

Microsoft
Research



Microsoft Research Asia **Faculty Summit 2012**



Cloud Computing Technology in Scientific Research

Yuanchun Zhou (zyc@cnic.cn)

Geng Shen(shengeng@cnic.cn)

Yuanke Wei(weiyuanke@cnic.cn)

Computer Network Information Center
Chinese Academy of Sciences

October, 27, 2012, Tianjin

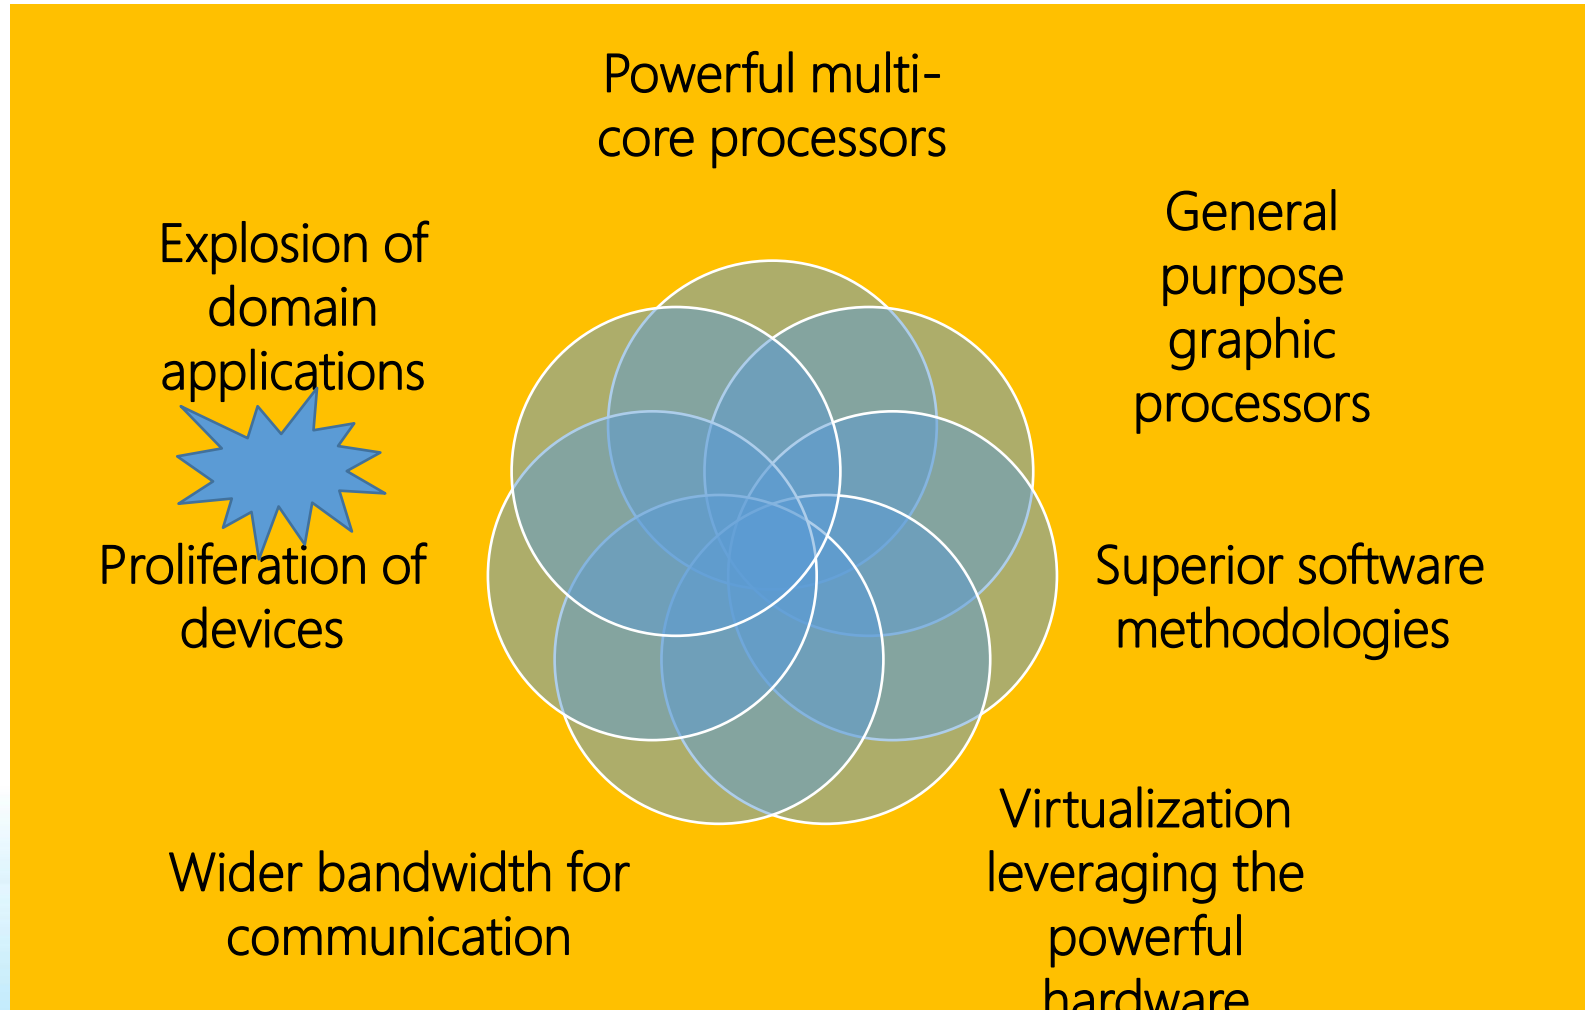


Outline

- What is Cloud Computing
- A case-study of Scientific Research of the Cloud
- Conclusion



