

Chris C. Kemp

**CLOUD COMPUTING
FOR A
UNIVERSE OF DATA**

We have Challenges

- Many, many, many different platforms
- Very high operating cost
- Long provisioning times
- Many, many, many websites
- Very large attack surface
- Confusing to outside users
- Lock-in, portability, and interoperability issues



What IS Cloud Computing?

[Cloud computing security: Who knew?](#)

Computerworld - 11 hours ago

By Scott Bradner April 27, 2009 (Network World)
less than perfect agreement on just what



[DMTF Group to Create Cloud Computing Spec](#)

eWeek - 13 hours ago

The Distributed Management Task Force is creating a group to develop management standards for **cloud computing**. ...

[Open Standards For The Cloud](#) SYS-CON Media (press release)
DMTF creates open cloud incubator ITworld.com

[SAS to build USD 70 million cloud computing facility](#)

Al-Bawaba - Apr 26, 2009

SAS, the leading provider of business intelligence software, will build a USD 70 million cloud computing facility in North Carolina. Demand spurs \$70 million



[Hosting.com Partners with VMware to Deliver Cloud Trial C](#)

HostSearch.com - Apr 26, 2009

As part of the VMware vCloud initiative, developers, IT professionals, SIs and MSPs are invited to experience **Cloud Computing** firsthand with 30 days ...

[Operating system for building the internal cloud](#) Help Net Security
EMC integrates with VMware vSphere 4 for next generation AME 1-6

[Fourteen universities to study cloud computing with \\$5M NSF grant](#)

Geek.com - 10 hours ago

The program works with IBM and Google to study how to look at the infrastructure requirements to support cloud computing

[Fujitsu Launches Cloud Services](#)

Japan Corporate News (press release) - Apr 26, 2009

Fujitsu is seeking to address this issue by providing its Trustee Cloud Services. The service has the following features. ...

[Heading Back to Cloud Computing](#)

SYS-CON Media (press release) - 2 hours ago

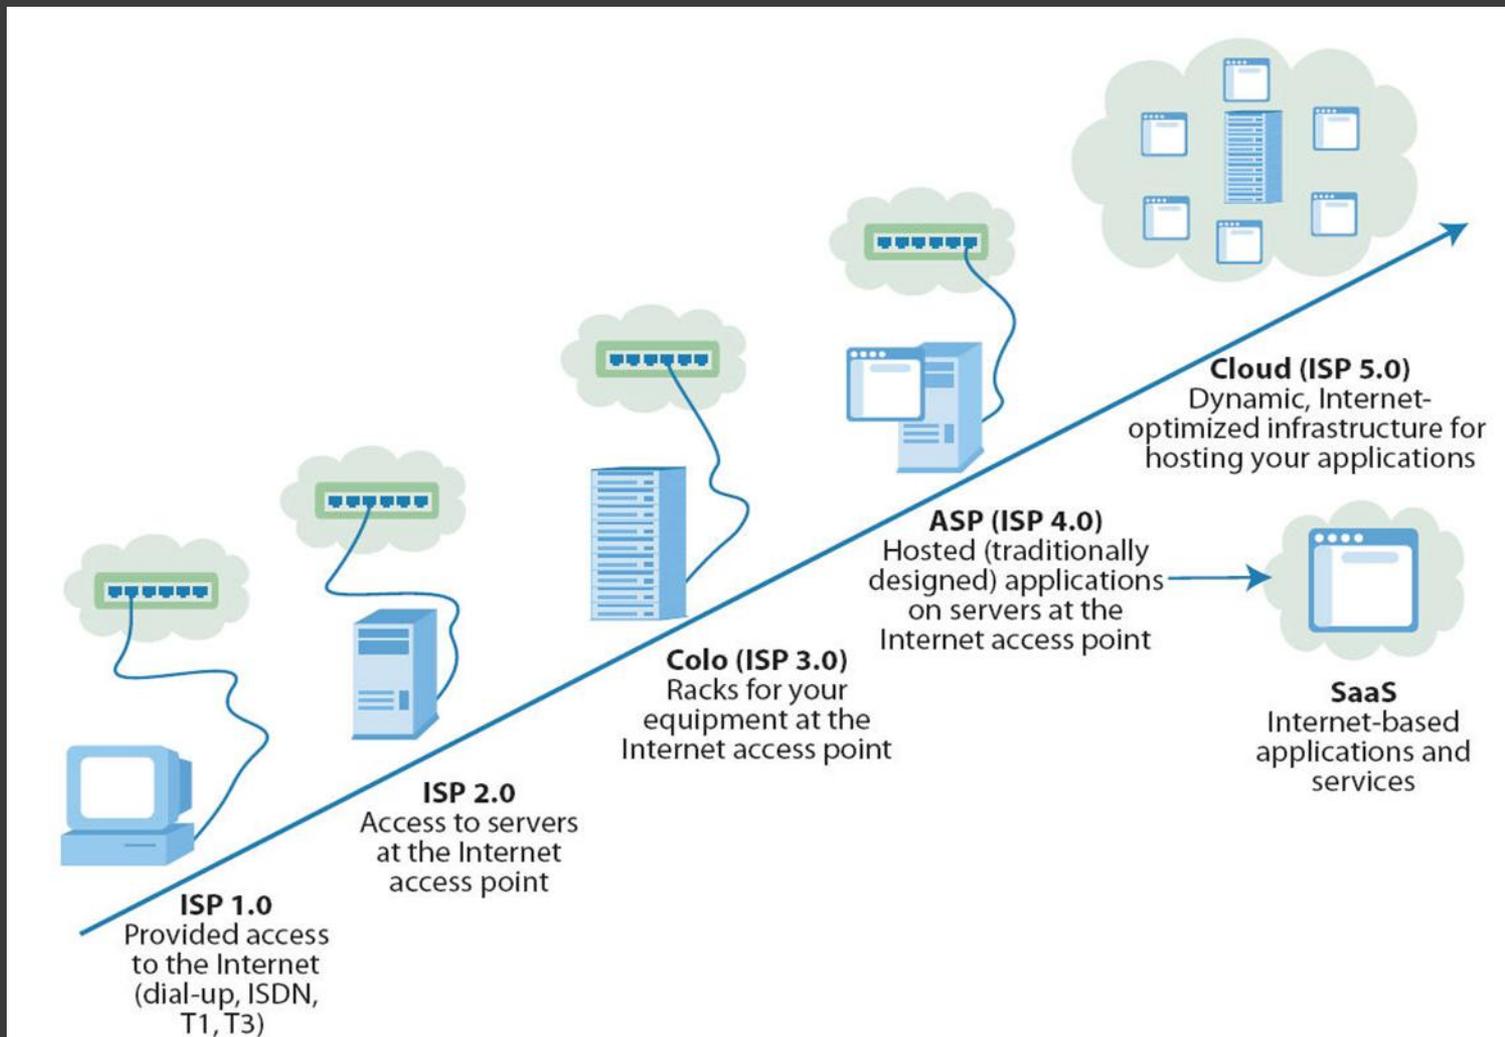
The merger does make Sun very relevant again, as its hardware is still a major player in the coming **Cloud Computing** deployment battle. ...

