



**ACTiCLOUD**

# ACTivating Resource Efficiency and Large Databases in the CLOUD

## Concepts and Objectives

Georgios Goumas

[goumas@cslab.ece.ntua.gr](mailto:goumas@cslab.ece.ntua.gr)

# ACTiCLOUD INFORMATION

EU H2020 project, Grant Agreement N°: 732366

Start date: 1 Jan 2017

Duration: 36 months

## Partners:



NUMASCALE


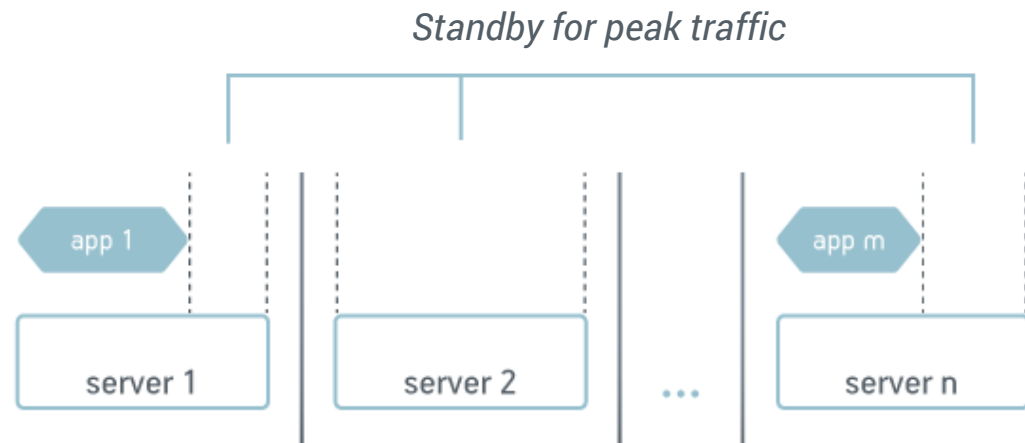


Coordinator: ICCS

# OBSERVATIONS THAT MOTIVATED ACTiCLOUD

- Severe resource **misuse** in cloud datacenters
- Misuse meaning:
  - Resource **waste** due to very low server utilization (10-20%)
  - Resource **fragmentation**
  - Resource **unavailability** (or very high pricing) for resource-hungry applications
    - driving use case: *in-memory databases*
  - Resource **contention** due to interference between co-executing applications

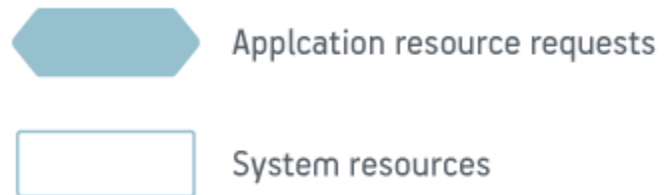
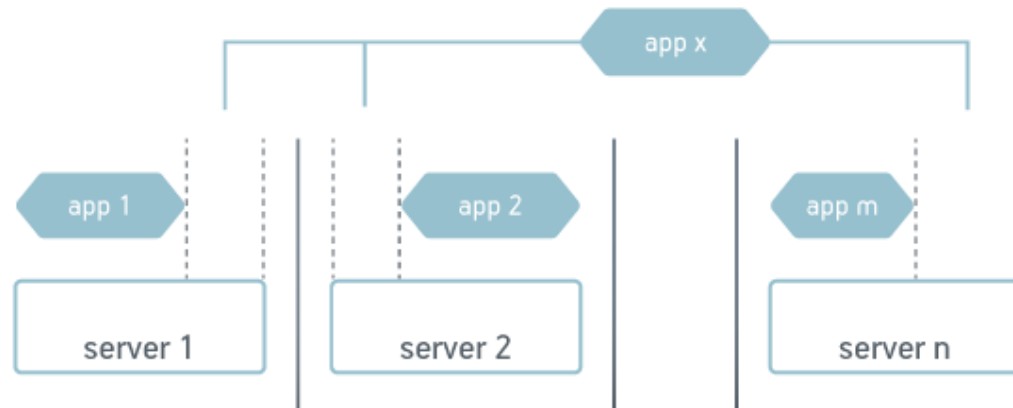
# RESOURCE WASTE



Wasteful allocation by conservative policies to cope with:



- Peak traffic
- Unpredicted collocation behavior

# RESOURCE FRAGMENTATION



Plenty of unallocated resources in small fragments

# RESOURCE UNAVAILABILITY



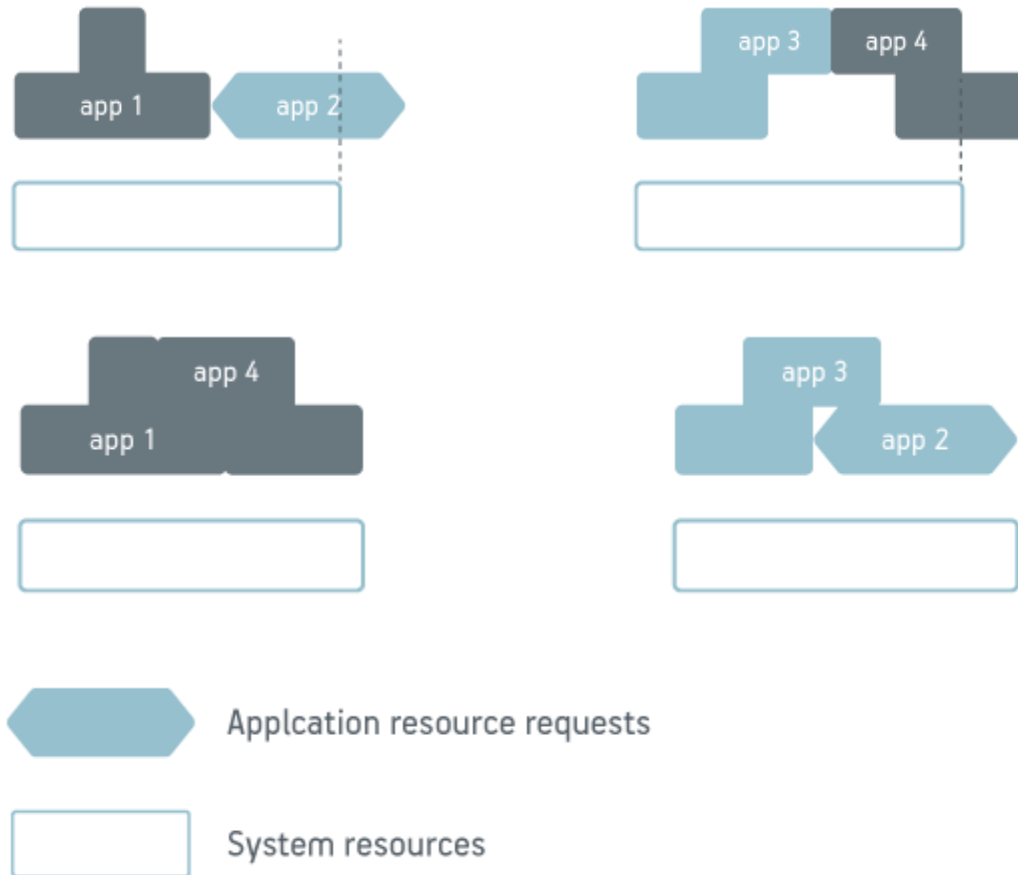
Typical cloud servers cannot service resource-hungry applications