A Golden Era in Computing





Global Data, Global Science

- Sloan Digital Sky Survey (SDSS) year 2000
 - about 200 GB every night
 - amassed more than 140 TB
- LHC(Large Hadron Collider)
 - 15 PB / year, Total data: Several hundred PB total
 - Analysed by thousands of physical scientists world-wide
- Walmart
 - handles 1 million customer transactions every hour
 - estimated to contain more than 2.5 PB data
- FaceBook
 - handles 40 billion photos from its user base





Challenges

- Alignment with the needs of the business / user / non-computer specialists / community and society
- Need to address the scalability issue: large scale data, high performance computing, automation, response time, rapid prototyping, and rapid time to production
- Need to effectively address
 - ✓ ever shortening cycle of obsolescence
 - ✓ heterogeneity
 - ✓ rapid changes in requirements
- Transform data from diverse sources into intelligence and deliver intelligence to right people/user/systems
- What about providing all this in a cost-effective manner?



Answer: The Cloud Computing?

- Typical requirements and models:
 - ✓ platform (PaaS),
 - ✓ software (SaaS),
 - ✓ infrastructure (IaaS),
 - ✓ Services-based application programming interface (API)
- A cloud computing environment can provide one or more of these requirements for a cost
- Pay as you go model of business
- When using a public cloud the model is similar to renting a property than owning one.
- An organization could also maintain a private cloud and/or use both.



What is the cloud computing

- **Cloud Computing** is a general term used to describe a new class of network based computing that takes place over the Internet,
 - ✓ basically a step on from Utility Computing
 - ✓ a collection/group of integrated and networked hardware, software and Internet infrastructure (called a platform).
 - ✓ Using the Internet for communication and transport provides hardware, software and networking services to clients
- These platforms hide the complexity and details of the underlying infrastructure from users and applications by providing very simple graphical interface or API (Applications Programming Interface).



What is the cloud computing

- In addition, the platform provides on demand services, that are always on, anywhere, anytime and any place.
- Pay for use and as needed, elastic
 - ✓ scale up and down in capacity and functionalities
- The hardware and software services are available to
 - ✓ general public, enterprises, corporations and businesses markets



Cloud Deployment Model

NIST Deployment Models

Public Cloud

Cloud infrastructure made available to the general public.

Private Cloud

Cloud infrastructure operated solely for an organization.

Hybrid Cloud

Cloud infrastructure composed of two or more clouds that interoperate or federate through technology

Community Cloud

Cloud infrastructure shared by several organizations and supporting a specific community

... and one other

Virtual Private Cloud

Cloud services that simulate the private cloud experience in public cloud infrastructure



Typical Cloud Platform


Windows Azure



Amazon EC2



Google App Engine





Outline

- What is Cloud Computing
- **A case-study of Scientific Research of the Cloud**
- Conclusion



CHINESE ACADEMY OF SCIENCES

- CAS is a leading academic institution and comprehensive research and development center in natural science, technological science and high-tech innovation in China.
- It was founded in Beijing on 1st November 1949 on the basis of the former Academia Sinica (Central Academy of Sciences) and Peiping Academy of Sciences.





CHINESE ACADEMY OF SCIENCES

Today's
with legal ent
and national
1,000 field st
Its sta

The total number of graduates in the whole
academy has now reached 44,978, including 45.5%
PhD candidates.





Scientific Data Deluge in CAS

- Large scientific facilities produce huge data
 - +20 being operation
 - +20 under construction
- Long-Term field observation stations
 - +100 stations including Ecology, Environment, Space, etc.
- Long-Term Research data need to be archived and curation and sharing
 - 100+ institutes



Field observation stations



Large Scientific facilities



Computer Network Information Center(CNIC)

- is a supporting institute of CAS for the Construction, Operation and Services of Cyberinfrastructure

• China Science and Technology Network Center

• Supercomputing Center


• Scientific Data Center

• ARP Operation Support Center

• Internet-based Science Communication Center

• China Internet Network Information Center

154,
20%



■ Faculty and Staff
■ Graduate Students



ARP



Microsoft Research Asia
Faculty Summit 2012

CNIC

中国互联网络信息中心
China Internet Network Information Center



Data Centers Distribution of CNIC

● Scientific Data

- ~1PB
- Above 60 institutions
- Multiple Disciplines

● Storage Capacity

- ~ 22PB(50PB)
- 1 major center
- 1 archive center
- 12 middle-size center

● Computing Capacity

- ~ 5000(10000) CPU cores
- Dedicated design for DIC





New challenge and requirement in Modern scientific research



Cloud Computing can help ?

- HPC (scientific computing)
- BigData (scientific data)



Case 1: High-Performance Cloud Computing

- Scientific computing often requires the availability of a massive number of computers for performing large scale experiments. Traditionally, these needs have been addressed by using high-performance computing solutions and installed facilities such as clusters and super computers, which are difficult to setup, maintain, and operate.
- Cloud computing provides scientists with a completely new model of utilizing the computing infrastructure. Compute resources, storage resources, as well as applications, can be dynamically provisioned (and integrated within the existing infrastructure) on a pay per use basis. These resources can be released when they are no more needed.



Case 1: High-Performance Cloud Computing

Move the computing work to cloud

An example of SaaS:

CAS Data Cloud Computing Service

We provide some simple scientific computing service. People can submit Application Wares jobs through a website, then our scheduler assign adequate resources for this job.



Case 1: *High-Performance Cloud Computing*

Move the computing work to cloud

Application Wares

Gauss Computation

Group:

You are the group manager

Description:

The application is now serving for Prof. Hao Wen's Research Group from Institution of Information Engineering, CAS. Gaussian 09 is the latest in the Gaussian series of programs. It provides capabilities for electronic structure modeling. Gaussian 09 is licensed for a wide variety of applications. All versions of Gaussian 09 contain every scientific/modeling feature, and none impose limitations on calculations other than your computing resources and patience.

