Clouds and the Network

Martin Swany

Indiana University, Informatics and Computing, InCNTRE

Cloud Computing

- Computing resources offered as a service
  - Multics (precursor to UNIX) was created to support a "computing utility"
- Advances in virtualization make service-oriented computing even more flexible
  - Key enabler for modern Cloud Computing
- For enterprise and scientific computing, there are many potential benefits
  - Better resource, power utilization
  - Access to large pools of resources without ongoing maintenance
- Ubiquitous, high performance network connectivity is critical to this vision

InCNTRE

Models of Cloud Computing

- Software as a Service (SaaS)
  - Generally, this implicitly includes network costs
- Infrastructure as a Service (IaaS)
  - Create your own
- For many applications, the ability to use cloud computing and storage assumes that they can communicate effectively
- The network is often pay-as-you-go
  - Amazon and others charge for bytes transferred in and out of the cloud (but not from the S3 storage system)
  - Profile of applications that are effective in (commercial) clouds is potentially limited
  - Paying for what you get is not the same as getting what you want

InCNTRE

Cloud Computing and the Network

- One key promise of Cloud is resource mobility
  - Economics – who has the best price?
  - Strategy – who has the most availability?
- Are Cloud resources fungible?
  - Fungible: interchangeable, equivalent
- Today compute is “tethered” to the data
  - We ask whether to move the compute to the data or the move the data to the compute
  - The cloud vision requires far more flexibility
- Network as a Service (NaaS) is the missing piece!

InCNTRE

Network (Connectivity) Service

- Basic connectivity is necessary for virtual infrastructure
  - VMs need to be on the network
- Cloud control frameworks have various ways to manage this
  - Generally create Virtual LANs (VLANs)
- VLANs are implemented via a “tag” in each packet
- Can be accomplished in both virtual and physical switches

InCNTRE

OpenStack Quantum

- OpenStack will be discussed later today
- Quantum is a component of OpenStack designed to make network connectivity explicit
- Create Networks, Ports and Attachments inside cloud provider infrastructure

InCNTRE

Dynamic Circuit Networks

- Implement “Bandwidth on Demand”
- Provide on-demand point to point circuits across Wide-area Networks
  - Often a few hours to a few days
- Designed for bandwidth intensive applications that require more deterministic network performance for short periods of time
  - E.g. experiments using the Large Hadron Collider, telemedicine, radio astronomy

InCNTRE

Dynamic Circuit Networks

- In active development
- Interoperable control plane standards being developed by ESnet, Internet2, GEANT and others as part of the Open Grid Forum (OGF) and Global Integrated Lambda Facility (GLIF) organizations

InCNTRE

Dynamic, Virtual Networks

- Networks will likely follow the trends that led processors to “multicore”
  - Single-link speeds will grow more slowly, while bundles of parallel links proliferate
  - Like with processor speed, many hard problems can be avoided by waiting…
- The solution is dynamic networks
- Applications, infrastructure allocate links on demand

InCNTRE

Network Virtualization

- Network Virtualization Overlays
  - Internet Engineering Task Force (IETF) Working Group
- Multi-tenant data center must separate customers
  - Coke and Pepsi
- Bandwidth reservation and Quality of Service are key issues

InCNTRE

Software Defined Networking

- Software Defined Networking embodies the idea that network infrastructure needs to be programmable and fully virtualized
- The OpenFlow protocol, developed at Stanford is transforming the networking industry
  - Fine-grained control of network hardware
- Rather than playing “tricks” with tagging and policy SDN enables explicit control

InCNTRE

Performance Matters

- Cloud computing is enabled by high-performance network connectivity
- Network overhead
  - How much work does the system need to do?
- Network proximity
  - Latency

InCNTRE

Host Overhead

- Significant effort for many years to mitigate overhead
  - Layering is good design, but has performance impact
- Data copying in the host is overhead
  - Segmentation to make a packet ready for the network
- Ideally a given bit of data is never copied

InCNTRE

Zero Copy

- Advanced APIs to manage network functionality
  - OpenFabrics “verbs” API
  - Critical, but potentially difficult to use
- What about Virtualization?
  - Think about what happens in a vswitch
- How do we provide isolation without significant performance impact?

