

Distributed Systems

Lesson 9

University of New York in Tirana
Master of Science in Computer Science
Prof. Dr. Marenglen Biba

Lesson 9

01: Introduction

02: Architectures

03: Processes

04: Communication

05: Naming

06: Synchronization

07: Consistency & Replication

08: Fault Tolerance

09: Security

10: Distributed Object-Based Systems

11: Distributed File Systems

12: Distributed Web-Based Systems

13: Distributed Coordination-Based Systems

Client-Server Architectures (1)

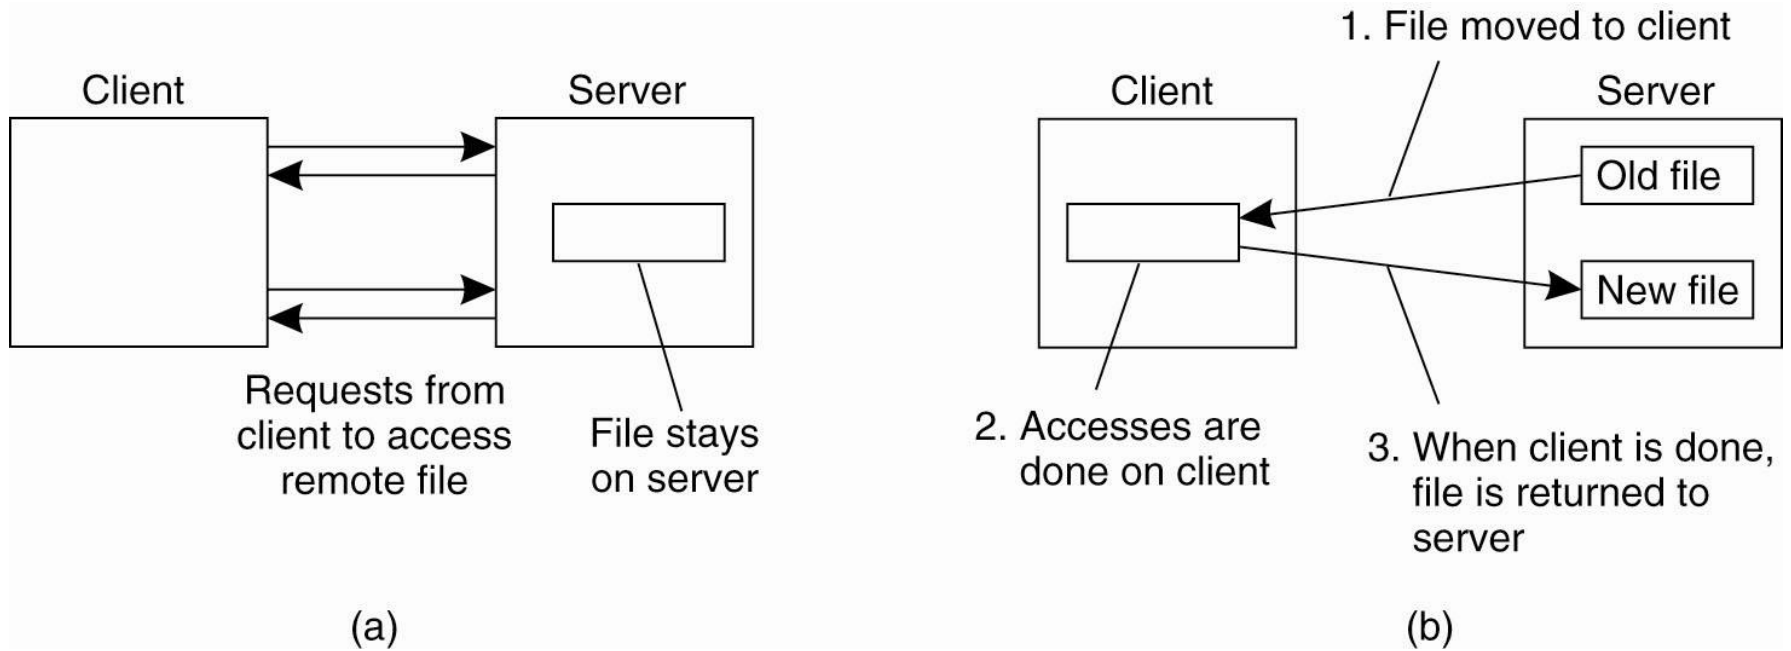


Figure 11-1. (a) The **remote access model**.
(b) The **upload/download model**.

The NFS architecture

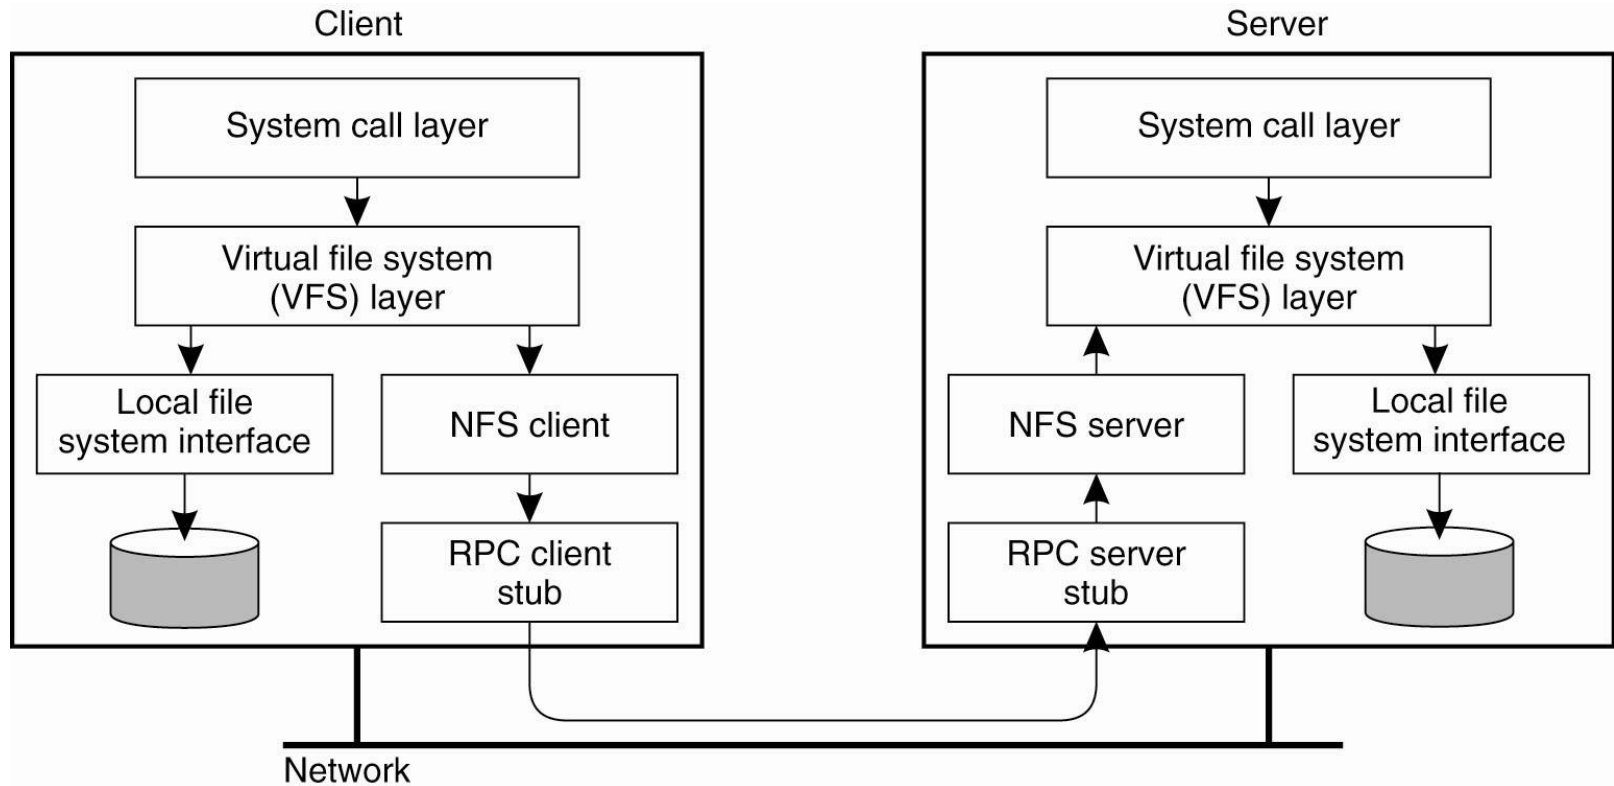


Figure 11-2. The basic NFS architecture for UNIX systems.

