Deploying at scale with PaaS

QCon
SAN FRANCISCO 2014

WSO2
lean • enterprise • middleware

Last Updated: October. 2014

About Me

Lakmal Warusawithana

Vise President, Apache Stratos

Director - Cloud Architecture, WSO2 Inc

lakmal@apache.org / lakmal@wso2.com

© WSO2

Agenda

- Introduction to PaaS
- Introduction to Apache Stratos
- Apache Stratos Architecture
- Discuss few Apache Stratos features
  - Multi-factored auto scaling
  - Scalable and dynamic load balancing
  - Smart policies
  - Multi tenancy
  - Cloud bursting
  - Logging, metering and monitoring
- Apache Stratos with Docker
- Introduction to Containers, Docker, CoreOS, Kubernetes ...etc
- Demo - using Apache Stratos 4.1.0-m2 developer preview

© WSO2

Conventional Software

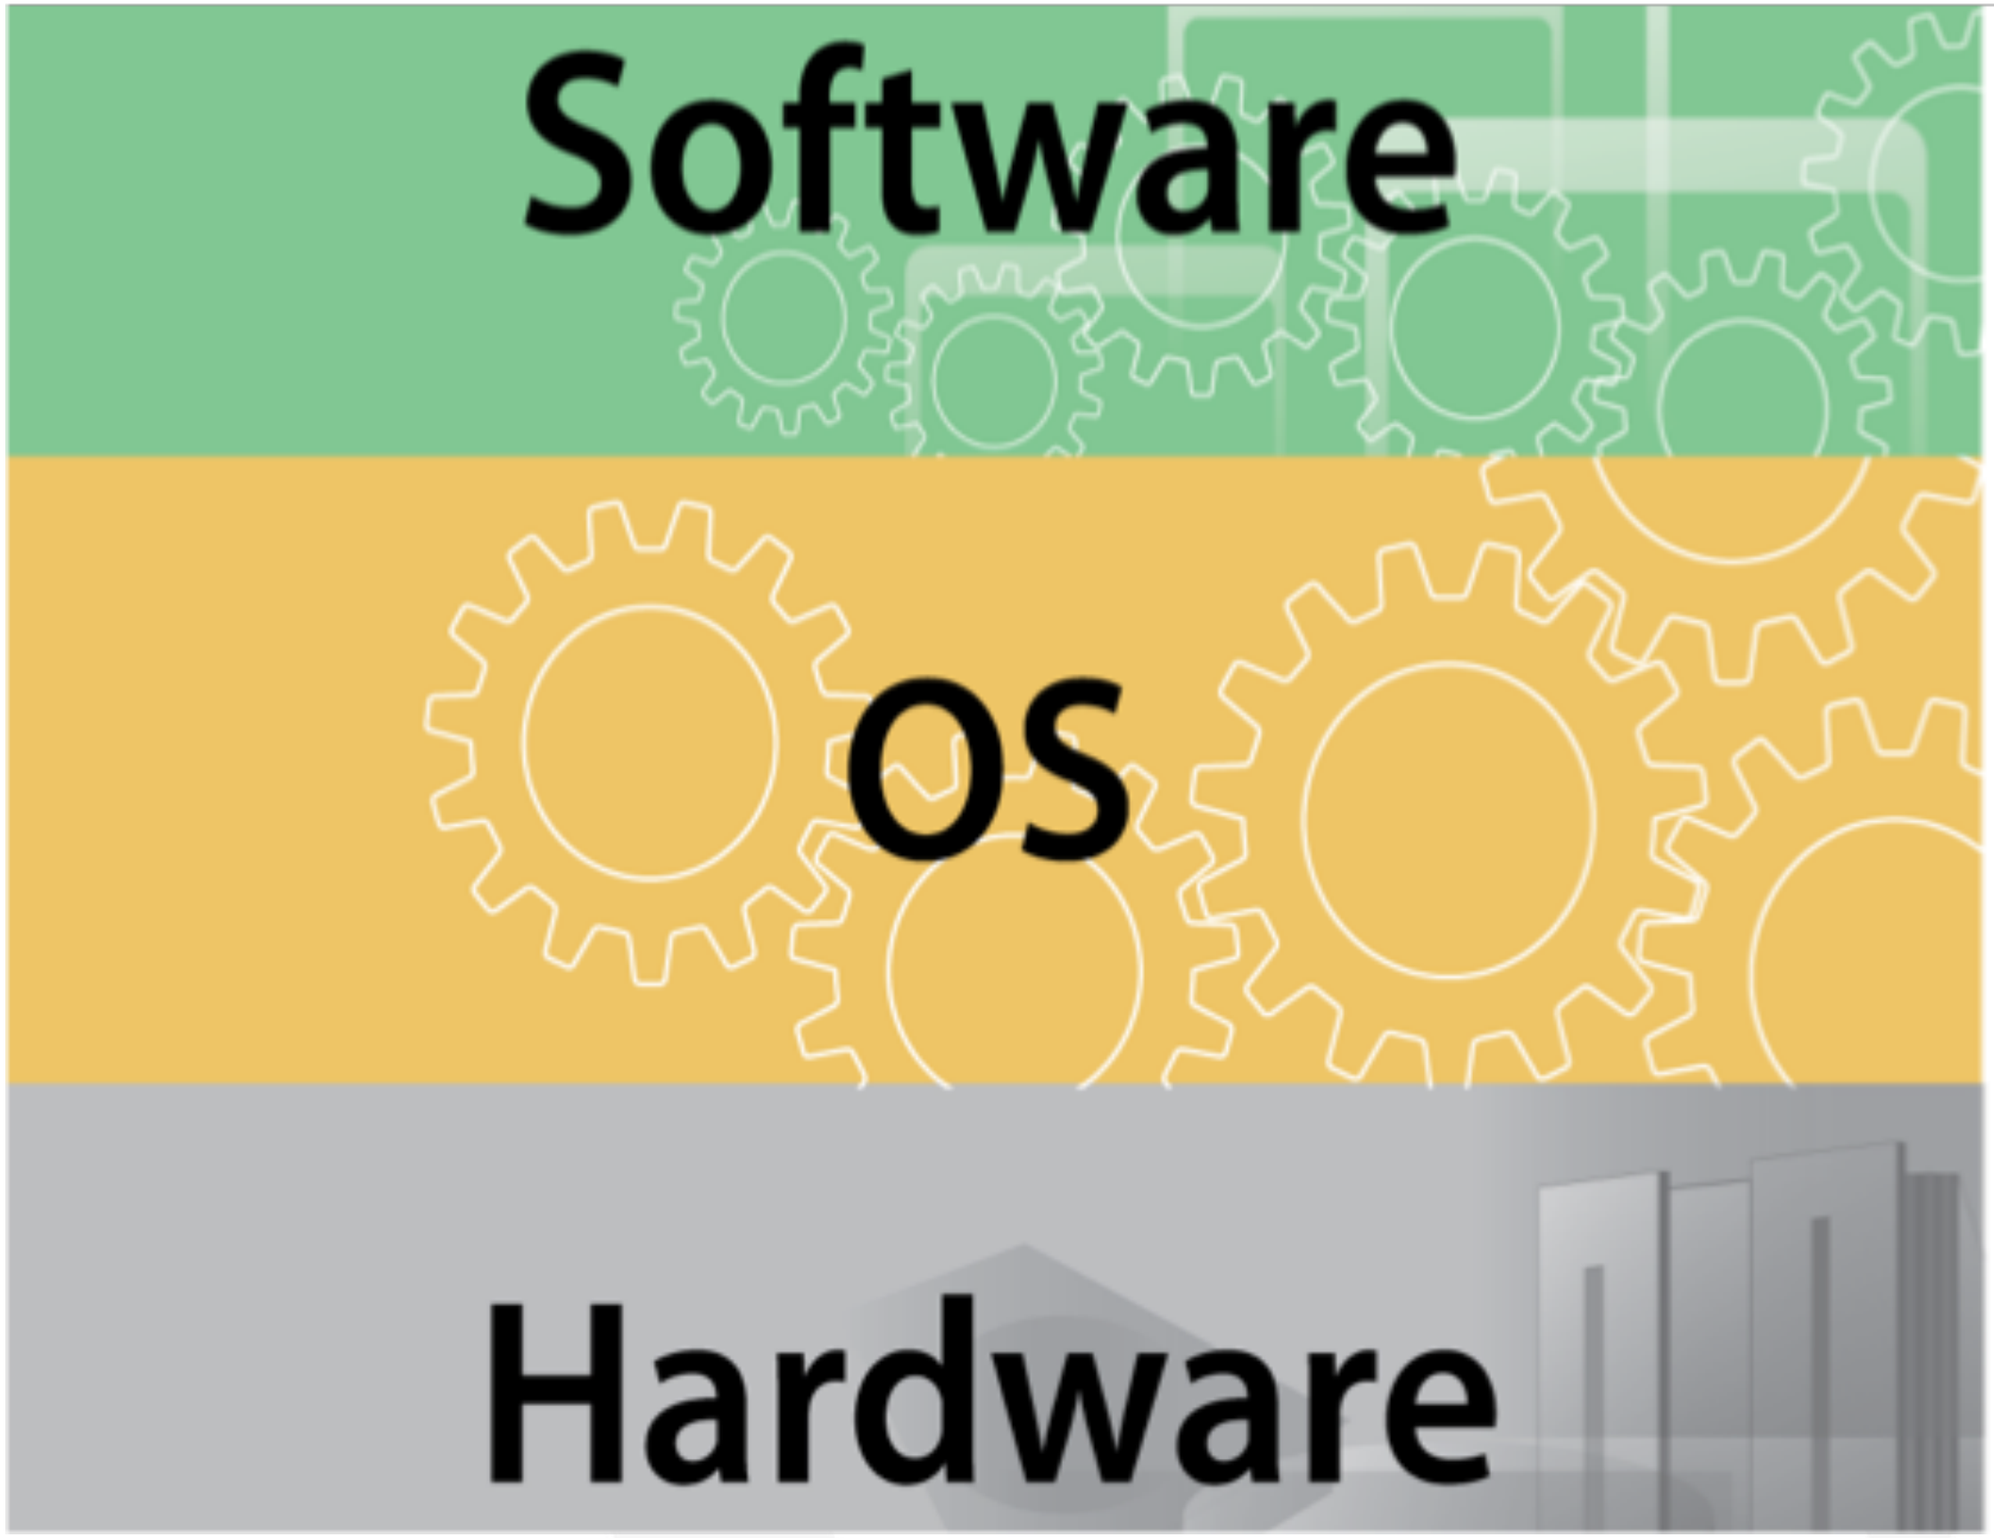

© WSO2

What is the Middleware Platform?

