Co (ne)přináší šifrování DNS

Do53, DoT, DoH, DoC, TRR & Firefox

Petr Špaček • petr.spacek@nic.cz • 14.11.2019

cz.nic | CZ DOMAIN REGISTRY

Obsah

- "Co znamená DoH ve Firefoxu?"
- Architektura
- Výkon a náklady
- Soukromí
- Závěr

DoH ve Firefoxu

- **Kombinace 4 změn najednou!**
- DNS-over-HTTPS (DoH)
  - DNS dotazy zabalené v HTTPS provozu
- Trusted Recursive Resolver (TRR)
  - resolver vybraný Mozillou/Firefoxem
- DNS v Cloudu (DoC)
  - resolver mimo hranice místní sítě
- DNS v aplikaci (add)
  - obchází stub resolver operačního systému

DoH ve Firefoxu

- Speciální "canary" doména
  - use-application-dns.net
  - Odpověď bez A/AAAA ⇒ vypne DoH
  - Prodlevy v řádu minut (!)
- https://support.mozilla.org/en-US/kb/canary-domain-use-application-dnsnet

Architektura

Tradiční DNS architektura

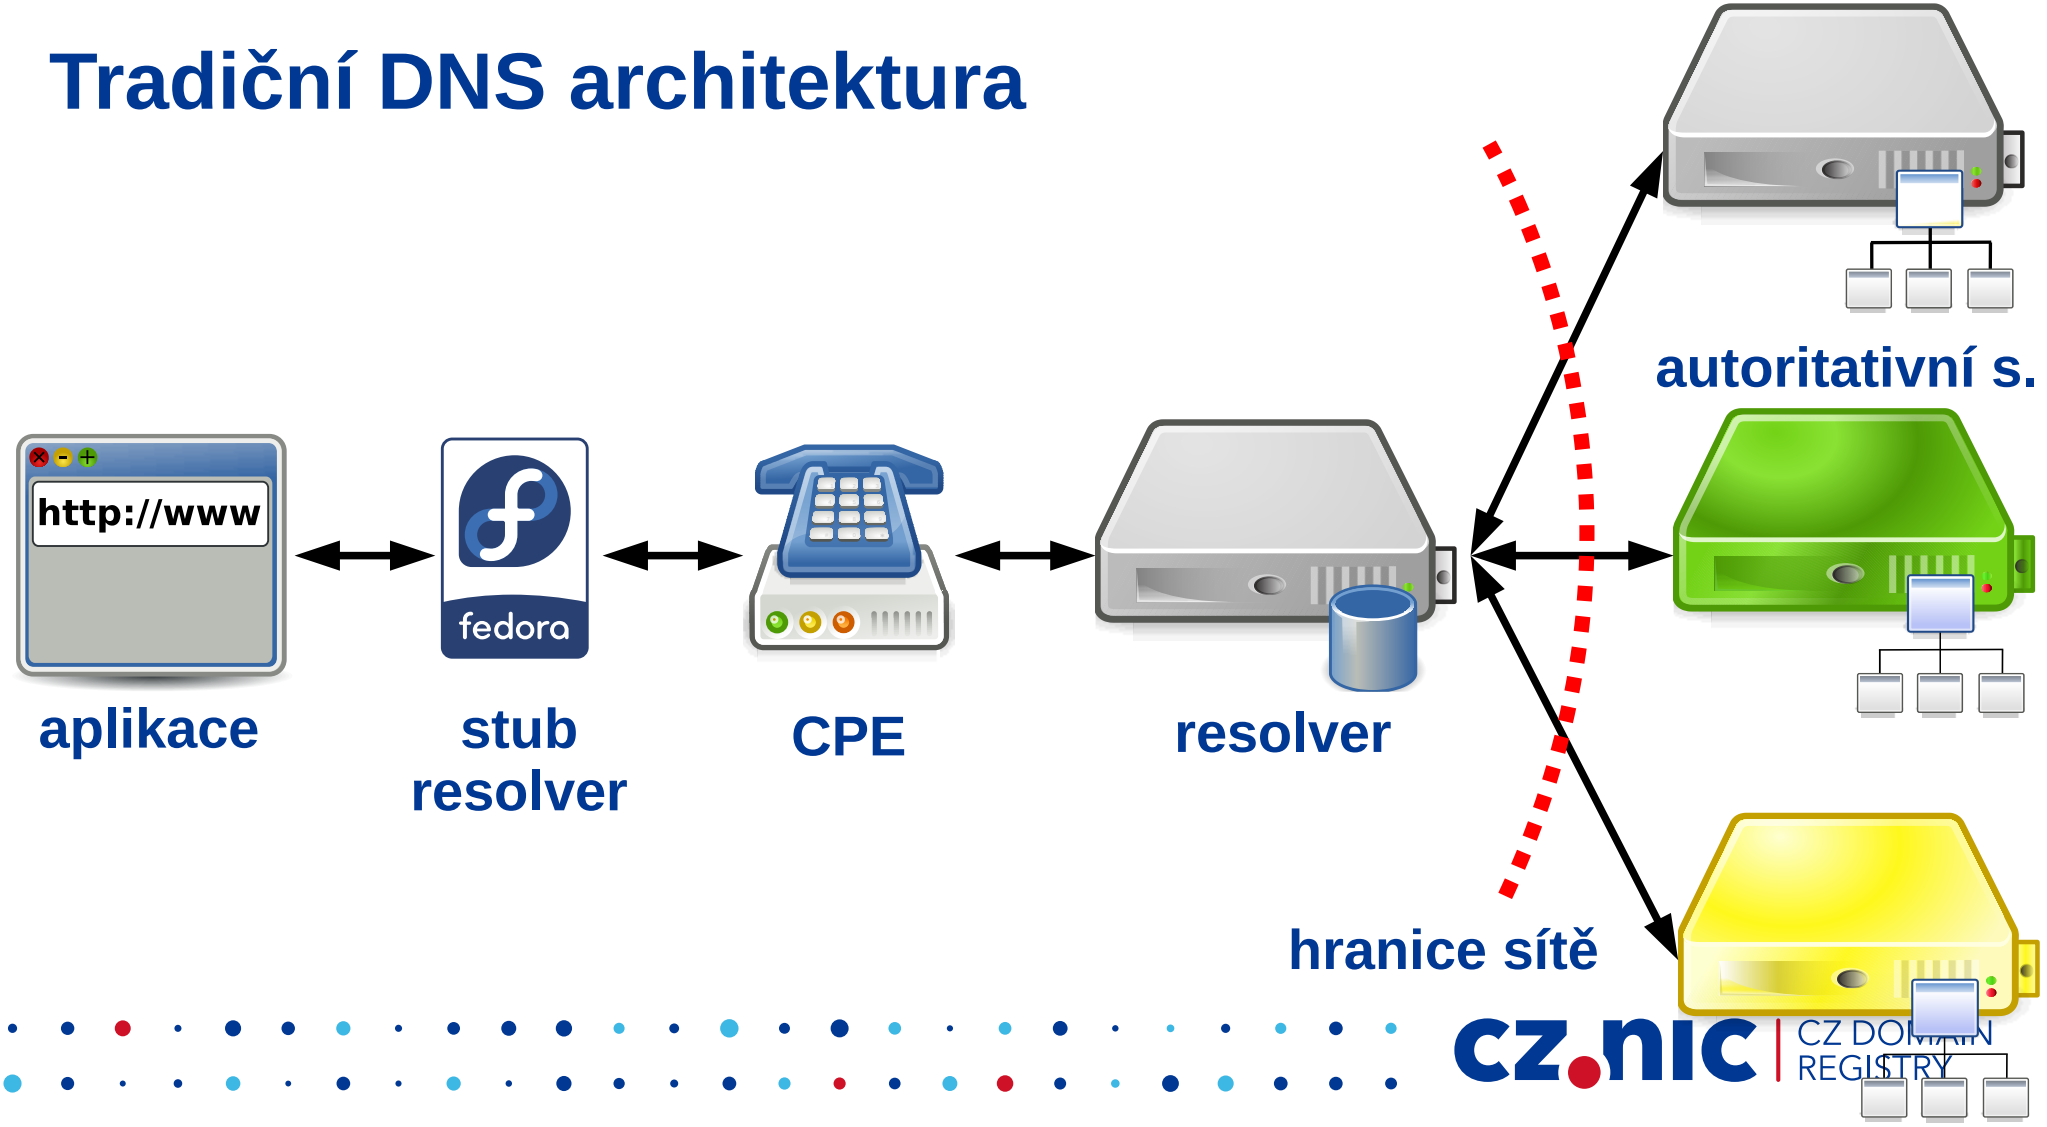

DNS v Cloudu

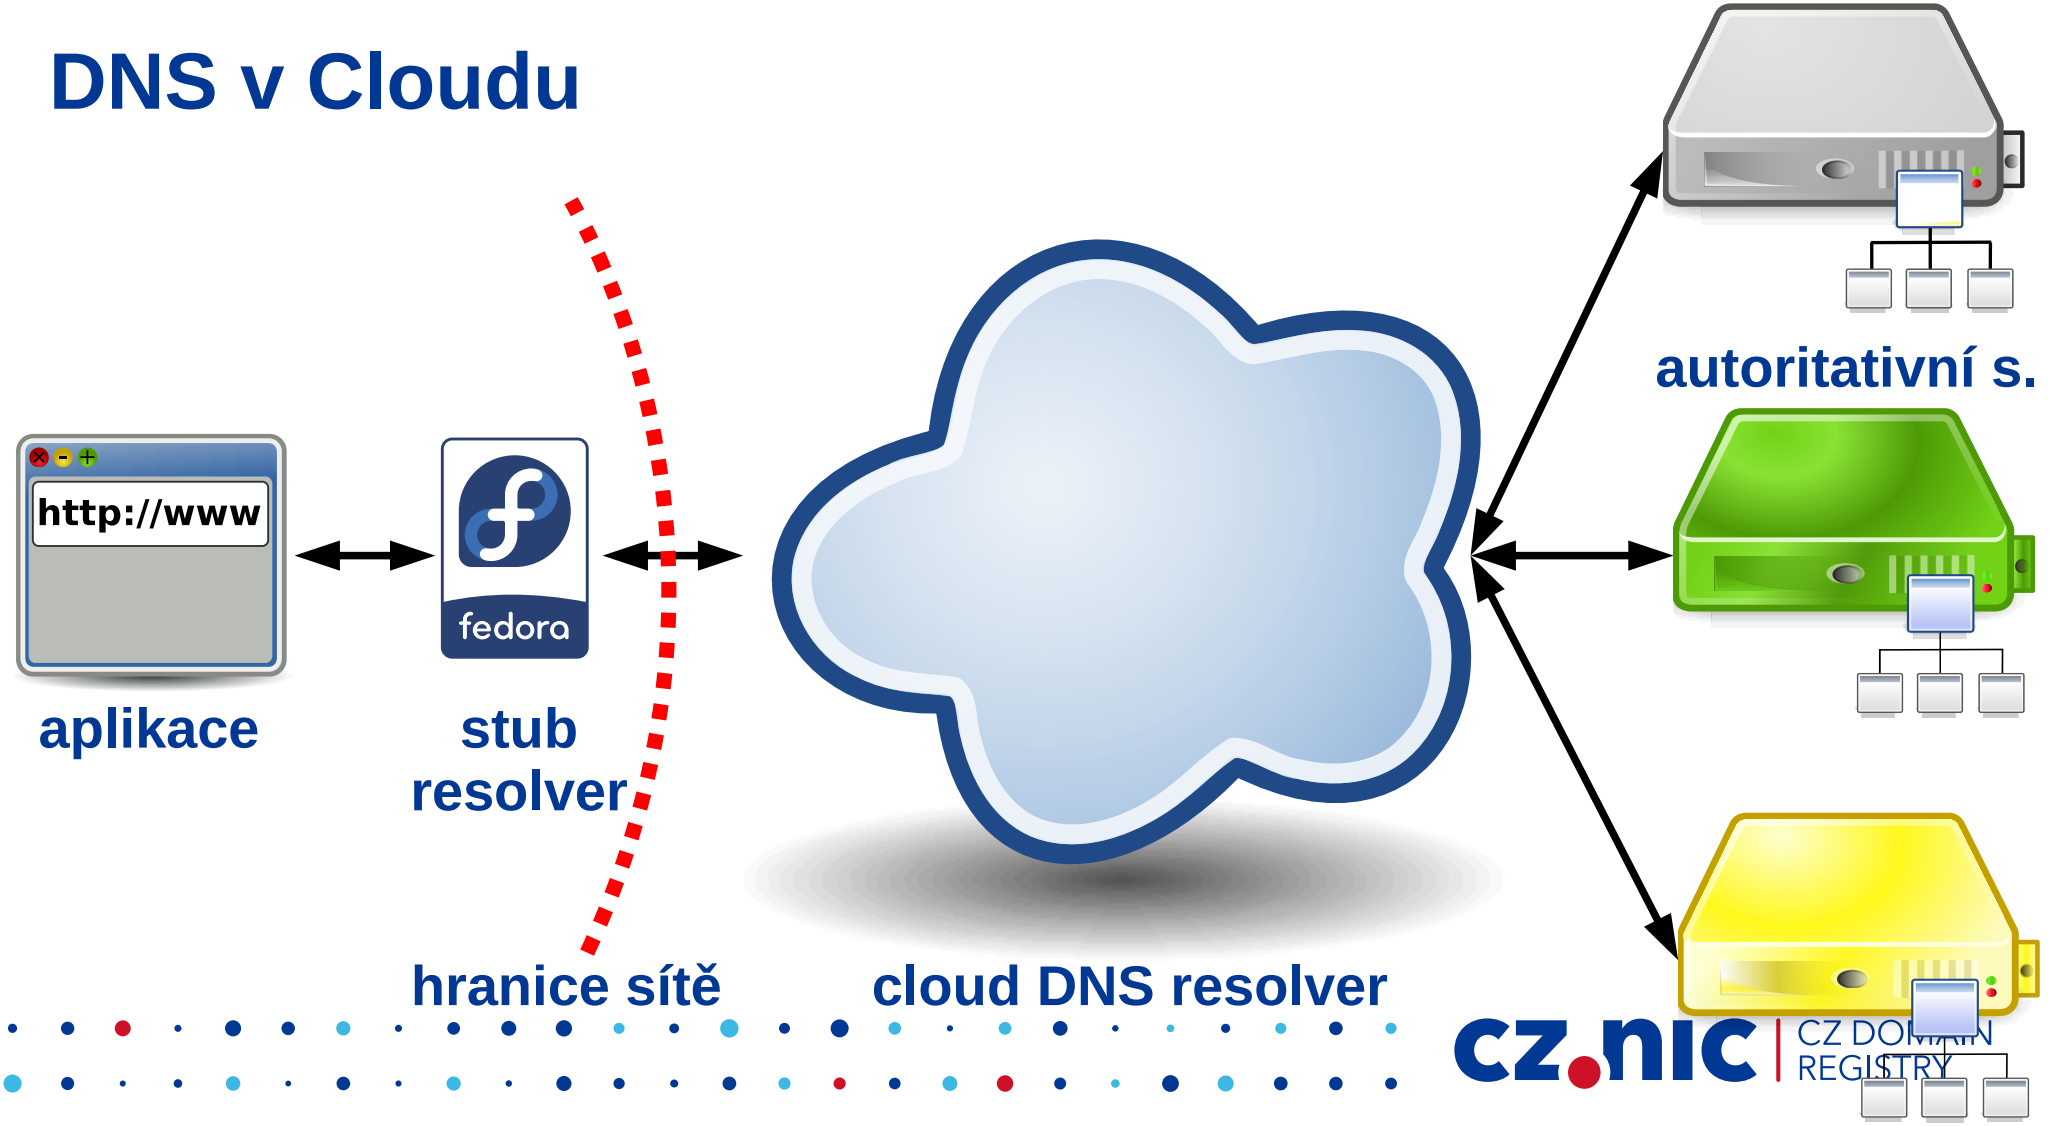
http://www
fedora
aplikace
stub resolver
hranice sítě
cloud DNS resolver
autoritativní s.
cz.nic
CZ DOMAIN REGISTRY

