



Eucalyptus
Systems

*Eucalyptus Open-Source Cloud
Computing Infrastructure - An Overview*

Eucalyptus Systems, Inc.

August 2009



EUCALYPTUS - Elastic Utility Computing Architecture Linking Your Programs To Useful Systems - is an open source software infrastructure for implementing on-premise clouds on existing Enterprise IT and service provider infrastructure. This whitepaper explains Eucalyptus from a technical as well as a business point of view and as such is intended to provide both an overview of the Eucalyptus architecture and the business value that Eucalyptus delivers.

Eucalyptus enables hybrid cloud and private cloud deployments for enterprise data centers and requires no special purpose hardware or reconfiguration. Leveraging Linux and web service technologies that commonly exist in today's IT infrastructure, Eucalyptus allows customers quickly and easily to create computing clouds "on premise" that are tailored to their specific application needs. At the same time, Eucalyptus supports the popular AWS cloud interface allowing these on-premise clouds to interact with public clouds using a common programming interface. Along with managing virtual machines, the technology supports secure virtualization of the network as well as the storage infrastructure within the cloud

environment. Eucalyptus is compatible with and packaged for multiple distributions of Linux including Ubuntu, RHEL, OpenSuse, Debian, Fedora, and CentOS and will work with a variety of hypervisors and virtualization technologies.

At Eucalyptus Systems, we develop enterprise-grade technology solutions that build upon the open-source Eucalyptus software core. Eucalyptus technology is quickly becoming the standard for on-premise cloud computing, delivering the cost efficiencies and scalability of clouds with the security and control that comes with an organization's own IT infrastructure. Enterprise Eucalyptus can provide capabilities such as end-user customization, self-service provisioning, legacy application support, customized "service level agreements" (SLAs), cloud monitoring, metering, and support for auto-scaling, as a highly available cloud platform.