IDC on Cloud Computing:

“This is about the IT industry’s new model for the next 20 years.”

Vernon Turner, head of enterprise infrastructure, consumer and telecoms research.

Cloud as “Evolution of Hosting”



Source: Forrester Research

Cloud as “Bow-Wave of Change”

Virtualization
Social Networks
Mashups, DIY Tools
Exponential Change
Tech-savvy customers
Tech-savvy employees
Consumerization / Web 2.0
Cloud Computing
Public Infrastructure
Global Collaboration
Cultural Change
Security
Mobility
Green



- Employees
- Missions / Data Products
- Business Operations
- Partners / Collaborators
- The Public

“...reliable services delivered through data centers ... accessible anywhere that has access to networking infrastructure.

The Cloud appears as a single point of access for all the computing needs...”

Architecture of the Cloud

- ⦿ “The canonical cloud architecture that has evolved revolves around dynamically scalable CPUs consuming asynchronous, persistently queued events.”
- ⦿ <http://highscalability.com/canonical-cloud-architecture>

Cloud computing:

- Processes-as-a-service
- Software-as-a-service
- Framework-as-a-service
- Storage-as-a-service
- Platform-as-a-service
- Infrastructure-as-a-service

- Everything-as-a-service

The Cloud - Under the Hood

Process-as-a-service

APIs

Apps

Partners

Software-as-a-service

Auth

Data

Queues

Platform-as-a-service

Virtualized
Environments

Scalable
Storage

Rich Network
Connectivity

NIST Definition

- Cloud computing is 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.

(Excerpts from Gartner's "Five Attributes of Cloud Computing")

Cloud has Aspects, not Form

Service-based

- ⦿ Abstracted from the implementation
- ⦿ Completely automated
- ⦿ Near real-time delivery (seconds or minutes)

Scalable and Elastic

- ⦿ Resources are drawn from a common pool
- ⦿ Dynamically allocated to meet demand
- ⦿ Dynamically released when appropriate
- ⦿ Fully automated

Shared

- ① Common resources build economies of scale
- ① Common infrastructure runs at high efficiency

