




Web Infrastructure for the 21st Century

Pablo Rodriguez




Web Infrastructure: Web Caching



“With 25 years of Internet experience, we’ve learned exactly one way to deal with the exponential growth: Caching”

1997, Van Jacobson

Web Infrastructure: Content Distribution Networks



- 40,000 servers, 900 PoPs, 71 countries
- **300 Gbps**



- 25 PoPs, Hundreds of servers per presence
- **1,000 Gbps**

Struggled to cope with flash crowd events

The screenshot shows the CNN.com homepage with the following elements:

- Header:** CNN.com logo, "MEMBER SERVICES", and "MAKE CNN.com YOUR HOME PAGE".
- Navigation:** "SEARCH" bar with "The Web" and "CNN.com" options, and "Enhanced by Google".
- Main Content:**
 - Top banner: "Informed Insightful In-Depth" and "Start Your FREE Trial To TIME! CLICK HERE >>"
 - Left sidebar: "Home Page" (Updated: 10:40 a.m. EDT (14:40 GMT) April 29, 2004), "World", "U.S.", "Weather", "Business at CNNMoney", "Sports at ESPN", "Politics", "Law", "Technology", "Science & Space", "Health", "Entertainment", "Travel", "Education", "Special Reports", "Paula Zahn NOW WEEKNIGHTS 8 PM ET", "SERVICES" (Video, E-mail Services, CNNtoGO, Contact Us, SEARCH), and "Web CNN.com".
 - Center: "Battle rages in Fallujah" with a video player showing a helicopter in flight.
 - Right sidebar: "MORE TOP STORIES" (Dozens dead in Thai gunbattles, Enemy combatants' case goes to high court, CNNMoney: Comcast drops Disney offer, Wreckage of crashed Black Hawk found in South Carolina, 'Nightline' to honor U.S. military dead, Toddler found alive 30 hours after crash killed father, Boy, 12, charged in 8-year-old girl's death, Life-saving kangaroo to receive valor award, AMERICA VOTES 2004), "CNNRADIO" (Listen to latest updates), "VIDEO" (Mr. Hamburger, meet Mr. Fries), "WATCH CNN TV" (Paula Zahn Now 8 p.m. ET), "BUSINESS at CNNmoney", and "STOCK/FUND QUOTES" (MARKETS: 10:10am ET: 428, DJIA: +57.69 10410.47, NAS: +21.81 2010.72, S&P: +7.28 1130.83).
- Footer:** "TIME 100", "THE MORNING GRIND", "HEAD-TO-TOE FLOW", and "HEALTH FEELS".

CNN.com BREAKING NEWS



Sep 11
attacks

AMERICA UNDER ATTACK

At 8:45 a.m. EDT, the first of two airliners crashed into the World Trade Center, opening a horrifying and apparently coordinated terrorist attack on the United States, which saw the collapse of the two 110-story towers into surrounding Manhattan streets and a later attack on the Pentagon.

[DEVELOPING STORY >>](#)

COMPLETE COVERAGE:


- [Chronology of terror](#)
- President Bush: U.S. will "punish those responsible" | [Statement](#)
- World Trade Centers [collapse](#) after planes hit, 10,000 emergency workers head to scene
- Plane hits [Pentagon](#), part of the Pentagon collapses
- American, United both confirm losing two planes each. The American planes alone held about 150 people
- U.S. Navy vessels deploy off East Coast
- Federal buildings, United Nations evacuated
- [FAA grounds](#) all U.S. flights, sends trans-Atlantic flights to Canada
- Sen. John McCain calls attacks 'act of war'
- Israel evacuates embassies
- Non-essential NATO employees asked to leave Brussels HQ
- Taliban issues [statement](#) to tell U.S. 'Afghanistan feels your pain'
- Kennedy Space Center, LAX, [Disney](#) Florida parks closed
- [World shock](#) over U.S. attacks
- [Evacuations amid world chaos](#)
- Middle East, Asia condemn attacks

18th WWW Conference Keynote

© 2009 Telefónica Investigación y Desarrollo, S.A. Unipersonal

Telefónica

6 fold growth in four years



Source: IDC, 2007



Cloud

Social

18th WWW Conference Keynote

© 2009 Telefónica Investigación y Desarrollo, S.A. Unipersonal

Much today's Web infrastructure to distribute content has been
an after-thought

Roadmap

- Part 1: Locating and Managing Content
- Part 2: Distributing Content
- Part 3: Data Centre Clouds
- Part 4: Clouds for Online Social Networks

Three waves of Networking




1930

1960

1990

Internet Design

- Internet was designed as a thin layer so that ANY application could run on top
- It was not optimized for any particular application, in particular not for Content
- and that created some problems...



