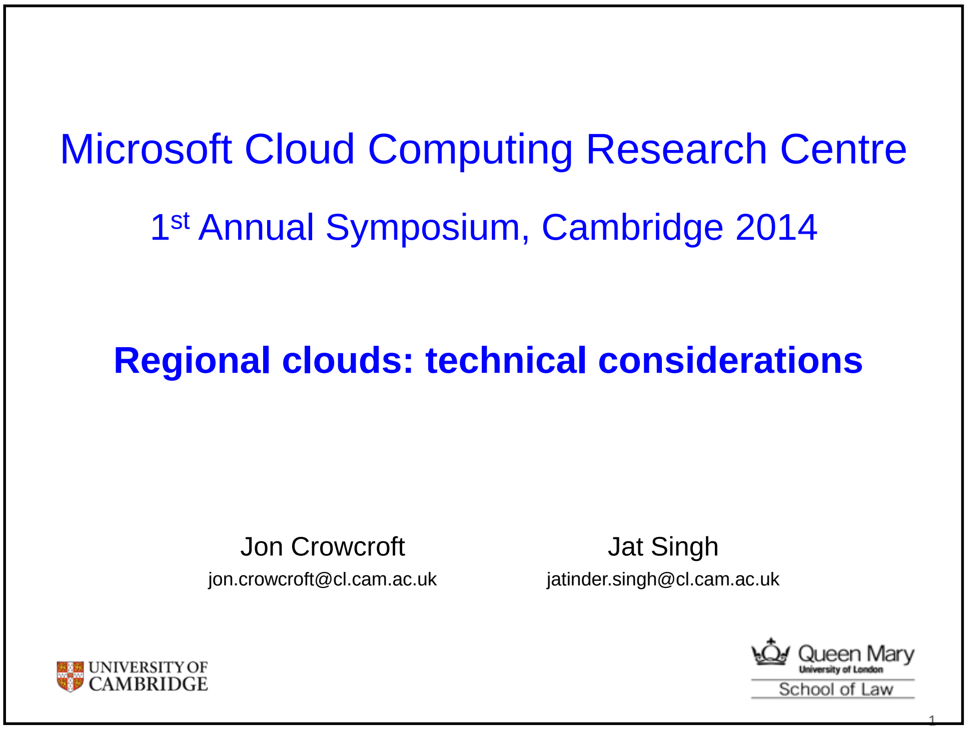

Microsoft Cloud Computing Research Centre

1st Annual Symposium, Cambridge 2014

Regional clouds: technical considerations

Jon Crowcroft
jon.crowcroft@cl.cam.ac.uk

Jat Singh
jatinder.singh@cl.cam.ac.uk

UNIVERSITY OF CAMBRIDGE

Queen Mary University of London School of Law

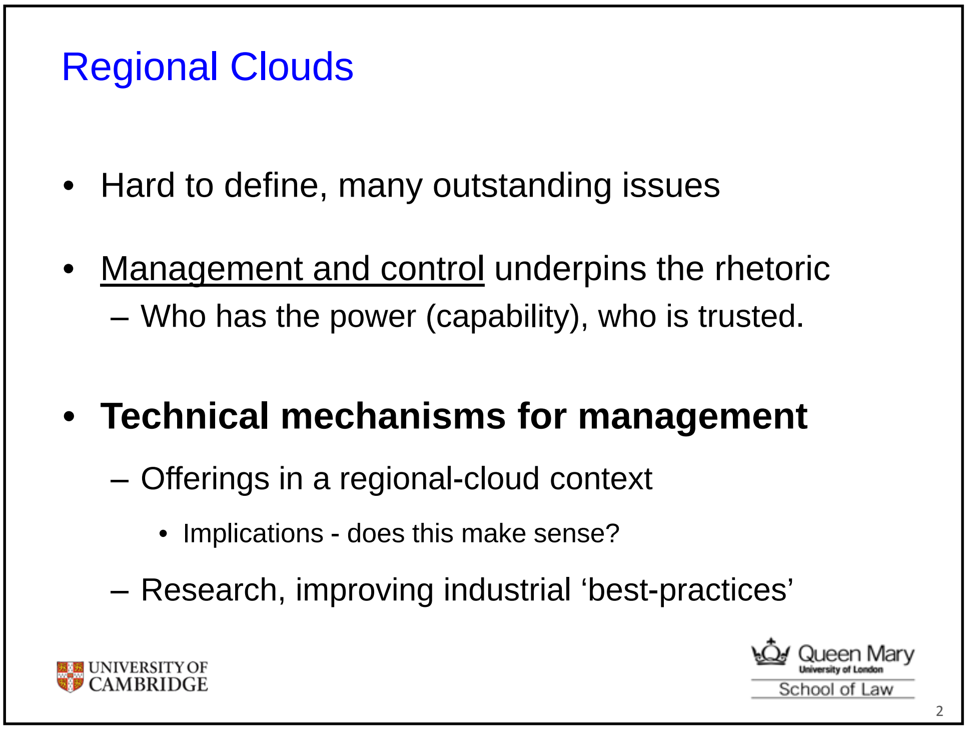

Regional Clouds

- Hard to define, many outstanding issues
- Management and control underpins the rhetoric
  - Who has the power (capability), who is trusted.
- **Technical mechanisms for management**
  - Offerings in a regional-cloud context
    - Implications - does this make sense?
  - Research, improving industrial ‘best-practices’

