

LARGE SCALE INTERNET SERVICES

2110414 Large Scale Computing Systems
Natawut Nupairoj, Ph.D.

Outline

2

- Overview
- Background Knowledge
- Architectural Case Studies
- Real-World Case Study

3

Overview

Overview

4

- Internet services become very essential and popular
 - Google serves hundreds of millions of search requests per day
- Main requirements
 - Availability
 - Scalability



Internet Service Application Characteristics

5

Table 1. Comparison of Internet service application characteristics with those of traditional applications.

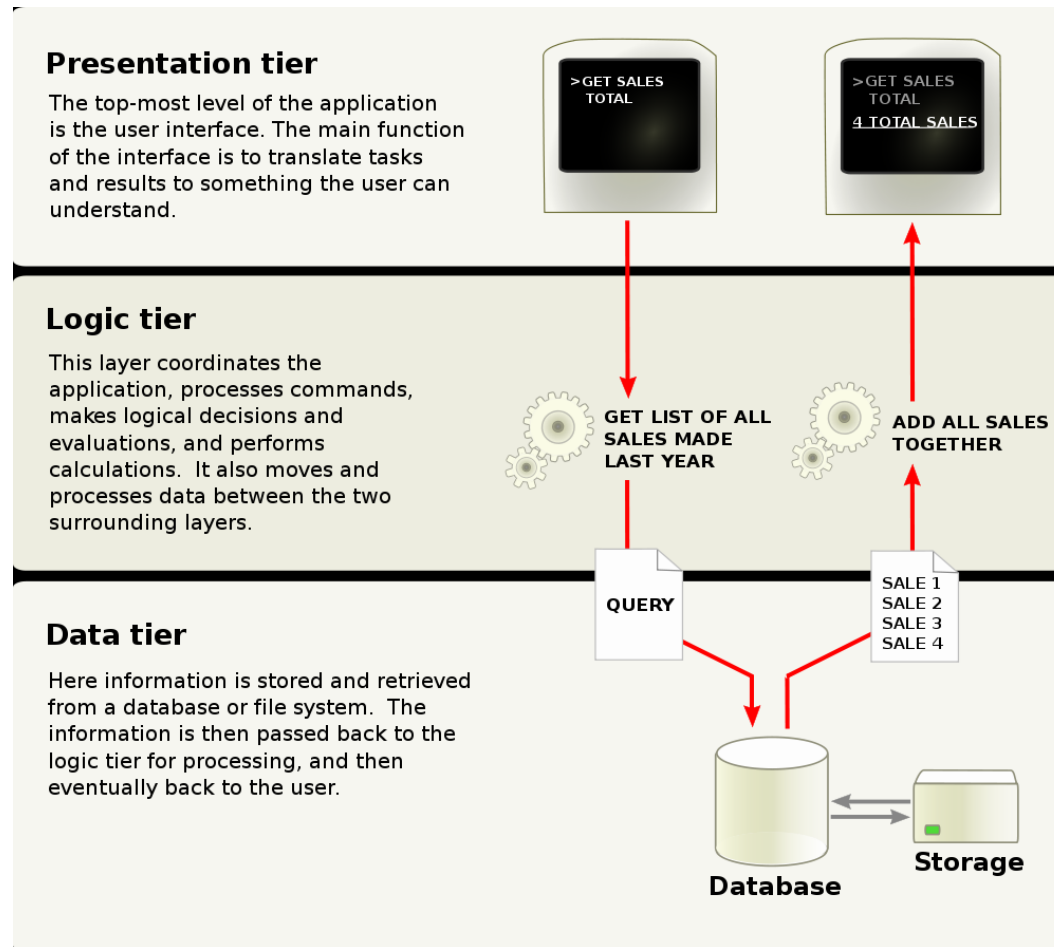
Characteristic	Traditional desktop applications	Traditional high-availability applications	Large-scale Internet service applications
Dominant application	Productivity applications, games	Database, enterprise messaging	E-mail, search, news, e-commerce, data storage
Typical hardware platform	Desktop PC	Fault-tolerant server or failover cluster	Clusters of hundreds to thousands of cheap PCs, often geographically distributed
Software model	Off-the-shelf, standalone	Off-the-shelf, multitier	Customized, multitier
Software release frequency	Months	Months	Days or weeks
Networking environment	None (standalone)	Enterprise-scale	Within cluster, on the Internet between data centers, and to and from user clients
Operator	end user	Corporate IT staff	Service operations staff, data center/collocation site staff
Typical physical environment	Home or office	Corporate machine room	Collocation facility
Key metrics	Functionality, interactive latency	Availability, throughput	Availability, functionality, scalability, manageability, throughput