Cluster-Based Distributed File Systems (1)




Figure 11-4. The difference between (a) **distributing whole files** across several servers and (b) **striping files** for parallel access.

Cluster-Based Distributed File Systems

Google File System

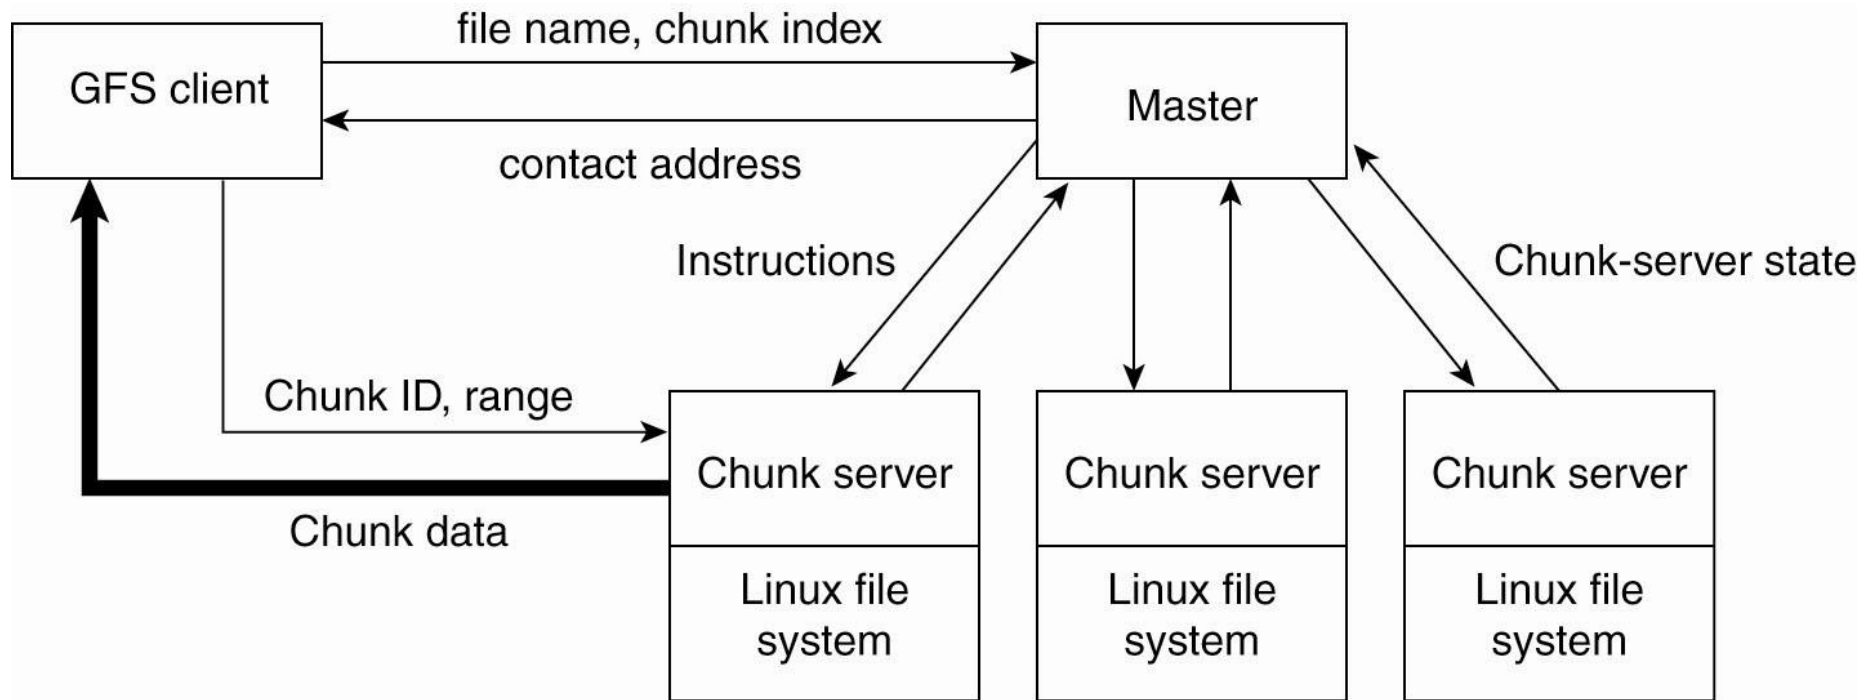


Figure 11-5. The organization of a Google cluster of servers.

Remote Procedure Calls in NFS

Open Network Computing RPC (ONC RPC) protocol





Figure 11-7. (a) Reading data from a file in NFS **version 3**. (b) Reading data using a compound procedure in **version 4**.

Mounting remote file systems in NFS

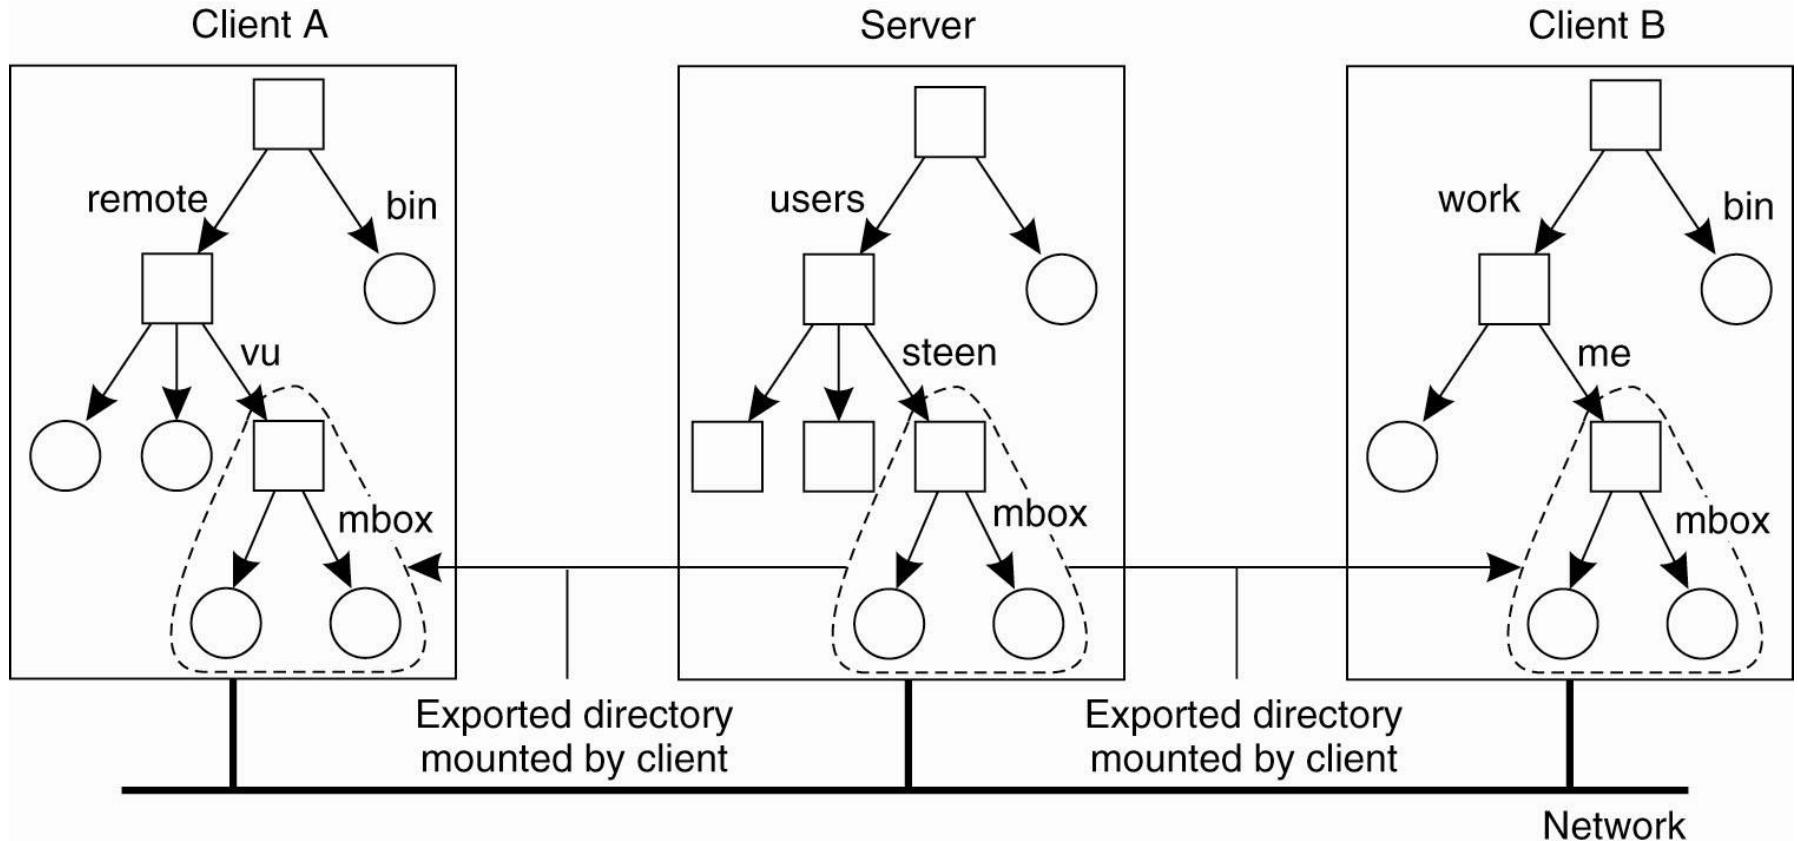


Figure 11-11. Mounting (part of) a remote file system in NFS.

