

Specification of the Cloud system

Technical concepts for the cloud environment

Toulouse, June 2nd 2016

Massimo Ruo Roch

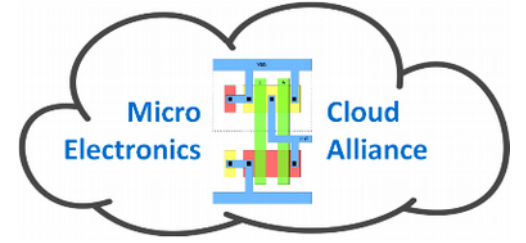
DET-PoliTO

massimo.ruoroch@polito.it



Co-funded by the
Erasmus+ Programme
of the European Union

Speech goals

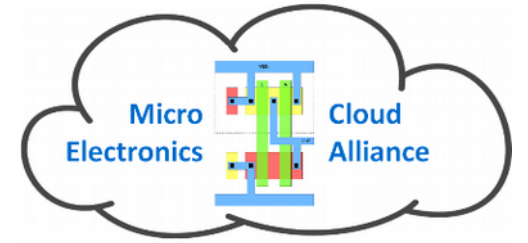


- Understand the MECA approach
- Be ready to implement a private cloud at Home
- Be ready to interface with the Partners



Co-funded by the
Erasmus+ Programme
of the European Union

MECA goals

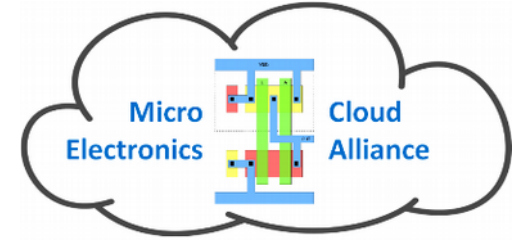


- Designing a cloud infrastructure tailored to support Microelectronics teaching
- Disseminate cloud technology knowledge to partners
- Implement an EU distributed cloud
- Deploy educational applications and teaching examples based on it



Co-funded by the
Erasmus+ Programme
of the European Union

Cloud.... What's this?



Simple answers: servers, accessible from network

They exist from 30 years, but their usage is just for experts

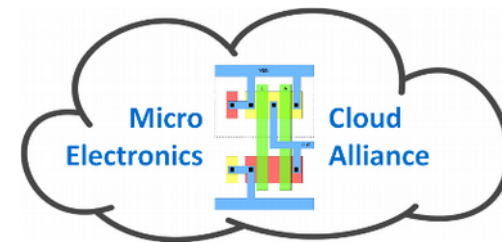
NOW

- They are accessible to everybody (if you have money....)
- No detail about data location
- No detail about server location
- No worries about maintenance, power consumption, heating, etc...



Co-funded by the
Erasmus+ Programme
of the European Union

Cloud.... How to make it?



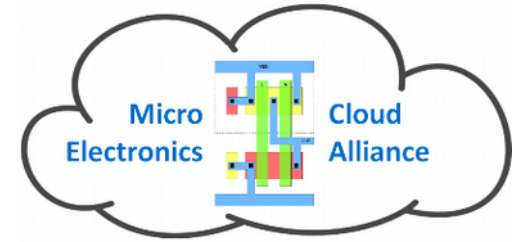
Assembly of available IT technologies

- Hardware technologies:
 - Virtualization extensions in modern CPU's
- Software technologies:
 - Web 2.0 (HTML5, Ajax)
 - Java, Application servers
 - JSON, REST
- Communication technologies
 - Gb networks/devices
 - Software Defined Networks (SDN)



Co-funded by the
Erasmus+ Programme
of the European Union

Available Technologies (1)



Virtualization platforms

Hosting several virtual machines (VM's) on a single host

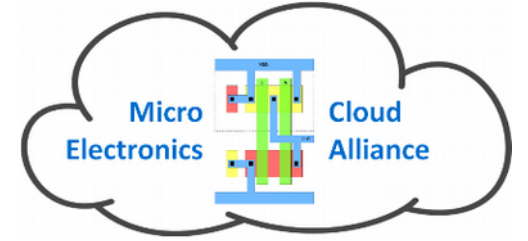
Pros:

- Hardware dependency removal
- Easy remote management (even power-off/power-on) of VM's
- Configuration change possible (add/remove memory, CPU, disk)
- VM's can be manually moved to a different host for maintenance or reconfiguration
- More efficient usage of hardware resources



Co-funded by the
Erasmus+ Programme
of the European Union

Available Technologies (2)



Virtualization platforms

Hosting several virtual machines (VM's) on a single host

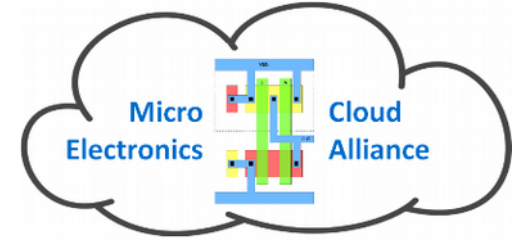
Cons:

- Specific knowledge on host and hypervisor feature needed
- Manual intervention of sysadmin needed
- Difficult to distribute privileges to users without security fallback
- Manual configuration of networking infrastructure



Co-funded by the
Erasmus+ Programme
of the European Union

Available Technologies (3)



The solution:

Further abstraction of hardware resources used by VM's

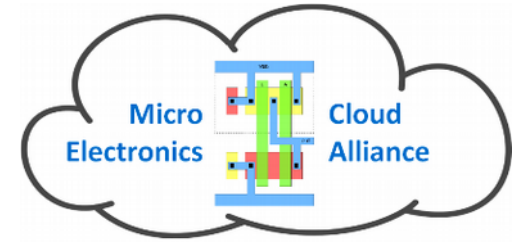
Three different approaches are up to now available:

- Infrastructure as a Service (IaaS)
- Platform as a Service (PaaS)
- Software as a Service (SaaS)



Co-funded by the
Erasmus+ Programme
of the European Union

Available Technologies (4)



IaaS

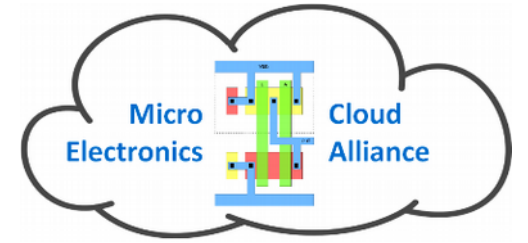
A management software application allows to configure underlying hardware (hosts and networking), using a uniform user interface

- A web application is typically used
- API is available to automate tasks



Co-funded by the
Erasmus+ Programme
of the European Union

Available Technologies (5)

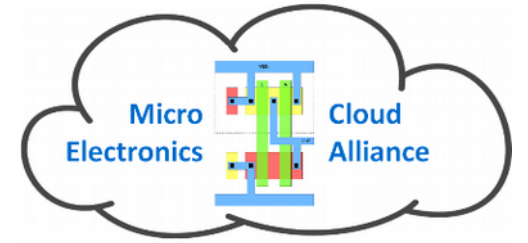


IaaS

- Hierarchical privilege structure:
 - Top administrators build the infrastructure
 - Domain sysadmin's receives privileges to create VM's using infrastructure resources, but with limits on maximum resource allocation
 - VM's sysadmin can just manage single instances (start/stop/terminate/users management, etc.)



