

Effectively and Securely Using the Cloud Computing Paradigm

Peter Mell, Tim Grance

NIST, Information Technology Laboratory



NIST Cloud Research Team



Peter Mell
Project Lead
301-975-5572
peter.mell@nist.gov

Tim Grance
Program Manager
301-975-4242
grance@nist.gov

Lee Badger
301-975-3176
mark.badger@nist.gov

Erika McCallister
301-975-5144
erika.mccallister@nist.gov

Karen Scarfone
301-975-8136
karen.scarfone@nist.gov

Caveats and Disclaimers

- This presentation provides education on cloud technology and its benefits to set up a discussion of cloud security
- It is NOT intended to provide official NIST guidance and NIST does not make policy
- Any mention of a vendor or product is NOT an endorsement or recommendation

Citation Note: All sources for the material in this presentation are included within the Powerpoint “notes” field on each slide

Agenda



- Part 1: Effective and Secure Use
 - Understanding Cloud Computing
 - Cloud Publications
 - Cloud Migration Paths
 - Computing Cloud Computing Security
- Part 2: Cloud Resources, Case Studies, and Security Models
 - Thoughts on Cloud Computing
 - Foundational Elements of Cloud Computing
 - Cloud Economics
 - Cloud Computing Case Studies and Security Models
 - Cloud Computing and Standards

Part I: Effective and Secure Use



Understanding Cloud Computing



Origin of the term “Cloud Computing”

- “Comes from the early days of the Internet where we drew the network as a cloud... we didn’t care where the messages went... the cloud hid it from us” – Kevin Marks, Google
- First cloud around networking (TCP/IP abstraction)
- Second cloud around documents (WWW data abstraction)
- The emerging cloud abstracts infrastructure complexities of servers, applications, data, and heterogeneous platforms
 - (“muck” as Amazon’s CEO Jeff Bezos calls it)

A Working Definition of Cloud Computing

- Cloud computing is a pay-per-use model for enabling convenient, on-demand network access to a shared pool of configurable and reliable computing resources (e.g., networks, servers, storage, applications, services) that can be rapidly provisioned and released with minimal consumer management effort or service provider interaction.
- The cloud model is comprised of five **key characteristics**, three **delivery models**, and four **deployment models**.

5 Key Cloud Characteristics

- On-demand self-service
- Ubiquitous network access
- Resource pooling
- Rapid elasticity
- Pay per use

3 Cloud Delivery Models

- Cloud Software as a Service (SaaS)
 - Use provider's applications over a network
- Cloud Platform as a Service (PaaS)
 - Deploy customer-created applications to a cloud
- Cloud Infrastructure as a Service (IaaS)
 - Rent processing, storage, network capacity, and other fundamental computing resources
- To be considered “cloud” they must be deployed on top of cloud infrastructure that has the key characteristics

4 Cloud Deployment Models

- Private cloud
 - enterprise owned or leased
- Community cloud
 - shared infrastructure for specific community
- Public cloud
 - Sold to the public, mega-scale infrastructure
- Hybrid cloud
 - composition of two or more clouds
- Two types: internal and external

Common Cloud Characteristics

- Cloud computing often leverages:
 - Massive scale
 - Virtualization
 - Free software
 - Autonomic computing
 - Multi-tenancy
 - Geographically distributed systems
 - Advanced security technologies

Cloud Computing Publications



Planned NIST Cloud Computing Publication



- NIST is planning a series of publications on cloud computing
- NIST Special Publication to be created in FY09
 - What problems does cloud computing solve?
 - What are the technical characteristics of cloud computing?
 - How can we best leverage cloud computing and obtain security?

