417

- DNS - admin-

originally

all names maintained in hosts.txt at SRI.

Problems?

- does not scale wrt # names / hosts

- spof - linear cost of retr.

- consistency? - censor

DNS : Domain Name System

RFC 1034, 1035

Design Goals

- General Purpose
  consistent namespace

- Distributed maintainance
  via Delegation

- Server of data controls
  tradeoff between
  cost / accuracy

DNS Namespace

- variable depth rooted tree
- each node has a label (name)

Resource Records

data associated with names.

Name Servers

- info. repositories

Resolvers

- extract info. frm NS in response to client queries

DNS from the perspective of...

user : DNS accessed via resolvers

Resolver : DNS is composed of an unkaown # of NSs.

Each NS stores some part of the namespace.

NS: DNS consists of partitioned info. Each partition is called a zone.

DNS namespace

- variable depth rooted tree

- each node has a corresponding resource set (≥ 0 RRs)

- each node has a label (0 - 63 bytes)

- sibling nodes cannot have the same label

- root label : $\emptyset$

DNS Namespace

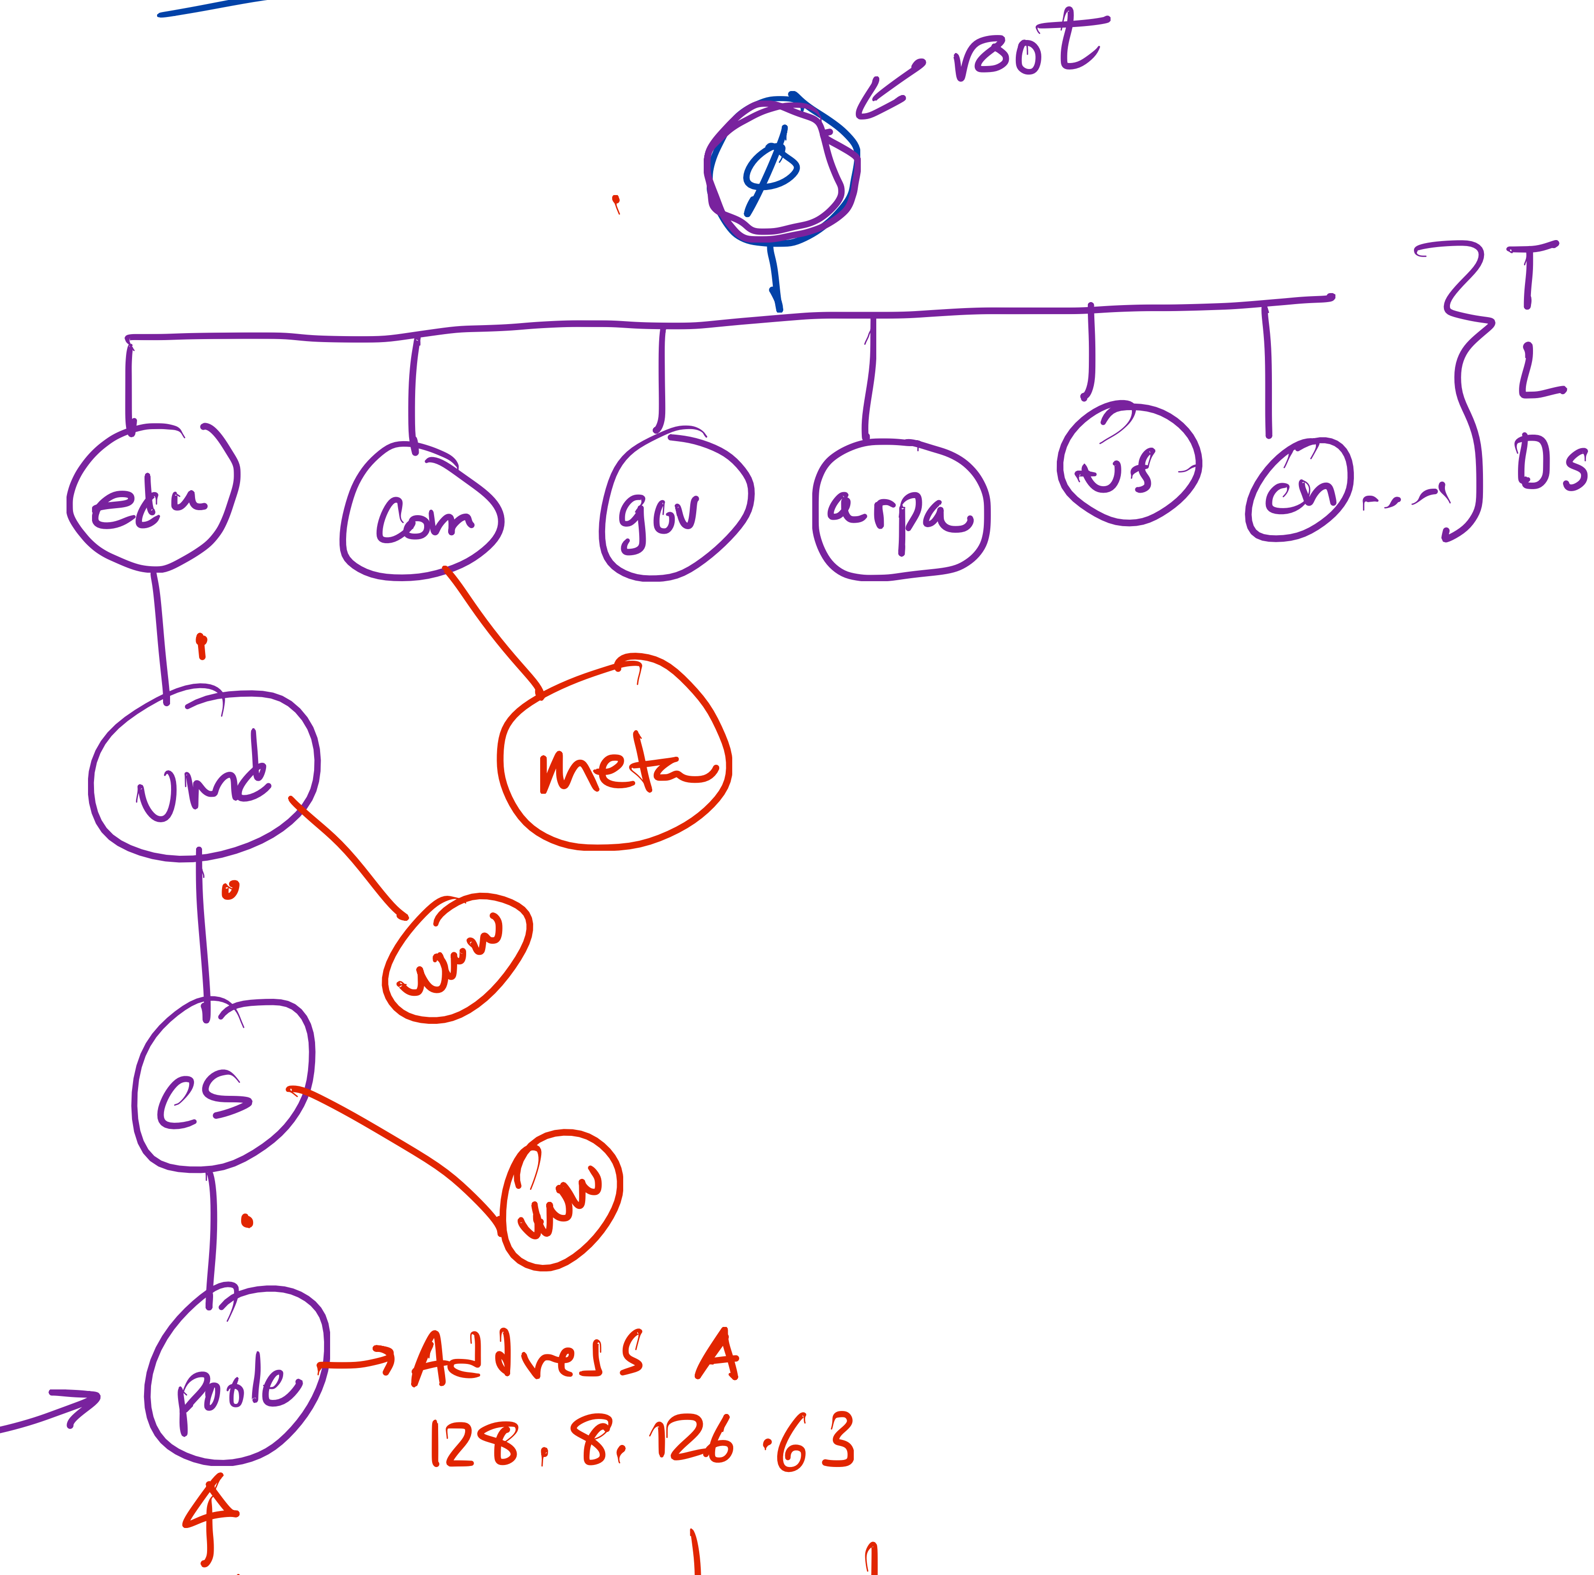

poole.cs.umd.edu.

