From Internet Data Centers to Data Centers in the Cloud

This case study is a short extract from a keynote address given to the Doctoral Symposium at Middleware 2009

by Lucy Cherkasova of HP Research Labs Palo Alto.

The full keynote is on the course materials page

The keynote focus is ***performance modelling***

- Data Centers Evolution
  - Internet Data Center
  - Enterprise Data Centers
  - Web 2.0 Mega Data Centers

Data Center Evolution

- Internet Data Centers (IDCs first generation – per company)
  - Data Center boom started during the dot-com bubble
  - Companies needed fast Internet connectivity and an established Internet presence
  - Web hosting and co-location facilities for company's services
  - Challenges in service scalability, dealing with flash crowds, and dynamic resource provisioning
    - New paradigm: everyone on the Internet can come to your web site!
  - Mostly static web content
    - Many results on improving web server performance, web caching, and request distribution
  - Web interface for configuring and managing devices (products sold by company)
  - New pioneering architectures such as
    - Content Distribution Network (CDN),
    - Overlay networks for delivering media content

Content Delivery Network (CDN)

- High availability and responsiveness are key factors for business Web sites
- “Flash Crowd” problem
- Main goal of CDN’s solution is
  - overcome server overload problem for popular sites,
  - minimize the network impact in the content delivery path.
- CDN: large-scale distributed network of servers,
  - Surrogate servers (proxy caches) are located closer to the edges of the Internet a.k.a. edge servers
- Akamai is one of the largest CDNs
  - 56,000 servers in 950 networks in 70 countries
  - Deliver 20% of all Web traffic