Special application in mind: in-memory databases

We use the term “fat VMs”

Typical scarce resource: main memory

# RESOURCE CONTENTION (INTERFERENCE)



Cloud resource allocators are interference-unaware

Interference can severely degrade performance and QoS

Lack of interference awareness and predictability is a main reason for low server utilization

# ACTiCLOUD STRATEGIC OBJECTIVES

- **Strategic Objective 1 (S01):** Effective utilization of cloud resources.
  - **S01.1:** Resource **efficiency**
  - **S01.2:** Performance **stability**
- **Strategic Objective 2 (S02):** Deployment of resource demanding applications in the cloud (*special focus on database applications*)
  - **S02.1: Scalability** in resource provisioning
  - **S02.2: Elasticity** in resource provisioning



# ACTiCLOUD APPROACH

**Step1:** Break the PC architecture barrier

Base on the technologies provided by Numascale and Kaleao for **resource disaggregation**



# ACTiCLOUD APPROACH

**Step 2:** Extend hypervisor technologies to pool resources at the rack scale


Extend OnApp's MicroVisor

- Memory management
- Hyper-converged storage
- Inteconnect optimization
- Fault tolerance



Addresses the problems of resource **fragmentation** and resource **unavailability**

# ACTiCLOUD APPROACH



**Step 2:** Extend hypervisor technologies to pool resources at the rack scale


Extend OnApp's MicroVisor

- Memory management
- Hyper-converged storage
- Inteconnect optimization
- Fault tolerance




Addresses the problems of resource **fragmentation** and resource **unavailability**