Available Technologies (6)



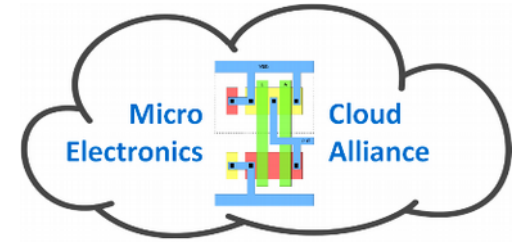
IaaS

- Creation of 'standard' services:
 - Host templates with pre-configured hardware
 - Security templates with predefined firewall rules
 - Volume templates with pre-installed OS's



Co-funded by the
Erasmus+ Programme
of the European Union

Available Technologies (7)



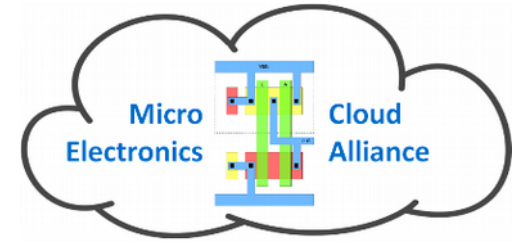
IaaS

- Drawbacks:
 - Difficulty to forecast needed templates
 - Application software changes can massively invalidate previous work
 - Auto scaling and load balancing not fully automated



Co-funded by the
Erasmus+ Programme
of the European Union

Available Technologies (8)



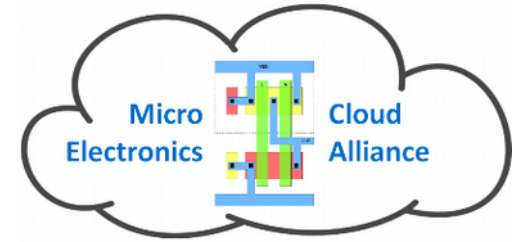
PaaS

A management software application allows to build custom VM's on-the-fly, assembling OS's, server software and application software, selected from a catalog

- User specifies 'recipes', to 'cook' the right configuration
- PaaS management interface (Web or API based) create VM's on demand, according to the defined recipes
- IaaS underlying system details are typically hidden



Available Technologies (9)



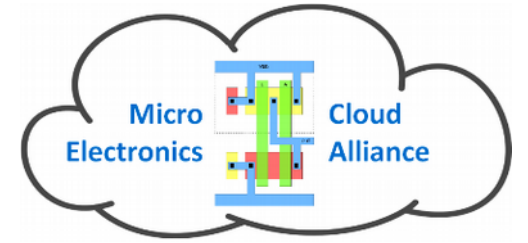
PaaS

Example:

- Recipe for IDE Web application for software development, 3-tiers model:
 - (1 – 2) CentOS 7 VM minimal + Postgresql
 - (1 – 8) CentOS 6 VM minimal + Glassfish
 - 2 OpenBSD VM + load balancer



Available Technologies (10)

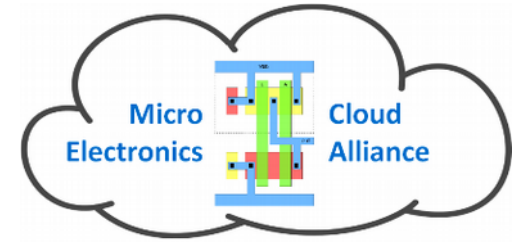


PaaS

- Numbers in brackets mean a variable number of instances is needed according to load
- PaaS receipts allow to specify maximum/minimum load on each VM needed to trigger auto-replication or auto-termination of instance replicas



Available Technologies (11)



SaaS

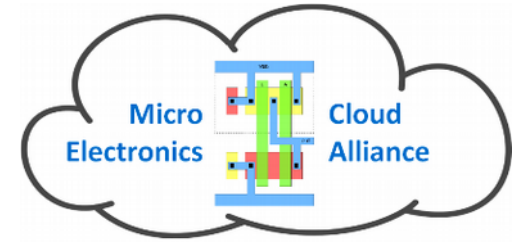
An application software is developed, but it is distributed on a pay-per-use business model.

- Typically based on a PaaS or IaaS underlying infrastructure, but these details are hidden to user.
- Example: Dropbox, videoconferencing, Webmail, ERP, CRM



Co-funded by the
Erasmus+ Programme
of the European Union

MECA Implementation (1)



IaaS and PaaS layers seem to be reasonably sufficient to build MECA computing environment.