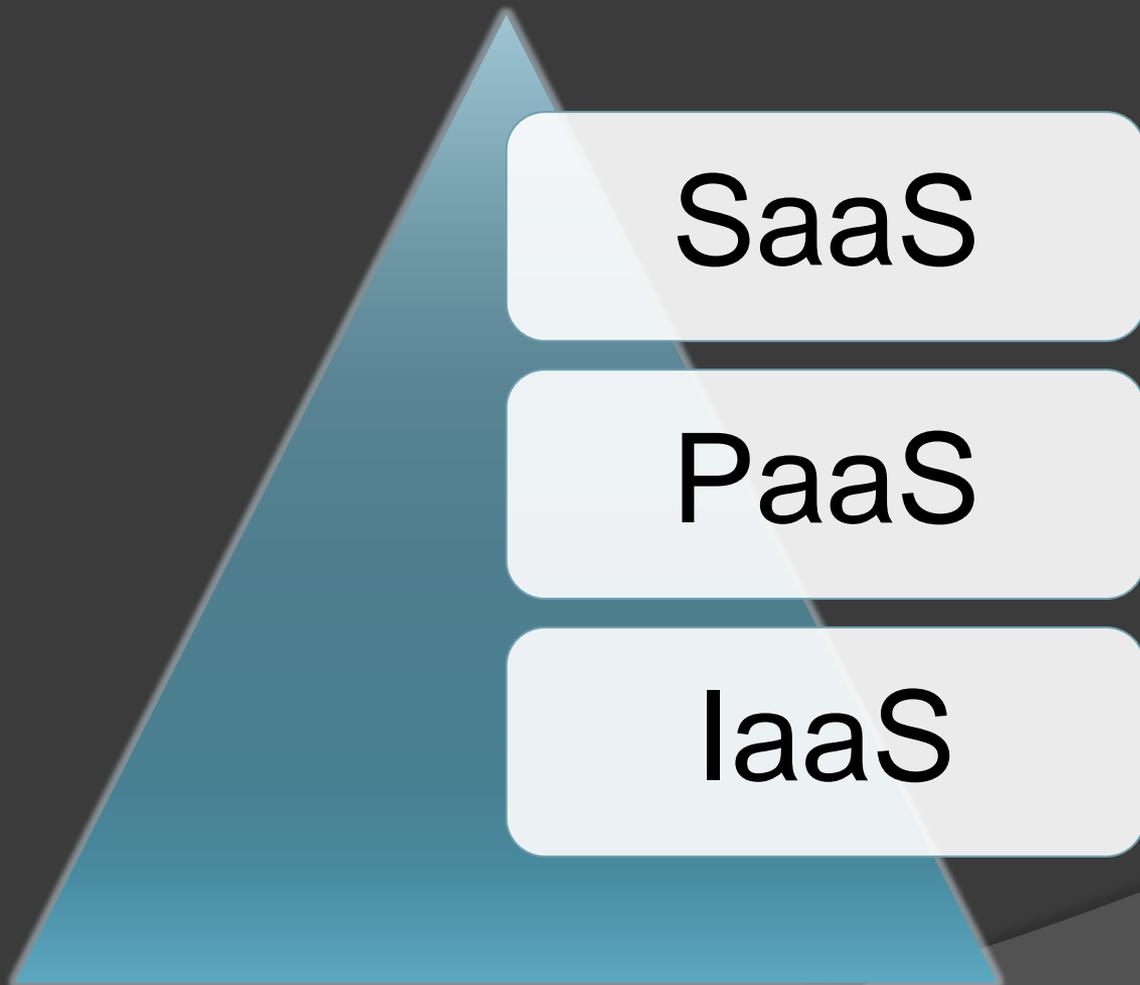
Metered by Use

- Consumers pay for services used
- Underlying hardware costs are irrelevant

Uses Internet Technologies

- ⦿ Open standards and APIs
- ⦿ Almost always IP, HTTP, and REST

Cloud Components



SaaS

PaaS

IaaS

Infrastructure (as-a-Service)

- ◎ Similar to:
 - Utility Computing or Grid Computing
 - Old-school “Time-sharing” on Mainframes
- ◎ Often uses:
 - Virtualization
 - Shared storage (SAN or Cluster)
- ◎ Target User: System Administrator

Infrastructure (as-a-Service!)

- ⦿ Doesn't matter what it uses
- ⦿ Clouds are service-based, e.g. abstracted
- ⦿ Implementation is hidden, changeable

Common IaaS

- Amazon EC2, S3 and EBS
- Linode
- Rackspace Cloud
- Terramark Enterprise Cloud*
- IBM CloudBurst

Common Use Cases

- Hardware Consolidation
- Short-term IT needs
- Development infrastructure
- Network storage
- Supplemental peak capacity

Platform (as-a-Service)

- ⦿ Interacts at the Source Code level
- ⦿ Total abstraction from infrastructure
- ⦿ “Infinitely” scalable
- ⦿ Target User: Software Developer

Common PaaS

- Google AppEngine
- Salesforce.com
- Microsoft Azure

PaaS Pros and Cons

⦿ Advantages:

- Much easier to use than IaaS
- Applications can be more secure
- Achieves higher efficiencies than IaaS alone

⦿ Disadvantages:

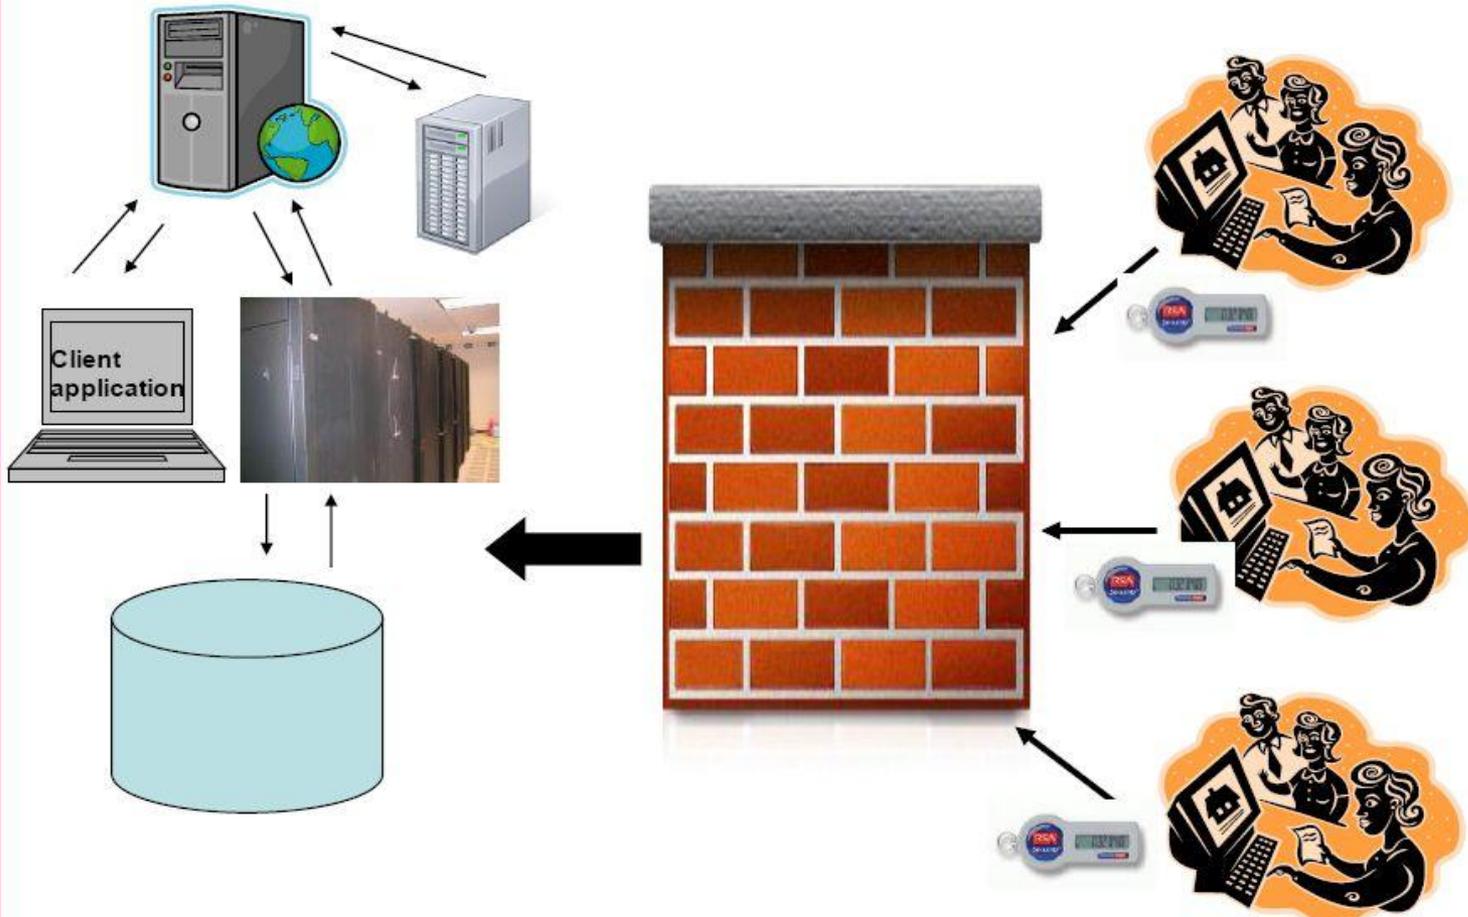
- Usually a language-specific platform
- Limited set of services (DB type, Queue, etc)

