

A Modular Email Architecture Using Open Source Components



Presented to MVMRUG
April 22, 2005

Scott Courtney
Sine Nomine Associates

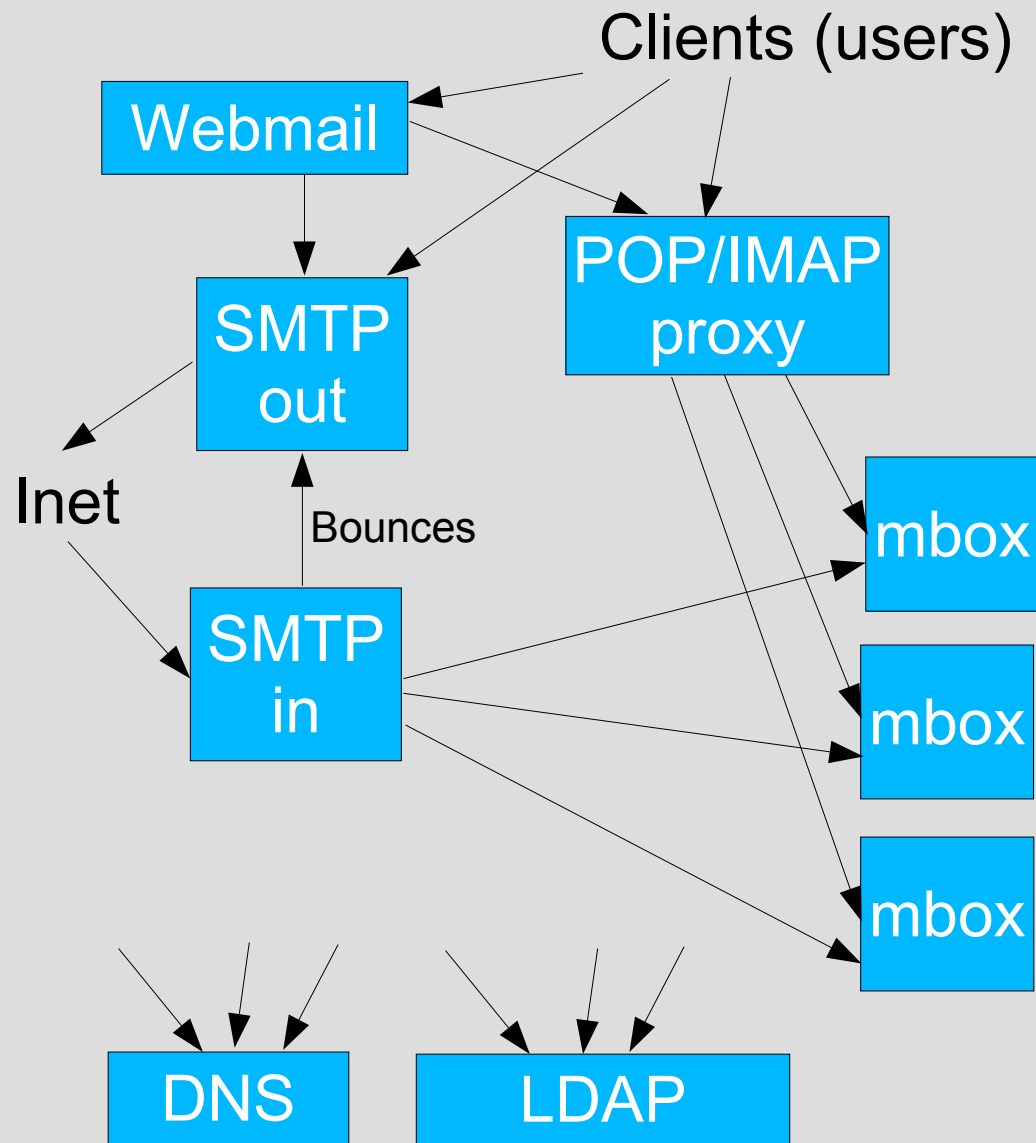
Functional Goals

- SMTP inbound and outbound mail
- POP3 and IMAP4 mailbox delivery
- Webmail
- Spam blocking
- Virus scanning and blocking
- Multiple virtual domains
- Scalable to 500K users

