lbnamed: A Load Balancing Name Server in Perl

Roland J. Schemers III

SunSoft, Inc.

schemers@eng.sun.com

Stanford Environment

- same environment on all workstations
  - home directory, mail, applications
- “cloned” workstations

58 Sparcs 31 Alphas 15 Indigos

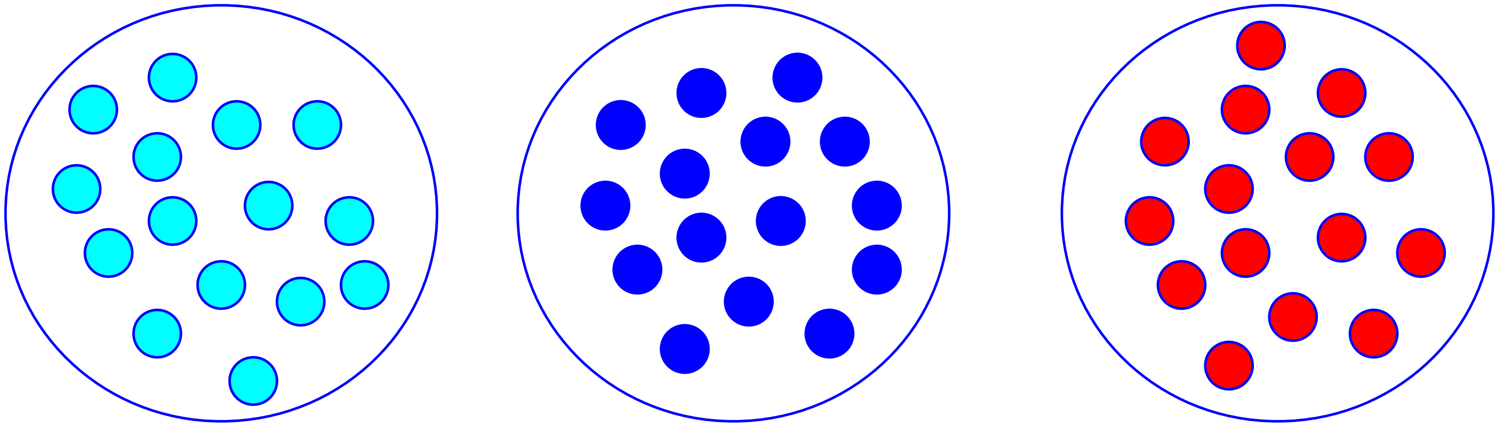

Goals

- load balancing
  - best effort
- multiple groups
  - best “elaine”, “sparc20”, “solaris”, etc...
- dynamic/exotic domains
  - best.stanford.edu
  - schemers.passwd.ns.stanford.edu
  - get information from a dbm file, etc...

Background

- *Load balancing without help from a DNS server*
  - *Users must first login to a workstation and run a program which points them at the best workstation*
- *Load balancing with help from a DNS server*
  - *Users use a special domain name which already points at the best name*

DNS

- nothing magical about the DNS protocol
  - simple question/response format
- create a dynamic response based on question
  - elaine.best.stanford.edu
- if it looks like a duck...
  - any process can listen on port 53