Software (as-a-Service)

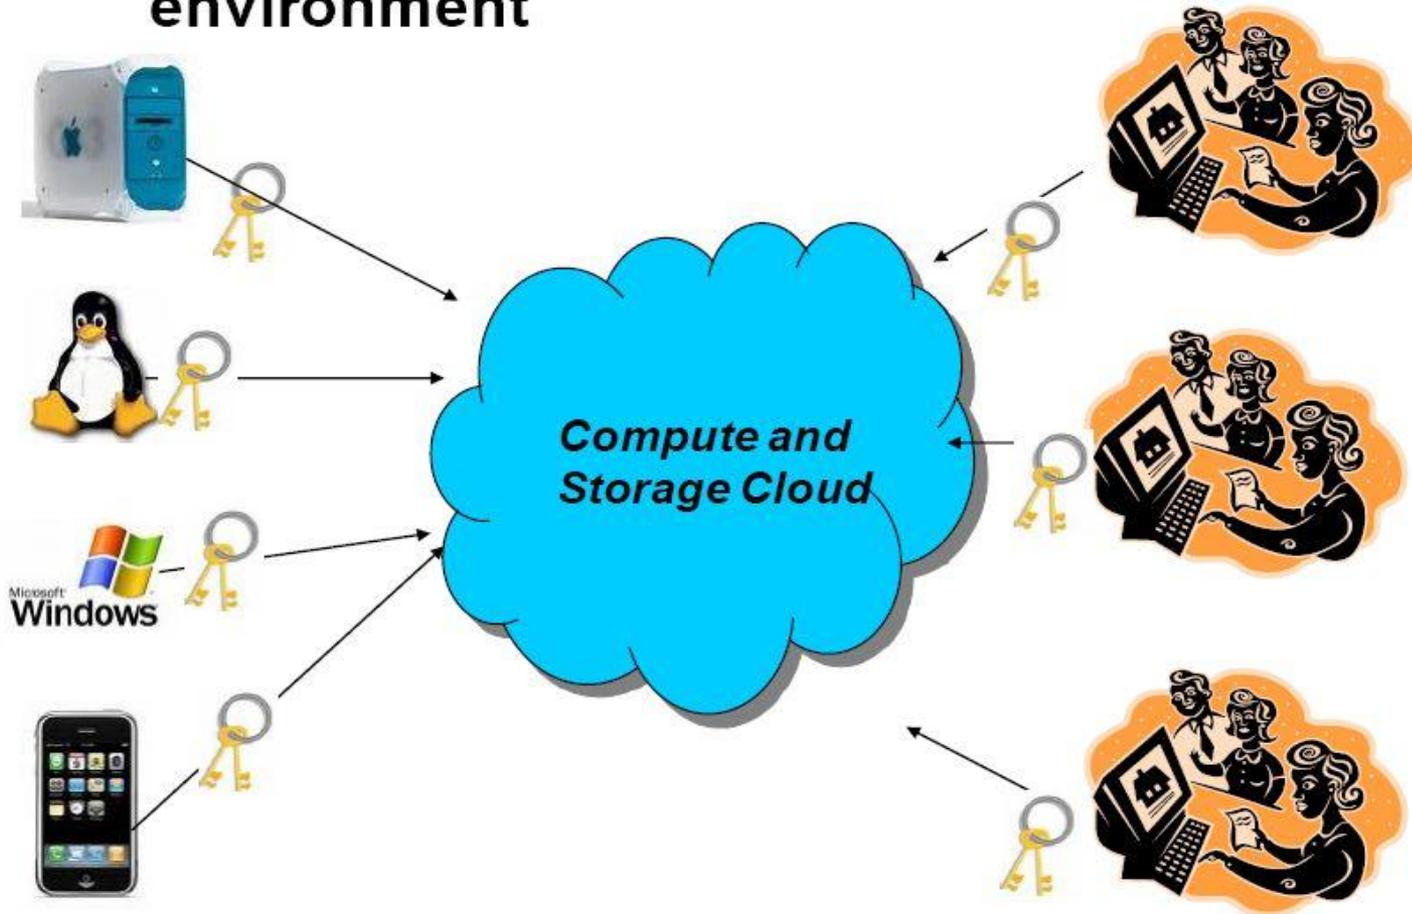
- ⦿ Can include almost *any* IT service:
 - Email
 - Web Hosting (Google Sites)
 - Blogs, Wikis, Forums, etc.
 - Source code control
 - Telephony
 - Office and Productivity Software
- ⦿ Target User: End-User

Cloud Security

To Collaborate in this environment today



Future collaboration in this environment



When is it coming?

