THE USENIX SIG FOR

[sage]

SYSADMINS

ISBN-13: 978-1931971492
ISBN-10: 1931971498
90000
9 781931 971492

15 Nick Christenson and Brad Knowles Internet Postmaster: Duties and Responsibilities

15

*Short Topics in*
System Administration

*Jane-Ellen Long, Series Editor*

# Internet Postmaster: Duties and Responsibilities

*Nick Christenson and*
*Brad Knowles*

THE USENIX SIG FOR

[sage]

SYSADMINS

## Booklets in the Series

**#15: Internet Postmaster: Duties and Responsibilities**
Nick Christenson and Brad Knowles

**#14: System Configuration**
Paul Anderson

**#13: The Sysadmin's Guide to Oracle**
Ben Rockwood

**#12: Building a Logging Infrastructure**
Abe Singer and Tina Bird

**#11: Documentation Writing for System Administrators**
Mark C. Langston

**#10: Budgeting for SysAdmins**
Adam Moskowitz

**#9: Backups and Recovery**
W. Curtis Preston and Hal Skelly

**#8: Job Descriptions for System Administrators, Revised and Expanded Edition**
Edited by Tina Darmohray

**#7: System and Network Administration for Higher Reliability**
John Sellens

**#6: A System Administrator's Guide to Auditing**
Geoff Halprin

**#5: Hiring System Administrators**
Gretchen Phillips

**#4: Educating and Training System Administrators: A Survey**
David Kuncicky and Bruce Alan Wynn

**#3: System Security: A Management Perspective**
David Oppenheimer, David Wagner, and Michele D. Crabb
Edited by Dan Geer

**#2: A Guide to Developing Computing Policy Documents**
Edited by Barbara L. Dijker

**#1:** *See #8 above*

15 *Short Topics in*
**System Administration**

*Jane-Ellen Long, Series Editor*

# Internet Postmaster: Duties and Responsibilities

**Nick Christenson and Brad Knowles**

**Published by the USENIX Association**
**2006**

## About SAGE

SAGE is a Special Interest Group of the USENIX Association. Its goal is to serve the system administration community by:

- Offering conferences and training to enhance the technical and managerial capabilities of members of the profession
- Promoting activities that advance the state of the art or the community
- Providing tools, information, and services to assist system administrators and their organizations
- Establishing standards of professional excellence and recognizing those who attain them

SAGE offers its members professional and technical information through a variety of programs. Please see http://www.sage.org for more information.


ISBN 1-931971-49-8

To purchase additional copies and for membership information, contact:

The USENIX Association
2560 Ninth Street, Suite 215
Berkeley, CA USA 94710
orders@usenix.org
http://www.usenix.org/

First Printing 2006

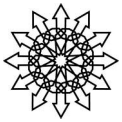

**Contents**

## *Acknowledgments*

The authors received the benefit of insight from a number of outstanding technical reviewers for this booklet. Specifically, we'd like to thank Strata R. Chalup, Esther Filderman, Jim Hickstein, Jim Lawson, and especially Gregory Shapiro for their comments and suggestions, which were extremely helpful. Of course, we take full responsibility for any errors, omissions, or just plain silliness.

We'd also like to thank the USENIX staff for helping us develop our raw ideas into this fully formed booklet. We'd like to explicitly thank copy editor Steve Gilmartin and, especially, series editor and jack-of-all-trades Jane-Ellen Long for making the process of writing this booklet as easy as we could possibly imagine.

# Introduction

Electronic mail was the first Internet "killer application." Ask people what they do first when they connect to the Internet, and most of them will answer, "Check my email." Most organizations consider email to be a critical communications channel, as important as or even more important than the telephone or physical mail. Consequently, to maintain an email service as a robust and effective means of communication, organizations need to take the administration of their email systems seriously.

We can divide the tasks required to properly maintain an email service into two distinct types of roles: the email administrator and the postmaster. Many times, these two functions are performed by the same person or people, but in larger organizations they may be split out into separate roles, or may even be subdivided further.

We consider the role of the email administrator to be to perform the technical tasks necessary to keep the email service functioning properly. Such duties include, but are not limited to, ensuring that the service is up and accessible, keeping email queues clear, and maintaining and updating the service's hardware and software.

On the other hand, we consider the postmaster role to include the personal interactions and policy work that surround the email service. These duties are as important as those of the email administrator, but they are less often discussed in the system administration literature. The purpose of this booklet is to discuss these roles and provide some suggestions on the best current practices (BCP) for those who hold postmaster positions for Internet-connected organizations.

We don't know for certain when the term "postmaster" was first used as it applies to electronic mail. The first mention of this term in this context that we could find is in section 6.3 of RFC 822, "Standard for the Format of ARPA Internet Text Messages" [Crocke82]. This document requires that every domain should have a mailbox called "postmaster." The RFC goes on to specify that the user portion of the "postmaster" email address must be case insensitive and should be read by someone who is responsible for the domain's email service. Further, it specifies that this is the proper channel for inquiries regarding problems in email communication with that domain. As far as we're concerned, this is the origin of the postmaster role, and the term's definition extends from the notion that the postmaster is the person who reads an email server's postmaster mailbox.

## What This Booklet Is About

This booklet suggests guidelines for defining the duties and responsibilities of the postmaster at various types of Internet-connected sites. The words "duties" and "responsibilities" in the title of this booklet have been chosen carefully by the authors. The role of the postmaster encompasses tasks that must be conscientiously performed at regular intervals to satisfy the demands of one's customer base and ethical obligations both to one's customers and to the Internet community.

The specifics of these duties and responsibilities can vary depending on the exact nature of the organization in question. Even among organizations of similar type and purpose, we would expect that the job functions of the postmaster can vary considerably. This booklet will explore some of the disparate aspects of this position, but each organization will have to determine which of the suggestions included here are most appropriate for adoption in their particular circumstances.

One of the difficulties in being postmaster is reconciling conflicts between the duties and the responsibilities of this position. It is by no means trivial to balance the objectives of the organization with the ethical obligations called for by the position. These sorts of conflicts are never easy to resolve, and no document can cover these situations exhaustively, but we hope that this booklet will help to place some of the situations that each postmaster faces in a useful context.

## Future Relevance

The duties and responsibilities of the Internet postmaster have changed dramatically over the past 25 years, and we see no reason to think that this role won't continue to change as the Internet evolves. The authors believe that the suggestions made in this document represent the consensus of the industry's best current practices for postmasters. Just as we believe that what would be considered BCP has changed over the past few decades, we recognize that BCP in the future will likely be quite different as well. Therefore, as the Internet and email service evolve, we're confident that some of the specific suggestions made in this booklet will need to be adjusted.

Even now, the situations faced by postmasters at distinct sites may render suggestions that are perfectly valid for one organization inappropriate for another. Nothing we say here serves as a replacement for the judgment of an experienced postmaster who has a thorough understanding of the circumstances surrounding their particular situation.

This is an issue we have had to deal with as we reviewed the older Internet RFCs. Many of these documents have not been updated to reflect a time in which spam and email-borne malware are commonplace. Consequently, in order to adapt the valuable principles contained in these earlier documents we must attempt to divine the intent of the authors and successfully adjust them to our present circumstances. We expect that not everyone will completely agree with our recommendations on how best to do this.

## Things This Booklet Does Not Cover

While this booklet covers many aspects of email service and operation, there are two areas that we specifically do not intend to cover. We feel it is worthwhile to explicitly list those here in case the reader becomes distracted by their absence later on.

While this booklet does discuss the legal aspects of postmaster duties, we do not provide legal advice. Although we have significant experience supporting email services, neither of us has any specific expertise in legal matters. Further, legal regulation on this topic is so dynamic and so location-dependent that any attempt to capture its state would be futile.

In this book we point out some areas in which postmasters may become embroiled in legal issues, but we do not want to suggest specific resolutions to these issues. Instead, the person who encounters these situations in the postmaster role should seek out advice from legal professionals who understand the specific situation in which any particular organization finds itself. In fact, we strongly recommend that those in the postmaster role should initiate and maintain a dialogue with an organization's legal counsel in order to help stave off such problems before they occur.

Another set of topics that will not be covered in this booklet is that of solutions to problems postmasters are likely to face based on specific email software. Each of these almost certainly has several books, Web sites, and other information sources available that discuss how to resolve implementation issues. Consequently, there's no way we could do justice to the many platforms that exist in a document of this size, so we won't even try. This booklet focuses on making qualitative suggestions to postmasters. We'll leave the task of making implementation suggestions for email administrators to other sources.

At times, we will refer to the documentation related to a specific software solution or methodology, but only if that source has something relevant to say about providing email services that is more general than the solution endorsed by that particular platform. Nothing we say here should be construed as endorsement or condemnation of any specific electronic mail solution.

# 1. Postmaster Basics

This chapter discusses some of the overriding issues associated with being an Internet postmaster. Subsequent chapters will deal with specific situations, but we wanted to begin with a broad overview of what we think the general mind-set of someone in the postmaster position should be.

As should be clear from the title of this booklet, we believe that the position of postmaster involves not only a set of tasks to be performed but also a set of responsibilities. These responsibilities have a strong ethical component, so we believe it would be most appropriate to discuss these in light of the SAGE System Administrators' Code of Ethics, ratified in September 2003. The Code of Ethics can be found at http://www.sage.org/ethics/.*

## The System Administrators' Code of Ethics

The System Administrators' Code of Ethics contains ten points. In this section, we will examine each of these points as we believe they apply to the role of Internet postmaster.

### Professionalism

Postmasters have a duty to the organizations they represent. The best interests of the organization for which one works should be considered at all times.

Postmasters deal with issues that need to be addressed from all types of customers and external domains. It is important not to let one's personal feelings interfere with the performance of one's job. Each situation should be evaluated and addressed solely within the context of the postmaster role.

### Personal Integrity

Everybody makes mistakes, and sometimes these will negatively impact what others are trying to accomplish. If this happens, don't lie about the situation, don't exaggerate, and don't try to deflect blame. If the problem is due to one's own error, admit it and then focus on addressing the situation.

It may be in the organization's best interest to not fully disclose the complete scope of certain problems. There is a fine line between ethical spin and unethical deception, and each person must walk that line themselves. Sometimes a conflict between profes-

* Throughout this booklet, any trailing punctuation—e.g., the period at the end of the sentence above—should not be read as part of the URL.

sionalism and personal integrity arises that cannot be easily resolved. However, these conflicts almost always can be settled without compromising the truth. In any case, one should always strive for the most ethical solution to any such dilemma.

As we all have many associations in our lives, it is always possible that conflicts of interest will arise. It is important to disclose these situations when they come up. If a situation develops where there is a risk of bias, consult with a colleague or supervisor about the situation. When an additional party is made aware of the situation before action is taken, one reduces the risk of bias and also provides some protection against the appearance of impropriety, and this is also important.

**Privacy**

As with other forms of communication, electronic mail should be afforded as much privacy and protection from prying eyes as possible. In many organizations and situations, people other than the sender and recipients will justifiably need to view a given piece of email. However, even when this occurs, the scope of the examination should always be as narrow as possible. Even if a broad examination of email is allowed by policy, such an imposition should only be performed if it is absolutely necessary and to the most limited extent possible.

On occasion, an otherwise private communication will be accidentally disclosed to someone other than the intended recipient. If this happens, the sender should be informed and the confidentiality of the message should be kept, as long as doing so does not violate laws, professional responsibility, or some higher ethical necessity.

At the same time, this does not justify a cavalier attitude toward email privacy. Considerable effort is justified in minimizing the risk that unauthorized disclosure of electronic mail incurs.

**Laws and Policies**

It is important to keep abreast of legal developments surrounding electronic mail. It is also important to periodically discuss these matters with one's legal counsel, both to understand what the situation is now and to plan for possibilities the future might bring.

Internet postmasters should take time to familiarize themselves with the laws and regulations that apply to the electronic mail handled by the organization they represent. Postmasters should also possess a thorough knowledge of their organization's email policies. Postmasters will likely be consulted on the wording and application of these policies, so it is critical that they be able to speak fairly and authoritatively about them.

**Communication**

It is a postmaster's responsibility to keep customers informed accurately and in a timely manner about the status of the email service. It is equally important for postmasters to keep themselves apprised of how the performance of the email service is perceived by their customer base.

Much of the communication with the postmaster occurs through the postmaster mailbox. This mailbox should be read frequently. Every attempt should be made to

identify every legitimate request for action that is made through this channel. These requests should be taken seriously and, if appropriate, responded to and handled in a reasonable amount of time.

**System Integrity**

Email is a vital service in most organizations. There are many best practices of system administration that can contribute to the robustness of any information service. Be cautious when making changes or upgrades. Carefully test all changes after they have been made. Meticulously document the system so that other people can understand its operation.

