

Introduction of New Kyushu University Primary Mail Service for Staff Members and Students

Kasahara, Yoshiaki
Kyushu University

Ito, Eisuke
Kyushu University

Fujimura, Naomi
Kyushu University

<https://hdl.handle.net/2324/2557140>

出版情報 : Proceedings ACM SIGUCCS User Services Conference. 42, pp.103-106, 2014-11-02.
Association for Computing Machinery

バージョン :

権利関係 :



Introduction of New Kyushu University Primary Mail Service for Staff Members and Students

Yoshiaki Kasahara

Kyushu University

6-10-1 Hakozaki, Higashi-ku

Fukuoka 812-8581, Japan

+81 92 642 2297

kasahara.yoshiaki.820@m.kyushu-u.ac.jp

Eisuke Ito

Kyushu University

6-10-1 Hakozaki, Higashi-ku

Fukuoka 812-8581, Japan

+81 92 642 4037

ito.eisuke.523@m.kyushu-u.ac.jp

Naomi Fujimura

Kyushu University

4-9-1 Shiobaru, Minami-ku

Fukuoka, Japan

+81 92 553 4434

fujimura.naomi.274@m.kyushu-u.ac.jp

ABSTRACT

In the end of fiscal year 2013 (March 2014), Kyushu University Information Infrastructure Initiative introduced new Primary Mail Service for Staff Members and Students. The previous service for staff members had been built using proprietary mail and LDAP appliance, but tight user license and the nature of proprietary system caused several troubles. The previous service for students had been built using open source software (Postfix and Dovecot), but there were some issues with the old implementation of LDAP authentication. With these experiences, we decided to design a new system by merging both system using open source software including Postfix, Dovecot, OpenLDAP, and Squirrelmail. We also extended alias address setting service (previously available for students only) to all the members including staff members. In this paper, we explain the design and implementation, user migration, current status and future works of our new mail service.

Categories and Subject Descriptors

H.4.3 [Information Systems Applications]: Communication Applications – *Electronic mail*

Keywords

University mail service, On-premises service, Service Migration.

1. INTRODUCTION

Information and communication service is indispensable for education and research activities in Universities. Among various services, electronic mail is one of the fundamental services even before the Internet had been widely available in the world. Kyushu University has been providing on-premises email services using its own domain name for our University staff members and students for several years. Email service for students was started

in 1995[1][2], and another service for University staff members started in July 2009 [3].

In the end of fiscal year 2013 (March 2014), Kyushu University Information Infrastructure Initiative introduced new Primary Mail Service for Staff Members and Students. This time, we decided to design a new system by merging two separate mail systems for staff members and students using open source software including Postfix, Dovecot, OpenLDAP, and Squirrelmail. We also extended name-based alias address setting service (previously available for students only) to all the members including staff members.

In this paper, we explain the design and implementation, user migration, current status and future works of our new mail service. The rest of this paper is organized as follows. In Section 2, we discuss the previous mail system. In Section 3, we introduce the overview of our new mail system. In Section 4, we explain about the migration of users from the old system to the new system. In Section 5, we introduce some notable troubles we encountered after migration, current status, and future works. Finally, we conclude the paper in Section 6.

2. PREVIOUS SYSTEM

First of all, we will give an overview of Kyushu University and University-wide mail services provided by Information Infrastructure Initiative.

2.1 Number of Users

Table 1 shows the approximate number of IDs issued for the University-wide authentication service (also provided by Information Infrastructure Initiative). The main members of our University are students, faculty, and staff. The students consist of undergraduate and graduate curricular. There are also some non-

Table 1. The number of IDs in Kyushu University (Mar 2013)

Role	Total No. of IDs (approx..)
Curricular students	19,000
Non-curricular students	500
Faculty and staff members	9,000
Temporary staff etc.	800
Total	29,300

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright is held by the owner/author(s).