Security in NFS

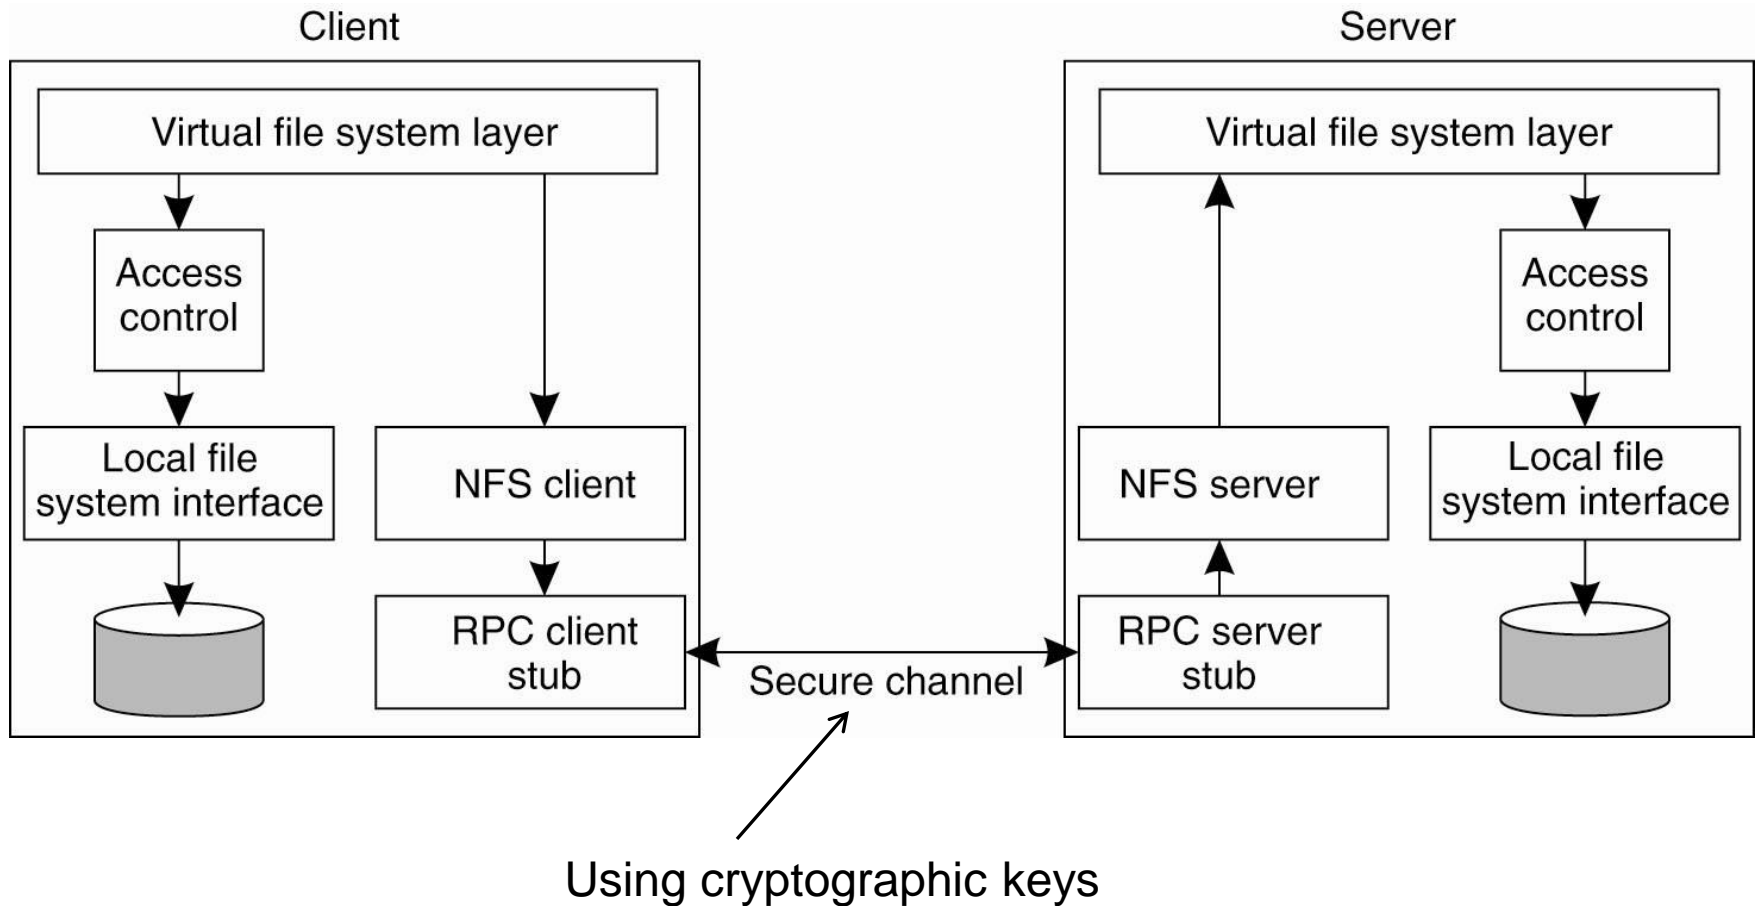


Figure 11-28. The NFS security architecture.