**Education**

It is the postmaster's responsibility to understand what is happening in the email world, based both on the perception of one's customer base and on what is happening out on the Internet. What are the current threats? What technical advances are being discussed or made? What problems might one's own organization be causing for other sites? As appropriate, this information should be made available to one's customer base.

**Responsibility to Computing Community**

Each Internet domain is a neighbor to every other as part of a vast community. As with physical neighborhoods, this entails certain responsibilities that an ethical member of that community should be willing to meet.

The postmaster has a responsibility to address good-faith complaints made regarding interactions with one's own domain. It is also a responsibility to "play nice" with the rest of the Internet. For example, don't intentionally configure software to "slam" other domains with email, especially if one runs a high-volume service.

Deploy software and systems that respect the Internet standards and are designed to be compatible and cooperative with other software on the Internet. Support those software vendors and developers that build standards-compliant systems.

**Social Responsibility**

The Internet is a community only to the extent to which people become involved in making it one. There are many avenues by which a person can help to improve this community. Participating in mailing lists, technical groups (e.g., the IETF), and USENET news groups by contributing the knowledge that one has accumulated is a good way to give back to the Internet community. Certainly we ourselves have benefited from the advice we have received from folks who took time to answer our questions out of the kindness of their hearts, and we feel a sense of responsibility to provide the same service when we are able.

Internet search engines provide amazing access to a wide variety of often useful information sources. The other side to this is that all people have to do to provide value back to the Internet is to post an article on some Web page, and interested parties will eventually find it. The barrier to entry has never been lower for sharing the wisdom one has accumulated as a consequence of solving some difficult problem.

Legislative bodies all over the world are wrestling with the issues the Internet pres-

ents. Sometimes they make a step forward, sometimes they move society in the other direction. They need the assistance of a community that deeply understands the technical issues the future presents.

At work, it is important to obey the laws that affect one's daily activities. If it is legal to do so, one has permission, and one's beliefs coincide with the interests of the organization for which one works, lobby governmental representatives on issues of the day. If it is not legal or appropriate to do this from work, then do so from a non-work account, making it clear that this petition is being made as a private citizen.

**Ethical Responsibility**

The SAGE System Administrators' Code of Ethics states an SA's ethical responsibilities:

- I will strive to build and maintain a safe, healthy, and productive workplace.
- I will do my best to make decisions consistent with the safety, privacy, and well-being of my community and the public, and to disclose promptly factors that might pose unexamined risks or dangers.
- I will accept and offer honest criticism of technical work as appropriate and will credit properly the contributions of others.
- I will lead by example, maintaining a high ethical standard and degree of professionalism in the performance of all my duties. I will support colleagues and co-workers in following this code of ethics.

We don't have anything to add to these statements. We believe they nicely sum up the responsibilities of any system administrator, and that they are as applicable to the Internet postmaster as they are to any other position in any organization.

## Looking at RFC 1173

RFC 1173 is titled "Responsibilities of Host and Network Managers: A Summary of the 'Oral Tradition' of the Internet" and was written by J. Van Bokkelen [VanBok90]. Even though it was published in 1990, this is still the most extensive discussion of the role of the Internet postmaster in any Internet standards document. Moreover, the tone in which this document addresses its topics wonderfully represents the best of the spirit of the Internet. Consequently, we believe it deserves special examination here.

Section 4 is titled "postmaster@foo.bar.baz," and we reproduce that brief section here in its entirety:

> 4. postmaster@foo.bar.baz
>
> Every Internet host that handles mail beyond the local network MUST maintain a mailbox named "postmaster." In general, this should not simply forward mail elsewhere, but instead be read by a system maintainer logged in to the machine. This mailbox SHOULD be read at least 5 days a week, and arrangements MUST be made to handle incoming mail in the event of the absence of the normal maintainer.

> A machine's "postmaster" is the normal point of contact for problems related to mail delivery. Because most traffic on the long-haul segments of the Internet is in the form of mail messages, a local problem can have significant effects elsewhere in the Internet. Some problems may be system-wide, such as disk or file system full, or mailer or domain name server hung, crashed or confused. Others may be specific to a particular user or mailing list (incorrect aliasing or forwarding, quota exceeded, etc.).
>
> In either case, the maintainer of a remote machine will normally send mail about delivery problems to "postmaster." Also, "postmaster" is normally specified in the "reply-to:" field of automatically generated mail error messages (unable to deliver due to nonexistent user name, unable to forward, malformed header, etc.). If this mailbox isn't read in a timely manner, significant quantities of mail may be lost or returned to its senders.

Let us examine this passage in detail and consider how one ought to apply its statements, especially with regard to today's Internet.

The first paragraph states that the "postmaster" mailbox should exist. This is a pretty straightforward statement, but it restates what has been said in RFC 822 and other places, namely, that this is a focal point for communication between domains on the topic of electronic mail.

RFC 1173 states that for each server that receives email, one should log on to that machine and read the postmaster mailbox locally rather than forwarding that email elsewhere. We believe that these days this is no longer a reasonable requirement in many environments. However, behind this statement are some good points that should be taken seriously by every postmaster.

Many organizations support a very large number of servers that are configured to receive email. This quantity by itself may make it impractical to log on to each machine individually to read each postmaster mailbox. If we elect not to follow RFC 1173's advice because of the quantity of servers we would have to support in this manner, we might be well advised to ask ourselves if the problem might not be in our architecture, rather than in Van Bokkelen's advice.

That most UNIX or UNIX-like systems come configured to receive email from the outside world doesn't mean that this is a good configuration in which to deploy most computers. Many experienced system administrators, including the authors of this booklet, have opined that by default computers should be configured to not accept email from other servers. As just one example of this advice, we cite *The Practice of System and Network Administration* by Limoncelli and Hogan [LimHog02]:

> Avoid delivering mail to people's desktops and make sure their mail clients are configured to send email by contacting a mail relay rather than routing mail themselves. Desktops should not even listen on the SMTP port. Servers that are not part of the email service should be configured the same way as the desktops. [section 19.1.4, p. 409]

Configuring computers in the manner advised in this excerpt leads to a more secure network that is simpler and easier to maintain. It also reduces the number of machines to which postmaster email may be sent.

While following this practice would greatly reduce the number of hosts that Van Bokkelen would have postmasters log in to, we still believe that his specific advice on this topic is outdated. Heck, there are doubtless professional system administrators out there who have never read email from the command line in their careers. Also, if one maintains a significant number of servers that do receive email, it may be expedient to forward all the postmaster email to a central point. However, there are good reasons why RFC 1173 admonishes against this practice, and they should be carefully considered by any responsible postmaster.

First, if many postmaster mailboxes are aggregated in a central location it is important to be able to easily determine to which email server any particular piece of postmaster email was originally directed. It does no good for an organization to be informed that they have a problem with their email service if they can't tell to which particular server the warning refers.

Generally, we believe it is perfectly acceptable to aggregate postmaster mail for multiple domains that are served by the same computer into one mailbox, as long as the postmaster can resolve to which domain each postmaster email refers. We also believe that it is perfectly acceptable to aggregate postmaster email for multiple servers that serve the same domain, as long as the postmaster can resolve to which server each email refers. Generally, we do not recommend aggregating the postmaster mailbox for multiple domains on multiple servers, although we would relax our stance on this if special tools were available to help keep track of the particular server and domain.

A second important issue is that if the postmaster email is forwarded off of a given server, how will we recognize when that server's email sending capability is broken? This is a distinct advantage to reading all postmaster email local to the server to which it was sent. If postmaster email is forwarded, then it is important to create some mechanism by which such a problem would be detected.

One way to do this would be to have the server itself generate periodic email to send to the postmaster. Some sort of periodic update or usage summary works fine for this. But how often should such a test message be sent? We would recommend roughly the same interval as that at which the postmaster mailbox is read. For example, if postmaster email is read daily and postmaster email is centrally aggregated from a given server, then the ability of that server to successfully send email should be tested at least daily as well.

RFC 1173 specifies that email should be read at least five days a week. If we were to encapsulate our recommendations into a single number, then this is probably a pretty reasonable one to select. We recommend that an organization ensure that its postmaster email be read at least once a day on any day in which that organization does business.

For most educational, governmental, and business organizations, this is a reasonable schedule. If an organization conducts significant business on every or nearly every day (e.g., an ISP or an online retailer), then every workday will be pretty much every day, and postmaster email should probably be read every day. For many very busy sites that critically depend on email as part of their business, merely once a day, every day may not suffice. The postmaster mailbox will need to be read more frequently, perhaps even continuously.

For small organizations where email is not a mission-critical service, it's probably not necessary to examine the postmaster mailbox every day. In some cases, reading it every few days or even once per week may be entirely reasonable. However, as the RFC suggests, if there will be periods in which the usual postmaster will be out of contact for longer than the normal frequency with which the postmaster mailbox will be read, someone should be appointed to the position of "substitute postmaster" who will become aware of email situations at the usual interval until the primary postmaster returns.

RFC 1173 mentions long-haul delivery problems. These are certainly less prevalent and less distance sensitive than they were when the RFC was written, but they do still occur. These situations aren't just about routing problems and line cuts but can also be caused by DNS errors, denial-of-service attacks, and botched upgrades. Moreover, even though the Internet and its services are generally more robust than they were in 1990, local issues can still have global implications.

If one experiences these problems and they can't be resolved on one's own end, the RFC informs us that contacting the postmaster at the other site is an entirely appropriate course of action. Therefore, postmasters should expect email regarding these sorts of situations, as well as many others, to land in their inboxes. Some of these queries will be of critical importance and should be treated as such.

Whether making or receiving postmaster queries, it's important to remember that what might be a critical problem to one site may be incidental to another. Yes, this can be maddening, and every postmaster should strive to respect the level of concern demonstrated by their counterparts. At the same time, postmasters should be aware that the people with whom one is communicating may not have the same priorities.

As an example, let's suppose that a single relatively small domain is unable to send email to one of the large ISPs or email service providers for some period of time due to configuration at the ISP's end. This problem is likely to be a relatively low priority for the ISP, as it may affect a tiny portion of its user base, and other issues may have a larger impact. On the other hand, a significant percentage of the small domain's email may be addressed to this large service, and the backlog may be causing the small domain's email service considerable distress.

In this case, it's important that the small domain's postmaster realize that the large service may have higher priority issues to deal with, while the ISP should understand that what's a mere annoyance to it might be a serious problem at a site where fewer resources are available to handle exceptional situations.

Postmasters should also expect to receive automated responses to failed email deliveries in their mailboxes. Not every one of these indicates a genuine problem. Many will be due to "double bounces" caused by misconfigured software, spammers, or malware. It is important to try to reconcile problems caused by the first of these. The latter two, once identified, are usually best discarded unless they are internally generated, in which case they should be addressed with dispatch.

In many cases these automated messages will be repeated, sometimes many times. Having identified the cause of the problem and taken appropriate steps, if any, to rectify it, the postmaster can reasonably ignore the automatically generated duplicate messages as long as one can be confident both that (1) the messages were automatically generated and (2) they are, indeed, duplicate reports of a problem already addressed.

Van Bokkelen concludes with the statement, "If this mailbox isn't read in a timely manner, significant quantities of mail may be lost or returned to its senders." Postmasters owe it to their customers and to the rest of the Internet to take their roles seriously. If email is important, then these duties are important and should be treated as such. No legitimate email should ever be lost. All legitimate email should be delivered, properly bounced, or end up in the postmaster mailbox for human resolution.

Each postmaster should try to set aside regularly scheduled time to handle postmaster duties. Scheduling these activities gives them weight. If management isn't informed that these tasks are important, they won't know to treat them as such. If they refuse to make postmaster duties a priority, they should be informed that the quality of email service is likely to suffer as a result. If they still don't believe that devoting time to these issues is worthwhile, so be it. Management gets to make these sorts of decisions, but postmasters need to make sure they've been diligent and honest in informing management of the risks.

We recommend that postmaster mailbox triage be performed early in one's work shift. Routine email issues may be eclipsed by more pressing responsibilities later in the day, but a person can't really identify the biggest problem they face until they have a complete list of those problems. If the postmaster issues are relatively benign, it's entirely appropriate to address them after higher priority situations have been resolved. However, if a considerable amount of time will pass before they can be addressed, it might be best to inform those who are impacted by these issues about what problems they might encounter and roughly when they might expect a resolution to these issues.

## Interactions with Other Services

Electronic mail does not exist in isolation. To function properly, this service depends on many other services. These systems may or may not be under control of the postmaster, depending on the circumstances. In fact, in some organizations, the postmaster and the maintainers of many of these services may not be organizationally closely tied together. Since email depends on these services in order to function properly, we feel it is worth enumerating the interfaces between these services and the email system. Even if the maintainers of these services don't answer to the same bosses, in order for post-

masters to do their jobs properly, good lines of communication must be maintained so that problems that cross organizational boundaries can be resolved quickly and efficiently.