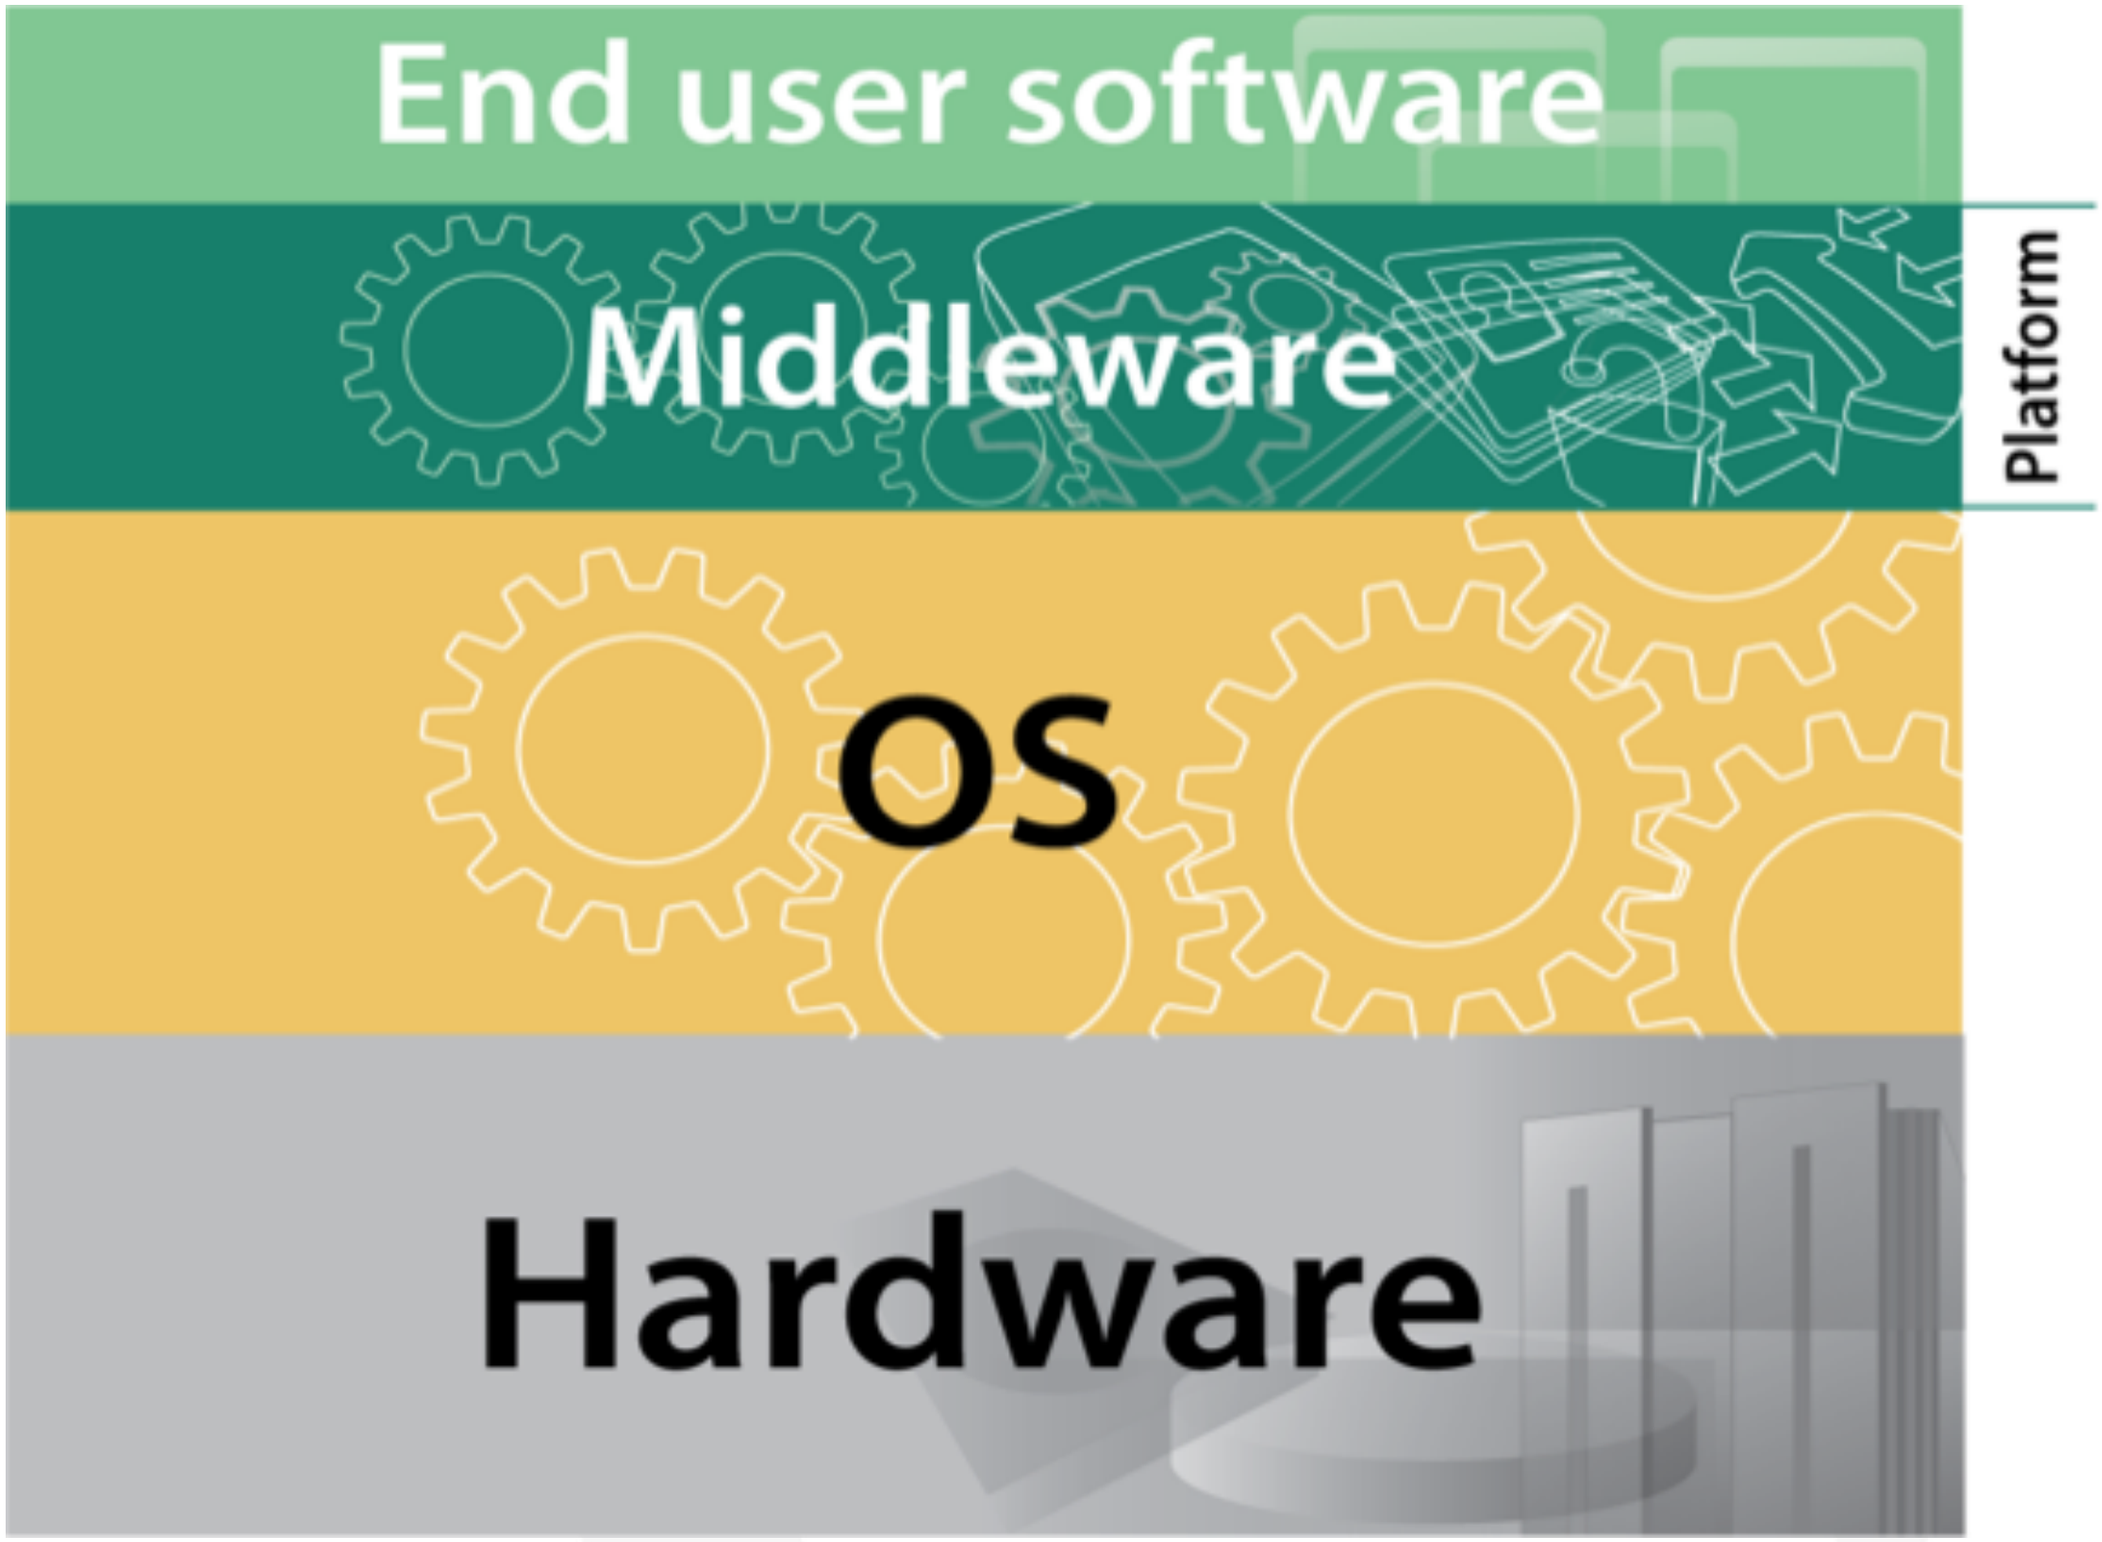

© WSO2

What is Cloud Service

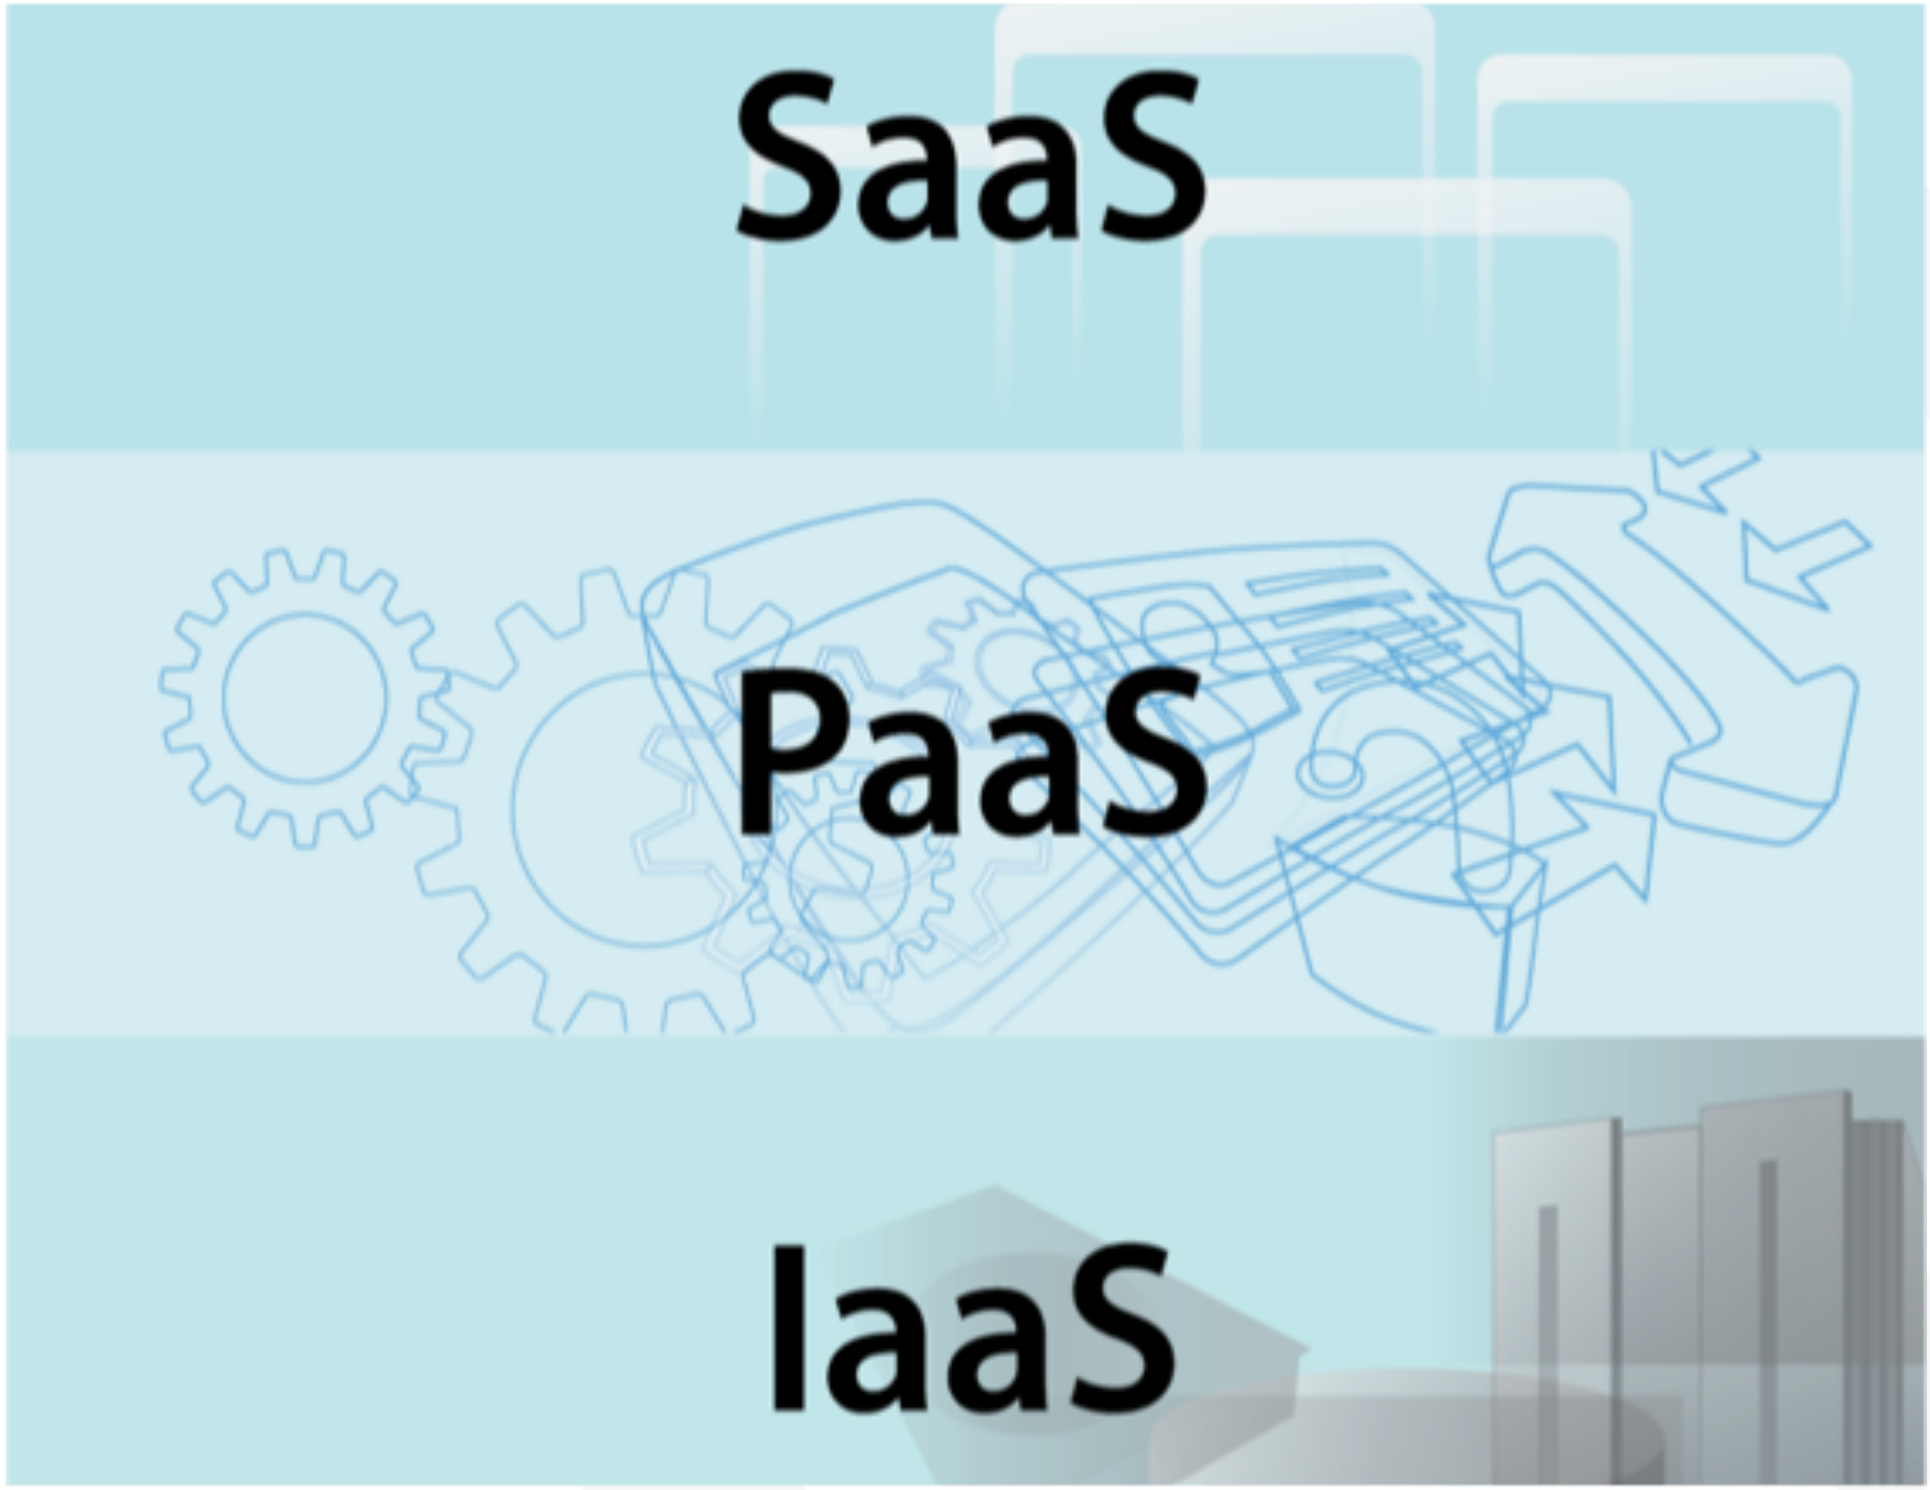

© WSO2

Apache Stratos

- Apache Stratos is a highly-extensible Platform-as-a-Service (PaaS) framework that helps run Apache Tomcat, PHP, and MySQL applications and can be extended to support many more environments on all major cloud infrastructures

- Stratos initially develop by WSO2 and last year donated to Apache Software Foundation

- After successfully complete the incubating process Stratos now graduated as Top Level Project