Fully Qualified Domain Name (FQDN)

FQDN

Resource Records

{type, class, data}

IN, CHAOS, HESIOD, ...

abstract resources:

hostname, mailbox, nameserv...

A AAAA MX MF MB

CNAME NS SOA

HINFO PTR TXT

Zones:

A complete description of all data in a contiguous section of the namespace that is administered by a single authority.

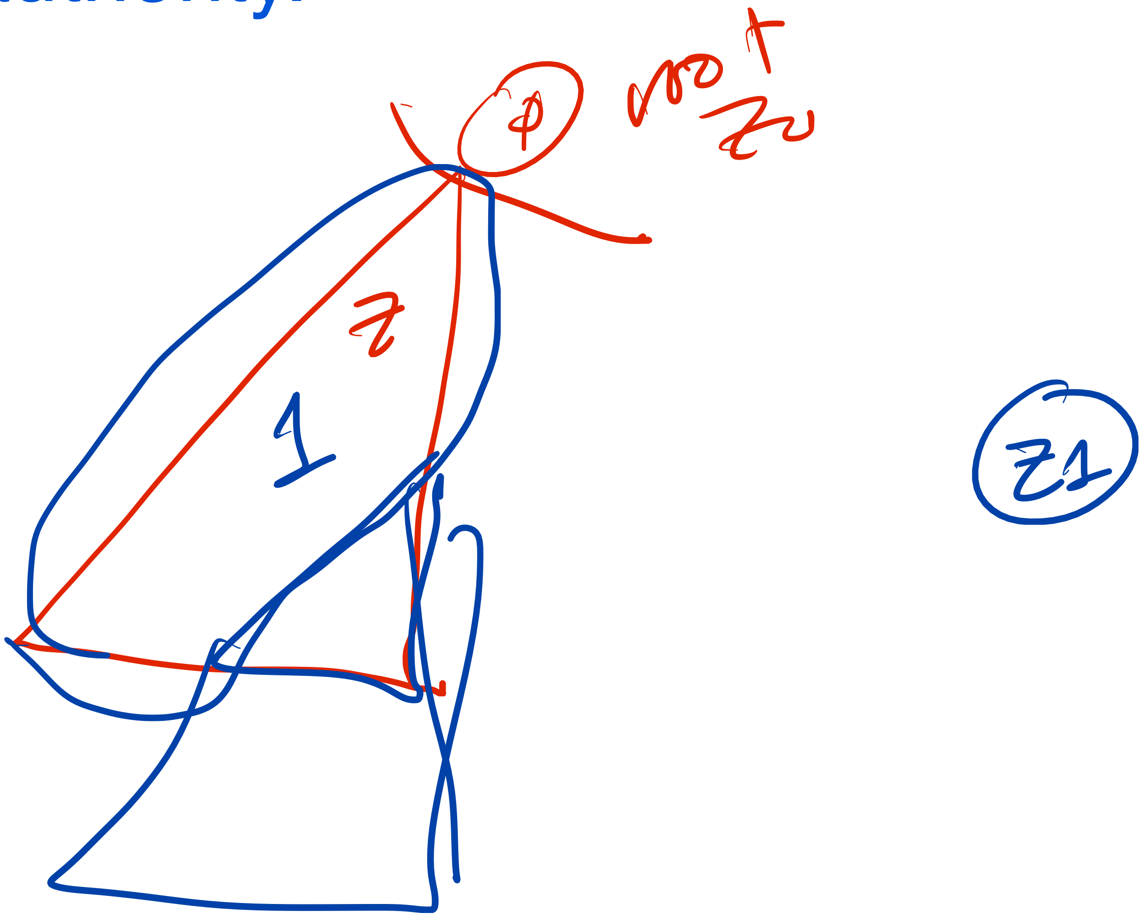

Zone example

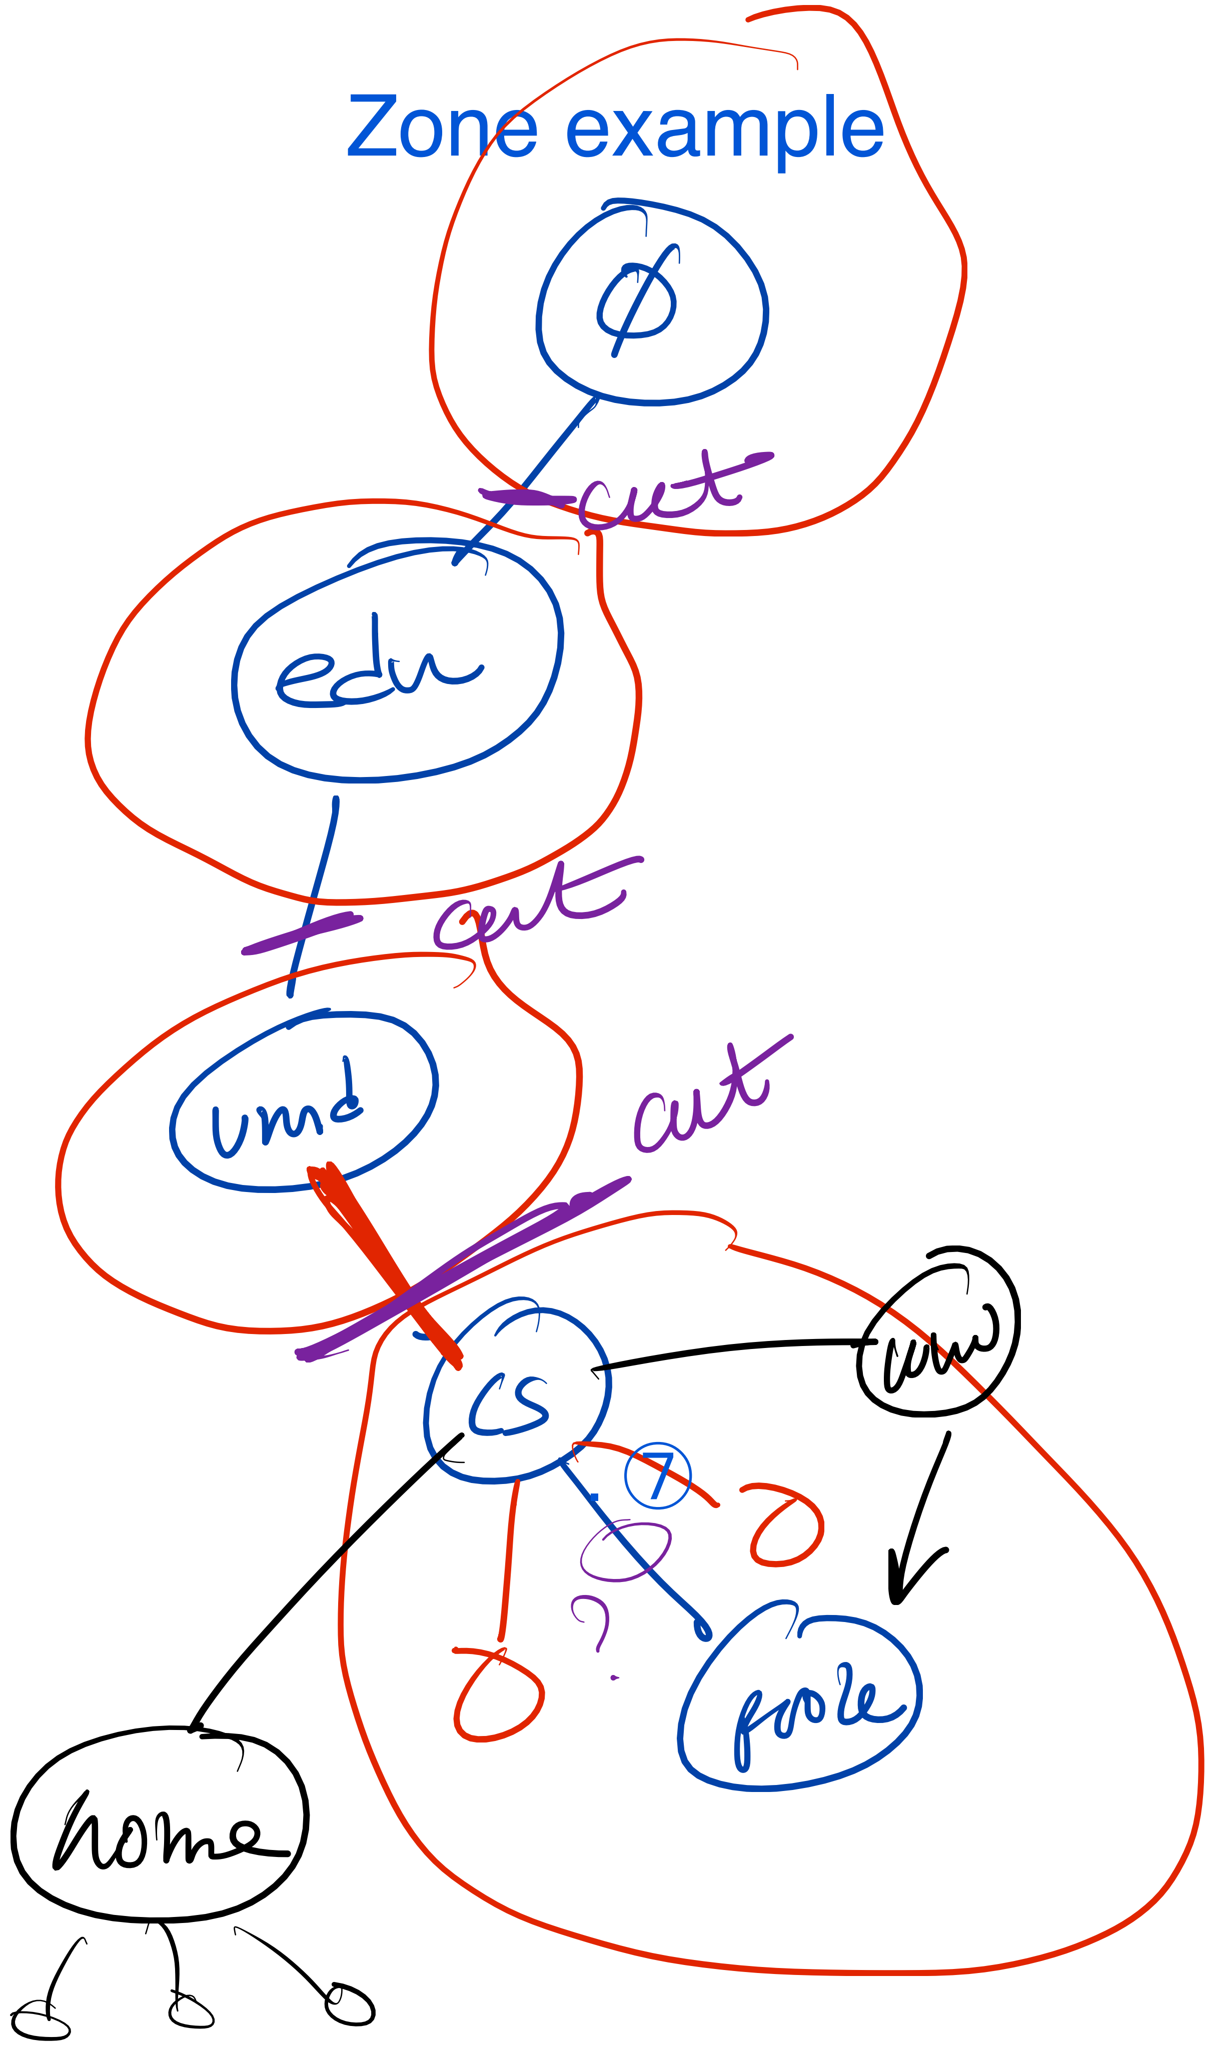

SoA record:

authority data that describes the top node in a zone.

cut: denotes zone boundaries
- may occur between any two nodes in the namespace.

Zones designate administrative boundaries.

- an organization gets control of a zone by persuading parent organization to _delegate_ a subzone consisting of a _single_ node.

Parent does this by inserting a single RR into its zone that designates a zone division

- new zone can grow / delegate independently.

Zone dB

umd.edu. NS ns1.umd.edu.