Clouds on the Horizon

Gartner on Cloud:

- "By 2011, early technology adopters will forgo capital expenditures and instead purchase 40 percent of their IT infrastructure as a service. Increased high-speed bandwidth makes it practical to locate infrastructure at other sites and still receive the same response times. ...as service-oriented architecture (SOA) becomes common, 'cloud computing' will take off, thus untying applications from specific infrastructure."

Industry Quotes

- Google's CEO Eric Schmidt has called such computing "transformative. It is the new model of computer architectures."
- IBM's Vayghan adds: "This is not something coming 20 years from now. There already are many products."

Government Clouds

Fiscal 2010 budget:

- “The Federal technology environment requires a fundamental re-examination of investments in technology infrastructure... Pilot projects will be implemented to offer an opportunity to utilize more fully and broadly departmental and Agency architectures to identify enterprise-wide common services and solutions, with a new emphasis on cloud-computing.”

Gartner on Hybrid Clouds:

- Clients should prepare for myriad hybrid models, ranging between pure private cloud implementations and pure public cloud offerings. Limit the contract duration with external cloud providers to as short a term as possible (preferably one year), because more choices will be available at the end of the current contract cycle.

How can NASA benefit?

- Case study: Nebula

NASA has Challenges

- Missions are focused on the Mission
- Scientists are focused on the Science
- Large-scale infrastructure requirements
- Too much is spent on infrastructure
- Missions Completely Fail (OCO)
- Missions Completely Succeed (Rovers)
- Politics impact Missions (Triana)



Kepler Mission: A search for habitable planets

- Kepler Home
- Overview
- News & Schedule

[Kepler Home](#) > [News & Schedule](#)

Mission News and Schedule

[News Releases](#) --||-- [Schedule](#) --||-- [Milestones/Mission](#)



NASA National Aeronautics and Space Administration

AMES RESEARCH CENTER

NASA ASTROBIOLOGY INSTITUTE

Early Habitable Environments
Evolution of Complexity
Principal Investigator: David J. Des Marais

AMES TEAM NEWS AND HIGHLIGHTS

NASA SELECTS "AMES RESEARCH CENTER TEAM" FOR ASTROBIOLOGY INSTITUTE

NASA awarded five-year grants, averaging seven million dollars each, to 10 research teams from across the country to study the origins, evolution, distribution, and future of life in the universe. According to Carl Pilcher, director of the NASA Astrobiology Institute (NAI), "The research of the new teams reflects the increasing maturity of astrobiology. They are focused on fundamental questions of life in the universe, but their work has implications for..."



DASHlink

Discovery in Aeronautics Systems Health

DASHlink is a virtual laboratory for scientists and engineers to disseminate results and collaborate on research problems in health management.



HOME

[Log In To MyNASA](#) | [Sign Up](#)

[For Public](#) | [For Educators](#) | [For Students](#) | [For Media](#)



NASA K12

Home

Atlantis Heads for Wednesday Rendezvous




NSPIRES

NSPIRES Home

NASA Research Opportunities

Supporting research in science and technology research announcements in a wide range of proposals submitted in response to the research proposals and conducting award

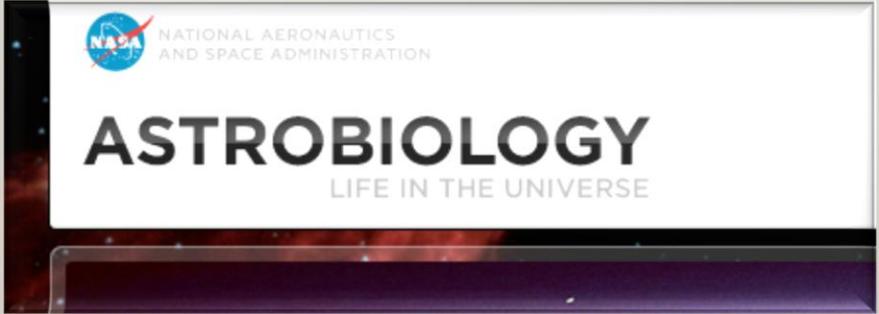
NASA Research

> [Solicitations](#)

Search for and view open, closed, past and future NASA research announcements. Text of the [solicitation announcements](#) viewed and downloaded.