UNIVERSITY OF CAMBRIDGE

Queen Mary University of London School of Law

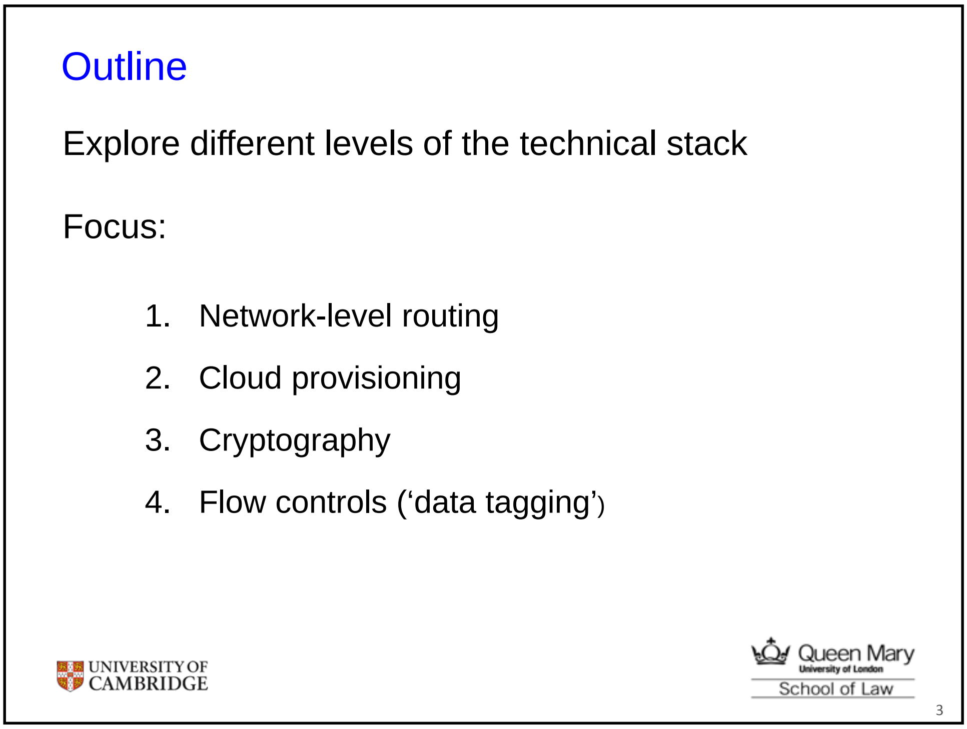

Outline

Explore different levels of the technical stack

Focus:

1. Network-level routing
2. Cloud provisioning
3. Cryptography
4. Flow controls (‘data tagging’)

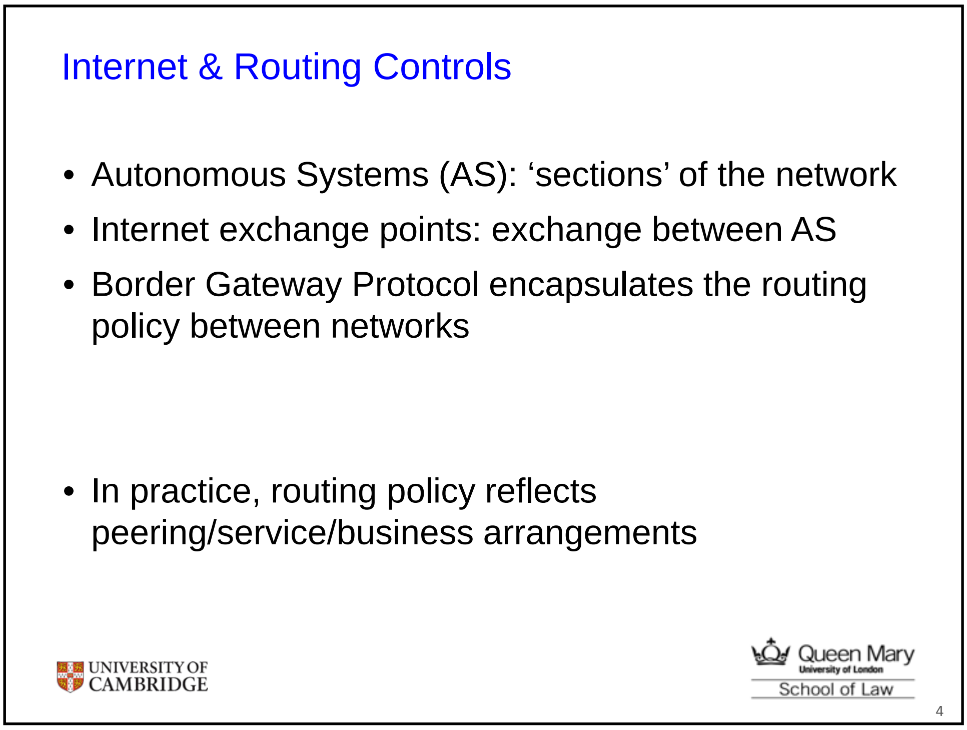

Internet & Routing Controls

- Autonomous Systems (AS): ‘sections’ of the network
- Internet exchange points: exchange between AS
- Border Gateway Protocol encapsulates the routing policy between networks

- In practice, routing policy reflects peering/service/business arrangements

Internet & routing controls (regional clouds)

- Cloud providers manage their infrastructure
  - Many already account for geography for better service provisioning (performance, latency, etc.)
  - Bigger providers already involved in peering arrangements
- Technically feasible with right *incentives* to ensure that data is routed within a geographical boundary

  E.g. economic benefits, regulation, …
- But such an approach is blunt
  - applies to all traffic, regardless

Cloud provisioning: service levels

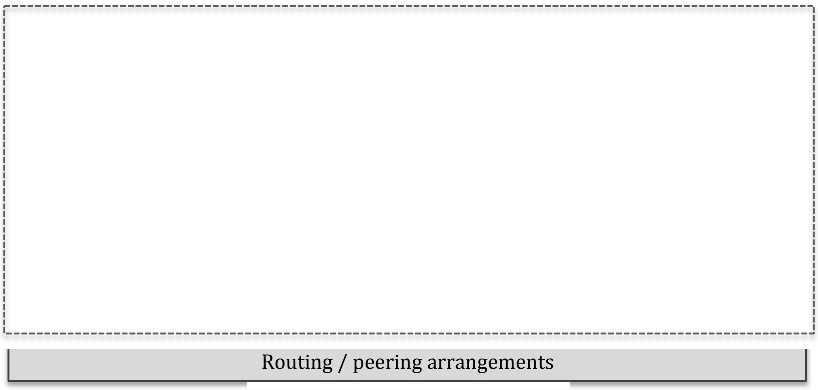

- Provider manages that below, tenants above
- Different management concerns for each service offering

Cloud provisioning: service offerings

- Already work on tailoring services to particular constraints
  - Differential privacy: tailor query results to not reveal too much private information
- Already offer services based on user/tenant locale
  - Not only for performance, but also security, rights management, etc. (e.g. iPlayer)
- Providers already manage their infrastructure
  - Customising service and content for regional concerns
- Thus, already the capability to tailor services for particular regional and/or jurisdictional concerns

Cloud provisioning: Unikernels

- Cloud exists to leverage shared infrastructure
- *Isolation* is important:
  - VMs – Separate for tenants, complete OS, managed by hypervisor
  - Containers – shared OS, isolated users
- Deployment heavy, isolation overheads, …
- Future? *Unikernels*:
  - library OS, build/compile a VM with only that required
  - Hypervisor managed, removes user-space isolation concerns