Where vs What

Container vs Content

- It is a serious mistake to point to the container, not the content

“I urged them to remove some of the technical mistakes of the language, the predominance of references...”

Turing Award lecture, Tony Hoare

- You have security issues; reasoning issues; you have robustness issues...

Problems...

■ Search

- Search relies on links, if content/links change/disappear search suffers

■ Distribution

- Routers waste capacity copying the same bits millions of times.

■ Replication

- If content is split, it is hard to obtain

Management Problems...

■ Security

- Authenticity, Chain of Custody/Transformation, Revocation!

■ Policy


- Lack of data control, data embargo, privacy and access rules

■ Traceability

- How many hits did my content get?


Content Networking

- Content networking paradigm
 - Van Jacobson
- Content indexed by keys
 - What you want, not where to get it from
- Data is self-certified
 - Secure the data, not the channel
- Storage everywhere
 - Why not adding 1TB to each router



But should it be a Revolution or an Evolution?

IP = the Internet Kernel



Revolution: The Internet as a database?

- Routing is a simple form of searching
- Today you route through paths to reach hosts
- The network could support more complex ways of routing (e.g. find me all files similar to X)

Evolution: Build as an overlay?

- Lots of things to learn from P2P networks
- P2P Naming, DHTs, Chunk retrieval, Swarming
- What if every Web file becomes a P2P Swarm?

Roadmap

- Part 1: Locating and Managing Content
- Part 2: Distributing Content
- Part 3: Data Centre Clouds
- Part 4: Clouds for Online Social Networks

Is the Internet the preferred medium for distributing bulk (delay tolerant) digital content?



Not beyond a certain size ...

e.g. movies, home videos, data backups, data replication

Currently, served by:

- Dedicated networks
- Parcel delivery



10M+ users


1.5 Million DVDs per day
2.5 PB/day!

All US P2P traffic
14 PB/day (cisco)

Postal still carries vast amount of
multimedia traffic



How well is the current Internet dealing with large content transfers...?



**Current bulk data demand is probably
higher than what the Internet can handle**



Internet connection
Blocktastic?

BitTorrent transfers
Slowtastic?

Comcast has violated Net Neutrality.

It's time for the FCC to act. We've filed a
legal action. Please add your voice.

**Tell the FCC: Stop Comcast.
Save Net Neutrality. »**


**SAVE THE
INTERNET**
com

The Effect of Distance

Distance from Server to User	4GB DVD Download Time
Local: <100 mi.	12 min.
Regional: 500-1,000 mi.	2.2 hrs.
Cross continent: -3,000 mi.	8.2 hrs.
Multi-continent: -6,000 mi.	20 hrs.

[Tom Leighton, 2008]


The further we travel in the network, the more bottlenecks we will see



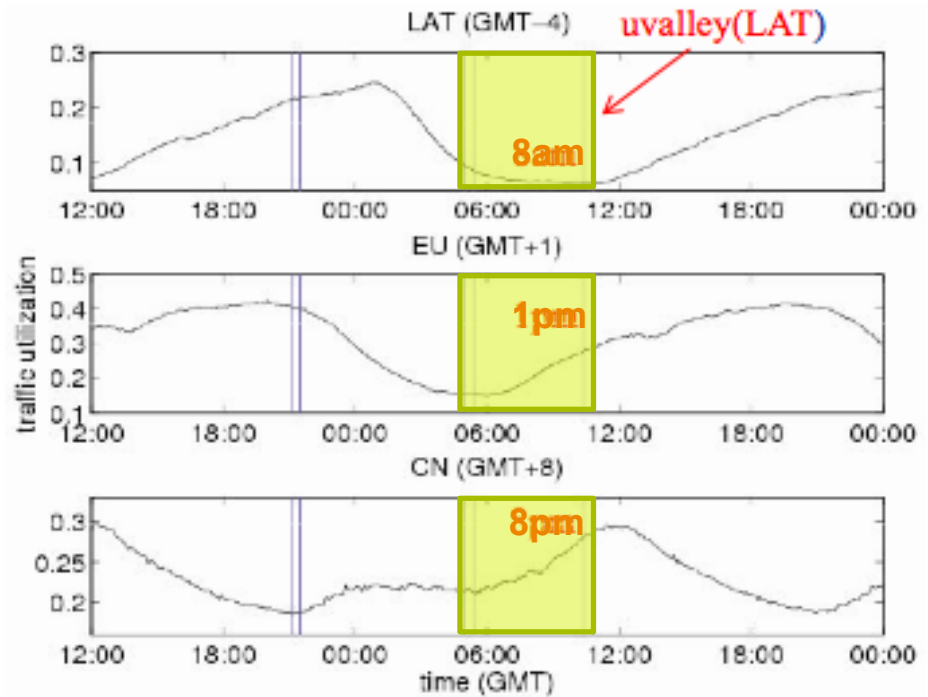
... and they are time dependant

Non-overlapping valleys

Available rate



Sender in LAT Receiver in EU or China




The Real Problem


Internet:

short burst/instantaneous ✓

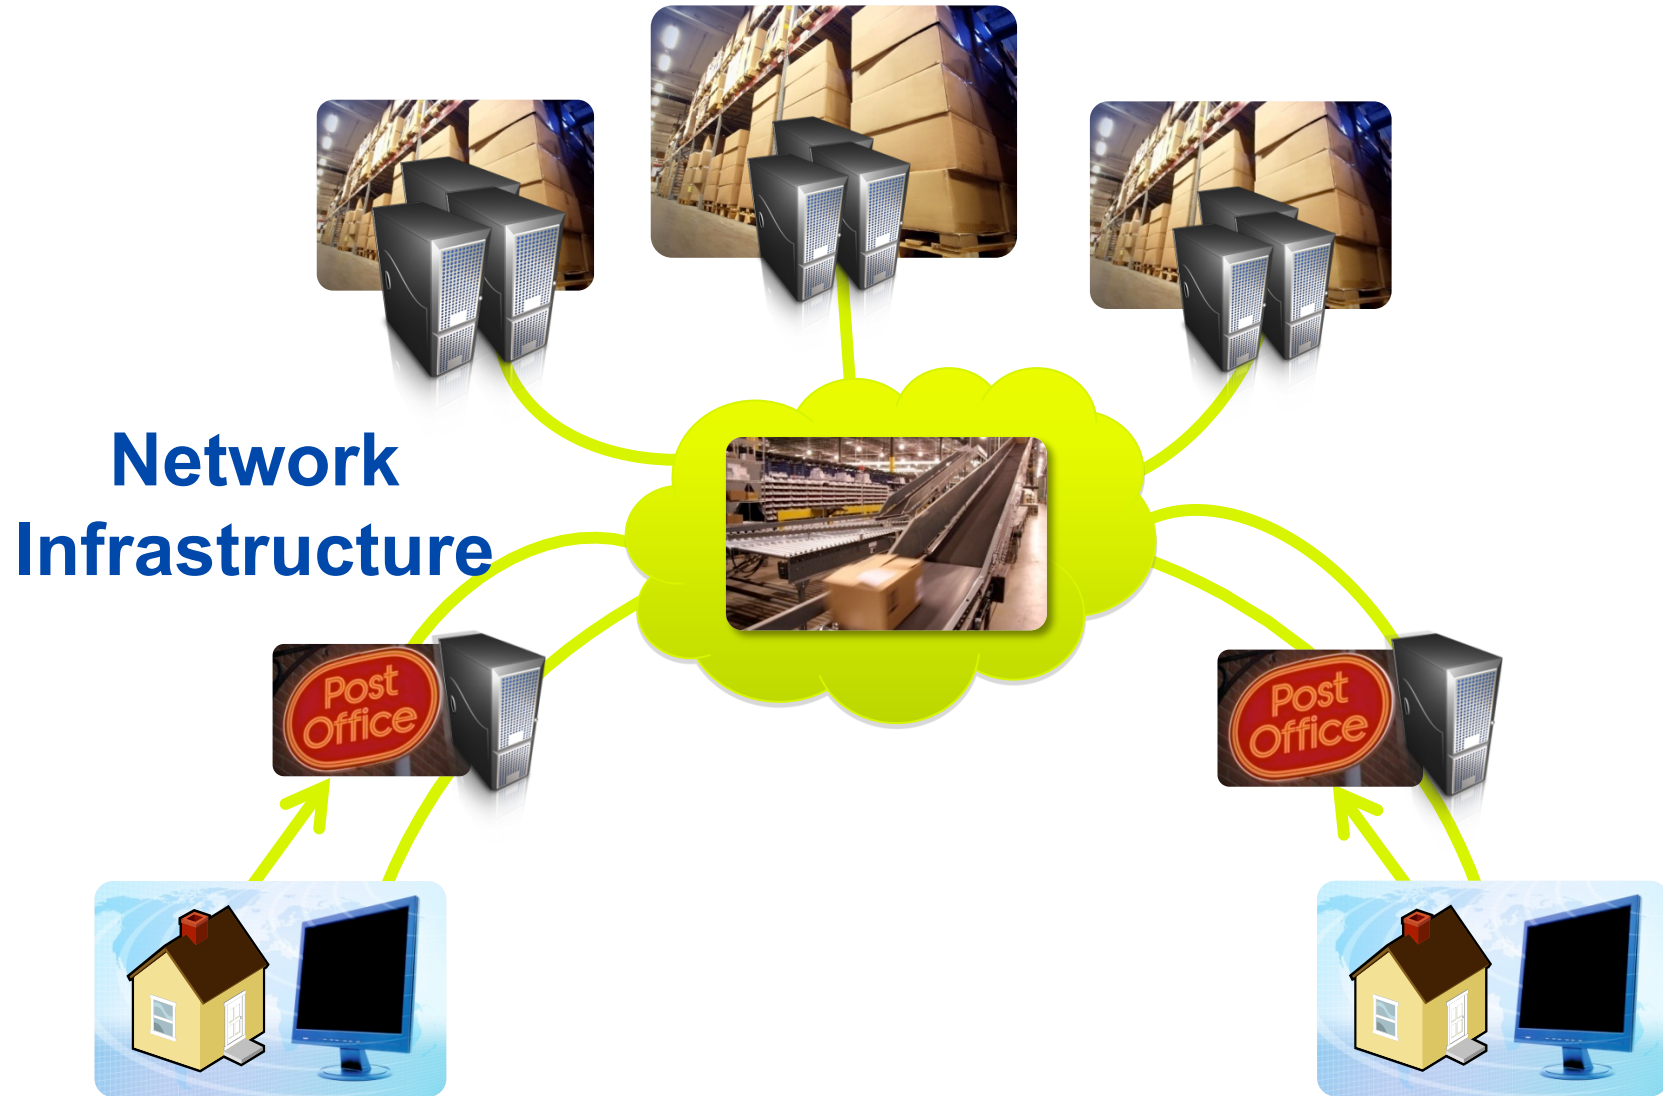
bulk/delay tolerant ✗



Current Internet



Internet Postal Service



Roadmap

- Part 1: Locating and Managing Content
- Part 2: Distributing Content
- Part 3: Data Centre Clouds
- Part 4: Clouds for Online Social Networks

Clouds: Hosting Web Content and Services

- Economies of scale
(Hamilton, 2008)

Resource	Cost (medium scale)	Cost (large scale)	Ratio
Network	\$95 / Mbps / month	\$13 / Mbps / month	~7x
Storage	\$2.20 / GB / month	\$0.40 / GB / month	~6x
Admin	≈140 servers/admin	>1000 servers/admin	~7x

- Virtualization technologies
- Off-peak capacity

... Cloud Computing


Data Centers:

100,000s of servers spread over
100,000s of square feet drawing
10 to 20MW of power





Cooling and Electricity are becoming more important than server's cost

Annual Amortized Costs in the Data Center for a 1U Server



Highly Distributed Data Centers




Challenges to host Applications

- Some candidates for distributed data centers:
 - P2P Video Delivery
 - Voice/Video Conferencing
 - Multi-player games
 - **Online Social networks...?**

Roadmap

- Part 1: Locating and Managing Content
- Part 2: Distributing Content
- Part 3: Data Centre Clouds
- Part 4: Clouds for Online Social Networks

- OSN are changing the way people interact on the Web
- But also changing its infrastructure
- The Cloud needs to become more “socially” aware



Some quick facts

- Facebook has grown from 100M to 200M in less than 8 months
- Twitter Feb to March growth of 1,230%
- The first twitter celebrity Ashton Kutcher with 1MM+ followers.
- Oprah got 100K+ followers on Twitter in 4 hours