© WSO2

Apache Stratos Layered Architecture

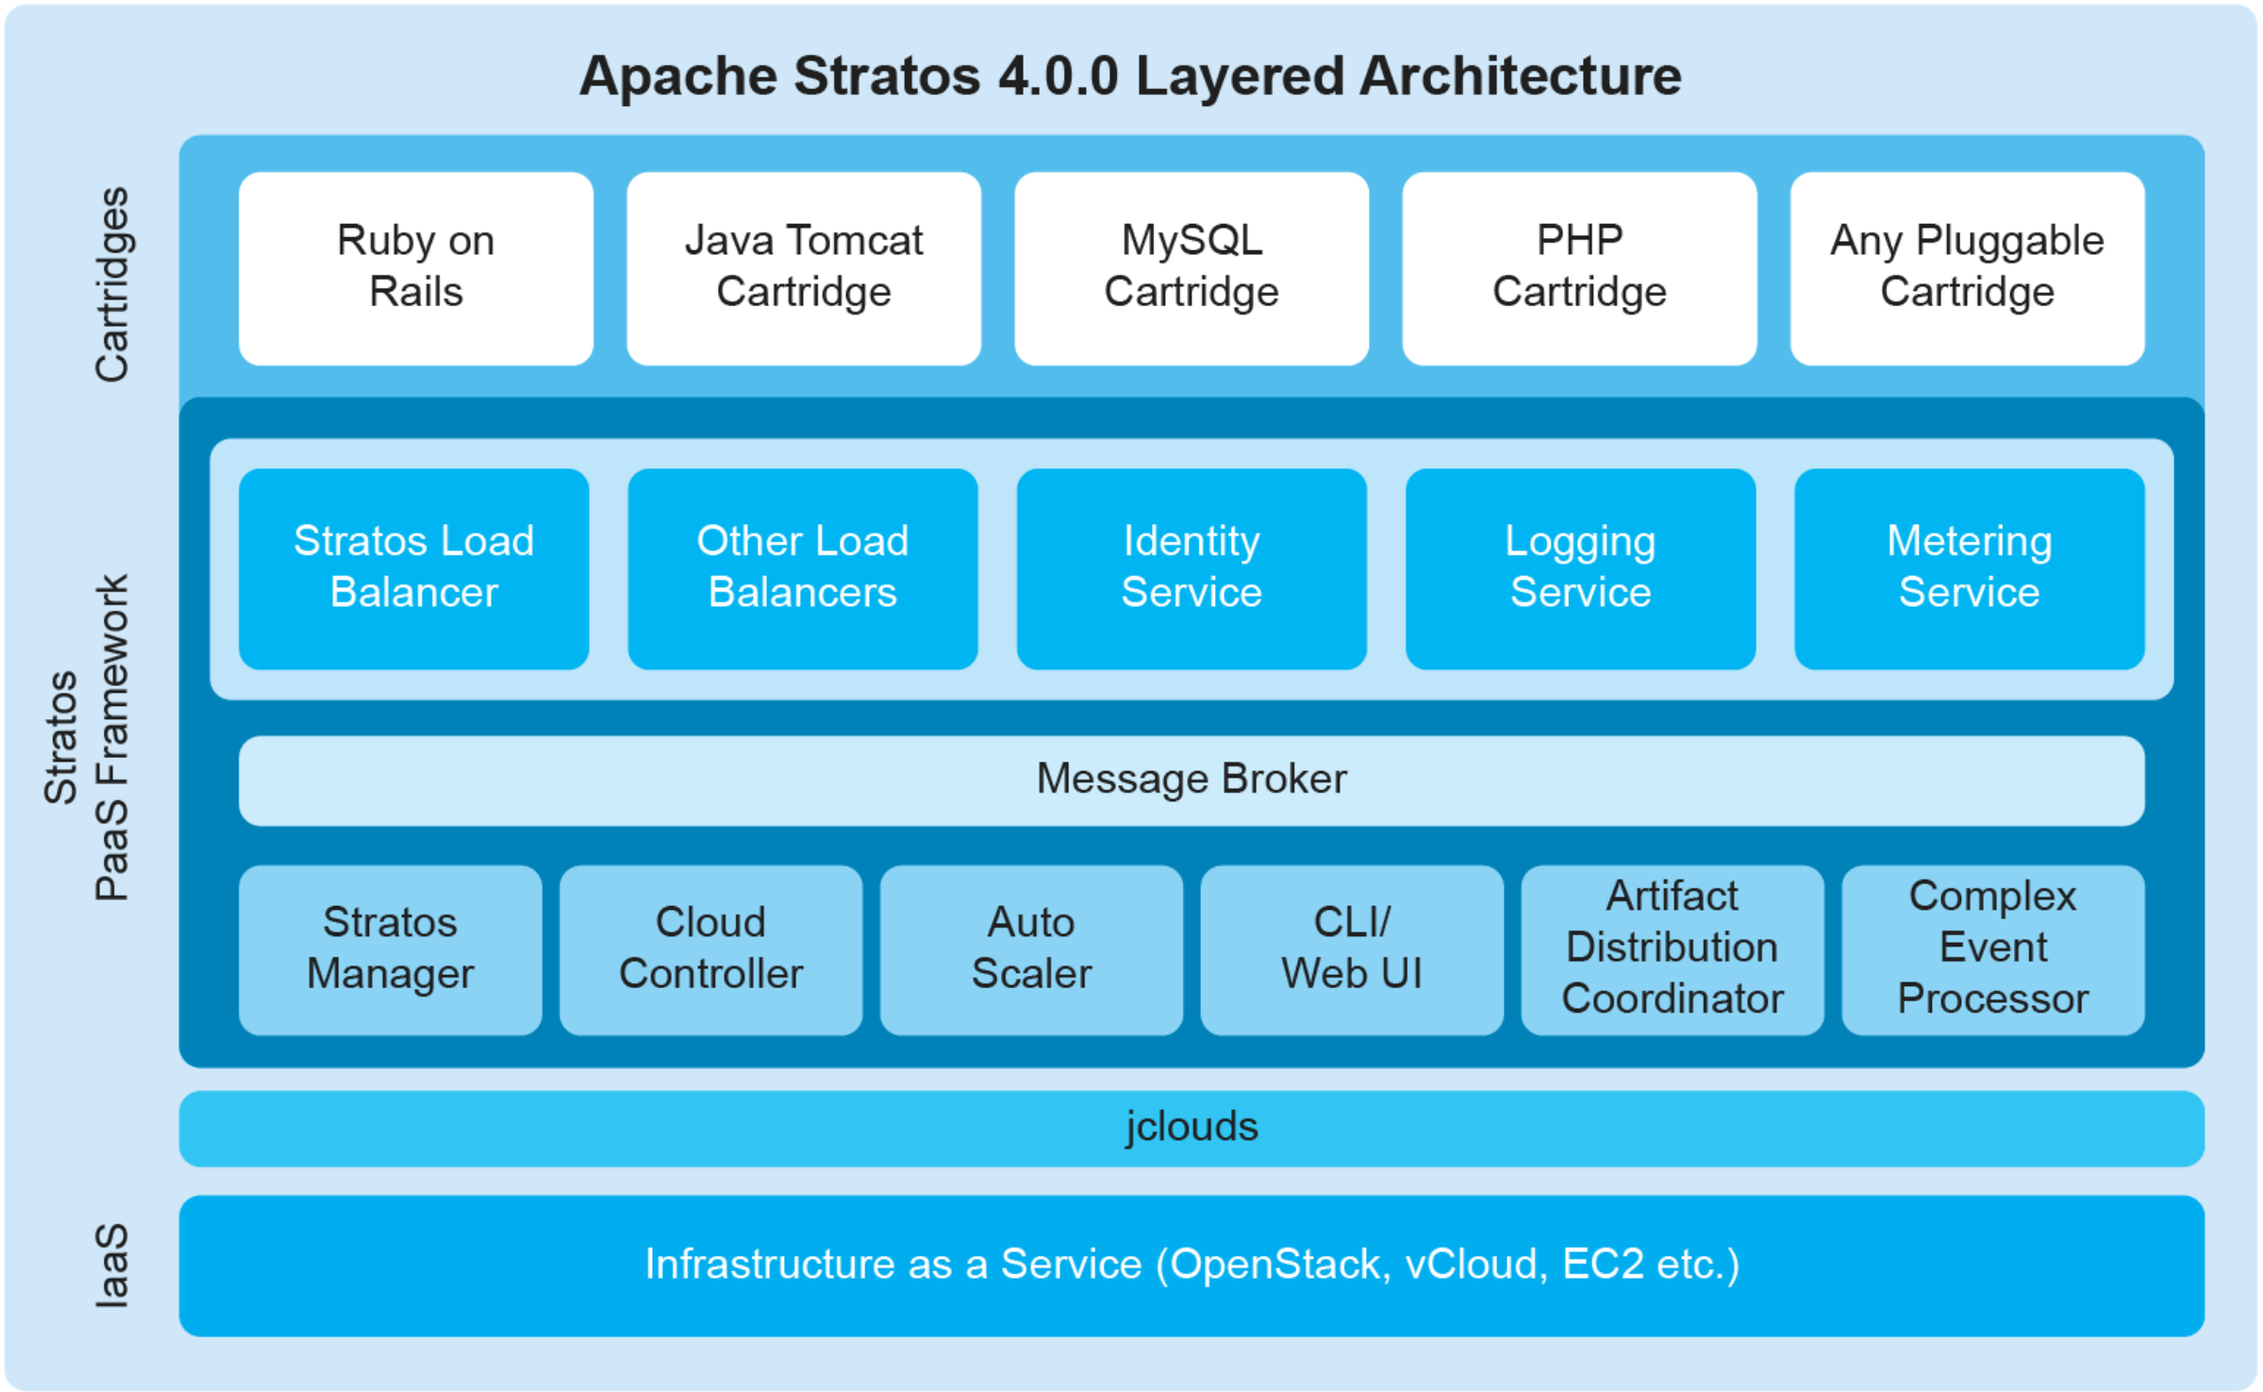

© WSO2

Apache Stratos L1 Architecture for VM based Cartridges

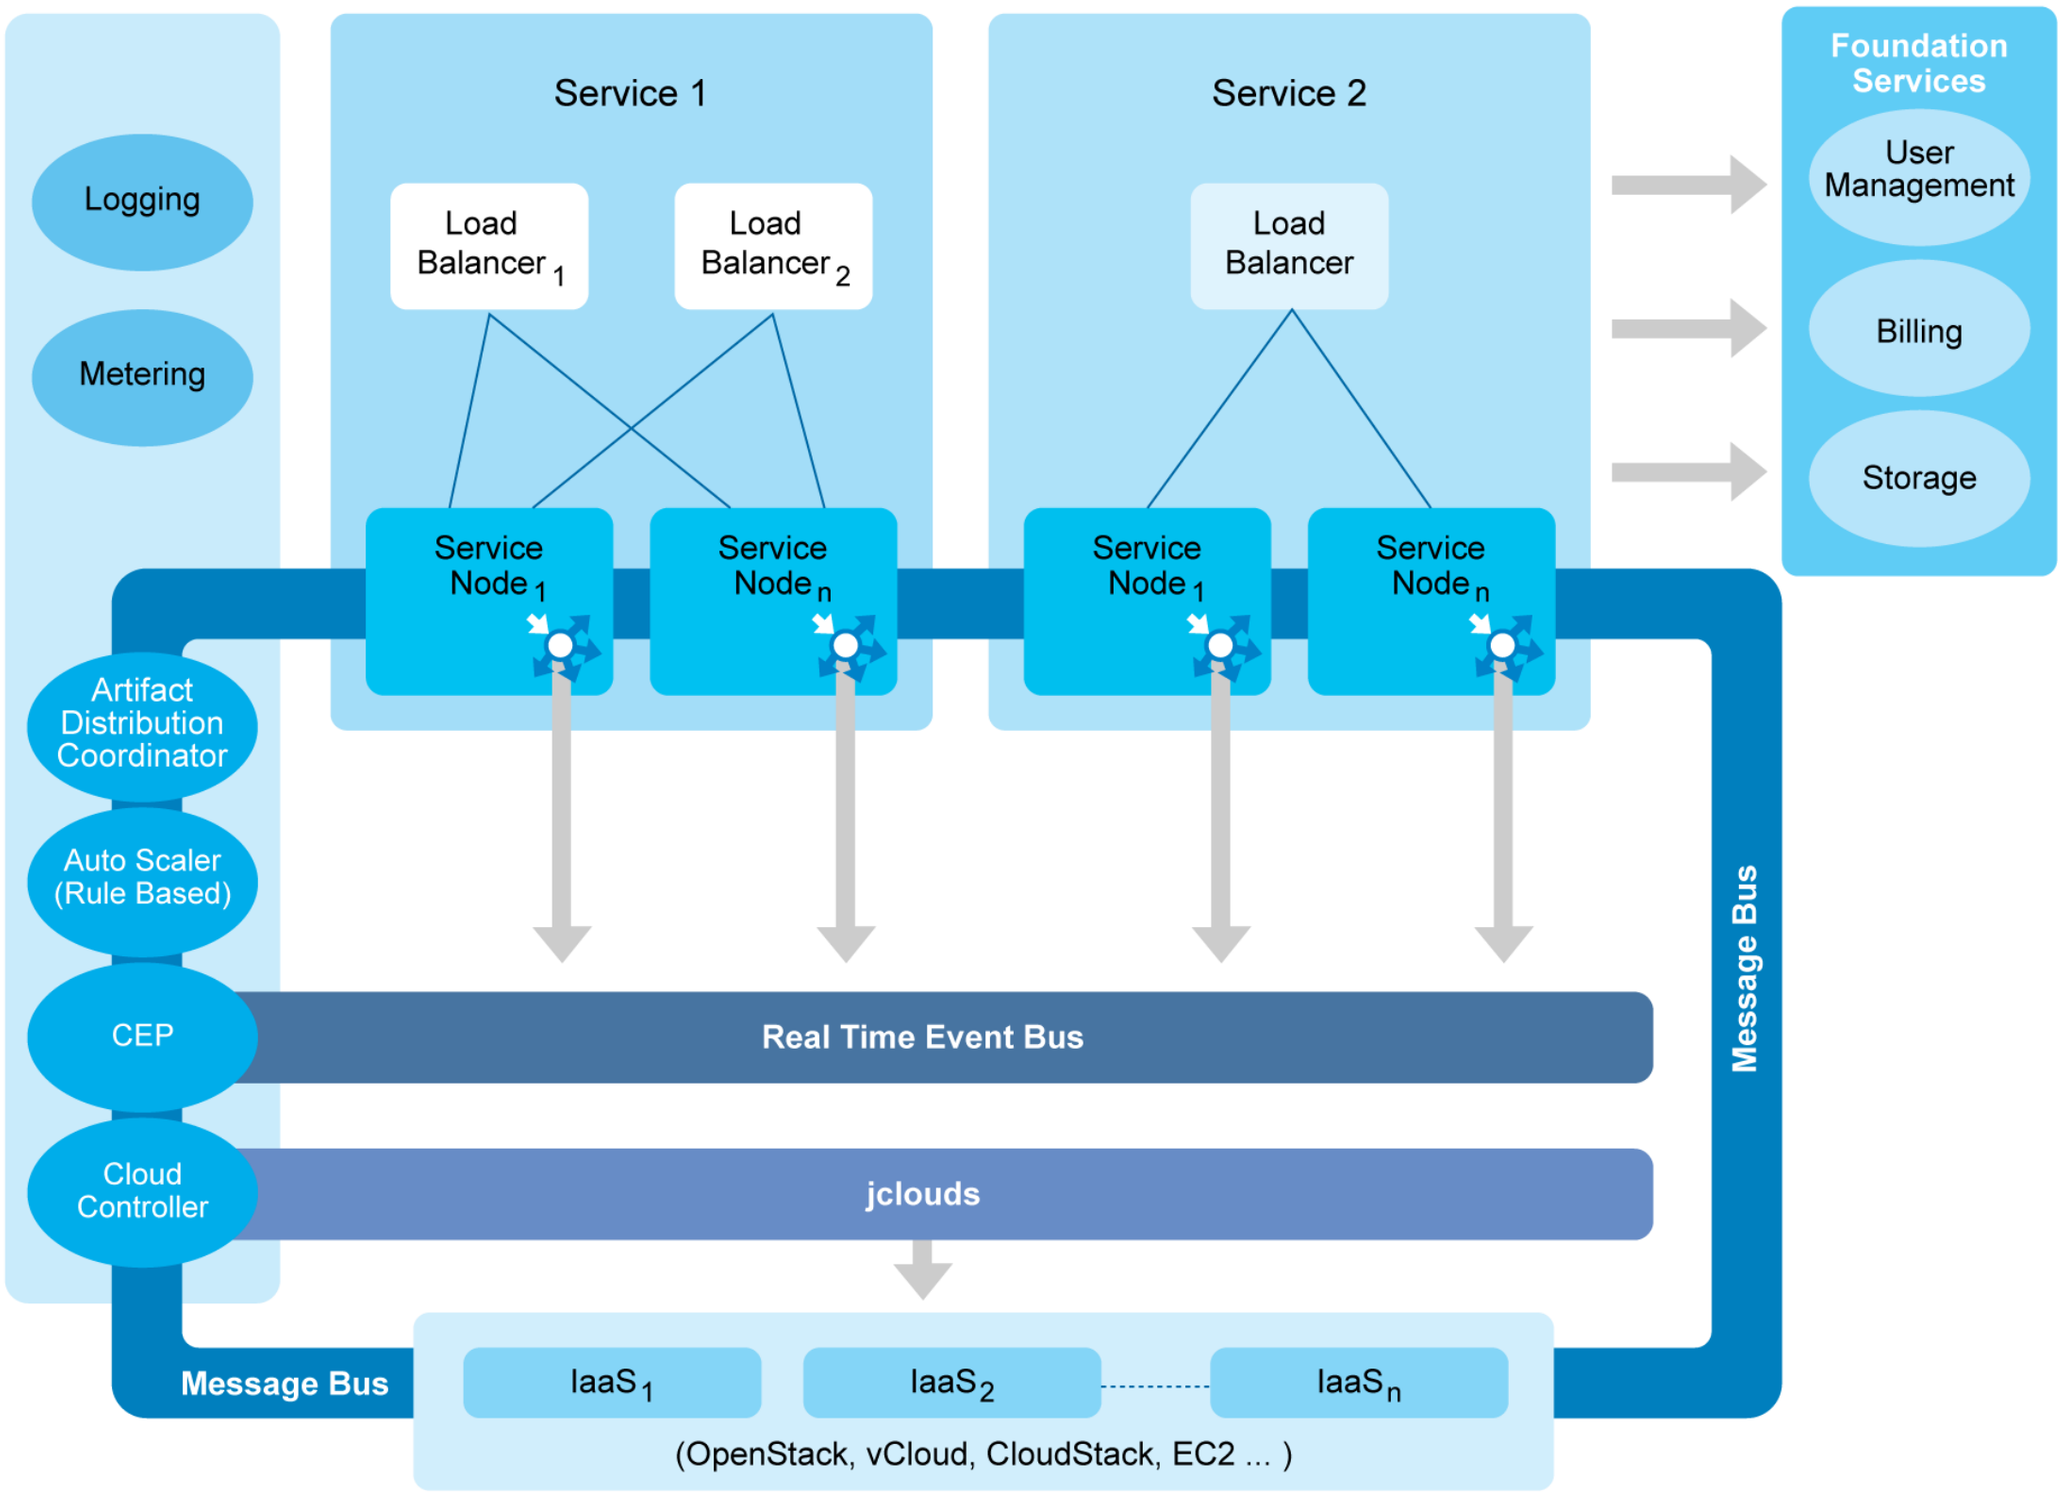

© WSO2

Apache Stratos Cartridges

Frameworks

- spring
- Struts
- Joomla!

App Servers

- Apache Tomcat 7
- JBoss
- APACHE HTTP SERVER
- WSO2 Application Server

Languages

- php
- Java
- Microsoft .net

Databases

- MySQL
- PostgreSQL
- Cassandra
- mongoDB

Legacy Apps

© WSO2

Multi-factored Auto Scaling

What is it?

- Scaling algorithm can use multiple factors. such as
  - Load average of the instance
  - Memory consumption of the instance
  - In-flight request count in LB

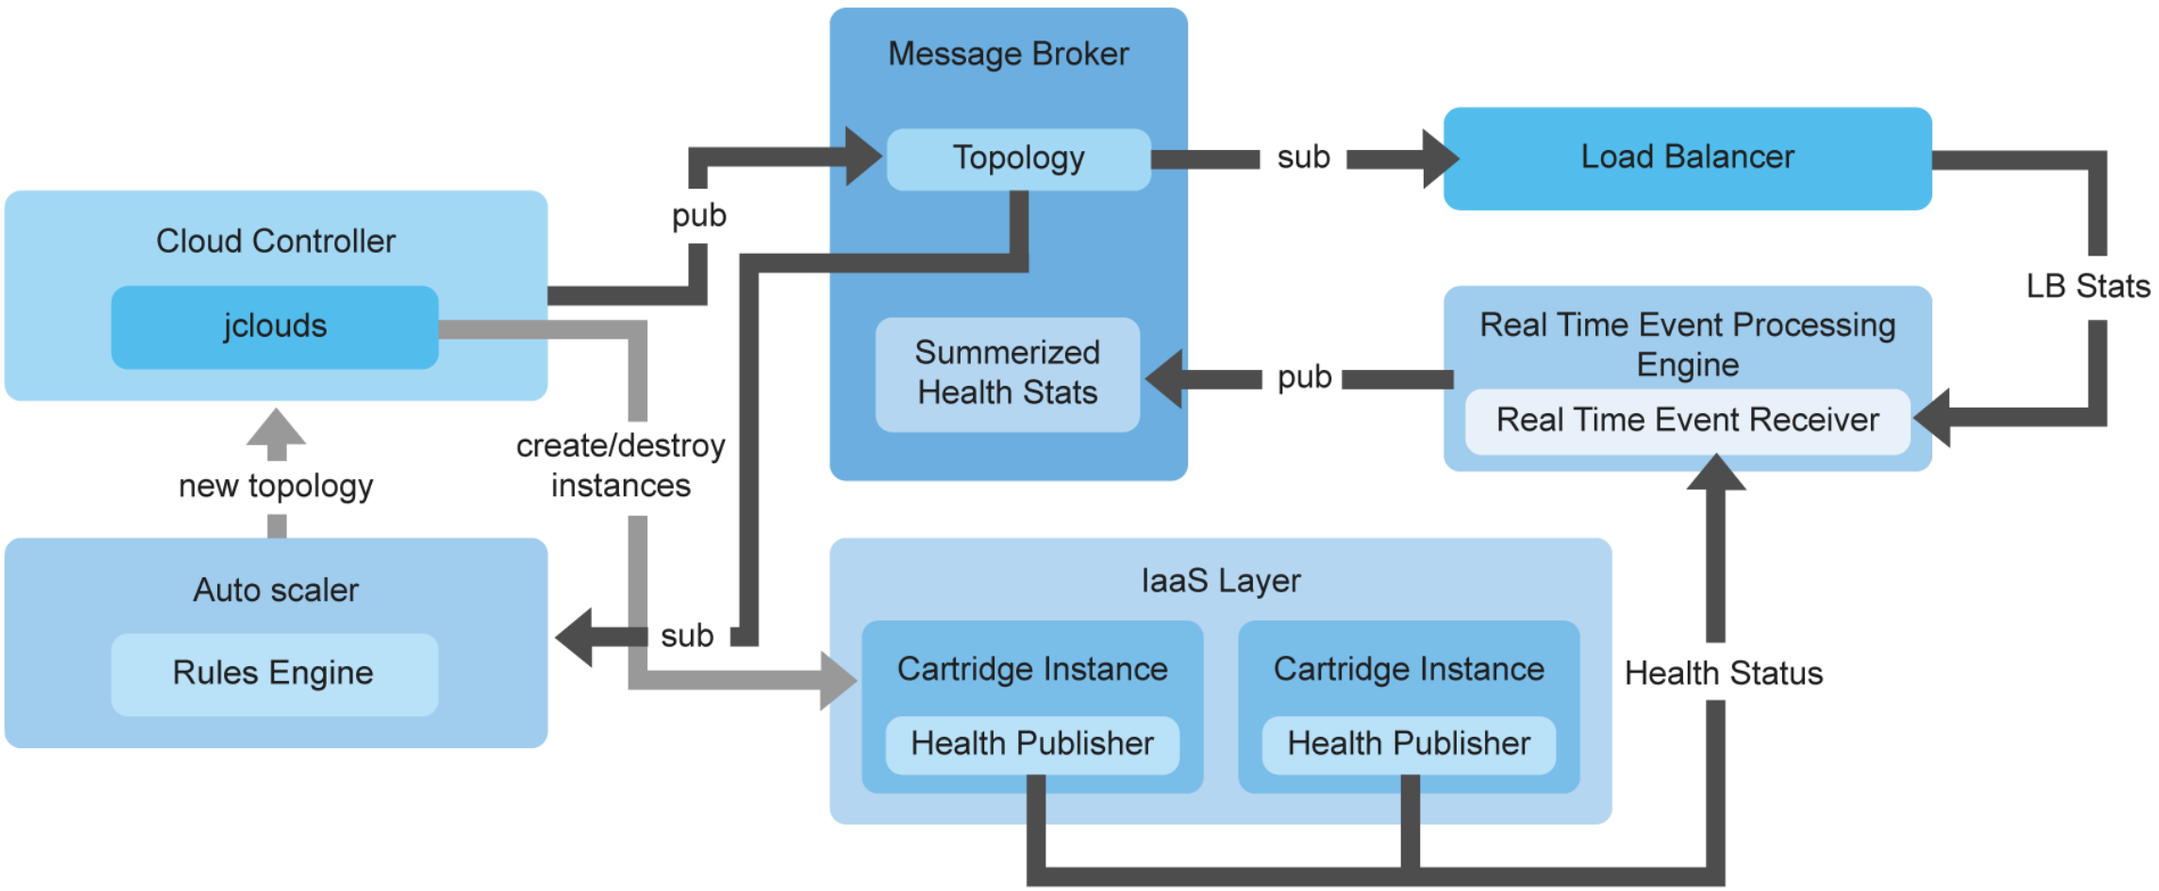

© WSO2

Multi-factored Auto Scaling...

- Capable of predicting future load
  - Real time analysis of current load status using CEP integration
  - Predict immediate future load based on CEP resulting streams
  - Predicting equation $s=ut + ½ at^2$
  - s=predicted load, u=first derivative of current average load, t= time interval , a=second derivative of current load

Why should one care?

- Maximise resource utilization
- Easy to do capacity planning
- Dynamic load based resource provisioning
- Optimizing across multiple clouds

© WSO2

Scalable and Dynamic Load Balancing

How Scalable it is?

- In theory infinite
  - horizontal scaling
  - limited by resource (instance capacity) availability

How Dynamic it is?

- Load Balancers are spawned dynamically
  - LB too is a cartridge
- In case of multi-cloud, multi-region, LB can scale per cloud/region
- Per service cluster LB

© WSO2

Scalable and Dynamic Load Balancing..

