WELCOME TO DOS CLASS

Hermann Härtig

Senior Professor

Head of OS group 1994 - 2019

ORGANISATION

- Lecturer in charge of DOS:
  Dr. Carsten Weinhold, Barkhausen Institute TUD
- Several lectures presented by research-group members.
- register for mailing list !!!   see Website
  - only way to inform about short-term issues
  - for questions and discussions and lecturing methods
  - must use: “tu-dresden.de” mail-adresses

- Name no more precise →
  Interesting/advanced Topics in Operating Systems
  - scalability
  - systems security
  - modeling
- Some overlap with „Distributed Systems“ (Prof Schill) and some classes by Prof Fetzer
- In some cases no written material (except slides)

TECHNISCHE UNIVERSITÄT DRESDEN

Faculty of Computer Science Institute of Systems Architecture, Operating Systems Group

1.0) DOS ORGANISATION

1.1) SCALABILITY IN COMPUTER SYSTEMS

1.2) EXAMPLE: DNS/BIND

HERMANN HÄRTIG, DISTRIBUTED OPERATING SYSTEMS, SS2020

GOAL OF ALL LECTURES ON SCALABILITY

Topics:

- scalability: terminology, problems, principle approaches
- case studies, all layers of compute systems

Goal:

- understand (some of the) important principles how to build scalable systems

TODAY

Outline:

- scalability … and simple model to reason about 1 aspect
- names in Distributed Systems:
  purposes of naming, terminology. (DNS)
- application of scalability approaches on name resolution

Goal:

- understand some of the important principles how to build scalable systems … using DNS as example

MORE CASE STUDIES LATER IN THE CLASS

- memory consistency
- locks and advanced synchronization approaches
- file systems
- load balancing (MosiX) and HPC(MPI)

GENERAL DEFINITION: SCALABILITY

Scalability:

Scalability is the property of a system to handle a growing amount of work by adding resources to the system.

Wikipedia (2019) and many other sources

SCALABILITY: WEAK ./. STRONG

ability of a system to use growing resources ...

- weak:
  to handle growing load, larger problem, ...
- strong:
  accelerate existing work load, same problem

problems

THINK and PAUSE!

PAUSE THE VIDEO HERE
AND THINK before continuing

PROBLEMS

- performance bottlenecks / Amdahl's Law
- failures / abuse
- administration

RESOURCES AND PERFORMANCE

- processors
- communication
- memory (remember basic OS course: “thrashing”)

$$\text{Speedup:} \frac{\text{original execution time}}{\text{enhanced execution time}}$$

SIMPLE MODEL: AMDAHL’S LAW

$$\text{Speedup:} \frac{\text{original execution time}}{\text{enhanced execution time}}$$

Parallel Execution

red: cannot run in Parallel

green: runs *perfectly* parallel

unlimited processors maximum speedup: blue/red

AMDAHL'S LAW

Parallel Execution, N processors

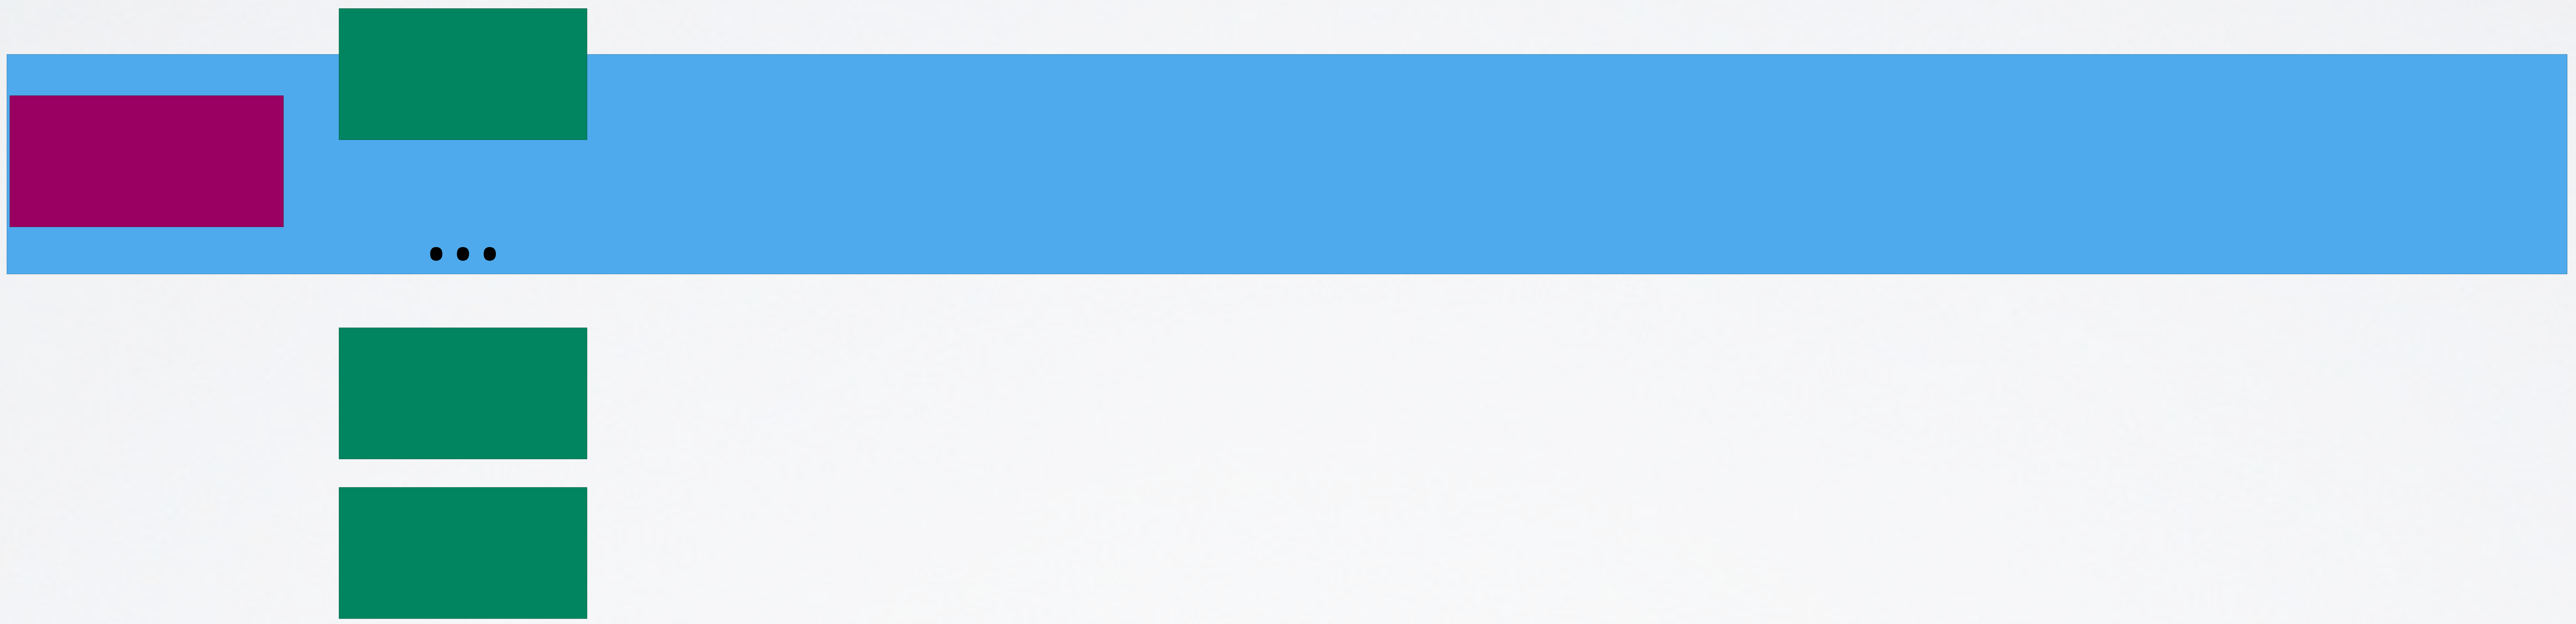

red: cannot run in Parallel

green: runs perfectly parallel

N processors maximum speedup: blue/(red+green/N)

Parallel Execution, N processors

red: cannot run in Parallel

green: runs perfectly parallel

maximum speedup: blue/(red+green/N)

AMDAHL'S LAW

$$\text{Speedup:} \frac{\text{original execution time}}{\text{enhanced execution time}}$$

- P: section that can be parallelized
- 1-P: serial section
- N: number of CPUs

$$\text{Speedup}(P,N) = \frac{1}{\left(1 - P + \frac{P}{N}\right)}$$

- if N becomes VERY large, speedup approaches: 1/(1-P) principle approaches

THINK and PAUSE!

PAUSE THE VIDEO HERE
AND THINK before continuing

THE "RPC" PRINCIPLES

- partitioning
  split systems into parts that can operate independently/parallel to a large extent
- replication
  provide several copies of components
  - that are kept consistent eventually
  - that can be used in case of failure of copies
- locality (caching)
  maintain a copy of information that is nearer, cheaper/faster to access than the original

MORE PRINCIPLES

- identify and address bottlenecks
- specialize functionality/interfaces
- right level of consistency
  caches, replicates, … need not always be fully consistent
- lazy information dissemination
- balance load (make partitioning dynamic)

1.0) DOS ORGANISATION
1.1) SCALABILITY IN COMPUTER SYSTEMS
1.2) EXAMPLE: DNS/BIND

HERMANN HÄRTIG, DISTRIBUTED OPERATING SYSTEMS, SS2020

MY FIRST EMAIL

- UUCP/MMDF (cum grano salis):
  - ira!gmdzi!oldenburg!heinrich!user (path to destination)
  - user@ira!heinrich%gmdzi
    (mixing identifiers and path information)

A BIT OF HISTORY

- ARPA-Net at the beginning:
  - a single file: hosts.txt
  - maintained at Network Information Center of SRI (Stanford)
  - accessed via ftp
  - TCP/IP in BSD Unix
    => chaos name collisions, consistency, load
- DNS: Paul Mockapetries (84) ...

DOMAIN NAME SYSTEM

NAMES, IDENTIFIERS, ADDRESSES

- names
  - symbolic, many names possible for 1 entity
  - have a meaning for people
- identifiers
  - identifies an entity uniquely
  - are used by programs
- addresses
  - locates an entity
  - changes occasionally (or frequently)

NAME RESOLUTION

- name resolution:
  map symbolic names to a set of attributes such as:
  identifiers, addresses, alias names, security properties
  encryption keys

- Principle interface:
  - Register (Context, Name, attributes, ...)
  - Lookup (Context, Name) -> attributes