**Email Administrators**

Organizations that support a large number of email customers will often split the postmaster and email administrator job functions. Even so, these two groups will still need to communicate efficiently and often in order to present a coordinated service to their customers.

A typical division of responsibilities between these two groups tends to assign communications with customers and external sources, as well as the more routine email maintenance activities (e.g., creating, deleting, and modifying account details), to the postmaster person or group. By contrast, the email administrators would focus on architectural issues, software, and technical changes that would need to be made that affect the servers themselves.

It should go without saying that everything that either group does of any significance must be coordinated with the other. If a significant change will be made to the operation of the email service, the email administrators need to inform the postmasters of all the details so they can communicate the situation to those who would notice an outage. As those performing the postmaster functions are made aware of problems people are having with the email service, this data needs to be aggregated and passed on to the email administrators for action.

**Other System and Network Administrators**

In some organizations, a small number of people are collectively responsible for every aspect of the information technology infrastructure. In others, service maintenance is stratified into fine-grained structures. At some sites we might find many different roles for IT professionals, including application administrators, who manage specific software packages such as email services; system administrators, who are responsible solely for the underlying operating systems; hardware administrators, who maintain the equipment on which the applications and OSes reside; network administrators, who handle data as it moves from server to server; and facilities managers, who are responsible for the maintenance of the data centers and other server environment issues.

If postmasters do not have the ability to directly effect changes to the email server operating systems, hardware, networking, or facilities, they will need to communicate their needs carefully to those who are permitted to make such changes. Similarly, since any change to the underlying system can influence the primary applications running on a server, changes to the foundations on which an application runs cannot be safely adjusted without coordination from the application owners.

**Name Service Administrators**

Of all the external Internet services, the one that electronic mail most relies upon must be the Domain Name System. Internet email requires DNS to route messages between

domains across the Internet. If DNS isn't functioning, email doesn't flow. Email and DNS are so tightly bound together that email is the only Internet service with its own widely deployed DNS record type, the MX record.

In order for email to be routed properly, a domain must have its DNS records set up correctly and have access to the DNS information for external domains. Internally, DNS or another service, such as LDAP or Active Directory service, may be used to route email. If a name service is down or a mistake is made in an organization's domain records, the results on the email service for that domain could be disastrous. Moreover, because DNS requests are cached around the Internet, even after an error is corrected it may take a considerable amount of time for the after-effects to subside.

A relatively new way that DNS and email interact is in the new anti-spam techniques that are being adopted by the Internet community. Sender Policy Framework (SPF) [WonSch06] and DomainKeys Identified Mail (DKIM) [Allman05] both store vital information in DNS TXT records. If these entries are misconfigured, domains around the Internet will likely start rejecting legitimate email. If either of these email authentication techniques are adopted, it is crucial that a site that deploys them ensures the integrity of these DNS records.

DNS changes that may affect mail routing need to be carefully coordinated with the postmaster. As an email service evolves, the postmaster and email administrator will make requests of the name service administrator that will need to be deployed carefully.

**Network Security**

Network security has become such a critical and complex set of tasks that in all but the smallest organizations it is typical that these functions will be split out into their own group. This group's functions typically include malware detection, intrusion detection, and system auditing. Often they include handling abuse complaints and spam prevention as well. These are all areas in which network security and postmaster duties intersect.

As we all know, email is one of the top vectors by which malware may be introduced into an Internet-attached network. Consequently, most organizations provide some sort of protection against such attacks, and in many organizations different groups are responsible for providing email service and email protection methods. Needless to say, these two groups must communicate upgrades and changes that may affect either email routing or email filtering.

The network security group may also be responsible for preserving the integrity of company confidential information. Therefore, they, or some other group, may be performing stateful examination of outgoing and/or incoming email for messages that violate various organizational policies. Changes to these mechanisms also need to be closely followed by the postmaster.

**Internal User Databases**

Internal email routing and email acceptance decisions are all partially based on some form of user identification. In many cases, the primary copy of the database that lists

valid internal users and their properties is not stored locally on the email servers, and it is often not maintained by the same people responsible for the email service. Clearly, if this information becomes unavailable or is unreliable, email routing will be significantly affected.

There are many ways this information can be stored and accessed. Along with the ubiquitous DNS, three other common mechanisms used in email routing are NIS, Active Directory, and LDAP. An email system that uses any of these services to assist in the sending or reception of email has yet one more external dependency, such that if problems occur with this service, email flow will be disrupted.

It may be possible to mitigate these potential effects by maintaining an image of this central database on or in closer network proximity to the email servers that depend on it. Even though this will likely increase the robustness and improve the performance of access to this data, doing so is not without disadvantages.

Chief among these disadvantages is the potential for the image to become outdated. Consequently, special effort needs to be made to ensure that this doesn't happen and to minimize the problems that result if database drift does occur.

Another disadvantage is that these databases often contain additional information that an organization might consider especially sensitive. For example, an LDAP database may include internal information about an organization's membership (e.g., phone numbers, titles, addresses, etc.), and LDAP, Active Directory, and NIS databases may all contain password information. Every additional computer that has access to such information reduces the effective security of that information. This is especially true if an email server with access to this information resides on an untrusted network, such as an email gateway that is part of an organization's Internet firewall.

If a necessary internal database is unavailable to an email server on a temporary basis, the correct action for that service is to locally queue email to the extent possible and to return a 4xx SMTP code indicating a temporary failure to those connections for which messages cannot be queued. If such a local resource becomes temporarily unavailable and email bounces as a consequence, that server is almost certainly configured incorrectly.

**Legal**

It's an unfortunate reality that email service now intersects a great deal with legal issues. As electronic data storage and communication have become more pervasive and more critical to how our society functions, it is inevitable that information technology would attract the attention of those who make laws and regulations. Every jurisdiction of which we are aware has laws covering electronic mail; as a consequence, there are now legal ramifications to running an email service.

As this is a new arena for legislators, the laws and their interpretations as they apply to email are changing rapidly. Moreover, each organization operates within a unique legal framework. Since postmaster is primarily a technical job, we expect that postmas-

ters' area of expertise will predominantly be in information technology. While they may have a solid grasp of the legal issues surrounding their job, they would probably need to turn to professionals for guidance, updates, and detailed expertise on legal matters.

Many postmasters have to deal with strong regulatory requirements in their industry, especially in the realms of government, medicine, and finance. Many email services will have requirements placed on them for email retention, duplication, archiving, and privacy. Postmasters will have to tailor the service that they provide to meet these requirements.

Most policies written by or in consultation with the postmaster should be examined by legal experts before they are implemented within an organization. The postmaster and legal counsel should communicate on a regular basis regarding new laws and regulations and discuss new interpretations of existing laws.

This should be a two-way exchange, not merely blind dictates by the legal department imposed on the IT staff. It is important that the lawyers understand what challenges are being faced by the folks who maintain the email service as well as how various threats change over time. It's also important for them to understand the costs and timetables various possible changes of the services might entail.

Finally, any impasse between these two groups over what should be done to address a particular situation should be resolved by a third party who can arrive at an appropriate cost-benefit analysis of all available options. This person or group of people is likely to be very high up in an organization's food chain. It is important that this arbitrator be assigned in advance of any possible conflict and be kept in the loop about the status of both the legal and the technical aspects of the email system. They should not be caught flat-footed if a situation arises that requires them to make a decision.

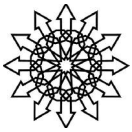

# 2. Email Policy Considerations

Before discussing email policy issues, we feel obligated to point out the obvious—that appropriate policies will vary greatly, depending on the needs of the organization for which they are written. Further, there are as many sets of organizational requirements as there are organizations. Consequently, we cannot presume to write email policies where "one size fits all." Instead, what we hope to do in this chapter is to list many, but by no means all, key issues that may be appropriate to include in email policy documents and to provide some suggestions of where the policy writer can look to receive additional inspiration. It is up to each individual organization to consider each of these issues and decide how or whether it should be included in their own documents.

Nonetheless, the policies produced by similar sorts of organizations are likely to cover many of the same issues. Therefore, it is worthwhile to speak generally about several types of organizations and our thoughts on what the general priorities of those organizational types are likely to be. For a high-level examination of the needs of various types of organizations we divide them into five types: corporate, service provider, academic, government, and loose organization. We will consider the unique aspects of each type of organization in its own section.

## Organizational Policies

### Corporate

In the corporate environment, email is a tool to be used to advance the goals of the company. Generally, it is reasonable for a company to think of email as corporate property, as are its equipment, paper files, and other data. Email should be used for business purposes, and it is perfectly acceptable for other uses to not be supported.

That doesn't mean that we recommend setting a policy that prohibits sending non-business email. Personally, we see no problem with co-workers sending email asking if a colleague would like to go to lunch, or whether enough people are available to field a company softball team. However, email of a personal nature that would be embarrassing to the sender, recipient, or company as a whole should not pass through corporate email servers.

Once a decision has been made regarding how strict a company's email policy will be, it is important that these expectations be clearly communicated to all employees. If a company wants to permit a worker's company email address to be used for personal

correspondence, that's fine. If a company wants to be draconian about not using a corporate email account for unofficial communication, that's entirely justifiable. What is most important is making sure all employees are properly informed about what behavior is expected of them.

In a corporate environment, it is reasonable to create a policy in which those who use email services have little or no expectation of privacy or confidentiality, but it is critical that this be communicated unambiguously. Encryption may or may not be employed within an organization as appropriate, but it is entirely reasonable for that company to require the escrow of all keys used to encrypt company email so that messages can be recovered even if neither the original sender nor the recipient is available to decode the message at some future time.

Many companies operate under specific legal or regulatory guidelines that require email to be handled in a certain way. If so, these guidelines and their impact should be clearly explained to each employee, and email policies should reflect these practices.

A useful principle for email that is sent to networks external to the company is that nobody should send anything that the corporation wouldn't be willing to see printed in the local newspaper. Even internal email could someday be subpoenaed, so before sending a message, employees should consider how they'd feel if they were called upon to justify the message they are about to send in a deposition. If at all possible, private or sensitive information should be encrypted or otherwise protected.

Determining the appropriateness of an email message is no different from any other form of communication, whether it be telephone, fax, postal mail, memo, or person to person. Each of these should be given as much protection as is appropriate. The ramifications of any form of communication should be considered before any message is sent.

**Service Provider**

For the service provider, email should be considered to be a service that is provided to valued customers. Generally, email is the personal property of the customer who sends and receives it, and it should be protected as such. Most providers consider email to be a service that can be used for anything that is not illegal or in violation of social conventions. What is considered a "violation of social conventions" should be clearly spelled out in an Acceptable Use Policy (AUP). It is customary to prohibit the use of a service provider's email service to send spam, spread malware, threaten or harass, attack another site, etc.

Even though safeguarding a customer's email should be a priority, for a service provider email is a bulk service. The ratio of messages handled to people providing the email service will be quite high. Consequently, the level of attention that can be given to tracking any single piece of email will often be much lower than it would be in other environments. This email service is generally provided on a "best effort" basis.

A service provider will be unwilling to make guarantees about the security and privacy of the email that passes through its servers. There is also typically no guarantee

that email won't get lost, and even if email loss is the fault of the service provider, usually no recourse is available. This is reasonable as long as the customers are made aware that they are ultimately responsible for protecting their own data.

In jurisdictions where email encryption is legal, the issues surrounding this practice rest entirely with the end user. We would not recommend that most service providers enter the key escrow business. In jurisdictions where encrypting email is illegal, the email use policy should state that sending or receiving encrypted email through this service is not allowed. The service provider may petition their government as to whether they believe such laws are appropriate or not, but at the same time they should inform their customers that they will do whatever local laws require of them.

Service providers may be subject to specific laws and regulations that require such things as log retention or even email archiving. Again, it is entirely fair for the organization to lobby for changes to these laws, but the company has a duty to its shareholders to stay in business, and that requires obeying the law. It is the service provider's duty, though, to make the customer base aware of the ways in which their email records may be made available to law enforcement.

Similarly, a service provider may be served a subpoena for specific email records. In this case we believe it is the service provider's duty to act as an advocate for the person whose data are requested. Any subpoena should be construed as narrowly as is reasonable, and, unless prohibited by law, the customer should be informed if their data has been subpoenaed. In many cases a service provider will be required to walk a fine line between their duties to their customers and their responsibilities to society and law enforcement. Since law enforcement can act as its own advocate, we believe that service providers should lean toward being a proponent for their customers while still fulfilling their duties under the law.

