

Complementing

~~Challenging~~ BIND and NSD

The World Premiere of a New
Open Source DNS Server



Peter Janssen, EURid.eu
Hamburg, Domain Pulse, Feb. 14 2012



Who is EURid?

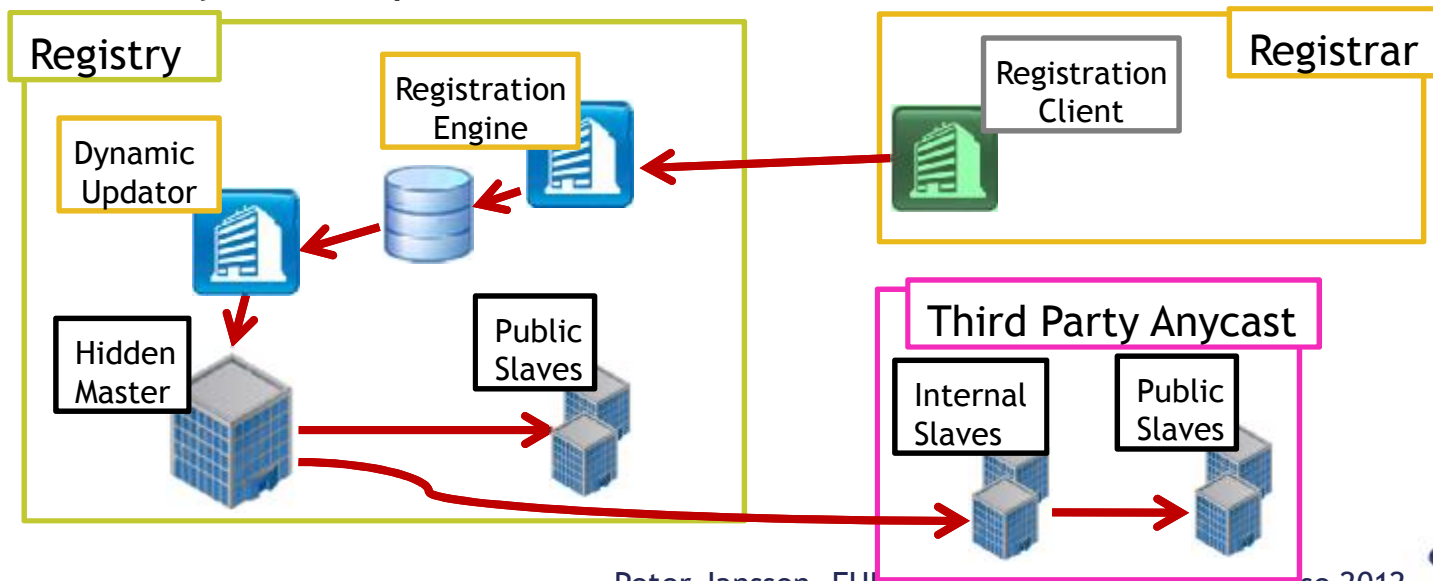
- Belgian not for profit organisation
- Under contract with the European Commission
- Offices in Brussels, Prague, Stockholm and Pisa
- Manage .eu TLD space
 - Registration services
 - “Run” the .eu zone
 - Support services (phone, email, chat, ...) in all 23 official languages of the EU
- Domain names
 - First Come, First Serve,
 - IDN, MYR, DNSSEC
 - Individual or organisation residing in the European Union
 - Currently over 3.5 million names

The .eu zone

- Name servers
 - “Own maintained” unicast nodes
 - “Third party” anycast meshes
- Update mechanism (Registration -> DNS)
 - Dynamic updates

The .eu zone

- Name servers
 - “Own maintained” unicast nodes
 - “Third party” anycast meshes
- Update mechanism (Registration -> DNS)
 - Dynamic updates



“Challenges”

- TLD public authoritative slaves
 - Geographical spread
 - Network topological spread
 - Hardware diversification
 - OS diversification
 - Name server software diversification
- Zone Update mechanism
 - “Listening” on core database
 - Sending “dynamic update” messages to name server
 - Room for improvement/de-complexification

A new DNS implementation

- Primary Design goals
 - “drop-in” replacement for BIND and NSD
 - Standards compliant
 - Only “RFCs” as design goals
 - Common sense/practice where RFCs are “open for interpretation”
 - Portable, clean, fast
 - Authoritative, DNSSEC support
 - AXFR/IXFR support (master and slave)
- Secondary goals
 - Recursive caching resolver?
 - Dynamic update API
 - Dynamic provisioning
 - “higher level storage” backend (sql db, ...)