Getting Started

This site provides information about NASA research opportunities. Information is intended to assist you in becoming registered with NSPIRES. See [Registration](#)

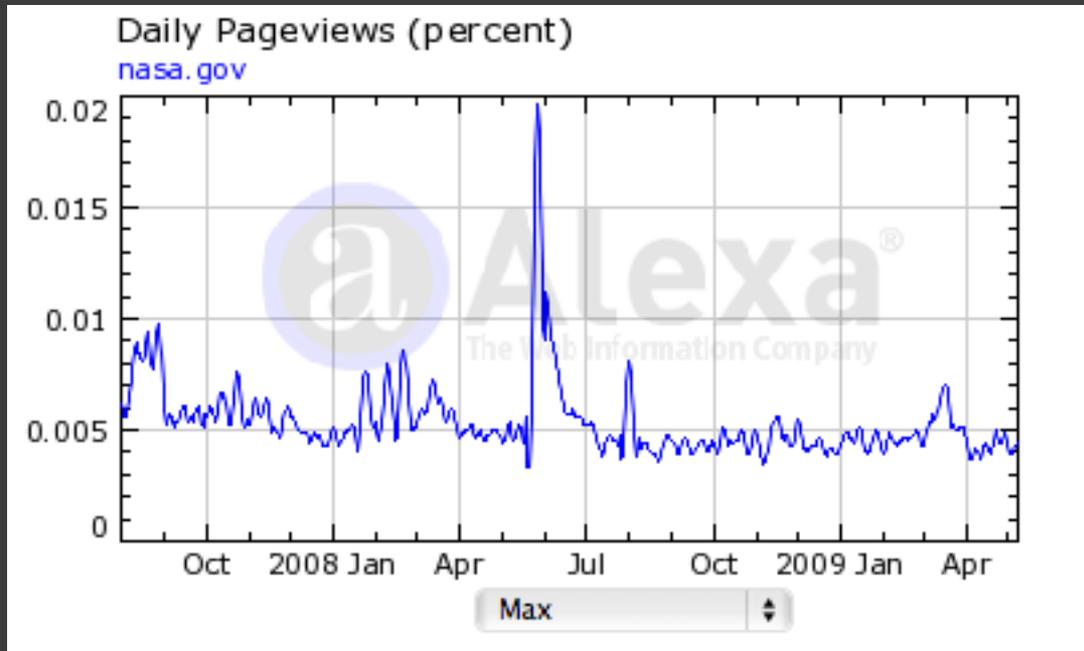


NATIONAL AERONAUTICS AND SPACE ADMINISTRATION

ASTROBIOLOGY

LIFE IN THE UNIVERSE

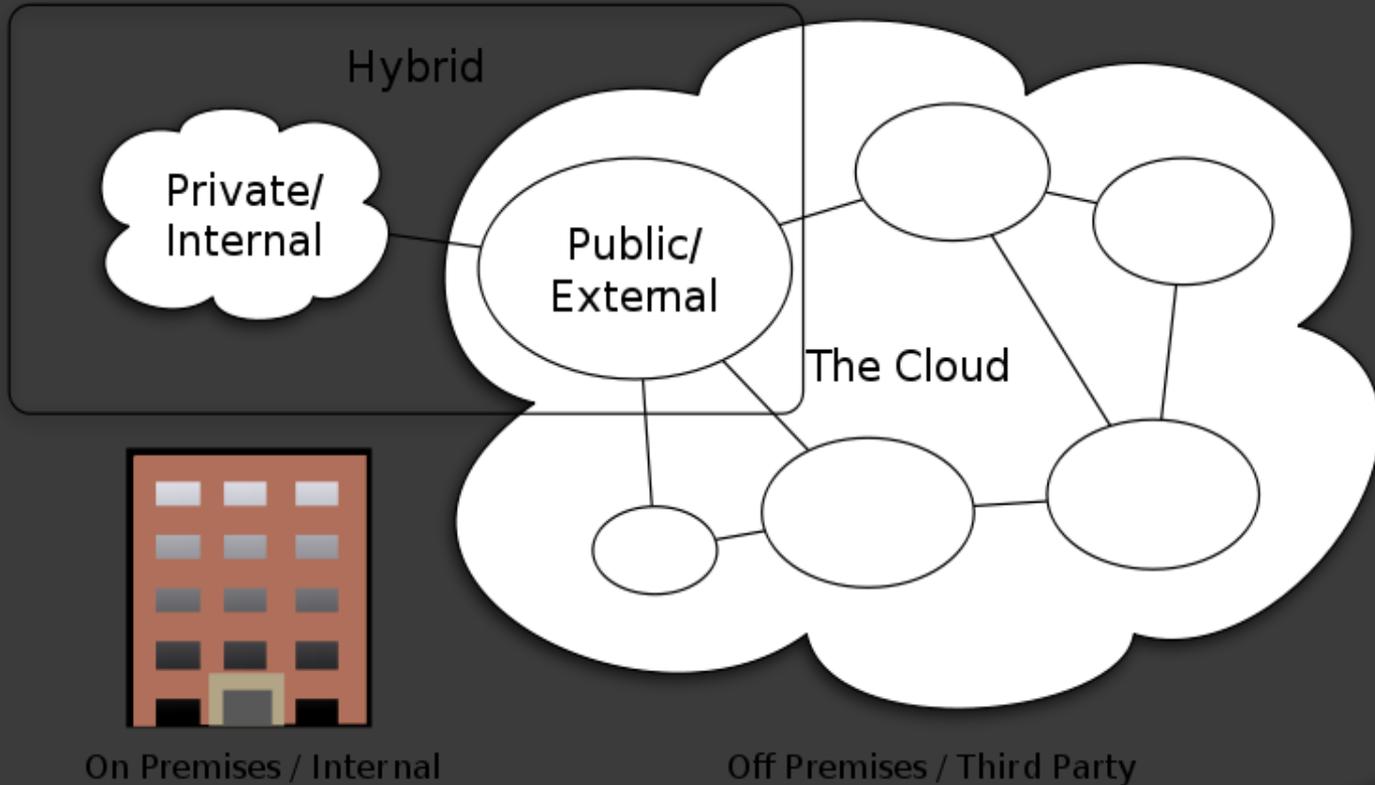
We have “Meta-Challenges”



