



NASA's Nebula Cloud Computing Initiative

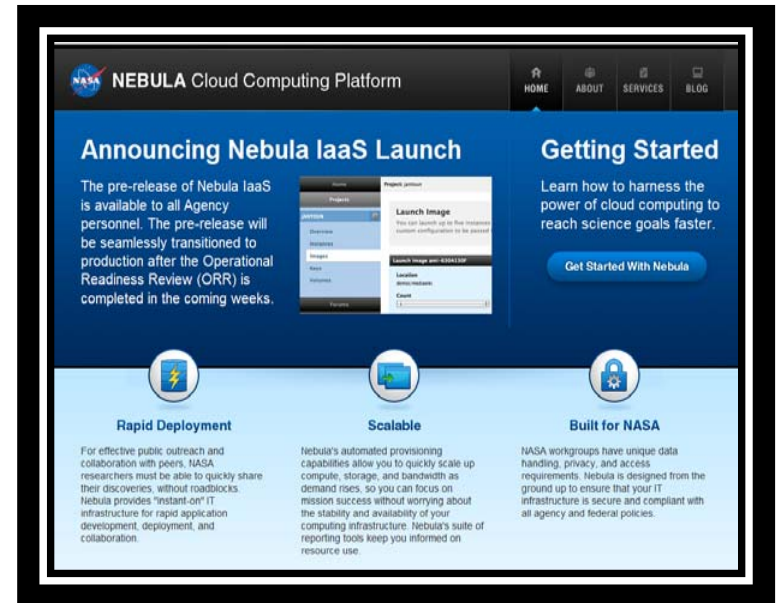
Cloud Innovation at NASA

James F. Williams
CIO, NASA Ames Research Center
22 February 2012



NASA Nebula Cloud Computing

- ▶ One of the first cloud computing platforms built for the Federal Government by the Federal Government
- ▶ Publicly launched IaaS with the White House as the first customer
- ▶ Basis of OpenStack Compute, aka “Nova”
- ▶ Today have over 400 Nebula accounts across nine NASA Centers + Jet Propulsion Lab + Headquarters



A Lot Has Been Said About Nebula...

Nebula, NASA's cloud-computing platform, is helping NASA to engage the public through the viewing and exploration of the Moon and Mars in unprecedented resolution.

Vivek Kundra, U.S Chief Information Officer

What do you get when you combine cloud computing and data center containers? You get NASA's Nebula, the space agency's new data powerhouse, which provides on-demand computing power for NASA researchers.

datacenterknowledge.com

December 2009.

"NASA, they're sort of the summa cum laude when it comes to open source"

June 2011 Network World

"Nebula could someday be the Linux of Cloud"

Former NASA Ames IT Official
2010



"Rackspace and NASA are taking an amazing step towards my vision of an open cloud future."

TIM O'REILLY, OREILLY.COM



Slide 3

R2 Add/change quotes. Audience is different. Use less humor.

Open source summa....
Raymond G. O'Brien, 7/5/2011

NASA Nebula

How it Began



The Challenge: Providing Suitable Hosting Options for NASA's Dynamic Websites

www.nasa.gov

Public portal targeted at NASA's static sites

But still have 1000's of other NASA websites, many of them dynamic



Photo Credit: www.starwars.com



There's a better way to do this...

- ▶ Problem: How do we get web developers to stop building out their own independently managed site hosting infrastructure
- ▶ Solution: Give developers a better alternative to status quo

NASA.net was born

Funded as local investment by NASA Ames with goal of eventually providing platform services to entire Agency



NASA.net

- ▶ Setting: Basement of NASA Ames Research Center (ARC) Building 200, in an old conference room
- ▶ Imagine: Small team of developers working delivering Platform-as-a-Service
 - ▶ Code hosting
 - ▶ Continuous integration
 - ▶ Bug Tracking
 - ▶ Best Practices in code development



But after working on Platform as a Service prototypes...

- ▶ Quickly determined that in order to run a web application framework properly as an Agency service, we needed dynamically scalable infrastructure capability



Over Indian Food in Mountain View, CA



Photo Credit: <http://www.tandooribistrosj.com/>

- ▶ The NASA.net Chief Architect pitched the “cloud” idea to us
- ▶ Decision to build out an IaaS capability just to support NASA.net
- ▶ A timeout was called on NASA.net to implement the pre-requisite IaaS capability
- ▶ Nebula was born



Problem: Sustained Project Funding

Everyone: “Great Idea!”



Project Team: “Thanks!
But we need to talk”



Why?