What is unique about Stratos

- Cartridge based LB model
- Can bring any third-party LB
  - HAProxy, nginx, AWS ELB
  - As easy as plugging into LB extension API

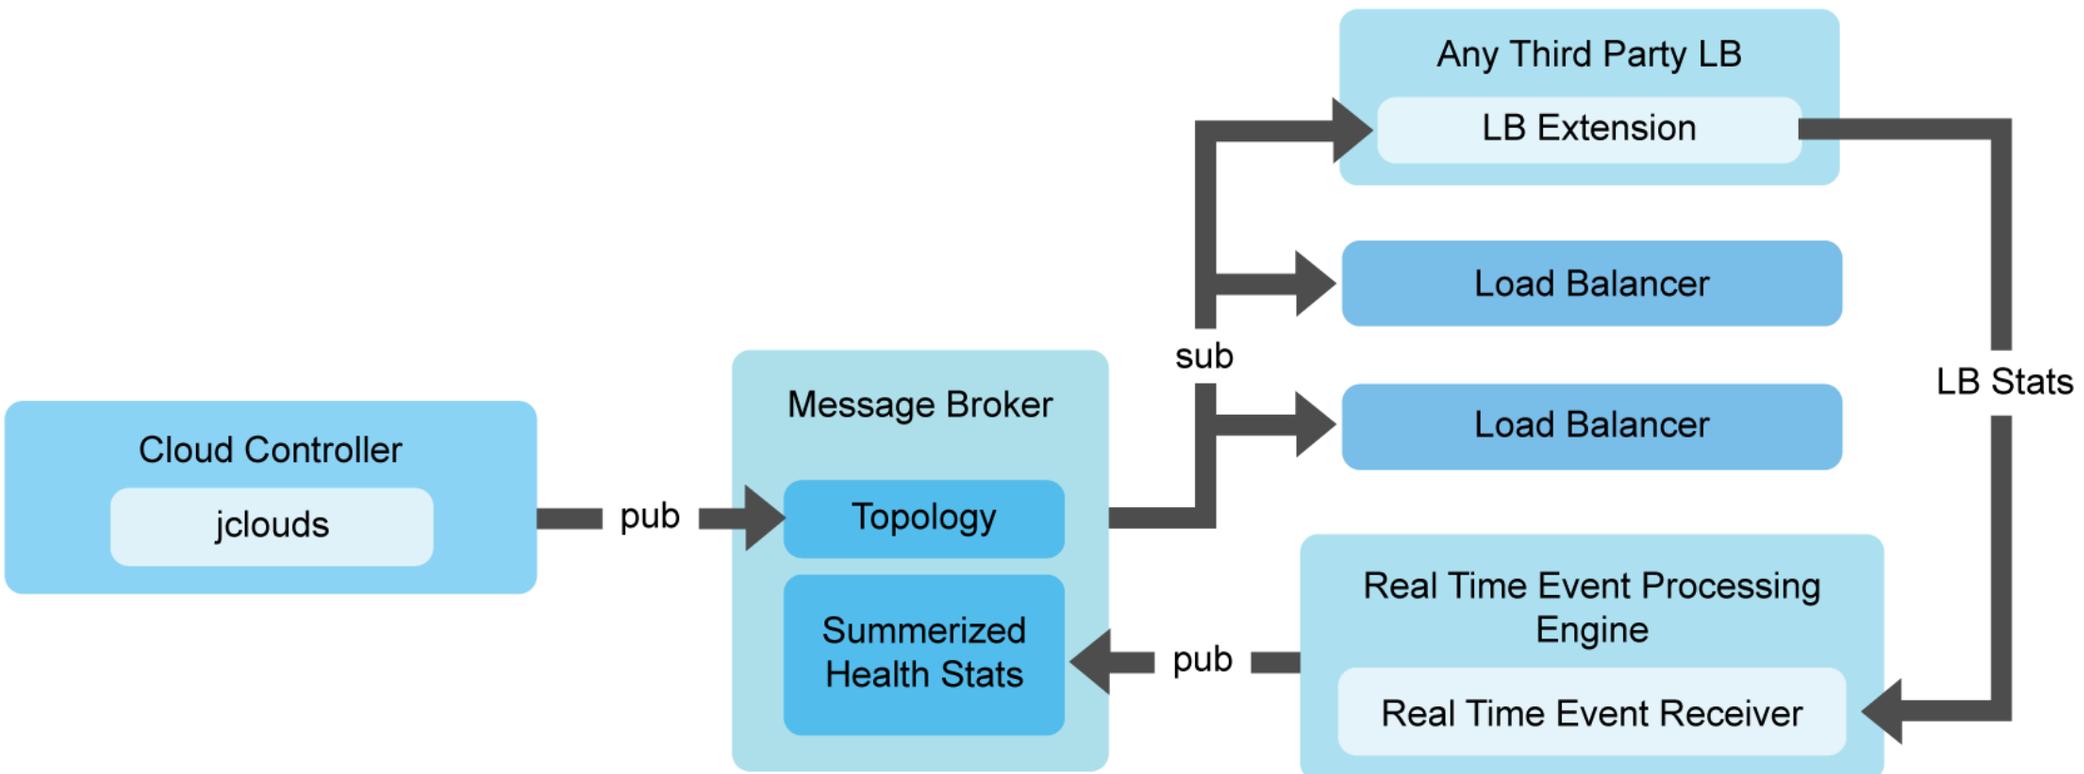

© WSO2

Smart Policies

What are the smart policies?

- Auto scaling
- Deployment

Auto scaling policy

- Define thresholds values pertaining scale up/down decision
- Auto Scaler refer this policy
- Defined by DevOps

Deployment policy

- Defined how and where to spawn cartridge instances
- Defined min and max instances in a selected service cluster
- Defined by DevOps based on deployment patterns

© WSO2

Smart Policies

Why should one care?

- Can provide cloud SLA

What are the advantages?

- Make DevOps life easy
  - help keep to SLA
- Make SaaS app delivery life easy
  - do not have to worry about availability in application layer

© WSO2

Multi-tenancy

What MT model does it support?

- Container MT
  - virtual Machine, LXC, Docker
- In-container MT
  - within VM/LXC/Docker tenancy

What is unique?

- Can have high tenant density

What are the advantage of this model?

- Optimizing resource utilization
  - by sharing resource such as CPU, memory across tenants
  - low footprint, based on utilization/usage of the tenants app
- No need dedicated resource allocation for tenants

© WSO2

Cloud Bursting

What is it?

- Expanding/provisioning application into another cloud to handle peak load.

Why Should one care?

- Resource peak time can be off-loaded to third party clouds/resources

What is unique about it?

- Can off-load to any cloud
  - Private, Public and Hybrid
- Easy to managed with the model of LB per busting cloud

© WSO2

Cloud Bursting...

What are the advantages?

- Make DevOps life easy
- Low TCO, and higher utilization existing dedicated resources

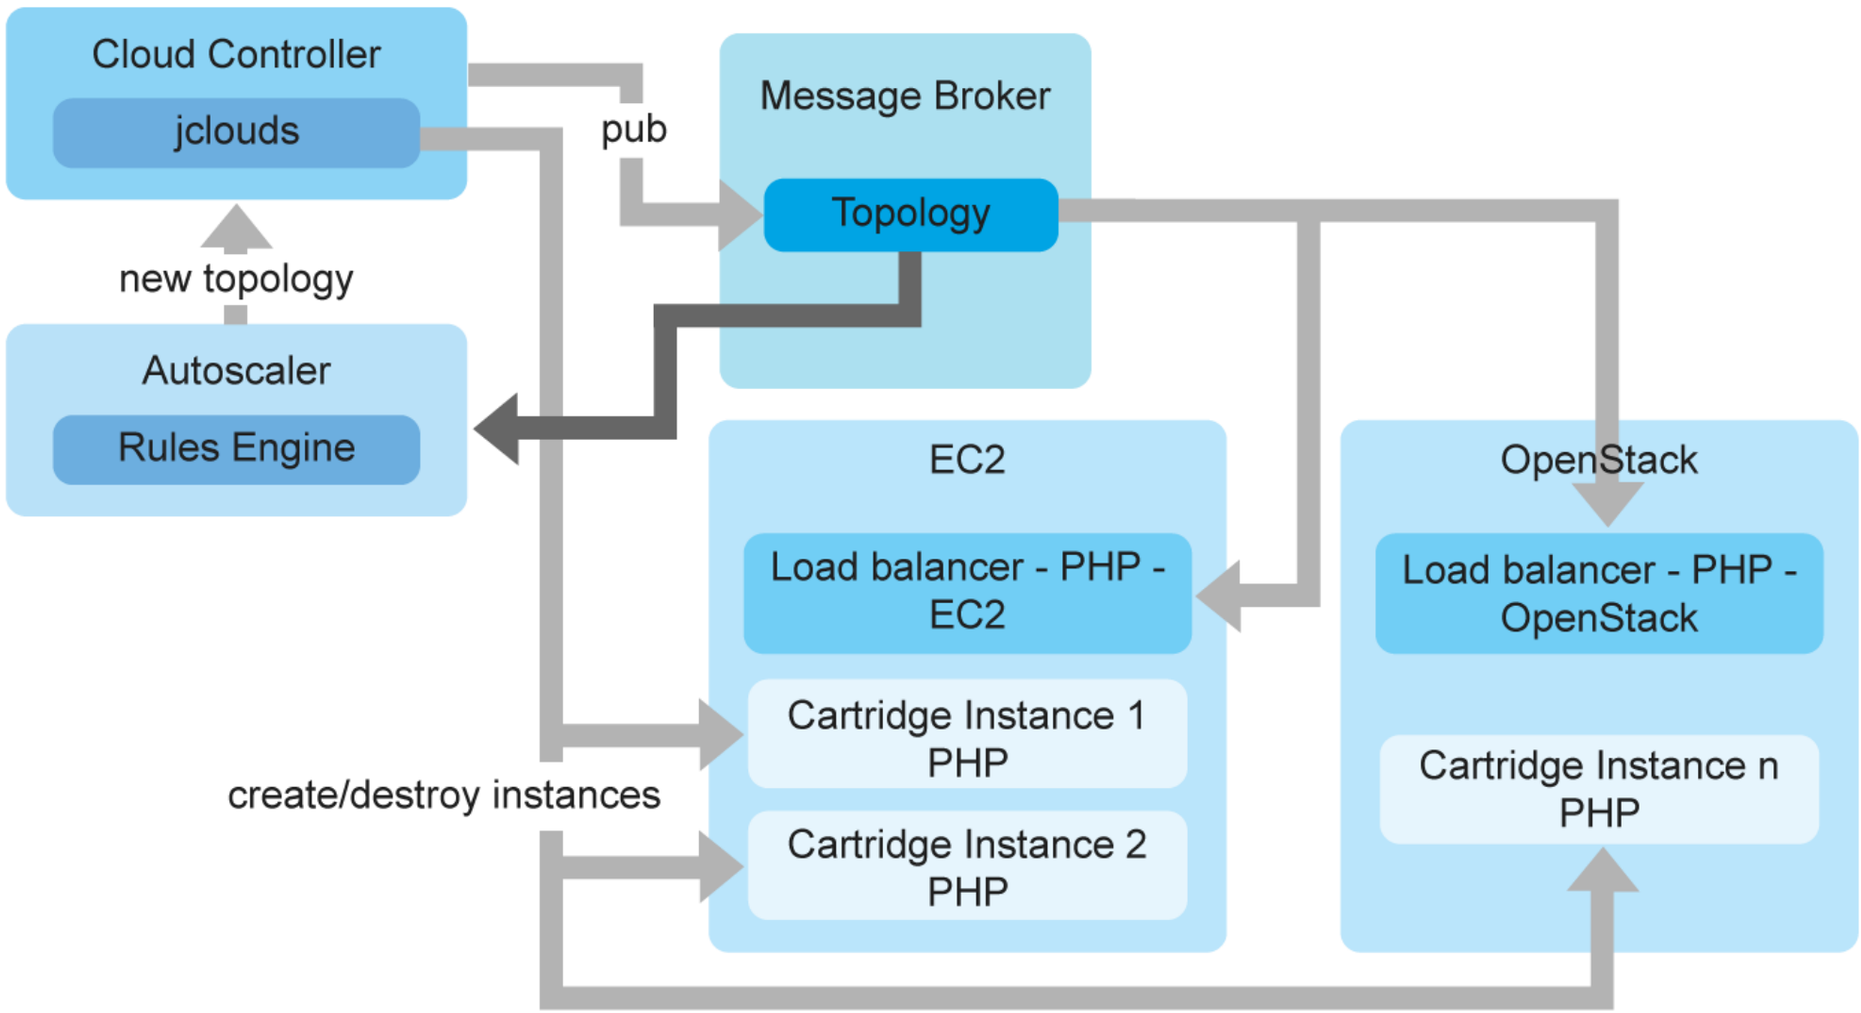

© WSO2

Logging, Metering and Monitoring

What details are?

- Instance up/down time
- Each and every instances health status
  - application health, load average, memory consumption
- Application logs

Why should one care?

- Centralize view for all logging, metering and monitoring

What are the advantages?

- DevOps life easy
  - centralize log viewer
  - centralize dashboard
- Easy to throttling

© WSO2

Apache Stratos with Docker

- Why just not with only Docker?
- Apache Stratos next release is mainly into
  - Docker based cartridge support
  - integration with CoreOS
  - integration with Kubernetes
  - integration with flannel
  - integration with discovery service
  - build in docker registry

© WSO2

What is Containers?

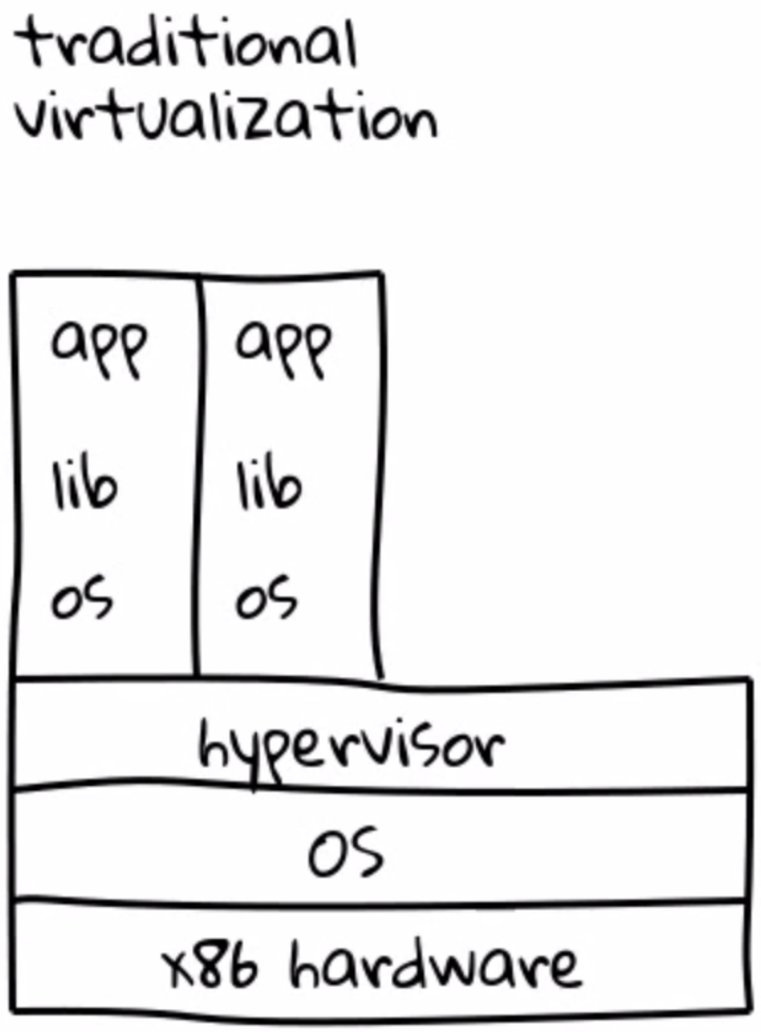

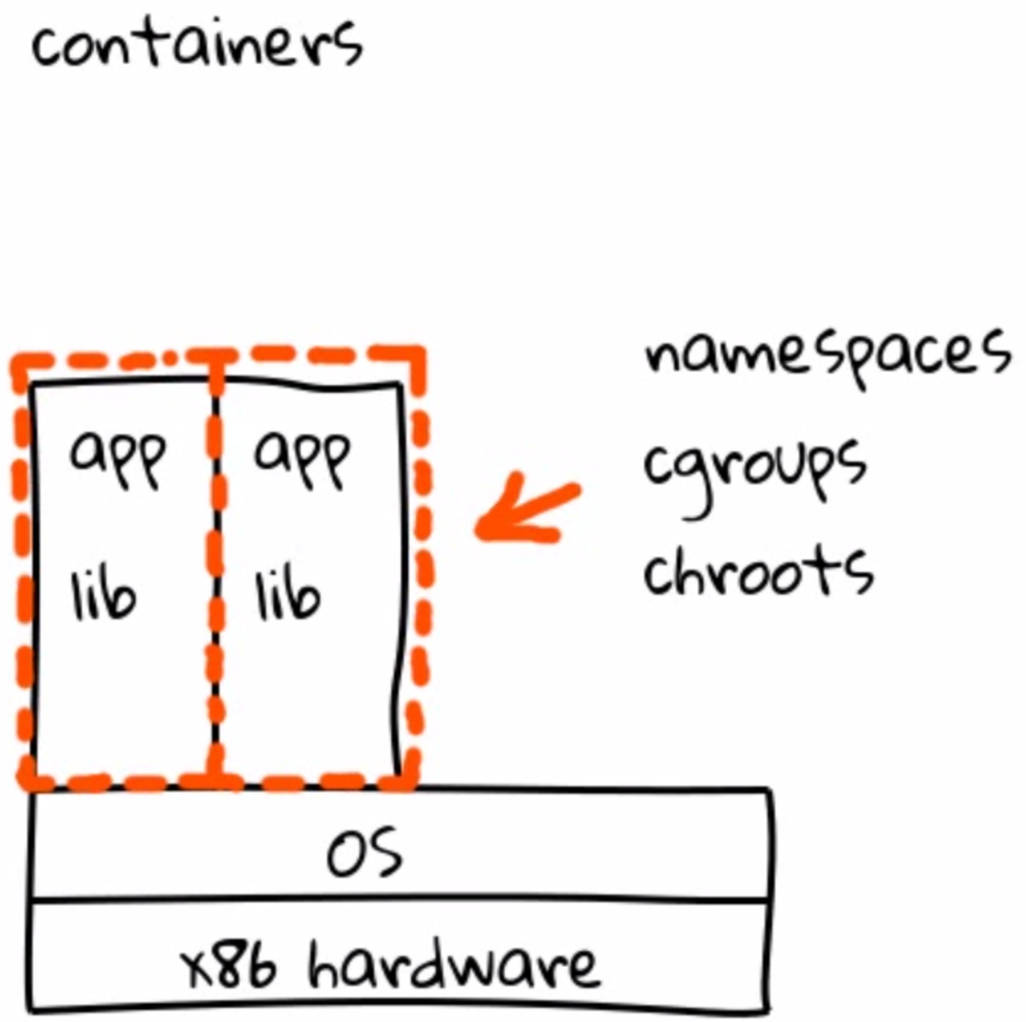

© WSO2

What is Docker?