Retrieving a Web Page

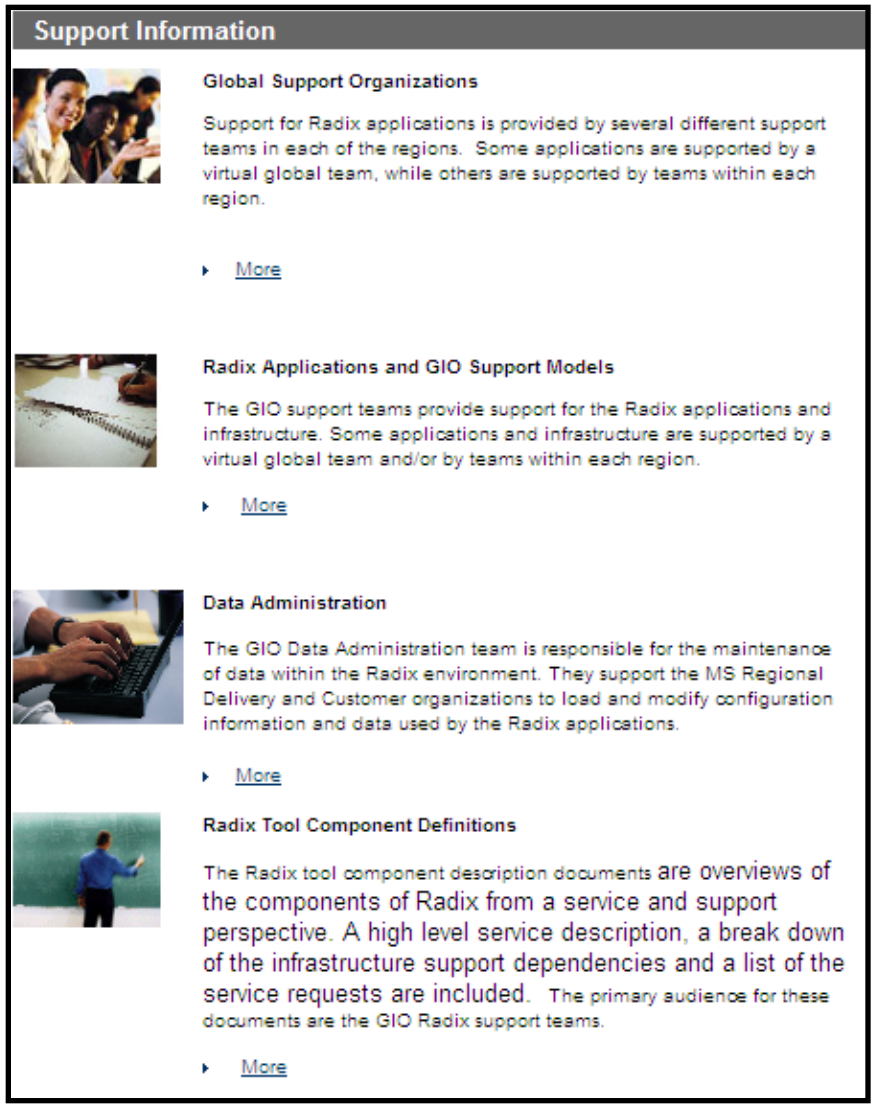

Web page is a composite object:

- HTML file is delivered first
- Client browser parses it for embedded objects
- Send a set of requests for these embedded objects
- Typically, 80% or more of bytes of a web page are images
- 80% of the page can be served by a CDN.

CDN's Design

- Two main mechanisms
  - URL rewriting
    - <img src =http://www.xyz.com/images/foo.jpg>
    - <img src =http://akamai.xyz.com/images/foo.jpg>
  - DNS redirection
    - Transparent, does not require content modification
    - Typically employs two-level DNS lookup to choose most appropriate edge server (***name -> list of edge servers, selected list item -> IP address***)

