

Indice

Hora 1

1. Problemática de sistemas de información un servidor-un sitio
2. Mirrors
3. Proxy caches
4. Redes de distribución de contenidos (CDN)
 - 4.1 CDN localización
 - 4.2 CDN redirección

Hora 2


- 4.3 CDN replicación
- 4.4 Estructura CDN
- 4.5 ESI
- 4.6 Proveedores CDN
 - 4.6.1 Akamai
- 4.7 Efecto de la versión HTTP en CDNs

Referencias

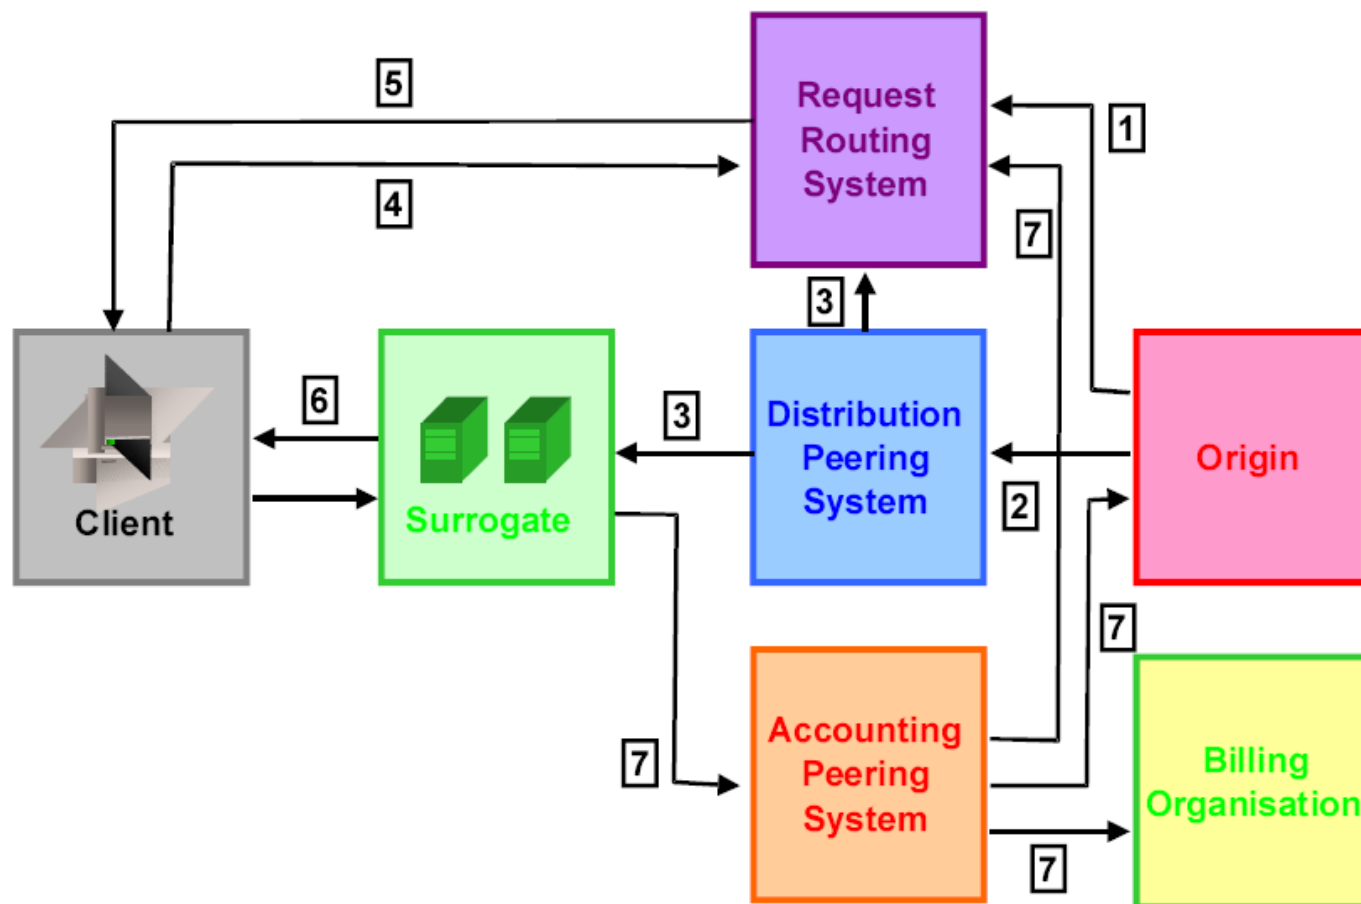
4.3 CDN replicación

- ▶ Un CDN replica los contenidos del servidor original de distintas formas:
 - Push-based
 - El servidor original descarga los contenidos a replicar en un sistema de distribución encargado de hacer las réplicas en los servidores surrogate del CDN.
 - Pull-based
 - El cliente es redireccionado al servidor surrogate del CDN que le corresponda y en caso de no tener los contenidos solicitados por éste los solicita al servidor original o a otros servidores surrogate del CDN, cacheándolos para futuras referencias. Trabaja a modo de la caché de un reverse proxy.

4.4 Estructura CDN



Estructura CDN



Fuente: "CDN Peering Architectural Overview" IETF Internet draft

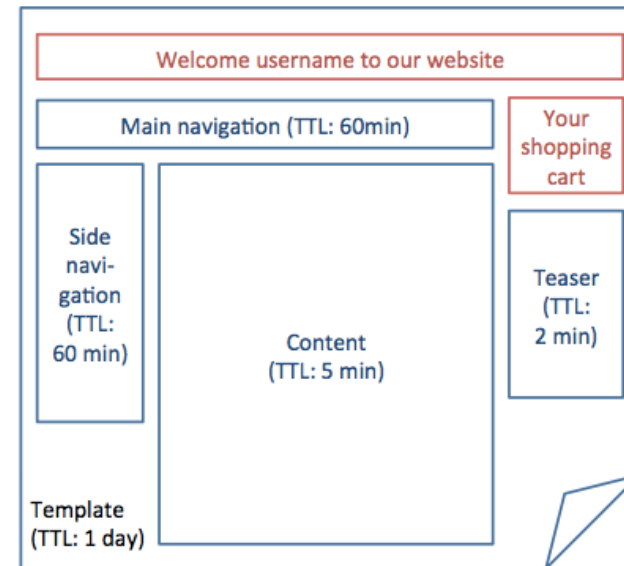
Estructura CDN

- ▶ 1. El servidor original delega su espacio de nombres URI para ser distribuido por el CDN a través del sistema de enrutado de peticiones.
- ▶ 2. El servidor original publica el contenido en un sistema de distribución.
- ▶ 3. El sistema de distribución replica el contenido en todos los servidores surrogate del CDN.
- ▶ 4. El cliente solicita un contenido con su URI como si se refiriera al servidor original.
- ▶ 5. El sistema de enrutado de peticiones redirecciona la petición al servidor CDN (surrogate) que convenga.
- ▶ 6. El servidor CDN atiende la petición del cliente.
- ▶ 7. La transacción es controlada por un sistema de contabilidad del CDN encargado del tema estadístico de utilización y facturación.