# ACTiCLOUD APPROACH

**Step 3:** Enhance the hypervisor and the cloud manager (OpenStack) to handle resources effectively

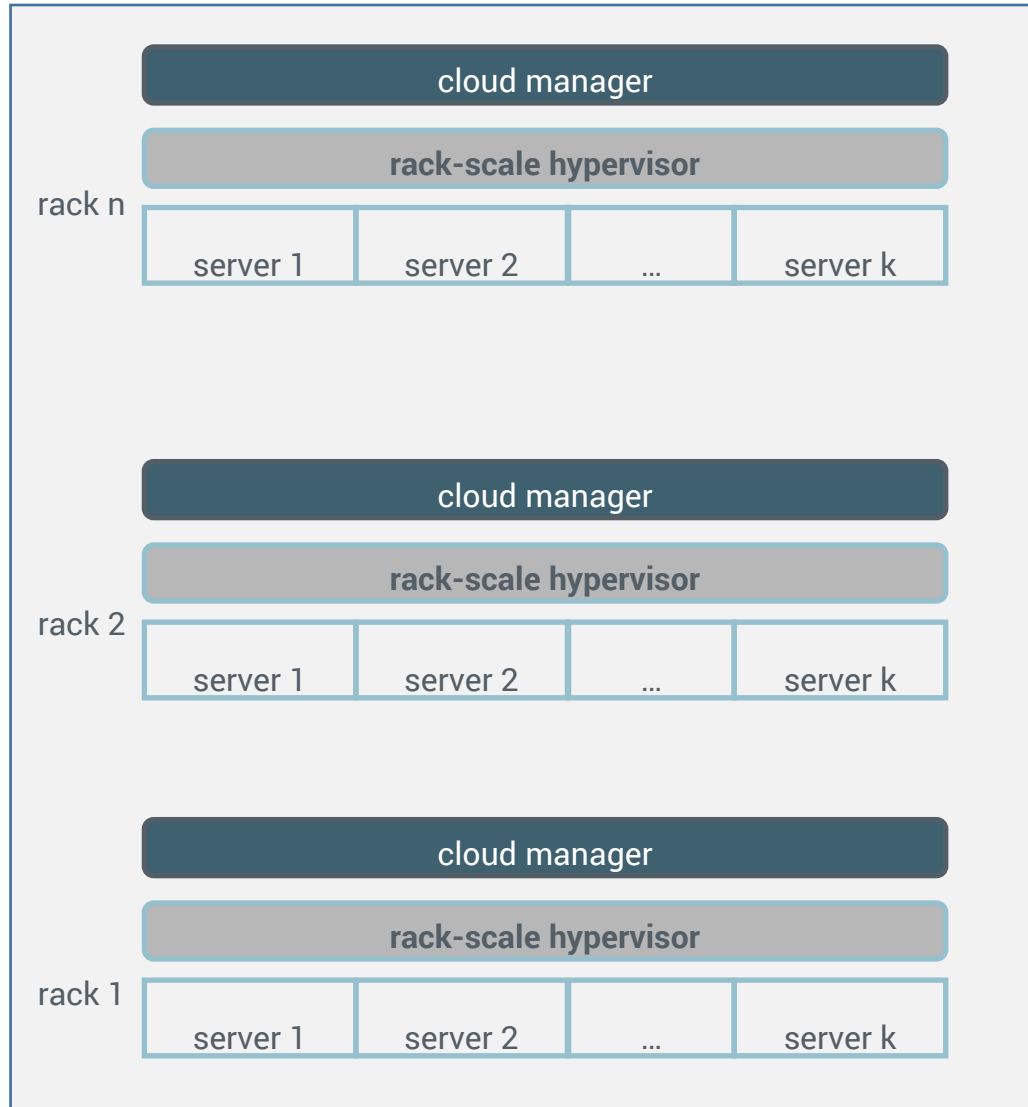


# ACTiCLOUD APPROACH



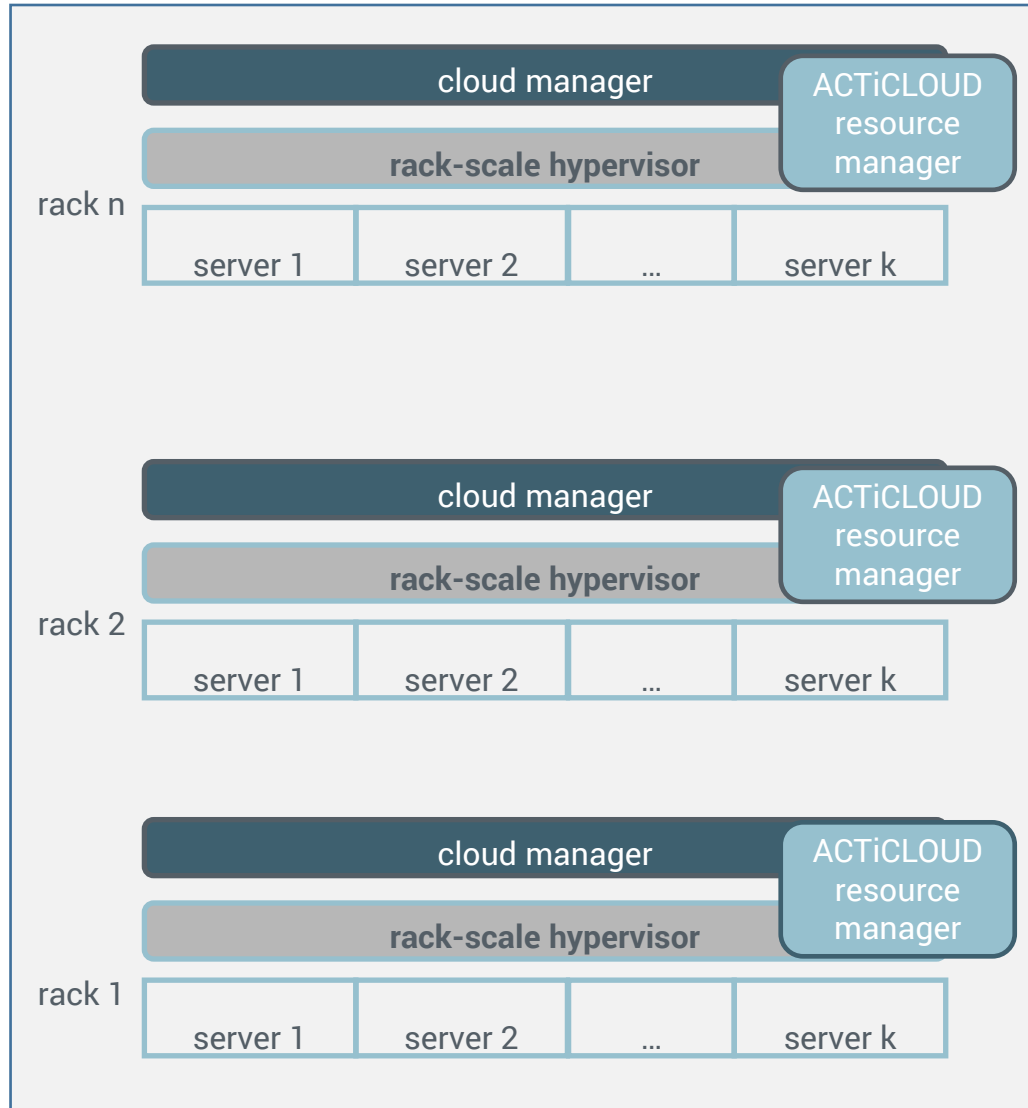
**Step 3:** Enhance the hypervisor and the cloud manager (OpenStack) to handle resources effectively

# ACTiCLOUD APPROACH



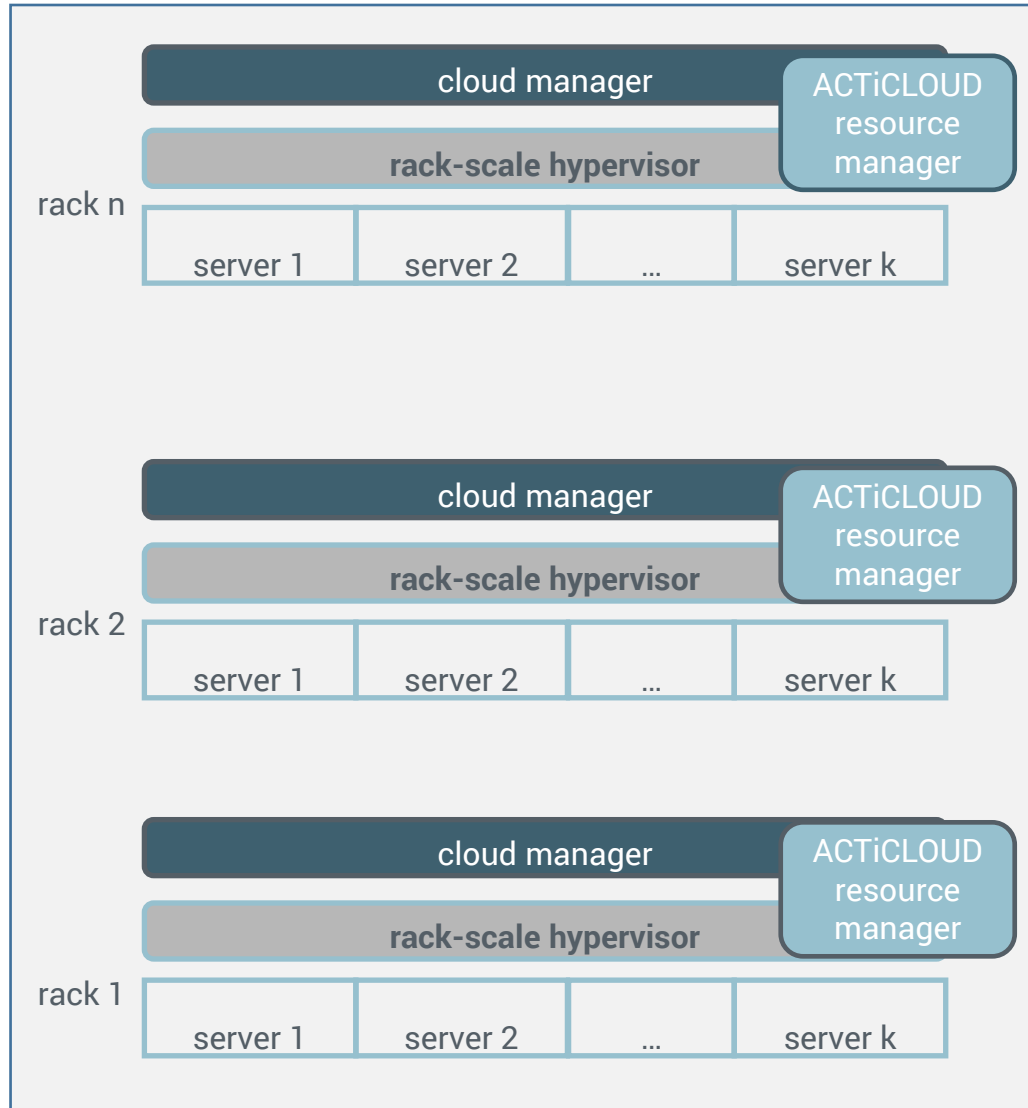
**Step 3:** Enhance the hypervisor and the cloud manager (OpenStack) to handle resources effectively