DNS DOMAINS

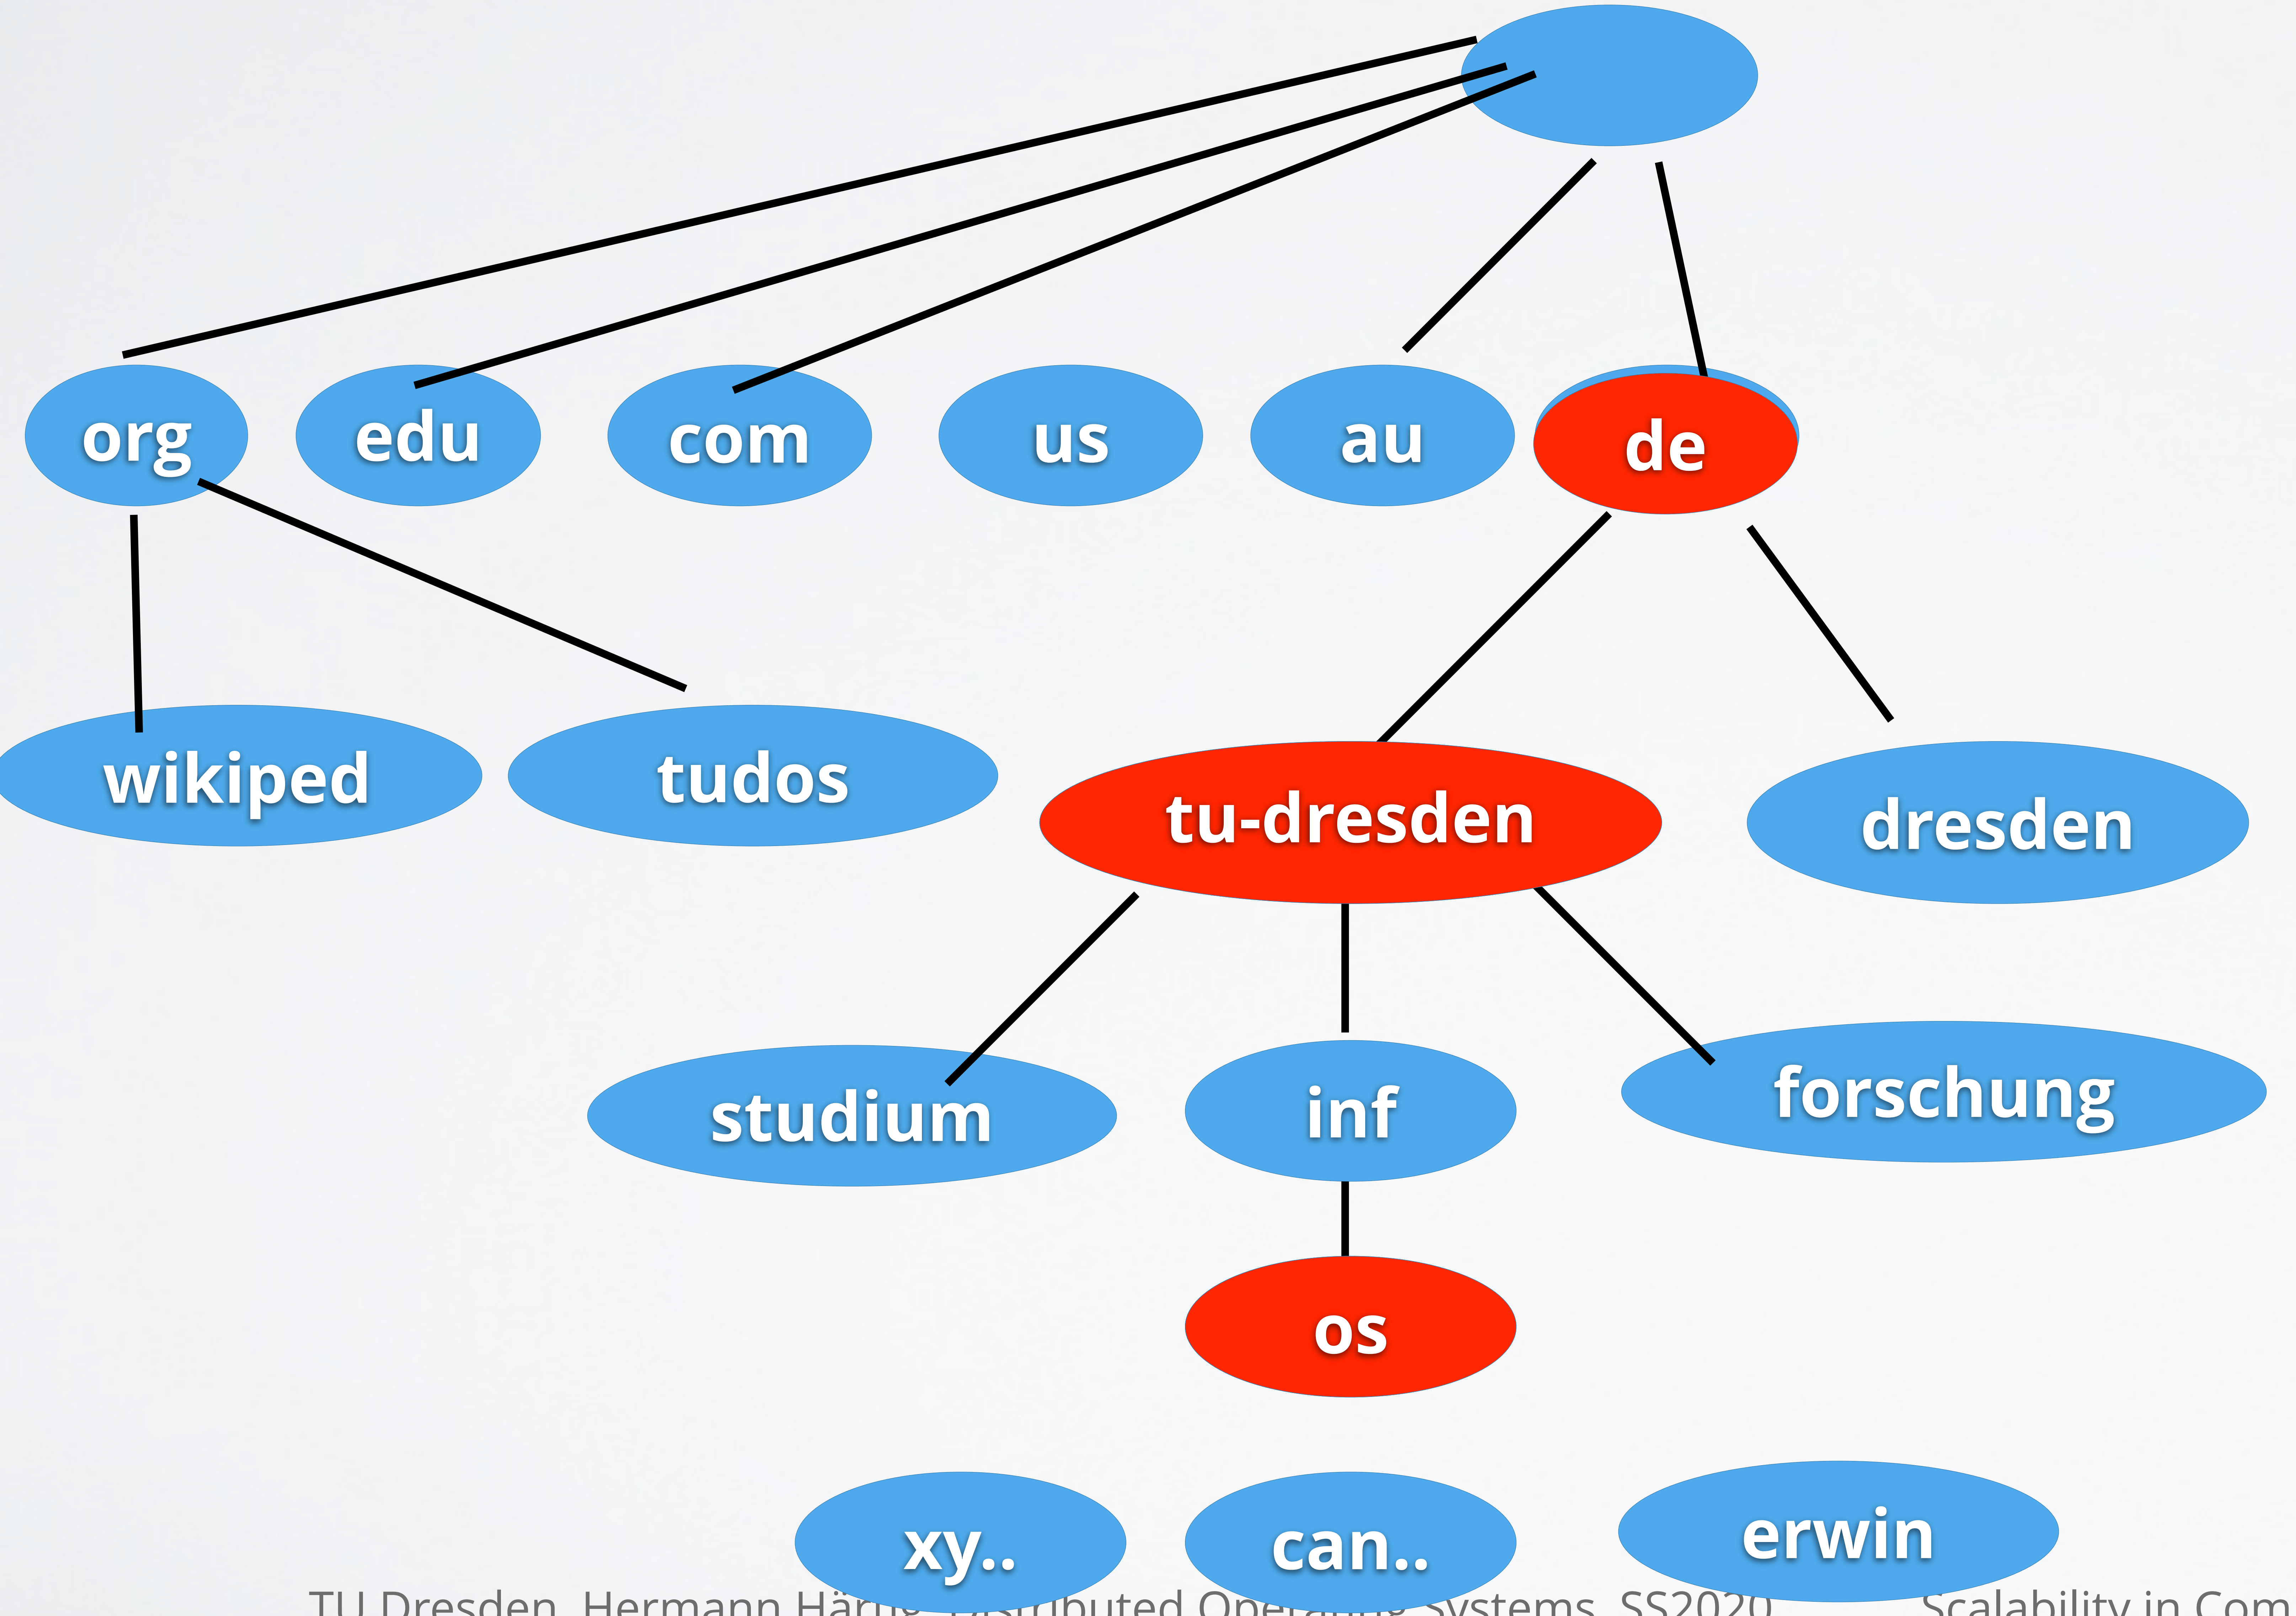

domain
subtree in DNS hierarchy:

- de
- tu-dresden.de
- os.inf.tu-dresden.de
- tudos.org and os.inf.tu-dresden.de are aliases

PARTITIONING: ZONE

- zone: Subset of a domain
  over which an authority has complete control.
  controlled by a name server.
  Subzones can be delegated to other authorities.
- navigation
  querying in a set of cooperating name servers