SIGUCCS '14, November 02–07 2014, Salt Lake City, UT, USA
ACM 978-1-4503-2780-0/14/11.
<http://dx.doi.org/10.1145/2661172.2662965>

curricular students such as research students and special register students. In addition to students, faculty, and staff members, there are other members with various roles such as research fellows, temporary staff, visiting researchers, and so on. The number also represents how many people can use “Primary Mail Services”.

2.2 Previous Mail Service for Staff Members

The previous mail service for staff members (called “Kyushu University Primary Mail Service”) was started in July 2009. Before that we didn’t have a centralized mail service for staff members. We had established authentication infrastructure and assigned unique SSO-KID to every employee in Kyushu University from Sept. 2007[4], and this enabled us to implement centralized mail service for staff members.

We had to build the entire system within a half a year. Since some of us had some experiences with Mirapoint mail appliance, we decided to purchase some of them and build mail service on top of them. Figure 1 shows the system overview. The number of equipment gradually increased and the final configuration included two message directors, four message servers (for user’s mailboxes) and two dedicated LDAP servers.

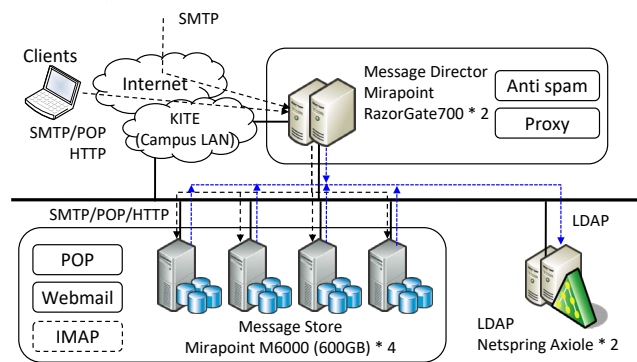


Figure 1. Overview of Previous Mail System for Staff Members

Generally speaking, the system offered average functionality for central mail system including standard email protocols (SMTP and POP3) and a decent webmail system. Due to the limited storage capacity, the mailbox quota was 300MB and messages were kept on the system for 90 days. For users who needed more capacity, we provided a premium service (\$10/month) with 1GB mailbox quota without message expiration and IMAP support.

One of the largest issues with this mail system was that all the components were proprietary appliance systems with per-user licenses, which caused rather high per-year maintenance cost. Because the total number of employees increased gradually due to reasons such as increase of temporal staff and the need to support more active users temporarily during the beginning of a fiscal year (overlapping newcomers and retired employees), we experienced failure of adding new users due to insufficient user licenses a couple of times in 4 years. The only way to resolve the problem was to pay the additional charge to add more user licenses. License fees also prevented us from enabling full functionality of the system because we couldn’t afford purchasing separate license fee for all the users. For example, we only purchased 100 IMAP licenses for premium users.

Another issue was that the appliance systems were black boxes for us. For example, the Mirapoint equipment was developed in US, so when we had problems, sometimes it took a long time to

investigate and fix the source of the problems. We thought many times that if we could have access to the source code of the system, because some of problems seemed quite trivial to fix. We also had a performance issue, which we couldn’t resolve during the lifetime of our service. The system could only process and receive about 2 incoming messages per second. We believed it was too slow considering the specification of the system. When we would send a broadcast email messages to all the employees (about 10,000 users), it took almost 1 hour and accumulated more than 2,000 message in the process queue during the broadcast, which caused a disturbing delay to all the message deliveries. We couldn’t find a reasonable explanation of the delay, partly because we couldn’t obtain enough diagnostic information from appliance equipment to point to the culprit.

2.3 Previous Mail System for Students

A centralized mail service for students had been started in 1995 as a part of educational information system. After some system replacements, the service we replaced this time had been started in 2011. At that time the system was (mostly) separated from educational system and became a part of our Primary Mail Service. The system used open source software such as Postfix and Dovecot on top of a virtualized IA server running Red Hat Enterprise Linux. The system offered standard mail protocols such as SMTP, POP3, and IMAP, but no webmail was provided. The mailbox quota was 300MB without expiration. Figure 2 shows the system overview.

