Eco4Cloud

Architecture and Requirements

Eco4Cloud Architecture

Eco4Cloud improves the economics of virtualized data centers with an intelligent software platform, which reduces the energy bill and increases efficiency.

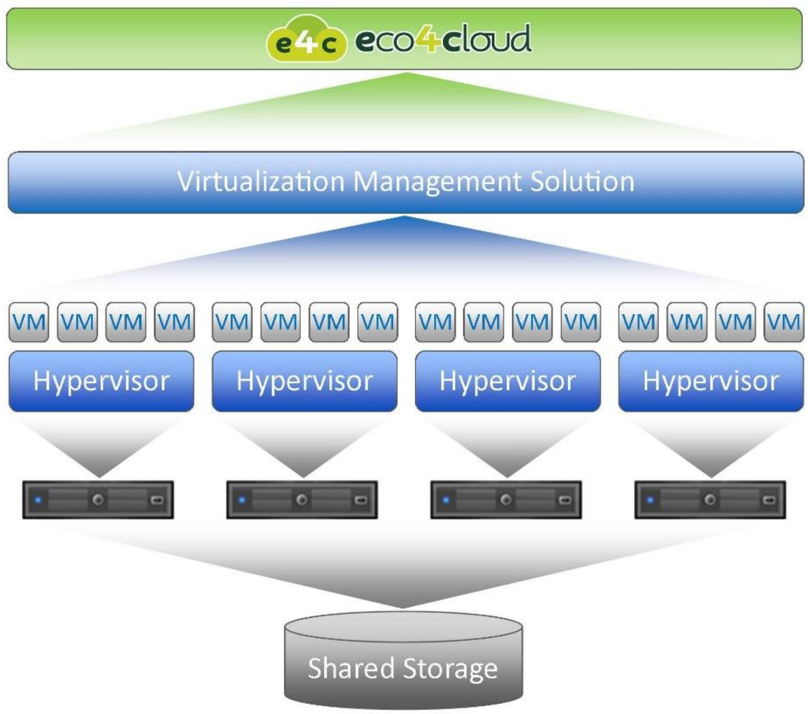

Figure 1. Eco4Cloud Architecture

As shown in Figure 1, the E4C software works on top of a virtualization platform, and uses the exposed API to connect itself to the platform, query the status of IT equipment, both physical and virtual, and issue actions within the data center in order to make it more efficient.

Eco4Cloud works as a kind of plugin to the virtualization platform, so several security measures have been implemented, in order to make it easy to activate/deactivate/customize Eco4Cloud without affecting SLAs nor experiencing any downtime.

In the largest deployment scenarios, an alternative architecture is available (shown in figure 2). An Eco4Cloud deployment can be composed of a single master node, containing the web dashboard, and several slave nodes, each one hooked to its own virtualization manager, so that each virtualization manager in the data center is covered by an Eco4Cloud node, but all the data center is visible using a single web dashboard.