POTENTIAL ZONES

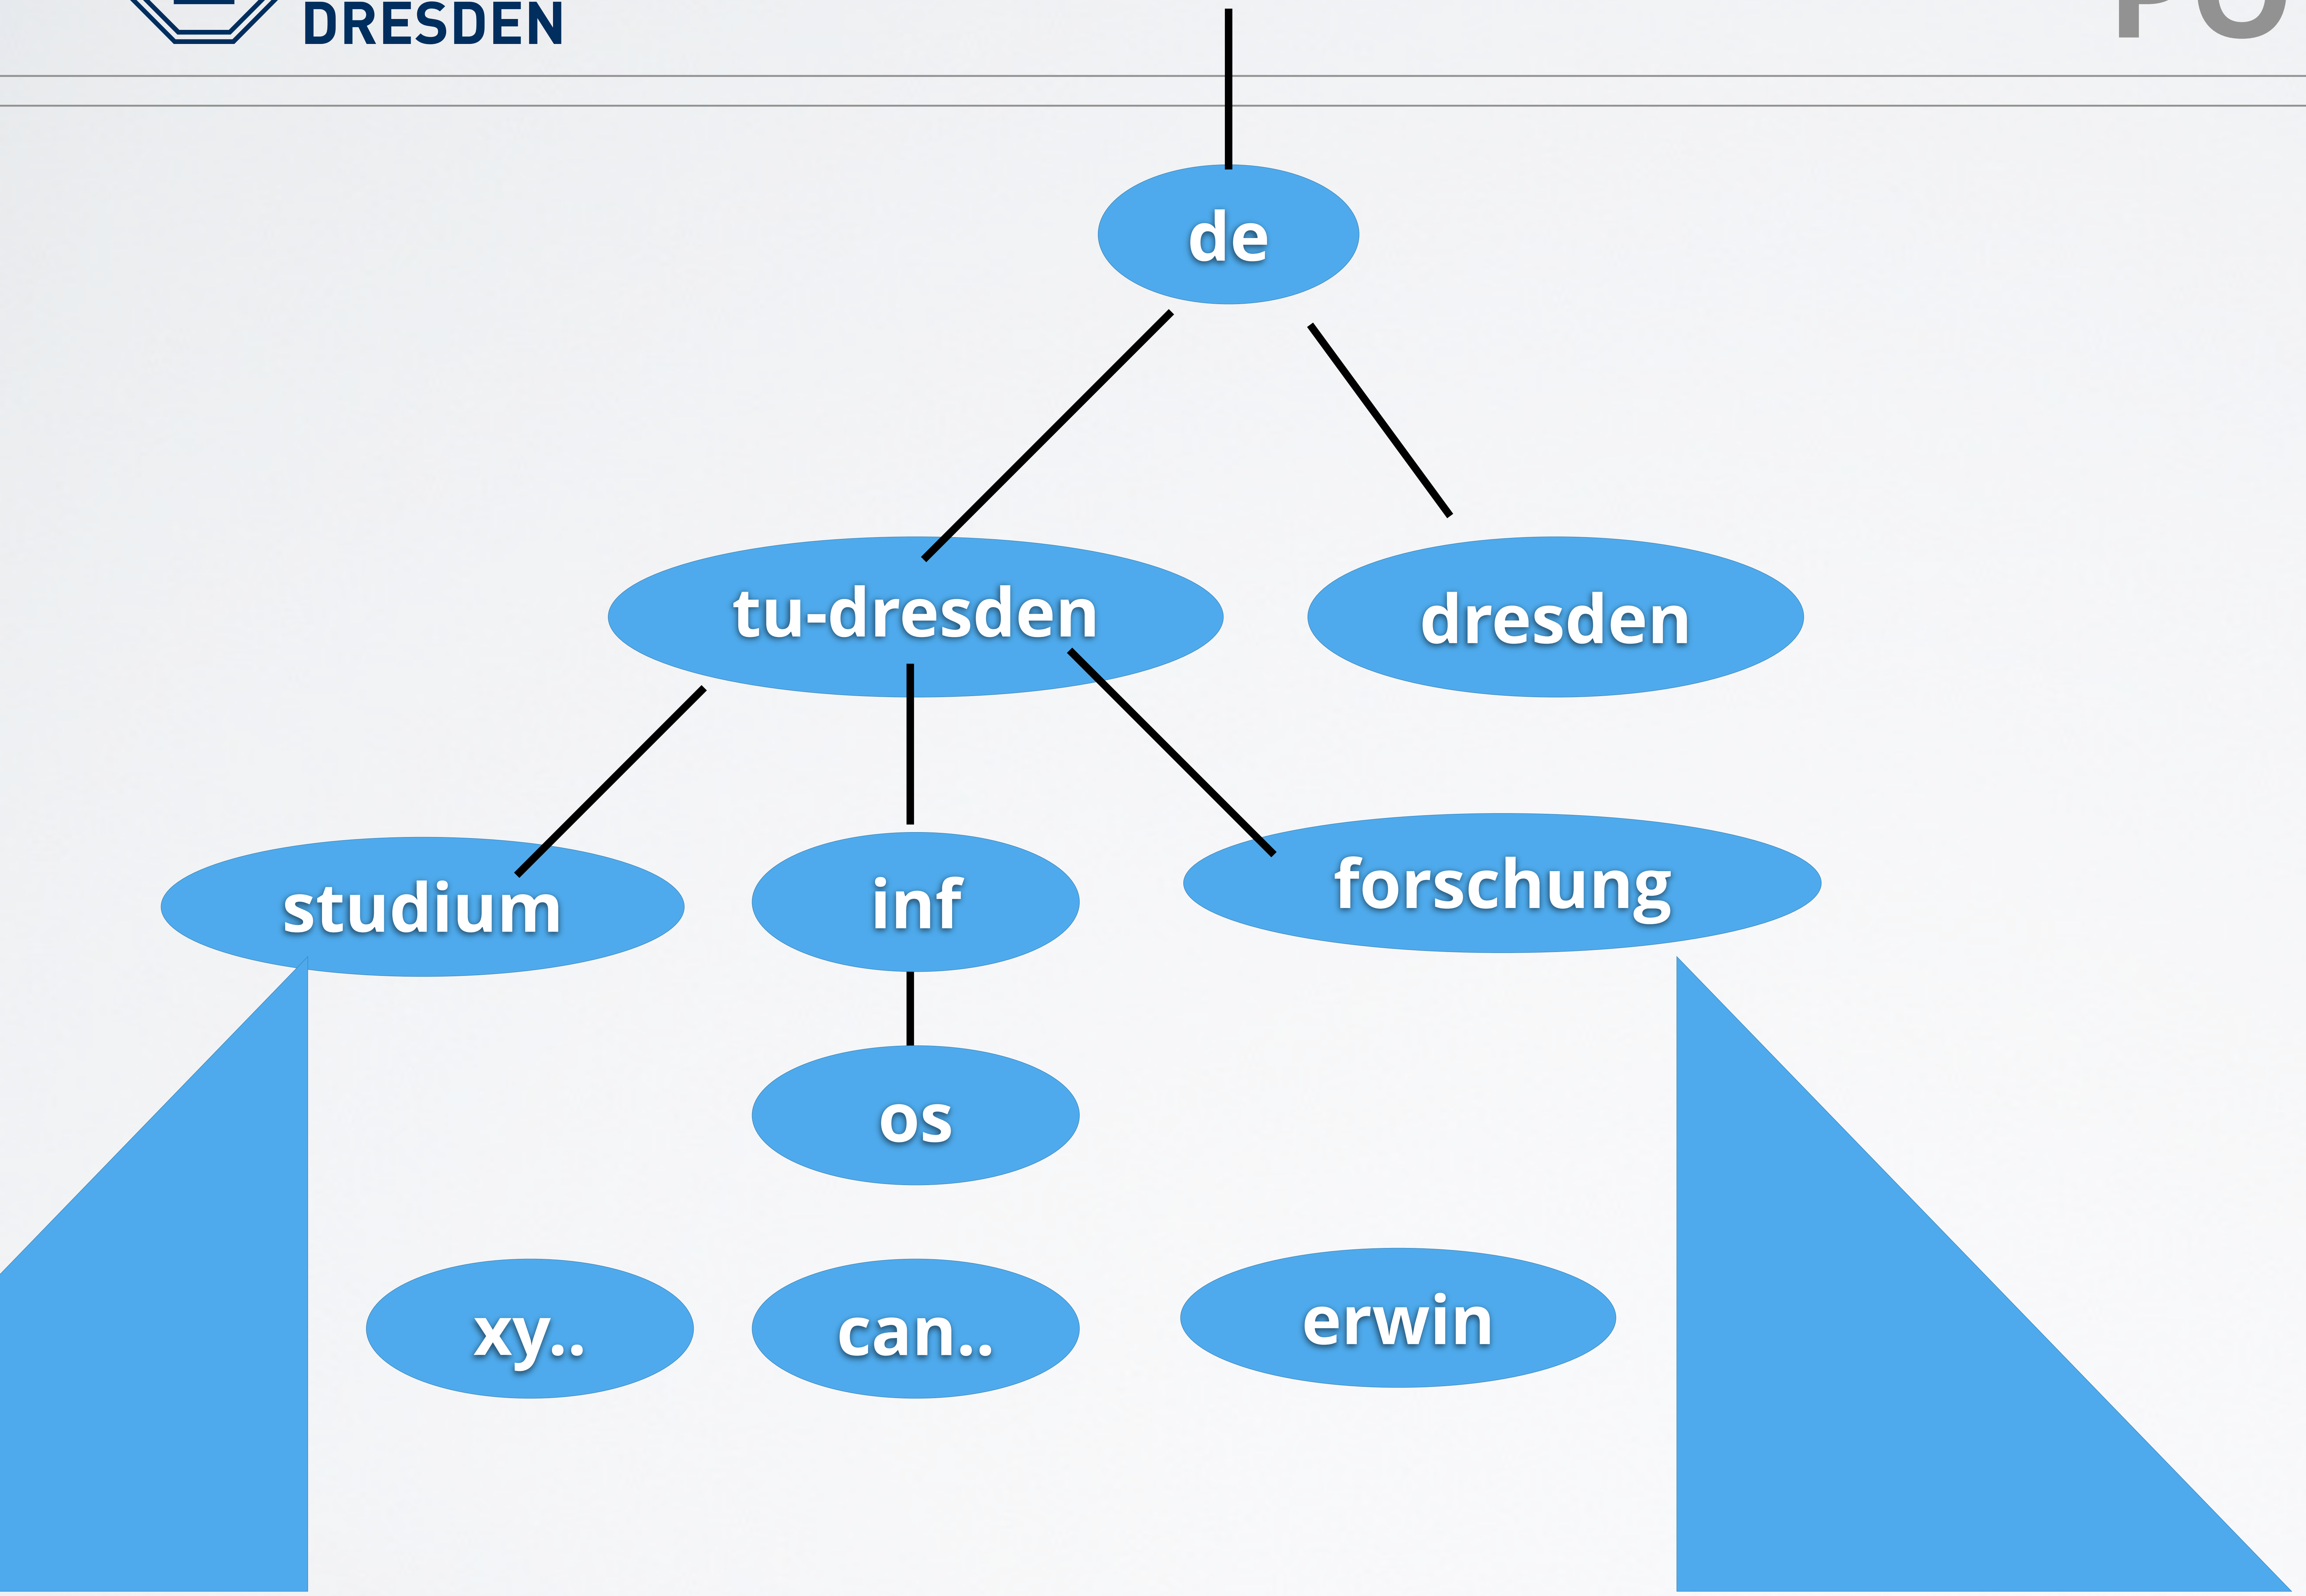

POTENTIAL ZONES