Secure RPCs

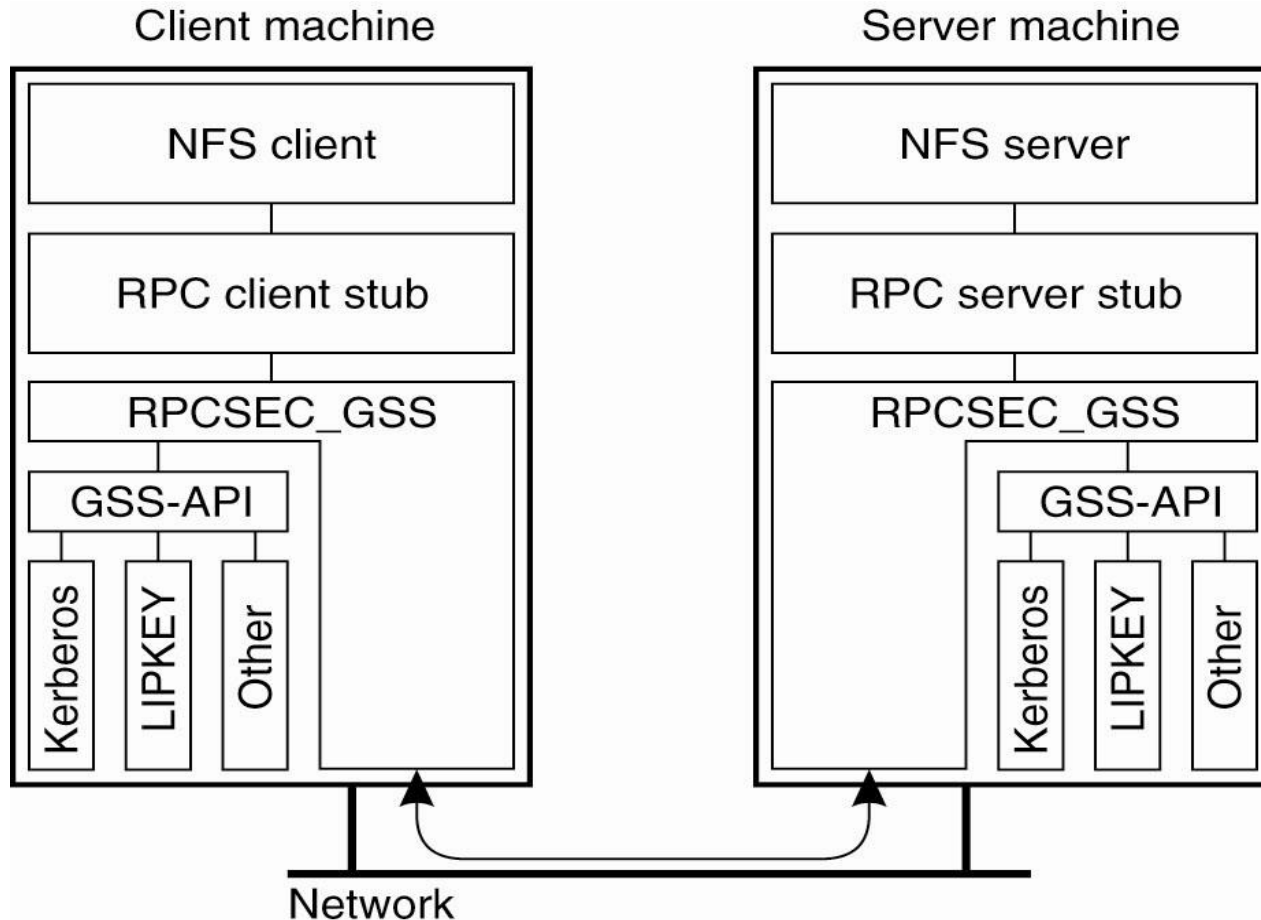


Figure 11-29. Secure RPC in NFSv4. RPCSEC_GSS is a **general security framework** that can support a myriad of security mechanisms.

GSS-API

- The Generic Security Service Application Program Interface (GSSAPI, also GSS-API) is an application programming interface for programs to access security services.
- The GSSAPI has been standardized for the C (RFC 2744) and Java (JSR-072) languages.

End of PART I

- Readings
 - Distributed Systems, Chapter 10

PART II

- 01: Introduction
- 02: Architectures
- 03: Processes
- 04: Communication
- 05: Naming
- 06: Synchronization
- 07: Consistency & Replication
- 08: Fault Tolerance
- 09: Security
- 10: Distributed Object-Based Systems
- 11: Distributed File Systems
- 12: Distributed Web-Based Systems
- 13: Distributed Coordination-Based Systems

Lesson outline

- PART I - Distributed Web-Based Systems
- PART II – Lab Session

Traditional Web-Based Systems




Figure 12-1. The overall organization of a traditional Web site. See the Web as a huge distributed system where services rather than just documents are being offered.

Web Documents

Type	Subtype	Description
Text	Plain	Unformatted text
	HTML	Text including HTML markup commands
	XML	Text including XML markup commands
Image	GIF	Still image in GIF format
	JPEG	Still image in JPEG format
Audio	Basic	Audio, 8-bit PCM sampled at 8000 Hz
	Tone	A specific audible tone
Video	MPEG	Movie in MPEG format
	Pointer	Representation of a pointer device for presentations
Application	Octet-stream	An uninterpreted byte sequence
	Postscript	A printable document in Postscript
	PDF	A printable document in PDF
Multipart	Mixed	Independent parts in the specified order
	Parallel	Parts must be viewed simultaneously

Figure 12-2. Six top-level Multipurpose Internet Mail Exchange (MIME) types and some common subtypes.

Multitiered Architectures

CGI: Common Gateway Interface

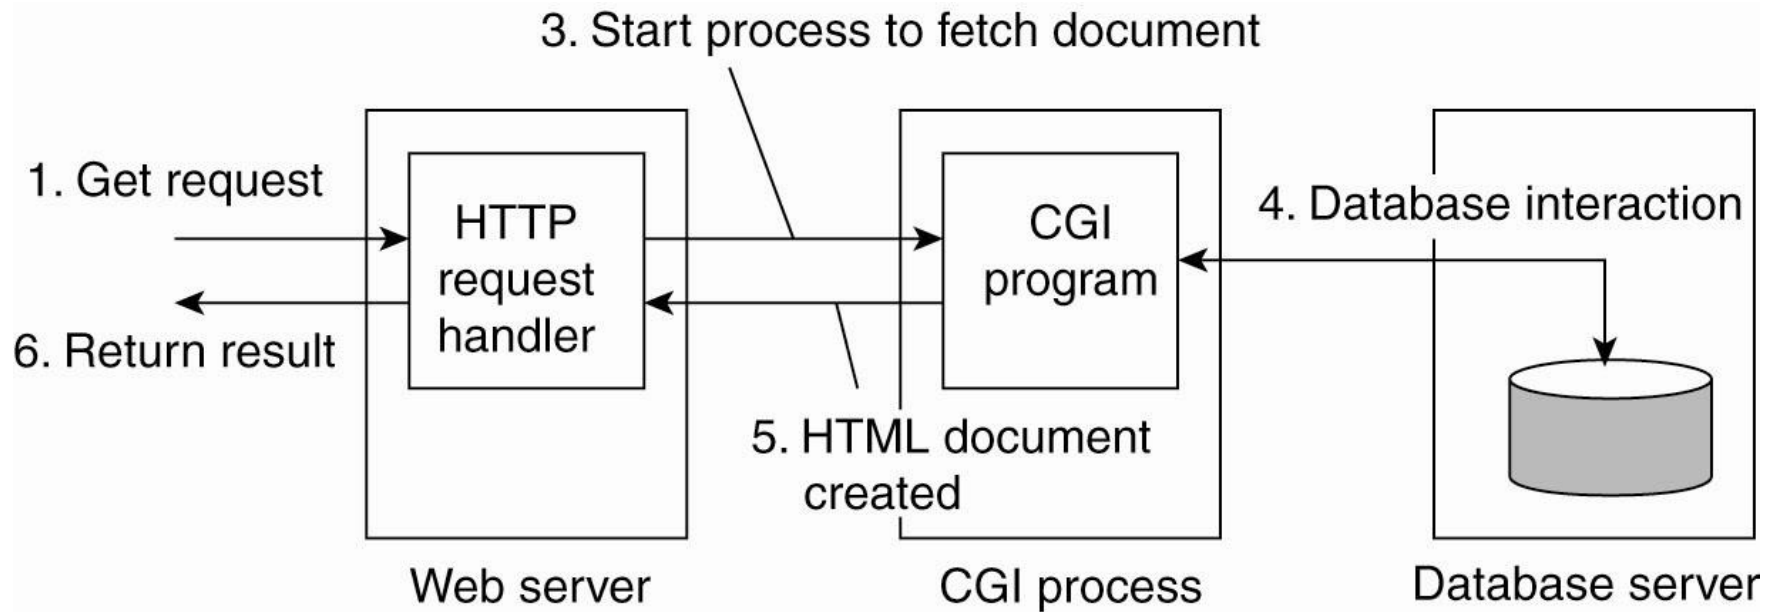


Figure 12-3. The principle of using server-side CGI programs.

Server-side scripts

- Servers nowadays do much more than just fetching documents.
- One of the most important enhancements is that servers can also **process a document** before passing it to the client.
- In particular, a document may contain a **server-side script**, which is executed by the server when the document has been fetched locally.
- The **result of executing a script is sent** along with the rest of the document to the client.
- **The script itself is not sent.**
 - In other words, using a server-side script changes a document by essentially **replacing** the script with the results of its execution.

Three-tiers and servlets

- As server-side processing of Web documents increasingly requires more flexibility, it should come as no surprise that many Web sites are now organized as a three-tiered architecture consisting of a **Web server, an application server, and a database**.
- **Application server** runs all kinds of programs that may or may not access the third tier consisting of a database.
 - For example, a server may accept a customer's query, search its database of matching products, and then **construct a clickable Web page** listing the products found.
- In many cases the server is responsible for running Java programs, called **servlets**, that maintain things like shopping carts, implement recommendations, keep lists of favorite items, and so on.

Servlets

- A Servlet is a **Java class** in Java EE that conforms to the Java Servlet API, a protocol by which a Java class may **respond to HTTP requests**.
- They are not tied to a specific client-server protocol, but are most often used with this protocol.
- The word "Servlet" is often used in the meaning of "HTTP Servlet".
- We may use a servlet to **add dynamic content** to a Web server using the Java platform.
- The generated content is commonly HTML, but may be other data such as XML.
- Servlets are the Java counterpart to non-Java dynamic Web content technologies such as CGI and ASP.NET.

Facelets

- Facelets is an open source Web template system and the default view handler technology for JavaServer Faces (JSF).
- JavaServer Faces (JSF) is a Java-based Web application framework intended to simplify development integration of web-based user interfaces.