Migration Paths for Cloud Computing



Migration Paths for Cloud Adoption

- Use public clouds
 - Option 1: as is with multi-tenancy
 - Option 2: no multi-tenancy of servers or storage + enhanced organization defined security
- Develop private clouds
 - Procure an external private cloud
 - Migrate data centers to be private clouds (fully virtualized)
- Use hybrid-cloud technology
 - Leverage a private and public cloud architectures
 - Workload portability between private and public clouds
- Build or procure community clouds

Possible Effects of Cloud Computing



- Small enterprises use public SaaS and public clouds and minimize growth of data centers
- Large enterprise data centers may evolve to act as internal clouds
- Large enterprises may use hybrid cloud infrastructure software to leverage both internal and public clouds
- Public clouds may adopt standards in order to run workloads from competing hybrid cloud infrastructures

A proposal: The Cloud Interoperability Profile

- We need to define minimal standards
 - Enable cloud integration, application portability, and data portability
 - Avoid over specification that will inhibit innovation
- Let's create a blueprint for cloud design
 - Specifies versions of standards
 - Separately addresses different cloud models
- Example: WS-I Basic Profile for SOA
- Let's call it the "Cloud Interoperability Profile (CIP)" (pronounced 'sip')

Strategies for Enabling Effective Migration

- Provide guidance on secure and effective use
- Research cloud standards
- Create a cloud interoperability profile

Cloud Computing Security



Cloud Security Roles and Responsibilities

- Cloud provider
- Cloud user and data owner
- Similarities and differences to traditional outsourcing
- Some key issues:
 - trust, multi-tenancy, encryption, compliance

Former Intel CEO, Andy Grove: “only the paranoid survive”

Cloud Security Advantages

Part 1



- Data Fragmentation and Dispersal
- Dedicated Security Team
- Greater Investment in Security Infrastructure
- Fault Tolerance and Reliability
- Greater Resiliency
- Hypervisor Protection Against Network Attacks
- Possible Reduction of C&A Activities (Access to Pre-Accredited Clouds)

Cloud Security Advantages

Part 2



- Simplification of Compliance Analysis
- Data Held by Unbiased Party (cloud vendor assertion)
- Low-Cost Disaster Recovery and Data Storage Solutions
- On-Demand Security Controls
- Real-Time Detection of System Tampering
- Rapid Re-Constitution of Services
- Advanced Honeynet Capabilities

Cloud Security Challenges

Part 1



- Data dispersal and international privacy laws
 - EU Data Protection Directive and U.S. Safe Harbor program
 - Exposure of data to foreign government and data subpoenas
 - Data retention issues
- Need for isolation management
- Multi-tenancy
- Logging challenges
- Data ownership issues
- Quality of service guarantees

Cloud Security Challenges

Part 2



- Dependence on secure hypervisors
- Attraction to hackers (high value target)
- Security of virtual OSs in the cloud
- Possibility for massive outages
- Encryption needs for cloud computing
 - Encrypting access to the cloud resource control interface
 - Encrypting administrative access to OS instances
 - Encrypting access to applications
 - Encrypting application data at rest
- Public cloud vs internal cloud security
- Lack of public SaaS version control



Additional Issues

- Issues with moving PII and sensitive data to the cloud
 - Privacy impact assessments
- Using SLAs to obtain cloud security
 - Suggested requirements for cloud SLAs
 - Issues with cloud forensics
- Contingency planning and disaster recovery for cloud implementations
- Handling compliance
 - FISMA
 - HIPAA
 - SOX
 - PCI
 - SAS 70 Audits


Speculation on FISMA and Clouds

- Cloud Infrastructure: General Support System (GSS) or Major Application (MA)?
- FIPS 199 levels limit types of cloud applications
- Physical cloud sites: Sub-systems?
 - Must address NIST SP 800-53 controls
 - Inherits all but physical cloud security controls from the GSS?
- Do we need interconnection agreements?
- Can we pre-certify public clouds for Federal use?

Part II: Cloud Resources, Case Studies, and Security Models



Thoughts on Cloud Computing



Thoughts on Cloud Computing

- Galen Gruman, InfoWorld Executive Editor, and Eric Knorr, InfoWorld Editor in Chief
 - “A way to increase capacity or add capabilities on the fly without investing in new infrastructure, training new personnel, or licensing new software.”
 - “The idea of loosely coupled services running on an agile, scalable infrastructure should eventually make every enterprise a node in the cloud.”