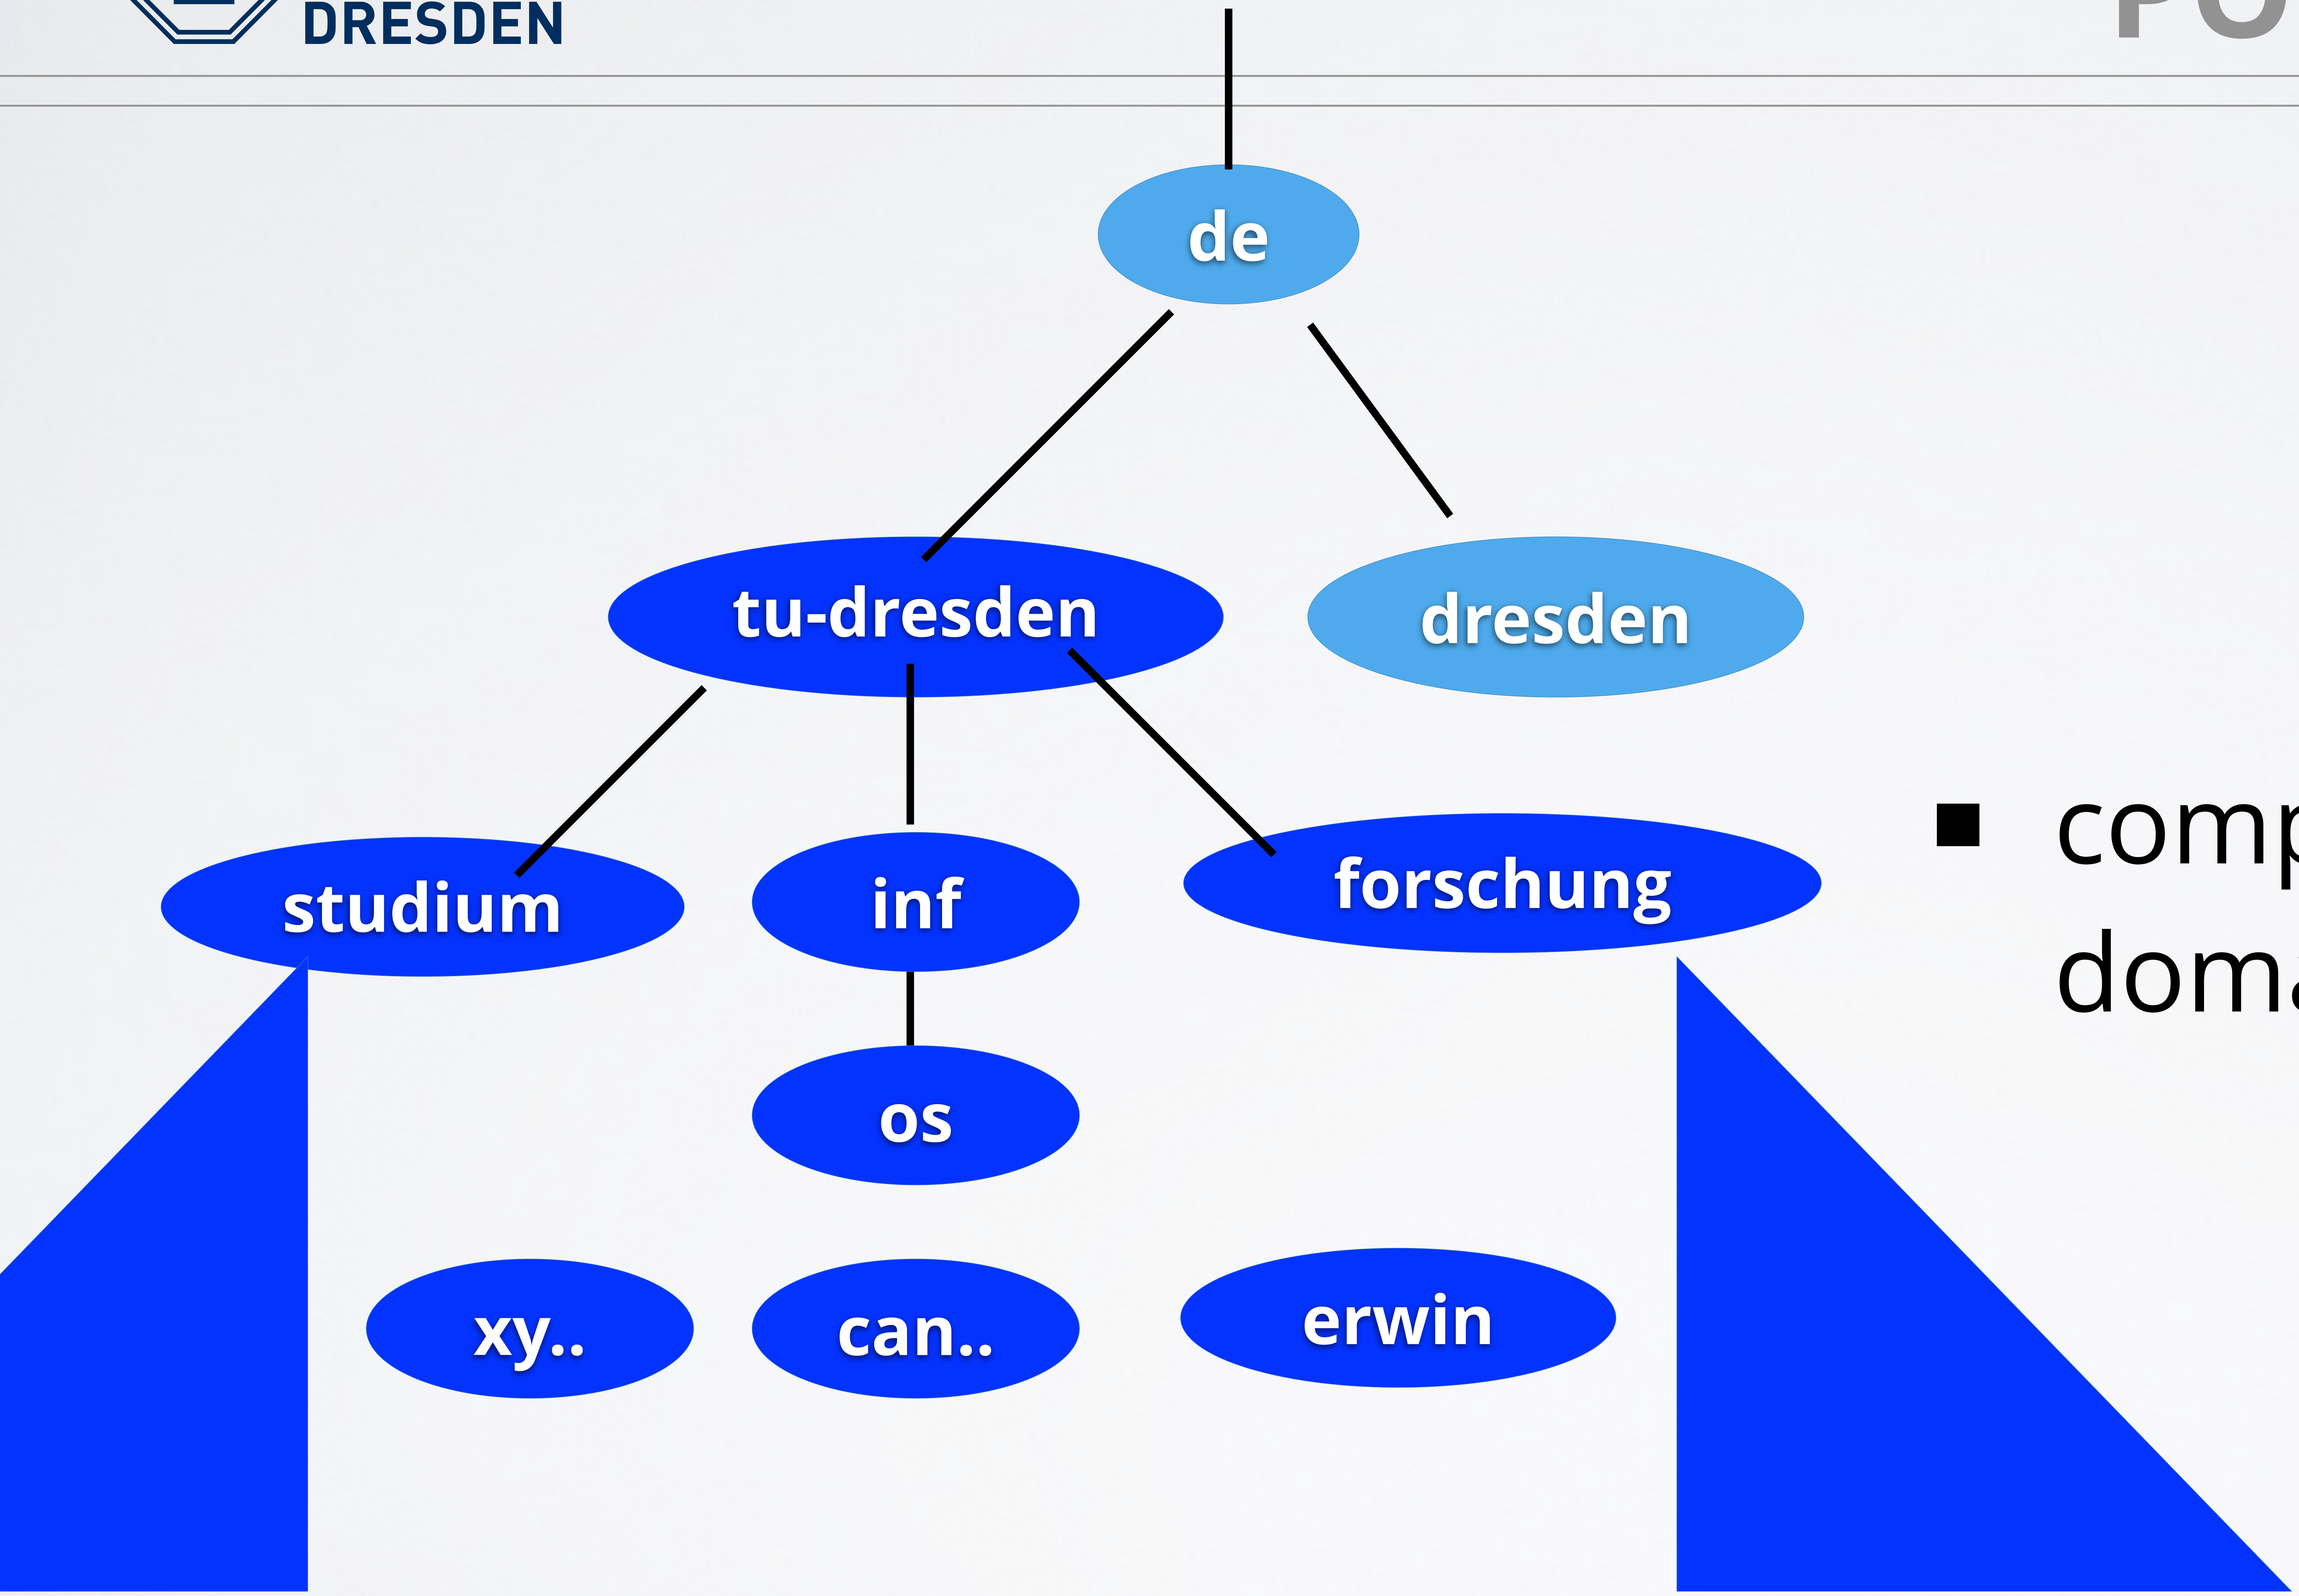

- complete tu-dresden domain

POTENTIAL ZONES

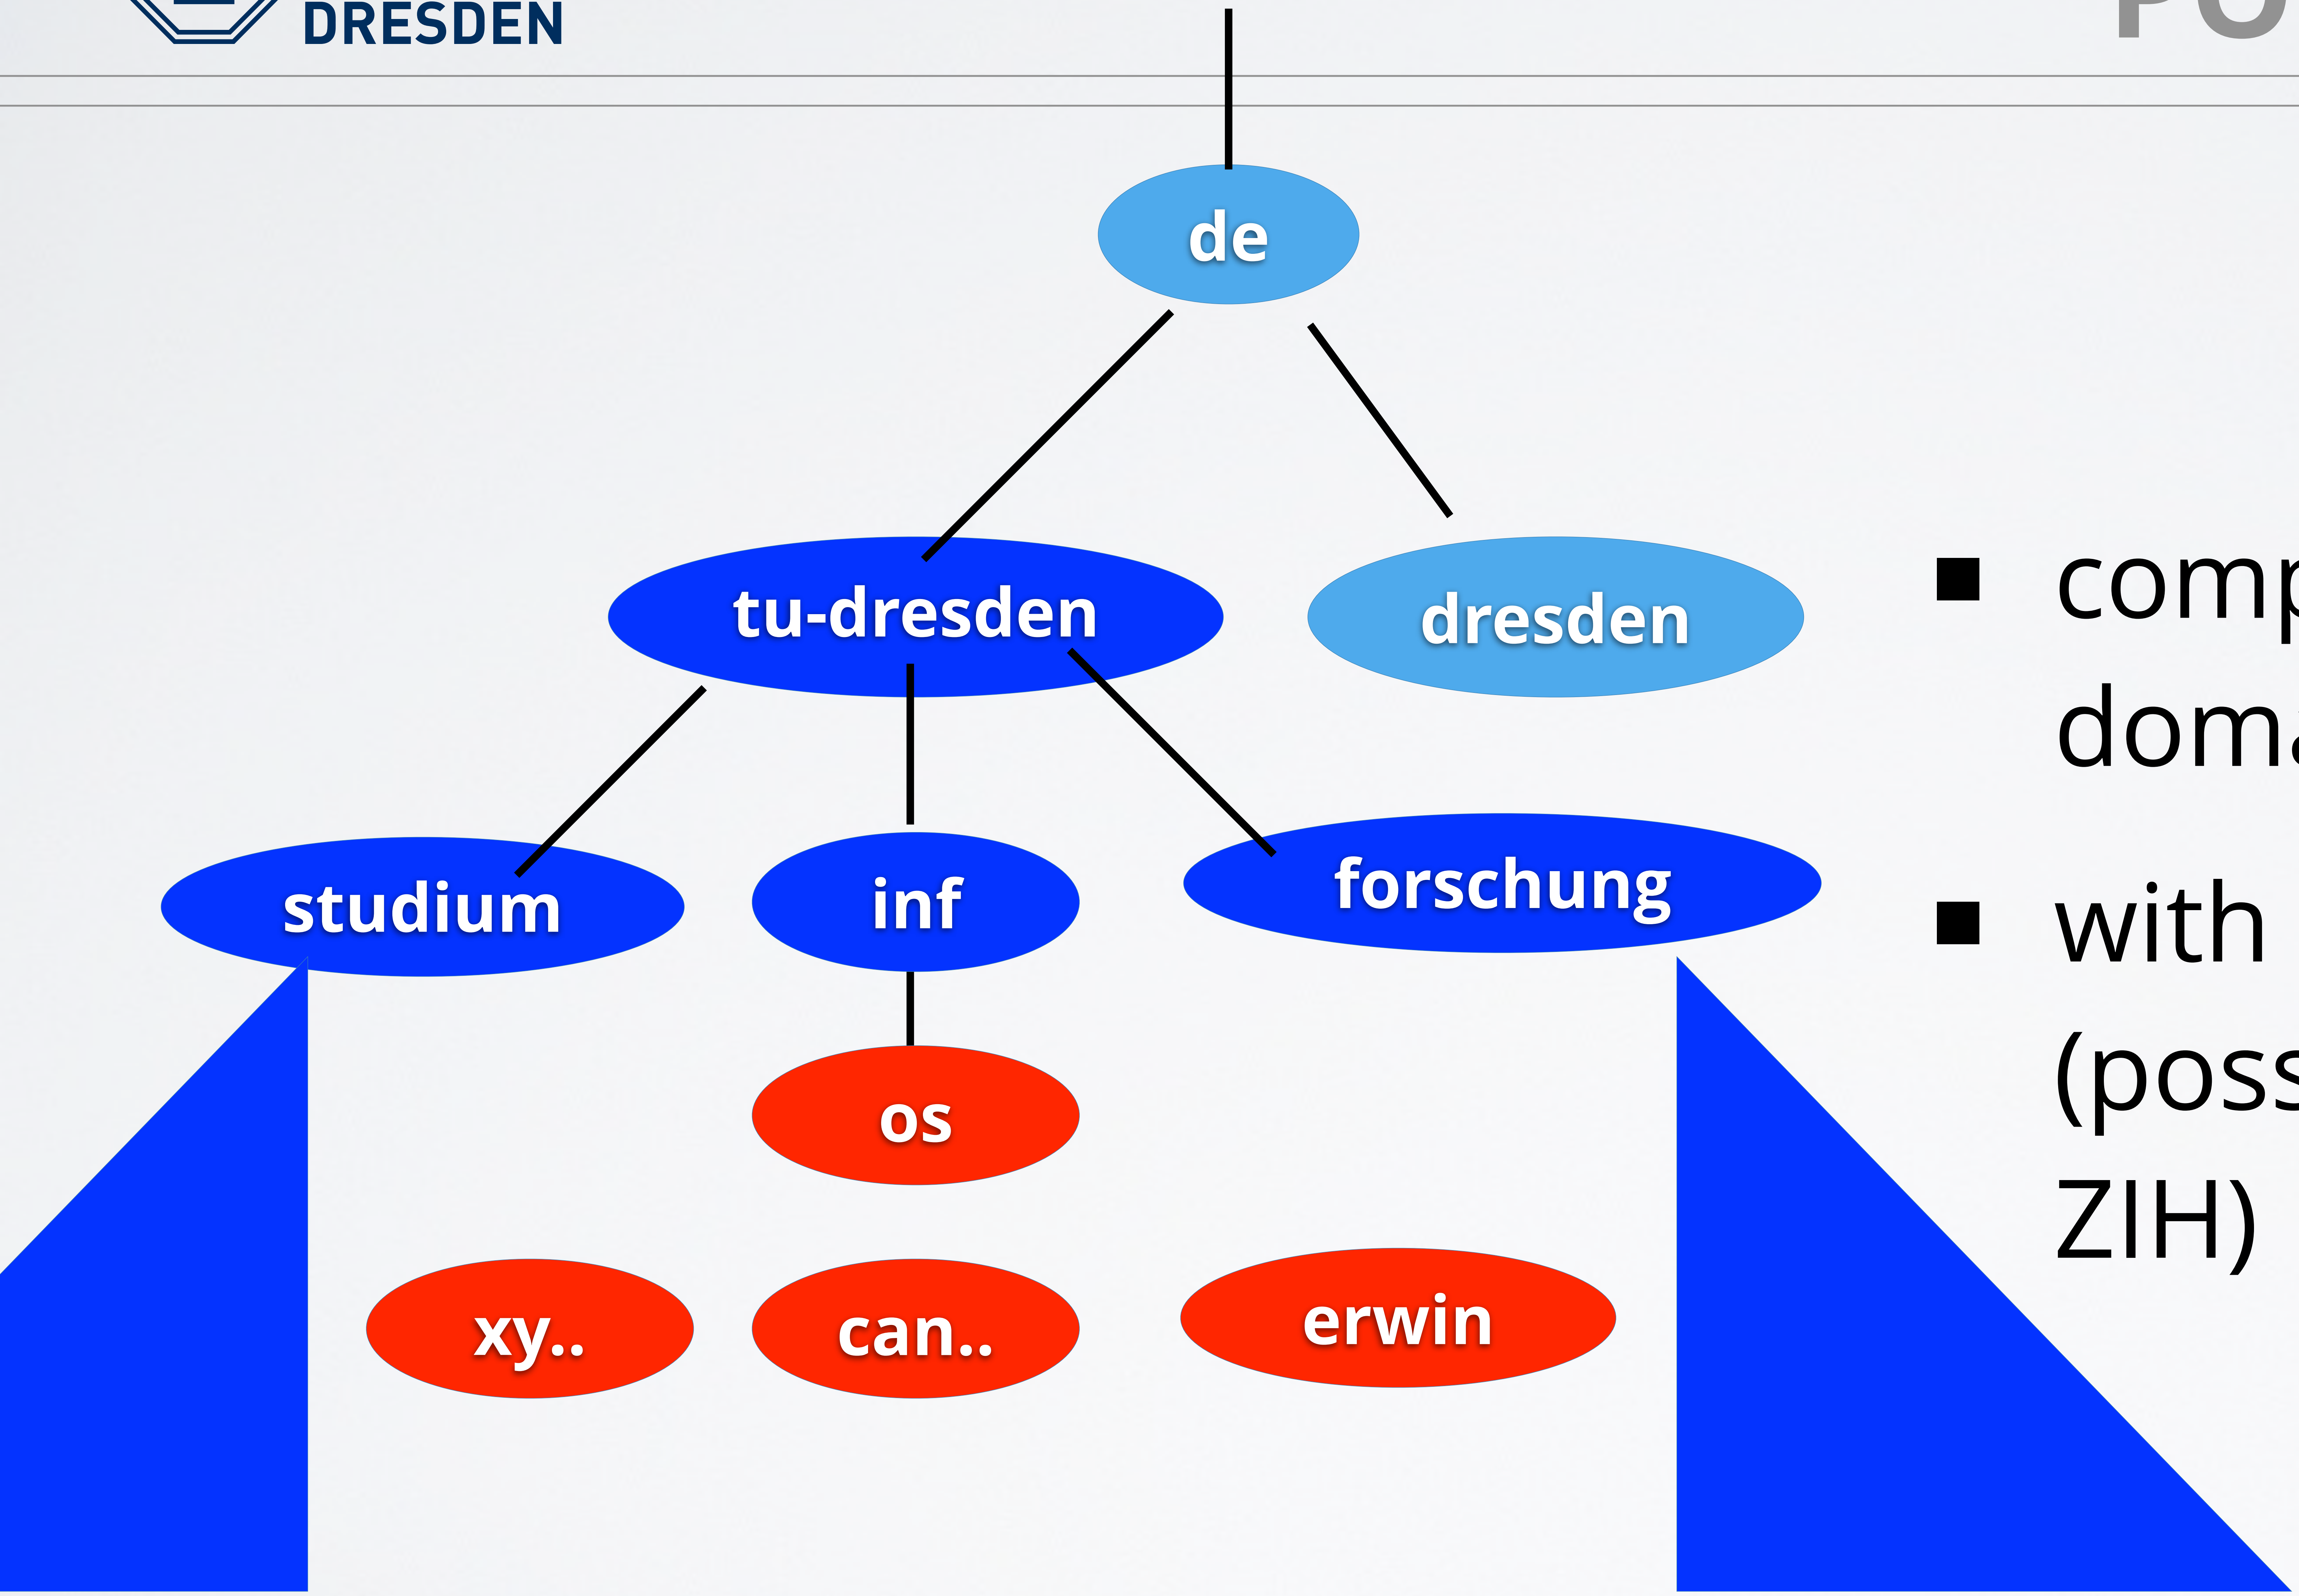

- complete tu-dresden domain
- with sub zone os (possible but unliked by ZIH)

