

CSE 486/586 Distributed Systems

Content Distribution Networks

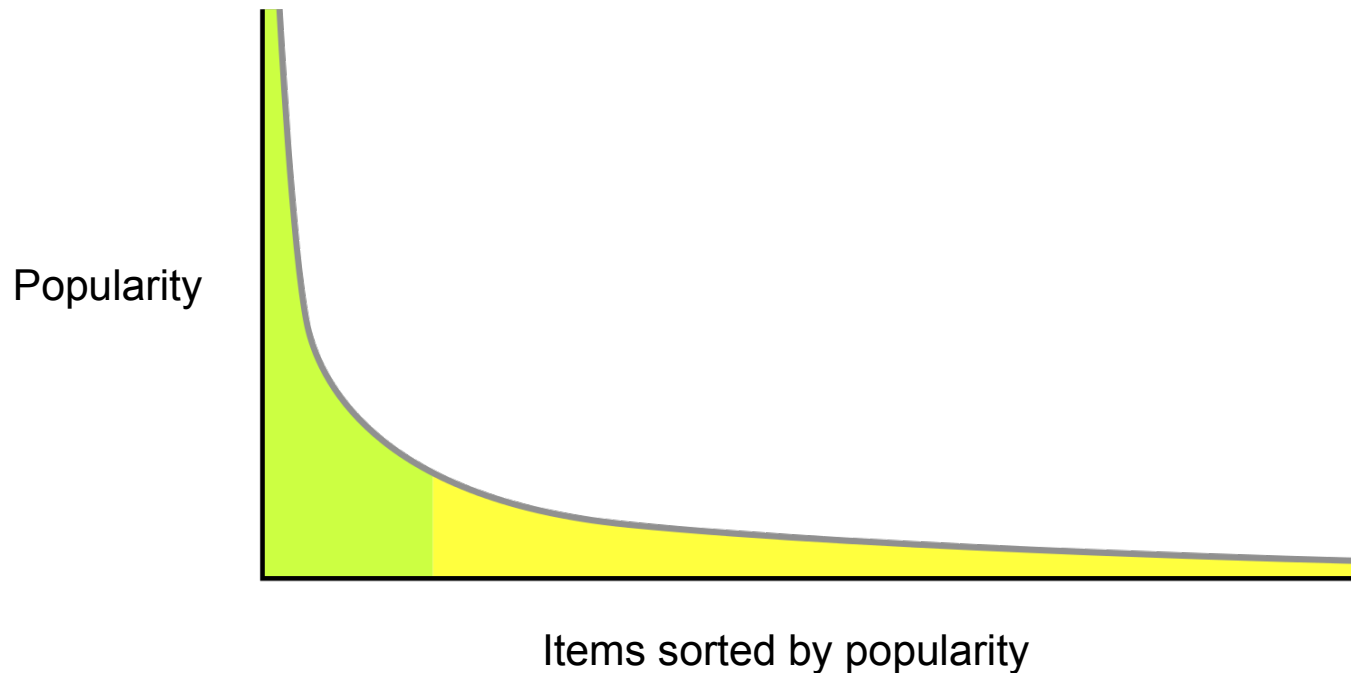
Slides by Steve Ko
Computer Sciences and Engineering
University at Buffalo

Understanding Your Workload

- Engineering principle
 - Make the common case fast, and rare cases correct
 - (From Patterson & Hennessy books)
 - This principle cuts through generations of systems.
- Example?
 - CPU Cache
- Knowing common cases == understanding your workload
 - *E.g.*, read dominated? Write dominated? Mixed?

Content Distribution Problem

- **Power law** (Zipf distribution)
 - Models a lot of natural phenomena
 - Social graphs, media popularity, wealth distribution, etc.
 - Happens in the Web too.




Content Distribution Workload

- What are the most frequent things you do on Facebook?
 - Read/write wall posts/comments/likes
 - View/upload photos
 - Very different in their characteristics
- Read/write wall posts/comments/likes
 - Mix of reads and writes so more care is necessary in terms of consistency
 - But small in size so probably less performance sensitive
- Photos
 - **Write-once, read-many** so less care is necessary in terms of consistency
 - But large in size so more **performance sensitive**

Facebook's Photo Distribution Problem

- “Hot” vs. “very warm” vs. “warm” photos
 - Hot: Popular, a lot of views
 - Very warm: Somewhat popular, still a lot of views
 - Warm: Unpopular, but still a lot of views in aggregate