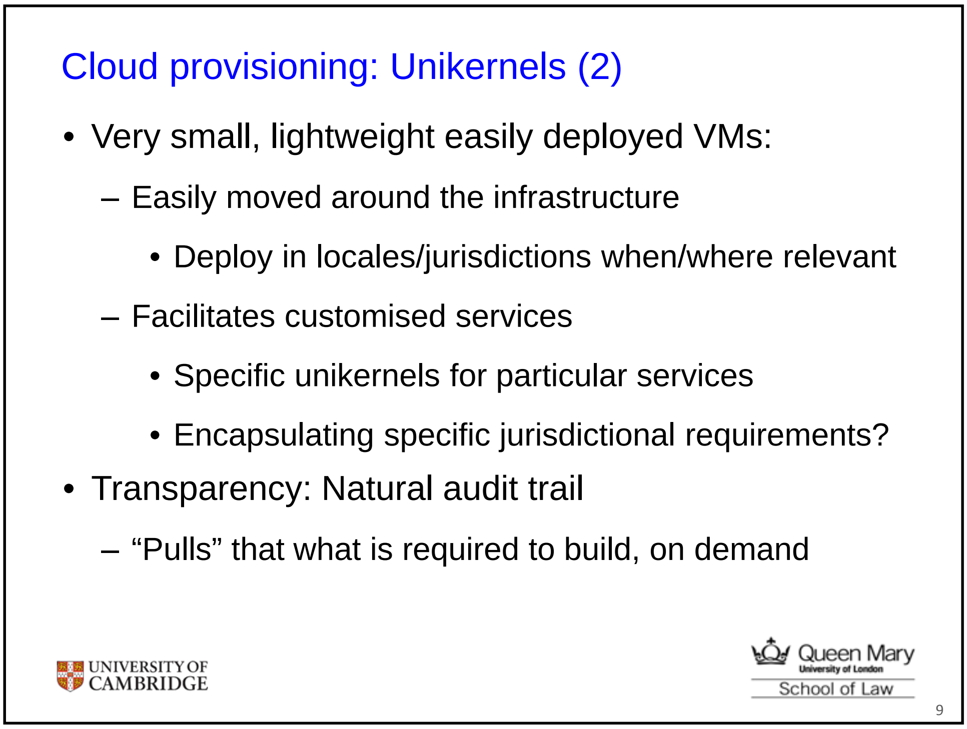

Cloud provisioning: Unikernels (2)

- Very small, lightweight easily deployed VMs:
  - Easily moved around the infrastructure
    - Deploy in locales/jurisdictions when/where relevant
  - Facilitates customised services
    - Specific unikernels for particular services
    - Encapsulating specific jurisdictional requirements?
- Transparency: Natural audit trail
  - “Pulls” that what is required to build, on demand

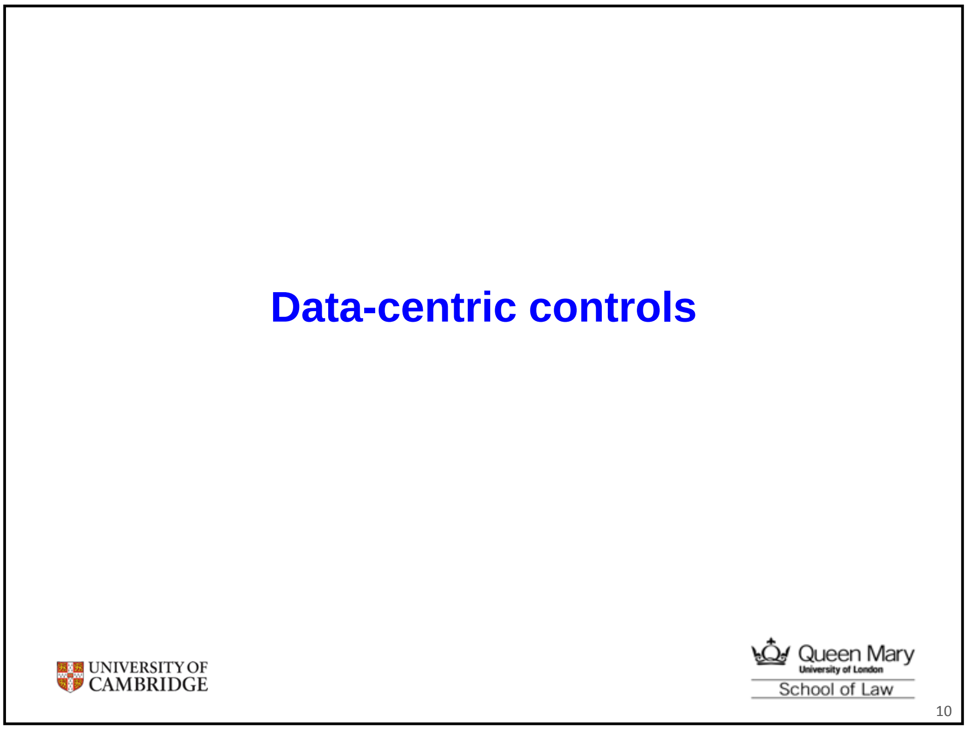

Data-centric controls

Cryptography

- Range of purposes:
  - Data protection: storage, transit, comm. channels
  - Authentication, certification, attestation, etc.
- Encryption
  - Unintelligible, except those with the keys
    - ***encrypt****(plaintext, key)* => ***ciphertext***
    - ***decrypt****(ciphertext, key)* => ***plaintext***
- *Regional Q: Who can (potentially) access the keys?*

Client-side encryption

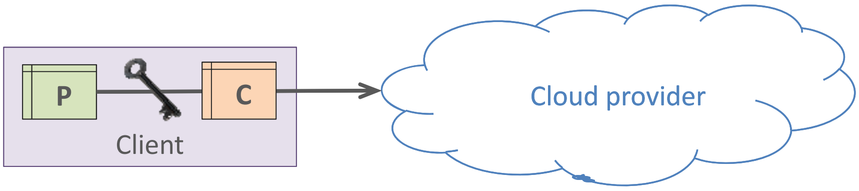

- **Cloud services**
  - Computation generally on *plaintext*
  - Fully homomorphic encryption not practicable (yet)
  - Encrypted search, privacy-preserving targeted ads

