The Database Software

Introduction

- Why ?
  - old style updates took ~2 hours a day
  - many mods to keep it running
  - object checking too loose
  - almost in “don’t touch” state
- Design decisions
  - decrease man-power needed for updates
  - lower turn-around times
  - strengthen object checking
  - easily extendible
  - guarded attributes provision

Features

- written in Perl
  - highly portable
  - nice DBM features
  - easy to make changes
- “on-the-fly” updates
  - updates are immediately visible in whois
  - turn-around time is minutes
- no human intervention
- guarded attributes provision
- improved syntax checking
- fast whois server
- can manage multiple databases (same syntax)

Changes

- `ripe-dbm@ripe.net` -> `auto-dbm@ripe.net`
- `assign@ripe.net` -> `auto-assign@ripe.net`
- strong syntax checking, more bounces
- whois server
  - more queries/second
  - strict source seperation
  - extra flag (`-s source`)
- guarded attributes
  - added once a day (6 AM currently)
  - blocks are split if necessary
  - cannot be changed through normal updates

Status

- code status
  - some 30 perl modules
  - 3500 lines of code
  - design: Marten + Daniel
  - implementation: Marten + Tony

- operational status
  - beta-tests for some weeks (NCC+others)
  - running production at NCC since last Friday
  - ~10 bugs first day, nothing fatal
  - fairly happy so far
  - distribution will be in November (sorry)