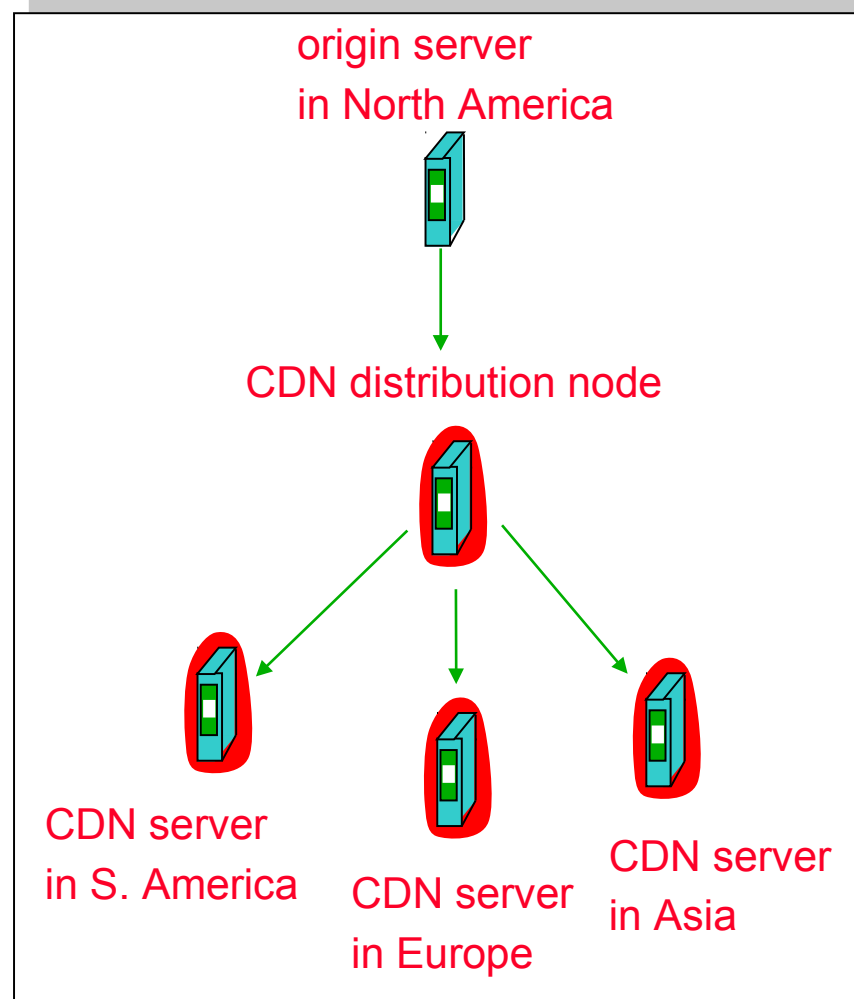
“Hot” Photos

- How would you serve these photos?
- Caching should work well.
 - Many views for popular photos
- Where should you cache?
 - Close to users
- What’s commonly used these days?
 - CDN
 - CDN mostly relies on DNS
- “Very warm” and “warm” will be covered later

Content Distribution Networks (CDNs)