Release 1.0 (Roadmap)

- Supported platform : Linux, *BSD, OS X
- Authoritative
- Load/parse zone files (BIND style files)
 - Include, *, @, /
 - Resource record types
 - SOA, A, AAAA, NS, CNAME, PTR, DNAME, HINFO, MX, TXT
 - SPF, SRV, NAPTR
 - DNSKEY, DS, RRSIG, NSEC, NSEC3, NSECPARAM
- Zone transfer
 - Master & Slave, AXFR / IXFR
 - Notify, TSIG
- Nsupdate (add, remove RR)
- DNSSEC
 - RSASHA1(5,7)
 - Online re-signing
 - *key roll-over (manual)*

Release x.y (Roadmap)

- Release 1.1
 - DSA/SHA1, DSA-NSEC3-SHA1, RSASHA256, RSASHA512
- Release 1.2
 - Full client to “control” the name server daemon (1.0 does stop and reload)
 - Dynamic zone management (add/remove zones on the fly)
- Release 2.0
 - Caching resolver
 - Validating
 - Stand alone libraries; perl wrapper around dnsresolv
- Release 3.0
 - Sql backend API (EURid will use this with its dynamic updator)
 - ...

Roadmap vs. current state



- General “code clean up”
- Extensive testing + (Bug fixing ⇔ Documentation)
- Documentation
 - RFC’s that are “implemented”
 - Reference manual
 - Man pages