Architectural Goals of System

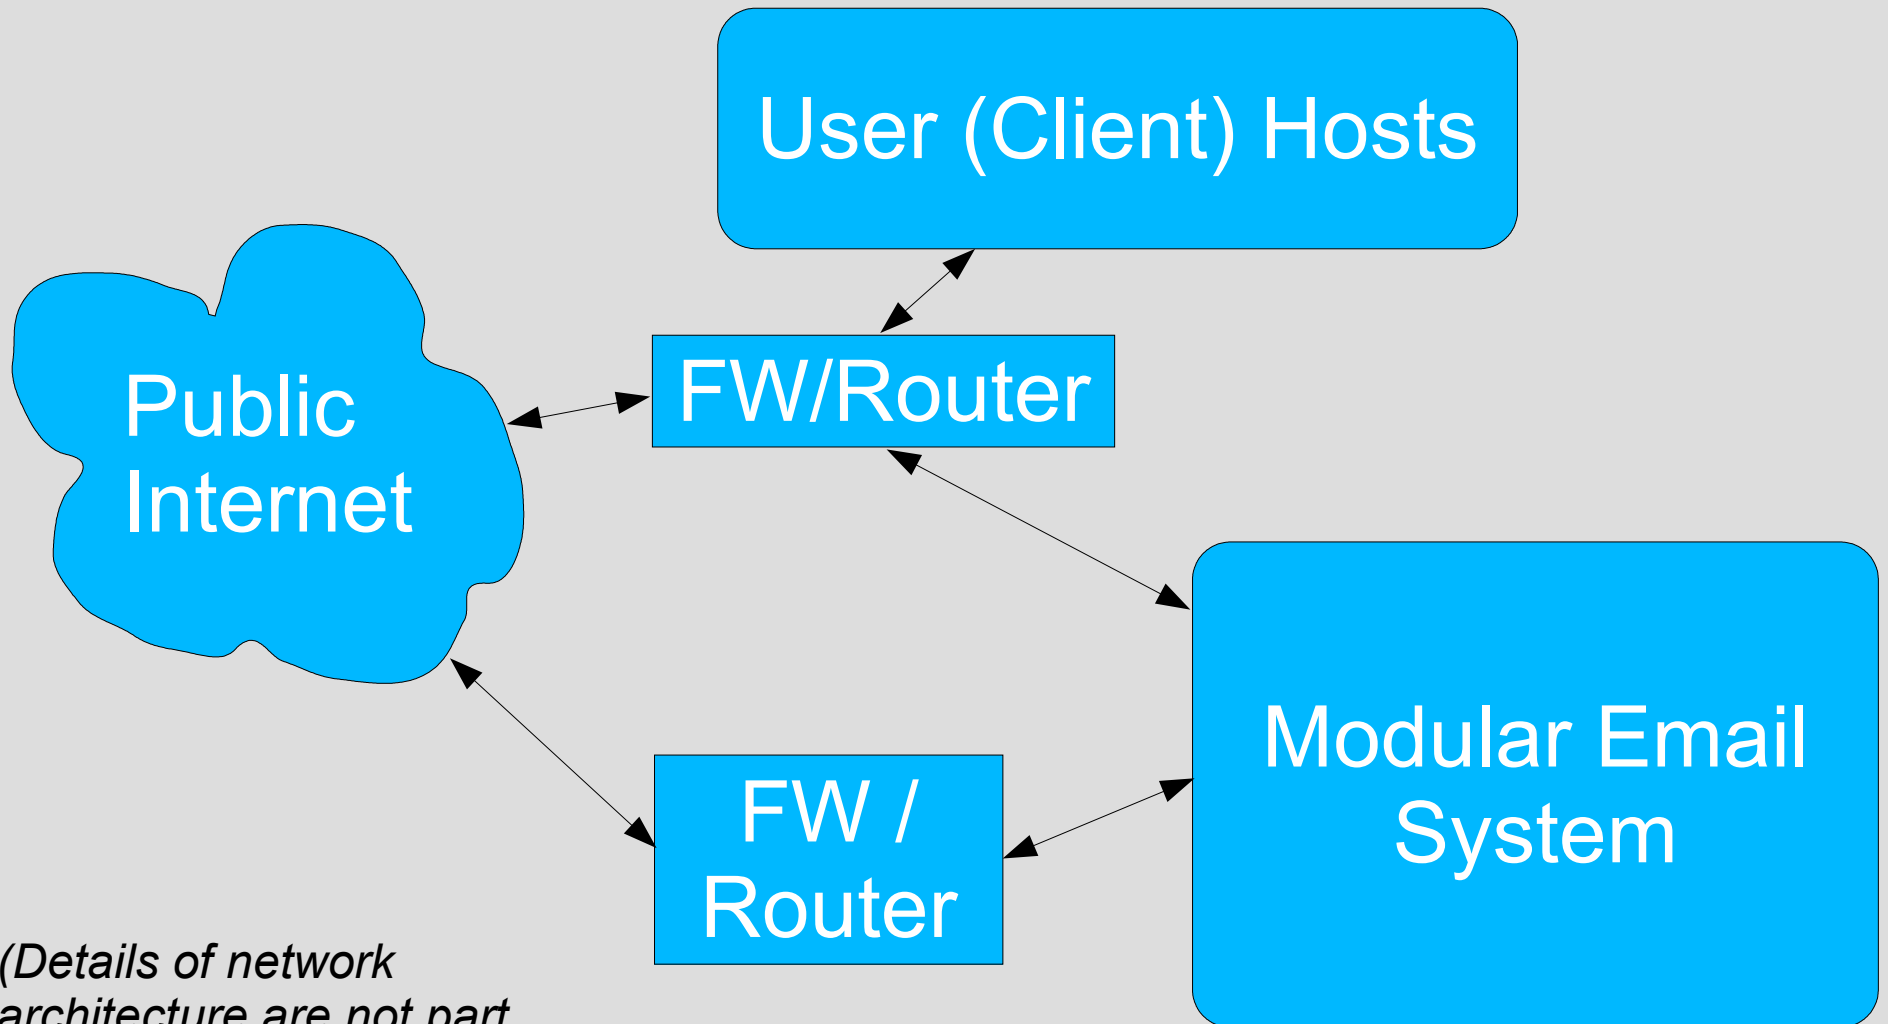
- Centralized user authentication, profiles
- Virtual user accounts (no shell accounts)
- Modular deployment from server templates
- Horizontally scaling
- Near-instantaneous failover where practical
- Centralized backup
- Centralized logging
- Deploy on low-cost Intel hardware or VM/Linux, or mix platforms