Web Proxies



Figure 12-6. Using a Web proxy when the browser does not speak FTP.

SQUID

<http://www.squid-cache.org>

Web Server Clusters (1)

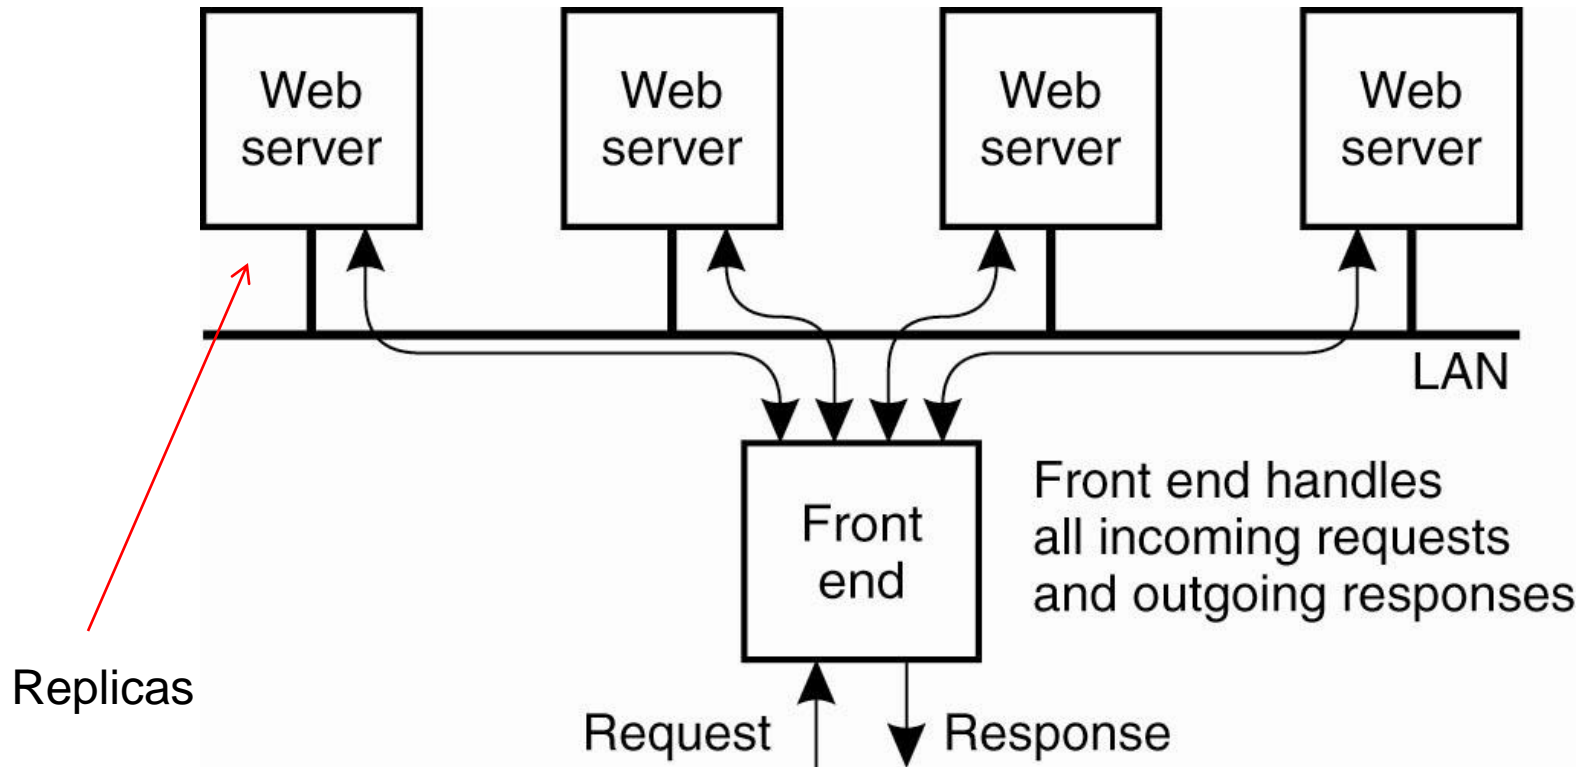


Figure 12-8. The principle of using a server cluster in combination with a front end to implement a Web service.

Web Server Clusters (2)

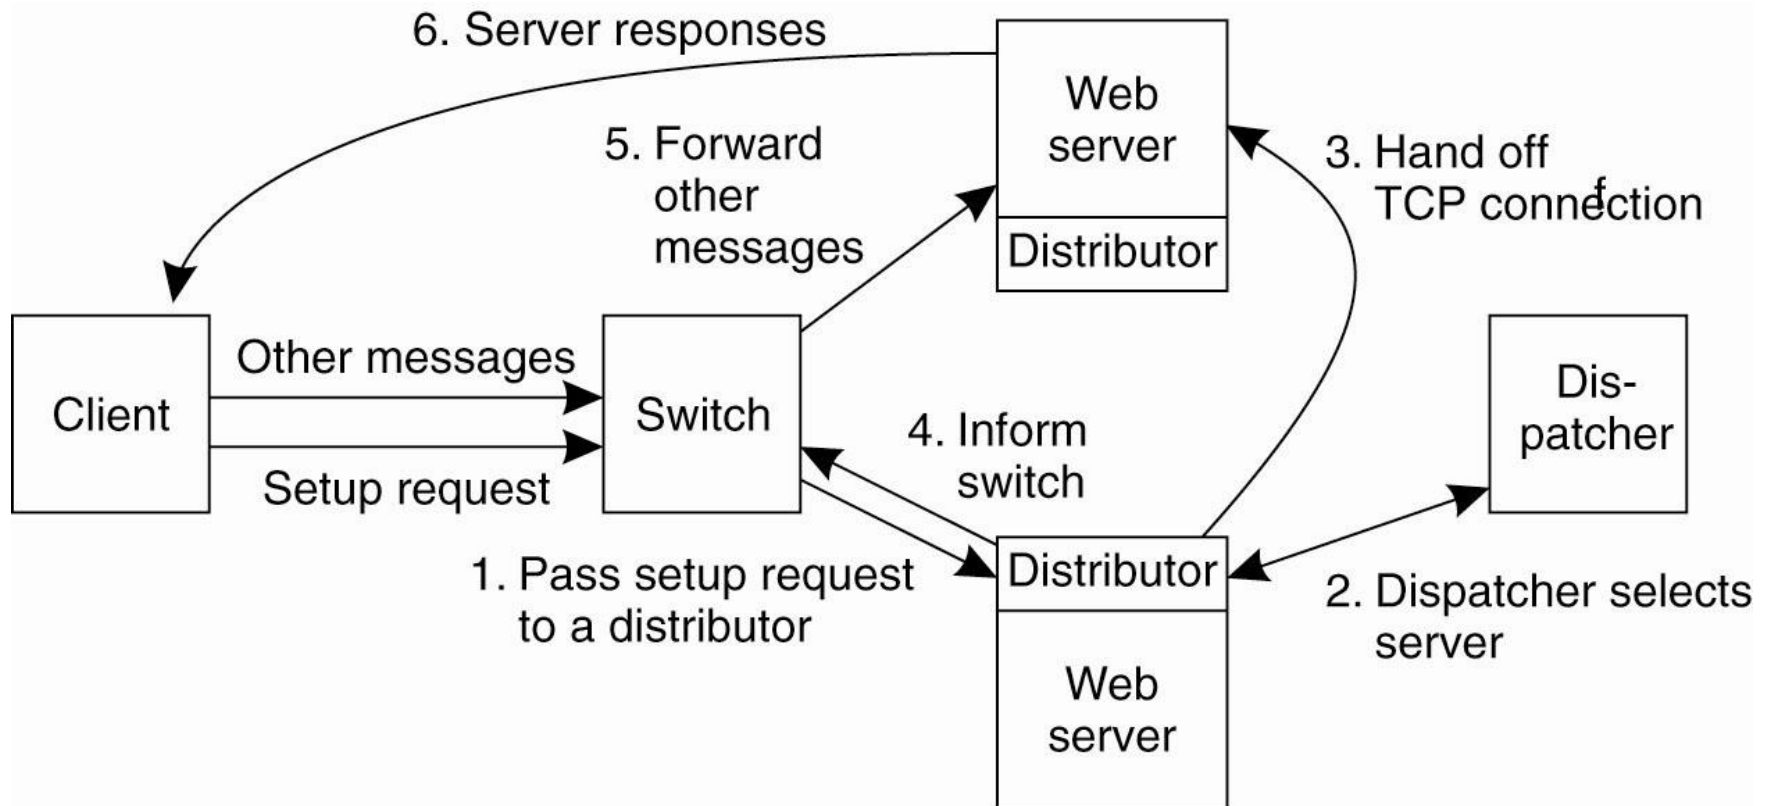


Figure 12-9. A scalable content-aware cluster of Web servers.