- Nielsen Online's latest research shows that OSN is now more popular than email






The Web

P2P

OSN

The Internet




New Design Challenges

- Can the social networks predict which videos are likely to be seen by consumers, at what times, where?
- How to design infrastructure to empower OSN celebrities to have their own broadcast channel?
- How do you handle security issues? OSN celebrities could produce DDoS over any website.

Hosting Social Networks in Distributed Clouds

- One operation results in a social cascade
- LinkedIn: 22M users, Facebook with 200M+
- Data structures no longer fit in the memory of a single server
- Data partition is a must:
 - How to minimize inter data center communication
 - How to ensure consistency and small latencies



Roadmap


- Part 1: Locating and Managing Content
- Part 2: Distributing Content
- Part 3: Data Centre Clouds
- Part 4: Impact of Online Social Networks

Final Thoughts

- Things are getting more complex, is it time to rethink some designs and move into firmer grounds?
- Going forward, do we need to think of the Web infrastructure design to be more embedded with the Internet design? can developments at these two levels move independently?
- What elements are matured enough to be pushed into the lower layers and become basic services, like routing is? (e.g. locating content, content distribution)

Conclusions

- The Web has often pushed the Internet infrastructure to its limits
- Locating, Managing, Distributing content still pose challenges
- Both Social Networks and Greener Clouds will re-shape the Web infrastructure, once more...



Telefonica
