

A Distributed Multicast DNS System

Dan Johnson

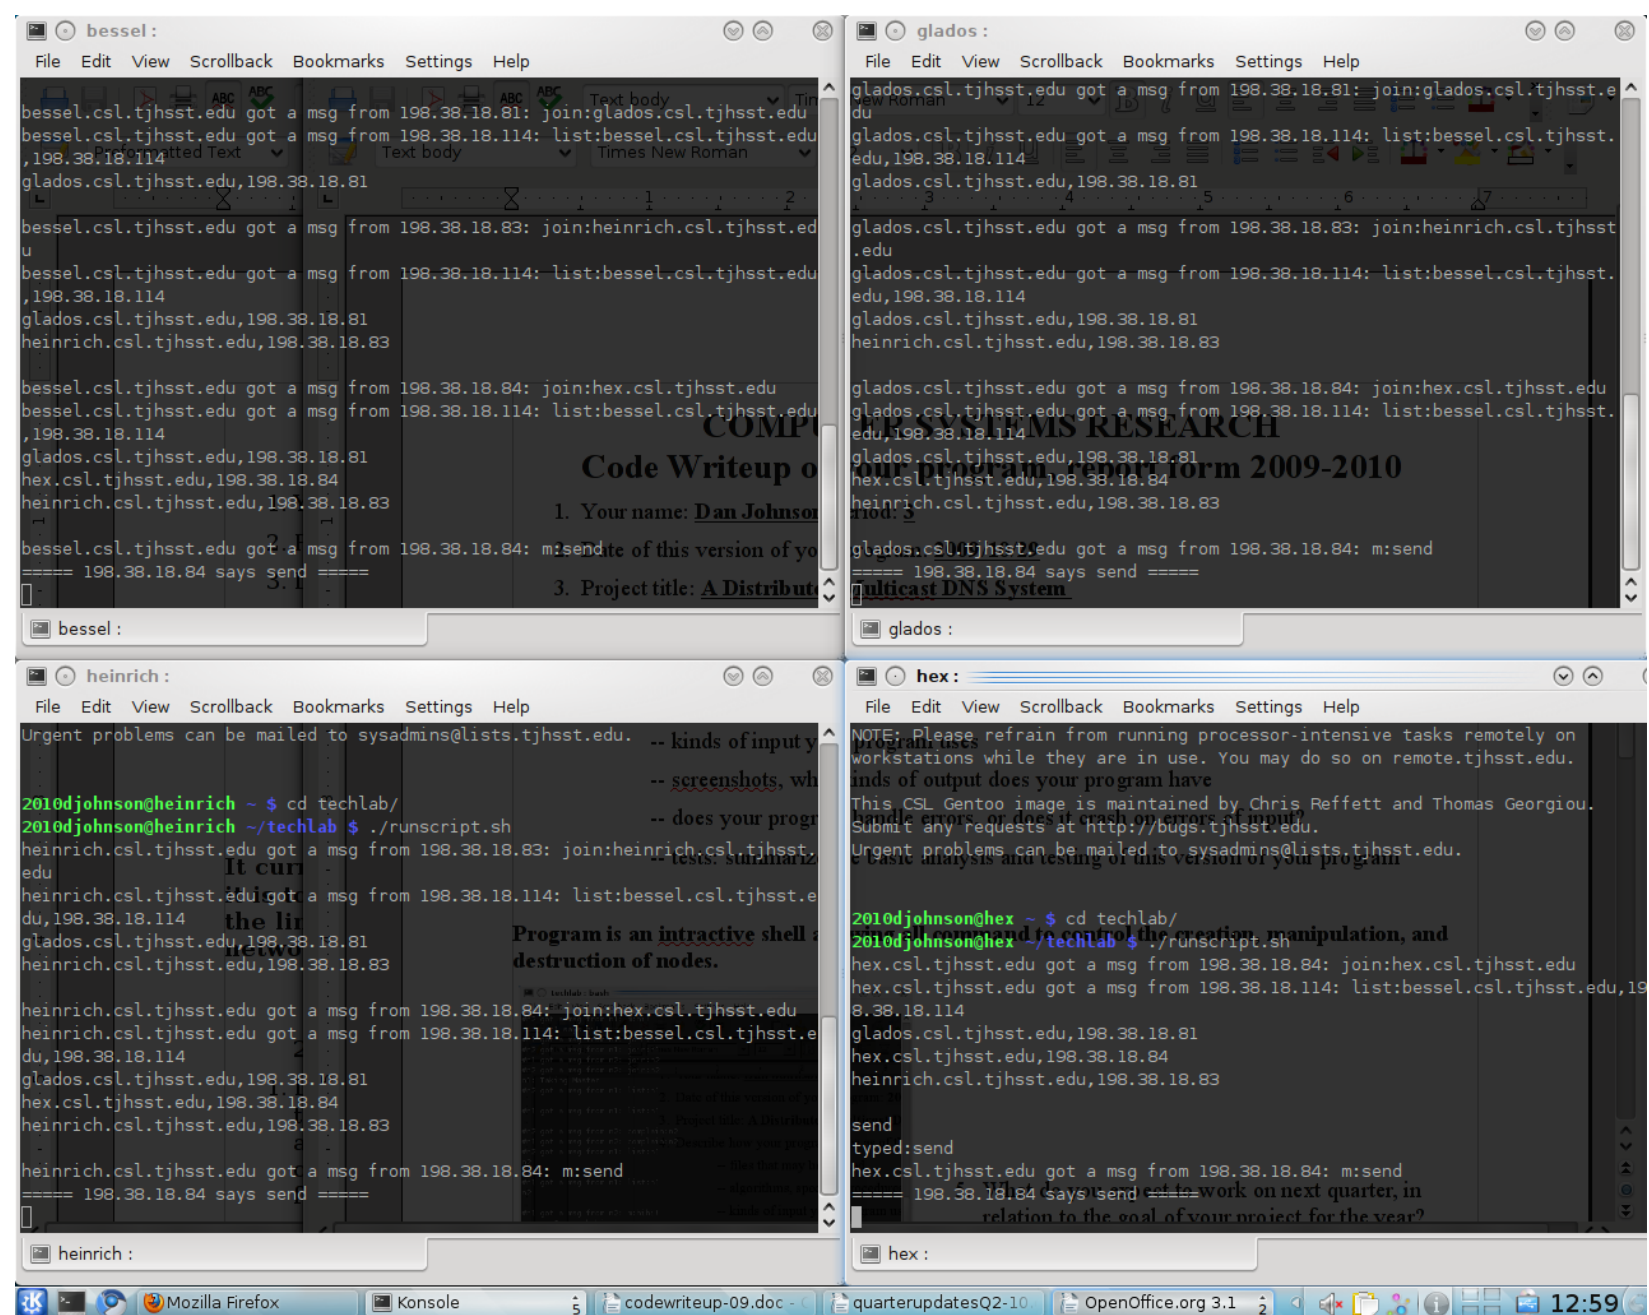
TJHSST Computer Systems Lab 2009-2010

Abstract:

There exists a need for less dependance on the single point of failure that is DNS servers. It is quite possible for a computer to be completely connected to the internet, but run slowly or not at all in the case of a malfunctioning DNS server. My goal is to allow a group of workstations as present on a medium-size subnet to survive complete loss of a DNS server through collaboration. The syslab workstations provide an example of an appropriate size.

Introduction:

I will be using python for the initial proof-of-concept simulation and to make sure the protocol will allow for all necessary features. After the simulated nodes can function properly, I will translate the protocol and implementation into a network version. Eventually, I will implement a NSS (Name Service Switch) module in order to allow a native linux system to take advantage of these features, or implement the regular DNS protocol and allow the current host to query us on the localhost address (127.0.0.1 or ::1)



The image shows four terminal windows arranged in a 2x2 grid, illustrating network communication between hosts. The top-left window is titled 'bessel:' and shows messages from other hosts like 'glados.csl.tjhsst.edu' and 'heinch.csl.tjhsst.edu'. The top-right window is titled 'glados:' and shows messages from 'bessel.csl.tjhsst.edu' and 'heinch.csl.tjhsst.edu'. The bottom-left window is titled 'heinch:' and shows messages from 'bessel.csl.tjhsst.edu' and 'glados.csl.tjhsst.edu'. The bottom-right window is titled 'hex:' and shows messages from 'bessel.csl.tjhsst.edu' and 'glados.csl.tjhsst.edu'. The messages include 'join:' and 'm:send' commands, indicating a distributed network protocol in action. A 'Code Writeup' document is visible in the background of the bottom-left window.

Benefits:

The benefits of offloading routine and emergency duties from the nameserver has several practical benefits. First, in the event of a nameserver outage, not all systems need to fail. While non-cached entries may not be available, those that have seen high use (google.com, for example) will still be available. This helps to eliminate one instance of a single point of failure. With a sufficient number of hosts, processing queries on the main nameserver can lead to performance issues. By dividing responsibility for name lookups among hosts, the speed and scalability of lookups can be improved.

Hopefully after this project is completed it will be good enough to put into production in UNIX computer labs around the world. With enough effort and review, it should be possible to gain acceptance into the community, assuming the security requirements and social requirements are met.