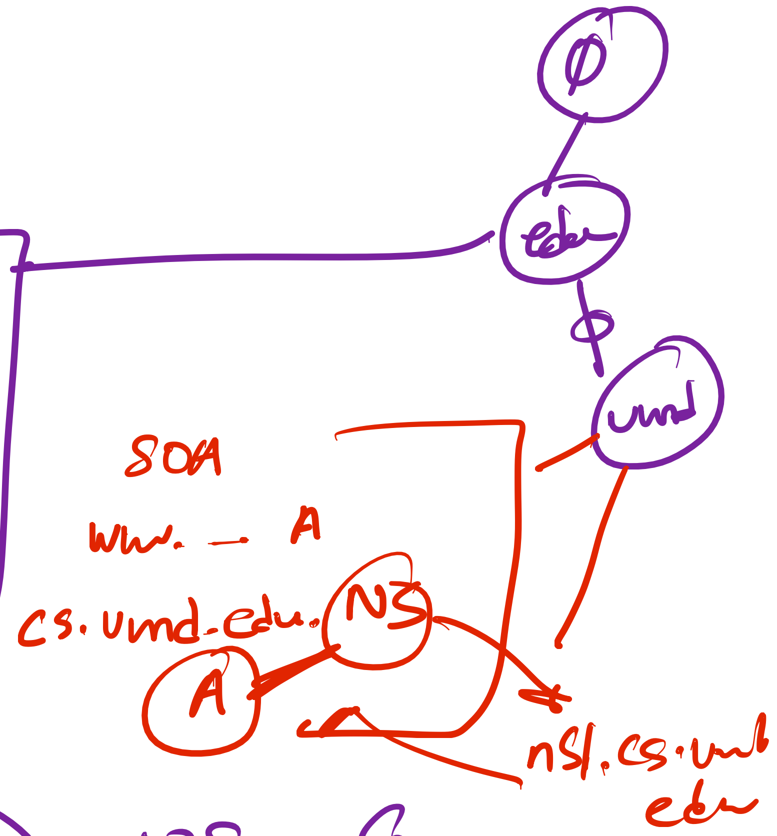

ns1.umd.edu. A 128. 8. ___

~~AAAA~~

Glue records

- used at zone boundaries

(NS, A) record pair for delegated namespace

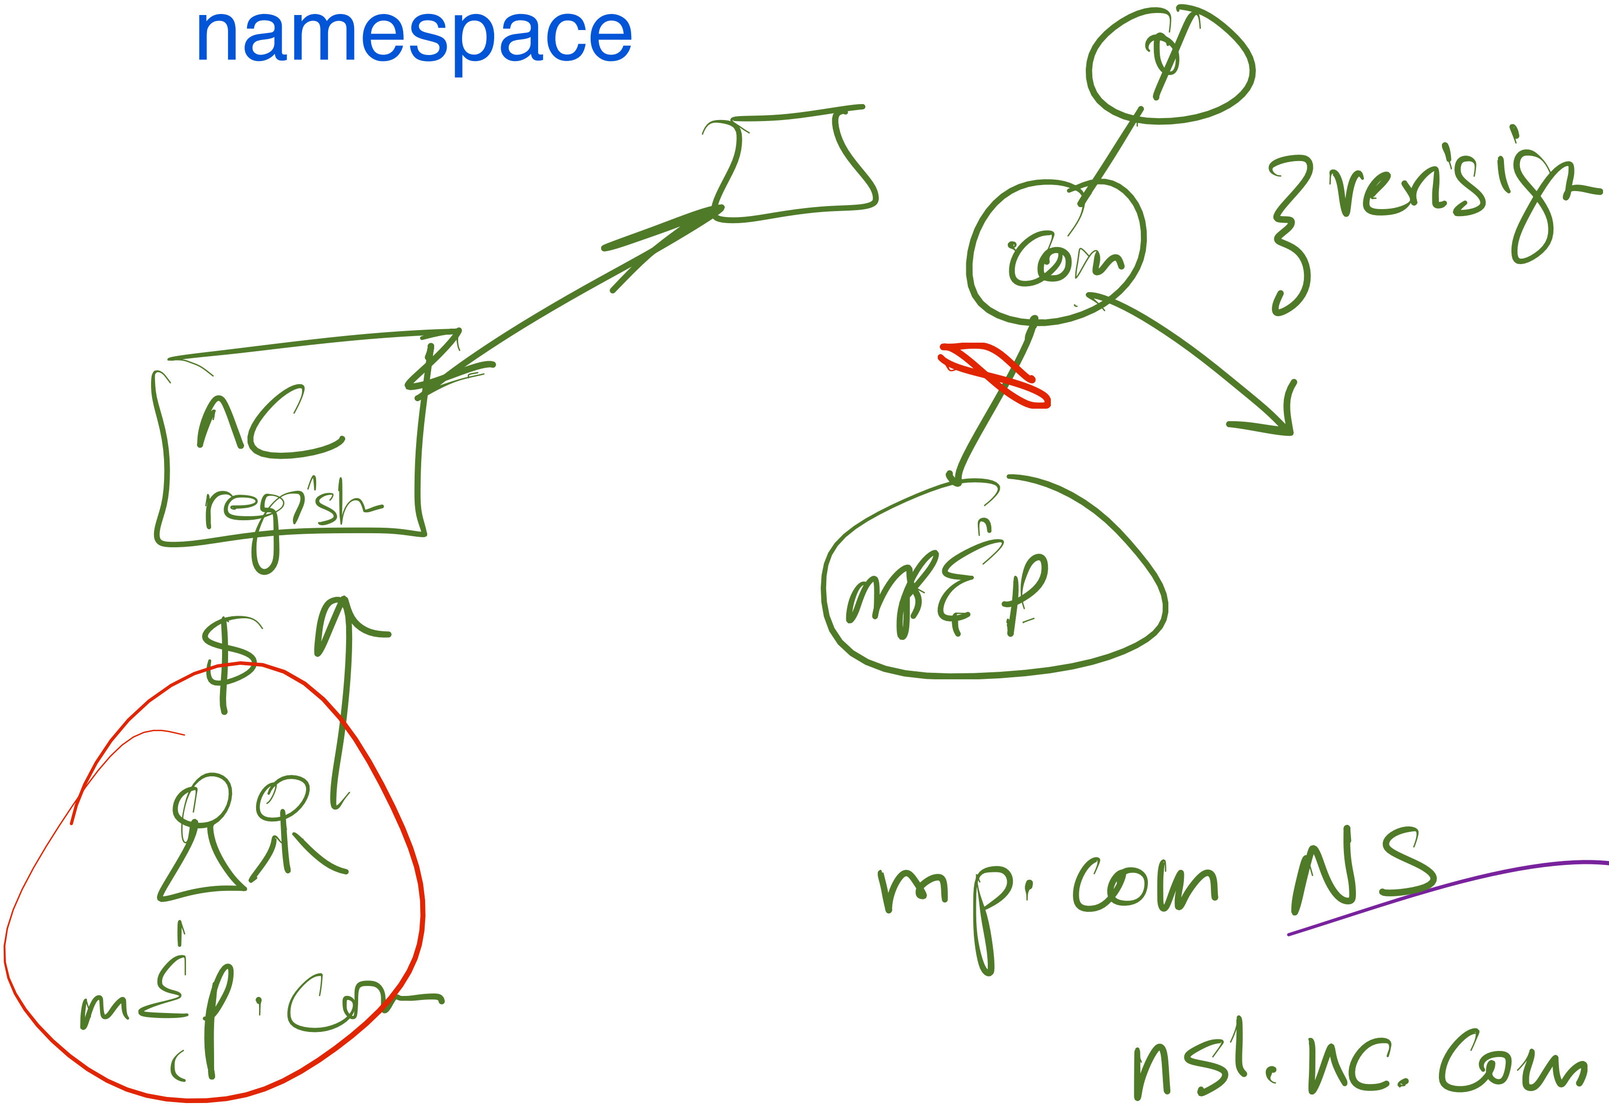

mp.com NS

nsl.nc.com

mp.com A ___

Each zone must have two power independent NSs serving zone data (one primary, one secondary).

A particular NS (hardware) can serve zone data for any number of zones.

Authoritative Answer:

an answer from a NS about its OWN zone

DNS wire protocol (port 53 UDP/TCP)

