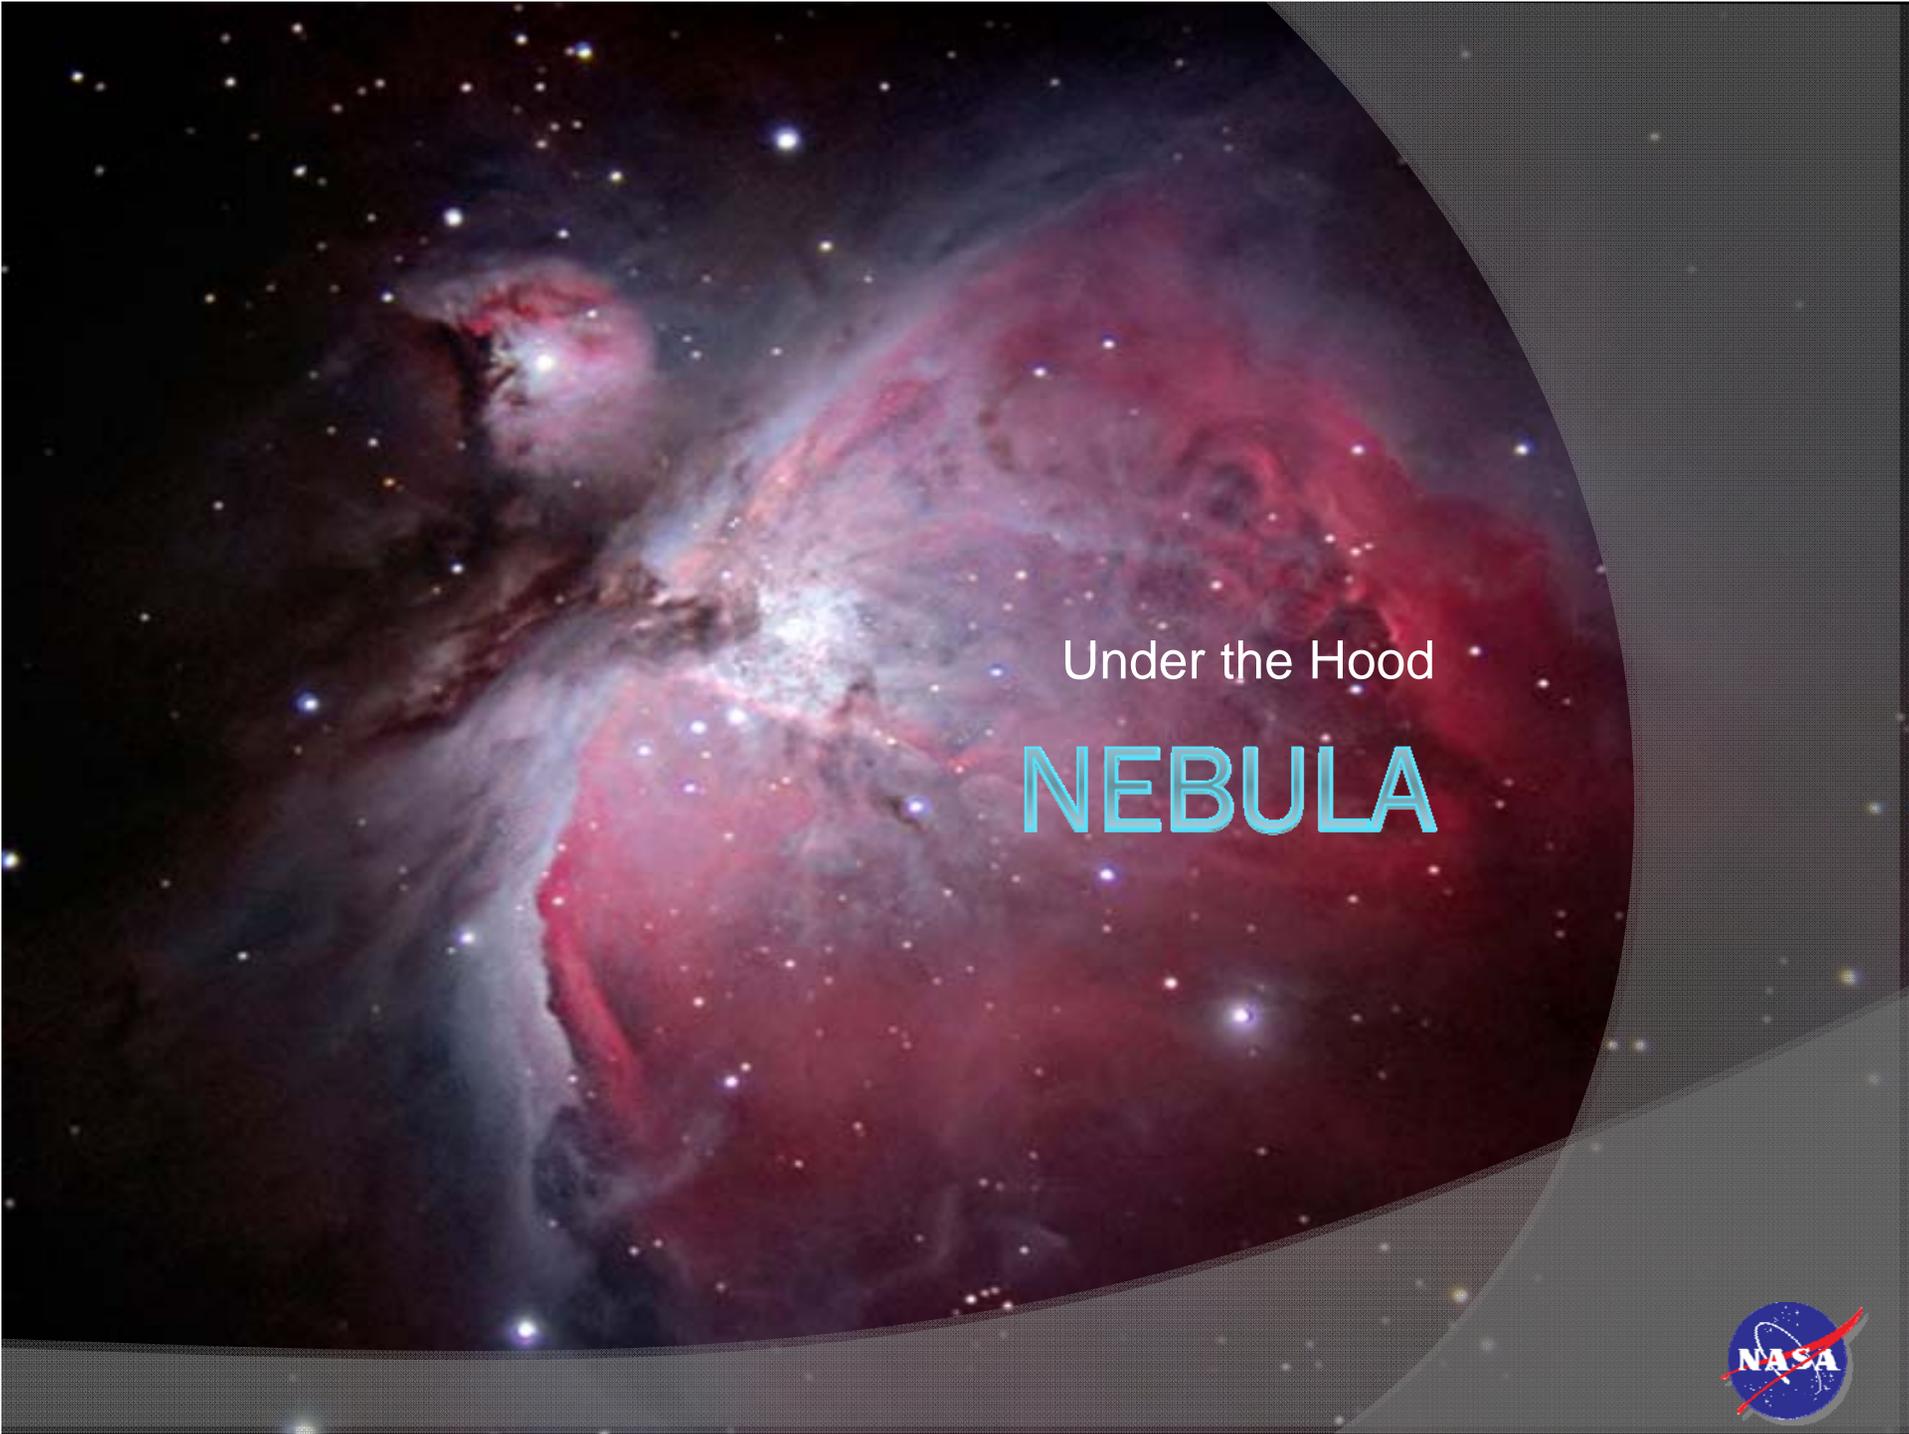


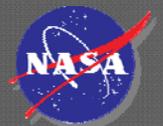
“A cloud does not know why it moves in just such a direction and at such a speed...It feels an impulsion...this is the place to go now.

But the sky knows the reasons and the patterns behind all clouds, and you will know, too, when you lift yourself high enough to see beyond horizons.”





