

**performing
databases**

Gateway to Hell

-

Networking Basics for DBAs, DEVs and DEVOPs

Martin Klier  and Jan Schampera
Performing Databases GmbH
Mitterteich / Germany

Speaker

- Martin Klier
- Solution Architect and Database Expert
- My focus:
 - Performance + Tuning
 - Highly available systems
 - Cluster and Replication
- Linux since 1997
- Oracle Database since 2003



SYMPOSIUM ^{L2}
Proud Member of symposium42

Speaker

- Jan Schampera
- Specialist for Heterogenous Integration
- My focus:
 - Real-world Database Challenges
 - Scripting and Interfaces
 - All-stack troubleshooting


- Linux since 199x
- Hobby Developer
- Performing Databases' Swiss Army Knife



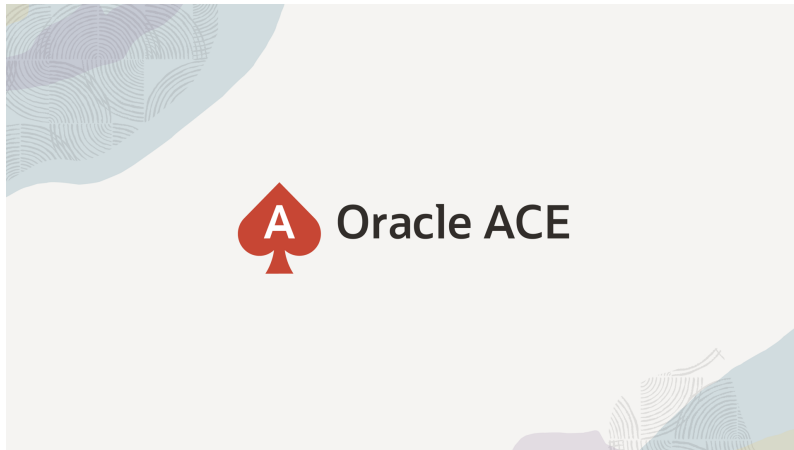
Performing Databases

- Three Experts for Database technology
 - Concepts and Project Competence
 - Architecture- and System planning
 - Licensing
 - Implementation and Troubleshooting
- Get in touch
 - Performing Databases GmbH
Wiesauer Strasse 27
95666 Mitterteich // Germany
 - <http://www.performing-databases.com>
 - Twitter: @PerformingDB





... it's all about Community!



