An Introduction to Cloud Computing and NASA Cloud Services from a Science User Perspective

Andrew L. Molthan
NASA Short-term Prediction Research and Transition (SPoRT) Center
Earth Science Office
Marshall Space Flight Center, Huntsville, AL

SPoRT Seminar Series at the National Space Science and Technology Center (NSSTC)
Thursday, August 25th, 2011

What is cloud computing?

- According to all-knowing Wikipedia:
  - "Cloud computing is the delivery of computing as a service rather than a product, whereby shared resources, software and information are provided to computers and other devices as a utility and over a network."
- But, what does this mean? Specifically,
  - "…the delivery of computing as a service…"
  - "…provided … as a utility and over a network …"

Overall Concept

- Rather than having individual machines hosted at the user's location, resources are aggregated in fewer points with large resource pools.
- Each pool is accessed through the internet.

Example:

The NASA Nebula Cloud Computing Platform currently hosts large volumes of IT hardware in two locations:

- NASA Ames Research Center
- NASA Goddard Space Flight Center

Overall Concept

- Cloud computing focuses on three primary **services** provided to the user community:
  - Software as a Service (SaaS)
    - Deliver software applications through cloud resources
  - Platform as a Service (PaaS)
    - Deliver web applications through cloud resources
  - Infrastructure as a Service (IaaS)
    - Deliver computing infrastructure through virtual platforms comprising various configurations and operating systems.
    - The remainder of this talk will discuss IaaS applications.

Infrastructure as a Service (IaaS)

- The definition of cloud computing included:
  - "the delivery of computing as a service"
  - "as a utility and over a network"
- With these concepts in mind, IaaS provides:
  - Service:
    - Configuration of cloud resources into individual machines driven by user requirements.
    - Rapid provisioning and decommissioning of virtual systems.
  - Utility:
    - Access to virtual machines through any internet connection.

Infrastructure as a Service (IaaS)

Your PC

FIREWALL

VPN

Request hardware configuration from Nebula
"I need 4 CPUs, 16 GB of RAM and 100 GB disk."

Nebula processes create the virtual machine with basic, selected OS.

Save the configuration for later as an "image".

User logs in to modify machine, install needed software, add storage, etc.

Machine executes software to perform desired tasks.

Decommission machine and return hardware to the Nebula pool.

Elastic Storage (IaaS)

- Virtual machines are often preconfigured with X CPUs, Y GB of RAM and Z of disk.
  - Often, additional disk is needed.
  - And, once terminated, content of the virtual machine is lost.
- Instead, use additional "elastic" storage.
  - User can request large volumes to "mount" to virtual machines.
  - Content on elastic storage is retained even after the virtual machine is terminated.

Infrastructure as a Service (IaaS)

VPN

Request hardware configuration from Nebula
"I need 4 CPUs, 16 GB of RAM and 100 GB disk."

Nebula processes create the virtual machine with basic, selected OS.

Save the configuration for later as an "image".

User logs in to modify machine, install needed software, add storage, etc.

Machine executes software to perform desired tasks.
Store output on elastic disk.

Decommission machine and return hardware to the Nebula pool.
Elastic storage remains.

Real World Example: WRF

Virtual Machine
4 CPUs, 8 GB RAM, 100 GB

- Supports WRF code execution, downloading of input data sets, and processing.

- Stores model output
- Large static input data sets
- Postprocessed output files
- Short-term forecast archive
- Configuration for run domains

After model execution:

- **When the virtual machine is terminated, all contents are *lost*.**
- **The user must create an *image* to save the contents of a virtual machine.**

- **Therefore, the compromise is…**
  - **Use the resources of the virtual machine for as little time as possible.**
  - **Store output on elastic storage.**

Infrastructure as a Service (IaaS)

- Summary of the concept:
  - The "cloud" is a large resource of hardware that can be configured into large numbers of individual machines.
  - Users can customize those machines for their purpose: operating system, software, etc.
  - The final machine(s) are used for the time period they are needed, then saved as an image for later use, or terminated.

Strengths of IaaS

- Availability of a hardware pool means no procurement delays.
- User-configurable machines with root access for software installation.
- Provides an "instant on" sandbox to test requirements.
- Ideal solution when additional compute capacity is needed for a short time.

Limitations of IaaS

- Care must be taken to retain your results.
- Once terminated, all material on the virtual machine is lost, unless imaged.
  - Solution: use external "elastic" or "object" storage
- Hardware configuration for multi-processor applications is less ideal than a cluster.
- Communication between multiple virtual machines may be complex.
- VPN is an extra hassle but is being "worked".