Content replication

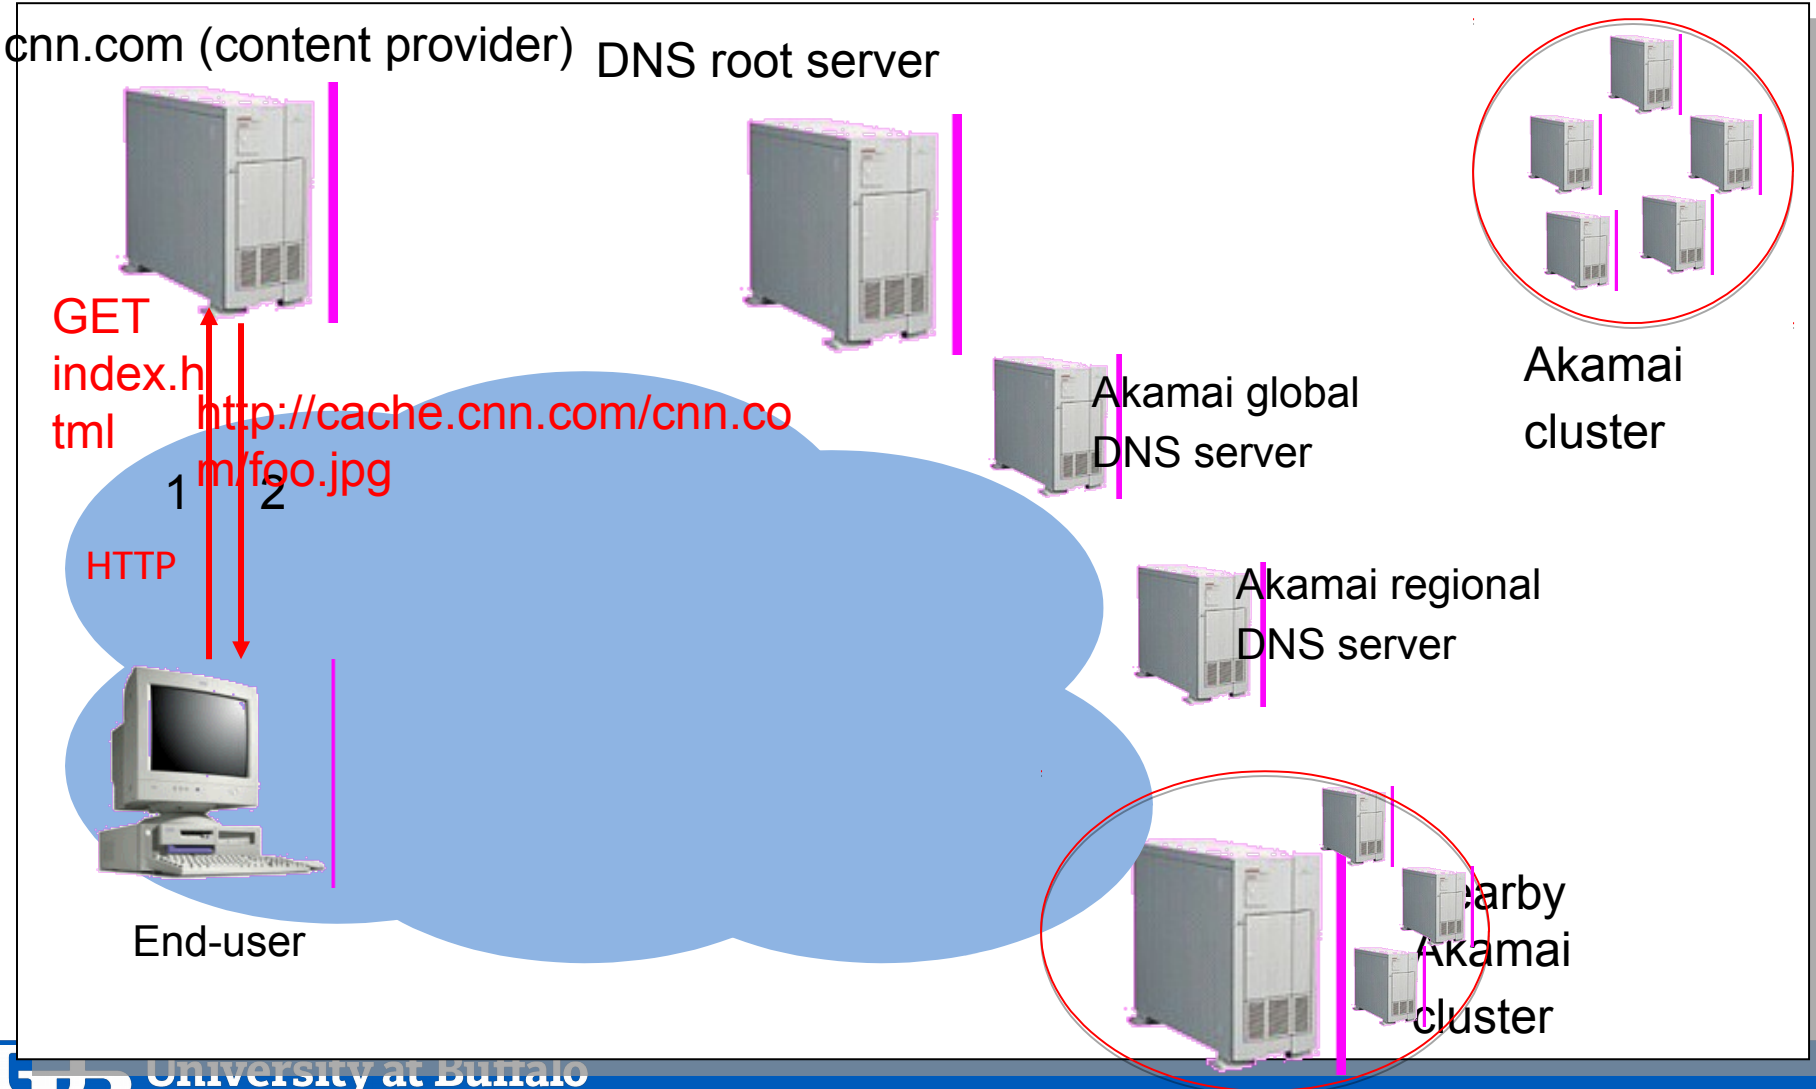
- CDN company installs thousands of servers **throughout the Internet**
 - In large datacenters
 - Or **close to users**
- CDN replicates customers' content
- When provider updates content, CDN updates servers