Under the Hood
NEBULA



Overview

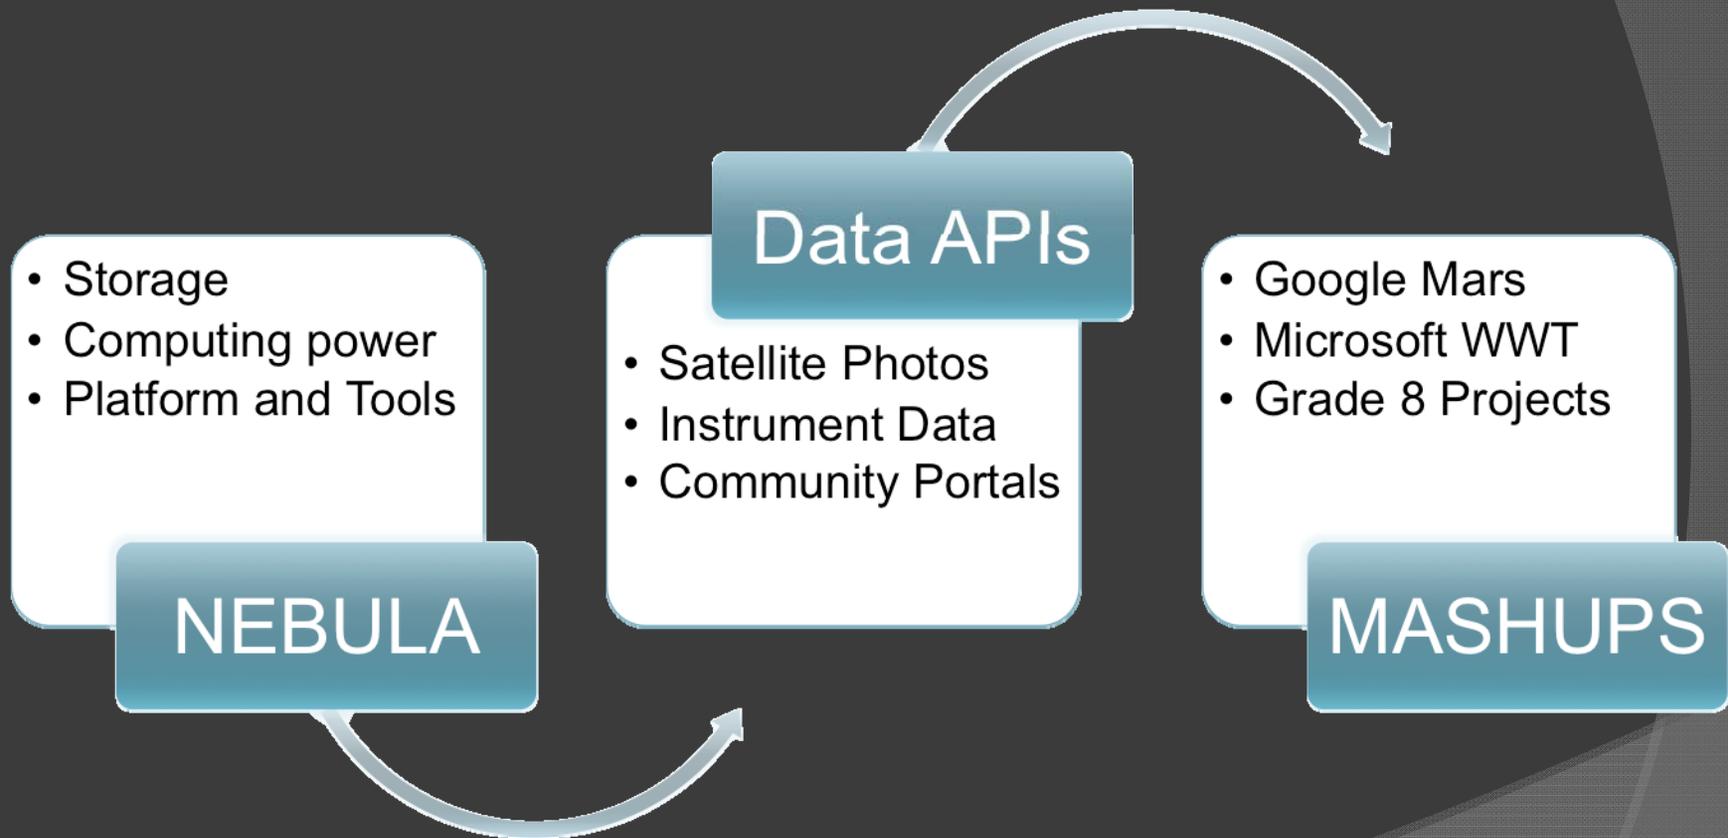


Vision: 'Zettabytes of structured data...'

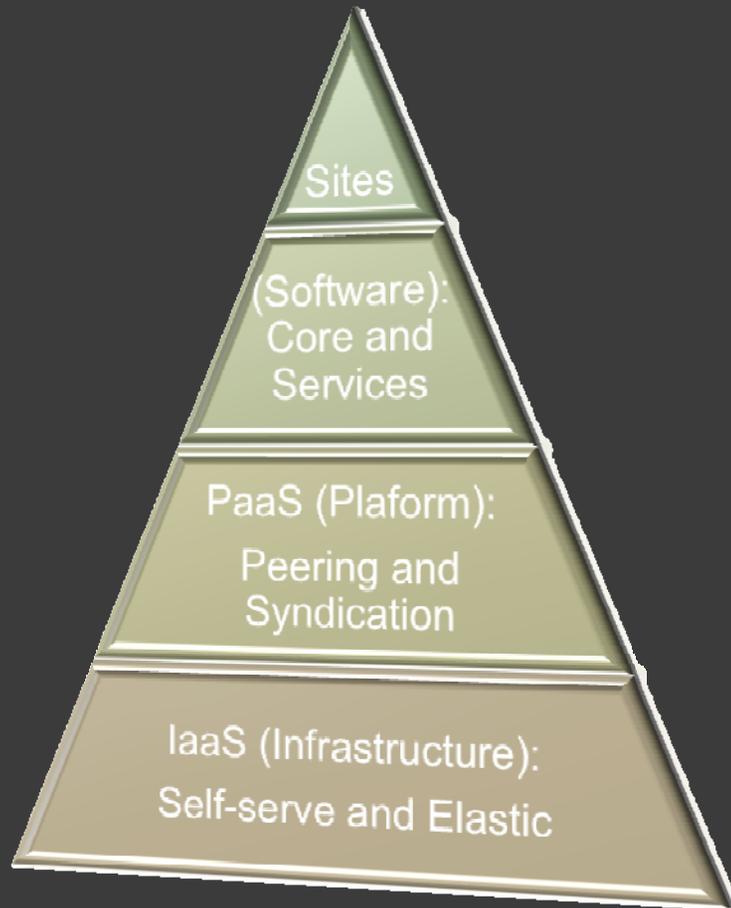
- Put all of NASA's public data on the Web
- Power space "Mashups" by everyone: from Google, to Grade-schoolers
- A flexible, reusable framework



