

Cloud Computing

UCD IT Services Experience



Background - UCD IT Services



Central IT provider for University College Dublin

- 23,000 Full Time Students
- 7,000 Researchers
- 5,000 Staff



Background - UCD IT Services



- Diverse Locations
 - Belfield (132 Ha Site)
 - Michael Smurfit School of Business, Blackrock
 - Lyons Estate, Newcastle
 - Teaching Hospitals

Background - UCD IT Services



- Diverse disciplines:
 - Arts & Celtic Studies
 - Business & Law
 - Engineering, Mathematical & Physical Sciences
 - Human Sciences
 - Life Sciences



Research IT Services



- Set up in late 2005
- Small, agile team of five
 - Close working relationships with researchers


“to support and facilitate researchers in their use of IT as an enabler for their research activities”

Research IT Services Strategy

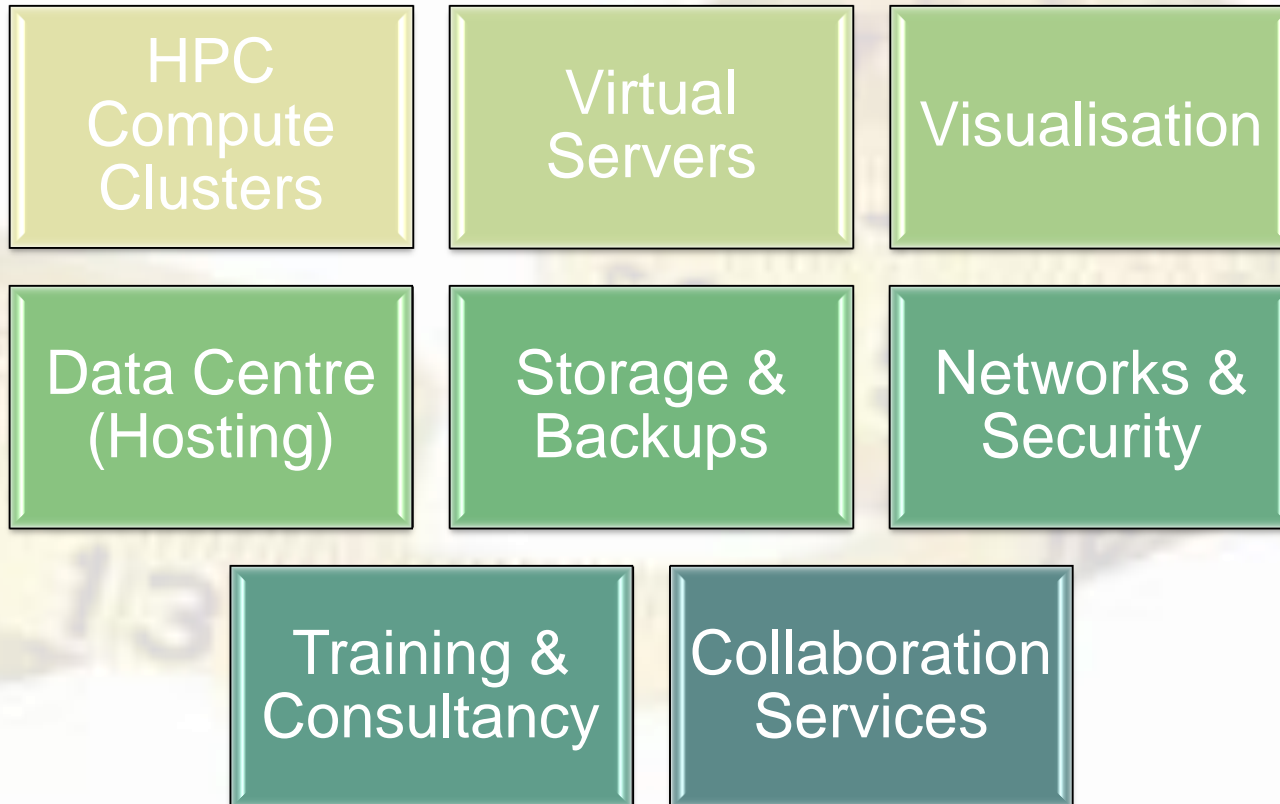


- Requirements report
- Research IT Steering Group
- Researcher Survey
- Customer Stated Requirements





Research IT Services



Maintaining Services



- Hardware Maintenance
- Software Maintenance
- 3 – 5 Year Cycle

What happens when....?