DNS z aplikace do Cloudu

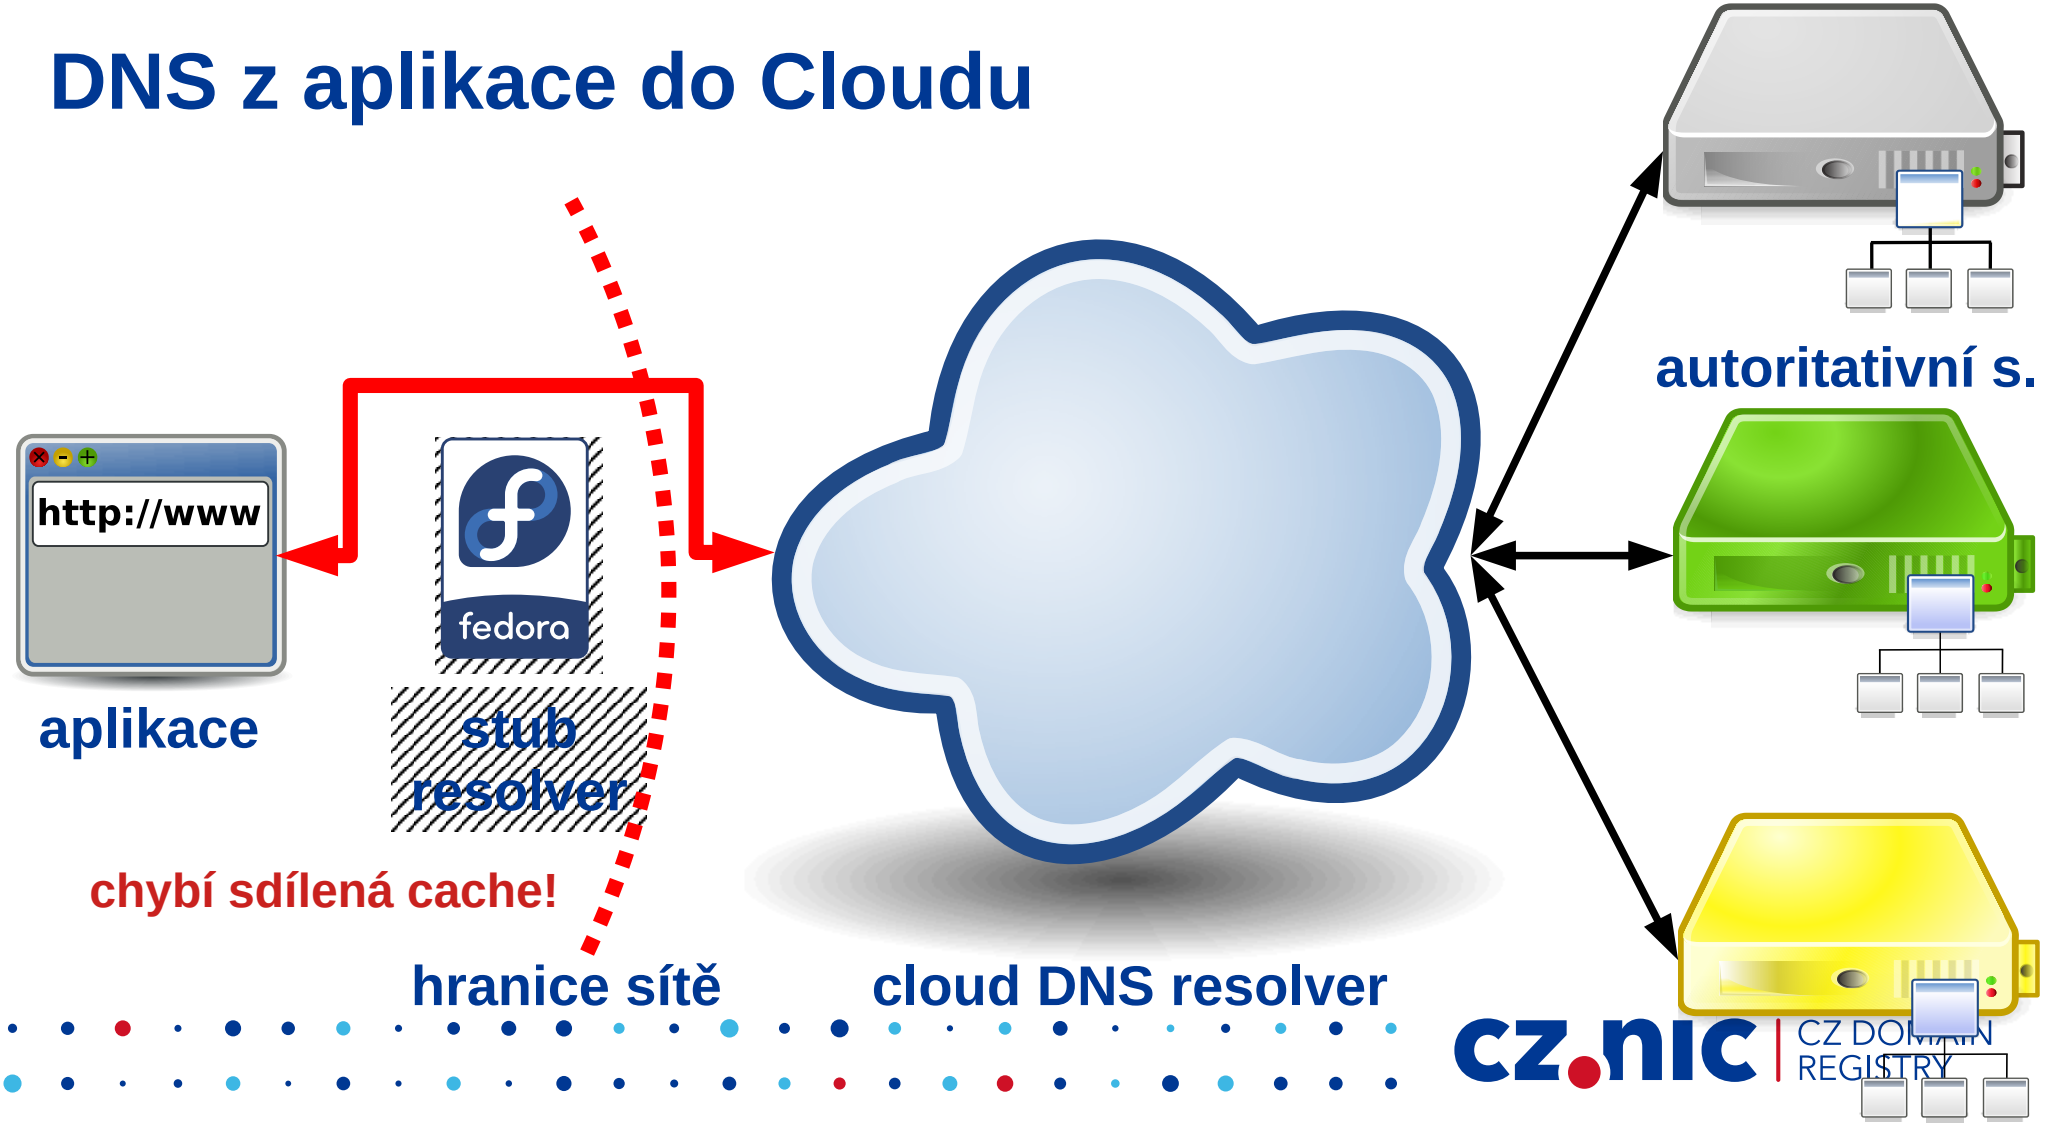

Víc cloudů víc ví!

