

Operating Systems

Fall 2014

Cloud Computing and Data Centers

Myungjin Lee

myungjin.lee@ed.ac.uk




Google data center locations



A closer look



Inside data center



- A datacenter has 50 - 250 containers
- A container has 1,000 - 2,000 servers
- A server has two processors, 2 disks, tons of memory, battery backup
- Processors are chosen for power efficiency, not performance

Some facts about data centers

- Google has **~0.9 million servers** in its all DCs
 - **260M watts** of power = 0.01% of global energy
- Facebook processes **750TB of data** every day
 - Around 7PB of photo storage from its facility every month
- Amazon serves **~40 PB of videos** per month
 - Around 450,000 servers
- Microsoft has 1,000,000 servers
 - It has so far spent around **\$23 billion**

What do these numbers imply?

- Fueling the Internet



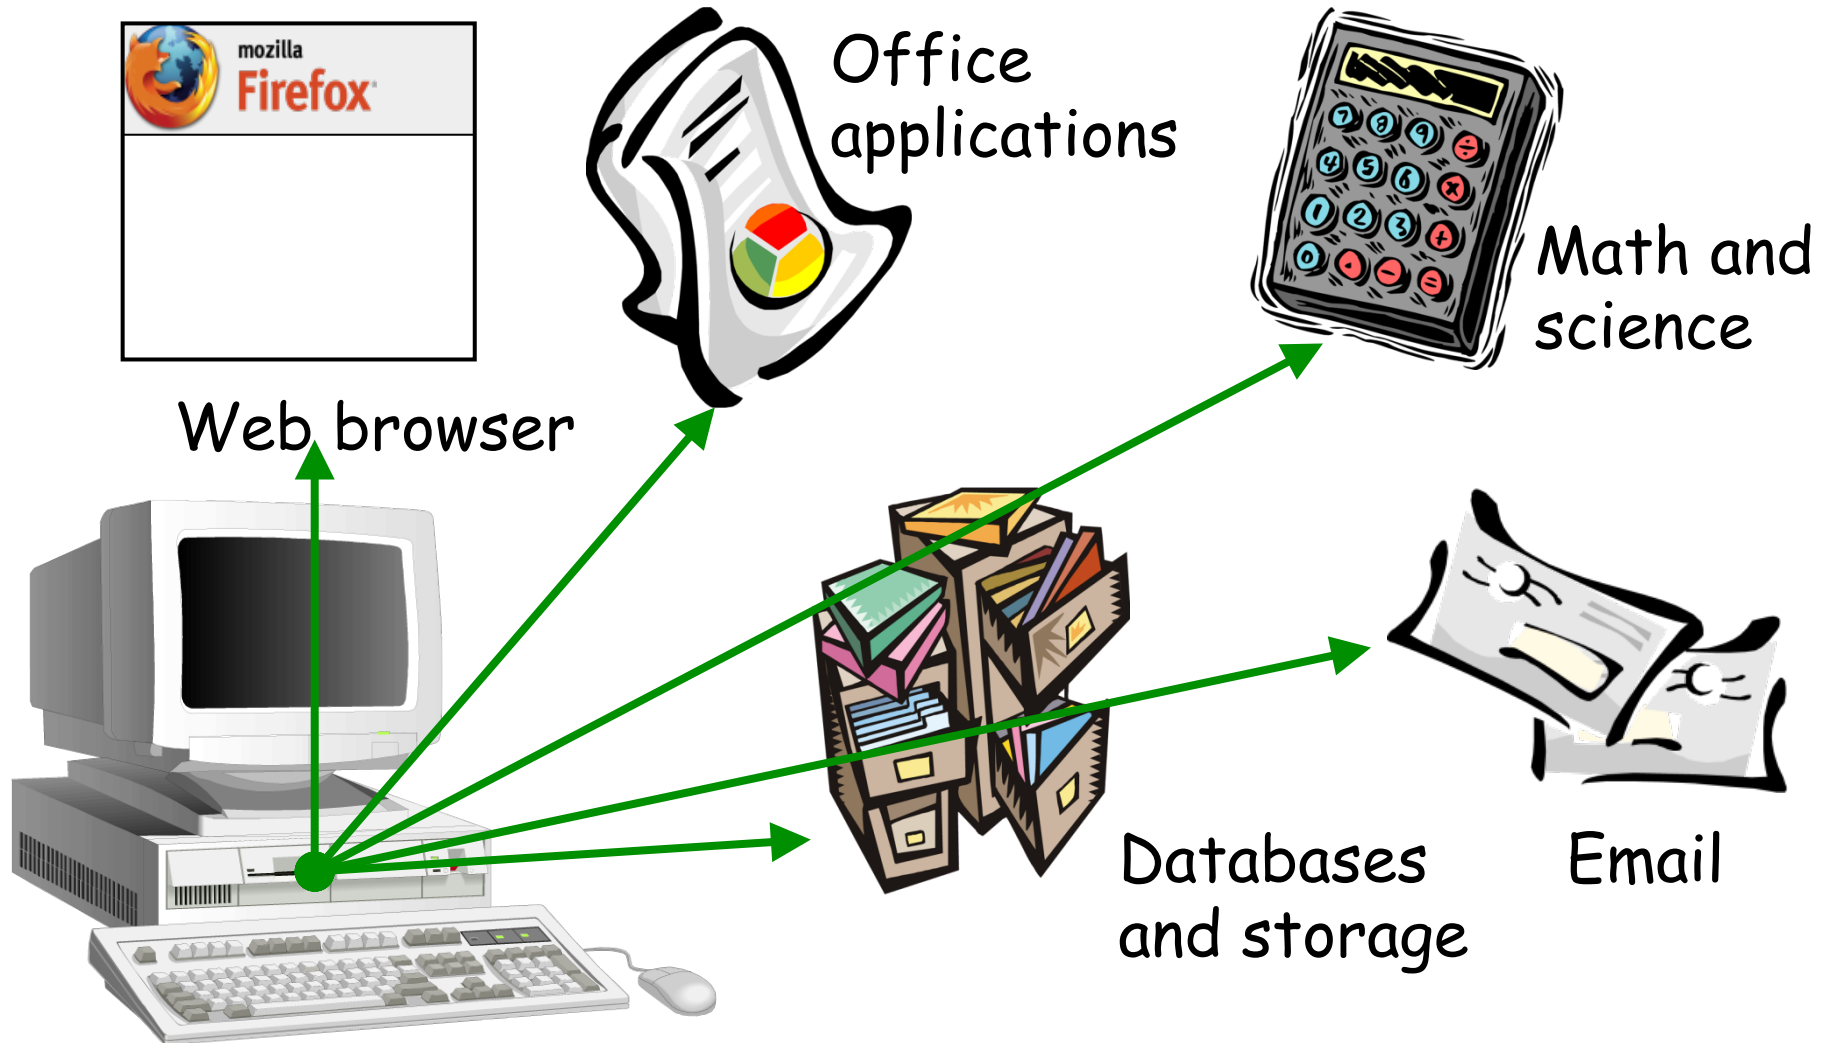
- Too big to fail

Google Outage on 17th Aug 2013

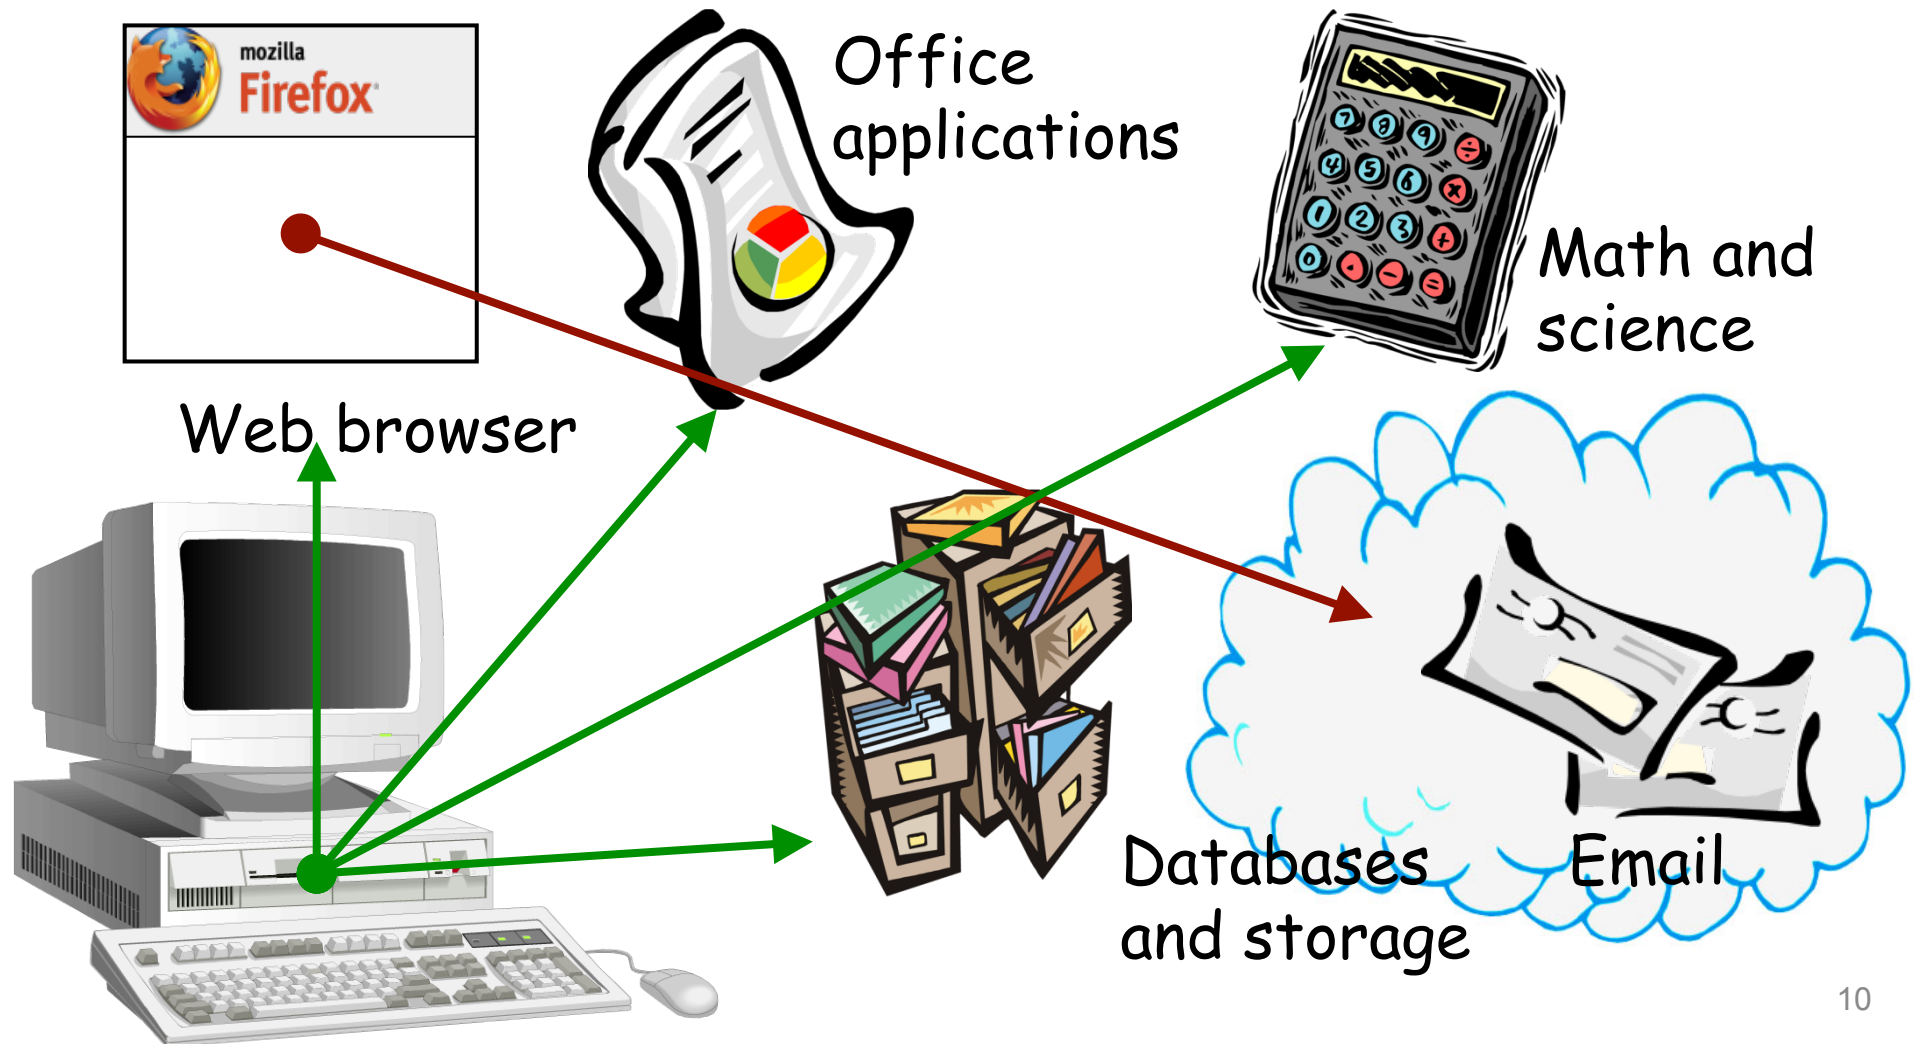
40% drop in Internet traffic

\$545,000 revenue loss for 5 min

Personal computing