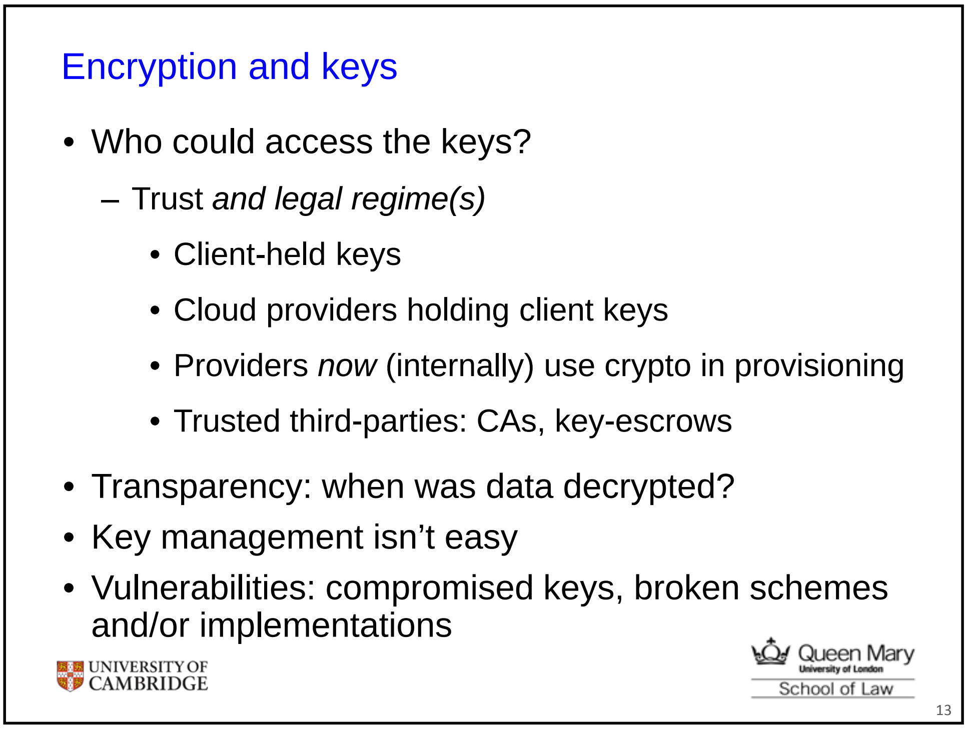

Encryption and keys

- Who could access the keys?
  - Trust *and legal regime(s)*
    - Client-held keys
    - Cloud providers holding client keys
    - Providers *now* (internally) use crypto in provisioning
    - Trusted third-parties: CAs, key-escrows
- Transparency: when was data decrypted?
- Key management isn’t easy
- Vulnerabilities: compromised keys, broken schemes and/or implementations

