

Performance optimisation of webmail

Katerina Mparmpopoulou Periklis Stefopoulos

Supervised by: *Michiel Leenaars*

University of Amsterdam
System & Network Engineering

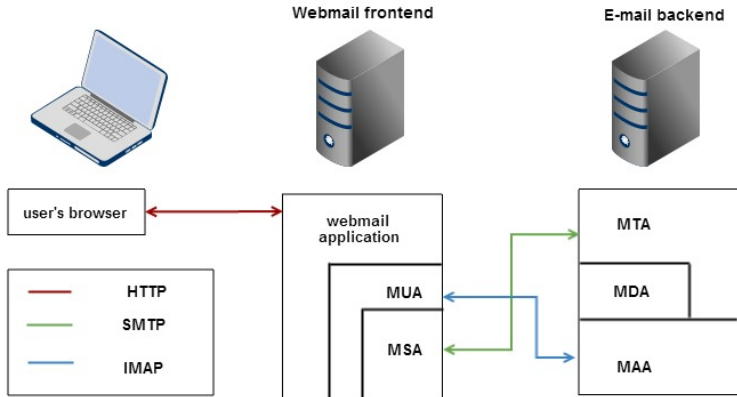
July 3, 2013



Outline

- 1 Introduction
- 2 Research question
- 3 Approach
- 4 Experimental Results
 - User Experience
 - System Performance
- 5 Conclusions

Webmail System Components



E-mail Components	Examples
Mail User Agent (MUA)	Afterlogic
Mail Submission Agent (MSA)	Afterlogic
Mail Transfer Agent (MTA)	Postfix
Mail Delivery Agent (MDA)	Postdrop
Mail Access Agent (MAA)	Cyrus

Research purpose

Introduction

Research question

Approach

Experimental Results

User Experience System Performance

Conclusions

- Hundreds of millions of end users depend on Webmail technologies
- Open source web frontends to mail servers are an often neglected area of development
- Better understanding the performance of web mail applications is a prerequisite to better tuning these applications

Research Question

What are the bottlenecks, in terms of performance, of current Webmail implementations and which could be the most optimal solution?

Experimental Environment

Webmail Frontends

- Squirrelmail
- Roundcube
- Horde IMP
- Afterlogic Webmail Lite

Webmail Backends

- **Courier** - Postfix - Amavis - ClamAv - SpamAssassin
- **Dovecot** - Postfix - Amavis - ClamAv - SpamAssassin
- **Cyrus** - Postfix - Amavis - ClamAv - SpamAssassin

Experiments

5 users with different mailbox size

- 1500 messages with only text
- 1500 messages with attachments plus text
- 3000 messages with only text
- 4500 messages with only text
- 6000 messages with only text

3 different user actions

- Log in to Webmail
- Searching a keyword from the "Subject" field
- Searching a keyword from the entire message content

Metric	Extraction method
Latency	tcpdump/Wireshark
CPU time ¹	systat/sar
Unique Set Size (USS)	smem
Proportional Set Size (PSS)	smem

Table: Benchmark metrics and their extraction method

¹ $CPU\ time = CPU\ utilization * elapsed\ time * number\ of\ CPUs$





Figure: average fetching time during login


Searching from the "Subject"

Dovecot IMAP backend




Searching from the "Subject"

Cyrus IMAP backend



Searching from the "Subject"

Courier IMAP backend



Searching from "Subject"

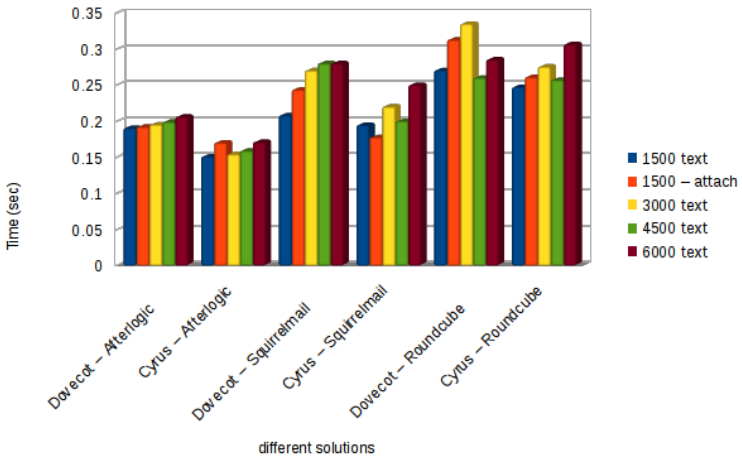



Figure: Comparison of the most effective Solutions regarding searching from "Subject"


Searching from the Entire Message

Courier IMAP backend




Searching from the Entire Message

Dovecot IMAP backend



Searching from the Entire Message

Cyrus IMAP backend



Searching from the Entire Message





Figure: Comparison of the most effective Solutions regarding searching from Entire Message content


Resources consumption when user searches from the "Subject"

Dovecot and Cyrus IMAP backends



Resources consumption when user searches from the "Subject"

Courier IMAP backend



Resources consumption when the user searches from the Entire Message




Figure: users search a keyword from the Entire Message Content: CPU time consumption for Dovecot and Courier IMAP backends

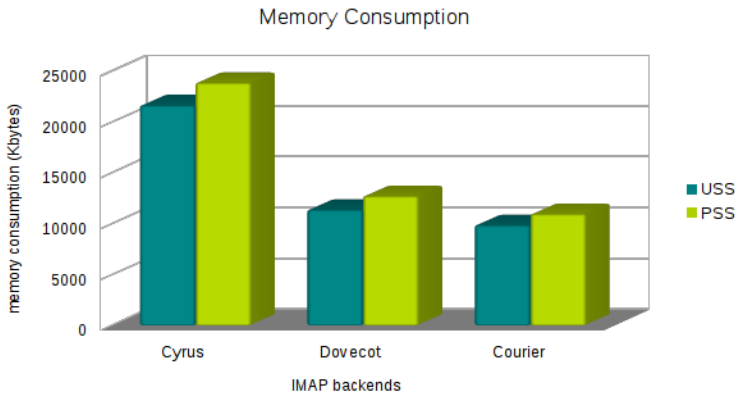


Figure: average memory consumption

Conclusions

- **Afterlogic** achieves shortest latencies in searching
- **Horde IMP** has the longest response times
- **Dovecot and Cyrus** carry out the search from "Subject" request efficiently and with relatively the same latency
- **Cyrus** is the IMAP backend that performs the best during search from the entire message
- **Cyrus** has the lowest CPU utilization and the highest average memory consumption for all IMAP functions, followed by Dovecot

Answering the Research Question

What are the bottlenecks, in terms of performance, of current Webmail implementations?

- The major bottleneck in an integrated webmail system is the IMAP backend

which could be the most optimal solution?

- the solution of using Cyrus IMAP combined with Afterlogic Webmail Lite performs better in terms of both user experience and system overall performance

Questions



Introduction

Research
question

Approach

Experimental
Results

User Experience
System
Performance

Conclusions