Cloud email accessed through the browser



... with the cloud provider's domain name ...

Gmail - Inbox - lazowska@gmail.com - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://mail.google.com/mail/?shva=1#inbox

Google Reader UW CSE News > Edit Post - WordPress Gmail - Inbox - lazowska@gmail.com

Gmail Calendar Documents Photos Reader Sites Web more ▾

lazowska@gmail.com | Settings | Older version | Help | Sign out

Search Mail Search the Web

Compose Mail

Inbox

Starred ☆

Sent Mail

Drafts

Personal

Travel

6 more ▾

Contacts

Tasks

Chat

Search, add, or invite

Ed Lazowska

Set status here ▾

CCIE Training + JOB - www.lanwanprofessional.com - Includes CCNA, CCNP, CCIE Training \$90,000 - \$150,000+ | IT F Sponsored Link

Archive Report spam Delete Move to ▾ Labels ▾ More actions ▾ Refresh 1 - 2 of 2

Select: All, None, Read, Unread, Starred, Unstarred

<input type="checkbox"/>	☆ me	PNWGP Board meeting of Saturday October 11 - DRAFT minutes of our b	10/11/08
<input type="checkbox"/>	☆ me	Saturday conference call - AI, The PNWGP Board is Ed Lazowska, David l	10/6/08

Select: All, None, Read, Unread, Starred, Unstarred

Archive Report spam Delete Move to ▾ Labels ▾ More actions ▾ Refresh 1 - 2 of 2

start Gmail - Inbox - lazo... Computer Manage... Microsoft Word Inbox - Microsoft ... Cloud computing 7:23 PM

... or with your own

Lazowska/Downs Family Mail - Inbox (1) - ed@lazowska.org - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://mail.google.com/a/lazowska.org/?account_id=ed%40lazowska.org#inbox

Google Reader UW CSE News > Edit Post — WordPress Lazowska/Downs Family Mail - In...

Start Page Mail Calendar Documents Sites

ed@lazowska.org | Manage this domain | Settings | Older version | Help | Sign out

Search Mail Search the Web Show search options Create a filter

Compose Mail

Inbox (1)

Sent Mail

Drafts

All Mail

Spam (13)

DAIO

PNWGP

13 more

Contacts

Tasks

Chat

Search, add, or invite

Ed Lazowska

Discovery News Top Stories : Discovery Channel - Oldest Known Paralyzed Human Discovered - 14 hours ago Web Clip

Archive Report spam Delete Move to Labels More actions Refresh 1 - 50 of 479 Older Oldest