Operation:

Submit Single Job

Submit Batch Jobs

Submit Zip Jobs

People can submit batch of jobs at one time.

When Huge amounts of jobs come, our computing resource get an effective utilization .



Case 1: *High-Performance Cloud Computing*

Move the computing work to cloud

Home Application Supermarket Submit Job My jobs My Submitted zip

Home

🔍 ⚙️

job statistics

Gausso9Job

- History : 10113
- Unknown : 3
- total : 10116 in all

BeastIBJob

- History : 15
- total : 15 in all

My Jobs

JobID	Command	Arguments
234501		/panfs/MPIBlade_Users/ib/zdchen_ib/beast1.7.1/BEAST_Cayratia_small.xml
234498		/panfs/MPIBlade_Users/ib/zdchen_ib/beast1.7.1/BEAST_Cayratia_7.25.xml
231143		/panfs/MPIBlade_Users/ib/zdchen_ib/beast1.7.1/combine_new8.xml
231142		/panfs/MPIBlade_Users/ib/zdchen_ib/beast1.7.1/combine_new7.xml

Management and statistics of submitted jobs are extended functions.



Case 2: Public Cloud of CAS Based on OpenStack

Opensource help us to establish our own cloud

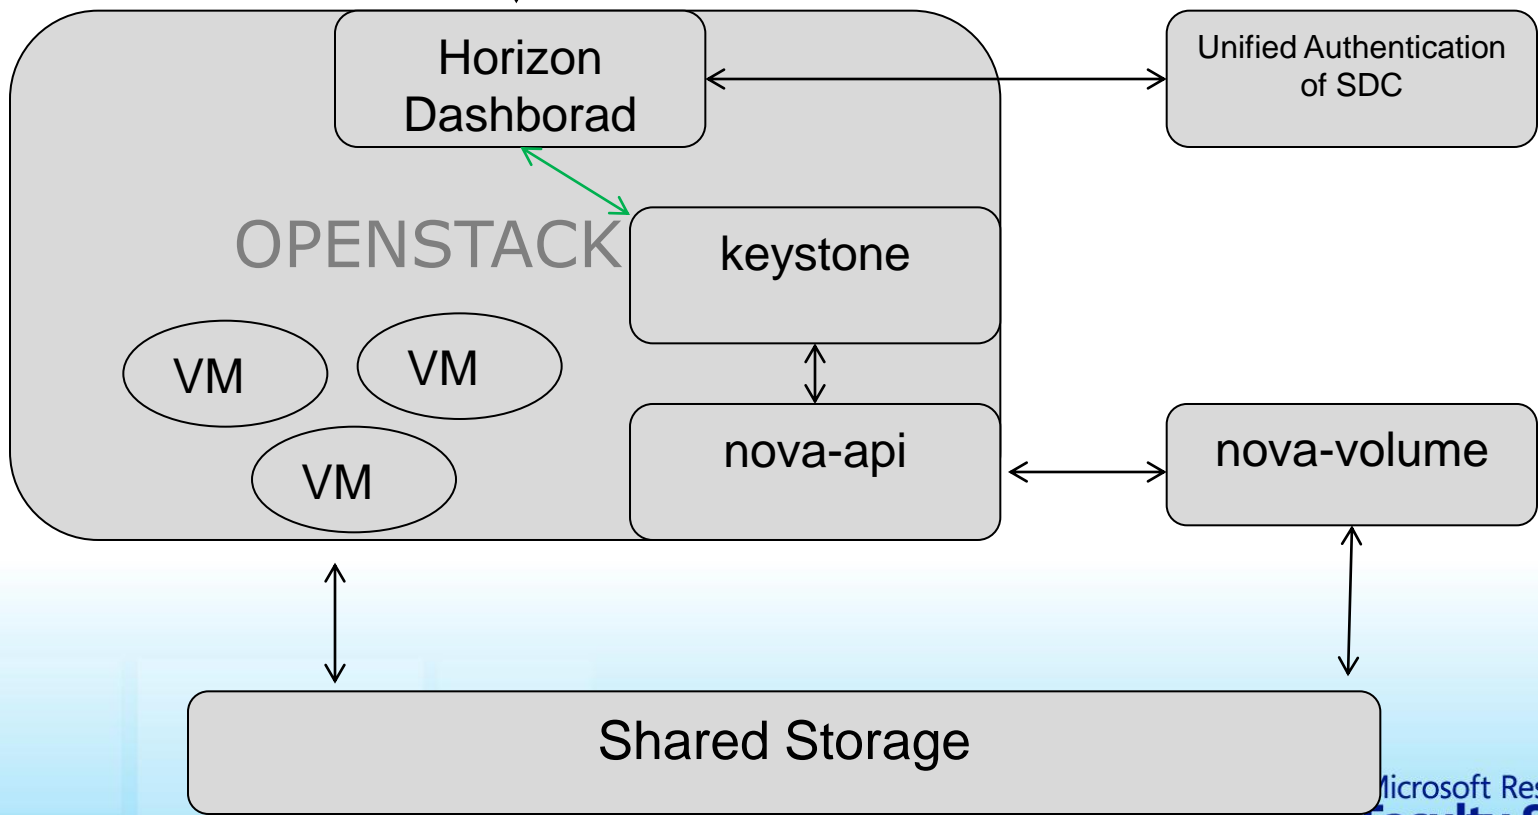
*“**OpenStack** is a global collaboration of developers and cloud computing technologists producing the ubiquitous open source cloud computing platform for public and private clouds. The project aims to deliver solutions for all types of clouds by being simple to implement, massively scalable, and feature rich. The technology consists of a series of interrelated projects delivering various components for a cloud infrastructure solution.”*