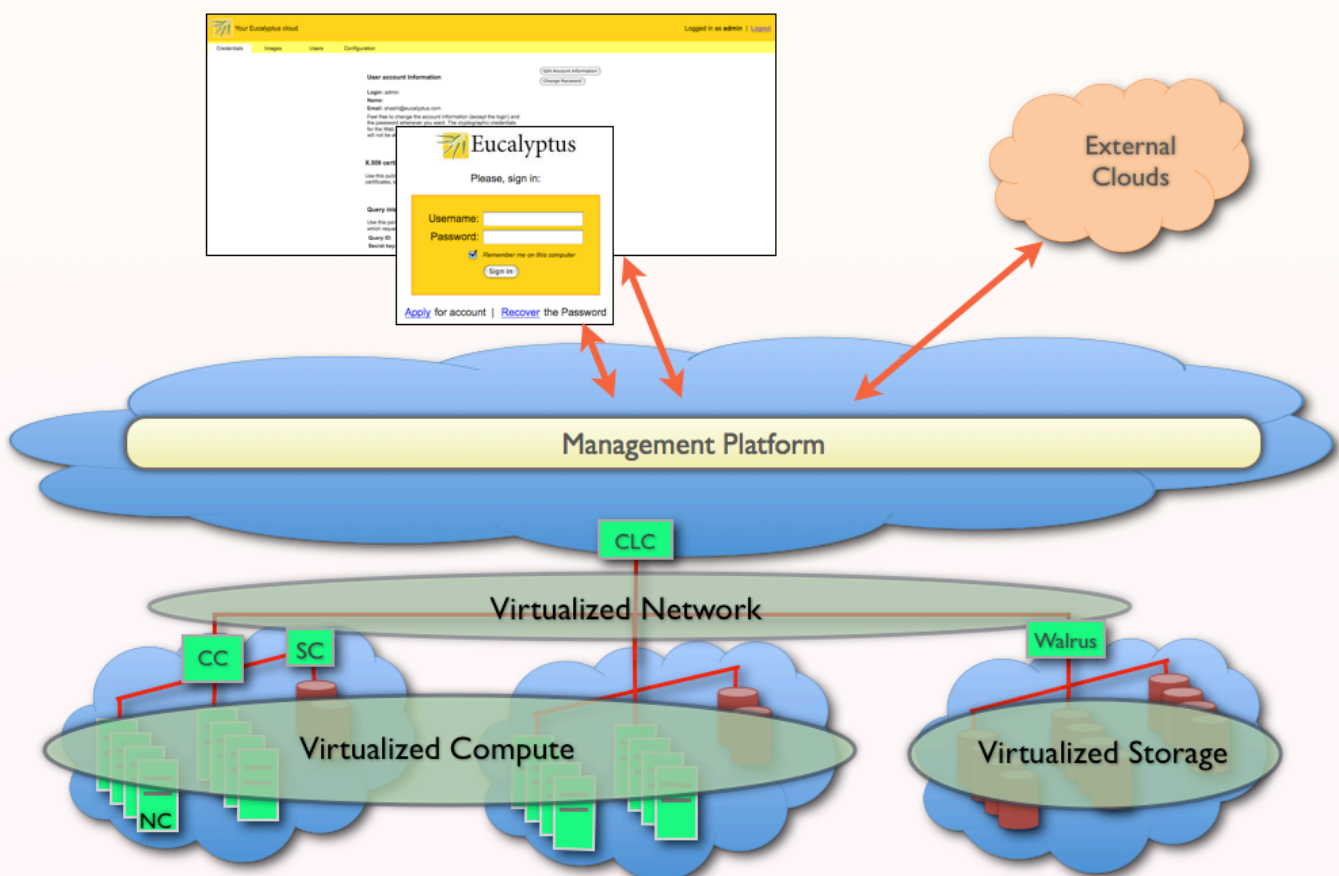
Technology

Eucalyptus was designed from the ground up to be easy to install and as non-intrusive as possible. The software framework is a highly modular cooperative set of web services that interoperate using standard communication protocols. Through this framework it implements virtualized



machine and storage resources that are interconnected by an isolated layer-2 network. From a client application and/or user perspective, the cloud API is

compatible with Amazon's AWS (both SOAP and REST interfaces are supported) although other interfaces are available as customizations.



Eucalyptus Systems, Inc. - Confidential

Figure 1. Conceptual Representation of the Eucalyptus Cloud. CLC is the Cloud Controller which virtualizes the underlying resources (servers, storage, and network). The Cluster Controllers (CCs) form the front-end for each cluster defined in the cloud. NCs are the machines on which virtual machine instances run. The Storage Controller (SC) provides block storage service (similar to Amazon EBS) while the Walrus storage system spans the entire cloud and is similar to the Amazon S3 in functionality. A Management Platform provides a one-stop console for the cloud administrator to configure and manage the cloud. The Management Platform also exports various interfaces for the administrator, project manager, developer, and other users, with customizable levels of access and privileges.



Eucalyptus Components

Each Eucalyptus service component exposes a well-defined language-agnostic API in the form of a WSDL document containing both the operations that the service can perform and the input/output data structures. Inter-service authentication is handled via standard WS-Security mechanisms. There are five high-level components, each with its own Web-service interface, that comprise a Eucalyptus installation (Figure 1). A brief description of the components within the Eucalyptus system follows.

Cloud Controller (CLC) is the entry-point into the cloud for administrators, developers, project managers, and end-users. The CLC is responsible for querying the node managers for information about resources, making high level scheduling decisions, and implementing them by making requests to cluster controllers. The CLC, as shown in Figure 1, is also the interface to the management platform. In essence, the CLC is responsible for exposing and managing the underlying virtualized resources (servers, network, and storage) via a well-defined industry standard API (Amazon EC2) and a Web-based user interface.

Cluster Controller (CC) generally executes on a cluster front-end machine, or any machine that has network



NASA builds a private cloud computing platform.

The NASA NEBULA cloud is built with Eucalyptus at the core. Powered by Eucalyptus for its virtual machine management, NEBULA cloud functions in Software-as-a-Service, Platform-as-a-Service, and Infrastructure-as-a-Service capacities. (<http://nebula.nasa.gov/services>)

connectivity to both the nodes running NCs and to the machine running the CLC. CCs gather information about a set of VMs and schedules VM execution on specific NCs. The CC also manages the virtual instance network and participates in the enforcement of SLAs as directed by the CLC. All nodes served by a single CC must be in the same broadcast domain (Ethernet).

Node Controller (NC) is executed on every node that is designated for hosting VM instances. NCs control the execution, inspection, and termination of VM instances on the host where it runs, fetches and cleans up local copies of



instance images (the kernel, the root file system, and the ramdisk image), and queries and controls the system software on its node (host OS and the hypervisor) in response to queries and control requests from the cluster controller. The Node controller is also responsible for the management of the virtual network endpoint.

Storage Controller (SC) implements block-accessed network storage (e.g. Amazon Elastic Block Storage -- EBS) and is capable of interfacing with various storage systems (NFS, iSCSI, etc.). An elastic block store is a Linux block device that can be attached to a virtual machine but sends disk traffic across the locally attached network to a remote storage location. An EBS volume cannot be shared across instances but does allow a snap-shot to be created and stored in a central storage system such as Walrus, the Eucalyptus storage service.

Walrus (put/get storage) allows users to store persistent data, organized as eventually-consistent buckets and objects. It allows users to create, delete, list buckets, put, get, delete objects, and set access control policies. Walrus is interface compatible with Amazon's S3, and supports the Amazon Machine Image (AMI) image-management interface, thus providing a mechanism for



*Ubuntu Enterprise
Cloud is powered by
Eucalyptus.*

Ubuntu Enterprise Cloud (UEC) brings Amazon EC2-like infrastructure capabilities inside the firewall. The Ubuntu Enterprise Cloud is powered by Eucalyptus. This solution is designed to simplify the process of building and managing an internal cloud for businesses of any size, thereby enabling companies to create their own self-service infrastructure. The next LTS release, Karmic Koala, is actually named after Eucalyptus. (<http://www.ubuntu.com/products/whatisubuntu/serveredition/cloud/UEC>)

storing and accessing both the virtual machine images and user data.