SYMPOSIUM 42

Good communication is the bridge between Confusion and Clarity

~ Nat Turner, Slave Rebel

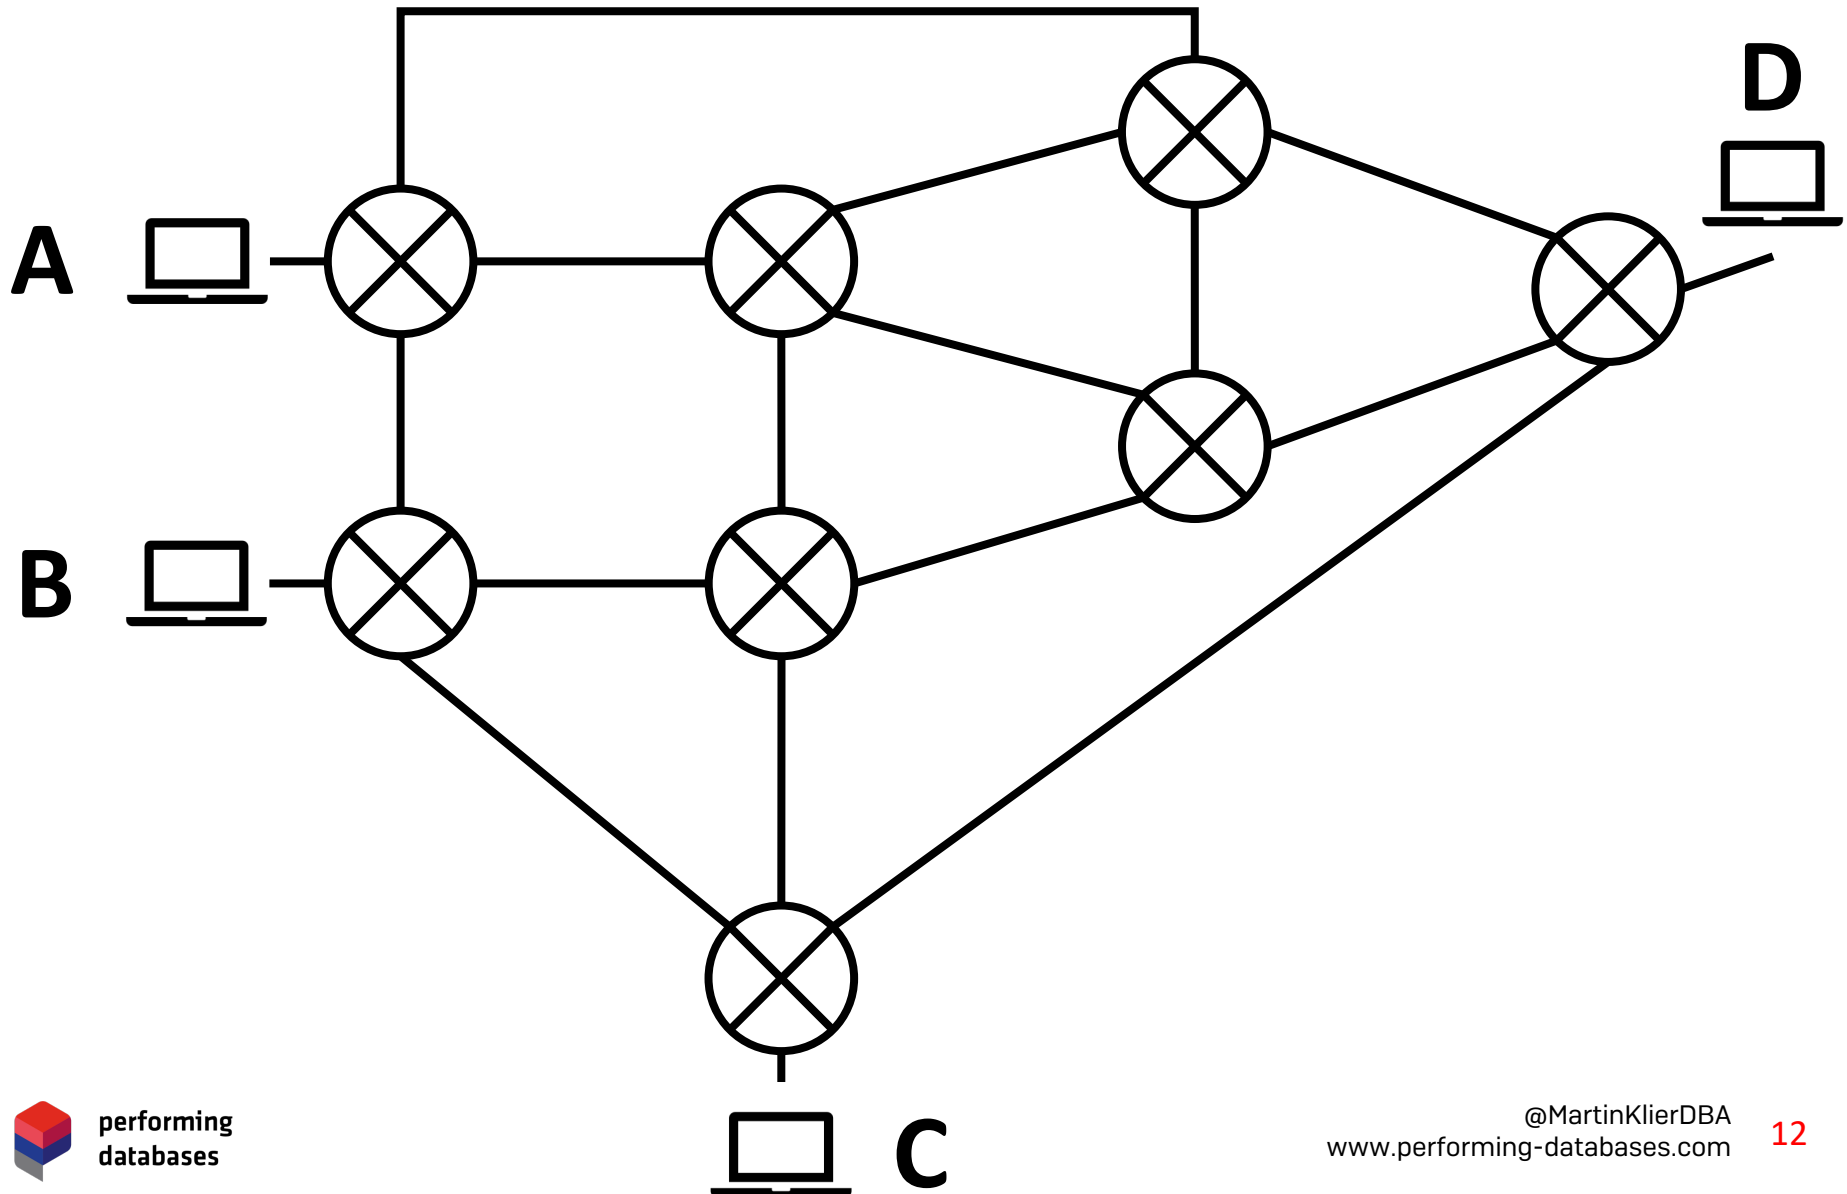
The Idea of an “Internet Protocol” (IP)