6

Background Knowledge

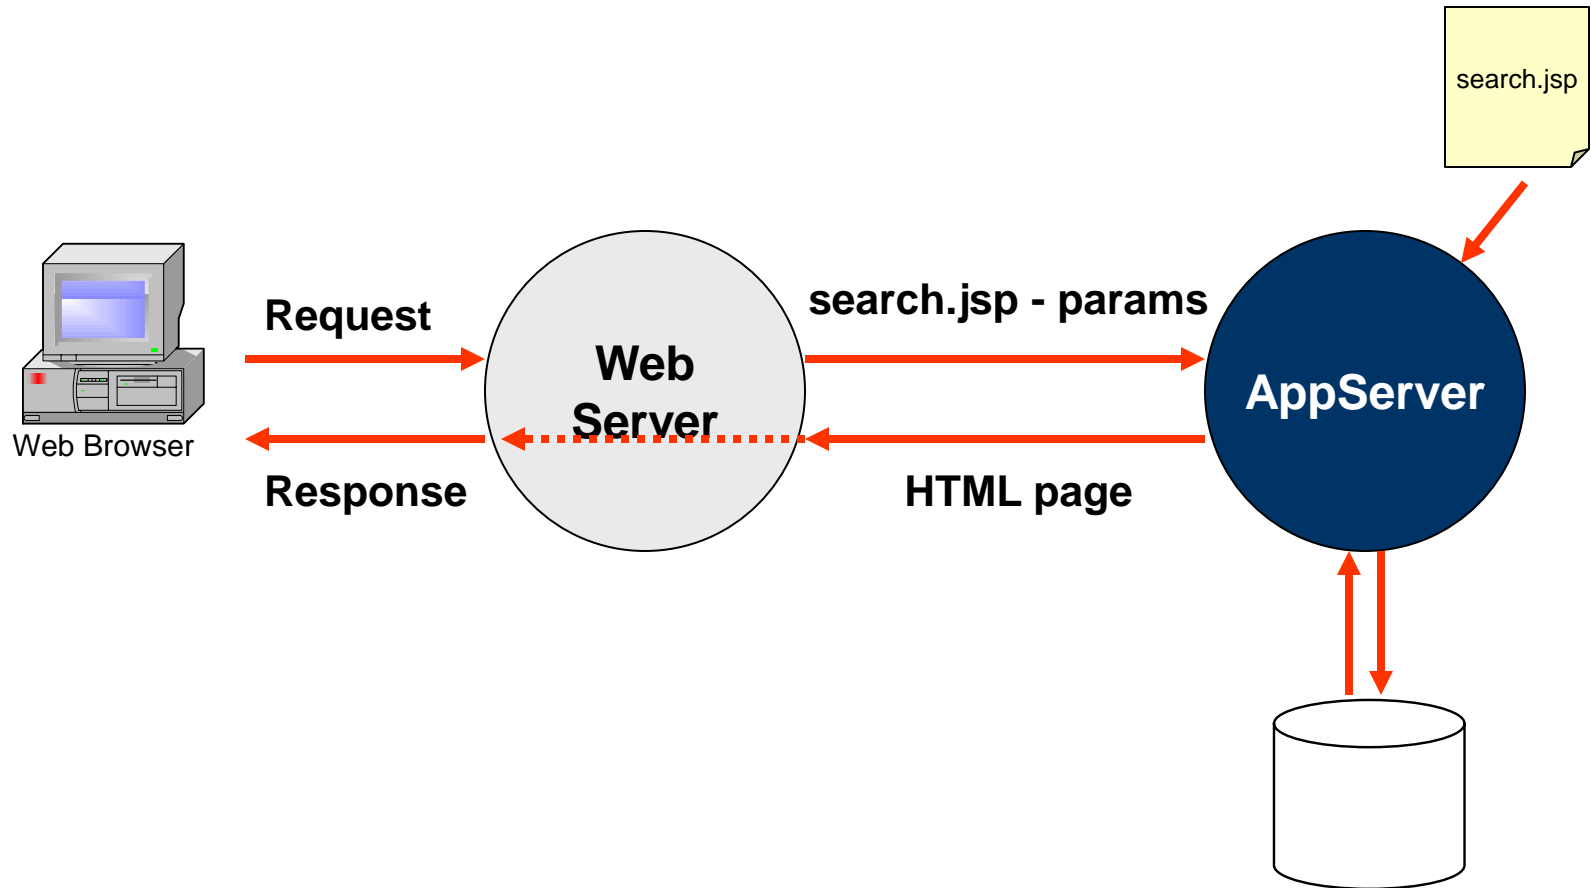
Multi-Tier Architecture

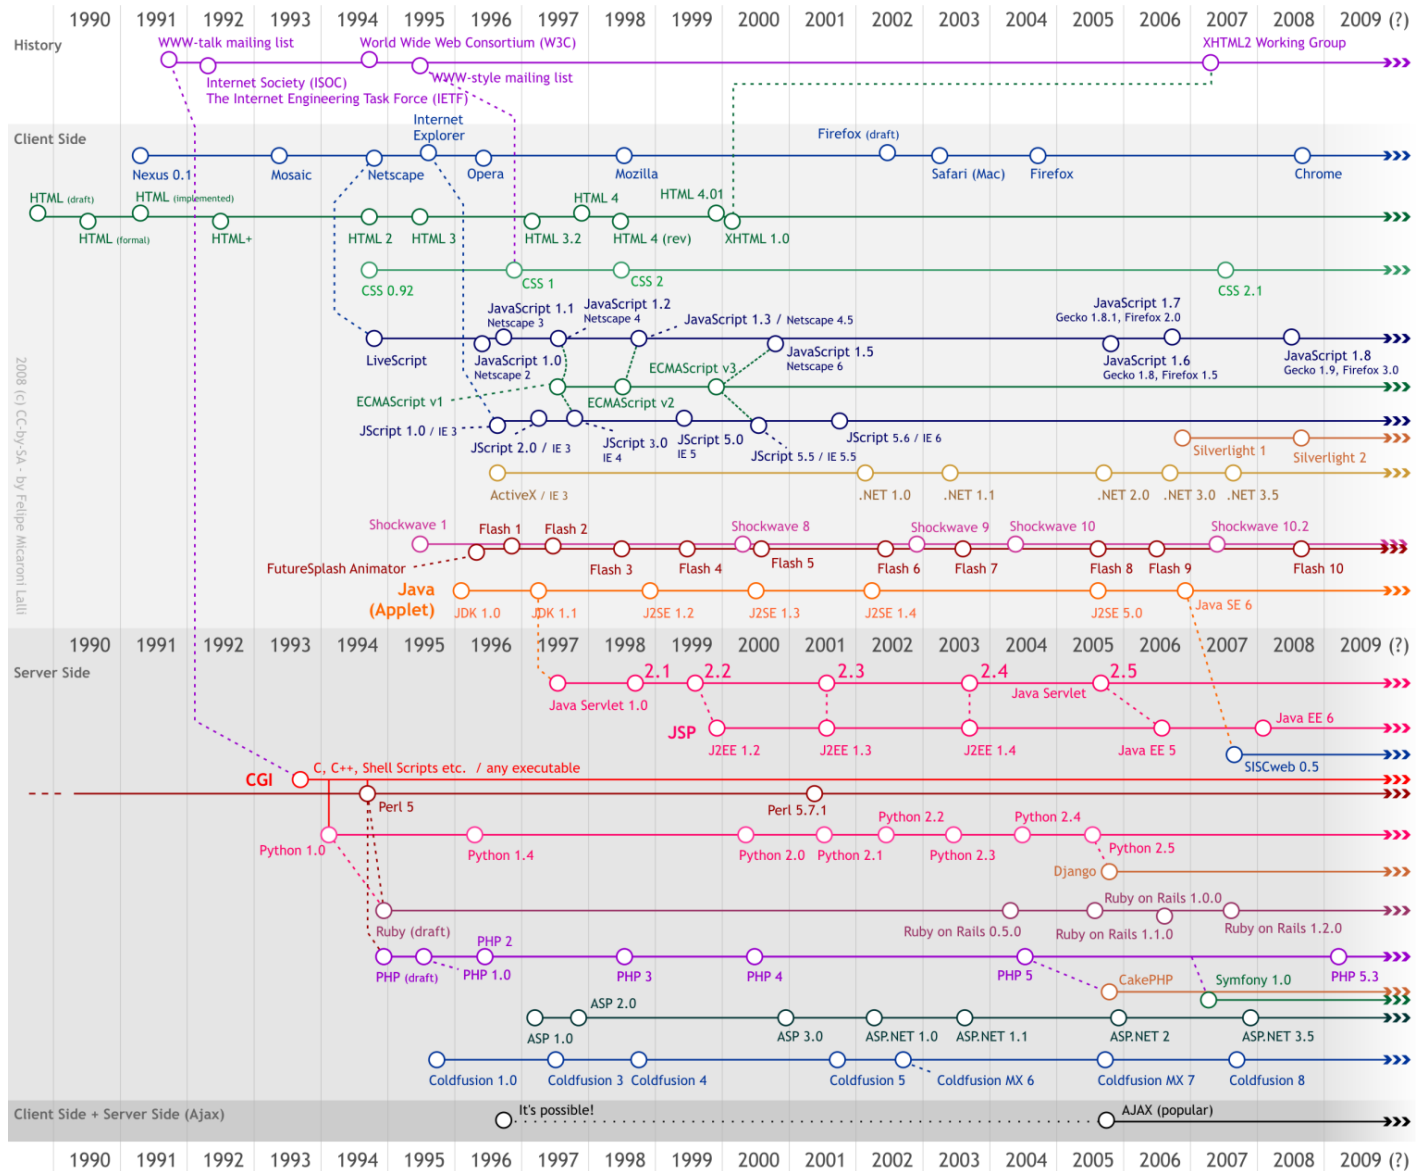
7



Web Based Architecture Revisited

8






System Availability

- How to ensures a certain absolute degree of operational continuity during a given measurement period
- Availability includes ability of the user community to access the system, whether to submit new work, update or alter existing work, or collect the results of previous work
- Model of Availability
 - ▣ Active-Standby: HA Cluster or Failover Cluster
 - ▣ Active-Active: Server Load Balancing

HA Cluster



- Redundant servers and other components