Solution Overview

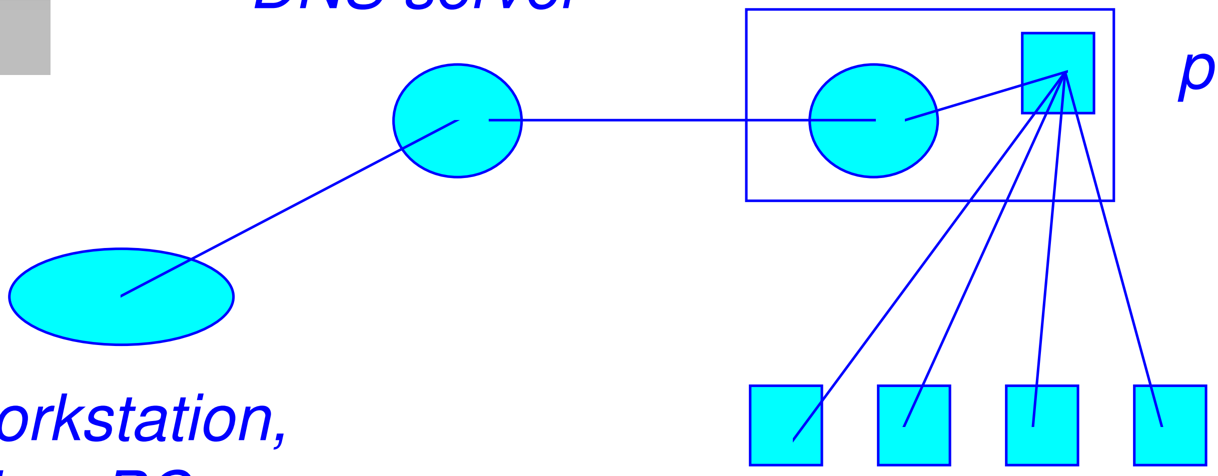

DNS server

lbnamed

poller server

workstation,
Mac, PC,
terminal server,
etc....

poller clients on
workstations

Polling for “load” information

- *poller client*
  - *runs on each workstation*
- *poller server*
  - *runs on the same host as lbnamed*
  - *tells lbnamed when new config file is ready*
- *polled information*
  - *load averages (1, 5, 15)*
  - *total users, unique users, someone on console*
  - *boot time, current time, utmp mtime*

Poller Config file

```
# host    weight      groups
elaine1     1       elaine sparc sunos sweet
elaine2     1       elaine sparc sunos sweet
amy1        1       amy alpha sweet
amy2        1       amy alpha sweet
```

- Key Features
  - no need to put in IP address
  - a host can be in multiple groups
  - groups can be used for CPU type, OS type, location, etc.

lbnamed

- *listens on port 53 (tcp and udp)*
- *reads config file produced by poller*
- *responds to requests from DNS clients*
  - see if dynamic group name exists
  - sort weight information for group
  - return the host with the lowest weight
  - update weight of returned host

existing DNS servers

- lbnamed is not a replacement for BIND
- add NS records to delegate dynamic domains
- In configuration file for foo.com:

```
best    IN    NS    lbnamed1.foo.com.
best    IN    NS    lbnamed2.foo.com.
```

requests for names in the “best.foo.com” domain would be sent to lbnamed1 or lbnamed2

Perl 4 version

- *Why Perl?*
  - *easy to prototype*
  - *make changes quick*
- *Planned on re-writing it in C*
  - *ended up working out in Perl*
  - *I moved on to the next hack*

Perl 5 version

- *Goals for Perl 5 version*
  - *extensible*
- *Major features*
  - *DNS “toolkit”*
  - *static domain data*
  - *dynamic domain data*

DNS “toolkit”

- *Perl functions to manipulate DNS packets*

```
use DNS;
$answer = dns_answer(QPTR, T_TXT,
                     C_IN, 60, rr_TXT(“hello”));
$data = rr_CNAME(“foo.sun.com”);
$data = rr_MX(20,”sun.com”);
```

Static domains

- *data that doesn’t change dynamically*
- *lbnamed looks for static data first, then dynamic data*
- *examples*

```
LBDB::add_static(“best.stanford.edu”, T_SOA,
    rr_SOA($hostname, $hostmaster, $serial, $refresh,
                                        $retry, $expire, $min));
LBDB::add_static(“foo.best.stanford.edu”, T_A,
        rr_A($ipaddress));
```

Dynamic domains

- *data that changes dynamically*
- *a dynamic domain is registered with lbnamed*
- *all requests for dynamic domain are passed to a function which handles them*

Dynamic domain examples

- *time.stanford.edu*
  - *Return current time when name is queried*
- *Register the domain handler:*

```
LBDB::add_dynamic("time.stanford.edu" =>
                                  \&handle_time_request);
```

- *handle_time_request will get called when a request is made for the name "time.stanford.edu"*