Performance measurements

- The setup

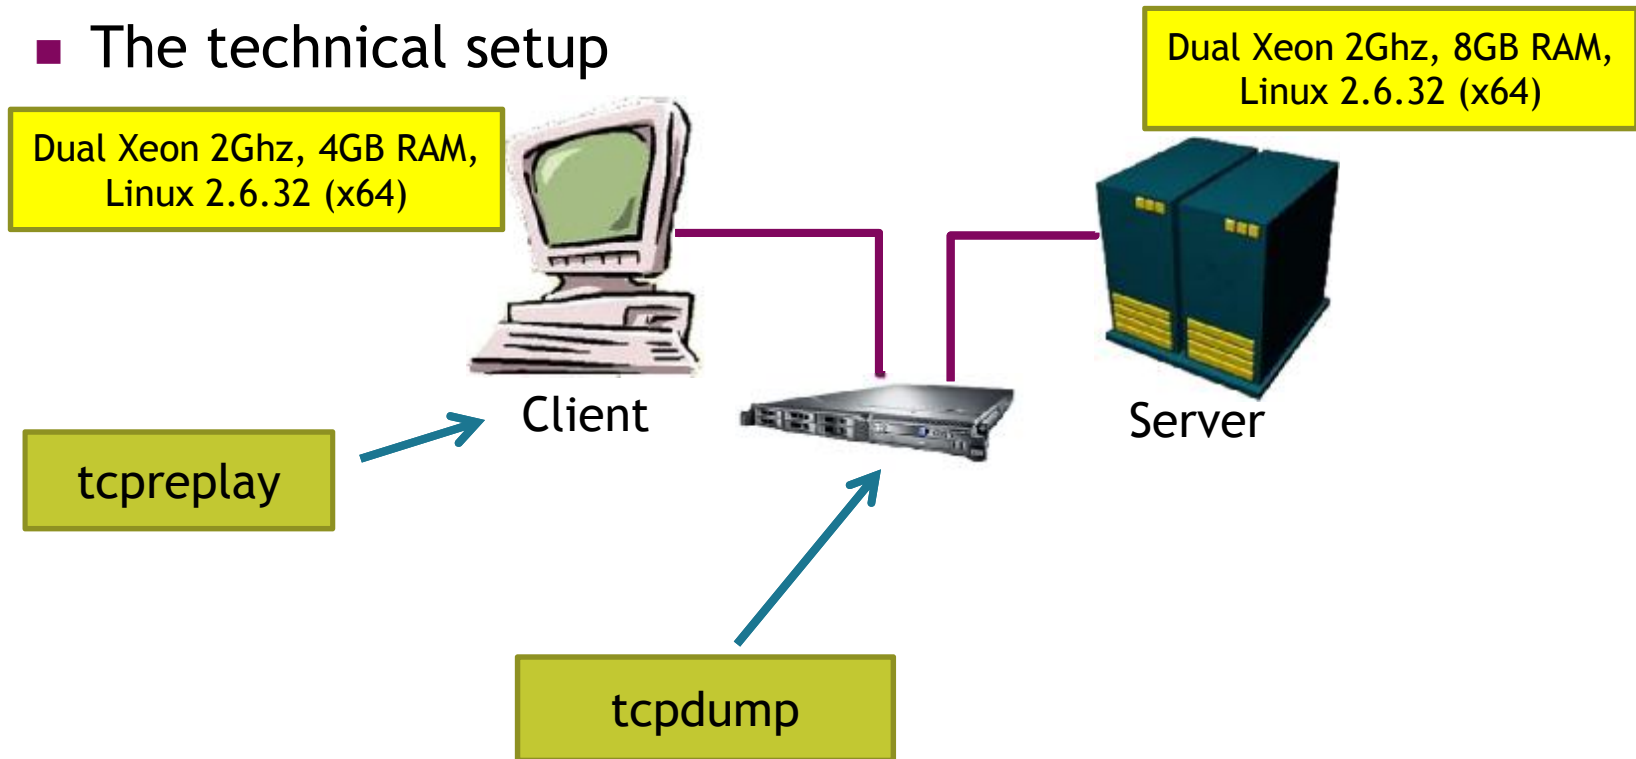


www.somedomain.eu A (no DNSSEC)
Random query list
50% existing domain names

.eu zone (3.3M domain names,
7.7M records)
NSEC3 signed

Performance measurements