4.5 ESI

- ▶ ESI, Edge Side Includes
 - Lenguaje de marcado que permite ensamblar contenido dinámico
 - Permite dotar de dinamicidad a los servicios de un CDN
 - Creado por un grupo de empresas con Akamai a la cabeza. No es estándar del W3C
 - Las etiquetas ESI son directivas que permiten instruir al surrogate server sobre las acciones a realizar. Permite
 - Incluir fragmentos de página
 - Reaccionar a variables de cookies o cabeceras HTTP
 - Establecer condiciones en sus acciones
 - Manejar condiciones de error

ESI



```
<esi:choose>
<esi:when test="$(HTTP_COOKIE{group})=='Advanced'"> <esi:include
src="http://www.example.com/advanced.html"/>
</esi:when>
<esi:when test="$(HTTP_COOKIE{group})=='Basic User'"> <esi:include
src="http://www.example.com/basic.html"/>
</esi:when>
<esi:otherwise>
<esi:include src="http://www.example.com/new_user.html"/>
</esi:otherwise>
</esi:choose>
```

ESI

```

<table>
<tr>
<td colspan="2">
<esi:try>
<esi:attempt>
<esi:include>
src="http://www.myxyz.com/news:
onerror="continue"/>
</esi:attempt>

<esi:except>
<!--esi
This spot is reserved for your comp
advertising. For more info
<a href="www.myxyz.com">click he
-->
</esi:except>
</esi:try>
</td> </tr>
</table>

```

Welcome To My XYZ.com

[Shopping](#)
[News](#)
[Sports](#)
[Fun](#)
[XYZ.com](#)

Local Weather

[TTL-]

Stock Quotes

[TTL-]

What's New at XYZ?

```

<table>
<tr>
<td colspan="2">
<esi:try>
<esi:attempt>
<esi:include>
src="http://www.myxyz.com/news/top.html"
onerror="continue"/>
</esi:attempt>
<esi:except>
<!--esi
This spot is reserved for your company's
advertising. For more info
<a href="www.myxyz.com">click here</a>
-->
</esi:except>
</esi:try>
</td> </tr>
</table>

```

Today's News

[Breaking Headlines TTL-]

4.6 Proveedores CDN

Adero
(www.adero.com)



Akamai
(www.akamai.com)



Cidera
(www.cidera.com)



Pushcache
(www.pushcache.com)



EpicRealm
(www.epicrealm.com)



CacheWare
(www.cacheware.com)



Clearway
(www.clearway.com)



IBeam
(www.ibeam.com)



Mirror Image
(www.mirror-image.com)



Digital Island
(www.digitalisland.com),
includes the former Sandpiper
Networks



Proveedores CDN

Table 1. Content delivery network provider characteristics.

| CDN service provider | Service type | Content distribution | Fees | Customers |
|--|--|--|--|---|
| Akamai www.akamai.com | Multi-ISP, partial-site request servicing, peering | More than 12,000 surrogate servers spanning 1,000 networks in 62 countries | US\$1,995 per month for each Mbps of delivered content | Covers 70 percent of the market, with more than 3,600 customers including Apple, CNN, MSNBC, Reuters, and Yahoo |
| Adero www.webvisions.com/ adero/ | Multi-ISP, full-site request servicing, peering | Surrogate servers in more than 30 countries | Depends on resellers (CDNs that buy Adero services) | Serves 30 customers, including resellers Exodus and UUNET |
| Digital Island www.sandpiper.net | Multi-ISP, partial-site request servicing, peering | 2,500 surrogate servers spanning 327 networks in 35 countries | Starts at US\$1,500 per month | More than 900 customers including AOL, Canon, Cisco Systems, Microsoft, and Hewlett Packard |
| Mirror Image www.mirror-image.com | Multi-ISP, partial-site request servicing, peering | 22 surrogate servers in North America, Europe, and Asia | US\$2,100 per month for each Mbps of delivered content | More than 200 customers including Creative, Open Systems, and SiteRock |
| Inktomi www.inktomi.com | Single-ISP, full-site request servicing, peering | 10 surrogate servers across China | Starts at US\$4,000 per month | 13 CDNs including Adero and Digital Island and more than 200 Web sites |

Año 2003

Proveedores CDN

| CDN | Routing Methods | POPs per continent | | | | | |
|-------------------|-----------------------|---|---------------|--------|-----------------|----------------|--------|
| | | North America | South America | Europe | Asia | Oceania | Africa |
| Akamai | DNS/Proprietary | Akamai claims to have over a hundred thousand servers in 75 countries. Many of these servers are located within ISP facilities | | | | | |
| Amazon CloudFront | DNS/EDNS /Proprietary | 14 | 2 | 10 | 9 | 1 | 0 |
| Azure CDN | DNS | 8 | 1 | 8 | 7 | 1 | 0 |
| CacheFly | IP Anycast/DNS/EDNS | 11 | 1 | 10 | 9 | 2 | 2 |
| CDNetworks | DNS/Proprietary | 12 | 7 | 21 | 42 ¹ | 1 | 4 |
| EdgeCast | DNS/EDNS /Proprietary | 9 | 1 | 11 | 11 ² | 2 | 0 |
| Internap CDN | DNS/Proprietary | 7 | 0 | 2 | 2 | 0 | 0 |
| Level 3 CDN | DNS/Proprietary | 19 | 3 | 13 | 8 | 1 | 0 |
| Limelight | DNS/Proprietary | 11 ³ | 0 | 4 | 3 | 1 | 0 |
| MaxCDN | IP Anycast/DNS | 9 | 0 | 3 | 1 ⁴ | 1 ⁴ | 0 |

¹ Some CDNetwork Asia POPs are located in China and Russia and require special permission to use

² Some EdgeCast Asia POPs are located in China and require special permission to use

³ Limelight POP numbers and locations are estimated because specifics are not disclosed

⁴ MaxCDN POPs outside of North America and Europe require special provisioning and an additional monthly fee

Proveedores CDN

- Retardo muy diferente para algunas localizaciones sin servidores cercanos

CDN Latency - 2014 (milliseconds)

| Service | North America | Europe | South America | Africa | Asia | Oceania |
|-------------------|---------------|--------|---------------|--------|--------|---------|
| Akamai | 56.88 | 48.15 | 94.84 | 84.06 | 114.24 | 103.53 |
| Amazon CloudFront | 61.07 | 58.78 | 119.09 | 198.45 | 105.75 | 92.31 |
| Azure CDN | 53.63 | 56.71 | 160.10 | 160.83 | 118.95 | 94.15 |
| CacheFly | 54.71 | 62.63 | 105.07 | 177.78 | 138.21 | 76.11 |
| CDNetworks | 58.91 | 64.57 | 169.37 | 115.82 | 107.93 | 98.10 |
| EdgeCast | 52.49 | 52.58 | 170.06 | 184.82 | 126.22 | 79.06 |
| Internap CDN | 66.67 | 82.86 | 192.73 | 203.58 | 179.72 | 229.42 |
| Level 3 CDN | 67.30 | 69.31 | 149.99 | 100.79 | 179.24 | 215.00 |
| LimeLight | 53.74 | 59.97 | 155.32 | 173.33 | 104.95 | 121.35 |
| MaxCDN | 56.09 | 65.86 | 184.99 | 183.56 | 255.82 | 240.92 |


CDN Small File Throughput - 2014 (Mb/s)

| Service | North America | Europe | South America | Africa | Asia | Oceania |
|-----------------------|---------------|--------|---------------|--------|------|---------|
| Akamai | 2.07 | 2.01 | 0.77 | 1.60 | 1.07 | 0.65 |
| Amazon CloudFront | 2.20 | 1.56 | 0.85 | 0.53 | 1.19 | 1.04 |
| Azure CDN (Microsoft) | 2.82 | 2.00 | 0.67 | 0.99 | 1.37 | 1.03 |
| CacheFly | 3.88 | 3.21 | 2.11 | 0.73 | 1.97 | 2.13 |
| CDNetworks | 2.38 | 2.36 | 0.53 | 0.90 | 1.25 | 1.32 |
| EdgeCast | 3.37 | 2.88 | 0.91 | 0.48 | 1.63 | 1.93 |
| Internap CDN | 1.57 | 1.34 | 0.74 | 0.44 | 0.68 | 0.39 |
| Level 3 CDN | 2.05 | 1.91 | 0.98 | 0.83 | 1.46 | 0.47 |
| LimeLight | 2.50 | 1.60 | 0.61 | 0.55 | 1.02 | 0.97 |
| MaxCDN | 2.78 | 2.44 | 0.40 | 0.62 | 0.53 | 0.55 |

