

Low-cost Server Hosting at UC Santa Cruz

Project title: Low-cost Server Hosting at UC Santa Cruz

Submitter: Carol Jordan
ITS Data Center Manager, UCSC
831-459-2631, carol.jordan@ucsc.edu

Project leader: Cliff Pearson, ITS Unix Systems

Project team: Josh Sonstroem, ITS Unix Systems
Jerry Chen, Intern

Introduction

UC Santa Cruz has developed a unique and innovative new service - **Nebula**. Nebula is designed to provide the best of both cloud hosting and virtualization, and we've just finished up a very successful first year offering Nebula as a production service.

Cloud hosting is great. Users can create, configure, and destroy servers at will. Application deployments grow or shrink automatically as utilization changes. There are no up-front capital expenditures. But, there can be a steep learning curve. A credit card or a purchase order is needed to get started. To take advantage of automatic re-sizing, an application needs to be re-factored for the cloud, not just moved to the cloud.

Virtualization is great. Our virtual-hosting service has a robust architecture, and is managed to meet the availability and security requirements of our campus's enterprise applications. The complexity of hardware management is concentrated in a small team, so server administrators never worry about hardware once they've gone virtual. But, the cost of that infrastructure means that the typical virtual-hosting service is not affordable for applications with lesser security and availability requirements.

UCSC has been able to bridge the gaps between Cloud hosting and Virtualization by offering the Nebula service.

To highlight how unique and innovative Nebula is: the primary engineer, Cliff Pearson, has been invited to speak at the Interop Container Summit in Las Vegas this year, to share with the rest of the world about what we've been able to do here at UCSC.

An example of a recent customer success with Nebula is the Transportation and Parking Services department. The department has been able to provide a system that tracks campus shuttles in real time, using Nebula to turn great work by students into a formal, ongoing service.

You can learn more about the service, SlugRoute (<http://slugroute.ucsc.edu>), in this recent article from Santa Cruz Tech Beat:

<http://www.santacruztechbeat.com/2016/04/28/ucsc-launches-slugroute-mobile-app-track-shuttles-real-time/>.

Nebula is also used by many of the local IT staff who provide direct academic and research support. With Nebula, they can take care of their clients' requests quickly, and provide them with the benefits of Data Center hosting without the typical costs or administrative overhead.

Project features

- The Nebula service is free to customers.
- A key value of the Nebula service is the ability to provide broader access to Data Center's resources, such as generator-backed power, regular security scans, and experienced engineers to manage the underlying operating system and hardware. We view this service as a lower-cost and lower-availability complement to our existing virtual-hosting service.
- Customers can provision, manage, and destroy their own virtual servers. The self-service portal is simple to use.
- Customers can use Nebula to incubate innovative work by our talented students, like SlugRoute. Nebula gives student workers a safe, and inexpensive, place to try new things.

Our collaboration

We've been successful not just because of the value that Nebula brings, but also because of our collaborative approach in developing the service.

- **Outreach:** We continue to provide a lot of hands-on labs, write tutorials, and work side-by-side with departments across campus as they learn about the service. We've been to lots of team meetings and all-hands sessions.
- **Community:** We've created a self-supporting community of users, who are generous in helping each other.
- **Autonomy:** We developed the parameters for the service in collaboration with the IT liaisons in the academic divisions. This group regularly reinforced the importance of providing autonomy to Nebula customers, so that they can create, manage, and destroy their servers on their own.

