaptain has been designed to provide the customer with a high-performance infrastructure to host the most modern distributed data analysis platforms.

Kaptain is a **flexible infrastructure**, where a Apache Spark cluster can easily coexist, for example, with second generation solutions such as KubeFlow, Dask or Ray. It is a solution that **follows the evolution of data-driven applications** without any compromise in terms of usability and performance.

Kaptain is a single high-performance infrastructure to put on-line a high-performance S3 object storage, host distributed platforms for ETL of huge amounts of data, train complex Data Models based on Machine and Deep Learning algorithms, thanks to the support of GPU computing, and make the “tools of the trade” accessible (workflow manager, sql and no-sql database, tools for advanced visualization and for CI / CD, ...), **increasing the productivity of Data Scientists and Data Engineers**.



KAPTAIN VERSIONS

PARAMETRAL / VERSION	HYPERCONVERGED	CONVERGED	DISTRIBUTED
HYPERCONVERGED NODE (MASTER + WORKER + STORAGE ROLE)	3x	0x	0x
CONVERGED NODE (WORKER + STORAGE ROLE)	0x	3x to ∞	0x
MASTER NODE	0x	3x	3x
WORKER NODE	0x	0x	3x to ∞
STORAGE NODE	0x	0x	3x to ∞

KAPTAIN SIZING

PARAMETRAL VERSION	INTEL	INTEL	INTEL	AMD	AMD	AMD
ROLE	Master or Storage node	Hyper Converged, Converged, Worker or Storage node	Hyper Converged, Converged or Storage node	Master or Storage node	Hyper Converged, Converged, Worker or Storage node	Hyper Converged, Converged or Storage node
FORM FACTOR	1 RU	2 RU / 4 N	2 RU / 4 N	1 RU	2 RU / 4 N	1 RU
CPU	2x 8 cores	2x 20 cores 2,1GHz	2x 8 cores 2,1GHz	1x 8 cores 3,10GHz	2x 24 cores 2,30GHz	1x 8 cores 3,10GHz
RAM	192GB (12x16GB)	384GB (12x32GB)	192GB (12x16GB)	128GB or 256GB (4x or 8x 32GB)	256GB (16x16GB)	256GB or 256GB (16x 16GB)
OS DISKS	2x 960GB SSD	2x 960GB SSD	2x 480GB SSD	2x 960GB SSD	2x 960GB SSD	2x 960GB SSD
DATA DISKS	Up to 6x 2TB SSD	Up to 6x 2TB SSD	Up to 4x NVME 8TB	Up to 6x 4TB SSD	Up to 4x 2TB SSD	Up to 4x NVME 8TB
RAID	YES	YES	YES	YES	Software	Software
NETWORK	Up to 2x SinglePort 100Gb/s	1x SinglePort 100Gb/s	1x SinglePort 100Gb/s	Up to 2x SinglePort 100Gb/s	1x SinglePort 25Gb/s	1x Single Port 100Gb/s

E4

COMPUTER
ENGINEERING

VIA MARTIRI DELLA LIBERTÀ, 66
42019 SCANDIANO . RE . ITALY

TEL. +39 0522 991811
FAX +39 0522 991803

INFO@E4COMPANY.COM
WWW.E4COMPANY.COM