CDN Architecture

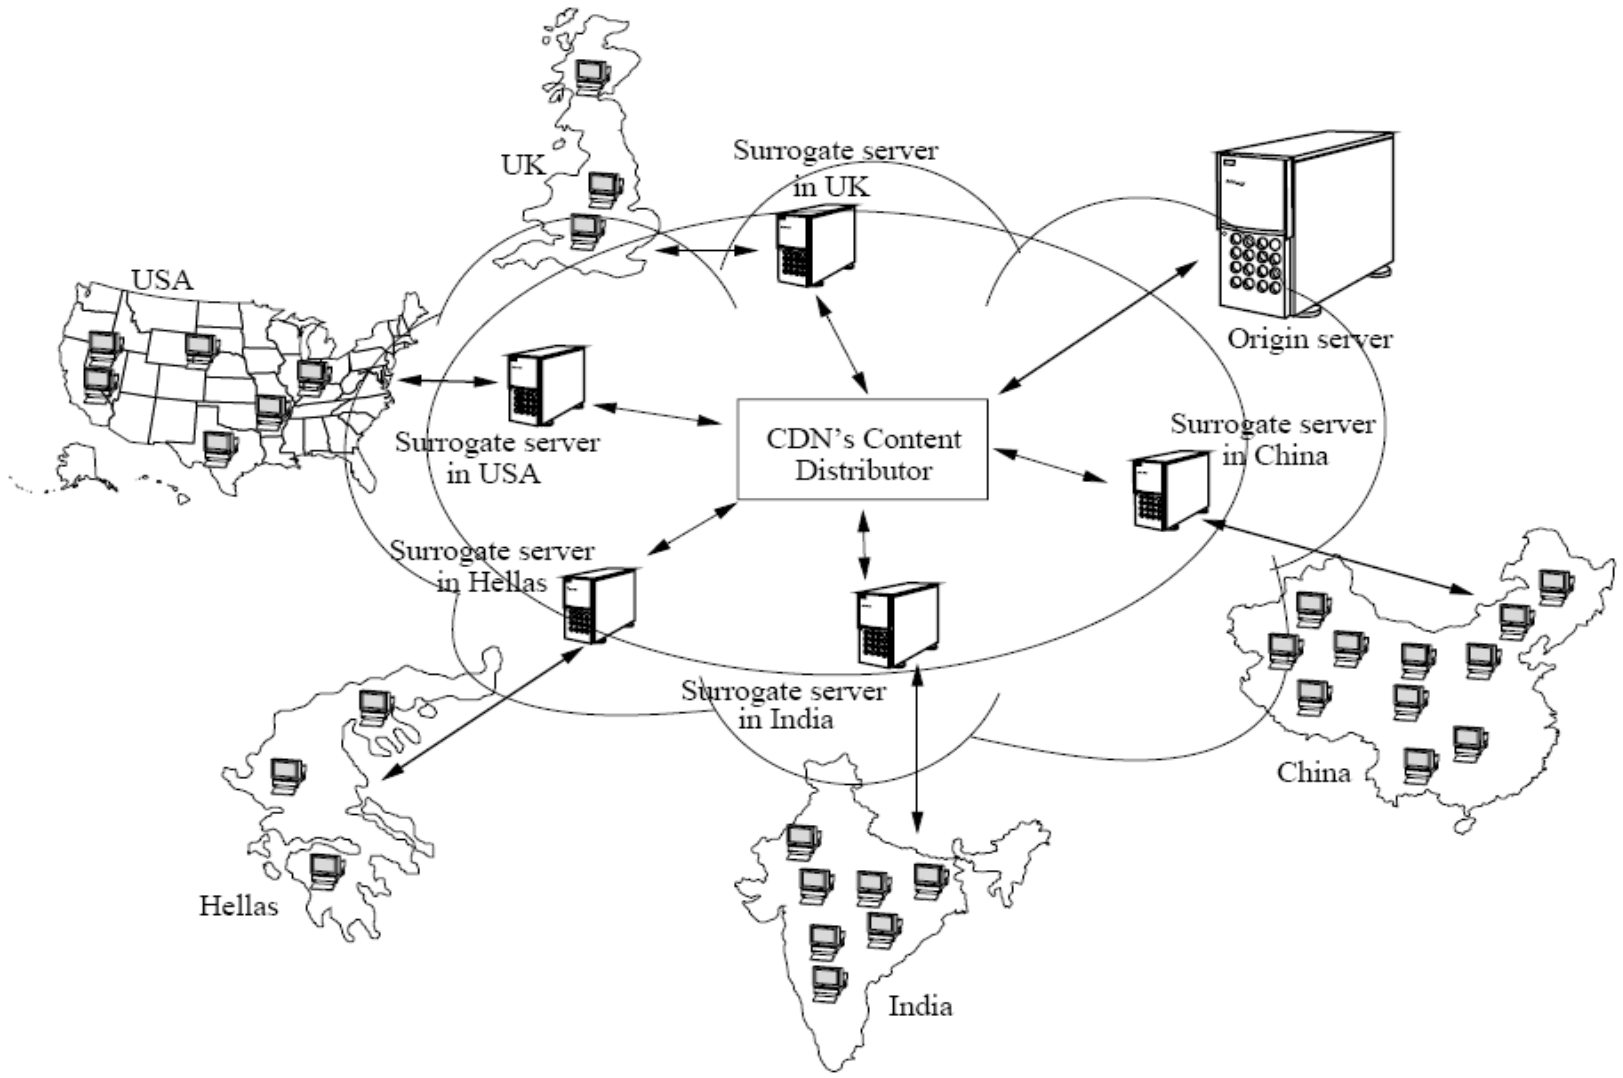

CDN Research Problems

- Efficient large-scale content distribution
  - large files,
  - video on demand, streaming media

    low latency, real-time requirement
    - FastReplica for CDNs
    - BitTorrent (general purpose)
    - SplitStream (multicast, video streaming)

