Science & Technology Facilities Council

Quality Aspects in CLOUD Computing

Keith G Jeffery
Science and Technology Facilities Council
Rutherford Appleton Laboratory, Harwell Oxford
OX11 0QX UK
e-mail: keith.jeffery@stfc.ac.uk

STRUCTURE

- **Introduction – Who?**
- The Pervasiveness of ICT
- CLOUD Computing – Hype and Characterisation
- CLOUD Computing – Challenges
- Underlying ICT Challenges
- Conclusion

Rutherford Appleton Laboratory

Science & Technology Facilities Council

Brief Biography

Science & Technology Facilities Council

- 1968 Degree Geology
- 1969 First relational database in UK (G-EXEC)
- 1971 PhD Geology (using much computing)
- 1971: first relational data system in Europe
- 1976 first demonstration of heterogeneous distributed geoscience database systems interoperating all on one machine – Filematch standard
- 1978 distributed environmental information system
- 1983-89 First heterogeneous distributed system for international research information
- 1991 initial interoperation standard for research information – CERIF91
- 1992- led team: first webserver outside CERN
  - Initial GUI, Initial CGI interface
  - SVG, SMIL, SKOS
- 1999 second (better) interoperation standard for research information – CERIF2000
- 1999 paper proposing the UK e-Science programme
- 2000 Institutional Repository at CLRC
- 2004-8 EC GRIDs expert group leader
- 2009- EC CLOUDs expert group leader
- 2009-10 design authority for RCUK SSC
- 2010 CERIF adopted as standard for UK universities and funders

- University of Exeter
- British Geological Survey: team leader
- Natural Environment Research Council team leader
- Science and Engineering Research Council - team leader
- Head of Research Division
- Director, IT
- Director IT & International Strategy

Science & Technology Facilities Council

Associations

ERCIM
European Research Consortium for Informatics and Mathematics

W3C®

Very Large Data Bases

euroCRIS
Current Research Information Systems

APA
ALLIANCE FOR PERMANENT ACCESS

EDBT

EOS

ERF
European Association of National Research Facilities
Open to International Access

=: SOFSEM :=

EC CLOUDs Expert Group

Report 1: Work 2010,
Event January 2011

http://cordis.europa.eu/fp7/ict/ssai/docs/cloud-report-final.pdf

Report 2: Work 2011,
Event May 2012

Science & Technology Facilities Council

So?

This background gives you some idea of 'where I'm coming from'

- Advanced research problems requiring ICT solutions
- Research not 'blue sky' but practical (yet leading edge)
- International working – consultancy, reviewing, expert
- Strategic thinking for / using blue sky research to plan roadmaps for ICT R&D
- Design authority for large industrial-scale projects

Science & Technology Facilities Council

So?

This background gives you some idea of ‘where I’m coming from’

- Advanced research problems requiring ICT solutions
- Research not ‘blue sky’ but practical
- International working
- Strategic thinking for / using blue sky research to plan roadmaps for ICT R&D

And what I am going to talk about is the ICT of the future that we shall all be using and/or developing

And the research challenges we have to overcome to make it happen **with quality**

Science & Technology Facilities Council

STRUCTURE

- Introduction – Who?
- **The Pervasiveness of ICT**
- CLOUD Computing – Hype and Characterisation
- CLOUD Computing – Challenges
- Underlying ICT Challenges
- Conclusion