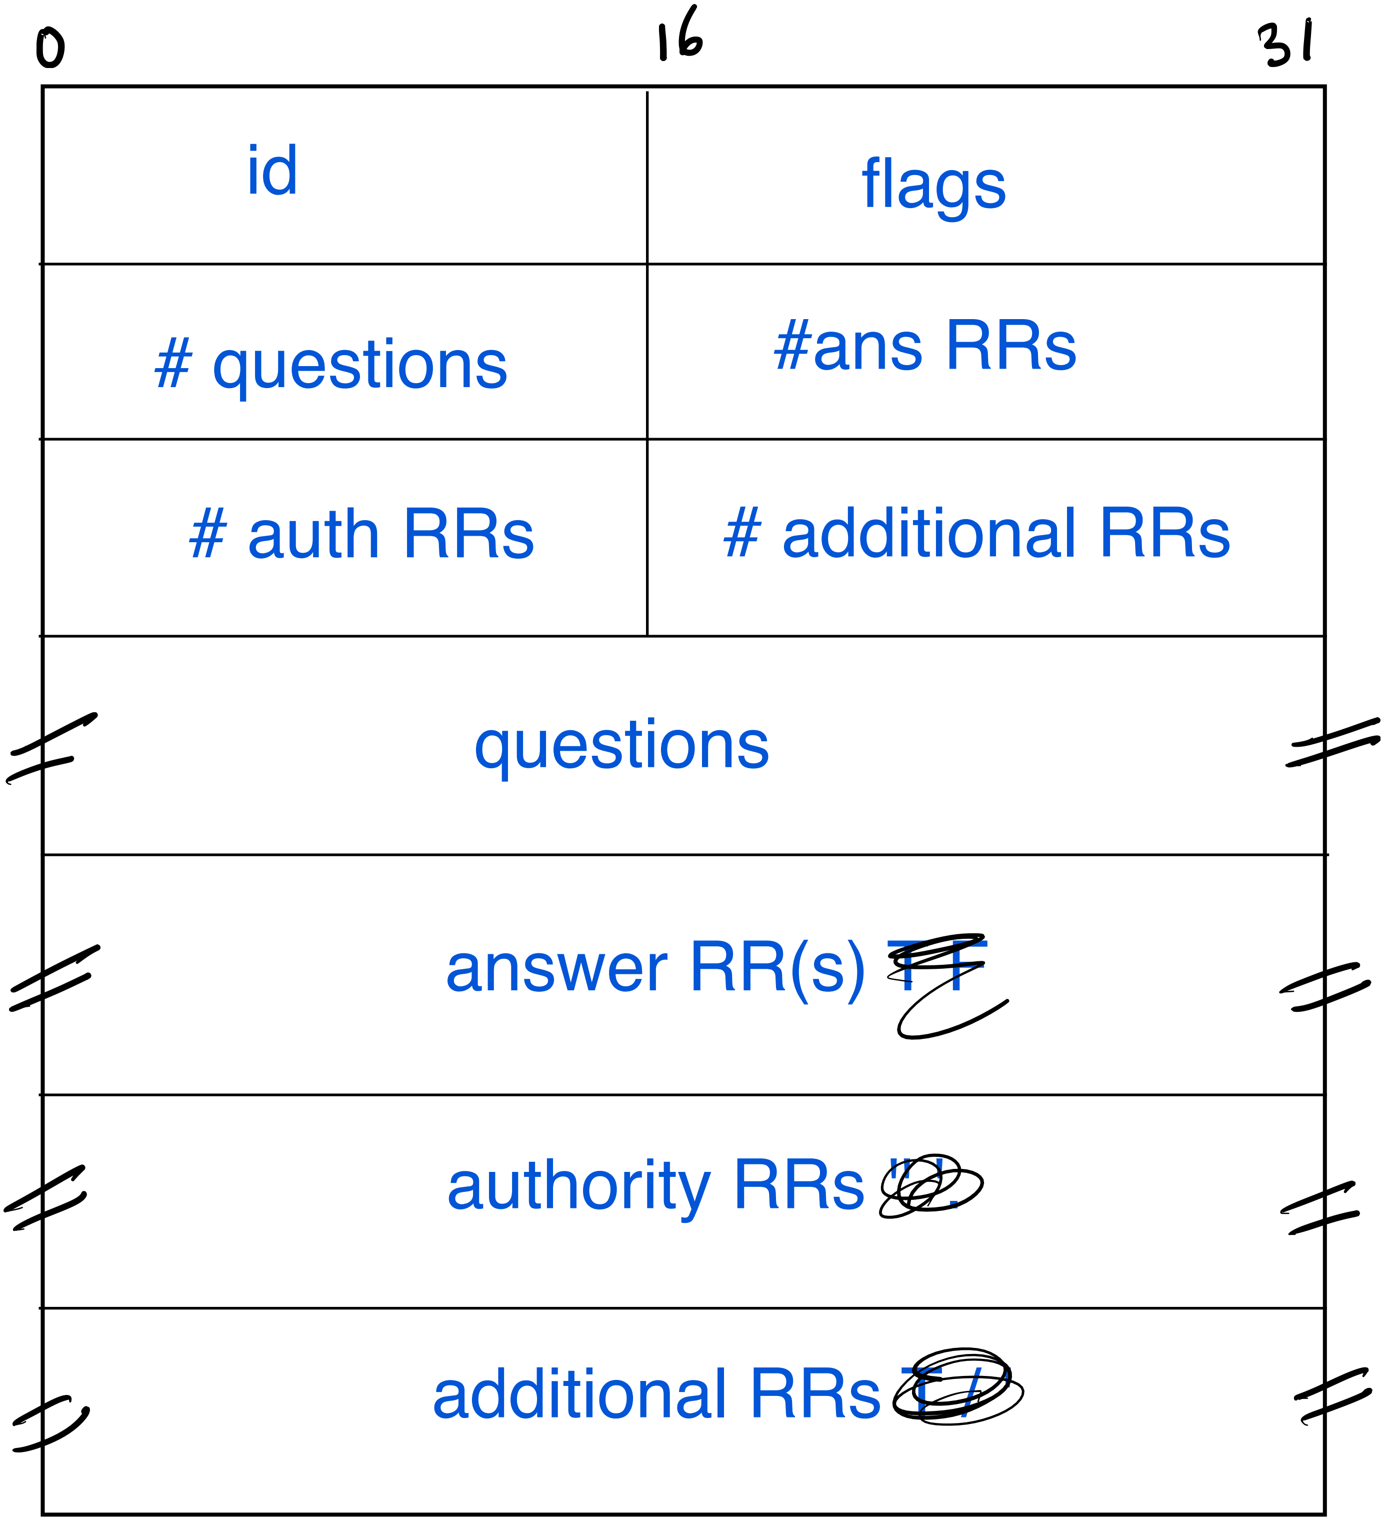

0 16 31

id | flags

questions | #ans RRs

auth RRs | # additional RRs

questions

answer RR(s) ~~TF~~

authority RRs

additional RRs

flags

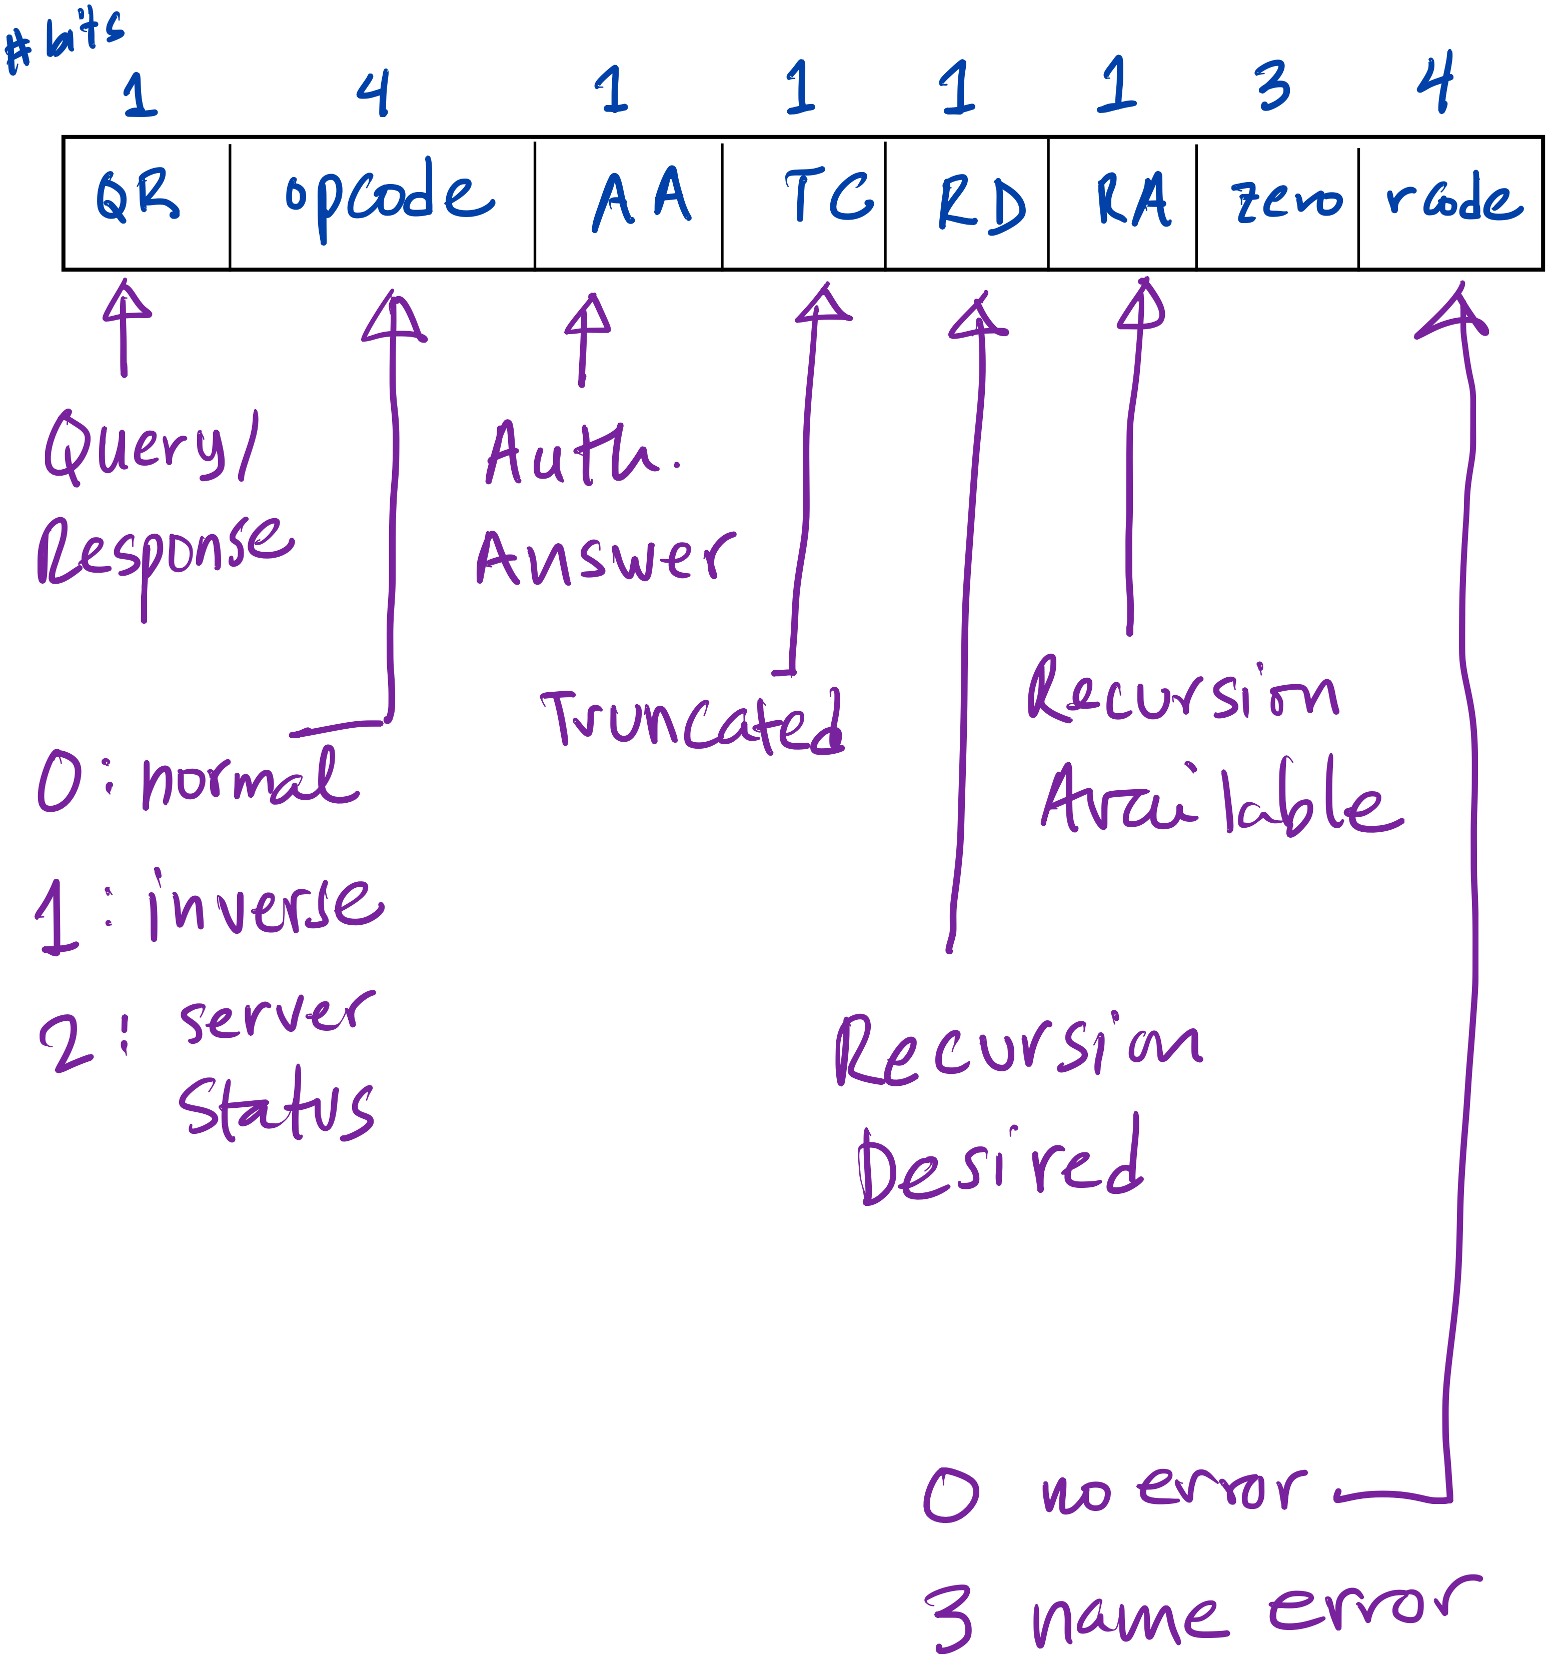

# bits	1	4	1	1	1	1	3	4
	QR	opcode	AA	TC	RD	RA	zero	rcode