FastReplica: Distribution Step

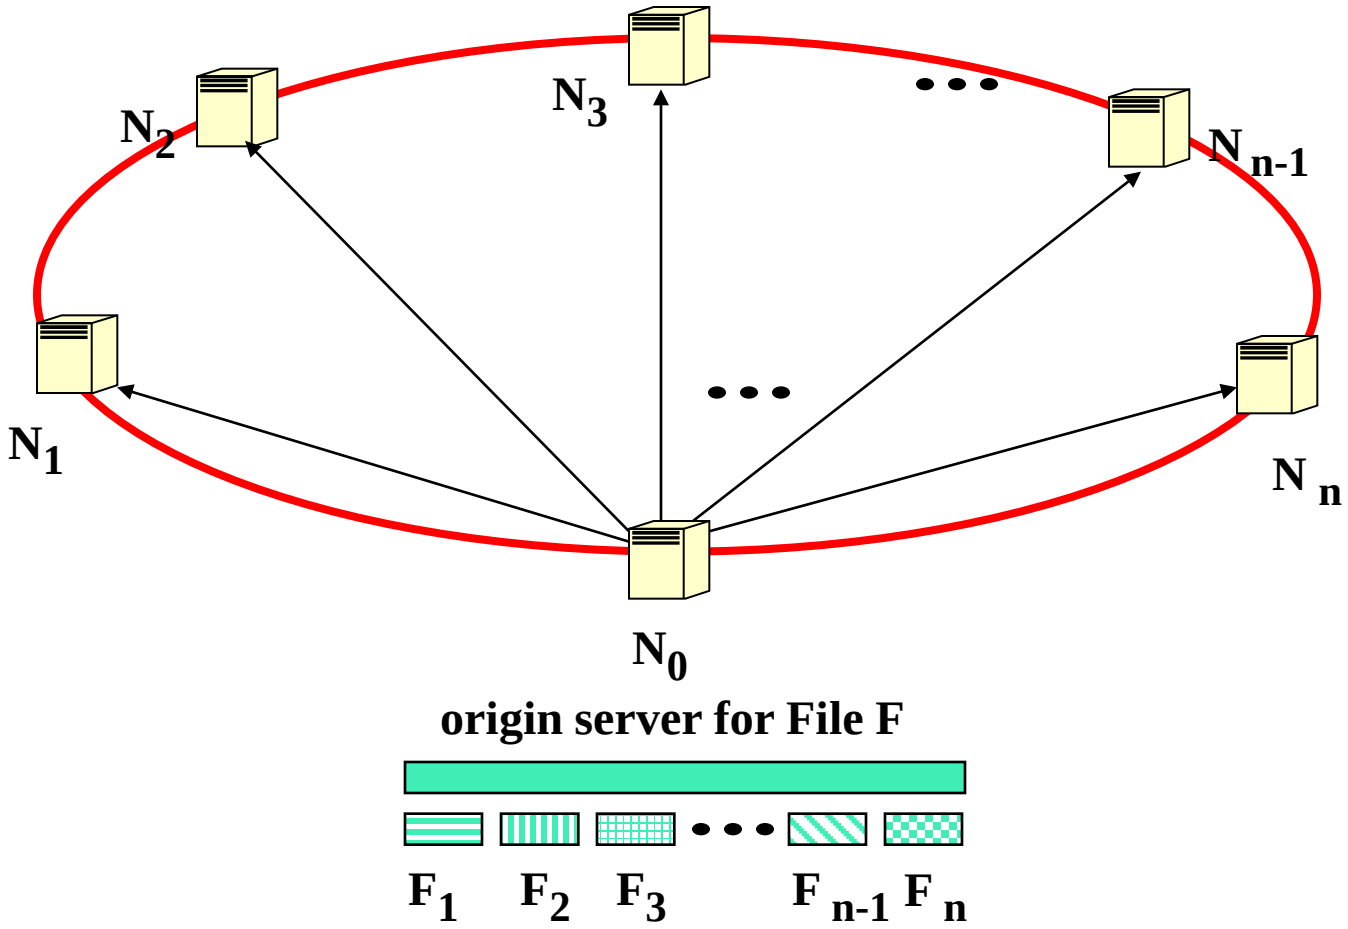

L. Cherkasova, J. Lee. *FastReplica: Efficient Large File Distribution within Content Delivery Networks*
Proc. of the 4th USENIX Symp. on Internet Technologies and Systems (USITS'2003).