Internet Users by Region

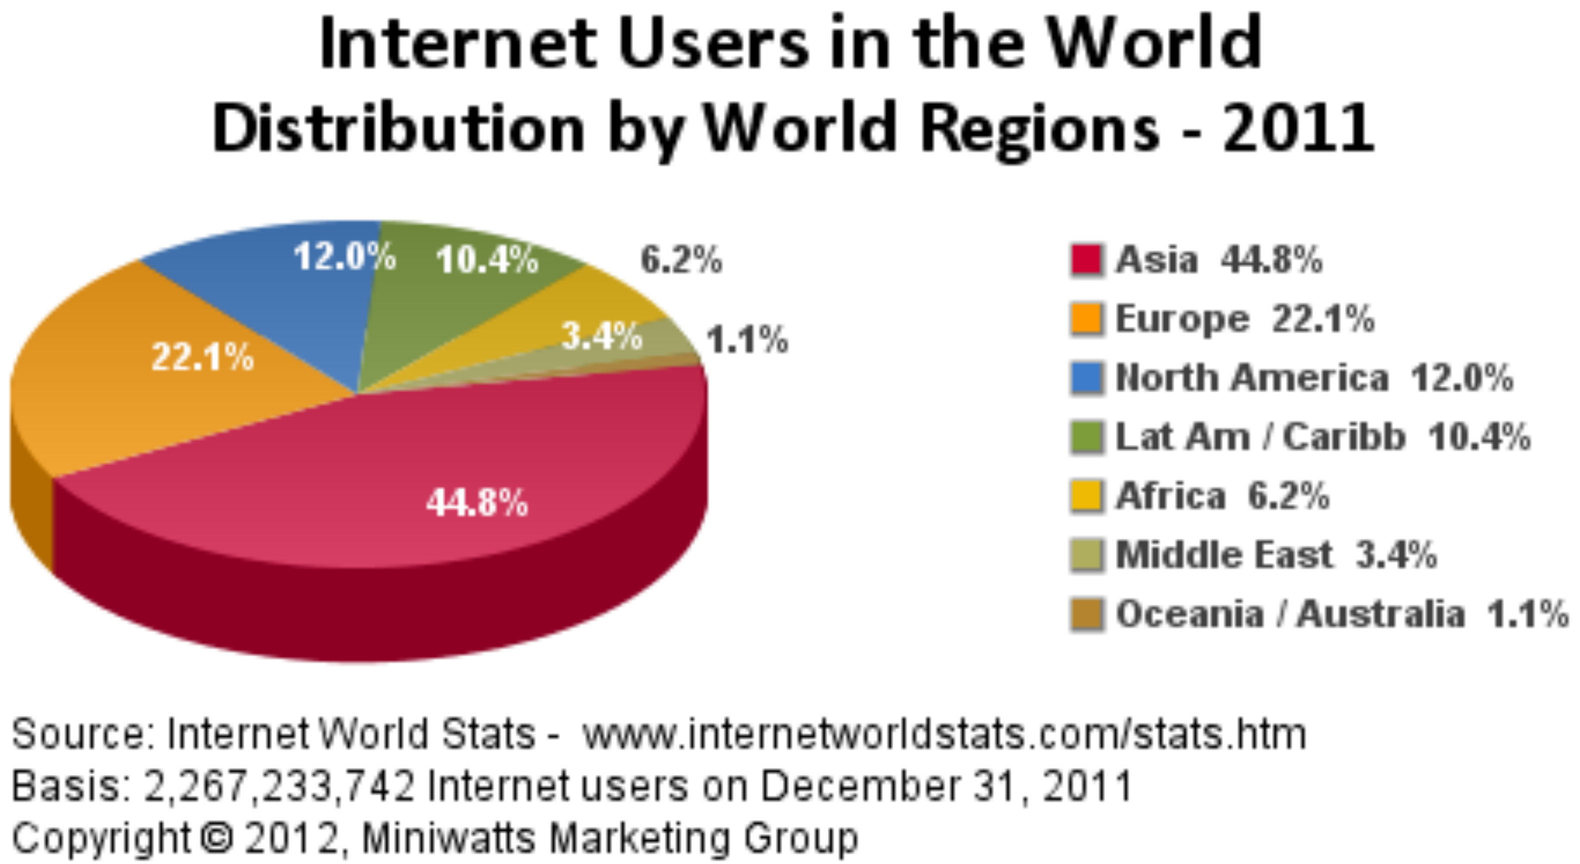

The ASIA Timebomb

Science & Technology Facilities Council

ASIA has largest population, largest number of users but relatively low penetration