Testimonials

We wanted to put this technology right into the hands of the people who could make the most of it, to easily provision and manage systems for themselves, their departments, or labs and research groups. They've told us that we've succeeded by helping them to succeed --

- One of our best testimonials was from a local IT specialist who said, "*It would be great for central IT to provide more services like this -- just give us the tools and let us do things for ourselves.*"

- *“Without this service, I would still be king of the recycled servers and unable to fix little problems...It is also providing some professional development for my staff and students where I don't have to worry about sysadmin issues.”* -- Peter McMillan, Director, Client Relationship Management
- *“I just talked to a client about it yesterday, before [Nebula] she would be one of those people who would set up a box in her office and run a LAMP stack on it, but...we can now just do it in the cloud in a VM and if the project idea does not get off the ground, as research ideas often don't, there's no cost and nothing to try and repurpose. Just delete the VM and go back to work. If the project takes off the infrastructure allows us to scale it up.”* -- Paul Sosbee, Academic Computing Expert for Social Sciences Computing
- The Unix team uses Nebula internally for short term software development projects and proof-of-concept work. The team has been able to decommission old, inefficient hardware by using Nebula to deploy systems-management tools and software repositories.

Technology used in the project

Nebula is built on Joyent Triton, which is based on the open-source fork of Solaris. This let us leverage the deep engineering experience that our staff already has, and let us use tools we were already familiar with, including ZFS, DTrace, and operating-system virtualization through local zones. With an upgrade this year, we are now able to offer servers running more familiar operating systems such as CentOS or Ubuntu.

Initial implementation

We began with a proof-of-concept on reclaimed hardware. At that point, our intention was to demonstrate the benefits to the campus community. It allowed potential customers to try out their applications and get a sense for how the service works. It also provided the team with insight into the features and architecture the community wanted, ultimately helping us develop the support model we use today.

Current implementation

With the success of our proof-of-concept, we secured funding for a formal pilot project with a three-tiered model offering a variety of paid and unpaid levels of support.

Using specifications developed by Joyent, we purchased three purpose-built Supermicro servers, each with 128GB of memory and 16 CPU cores. One of the three machines (called the “head node”) is reserved for internal use. Customers' servers live on one of the two “compute nodes.”

This configuration currently supports more than 100 customer VMs, utilizing between 256MB and 16GB of RAM each. The current infrastructure has capacity to support around 300 customer VMs.

After one pilot year, we received approval to proceed as a production service for fiscal year 2015-16.

Timeframe of implementation:

June 2012: the seed is planted

September 2012: the Unix team starts working with Nebula for internal tools

November 2013: the formal support contract is signed with Joyent

May 2014: the campus pilot is approved and the hardware is ordered

April 2015: the formal, production service is approved

June 2016: milestone -- one year of formal service!

Relevant URLs:

UCSC's cloud-computing service definition:

<http://its.ucsc.edu/data-center/dc-services/joyent.html>

Joyent private cloud:

<https://www.joyent.com/private-cloud>

Conclusion

We think Nebula is a fine candidate for a Sautter Award!

We're **innovative** - Nebula is a completely new way for Information Technology teams to provision and manage their own servers. We took a small, internal service and found a way to make it available campus-wide. People were asking for a less expensive way to host divisional or departmental services, and we delivered a solution that not only does this, but allows them to retain autonomy over their own virtual infrastructure.

We're **collaborative** - Since this is a new kind of service, we developed a new kind of partnership with the distributed IT community. We worked closely with UCSC academic divisions to shape the service for adoption across campus by a wide variety of departments. We also developed a support community where users of the service can document their solutions, share them with others, and support newer users as they build out their own applications and architectures. Finally, we spent a lot of time working with our Infrastructure Security and Networking teams to establish a strong, flexible, and secure infrastructure to support the service.

We're **improving operational efficiency** - Nebula lets people across campus provision and manage servers in a way that is even easier (and more sustainable) than running a server under their desks. The user interface is simple and easy to use -- and it lets people learn about cloud computing without the complexity and expense of public cloud hosting. With this project, we can now extend the benefits of the data-center environment more broadly than we ever could before. And because we can host lots of virtual servers on relatively inexpensive

hardware, this service makes very efficient use of servers -- reducing our carbon footprint and helping us to achieve our sustainability goals.

Nebula is **shareable** - We can federate our UCSC-based cloud with other campuses; users could then specify the campus (availability zone) where they want their virtual servers to run.