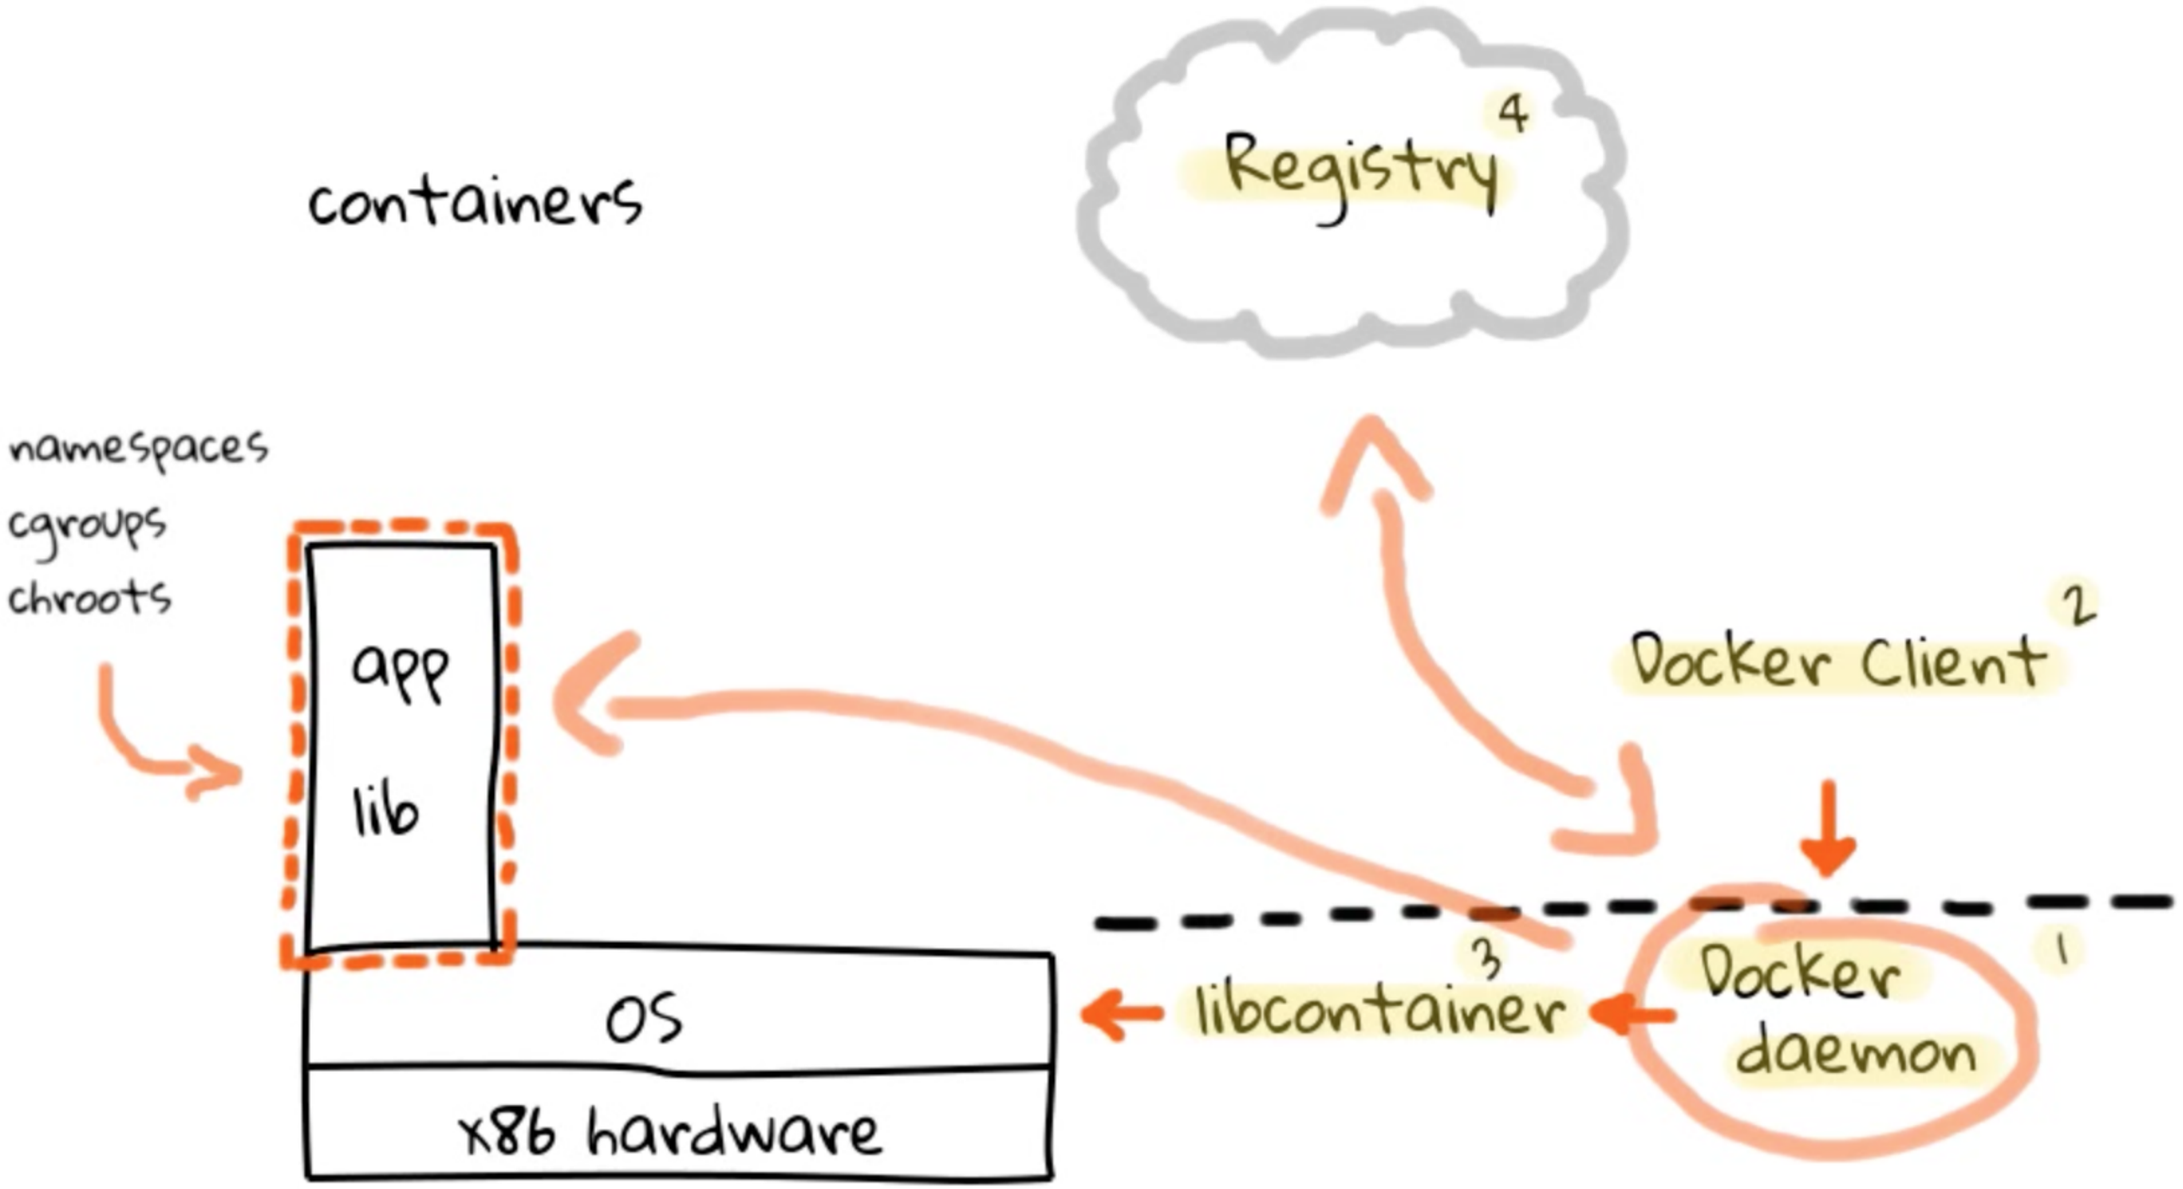

© WSO2

Docker Filesystem

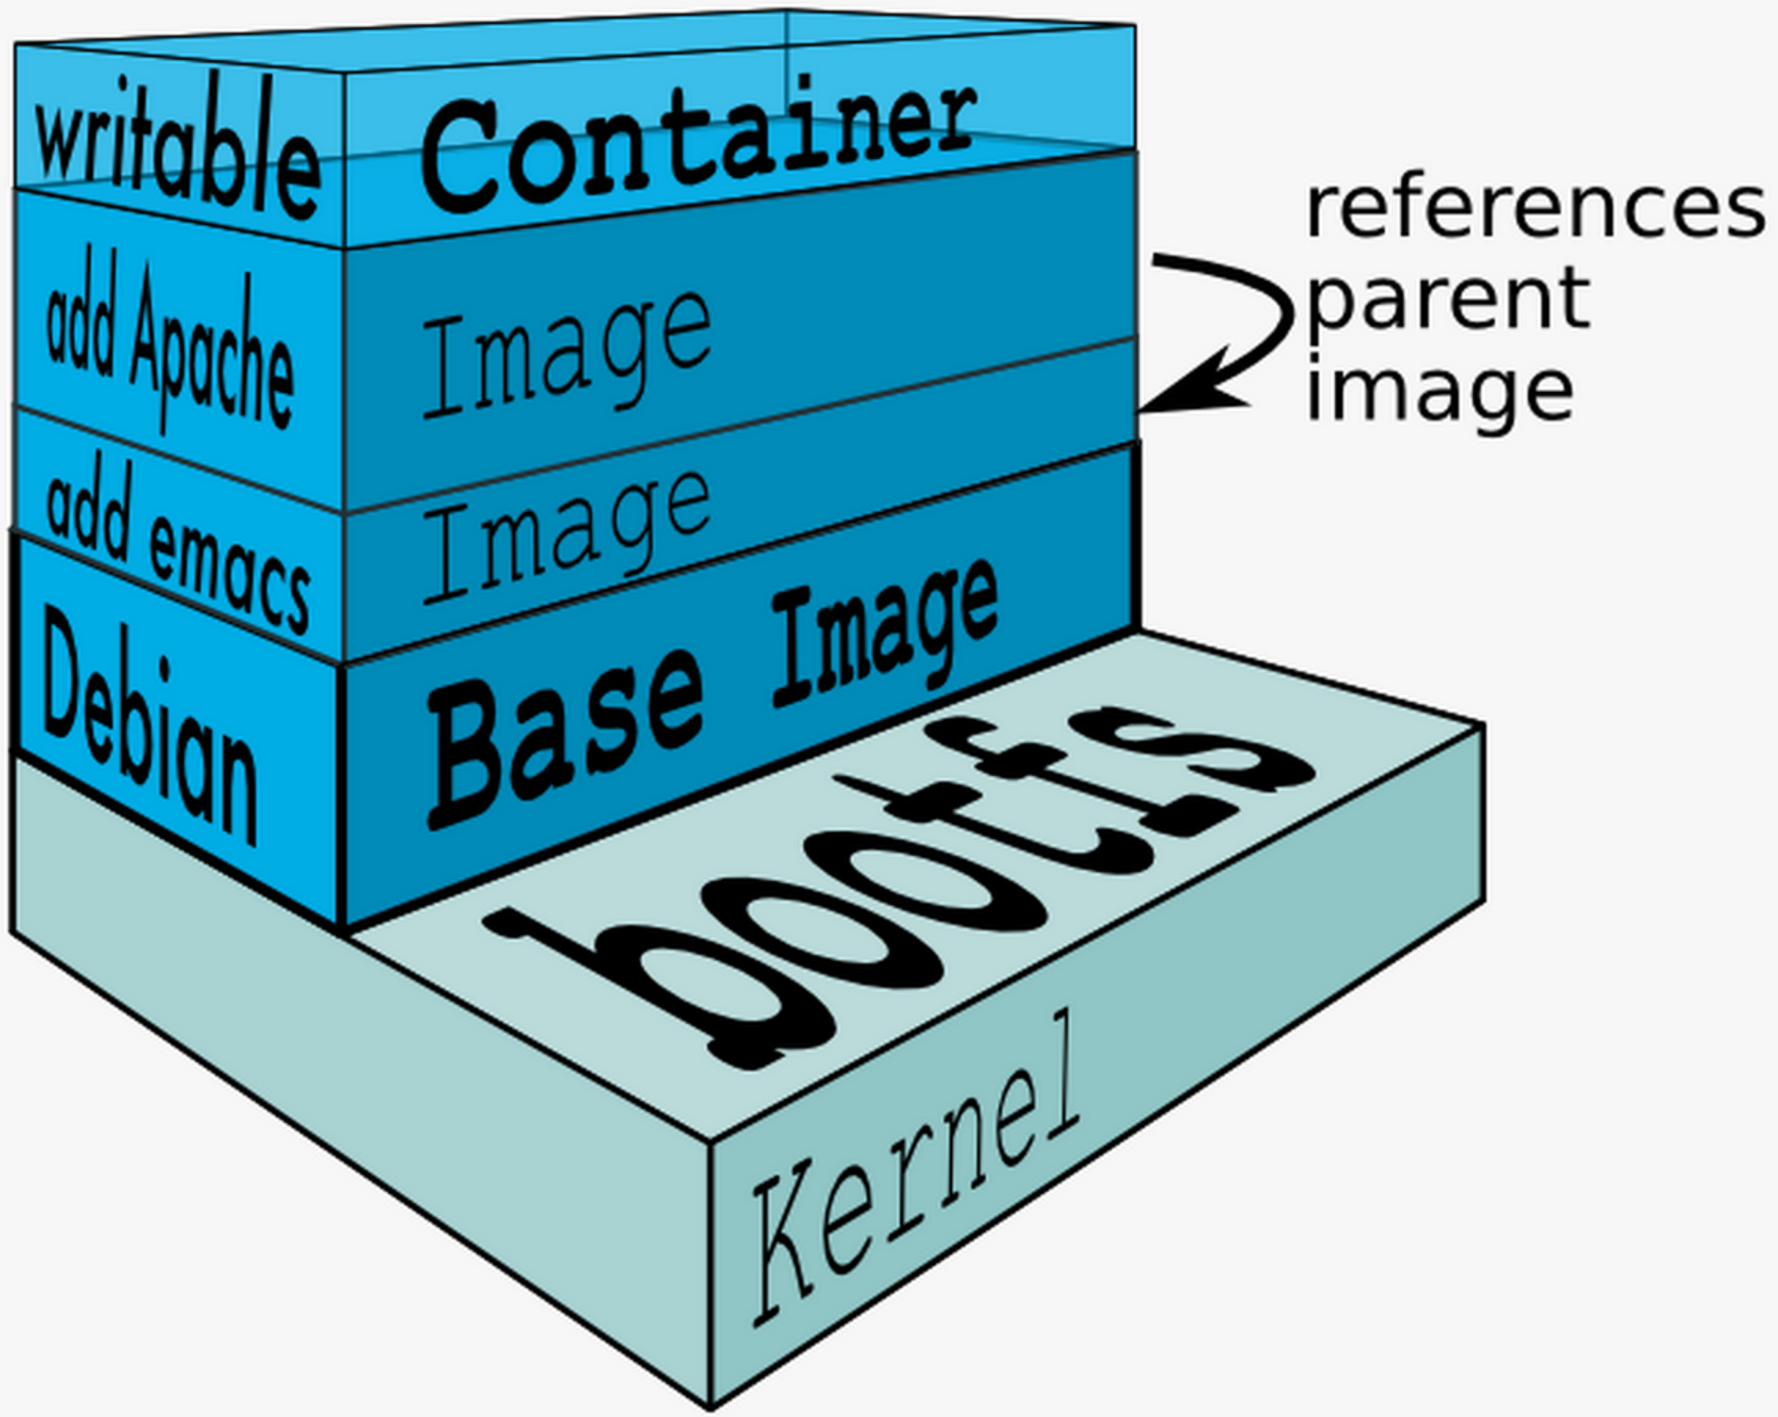

© WSO2

What is CoreOS?