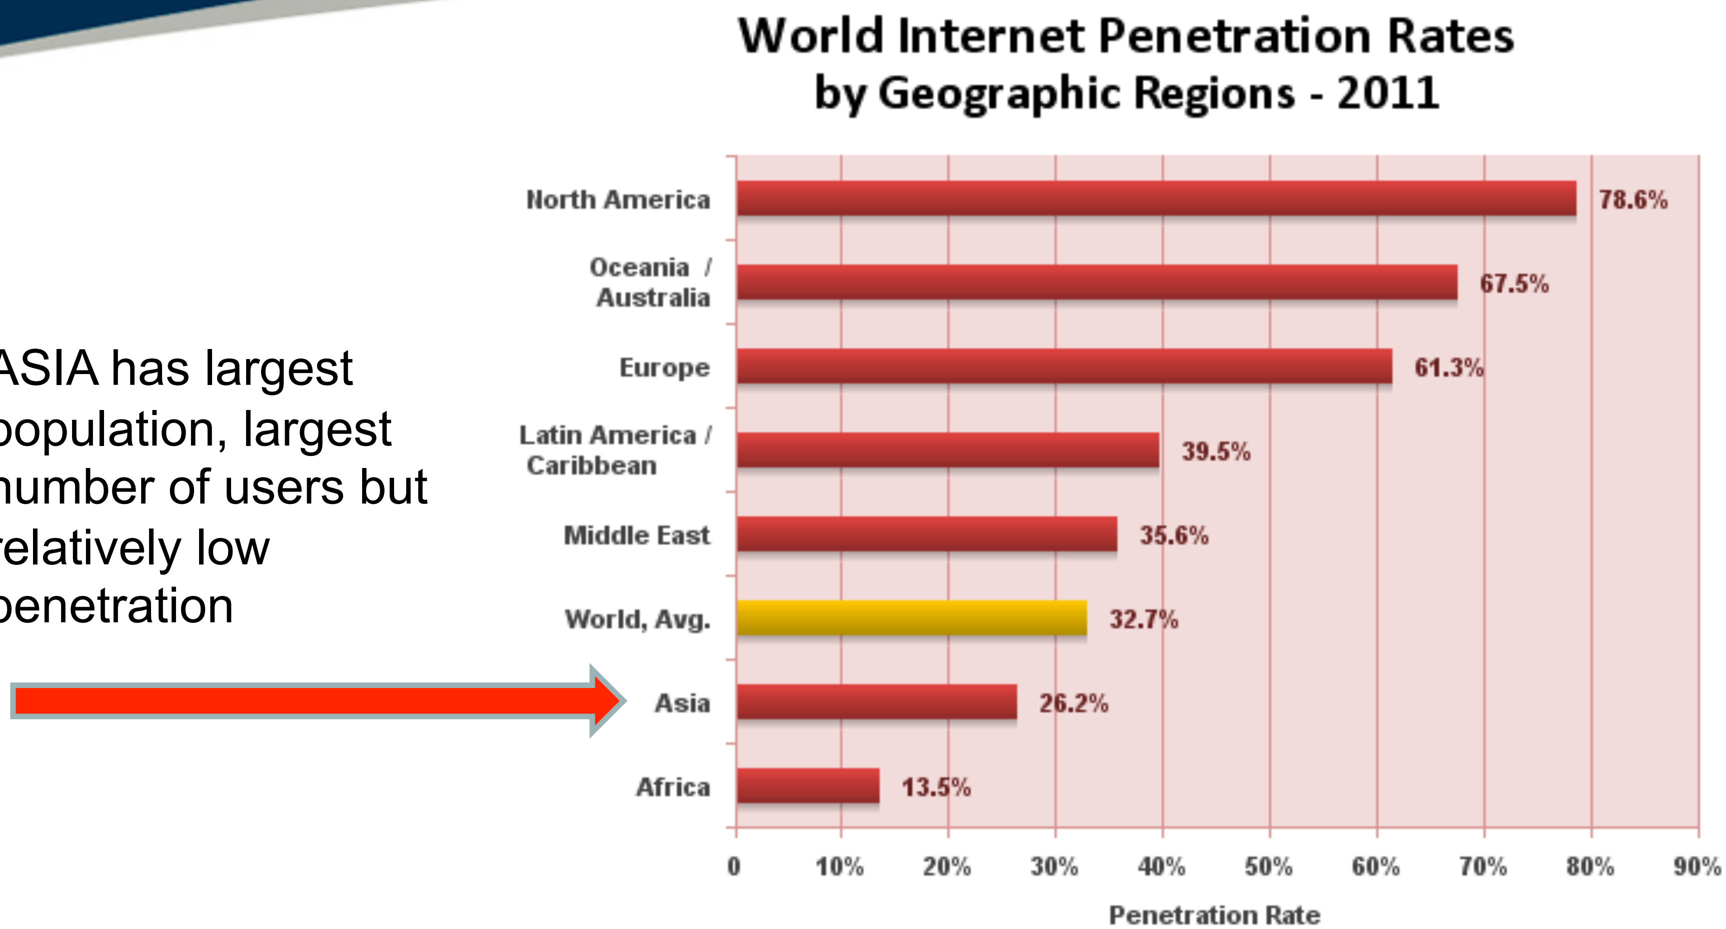

Source: Internet World Stats - www.internetworldststs.com/stats.htm
Penetration Rates are based on a world population of 6,930,055,154 and 2,267,233,742 estimated Internet users on December 31, 2011.

Issues

Elastic scalability
Trust & security & privacy
Manageability
Accessability
Useability
Representativity

Issues

Science & Technology Facilities Council

Elastic scalability
Trust & security & privacy
Manageability
Accessability
Useability
Representativity

QUALITY

Social Context

- The number of computers will vastly out-number humans on the planet very soon;
- Everything will be computerised;
  - Sensor networks
    - Home, healthcare, environment, industrial processes, transport systems….
  - Control systems
    - Industrial, transport, home (central heating)…

We rely on these systems

They must have QUALITY

So…

- This is the ‘internet of things’ or ‘future internet’
- We need to :
  - Manage the huge numbers, sizes
  - Integrate the different kinds of systems
  - Into one environment leading to human decision-making
    - whether managing a business, shopping, media choice, social interaction
- But there is a problem…in last 20 years
  - Data storage density increased ~10**18
  - Processor power increased ~10**15
  - BUT broadband capacity increased ~10**4

- This has implications for Information Systems Engineering!
- In fact the requirement and limitations challenge the very basis of traditional computer science / ICT

STRUCTURE

- Introduction – Who?
- The Pervasiveness of ICT
- **CLOUD Computing – Hype and Characterisation**
- CLOUD Computing – Challenges
- Underlying ICT Challenges
- Conclusion

Definition

cloud (klaʊd) *an elastic execution environment of resources involving multiple stakeholders and providing a metered service at multiple granularities for a specified level of quality (of service).*