The Case for Cloud:

Research IT HPC Facilities



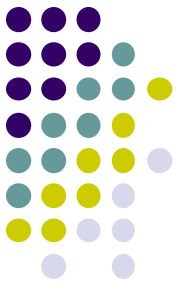
- Shared Service
 - Free
 - Familiarisation
- Community Service
 - Contribution
 - Additional resources, increased memory
- IBM Blade Technology



Real Issues

- Capital intensive process
- Physical HPC Cluster is inelastic
 - Users sometimes need more resources than are available
 - Adding resources can be a slow process
- Physical space at a premium


Solutions



- Pre-order hardware
 - Capital Intensive
 - Power
 - Cooling
 - Physical space




Solutions



- Buy compute time on other facilities
 - Can be Expensive
 - Non-dynamic
 - National HPC Service
 - ICHEC Stokes
- Investigate “Cloud”





Moving to Cloud computing is not all about the technology, it is a behavioural rather than technological change.




Barriers to Adoption

- Organisational
- Behavioural
- Technical



Barriers to Adoption:


Organisational & Behavioural



- Culture gives people
 - Stability
 - Security
 - Understanding
 - Ability to respond to a given situation
- Change causes people to fear that all those things will be lost to them

Barriers to Adoption:

Technical



- Technical
 - Shortage of suitable vendors
 - Established vendor inflexibility
 - Vendor Lock In
 - Data Security & Privacy

Barriers to Adoption:

Shortage of suitable vendors




- Vendor Identification Criteria
 - Infrastructure as a Service
 - Hardware within the EU
 - Current use based pricing model
 - Invoice based billing



Barriers to Adoption:


Shortage of suitable vendors



- Vendor Identification Criteria (cont.)
 - Customised / Customisable systems
 - Connect via VPN technology
 - Test / Free period of use for Proof of Concept



UCD Cloud



Main purpose to enable HPC Cloud Bursting


- Phaeton HPC Cluster

- Physically located in Research IT Data Centre
- IBM Intel hardware
- Red Hat Enterprise Linux 5.