Web Proxy Caching

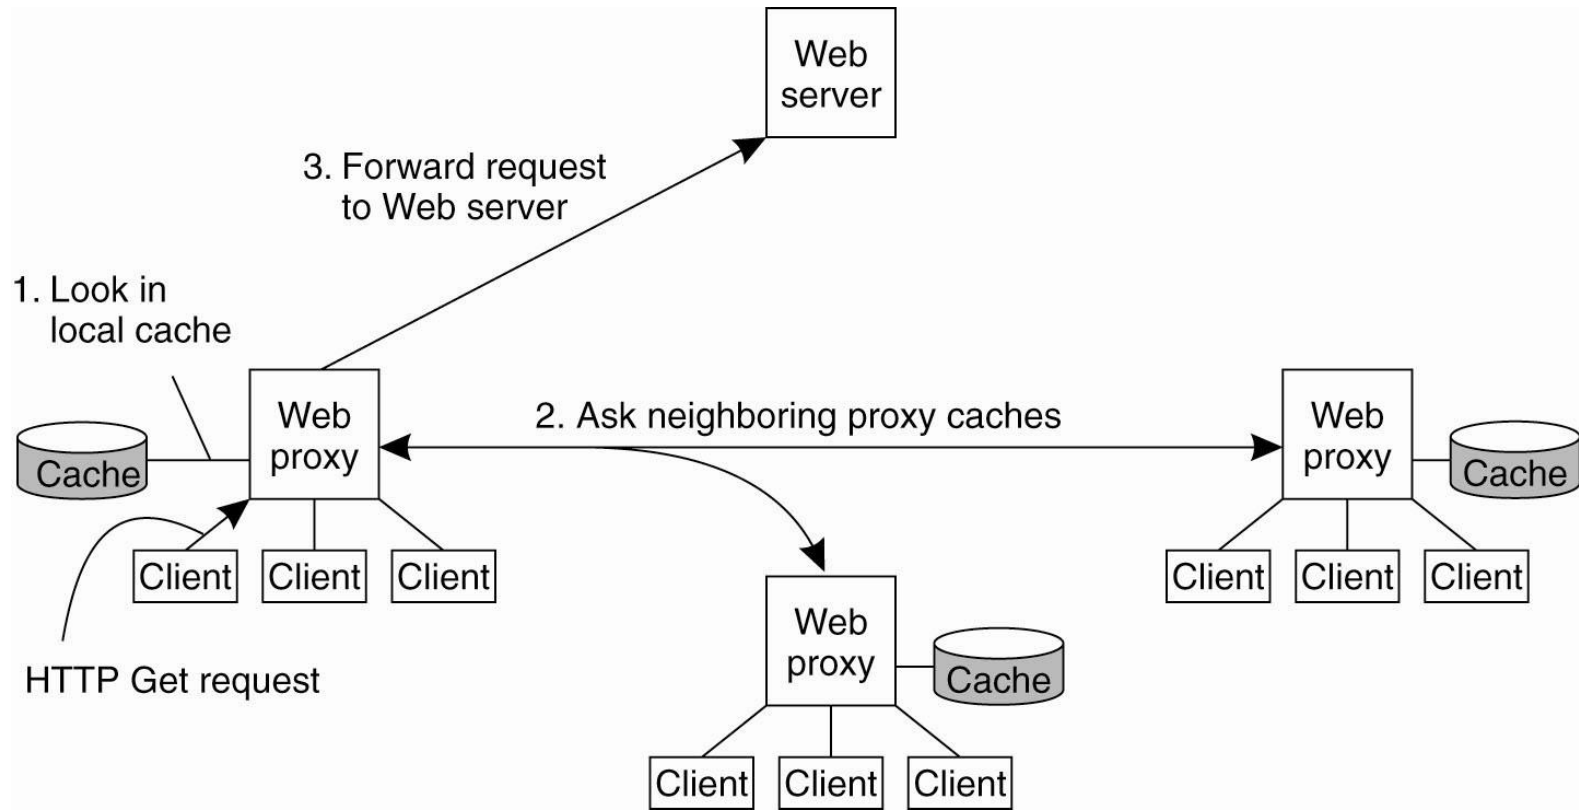


Figure 12-17. The principle of cooperative caching.


Real-world dimensions: what are we talking about?

- 15 years since its creation, distributed computing platform Akamai now delivers 15-20 per cent of the world's internet traffic by request.
- Whether you've heard of Akamai or not, you've already used their products. Akamai owns and operates the world's largest distributed computing platform. On **250,000 servers in more than 80 countries**, it transparently mirrors content from its customer's servers.
- State of the internet report 2015
 - <https://www.akamai.com/us/en/about/news/press/2015-press/akamai-releases-second-quarter-2015-state-of-the-internet-report.jsp>
 - *More than 700 million unique IP addresses connected to Akamai Intelligent Platform*
- AKAMAI: founded by Prof. Tom Leighton and his student Daniel Lewin in 1998 at M.I.T.

Akamai Customers

- **Apple** uses Akamai as their primary content delivery network for a wide range of applications including software downloads from Apple's Website, QuickTime movie trailers, and the iTunes Store.
- In September 1999, **Microsoft** and Akamai formed a strategic relationship to incorporate Windows Media technology in Akamai's FreeFlow service, as well as to facilitate the porting of the FreeFlow product to the Windows platform; this relationship exists to this day.
- The **BBC iPlayer** uses Akamai to stream its recorded and live programs, focused through an XML playlist.
- The official U.S. government **White House** website (WhiteHouse.gov) uses Akamai Technologies for hosting video clips of President Barack Obama's Web addresses on their own in-house servers.
- The entire China Central Television website (CCTV.com), including its streaming video, has been hosted on Akamai's edge servers since late 2009
- **MIT OpenCourseWare** utilizes Akamai's EdgeSuite for its content delivery network.
- Other customers include **Facebook, Twitter, AMD, Wedubox, Hilton Worldwide Hotels, Amazon.com** (for their dynamic contents), **Adobe Systems, Netflix.**

