

Infrastructure as a Service (IaaS) for NASA's Mission Directorate
 NASA's Nebula Pioneers a New Frontier for Cloud Computing
 Keith Shackelford and James Williams
 Ames Research Center
 National Aeronautics and Space Administration
 Moffett Field, Ca 94035
keith.c.shackelford@nasa.gov

Abstract

Scientists focusing on space research and space-based Earth-research are poised to realize the benefits of cloud computing. Features of cloud computing, such as "instant-on" IT infrastructure, automated provisioning capabilities, and quick scale-up services, are ideally suited for organizations seeking to process large datasets fast, easily share them with colleagues and ultimately store them securely at a good price. In 2009 NASA started work on developing its own cloud, Nebula, with the view to provide a cutting-edge, dynamically scalable computing platform capable of meeting Federal/Agency IT standards. This paper examines the Nebula team's process of exploring the cloud computing paradigm ultimately arriving at the decision to build NASA its own private cloud for science and mission use. Examples of current and future uses of Nebula are provided as insights into how cloud technology can be used to advance earth and science research.

1. Introduction

The future of cloud computing is assured; in the next 3 to 5 years more than 80 percent of the world's computing and data storage could occur "in the cloud."¹ Playing a critical role in determining the development of the Internet's next generation, cloud computing has evolved into a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.² Scientists studying space as well as scientists studying conditions on Earth from space are re-thinking their strategies to process, share and store large datasets with the advent of this technological advancement.

Leading IT researchers rank the cloud model as the most significant strategic technology for 2011 and growth is expected to continue.³ Decisions made now by businesses and governments will define how the cloud computing model will be delivered, what standards will be established and how security and privacy matters are to be resolved. Space research organizations need to move with immediacy to better understand cloud computing's capabilities, plan strategically, and voice their needs and concerns to the architects of this latest technological trend. The cloud's applications to research in space and the conduct earth science research from space are becoming apparent as the utility of the cloud quickly becomes ubiquitous in the IT world.

2. Building Nebula for Science

In September 2010, Nebula Infrastructure Service 1.0 was released in Beta to over 200 users in all major NASA Centers. NASA's believed Agency scientists would benefit from the flexibility, scalability and customization offered by the development of a private NASA cloud. Nebula, as it would be named, could enable scientists to work more efficiently, share their data with greater ease and this would also help the Agency save money. Since cloud technologies minimize the investment necessary to address a given aggregate workload, cloud computing maximizes necessary computing resource utilization. NASA's Ames Research Center calculated it would yield a significant 85 percent savings⁴ if it could deliver the equivalent amount of servers to its scientists using a cloud platform.

Over the past several years, NASA has undertaken an effort to evaluate and prototype the application of the latest industry trends and capabilities in the delivery of Host Computing resources to NASA customers. Since industry is already delivering General Utility Computing Services to commercial and federal

customers, and the Science Mission Directorate is already providing High Performance Computing Services, an approach to Cloud and Grid Computing Services is needed. To that end, NASA's Nebula Project focused on the unique challenges of delivering Cloud Computing Services to the Science Community. The following subsections contain a brief explanation of some of the scenarios that require Cloud Computing capabilities in the Science Community.⁵

2.1. Limited Resource Availability

Often, research scenarios are intentionally limited to account for the constraints of time, cost, and host computing availability. This results in studies that are not as comprehensive as scientists would like them to be. Decisions may be made on insufficient or incomplete information that can have longer-term risks and consequences. In a similar situation, scientists often need to mature and approach to a specific research problem through an initial prototype, to a pilot, and then into a full scale production system. In the early stages of their effort, they typically make do with what Host Computing resources they have readily available and then attempt to scale up to alternative platforms as their approach gains merit. Because they do not start out with a platform that is designed to scale up and down, time, money, and resources can be wasted porting approach to different platforms. There are also times when scientists must suspend the development of a particular approach and bring it back up as time, money, and scope permit.