■ The technical setup



Performance measurements

■ The setup

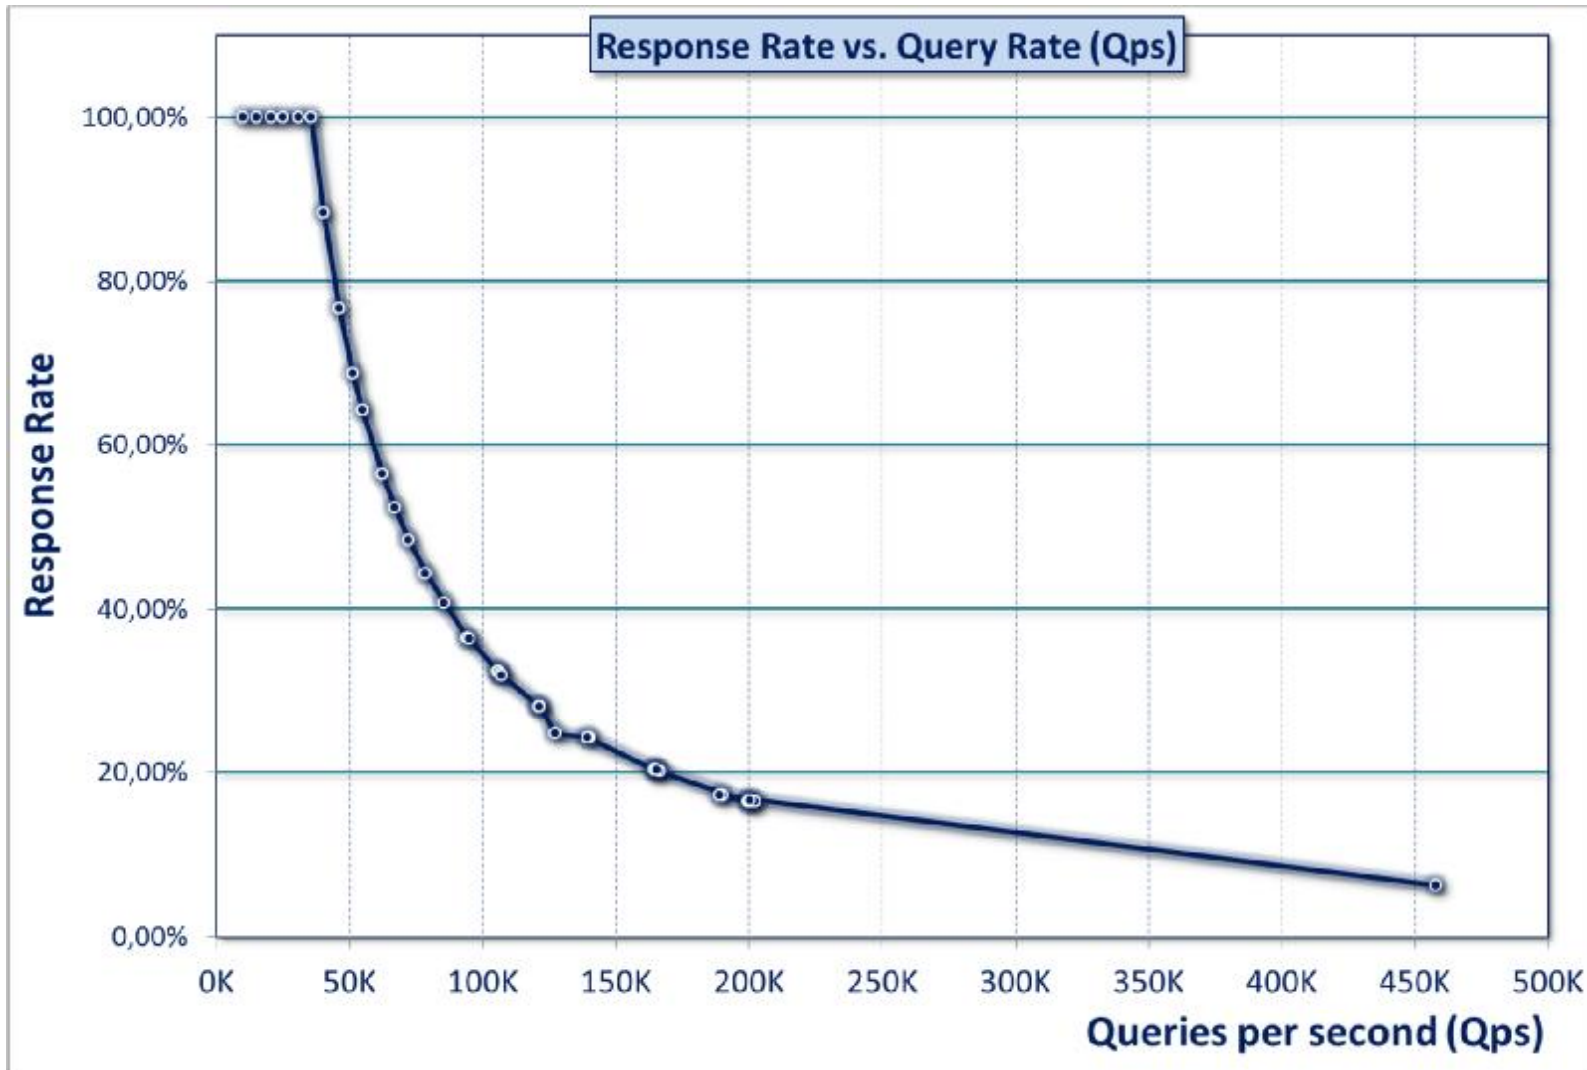
- Pre-created PCAP file
- Tcpreplay -p *speed* 'pcapfile' : to generate queries
- Tcpcap : to capture responses
- Runs of 30 seconds