FastReplica: Collection Step

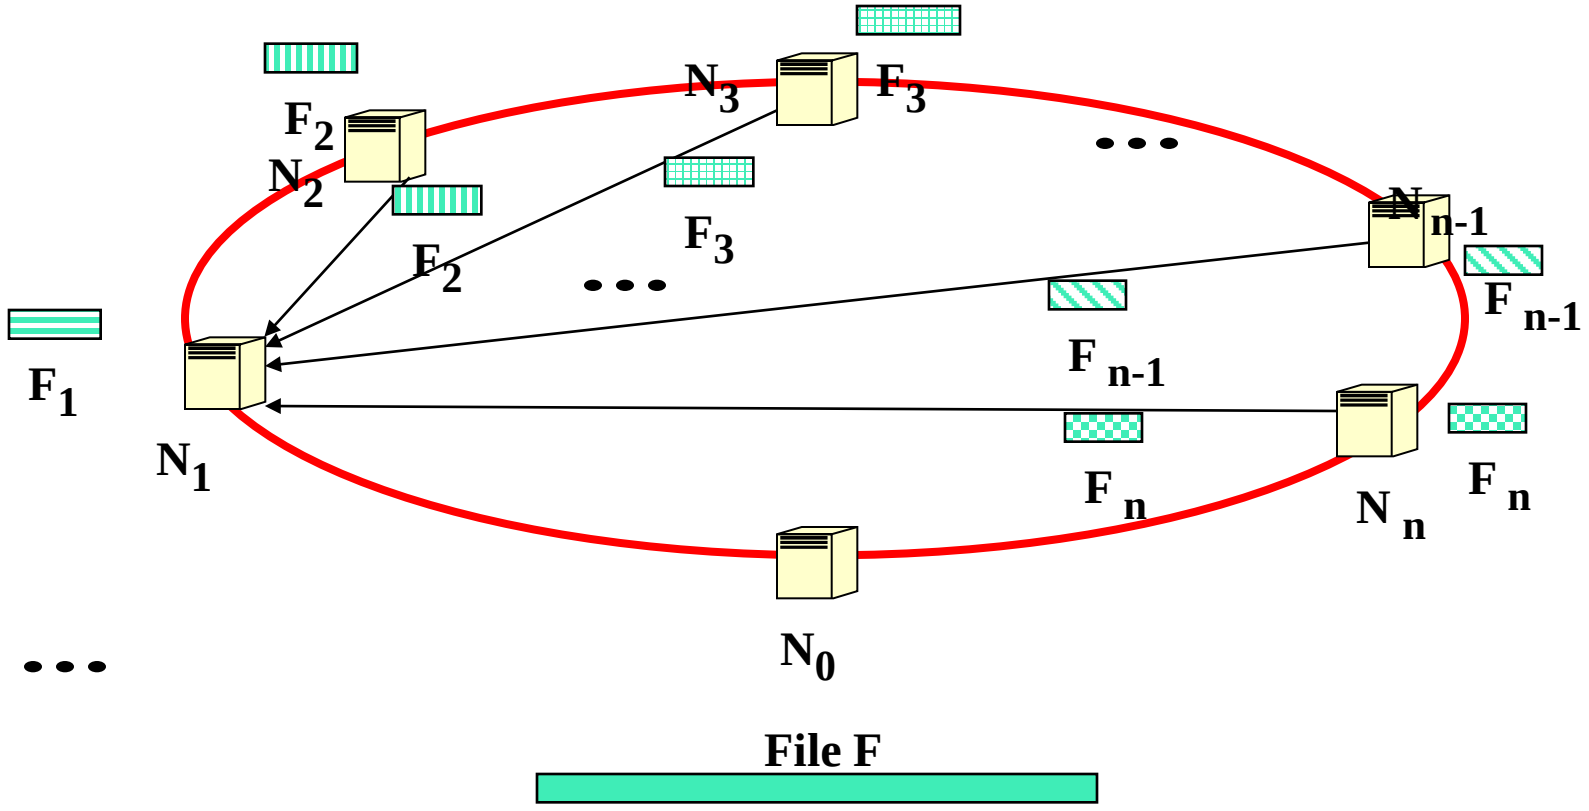

Remaining Research Problems

Some (2009) open questions:

- Optimal number of edge servers and their placement
  - Two different approaches:
    - *Co-location:* placing servers closer to the edge (Akamai)
    - *Network core*: server clusters in large data centers near the main network backbones (Limelight and AT&T)
- Content placement
- Large-scale system monitoring and management
  - to gather evidence as a basis for design decisions

Data Center Evolution

- Enterprise Data Centers
  - New application design: multi-tier applications - database integration, see next slide
  - Many traditional applications, e.g. HR, payroll, financial, supply-chain, call-desk, etc, are re-written using this paradigm.
  - Many different and complex applications
  - *Trend: Everything as a Service*
    - Service oriented Architecture (SOA)
  - Dynamic resource provisioning within a large cluster
  - Virtualization (datacenter middleware)
  - Dream of Utility Computing:
    - Computing-on-demand (IBM)
    - Adaptive Enterprise (HP)

Multi-tier Applications

- Enterprise applications:
  - Multi-tier architecture is a standard building block