Thoughts on Cloud Computing

- Tim O'Reilly, CEO O'Reilly Media
- “I think it is one of the foundations of the next generation of computing”
- “The network of networks is the platform for all computing”
- “Everything we think of as a computer today is really just a device that connects to the big computer that we are all collectively building”



Thoughts on Cloud Computing

- Dan Farber, Editor in Chief CNET News
- “We are at the beginning of the age of planetary computing. Billions of people will be wirelessly interconnected, and the only way to achieve that kind of massive scale usage is by massive scale, **brutally efficient** cloud-based infrastructure.”



Core objectives of Cloud Computing

- Amazon CTO Werner Vogels
- Core objectives and principles that cloud computing must meet to be successful:
 - Security
 - Scalability
 - Availability
 - Performance
 - Cost-effective
 - Acquire resources on demand
 - Release resources when no longer needed
 - Pay for what you use
 - Leverage others' core competencies
 - Turn fixed cost into variable cost



Amazon CTO Werner Vogels
(Credit: Dan Farber)

A “sunny” vision of the future

- Sun Microsystems CTO Greg Papadopoulos
 - Users will “trust” service providers with their data like they trust banks with their money
 - “Hosting providers [will] bring ‘brutal efficiency’ for utilization, power, security, service levels, and idea-to-deploy time” –CNET article
 - Becoming cost ineffective to build data centers
 - Organizations will rent computing resources
 - Envisions grid of 6 cloud infrastructure providers linked to 100 regional providers

Foundational Elements of Cloud Computing



Foundational Elements of Cloud Computing



Primary Technologies

- Virtualization
- Grid technology
- Service Oriented Architectures
- Distributed Computing
- Broadband Networks
- Browser as a platform
- Free and Open Source Software

Other Technologies

- Autonomic Systems
- Web 2.0
- Web application frameworks
- Service Level Agreements

Web 2.0

- Is not a standard but an evolution in using the WWW
- “Don’t fight the Internet” – CEO Google, Eric Schmidt
- Web 2.0 is the trend of using the full potential of the web
 - Viewing the Internet as a computing platform
 - Running interactive applications through a web browser
 - Leveraging interconnectivity and mobility of devices
 - The “long tail” (profits in selling specialized small market goods)
 - Enhanced effectiveness with greater human participation
- Tim O'Reilly: “Web 2.0 is the business revolution in the computer industry caused by the move to the Internet as a platform, and an attempt to understand the rules for success on that new platform.”

Software as a Service (SaaS)

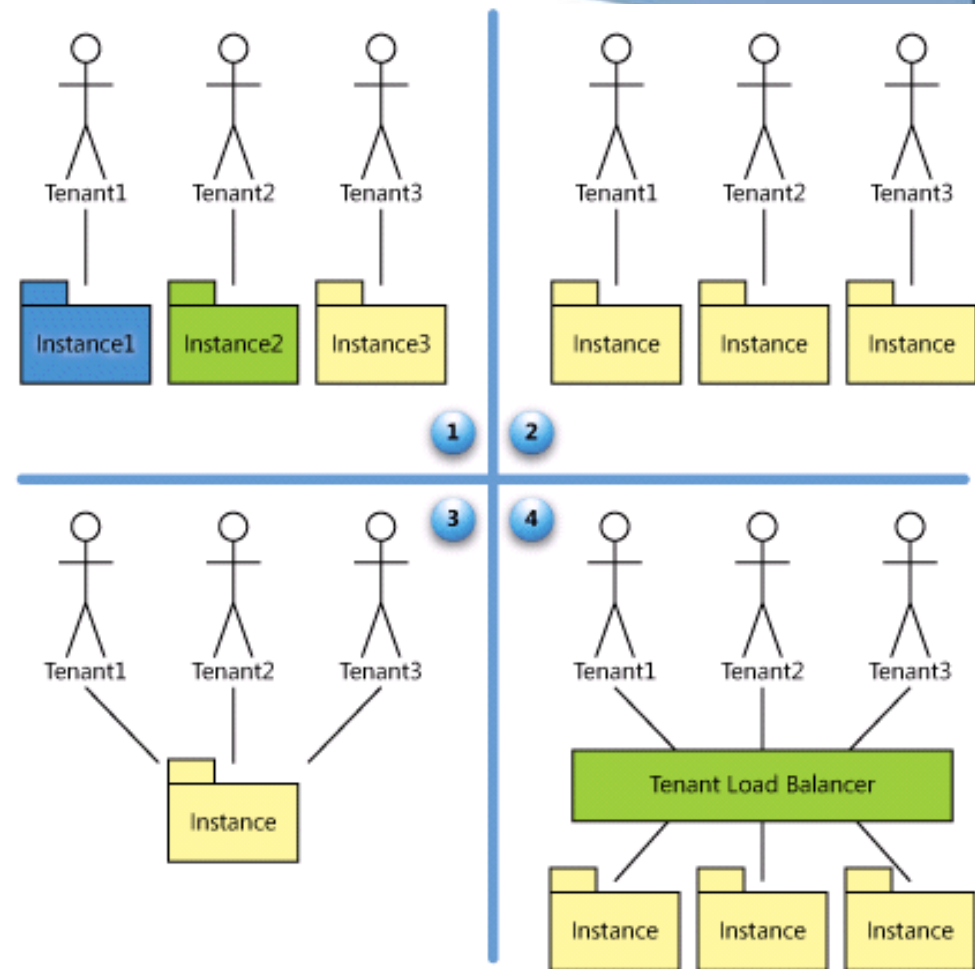
- SaaS is hosting applications on the Internet as a service (both consumer and enterprise)
- Jon Williams, CTO of Kaplan Test Prep on SaaS
 - “I love the fact that I don't need to deal with servers, staging, version maintenance, security, performance”
- Eric Knorr with Computerworld says that “[there is an] increasing desperation on the part of IT to minimize application deployment and maintenance hassles”