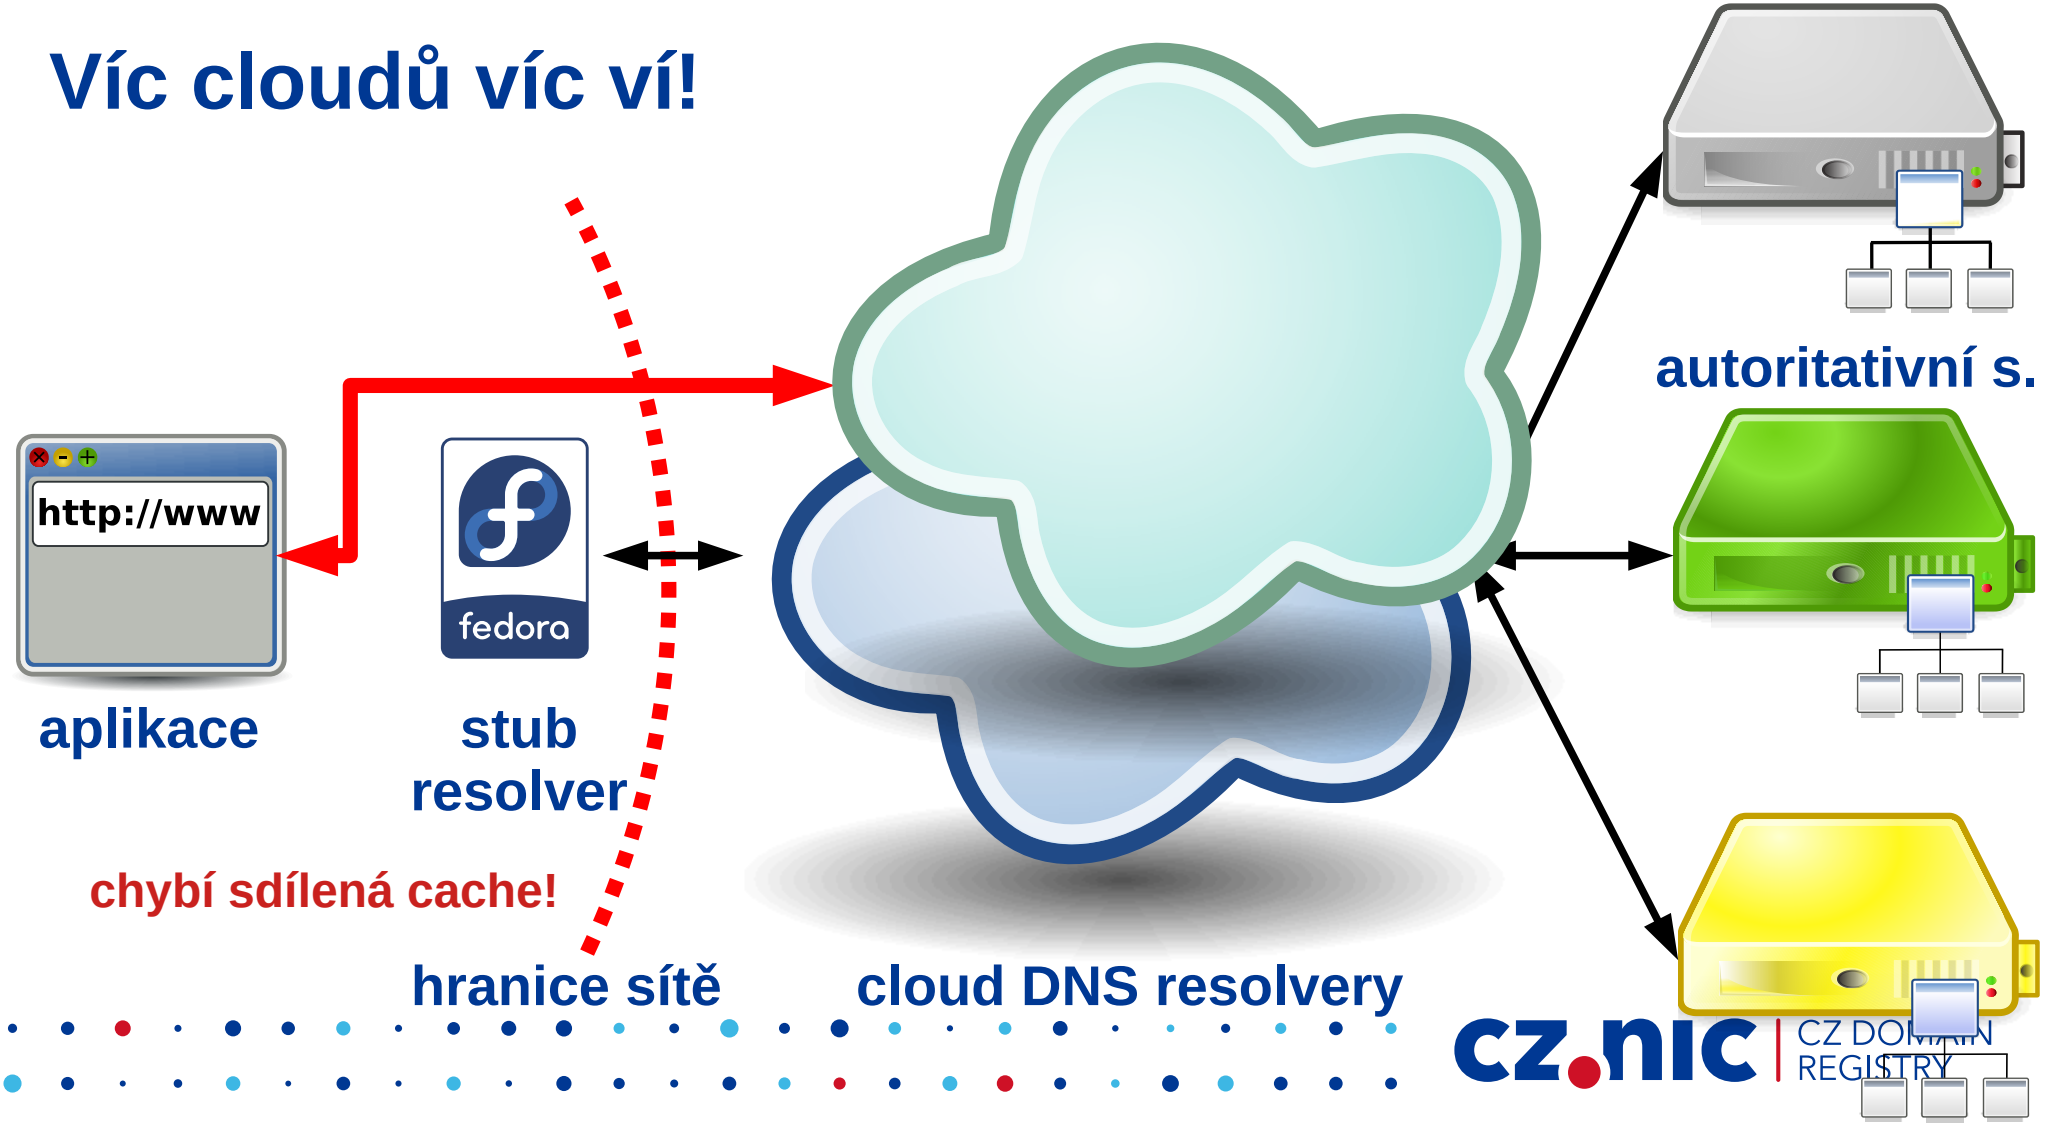

Soumrak tradiční DNS architektury

- 100 000 resolverů/denně
  - > 1000 dotazů na CZ domény
- Kolik je/bude velkých cloudů?
- DNS v Cloudu ⇒ **centralizace**
- Centralizace ⇒ lákavější cíle

Výkon a náklady

Latence DNS resolvingu

1. Cache hit rate
2. Latence mezi klientem a resolverem
3. Server selection algoritmus (v resolveru)
4. **Komunikační protokol**
   - **pozor na navazování spojení!**

významnost

Náklady na provoz DNS resolveru

1. Cache hit rate
2. **Komunikační protokol**
   - **navazování spojení je drahé!**
3. Efektivita implementace
4. Vše ostatní

významnost

Co je rychlejší? To záleží ...

- Cloud × **malý** lokální resolver
- Cloud × **velký** lokální resolver
  - Přednaplnění cache …
- Rychlá × pomalá linka
  - latence
  - propustnost