From report EC Cloud Computing Expert Group January 2011

Gartner Hypecycle

Science & Technology Facilities Council

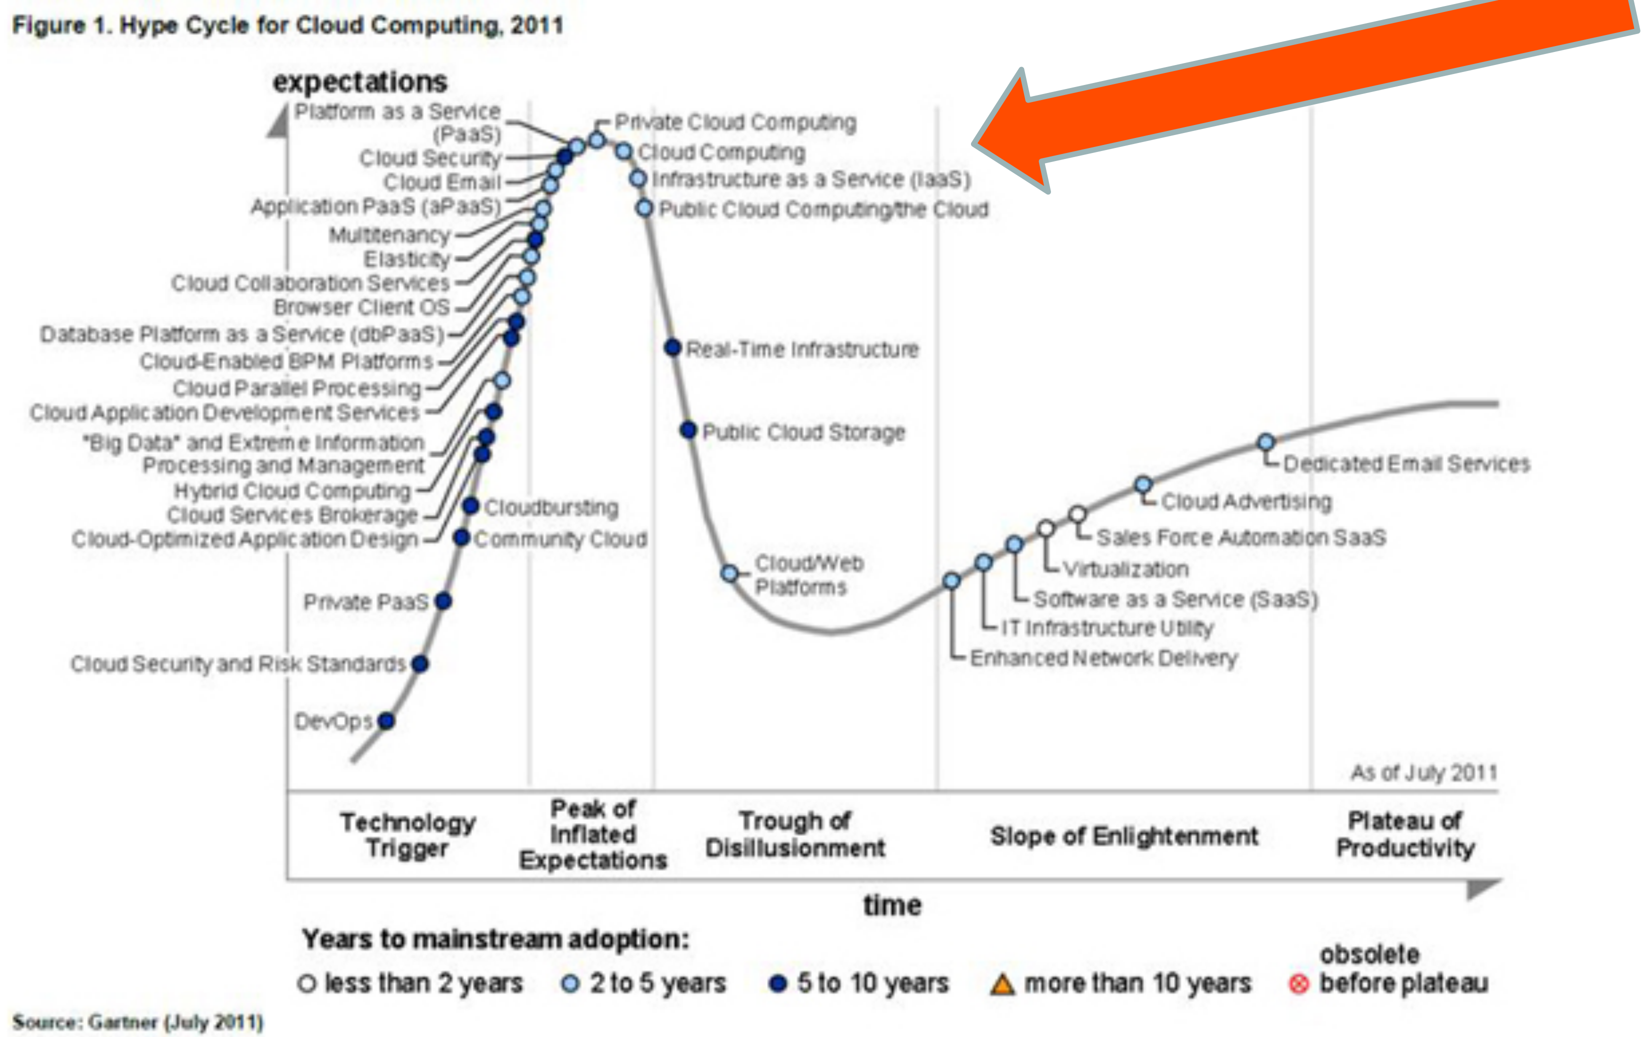

Hardware

- A very large number of processors
  - Clustered in racks as blades
- In one major computer centre
  - May be replicated for business continuity
- With massive online storage
  - RAID for resilience
- And excellent communications links
  - For access

Economies of scale – both purchasing and operation

Energy economies in location

Staffing economies in location

Customer View

- Low cost of entry for customers
- Device and location independence
- Capacity at reasonable cost (performance, space)
- Cloud Operator manages resource sharing balancing different peak loads
- Elastically scalable as demand rises (or falls) from user
- Security due to data centralisation and software centralisation
- Sustainable and environmentally friendly – concentrated power

- ➔ it is a service and the user does not know or care from where, by whom, and how it is provided
- ➔ as long as the SLA (service level agreement) QoS (**quality of service**) is satisfied
- ➔ it is a 'computing utility' (IaaS, PaaS, SaaS….'XaaS')

Cloud Computing

Ownership

- Private Cloud: in-house cluster run using CLOUD middleware;
- Public Cloud: outsourced computing to commercial provider – proprietary;
- Hybrid Cloud: linked Private and Public CLOUDs