CACHING

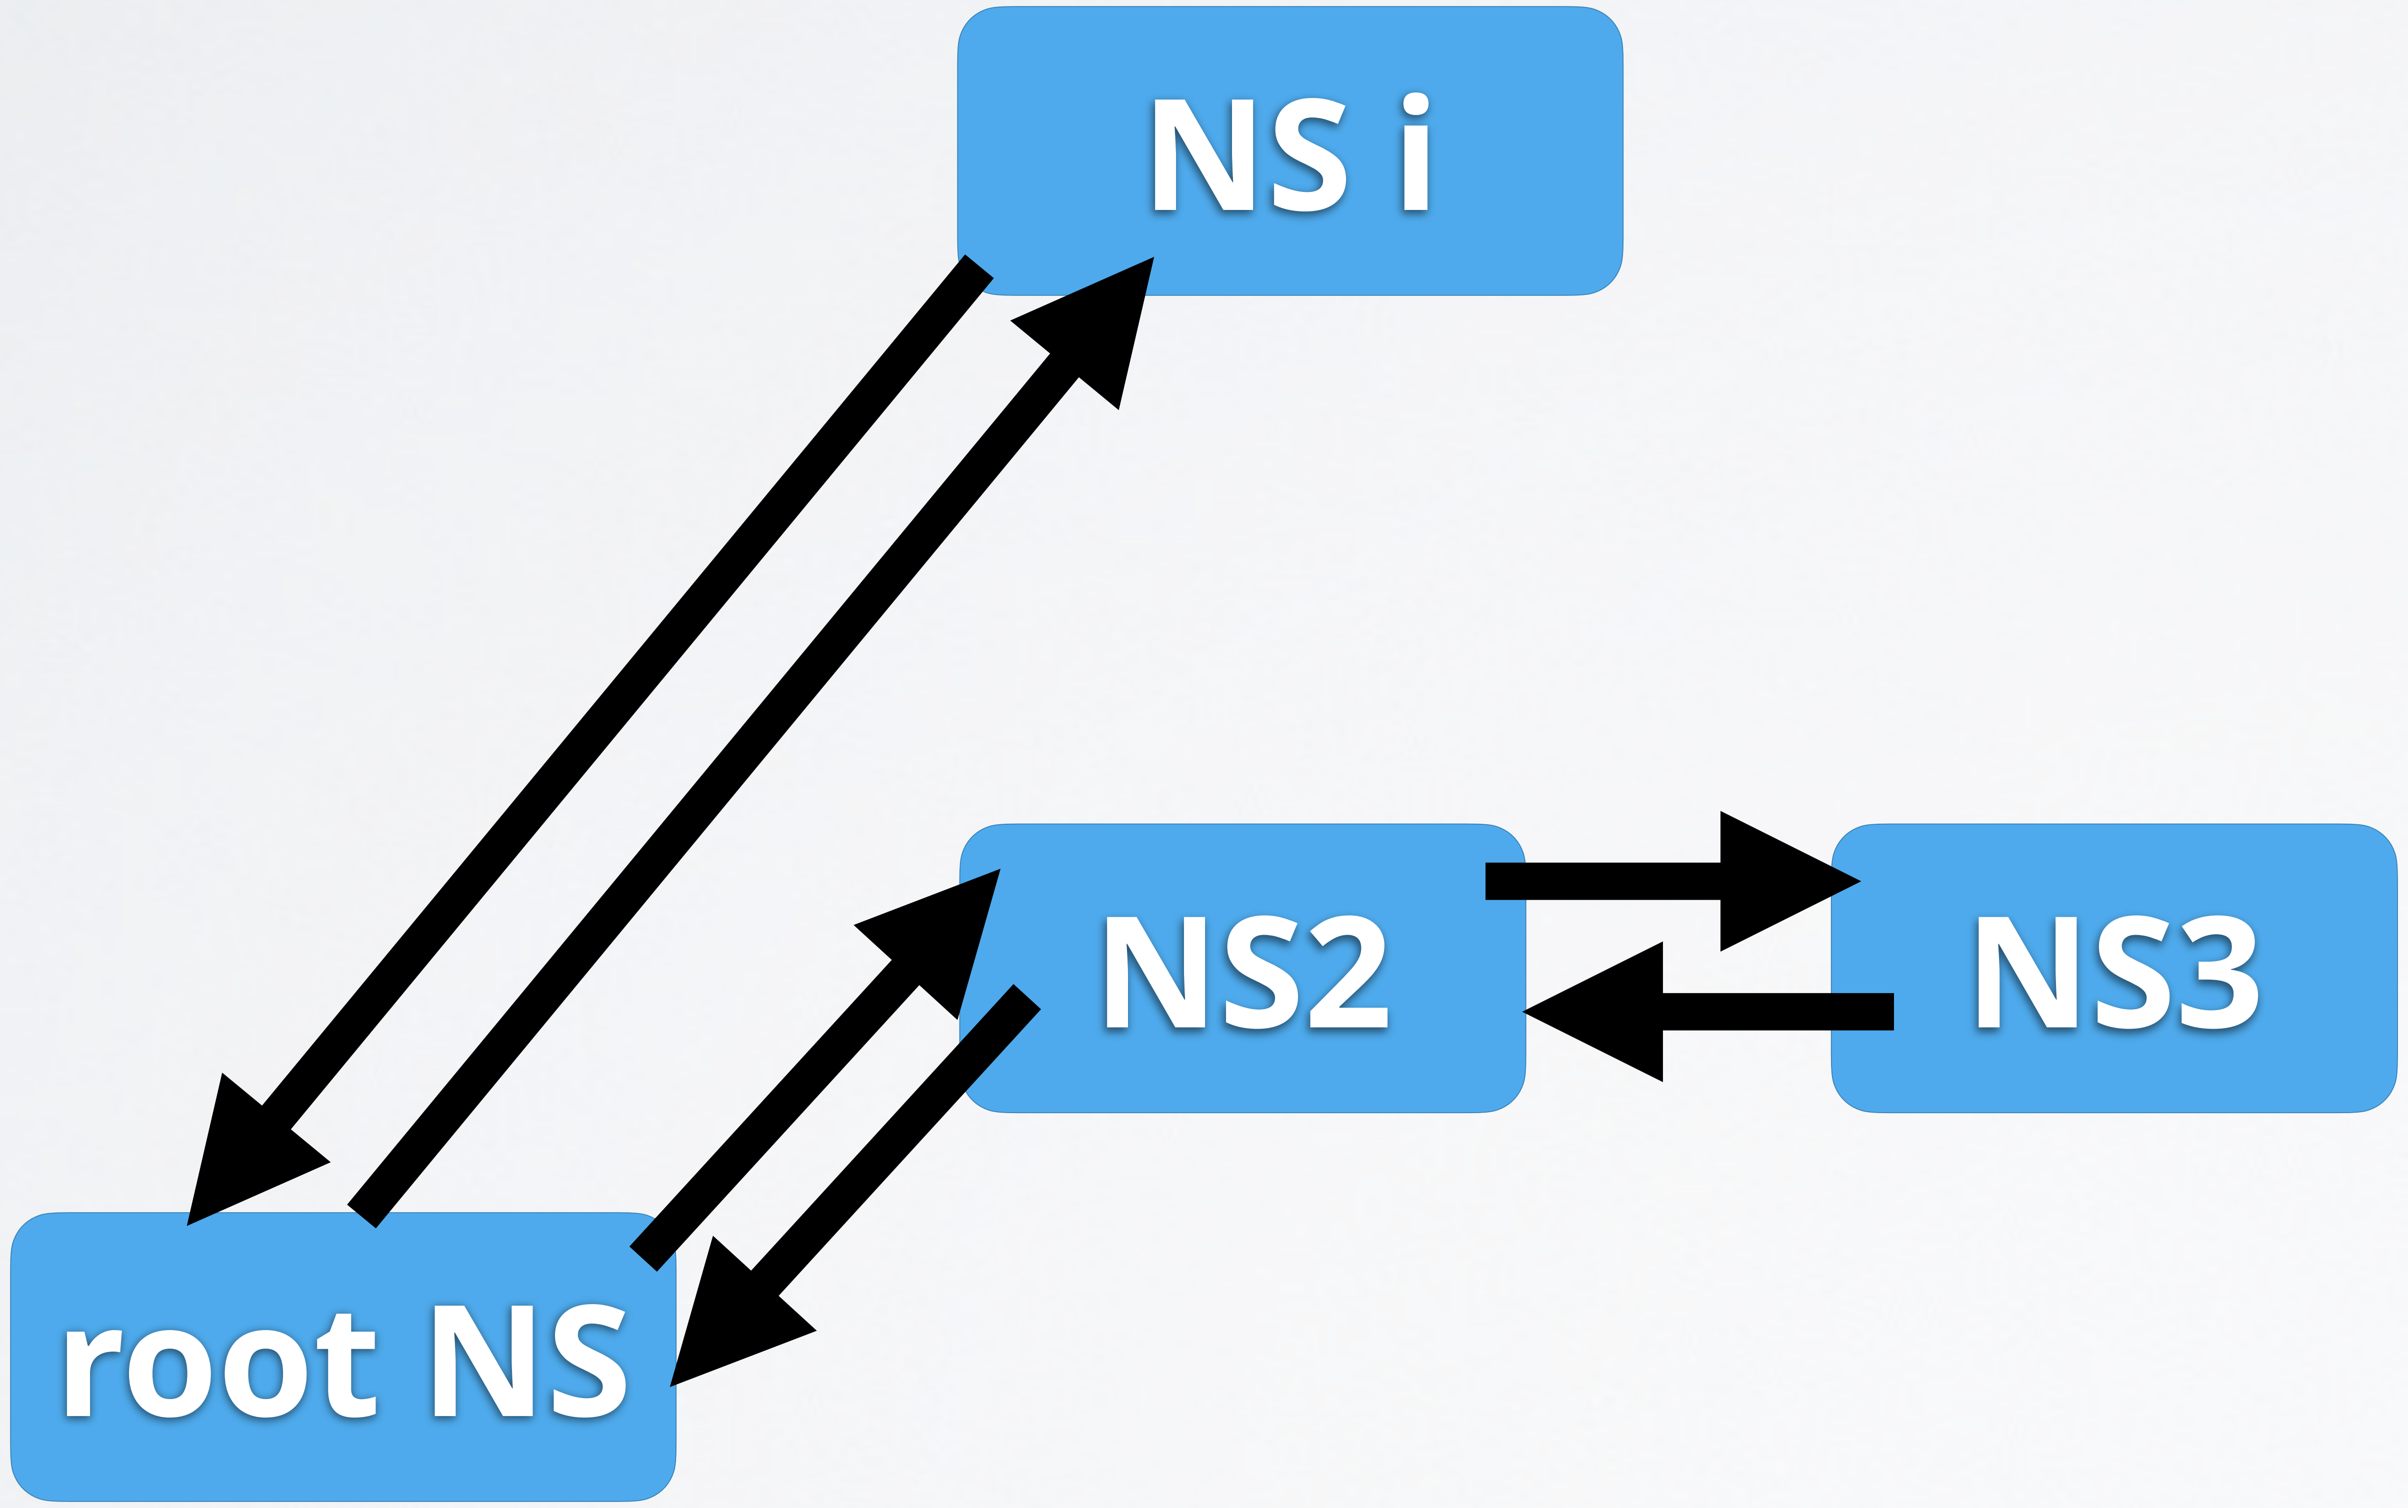

CACHING

- remember intermediate results
- @ root NS makes no sense! (overload)
- @ NS i !!!