Soukromí

Srovnání protokolů

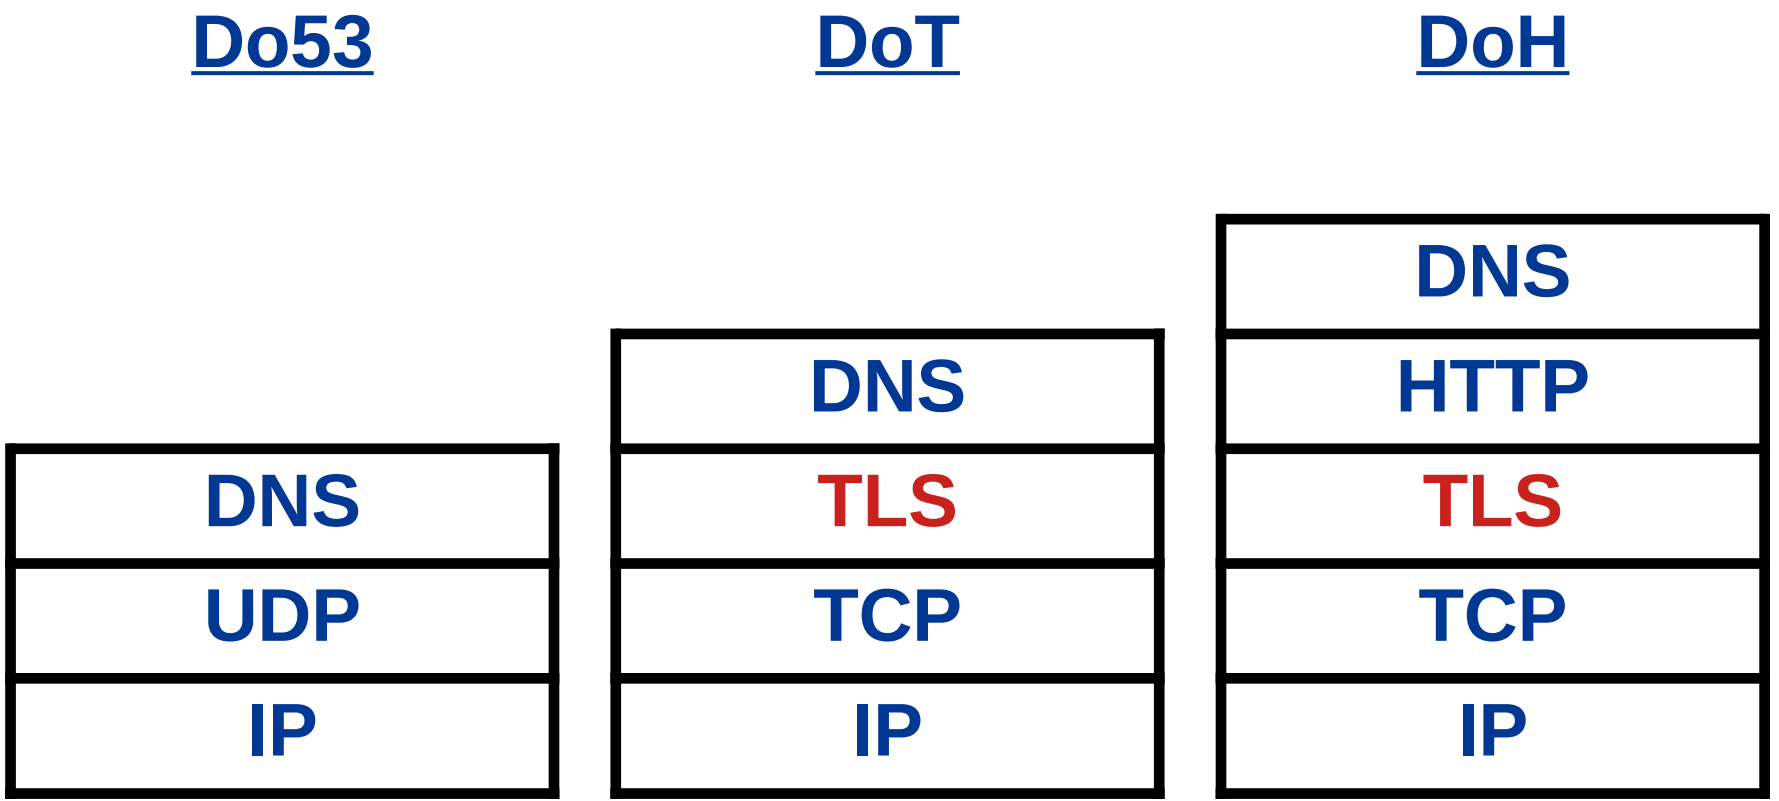

Načtení webu

- Jedna web stránka
  - ⇒ **mnoho** TCP spojení
- ihned.cz – 317 HTTP requestů
- 30 IPv4 adres
- 31 IPv6 adres

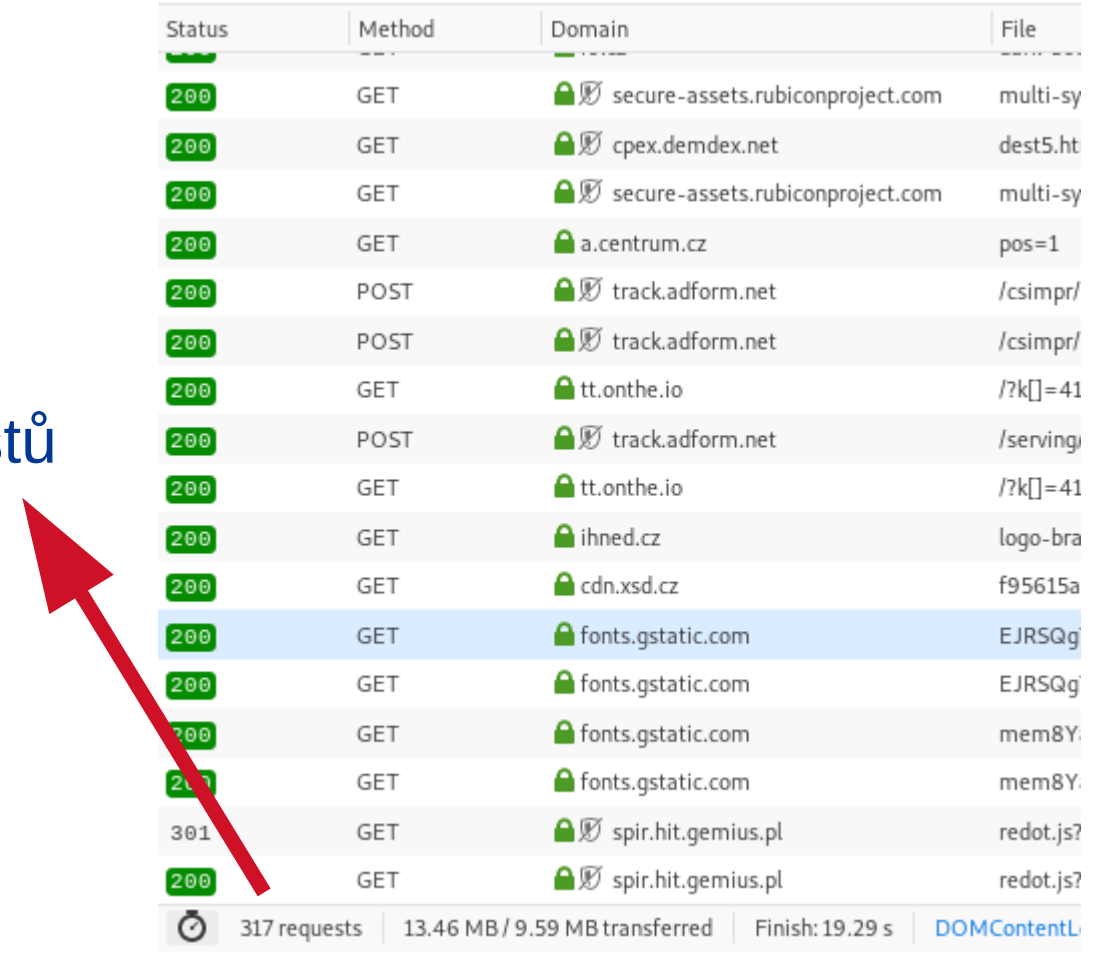

Status	Method	Domain	File
200	GET	secure-assets.rubiconproject.com	multi-sy
200	GET	cpex.demdex.net	dest5.ht
200	GET	secure-assets.rubiconproject.com	multi-sy
200	GET	a.centrum.cz	pos=1
200	POST	track.adform.net	/csimpr/
200	POST	track.adform.net	/csimpr/
200	GET	tt.onthe.io	/?k[]=41
200	POST	track.adform.net	/serving
200	GET	tt.onthe.io	/?k[]=41
200	GET	ihned.cz	logo-bra
200	GET	cdn.xsd.cz	f95615a
200	GET	fonts.gstatic.com	EJRSQg
200	GET	fonts.gstatic.com	EJRSQg
200	GET	fonts.gstatic.com	mem8Y
200	GET	fonts.gstatic.com	mem8Y
301	GET	spir.hit.gemius.pl	redot.js?
200	GET	spir.hit.gemius.pl	redot.js?

317 requests | 13.46 MB / 9.59 MB transferred | Finish: 19.29 s | DOMContentL

Získání web objektu

1. DNS dotaz **www.example.com.** AAAA?
2. DNS odpověď **www.example.com.** AAAA 2001:db8::1
3. TCP spojení 2001:db8::1
4. TLS Server Name Indication = **www.example.com**
5. HTTPS Host: **www.example.com** GET /
6. HTTPS Host: www.google-analytics.com
   Referer: **www.example.com**, Cookie: clientid

Soukromí: Kdo se dívá "doma"?