AKAMAI: Delivery

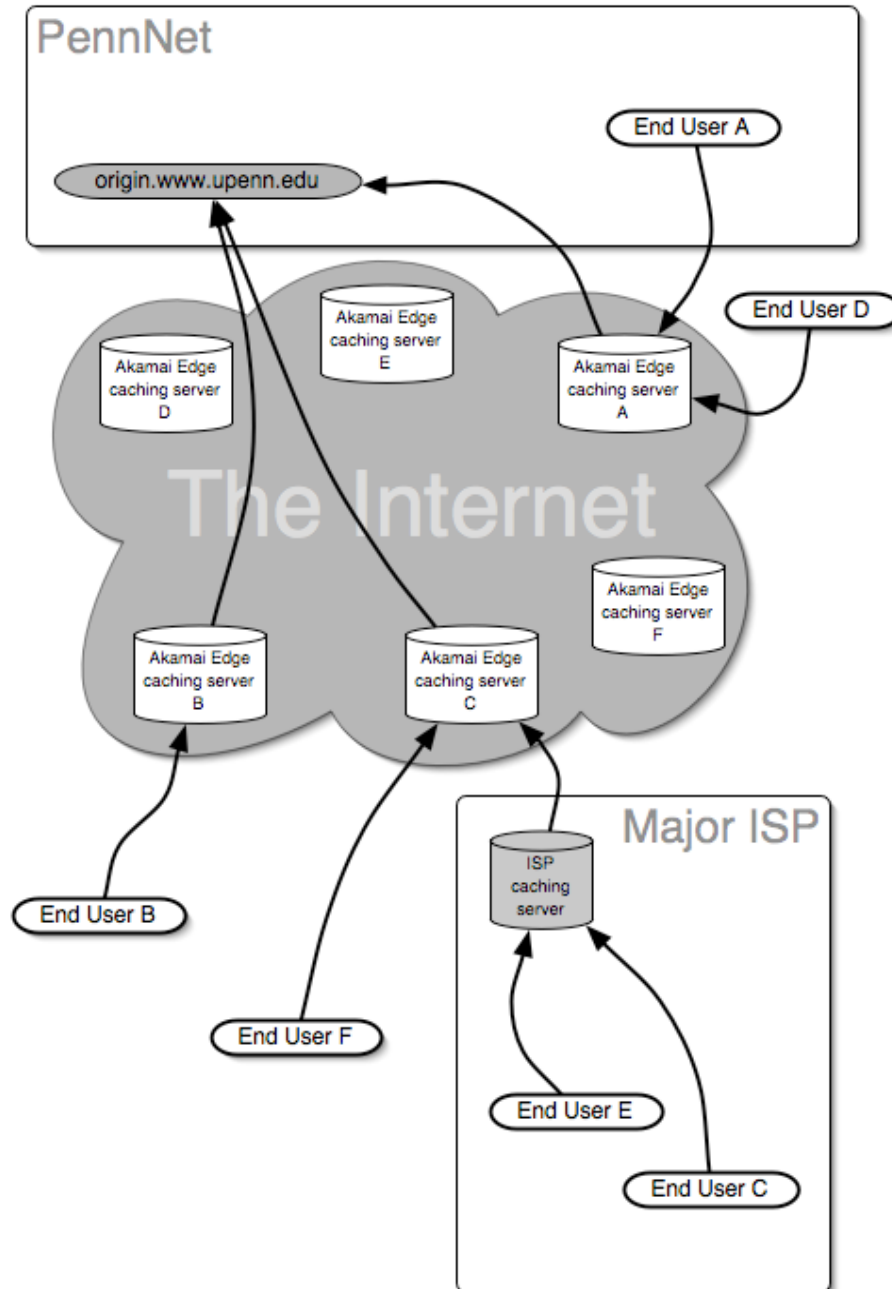


Replicated web application



Today Akamai handles tens of billions of daily Web interactions for companies like Audi, NBC, and Fujitsu, and organizations like the U.S. Department of Defense and NASDAQ -- powering brand new business models that serve the changing online economy. (From AKAMAI Website)

AKAMAI Caching



Replication for Web Hosting Systems

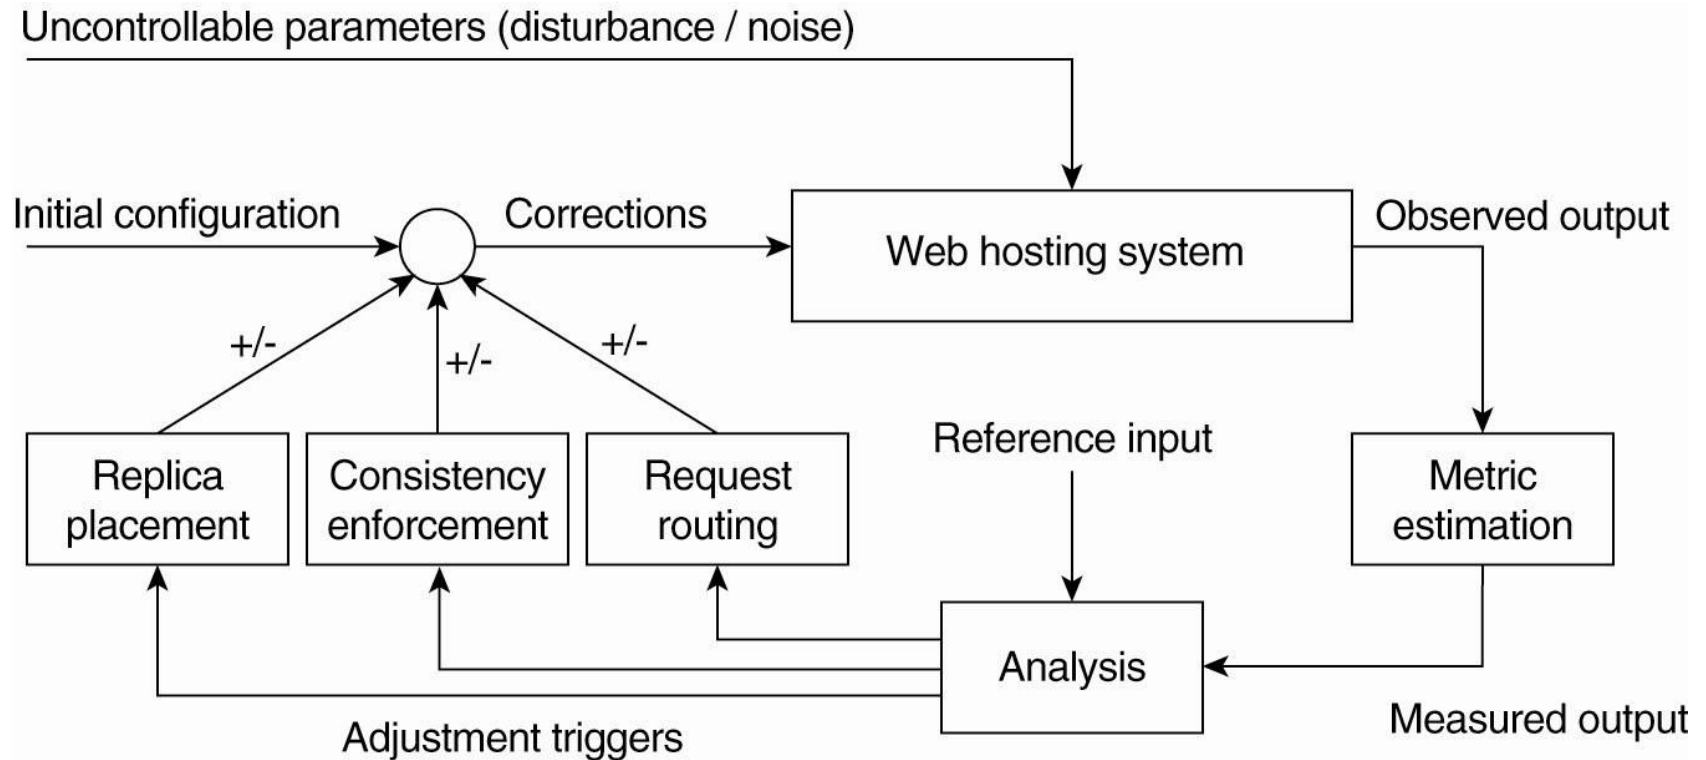


Figure 12-18. The general organization of a CDN as a feedback-control system (adapted from Sivasubramanian et al., 2004b).

Metric Estimation

- Latency metrics
 - Time for fetching a document
- Spatial metrics
 - Distance in terms of network-level routing hops
- Network usage metrics
 - Consumed bandwidth
- Consistency metrics
 - To what extent the replica is deviating from its master copy
- Financial metrics
 - How well a CDN is doing in terms of ROI (return of investment)