Three Features of Mature SaaS Applications

- **Scalable**
 - Handle growing amounts of work in a graceful manner
- **Multi-tenancy**
 - One application instance may be serving hundreds of companies
 - Opposite of multi-instance where each customer is provisioned their own server running one instance
- **Metadata driven configurability**
 - Instead of customizing the application for a customer (requiring code changes), one allows the user to configure the application through metadata

SaaS Maturity Levels

- Level 1: Ad-Hoc/Custom
- Level 2: Configurable
- Level 3: Configurable, Multi-Tenant-Efficient
- **Level 4: Scalable, Configurable, Multi-Tenant-Efficient**



Utility Computing

- “Computing may someday be organized as a public utility” - John McCarthy, MIT Centennial in 1961
- Huge computational and storage capabilities available from utilities
- Metered billing (pay for what you use)
- Simple to use interface to access the capability (e.g., plugging into an outlet)

Service Level Agreements (SLAs)

- Contract between customers and service providers of the level of service to be provided
- Contains performance metrics (e.g., uptime, throughput, response time)
- Problem management details
- Documented security capabilities
- Contains penalties for non-performance

Autonomic System Computing

- Complex computing systems that manage themselves
- Decreased need for human administrators to perform lower level tasks
- Autonomic properties: Purposeful, Automatic, Adaptive, Aware
- IBM's 4 properties: self-healing, self-configuration, self-optimization, and self-protection

*IT labor costs are 18 times that of equipment costs.
The number of computers is growing at 38% each year.*


Grid Computing

- Distributed parallel processing across a network
- Key concept: “the ability to negotiate resource-sharing arrangements”
- Characteristics of grid computing
 - Coordinates independent resources
 - Uses open standards and interfaces
 - Quality of service
 - Allows for heterogeneity of computers
 - Distribution across large geographical boundaries
 - Loose coupling of computers

Platform Virtualization

- “[Cloud computing] relies on separating your applications from the underlying infrastructure” - Steve Herrod, CTO at VMware
- Host operating system provides an abstraction layer for running virtual guest OSs
- Key is the “hypervisor” or “virtual machine monitor”
 - Enables guest OSs to run in isolation of other OSs
 - Run multiple types of OSs
- Increases utilization of physical servers
- Enables portability of virtual servers between physical servers
- Increases security of physical host server

Web Services



- Web Services
 - Self-describing and stateless modules that perform discrete units of work and are available over the network
 - “Web service providers offer APIs that enable developers to exploit functionality over the Internet, rather than delivering full-blown applications.” - Infoworld
 - Standards based interfaces (WS-I Basic Profile)
 - e.g., SOAP, WSDL, WS-Security
 - Enabling state: WS-Transaction, Choreography
 - Many loosely coupled interacting modules form a single logical system (e.g., legos)