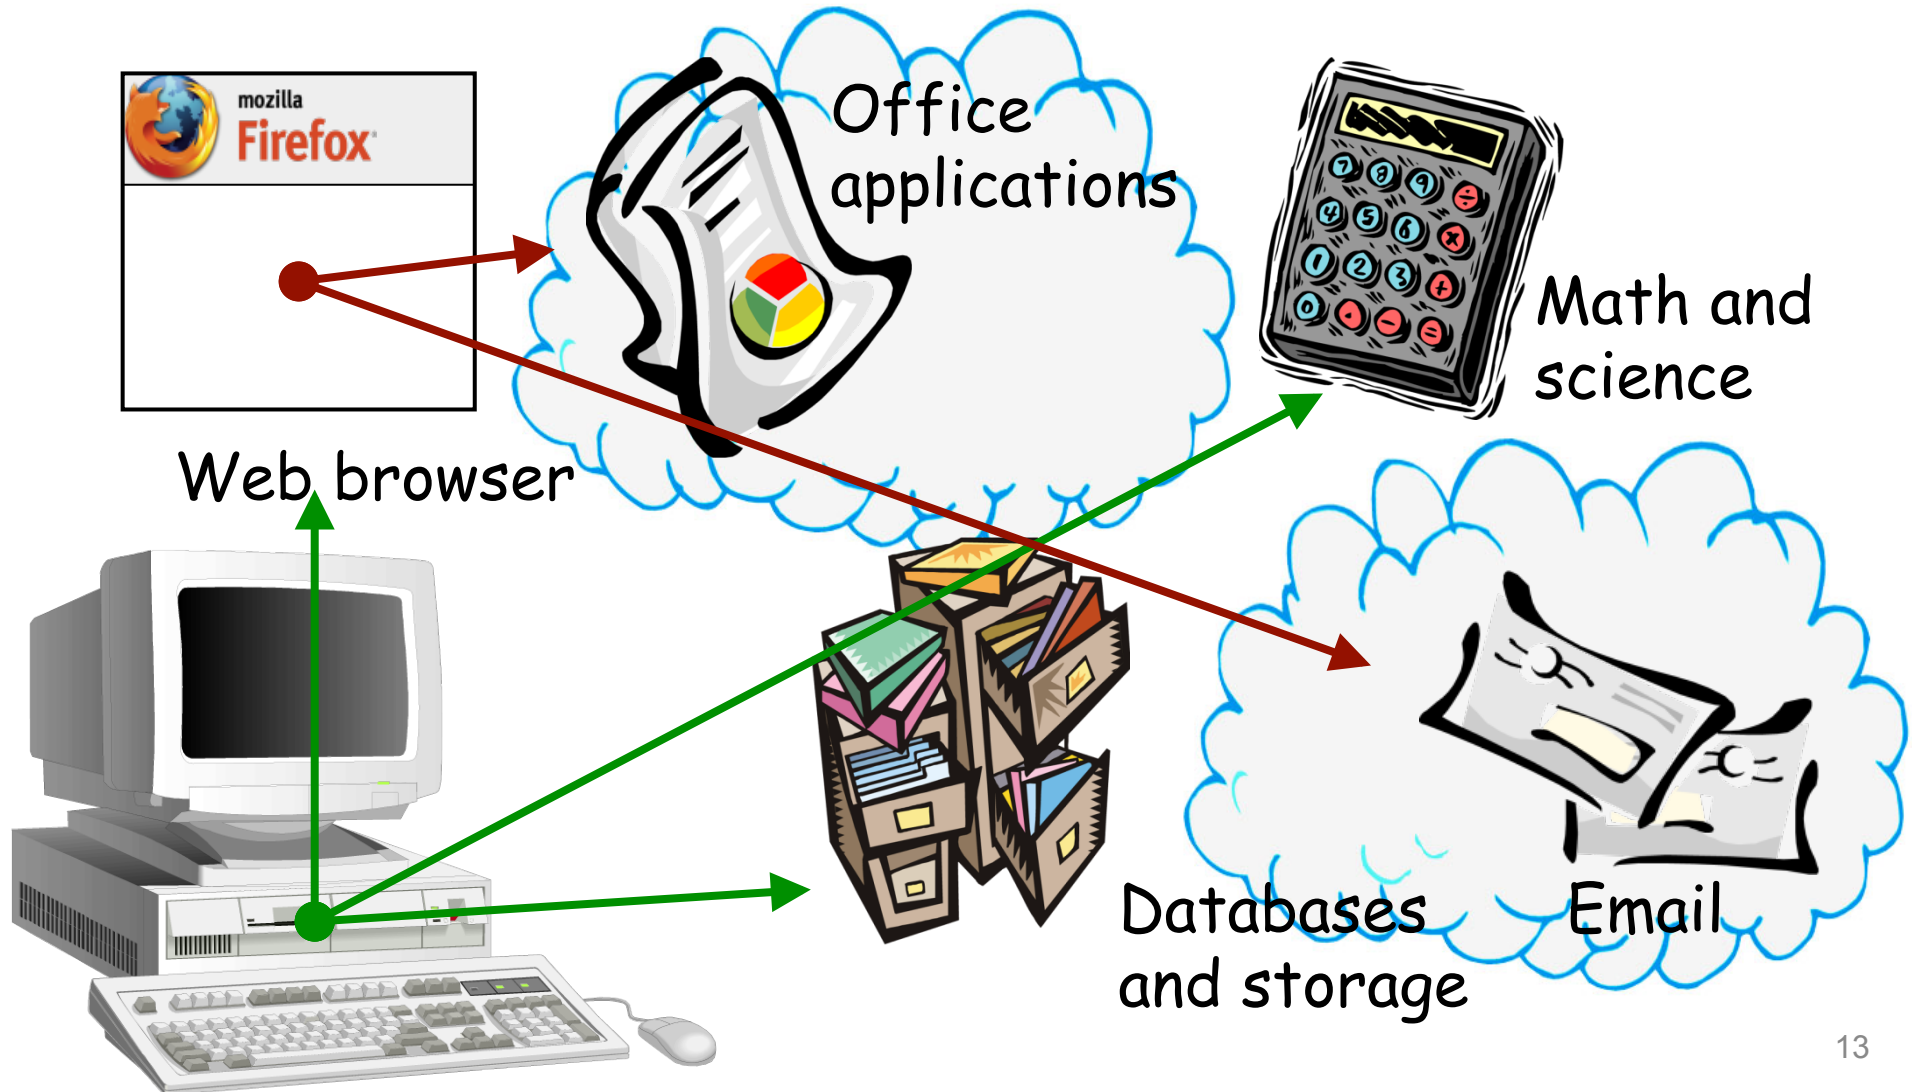
Select: All, None, Read, Unread, Starred, Unstarred

<input type="checkbox"/>	★ Maura, me, Maura (6)	Log Ins- one more question - Hi Ed, I hope you enjoyed your vacation ar	1:22 pm
<input type="checkbox"/>	★ auto-confirm@amazon.com	Your Order with Amazon.com - Amazon.com Logo Thanks for your order, E	Aug 5
<input type="checkbox"/>	★ Valerie, me (3)	Deutsche Investment Settlement - Hi Ed, We sent out letters of instruction	Aug 5
<input type="checkbox"/>	★ David, me, Hank, van (8)	FW: Lawyer for parenting plan issues - more than you know. I again got dr	Aug 5
<input type="checkbox"/>	★ Amazon.com	Shipping update for your Amazon.com order 002-8432582-6117803 - Dear	Aug 4
<input type="checkbox"/>	★ Amazon.com	Your Amazon.com Purchase from Miriam Holdings LLC - Dear ed@lazows	Aug 4
<input type="checkbox"/>	★ me, Kroll (11)	UW CSE Kindle DX allocation - FYI... here are the serial numbers B00418:	Aug 4
<input type="checkbox"/>	★ Amazon.com	Your Amazon.com Order (D01-4722560-1980766) - Amazon.com logo your	Aug 2
<input type="checkbox"/>	★ Amazon.com	Your Amazon.com order has shipped (#002-5339411-4841016) - Greetings	Aug 1
<input type="checkbox"/>	★ auto-confirm@amazon.com	Your Order with Amazon.com - Amazon.com Logo Thanks for your order, E	Aug 1
<input type="checkbox"/>	★ Maura Northen	Deutsche Settlement Check - Dear Ed, You have received or should receiv	Jul 31

Done

start Lazowska/Do... Computer Ma... Microsoft Word Inbox - Micro... Cloud computing Editor - Photo... 7:31 PM