SHALL WE PLAY A GAME

IP's Mission: Survive Hell



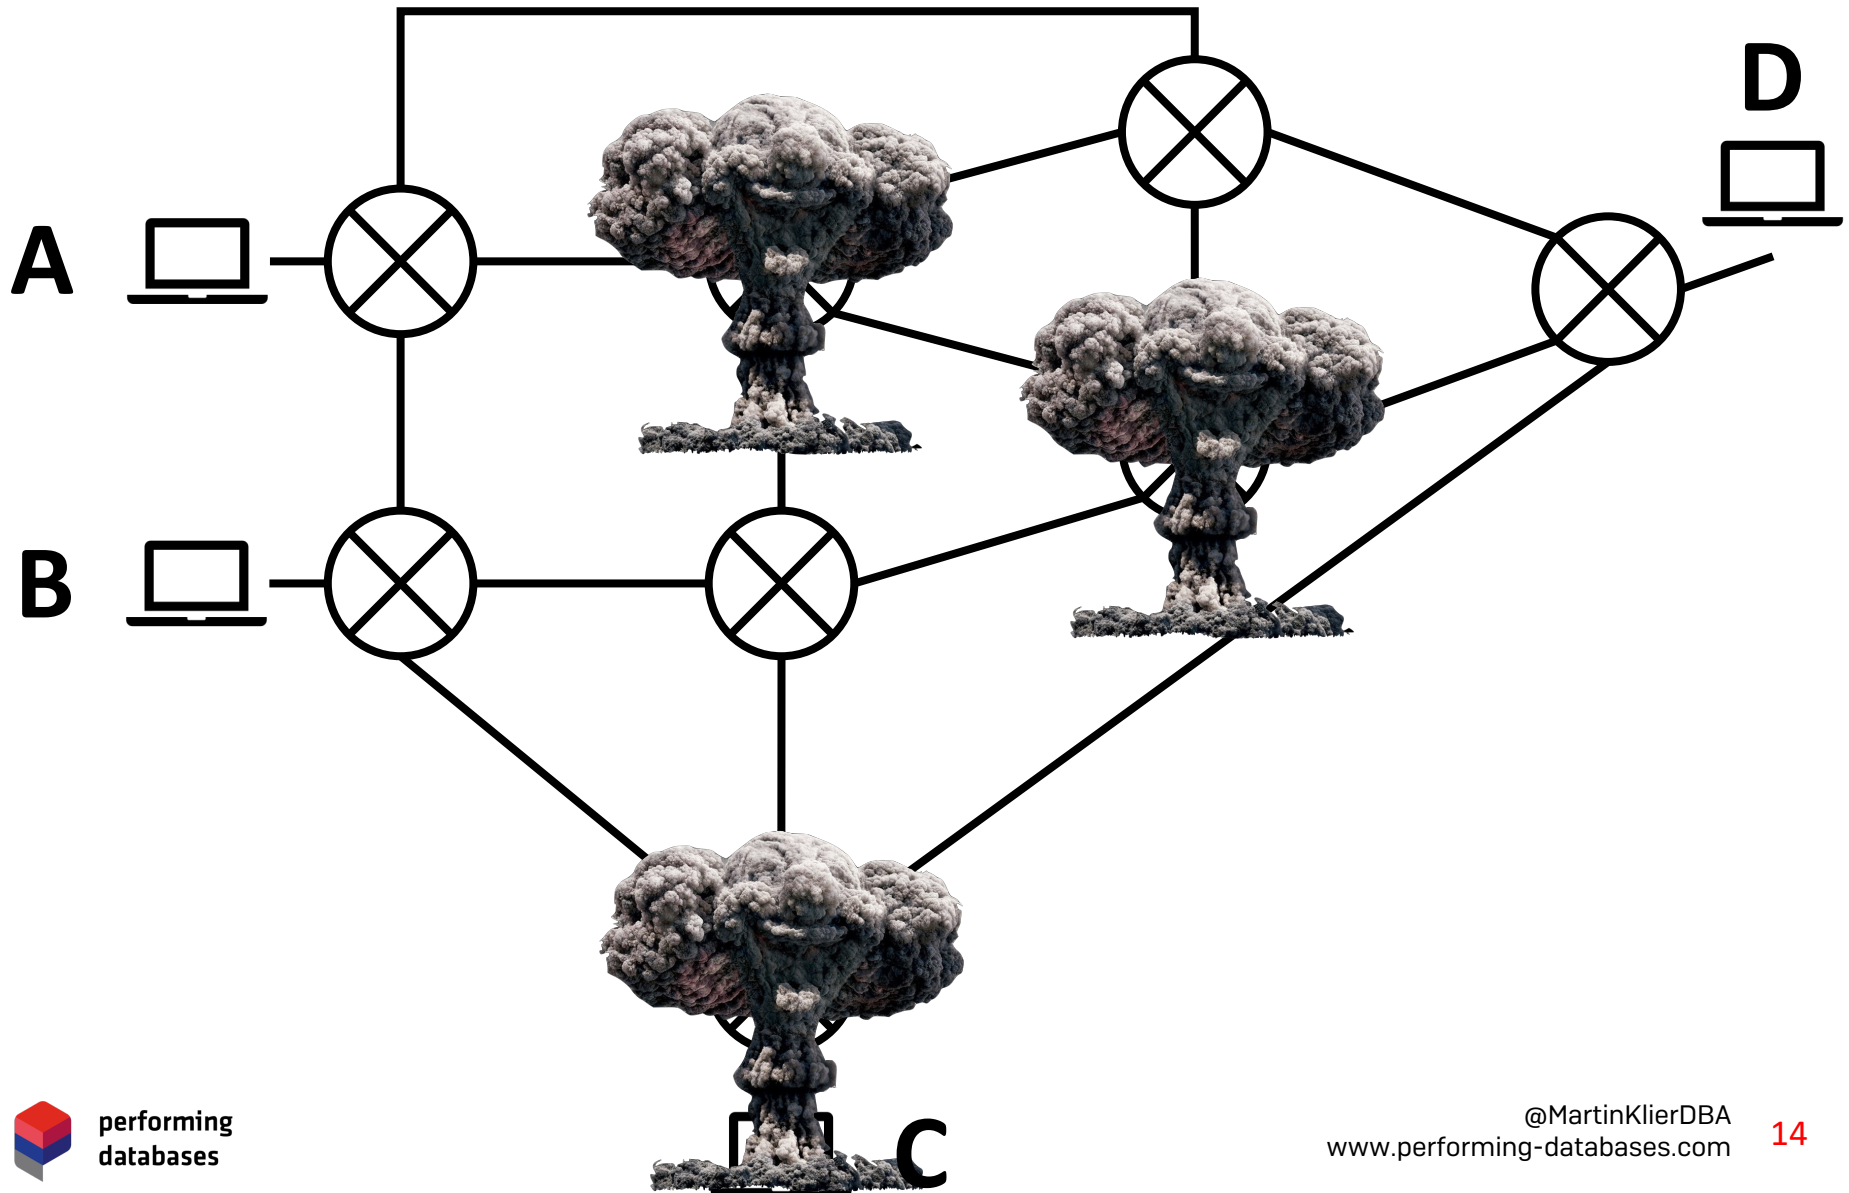
IP's Mission: Survive Hell




IP's Mission: Survive Hell



IP's Mission: Survive Hell




Disclaimer




The only winning move is not to play.

First Things First: Basics

ISO/OSI's 7 Layer Model



Away!
Pizza
Sausage
Throw
Not
Do
Please




IP and all the stuff...

The IP Address

82.128.9.133

082:128:009:133




(IPv4 example, concept in IPv6 is similar)

The IP Network

82.128.0.0/16

082.128.000.000




= the “network address”

The IP Network

82.128.255.255/16

082.128.255.255




= the "broadcast address"

The IP Network Mask

Tell the machine where the
network part ends

255.255.000.000



~~The IP Network Classes~~

(worked without network masks ...)