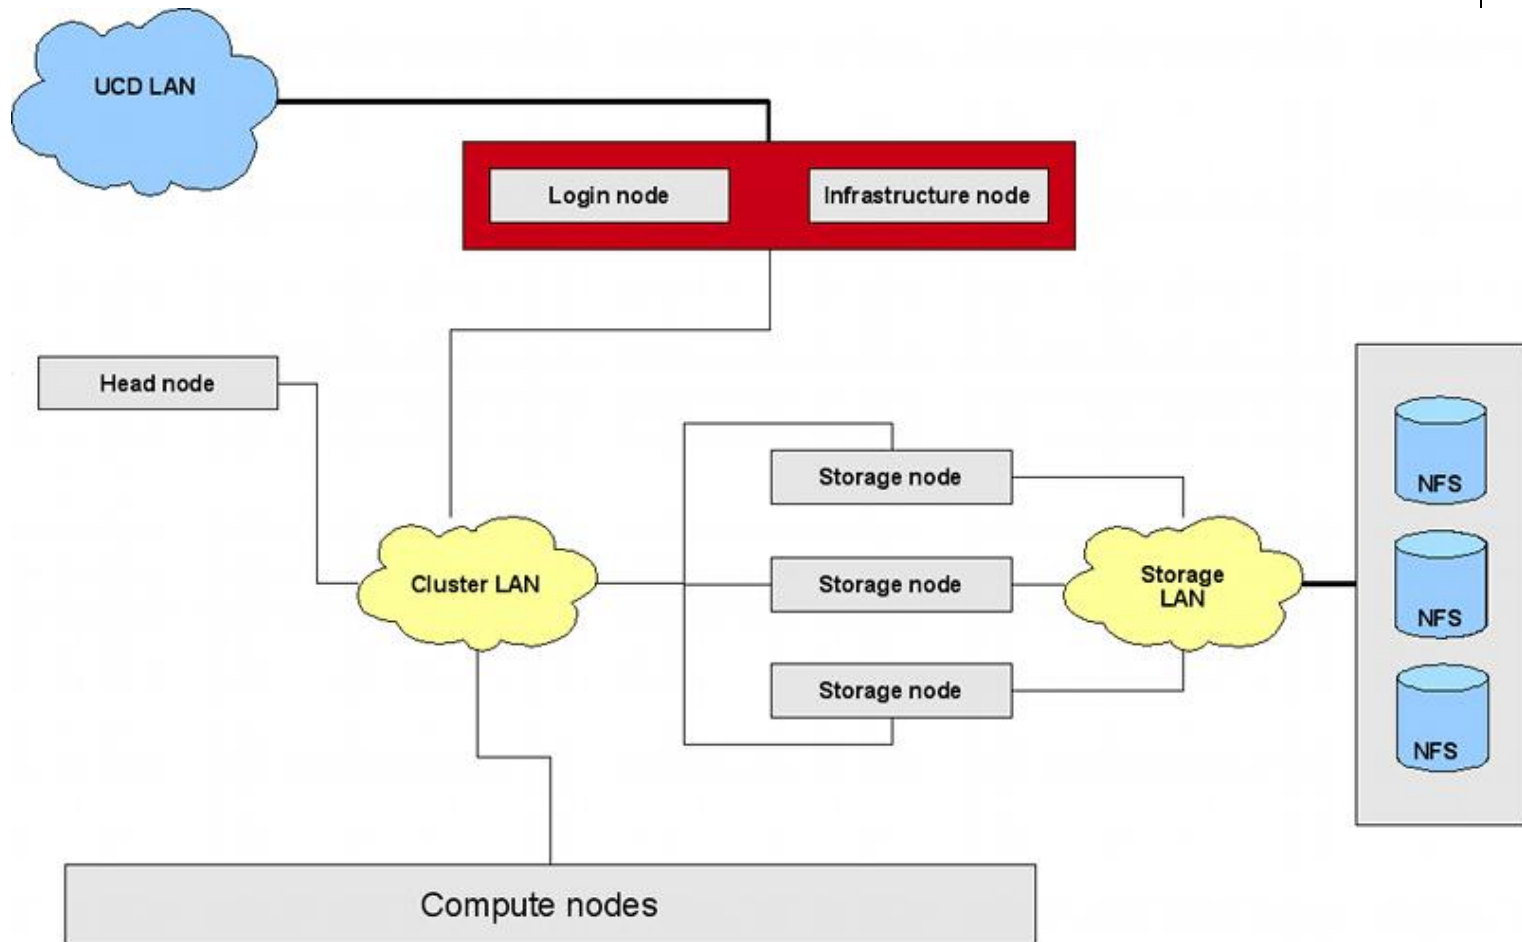
- Worker Node Hardware

- *IBM HS21 XM Dual Xeon Quad Core E5430 Blade Servers with 32GB RAM*
- *HS22 Quad Quad Core Blade Servers with 64GB RAM*





Phaeton HPC Cluster Architecture



UCD Hybrid Cloud



- Private Cloud
 - Ubuntu Enterprise Cloud / Eucalyptus (KVM)
 - HP hardware
 - VM Ware vSphere (ESX)
 - HP Blade Servers