FY07:
Plan FY09 Budget



FY09:
Line item for
cloud?

Planning cycle would
have required
advocacy for project in
2007



About the same time, White House “cloud computing” initiative was to be announced

White House CIO to Outline Cloud Computing Plan at NASA Ames

Posted by Doug Messier on September 12, 2009, at 6:42 pm in NASA and NASA Ames

Tags: cloud computing, lori garver, nasa, nasa ames, vivek kundra.

Comments: no responses

0 Comments

NASA PRESS RELEASE

NASA will conduct a news conference at NASA’s Ames Research Center at 10 a.m. PDT on Tuesday, Sept. 15, 2009, featuring Vivek Kundra, White House federal chief information officer. He will outline his vision for a new federal government cloud computing initiative.

NASA Deputy Administrator Lori Garver and top Silicon Valley information technology leaders are scheduled to attend the news conference. The event will be broadcast live on NASA Television.

For NASA TV streaming video, downlink and scheduling information, visit:

<http://www.nasa.gov/ntv>

Editor’s Note: During her speech during the ISU graduation ceremony at Ames two weeks ago, Garver mentioned she had been quite impressed by a briefing earlier in the day on the cloud computing work that had been taking place at the center.



Vivek Kundra



Shortly Afterwards, White House Issued Cloud Guidance to Agencies

In coordination with the data center consolidations, agencies should evaluate the potential to adopt cloud computing solutions by analyzing computing alternatives for IT investments in FY 2012. **Agencies will be expected to adopt cloud computing solutions where they represent the best value at an acceptable level of risk.**

http://www.whitehouse.gov/sites/default/files/omb/assets/memoranda_2010/m10-19.pdf



Getting Resourceful: Hosting USASpending.gov 2.0 on Nebula

- ▶ Vivek Kundra thought it would be a great idea to host the White House's enhanced USASpending.gov website on Nebula
- ▶ We struck a deal to fund the development of USASpending.gov 2.0 on Nebula, providing immediate funding (round 2) to partially offset Nebula project costs
- ▶ The Nebula project was also formalized to a larger degree by adding project management resources
- ▶ The Nebula project now had pretty good air cover too
- ▶ However, the agreement generated developer discontent



Another Deal: Processing Mars Imagery for use with Microsoft's WWT


- ▶ NASA entered into a Reimbursable Space Act Agreement with Microsoft to deliver high resolution Mars imagery to Microsoft's very popular World Wide Telescope site
- ▶ NASA's image processing team chose to use Nebula for image preparation and delivery, providing additional Nebula funding

Microsoft and NASA Bring Mars Down to Earth Through the WorldWide Telescope

07.12.10

Today, Microsoft Research and NASA are providing an entirely new experience to users of the WorldWide Telescope, which will allow visitors to interact with and explore our solar system like never before. Viewers can now take interactive tours of the red planet, hear directly from NASA scientists, and view and explore the most complete, highest-resolution coverage of Mars available. To experience Mars up close, Microsoft and NASA encourage viewers to download the new WWT|Mars experience at <http://www.worldwidetelescope.org>.


Dan Fay, director of Microsoft Research's Earth, Energy and Environment effort, works with scientists around the world to see how technology can help solve their research challenges. Since early 2009, he's been working with NASA to bring imagery from the agency's Mars and Moon missions to life, and to make their valuable volumes of information more accessible to the masses.



A stunning new image of Mars now available in the WorldWide Telescope. Image credit: Microsoft/NASA
Click photo for larger image.



Then things really started to happen



Initial Path Chosen for IaaS

- ▶ Private cloud capability targeted at science-class workloads, not just for providing scalable web hosting platform
- ▶ Almost exclusive use of open source components
- ▶ Any IaaS code produced by NASA would be released as open source
- ▶ Commodity hardware with no dependence on proprietary management interfaces
- ▶ 10 GbE networking (LAN and WAN)
- ▶ Physical infrastructure resources located at a minimum of two geographically dispersed locations



Over One Weekend...