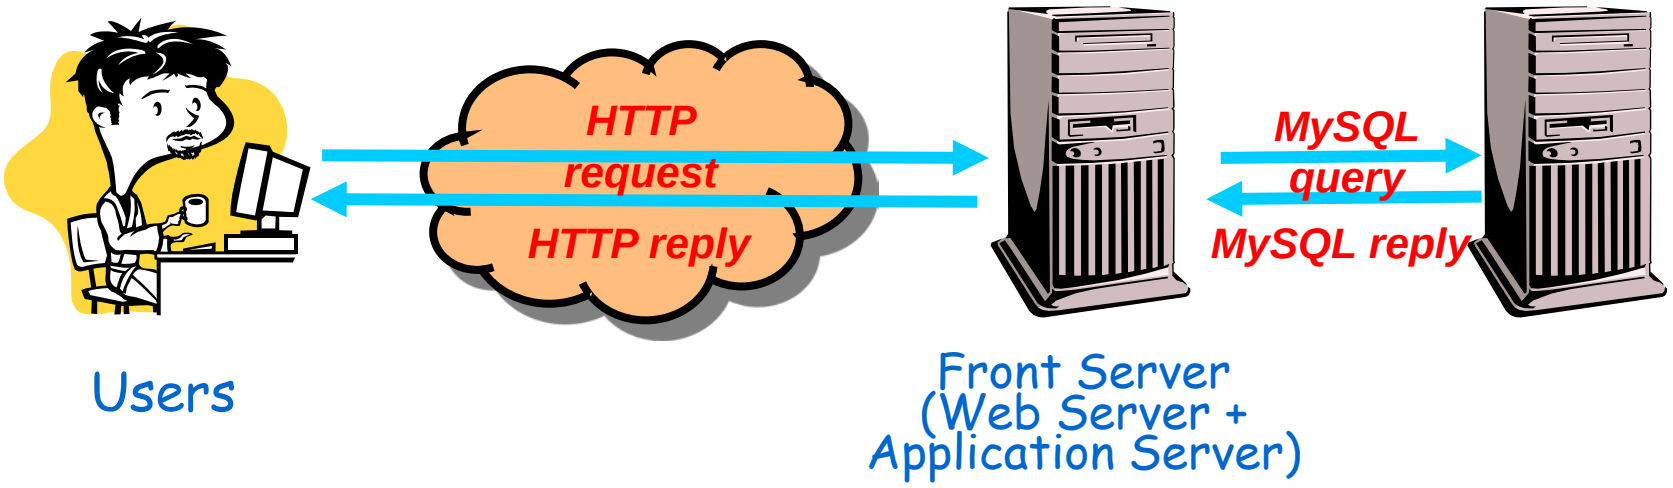

Example: Units of Client/Server Activity

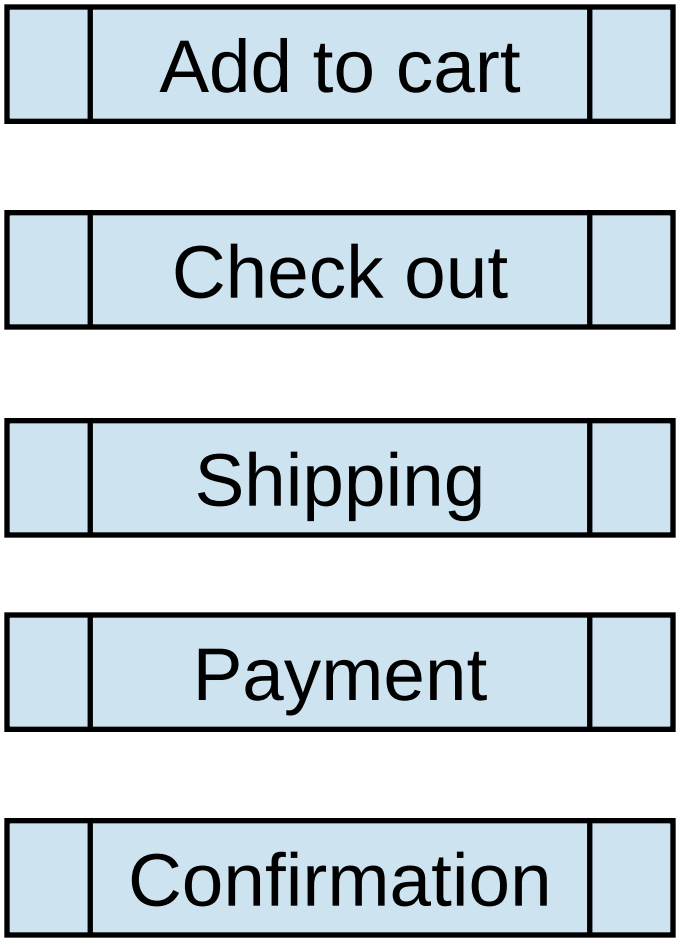

- Session:

  A sequence of individual transactions issued by the same client

- Concurrent Sessions = Concurrent Clients

- Think time:

  The interval from a client receiving a response to the client sending the next transaction

Data Growth

- Unprecedented data growth:
  - The amount of data managed by today's Data Centers quadruples every 18 months
- New York Stock Exchange generates about 1 TB of new trade data each day.
- Facebook hosts ~10 billion photos (1 PB of storage).
- The Internet Archive stores around 2PB, and it is growing at 20TB per month
- The Large Hadron Collider (CERN) will produce ~15 PB of data per year.

Big Data

- IDC estimate the size of “digital universe” :
  - 0.18 zettabytes in 2006;
  - 1.8 zettabytes in 2011 (10 times growth);
- A zettabyte is 10^{21} bytes, i.e.,
  - 1,000 exabytes or
  - 1,000,000 petabytes
- Big Data is here
  - Machine logs, RFID readers, sensors networks, retail and enterprise transactions
  - Rich media
  - Publicly available data from different sources
- New challenges for storing, managing, and processing large-scale data in the enterprise (information and content management)
  - Performance modeling of new applications

Worldwide Server Market:

Cost of Management and Power Ramps Dramatically

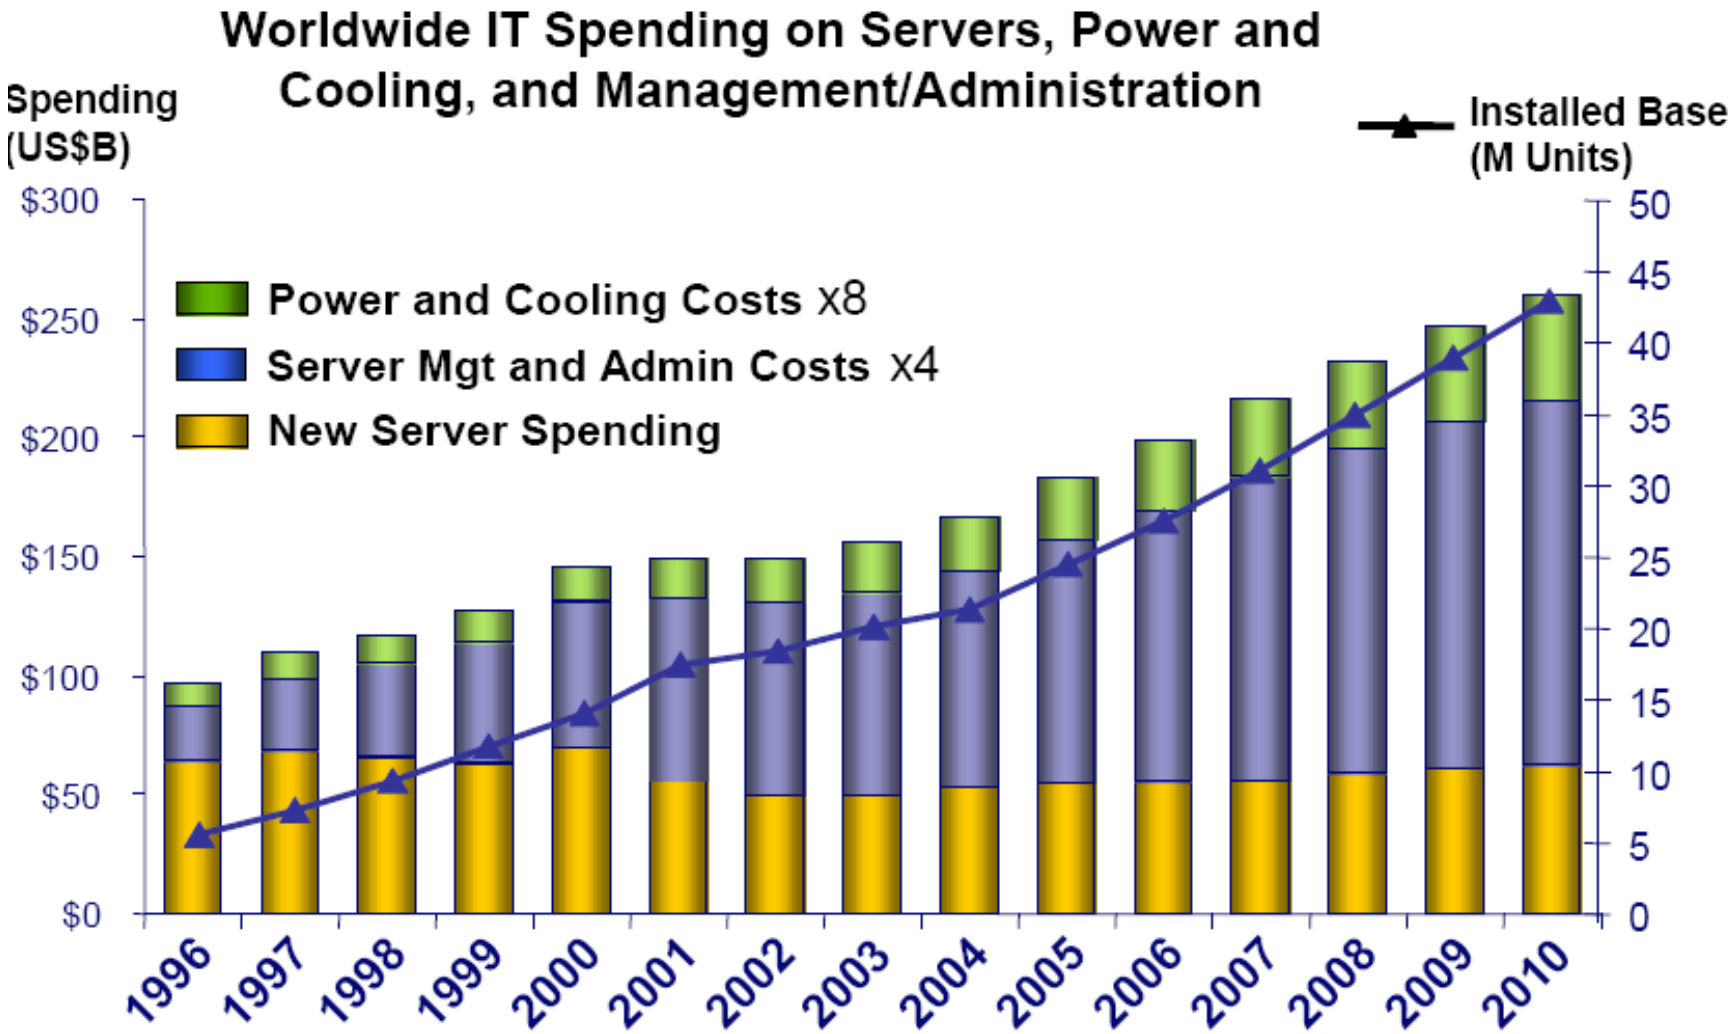

Source: IDC, 2008

Data Center Evolution

- Data Center in the Cloud
  - Web 2.0 Mega-Datacenters: Google, Amazon, Yahoo
  - Amazon Elastic Compute Cloud (EC2)
  - Amazon Web Services (AWS) and Google AppEngine
  - New class of applications related to ***parallel processing*** of large data
  - Google's Map-Reduce framework (with the open source implementation Apache Hadoop)
    - *Mappers* do the work on data slices, *Reducers* process the results
    - Handle node failures and restart failed work
  - One can rent ones own Data Center in the Cloud on a "pay-per-use" basis
  - Cloud Computing: Software as a Service (SaaS) + Utility Computing