UCD Hybrid Cloud



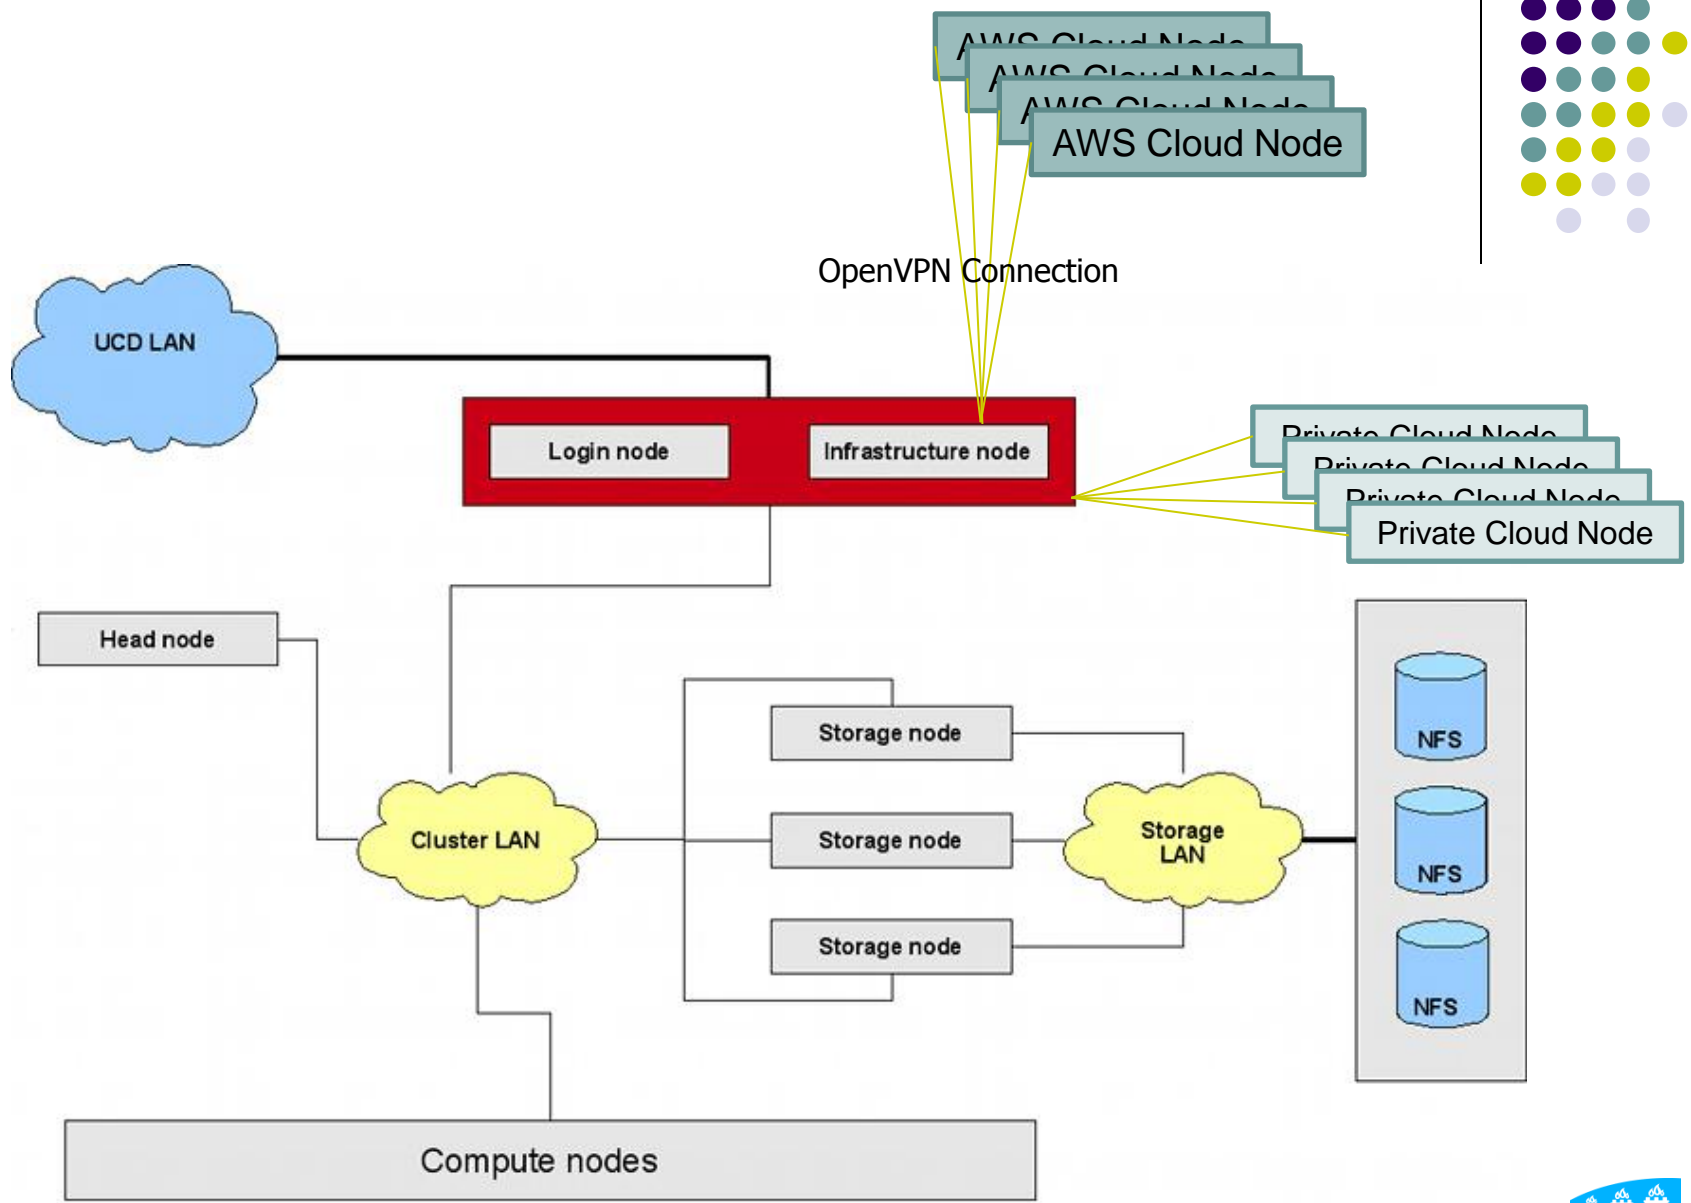
- Public Cloud
 - AWS
 - EU West & East locations (Dublin)
 - Limit of 128 nodes currently
 - UCD CentOS image
 - AWS Large Instance