Some service providers use data provided in customers' email messages for marketing purposes. While the authors of this booklet find this practice to be disturbing on a personal level, we do not believe it is unethical as long as the extent to which this is practiced is clearly spelled out to the customer in unambiguous language. Any practice about which one doesn't want one's customers to be aware should automatically be categorized in an ethical gray area.

At service providers there is often considerable tension regarding who is the primary customer. Is it the holder of the email account, law enforcement, or one's own marketing department? Some service providers send mixed messages on this issue, not only to their subscribers but also to their employees. As long as this occurs, there will generally be widespread unhappiness. We believe that it is better for the management of the service provider to send a strong message on this issue, whatever it may be, even though in practice few do.

One issue faced by service providers is that email used by the customers and email used by employees in furtherance of the service provider's business goals can become intermixed. We believe that this is inappropriate. If employees are encouraged to use

the company's services as customers, then we strongly recommend setting up two different accounts for each employee, one for use for their job functions and one for use as a subscriber.

We also recommend that corporate email be handled and stored on separate servers and be addressed to and from a different domain or a subdomain of the subscriber service. For example, an organization may consider using example.net as their subscriber email domain and example.com for their corporate email, although this can be confusing. Similarly, a service provider using the format: user@example.com for their customer email addresses may want to use an email address of the form worker@office.example.com for their corporate email.

While the authors applaud the sentiment that a company should "eat its own dog food" by relying on the email service they provide their customers for their own business communication, doing so presents some significant problems. Corporate email may have different requirements for archiving and retention than customer email. If that's not the case now, it may become the case in the future. Further, at some point corporate email may be subpoenaed for some legal action. It is appropriate to take steps to help protect the privacy of one's customer base by keeping the two sets of data as separate as possible.

**Academic**

In an academic environment, email is a service to its community to support the mission(s) of the institution. In general, academic email requirements are a mix of aspects from the commercial and the service provider realms. Of course, there are many official uses of email at academic institutions, but students, faculty, and staff are typically encouraged to use their academic computer accounts for personal use as well, as long as it doesn't interfere with "official business." These two sets of requirements can cause problems that are not easy to reconcile.

One of the conflicts faced in an academic environment is the issue of email privacy. Because a large segment of an academic email service will have a lot of student access and be rather open, we believe that any confidential email should not be sent or stored on the general-purpose email service. This includes salary and other HR records, medical information, grades, supervisor reviews, etc. It very well may be the case that the best solution is to have the institution's administration staff use an entirely separate system with tighter controls than that used by the students, faculty, and other staff.

In an environment where email is used for both personal and professional correspondence it can be tricky to know what sorts of accesses are appropriate. Just as one example, if a dean needs information that was sent to a professor, and that professor cannot be contacted, is it okay to have someone look through that professor's email for that information? Even if such an inquiry is justified on its face and well-intentioned, there is the distinct possibility that the person looking for this information will stumble across some personal correspondence, and this may have undesirable ramifications.

An academic institution's email policy should spell out how this sort of situation

would be resolved. Either email account holders have an expectation of privacy or they don't. We believe one can justify either of these two positions, but what is most important is that there be no surprises.

An academic institution may be subject to legal restrictions with which corporations and service providers do not have to contend. Some of those with email accounts may be minors. If the institution is operated or supported by a government, there may be free speech or other regulatory issues that need to be considered. It's possible that the content of certain email messages may be contrary to one of the institution's missions. Each of these possibilities needs to be considered when crafting academic email policies.

**Government**

The issues surrounding governmental email services are probably most like those concerning corporate services, but there are some important distinctions. The specific issues each governmental email service has to face will vary widely from jurisdiction to jurisdiction. Nonetheless, many of the users of these systems face similar constraints.

First, restrictions on the use of these systems are often by law rather than by company directive. Many governments restrict government property to official use only, and this typically includes computers and networks. We have never heard of anyone getting fired for sending email to a colleague saying, "Some game last night, eh?" even though this might violate the wording of the law. Still, we recommend that if a literal interpretation of a regulation prohibits this sort of communication, that this prohibition, draconian though it may be, be stated as policy and followed by employees.

Much as in a corporate environment, email users in a government environment typically have little or no expectation of privacy in their email. Again, though, this should be explicitly communicated with the employees.

Encryption may be used in some government organizations, and under some circumstances it may even be mandatory. Expect that some sort of key escrow system will need to be adopted in this environment.

A government email service may have requirements placed on it for message retention. These requirements are typically satisfied by the underlying service, that is, the employees are not required to perform the document retention themselves, except, perhaps, to perform some categorization of the messages they send and receive.

**Loose Organization**

Some organizations set up email largely as a convenience for their membership. These sorts of organizations typically have a small number of email users and include clubs, social groups, housemates, or those working on a network-based community project. For these organizations intra-group email is used primarily for communication within and without the group on group topics and for group identification purposes.

The size and scope of these groups often do not require a great deal of formalism regarding allowed email actions, but in these cases privacy and personal responsibility

are usually paramount. At the same time, attention to these sorts of email services is often sporadic and informal, so accidental disclosure may actually be more likely in this environment. If this is a risk, then this fact should be disclosed to those who would use this service.

Much as with a service provider, encryption issues are addressed by individual users, although if general encryption use is prohibited by law in the jurisdiction in which the server or the organization it represents resides, prohibiting its use by policy would be justified. In these sorts of environments there shouldn't be the sorts of conflicts of interest that often appear in service providers between customer benefit and marketing issues.

Loose organizations also have to deal with legal and regulatory issues. If the local laws for email service are relatively onerous, requiring, for example, all email services to retain activity logs for a long period of time, then the burden of legally maintaining this service may overcome the benefits to providing it. Subpoenas and the like will still need to be respected, although since providing this service is an incidental aspect of that organization, having to deal repeatedly with legal issues may be very time-consuming.

It isn't always easy to categorize an organization into one of these five types. In fact, some organizations may resist this taxonomy altogether, appearing to be one type to one customer but another type to another. Managing these systems can be especially tricky.

Let's consider one real example: email service for the usenix.org domain. Some professional organizations, such as the ACM, provide user@acm.org email addresses to their membership. In the case of the ACM, the email service is being run as a service provider. USENIX doesn't do this, however. USENIX email accounts for *;login:* columnists or conference chairs, for example, appear to function much as they would for a loose organization. For full- or part-time USENIX staff, their usenix.org email address may be more like any other corporate email account.

What sort of email use policy should such an organization have? It's not at all clear. The safest way to proceed would probably be to split up the email space between, for example, mail.usenix.org, which would be run as a loose organization, and usenix.org, which would be run as a small company. Then USENIX could adopt multiple policies, each of which would apply to a single group of users operating under a given domain.

## Policy Guidelines

Writing email policy documents isn't easy. Certainly, we can't presume to write such a document that would be applicable to everyone, or even anyone, in this booklet. However, we can provide some general advice for those who are setting about writing such policies. This is what we aim to do here.

Our first piece of advice is to obtain and read a copy of *A Guide to Developing Computer Policy Documents*, edited by Barbara Dijker [Dijker96] and published by

SAGE in the same Short Topics in System Administration series as this booklet. Dijker's booklet provides a great deal of good advice on the subject of writing policy documents, and it is unnecessary for us to repeat all of this advice here.

It is tempting to write these documents in a unilateral manner, reserving all rights to the organization providing the email service. We would urge some restraint in this matter, although certainly it is appropriate for each organization to protect its interests. We believe that some organizations tend to go overboard here, making policies more oppressive than is necessary.

In the process of writing these documents, someone should be involved whose purpose is to represent the interests of the user community. This could be the postmaster or it could be one or more sophisticated email users in the community. This doesn't mean that the use advocates should get everything they want either, but at the very least their participation may provide a valuable sanity check on the proceedings.

Policy documents should be explicit. They should state what is allowed, what is not allowed, and what the consequences will be if their dictates are not followed. Of course, it is impossible to cover every eventuality, and there is always the risk that situations will arise with extenuating circumstances, but specificity is generally a good thing.

We recommend that all but the most trivial policies should be examined by someone with legal expertise before they are rolled out. Obviously, the policies should adhere to relevant laws and regulations. A legal expert may also help by making suggestions that make the policy documents more precise, hopefully without obfuscating their meanings.

It's generally best if email customers are given a copy of the policy for their own reference along with something to sign indicating that they agree to abide by the policy. The signed page should then be filed away for future reference, probably by whoever is in charge of human resources in a corporate or government setting. In some environments, most notably for email service providers, obtaining a physical signature may not be practical. Consequently, click-through agreements are commonly used. This is an expedient that is often appropriate, although we believe it is an imperfect solution.

It is important to consider how an organization will handle updates to the policy. Certainly it is reasonable for a policy to state that it may be updated as necessary. When an update occurs it is important to disseminate these changes as soon as possible to the people who are affected by them. Whether it is the case that one can state in the policy that the signatory is automatically bound by future revisions to that policy is something each organization should take up with legal experts. We believe, though, that wherever practical, updated signatures should be obtained.

In any case, when an email policy is updated it is critical that those affected by it be informed of the change. The new, complete policy should be made available for reference, and a compact list that itemizes the changes should be provided as well.

On the one hand, it is important for policy documents to be as comprehensive as possible. If a certain situation isn't covered in a policy document, then it is fundamen-

tally unfair to expect a specific behavior from a signatory when that situation arises. On the other hand, brevity is also a virtue in policy writing. The longer the policy, the less chance it will be read, understood, and remembered. The primary goal of any policy document should always be encouraging willing compliance. Protecting against possible consequences should be a secondary consideration.

Balancing considerations of completeness and brevity is a challenge for policy document authors, especially if policies are being written by a committee or by someone without prior policy writing experience.

**Policy Document Issues**

The rest of this section provides our thoughts about some "big picture" issues that should be considered in email policy documents. The subsequent chapters of this booklet deal with issues faced by postmasters, some of which will be appropriate to mention in policy documents. Appendix 1 of this booklet contains two policy document checklists that can help make sure any given document covers many important issues.

At many sites it will be helpful to split the email policy design process into two separate documents. It may be appropriate to create an email use policy to be signed by all email users, while a separate document to be signed by those who have privileged access to the email service details what is and is not appropriate behavior in the course of performing one's duties. One of the benefits of dividing email policy into these two parts is that doing so can help keep the use policy short while still recognizing that a certain code of behavior is required of those who are charged with additional email service responsibilities.

It may also be appropriate to use the email administration policy as part of the basis for documenting email administration procedures. The policy document would discuss guiding principles for email service management, while the procedure documents would define the steps that should be taken when performing the position's job functions.

Let us consider an example of how email policy might become email administration procedure. Suppose our organization had a policy that listed a restricted set of circumstances under which an email administrator could examine the contents of a customer's email account. One of these circumstances might be that such examination was allowed if the email administrator had the customer's explicit permission to do so.

Let's assume that a customer contacted the email administration group about being unable to download new messages that they have received, and let us further suppose that the email administrators have determined the problem is caused by data corruption of that customer's inbox. A procedure might define the proper way to go about rectifying this situation, including obtaining explicit permission to edit the mailbox to repair the damage, potentially removing some or all of the offending message.

## Internet Email Policy Resources

A number of documents that may assist in the generation of email policies exist on the Internet. This section mentions some of these and where they can be found. Just typ-

ing in "email policy" to a popular search engine returns a staggering number of hits, including the policies of several organizations, experts willing to assist in email policy formations for a fee, tips for creating email policy documents, and many other resources, some useful, some not.

Most of the policy documents that are available for examination on the Internet come from educational institutions. Most of those that are not educational in nature are governmental. Several service provider policies are available as well, although most email policies exist as a part of a more general Acceptable Use Policy (AUP). Few corporate policies are available for public examination. This isn't surprising, as most of the companies who make these documents available electronically do so exclusively on their internal networks. At educational institutions the distinctions between internal and external networks are often blurred.

The SANS Institute has some email policy templates available at their Web site, http://www.sans.org/resources/policies/. The SANS resources aren't exhaustive, but they do contain some good ideas that are well worth consideration.

Some useful information is available from the Information Systems Audit and Control Association (ISACA) at their Web site, http://www.isaca.org/.

There are other worthwhile Web sites that contain information on this topic. Many of these change from time to time, so we won't list some of the other specific resources that we have found to be insightful here. However, there are many good information sources available to those who expend some effort to look for them.

We have one final piece of advice for those who are writing email policies: email, or any form of electronic data storage or access, rarely breaks new ethical ground. Sometimes email use results in a novel situation that is exclusive to that medium, but for each electronic breach of ethics there is usually an applicable physical analogy. We believe that the punishments for the two types of breaches should be commensurate.