255.000.000.000	8bit	Class A
255.255.000.000	16bit	Class B
255.255.255.000	24bit	Class C

Not used for Internet Routing any more

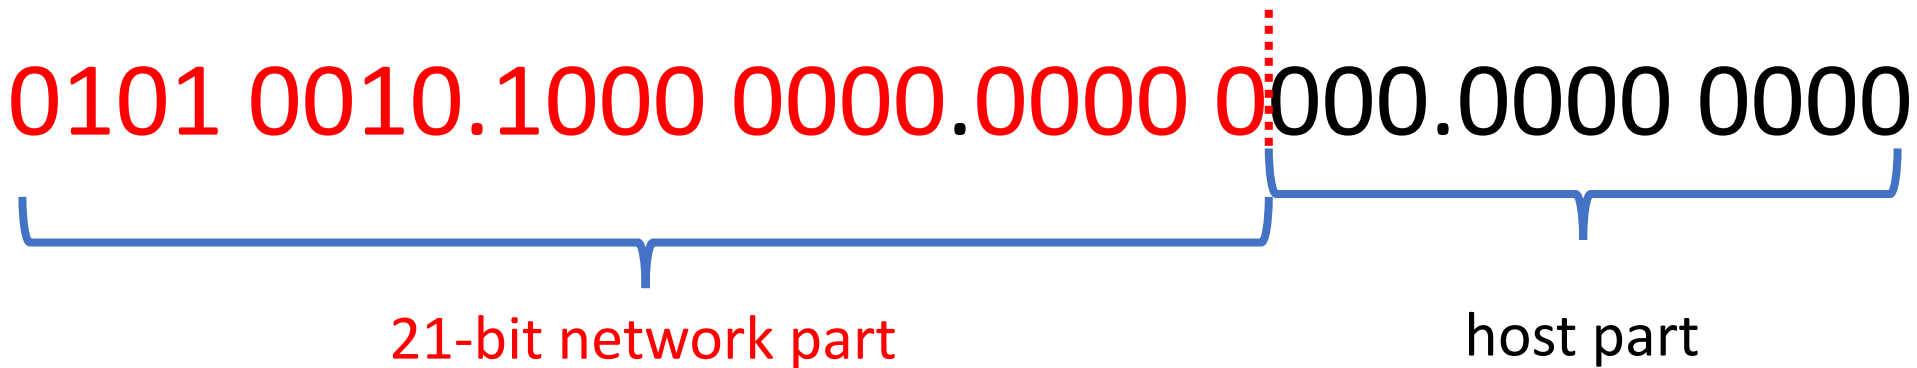
In practice often Octet-Delimiters = Network-Delimiters