Design Overview




- Arrows show direction of service request, not necessarily data flow
- "SMTP in" is actually three layers
- Webmail behaves like a user client, conceptually "outside the box"
- DNS and LDAP connect to everything (not all shown here)
- Admin and backup servers omitted for clarity

Network Overview



(Details of network architecture are not part of this presentation.)

SMTP Inbound -- Details

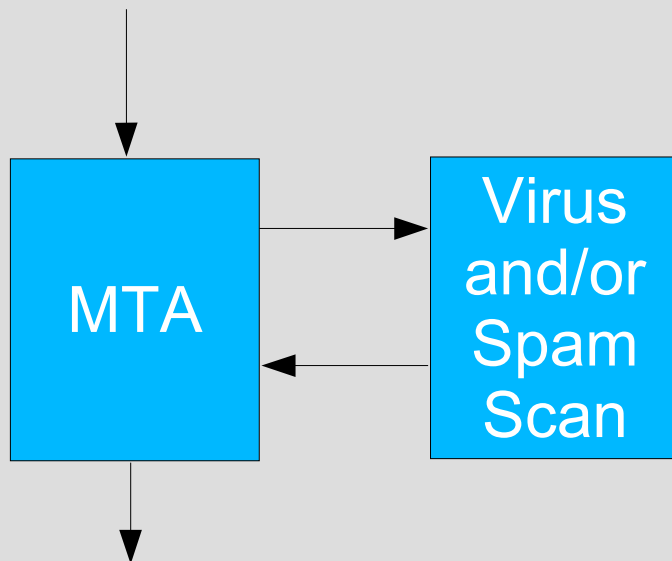


- Inbound connections distributed by simple round-robin DNS
- Filters can go offline for scheduled maintenance after letting queue drain
- Philosophy:
 - Block before scan
 - Scan before route
 - No spam, virus bounce
- Early *or* late lookup of valid destination

Virus and Spam Removal

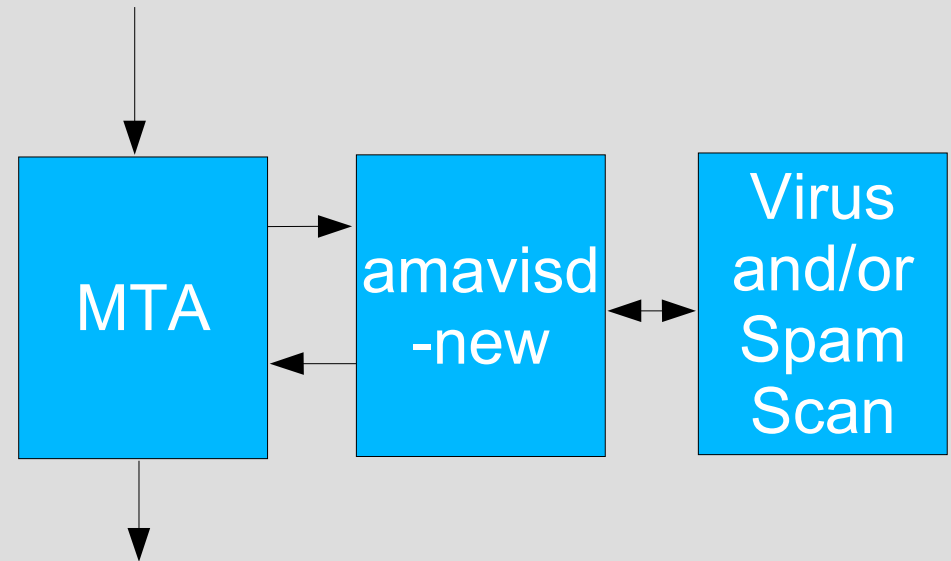
Method 1: Direct Daemon

- SpamAssassin, ClamAV, or similar tools listen on TCP ports
- May be on dedicated server separate from MTA
- Simpler configuration and less points of failure than with amavisd-new
- Scanning closely coupled to MTA configuration