Vision: Space 2.0, Powered by NASA



Full-Stack Cloud Platform



- ⦿ Integrated Strategy
 - Policy and Workflow
 - Governance
 - Community
- ⦿ Built on Open Source
- ⦿ Modular Design



laaS



Containerized Hardware

- Modular
- Mobile
- Highly efficient
- Rapid deployment



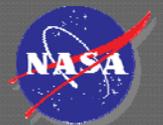
Self-Managed Virtualization

- ◎ XEN
 - Open Source, Powerful and Ubiquitous
- ◎ Eucalyptus
 - API-compatible with Amazon AWS
 - Basis of the Ubuntu Enterprise Cloud



STORAGE

- Lustre – HPC Cluster filesystem
- High performance (> 1000MB/sec)
- POSIX semantics
- 10GigE throughout the cloud
- Commodity hardware



NETWORKS

- ⦿ Settlement-free peering
- ⦿ NLR, I2, CENIC connections
- ⦿ Multiple CDN providers
- ⦿ 802.1Q-based logical separation
- ⦿ DNS and mDNS
- ⦿ 10Gb Ethernet – Commodity Equip.



Scalable Operations Mgmt

- Puppet and PXE-Kickstart
- Multicast DNS using Avahi (ZeroConf)
- Elasticity Mangement (Scalr)

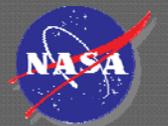


Highly Available

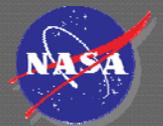
- ⦿ Master Nodes use:
 - DRBD (Network Replication) of RAID 1+0
 - Linux-HA Heartbeat with IPMI STONITH
- ⦿ Redundant networks using 802.1w



PaaS

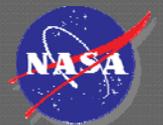


- ◎ *“...Dynamic languages are common because they are much easier to use than Java or .Net on the background, providing as much as 10x more developer productivity.”*



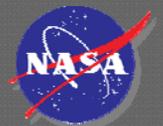
Software Framework

- ⦿ Django (<http://www.djangoproject.com>)
- ⦿ Python is for Science
- ⦿ “Django is a high-level Python Web framework that encourages rapid development and clean, pragmatic design.”



Database

- MySQL
- KVS (Tokyo Cabinet)



Message Queues FTW

- Open Standard – AMQP
- Best-of-Breed Implementation – RabbitMQ
- Erlang – Massively Concurrent
- ~ 400,000 Messages/sec – per server



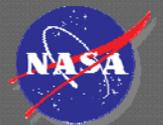
Central Search

- Built on SOLR+Lucene
- One Index to Rule them ALL
- Flexible, common taxonomy
- Master w/ Read-only Slaves



Single-Sign-On

- Federated LDAP (AD)
- eAuthentication – its not just cool, it's the law
- Simple JSON API – seamless to all apps



Analytics

- Custom Web Analytics Application
- Real-time Server monitoring, balancing
- Munin for Hardware, Ganglia for VMs
- SNMP and NetFlow for Traffic



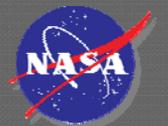
APIs

- Logically consistent RESTian URLs
- OAuth and OpenID
- hCard and other Microformats
- RSS / Atom, and PubSub
- Twitter, Facebook, Google Sitemaps, Pingomatic, flickr, s3, youtube...



CACHEING / CDN

- Varnish
- Memcache



Continuous Integration

- Subversion (future: GitHub.FI)
- Trac – Integrated ticketing and Wiki
- Selenium, Bitten – Automated Testing
- Fabric - Automated Deployment



Unified, logical URIs

- APPS.NASA.GOV – Public sites (think Google AppGallery)
- NEBULA.NASA.GOV – The Platform (think Google AppEngine or AWS.AMAZON.COM)
- CODE.(ARC).NASA.GOV –The Source and IDE



Private Cloud, Public Data



Status: Phase 1 Complete

- ✓ Requirements gathered
- ✓ Platform selected
- ✓ System Design baselined
- ✓ Dedicated hardware live
- ✓ Pilot Projects in Development

Phase 1

Phase 2

Phase 3