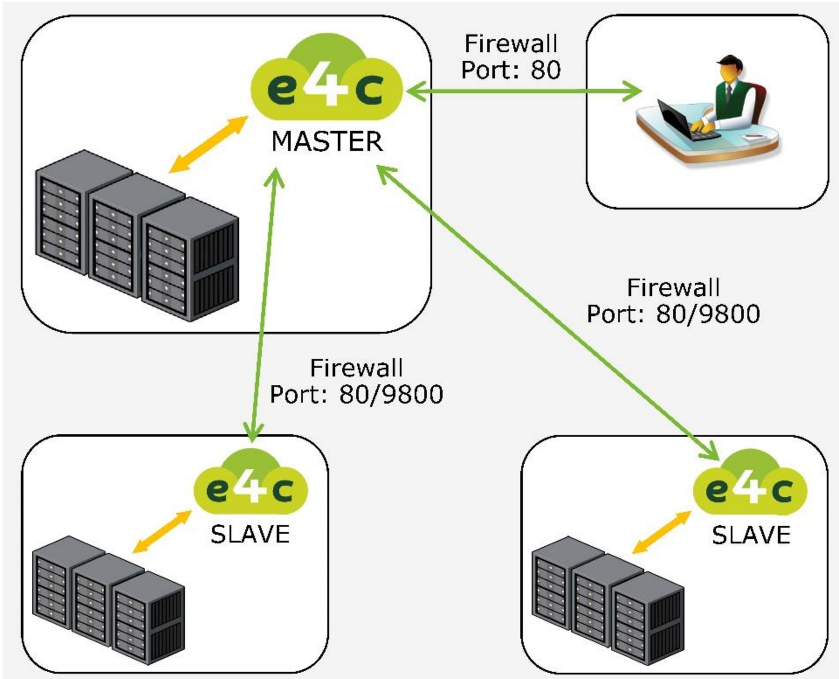

Figure 2. Eco4Cloud Master-Slave architecture

Eco4Cloud setup phases and passive / active mode

The Eco4Cloud setup consists of 4 phases:

1. Analysis – Eco4Cloud uses a monitoring tool to analyse the data center;
2. Simulation – Eco4Cloud estimates the impact of virtual machine consolidation in terms of energy savings and QoS improvement;
3. Installation and early monitoring – Eco4Cloud is installed on the data center, full on-site support is provided;
4. Remote monitoring – The data center is constantly monitored and product upgrades are provided.

Availability of energy/power measurement tools should be granted through the use of sensor-type technologies that allow the comparison of energy consumption before and after the installation of the Software. Eco4Clous is also able to leverage via software several energy consumption sensors, that modern servers are equipped with.

Alternatively, the testing activities can be performed by cross-matching the server data usage with the respective server vendors' data sheets. The data sheets in fact report the servers energy consumption, under particular usage scenarios, measured by the server manufacturer.

The analysis phase is also useful to run ROI simulations on the entire facility or facilities. In order to perform the Service, Eco4Cloud needs to be granted access to detailed info provided by the Client regarding of all the Client's hardware and software resources that will be subject to the analysis.

Once the analysis phase is completed, a second phase starts: Simulation. The simulation phase is run by the Eco4Cloud Team, comparing energy consumption before and after the consolidation of virtual machines. Due to the very nature of Eco4Cloud's business model, the obtained results are functional to any future potential commercial agreement.

The third phase is the installation and early monitoring. Once the software is ready, it is installed on the data center, and starts the virtual machines consolidation. Full on-site support is provided by Eco4Cloud.

The final step is the remote monitoring, the data center is constantly monitored and product upgrades are provided. Eco4Cloud shall be provided with an enduring VPN connection to the vApp, or, alternatively, a temporary network connectivity to the vApp, periodically provided.

During E4C “passive mode” (analysis – simulation phases) the software does not issue VM migrations, nor hosts shutdowns, the E4C team does not strictly need remote monitoring, only temporary access via ssh/remote desktop/teamviewer to the virtual appliance or a machine having network access to the virtual appliance. On the other hand, if –based on the initial results of the test in passive mode only– we also want to test/assess the E4C software behavior and ultimate impact in terms of dynamic consolidation and related energy efficiency, then of course the following phases are in fact required in order to generate the relevant outcomes and quantity metrics.

Requirements

The only hard requirement for Eco4Cloud in order to be installed in a data center is the presence of a virtualization platform managing the hosts and the virtual machines that must be consolidated. Once the hypervisor that best fits the data center requirements has been chosen, any physical host today is able to support a virtualization layer. Table 1 lists the main virtualization platforms with respective supported processors and operating systems which allow Eco4Cloud to perform.