# ACTiCLOUD APPROACH



**Step 3:** Enhance the hypervisor and the cloud manager (OpenStack) to handle resources effectively

# ACTiCLOUD APPROACH



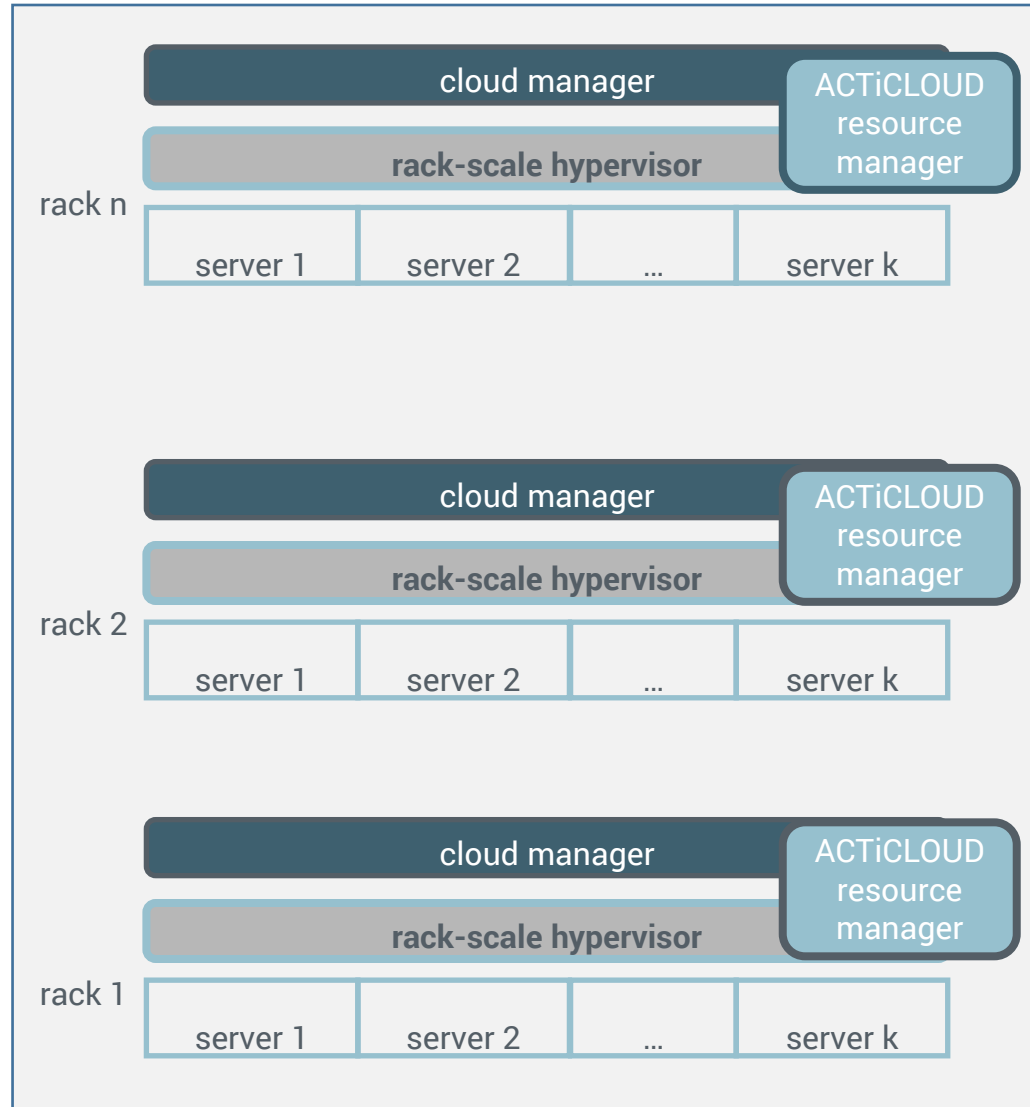
**Step 3:** Enhance the hypervisor and the cloud manager (OpenStack) to handle resources effectively

The ACTiCLOUD resource manager:

- **Monitors** available resources and application requirements
- **Analyzes** the current status and **predicts** the status of alternative states
- **Decides** on alternative resource allocations
- **Takes actions** (e.g. migrate, postpone, time-schedule)



# ACTiCLOUD APPROACH

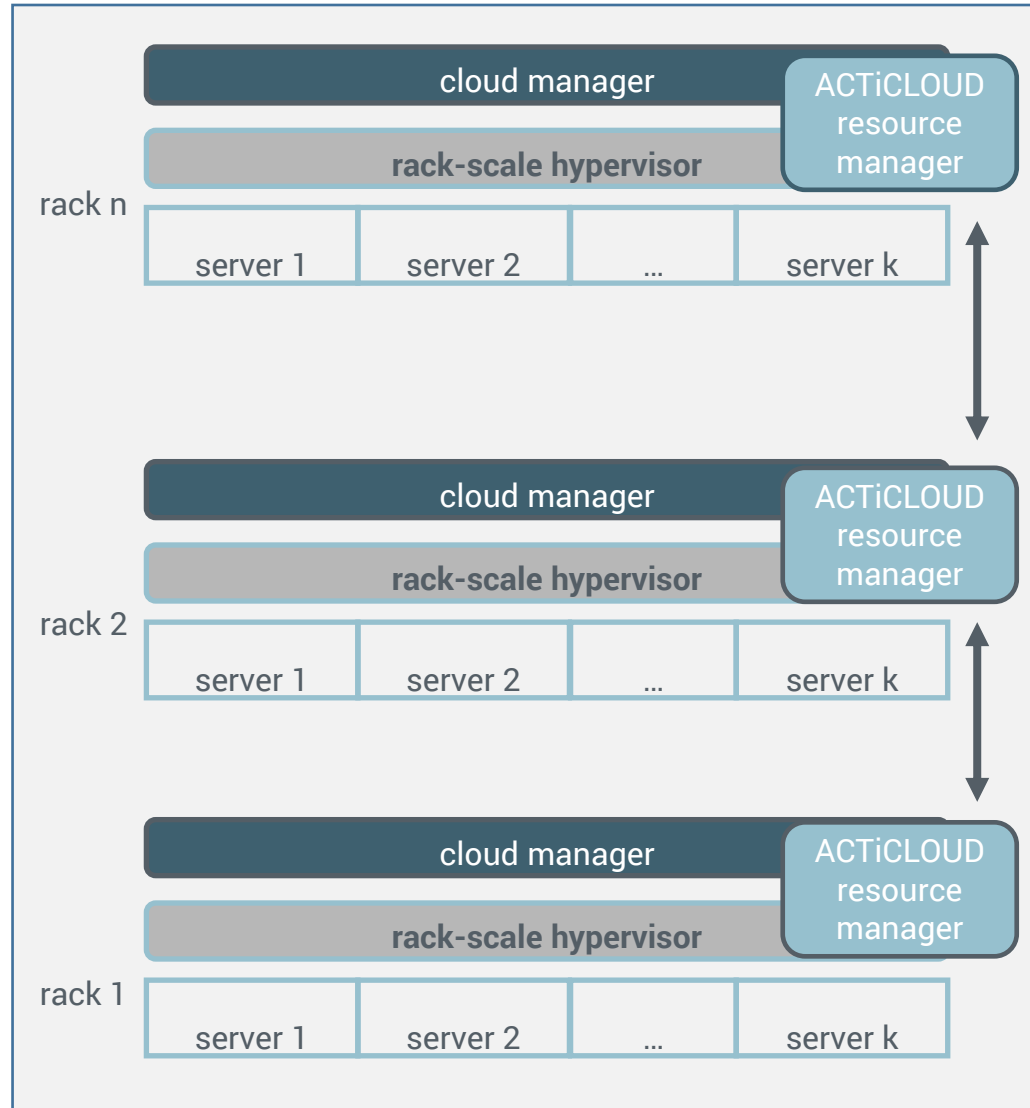


**Step 3:** Enhance the hypervisor and the cloud manager (OpenStack) to handle resources effectively

The ACTiCLOUD resource manager operates:

- At the **rack level** (in cooperation with the hypervisor)

# ACTiCLOUD APPROACH

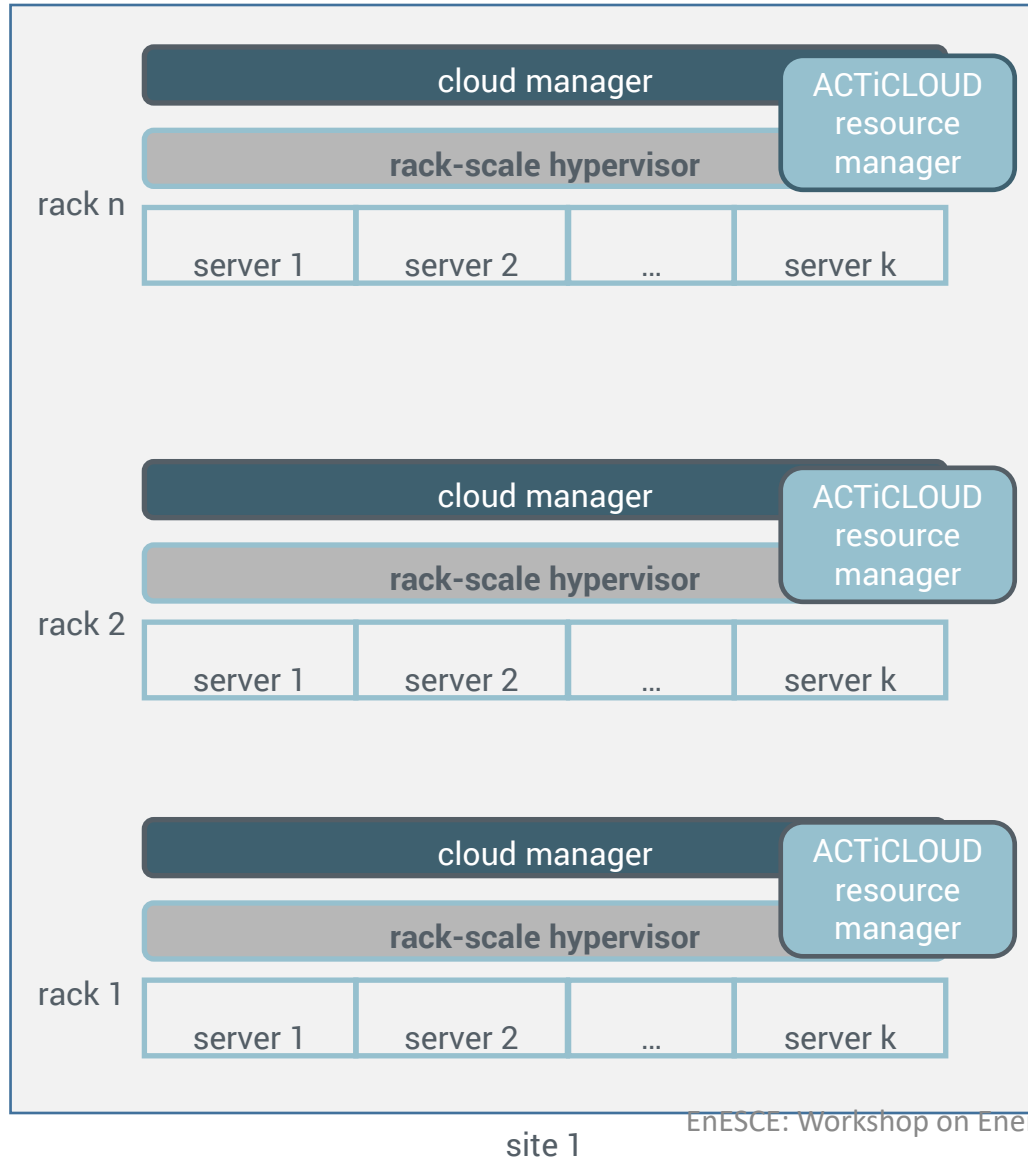


**Step 3:** Enhance the hypervisor and the cloud manager (OpenStack) to handle resources effectively

The ACTiCLOUD resource manager operates:

- At the **rack level** (in cooperation with the hypervisor)
- At the **site level** (in cooperation with OpenStack)