Scouting has been performed on existing solutions, according to the following criteria:

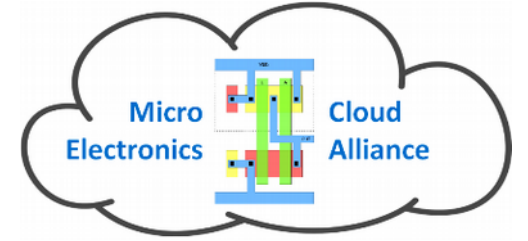
- Software cost minimization: OSS software preferred
- Framework reliability: as a production level system must be built, only OSS with reliable support and reasonable roadmaps have been taken in account. No experimental feature!

(...continued)



Co-funded by the
Erasmus+ Programme
of the European Union

MECA Implementation (2)



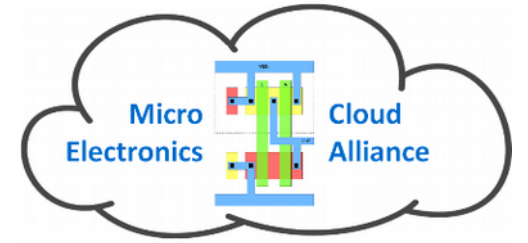
(continued...)

- Performances:
 - Hypervisor must guarantee negligible virtualization penalties, both from CPU and I/O point of view
 - Management software resource consumption must be negligible with respect to payload
- Deployment effort:
 - Fast and efficient installation/configuration of simple system is mandatory



Co-funded by the
Erasmus+ Programme
of the European Union

MECA Implementation (3)



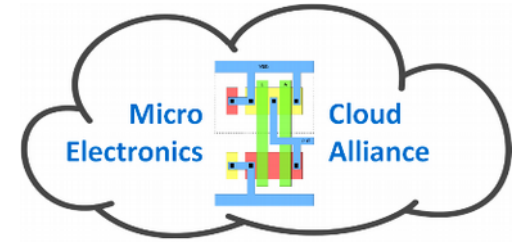
Final choices:

- IaaS layer: Cloudstack
- PaaS: Cloudify (optional, implemented only if needed)
- Educational Cloud distributed on academic/company hosted nodes



Co-funded by the
Erasmus+ Programme
of the European Union

CloudStack Main Features (1)



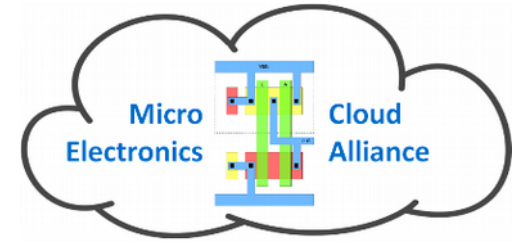
- Complete out-of-the-box IaaS solution
- Apache Software Foundation maintained project
- Java based Web application
- Tomcat 6 application server
- Pre-compiled packages for the main Linux distributions
- Support of different hypervisors, both free and commercial (BareMetal, Xen, KVM, Hyper-V, LXC, vSphere)

(...continued)



Co-funded by the
Erasmus+ Programme
of the European Union

CloudStack Main Features (2)



(continued...)

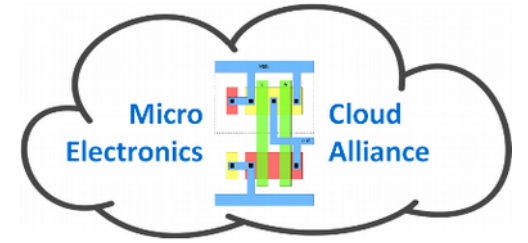
- Near linear scalability, up to tens of thousand hosts
- Supports geographically distributed datacenters inside the same cloud
- REST-like API, and AWS EC2 support
- High-availability through redundant Management Servers and DB replica

(...continued)



Co-funded by the
Erasmus+ Programme
of the European Union

CloudStack Main Features (3)



(continued...)