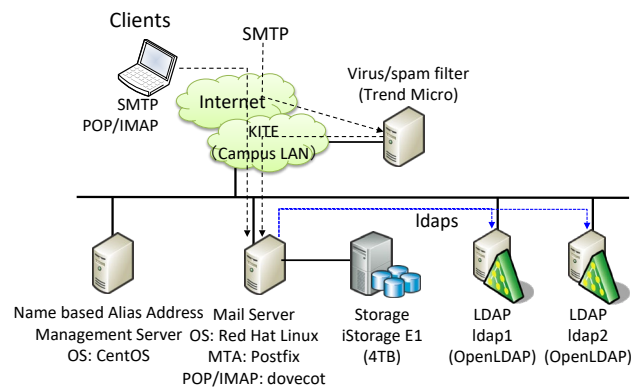


Figure 2. Overview of Previous Mail System for Students

Because the system employed standard open source software such as Postfix and Dovecot, it was easier to customize and introduce new services. One such service was name-based alias address setting service we reported in SIGUCCS 2012 [2]. On the other hand, the entire design was inherited from the previous system (as a part of educational system), which was not very efficient for an email-only system. For example, the user authentication and authorization was done via (outdated) pam_ldap and nss_ldap whose LDAP handling was very inefficient (no LDAP connection keep-alive and serialization). Due to this inefficiency we experienced LDAP service overload and other problems [5][6].

2.4 SSO-KID

As mentioned in section 2.2, Information Infrastructure Initiative provides unified user IDs called “SSO-KID” to employees. We provide various identity providers (including LDAP servers and Shibboleth IdP) to support using SSO-KID as a user

authentication credential for various information services in Kyushu University.

On the other hand, until April 2014 all the students used "Student ID" for user ID of various information services. Also the ID was used as a local-part of their email address. Student ID includes the department code and enrollment year, so it changes when a student proceeds from undergraduate to graduate school or changes departments, which hinders continuity of various information services. To overcome the problem, in April of 2014 we introduced student SSO-KID, which is static throughout the student life in Kyushu University, and started using the ID for our new mail service.

In March 2014, the identity management system would be completely replaced in parallel to mail system replacement. The mail system heavily relied on the proper operation of identity management system, so it was a big concern.

3. NEW MAIL SYSTEM

In the middle of 2013, we decided to merge these two mail services and introduce a new "Primary Mail Service for Staff Members and Students". Based on our experience with both systems, we decided to use open source software similar to the student mail system, but employ virtual user with Dovecot's native LDAP support instead of using support provided by the operating system. To have finer control to LDAP attribute, the system also contained dedicated OpenLDAP servers that would receive user information from the central identity management system. We used virtual machines proactively to implement all of them.

Since we would support more than 30,000 users, we decided to incorporate dedicated hardware load balancers with SSL accelerator. Also, we knew it was hard to build a good spam filtering system using only open source software, so we decided to use Barracuda spam firewall, which employed a per-server license policy. We didn't want to be annoyed by a per-user license policy anymore. Figure 3 shows the overview of the entire system.

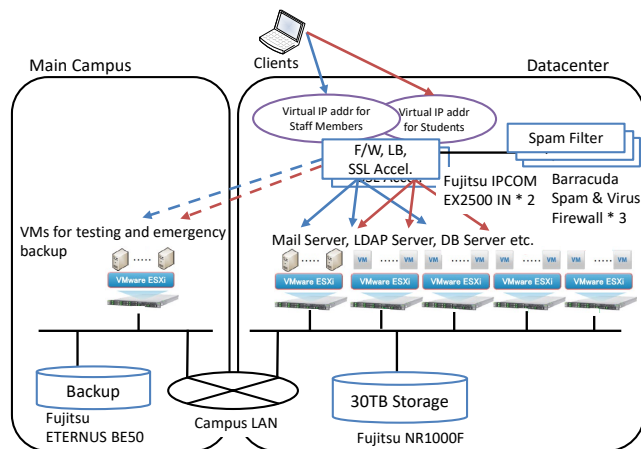


Figure 3. Overview of New Primary Mail System

As you see, the main function of the system resides in a datacenter. We had a separate project to borrow a rack in the datacenter and move the inter-campus core switch and border router there for network resiliency, and we also decided to place the new mail system next to the network rack. In case of emergency, we placed a minimum amount of server resource and a data backup appliance on our own campus. The racks in the datacenter are

considered as one of the remote campuses, and these are connected via 10GbE inter-campus LAN.