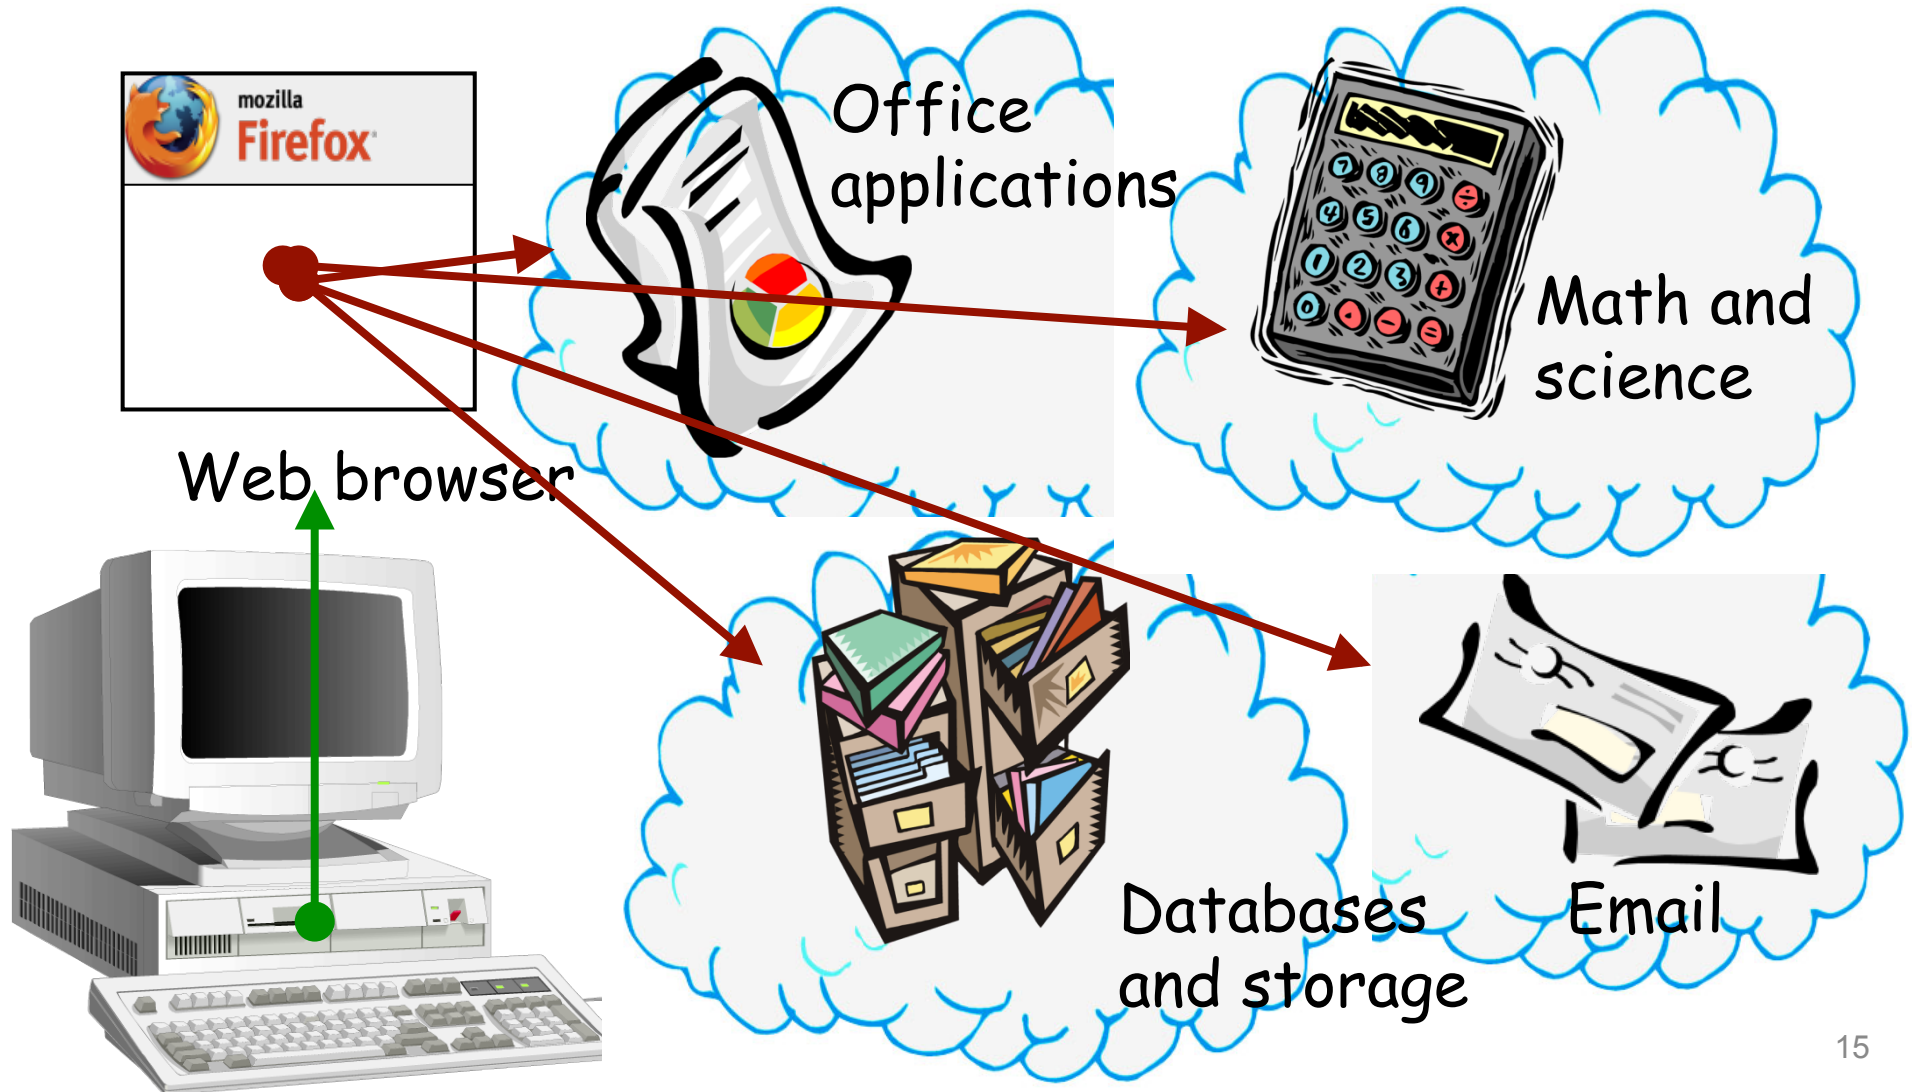
Why not office applications too?