Platform Architecture of Public Cloud in CAS



User





the Status of Public Cloud in CAS

- 9 nodes (one controller + 7 compute node + one storage node)
 - 32G RAM, 16cores, 4T DISK
 - OS: ubuntu 11.10
 - Hypervisor: KVM
- Openstack Release: ESSEX
- Network: FlatDHCP, multi-host
- 7 months
- Use Puppet to deploy and manage
- Maximum VMS: 300+



the Future Plan of Public Cloud in CAS

- Openstack HA
 - HA MySQL
 - DRBD Master/Slave
 - Multi Master
 - HA RabbitMQ
 - DRBD Master/Slave (Deprecated)
 - RabbitMQ Active/Active (Mirrored Queues)
 - HA Glance/Keystone
 - Glance API
 - Glance Registry
 - Keystone
 - Nova
 - Front-end API Servers - HW/SW load balancer
 - Other services
- HPC in Openstack
 - Dedicated configuration and optimization for VMS
- Schedule algorithm
 - Different types of application scenario
 - Customized requirements
- Monitor Components



Case 3: A Migratory Birds' Spatial Distribution Prediction System Based on *Hadoop*

Lots of algorithms used by scientific computing can be divided to adapted to the MapReduce model. Mapreduce can describes problems suitable for clouds. MapReduce covers many high throughput computing applications including "parameter searches". Many data analysis applications including information retrieval fit the MapReduce paradigm.



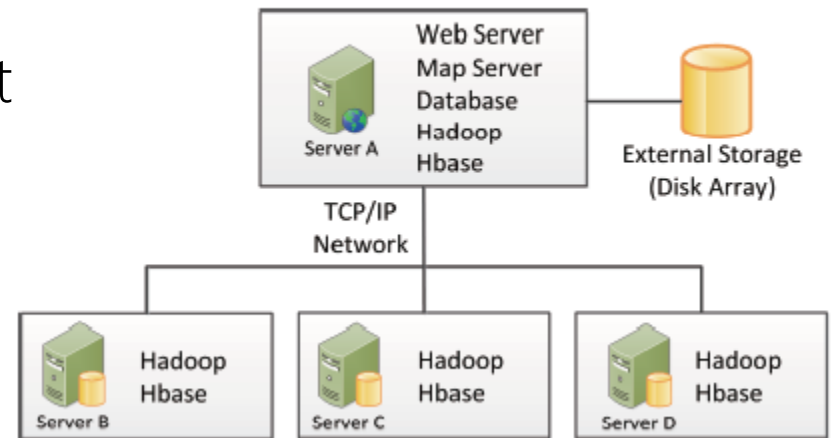
Case 3: A Migratory Birds' Spatial Distribution Prediction System Based on *Hadoop*

The Bird-SDPS uses birds' GPS tracking data and remote sensing data as input to build multiple distribution models, which are implemented by different programming languages. And the system provides online access and visualization functions. In order to store large dataset of remote sensing data, we design a hybrid storage structure based on Hbase (based on hadoop).



Case 3: A Migratory Birds' Spatial Distribution Prediction System Based on *Hadoop*

- we have to confront a much larger dataset (larger than 100 GB). The dataset is so large that it cannot be stored and accessed through relational database.
- Therefore, we employed HBase, which is a typical NoSQL database based on Hadoop framework, to store and process these text data.





Case 4: *Collaboration with Microsoft(Windows Azure)*



Windows Azure

- Enterprise-level on-demand capacity builder
- Fabric of cycles and storage available on-request for a cost
- You have to use Azure API to work with the infrastructure offered by Microsoft
- Significant features: web role, worker role , blob storage, table and drive-storage



Case 4: *Collaboration with Microsoft(Windows Azure)*

We got a certain amount of Windows Azure Resources and we are trying to use them properly. More in-depth cooperation will be carried out.

Windows Azure Collaboration Timeline

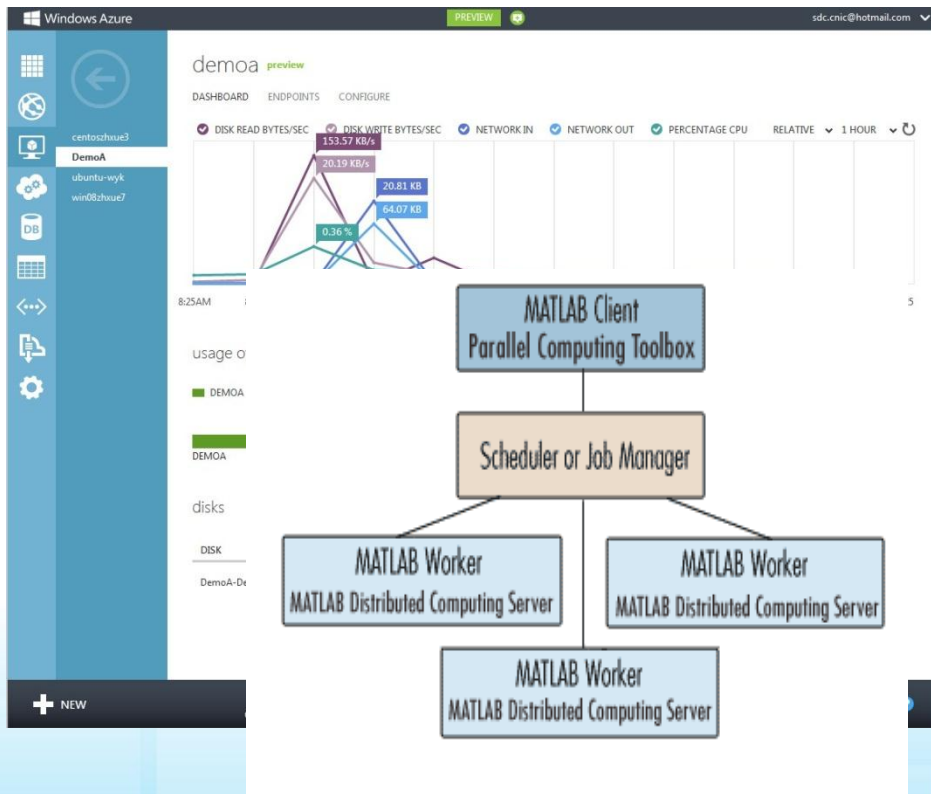
2012

- February 15 Effective Date of Microsoft Online Subscription Agreement
- February 22 Our account is ready for activation in the system
- April 02 Deploy our first Windows Azure testing service
- July 10 Our experimental mapreduce job running on Azure Hadoop
- October 22 We made a plan of further collaboration with MicroSoft