We tried to provide standard functionalities such as SMTP, POP3, IMAP4, SSL ports for all of these protocols submission port support, and STARTTLS.. The mailbox quota became 1GB without expiration. We also decided to provide webmail, but due to budget constraints, we decided to use free Squirrelmail, a decision which created problems later. While we used the same infrastructure and software to provide services for staff members and students, we used different domain names and separate virtual machines for each service.

For the previous student mail service, the name-based alias address system was developed in-house. This time we wrote detailed specifications and asked the system provider of the new mail system to re-implement it from scratch. The new alias address service supports both students and staff members, to solve our long-standing problem of handling middle names. The default mail address for a staff member was generated mechanically from his/her alphabetical name including first, last, and middle name, but some people didn't want to disclose their middle names for religious reasons.

4. MIGRATION

To migrate users to the new system, we needed to move their mailboxes, configurations, and, in some cases, asked them to change the configuration of their mail clients.

4.1 Mailbox

We faced a severe problem migrating staff members' email service, because there was no mailbox migration option from Mirapoint to other mail system. The vendor told us that it was impossible to export users' mailboxes from the system. Technically speaking, we could migrate mailboxes using IMAP, but we didn't have enough IMAP licenses to do that. To overcome the problem, we decided to operate both the old and new system simultaneously for about three months (until most messages expired on old system), and asked users to import their messages (retrieved by POP) to the new system via (now freely available) IMAP.

On the other hand, migration of students' email should be no problem, because both the old and new system use the same software. Of course we needed to test it because of differences between versions and a slight change of file hierarchy. We believed any issues would be resolvable, and the migration test finished without problem.

4.2 Configuration

For staff members, we could only migrate mail-forwarding settings from the old system to the new system. Mirapoint's enterprise webmail had various useful features such as an address book with many items and message filters, but none of these features were importable to the new service.

For students, we migrated mail forwarding settings and alias address settings. The new system employed LDAP-based virtual users, so we needed to convert ".forward" files to LDAP entries. We didn't provide webmail to students as part of the previous mail service, so we didn't need to worry about migrating webmail settings.

4.3 User side settings

As mentioned in section 4.1, we decided to operate the old and the new system for staff members simultaneously, so we had to ask users to change their mail client configuration. There were two options. One was to add a new IMAP user in addition to the previous POP user. By doing this, a user could easily migrate their messages to the new service. Another option was to change the POP server setting of the existing account. By doing this, a user could continue to use the same local mail folders. For the outgoing SMTP server, we kept compatibility with old service and changed the IP address of the server name.

Since we were able to migrate student settings, we announced that students didn't need to change their settings. Actually we overlooked a couple of the old server's configuration settings such as a loose connection security setting from the campus LAN. Some students still relied on such settings, so we had to make the new server a bit less secure than expected. Due to time constraints, we decided to address the issue later.

5. OPERATION ISSUES

In this section, we explain some problems we experienced after the system migration, current status and future works.

5.1 Troubles after Migration

We tried to test the system migration before the real migration as much as possible, but there was a simultaneous migration of the identity management system so we couldn't test the new ID system prior to the actual migration of our mail system. The mail system was the first and one of the biggest service providers using the new ID system, and there were several login failure complaints due to password migration failures (more than 3,500 users were affected). There were multiple causes of the failure and it took a couple of weeks to fully resolve the problems.