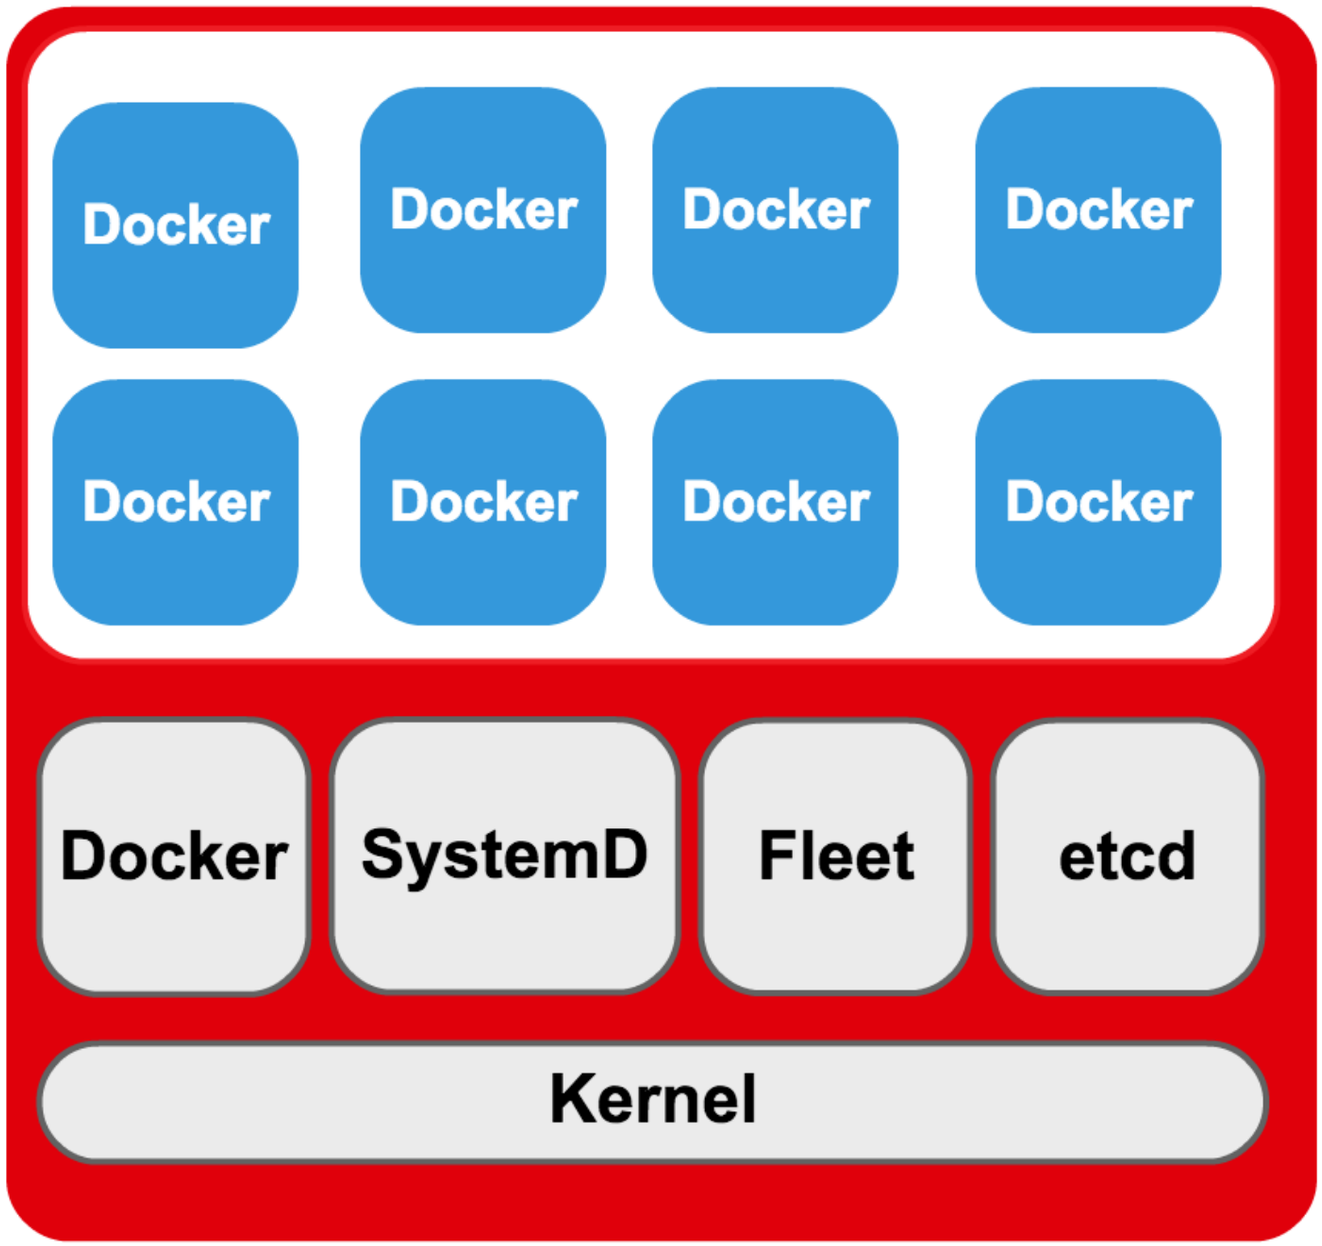

CoreOS Host

© WSO2

CoreOS Cluster

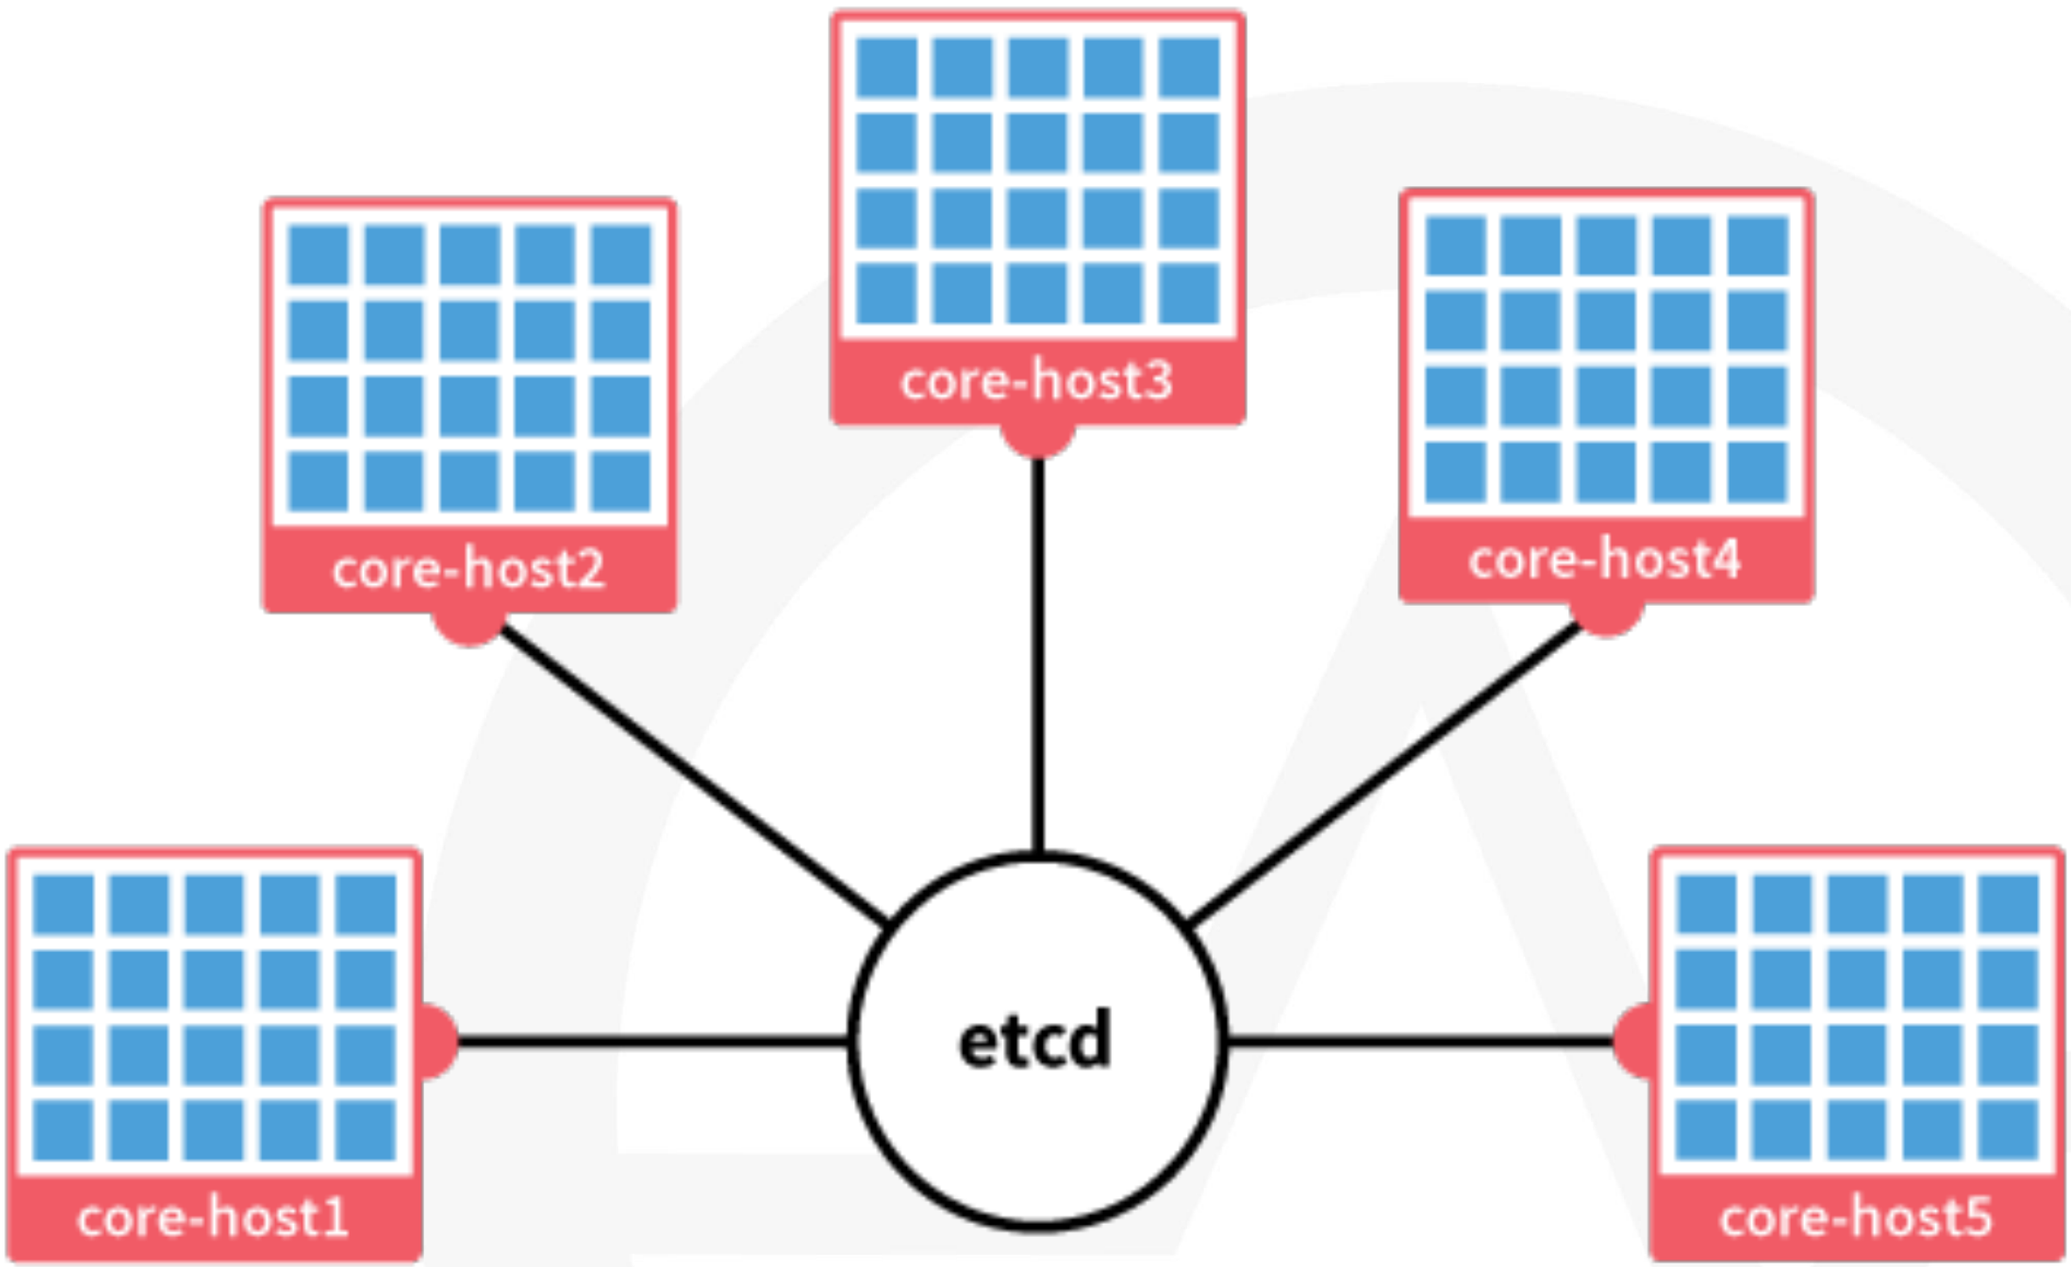

© WSO2

What is Kubernetes?