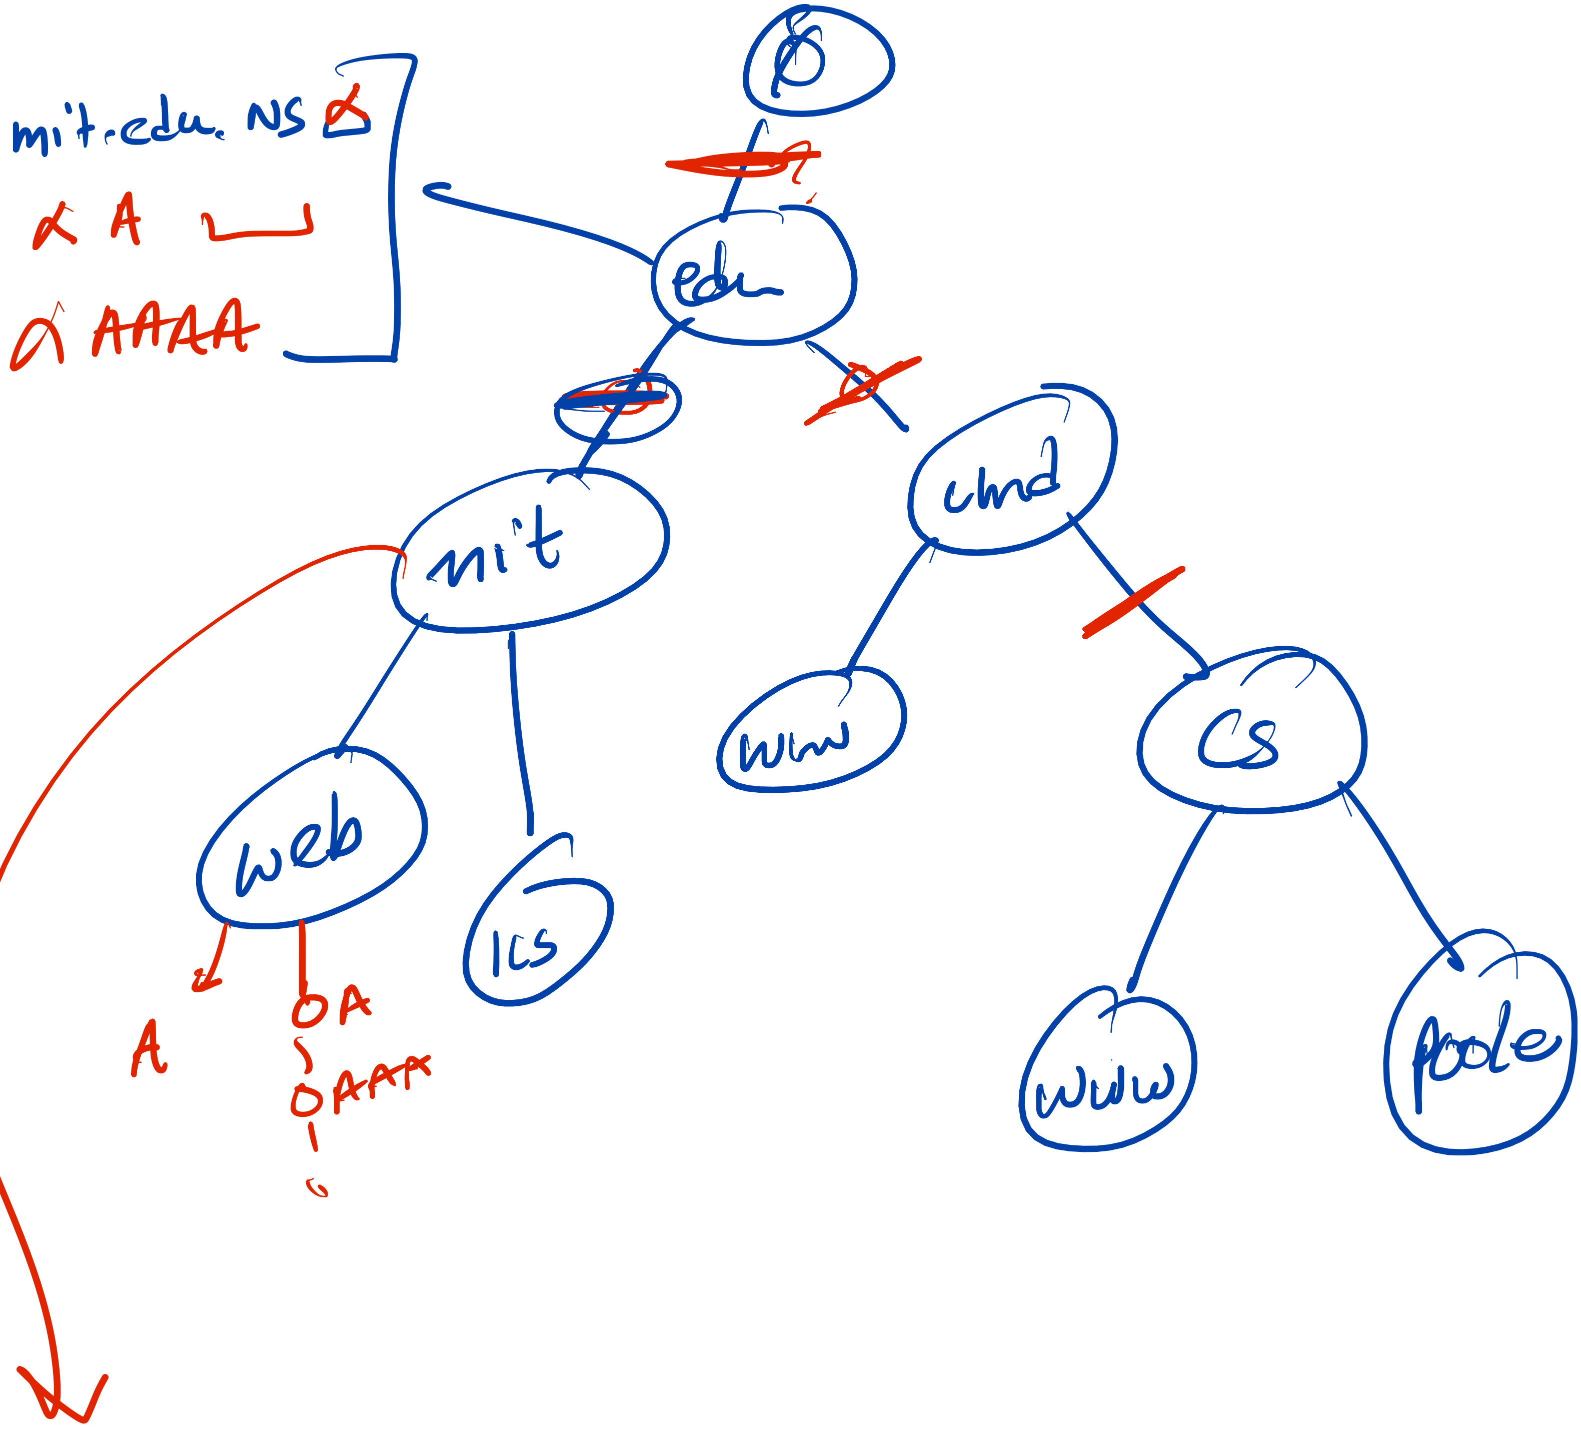

lcs.mit.edu. A ___

web.mit.edu A ___

web.mit.edu AAAA ___

RRs

Question

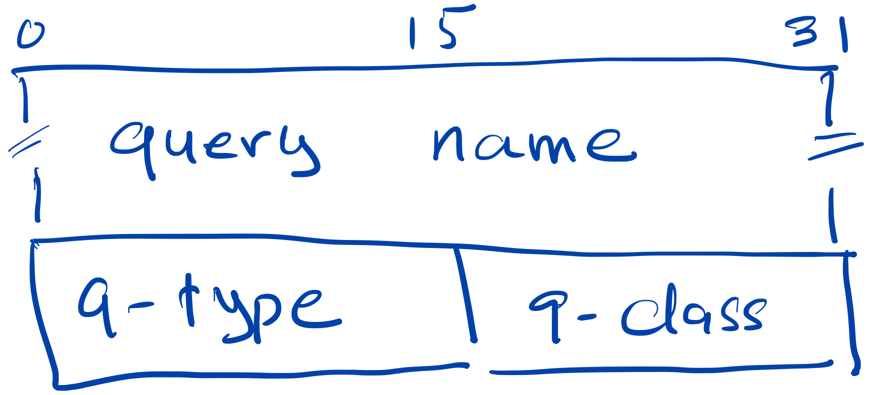

encoding

$\boxed{5}$ poole $\boxed{2}$ cs $\boxed{3}$ umd $\boxed{3}$ edu $\boxed{\emptyset}$

↑
FQDN

q-type:

A NS CNAME PTR MX

...