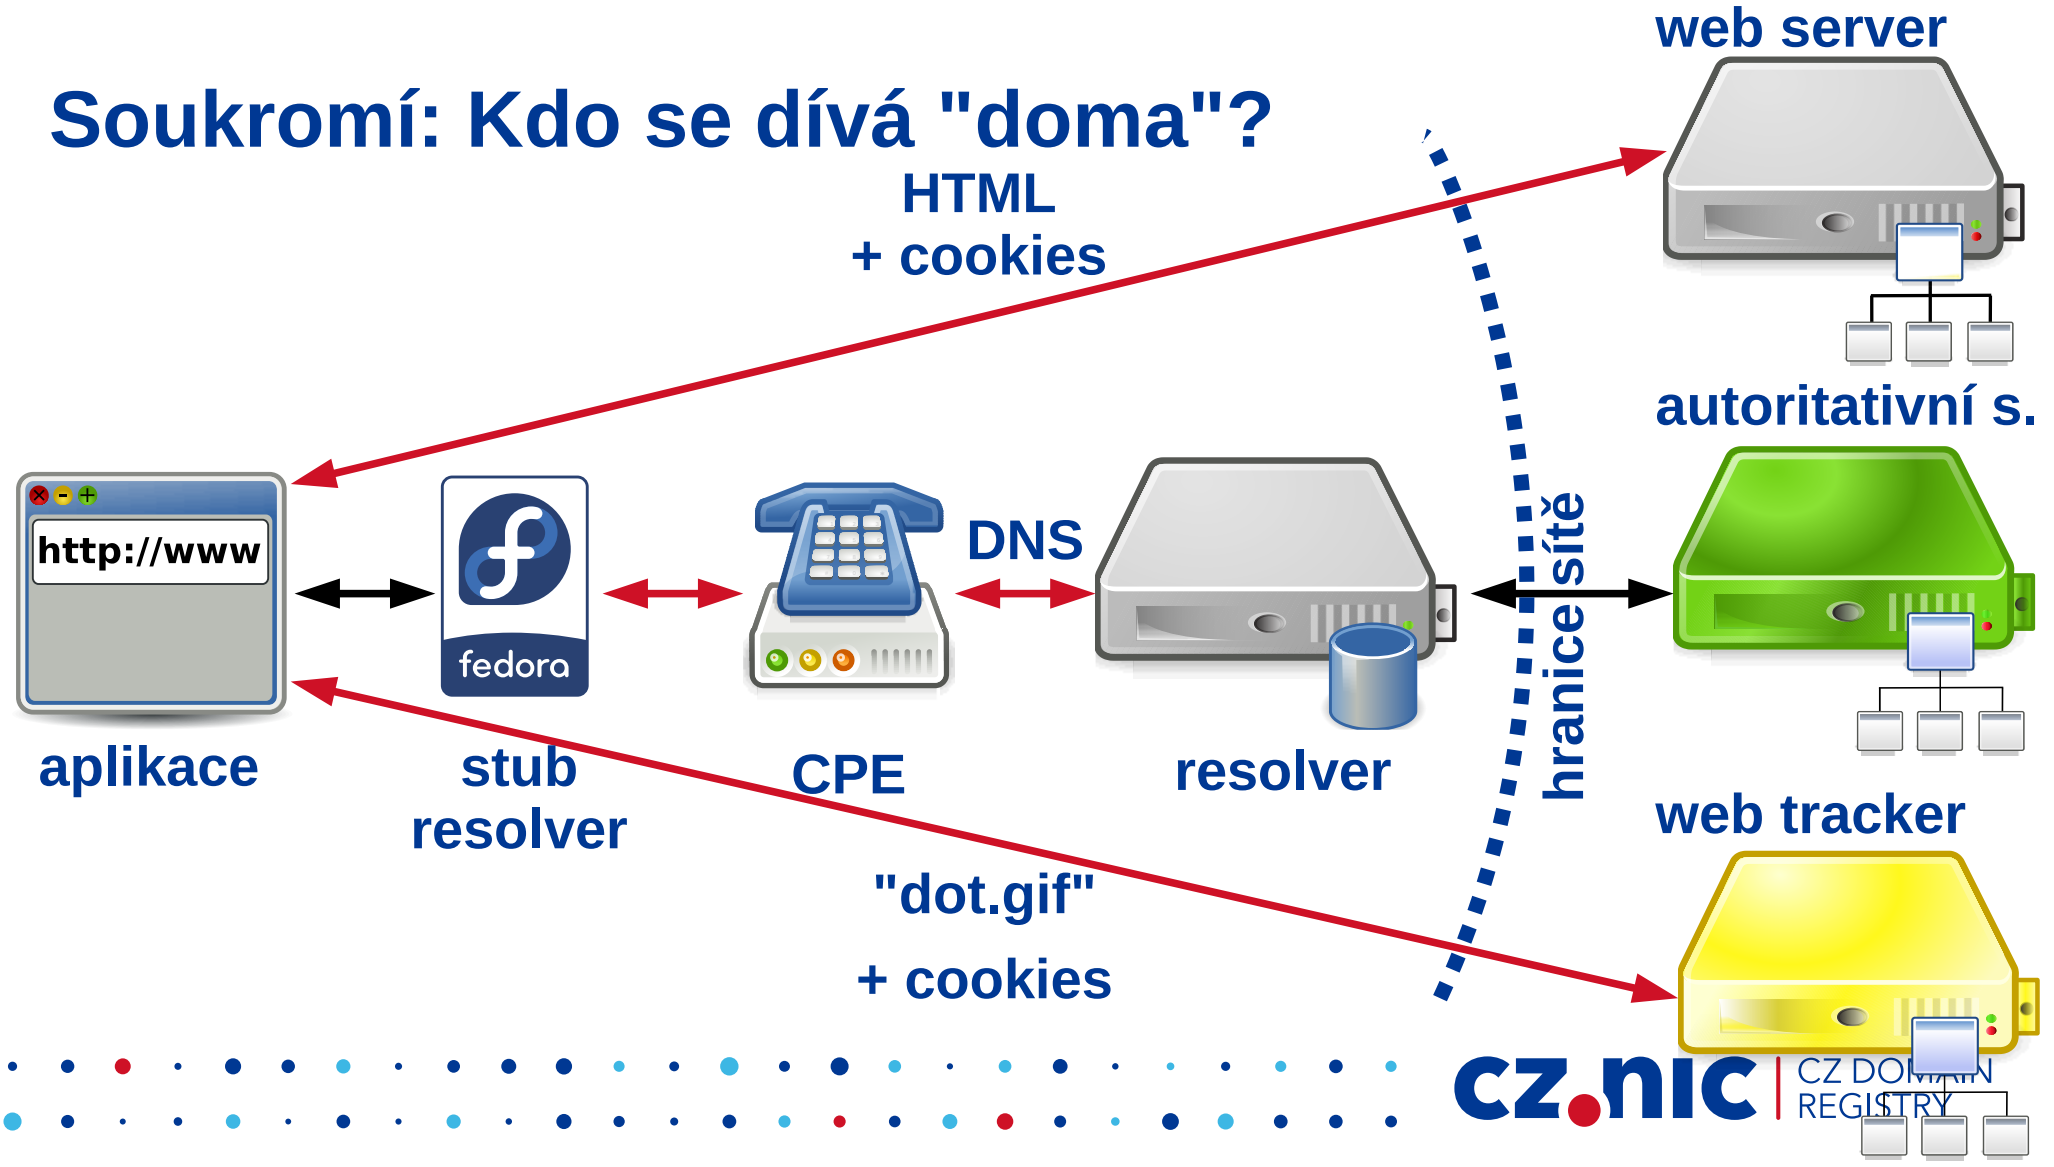

Šifrování: TLS

1. DNS dotaz **www.example.com.** AAAA?
2. DNS odpověď **www.example.com.** AAAA 2001:db8::1
3. TCP spojení 2001:db8::1
4. TLS Server Name Indication = **www.example.com**
5. ~~HTTPS~~ ~~Host: www.example.com GET /~~
6. HTTPS Host: www.google-analytics.com
   Referer: **www.example.com**, Cookie: clientid

Šifrování: TLS + Encrypted SNI

1. DNS dotaz **www.example.com.** AAAA?
2. DNS odpověď **www.example.com.** AAAA 2001:db8::1
3. TCP spojení 2001:db8::1
4. ~~TLS~~ ~~Server Name Indication = www.example.com~~
5. ~~HTTPS~~ ~~Host: www.example.com GET /~~
6. HTTPS Host: www.google-analytics.com
   Referer: **www.example.com**, Cookie: clientid

Soukromí: Kdo se dívá z Cloudu?