# ACTiCLOUD APPROACH

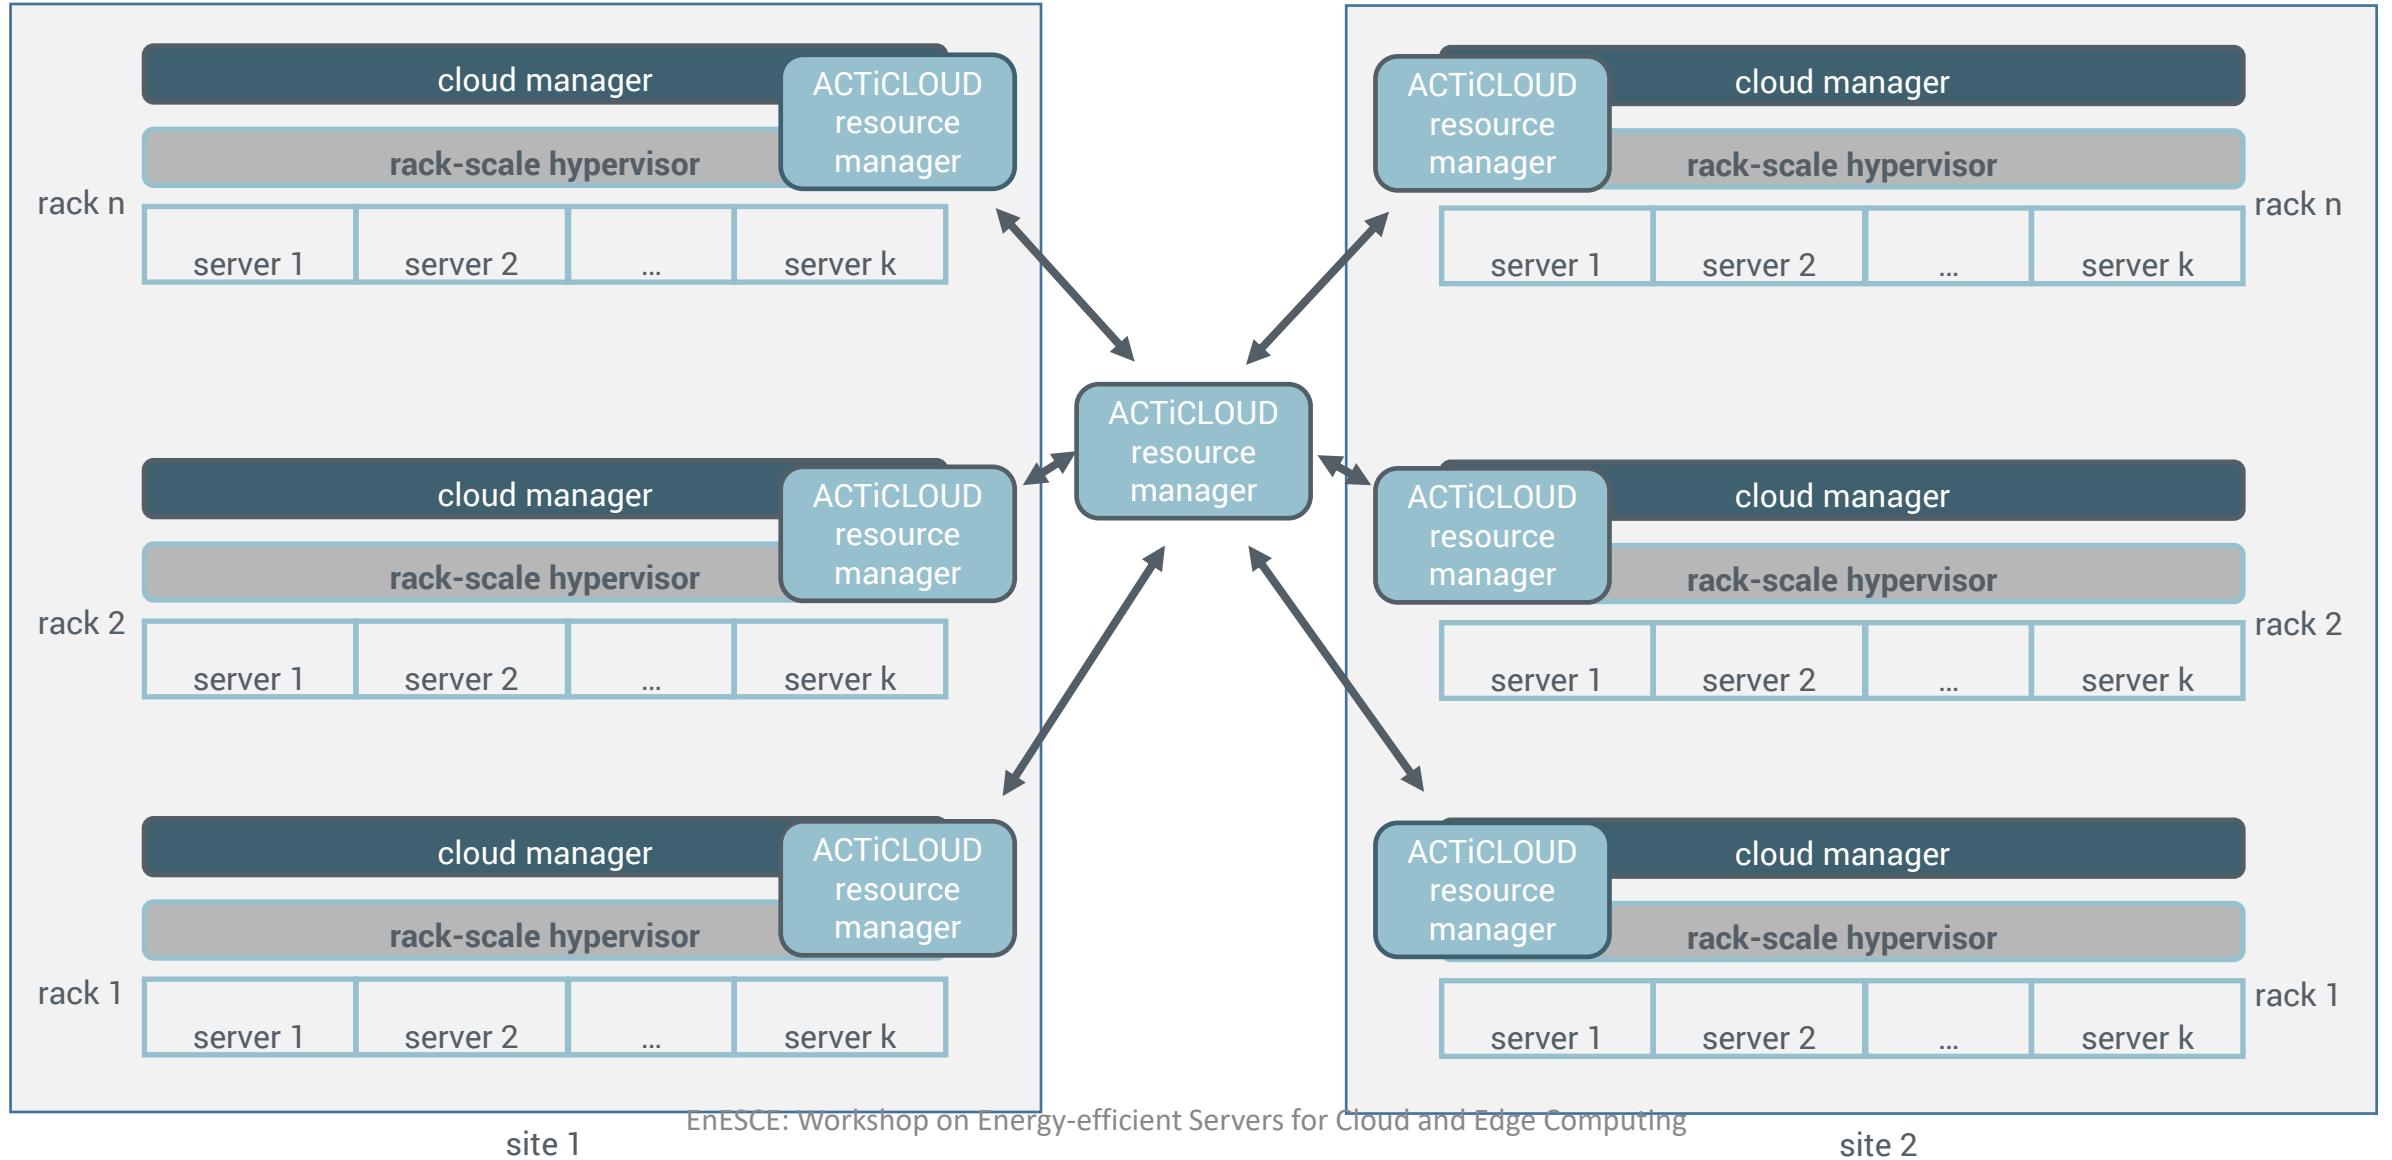


**Step 3:** Enhance the hypervisor and the cloud manager (OpenStack) to handle resources effectively


The ACTiCLOUD resource manager operates:

- At the **rack level** (in cooperation with the hypervisor)
- At the **site level** (in cooperation with OpenStack)
- At the **cross-site level** (by utilizing extensions of OpenStack, e.g. Tricircle)