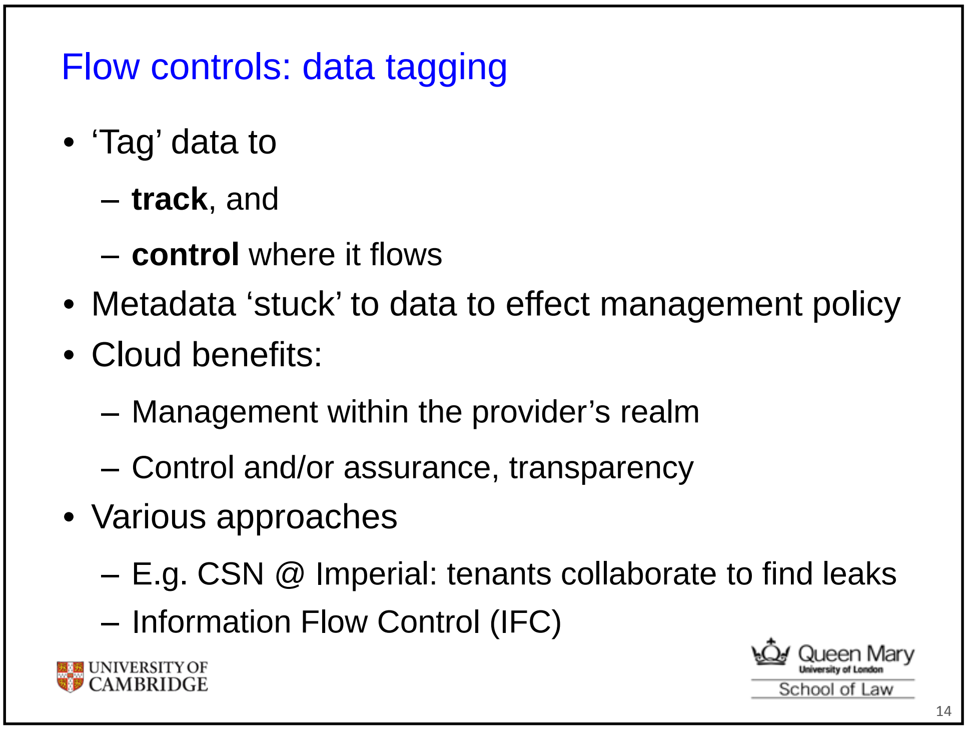

Flow controls: data tagging

- ‘Tag’ data to
  - **track**, and
  - **control** where it flows
- Metadata ‘stuck’ to data to effect management policy
- Cloud benefits:
  - Management within the provider’s realm
  - Control and/or assurance, transparency
- Various approaches
  - E.g. CSN @ Imperial: tenants collaborate to find leaks
  - Information Flow Control (IFC)

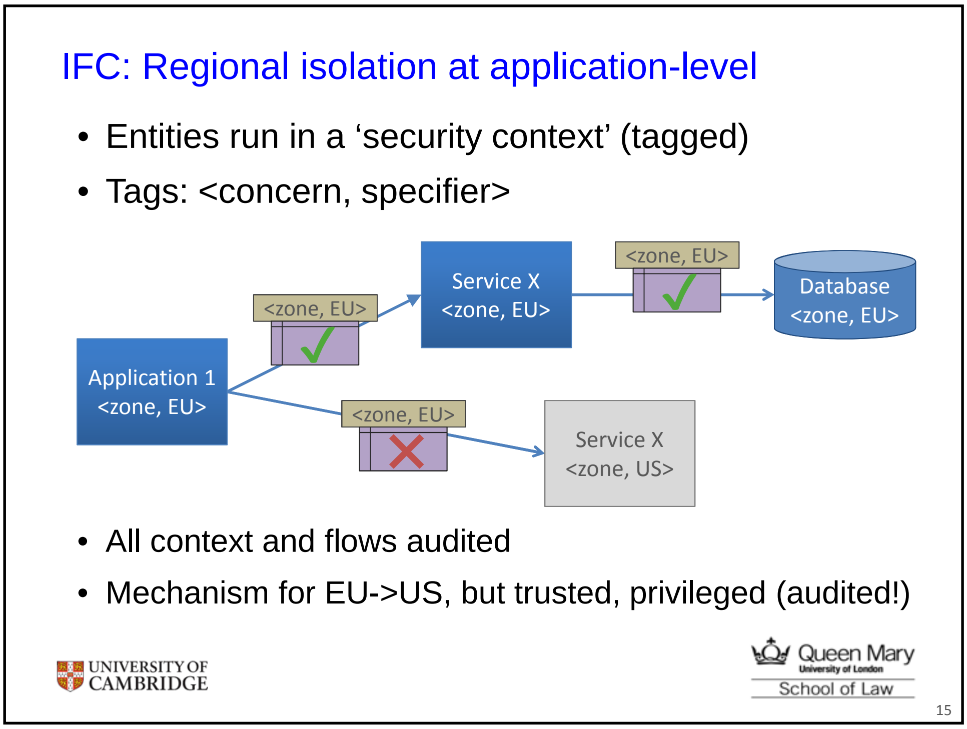

IFC: Regional isolation at application-level

- Entities run in a ‘security context’ (tagged)
- Tags: <concern, specifier>

Application 1 <zone, EU>

<zone, EU> ✓

Service X <zone, EU>

<zone, EU> ✓

Database <zone, EU>

<zone, EU> ✗

Service X <zone, US>

- All context and flows audited
- Mechanism for EU->US, but trusted, privileged (audited!)