Embedded documents

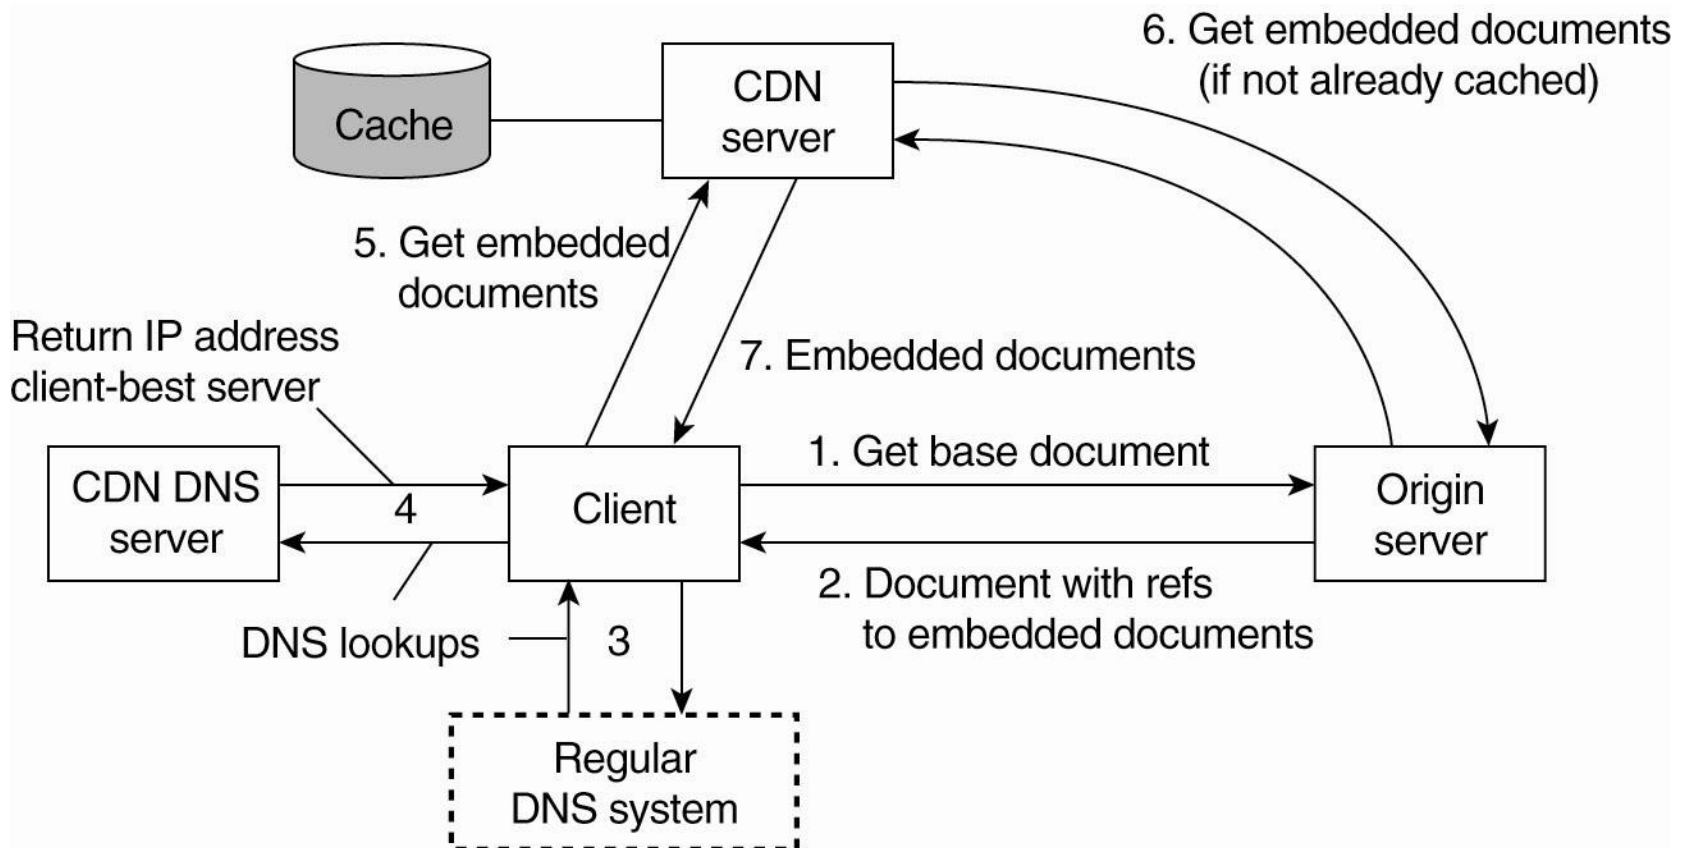


Figure 12-20. The principal working of the Akamai CDN.

Replication of Web Applications

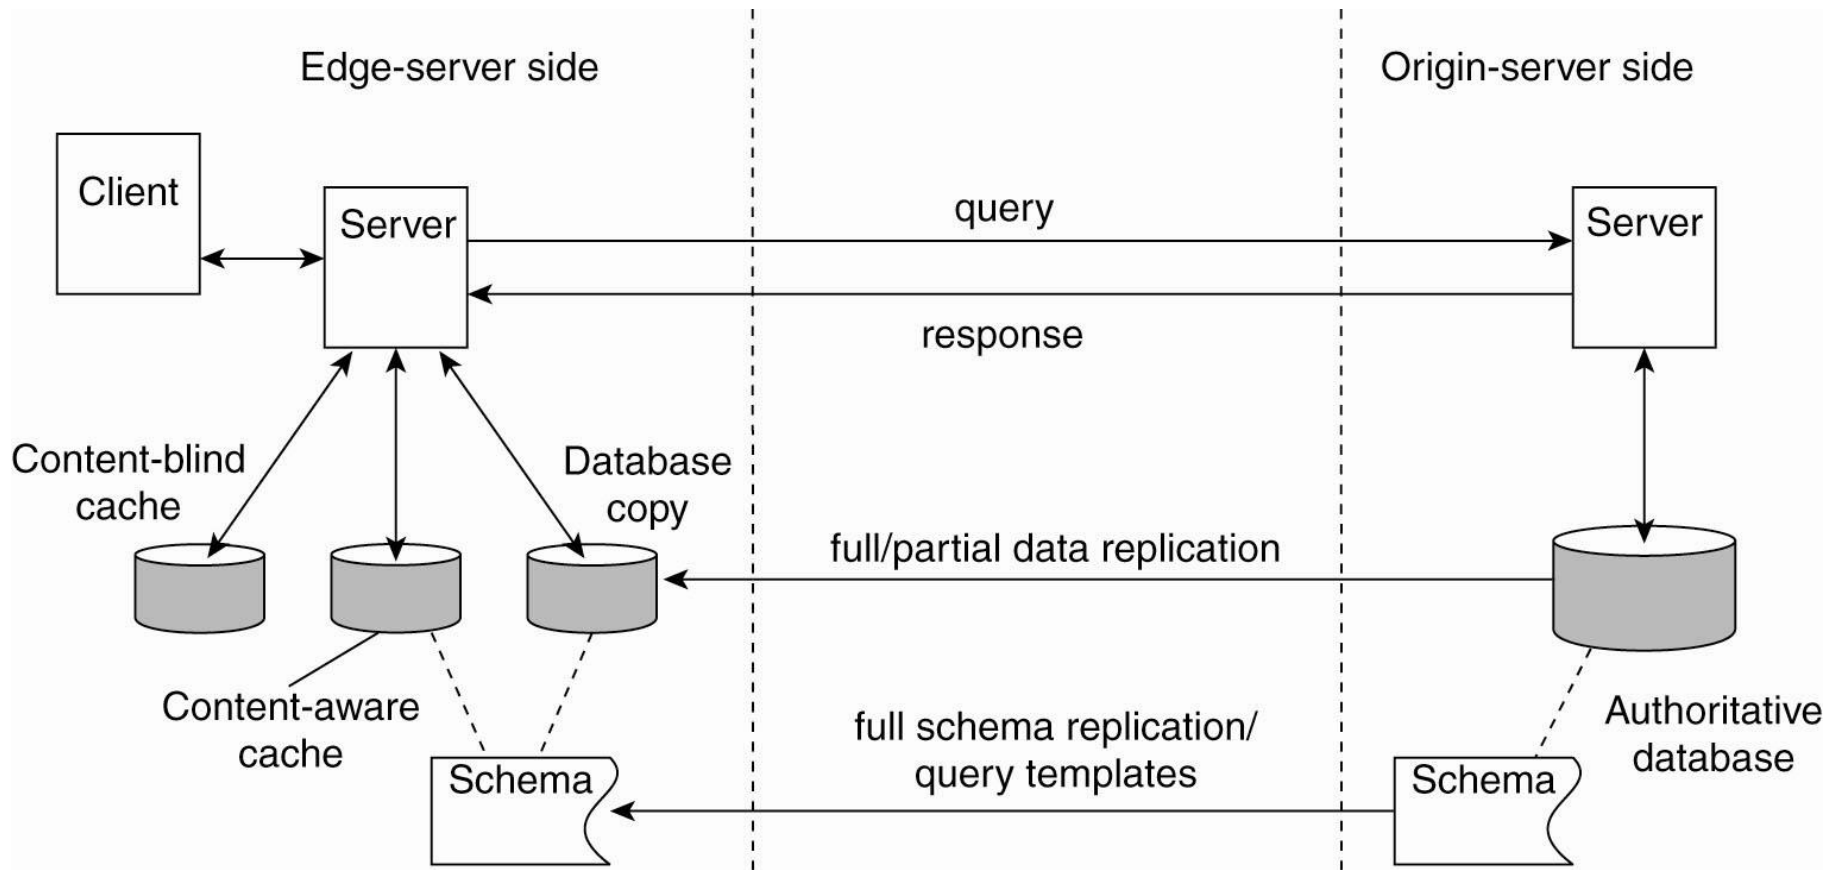


Figure 12-21. Alternatives for caching and replication with Web applications: **data or schema replication**


End of theory part

- Readings
 - Distributed Systems, Chapter 12.

Lab Session

- Web Server: Java EE 6 Web
- Web Interface Tier:
 - Facelets
 - Servlets
- Processing Tier
 - Java Beans
- Applications
 - Hello (facelet + Java Bean)
 - Hello2 (servlet)
 - GuessNumber (facelets + Java Bean)
- Banking App
 - Follow the manual given in class.

Web Banking with Servlets + EJB + MySQL



End of Lesson 9

- Readings
 - Distributed Systems, Chapter 12.
- Lab Manual on EJBs + Servlets + MySQL given in class