We all are lazy. 😊

Classless Inter-Domain Routing CIDR

82.128/21

082.128.000.000



= Network Mask: 255.255.248.0



Sub and Super



Subnetting

082.128.000.000/21

0101 0010.1000 0000.0000 0000.0000 0000

fixed / foreign-owned 21bits

flexible / my own 11bits

0101 0010.1000 0000.0000 0000.0000 0000

my subnet


my hosts

21+4 bits used for the network now:
82.128.0.0/25



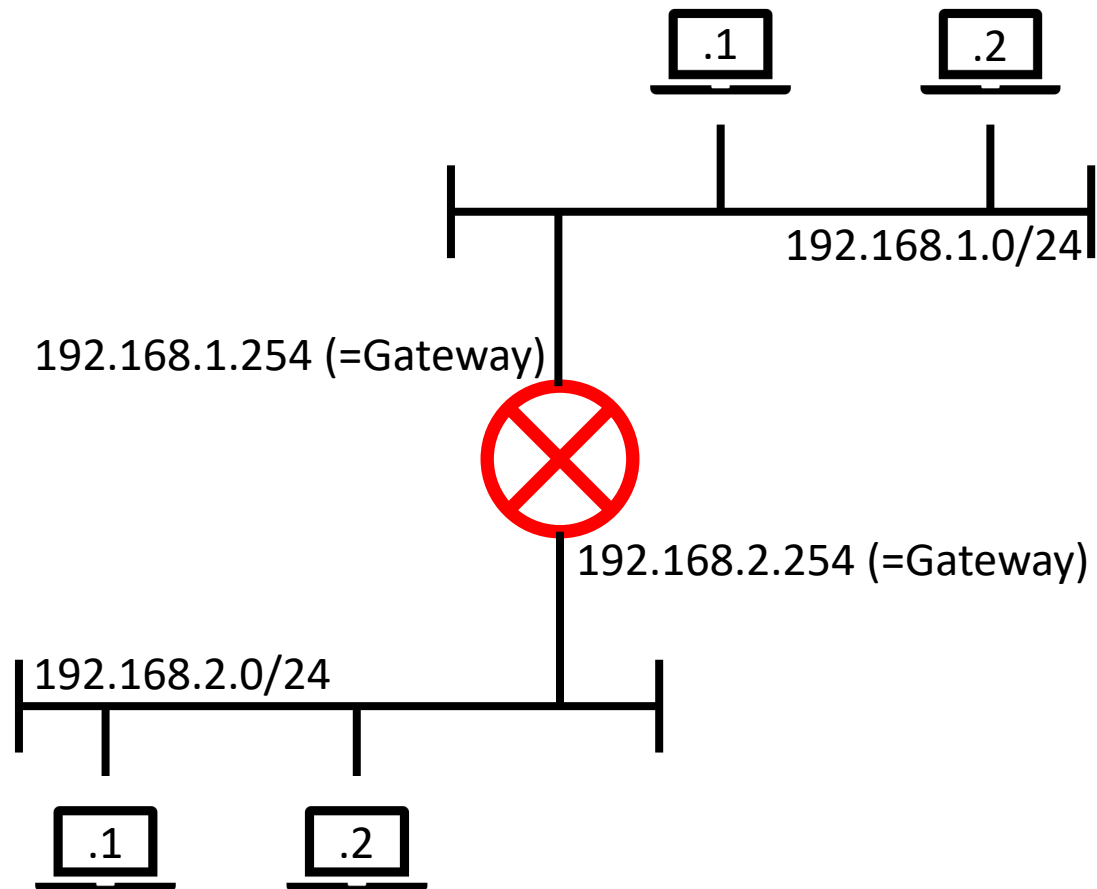
Supernetting

082.128.000.000/21




Routing

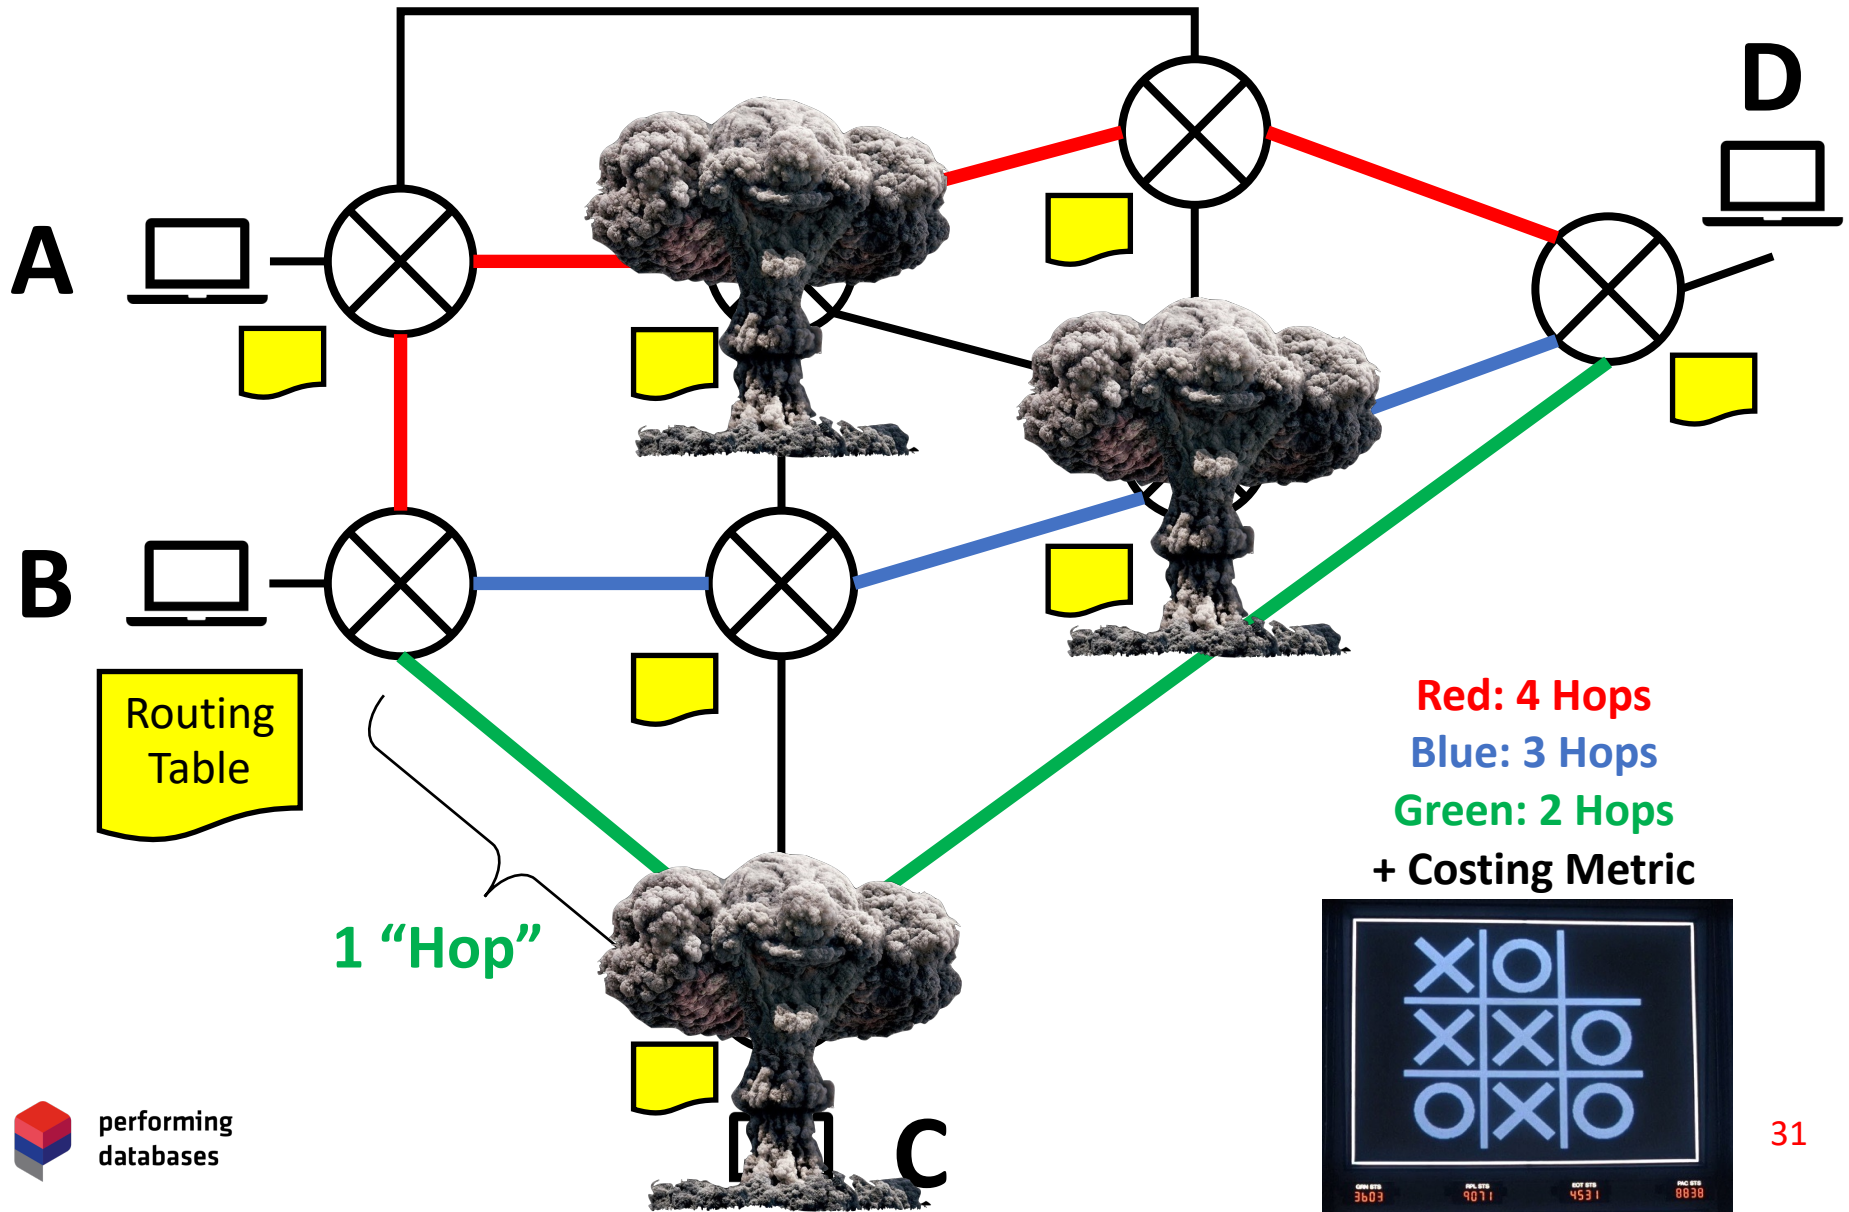
Routing: Connecting Networks



Routing in Layers



IP: Routing Around Hell



Finding Hops: traceroute

```
$ traceroute www.bing.com
traceroute to e86303.dscx.akamaiedge.net (2.21.133.203), 64 hops max, 52 byte packets
 1  router.performing-databases.com (10.11.12.254)  3.372 ms  0.302 ms  0.284 ms
 2  172.16.0.1 (172.16.0.1)  0.995 ms  0.508 ms  0.518 ms
 3  62.156.244.46 (62.156.244.46)  17.570 ms  17.099 ms  16.433 ms
 4  62.156.247.146 (62.156.247.146)  18.723 ms  17.507 ms  17.264 ms
 5  b-eh3-i.b.de.net.dtag.de (62.154.46.218)  18.863 ms  17.948 ms  17.900 ms
 6  80.157.206.70 (80.157.206.70)  178.710 ms  57.278 ms  305.652 ms
 7  a2-21-133-203.deploy.static.akamaitechnologies.com (2.21.133.203)  18.831 ms  17.479 ms  17.002 ms
```

```
$ traceroute www.bing.com
traceroute to e86303.dscx.akamaiedge.net (2.21.133.203)
 1  router.performing-databases.com (10.11.12.254)  3.372 ms  0.302 ms  0.284 ms
 2  172.16.0.1 (172.16.0.1)  0.995 ms  0.508 ms  0.518 ms
 3  62.156.244.46 (62.156.244.46)  17.570 ms  17.099 ms  16.433 ms
 4  62.156.247.146 (62.156.247.146)  18.723 ms  17.507 ms  17.264 ms
 5  b-eh3-i.b.de.net.dtag.de (62.154.46.218)  18.863 ms  17.948 ms  17.900 ms