Dynamic domain examples

```
sub handle_time_request {
   my($domain, $residual, $qtype, $qclass, $dm) = @_;
   my($date) = &get_ctime;
    $dm ->{‘answer’} .= dns_answer(QPTR, T_TXT, C_IN, 60,
                                                        rr_TXT($date));
    $dm->{‘ancout’} += 1;
}
```

Dynamic domain examples

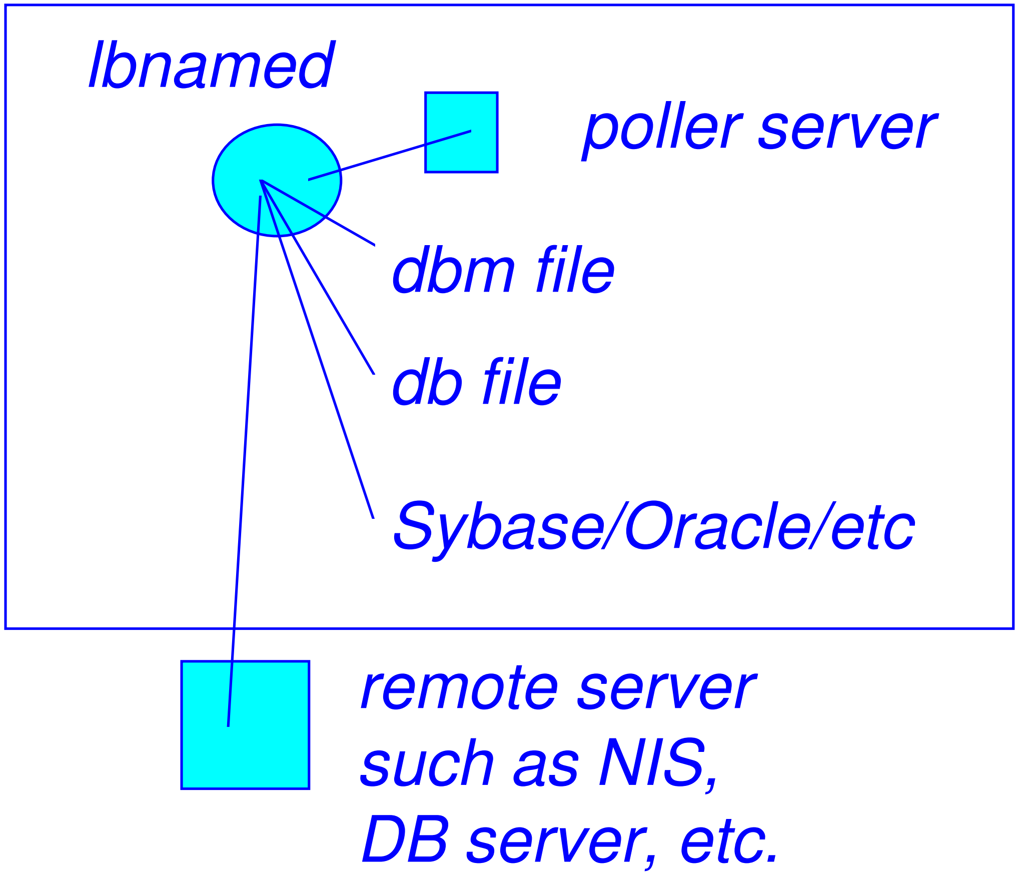

lbnamed

poller server

dbm file

db file

Sybase/Oracle/etc

remote server
such as NIS,
DB server, etc.

Conclusion

- *has worked out well*
- *new users don’t even know they are using load balancing*
- *sysadmins can add/remove machines and groups without disrupting users*
- *the poller automatically detects when a host is down an removes it from the generated config file*
- *problems have been minor*

Future Directions

- package up and document Perl 5 version
- recognize bad hosts
- factor in swap, memory, etc...
- allow client to specify weight
- add lbnamed functionality
  - have someone else do all the above!

lbnamed: A Load Balancing Name Server in Perl

Roland J. Schemers III

schemers@eng.sun.com

http://www-leland.stanford.edu/~schemers/dist/lb.tar