Summary of "…as a service"

- Computing hardware is provided "as a service" because it is rapidly available with minimal programmatic overhead.
  - No procurement delay
  - Create and terminate instances as needed, when needed, as long as resources are available
- Next: Cost and the "…as a utility" concept.

Utility and Cost

- One might ask:
  - "If I can create a virtual machine out of any hardware available, why would I ever destroy it?"
- Market concepts enforce constraints on the user by assigning rates for the use of cloud resources.
  - For example, the Amazon cloud has a complex rate structure for different use scenarios.

Amazon Costs

	Linux/UNIX Usage
Standard On-Demand Instances	
Small (Default)	$0.085 per hour
Large	$0.34 per hour
Extra Large	$0.68 per hour
Micro On-Demand Instances	
Micro	$0.02 per hour
Hi-Memory On-Demand Instances	
Extra Large	$0.50 per hour
Double Extra Large	$1.00 per hour
Quadruple Extra Large	$2.00 per hour
Hi-CPU On-Demand Instances	
Medium	$0.17 per hour
Extra Large	$0.68 per hour
Cluster Compute Instances	
Quadruple Extra Large	$1.60 per hour
Cluster GPU Instances	
Quadruple Extra Large	$2.10 per hour

Data Transfer IN	
All data transfer in	$0.000 per GB
Data Transfer OUT	
First 1 GB / month	$0.000 per GB
Up to 10 TB / month	$0.120 per GB
Next 40 TB / month	$0.090 per GB
Next 100 TB / month	$0.070 per GB
Next 350 TB / month	$0.050 per GB
Next 524 TB / month	Contact Us
Next 4 PB / month	Contact Us
Greater than 5 PB / month	Contact Us

In the "cloud computing" framework, users are encouraged to be efficient and avoid hoarding of resources through metered pricing rates.

Costs not shown: elastic storage, monitoring, web server load balancing, public IPs, and others.

Some examples for Linux use cases, as of 8/23/11

Nebula Costs

- Costs on Nebula are also charged based upon usage but with less micromanagement than Amazon.
- $0.12/CPU-hour
  - Multi-CPUs cost per CPU, e.g. 4 CPU is $0.48/hr
- $0.15/GB-month
  - "Object storage" provides triplicate backup at a rate of $0.45/GB-month
- Billing is "taken care of via an allocation, from which usage costs are decremented."

Utility and Cost

- Given that use of cloud resources are **metered**, cost effectiveness is a primary concern.
- Considerations:
  - Procurement delays in obtaining resources
  - Length of time the hardware is needed
  - Overall cost of cloud computing versus permanent purchase and maintenance of a new IT resource.

Real World Example: WRF

- An example of costs for an application such as running the WRF model on an instance.
- Amazon:
  - CPU: $0.68/hr * 3 hr = $2.04 per forecast
  - I/O: $0.12/GB * 10 = $1.20 per forecast
  - Many additional charges are not reflected here.
  - Total per forecast is $3.24, probably closer to $4.
- Nebula:
  - CPU: $0.96/hr * 3 hr = $2.88
  - I/O: N/A
  - Nebula does not "nickel and dime" to the degree of Amazon.
- Average per run: $3.
  - $12 per day (four cycles)
  - $360 per month
  - $4320 per year

Caveat: It is ***extremely difficult*** to estimate true costs without actually attempting the project on a small scale.

Cost Summary

- Nebula and corporate clouds (e.g. Amazon) provide services with a charge structure similar to utilities.
  - The meter runs and you pay for what you use.
  - Increased efficiency reduces cost.
- Nebula provides a far simpler charge structure than Amazon.
  - Differences in cost are difficult to ascertain without executing a small project.

Applications!

- After a lengthy introduction, on to the fun part… applications.
- Several concepts are being developed within SPoRT / SERVIR along with an SBIR proposal.
  - Establishing WRF forecast domains to support SPoRT projects and real-time SERVIR work
  - Hydrologic modeling in E. Africa for SERVIR
  - SATCAST algorithm for convective initiation
  - Future data dissemination for NPP

WRF on Nebula

- Goals:
  - Mimic operational configurations at WFOs to support local case study applications (SPoRT)
  - Develop a Mesoamerica forecast domain to support local weather predictions (SERVIR)
- Why Nebula?
  - Avoids maintenance requirements and cost for geographically distant infrastructure (SERVIR)
  - Add real-time modeling capabilities without disrupting other activities (SPoRT)