 6  80.157.206.70 (80.157.206.70)  178.710 ms  57.278 ms  305.652 ms
 7  a2-21-133-203.deploy.static.akamaitechnologies.com (2.21.133.203)  18.831 ms  17.479 ms  17.002 ms
```



Local Addresses

10.0.0.0/8

172.16.0.0/12

192.168.0.0/16

= not routed on the internet

Reserved private IPv4 network ranges^[2]

Name	CIDR block	Address range	Number of addresses	Classful description
24-bit block	10.0.0.0/8	10.0.0.0 – 10.255.255.255	16 777 216	Single Class A.
20-bit block	172.16.0.0/12	172.16.0.0 – 172.31.255.255	1 048 576	Contiguous range of 16 Class B blocks.
16-bit block	192.168.0.0/16	192.168.0.0 – 192.168.255.255	65 536	Contiguous range of 256 Class C blocks.

Socket Communication

The Socket

describes a connection unambiguously

ClientIP:ClientPort – ServerIP:ServerPort

My Macbook connected to an HTTPS Server on the Internet:
192.168.118.120:50479 connected to 149.154.167.41:443


```
$ netstat -tn | grep tcp4
tcp4      0      0 192.168.118.120.50479 149.154.167.41.443  ESTABLISHED
tcp4      0      0 192.168.118.120.50479 149.154.167.41.443  ESTABLISHED
```

Network Address Translation (NAT)

(simplified)

```
$ netstat -tn | grep tcp4
tcp4      0      0 192.168.118.120.50479 149.154.167.41.443  ESTABLISHED
tcp4      0      0 192.168.118.130.50082 149.154.167.41.443  ESTABLISHED
```

But wait – Local Address...?



Good spot for the Firewall!

“The single biggest problem
in communications is the illusion
that it had taken place.”

~ George Bernard Shaw, Pacifist, Nobel Prize- and Oscar Winner

=> Use TCP for peace!

Domain Name System (DNS)

Domain Name

www.oracle.com
herakles.ws27.performing-db.de

hostname sub domain domain top level root

“Full Qualified Domain Name” (FQDN)