Content Distribution Networks

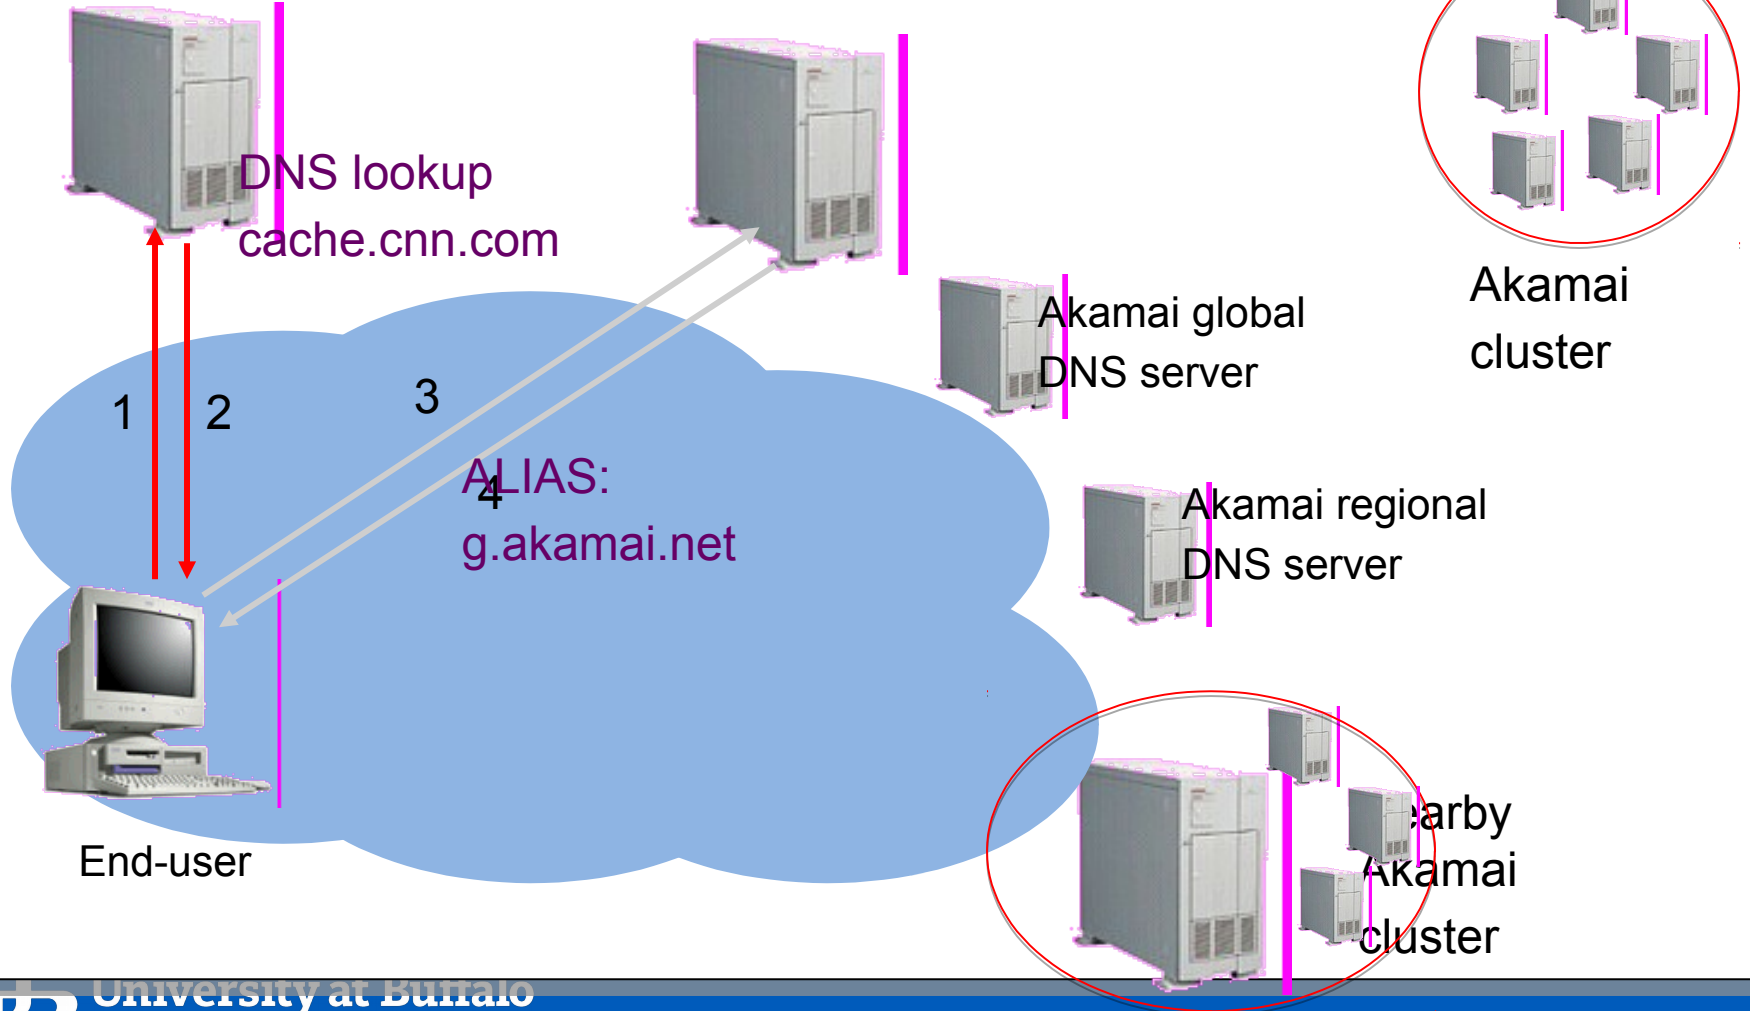
- Replicate content on many servers
- Challenges
 - How to replicate content
 - Where to replicate content
 - How to find replicated content
 - How to choose among replicas
 - How to direct clients towards a replica