 - ▣ Only one server is active (master)
 - ▣ One server is standing-by
 - ▣ Shared storages
- Pro:
 - ▣ Simple
 - ▣ Half software license costs
- Con:
 - ▣ Double hardware cost with single performance


Server Load Balancing

12

- Spread work between two or more computers, network links, CPUs, hard drives, or other resources, in order to get optimal resource utilization, throughput, or response time
- Approaches
 - DNS Round-Robin
 - Reverse Proxy
 - Load Balancer

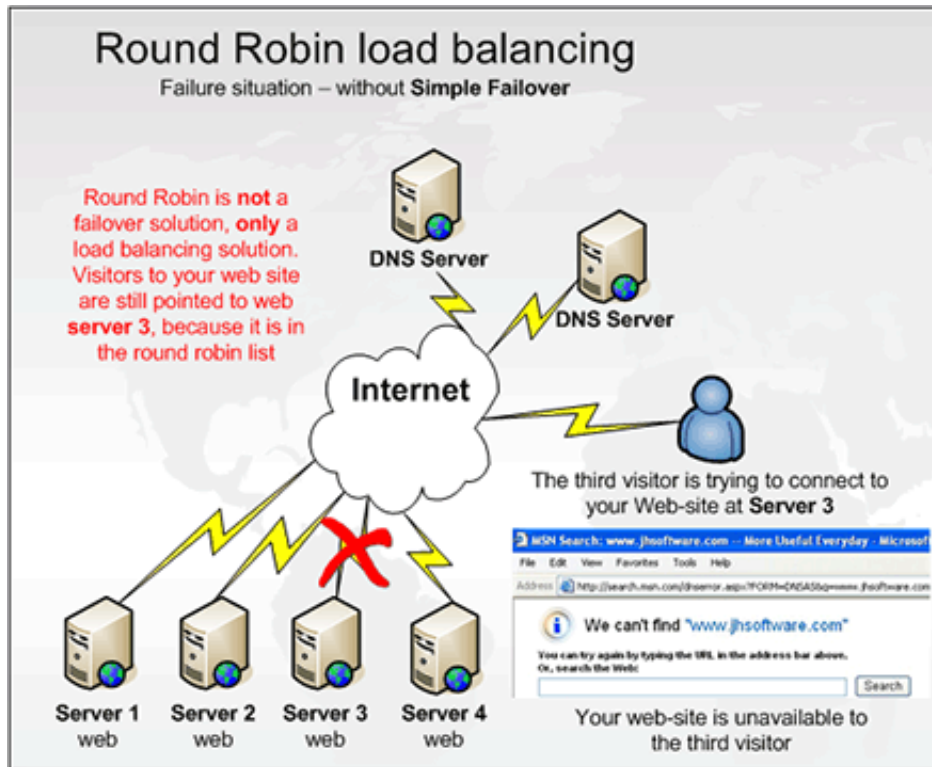
DNS Round-Robin

13



DNS Round-Robin


14



- Pro:
 - ▣ Inexpensive
- Con:
 - ▣ Load distribution, but not high availability
 - ▣ Problem with DNS caching