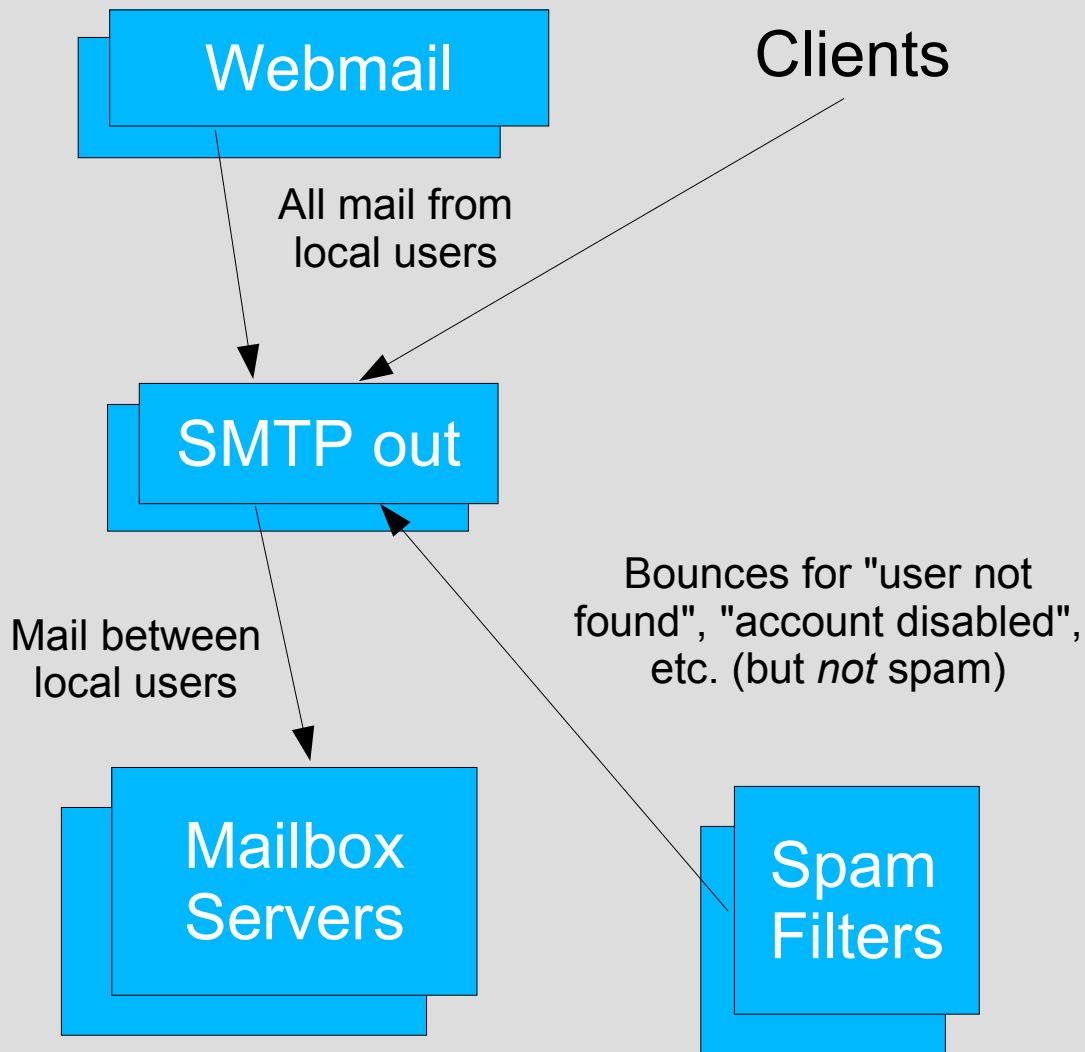


Method 2: amavisd-new

- amavisd-new listens on TCP port, runs spam and/or virus scan from exec() call or API function call, or shells out
- amavisd-new may be on separate server, but scanner(s) local to amavisd-new host
- Can change or add new scanners w/o touching MTA configuration