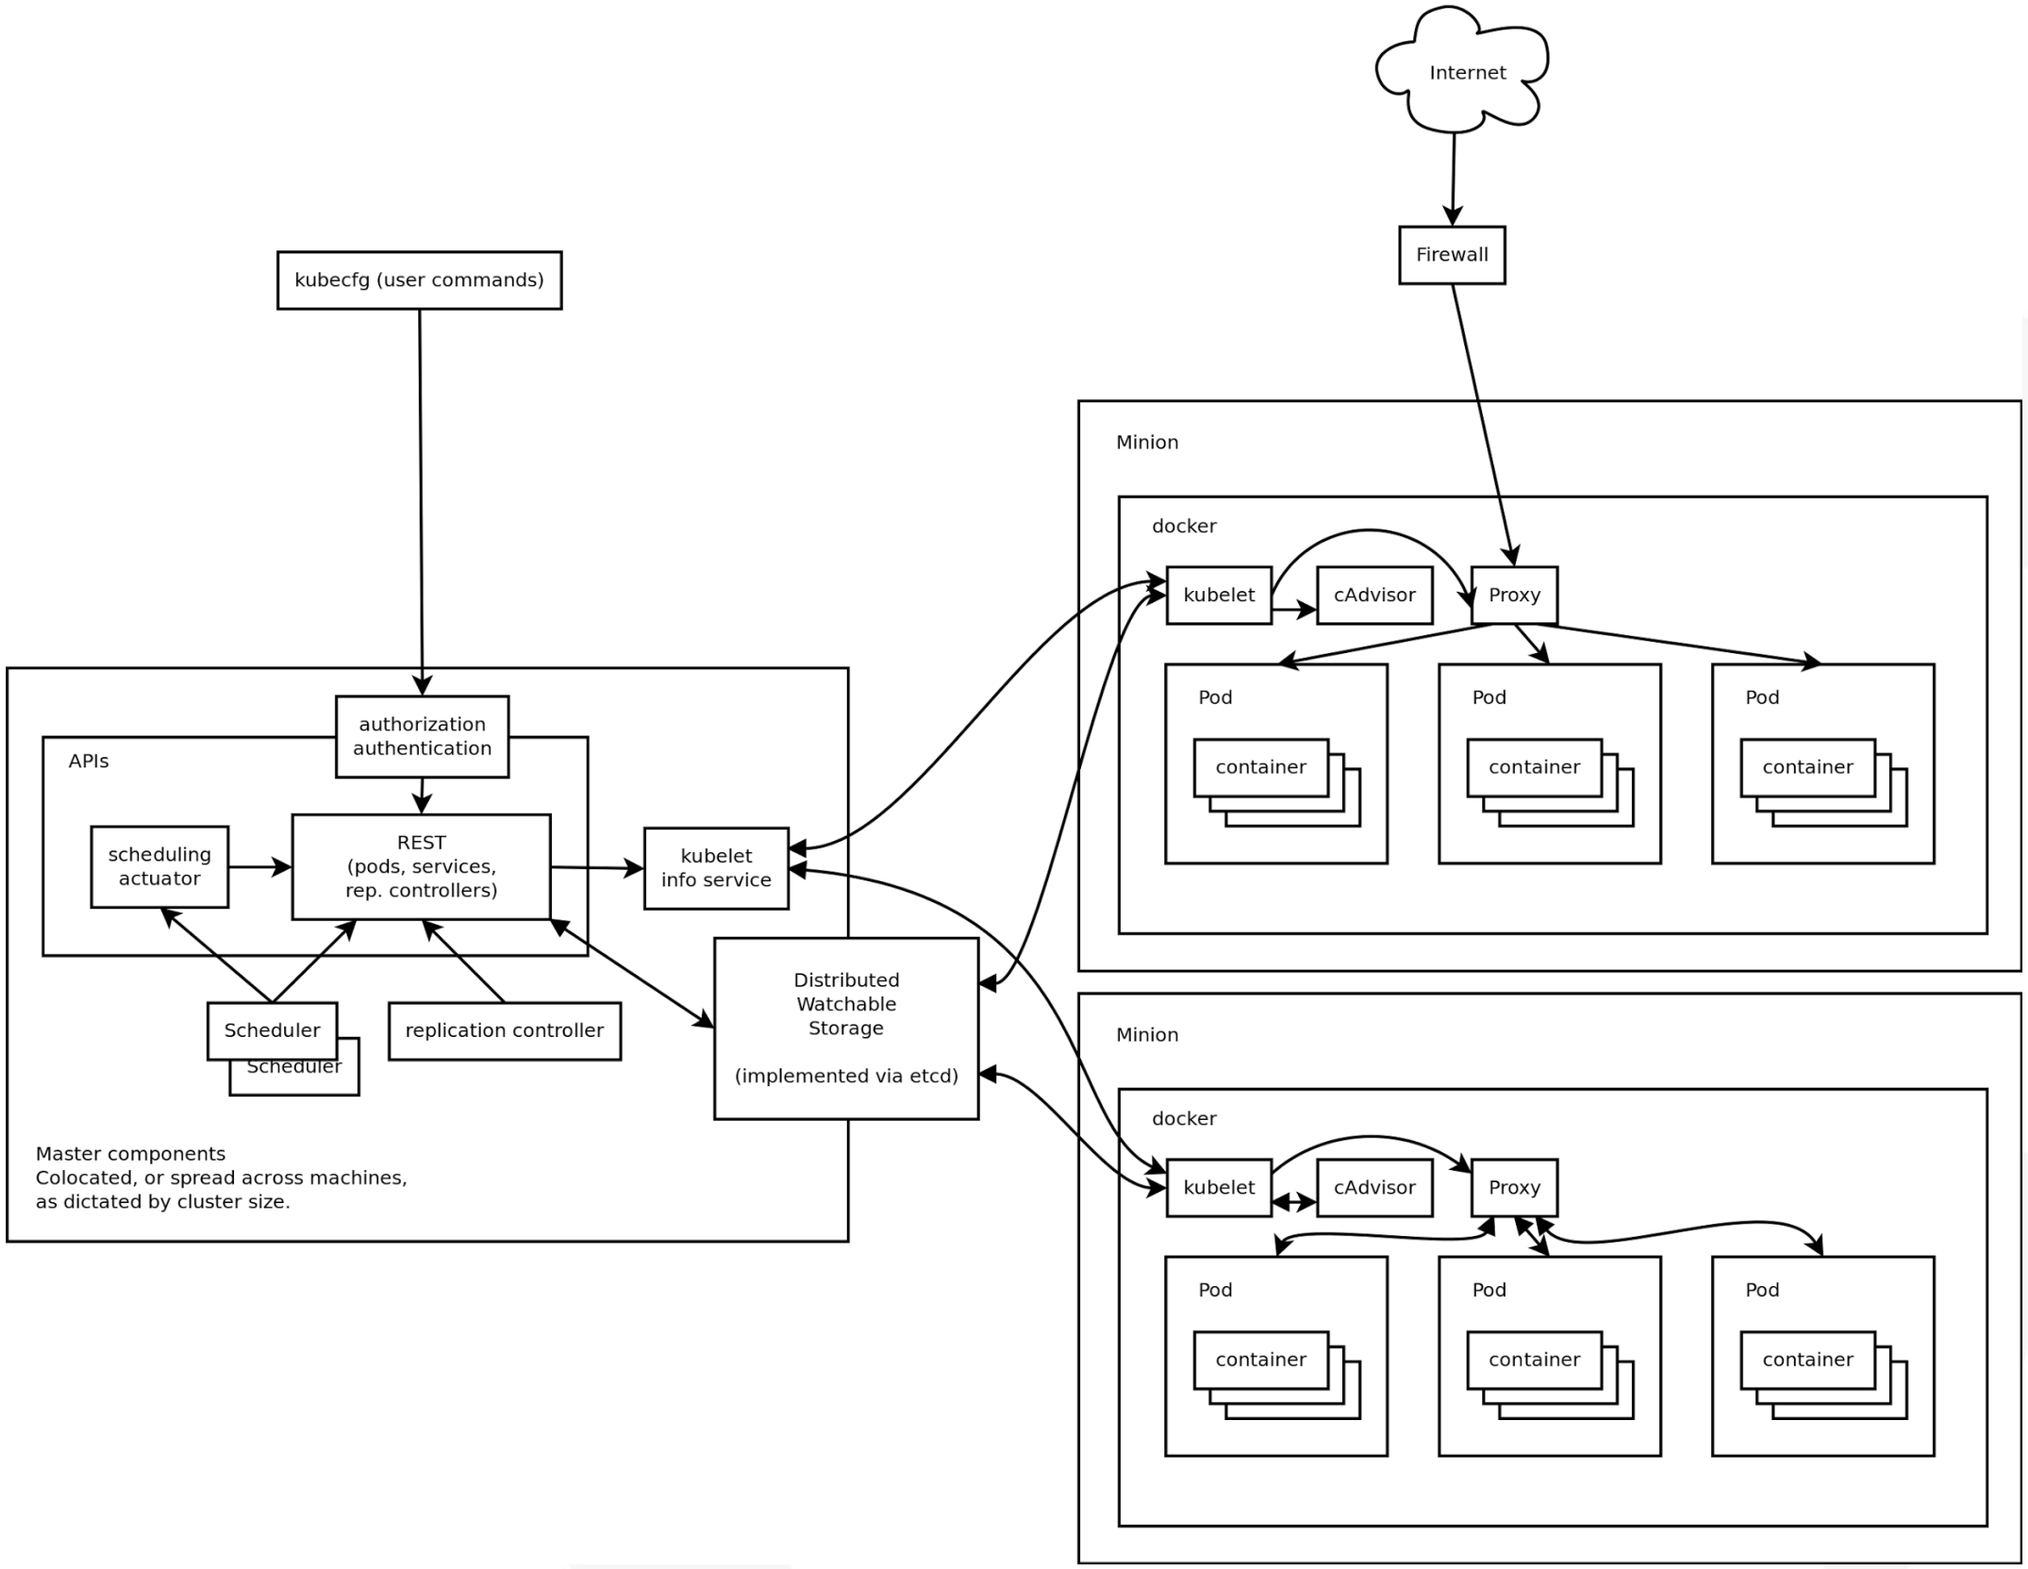

© WSO2

Kubernetes with CoreOS

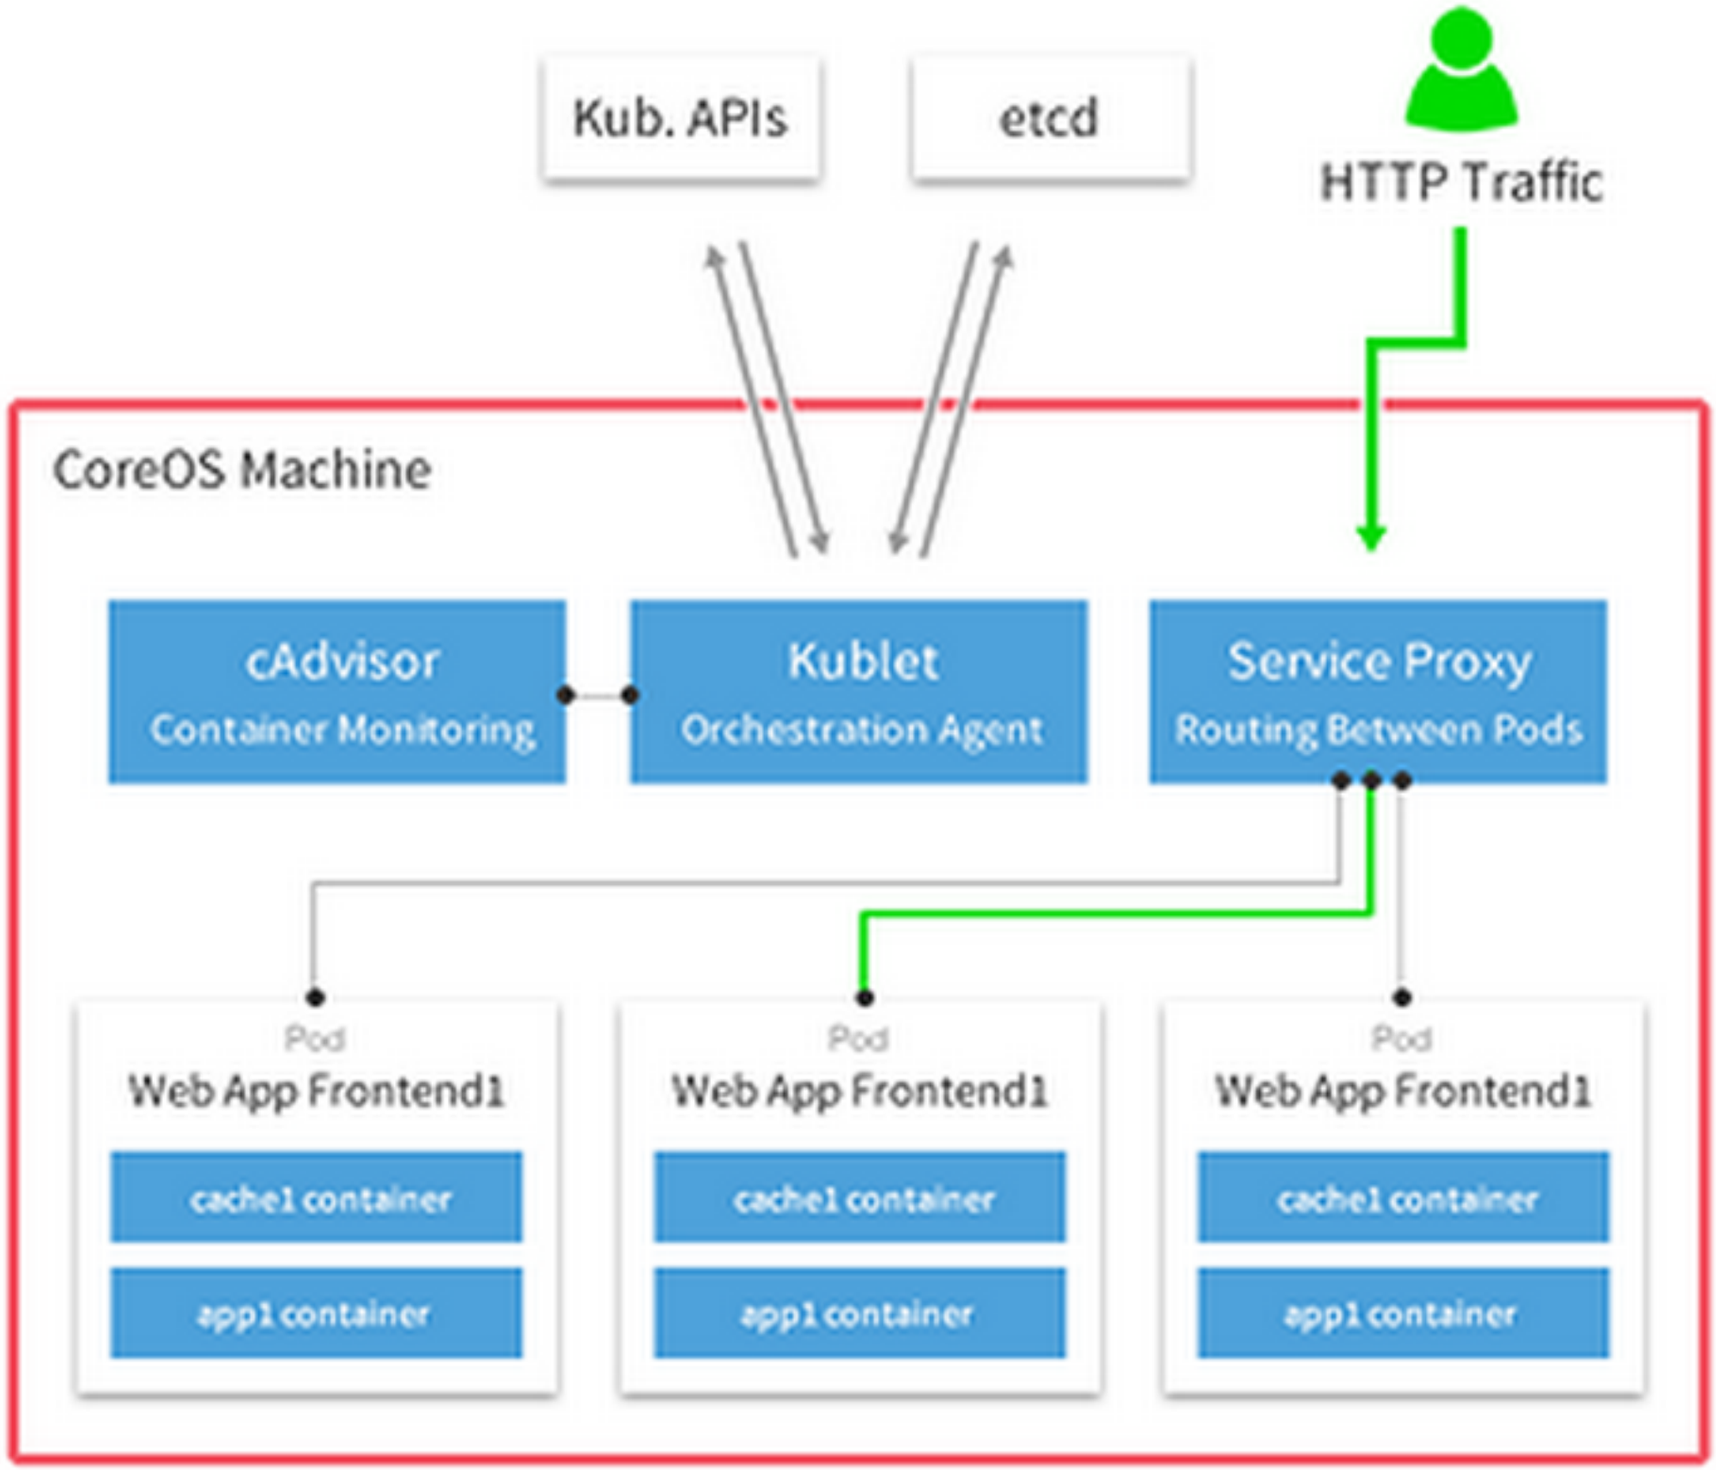

© WSO2

Kubernetes is spreading...

- **Microsoft** is working to ensure that Kubernetes works great in Linux environments in Azure VMs. Scott Guthrie, Executive Vice President of the Cloud and Enterprise group at Microsoft told us, “Microsoft will help contribute code to Kubernetes to enable customers to easily manage containers that can run anywhere. This will make it easier to build multi-cloud solutions including targeting Microsoft Azure.”
- **Red Hat** is working to bring Kubernetes to the open hybrid cloud. Paul Cormier, President, Products and Technologies at Red Hat, told us, “Red Hat has a rich history of contributing to and maturing innovative, open source projects. Through this collaboration with Google on Kubernetes, we are contributing to the evolution of cloud computing and helping deliver the promises that container technologies offer to the open hybrid cloud.”
- **IBM** is contributing code to Kubernetes and the broader Docker ecosystem to ensure that containers are enterprise-grade, and is working with the community to create an open governance model around the project.
- **Docker** is delivering the full container stack that Kubernetes schedules into, and is looking to move critical capabilities upstream and align the Kubernetes framework with Libswarm.
- **CoreOS** is working to ensure that Kubernetes can work seamlessly with the suite of CoreOS technologies that support cloud-native application development on any cloud.
- **Mesosphere** is actively integrating Kubernetes with Mesos, making the advanced scheduling and management capabilities available to Kubernetes customers.
- **SaltStack** is working to make Kubernetes a portable container automation framework that is designed for the reality of the platform-agnostic, multi-cloud world.

© WSO2

What is Flannel?