Case 4: *Collaboration with Microsoft(Windows Azure)*

The use of Azure



Open Azure to our graduate student to help them establish experiments environment. Also, we can get high-performance Windows/Linux hosts from Azure to establish distributed Matlab environment to run some algorithms.



Case 4: *Collaboration with Microsoft (Windows Azure)*

Hadoop Azure is an exciting product

Reduce to deduce!

Your Cluster: test2cnic.cloudapp.net

JavaScript: Idle
Hive: Idle

Interactive Console

Status: OK

Remote Desktop

ODBC Server

Open Ports

Manage Cluster

0
Job History

9
Samples

Downloads

Hadoop Azure make people can generate their own Hadoop cluster on cloud. All you need to do is submitting mapreduce job on the website. It will make everything easy and safe.

Your Tasks

+ New

Create Job

Microsoft
ResearchConnections



Microsoft Research Asia
Faculty Summit 2012



Case 4: *Collaboration with Microsoft(Windows Azure)*

Plans of Windows Azure academic collaboration

Windows Azure Theme - Project Proposal

▪ **Project Name:**

▪ **Principal Investigator(s):**
Name, title, Department, University or Institute, email, phone number, personal web URL

▪ **Abstract (150 to 300 words)**
Give a clear summary of the overall project on the problem you propose to solve, your approach for investing the problem, and the tangible outputs. State what your proposed project will contribute to the area you worked in and why this is relevant for MSRA to support.

▪ **Project Description (No exceed 5 single-space pages, 10pt. font or larger)**

- Description of the problem to be investigated and its technical importance
- Objectives to be achieved to contribute to its solution
- Proposed methodology to achieve the objectives, include diagrams or algorithms if it is helpful to describe your methodology clearly and concisely

▪ **Major expected outputs and plan for disseminating the results**
Note: publication of the results should acknowledge the sponsorship of MSRA

- Social Impact - [Benefit to large number of users or plans to solve the real world challenges in the environmental, social, research, government, and/or other areas]
- Benefit to Talent - [Training of high-quality research and or engineering talent]
- Publications - [Books, chapters, journal papers, conferences, white papers, and technical reports, and dissemination plan]
- Demonstration - [Demo video and/or website]
- Research Tools - [Research tools be developed as a result of this project? If so, are they being made available to other researchers in your domain?]
- Curriculum Materials - [Training courses, tutorial programs or other educational materials developed in association with the project]

we will issue a statement about an *academic* collaboration. You can fill in a simple proposal paper and send it to us. We will collect the proposal and do some associated review work with Microsoft, we can provide Window Azure resources to users who can offer their stated requirement.