For example, if an employee of a company were to download and display an inappropriate image on his or her computer screen, it would be fair in many organizations to subject that person to disciplinary action. In our opinion, though, this discipline should not be any more or less severe than if the same person left a magazine displaying the same image on his or her desk.

As another example, many organizations expend a great deal of energy trying to prevent the disclosure of company secrets through electronic mail. At the same time, some of these organizations do nothing to prevent the same secrets from leaving their property in paper form in somebody's briefcase. Consequently, we wonder if the extra energy that is expended searching email produces a good return on investment.

In both of these examples we mention an inappropriate use of company information or resources, one with an electronic manifestation and one with a physical manifestation. In the two cases the violation was approximately equally severe; therefore we feel that the two cases should have equivalent protections and punishments.

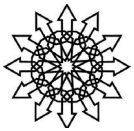

# 3. Internet Issues

An Internet postmaster has to deal with a wide variety of issues, many of which defy easy categorization. We attempt this, however, since we need to divide this booklet into digestible chunks. Therefore, in this section we discuss those postmaster responsibilities that involve interaction with the Internet.

The most obvious of these duties is one that we have touched on before: managing the postmaster mailbox. The postmaster email address must be accessible to the Internet as is required by RFCs 822, 1123, 1173, 2142, and 2821. As we have also mentioned previously, this mailbox must be read frequently, and the issues that are brought to the postmaster's attention should be handled expeditiously. How this mailbox is handled in actual practice will depend a great deal on the composition of the organization and its demands of the position.

For a small organization where postmaster issues occupy considerably less than a full-time job, the postmaster duty isn't much of a burden. Remembering to perform these duties is often a more difficult challenge than the tasks themselves.

Unless the "organization" represented by the postmaster consists of a single person, we recommend keeping the postmaster mailbox separate from other personal accounts, although it may make sense to combine it with other administrative email. The primary reason for keeping it separate is so that if the person who usually reads the postmaster email is not able to perform these duties for a while, no email needs to be redirected and nobody needs to access another person's personal account. Not redirecting email is a good thing because it prevents issues associated with the office from being scattered around several different email inboxes. Additionally, this keeps records of the transactions the postmaster makes centralized for easy reference.

As we mentioned earlier, if postmaster email is centralized, it is essential that the postmaster ensure that each email server can successfully send email. If the volume of postmaster email is large enough, the absence of some or all of the postmaster email might be a clue that something is amiss. Otherwise, ensuring that the server is capable of sending email can be accomplished in many ways, perhaps the easiest of which is to have each server send an automated periodic mailing to the centralized postmaster mailbox. This mailing can be made more useful by including statistics on the service or a summary of the email logs.

For some organizations the postmaster duties entail more than can be accomplished by one person. For these organizations, absences are relatively easy to cover since the

job will be performed by a pool of people. Because of the volume, it is almost certainly necessary to centralize postmaster email. For organizations in this position we can think of three general ways to handle the postmaster mailbox:

1. The hot seat. Assign one person to just go through the mailbox and delegate resulting tasks to other people in the postmaster group. This delegation process may be general, or certain types of tasks can be given to specialists operating in subgroups as appropriate. The person in the hot seat keeps track of who is doing what, whether all servers have reported in, and other "big picture" issues pertaining to the service. This works fine as long as the load on this one person isn't overwhelming.
2. Divide by servers. Have some of the staff handle postmaster email for some of the servers, dividing the problem horizontally. This scales nicely, but some of the problems that are reported are likely to be duplicates that end up on multiple postmasters' "To Do" lists. This is inefficient and can lead to confusion, both within the postmaster group and for external contacts. It is conceivable that one person wouldn't be able to handle all the postmaster email coming in to a single server during a work shift. If that's the case, then this method won't scale very well.
3. Round robin. Each email that arrives for the postmaster is assigned to the queue of one of the postmasters. This can be done randomly, in a round-robin fashion, or based on which queue is the least full at any one moment. This method requires some sort of ticket-tracking system in order to assign and manage the team members' work queues. In either case, this method has the advantage of scaling better than any other method, but it is much more complex to implement and presents the same problems of maintaining context and avoiding duplication of effort as method 2.

Of these, method 1 is the most efficient as long as it's practical. If volumes get high enough, an organization may not be able to handle the load as it passes through one person, regardless of how much additional work is offloaded to other people. If this happens, the organization must adopt another strategy. It's important to recognize the warning signs that this is about to occur before the workload becomes overwhelming for a single person, otherwise life during the transition can be inappropriately stressful. Keep an eye out for this: it's the sort of situation that can sneak up on an organization that is not prepared for this eventuality.

For those organizations where postmaster and related duties occupy more than one full-time equivalent employee, it probably cannot be overemphasized how important automated tools will be to managing these tasks. Investing in developing a powerful toolkit can easily reduce the required number of personnel by a factor of two or more. Of course, the development of such a toolkit will be dependent on the organization

and service software used. It's also a more technical question than will be addressed in this booklet.

If the postmaster duties are shared among multiple people, it is important that responses to that organization's postmaster email return to the person who additionally addressed the issue, not to the general postmaster email address. Once one person has taken responsibility for a given issue, it is important that they handle follow-ups as well in order to maximize efficiency and minimize confusion.

One contentious issue is that of email filtering on the postmaster mailbox. Our position is that if postmaster email can be positively identified as malware, it is allowable to discard that message unread. We don't have any formal recommendations for what the threshold should be for positive identification, but the qualifications should be formalized and the message should be identified as malware to a high degree of certainty. This is the only condition, however, under which postmaster email ought to be filtered.

## Spam and the Postmaster

Spam remains a terrible scourge on the Internet, rendering many email addresses and some email servers essentially useless for legitimate communications. However, despite the undeniable reality of the situation, while we admit to the manifest need for spam filtering of email in general, we strongly discourage the use of spam filters on postmaster mailboxes.

The postmaster email address is in many ways the final recourse for email problems on the Internet. If one can find no other solution to an email problem, it's time to contact the postmaster. Consequently, if legitimate postmaster email doesn't get through, it's a disaster. Because of this, we believe that a responsible organization will want to minimize the chance of rejecting legitimate postmaster email.

We believe that many sites can easily justify blocking email from certain email addresses, domains, or IP address ranges under extreme circumstances. However, we also believe that in general email to postmaster should be allowed to go through even if these strong measures have been employed. What happens if a formerly abusive domain or IP address range changes hands? If the postmaster mailbox is closed off, how will they inform the domain that's blocking them that doing so is no longer appropriate? What if the abusive servers had been hijacked and the problem has now been addressed? How would the blocking domain be made aware of this? For these reasons, we believe that the threshold for blocking email sent to postmaster should be especially high.

Some sites will claim that they get so much spam and malware addressed to postmaster that they are justified in filtering that email. If they did not filter, they would not have the resources to deal with the volume of email that the postmaster would receive. As a practical matter, we understand this concern and realize that in some cases not filtering postmaster email would mean that not all postmaster messages would be

read, and that maximizing the number of legitimate requests that can be handled is the top priority. However, if this means that there is even a slight chance that legitimate postmaster email won't get through, a professional postmaster should not be satisfied with this level of service.

The best solution for an organization lacking the resources to handle all of its postmaster email is to make more resources available to the task. If resources to handle all the postmaster email are not available and will not be made available, then out of necessity some triage on the messages will have to occur. We expect that a great number of organizations find themselves in this position, and we understand their plight. However, a responsible postmaster should arrive at this conclusion reluctantly.

We understand the reasons why organizations want to filter postmaster email. We appreciate the rationalizations that sites will come up with regarding the costs of adopting the policy we recommend. We're very sympathetic with the problems faced by the organizations that handle the largest volumes of email on the Internet, because we've been involved with them. We understand the personnel and equipment pressures these organizations face in handling the deluge of spam and malware. But while these organizations have our sympathies, ultimately we believe that filtering postmaster email is not in the best interests of the Internet as a whole. The RFCs do not list exceptions to the rule that the postmaster mailbox should be available to the Internet, and we encourage organizations to do everything in their power to avoid the need to perform such filtering.

At the same time, we will admit a limited exception for those sites whose email behavior is truly egregious. If a site is sending truly unconscionable quantities of email addressed to postmaster, then we believe that temporarily blocking those messages is justifiable for a limited amount of time under the following circumstances:

1. The blocking is as narrowly defined as possible.
2. The list of blocked sites is managed by the postmaster at the blocking site. Such a block list *must not* be maintained by an external organization.
3. Attempted connections are monitored from the blocked sites, and the list of blocked sites is reviewed on a periodic basis (as frequently as the postmaster mailbox is read).
4. The block is lifted *immediately* if the traffic volume falls to manageable levels, even if the offending site hasn't completely mended its ways.

There are other conditions under which a site may feel it is justified in blocking email bound for postmaster; however, we believe that sites should think twice before doing so.

Several new email validation techniques that may assist in combating spam include the Sender Policy Framework (SPF) (codified in RFC 4408 [WonSch06]) and DomainKeys Identified Mail (DKIM) [Allman05]. Other similar techniques have been proposed as well, but these two seem to receive the broadest support at the present time.

Sender Policy Framework is a new, DNS-based mechanism for matching domain portions of email addresses with severs permitted to send email on behalf of that domain. If the domain's SPF records don't match the IP address from which an email message originates, the recipient is allowed to reject that email.

DomainKeys is a protocol that uses public key cryptography to sign certain key email header fields as well as the body of an email message. The public keys are made available through DNS, so the recipient email server can easily obtain them and verify that the messages were indeed generated by someone with access to that domain's keys. If the signatures in a message aren't valid for the domains they represent, then the recipient is allowed to reject that email.

Much as is the case with other potentially problematic email, email sent to the postmaster from IP addresses excluded by SPF or with improperly signed DKIM message fields should not be blocked. It's entirely possible that a domain has misconfigured their SPF or DKIM configuration and is sending email to postmaster in order to determine why their email isn't getting to its intended recipient. This is certainly the sending domain's problem, but it's neighborly to assist if possible.

Another case where blocking is justifiable but can still be problematic is for those ISPs who publish lists of their dynamic IP blocks, encouraging sites to reject email originating from those IP addresses. Again, we think that blocking email on this basis is entirely appropriate for every destination email address except for postmaster. We have known ISPs who have made errors when they have published their lists and have included the IP addresses of valid email servers in these records. Further, we have encountered situations where domains have added competitors' address ranges to these lists to be a nuisance. If legitimate email is blocked through the accidental or malicious addition of IP addresses to these lists, someone may be entirely justified in trying to contact the postmaster at a blocking site to try to determine why legitimate email that ought to be reaching its destination isn't.

While it's certainly the case that postmaster mailboxes can receive an incredible amount of spam, we believe that few spammers and malware senders out on the Internet try to specifically target postmaster mailboxes. They generally know they can get "more bang for their buck" by sending to other live addresses, as the postmaster tends to be more cautious and savvy when it comes to email issues. However, the overriding characteristic of these people is that they are indiscriminate, so if postmaster@example.com shows up on their mailing lists, it's extremely unlikely they'll go to the trouble of removing it.

Consequently, it would be prudent for sites to take steps to reduce the number of mailing lists on which their postmaster address will wind up. This is especially true since, as we've said, we'd like to avoid blocking email sent to this address if at all possible.

We recommend not publishing the postmaster@example.com email address on any Web pages, and the only outgoing email that should come from this address is that

which is automatically generated by the email system itself. Don't send email to mailing lists or post to USENET news groups using this (or any other administrative) account.

Don't list the postmaster email address in Internet databases, such as WHOIS records. Instead, create other administrative accounts specifically for publication in these databases, and set up an alias to redirect this email to go to whomever is appropriate. Since these are not the postmaster addresses themselves, it's reasonable to filter the most egregious spam before reaching these mailboxes. The postmaster email address is one of the very few on the Internet that is useful without having to explicitly notify people of its existence. There's no sense in publicizing it any more than is necessary.

Even if legitimate email is rejected when sending to, say, abuse@example.com, at the very least the message can be re-sent to postmaster@example.com in order to correct this error. It's bad, but not a disaster, if legitimate email sent to other administrative accounts is mistakenly rejected, but it is a disaster to reject legitimate postmaster email.

A general Internet maxim is that a well-behaved site should be "liberal in what they accept and conservative in what they send." In these days of rampant mail abuse, we can't be as open as we might like to be. However, since the postmaster mailbox is supposed to be available as the recourse of last resort for legitimate email that can't get through, we believe that this mailbox should firmly adhere to this dictum. That may make the postmaster's job more difficult, but we believe that it is the appropriate course of action.

Again, we understand that there are many organizations that due to spam volumes or insufficient resources handling postmaster duties simply cannot presently afford to not filter postmaster email. We don't want to imply that these organizations are evil or negligent in any way. However, we do want to impress upon email professionals just how important it is to accept all legitimate postmaster email. If an organization can't process it all, well, that's understandable, but this should be viewed as something which should be rectified if at all possible.