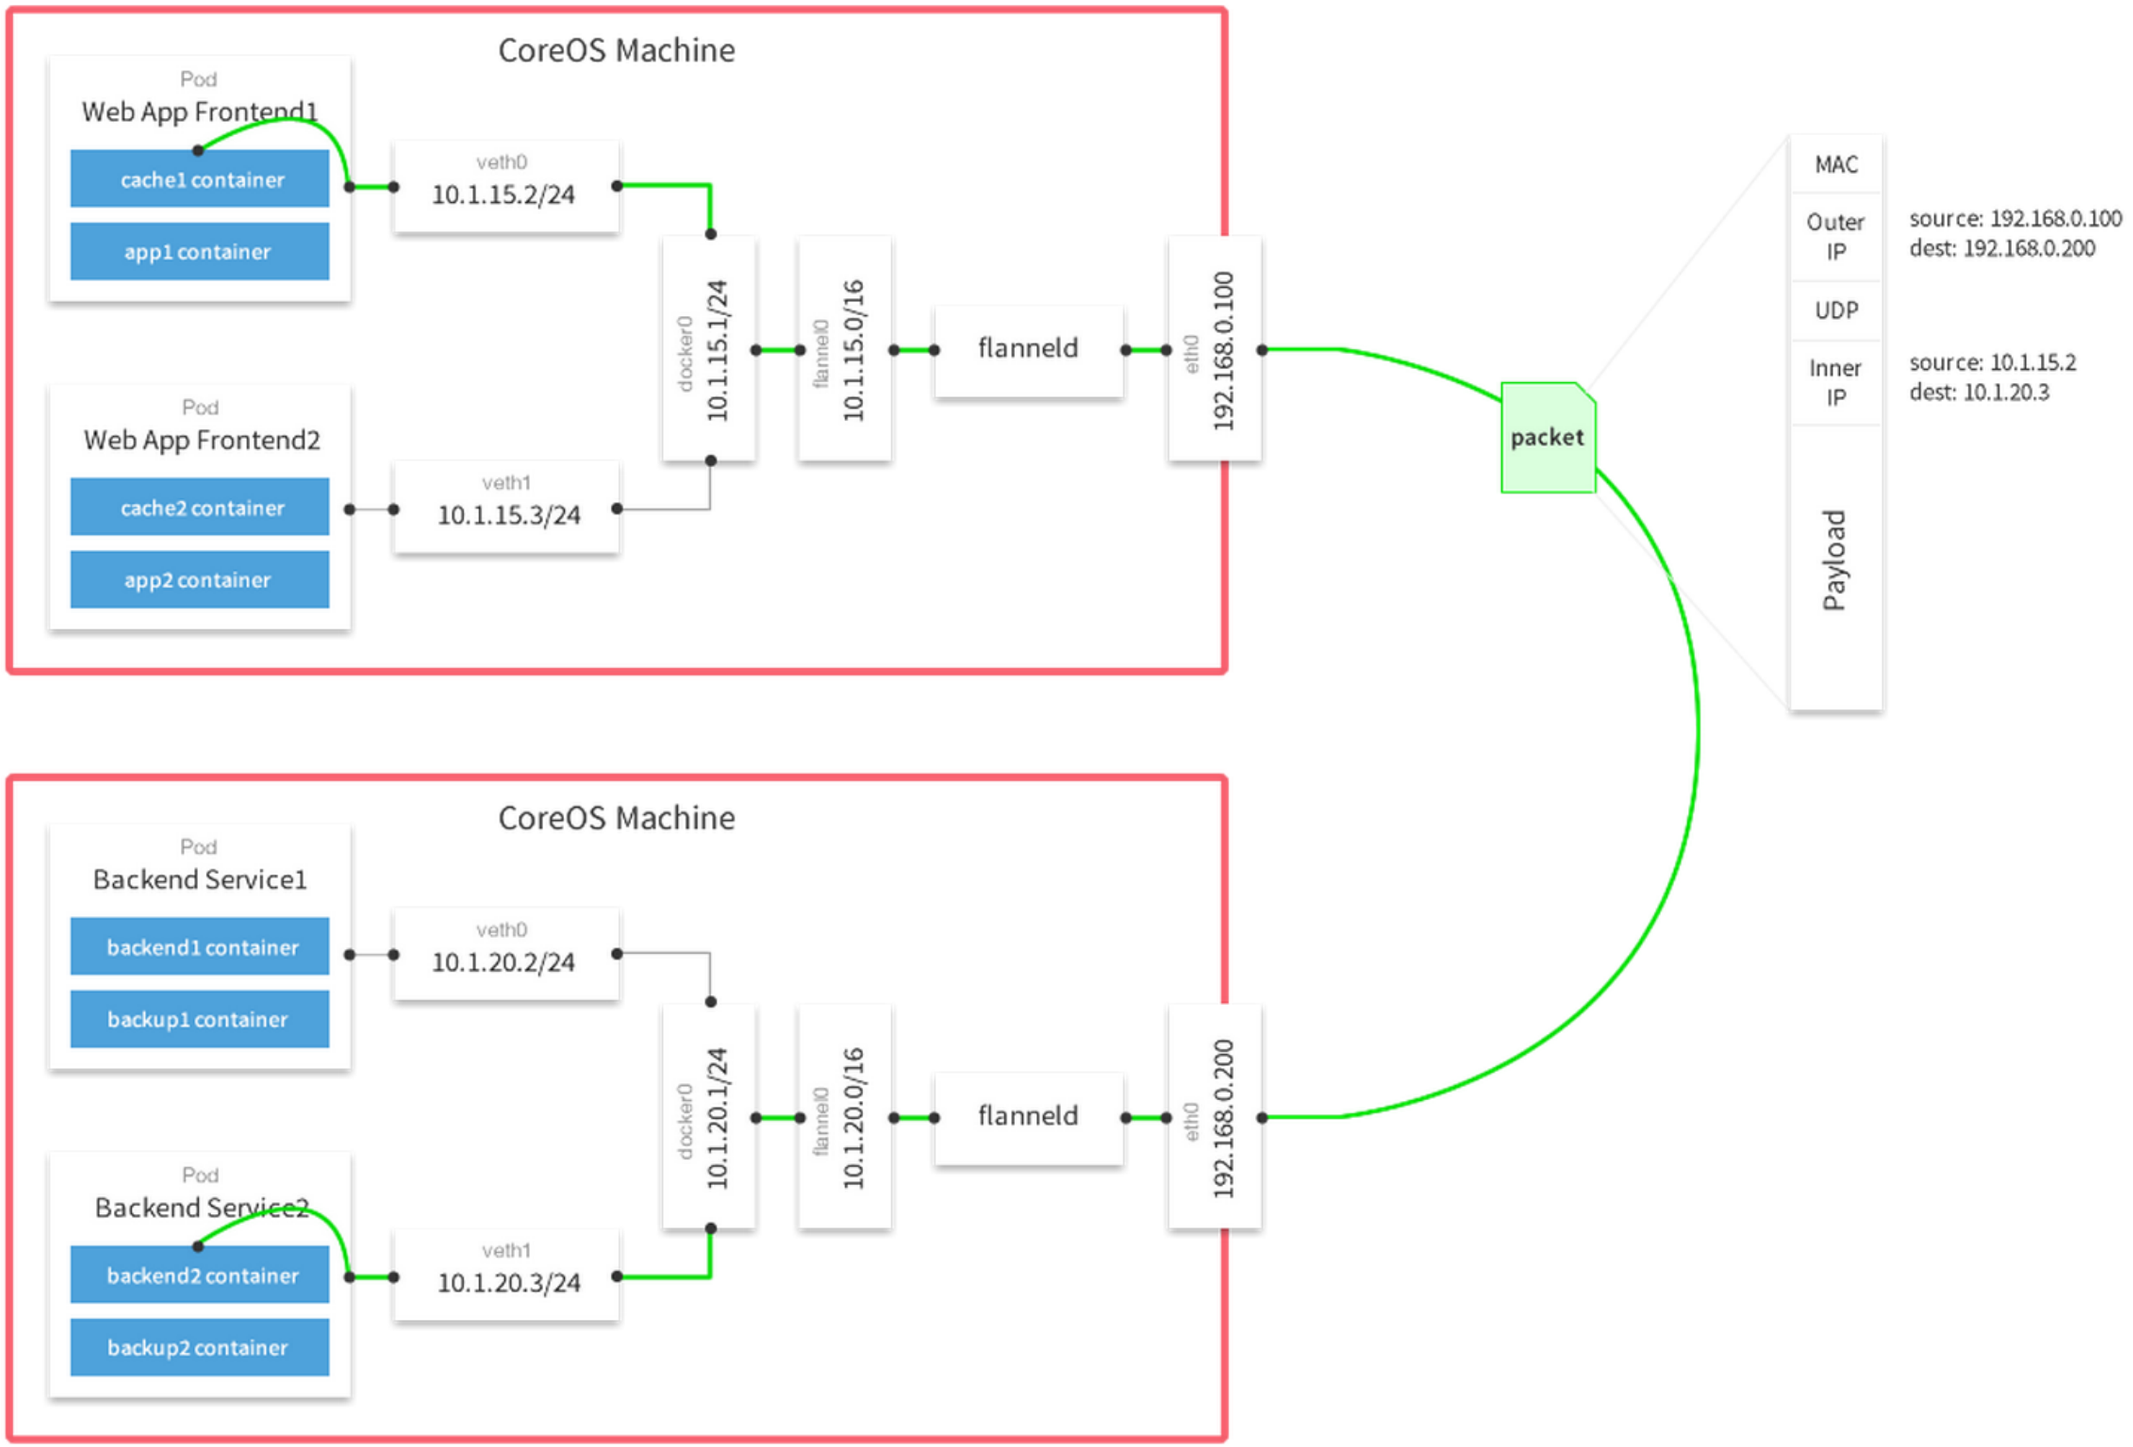

© WSO2

Two level of scalability

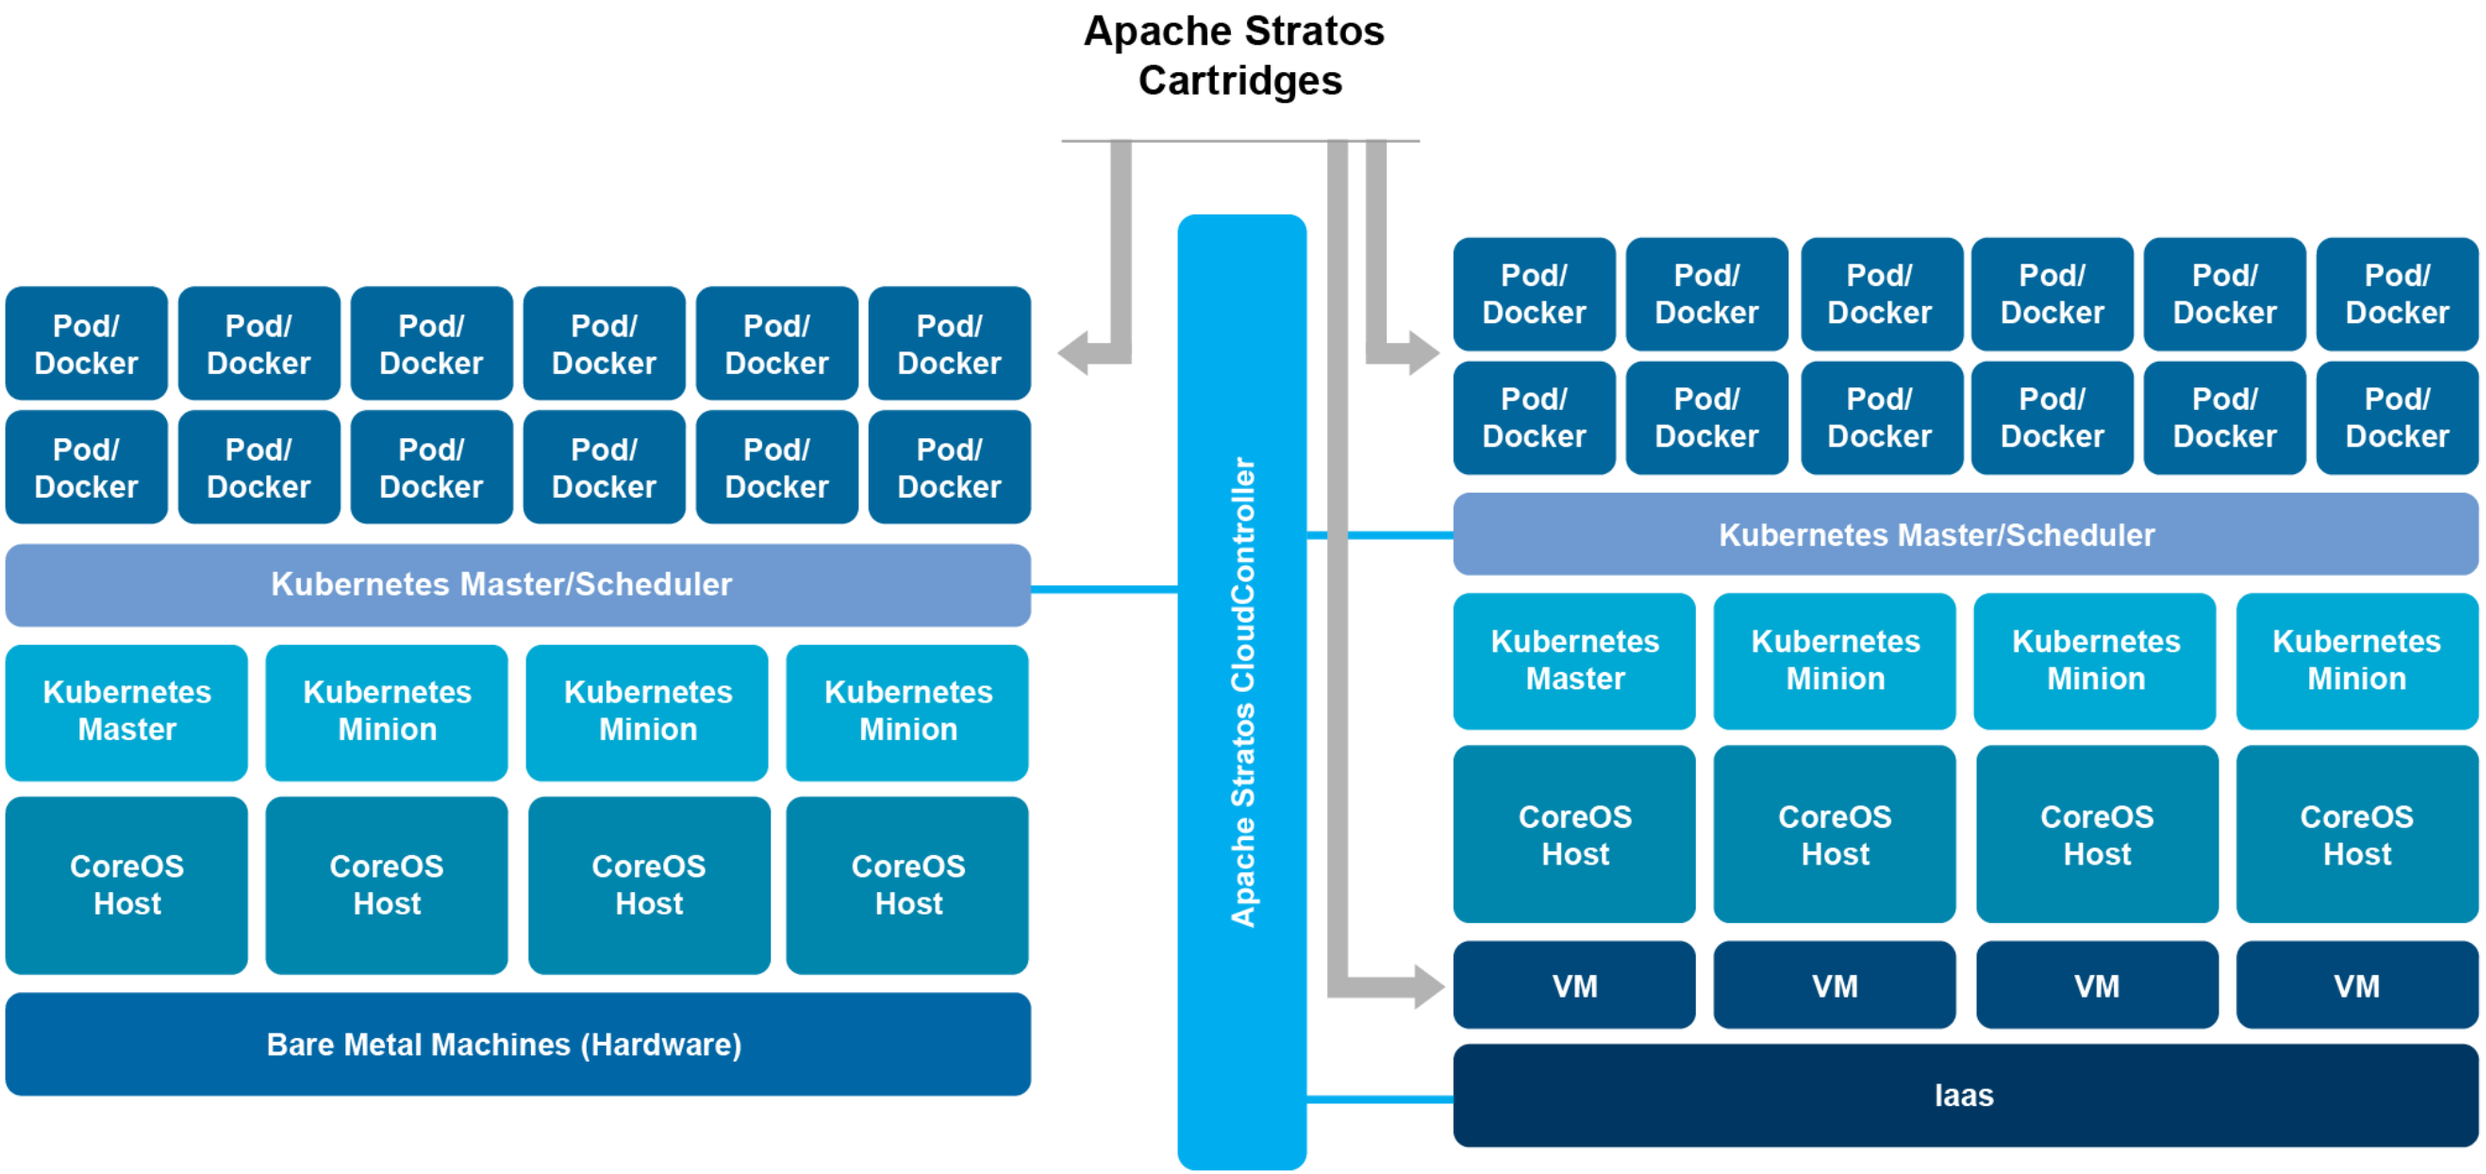

© WSO2

Apache Stratos L1 Architecture for Docker based Cartridges

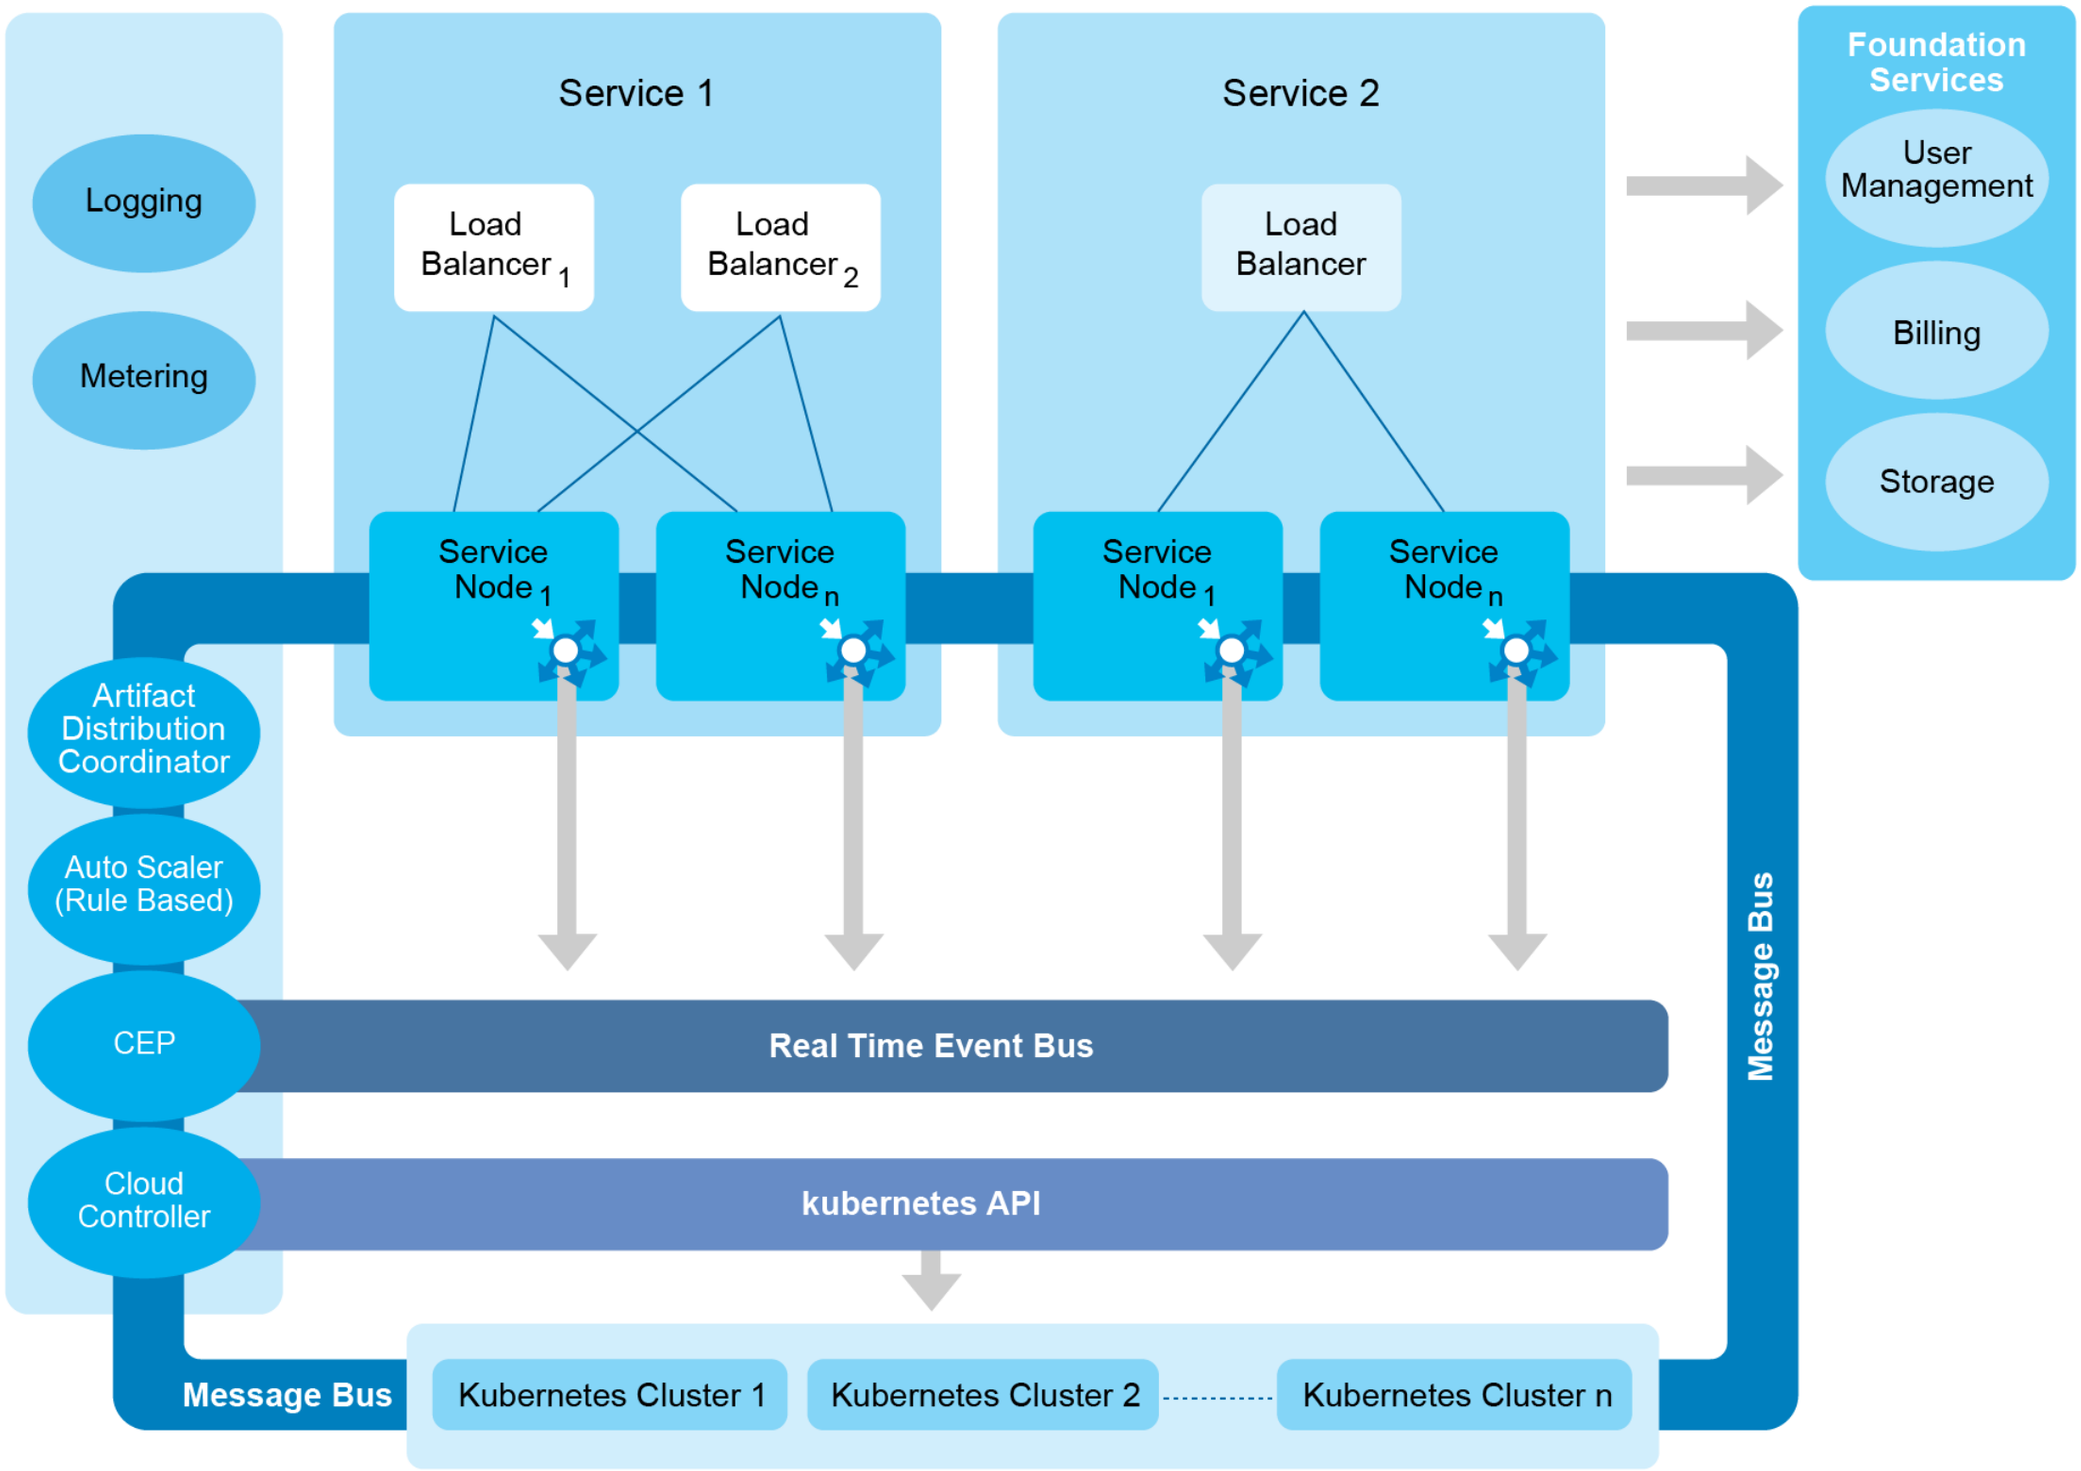

© WSO2

Demo - Apache Stratos 4.1.0-M2 developer preview

- Setup with 3 node CoreOS cluster, Discovery service, Kubernetes master and 3 minions and flannel
- Configure Stratos
- Register Kubernetes-CoreOS host cluster to Stratos
- Deploy Docker based PHP Cartridge
- Deploy PHP application using PHP Cartridge
- Automated artifact updates
- Manual Scaling
- Autoscaling based on load avarage

© WSO2

More Information !

- https://sysadmincasts.com/episodes/31-introduction-to-docker
- https://www.youtube.com/watch?v=tsk0pWf4ipw
- http://stratos.apache.org
- http://lakmalsview.blogspot.com/2013/12/sneak-peek-into-apache-stratos.html
- https://cwiki.apache.org/confluence/display/STRATOS/4.1.0+Stratos+M2+Developer+Preview
- https://github.com/coreos/etcd
- https://github.com/coreos/flannel

© WSO2

North America

CISCO, ATC, ACQUITY GROUP, AMERICAN FAMILY INSURANCE, Children's Hospital Boston, CAA, SAIC, Institut national de santé publique Québec, MEMC, Concur click. done., Trimble, UMUC, JPMorganChase, Prudential, VERIO, AAA, Blue Cross of Idaho, dg, Intermountain Healthcare, BNY MELLON, FRONT PORCH DIGITAL, imshealth, BOEING, StubHub!, KAISER PERMANENTE, UBS, NOEL GROUP, Expedia, TRACTOR SUPPLY CO, ebay, FLEXERA, REGIONS, barclaycard, neustar, west, NATO OTAN, SPECTRUM HEALTH, GSA, Capital Group Companies, intel, GE, Rakuten Global Market, FRONT PORCH DIGITAL

Europe

BBVA, BRAND IMPLEMENTATION GROUP, CINECA, IBERIA, Berlin Airports, VOLVO, suva, Truvo, csi piemonte, GAD, digipolis, Soft Design, A.M.A., Alfa-Bank, Tag:, SECURITAS, RosEvroBank, orange, Salzburg AG, HILTI, BRITISH AIRWAYS, O2, Alcatel-Lucent, eurotaxGLASS's, Xchanging, ALLIANCEBERNSTEIN, HOUSE OF FRASER, bpifrance

Middle East and Asia-Pacific

COMMERCIAL CREDIT, Sri Lanka Telecom Mobitel, Cognizant, ICTA ideas actioned, ODEL, Dialog, Sapura, ELH, onesteel, Australian Government Department of Health and Ageing, JESAI 杰赛科技, VIB

South America

SIF, BI&P Banco Industrial & Partners

Contact us !

© WSO2