	Name	Folders / Sharing	Date
YESTERDAY			
<input type="checkbox"/>	profit and loss 2009	Not shared	Aug 5 me
<input type="checkbox"/>	Lazowska_To_Do_List	Not shared	Aug 5 me
<input type="checkbox"/>	UW CSE Kindle DX allocation	Shared by me to 1 viewer	Aug 5 me
EARLIER THIS YEAR			
<input type="checkbox"/>	travel log 2009	Not shared	Jul 10 me
<input type="checkbox"/>	March 25 Invitation List	Shared from Greg to everyone, 6 collaborators, 9 viewers	Mar 24 Greg
<input type="checkbox"/>	March 17 NAE Regional Meeting	Shared by me to 1 collaborator	Mar 18 me
<input type="checkbox"/>	CSE 490H: Geo Data Store Proposal	Shared from Tannewt to 5 collaborators	Mar 14 Tannewt
<input type="checkbox"/>	Surface Transportation 2.0	Shared from Henry.c.kelly to 4 collaborators	Jan 30 Henry.c.kelly
<input type="checkbox"/>	CSE490H Capstone Project Proposal	Shared from Rylan to 2 collaborators, 3 viewers	Jan 27 Rylan

Select: All 9, None Showing items 1-9 of 9

Why not everything else?





Search e.g. "templates" or "datastore"

Google App Engine

Home Docs FAQ Articles Blog Community Terms Download



Run your web apps on Google's infrastructure.

Easy to build, easy to maintain, easy to scale.

Getting Started

1. [Sign up](#) for an App Engine account.
2. [Download](#) the App Engine SDK.
3. Read the [Getting Started Guide](#).
4. Check out sample apps in the [App Gallery](#).

Watch and Learn

An Early Look at Java™ Language Support New!

App Engine is unveiling its second language runtime, integration with Google Web Toolkit, and a new Java solution for AJAX web applications and we're eager to get your help and input. Who [sign up](#), but we'll be including more features soon.

- Get the full scoop in our [blog post](#).
- Click over to YouTube to watch our [video](#).
- See our docs for other new features [data](#).




Get an overview of the new sample app. [Watch Now](#)

Azure Services Platform

Search Microsoft Web Microsoft

Home About Solutions Services Resources Community Sign In

Experience the Azure Services Platform



Build new applications in the cloud - or use interoperable services that run on Microsoft infrastructure to extend and enhance your existing applications for you.

- ▶ Azure for web developers
- ▶ Azure for corporate developers


Explore Azure Services

The Azure Services Platform provides a wide range of internet services that can be consumed from both on-premises environments or the internet.



About AWS Products Solutions Resources Support Your Account

Contact Us Create an AWS Account



Hadoop + The AWS Cloud

Introducing Amazon Elastic MapReduce—the Hadoop-based infrastructure service that lets you build and deploy large-scale data processing applications in the cloud.

[Learn More...](#)

Get Started

Sign up for a free AWS account.

[Sign Up Now](#)

Developers

Simply sign up & start developing in the cloud with these resources and tools:

- [Technical Documentation](#)
- [Cloud Architectures Whitepaper \(pdf\)](#)
- [Amazon Machine Images](#)
- [AWS Community Forums](#)

Business Managers

Learn how Amazon Web Services enables you to reach business goals faster:

- [AWS Solutions for Enterprise Customers](#)
- [Security Whitepaper \(pdf\)](#)
- [Case Studies & Customer Testimonials](#)
- [AWS Blog](#)

Explore Products

Infrastructure Services

- Amazon Elastic Compute Cloud (Amazon EC2)
- Amazon SimpleDB
- Amazon Simple Storage Service (Amazon S3)
- Amazon CloudFront
- Amazon Simple Queue Service (Amazon SQS)
- Amazon Elastic MapReduce
- AWS Premium Support

Payments & Billing

On-Demand Workforce

Alexa Web Services

News & Events

What's New? Media Coverage Events

- May 07, 2009 Amazon CloudFront Adds Access Logging Capability
- Apr 29, 2009 AWS Goes To School With Programs For Educators, Researchers, and Students
- Apr 22, 2009 Amazon EC2 Running IBM Now Available
- Apr 15, 2009 Amazon EC2 Reserved Instances Now Available in Europe
- Apr 09, 2009 Announcing Amazon SQS WSDL Version 2009-02-01 and Amazon SQS in Europe

Consider ...

- Sharing is easy
- Someone else does backup
- Someone else handles software updates
- There's 7x24x365 operations support, auxiliary power, redundant network connections, geographical diversity
- Scalability – both up and down – is instantaneous
- Many fewer demands on the local operating system and machine

Amazon Elastic Compute Cloud (EC2)

- \$0.68 per hour for
 - 4 cores of 2.5 GHz 64-bit 2007 Xeon or Opteron
 - 15 GB memory
 - 1.69 TB scratch storage
- Need it 24x7 for a year?
 - \$3900

- \$0.085 per hour for
 - 1 core of 1.2 GHz 32-bit Intel or AMD
 - 1.7 GB memory
 - 160 GB scratch storage
- Need it 24x7 for a year?
 - \$490

- This includes
 - Purchase + replacement
 - Housing
 - Power
 - Operation
 - Reliability
 - Security
 - Instantaneous expansion and contraction
- 1000 processors for 1 day costs the same as 1 processor for 1000 days!

The nuts and bolts of data center



- Networks
- Servers
- Storages
- Software
- Power systems
- Cooling systems
- ...



How should we design a data center network?