## Other Mandatory Email Addresses

Even though the topic of this booklet is postmaster duties, during the evolution of the Internet several other email addresses that are mandatory or nearly so have come into prominence. Since these have become so pervasive and because they share many characteristics with the postmaster mailbox, we feel it is prudent to spend some time mentioning them here.

Most of these addresses are specified in RFC 2142, "Mailbox Names for Common Services, Roles and Functions" [Crocke97]. Not all of these are necessary for a given Internet domain or email server. However, most of them are probably worthwhile for any medium-sized or larger Internet organization. We will consider aspects of some of the more popular and necessary RFC 2142 email addresses here.

**abuse**
This email address is where we should expect the Internet to send complaints against inappropriate data emanating from that domain. We believe that these days every domain, no matter how small, should support this email address. It, or better yet an alias for it, should be listed as the abuse contact in WHOIS records for each registered IP address block. The email sent to this address should be read by someone empowered to take appropriate action against perpetrators of network abuse, including the sending of viruses, spam, or unauthorized network scans.

In addition to the postmaster mailbox, we believe this is the one address on which content filtering should not occur. The reason is that if someone is sending a spam message or virus sample that originated from one's network, we want notification of that event to get through. If this email address is filtered, the message containing a perfectly valid alert may be bounced or discarded, and that will likely irritate the sender, who is just trying to be helpful.

Sometimes email issues will be sent to this address. If different people are reading the abuse and postmaster mailboxes, it often can be tough to know who should handle a given issue, but it is up to each organization to determine where these boundaries ought to be. Moreover, some of the issues sent to this address will need to be handled by computer security, legal, or other departments.

**hostmaster**
As the postmaster is to email and the webmaster is to Web services, the hostmaster is to DNS. Every domain should have a corresponding hostmaster, although when DNS service is outsourced it is certainly reasonable to expect the outsourcing company to support hostmaster email for that domain as well.

An email point of contact for each DNS zone is expected to be listed in each zone file and available as part of that zone's SOA record. We believe that under most circumstances, hostmaster at that domain (in the format hostmaster.example.com) is the proper email address to use. We also believe it is allowable to perform reasonable content filtering on email arriving at this address.

**security**
This email alias is less widely adopted than abuse, but it is still very commonly used as a mailbox to which information on Internet security matters may be sent, typically those that may not involve specific incidents that might be more properly sent to abuse. We recommend setting this email alias for all Internet domains, although many sites will forward this email to the same mailbox as the abuse account, or vice versa.

**usenet**
If a site runs a USENET news server, this email alias should be defined for the domain as a whole (e.g., usenet@example.com), and for all USENET news servers (e.g., usenet@nntp1.example.com). If a site does not support USENET news, then it isn't necessary to support this email address at all.

Email to this address should be read by whomever is responsible for supporting

news services. Common issues include USENET news spamming issues—although these are probably best sent to the abuse address—new newsgroup requests, newsfeed problems, and information about inappropriate or illegal content.

Some issues sent to this address will be best addressed in conjunction with other groups within an organization. Therefore, the person reading this mailbox will need to perform appropriate triage on the issues raised and forward the messages to the appropriate person to deal with the situation.

The news email address is often used in place of the USENET address. We suggest that news administrators use one as their primary address and make the other an alias for the primary. We don't believe there are problems with providing email content filtering on this address.

**webmaster**

The webmaster email address is commonly used to address problems with World Wide Web sites. Some organizations split up responsibilities for the technical operation of the Web server and Web content, and split the email contact for each of these responsibilities. If a site expects even the slightest chance that people will separate out their queries based on this convention, the addresses in question and the circumstances they cover should be prominently indicated on the Web site. At the very least, we advocate that any Web server should support webmaster email for that domain. Additional email addresses for different Web-based functions on that server may be deployed at the organization's option.

Because these email addresses are typically listed for public viewing on an organization's Web site, they are prime candidates for harvesting by email-collecting robots for addition to spam lists. Consequently, we have no problem with the idea that email bound for these addresses be filtered for content.

**Others**

Other email addresses are listed in RFC 2142, and these should be adopted by each organization as appropriate. See the RFC for a complete description of these email addresses. However, we believe that the ones we have included here are the ones that are most likely to be deployed, as well as the ones that usually will be called upon to handle most Internet issues.

## Misdirected Queries

Even though this booklet considers postmaster duties only in conjunction with electronic mail, some people may consider the postmaster to be their point of contact for queries about all Internet services. Because of this, the postmaster mailbox may receive queries about services other than email. We feel that the best practice would be to both forward the message to the appropriate contact and to respond to the sender, letting them know that the message has been forwarded and to whom future queries of that nature should be addressed. It would be considered acceptable to either forward or respond with the proper address, but doing both requires only marginal additional effort.

## Handling Email Aliases

At this point it would seem to be appropriate for us to remark on handling the aggregation of these email aliases for multiple servers. First, as we discussed with the postmaster mailbox, aggregating service-related mailboxes from multiple servers carries with it a set of risks. These risks occur for all administrative email traffic, but they are more pronounced for postmaster than for other addresses. We expect, though, that most organizations that support multiple servers performing the same or similar services will feel that the benefits of centralizing service-related email will significantly outweigh any shortcomings.

If this centralization is performed, we recommend that each email address supported in this manner be centralized first and then rerouted via alias to the actual account that will read this email. Here's an example to make this clear:

Assume that we're considering a small site with a single system administrator who handles webmaster, hostmaster, postmaster, and other duties for multiple servers at a given site. We recommend that for each Web server webmaster@www1.example.com first be forwarded to webmaster@example.com. This may be accomplished by MX records or by running an SMTP server on each Web server as appropriate. On the central server one may decide to read email as webmaster, an alias may send it to another account, or several different system email addresses may be aggregated in a single account, which we might name "admin," for example. By aggregating email by service before changing the account name, the final destination mailbox can be modified by changing only a single email alias rather than having to make this change on several servers. This makes the overall configuration of the email service simpler and easier to maintain.

It is worth noting that email should never be read from an account with special privileges, for example, root. Instead, that email should be forwarded to another account that does not have special privileges. On all but the smallest sites, and on any site where more than one person might read system email, we recommend the creation of a special account with no special privileges, e.g., admin, for this purpose.

Of course, at sites where several people assume these administrative duties, it's important to be able to determine who has done what. This can be tricky, although not impossible, with a shared account. For these reasons, many large sites prefer to have an individual account for each administrator performing a specific duty, in order to make it easier to track responsibilities. If this is the case, then there needs to be some way to access the data for which another administrator is responsible in case the responsible party is unavailable. Either shared or individual accounts can be made to work as long as the ramifications of any particular strategy are carefully considered.

## Problem Email

Every now and then a postmaster will receive a message from someone who can most politely be described as difficult. There seem to be a truly astounding number of peo-

ple in the world who both seem to enjoy complaining and have a nearly unlimited amount of free time to do so. We all have different definitions of what might constitute an unreasonable person, but we all occasionally encounter people on the Internet who exceed our thresholds. We believe it is worth spending a little time discussing how best to deal with these sorts of encounters.

If the postmaster receives email from someone whom one suspects will be difficult, we recommend treating the first message at face value. It is better to receive a few superfluous email messages from a crank than to be dismissive of what later turns out to be a legitimate complaint. Reply to these messages in a professional manner indicating that their message is taken seriously, just as one would to any other request.

If a complainant appears to be abusive, respond to the potentially legitimate issues only. Don't become insulting in return, and don't address the extraneous issues. It does no good to exacerbate the situation, and there's no benefit in encouraging additional abusiveness or, frankly, additional correspondence.

Carefully craft a clear, unambiguous response. Don't mislead or exaggerate the situation. Of course, these are good recommendations for composing outgoing postmaster email even if the requester isn't difficult. Don't ask any questions if one does not want to invite further response.

Once the situation has been addressed, we recommend at least glancing at additional correspondence from this person just in case additional legitimate issues are raised. Unless this is the case, though, we recommend not responding to additional messages. Typically, the one thing these sorts of people want most is attention. The best bet in getting them to go away is to deprive them of this. If their correspondence gets truly abusive, then one would be justified in blocking the sender for a limited amount of time, although we find that even in these cases ignoring them is still the best course of action.

Sometimes the abusive sender isn't some stranger on the Internet but someone from within one's own organization. These situations are especially tricky to handle, although most of the guidelines listed above are still appropriate. The two main differences are (1) it may not be allowable to ignore additional correspondence after the situation has been addressed, and (2) unlike strangers on the Internet, it's often possible to do something about the offending behavior.

The best way to proceed under these circumstances will depend on the type of organization and the magnitude of the abuse. If the abuse is of a violent or sexually explicit nature, it's almost always appropriate to go straight to the organization's human resources department as well as one's own supervisor. If the abuse is less severe, the first step probably ought to be a carefully worded statement indicating that the offender's behavior is inappropriate, and if it continues it will be reported through the proper channels. As soon as something like this happens, though, get another person into the loop. Preferably this would be a supervisor, but if an authority figure isn't available, then an experienced colleague will certainly do. Having a second head that can help

one avoid doing anything either rash or that might be misconstrued is valuable. Having someone else as a witness to the proceedings is also important in case any offenders choose to file a false complaint themselves.

## Wildcard MX Records

This topic strays a bit from the other topics in this booklet, but it is an issue that a postmaster may face that is implementation-neutral, so we believe it is worth covering.

Wildcard MX records were a misguided attempt to simplify email management by allowing email on behalf of a domain using wildcards, such as *.example.com, to be routed to a single location as a "catchall." While such a development was well intentioned, often unintended and undesirable consequences result.

Our best advice on this topic is to not use wildcard MX records. They cause many more problems than they solve, they're difficult to implement properly, and when they are implemented, the rationale for doing so almost always boils down to laziness rather than necessity. Instead of using a wildcard MX record, create an explicit MX record for each host to be covered in the DNS zone file. This seems like more work, but it's only marginally so and will result in a DNS that is more robust and easier to maintain.

While their use is discouraged in RFC 1537, "Common DNS Data File Configuration Errors" [Beerte93], the admonishment is not strong enough for our tastes. We might use them for expediency on an unroutable test network, but we do not advocate their use with production services in any situation we can imagine.

Using these records can cause many problems, only some of which we list here. Wildcard MX records are difficult to get "right," to the extent that it is possible to get them "right" they can easily cause mail loops that can be confusing and difficult to fix, and they can cause significant spam control/relaying problems that are tricky to diagnose. If anyone is tempted to deploy wildcard MX records, we would urge them to reconsider this decision.

## Being Added to Anti-Spam Lists

There are many organizations that provide the service of classifying networks and domains as sources of legitimate email or not. Some of these lists are commercial, some are maintained by volunteers. On occasion, legitimate domains and networks are improperly classified as bad guys by these organizations. This section discusses what one can do if this happens to one's own organization, and makes some suggestions to minimize the possibility that this may occur.

In addition to legitimate reasons a network or domain may be classified as abusive, some reasons that sites have been added to these lists include: typographic errors, improperly classified email (email marked as spam that isn't), or personal agendas. Some services frequently add networks to their lists rather than just hosts, and this can cause collateral damage as legitimate services get blocked just because their IP addresses are close to those that are abusive. Sometimes a service provider assigns a previously

blocked IP address range to an unsuspecting new customer, and sometimes an entire service provider's address range is blocked because parts of the anti-spam community don't think they're doing enough to police their customers.

It's also possible that the complaints that trigger blacklisting are legitimate. Disabled or out-of-date anti-virus software or some system vulnerability may turn an innocent host into an unintentional source of spam. A well-designed system architecture and adherence to good security practices will reduce the chances that this may happen, but they can never be eliminated.

Another set of problems can occur if a domain is assigned IP addresses from a previously unassigned address block. Many sites try to maintain their own list of valid networks for various purposes, and some may block email as well as other network traffic originating from these "obviously" invalid IP addresses.

What should an organization do if they find their IP address range is added to a blocked list? Well, the first step should be to address any legitimate problems. If one's IP addresses have been assigned from a newly allocated address range, make sure that the service provider who assigned them is announcing them properly. If a server is accused of being an open relay, make certain it is not before trying to get that server unblocked. If one's own site is accused of sending spam, be certain that this has not happened before one denies that it did. If one can obtain a sample of the offending email, by all means do so. It may be that the circumstances surrounding the complaint can be easily cleared up.