How Akamai Works



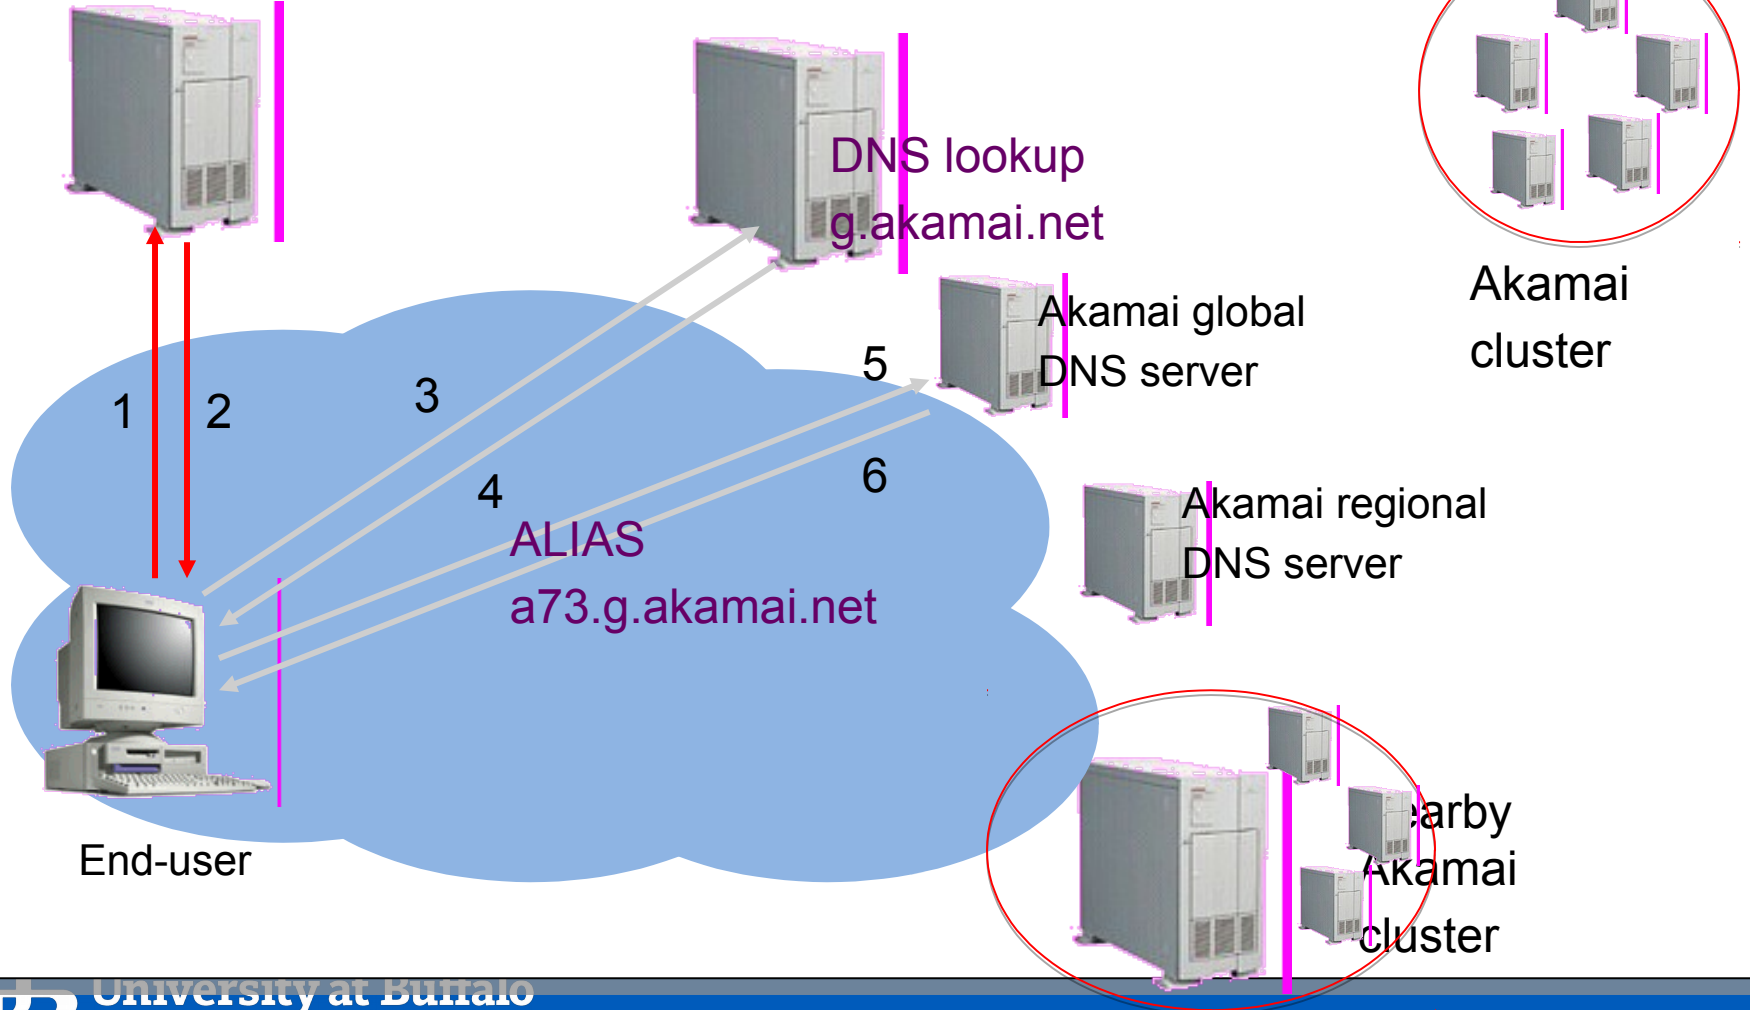
How Akamai Works

cnn.com (content provider) DNS root server



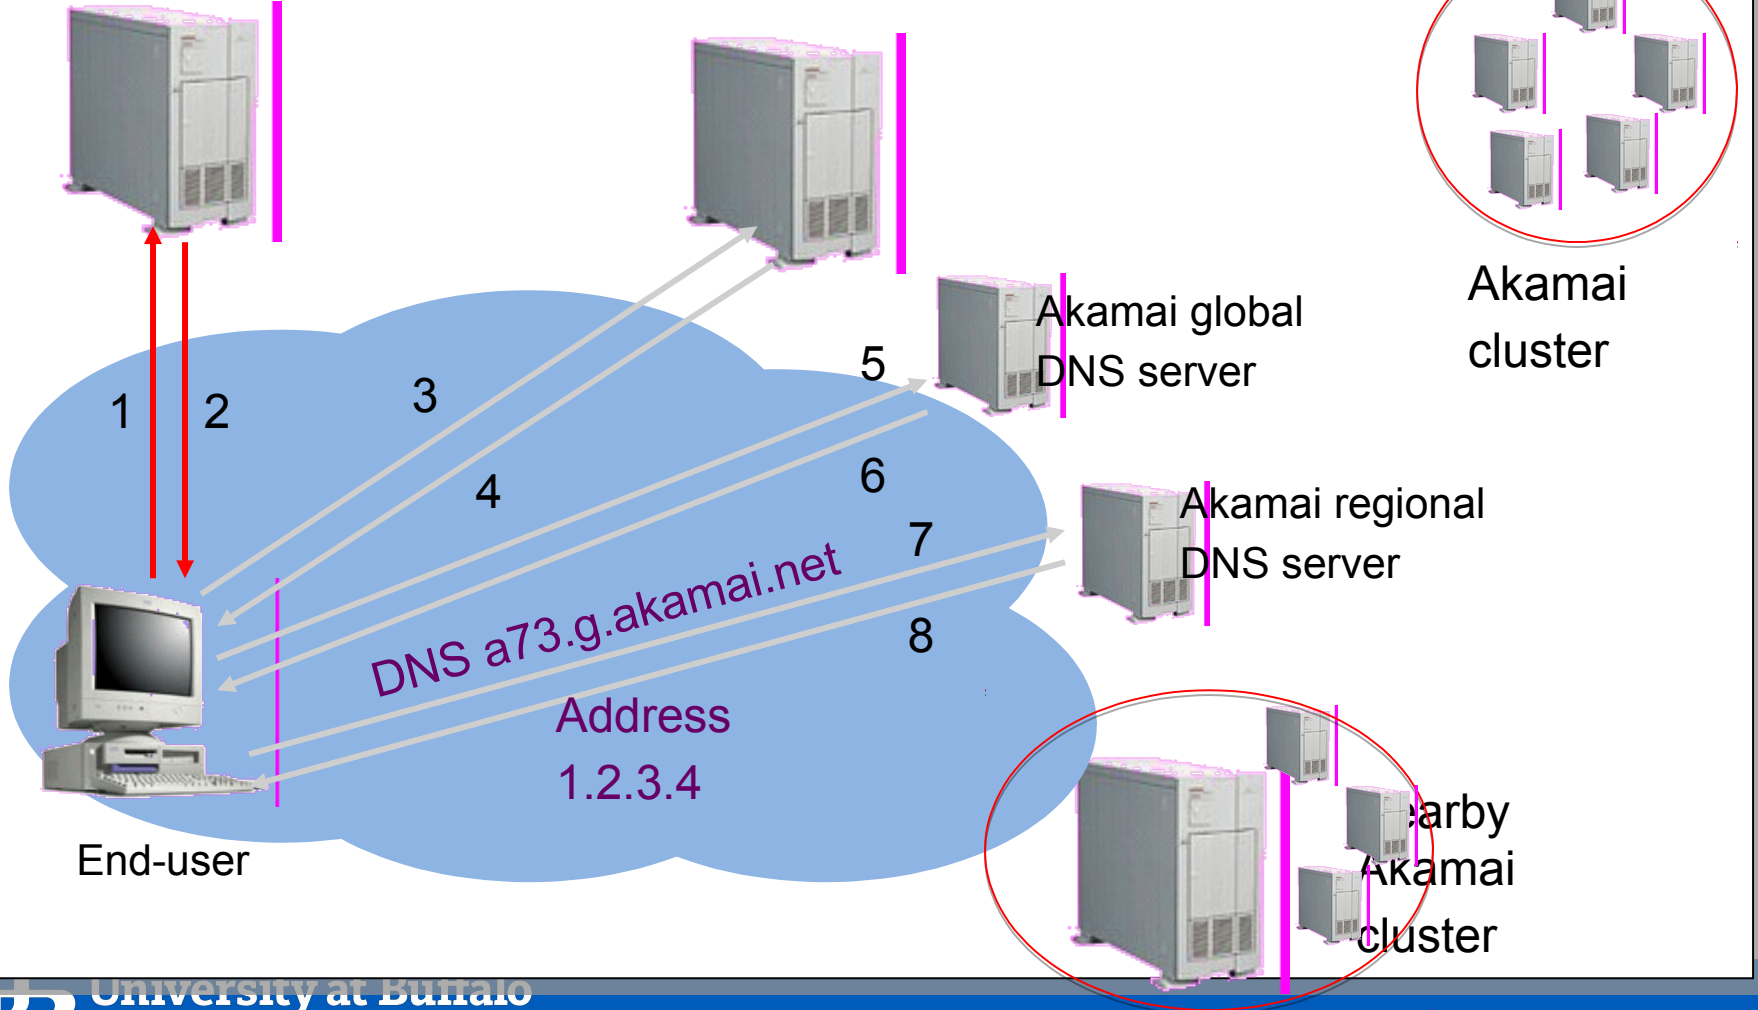
How Akamai Works

cnn.com (content provider) DNS root server




How Akamai Works