Reverse Proxy

15



Server Load Balancing

16



- Special equipment “load balancer” to distribute request to servers
- Clients will see only single “virtual” host based on “virtual” IP

Stateful vs. Stateless Servers

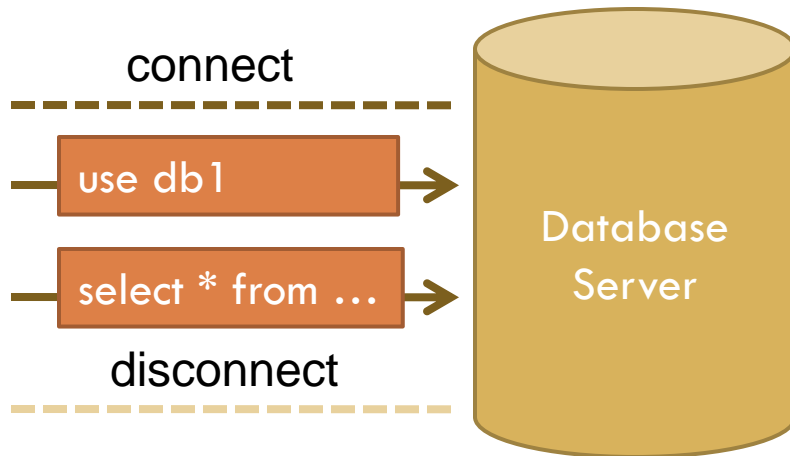
17

- Stateful server
 - ▣ server maintains some persistent data
 - ▣ Allow current request to relate to one of the earlier requests, “session”

- Stateless server
 - ▣ server does not keep data
 - ▣ A request is independent from earlier requests
 - ▣ Example: Web server, NFS

Stateful Servers

18

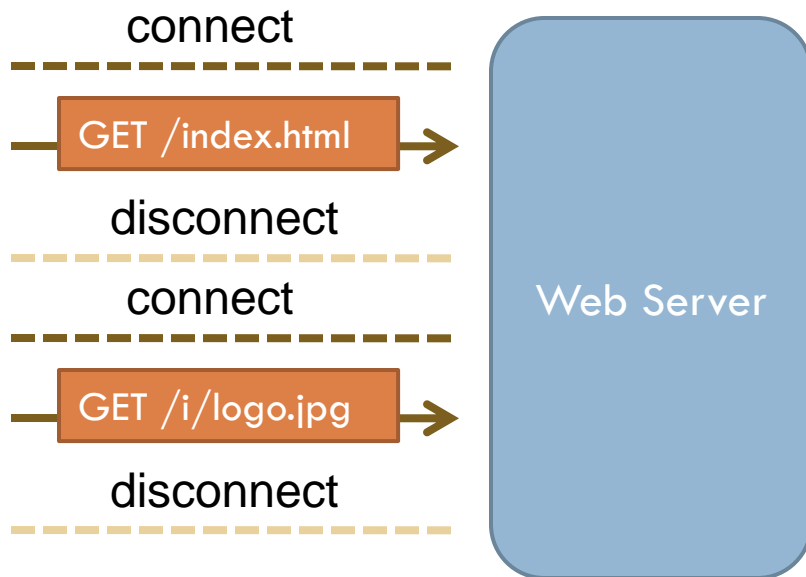


- Example: Database, FTP

- Server has to maintain some “session” information of each connection
 - ▣ Current request may depend on previous requests
- Consume server’s resources (memory, TCP port, etc.)
- Lead to limit number of clients it can service
- If connection is broken, the service is interrupted

Stateless Servers

19



- Server does not maintain information of each connection
- Connect-request-reply-disconnect cycle
- Consume less server's resources
- Lead to large number of clients it can service

- Example: Web server, NFS

Web Caching

20




- Utilize the fact that LAN has more bandwidth and less accessing latency than WAN

$$t = \text{accessing latency} + \text{data size} / \text{bandwidth}$$

- Web pages usually have some “popularity”
 - ▣ User usually goes back-and-forth between pages
 - ▣ Users tend to share the same interest (fashion)