Introducing Nebula



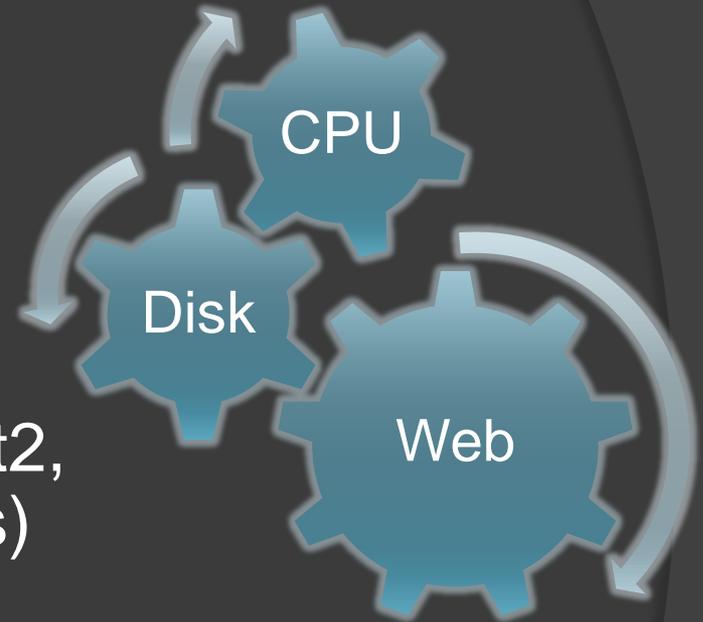
Nebula: A Hybrid Cloud



Cloud Computing Types

NEBULA – Built for Science

- ◎ Science-Class Cloud Computing
 - High CPU-to-Disk Ratio
- ◎ Built for Research
 - MAE-West Peering (Internet2, NLR, CENIC, 11 Tier-1 ISPs)
 - Massively Parallel, Loosely Coupled
- ◎ In a Federal Security Perimeter



NEBULA – Built for Collaboration

- ⦿ True Single-Sign-On, for the Public
- ⦿ Enterprise Search, across the Cloud
- ⦿ All Cloud apps live at apps.nasa.gov

NEBULA – Built for the Web

- Friendly URLs
- Designed for Search Engines, RSS, and aggregation
- It's RSSable, Tweetable
- What if NASA was on the first page of Google results for the term 'Space'?

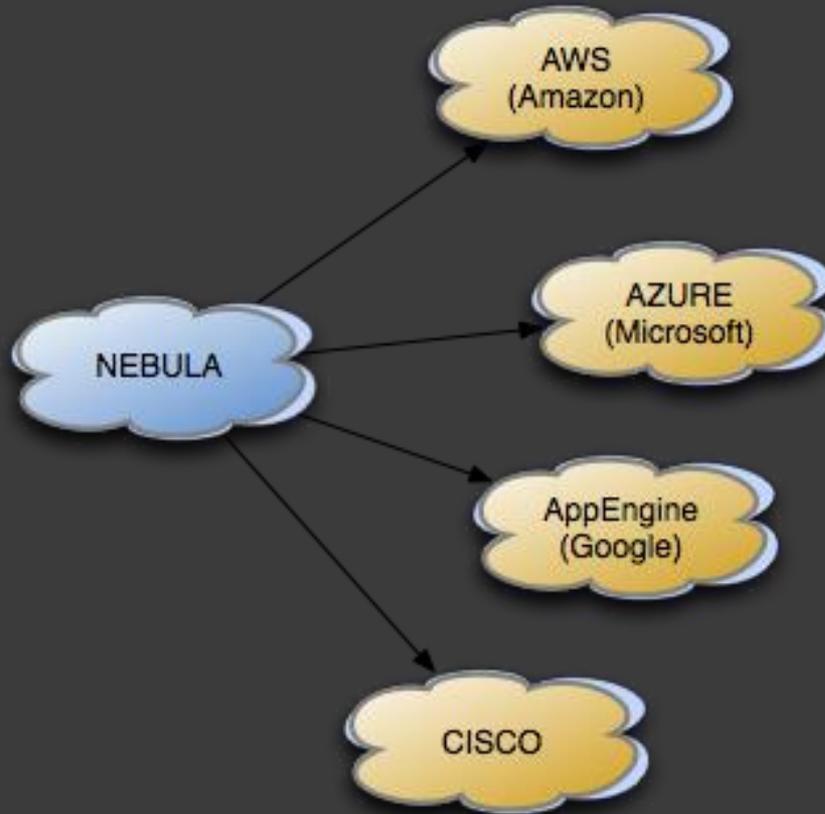
NEBULA – Built for Partners

- Your science partners can instantly connect from your NEBULA app, to their own research tasks within public Cloud Services (EC2, Azure, AppEngine)
- Your private fleet of Post-Docs can work on your data – at 10 cents an hour

NEBULA – Built for Government

- Policy compliant for contributions
- Consolidated moderation interface
- Everything-compliant (PII, First Amendment, COPPA, Section 508, etc)

NEBULA – Built for Partners



NEBULA – Built for Developers

- Integrated Development Environment
- Revision control
- Automated testing
- Continuous Integration
- Bug tracking

NEBULA – Built for Managers

- Interoperates with NASA Portal
- Controlled deployment provides simple content and security controls
- Consolidated analytics and built-in project tracking provide end-to-end visibility across multiple applications

NEBULA – Built for Community

- ⦿ Cloud means Turnkey
- ⦿ Dedicated Platform Staff
- ⦿ Engaged External Partners
- ⦿ Common KB shares code, tricks, tips

NEBULA – Built for Green

- ◉ Environmentally friendly hydro-power
- ◉ High-density data centers
- ◉ Virtualized, high-efficiency architecture