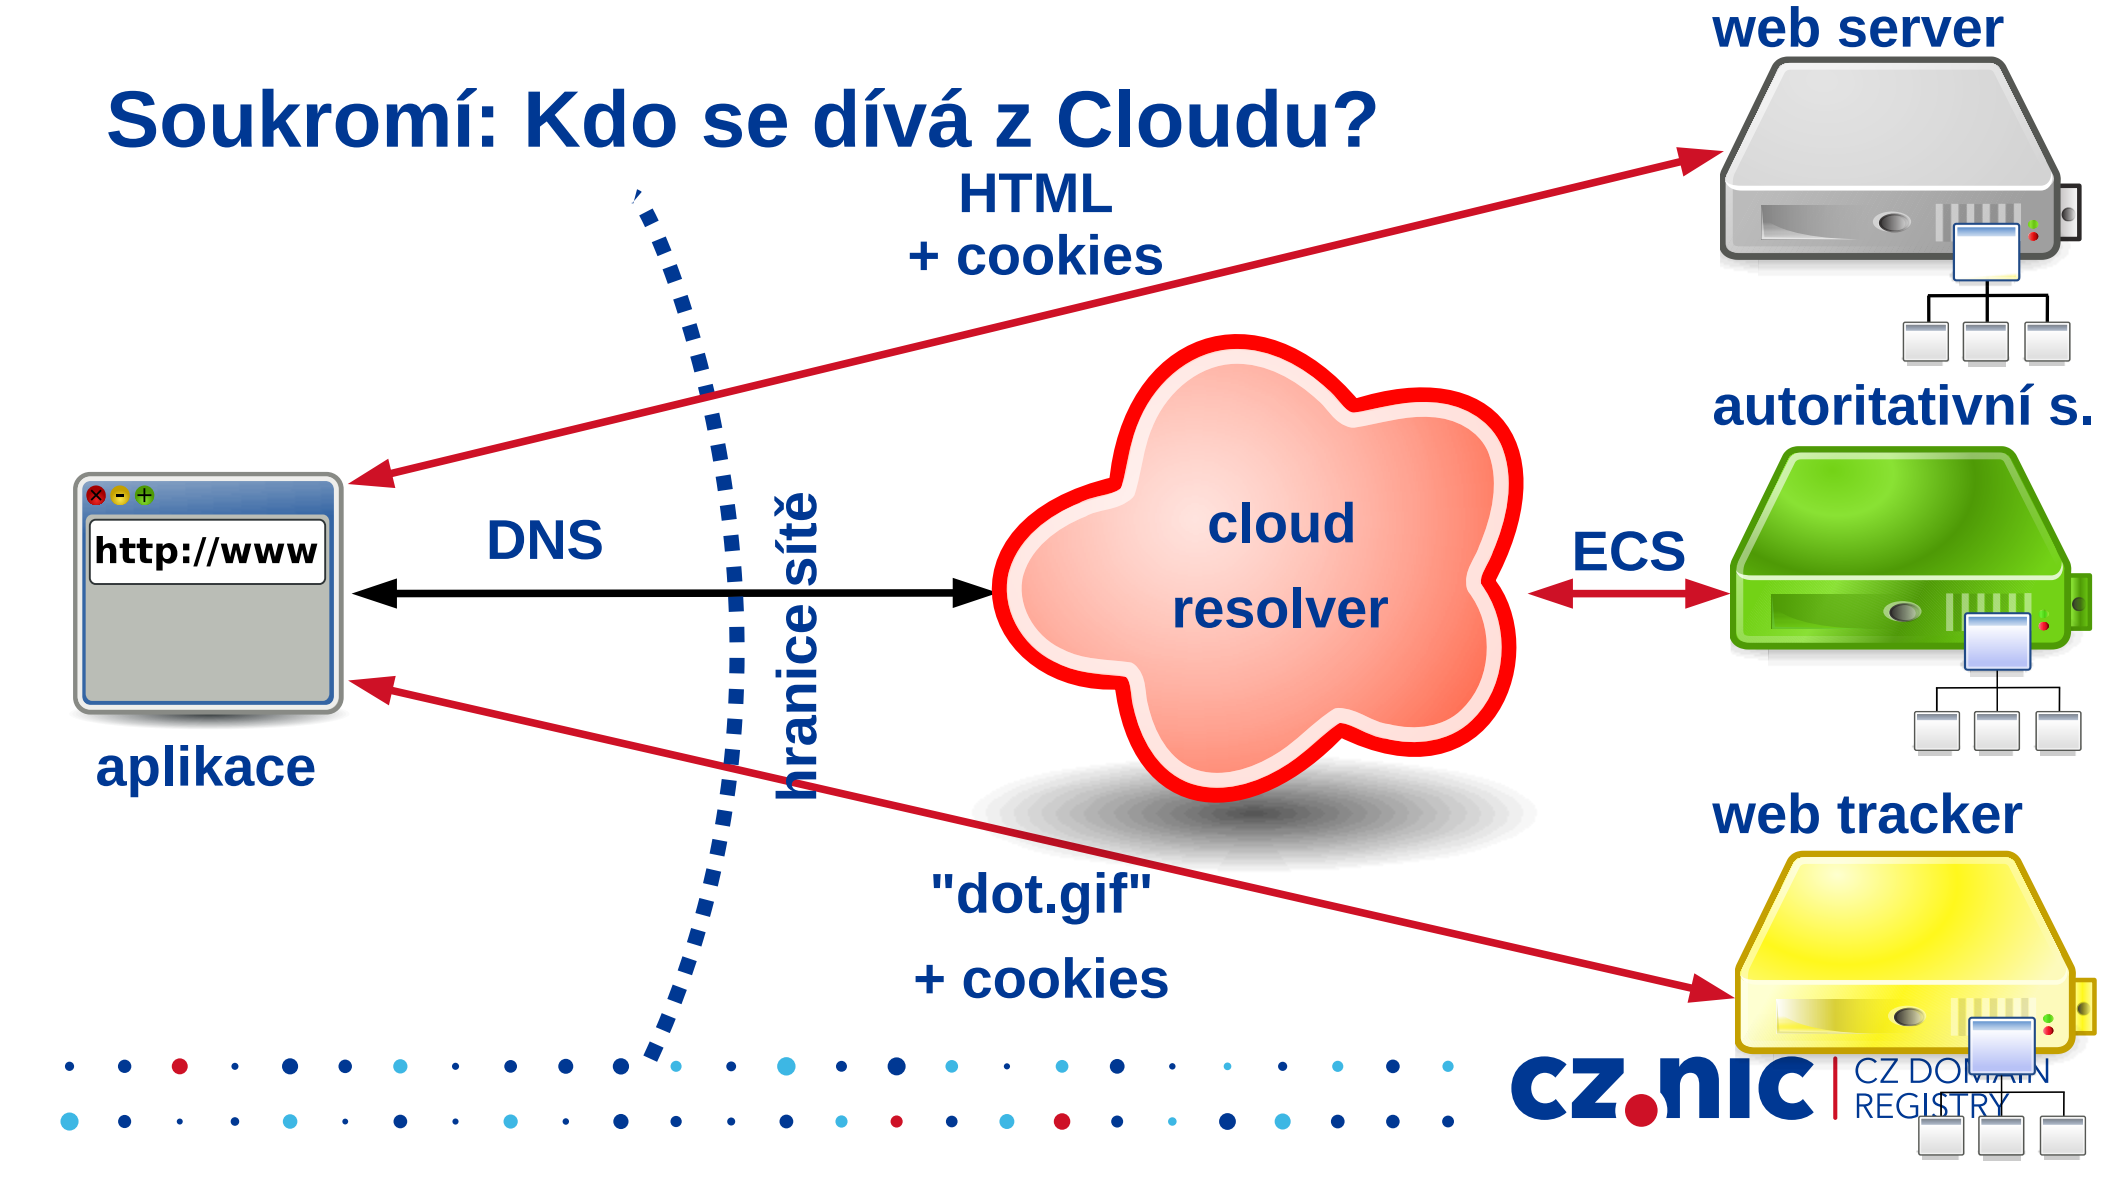

Šifrování: TLS + Encrypted SNI + DNS

1. ~~DNS dotaz~~ ~~www.example.com. AAAA?~~
2. ~~DNS odpověď~~ ~~www.example.com. AAAA~~ 2001:db8::1
3. TCP spojení 2001:db8::1
4. ~~TLS~~ ~~Server Name Indication = www.example.com~~
5. ~~HTTPS~~ ~~Host: www.example.com GET /~~
6. HTTPS Host: www.google-analytics.com
   Referer: **www.example.com**, Cookie: clientid

Šifrování: TLS + Encrypted SNI + DNS

- Třetí strana vidí TCP spojení
  - IP adresa cíle 2001:db8::1
- Web trackery **vidí vše beze změn**

Co se dá vyčíst z IP adresy?

What Can You Learn from an IP?

Simran Patil
University of Illinois at Urbana-Champaign
sppatil2@illinois.edu

Nikita Borisov
University of Illinois at Urbana-Champaign
nikita@illinois.edu

ABSTRACT

The Internet was not designed with security in mind. A number of recent protocols such as Encrypted DNS, HTTPS, etc. target encrypting critical parts of the web architecture, which can otherwise be exploited by eavesdroppers to infer users' data. But encryption may not necessarily guarantee privacy, especially when it comes to metadata. Emerging standards can protect the contents of both DNS queries and the TLS SNI extensions; however, it might still be possible to determine which websites users are visiting by simply looking at the destination IP addresses on the traffic originat-

protocol. Most widely deployed are DNS-over-TLS and DNS-over-HTTPS [1, 6], which protect the contents of both DNS queries and responses from eavesdroppers by relying on existing deployed end-to-end encryption protocols.[1] In the context of web browsing, however, a DNS request is followed by an HTTP or HTTPS connection to a web server located by the request. In the case of HTTP, the entire request is sent in plaintext. With HTTPS, TLS protects the majority of the communication section, yet some data is sent in the clear. Most importantly, the server name indication (SNI) extension [7] specifies the domain name of the web server,

Co se dá vyčíst z IP adresy?

4 CONCLUSION

Our measurements show that, in the context of web browsing, DNS and SNI privacy offers limited protection against an adversary who knows a plausible set of sites a user might visit (even if the set is quite large), and who performs forward lookups to infer the domain names and sites associated with given IPs. Using a crawl of Alexa top 1 million sites, we find that nearly half of all IPs involved in the crawl correspond to a unique domain name, and over 95% of sites have a unique set of IPs corresponding to the domains of all the sub-resources. We do identify a significant opportunity for content distribution networks (CDNs) to offer additional protection by coalescing more domains onto the same IP address.

Co se dá vyčíst z IP adresy?