Management Platform provides an interface to various Eucalyptus services and modules. These features can include VM management, storage management,



user/group management, accounting, monitoring, SLA definition and enforcement, cloud-bursting, provisioning, etc.

Eucalyptus Configuration

With these components, Eucalyptus can be configured to support a wide variety of infrastructure features and topologies. For example, four different networking modes are supported, each corresponding to a different level of security and infrastructure intrusiveness allowing system administrators to “tune” each cloud configuration to meet local policy and management needs. It is also possible to deploy Eucalyptus to include different hypervisors and virtualization technologies within a unified cloud exporting a single API. Thus, a Eucalyptus cloud can act as a platform for unifying a variety of technologies (each at a potentially different point in its data center lifecycle) within a single cloud.

Benefits

✓ Eucalyptus enables virtualization of servers, network, and storage in a secure manner, thereby reducing the cost, increasing the ease of maintenance, and providing user self-service.

Lilly

Lilly builds a private cloud with Eucalyptus.

Pharma giant Lilly, has deployed Eucalyptus to leverage their existing hardware resources to build an on-premise cloud that is interface compatible with Amazon AWS.

- ✓ The modular design of Eucalyptus enables a variety of user interfaces, bringing the benefits of virtualization technology to a broad range of users (admins, developers, managers, hosting customers) and provides a platform for service providers to devise profitable consumption-based pricing models.
- ✓ VM and Cloud snapshot features provide an exhaustive set of opportunities to improve cluster reliability, template manipulation, and automation. This makes the cloud easy to use, reduces learning time for the average user, and reduces the turn around time for projects.
- ✓ Leverages existing virtualization technology, supports Linux-based operating systems, and supports multiple hypervisors.



✓ Ease of cluster/availability-zone management provides the administrator/user with a host of opportunities to form logical groups of servers, storage, and networks on a per-project/per-user/per-customer basis.

✓ The core Eucalyptus framework will continue to remain open-source. This provides the users access to the source code as well as an opportunity to leverage the contributions from a world-wide community of developers.

✓ The ongoing effort to develop public-cloud compatible interfaces is a unique advantage for users who intend to cloud-burst to other public-clouds (also known as hybrid clouds -- a private on-premise cloud, in this case, a Eucalyptus cloud, working seamlessly with a public cloud).

✓ A vibrant ecosystem built around the Amazon AWS can be leveraged. For example, RightScale, CohesiveFT, Zmanda, rPath are just a few of the partners that deliver solutions for Amazon AWS that in turn work seamlessly with Eucalyptus.

What To Do Next?

Here is a simple three step process to deploy an on-premise cloud free of cost.

1. Identify a project in your organization which involves virtualization - data center virtualization, server consolidation,

streamlining IT management, enabling self-service provisioning, implementing SLAs, setting up an on-premise cloud leveraging your existing virtualization software, planning for a hybrid cloud in conjunction with Amazon EC2, converting your software infrastructure to a Software-as-a-Service (SaaS) based model, or simply experimenting with an open-source cloud platform that transforms your IT infrastructure into an EC2-like on-premise cloud. Any project in your organization, which intends to optimize IT resources and deliver improved IT services can use Eucalyptus.

2. Select an application you want to run

on your on-premise cloud - Simply choose your favorite application that you want to run on a cloud platform. Web services, database-centric applications, business intelligence software, productivity tools, data mining applications, media processing, rendering, numerical analysis, or other mission-critical applications can be deployed on a Eucalyptus cloud.

3. Set up a Eucalyptus test-bed to deploy

your on-premise cloud - Pick a handful of machines (at least two), choose your Linux distribution, download the open-source Eucalyptus software and follow the simple installation guide on our website to get your on-premise cloud up and running. You will instantly see



how your organization can benefit from an EC2-compatible on-premise cloud. If you require assistance with the installation, configuration, optimization, or customization of the Eucalyptus cloud, experts at Eucalyptus Systems are available. Please feel free to contact us to help you set up your proof of concept.

Conclusion

Focusing on the cost advantage for IT infrastructure and data center management, Eucalyptus provides a unique framework with a variety of interfaces to manage the resources. The hardware, network, and storage can easily be consolidated under the Eucalyptus cloud, hiding the heterogeneity in hardware, software stack, policies, and configurations.

Eucalyptus Systems develops enterprise-grade technology solutions built on the open source Eucalyptus software for private and hybrid cloud computing. Originally developed as part of an academic research project, Eucalyptus technology is quickly becoming the standard for on-premise cloud computing, delivering the cost efficiencies and scalability of cloud architecture with the security and control of deploying on an organization's own IT infrastructure. Eucalyptus Systems' mission is to support the open source Eucalyptus platform and to deliver private and hybrid cloud computing solutions for large-scale

enterprise deployments. For more information about Eucalyptus, please visit www.eucalyptus.com.

Contact Us:

Phone: 805-845-8000

Email: info@eucalyptus.com

Eucalyptus and Eucalyptus Systems are pending trademarks in the U.S. All other trademarks are property of their respective owners. Other product or company names mentioned may be trademarks or trade names of their respective companies.