SMTP Outbound -- Details



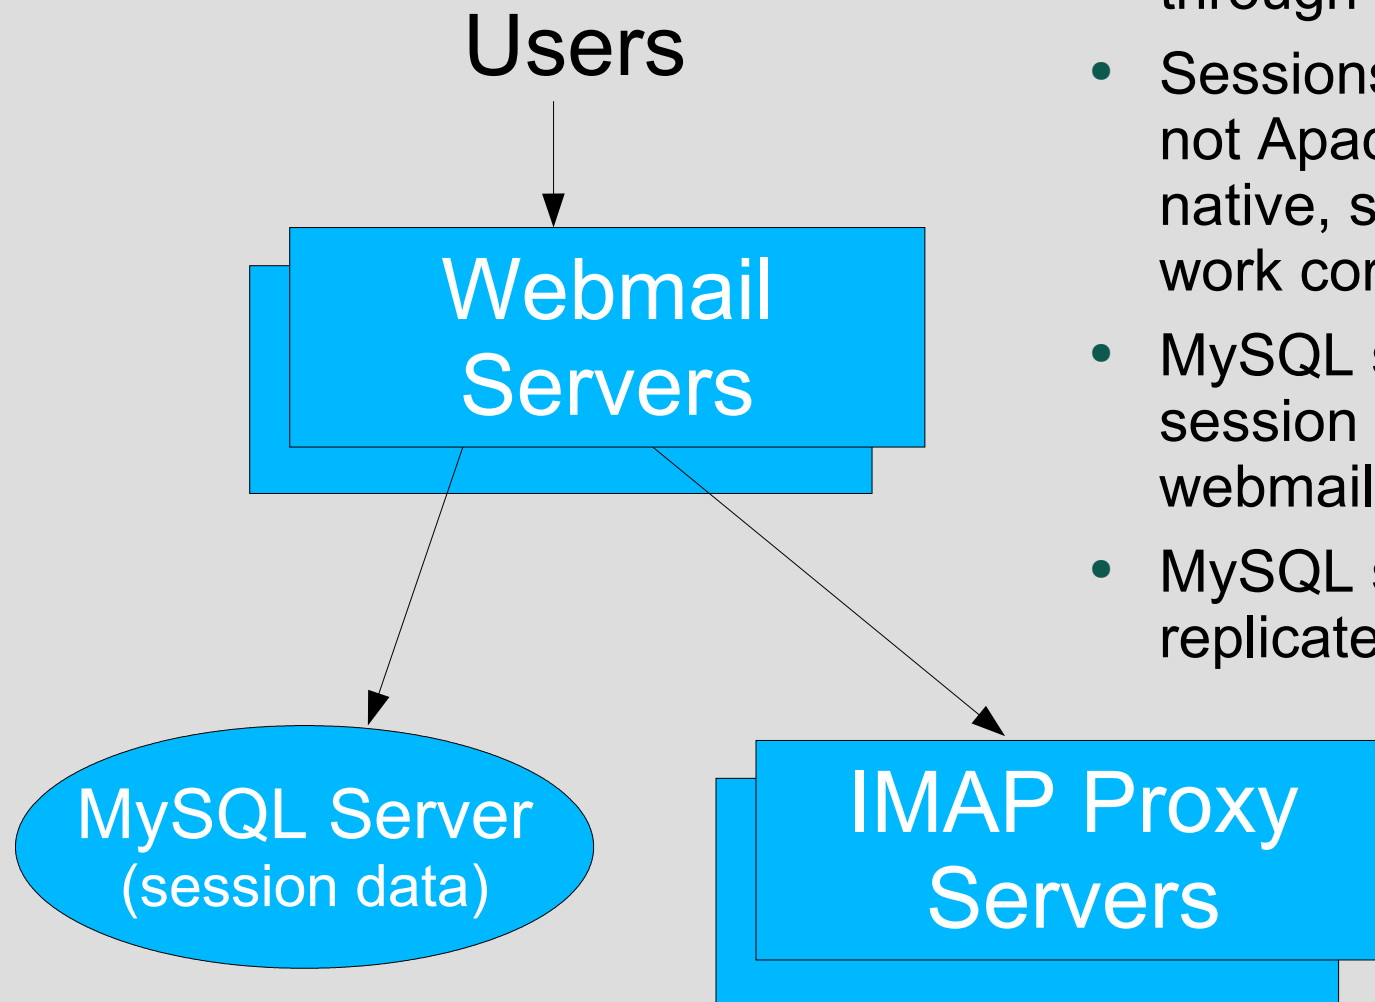
- All user-originated messages pass through virus and spam filters
- Outbound servers have local spam, virus filter
- Philosophy:
 - Outbound << Inbound
 - Nothing nasty out!
 - No spam, virus bounce
 - Decouple outbound queuing

POP3/IMAP4 Service



- Mailbox servers are just POP3/IMAP4, nothing else
- Routing to mailbox host based on LDAP attribute
- No direct connection to mailbox servers from untrusted hosts
- POP/IMAP proxy looks up mailbox location in LDAP, emulates (then passes through) authentication
- No user data whatsoever on proxy server -- admin is trivially simple