■ Eg. 50Kq/s

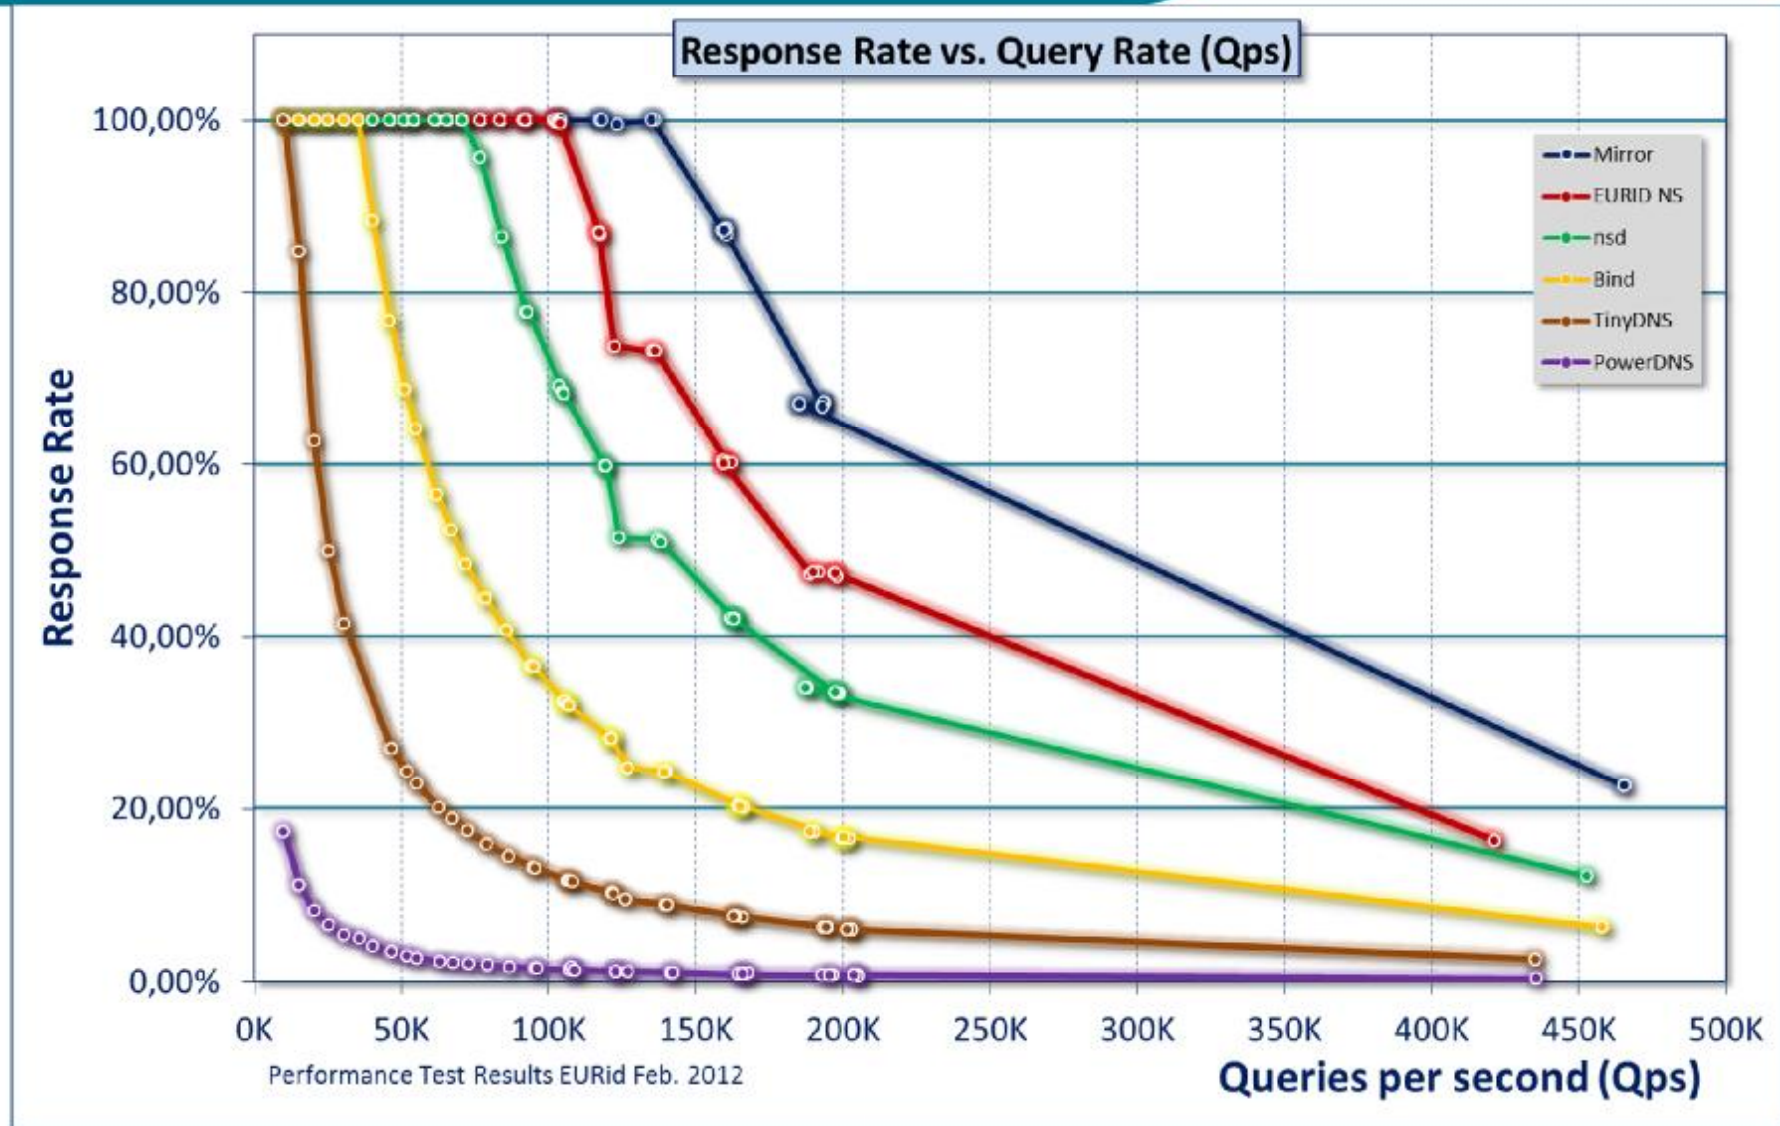
- $50\text{Kq/s} * 30\text{s} = 1.500.000$ packets -> pcap file
- Tcpreplay -p 50000
 - 29.7 seconds -> 50.505 q/s
- Tcpcap results
 - How many packets received