CDN Large File Throughput - 2014 (Mb/s)

| Service | North America | Europe | South America | Africa | Asia | Oceania |
|-----------------------|---------------|--------|---------------|--------|------|---------|
| Akamai | 15.11 | 13.71 | 4.41 | 1.83 | 7.39 | 5.73 |
| Amazon CloudFront | 10.11 | 8.17 | 4.49 | 1.13 | 4.79 | 2.47 |
| Azure CDN (Microsoft) | 15.54 | 13.33 | 3.40 | 2.35 | 5.62 | 4.88 |
| CacheFly | 18.62 | 18.51 | 3.82 | 2.57 | 8.26 | 9.38 |
| CDNetworks | 15.97 | 17.52 | 2.61 | 1.76 | 6.19 | 6.21 |
| EdgeCast | 17.45 | 15.17 | 2.79 | 1.18 | 7.39 | 6.98 |
| Internap CDN | 13.18 | 10.65 | 2.46 | 1.51 | 4.45 | 2.28 |
| Level 3 CDN | 12.32 | 9.86 | 3.46 | 1.45 | 4.22 | 1.95 |
| LimeLight | 14.40 | 9.48 | 3.14 | 0.92 | 6.33 | 4.53 |
| MaxCDN | 16.22 | 16.89 | 1.89 | 1.59 | 2.93 | 2.52 |

Proveedores CDN



Año 2001

4.6.1 Akamai

- ▶ Datos 2010
 - Servidores ~61.000
 - Redes ~1.000
 - Países ~70
- ▶ Clientes más importantes
 - Apple, BBC, FOX, GM IBM, MTV, NASA, NBC, NFL, NPR, Puma, Red Bull, Rutgers, SAP, ...
- ▶ Tráfico cursado
 - Centenares de billones de peticiones al día
 - De las 45 redes con más tráfico cursado, la mitad son de Akamai
 - Cursa el 15-20% del tráfico web a nivel mundial

Akamai

- ▶ Servidores en POPs de proveedores más importantes.
- ▶ Dispone de su propia red de servidores de DNS.
- ▶ Basado en combinación redirección DNS y reescritura
 - URL embebida: ARL (Akamai Resource Locator)

http://www.foo.com/a.gif



http : //a ^{Serial #} 836 ^{Akamai Domain} .g.akamaitech.net / ^{Type} 7 / ^{Serial #} 836 / ^{Provider Code} 123 / ^{Object Data} e358f5db0045 / ^{absoluteURL} www.foo.com/a.gif

- Serial: identifica a un conjunto de servidores CDN.
- Provider code: identificador del cliente (servidor original).
- Object data: hash del objeto, para detectar versiones del mismo.