2.2. Maturing Scientific Understanding

Often, programs and projects will implement information technology solutions to satisfy their current research needs. As the project matures and collects large volumes of data over time, new algorithms or techniques for processing data will be developed. These changes will cause scientists to re-process existing data. However, because of the increases in data volume over time, the project's Host Computing strategy may be sub-optimal for handling the large, temporary increase in host computing capabilities. The result of processing the data on the existing limited infrastructure capabilities means that months may be consumed rather than days or hours. Investing in excessive host computing capabilities to

account for this eventuality will mean that the project must "sink" funding into fixed assets that will largely go underutilized.

2.3. Science Data Analysis

After collecting data through mission activities, scientists perform research for the principle investigators to extract useful information and turn that data into knowledge. Often, they will correlate large sets of data from other mission activities to create new insights. However, the totality of data that NASA collects is enormous and presents a challenge not only from a storage perspective, but also from the perspective of processing the data to integrate and mine it for results. Scientists must expend time and money to find effective ways to address their needs. Sometimes, they must result to using capabilities that are not ideally suited for their specific problems. This can cause a reduction in quality, an increase in costs, risk, and a loss of time.

2.4. Procurement, Security and Cost Overhead

Scientists often compete for research grants against non-federal entities. Federal scientists are at a disadvantage because there are several legal, policy, and procedural challenges they must overcome that universities and commercial entities do not encounter. For example, the use of any Host Computing environment used by federal scientists must undergo Federal Information Security Management Act (FISMA) accreditation. This also drives the federal scientists to use particular technologies and controls like the Homeland Security Presidential Directive (HSPD) 12 for authentication. Without having a pre-configured system to address these, and other challenges, federal scientists find that they are at a disadvantage due to the cost, time, and technical overhead involved. A secure cloud computing infrastructure and platform that provides those federal security requirements would be a tremendous benefit to scientist.

3. NASA Selects Nebula for Science Cloud

In the early stages of research NASA considered the use of commercially available cloud services to deliver more robust applications for large-scale projects. The public clouds, however, as the exploration team discovered, target the needs of the broader customer base and did not meet NASA's scientific computing requirements with regard to science-class performance and capacity requirements. The Nebula development team also found that certain desirable cloud services or features were simply not available. Collaborating with the Goddard Space Flight Center (GSFC) innovation team, Nebula would need to perform the targeted research and development necessary to provide certain capabilities advantageous to NASA within the Nebula science-class cloud computing environment. The cloud computing for Nebula has several key requirements that are unique to the Science Community. They are listed in subsections below.

3.1. Nebula Unique Benefits to Science

Open Source brings numerous benefits to NASA software projects, including reduced or no software license costs, increased software quality, reduced development costs, faster development cycles, and reduced barriers to public-private collaboration through new opportunities to commercialize NASA technology. The Nebula development team realized numerous efficiencies could work ideal in an open source environment and welcomed the opportunity to build the project in cooperation with the Silicon Valley IT community.

Open Source advances NASA software projects, including increased software quality, reduced development costs, faster development cycles, and reduced barriers to public-private collaboration through new opportunities to commercialize NASA technology. Nebula has become an example of what open source can mean for the Agency following the successful co-founding of OpenStack and the accelerated launch of Nebula 1.0.

Nebula provides for increase processor power, performance and count. Nebula science users can request larger or smaller instance sizes based on their performance requirements. Also, Nebula provides scientists with the capability to know their specific chip set which is important in some science processing and this capability is not available in commercial cloud offerings.

Nebula networking provides faster node interconnect speeds than General Utility Computing and existing commercial cloud services currently

provide up to 10G networking. This allows for increased storage capacities and performance

The Nebula ObjectStore feature benefits scientists by providing them with a capability to store information at a greatly reduced cost because it scales linearly with size unlike other storage capabilities.

Nebula provides Federal Information Security Management Act (FISMA) Compliance of platforms for low and moderate data types. The infrastructure service provides scientists with a secure security compliant hosting service for compute and storage requirements. It also provides for incident response (IR) at both policy and technical levels, where, IR is built into the design of Nebula such that it isolates components and users while providing access to the Inspector General without affecting system performance.

4. Nebula Case Examples

NASA initiated several pilot projects prior to Nebula's pre-release to determine if a dynamically scalable science-class computing platform capable of complying with Federal Agency security standards could be delivered as anticipated. The Nebula team worked with system administrators throughout the Agency to work in conjunction with scientists to test the platform and see how it could help advance the processing, storage and sharing of project data.