Forward Resolution

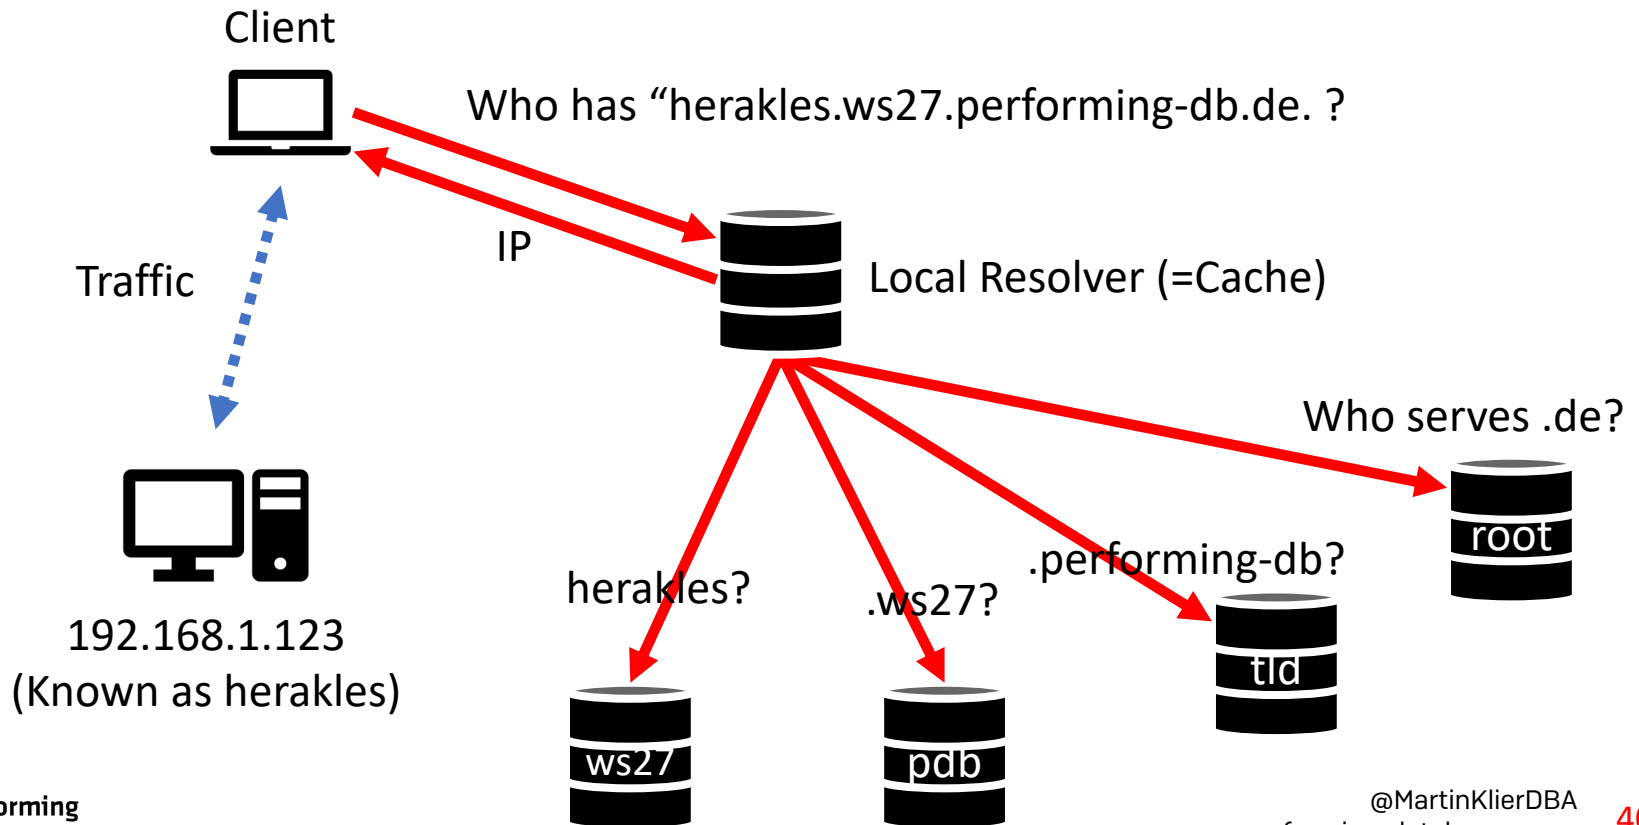
herakles.ws27.performing-db.de.

hostname

sub
domain

domain

top
level



Forward Resolution

```
$ dig www.performing-databases.com @8.8.8.8

; <<>> DiG 9.10.6 <<>> www.performing-databases.com @8.8.8.8
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 27924
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 512
;; QUESTION SECTION:
;www.performing-databases.com.  IN  A

;; ANSWER SECTION:
www.performing-databases.com. 600 IN  A  83.243.40.67

;; Query time: 81 msec
;; SERVER: 8.8.8.8#53(8.8.8.8)
;; WHEN: Wed May 24 22:37:52 CEST 2023
;; MSG SIZE rcvd: 73
```



Oracle SCAN uses Forward Resolution

```
;; QUESTION SECTION:
```

```
;scan.mydomain. IN A
```

```
;; ANSWER SECTION:
```

```
scan.mydomain. 600 IN A 10.0.0.102
```

```
scan.mydomain. 600 IN A 10.0.0.103
```

```
scan.mydomain. 600 IN A 10.0.0.101
```

If done by any TNS EZconnect client, this leads to a tnsnames.ora-like ADDRESS_LIST:

```
(ADDRESS_LIST = (ADDRESS = (Host = 10.0.0.102))  
                (ADDRESS = (Host = 10.0.0.103))  
                (ADDRESS = (Host = 10.0.0.101)))
```

Reverse Resolution

```
$ dig -x 217.237.148.70
```

```
; <<>> DiG 9.10.6 <<>> -x 217.237.148.70
```

```
;; global options: +cmd
```

```
;; Got answer:
```

```
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 29500
```

```
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1
```

```
;; OPT PSEUDOSECTION:
```

```
; EDNS: version: 0, flags:: udp: 512
```

```
;; QUESTION SECTION:
```

```
;70.148.237.217.in-addr.arpa. IN PTR
```

```
;; ANSWER SECTION:
```

```
70.148.237.217.in-addr.arpa. 15337 IN PTR d-lb-a01.isp.t-ipnet.de.
```

```
;; Query time: 311 msec
```


```
;; SERVER: 2a01:598:b037:d768::49#53(2a01:598:b037:d768::49)
```

```
;; WHEN: Wed May 24 23:06:07 CEST 2023
```

```
;; MSG SIZE rcvd: 93
```



The Black Hole



blackhole-1.iana.org (192.175.48.6)
blackhole-2.iana.org (192.175.48.42)
prisoner.iana.org (192.175.48.1)