- ▶ Nebula team members decided to spend a couple of days coding a new cloud controller from scratch. **This is what became “Nova.”**
- ▶ Some thoughts that drove design:
Monolithic is bad – each component should scale independently
 - ▶ APIs are good – you shouldn’t have to use web UI to configure
 - ▶ Simple things should be simple, hard things possible...
 - ▶ The hard parts (hypervisor, storage system, networking) is done by others. Build a cloud the same way the way you build a scalable web application...
 - ▶ **Nova has been empowering NASA users since May 2010**



Rackspace Called NASA



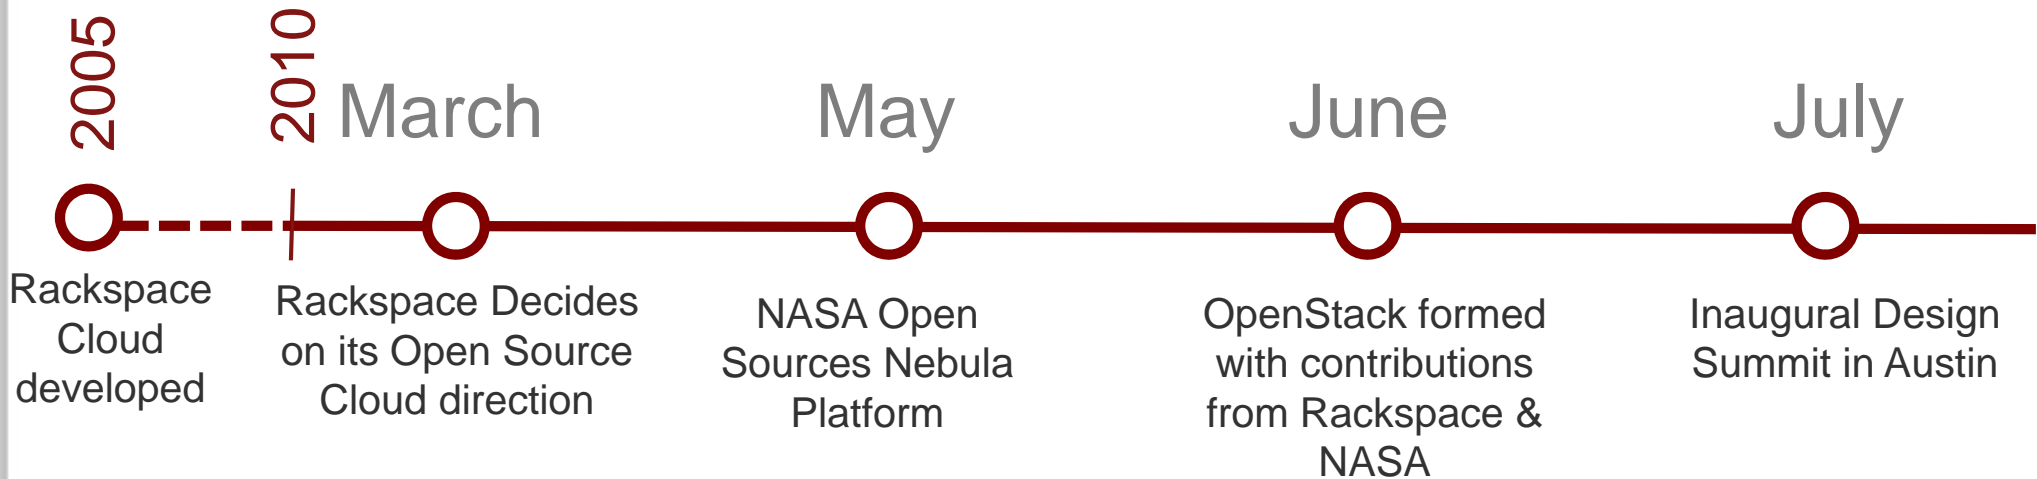
- ▶ “Wow. Can we meet your team of 400 developers?”



- ▶ “Sure. We got 8 developers. And not all of them full time. “



The Birth of Openstack Timeline



NASA DIY Cloud – Why?

- NASA already had a History of IT Innovation
 - Supercomputing (Performance, CFD Codes, PBS Scheduler)
 - One of the First Nodes on DARPAAnet
 - Early promoter of TCP/IP
 - Operate one of the Internet's 13 DNS Root Servers
 - Operate West Coast Federal Internet Exchange
- We had a need, a vision, the developer talent, and believed NASA could make a significant contribution to the future of cloud
- The OpenStack initiative and the rapid growth of its community subsequently validated this belief




NASA Nebula's contributions to OpenStack align with the Administrator's strategic goals

- ▶ Promote enhanced cooperation with international, industry, other U.S. government agency, and academic partners in the pursuit of our missions.
- ▶ Facilitate the success of a viable commercial space industry to provide assured U.S. access to low Earth orbit for cargo and crew and acquire, mature, and infuse commercial capabilities across all NASA activities



NASA Nebula + Openstack supports the Agency's Goal of Public Participation:

- ▶ Goal 6: Share NASA with the public, educators, and students to provide opportunities to participate in our mission, foster innovation and contribute to a strong National economy



Nebula is also a Key Component of NASA Open Government Plan

About Open Government



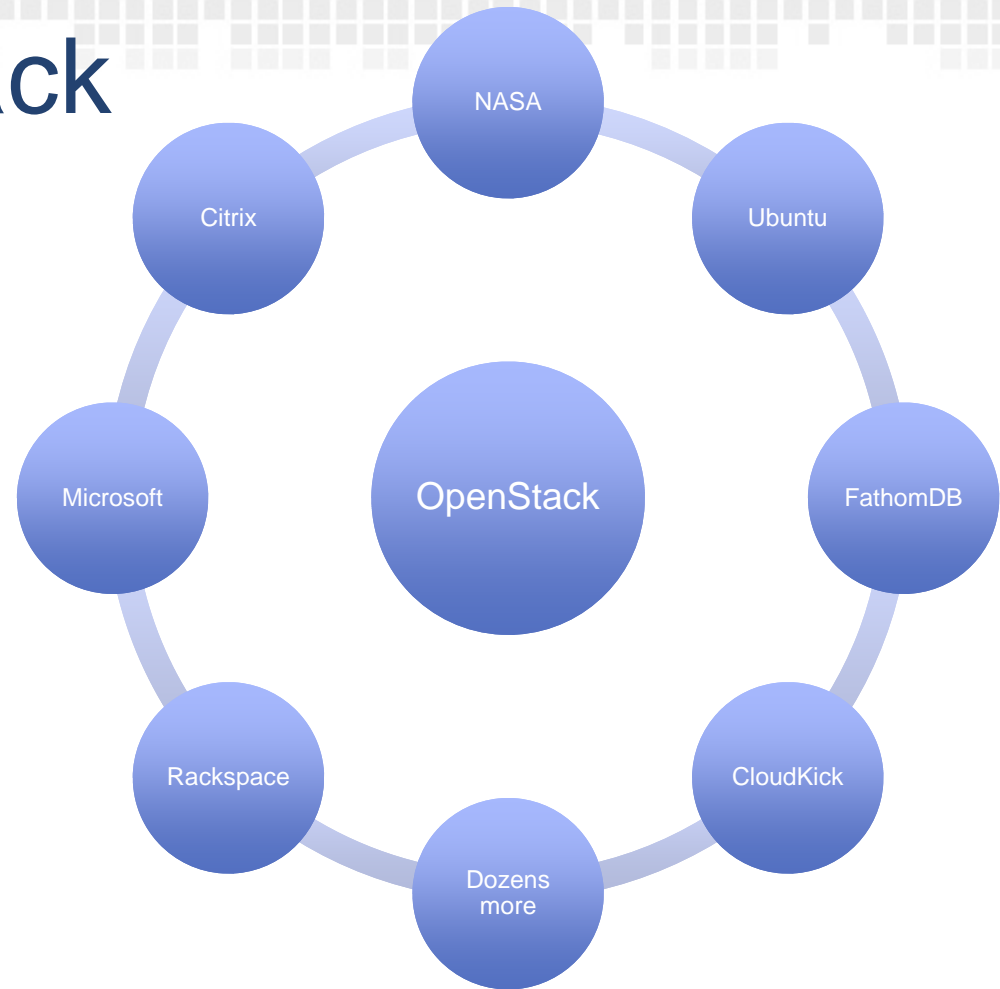
NASA's founding legislation in 1958 instructed NASA to "...provide for the widest practicable and appropriate dissemination of information..." The principles of Open Government have been embedded in our operations for 50 plus years. We recognize that open government is a process rather than a product, and have taken a continuous-learning approach. Please browse and read Version 1.0 of the [NASA Open Government Plan](#). We invite you to contribute your ideas, solutions, and comments on every element of our plan.

Highlights of NASA's Open Government Activities

- Flagship Initiatives: Three "Flagship" initiatives at NASA's will take Open Government to a new level:
 - **Policy:** NASA is working to make [open source software development](#) more collaborative at NASA to benefit both the Agency and the public.
 - **Technology:** [NASA Nebula](#), as the government's first open-source cloud computing platform, offers an easier way for NASA scientists and researchers to share large, complex data sets with external partners and the public.
 - **Culture:** The creation of a new [NASA Participatory Exploration Office](#) will infuse more public participation into NASA's mission.



Continue Openstack Collaboration




- Hundreds of contributions of large and small companies have improved OpenStack since release
- **NASA has benefited from OpenStack contributions** integrating ideas and patches into Nebula after appropriate review.