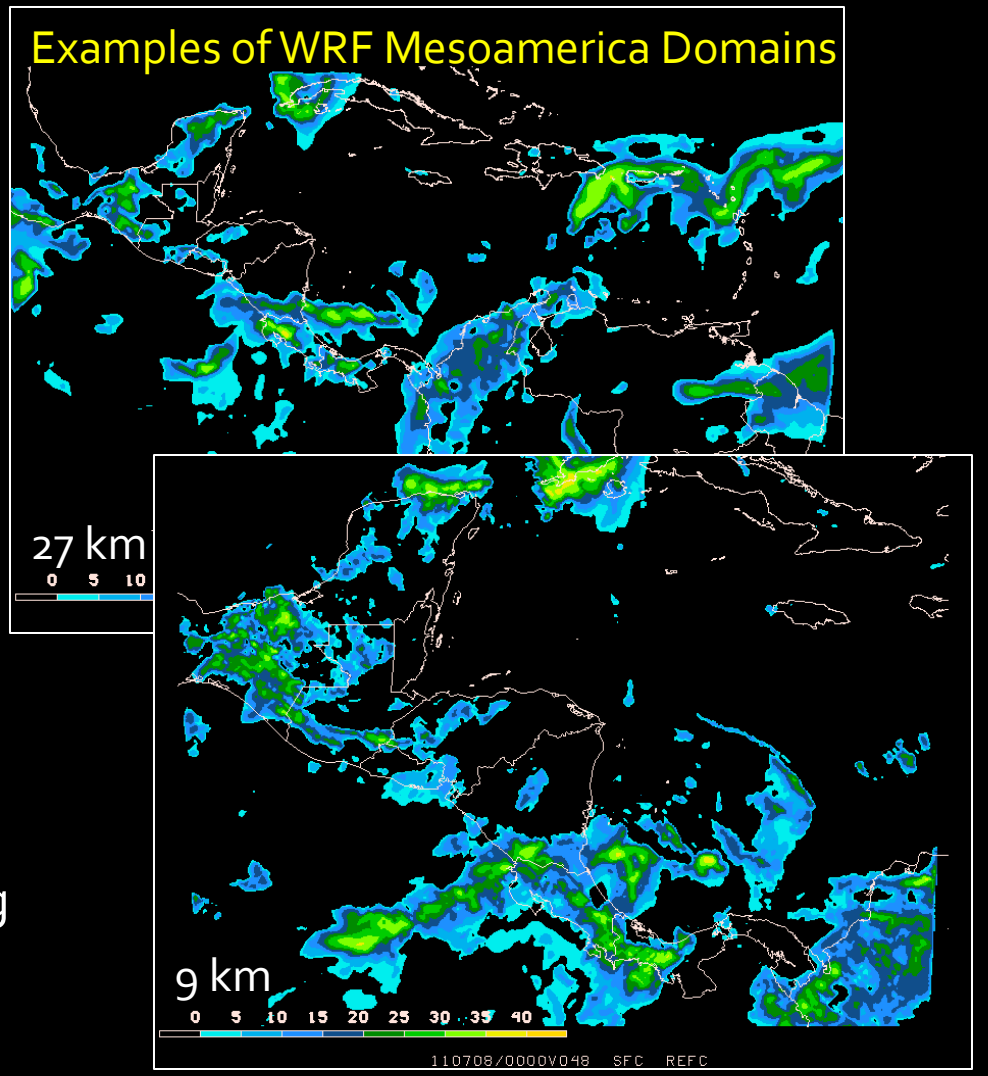

Hydrologic Modeling

- SERVIR has established a domain of the CREST hydrologic model for Kenya.
- Why Nebula?
  - Similar reason – alleviates maintenance of distant hardware.
  - Allows rapid development without delay in procurement.