Web Page Popularity


21



Source: <http://www.useit.com/alertbox/zipf.html>

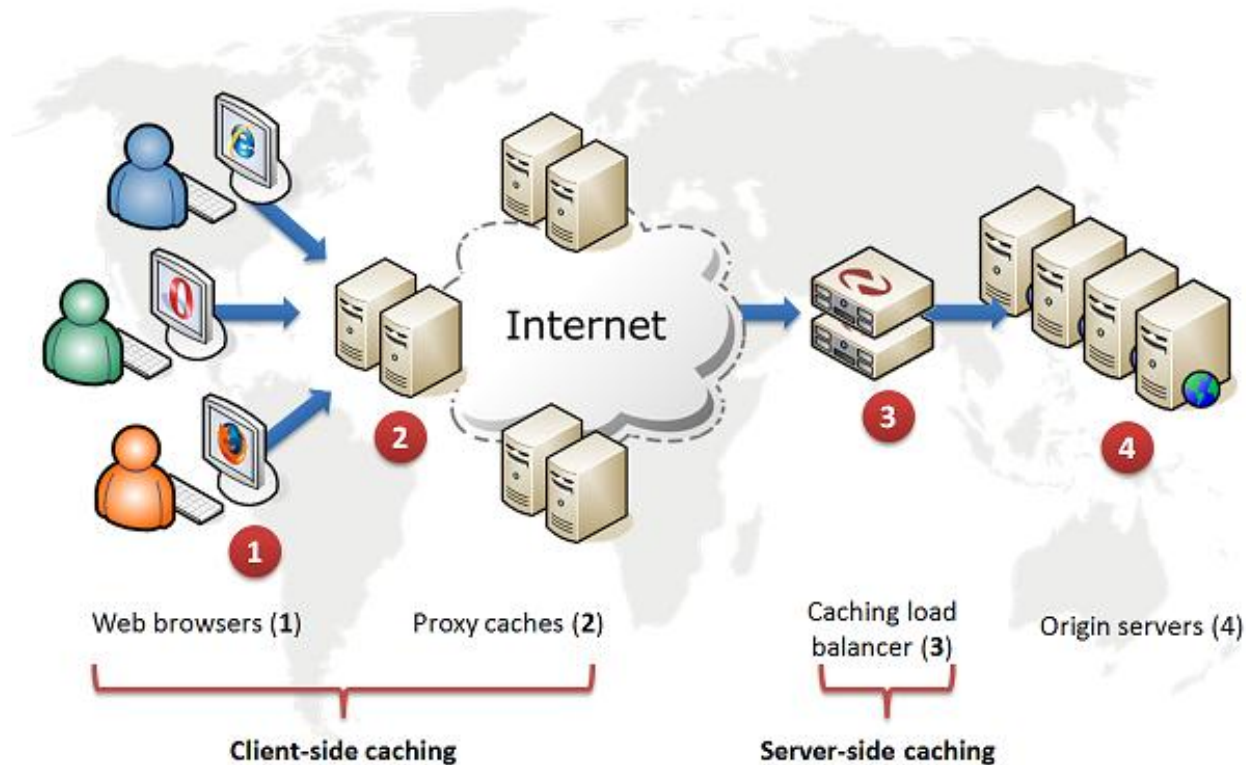
Web Caching Mechanism

22



Web Caching Location

23



Source: http://knowledgehub.zeus.com/articles/2009/08/05/cache_your_website_for_just_a_second

Architectural Case Studies

D. Oppenheimer and D. Patterson, “Architecture and Dependability of Large-Scale Internet Services”, IEEE Internet Computing, Sept-Oct 2002

Case Studies

25

Table 2. Characteristics of the large-scale Internet services examined.

Characteristic	Online	Content	ReadMostly
Hits per day	~100 million	~7 million	~100 million
Number of machines	~500, at two data centers	~500, at four data centers plus client sites	>2,000, at four data centers
Hardware	Sparc and x86	x86	x86
Operating system	Solaris	Open-source x86	Open-source x86
Relative read/write ratio	High	Medium	Very high (and users update very little data)

- Online - an online service/Internet portal (Hotmail)
- Content - a global content-hosting service (File sharing)
- ReadMostly - a high-traffic Internet service with a very high read-to-write ratio (Wikipedia)