VIRTUALIZATION PLATFORM	HOST CPU	HOST OS
VMware vSphere	x86, x86-64	No host OS
Microsoft Hyper-V	x86-64, Intel VT-x, AMD-V	Windows 2008/2012
Citrix XenServer	x86, x86-64, IA-64	Linux, NetBSD, Solaris
Openstack	x86, x86-64, IA-64	Linux, FreeBSD, illumos
Red Hat Enterprise	x86, x86-64, IA-64	R.H.E.V.
HP Integrity Virtual Machines	IA-64	HP-UX
Oracle VM Server	x86, x86-64, Intel VT-x, AMD-V/ SPARC	No host OS / Solaris 10 11

Table 1 – Hypervisors requirements for Eco4Cloud

The interaction between the Eco4Cloud software and the virtualization hypervisors occurs via the APIs exposed by the virtualization platforms. Through the APIs, Eco4Cloud can then gather both physical hosts and virtual machines performance data, instruct virtual machines live migrations and power-on/off physical hosts. To this extent, the functionalities of the Eco4Cloud software are then completely hardware-agnostic.

As Table 2 shows, all main virtualization platforms expose APIs even if these APIs are heterogeneous with respect to the adopted technologies such as Web Services, REST Services and System Calls. Eco4Cloud is independent from any specific virtualization platform, as it integrates all required libraries to interact with the respective hypervisors.

VIRTUALIZATION PLATFORM	MANAGEMENT	LIVE MIGRATION	POWER ON/OFF
VMware vSphere	VMware API	vSphere 3.0, or superior	IPMI
Microsoft Hyper-V	Microsoft WMI	Hyper-V 2008 R2, or superior	IPMI
Citrix XenServer	Citrix API	XenServer 4.0, or superior	IPMI
Openstack	Openstack API	Release 2011.1, or superior	IPMI
Red Hat Enterprise	R.H.E.V. API	R.H.E.V. 5.4, or superior	IPMI
HP Integrity Virtual Machines	VMGuestLib	Integrity VM 4.0, or superior	IPMI
Oracle VM Server	O.VM S. API	VM Server 2.1, or superior	IPMI

Table 2 – Live Migrations requirements for Eco4Cloud

The Live Migration feature (ref. Table 2) is currently available on any commercially available hypervisor; however it is recommended to check the actual version of the installed hypervisor, in order to confirm that the Live Migration feature is actually supported.

Finally, should the release of the virtualization platform in use be older than those listed in Table 2, i.e. the host power-on/off feature not be supported by the virtualization platform's APIs, then Eco4Cloud can still instruct the power-on/off operation independently through IPMI (Intelligent Platform Management Interface) libraries. IPMI is a standard protocol, commonly implemented in any modern data center.

Minimum System Requirements

- 4 GB of RAM memory;
- 1 processor;
- 10 GB of thin provisioned storage – SATA drive with a 15MBps throughput;
- 1 network card.

Recommended System Requirements

- 8 GB of RAM memory;
- 2 processor;
- 20 GB of thick provisioned, eager zeroed, storage – SSD drive with a 100MBps throughput;
- 1 network card.

Example: deployment in VMware environments

IMPACT ON PHYSICAL HOSTS

When deployed in a VMware environment, Eco4Cloud is a VMware virtual appliance with a single core and 4GB of configured memory. The vApp OS is Ubuntu 10.04 server edition, but it is available in any Linux distros. The vApp must be within the same vLAN where the vCenter resides, ideally vCenter and Eco4Cloud should be running on the same ESX host.

Eco4Cloud impact is minimal both on vCenter and physical host. In this paragraph, the impact on physical hosts is analyzed. Eco4Cloud employs the use of two specific performance counters: cpu ready time and ballooned memory, in order to monitor the consolidation process that leads to the deployment of virtual machines onto the physical hosts. This constant performance monitoring ensures that strict service level agreement rates can be adhered to whilst providing the flexibility in order to tune the consolidation processes in a customized way for a particular data center.

CPU ready time is defined as “Percentage of time that the virtual machine was ready, but could not get scheduled to run on the physical CPU. CPU ready time is dependent on the number of virtual machines on the host and their CPU loads.” on the VMware developers reference guide.

Regarding memory, in ESX, a balloon driver is loaded into the guest operating system as a pseudo-device driver. It has no external interfaces to the guest operating system and communicates with the hypervisor through a private channel. The balloon driver polls the hypervisor to obtain a target balloon size. If the hypervisor needs to reclaim virtual machine memory, it sets a proper target balloon size for the balloon driver, making it “inflate” by allocating guest physical pages within the virtual machine (Understanding Memory Resource Management in VMware ESX Server). The ballooning process, though, is time and CPU consuming, and should be avoided when possible, this is why ballooned memory has been chosen to analyze E4C impact on physical hosts.