UNIVERSITY OF CAMBRIDGE

Queen Mary University of London School of Law

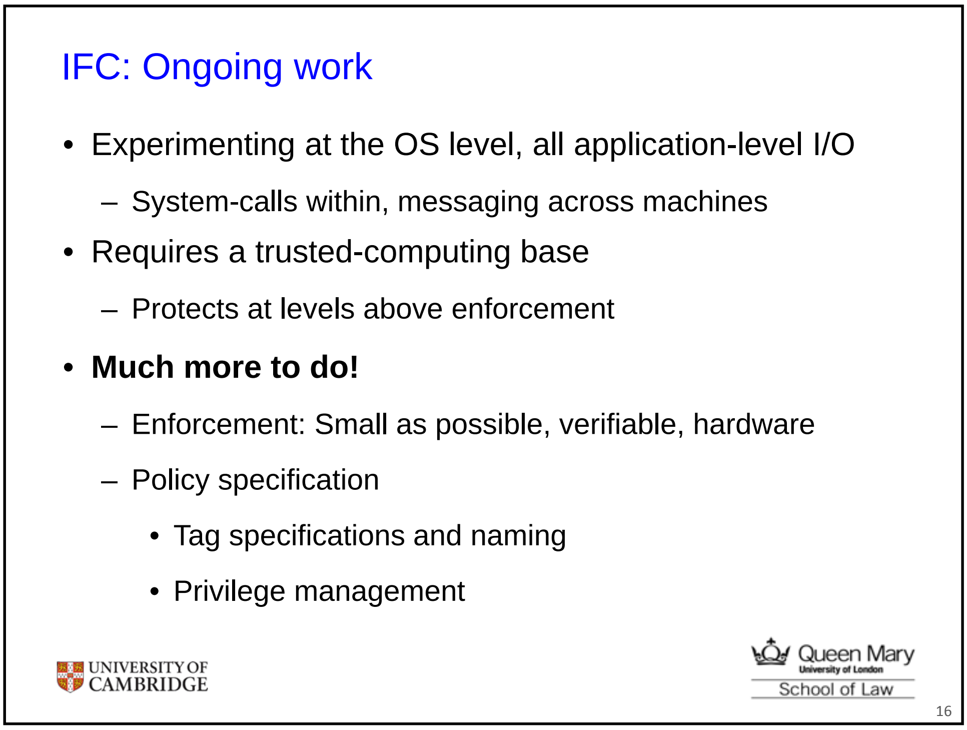

IFC: Ongoing work

- Experimenting at the OS level, all application-level I/O
  - System-calls within, messaging across machines
- Requires a trusted-computing base
  - Protects at levels above enforcement
- **Much more to do!**
  - Enforcement: Small as possible, verifiable, hardware
  - Policy specification
    - Tag specifications and naming
    - Privilege management

UNIVERSITY OF CAMBRIDGE

Queen Mary University of London School of Law

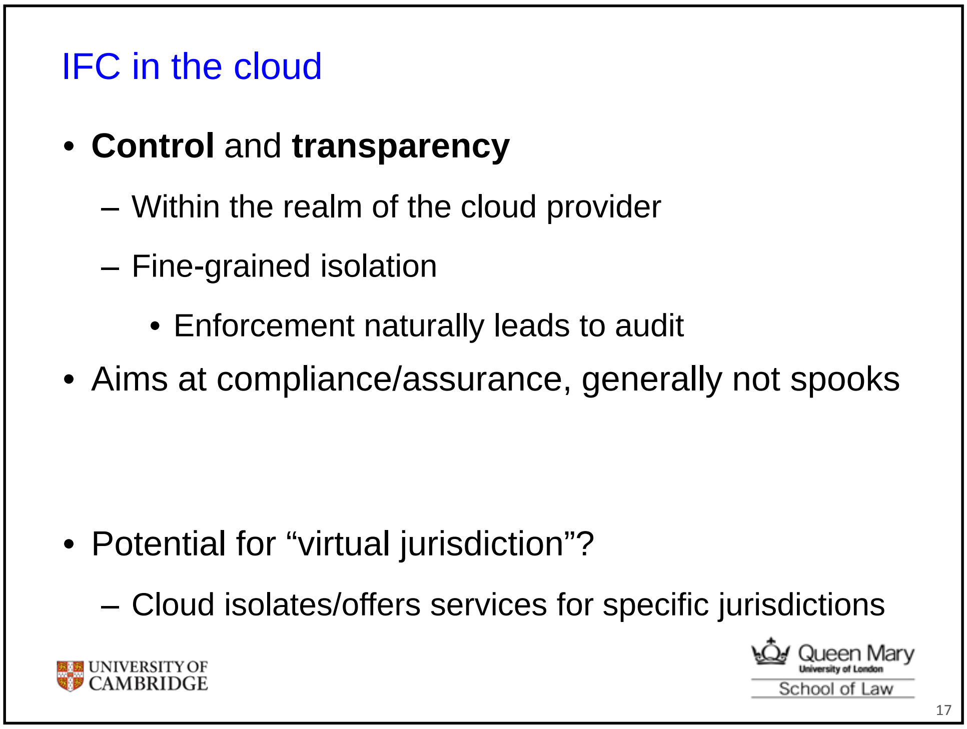

IFC in the cloud

- **Control** and **transparency**
  - Within the realm of the cloud provider
  - Fine-grained isolation
    - Enforcement naturally leads to audit
- Aims at compliance/assurance, generally not spooks
- Potential for “virtual jurisdiction”?
  - Cloud isolates/offers services for specific jurisdictions