$\underbrace{\text{AXFR} \quad \text{ANY}}_{\text{not RRs}}$

RR formatting

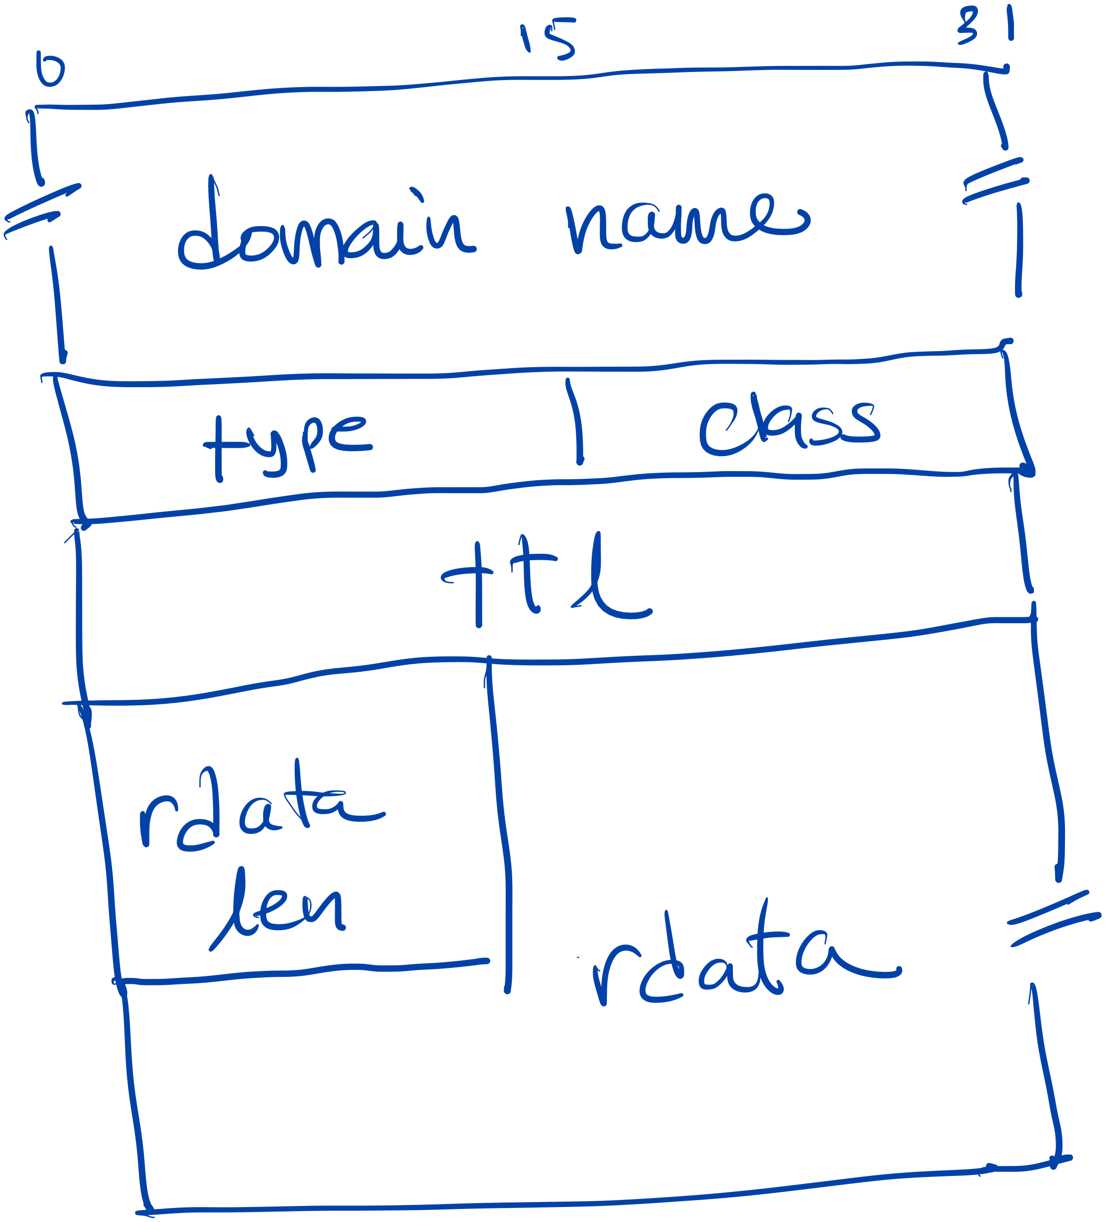

Compression

Count byte $\boxed{5}$ poole

legal range: 0 - 63

if 2 MSB of "count" are 1 : count is not a count at all, but a pointer

1	1	offset

14 bits

id field : offset = 0.