cnn.com (content provider) DNS root server




How Akamai Works

cnn.com (content provider) DNS root server

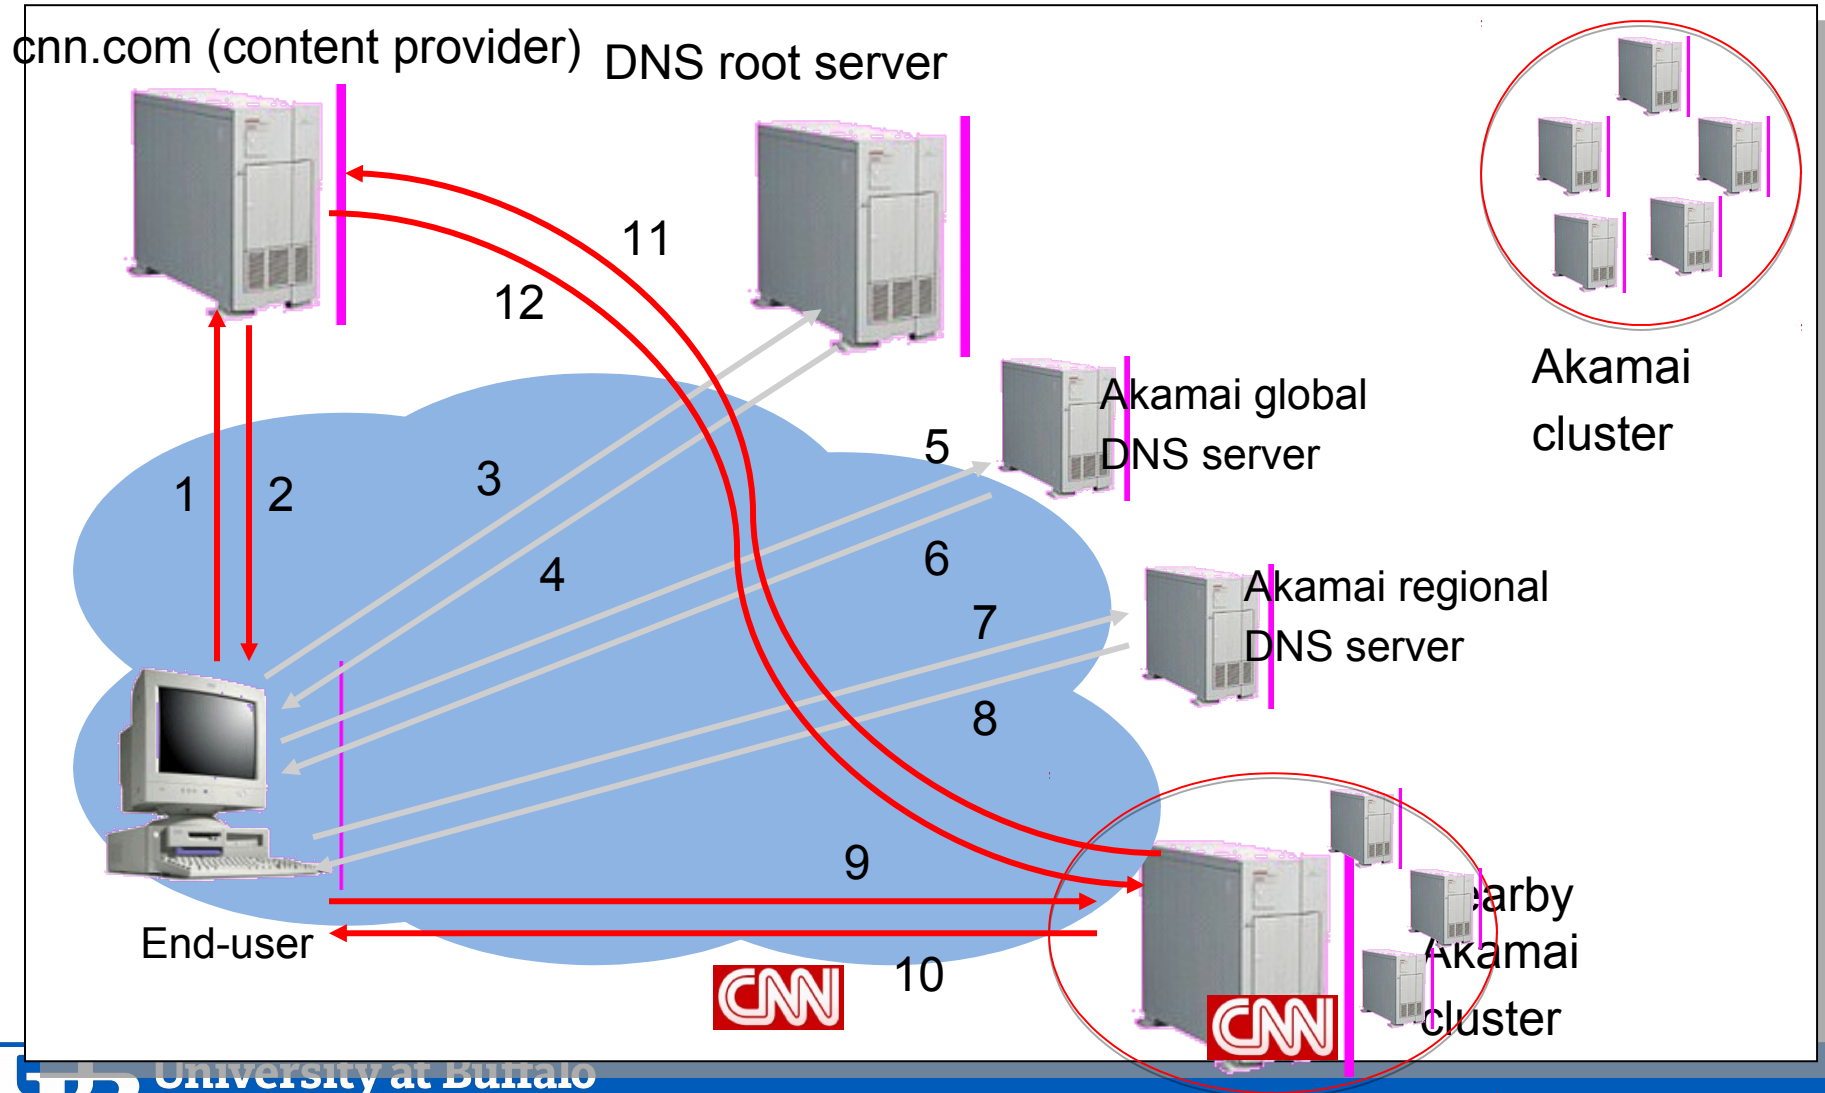


How Akamai Works

cnn.com (content provider) DNS root server



How Akamai Works



Summary

- Content Distribution Networks
 - Deliver **popular content** for content providers
 - Improve performance by **moving content closer to users**
 - Improve reliability by **duplicating content**
 - **Use DNS to redirect traffic** to local mirrors

Acknowledgements

- These slides originally by Steve Ko, lightly modified and used with permission by Ethan Blanton
- These slides contain material developed and copyrighted by Indranil Gupta (UIUC), Michael Freedman (Princeton), and Jennifer Rexford (Princeton).