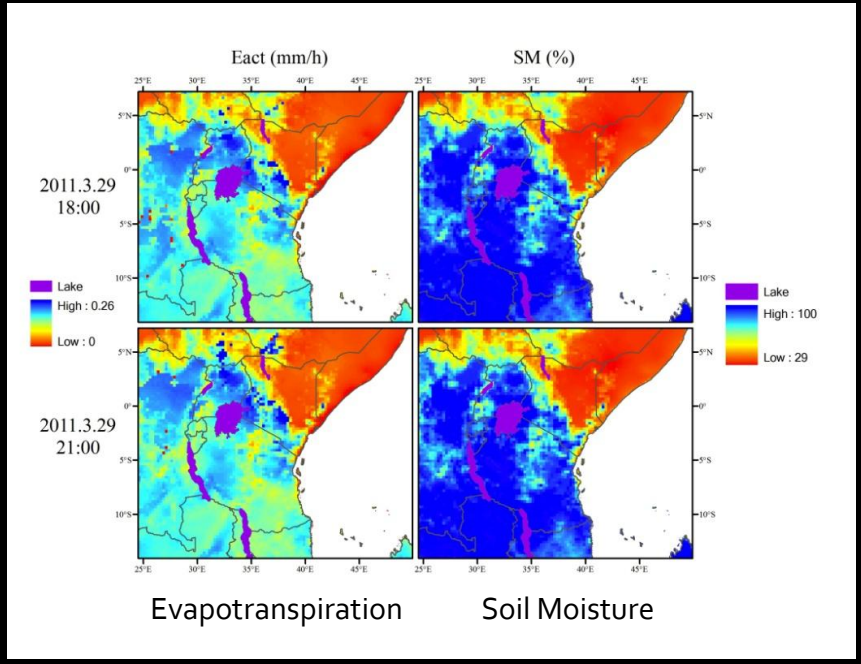

An example of outputs from the CREST model running over eastern Africa.

Convective Initiation

- SERVIR is applying the UAH "SATCAST" algorithm over Central America.
  - Improve lead time on convective storms and heavy precipitation.
- Nebula benefits:
  - Flexibility in virtual machine hardware availability.
  - Avoids maintenance of infrastructure in distant locations.

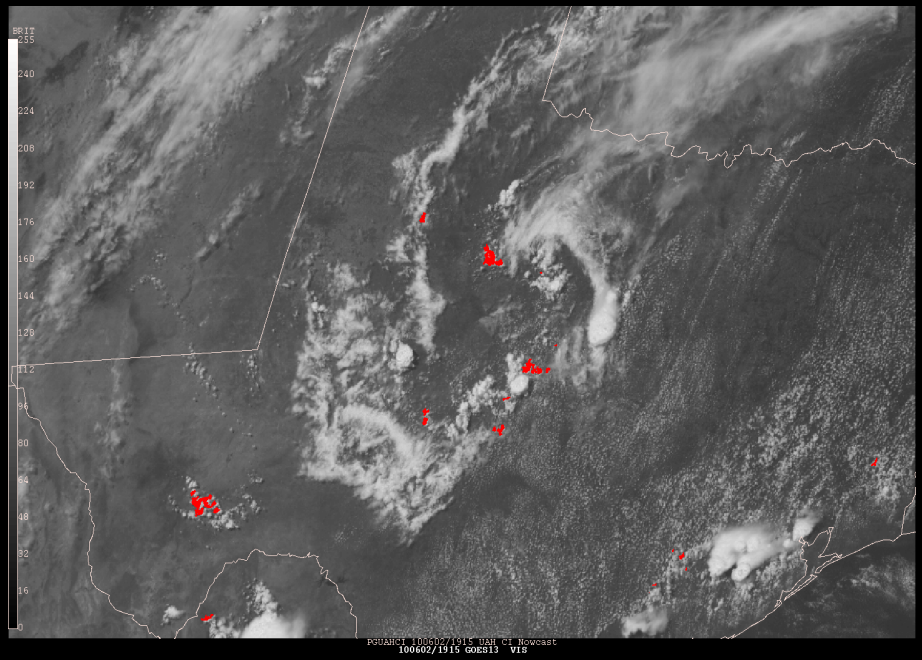

An example of convective initiation detections in Texas (red) based upon the UAH-SATCAST algorithm.

Satellite Image Processing

- SPoRT used Nebula to process (tile) high resolution ASTER imagery following the April 27, 2011 tornadoes.
- Nebula benefits:
  - Rapidly deployable virtual machine "sandbox" to test different tiling tools.
  - May serve as guidance for developing a future, local system.