# ACTiCLOUD APPROACH



# ACTiCLOUD APPROACH



**Step 3:** Enhance the hypervisor and the cloud manager (OpenStack) to handle resources effectively

Addresses the problems of resource **contention** and resource **waste** (sibling, remote sites can be used to offload peak traffic)

# ACTiCLOUD APPROACH: Where we stand


We are half way towards *Activating **Resource Efficiency** and **Large Databases** in the Cloud*

What comes next: *Handling large, in-memory databases*

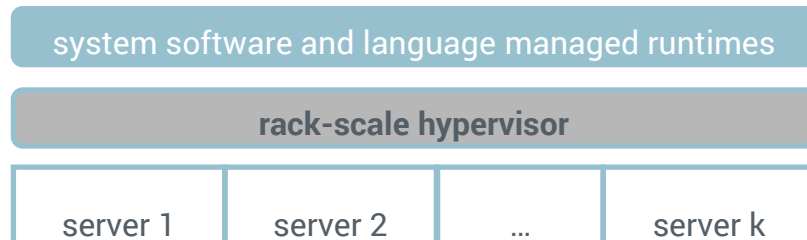
# ACTiCLOUD APPROACH

**Step 4:** Optimize system software and language managed runtimes

Memory managers across the stack  
JVM optimizations garbage collector  
Other system libraries (e.g. within the host VMs)



# ACTiCLOUD APPROACH

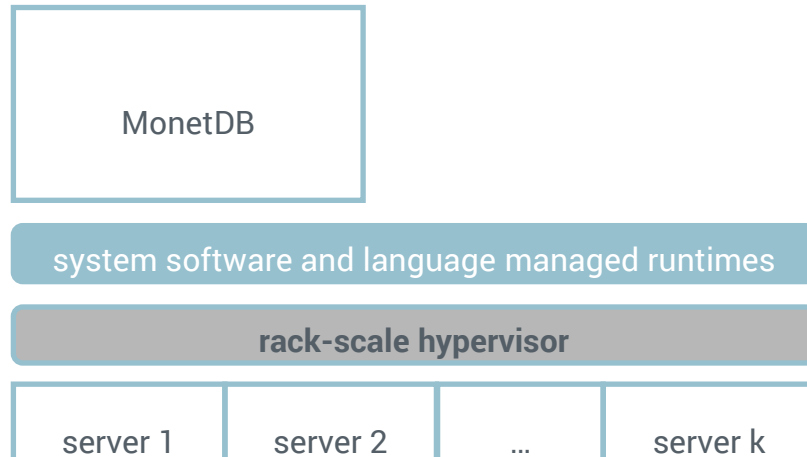


**Step 4:** Optimize system software and language managed runtimes

Memory managers across the stack  
JVM optimizations garbage collector  
Other system libraries (e.g. within the host VMs)

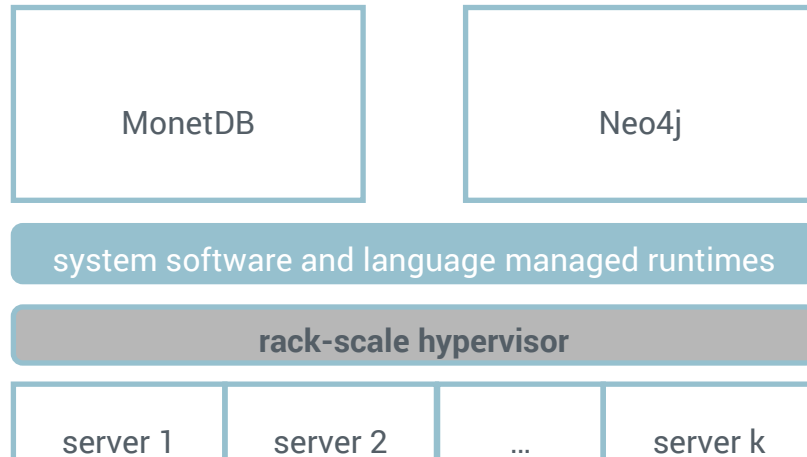


# ACTiCLOUD APPROACH



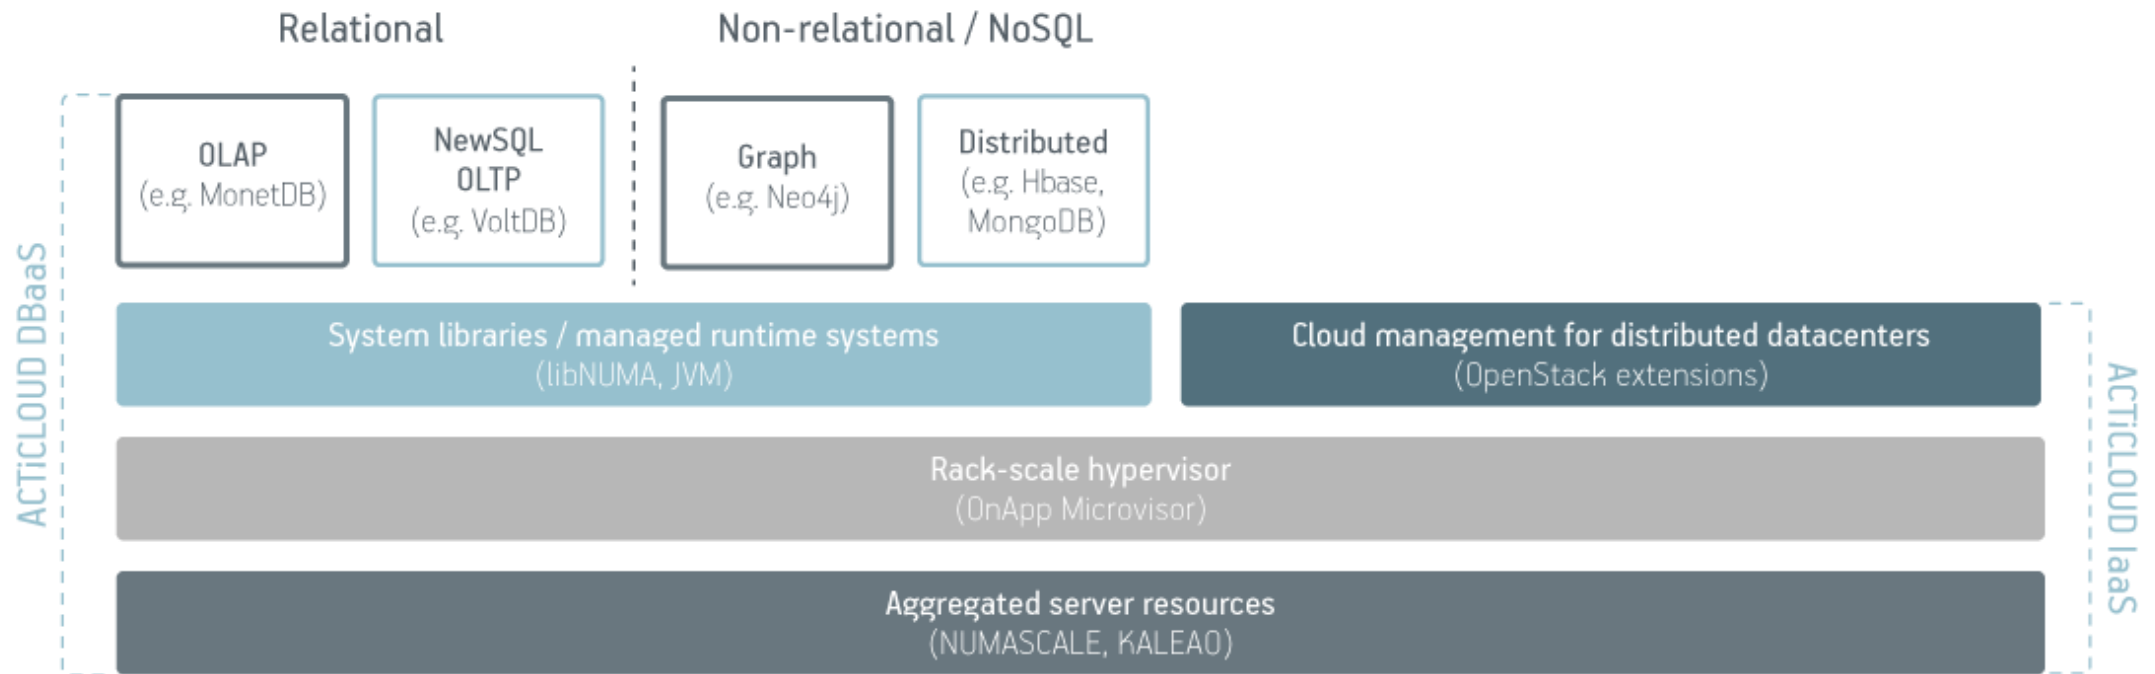
**Step 5:** Evolution of in-memory *column-store databases*