- Simran Patil and Nikita Borisov. 2019.
  What can you learn from an IP?.
  In Proceedings of the Applied Networking Research Workshop (ANRW '19).
  ACM, New York, NY, USA, 45-51.
  DOI: https://doi.org/10.1145/3340301.3341133
- https://irtf.org/anrw/2019/program.html
- Jednoduchá statistika!

Soukromí: Všichni do VPN!

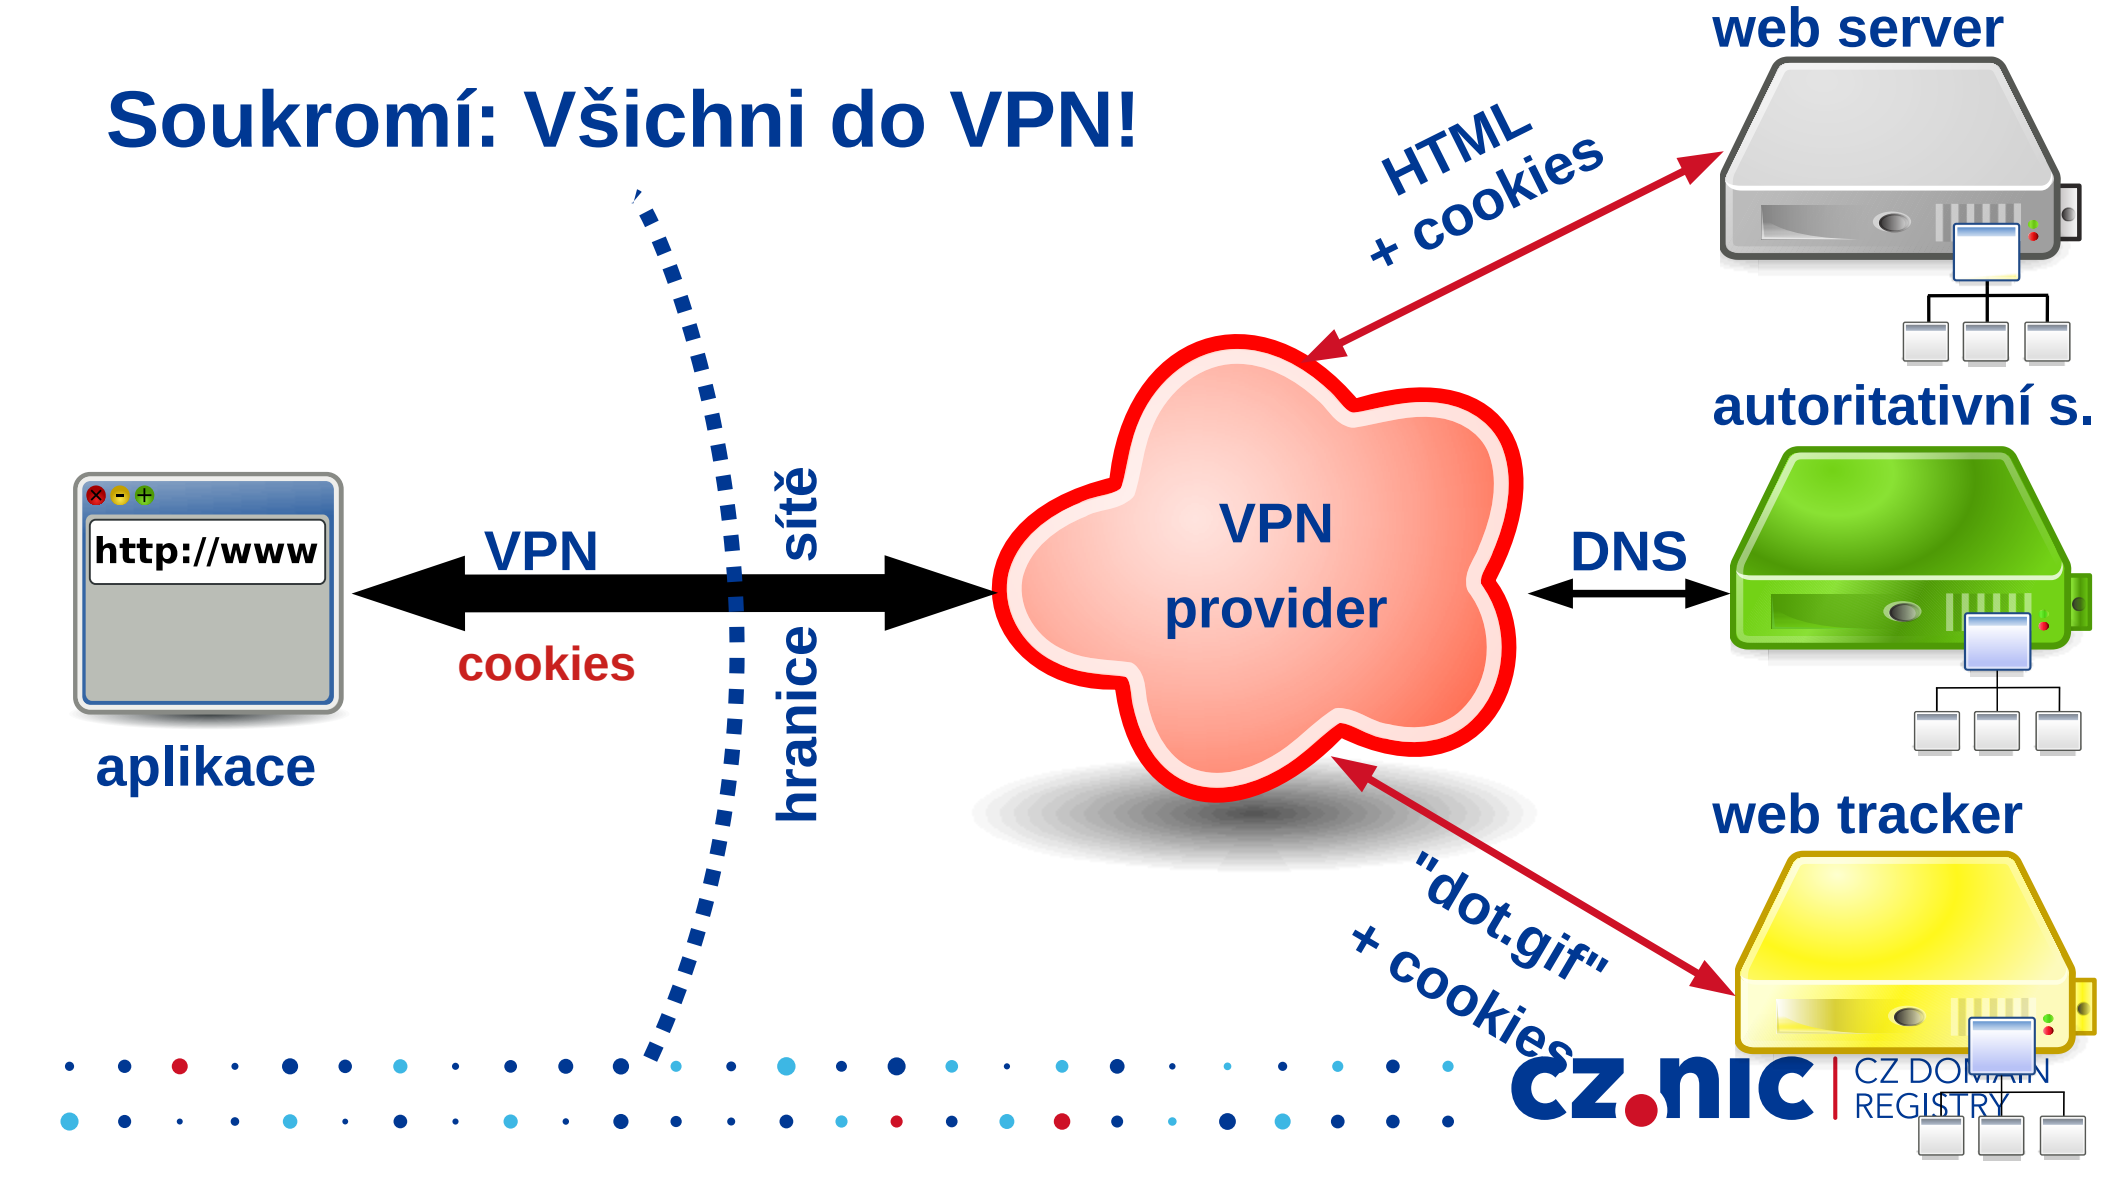

Shrnutí

- Šifrované DNS
  - **Typicky nezlepšuje** soukromí (falešný pocit bezpečí)
- Šifrované DNS v cloudu
  - **Zhoršuje** soukromí – zvětšuje počet citlivých míst
- Plná VPN
  - **Přesun důvěry** z ISP na VPN providera
- Šifrování DNS neřeší HTTP trackery