Cloud Computing

Offerings

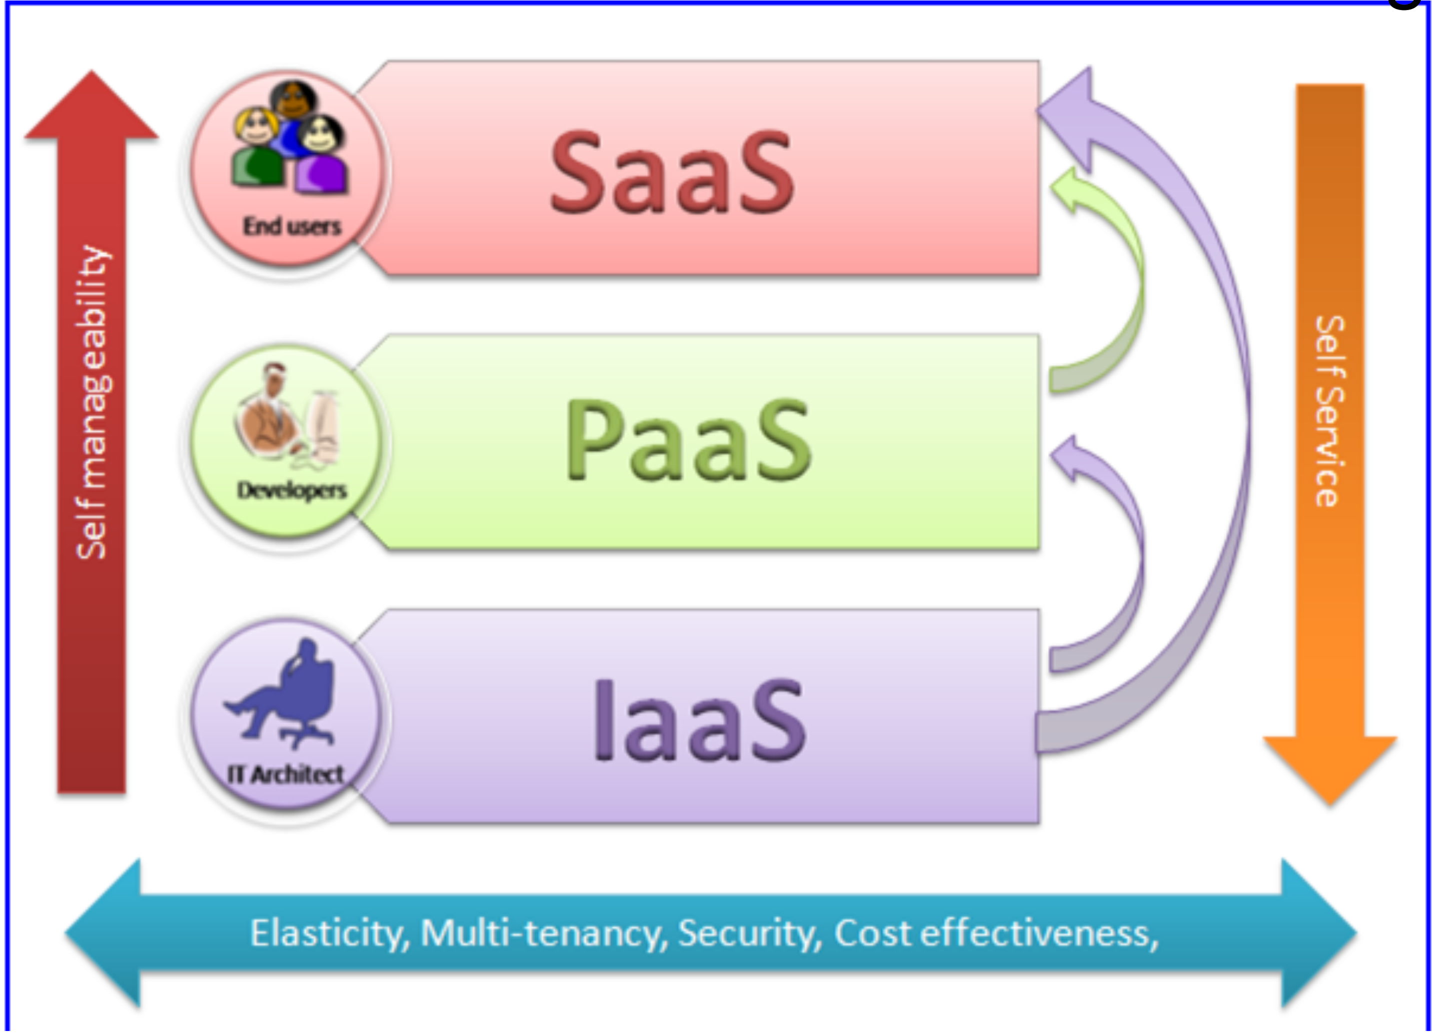

Acknowledgements to U Southampton

How does it work

- Multitenancy: Cloud resources (hardware) shared dynamically between customers;
- Each customer application in its own virtual machine
  - Isolation for security, privacy
  - Allows scheduling with respect to shared resources

Cloud Computing

- Obtains from GRIDs work:
  - resource sharing/scheduling
  - virtualisation of hardware and low-level software (under middleware)
  - resilience
  - trust, security, privacy
  - (more or less) self-*

Utility computing

Autonomic computing

Cloud Computing

Obtains from software/systems engineering:

Service-Oriented Architecture

with implications of interfaces, metadata, composition

But the real novelty is…

- Pay-As-You-Go (only for what you need);
- Accounting for ICT used by departments in an organisation;
  - Private cloud
  - Public cloud
- CAPEX to OPEX

Cloud Computing

Science & Technology Facilities Council

- Private Cloud Software : Eucalyptus
- Private/Hybrid Cloud software: Open Nebula
- Commercial examples of Public Clouds:
- Amazon EC2 Elastic Compute Cloud
- Google (Engine for Apps; Connect for Office)
- Microsoft Azure
- IBM SmartCLOUD
- (note all needed massive resource for infrequent use so could sell of excess capacity)

- Note Thomas J Watson in late fifties:

“total number of computers required in the world is five”

- are we reaching this goal?

