

# PowerDNS.COM BV

All about the PowerDNS nameserver  
and how you can use it.

bert hubert

[bert@powerdns.com](mailto:bert@powerdns.com)

<http://wiki.powerdns.com> - <http://www.powerdns.com>

<http://netherlabs.nl>

# What is PowerDNS?

- PowerDNS is a company
- PowerDNS is also the name of one of our products
- PowerDNS is a full-featured authoritative nameserver that can read from most databases, BIND zone files, LDAP but can also be scripted
- PowerDNS is a high-performance recursor

# PowerDNS Company

- Around since 2000
- Co-founded by V3 Redirection Services, well known from sites such as *come.to*, *go.to*, *start.at* and other short domains.
- Started [express.powerdns.com](http://express.powerdns.com), a hosted domain service.
- Does installations, support, custom enhancements (public, private)

# PowerDNS Software

- Open source since 2001 (GPL)
- Started out as a means to do geographic loadbalancing between servers
- Goals:
  - Secure (37,000 lines of code, bind 350,000)
  - Fast
  - Modular
- Linux, FreeBSD, NetBSD, OpenBSD, Solaris. Windows experimental.

# Some interesting users

- Top-level domains .MN, .MP (and .TK)
- All Wikipedia sites (Global Load Balancing)
- Schlund + Partner
- Tucows
- E164.org (ENUM)
- Siemens (ENUM trial)
- Register.com (visa.com, burgerking.com bayer-ag.com)
- Ascio (Speednames)
- ~~... about 8 million domains known~~

# PowerDNS features

- Implements RFC1034/1035 DNS protocol
- All common, and some uncommon, record types (including SRV, NAPTR, SPF)
- Full master / slave support - or use DB replication
- Serves data from SQL, LDAP, BIND &c
- Whole-packet caching for high performance
- Large amount of statistics

# PowerDNS features

- Can serve DNS from:
  - MySQL, PostgreSQL, SQL Server, Sybase, Oracle, DB2, ODBC...
  - SQLite
  - LDAP
  - Geographical information system
  - Any external program
  - Bind zone files!

# SQL Features

- Defaults to a simple schema:
- create table domains (  
id NUMBER,  
name VARCHAR(255) NOT NULL,  
master VARCHAR(20) DEFAULT NULL,  
last\_check INT DEFAULT NULL,  
type VARCHAR(6) NOT NULL,  
notified\_serial INT DEFAULT NULL,  
account VARCHAR(40) DEFAULT NULL,  
);



# SQL Schema

- CREATE TABLE records (  
    id                    number(11) not NULL,  
    domain\_id            INT DEFAULT NULL REFERENCES  
Domains(ID) ON DELETE CASCADE,  
    name                  VARCHAR(255) DEFAULT NULL,  
    type                  VARCHAR(6) DEFAULT NULL,  
    content              VARCHAR(255) DEFAULT NULL,  
    ttl                  INT DEFAULT NULL,  
    prio                  INT DEFAULT NULL,  
    change\_date          INT DEFAULT NULL,  
);

# SQL Queries can be modified

- Five basic queries:
  - basic-query
  - id-query
  - any-query
  - any-id-query
  - list-query
- Can all be modified for your own schema:
  - `gpysql-basic-query=select content,ttl,prio,type,domain_id,name from records where type='%s' and name='%s'`

# Other backends

- Pipebackend
  - A separate program is launched that receives queries on standard input and sends answers on standard output
  - Sample perl script provided
- geobackend - works with standard maps
- Randombackend
  - A new IP address for each question!

# Special features

- automatic provisioning:
  - An authorized IP address sends a notification for an unknown domain
  - PowerDNS checks if that domain is indeed present on that server, and that it lists itself as a nameserver
  - If so, the domain is added as a slave domain automatically
- Built-in webserver for operational insights

# BIND compatibility mode

- PowerDNS can read almost all named.conf files directly
- Reads and writes zonefiles as well
- Additionally, can read gzipped zones
- Loads faster
- Works faster
- Uses less memory
- Just as cheap though!

# Multiple modules

- Multiple backends can be active at the same time
- Each backend can be loaded multiple times, for example to connect to several databases
- Queried in order: first one that answers
- geobackend can 'overlay' a single record
- Same for pipebackend

# Things PowerDNS doesn't do

- DNSSEC
  - Too complicated in its current form. Perhaps TSIG soon. Dynamic TSIG?
- Views
  - Complicate a nameserver a lot
- Send out optional authority records
  - I think there is still 1 registry with problems
- Dynamic updates over DNS
  - Security risk - update your database!

# Safe migration

- The nameserver is the most vital part of operations: migration is scary
- Three tools:
  - dnswasher
    - anonymise tcpdumps for analysis
  - dnsreplay
    - replay recorded DNS traffic and compare
  - dnsscope
    - get statistics: latencies, errors, drops



# How to test

- Record existing DNS traffic using tcpdump
- Install PowerDNS on a separate server
- Migrate BIND zonefiles using zone2sql:
  - `zone2sql --named-conf=/etc/bind/named.conf | ..`
- Start PowerDNS, stop tcpdump
- Start dnsreplay:
  - `dnsreplay new-ip-address new-port < dumpfile`
- Will display statistics

# PowerDNS recursor

- Resists Denial of Service attacks
- Remembers broken nameservers
- Remembers broken domains
- Very fast - can handle traffic of 300.000 ADSL users on laptop over WiFi and DSL with 30% cpu load.
- More standards compliant than dnscache

# So.. Why PowerDNS?

- No embarrassing startup downtime
  - even the bind backend can serve while (re) loading!
- Most large ISPs will have DNS data in a database anyhow. Easier to update, easier to manage.
- Diversity - it makes sense not to rely on one program
- Probably less security risk (or use nsd)

# Future

- TSIG
  - dynamic TSIG?
- Stateful zone (de)provisioning
  - push zones to slaves
  - remove zones from slaves
- Improved Windows support
- DNS-based failover

# Support by PowerDNS.COM BV

- Installation services
- Private support
  - 8 hours, 24 hours, 48 hours response time
  - 4 hour phone support for large contracts
- Private new features (backends, management etc)
- Sponsor new features (half-rate)
  - xs4all, register.com

# Community Support

- [Wiki.powerdns.com](http://Wiki.powerdns.com)
  - user editable FAQ
  - pointers to other PowerDNS pages
    - LDAP has a separate page for example
- Mailing lists
  - Over 500 members
  - Questions often answered by other PowerDNS users
  - We answer public questions there as well