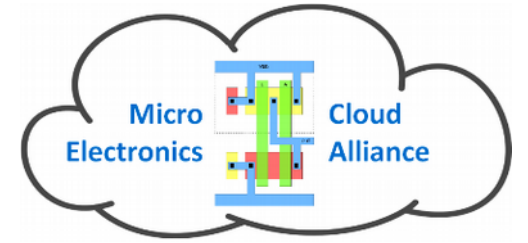
And, under the sheets.....

- Powerful, customizable logging system
- No false errors on log files
- Reduced log footprint



Co-funded by the
Erasmus+ Programme
of the European Union

CloudStack Architecture (1)

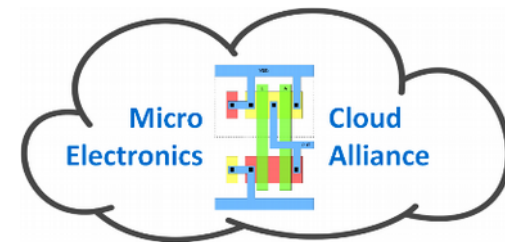


Management server(s)

- Provides a web interface for both administrator and end users.
- Provides API interfaces for both the CloudStack API as well as the EC2 interface.
- Manages the assignment of guest VMs to a specific compute resource
- Manages the assignment of public and private IP addresses.
- Allocates storage during the VM instantiation process.
- Manages snapshots, disk images (templates), and ISO images.



CloudStack Architecture (2)

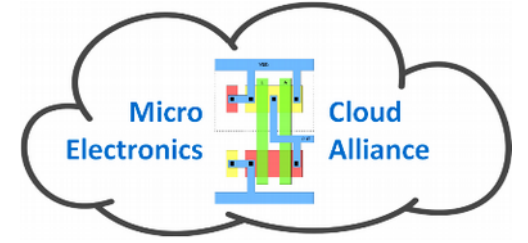


Infrastructure hierarchy

- **Regions:** A collection of one or more geographically proximate zones managed by one or more management servers.
- **Zones:** Typically, a zone is equivalent to a single datacenter. A zone consists of one or more pods and secondary storage.
- **Pods:** A pod is usually a rack, or row of racks that includes a layer-2 switch and one or more clusters.
- **Clusters:** A cluster consists of one or more homogenous hosts and primary storage.
- **Host:** A single compute node within a cluster; often a hypervisor.



CloudStack Architecture (3)

