BIND-DLZ

Presentation by: Lars-Göran Forsberg

LOOPIA

Who am i?

- Registrar Manager, Loopia
- DNS Service Development
- Registry/Registrar contacts

LOOPIA

Outline

- Short intro to BIND-DLZ
- Background of my experience
- Is BIND-DLZ up for consideration
- Production stable?
- Pros & Cons
- Its all about the system design
- Trial And Error
- Example, LoopiaDNS
- Summary
- Questions

LOOPIA

Intro to BIND-DLZ

- An extention to BIND, an zone storage backend
- Separate in 9.3, included in BIND release from 9.4
- Dynamically Loadable Zones
  - INSERT, MODIFY or DELETE records in a zone without reloading/restarting
- Can use a wide selections of databases as data source
  - Berkley DB
  - MySQL
  - PostgreSQL
  - LDAP
  - ODBC (!?!)

LOOPIA

Background of my experience

- Use of BIND-DLZ in a production environment
  - 250 000 zones
  - ~100 000 uniq customers depending on that the service is running
  - 3 years and running
- Use of BIND 9 hosting DNSSEC services
  - ~ 50 zones signed (.SE)
- Have had the pleasure of designing services with BIND-DLZ as counterpart

LOOPIA

Is BIND-DLZ up for consideration?

- What do you intend to do?
  - Number of zones?
  - Zone data update frequency?
  - Needed queries per seconds?
- If you have more than a few thousand of zones then BIND-DLZ is for you
- Maybe not something for your local root server operator

LOOPIA

Pros & Cons

- Provisioning, ever hear of SQL?
- Downtime!, no 35 minute restarts
- Performance, not NSD nor Vanilla BIND but do you need that?

LOOPIA

Production stable

- YES, without a doubt!
- Most problematic experiances is due to bad system design

LOOPIA

Its all about system design

- As with all systems, if you use it the wrong way it does not work
- Some development per data source needed but, it does pay out

LOOPIA

Trial And Error

- Different database back ends
  - Direct connection to MySQL
    - Slow
    - No effective local cache
    - Fast Updates
  - LDAP
    - Slow
    - No effictive local cache
    - Fast updates
    - Random crashes ☹
  - Berkley DB with a master data source
    - Fast
    - Moderat speed updates (once every time data replicate from master)
    - Stable ☺

LOOPIA

Trial And Error

- Performace tuning
  - No system alike
    - What kind of zones (customers)
  - Cache Versus Memory File System
    - Active zones or inactive zones
  - Cache size tuning
    - Standard size 256kb (45% hitrate)
    - 1mb, 10mb or 100mb?
    - Trying to find the best size of cache for us

LOOPIA

Example, LoopiaDNS

- LoopiaDNS, a free DNS-service
  - One account, Multiple domains
  - DNS-editor (A, AAAA, MX, CNAME, NS, ...)
  - Bulkedit features included (stresstest for provisioning)
- DNS System Design
  - One master data source (MySQL v5)
  - Replication Agent
    - Reads changes from MySQL to local Berkeley DB on NS Nodes.
    - Runs once per minute
  - Berkeley DB with 10 mb cache as local data source for BIND-DLZ

LOOPIA

System Design Chart

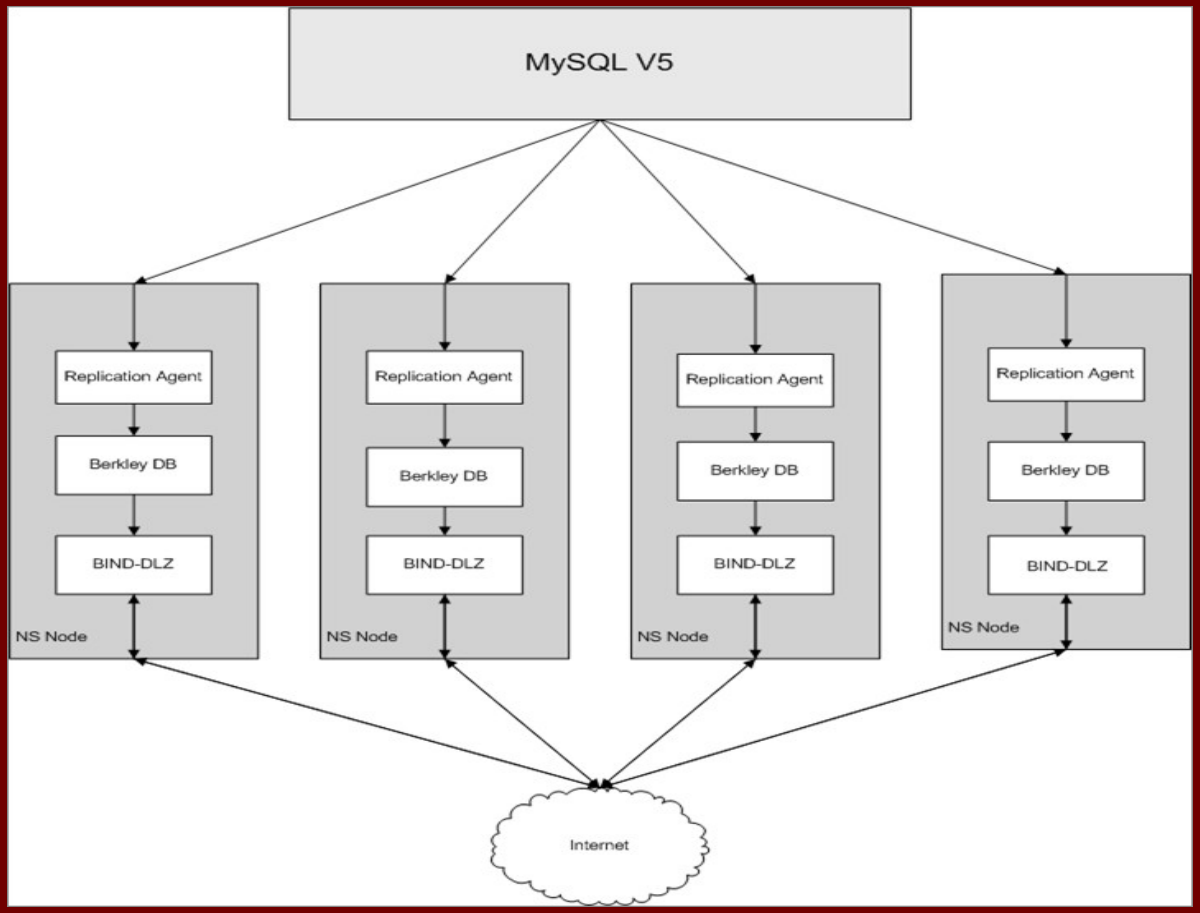

LOOPIA

Summary

- Performance?
  - Hardware
    - Dell servers
    - 2Gb Memory
    - Dual Core Intel Processor (3ghz)
  - Replaying one week of DNS-queries
    - ~ 95% cache hit rate, basically 95% queries goes to memory for answers
    - ~ 3% disk usage
    - ~ 85% processor usage
    - ~ 10000 queries per second / server

LOOPIA

Contact Me

- lars-goran.forsberg@loopia.se
- http://www.loopia.se

Questions

LOOPIA