Once one has the ammunition necessary to dispute a listing on these services, the next step is to contact the list maintainers. This can be easier said than done. Many of these organizations hide the true identities of their maintainers in order to avoid legal entanglements. Often these sites will not respond directly to inquiries. It can be tough to know if one's queries to them have been read or not. Even if the block is unjustified, there is nothing to be gained from revealing one's anger or becoming abusive. This is exactly what these sites expect from spammers, so one has a better chance of having one's case heard by exhibiting behavior different from what they expect.

Often, a site that maintains such a list will tell someone who has been unfairly added to it that all they need to do is to wait and their network will be removed from that list. Sometimes this will occur as predicted, but sometimes it will not. In any case, it's no fun to have legitimate email bounce while waiting for an indeterminate amount of time for a block to be lifted.

If one is being blocked, there are a few things that can be done to address the situation. If critical sites are blocking email, one can attempt to contact them. They may be persuaded to stop using that particular anti-spam service temporarily or permanently. One can attempt to contact the postmaster at their site to make this request. If they follow the advice in this booklet, email to that address will get through even if email bound for other recipients is rejected. If postmaster email bounces as well, then one could repeat the same request from another unblocked email account.

## Mitigating Blacklisted Servers

If this communication attempt comes to naught, one can attempt to relay email bound for the domains that are blocking connections through other servers. If an organization has access to other network blocks or other domains, it may be worthwhile to set up an email server on that network or within that domain to act as a legitimate relay. As necessary, forward affected email through that domain in order to circumvent the block. It may be worthwhile keeping such a machine online in order to provide this service on short notice just in case such a situation arises (again).

If another domain or network isn't available, it's quite likely that one's network service provider would agree to act as a relay, at least on a temporary basis. If they can't or won't provide this service, then one might be able to persuade a colleague at another trustworthy site to perform this service. Perhaps their interest would be heightened by promising to return the favor if or when the same thing happens to them. If so, it is important that you trust this site not to abuse the privilege. Otherwise you may find one's own network blocked for assisting them.

In any case, we recommend relaying only affected domains in this manner. Unblocked email should be sent directly from one's own email service. Details on how to selectively relay email in this manner are implementation specific. Consult the documentation for the email software that has been deployed for details on how to make this work.

There are several preemptive steps one can take to reduce the chances of accidentally being victimized by one of these services. First, before deploying a new IP address block for email service, check to make sure that this IP address range is not listed as abusive in any of the major anti-spam lists. If it is, perhaps when they are informed that these addresses have a new owner they will expedite removing them from their lists. If not, then it would probably be a good idea to deploy a different address range for email service. Of course, this is only an issue if servers on this network will be used for sending email out to the general Internet. Other uses for this network block are unlikely to cause a significant problem.

Many services publish lists of ISPs' dynamic IP address blocks that are assigned to customers for temporary use, as these networks are often a source of email malware and spam. An ISP can encourage these services to publish accurate information about them if they preemptively provide this information to these services. If anti-spam lists don't have to guess at this information, they're less likely to get it wrong. If one does make this information public, it's important to commit to keeping these lists updated.

Any time an ISP deploys network blocks for use by customers as part of dynamic IP address pools, we recommend not allocating them at a later date for use as part of an email service. One cannot count on the anti-spam services culling their lists of invalid entries in as timely a manner as one might like.

The Internet postmaster not only has to deal with Internet issues as part of the job function but, as we have said earlier, with an ethical responsibility to the Internet community. This chapter has discussed some of these duties as well as some of these responsibilities. That other organizations do not take this role as seriously as they should isn't a reason to shirk one's ethical and professional responsibilities.

# 4. Security Issues

Many postmaster duties overlap with issues surrounding system and network security. Of course, this overlap occurs for anyone working in any aspect of information technology. Information security is a topic that needs to be considered in any deployment or operation of the IT world, so it is natural and appropriate that it be discussed in this booklet.

## Spam Filtering and the User Community

We have already addressed the issue of spam filtering for the postmaster mailbox. In this section we consider this topic from the perspective of the general user community. The first issue to consider is whether spam filtering is to be provided at all and, if so, whether it will be optional or mandatory.

Asking Internet postmasters what constitutes best practices for spam filtering is going to elicit a wide variety of responses. There is no consensus on what false-positive rate is acceptable, whether potential spam should be considered private or not, or to what extent the customers should control the granularity of the spam filtering. Each organization will need to adopt its own set of policies that suit it best. What is paramount, however, is that these policies be communicated to the customers so that at the very least they know what to expect from the email service on which they rely.

With the exception of the smallest systems using the most carefully guarded email addresses, most email services on the Internet today require some form of spam filtering in order for email to retain its usefulness. Customers of any email service should be told whether spam filtering is provided as part of an email service and informed as to how it operates, at least at a high level. It is important that the customers understand what is happening to the email that is sent to them.

In corporate and most governmental settings, it is entirely reasonable for the organization to be unilateral about their spam control policies. After all, receiving spam in the first place, having people take the time to read it, and certainly dealing with some of the potential consequences if people respond to it wastes the organization's resources, and this costs money.

If spam filtering is deployed as a mandatory part of an email service, it is an organization's responsibility to keep its email users informed about what is being done to keep spam under control, but it is that organization's prerogative to take appropriate measures to increase the efficiency of email operations with an eye toward maximizing the productivity of the organization as a whole.

A service provider, and to a lesser extent an educational institution, will likely have a different set of goals. In these cases, dealing with spam still costs the organization time, effort, and other quantifiable resources such as storage and bandwidth, but here the primary purpose is the satisfaction of the email customer rather than maximizing productivity. The optimum level of spam filtering is determined not by the efficiency of operations but by the contentment of the customer base. Satisfying each of these two criteria may require two different levels of aggressiveness in filtering.

One of the big problems with customer satisfaction is that each customer may want something different. One way to address this is to provide a set of tunable parameters so that the customers can set their spam threshold to whatever most meets their needs. These systems are more complex to build and maintain and often more expensive to deploy than a more unilateral system, but they also provide the opportunity for higher customer satisfaction and so are entirely appropriate for use by service providers and for other organization types.

A customizable system can also be advantageous for corporate environments, but there are some pitfalls. If a service provider's customers tune their spam threshold settings incorrectly, they create their own problem. If someone does this in a corporate or governmental setting, they can cause problems for the entire organization (e.g., by classifying legitimate and important email sent to them as spam and, consequently, not getting important information in a timely manner). So it is perfectly reasonable for these organizations to restrict possible setting levels or even require IT intervention to make these changes as a sanity check on possible errors, provided that the user community is informed about how the system works and about how settings may be changed.

Suppose a site has deployed an anti-spam system that provides the user fine-grained control over determining which email messages are classified as spam and which are not. We recommend that such a site provide some rough guidelines on setting the system to assist those who may not want to involve themselves in the intricacies of the system. We expect that creating shortcuts for setting a person's filters to representative "permissive," "nominal," and "strict" settings will meet most users' needs. Doing so will reduce the chance that people's filter settings do not match their expectations.

## Effective Spam Filtering

Whenever possible, it should be determined if a message is spam before it is accepted by the mail server. If a message can be identified as spam during an SMTP session before message acceptance, then it can be rejected before further processing of the message consumes additional resources on the recipient server.

At the earliest opportunity following the determination that a message is spam, the recipient email server should return a 5xx error code indicating permanent delivery failure. If the email server supports the ESMTP-enhanced status codes [Vaudre96], a status code of 5.7.1 (indicating a permanent failure, delivery not authorized) is appropriate.

Another advantage to rejecting the message at the connection is that if this detec-

tion turns out to be a false positive, the sender will be immediately notified of the delivery status and can take action to contact the intended recipient or, failing that, the recipient site's postmaster to address the situation. Another benefit is that since the error status can be communicated as part of the SMTP transaction, no bounce message is generated, which may annoy someone if the sender's email address has been forged.

Not all spam can be identified during the SMTP transaction, however. Sites must address what they will do with those messages that have been accepted and then identified as likely spam. There are several possible courses of action that can be taken, including:

1. Tag the spam and deliver it normally.
2. Deliver the spam messages to a special mailbox.
3. Bounce the spam messages.
4. Discard the spam messages.

Option 3 is problematic since much of the truly egregious spam is sent with forged credentials. The real sender is unlikely to receive notification that the email has bounced and, in any event, is unlikely to care. More probable is that the bounce message will clog one's own outgoing queues and double bounce as undeliverable, or end up on the server of an innocent third party. The advantage of bouncing email is that if a false positive occurs, the sender will be notified of this fact. However, our experience is that, on most systems, for every false positive there will be hundreds or thousands of undesired spam messages, so bouncing spam after delivery requires a great deal of effort from email servers in general for a very limited return.

If an email service does bounce a great deal of the spam it receives, that domain may wish to consider rate-limiting these bounce messages. This is especially true for those sites who have deployed high-capacity email servers. By rate-limiting these responses a site can limit the collateral damage done to an innocent third party while still providing notification of non-delivery to legitimate senders.

Unfortunately, the consequences of simply discarding these messages, option 4, can be even worse. Even though fewer resources are consumed, if the filtering software does register a false positive and the message is discarded, neither the sender nor the recipient will be aware of this fact, and, depending on the message, someone may become understandably upset by this turn of events. Because of this, in general, we don't recommend blindly discarding messages under these circumstances.

It's important to understand that the guidelines presented here do nothing to prohibit mailbox owners from deleting messages in their mailbox that they believe are spam without reading them. While it can be argued that this has the same effect as discarding messages at the system level, there are key distinctions; chief among them is who is taking responsibility for discarding the message in question. Much as with postal mail, recipients should be permitted to do whatever they want with the mail they receive without the post office making such decisions for them.

Customers might want an email administrator to add a filter that discards email

addressed to them that matches a certain pattern; this is permissible since the recipients are specifying what is done with their email messages. Of course, this is only reasonable for email services where the administrators have the bandwidth to deal with such fine-grained user requests. This will probably only be the case in a loose organization or small company.

Where practical and where spam rates are high, option 2, segregating potential spam in a separate mailbox for each user, is usually the course of action we'd recommend. In practice, people rarely examine their spam mailboxes at all unless they are expecting something that didn't show up in their primary inbox. However, this largely solves the problems of the previous methods while keeping the spam messages from bothering customers on a day-to-day basis. Nobody is annoyed by meaningless bounces and if an important legitimate message gets trapped by the filter this message is at least recoverable if its absence is noticed.

If this solution is deployed, a site may want to adopt the policy of automatically deleting the oldest messages from each user's spam mailbox. At the very least, a site that saves but segregates its spam should keep an eye on how much storage space is being consumed by these messages. Spam may be discarded, oldest first, if it exceeds either some age or storage threshold. Customers should be made aware of these. Proper settings will depend on the resources available to the system and the rate at which false positives can be detected and recovered from the spam mailbox. For most sites that serve as a primary mailbox for their customers, a time duration on the order of a month will usually be sufficiently long to ensure that legitimate email can be discovered and recovered if it ever will.

One problem with the spam quarantine solution is for those sites where POP is the exclusive or nearly exclusive email access method. The POP protocol does not allow a single account to access multiple mailboxes. Spam quarantine methods work much better if the customers have IMAP or some sort of Web-based email access to their mailboxes.

If email access is primarily through the POP protocol, then the best mechanism for handling spam messages will almost certainly be option 1, tagging each spam message. This method relies upon client-side software to sort messages based on this tag into "spam" and "not spam" categories. This consumes some extra bandwidth downloading the spam messages via the POP protocol, but this is certainly a viable solution. Not every POP client has this capability, however, or is able to handle the same sets of tags.

## Reducing Spam vs. RFC Compliance

It is our experience that many email services are so driven to reduce the amount of spam they have to process that they are willing to violate the email protocols in order to accomplish this goal. Simply put, it is going too far to break standards compliance in order to reduce the amount of spam.

Consider a real-world example that one of the authors encountered. A site is maintaining a small, opt-in-only mailing list with several subscribers, including one using an

email service that will remain nameless. When a second subscriber at this particular site joins the mailing list, we find that not all messages get through to the two recipients at this site. After considerable investigation, it is discovered that when this site's SMTP server receives messages bound for multiple recipients, the message will be delivered to the first one but messages to subsequent recipients are discarded. Despite this, the server issues an SMTP "250 Recipient okay" acknowledgment for each recipient.

There is no way around it; doing this is simply wrong. This is a clear violation of the RFCs, and in this case it has the powerful effect of discarding email that the service's subscribers want. We cannot condone these sorts of extreme steps, which we presume are done in the name of spam prevention.

A much better approach, although we don't recommend this solution either, would be to issue a 4xx temporary failure message to each recipient after the first. Many true spammers won't try to resend if they receive temporary failure messages, while we expect that all legitimate mailers will. We believe that this still violates the RFCs, but as a practice it is certainly more defensible than the first behavior described. (RFC 2821 [Klensi01], section 4.5.4.1, states that it would be okay to limit the number of recipients per SMTP MAIL command by using temporary failure error codes if the sender tries to deliver to "very many addresses." We believe that under no circumstances can two addresses be construed as "very many.")