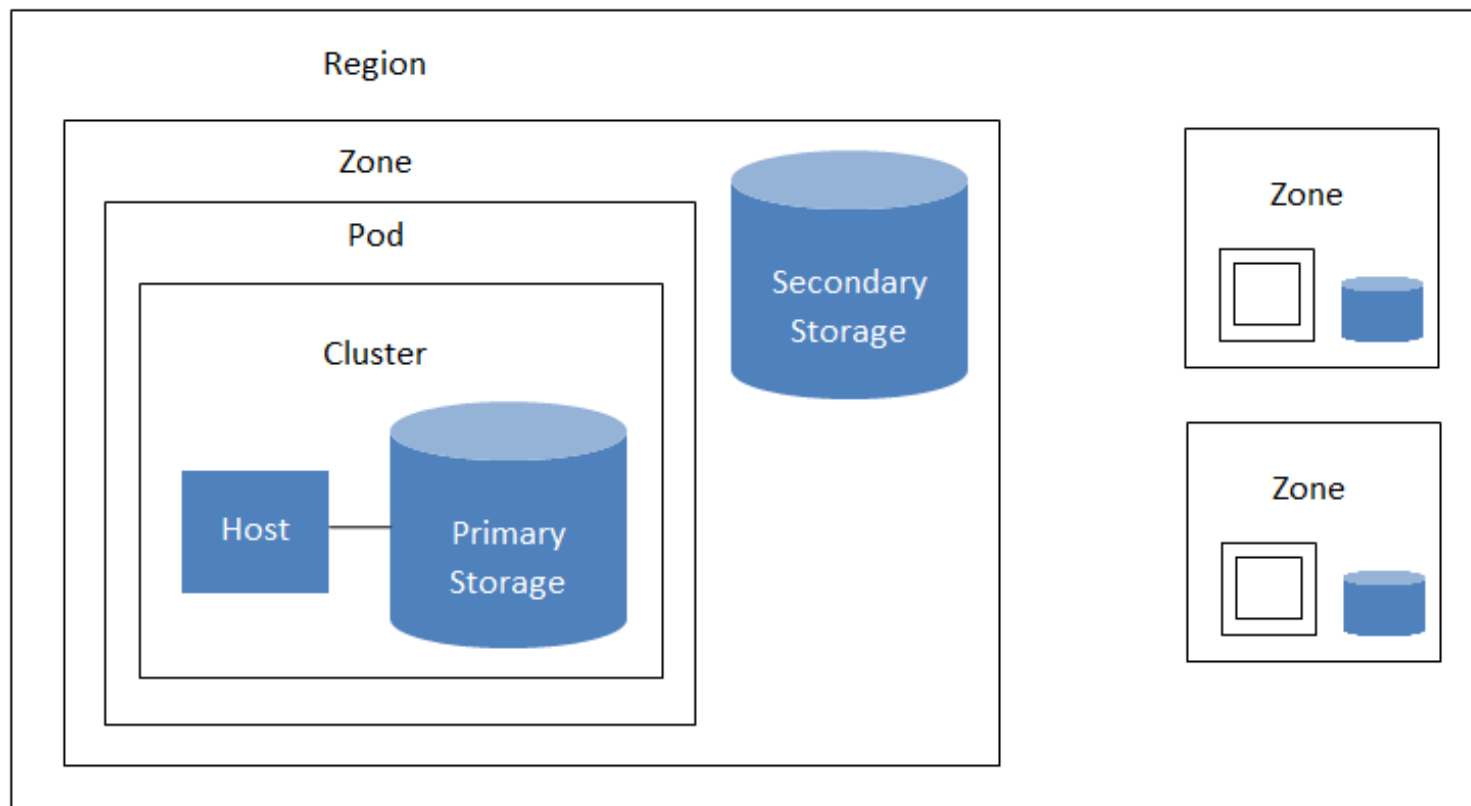
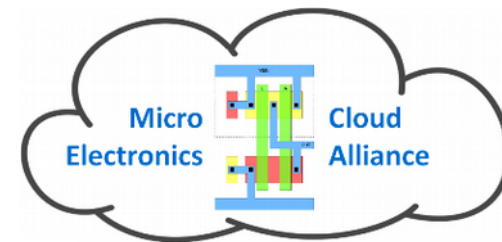


Storage types

- **Primary Storage:** A storage resource typically provided to a single cluster for the actual running of instance disk images. (Zone-wide primary storage is an option, though not typically used.)
- **Secondary Storage:** A zone-wide resource which stores disk templates, ISO images, and snapshots.



CloudStack Architecture (4)

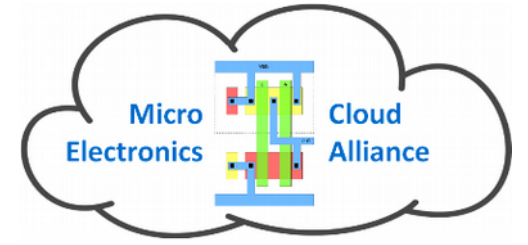


A region with multiple zones



Co-funded by the
Erasmus+ Programme
of the European Union

CloudStack Architecture (5)



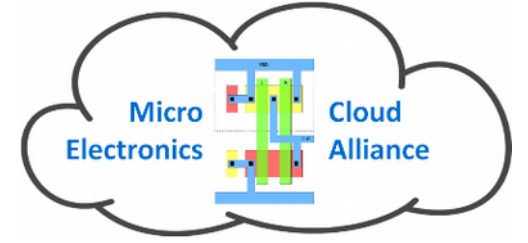
Networking

Two different networking schemes are available, according to resulting system complexity and requirements:

- Basic: Most analogous to AWS-classic style networking. Provides a single flat layer-2 network where guest isolation is provided at layer-3 by the hypervisors bridge device. All the traffic on a single network.
- Advanced: This typically uses layer-2 isolation such as VLANs, though this category also includes SDN technologies such as Nicira NVP, or OVS (OSS). Traffic separation according to payload type.