Interconnecting 10,000s of machines

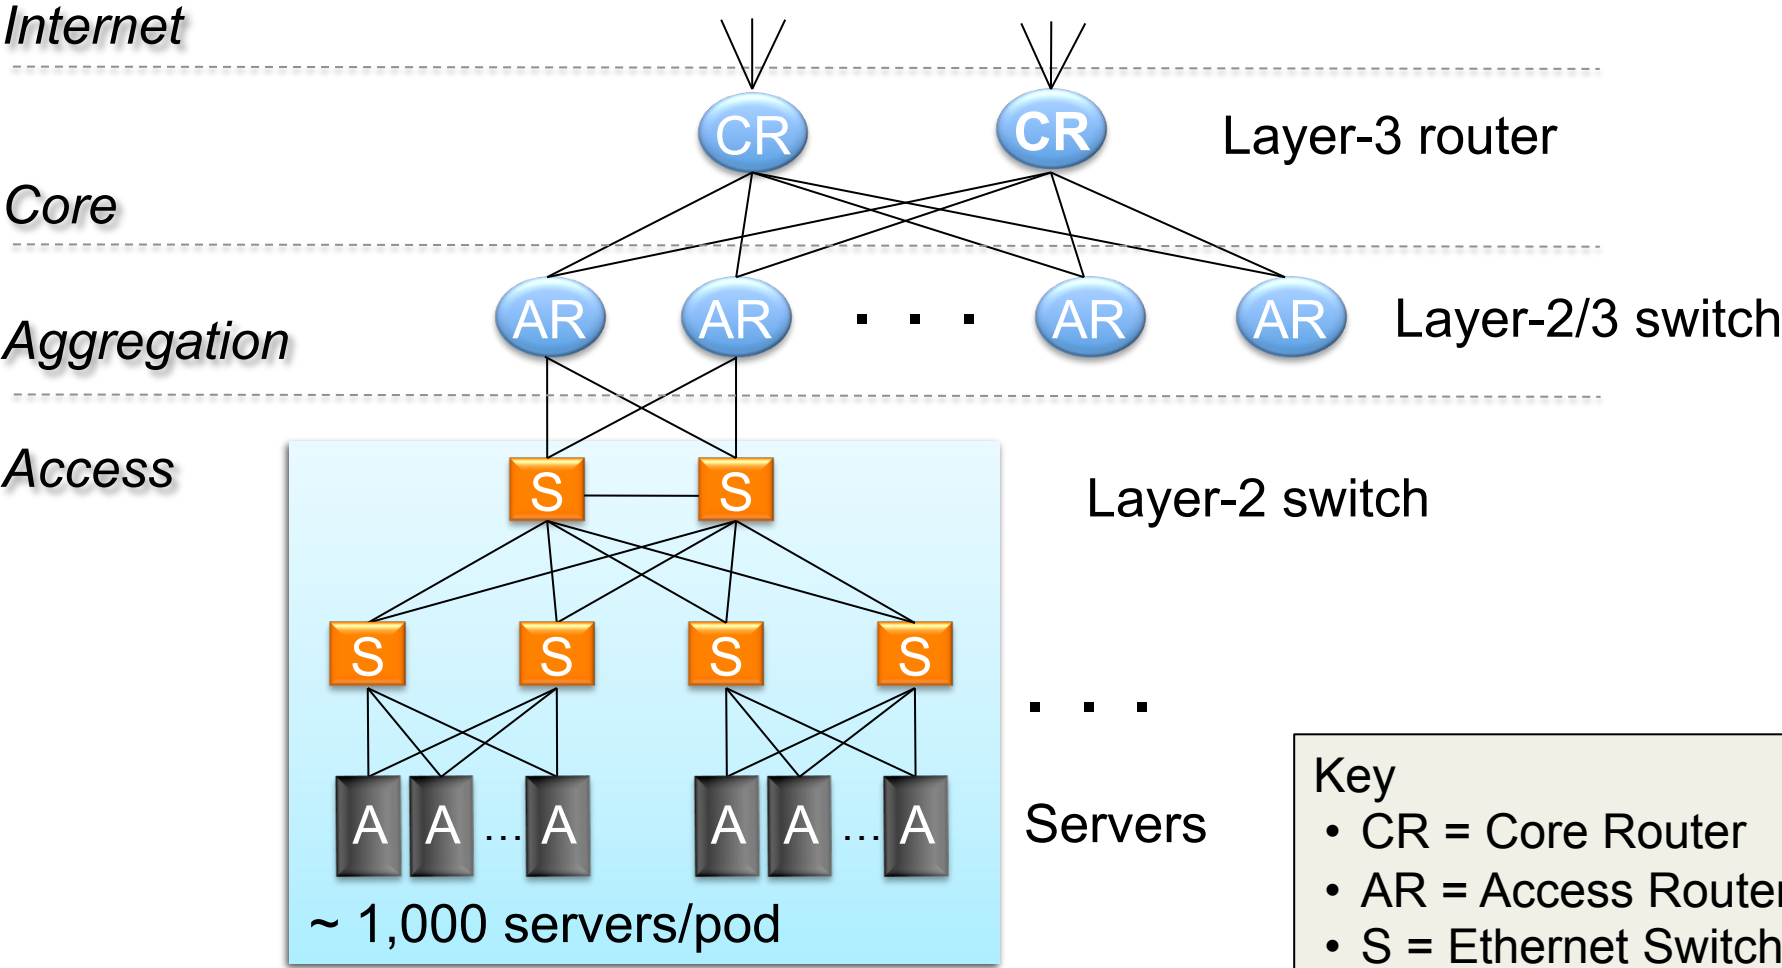
- Top-of-Rack architecture
 - Rack of commodity servers
 - Top-of-Rack Switch
- Aggregation of ToRs



Overall picture



Common data center network topology



Key

- CR = Core Router
- AR = Access Router
- S = Ethernet Switch
- A = Rack of servers

Source: Jen Rexford's COS 561 slides

Characteristics of data center (networks)