Service Oriented Architectures

- Service Oriented Architectures
 - Model for using web services
 - service requestors, service registry, service providers
 - Use of web services to compose complex, customizable, distributed applications
 - Encapsulate legacy applications
 - Organize stovepiped applications into collective integrated services
 - Interoperability and extensibility

Web application frameworks

- Coding frameworks for enabling dynamic web sites
 - Streamline web and DB related programming operations (e.g., web services support)
 - Creation of Web 2.0 applications
- Supported by most major software languages
- Example capabilities
 - Separation of business logic from the user interface (e.g., Model-view-controller architecture)
 - Authentication, Authorization, and Role Based Access Control (RBAC)
 - Unified APIs for SQL DB interactions
 - Session management
 - URL mapping
- Wikipedia maintains a list of web application frameworks

Free and Open Source Software



- External ‘mega-clouds’ must focus on using their massive scale to reduce costs
- Usually use free software
 - Proven adequate for cloud deployments
 - Open source
 - Owned by provider
- Need to keep per server cost low
 - Simple commodity hardware
 - Handle failures in software

Cloud Economics



Cost of Traditional Data Centers



- 11.8 million servers in data centers
- Servers are used at only 15% of their capacity
- 800 billion dollars spent yearly on purchasing and maintaining enterprise software
- 80% of enterprise software expenditure is on installation and maintenance of software
- Data centers typically consume up to 100 times more per square foot than a typical office building
- Average power consumption per server quadrupled from 2001 to 2006.
- Number of servers doubled from 2001 to 2006

Energy Conservation and Data Centers

- Standard 9000 square foot costs \$21.3 million to build with \$1 million in electricity costs/year
- Data centers consume 1.5% of our Nation's electricity (EPA)
 - .6% worldwide in 2000 and 1% in 2005
- Green technologies can reduce energy costs by 50%
- IT produces 2% of global carbon dioxide emissions


Cloud Economics

- Estimates vary widely on possible cost savings
- “If you move your data centre to a cloud provider, it will cost a tenth of the cost.” – Brian Gammage, Gartner Fellow
- Use of cloud applications can reduce costs from 50% to 90% - CTO of Washington D.C.
- IT resource subscription pilot saw 28% cost savings - Alchemy Plus cloud (backing from Microsoft)
- Preferred Hotel
 - Traditional: \$210k server refresh and \$10k/month
 - Cloud: \$10k implementation and \$16k/month

Cloud Economics

- George Reese, founder Valtira and enStratus
 - Using cloud infrastructures saves 18% to 29% before considering that you no longer need to buy for peak capacity

Cloud Computing Case Studies and Security Models



Google Cloud User: City of Washington D.C.

- Vivek Kundra, CTO for the District (now OMB e-gov administrator)
- Migrating 38,000 employees to Google Apps
- Replace office software
 - Gmail
 - Google Docs (word processing and spreadsheets)
 - Google video for business
 - Google sites (intranet sites and wikis)
- “It's a fundamental change to the way our government operates by moving to the cloud. Rather than owning the infrastructure, we can save millions.”, Mr. Kundra
- 500,000+ organizations use Google Apps
- GE moved 400,000 desktops from Microsoft Office to Google Apps and then migrated them to Zoho for privacy concerns

Are Hybrid Clouds in our Future?

- OpenNebula
- Zimory
- IBM-Juniper Partnership
 - "demonstrate how a hybrid cloud could allow enterprises to seamlessly extend their private clouds to remote servers in a secure public cloud..."
- VMWare VCloud
 - "Federate resources between internal IT and external clouds"

vCloud Initiative

- Goal:
 - “Federate resources between internal IT and external clouds”
 - Application portability
 - Elasticity and scalability, disaster recovery, service level management
- vServices provide APIs and technologies
- Appears to enable portability between cloud providers but locks applications into VMWare technologies
- BT, Rackspace, Verizon, SunGard, SAVVIS, T-Systems and 100+ partners