CloudStack Architecture (6)



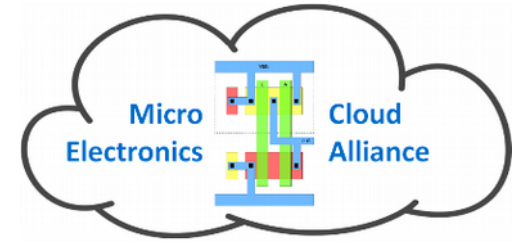
Network traffic classification

- Guest traffic
- Management traffic
- Public traffic
- Storage traffic



Co-funded by the
Erasmus+ Programme
of the European Union

CloudStack Architecture (7)

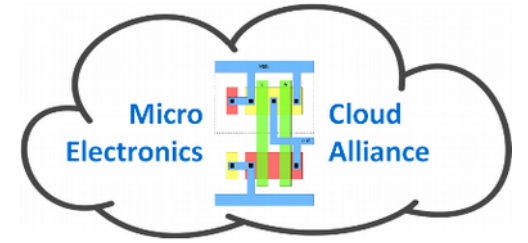


Network traffic classification

- Guest. When end users run VMs, they generate guest traffic. The guest VMs communicate with each other over a network that can be referred to as the guest network. This network can be isolated or shared. In an isolated guest network, the administrator needs to reserve VLAN ranges to provide isolation for each CloudStack account's network (potentially a large number of VLANs). In a shared guest network, all guest VMs share a single network.



CloudStack Architecture (8)

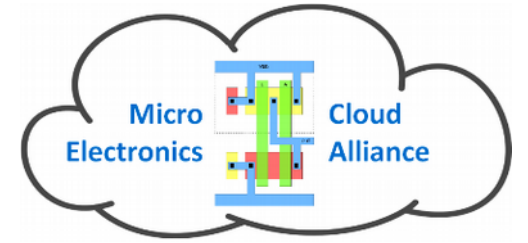


Network traffic classification

- Management. When CloudStack's internal resources communicate with each other, they generate management traffic. This includes communication between hosts, system VMs (VMs used by CloudStack to perform various tasks in the cloud), and any other component that communicates directly with the CloudStack Management Server. You must configure the IP range for the system VMs to use.



CloudStack Architecture (9)

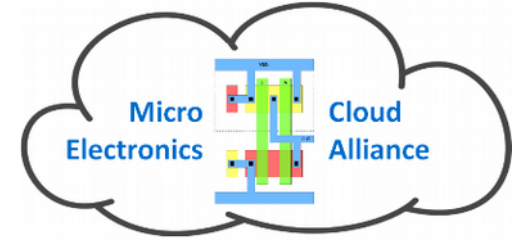


Network traffic classification

- Public. Public traffic is generated when VMs in the cloud access the Internet. Publicly accessible IPs must be allocated for this purpose. End users can use the CloudStack UI to acquire these IPs and to implement NAT/port forwarding between their guest network and the public network.



CloudStack Architecture (10)

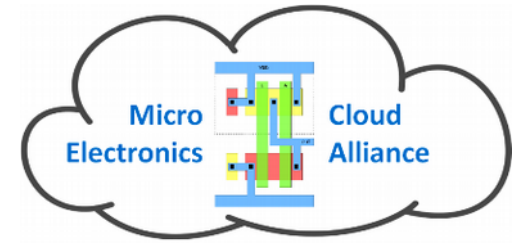


Network traffic classification

- Storage. While labeled “storage” this is specifically about secondary storage, and doesn’t affect traffic for primary storage. This includes traffic such as VM templates and snapshots, which is sent between the secondary storage VM and secondary storage servers.



Q & A



???



Co-funded by the
Erasmus+ Programme
of the European Union