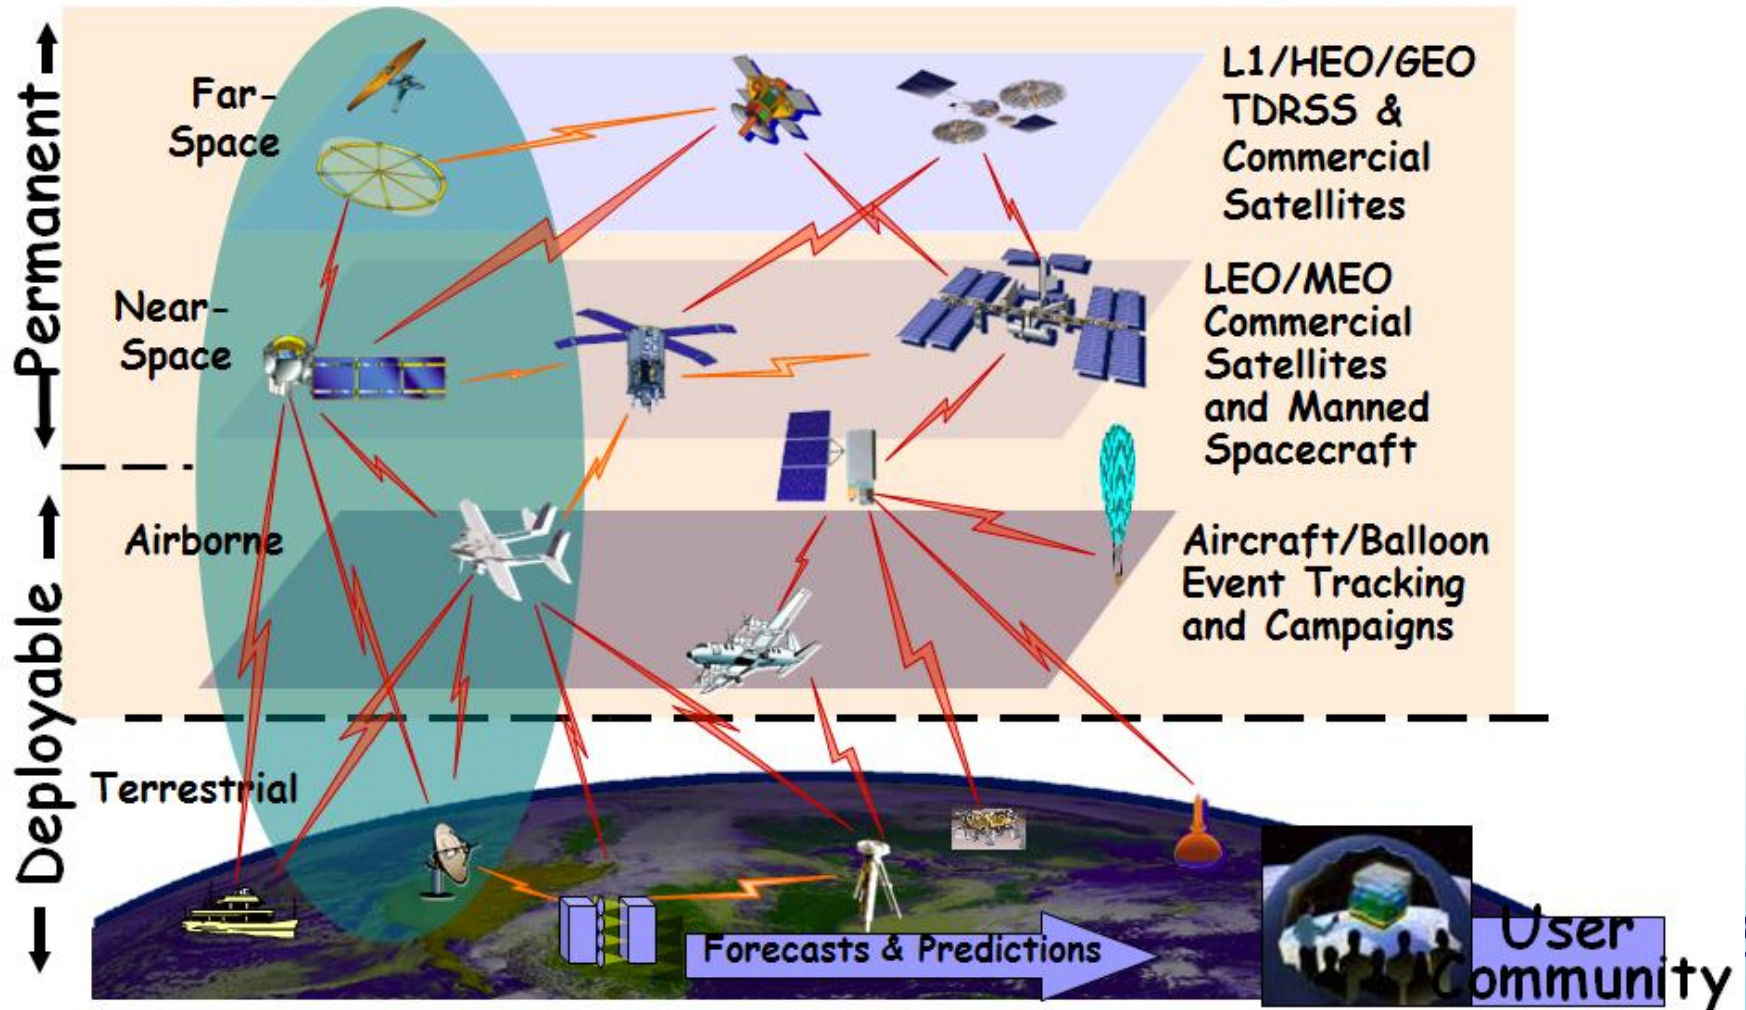




Case 5: Geospatial Data Cloud

Welcome to the Age of Data-intensive GIScience!



Vantage Points

Capabilities






Data-intensive GIS = principles and applications of geoinformatics for handling very large data sets