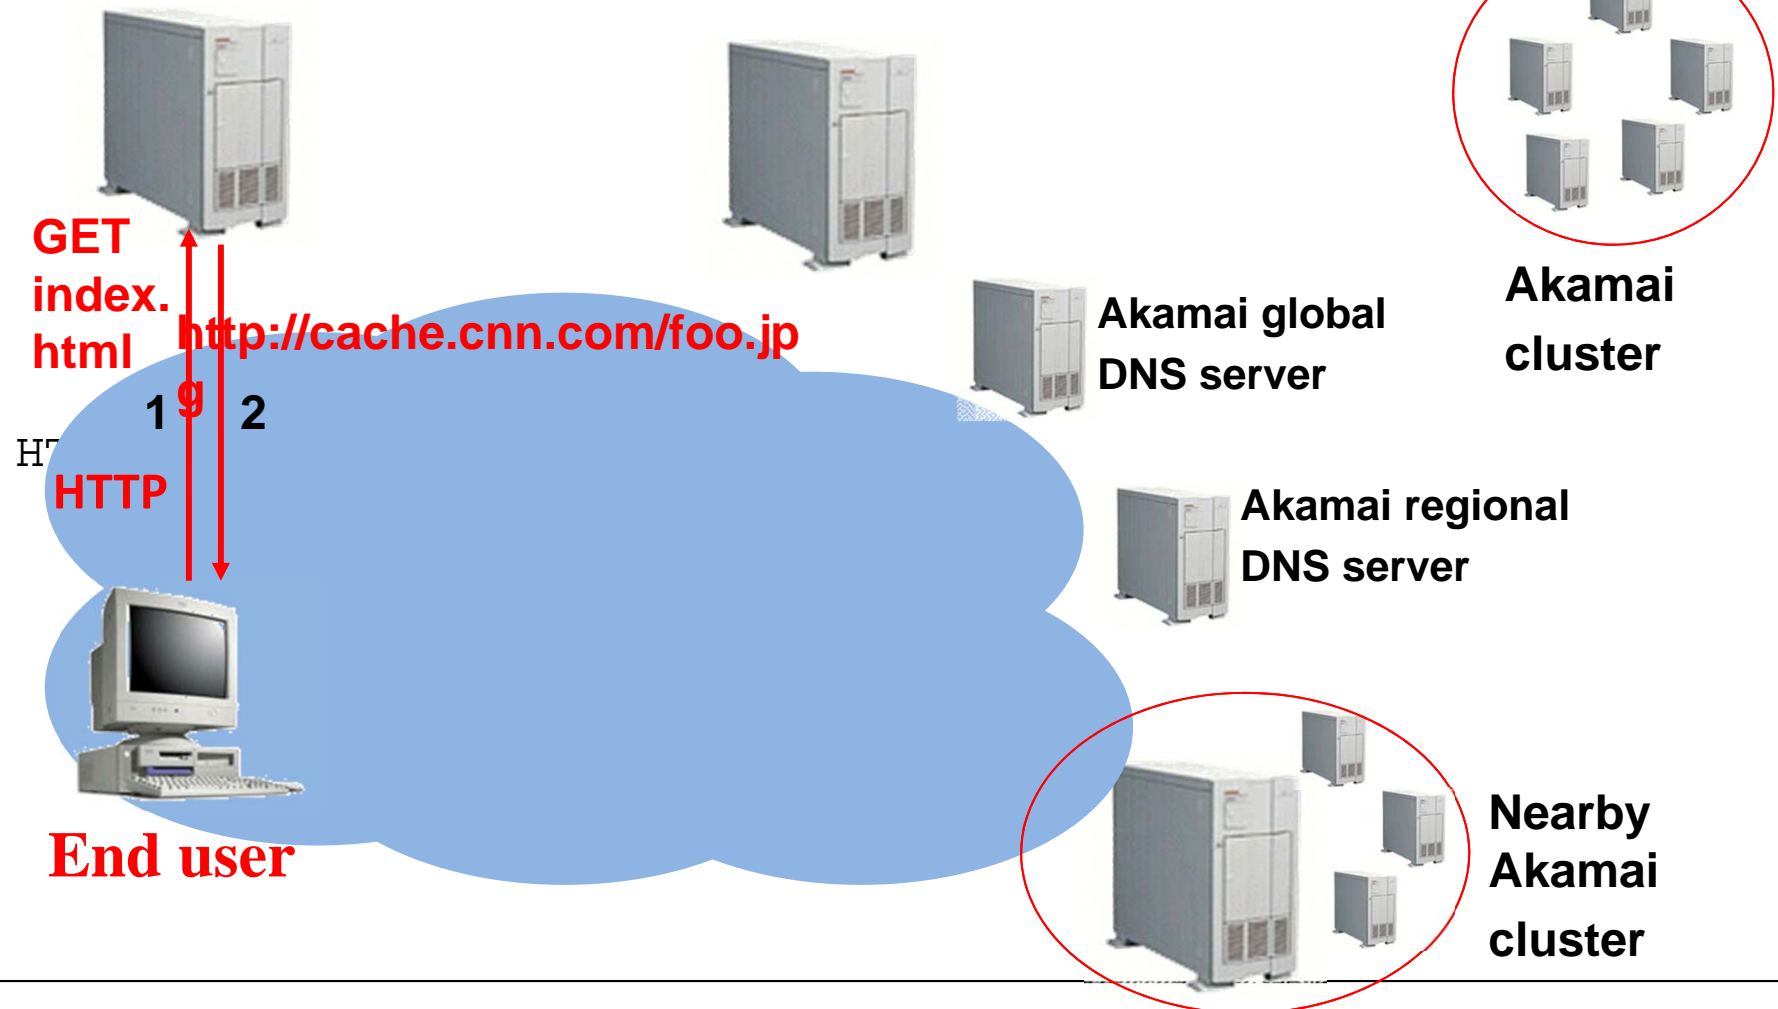
Akamai

◦ Redirección DNS

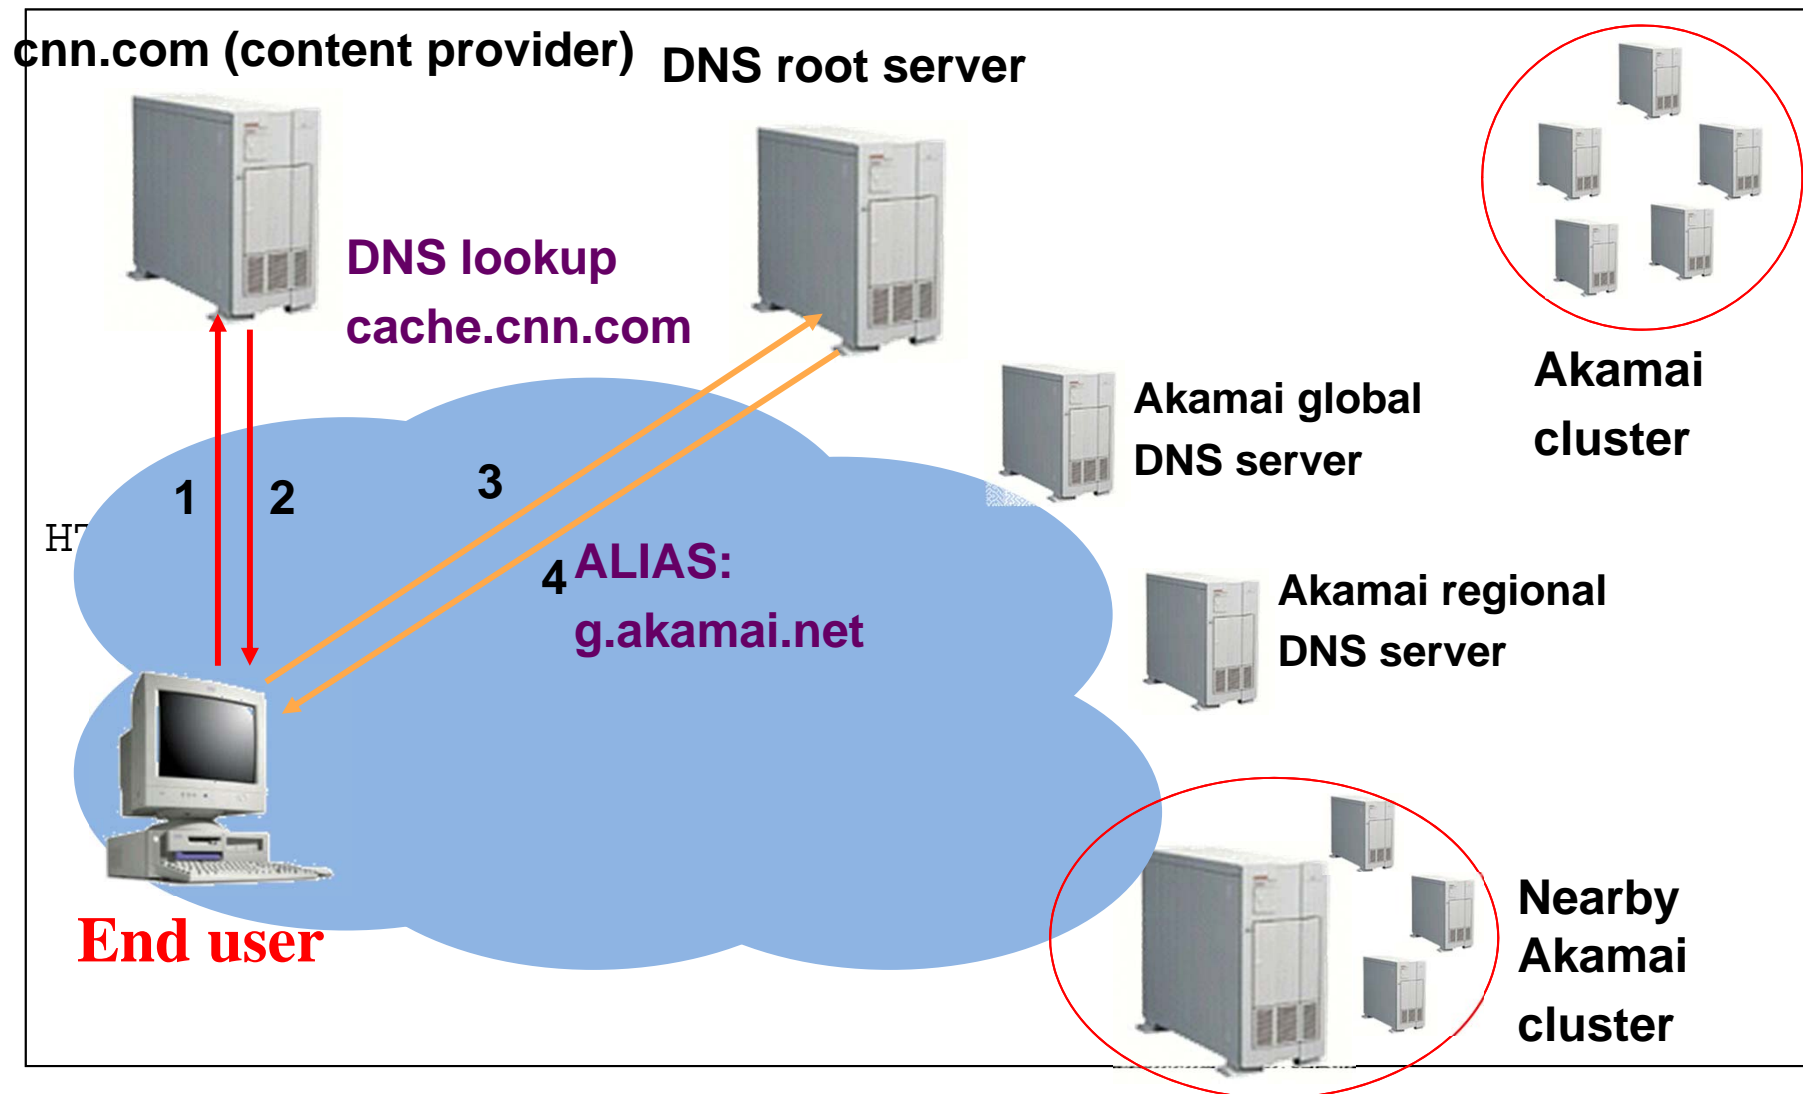
- Tiene en cuenta la IP origen que realiza la petición (localización).
- Utilizar una red de servidores DNS en 2 niveles:
 - high-level.akamai.net servers (HLDNS)
 - low-level.g.akamai.net servers (LLDNS)
- Los servidores raíz de Internet (*.net) se encargan de devolver al cliente las IPs de HLDNS de las que selecciona una,
 - DNS raíz, TTLs de alrededor de 2 días.
- Los servidores de alto nivel son los que reciben las peticiones de resolución (*.akamai.net) y la redireccionan a servidores de segundo nivel (*.g.akamai.net) que estén lo más cercano posible al cliente.
 - HLDNS, TTLs de alrededor de 20 minutos.
- El servidor de bajo nivel es el que hace la resolución final devolviendo la IP del servidor CDN “más cercano” para ese cliente. Basado en mapas de la red que incorporan el estado de la red y de los servidores, actualizados cada 2-10 segundos.
 - LLDNS, TTLs de alrededor de 20 segundos.
- TTLs compromiso entre disminuir el tráfico de peticiones DNS y ahorrarse el tiempo de resolución, con el de tener la información lo más actualizada posible.