Results of Eco4Cloud impact are comfortable and in a real use case are:

- CPU ready time – Eco4Cloud typically increases CPU ready time from 0.284% to 0.427%, which is well within limits. For example, VMware recommends to monitor CPU ready time when it becomes larger than 5% and then gives a warning when it becomes more than 10%;
- Ballooned memory – Eco4Cloud increases the use of ballooned memory from 0.978% to 1.193%, which is also well within VMware best practices, which specifically states that ballooning of this order is quite normal and not indicative of over commitment.

IMPACT ON VCENTER

Eco4Cloud uses VMware VI (vSphere) Java API to connect to vCenter. It is an open API, developed in Java, useful to connect to the Web Services exposed by VMware. The Web Services are the same exploited by VMware vClient software to monitor and administrate a vSphere environment.

Each time a vClient is opened, the software will update an inventory of resources querying the Web Services every 20 seconds.

Eco4Cloud software is 15 times lighter than a standard vClient instance, as it runs the same queries, but just once every 5 minutes. VMotions are issued only when needed, in active mode.

Table 3 shows the list of API queries run by Eco4Cloud.

Every 5 minutes	**When a vMotion is requested**
get service instance UUID	check VM - host compatibility
get user privileges	vMotion
get list of datastores	
get list of clusters	
get list of hosts	
get list of virtual machines	
get performances of datastores	
get performances of clusters	
get performances of hosts	
get performances of virtual machines	

Table 3 – API calls issued by Eco4Cloud

The "quickness" of the consolidation algorithm is tunable; i.e. the frequency of the vMotions can be increased or reduced.

Furthermore, it is possible to disable automatic consolidation, and allow manually each vMotion, or host standby/power on.

TECHNICAL DETAILS

IO

E4C vApp measured I/O is 400kb per minute, mostly on the database, which however is aggregated weekly just in order to reduce its size. The vApp uses a single database persisted on the virtual disk of the virtual appliance. An Eco4Cloud vApp uses performance intervals to perform data aggregation. Performance data are aggregated weekly, monthly and yearly, so that in the long term, the size of database is constant: a simple way to predict the size of the database is 15Mb * # of ESX hosts.

Differences between E4C and VMware DPM

The algorithm used by DPM is not scalable and does not allow the user to control and monitor the distribution of the workload, which instead Eco4Cloud features. E4C allows setting the following global parameters:

- the target utilization rate of a ESX host;
- the maximum utilization rate of ESX hosts.

Furthermore, E4C allows specifying a set of constraints to the consolidation algorithms, at individual VM level:

- a VM can be flagged as not migratable;
- a VM can be flagged as migratable outside its cluster.

As per hosts:

- a host shutdown can be flagged as enabled/disabled;
- minimum uptime and minimum downtime can be specified (default 120 minutes), in order to prevent repeated power on/off events.

As per clusters:

- consolidation can be enabled/disabled;
- a minimum number of always on servers can be specified;
- a default DRS behavior when e4c is paused/stopped can be entered.

vMotions issued by Eco4Cloud follow both affinity and anti-affinity rules that are set in the vCenter.

High Availability

E4C is a single core vApp. In VMware environments, this is enough to activate the HA, so if a vApp copy fails, another copy immediately restarts. Also, whenever E4C vApp does not communicate with the vCenter for a long period of time, a series of tasks are activated in order to put the system in a less consolidated state and allow it respond to load peaks more easily, even without E4C intervention (however consuming more power).

Eco4Cloud provides a failsafe mechanism. The Eco4Cloud VM consolidation daemon is monitored constantly. If it fails, two things happen. First, for protection a task is created within vCenter that starts the default "fully automated" DRS mode in order to monitor and avoid peak usage on any physical host. Second, the Eco4Cloud daemon will be restarted using a watchdog mechanism and as soon as it becomes available, it takes over the process.