Another issue was forwarding setting migration for students. We believed that interpreting ".forward" file was trivial, so we asked the vendor to do the conversion. Actually they didn't fully understand the syntax of ".forward" and failed to migrate settings of about 100 users. Maybe we should have done the conversion ourselves.

We had few candidates for webmail software other than Squirrelmail due to budget and time constraint, and we didn't realize until the migration that Japanese language support of Squirrelmail was mediocre. There was an unofficial patch to support Japanese language, but it didn't support UTF-8 so it couldn't display some kanji characters (especially for user's real name) properly. Also it didn't work properly when using different UI language other than Japanese.

There were other minor issues such as tweaking the spam filter configuration, SMTP message size limitation, enforcing encryption for LOGIN/PLAIN authentication protocol, failed migration of alphabetical name, and so on.

5.2 Current Status and Future Works

Currently most of the severe issues have been resolved, but we still have some outstanding issues to resolve. One of these issues is an account activation timing issue. Due to some miscommunication with the system provider, email accounts become active after each user "activates" their account via the activation web page. This is an issue for new users, especially

students, because as soon as Student IDs are assigned, professors can send email to a Student ID based email address. The delivery fails due to non-existent users if the student hasn't activated his/her account yet. It isn't an ideal situation and we expect the problem will be fixed soon.

Another ongoing issue is webmail. There are still several issues with Squirrelmail (mainly about internationalization). We now realize that users use webmail far more than we expected, so we may need to consider changing webmail software.

Due to budget and time constraint, we deliberately omitted some functionality such as message filters and vacation function. One consideration to solve both issues is adding "sieve" function using Dovecot Pigeonhole plug-in, but again internationalization is a big concern.

6. CONCLUSION

In this paper, we explained about our old email service, the design and implementation, user migration, current status and future works of our new mail service. Email service seems one of the mature technologies of the Internet, but we still need much detailed consideration to design and implement "sane" email service. The operation of the new mail system has just started, and we continue to improve the service.

7. ACKNOWLEDGMENTS

Our thanks to all the users using our mail service, and staff members of the Primary Mail Service working group and the Authentication Infrastructure working group to develop and maintain these systems in Information Infrastructure Initiative of Kyushu University.

8. REFERENCES

- [1] Fujimura, N., Togawa, T., Kasahara, Y., and Ito, E. 2011. Primary Mail Service for students based on their names. *IPSJ SIG Technical Report*, Vol. 2011-IOT-14, No. 10, 1-6.
- [2] Fujimura, N., Togawa, T., Kasahara, Y., and Ito, E. 2012. Introduction and experience with the Primary Mail Service based on their names for students. In *Proceedings of the SIGUCCS 2012* (Memphis, TN, October 17 - 19, 2012). ACM, New York, NY, 11-14. DOI=<http://dx.doi.org/10.1145/2382456.2382460>.
- [3] Ito, E., Kasahara, Y., and Fujimura, N. 2009. The current status of e-mail services for staff members in Kyushu University. In *Proceedings of ECIP 2009*, D3-4.
- [4] Ito, E., Kasahara, Y., and Fujimura, N. 2013. Implementation and operation of the Kyushu university authentication system. In *Proceedings of the SIGUCCS 2013* (Chicago, IL, November 3 - 8, 2013). ACM, New York, NY, 137-142. DOI=<http://dx.doi.org/10.1145/2504776.2504788>.
- [5] Ito, E., Kasahara, Y., and Fujimura, N. 2012. A study of LDAP load balancing for University ICT services. *IPSJ SIG Technical Report*, Vol. 2012-CSEC-57/2012-IOT-17, No. 11, 51-56.
- [6] Kasahara, Y., Ito, E., and Fujimura, N. 2013. Gulliver's toss: Google's chronic big load to university mail server and its sudden resolution. In *Proceedings of the SIGUCCS 2013* (Chicago, IL, November 3 - 8, 2013). ACM, New York, NY, 169-174. DOI=<http://dx.doi.org/10.1145/2504776.2504815>.