# ACTiCLOUD APPROACH

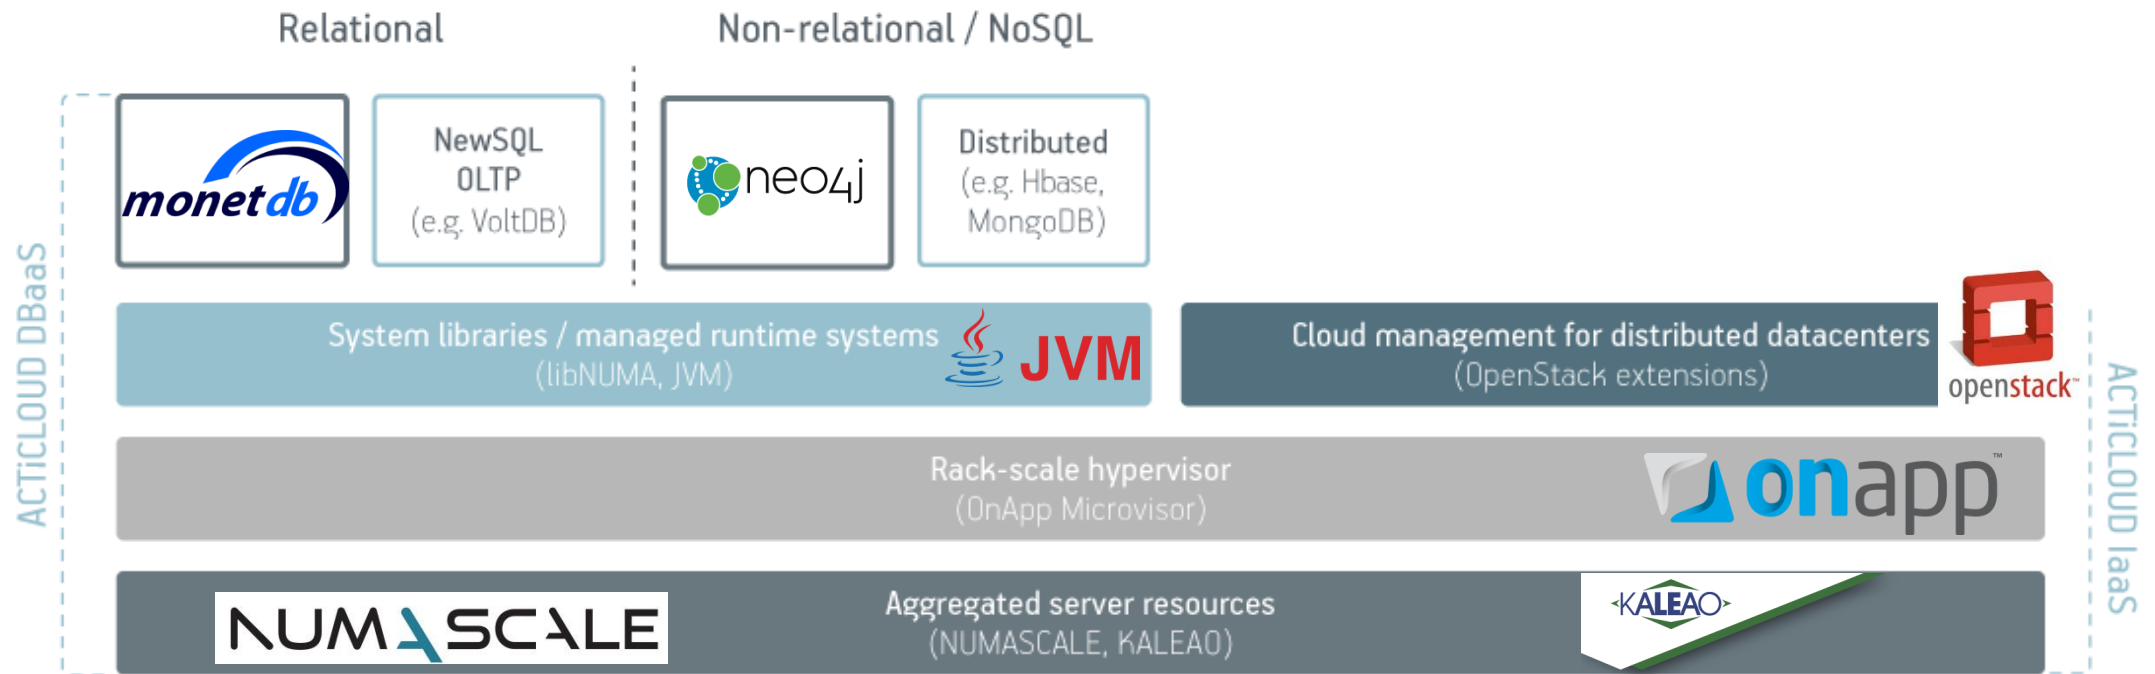


## Step 6: Evolution of large graph *databases*

# ARCHITECTURE AT A GLANCE



# ARCHITECTURE AT A GLANCE - technologies



This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement no 732366 (ACTiCLOUD)



# Questions?

More info:

- Georgios Goumas: [goumas@cslab.ece.ntua.gr](mailto:goumas@cslab.ece.ntua.gr)
- [acticloud.eu](http://acticloud.eu)
- [twitter.com/ACTiCLOUD](https://twitter.com/ACTiCLOUD)