Cloud Platforms, Compared

	AWS	AppEngine	Nebula
Storage	S3, EBS	SimpleStorage API	Luster
Database	SimpleDB	CouchDB	MySQL
Queue	SQS	--	RabbitMQ
Virtualization	EC2 (Xen)	--	Eucalyptus
Framework	--	Django	Django
User Accounts	--	Google Accounts	eAuthenticatio n
Search	--	Google Search	SOLR+Lucene
Networking	1GigE	1GigE	10GigE+
CDN/Cache	CloudFront		Varnish

Application Time-to-Market

- ⦿ Gov't standards at real-world speed
- ⦿ Rapid development means minutes

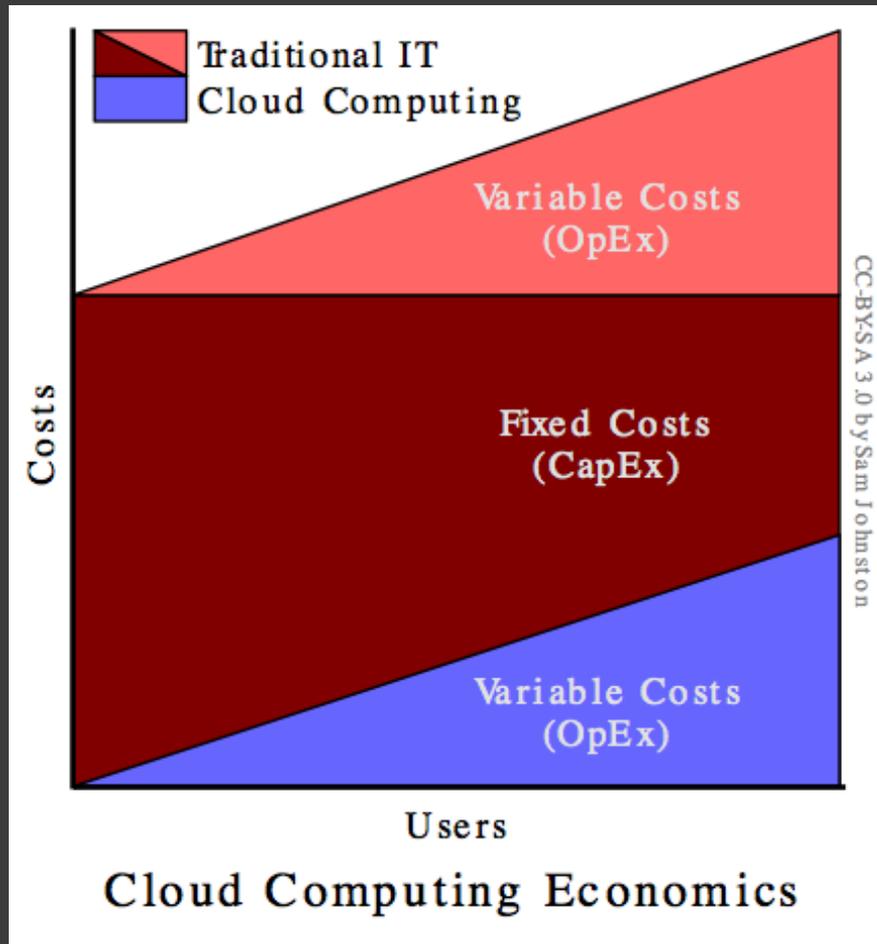
Application Time to Market

Current Web App Process	NEBULA Cloud Platform
Procure Server – 6-12 weeks	Procure VM – 60-120 seconds
Configure Server – 2-3 days	Included.
AWRS Filing – 2-3 days	Already done. (APPS.NASA.GOV)
Set up Source Control – 2-3 hours.	Included.
Security Plan – 3 weeks, min.	Included.
SSL Certificates – 2-4 weeks.	Included.
Develop Terms of Use – 6 months.	Included, for most collaboration.
eAuth Integration – 40 hours.	Included.
Develop Processes – 3-6 weeks.	Basic moderation included.
Set up backups – 2-3 hours.	Included.
24,192,000 Seconds.	60-120 Seconds.

Cost Efficiencies

- ⦿ No upfront capital expense for scientists
- ⦿ Just-in-time capacity management
- ⦿ 10^2 Cheaper (than what we do now)...

Cost Efficiencies



Continuous Improvement

- ⦿ Always getting better – free
- ⦿ Compliance with new policies – free
- ⦿ Security patches and upgrades – free
- ⦿ Surviving launch day – priceless.

NEBULA – Designed for the Future

- “Open standards are critical to the growth of cloud computing and open source software has provided the foundation for many cloud computing implementations.”
- NEBULA is 100% Open Source

NEBULA – the way to Data.gov

- Best practices in moderation, open collaboration
- Open and Public APIs, everywhere
- Feeds (RSS, Atom) power mash-ups
- Open-source platform, apps, and data
- Full transparency

By the Numbers

- Storage: Under \$2.00 / Gb, falling monthly
- Network: 10GigE to NISN, Internet2, CENIC, AT&T, Quest, L3, Verizon, Google, Amazon.com, Cisco, Microsoft
- Capacity: Currently .5 PB... adding 100TB / mo. based on customer requirements
- Service Level – Five 9s uptime (99.999%)

Pilot Projects

- White House USASpending.gov 2.0
- Microsoft World Wide Telescope (Mars / Moon)
- Google Earth Planetary Content (Mars)
- LMMP Program Data Processing (Moon)
- TOPS Earth Climate Modeling



NEBULA Timeline

- Under development since May 2008 (under codename NASA.NET)
- Now accepting a limited number of additional partners for Fall, '09
- Reserve your spot in the Cloud
- Full Launch, Q2 2010 (in Apps.gov)

Questions?

Chris.C.Kemp@NASA.gov

The Twitter logo, featuring the word "twitter" in a lowercase, rounded, blue font with a white outline.

ChrisCKemp

NASANebula