Characterisation

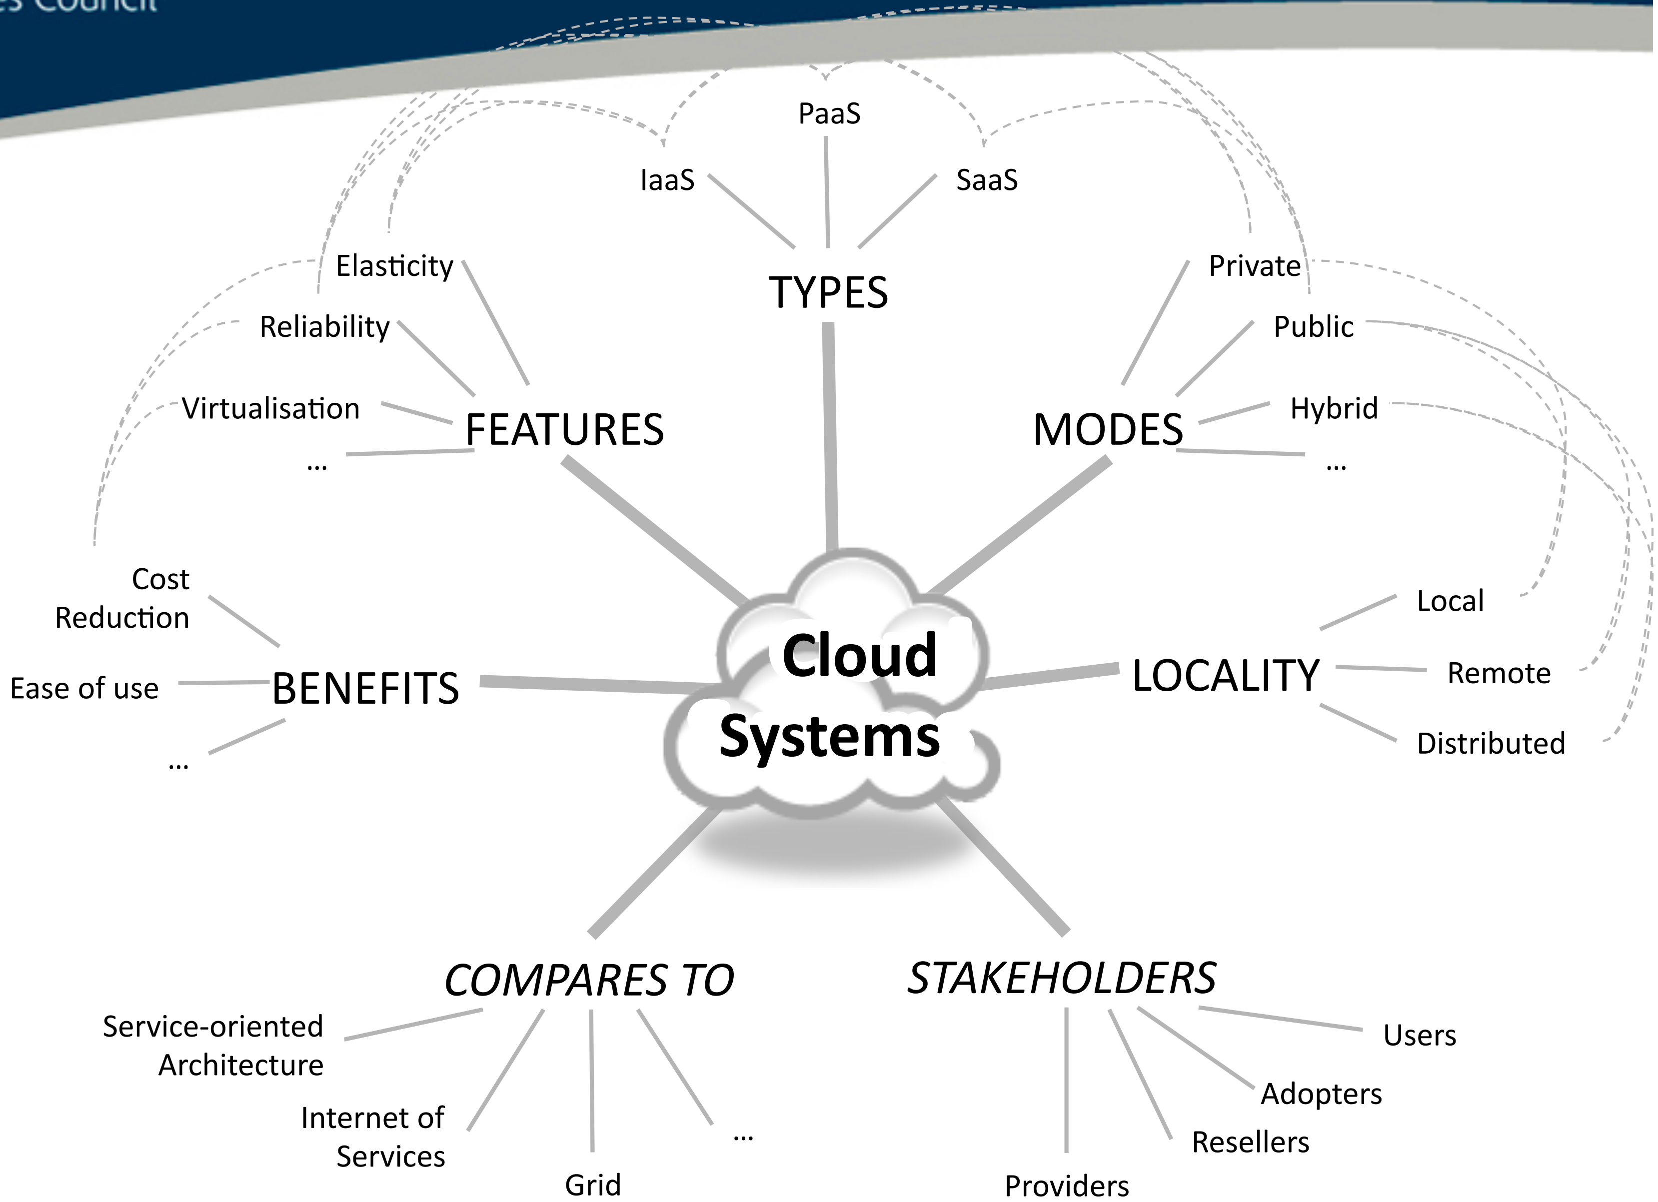

Science & Technology Facilities Council

Characteristics

Non-Functional	Economic/Legal	Technological
relate to qualities of cloud systems, rather than technological aspects. These include: •Elasticity •Reliability •Quality of service •Agility and adaptability •Availability	key driver behind (commerical) cloud systems. Typical interest rests on: •Cost reduction •Pay per use •Improved time to market •Return of investment •CAPEX to OPEX •"Going green" • business law • security • privacy	Arise from realising non-functional / economic concerns. Particular issues: •Virtualisation •Multi-tenancy •Security, privacy and compliance •Data management •APIs and / or programming enhancements •Metering •Tools in general