Site Architecture

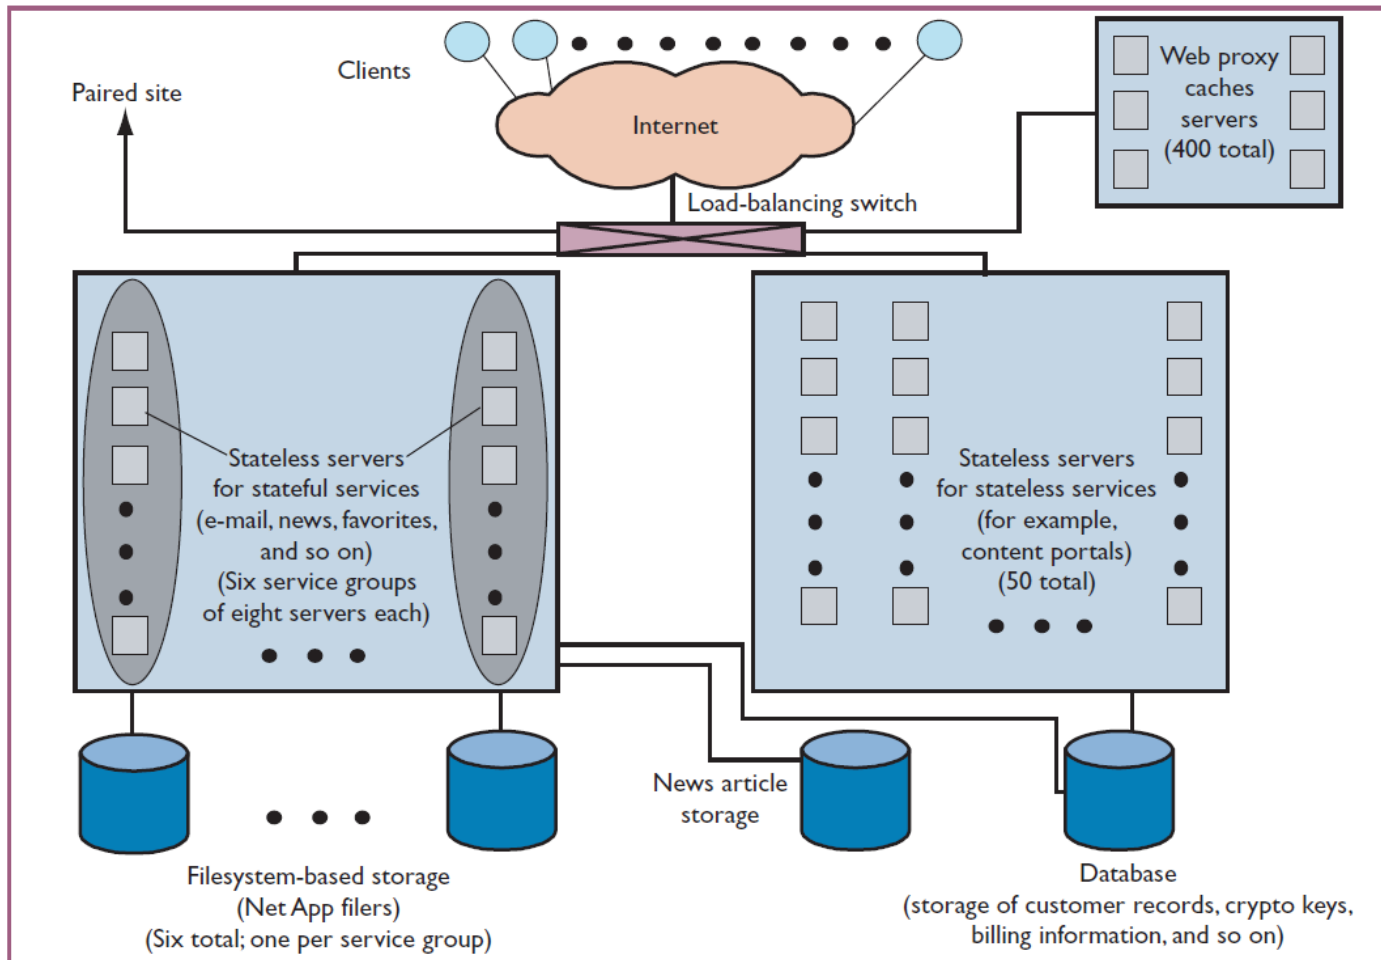
26

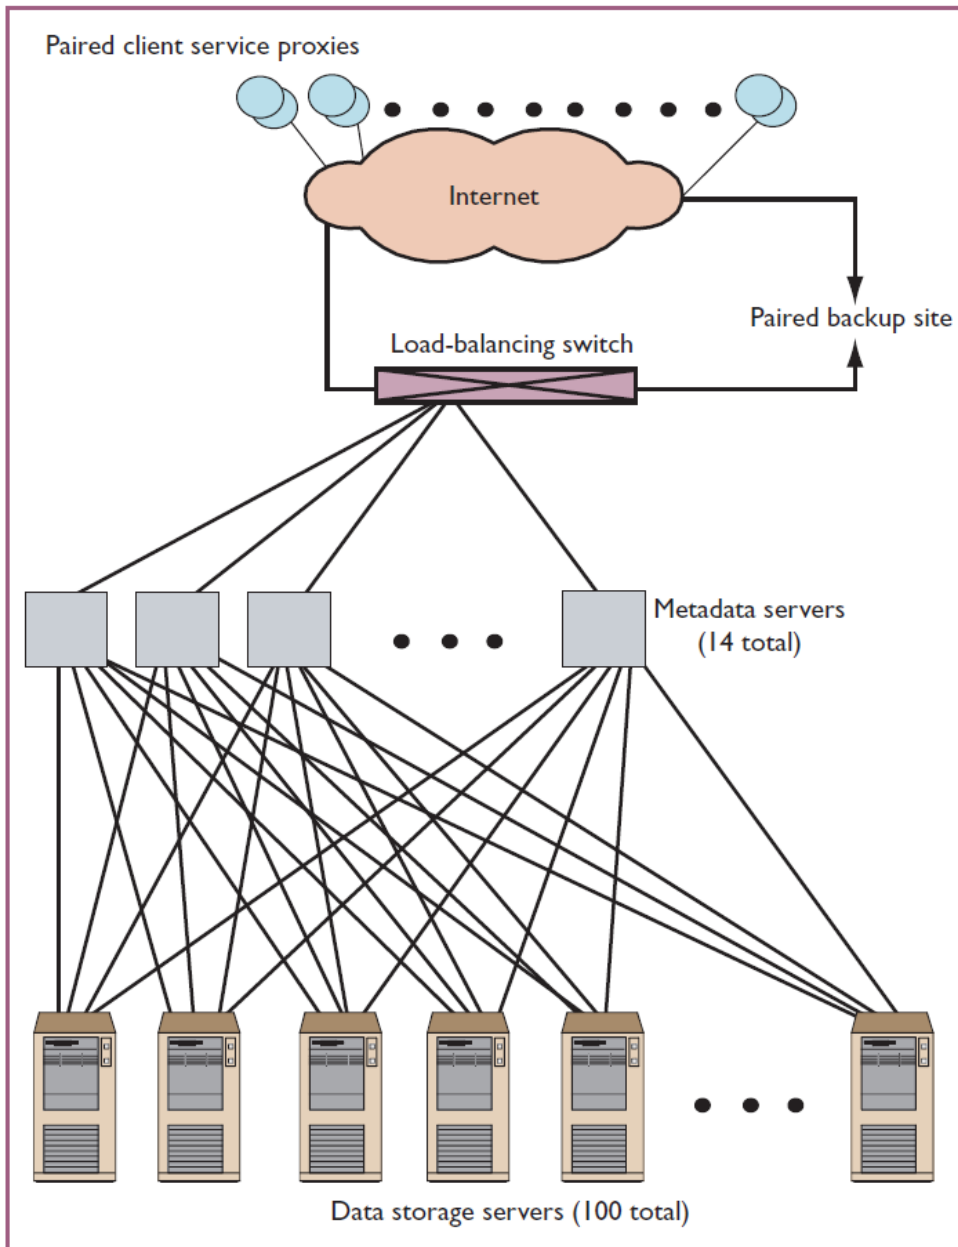
- Load balancing servers
- Front-end servers
 - ▣ Run stateless codes to service requests and gather data from back-end servers
 - ▣ Web server / AppServer
- Back-end servers
 - ▣ Provide persistent data (databases, files, emails, user profiles)
 - ▣ Should utilize RAID-based storages