InCNTRE

Latency and the Network

- The Transmission Control Protocol (TCP) is widely used
  - provides reliable transmission of byte streams over best-effort packet networks
  - Sequence number to identify stream position inside segments
  - Segments are buffered until acknowledged
- TCP suffers in the face of latency

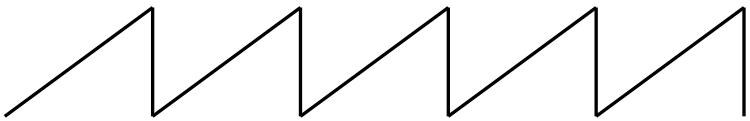

$$BW = \frac{mss}{rtt\sqrt{p}} * C$$

InCNTRE

Phoebus

- The Phoebus project aims to help bridge this network performance gap
- Phoebus is based on the concept of a “session” that enables multiple adaptation points in the network to be composed
- Phoebus provides a gateway for legacy applications
  - Use advanced networks without modification

Internet2 Phoebus Prototype

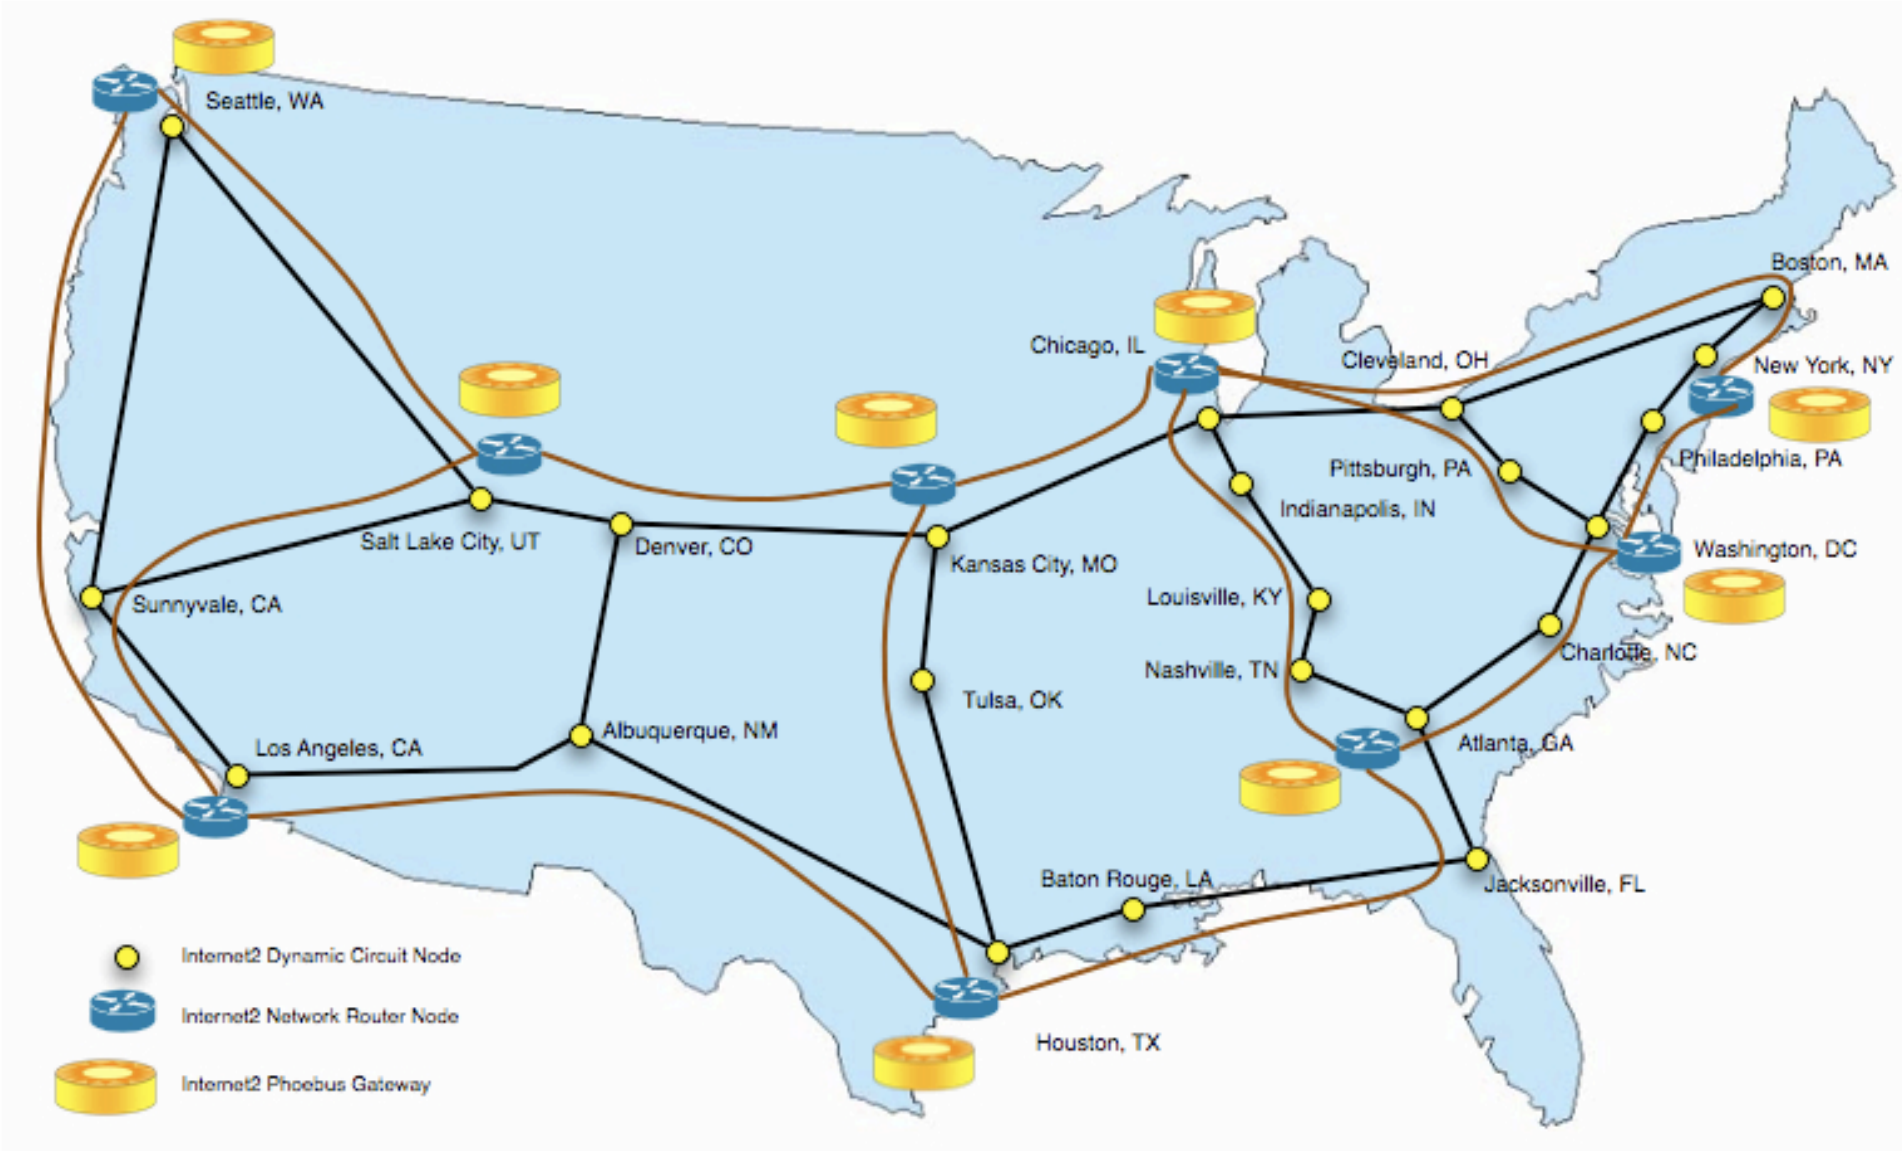

The Network is the Cloud

- Dynamic networks are the missing piece in the cloud picture
- The metaphor behind the name cloud computing seems to have been “using resources out in the cloud”
- Performance does not “just happen” and this aspect cannot be ignored

InCNTRE