4.1. Microsoft's WorldWide Telescope

In July 2010, Nebula was used to demonstrate delivery of high-resolution imagery of Mars to Microsoft's WorldWide Telescope. NASA worked with Microsoft Research, under the Space Agreement Act, to bring imagery from the Agency's Mars and Moon missions to life, and make volumes of never-before-seen data accessible to the world. This project pushed the boundaries of technology to deliver a smooth, seamless experience for the public. Nebula processed, and presently hosts, more than 100 terabytes of high-resolution images, the equivalent of 20,000 DVDs worth of information, for use in the WorldWide Telescope. Each HiRISE image is a gigapixel in size – containing about 100 times more information than a 10 megapixel camera. Roughly 13,000 of these images were stitched together to form a single coherent map of the Martian surface, then combined with elevation data using a stereo-image reconstruction process to give the public a near 3-D view of the red planet.

Once the pictures are taken, the data is prepped for processing and enters the Nebula queue. After Nebula processes the new data, it is immediately integrated into the WorldWide Telescope data set and made available to the public through the WorldWide Telescope client. The process from when it snaps a photo, to being accessible on your computer screen is roughly two weeks.⁶

4.2. WISE Project Uses Nebula to Help Find Galaxies

WISE, or Wide-field Infrared Survey Explorer, 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. Data collected by WISE provides an important catalog for the James Webb Space Telescope as well as a lasting legacy for the astronomical community. Nebula helps the WISE team process more data than could ever be accomplished on local servers.

The team encountered a short-term need for a large number of small servers. Making the best use of WISE data requires low-resolution galaxies to be processed by many small virtual servers to resolve and measure their source sizes, and differentiate point-like emissions (such as from active nuclei or nuclear star formation) from distributed emissions (such as from merger or spiral galaxies). Additionally, WISE surveys several large regions of the sky for distributed star formations in our own galaxy. Using instances with enormous computational capacity, a few large regions can be processed to tremendous depth.

4.3. IRODS: Nebula Demonstrates Use for Distributed Data Storage and Management

The NASA Center for Climate Simulation (NCCS) at Goddard completed a pilot project to develop an iRODS-based Data Management System (DMS) to deal with the massive amounts of observational and model data used in NASA's climate and weather studies. iRODS, the Integrated Rule-Oriented Data Management System, is an open-source, data grid software solution that allows for the management, sharing, searching, and distribution of large and diverse scientific data sets. Through iRODS, Nebula demonstrates the ability to “federate,” a feature in which separate iRODS instances can be integrated,

thereby giving users and applications a consolidated view of separately maintained data.

The NCCS provides compute engines, analytics, data sharing, long-term storage, networking, and other high-end computing services designed to meet the specialized needs of the Earth science community. The DMS team used Nebula to host the DMS prototype with a goal of managing and publishing climate simulation data using iRODS within a distributed set of Nebula instances. A collection of MERRA (Modern Era Retrospective-Analysis for Research and Applications) data was placed under iRODS control. The MERRA data was stored on the file system and registered with iRODS. The registration process stored metadata about each of the MERRA files in the iRODS database. The entire catalog of monthly MERRA products resulted in the ingestion of 360 files that occupy 47 GB. The data was shared between two instances. By eliminating the need to explicitly switch an iRODS client between distinct grids, federation allowed perusal or download of data from multiple iRODS repositories through a single interface. Upon completion, users could examine, search for, and download simulation data from either Nebula instance through a single iRODS web interface.

4.4. SERVIR and SPoRT Use Nebula to Advance Short-Term Weather Forecasting

NASA's Marshall Space Flight Center uses Nebula for two short-term weather forecasting programs. The SERVIR program integrates satellite observations, ground-based data, and forecast models to monitor environmental changes and improve response to natural disasters. SERVIR enables scientists, educators, project managers and policy implementers to better respond to a range of issues including disaster management, agricultural development, biodiversity conservation and climate change. The SPoRT Center transitions unique NASA satellite observations and modeling capabilities to NOAA's National Weather Service with a goal of improving the analysis and prediction of weather events occurring within the 0-48 hour time frame. High-resolution satellite imagery from NASA's Terra and Aqua instruments can be used in near real-time for event analysis or combined with weather model forecasts by providing improved representation of the land, water, and three-dimensional atmosphere.