Online Site

Front-end: functional partitioned
Back-end: single file, single database

27



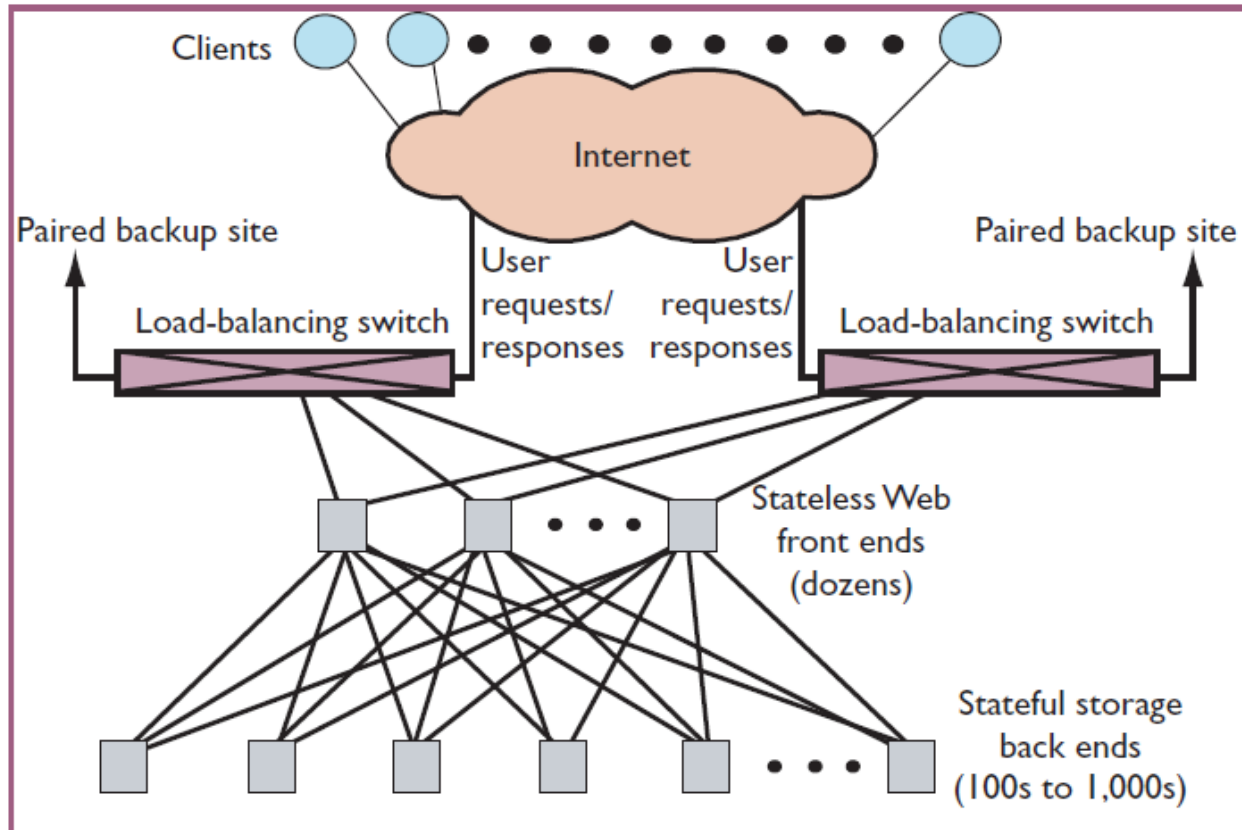


Front-end: all the same
Back-end: data partitioned

ReadMostly

Front-end: all the same
Back-end: full replication

29



Real-World Case Study: eBay

R. Shoup and D. Pritchett,

“The eBay Architecture”, SD Forum 2006

eBay

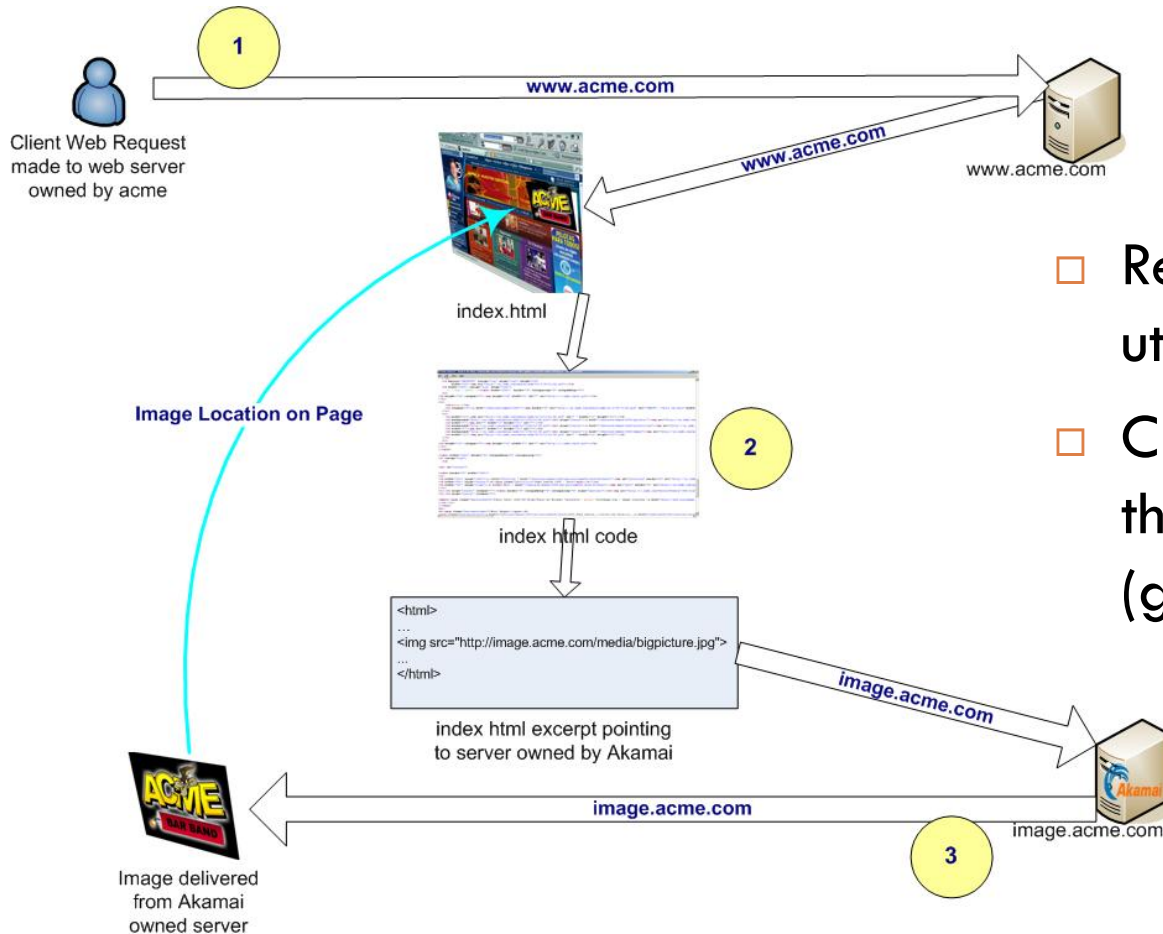
31

- Lots of workloads
 - 212 millions registered users
 - 1 billion page views a day
 - 2 petabytes of data
- Large number of servers
 - 15,000 AppServers (IBM WebSphere)
 - 100 database servers (Oracle)
 - Utilize Akamai (CDN) for static contents