Ejemplo Akamai

cnn.com (content provider) DNS root server

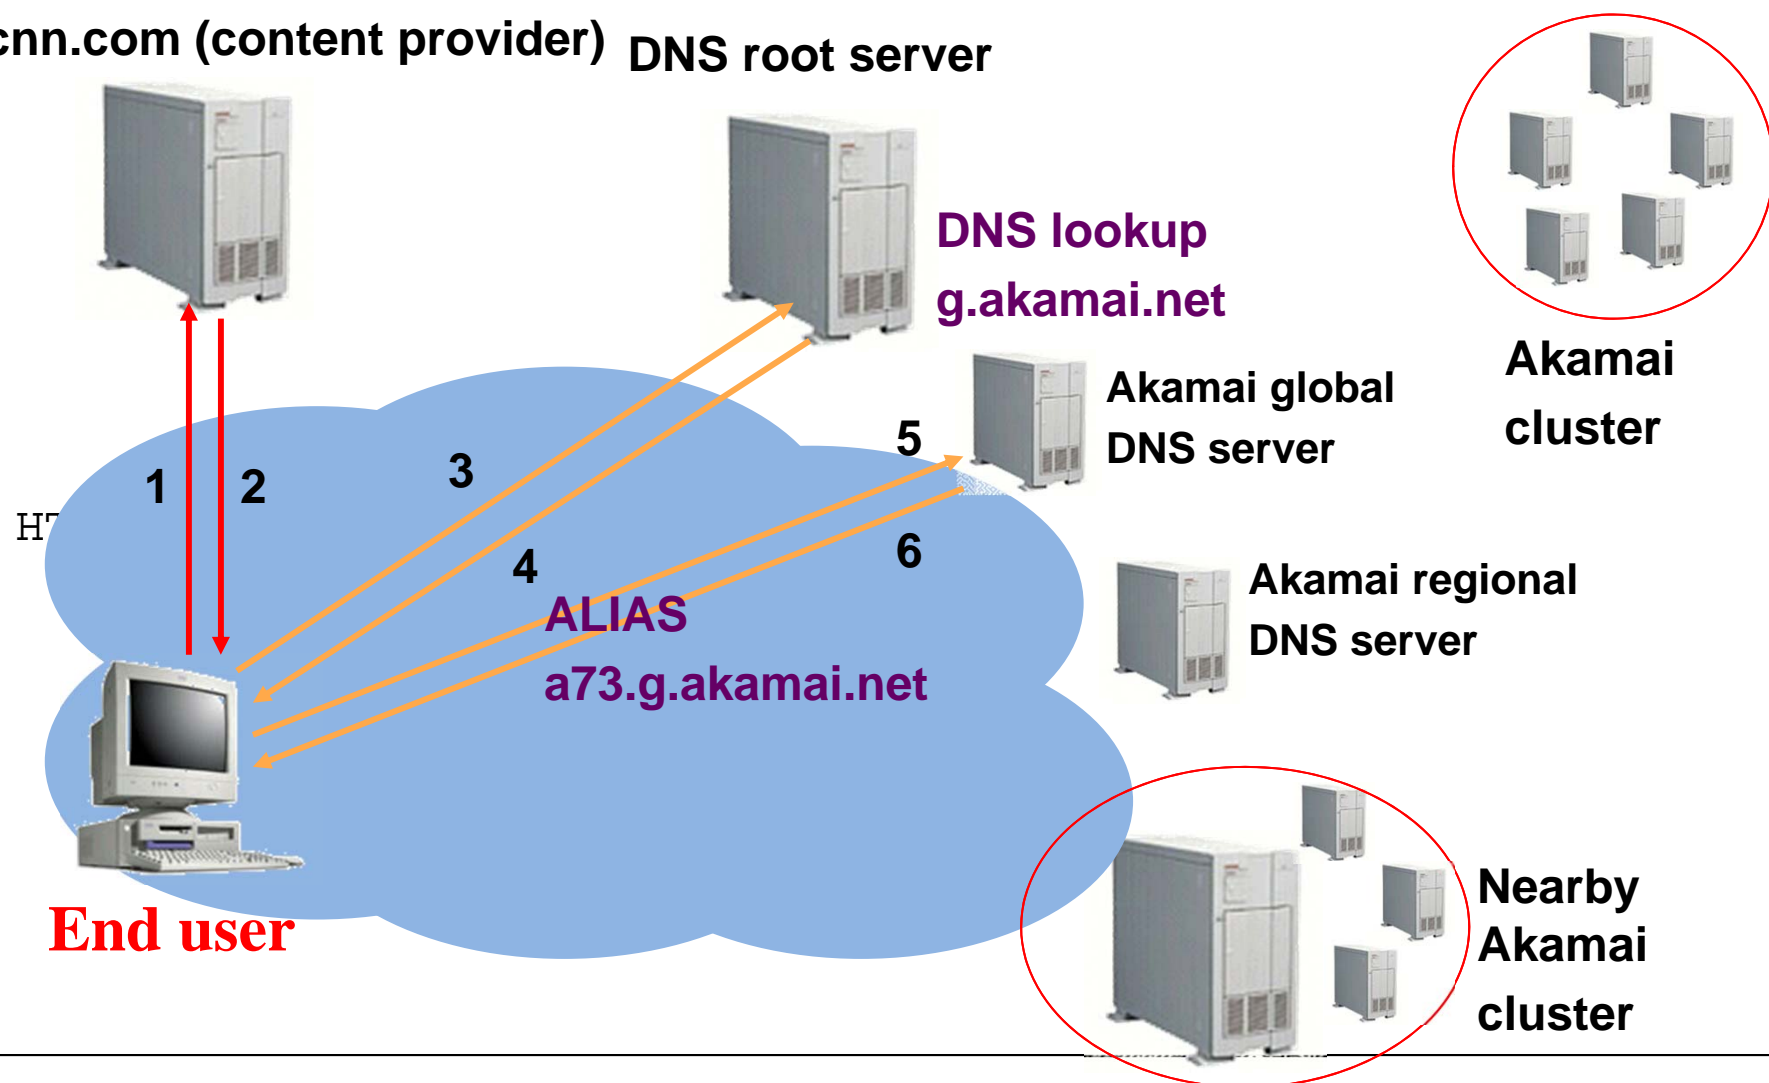


Ejemplo Akamai



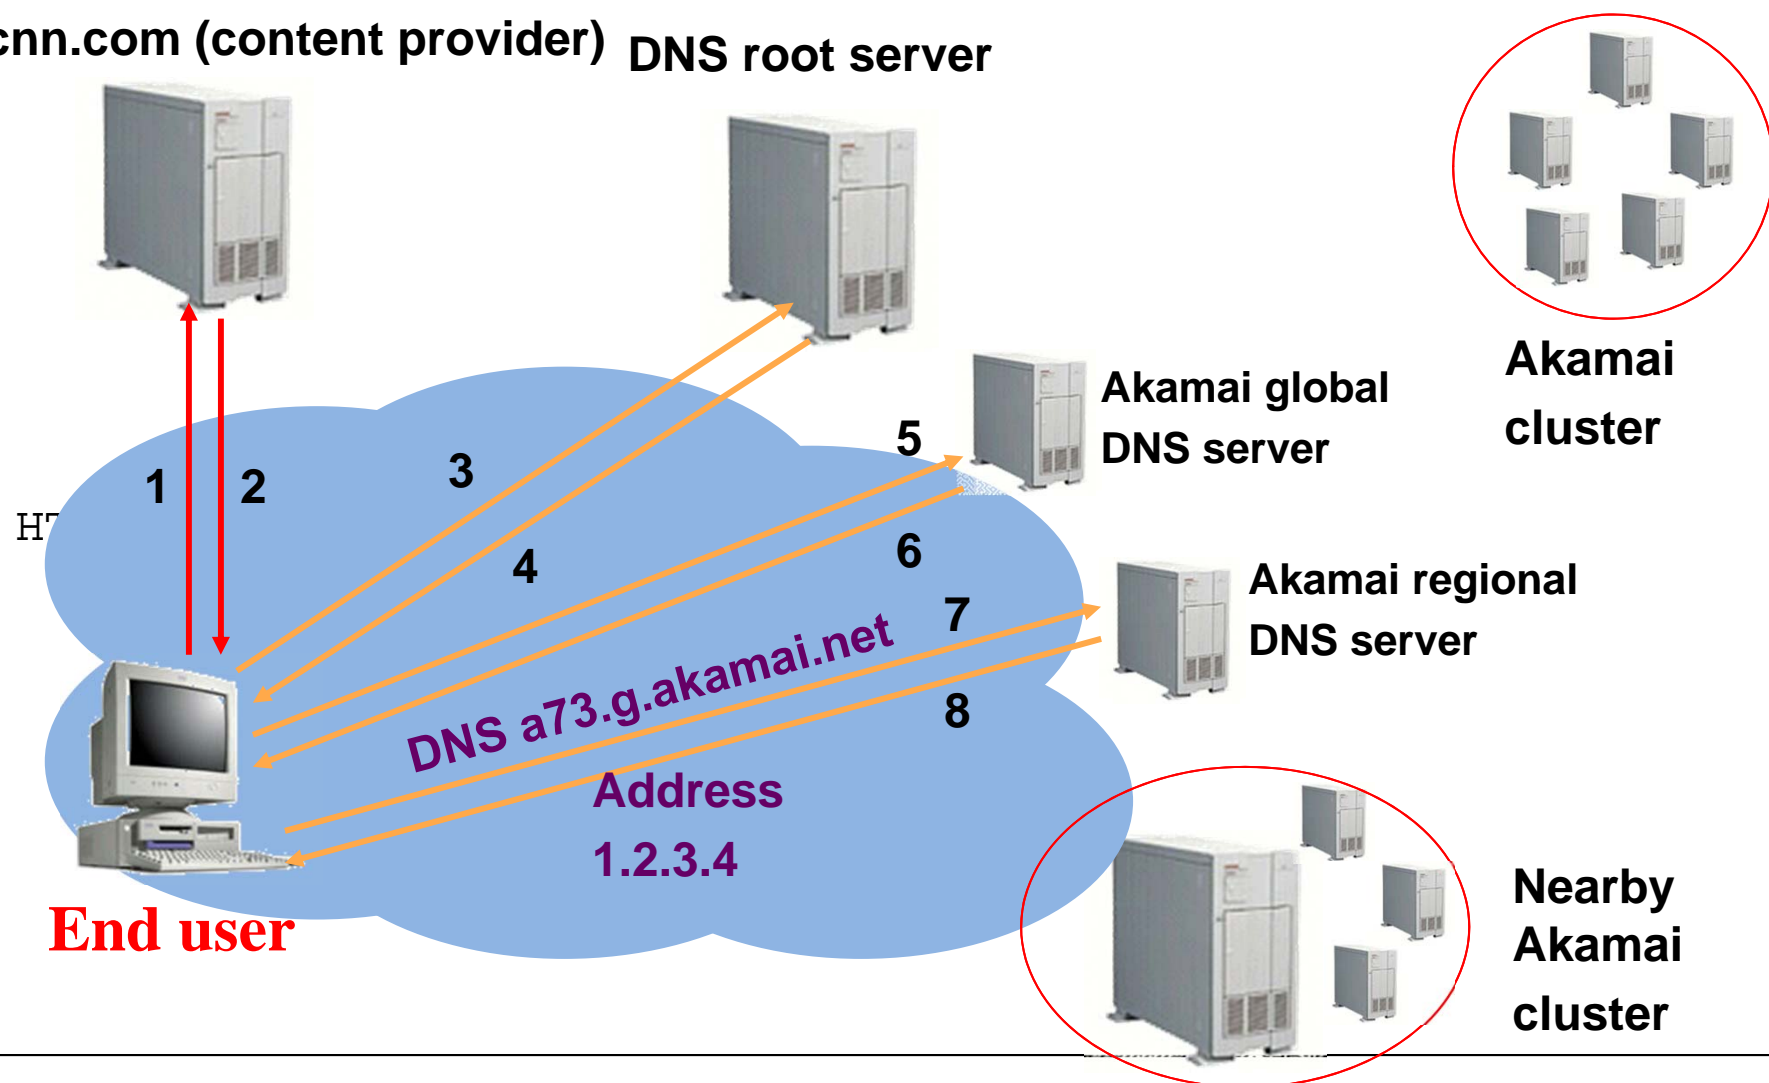
Ejemplo Akamai

cnn.com (content provider) DNS root server



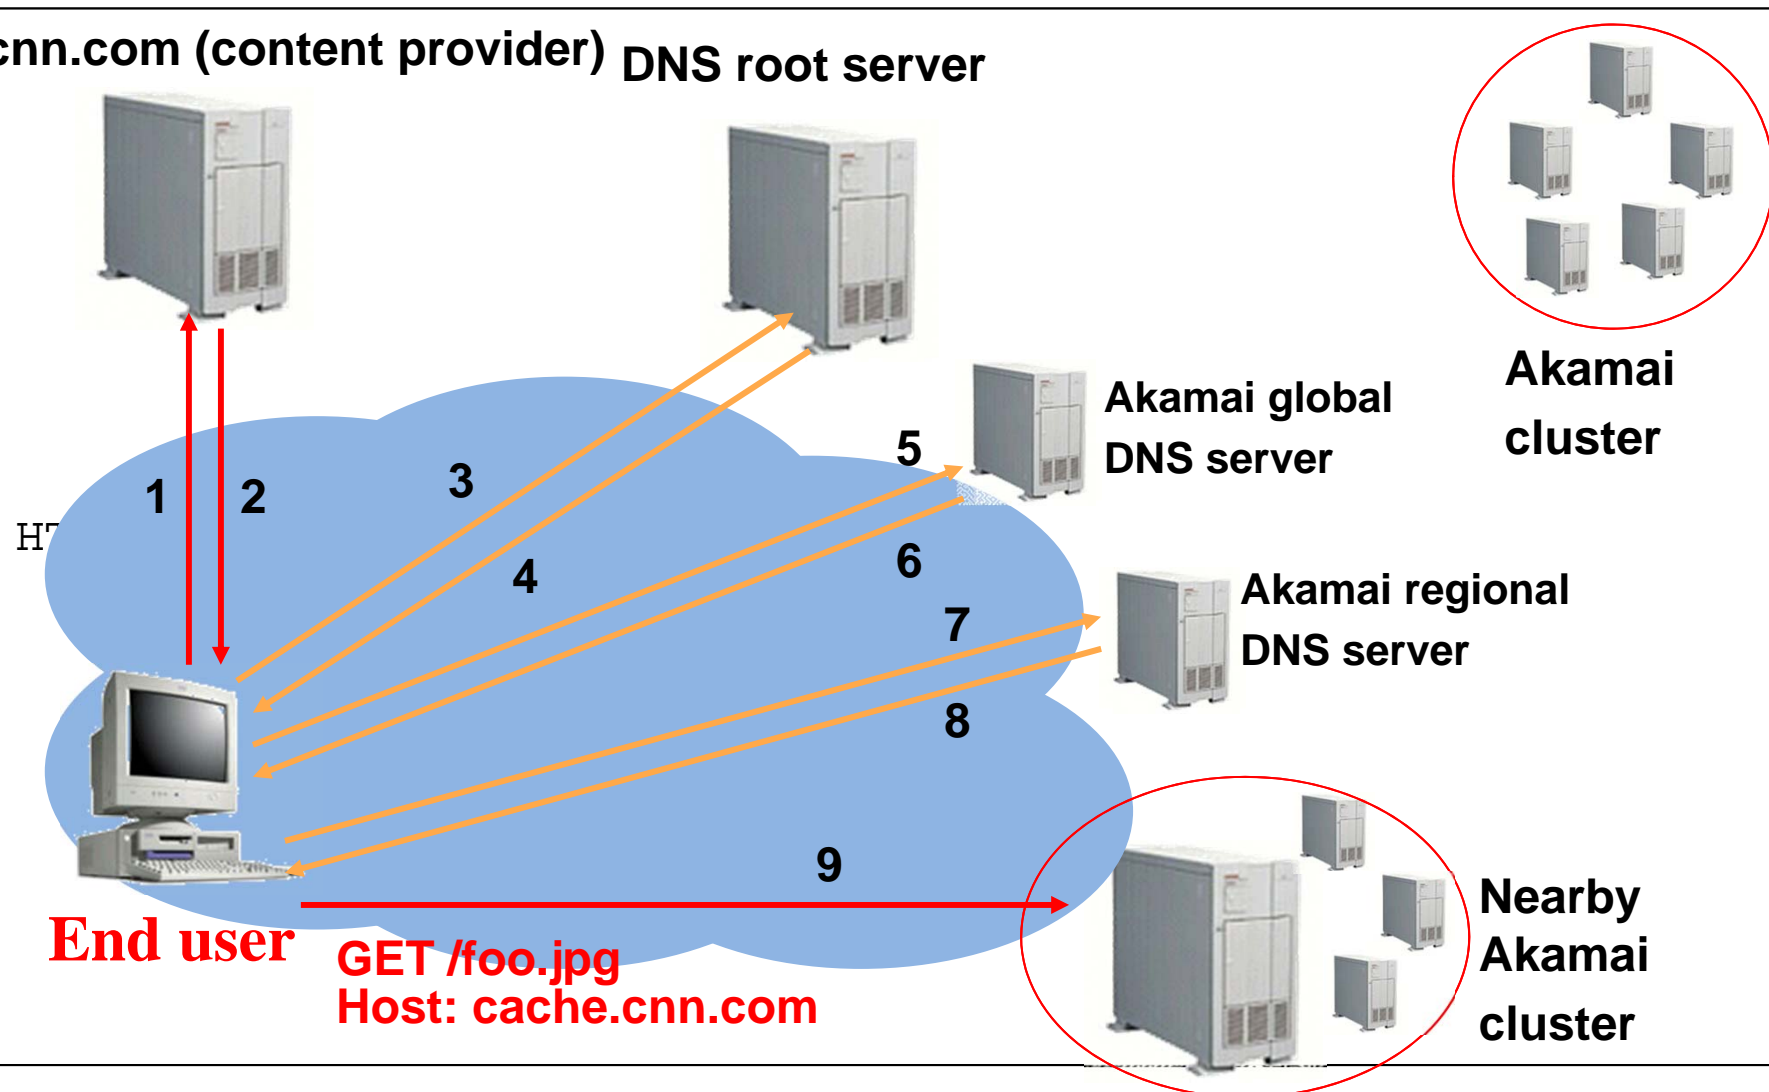
Ejemplo Akamai

cnn.com (content provider) DNS root server



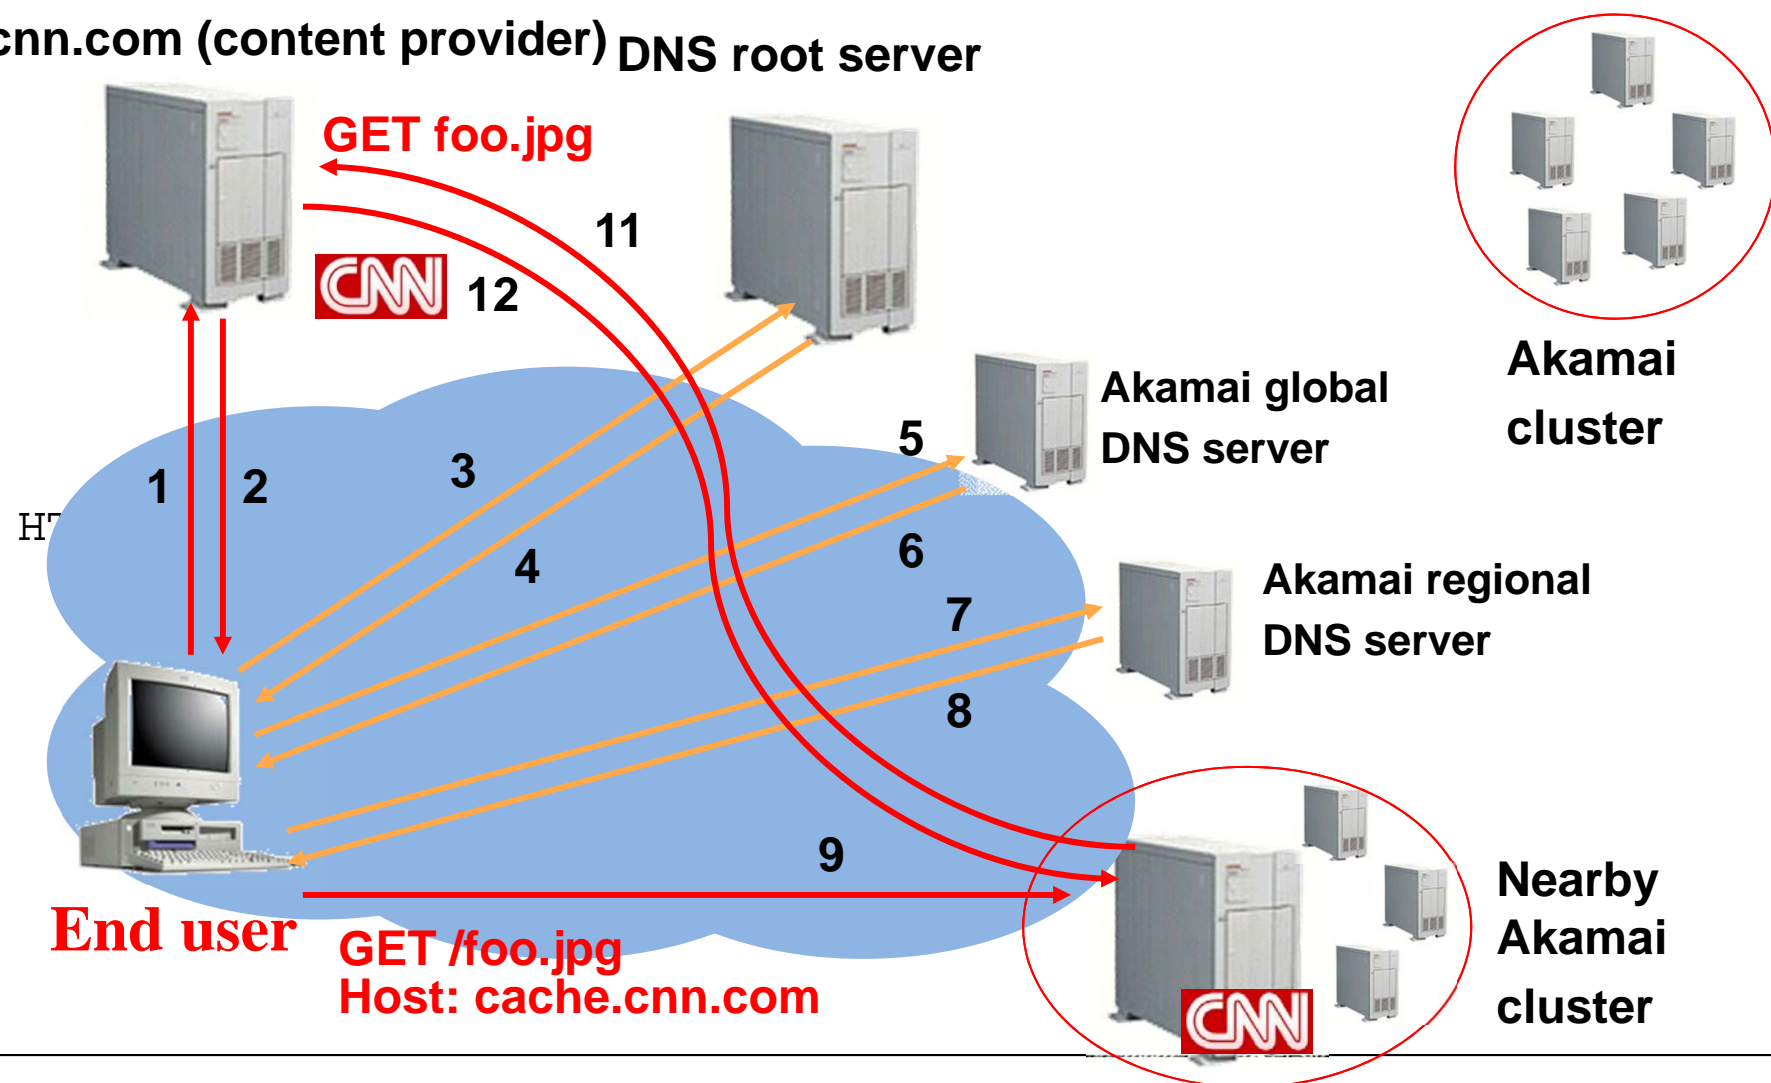
Ejemplo Akamai

cnn.com (content provider) DNS root server



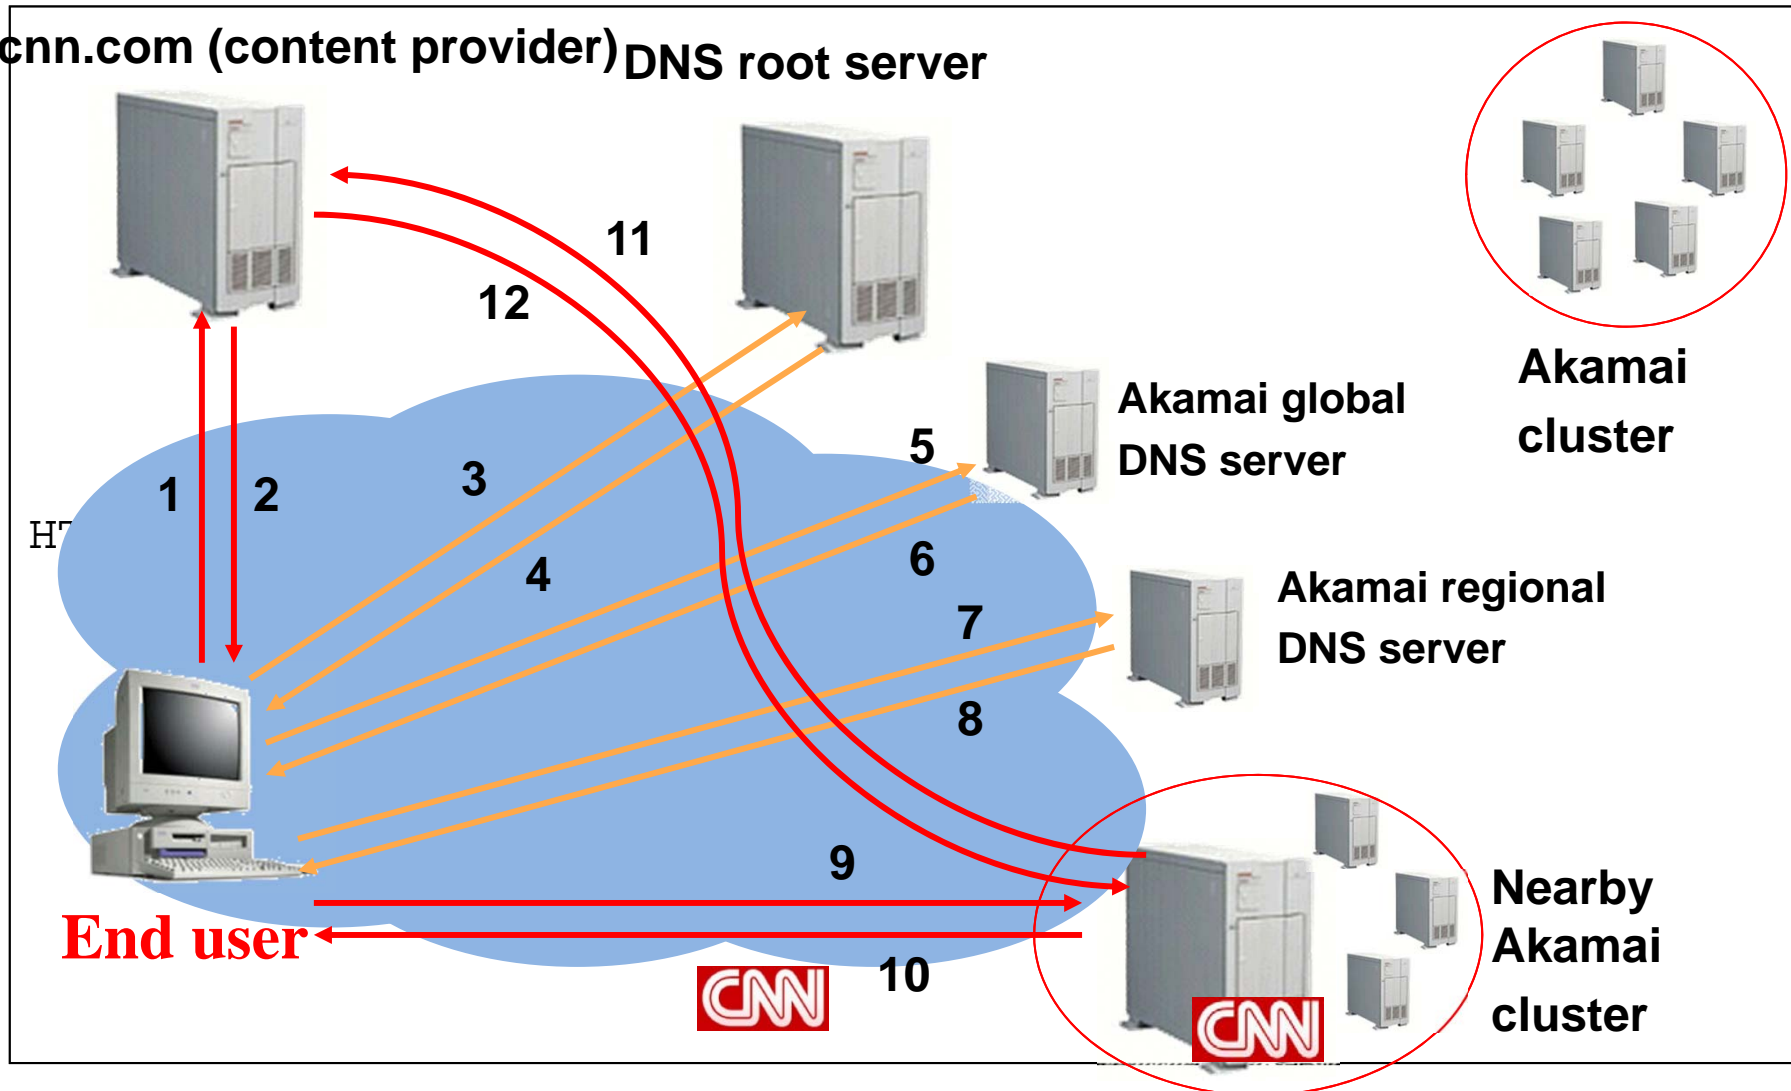
Ejemplo Akamai

cnn.com (content provider) DNS root server