State of the Art

Science & Technology Facilities Council

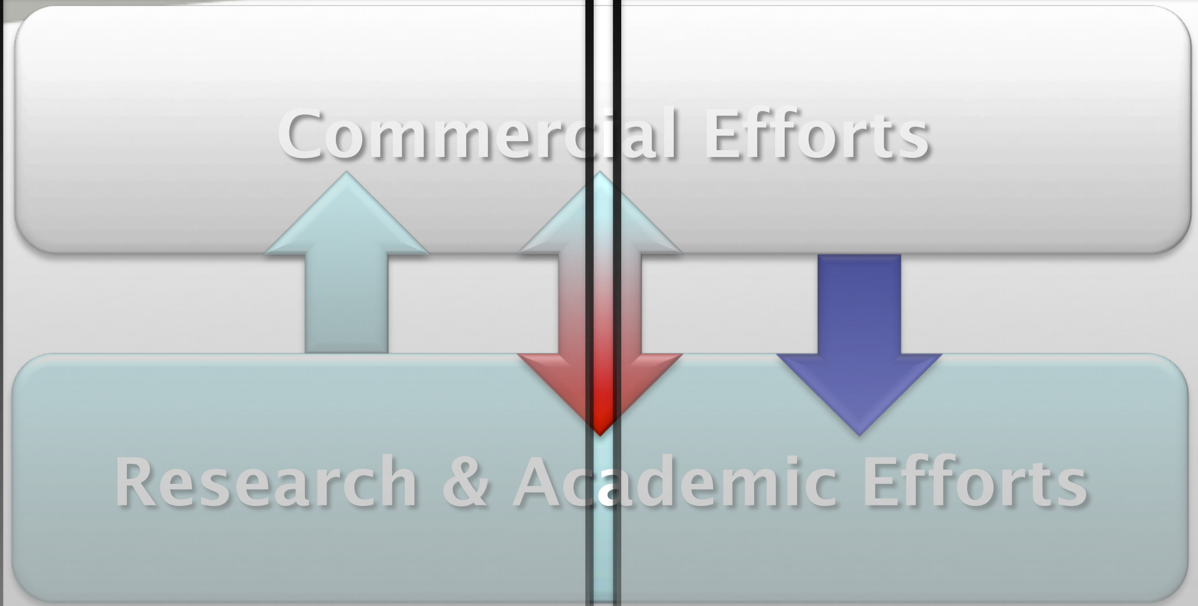

- Manageability and Self-*
- Data Management
- Privacy & Security
- Federation & Interoperability
- Virtualisation, Elasticity and Adaptability
- APIs, Programming Models & Resource Control

Technical Gaps

- Legislation, Government & Policies
- Economic Concerns

Non-Technical Gaps

Summary

So, CLOUD Computing is the bringing together of many pre-existing disparate ideas to form a novel offering.

All those disparate ideas had their own (and some common) challenges

Bringing them together creates complex combined challenges

STRUCTURE

- Introduction – Who?
- The Pervasiveness of ICT
- CLOUD Computing – Hype and Characterisation
- **CLOUD Computing – Challenges**
- Underlying ICT Challenges
- Conclusion

Cloud Studies

Above the Clouds: A Berkeley View of Cloud Computing

http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf

Science & Technology Facilities Council

Cloud Quotes

The interesting thing about Cloud Computing is that we've redefined Cloud Computing to include everything that we already do. . . . I don't understand what we would do differently in the light of Cloud Computing other than change the wording of some of our ads.

Larry Ellison, quoted in the Wall Street Journal, September 26, 2008

It's stupidity. It's worse than stupidity: it's a marketing hype campaign. Somebody is saying this is inevitable — and whenever you hear somebody saying that, it's very likely to be a set of businesses campaigning to make it true.

Richard Stallman, quoted in The Guardian, September 29, 2008

Research Issues

Business opportunities not currently realisable

Existing knowledge from preceding research and development can be harvested

Technical	Non-Technical
• Scale and elastic scalability • Trust, security and privacy • Data handling • Programming models and resource control • Systems development and management	• Economic aspects • Legalistic issues

But need to ensure quality

Target

Private CLOUD (inhouse, cluster)

interface

Common Service Environment with metadata and dynamic systems development and composition capability

interface	interface	interface
Public CLOUD	Public CLOUD	Other Private CLOUD

elasticity **e-business**

Quality

Science & Technology Facilities Council

Quality in business, engineering and manufacturing has a pragmatic interpretation as the *non-inferiority* or *superiority* of something; it is also defined as *fitness for purpose*. Quality is a perceptual, conditional and somewhat subjective attribute and may be understood differently by different people. Consumers may focus on the **specification quality** of a product/service, or how it compares to competitors in the marketplace. Producers might measure the **conformance quality**, or degree to which the product/service was produced correctly.