SPoRT has worked with the Weather Research and Forecasting model – Environmental Modeling System (WRF-EMS) developers to incorporate high-resolution sea surface temperatures from MODIS aboard Aqua and Terra, and high-resolution soil and moisture characteristics obtained from the NASA Land

Information System. Therefore, the WRF-EMS is an ideal conduit for the inclusion of NASA satellite and research products within an operationally relevant forecast and decision support system.

5. Further science capability and performance testing

In July 2011, NASA Science Mission Directorate (SMD) commission extensive and exclusive performance testing on a large portion of Nebula resources. This performance testing will ensure that Nebula can meet NASA mission requirements, and it will fully characterize Nebula for mission use. After performance testing period of approximately one month, SMD plans to utilize Nebula on a variety of their mission projects that best match Nebula performance characteristics. The following describes some the projects SMD proposed to move to Nebula based on SMD initial understanding of it's Nebula.

NEX, which provides remote user access to earth science modeling and resources and data, plans to utilize Nebula for elasticity. NASA's Climate in a Box Project, which is exploring the utility of "desktop" supercomputers in providing a complete, pre-packaged, ready-to-use toolkit of climate research products and on-demand access scale up and down on Nebula infrastructure, would utilize Nebula as an alternative delivery method to users and demonstrate the portability of it's applications. NASA's Climate in a Box Project, which is exploring the utility of "desktop" supercomputers in providing a complete, pre-packaged, ready-to-use toolkit of climate research products and on-demand access to an HPC system, would utilize Nebula for it's parallel processing requirements and portability. Again, these are just of few of the projects that SMD is planning to utilize Nebula for science processing and storage requirements.

6. Conclusion

In 2009, the Obama Administration announced the Federal Government's Cloud Computing Initiative to greatly reduce waste, increase data center efficiency and utilization rates, and lower operating costs, and support green technology.⁹ In the same month of Obama's mandate, the Nebula cloud container arrived at Ames Research Center and NASA entered a new era of information technology development. U.S. Chief Information Officer Vivek Kundra visited Ames the same month to share the Administration's vision on cloud computing for the federal government and cited Nebula as an example of the Government's ability to "leverage the most innovative technologies."⁸

From an outgrowth of a simple research exercise, Nebula developed into one of the space industry's first cloud computing environments. Nebula delivered to NASA more than what the Administration mandated, Nebula offers not only a cost-cutting, more efficient compute model, but a tool that enables scientists to process, share and store valuable data more effectively. During its short tenure, Nebula helped Agency scientists and administrators realize cloud computing's potential for processing data quicker, storing it more efficiently, and sharing it with colleagues and the public more effectively. The WorldWide Telescope project brought imagery from Mars to life so both scientists and the public could learn and appreciate our nearest planet. Nebula enabled the WISE program to better explore images in the sky using instances with enormous computational capacity to process large regions in tremendous depth. Additionally, Nebula's use on a series of Agency sponsored environmental/weather programs has enabled scientists to better understand issues impacting planet Earth.

[1] Michael R. Nelson. "Building an Open Cloud." Science. June 2009, vol. 34 no. 5935 pp. 1656 – 1657. <http://www.sciencemag.org/content/324/5935/1656.full>

[2] John Foley, "Private Clouds Take Shape." Informationweek.com. August 9, 2008.

[3] Gartner Research. "Gartner Identifies the Top 10 Strategic Technologies for 2011." October 19, 2010. <http://www.gartner.com/it/page.jsp?id=1454221>

[4] 2,600 servers x \$3,000 server x .15 utilization = \$1.15M (85% savings)

[5] Linda Cureton, "NASA Cloud Computing Services for Science," April 22, 2011.

[6] Information on the WorldWide Telescope can be found at <http://www.worldwidetelescope.org>

[7] Vivek Kundra, State of Public Sector Cloud Computing, May 20, 2010. <http://www.cio.gov/pages.cfm/page/State-of-Public-Sector-Cloud-Computing>

[8] NASA Announcement. NASA Ames Hosts White House CIO, http://www.nasa.gov/centers/ames/news/features/2009/cloud_computing.html