Microsoft Azure Services

 Windows Live  Microsoft Office Live  Microsoft Exchange Online  Microsoft SharePoint Online  Microsoft Dynamics CRM Online

Azure™ Services Platform

 Live Services

 Microsoft .NET Services

 Microsoft SQL Services


 Microsoft SharePoint Services

 Microsoft Dynamics CRM Services

 Windows Azure™

Source: Microsoft Presentation, A Lap Around Windows Azure, Manuvir Das

Windows Azure Applications, Storage, and Roles



Source: Microsoft Presentation, A Lap Around Windows Azure, Manuvir Das

Case Study: Facebook's Use of Open Source and Commodity Hardware (8/08)

- Jonathan Heiliger, Facebook's vice president of technical operations
- 80 million users + 250,000 new users per day
- 50,000 transactions per second, 10,000+ servers
- Built on open source software
 - Web and App tier: Apache, PHP, AJAX
 - Middleware tier: Memcached (Open source caching)
 - Data tier: MySQL (Open source DB)
- Thousands of DB instances store data in distributed fashion (avoids collisions of many users accessing the same DB)
- “We don't need fancy graphics chips and PCI cards,” he said. “We need one USB port and optimized power and airflow. Give me one CPU, a little memory and one power supply. If it fails, I don't care. We are solving the redundancy problem in software.”

Case Study: IBM-Google Cloud (8/08)

- “Google and IBM plan to roll out a worldwide network of servers for a cloud computing infrastructure” – Infoworld
- Initiatives for universities
- Architecture
 - Open source
 - Linux hosts
 - Xen virtualization (virtual machine monitor)
 - Apache Hadoop (file system)
 - “open-source software for reliable, scalable, distributed computing”
 - IBM Tivoli Provisioning Manager

Case Study: Amazon Cloud

- Amazon cloud components
 - Elastic Compute Cloud (EC2)
 - Simple Storage Service (S3)
 - SimpleDB
- New Features
 - Availability zones
 - Place applications in multiple locations for failovers
 - Elastic IP addresses
 - Static IP addresses that can be dynamically remapped to point to different instances (not a DNS change)

Amazon Cloud Users: New York Times and Nasdaq

(4/00)

- Both companies used Amazon's cloud offering
- New York Times
 - Didn't coordinate with Amazon, used a credit card!
 - Used EC2 and S3 to convert 15 million scanned news articles to PDF (4TB data)
 - Took 100 Linux computers 24 hours (would have taken months on NYT computers)
 - "It was cheap experimentation, and the learning curve isn't steep." – Derrick Gottfrid, Nasdaq
- Nasdaq
 - Uses S3 to deliver historic stock and fund information
 - Millions of files showing price changes of entities over 10 minute segments
 - "The expenses of keeping all that data online [in Nasdaq servers] was too high." – Claude Courbois, Nasdaq VP
 - Created lightweight Adobe AIR application to let users view data

Cloud Computing and Standards



A Perspective on Cloud Computing and Standards

- Cloud computing is a convergence of many technologies
 - Some have their own standards
- This convergence combined with massively scaled deployments represents “**leap-ahead**” capabilities
- We have a choice
 - Proprietary stovepipe clouds
 - Development of an expensive cloud integration market
 - Standards based clouds

Major Concerns and Needs

- Major concerns:
 - Vendor lock-in
 - High availability (whole clouds do fail)
 - Ability to use multiple cloud vendors simultaneously
 - Integration with internal data centers
- Major needs:
 - Cloud application and workload portability
 - Cloud data recovery (in usable formats)
 - Standardized cloud control interfaces
 - Standardized security interfaces

So where are we?

- Today cloud standards don't exist
 - GRID computing has many standards but isn't the same model
 - SOA/WS has standards but this only applies to the applications running on the cloud
- Standards will be vital to achieve success
- Hurdle: Many definitions/models
- Can't standardize what you can't define

Questions?

- Peter Mell
- NIST, Information Technology Laboratory
- Computer Security Division
- 301-975-5572
- mell@nist.gov

- Tim Grance
- NIST, Information Technology Laboratory
- Computer Security Division
- 301-975-4242
- grance@nist.gov