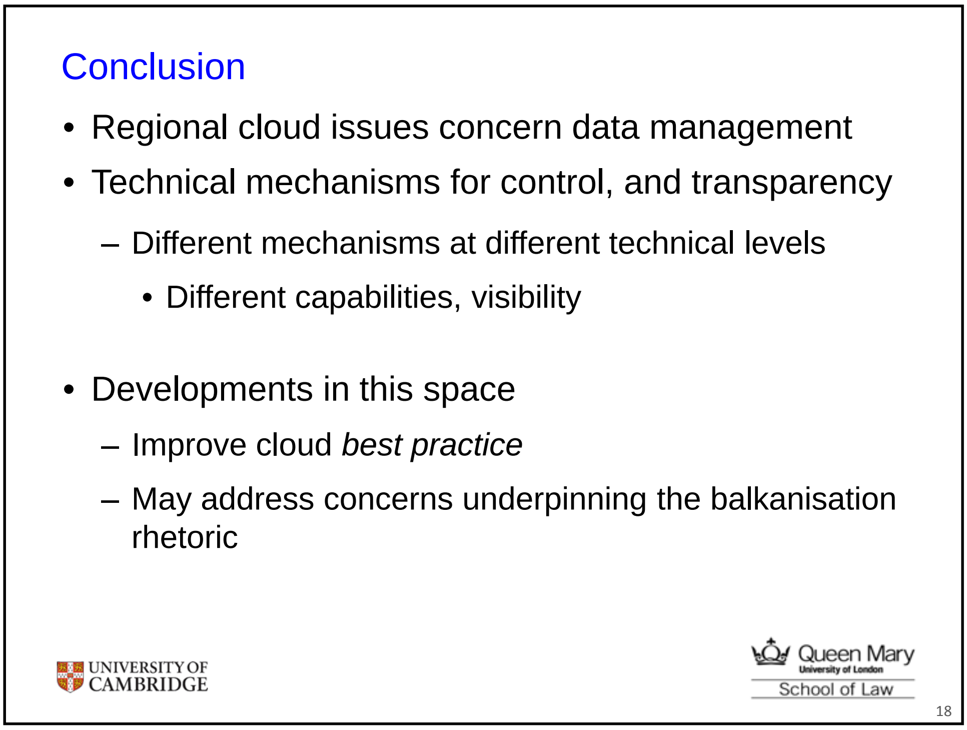

Conclusion

- Regional cloud issues concern data management
- Technical mechanisms for control, and transparency
  - Different mechanisms at different technical levels
    - Different capabilities, visibility
- Developments in this space
  - Improve cloud *best practice*
  - May address concerns underpinning the balkanisation rhetoric

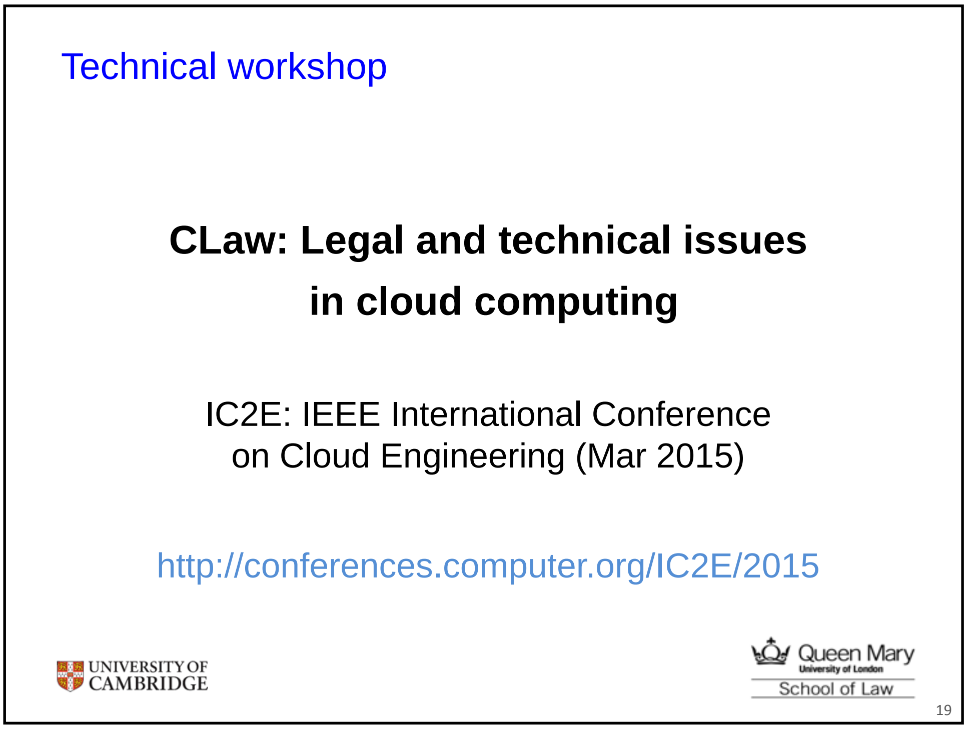

Technical workshop

CLaw: Legal and technical issues in cloud computing

IC2E: IEEE International Conference
on Cloud Engineering (Mar 2015)

http://conferences.computer.org/IC2E/2015

UNIVERSITY OF CAMBRIDGE

Queen Mary University of London School of Law

19