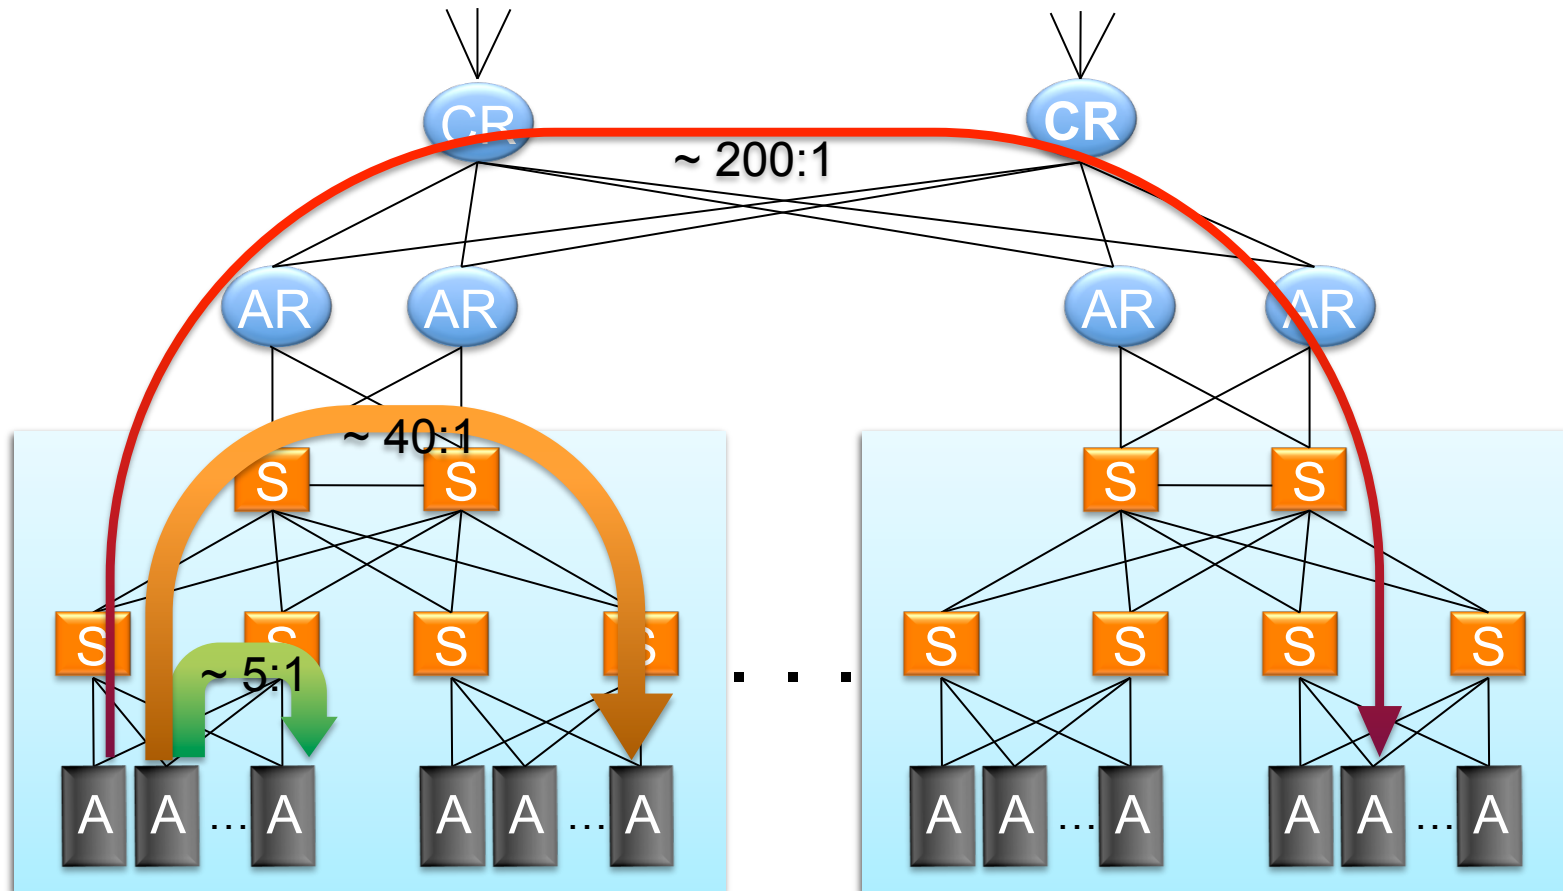
- Single ownership
 - Allows full control over an entire system
 - Less concern about standards and interoperability
- Less heterogeneous environments
 - Similar servers, storage, topology, software stack
- Multiple end-to-end paths
 - E.g. Clos topology, multi-rooted tree topology
- Low end-to-end delays when no congestion
 - Servers in a geographically small region
 - DC: 100s of μ s vs. Internet: 10s of ms to 100s of ms

Applications in data centers

- Web services
- Web search
 - Google Search, Microsoft Bing
- High performance computing (HPC)
- Big data analytics
 - Hadoop, MapReduce, Twitter Storm, etc.
- Machine learning
- Cloud applications
 - DropBox, Google Drive, etc.


Applications compete for data center resources

Capacity mismatch



Example: Fat-tree topology


Example: $K = 4$



Fat-tree topology

A set of $K/2$ ports used for upper level connectivity, another set for lower level connectivity


K -port switches/routers



Pod


Example: $K = 4$

Top Level: core routers

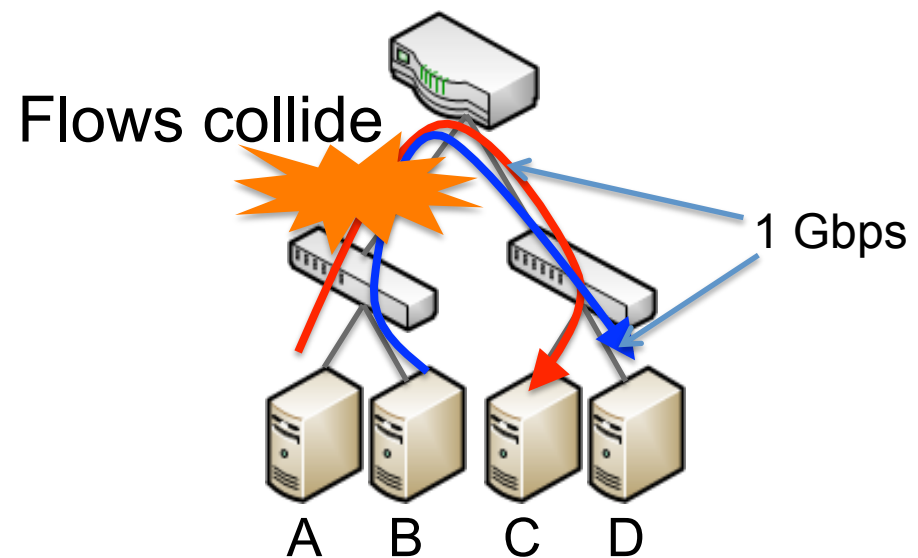


Benefit of multiple equal-cost paths

Each link = 1 Gbps; A talks to C; B talks to D



Throughput = 1 Gbps



Throughput = 500 Mbps

Deciding an end-to-end path of flows is an important scheduling task to fully utilize multiple paths

Exploiting multiple equal-cost paths

- Many approaches
 - ECMP (Equal-Cost Multi-Path) forwarding
 - Monsoon [PRESTO'08]
 - VL2 [SIGCOMM'09]
 - Hedera [NSDI'10]
 - Mahout [Infocom'11]
 - MPTCP (Multipath TCP) [SIGCOMM'11]
 - Packet Spraying [Infocom'13]
 - ...

Data centers are cool!



Google data center, Lenoir, North Carolina, US