Ejemplo Akamai

cnn.com (content provider) DNS root server



Ejemplo Akamai: Cache Hit




Akamai, eficiencia en la localización

[Johnson00]

Distribución acumulada del retardo.

- Gris: para cada servidor. Gran variabilidad.
- Rojo: media de todos los servidores.
- Azul: Akamai, no el mejor.



4.7 Efecto de la versión HTTP en CDNs

Mean Download Performance Range for Different Numbers of Images and Protocol Options (Jan. 2001)

| Protocol Option | Site | Mean Download Time Range (sec.) | | | |
|-----------------|-----------|---------------------------------|-----------|-----------|-----------|
| | | 6 images | 12 images | 18 images | 54 images |
| Parallel-1.0 | CDN | 0.26-0.76 | 0.40-1.23 | 0.58-1.53 | 1.49-3.31 |
| | US Origin | 1.63 | 2.45 | 3.40 | 8.42 |
| Serial-1.1 | CDN | 0.27-0.53 | 0.42-0.81 | 0.61-1.13 | 1.46-2.52 |
| | US Origin | 1.06 | 1.46 | 1.96 | 4.87 |
| Pipeline-1.1 | CDN | 0.26-0.50 | 0.37-0.67 | 0.47-0.88 | 1.09-2.04 |
| | US Origin | Partial Support | | | |

CDNs perform significantly better than origin sites, although reducing the number of images (e.g. due to caching) and using HTTP/1.1 options reduces the performance difference.

Referencias

- ▶ [Kurose]
 - Capítulo 7, sección 7.5
- ▶ [Peterson]
 - Capítulo 9, sección 9.4
- ▶ [Held] Gilbert Held. A Practical Guide to Content Delivery Networks. CRC Press, 2 edition, 2010, ISBN-13: 978-1439835883
- ▶ [Hofmann] Markus Hofmann (Author), Leland R. Beaumont. Content Networking: Architecture, Protocols, and Practice (The Morgan Kaufmann Series in Networking), 2005. ISBN-13: 978-1558608344