Webmail Service



- Webmail authentication is through IMAP, not LDAP
- Sessions *must* be in MySQL, not Apache or PHP or Horde native, so load balancing will work correctly
- MySQL server stores only session data and [optionally] webmail user preferences.
- MySQL server can be replicated for warm-failover

LDAP (My Diagrams vs. Reality)

My Diagrams

- One big LDAP server
- If LDAP goes down, the world comes to a grinding halt

Reality

- One *master* LDAP server
- Any other server that needs LDAP lookups has its own replicant locally
- Master server down?
No user profile changes but things keep working for all that matters.

Software Selection: MTA

Selected:

Exim

<http://exim.org/>

Why?

- Other Candidates
- Sendmail
 - QMail

- Extremely configurable
- Sophisticated and *readable* config syntax
- Config localized to one file (manageability)
- Secure, fast, reliable
- Integrates smoothly with external data sources
- Superb documentation

Software Selection: POP/IMAP

Selected:

Courier

<http://www.courier-mta.org/imap/>

Other Candidates

- WU-imapd
- (numerous others)

Why?

- Extremely configurable
- Modular authentication
- Widely supported
- Secure, fast, reliable
- Integrates smoothly with external data sources

Software Selection: IMAP/POP Proxy

Selected:
Perdition

www.vergenet.net/linux/perdition

Why?

- Does what is needed
- See "Other Candidates"

Other Candidates

- None, really

Software Selection: Webmail

Selected:

Horde

www.horde.org

(Plus IMP and Turba plugins)

Other Candidates

- Squirrel Mail
- EMU Webmail
- Courier Webmail
- etc.

Why?

- Extremely configurable
- Modular design with other "groupware" components available
 - Addressbook/contacts
 - Calendar
- Mainstream and widely supported
- Integrates smoothly with external data sources

Adding z/VM Linux to the Mix

Great on z/VM Linux

- Mailbox Servers +++
- POP/IMAP Proxy Servers
- Webmail Servers
- MySQL Server(s)
- Master LDAP (?)
- Inbound SMTP route (spin off from spam scan)
- Outbound SMTP -- *but*, spin off virus/spam scan to Intel

Maybe Stay on Intel

- Virus Scanning
- Spam Scanning
- Outbound SMTP if you keep virus/spam scan local to these hosts

Note that the servers that *don't* make sense on z/VM Linux are also the servers that store *no important data!*

Performance in the Real World

Configuration

- Intel P4-class hardware
~2.0 GHz
- Uniprocessor except
virus, spam scan and
master LDAP
- 2 GB RAM per host
- 100 Mb Ethernet
- ~40K users
- > 30K inbound/hour not
counting deleted spam,
virus

Performance

- Virus + spam \geq 2X
actual inbound mail
- About 8K~10K users
per mbox
- 2 servers at each other
pipeline stage can
handle full load (3 used
by choice to allow
scheduled outages)
- In production ~2 years
with no large outages

Remaining Challenges

- Mailbox servers are "single point of failure" to the specific users of a given server
 - Using z/VM Linux largely eliminates this problem
 - Can use tools like "heartbeat" for hot failover
- Spam and viruses are absolutely *rampant* on the Internet -- even state-of-the-art filtering isn't good enough yet
 - Sender Policy Framework (SPF)
 - CRM114, Bayesian filter -- good but complex
 - Authenticated SMTP (supported by this design!)

Conclusion

- System is horizontally and vertically scalable
- Platform-neutral, and most pieces work just fine on z/VM Linux
- *All* pieces work fine on z/VM Linux if you have a small number of users
- Load balancing accomplished with simple round-robin DNS and built-in replication features of OpenLDAP and MySQL
- Servers can go offline easily for scheduled maintenance
- Servers can be cloned from server-type templates

Contact Information



Scott Courtney, Senior Engineer
Sine Nomine Associates
scourtney@sinenomine.net
330.353.0403

<http://sinenomine.net/>