OpenStack & NASA

- ▶ How we are ensuring OpenStack works for NASA
 - Initially very active participant, ensuring NASA's requirements are in the DNA of OpenStack
 - As OpenStack matures NASA's involvement will shift from contributor to user. "As community stands up, we stand down."



Use Case: SERVIR.net

Background



- SERVIR – the Regional Visualization and Monitoring System
- Applies Earth observations and predictive models to support decision-making by government officials, managers, scientists, researchers, students, and the public.
- Improved monitoring of environmental threats, air quality, extreme weather, biodiversity, and changes in land cover.



- Transitions NASA satellite data and research capabilities to the weather forecasting community.
- Demonstrate societal benefit of NASA resources by improving short-term weather analysis and forecasting capabilities.



Use Case: SERVIR.net

Weather Forecasting

- Severe weather is a natural hazard of interest to both SERVIR and SPoRT.
- Use the Weather Research and Forecasting (WRF) Model to produce high-resolution, short-term forecasts.
- Instances can be used to:
 - Use one instance for a single region.
 - Share resources for a high resolution run or a larger forecast domain.
 - Provide rapid response to new events or research opportunities without impacting other resources.
- Other potential applications:
 - Air quality (GOCART)
 - Soil moisture (NASA LIS)
 - Hydrologic modeling



WRF Image

Instance #1



Image: SPoRT

Instance #2




Image: UCAR



Use Case: SERVIR

Modeling Capabilities with Nebula

- Application Concept:
 - Create Nebula images that are capable of supporting the research and operational goals of both SERVIR and SPoRT.
 - Potential benefits:
 - Rapid deployment of standard models to respond to natural disasters, without disrupting other activities.
 - Reduces the installation and maintenance of IT resources at remote or offsite locations.



Use Case: WISE Mission

- ▶ Wide Field Infrared Survey Explorer (WISE): Images the sky with greater than 8X redundancy
 - ▶ Helping NASA find the most luminous galaxies in the universe and the closest stars to the sun. 2.7 million images captured.
 - ▶ Issue: Encountered a short-term need for a large number of small servers and also needed a server with a large memory footprint, did not have access and could not justify cost for his needs alone


- ▶ Nebula Project #1: 2000 distant galaxies

- Increase resolution with processing
- 100 CPU hrs per galaxy
- We upped instance quota to get started

- ▶ Nebula Project #2: Some sky areas require huge RAM-based processing.

- We set up an 80 GB RAM instance
- Finished first phase on Nebula in two days.


WISE Mapping the Infrared Sky




<http://wise.ssl.berkeley.edu/documents/FactSheet.2010.1.4.pdf>



WISE Region Processing




WISE Cloud Processing



WISE Galaxy Processing

blue=W1 ($3.4\mu\text{m}$); **green=W3** ($12\mu\text{m}$); **red=W4** ($22\mu\text{m}$)



Nebula Value to WISE Project

- ▶ Hundreds of galaxies can be hi-res processed to provide higher angular resolution images for better studies of star formation, galactic structure, etc.
- ▶ Thousands of galaxies could be hi-res processed to resolve & measure source sizes, etc.
- ▶ Hundreds of thousands of galaxies can be hi-res processed to differentiate point-like (active galactic nucleus or nuclear star formation) from distributed (merger or spiral galaxy) emission.
- ▶ Several large regions nearby can be surveyed for distributed star formation in our galaxy.
- ▶ A few large regions can be processed to tremendous depth.
- ▶ The whole sky can be hi-res processed.



Key Lessons Learned.....

- ▶ Don't underestimate the effort to transition from an IT innovation effort to a institutionally accepted IT service (the technology is just a piece of it)
- ▶ Timing can make a big difference
- ▶ Innovation within an established institution can be disruptive in ways you that are very difficult to anticipate
- ▶ Great things can be done by a motivated open source development community
- ▶ Public-private collaborations can be very powerful
- ▶ NASA can make a significant contribution to the development of emerging information technologies



NASA and Nebula



- ▶ Regardless of the road and detours taken to get to this point.....
 - ▶ NASA does support Nebula ... At all levels. From the NASA Administrator to the stock room. A group of very committed people continue working together to make it happen....technical, business, management and political.
 - ▶ However, the team is having to prove itself at every level by going through the same reviews and funding justifications as other NASA initiatives and projects
 - ▶ The Nebula team looks forward to the day that it can point to the significant impact of its work both inside and outside NASA



Nebula - NASA's Cloud Computing Initiative