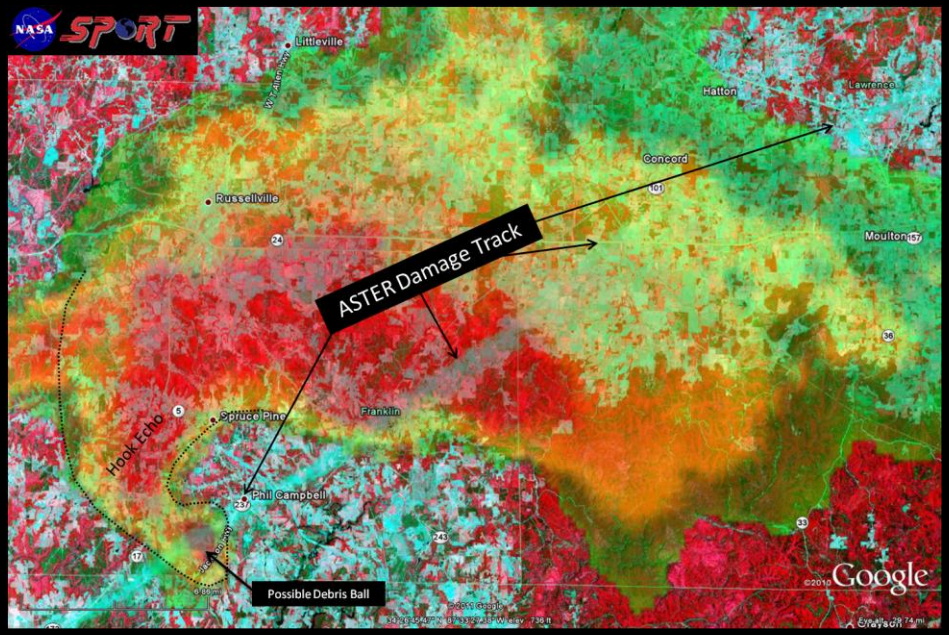

An example of the ASTER RGB image for the Phil Campbell tornado track. This high resolution Google Earth/KML formatted image was tiled using Nebula resources.

SBIR: NPP Data Processing

(Small Business Innovation Research)

- Global Science and Technology (GST), Inc. is working with SPoRT on an SBIR proposal using cloud computing.
- Nebula benefit:
  - Use cloud resources to spin up large IT needs focusing only on periods when polar orbiting data dissemination is required.
  - Otherwise, return resources to the cloud during low activity periods.

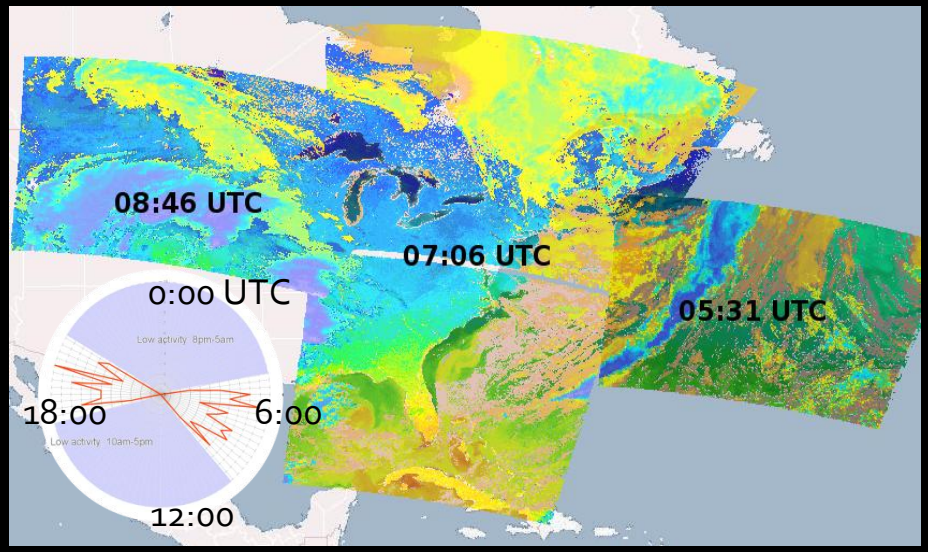

An example of polar orbiting satellite data (MODIS) and related CONUS processing and dissemination times.

Current Status

- The Nebula system is currently "down" for stress testing by SMD.
  - Part of TRL review
- Some applications have moved to Amazon during the downtime.
- VP61 staff will be involved in helping SMD with testing Nebula capabilities.

- Future goals:
  - Establish WRF system on Nebula to support SERVIR activities in Mesoamerica.
    - Weather forecasts
    - Air quality forecasts (…?)
    - SATCAST capabilities
  - Develop instances to support SPoRT needs
    - Occasional WRF runs
    - Satellite image processing

Cloud Resources at MSFC

- Under the CIO at MSFC, cloud resources are being developed.
  - MSFC provides access to Amazon resources along with security plan assistance and FISMA compliance.
- MSFC is developing a private virtual cloud resource – currently slated for operations in January 2012.
  - Currently unclear how the structure and interface will relate to Nebula or Amazon resources.

Summary

- Cloud computing concepts develop large computing resources, then charge users on an as-needed basis.
  - Reduce system procurement time
  - Reduce some maintenance overhead costs
  - Reduce energy costs by only paying when needed
- Although it is a different paradigm for the management of IT resources, there are potential applications and benefits for science users.
  - Thus far, have only used the IaaS concept to use virtualized hardware to support activities.

Questions?