DNS Reverse Lookup for Local Addresses

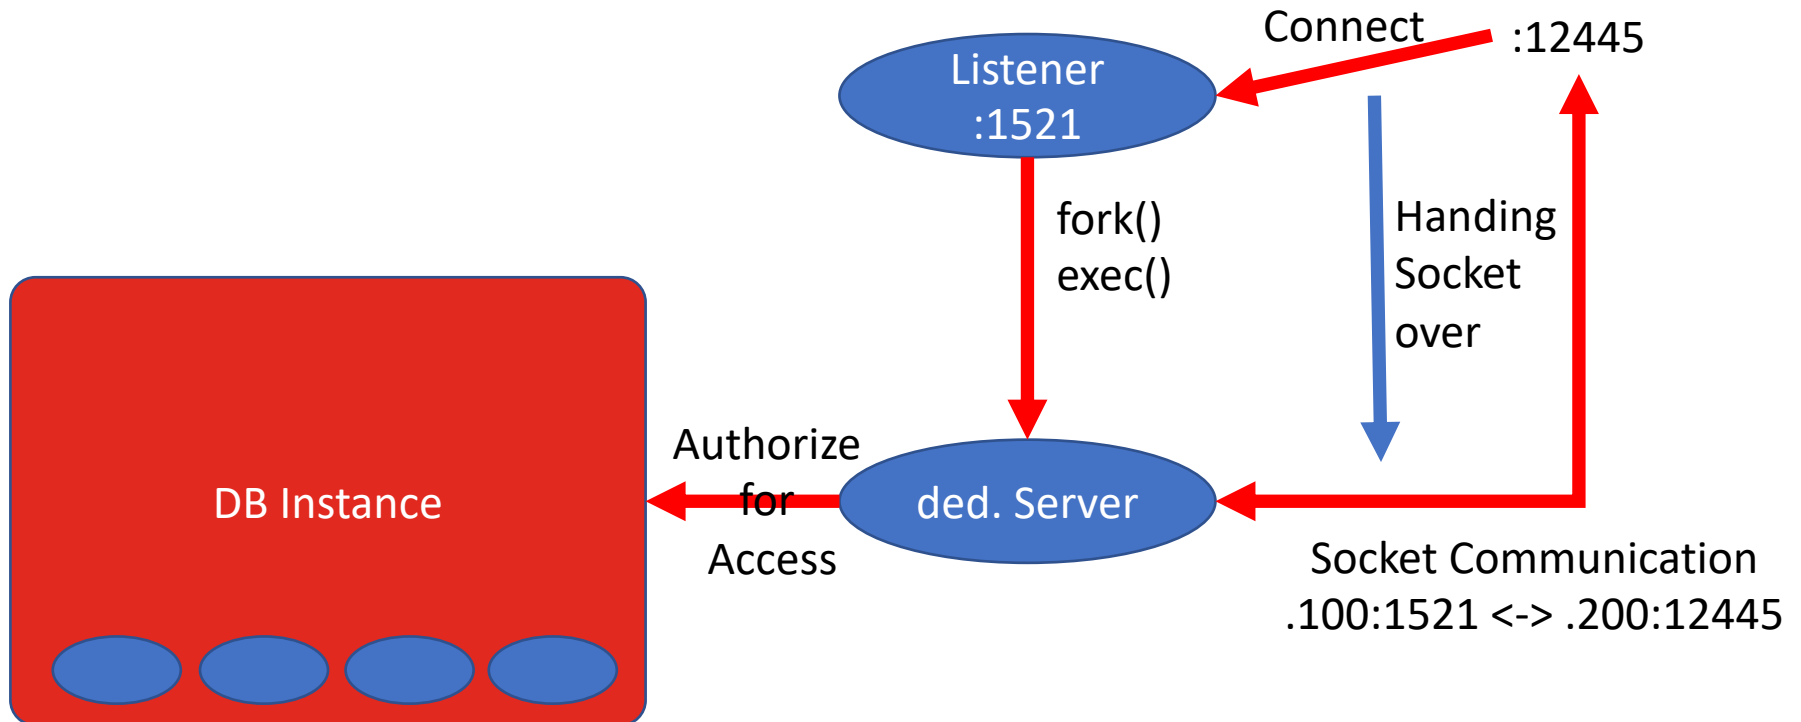
1% of Practice

Connecting via Listener

Database Server
192.168.0.100



Client
192.168.0.200



What we did today...

- 📖 Learned about the idea of IP to route around hell
- 📖 7 Layers of Pain
- 📖 IP Addresses, Networks, Subs and Supers
- 📖 Translating Network Addresses for your safety
- 📖 Domain Names and where to find them
- 📖 The Black Hole of the Internet
- 📖 Establishing Listener Connections
- 📖 DO NOT PLAY!

GREETINGS PROFESSOR FALKEN

HELLO

A STRANGE GAME.

THE ONLY WINNING MOVE IS
NOT TO PLAY.

HOW ABOUT A NICE GAME OF CHESS?



Meet & Greet

martin.klier@performing-db.com

www.performing-databases.com

Many national // international events

The logo for DOAG (Deutsche ORACLE-Anwendergruppe e.V.) consists of the letters 'DOAG' in a bold, red, sans-serif font.

Deutsche ORACLE-Anwendergruppe e.V.

DOAG Database Conference Düsseldorf

May 15-16, 2024

CALL FOR PAPERS OPEN till Feb 7

The logo for utoug (Utah Oracle Users Group) features a stylized blue mountain range icon to the left of the text 'utoug' in a bold, blue, sans-serif font. Below 'utoug' is the text 'utah oracle users group' in a smaller, blue, sans-serif font.

UTOUG Training Days

March 21-22, 2024

Salt Lake City, UT, USA


The logo for aioug (All India Oracle Users Group) features the text 'aioug' in a bold, black, sans-serif font with a green dot above the 'i'. Below 'aioug' is the text 'ALL INDIA ORACLE USERS GROUP' in a smaller, black, sans-serif font. To the right of 'aioug' is the text 'All India Oracle Users Group' in a larger, black, sans-serif font, and below that is the tagline '“Serving the Oracle Community”' in a smaller, italicized, black, sans-serif font.

AIOUG OCYatra Replay Webinar

March 15, 2024

CU at 15:40 back here
“performance is Rarely an Accident”!?





**performing
databases**