CACHING

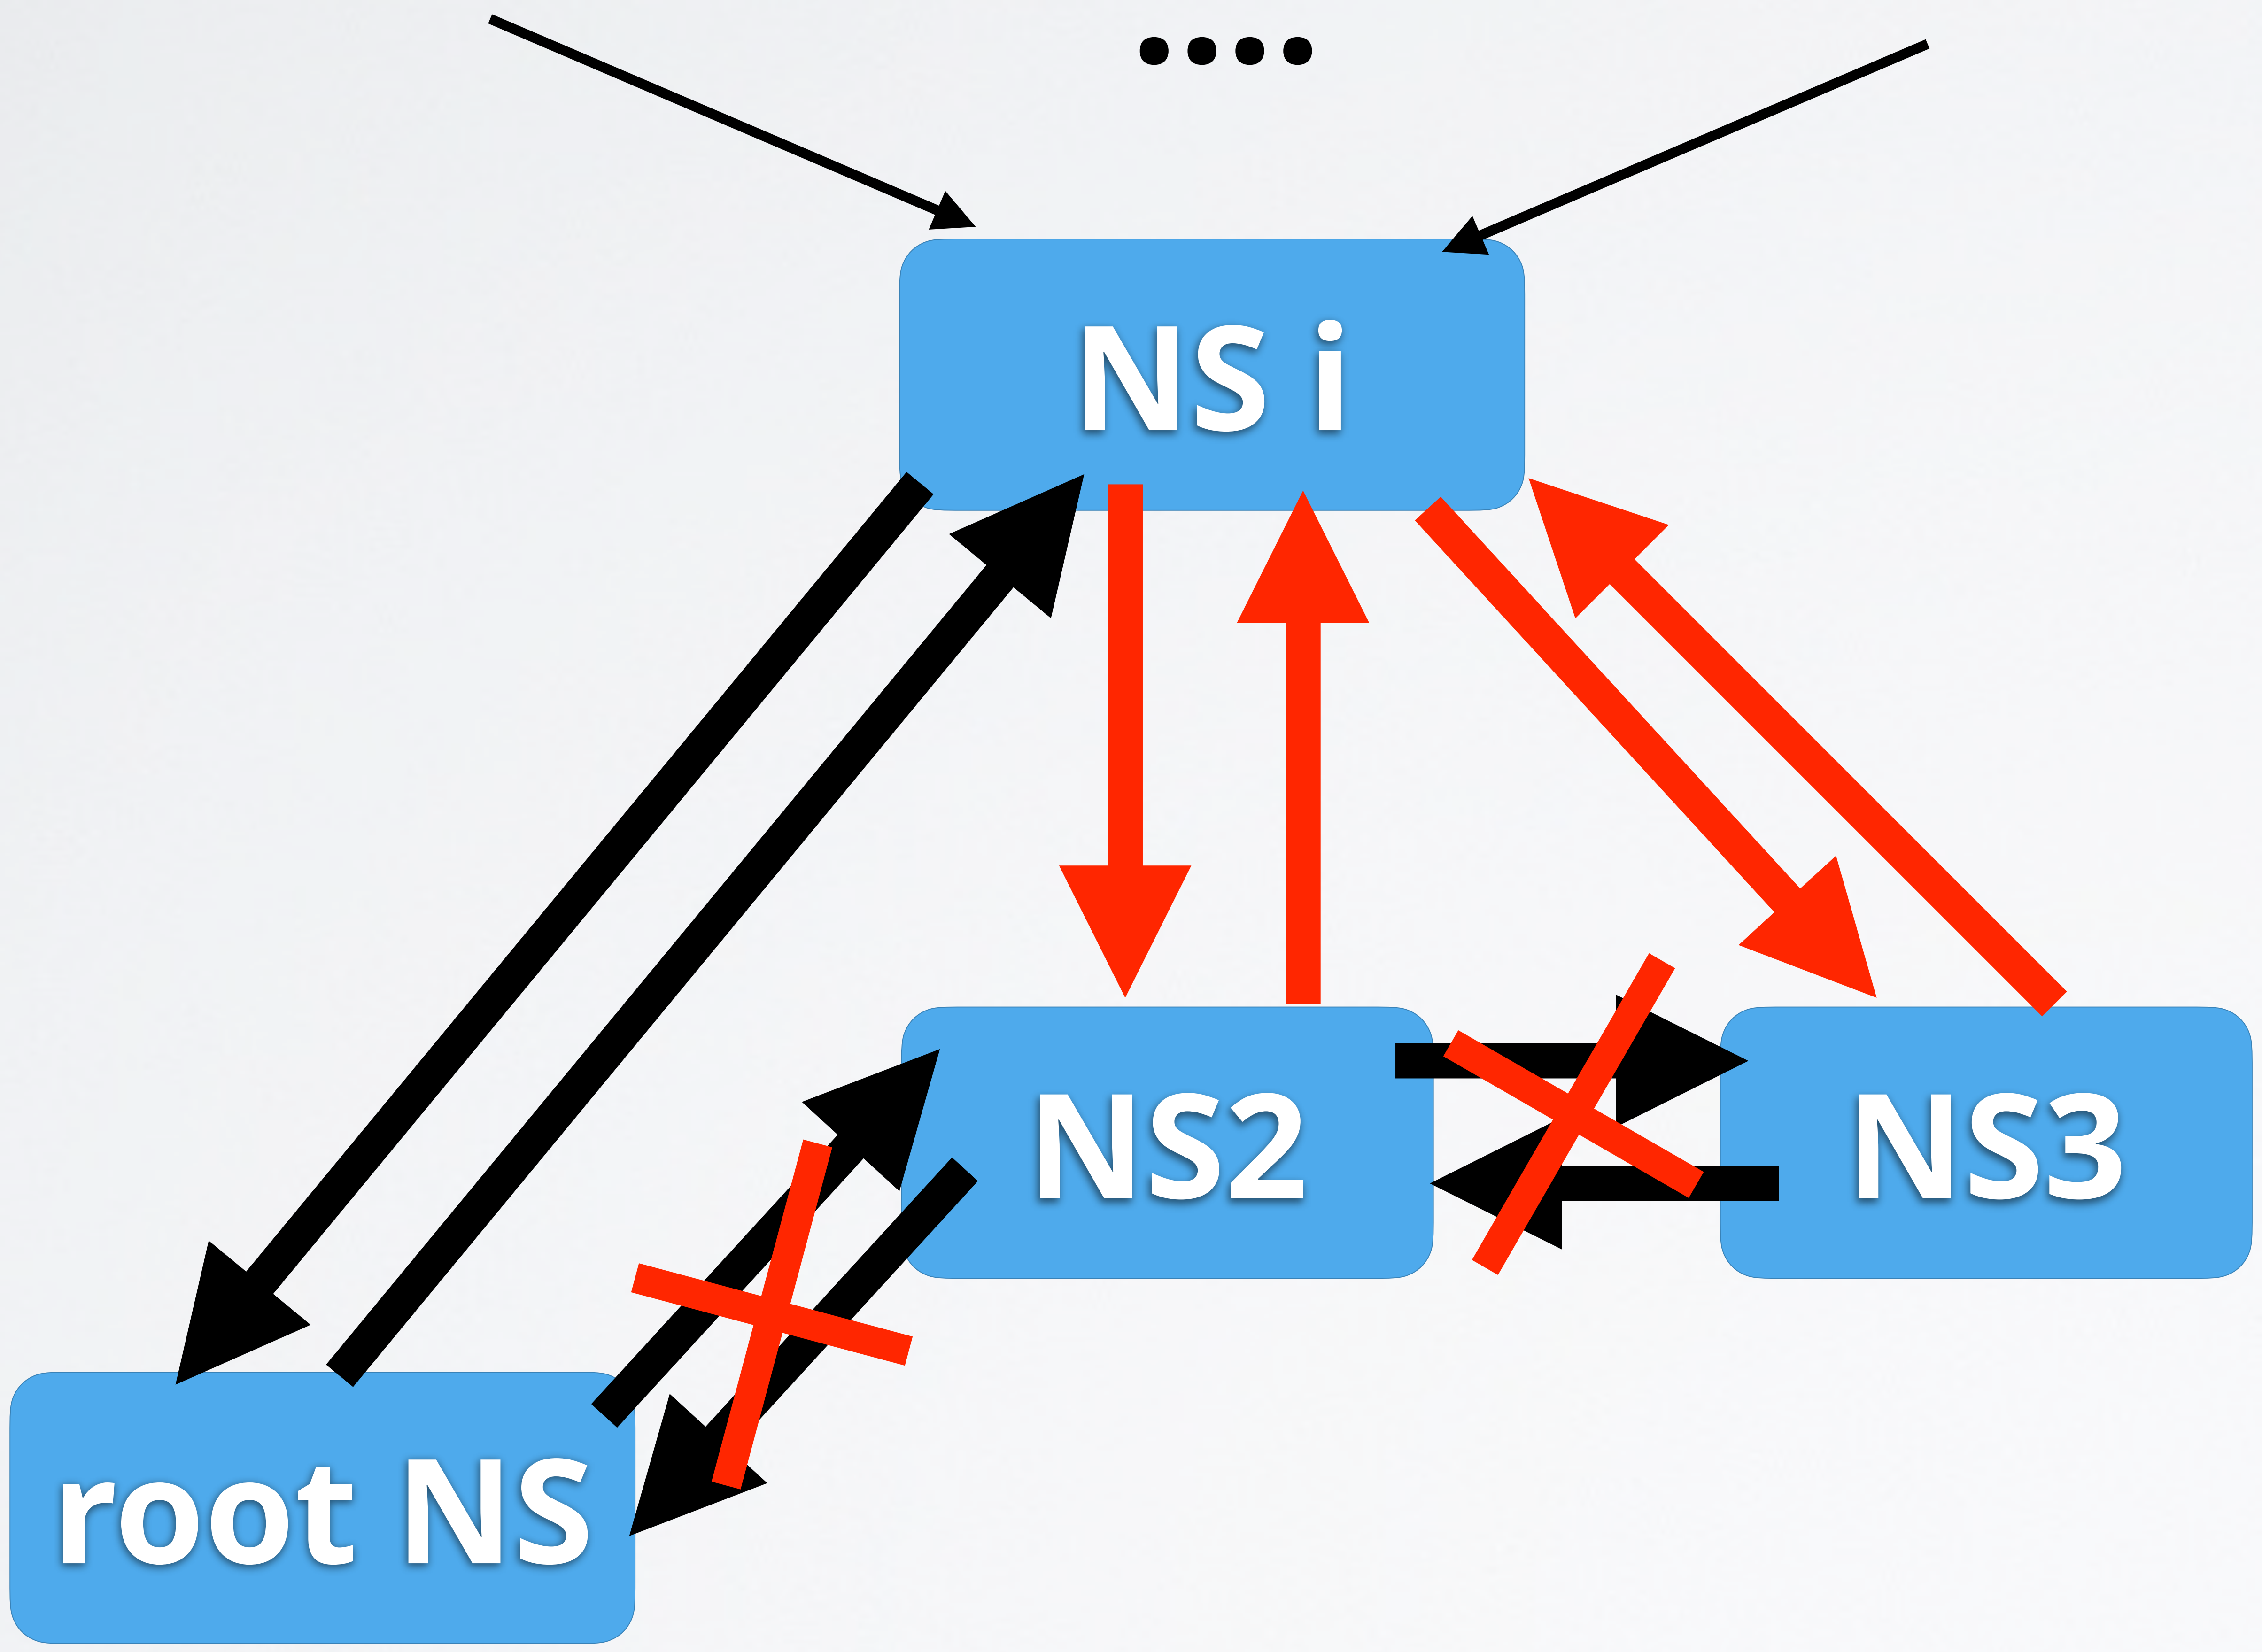

RECURSIVE ./. ITERATIVE

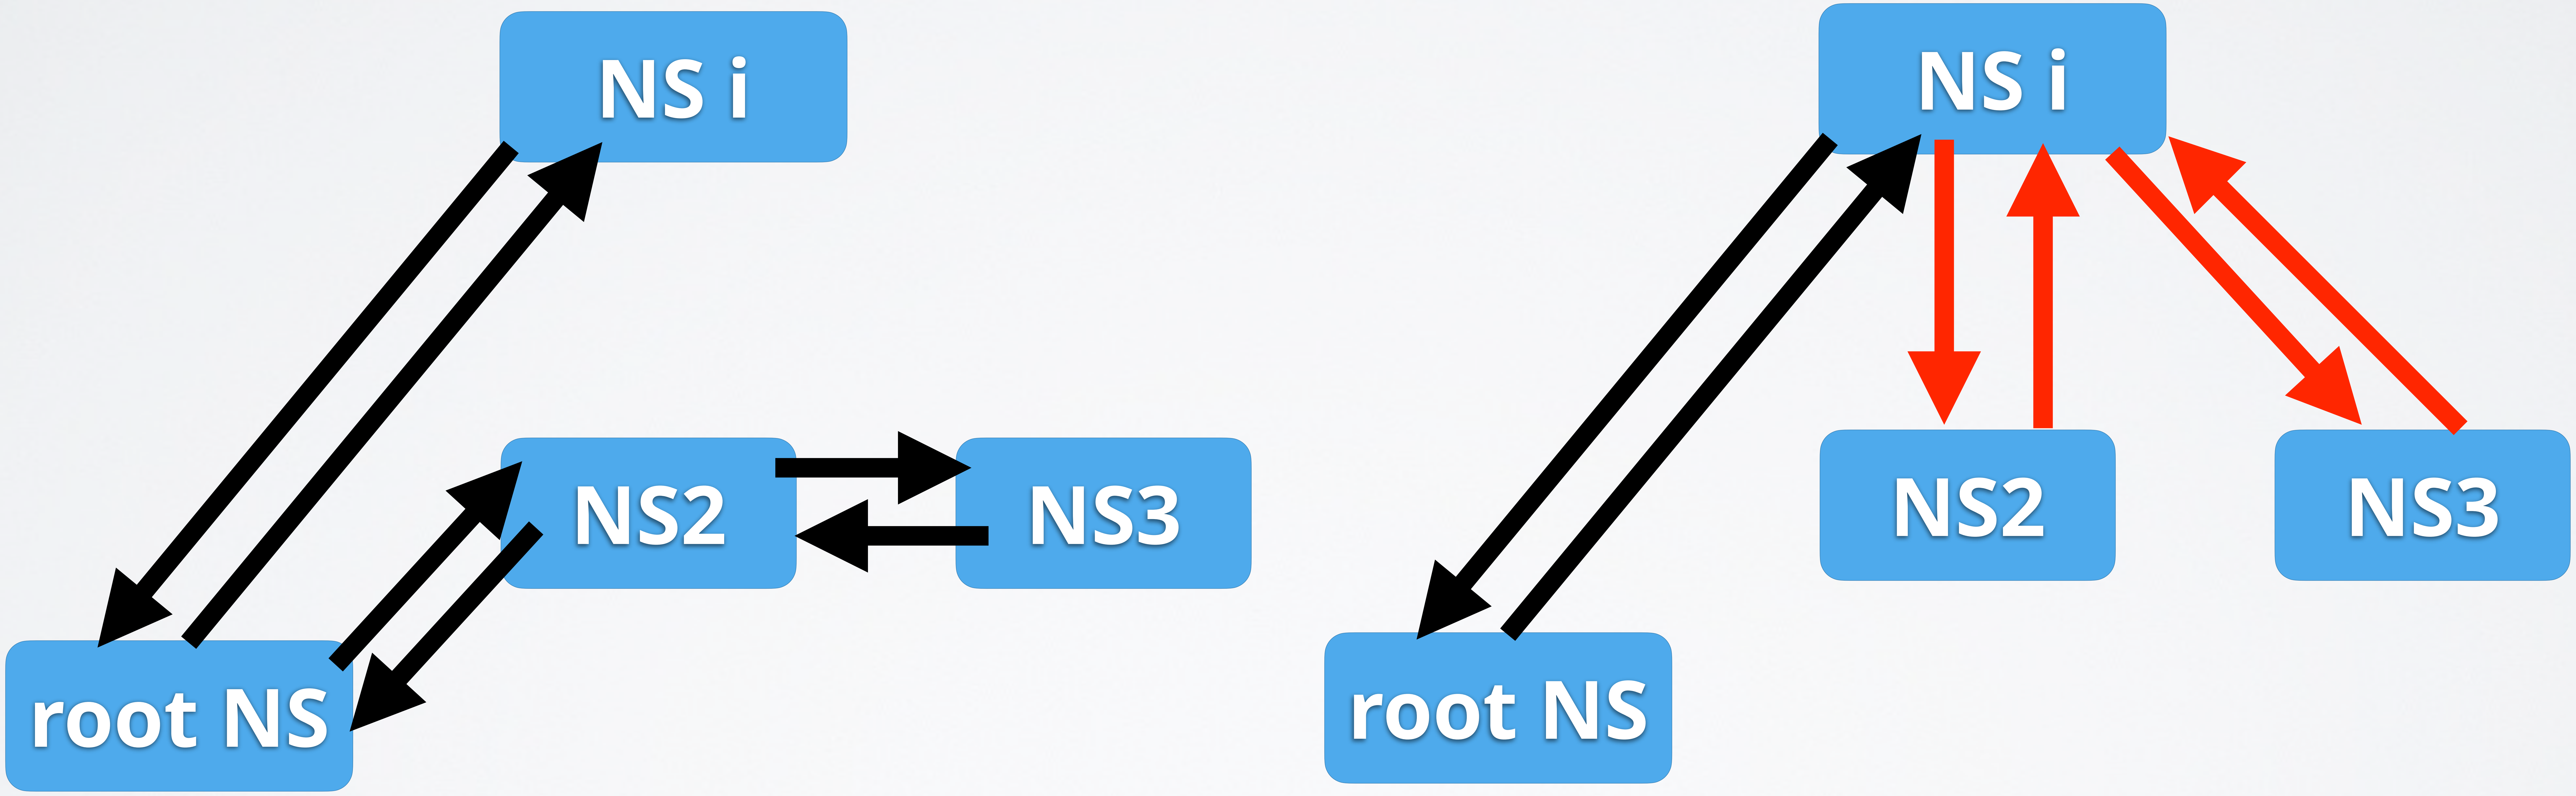

REPLICATION

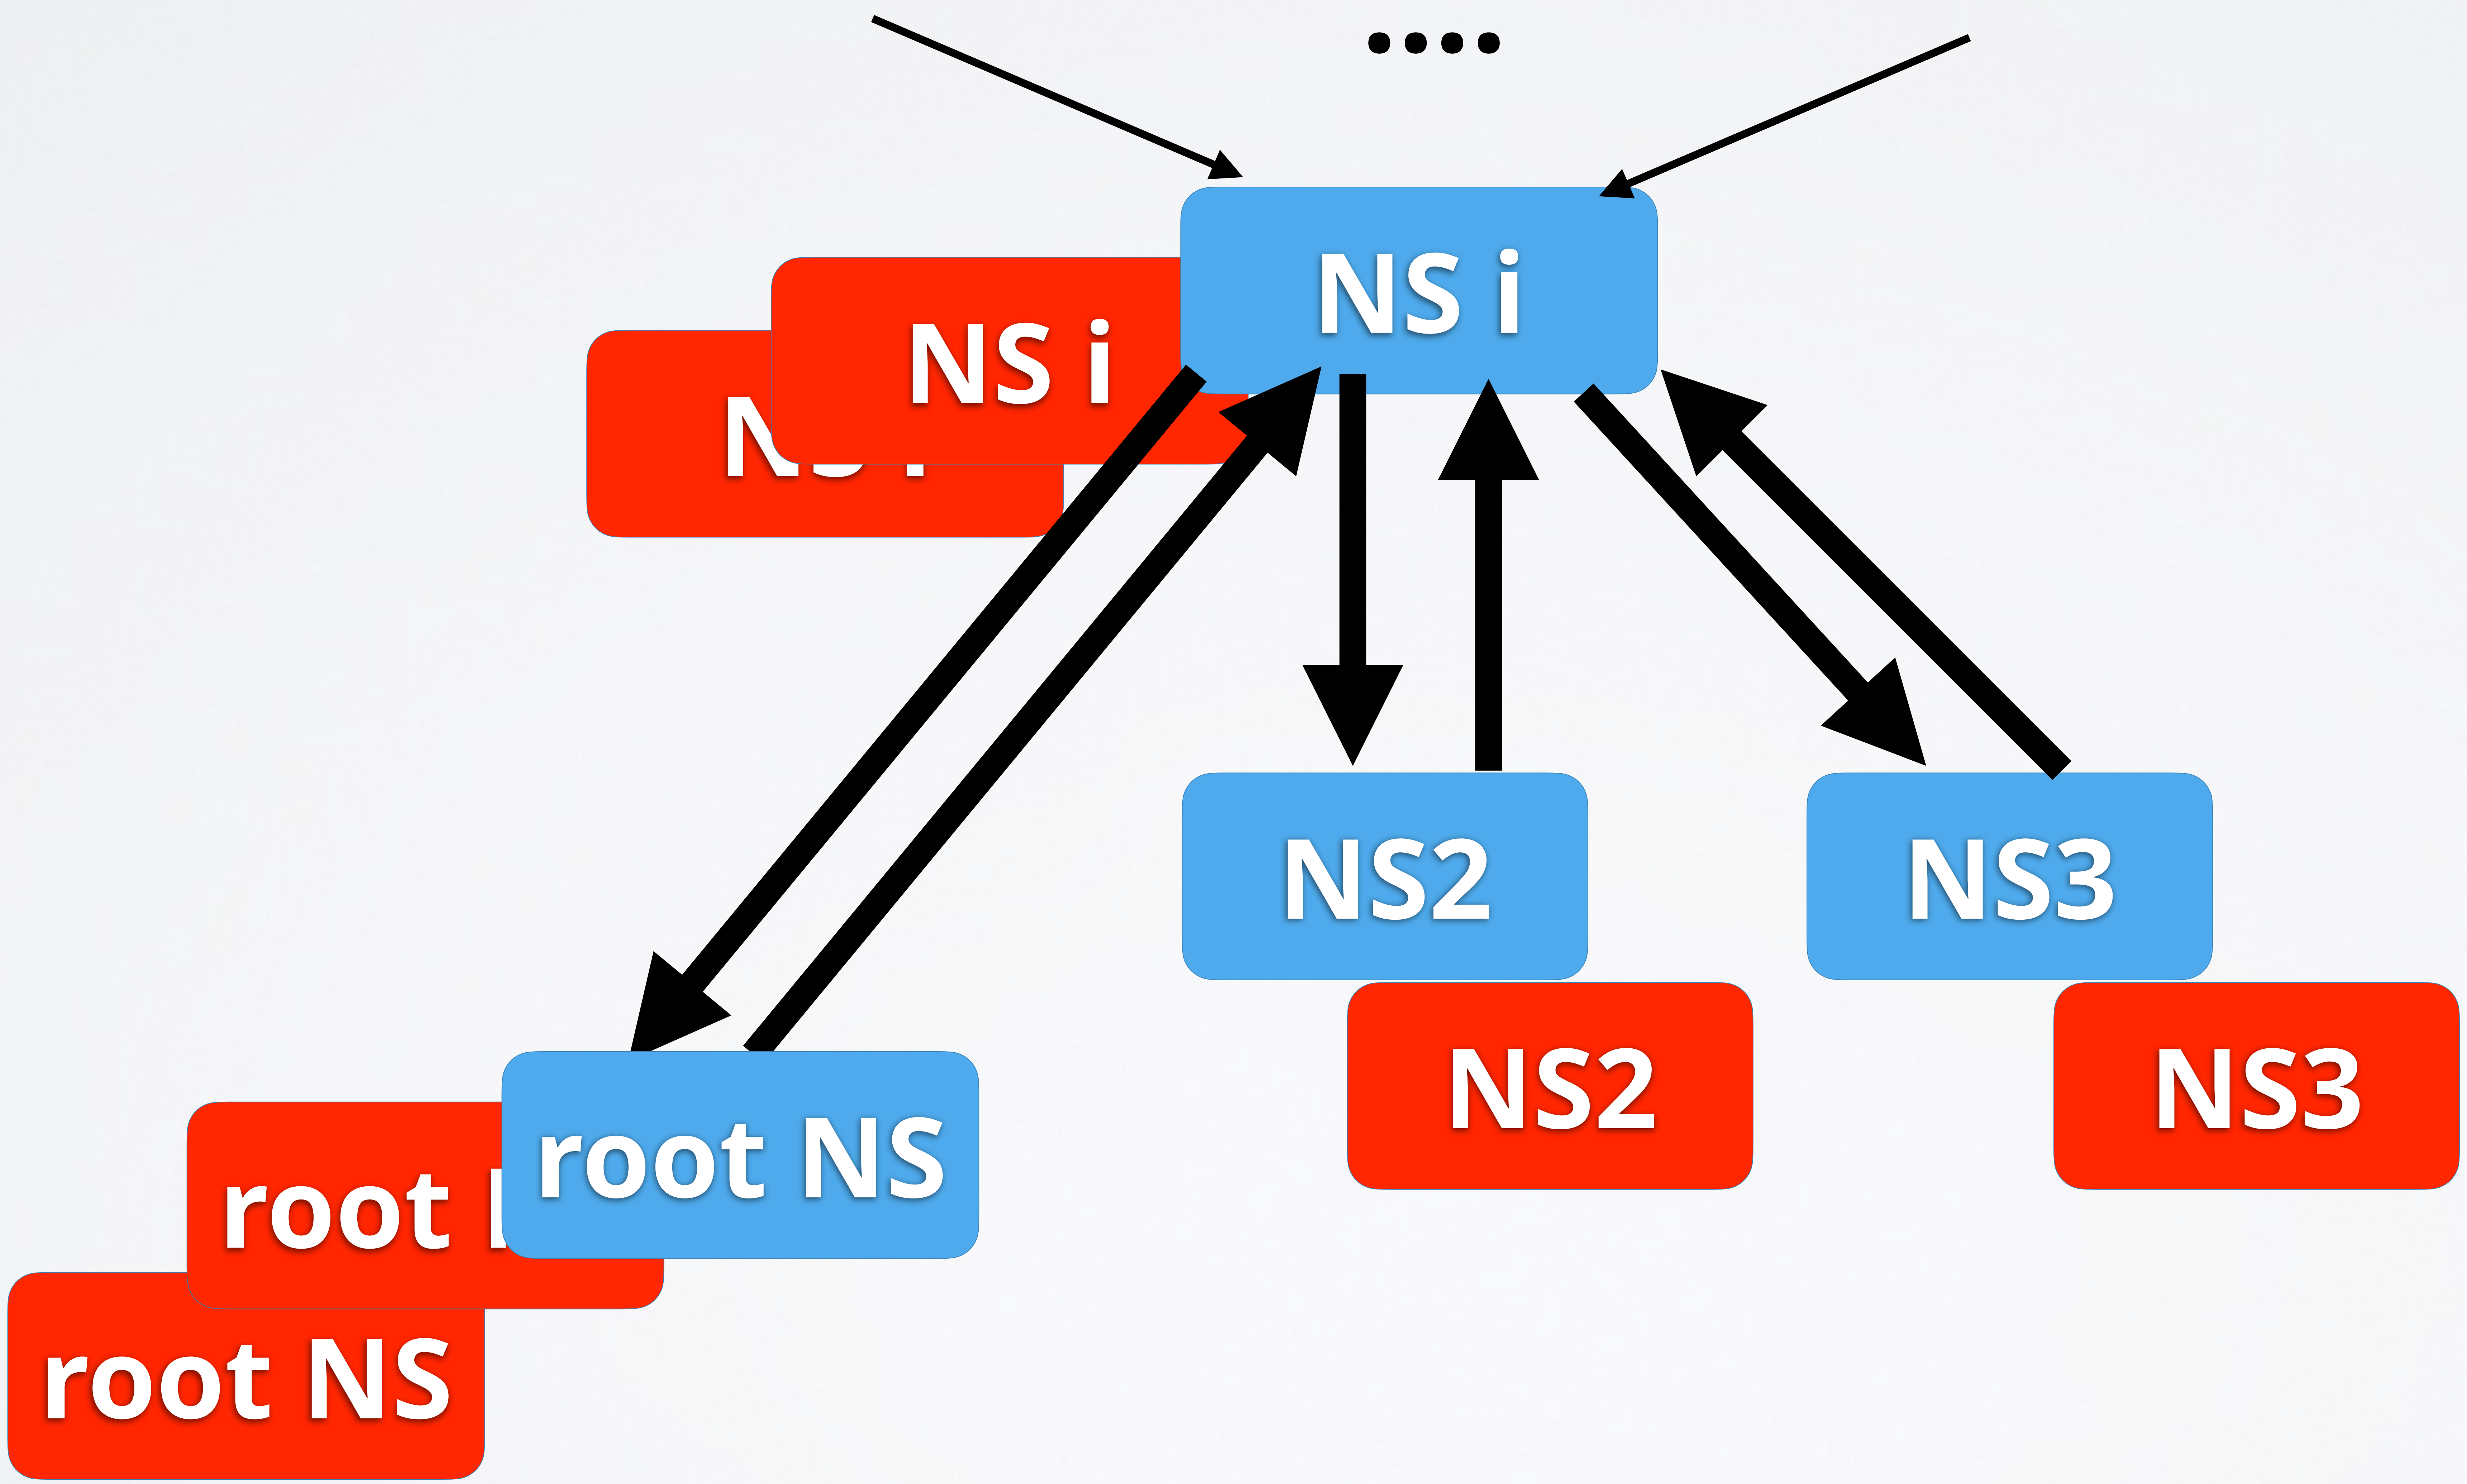

RPC: REPLICATION

- 2 techniques for replication:
  - several IPs/names
  - “any cast" (send packet to one of many servers with same IP)
- 13 root name server IPs, several hundreds of any cast
- each zone has at least
  one primary and one secondary IP

RESOURCE RECORDS

name =>

Record type	Interpretation	Content
A	address	IPv4 address
AAAA	address	IPv6 address
NS	Name server	DNS name
CNAME	Symbolic link	DNS name of canonical name
SOA	Start of authority	Zone-specific properties
PTR	IP reverse pointer	DNS name
HINFO	Host info	Text description of host OS
...	...	...

SUMMARY

- main problems for scalability ….
- simple model: Amdahl’ law)
- few principle approaches …
- DNS as fine example … more examples to come
  study DNS it in your first exercise

- **register in mailing list**

LITERATURE

- Paul Albitz & Cricket Liu
  DNS and BIND
  O´Reilly & Associates, Inc.
- Mark Hill, Michael Marty
  Amdahl's Law in the Multicore Era IEEE
- Couluris, Tollimore, Kindberg
  Distributed systems