q/s	Real q/s	Sent	Received	%
50K	50.50K	1.5M	1.5M	100
75K	74.47K	2.25M	2.2M	97.7
100K	101.2K	3M	2.8M	93.3

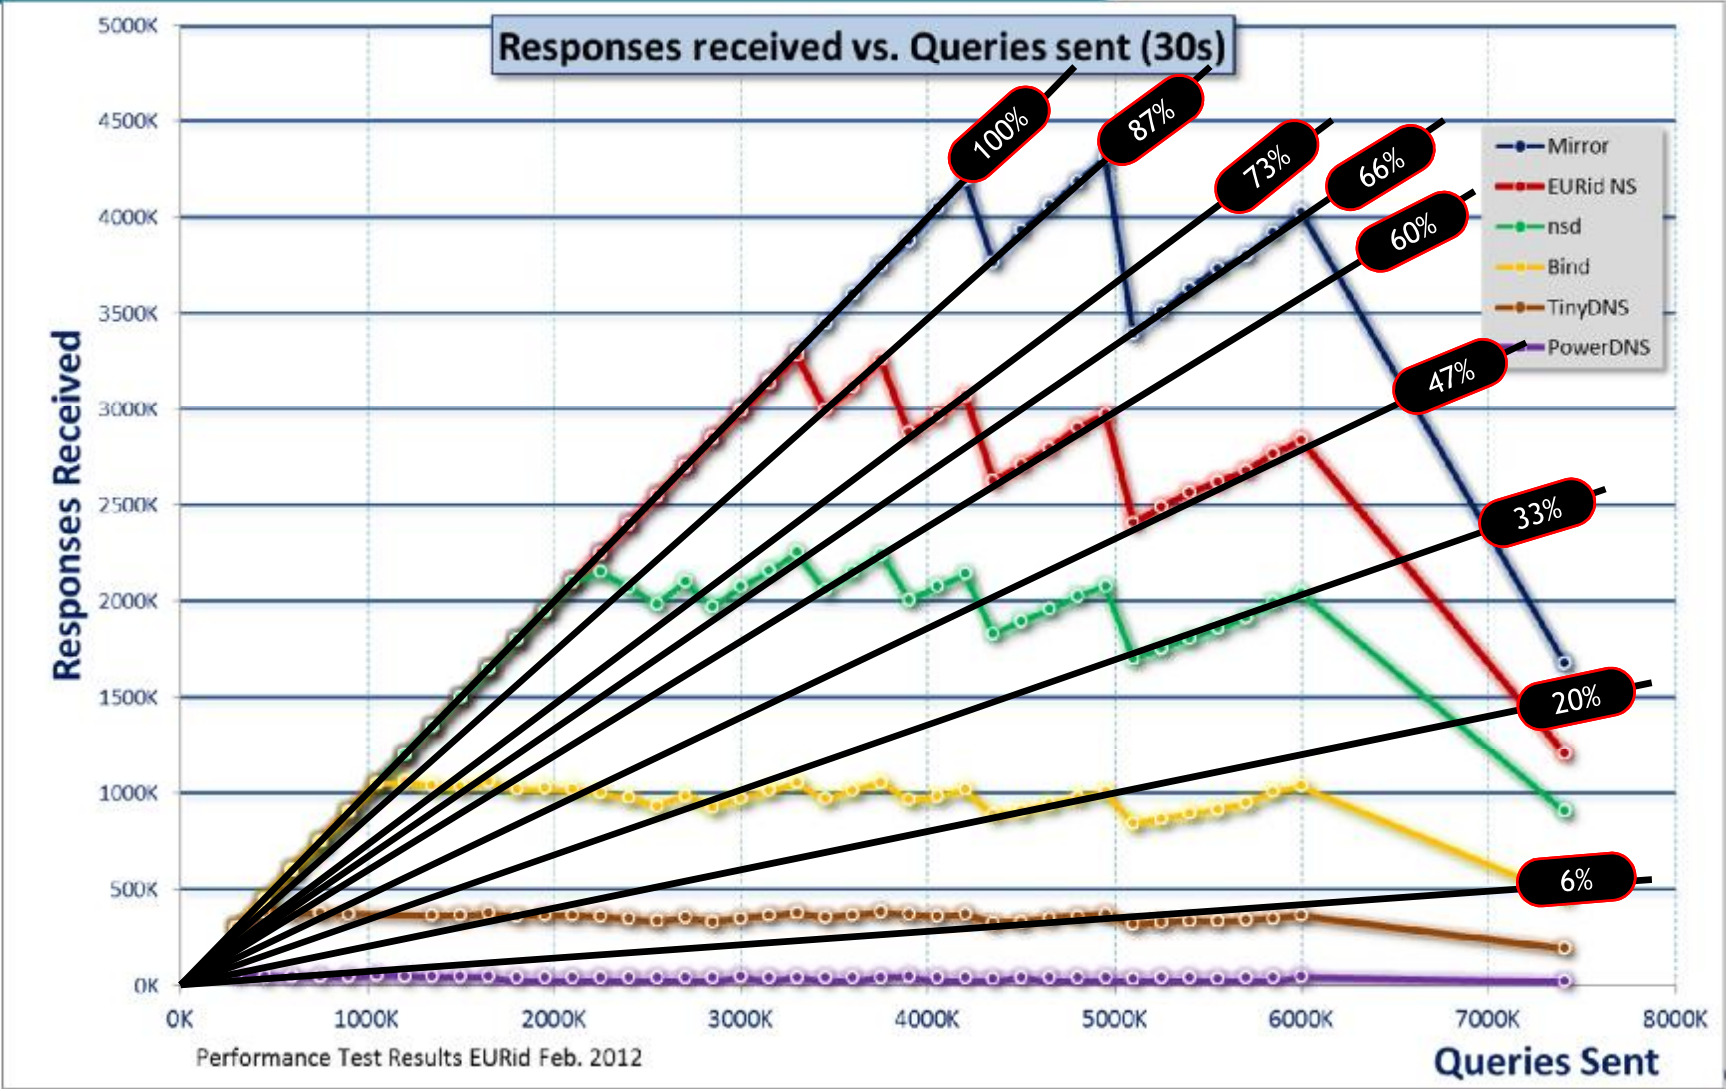
Performance measurements



Performance measurements



Performance measurements



Performance measurements

- .com zone file (198 million lines)

	Prepare	Load	Ready
BIND 9		38m40s	38m40s
NSD	36m58s	12m03s	40m01s
EURid NS 0.8		8m26s	8m26s

- 100.000 zones (7 RR)

	Prepare	Load	Ready
BIND 9		37s	37s
NSD	8s	2s	10s
EURid NS 0.8			4s

(Dual Xeon 2.1Ghz, 48Gb, Linux Debian)

Code & stats

- Written in C (clean room implementation)
- Libraries
 - 30K lines code + 5.5K lines header
- Main program
 - 8.9K lines code + 1.1K lines header
- Dev effort
 - 2.6 man year

Availability

- Release 1.0 (February 2012)
 - Binary packages
 - CentOS 32 and 64 bit
 - Debian 32 and 64 bit
 - freeBSD 64 bit
 - osX(Lion) 64 bit

- Release 1.2 (June 2012)
 - “BSD” open source license

One more thing

- Does it have a name?
- All the good ones were taken....
- Aber sicher, es hat einen Namen....

Does it have a name?

Y:A:D:I:F:A

One slide to say it all.....



<http://www.yadifa.eu>

info@yadifa.eu

Mailinglists :

- yadifa-announce
- yadifa-users

yadifa.eu.

NS ns.yadifa.eu.

NS yadifa.eurid.eu.

yadifa.eu DNS is being served by YADIFA!