UCD Hybrid Cloud



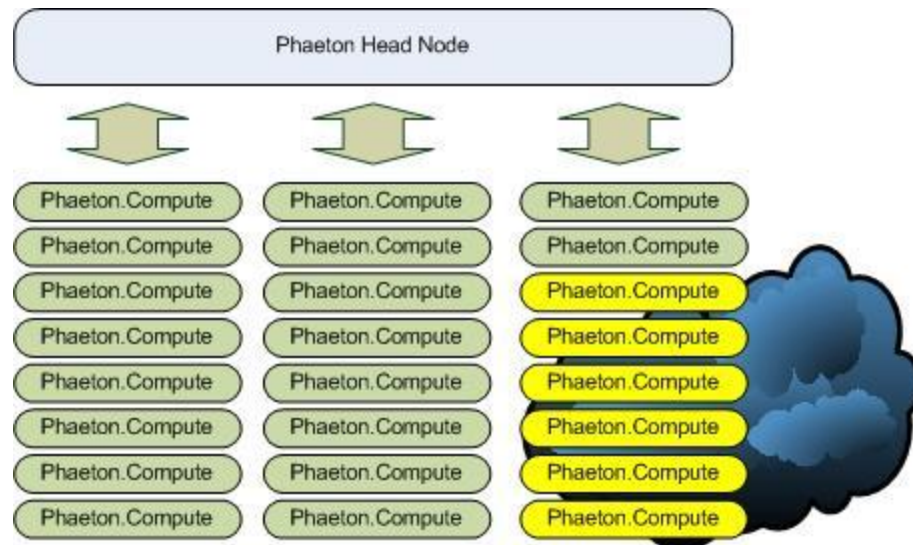
- Other Vendors
 - Fort (*Dublin*)
 - PoC completed
 - IBM (*Germany, UK*)
 - PoC commencing
 - OpSource (*US*)
 - PoC completed
 - Digiweb (*Dublin*)
 - PoC in progress






Cloud & Node Management

- All nodes are managed by Phaeton Head Node
 - Invisible to users
 - MOAB from Adaptive Computing / Torque
 - Specific queues for cloud based resources




Authentication

- Phaeton Login Node
 - LDAP Authentication
 - Federated Access (Edugate)
 - UCD Login ID
 - User's identity has already been established
 - Pre-assigned individual user SSH Keys to access nodes



Encrypted Communications


- Encrypted Communications
 - OpenVPN
 - OpenSSL
 - Performs and scales well in network-IO light workloads
 - Natively cross-platform
 - 128 bit AES encryption for the data link
 - 160 bit SHA1 checksum for both the data and control channels