Challenges for data-intensive GIScience

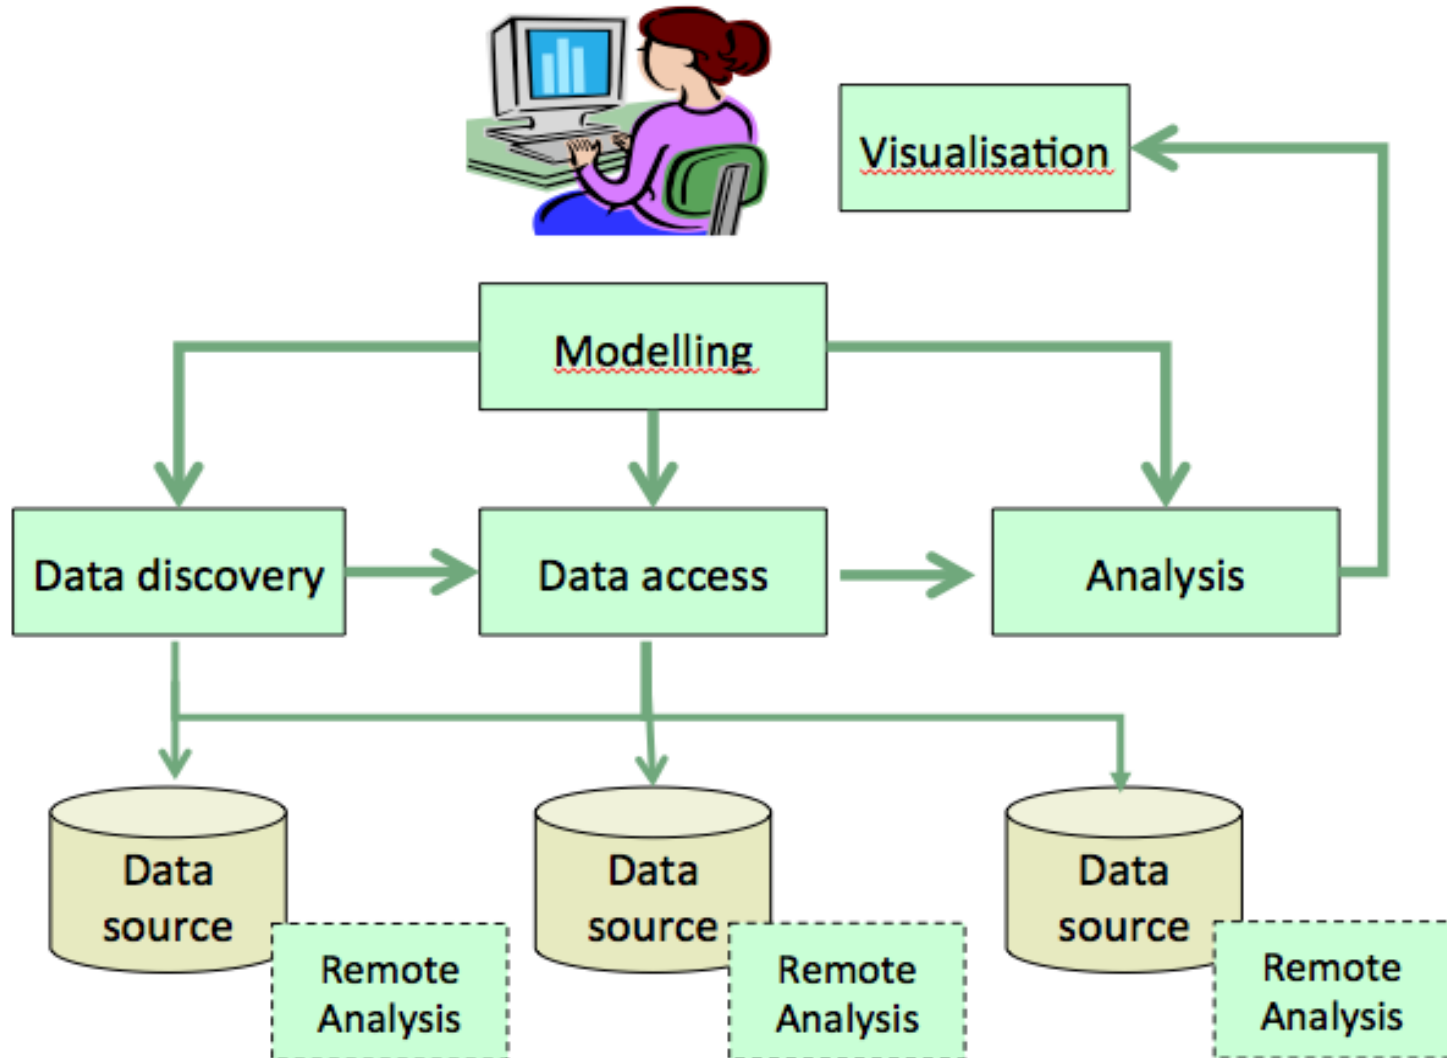


Which data is out there?
How to organize big spatial data?
How to get the data I need?

How to model big data?
How to access and use big data?



GIS technology for big data





Geospatial Data Cloud

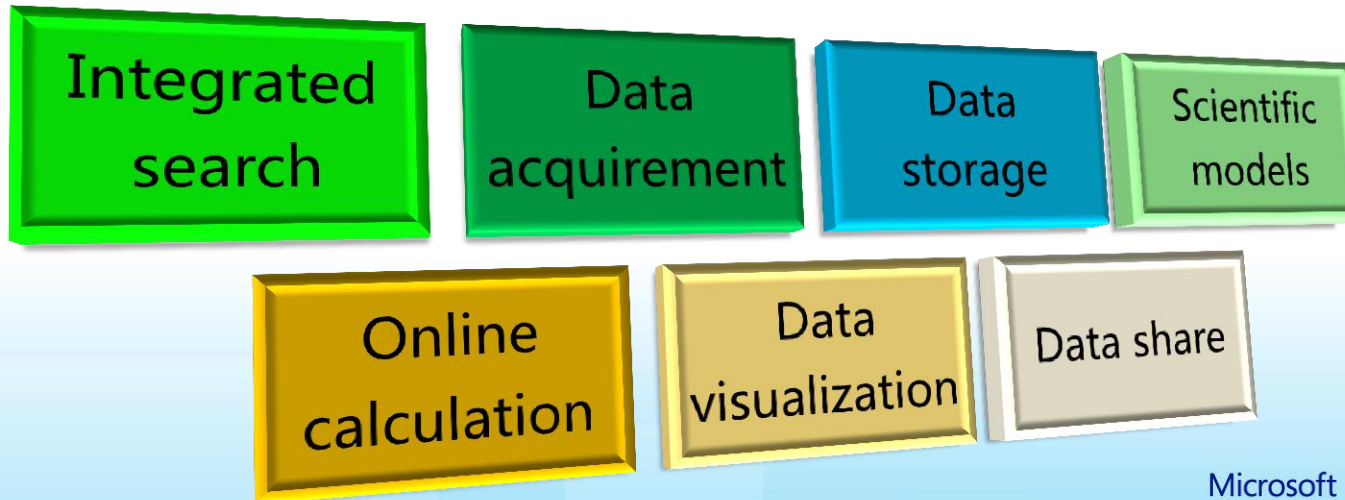
- The management, processing and utilization of these mass data is a big and difficult challenge for scientists
 - Distributed: The RS data was stored in different data centers around the world
 - Spatial-temporal: Scientists want to search their focused data quickly by spatial and temporal characteristics
 - Massive data entities, Petabytes or more : A great challenge to view, store, move, share these massive data for scientists

To solve these problems, the recent cloud computing and cloud storage technology can meet these requirements



Geospatial Data Cloud--Objectives

- Based on *CASDC*, Constructing a highly scalable Geospatial Data Cloud for the Geographical research, Remote Sensing application and related research
- To meet the requirements of:





Geospatial Data Cloud-- **Key technologies**

- Aggregating open data resources
 - Crawling metadata and cache data entities from the Internet
 - Reorganizing raw data and push them to the cloud database
- Providing an user-friendly search engine
 - Users can find their favorite data records easily and quickly
 - Online open data entity, download directly
 - Offline open data entity, submitting application form. Crawling tools fetch and cache data from Internet, then provide for downloading
 - Commercial data entity, providing quick link to the data provider



Geospatial Data Cloud-- **Key technologies**

- Integrating typical analytical models and GIS tools:
 - Encapsulating models as services in the cloud computing environment
 - Providing online calculation for scientists: NDVI, EVI, Land Surface temperature, projection, clip, resample, format conversion et al.
- Providing private cloud storage for scientists
 - Free, secured, and permanent cloud storage
 - Users can manage their private data
 - Permanent store model results
 - Share data between users



Geospatial Data Cloud-- **Key technologies**

- Efficient mass data management strategy
 - All earth observation data stored in the parallel cloud databases
 - 💡 System Architecture: postgresql + postgis + pgpool
 - 💡 The data table is divided properly, and dispersed on database nodes
 - Establish attribute indexes and geospatial indexes enables fast retrieval become possible



Geospatial Data Cloud-- **Key technologies**

- Web crawling technique for geospatial data
 - Developing web crawling tools to fetch data from Internet
 - Rearranging the raw metadata, push to the cloud database for search service
 - Automatically fetch metadata & data entities from NASA and USGS web sites



Geospatial Data Cloud-- **Key technologies**

- Online geospatial model calculation
 - Analysis models and GIS tools were wrap as web services
 - Providing online parameter pages for users to submit tasks
 - Users can use private data and platform data for model calculation
 - All user tasks were dispatched by task scheduler asynchronously




Geospatial Data Cloud-- **Key technologies**

- Data localized task scheduling
 - Big Data File System
 - 💡 A smart distributed file system for big data computation
 - 💡 High scalable, dynamically expanding storage space
 - 💡 Smart Data Replication
 - 💡 Automatically balance data distribution when a computer node breaks down



Geospatial Data Cloud-- **Key technologies**

- Data localized task scheduling (cont.)
 - File Allocation Strategy
 - 💡 A file has at least 2 copies, the file stored as data chunks, the first copy distributed in a node, and other copies distributed in other nodes
 - 💡 With this strategy, high read speed up to 800MB/s, VS single disk only up to 60MB/s





Geospatial Data Cloud-- **Key technologies**

- Data localized task scheduling (cont.)
 - Task Scheduling Strategy
 - ✦ Goal: Minimized the data transmission through network
 - ✦ Data localized task scheduler:
 - ✦ First, find the node location of data chunks of all model input files
 - ✦ Then dispatch the computation task to the corresponding node



Geospatial Data Cloud-- Online data resources


- Original Data from Internet: 76 datasets
 - ✎ Several data resources. Including LANDSAT, MODIS, MODIS_L1B, EO-1, DEM, NCAR, AVHRR.....
 - ✎ The amount of data reached **220 TB**, the number of metadata reached to 6 million
- Processed data products: 40 datasets
 - ✎ Including Land Surface Temperature, NDVI, EVI of 5 days, 10 days, 15 days and 30 days ...
- All these open data are providing online services for free
- All data covered China and the surrounding region




Geospatial Data Cloud-- Online data resources

- Data Products produced by CNIC


- ✓ The maximum value of Land Surface Temperature, NDVI, EVI in 5 days, 10 days, 15 days and 30 days, from 2000 – 2010
- ✓ Calculated from original modis hdf files



The first ten days of July 2010,LST



The middle ten days of July 2010,LST



The last ten days of July 2010,LST



Geospatial Data Cloud- System Introduction

- Homepage(www.gscloud.cn)

地理空间 Geospatial Data Cloud

网站首页 高级检索 数据目录 模型服务 平台信息

登录 | 注册 我的数据云

当前位置: 首页 -> 检索结果

共 1113 条记录, 已选择 0 条

数据	缩略图	操作
有		更多
有		更多
无		更多
有		更多
有		更多
无		更多
有		更多
无		更多
有		更多
无		更多
有		更多
无		更多

坐标: 75.73870, 38.58110

7 LE71480342003144ASN00 148 34 2003-05-24 77.13 37.45 53 无

8 LE71480332003144ASN00 148 33 2003-05-24 77.56 38.88 26 无

联系电话: 010-58812534 E-Mail: data@cnic.cn

Search data



Outline

- What is Cloud Computing
- A case-study of Scientific Research of the Cloud
- Conclusion



Conclusion

- The use of the cloud provides a number of opportunities
 - ✓ It enables services to be used without any understanding of their infrastructure.
 - ✓ Cloud computing works using economies of scale:
 - ✓ Data and services are stored remotely but accessible from “anywhere”.



Conclusion

- In parallel there has been backlash against cloud computing:
 - ✓ Use of cloud computing means dependence on others and that could possibly limit flexibility and innovation:
 - The others are likely become the bigger Internet companies like Google and IBM, who may monopolise the market.
 - Some argue that this use of supercomputers is a return to the time of mainframe computing that the PC was a reaction against.
 - ✓ Security could prove to be a big issue:
 - It is still unclear how safe out-sourced data is and when using these services ownership of data is not always clear.
 - ✓ There are also issues relating to policy and access:
 - If your data is stored abroad whose policy do you adhere to?
 - What happens if the remote server goes down?
 - How will you then access files?
 - There have been cases of users being locked out of accounts and losing access to data.



Conclusion

- should be able to run a variety of scientific applications on the cloud
- Microsoft and CNIC will issue a **statement about an academic collaboration**
 - ✓ You can fill in a simple proposal paper and send it to us. We will collect the proposal and do some associated review work with Microsoft, we can provide Window Azure resources to users who can offer their stated requirement

The End
Thanks !

Please contract me:
Yuanchun Zhou(zyc@cnic.cn)