[Wikipedia]

Science & Technology Facilities Council

Quality in CLOUD Computing

- Quality is related to the satisfaction of **functional** and **non-functional** requirements.
- Key questions include how to:
  - **define** all these aspects of quality;
  - how to **measure** them;
  - how to **monitor** their value during the whole process from requirements to execution completion;
  - how to **utilise** the information generated to improve quality end-to-end.
  - ➔ implies self-* technology (autonomic computing)

Quality Challenges(1)

- **quality of the service provision** (quality of service, service level agreement management);
  - Trust, security, privacy
- **quality in modelling business requirements** in a Cloud environment;
  - Advanced data and process representations
- **quality in the languages used** to express those requirements and to execute them;
  - Improved declarative business languages
- **quality in the services** in both what they provide and how they provide it;
  - reliability
- **quality in service composition** including dynamicity;
- **quality in scheduling execution** including partitioning and parallelism;
  - Autonomic computing

Quality Challenges(2)

- **quality in monitoring execution** and triggering actions to maintain quality of service;
  - Autonomic computing
- **quality of business continuity** provision;
  - Automated backup/recovery and reassignment
- **quality of virtualisation** – including seamless execution transfers to assure timeliness;
  - Autonomic computing
- **quality** of the Cloud infrastructure and platform provision in terms of **'green ICT'**;
  - Environmental conscious design and execution monitoring

- Unlike other areas of ICT there is not yet for Clouds a quality of conformance since there is no agreed standard (unlike, for example, WWW).

Science & Technology Facilities Council

What Stops us..

- I have listed a whole lot of CLOUD specific problems already such as:
  - Interoperability / vendor lock-in
  - Security, privacy
  - Quality of service / service level agreements
  - Legislation
- But CLOUDs throws into sharp relief many underlying computer science / informatics problems
  - Which are required to solve the problems of CLOUDs
- And the solution to each must involve QUALITY

STRUCTURE

- Introduction – Who?
- The Pervasiveness of ICT
- CLOUD Computing – Hype and Characterisation
- CLOUD Computing – Challenges
- **Underlying ICT Challenges**
- Conclusion

Science & Technology Facilities Council

Research Challenges: 1

Metadata

- the need for metadata related to services, data/information/knowledge, resources;
- what is data, what is metadata?
- kinds of metadata and their use;
- representation and structure - syntax;
- semantics (meaning);
- Quality of representation of the objects

Research Challenges: 2

Management of state

- detection of state across millions of individual nodes;
- maintenance of state across many nodes;
  - transactions and locking;
  - roll-back and compensation;
- Quality of the ‘state picture’ compared with real world

Research Challenges: 3 (1)

Data representativity

- data structures representing real-world inter-relationships;
  - data attribute value encoding (character set, media encoding), types, lengths;
  - data attribute value language;
  - fully connected graphs – the death of the hierarchy;
  - the time-machine: temporal duration of the inter-relationships;
  - certainty, probability of the inter-relationships
  - Incomplete and inconsistent information

Research Challenges: 3 (2)

Data representativity

- Interoperation
  - – reconciliation of different data structures representing a similar real-world domain;
- data location / locality and replication
  - – for business continuity;
- Quality of representation of the real world

Science & Technology Facilities Council

Research Challenges: 4

Data quality, veracity and permanency

- detection of quality against metadata parameters e.g. precision, accuracy;
- provenance;
- temporal recording;
- data curation across media and policy evolution;

Research Challenges: 5

Trust, security and privacy

- policies declared, enforced and monitored through restrictive metadata;
- policy reconciliation for interoperation;
  - With quality assurance

Research Challenges: 6

Management of service levels and quality of service

- policies declared, enforced and monitored through restrictive metadata;
- service level negotiation (e.g. lower price for lower performance);
  - With quality assurance

Research Challenges: 7

Systems design, development, maintenance and decommissioning

- based on strong separation of:
  - services (processes),
  - data, information and knowledge
- assuming self-composition, self-managing and adjusting, self-maintaining properties;
- assuming mobile code properties
- Quality of the delivered product

A final challenge?

- Is the von Neumann architecture still valid?
  - Should we not optimise communications over other priorities?
- Do we need to write programs?
  - Should we not just compose (dynamically – software 'robots') from services as components (like other branches of engineering)?
- Will social / legal changes ever catch up with technology?
  - Can technology propose them?

A final challenge?

- Is the von Neumann architecture still valid?
  - Should we not optimise communications over other priorities?
- Do we need to write programs?
  - Should we not just compose (dynamically – software ’robots’) from services as components (like other branches of engineering)?
- Will social / legal changes ever catch up with technology?
  - Can technology propose them?

To improve QUALITY

STRUCTURE

- Introduction – Who?
- The Pervasiveness of ICT
- CLOUD Computing – Hype and Characterisation
- CLOUD Computing – Challenges
- Underlying ICT Challenges
- **Conclusion**

Conclusion

Being ever optimistic, I believe these challenges will be met – with QUALITY

But has some machine passed the Turing test and nobody noticed?

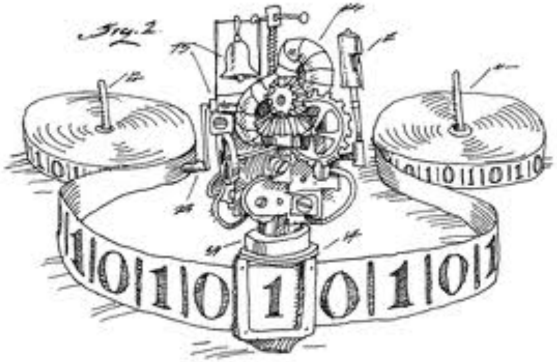

(2012 is the year commemorating Alan Turing birth centenary)

Prof. Keith G Jeffery
CEng, CITP, FBCS, FGS, HFICS
Director, International Relations

STFC Rutherford Appleton Laboratory
keith.jeffery@stfc.ac.uk

Acknowledgements to Lutz Schubert, HLRS, Stuttgart; rapporteur EC CLOUDs expert Group