Data Concerns

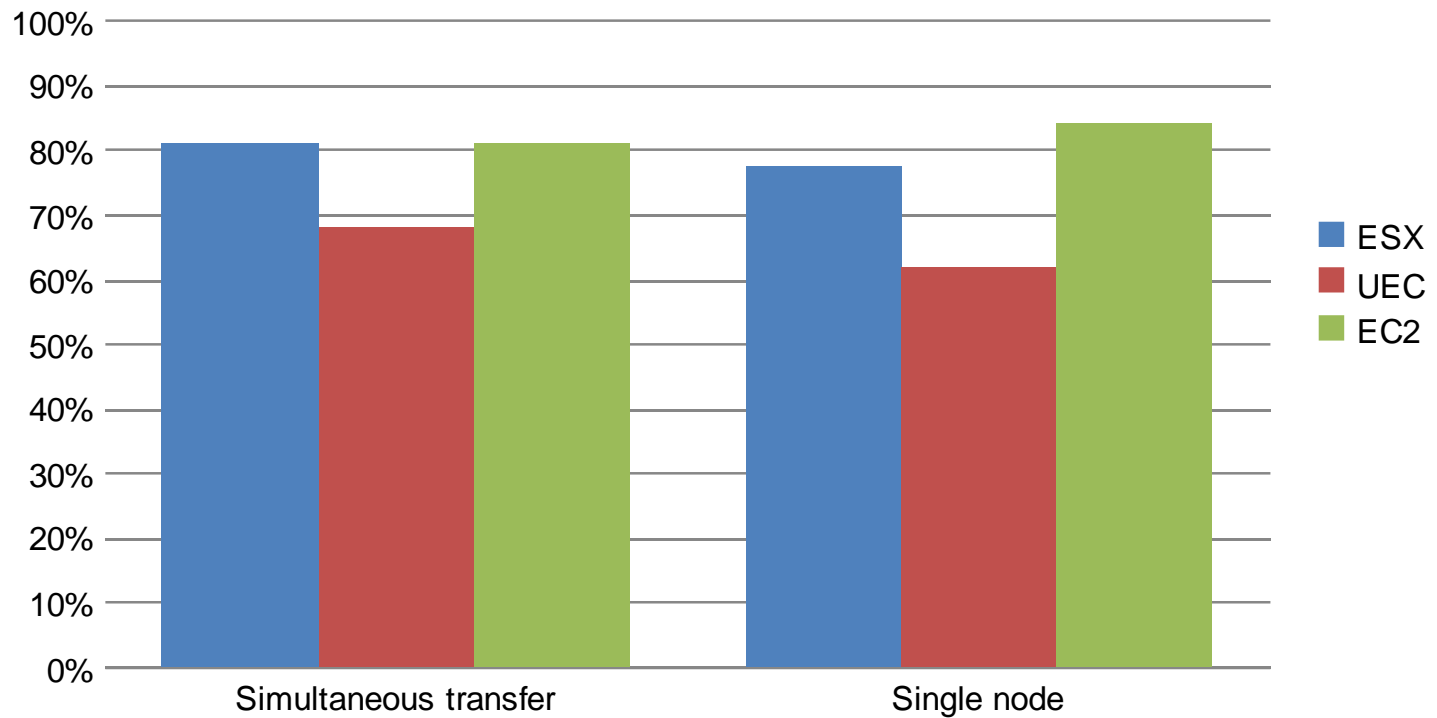
- Amazon gives choice of location
 - EU-WEST-1A & EU-WEST-1B
- Data not "stored" in the cloud
- Computation only in the cloud




Data Transfer Rates



NFS transfer
100MB of random data
relative performance [local node = 100%]



Experience so Far



- Positive overall experience
- HPC
 - Each cloud instance is a fully-capable HPC worker node
 - “Cloud” is transparent to users
 - Proven viable solution

Successful Research Output




- *First Major Project Output:*
 - *“The enumeration of Costas arrays of order 28”*
 - *Seven years of CPU time reduced to three months real time*
 - *<http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=05592926>*

Cloud Future



- HEI Cloud Broker
- Removal of Traffic Costs to commercial solutions
- Academic & Research Cloud



Thank You

- Ruth.Lynch@ucd.ie, Fred.Clarke@ucd.ie
- <http://www.ucd.ie/itservices/researchit/>