Compression Example

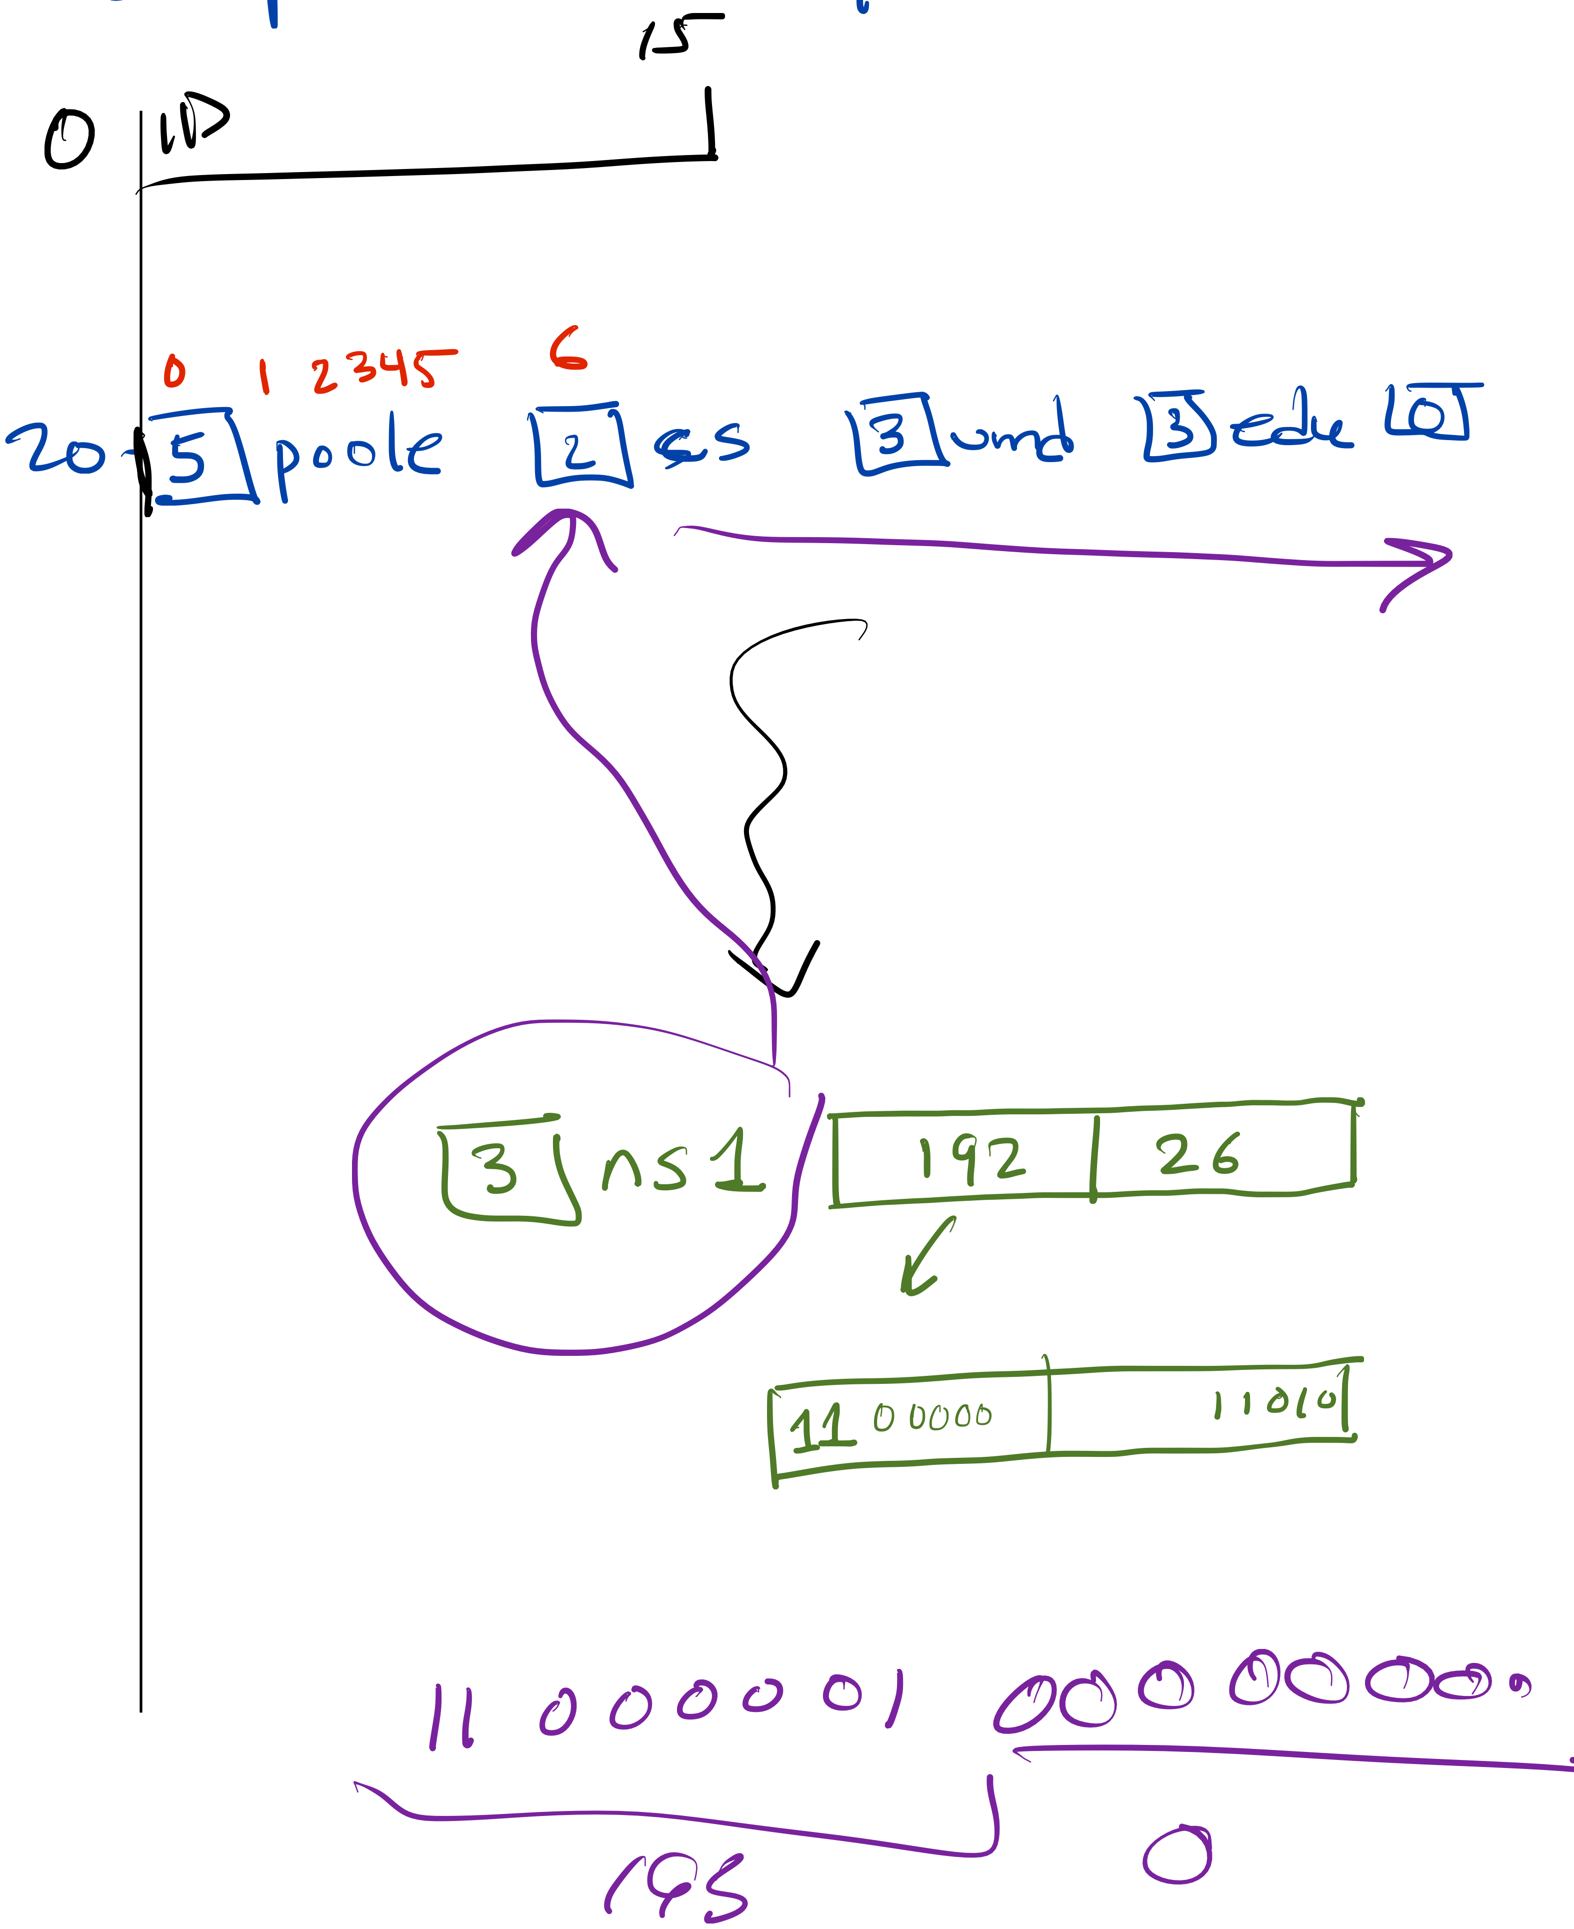

PTR queries

"magic"

in-addr.arpa. suffix

128.8.126.63

↓

63.126.8.128.

in-addr.arpa

PTR query

↑

RR