CDN: Akamai

Source: http://en.wikipedia.org/wiki/Akamai_Technologies

32



- Reduce bottlenecks by utilizing geographic
- Client gets contents from the nearest servers (geographically)

eBay Architecture Design Principles

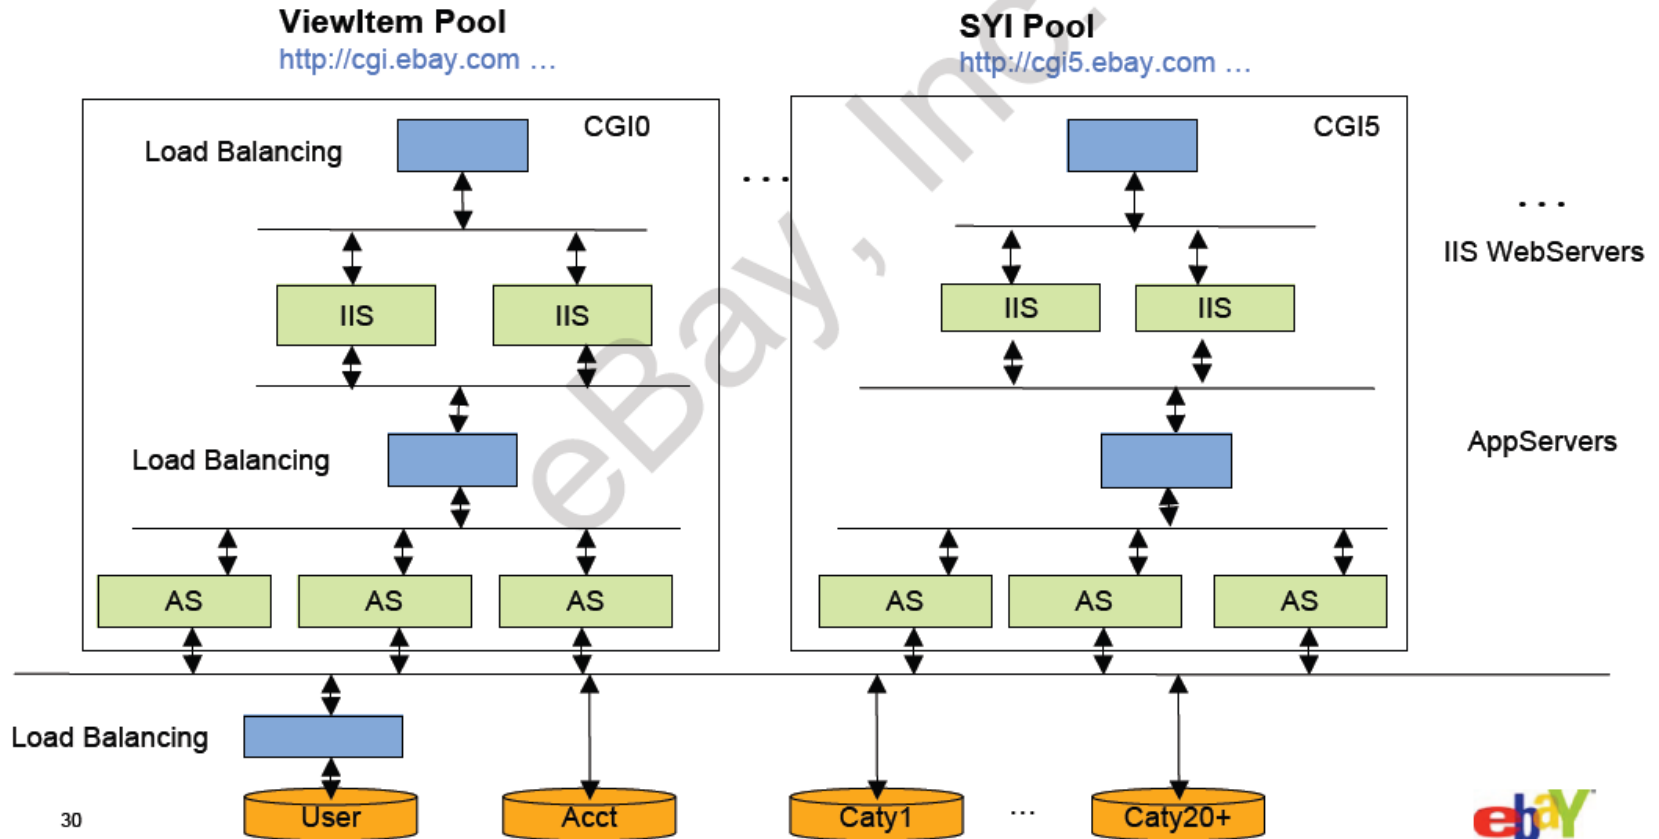
33

- Application Tier
 - ▣ Segmented by function
 - ▣ Horizontal load-balancing
 - ▣ Minimize dependencies
- Data Tier
 - ▣ Data partitioned by functional areas
 - ▣ Minimize database work
 - No stored procedure / business logic in database
 - Move CPU-intensive work to applications (no join, sort, etc.)
 - AppServers are cheap, databases are bottlenecks

eBay Architecture

Source: R. Shoup and D. Pritchett, "The eBay Architecture", SD Forum 2006

34



30


© 2006 eBay Inc.

References

35

- D. Oppenheimer and D. Patterson, “Architecture and Dependability of Large-Scale Internet Services”, IEEE Internet Computing, Sept-Oct 2002
- S. Hanselman, “A reminder on "Three/Multi Tier/Layer Architecture/Design" brought to you by my late night frustrations”,
<http://www.hanselman.com/blog/AReminderOnThreeMultiTierLayerArchitectureDesignBroughtToYouByMyLateNightFrustrations.aspx>, June 2004
- R. Shoup and D. Pritchett, “The eBay Architecture”, SD Forum 2006