## Using External Information for Spam Identification

Keeping track of the methods and IP addresses from which spam originates is more than a full-time job these days. Because few of us have the time required to keep up with the spammers, it is beneficial to leverage the work of others. There are many publicly available lists of IP address ranges, domains, body text patterns, headers, and other clues that can be used to help classify email as spam or not. Some of these databases are commercial, some are available by special request, and some are open to everyone. Many organizations use these information sources to help reduce the negative impact of spam on their email servers.

Many of these services exist, and we certainly can't list all of them, or even all types. They also come and go over time for various reasons, so even if we could list them, our list would soon fall out of date. Each of these sites has their good points and bad points, though, and while we neither endorse nor decry any of them, we believe it is important that organizations who look to these information sources for assistance do so with their eyes open.

On the plus side, using these external rule sets to help classify whether email is spam or not can certainly reduce the amount of spam that sites collect in their inboxes. Moreover, once they are set up, these methods generally require very little continuing effort on the part of the subscribing organizations. If one is content to be self-supporting in these efforts, there are plenty of freely available anti-spam databases that are frequently updated. If an organization wants more support, then commercial solutions present a very viable option.

Each of these solutions adds a little overhead to message processing—in some cases, lookups in local or remote databases, while some require a bit of additional computational effort. Overall, the extra resources required to access these data sources are small compared to the disruption that a large amount of spam creates.

Using these databases does have downsides, however. Many sites adopt a "more the merrier" approach to these methods, figuring that if subscribing to one of these data sources reduces some spam, then blocking every email message that any of them identifies as spam will permit the least amount of spam to get through. This is probably a true statement, but it's not a practice we recommend.

First, the more stringent one becomes about blocking spam, the higher the risk of false positives. There exist people who dislike spam so much that they are unfazed by the possibility that legitimate email might be blocked as well. That's fine for those who have made this conscious choice, but not everyone shares this sentiment.

Second, adding additional blocking methods will provide diminishing returns in identifying spam messages until, at some point, additional methods will consume more resources than they save by identifying incrementally more spam. Adding another identification method beyond the point of diminishing returns only consumes resources and increases system complexity without providing any benefit.

Another downside to some of these services is that their subscribers are dependent on the decisions of those who maintain these databases. On occasion some of these services have succumbed to sloppiness and poor judgment. In the case of freely available services, the maintainers will compile the lists in a way that satisfies their own agendas. This, of course, is their prerogative, but a subscribing site may find at some point that their needs and what the service is providing diverge. Because subscribing sites have little influence over the process, they may find themselves stuck between two unpleasant alternatives: keeping on with the service even though it doesn't meet their needs, or dropping it and suffering under a potential deluge of spam. With commercial solutions, this is less likely to happen, as the data provider has more reason to want to please the customer, but, of course, these cost money.

Few happy, content people make the decision to start an anti-spam database. Most of the founders of these services do so because they're so upset at dealing with the problem that they'd rather expend energy coming up with a solution than live with the status quo. As a consequence, the process of maintaining anti-spam databases tends to attract certain personality types, and these people are often militant about the problem of spam.

Stopping spam is good, but often it seems that these organizations tend to take an especially hard line against spam. There isn't anything necessarily right or wrong about this, but it can be difficult for an organization that wants to (and can afford to) take a more pragmatic approach and be sure they aren't blocking legitimate email to find services that share their philosophy. Again, commercial solutions are likely to be more service-oriented, so they may prove to be more receptive to tailoring their system to meet the needs of the customer.

## Living with Anti-Spam Solutions

On balance, using external data sources to assist in the identification of spam can be of great benefit in helping to keep email useful as a communications mechanism. However, we believe that using these services as a sort of "fire and forget" weapon against spam is ill-advised. We have some recommendations on how to evaluate these decisions.

First, understand the objectives of the people behind the lists one considers using. Make sure that their goals are compatible with those of one's organization. Do these organizations block sites based on criteria other than the generation of spam? Are they concerned with collateral damage issues? It is also a good idea to try to determine if an organization's goals have remained consistent over time.

We would also recommend doing some research on the histories of these organizations to see what their critics have to say about them. There isn't a Better Business Bureau for anti-spam databases, but doing some research on the Internet will likely provide some idea as to what the complaints against various services have been and whether these objections should be a concern or not.

Second, we recommend confining the number of separate services concurrently used to the minimum number that will achieve an organization's spam reduction goals. Keeping the number down makes a service less complex, and that makes it more robust and more maintainable. It's also a good idea to know which service is the one that just missed the cut-off list. It may be that a service that is being employed suddenly ceases to exist or adopts tactics contrary to an organization's email goals. In this case, it's nice to already know before this situation arises which other list will replace the one that no longer meets one's needs.

Third, we recommend keeping tabs on what this service is doing, who they're blocking and why. It is worth the effort to determine which groups are unhappy with the way things are run at a service one is using. We use these services because we don't have the time and energy to maintain these databases ourselves, but that doesn't excuse us from watching those whom we trust to make our email service more useful.

Fourth, we favor the idea that many anti-spam software packages have adopted of using several sources for evaluating messages and then assigning a scoring system to potential spam. In this way, it requires several criteria to identify a given message as spam before it is considered as such. This reduces the possibility that a personal agenda or mistake by a single service will cause the rejection of legitimate email.

Note that the second and fourth mechanisms listed here at least partially contradict each other, although we believe that to the extent that they are compatible they both provide useful recommendations.

There are several ways to reconcile some of these contradictions. As an example, where possible, we'd recommend using one anti-spam package that performs multiple checks against an incoming message rather than use several unaffiliated packages to perform the same set of checks. Moreover, when multiple services are used, we would recommend that the checks that each performs should overlap as little as possible.

We recommend that non-trivial organizations considering the use of external services to assist in the identification of spam email not leave the decision up to the whim of a single email administrator who is willing to accept this responsibility. The flow of email affects everyone in an organization, and the methods by which this data stream is filtered should be carefully considered, actively monitored, and periodically reviewed by more than one individual.

## Malware Filtering

Email viruses, worms, and Trojan horses have become a true scourge on the Internet, several times causing networks and email services at many sites to become essentially useless. Any Internet organization needs to carefully consider its response to malware.

For most significant organizations it makes sense to try to filter out email-borne malware at the email gateway. This is especially true if one's organization is populated by email customers who don't have a great deal of technical sophistication and have deployed systems widely that are especially susceptible to these forms of attacks.

We expect nearly every corporate and governmental organization to consider email filtering for malware to be an indispensable part of that organization's information security plan. For service providers, malware filtering of customer email is probably not a business necessity, but so many services provide this capability that to not do so would probably be a significant competitive disadvantage.

In corporate and governmental environments, it makes sense to make the use of malware filtering mandatory for everyone. In an educational setting, we'd probably lean toward making its use mandatory as well, although we can imagine specific circumstances in which it might be reasonable for certain people to request that this be turned off on an individual basis. In a service provider setting we would recommend that malware filtering be optional, with a default setting of "on."

So far we have been talking about malware filtering that works by making positive identification of characteristics present in email messages, such as virus signatures, certain precise email headers, or other patterns that can conclusively identify a given email message as hostile. This covers a tremendous amount of malicious email, but the process by which anti-virus vendors discover, fingerprint, and send out signatures for brand-new viruses takes time. During the time between an email worm's first release and the availability of a signature for it, a site may be open to attack by these means. This is especially true if the attack uses a so-called "zero-day" exploit, that is, one for which no patch has been released.

No matter how frequently an organization updates their lists of virus signatures, these methods cannot provide protection against all zero-day attacks. Filtering software is available that examines incoming attachments for the sorts of system calls that malicious software is likely to require in order to take control of another computer, but these methods are relatively new and will likely always be imperfect. Still, for those sites that are truly worried about zero-day attacks, these solutions are worth considering.

In order to guard against zero-day attacks, many sites resort to blocking, or at least

quarantining, email that contains dangerous attachment types (e.g., those with .exe, .zip, .scr extensions). An organization that doesn't routinely send around these sorts of files via email is certainly prudent in blocking these types of attachments. Many corporate and government organizations have already implemented these measures. As long as the user community is notified as to what is going on and steps have been taken so that necessary data can be exchanged, we believe this is an entirely appropriate response to this issue.

For service providers, we also think that it is valuable to provide this sort of filtering capability, but it is important that customers be told exactly what is going on, and that they have the ability to either switch filtering off if desired or examine the area in which these messages are quarantined in order to be able to recover any legitimate messages that may have been blocked by these techniques. It is less appropriate for an email service to require these sorts of filters as a service provider than it is for an employer to force them on an employee.

Handling this situation in an educational environment is a sticky problem. We can understand what would drive any given institution to decide to go either way on this issue, and we think that, at the very least, optional blocking based on email attachment type is a good feature to provide customers. However, we believe that any choice on this issue can be adequately justified, as long as the policy is clearly communicated to the customer base.

We believe that if an undesired message such as spam or malware is identified as such by some filtering software, it is generally best to deliver that message somewhere where the recipient can recover it but is typically not bothered by it. We believe the second best policy is to bounce the message on the off chance it turns out to have been legitimate. However, if a message has been positively identified as malware, this is the one situation in which we believe a site is not only justified in acknowledging receipt and silently discarding the message, but that doing so by default is best practice.

Of course, where possible, it's even better to reject receipt of the message during the SMTP connection than it is to either accept and discard or accept and quarantine the message. As was mentioned before, doing so consumes no resources for the recipient, informs legitimate senders that the message has been rejected, and doesn't create the potential of a bounce message being sent to a forged originator. Because rejections that consume minimal resources have to occur at an organization's SMTP gateway, it makes sense to focus one's anti-spam and anti-malware efforts at that point.

Note that we believe that discarding is justified only if the message contains malware, and only if it can be positively identified as such. Unless the recipient happens to be a computer security professional, one can be virtually certain that the message is not desired by the recipient, and that bouncing it is unlikely to provide valuable information to the ostensible sender. Under these very narrow circumstances, we can advocate discarding the message. A message accepted by an email system for delivery but tagged as malware due to an attachment type or some loose pattern in the message should not, we believe, simply be discarded.

One practice that we find especially distasteful is anti-malware software companies using bounced email messages as an advertising mechanism for their services. They usually know the ostensible sender didn't really send the message, but they use this as an opportunity to let the forged sender know that it was their wonderful product that blocked yet another horrible message. Using malware as an excuse to spam is inappropriate. Admittedly, there's a fine line between a legitimate bounce message and using malware as an excuse to spam, but certain software packages have strayed way into the gray area.

There is one other situation in which we believe that discarding messages is desirable even if one is not 100% certain that they contain malware, and that is during a massive zero-day malware attack. For example, recall the email worm designated Sobig.F [Law03], which hit the Internet in August 2003. This virus propagated so quickly and deluged systems with so many messages that many of them were rendered useless. In our own mailboxes, we found not only thousands of these messages, but nearly an equal number of bounce messages, forged to appear to come from our addresses. When this sort of situation occurs, we believe it is appropriate, but only on a temporary basis, to discard rather than bounce email messages that match the rough pattern of an ongoing malware outbreak. Even if this includes some legitimate messages, the sender should understand the special circumstances and realize that there was a good chance the recipient would have deleted such a message unread even if it got through.

One final note on malware and bounce messages. Largely because of these sorts of outbreaks, we believe that when bouncing messages with attachments, best practice is to include the regular message body but not include the attachment itself in the Delivery Status Notification (DSN) message. An indication should be made as to the name and file type that was included in the original message, but we believe that at the present time it does more harm than good to include these attachments in a bounce notification. Doing so consumes bandwidth, storage, and computing resources to no useful purpose.

## Relaying

A decade ago, a majority of the email servers on the Internet were configured to promiscuously relay email by default. These days open relays, especially those on uncompromised machines, are much rarer. In fact, every email server software package of which we're aware now ships with open relaying turned off by default. Nonetheless, open relays still exist, and we believe it is worthwhile to say a few words about the topic of relaying here.

Relaying can be permitted based on IP addresses or DNS information, or it could be based on authentication techniques such as POP-before-SMTP and SMTP AUTH. Allowing relaying based on authentication techniques to specified individuals or domains is always appropriate. We'll focus our discussion here on the issues involved in non-authenticated relaying.

Let's first define the process of email relaying. Suppose server A receives an email

message from another server. Let's further suppose that this message is not intended for a recipient local to server A. If server A sends this message on to another email server then it is relaying this email. Because of the spam issue, email relaying is rightly viewed as a scary topic these days. However, many organizations, especially those with large and complex email services, need to relay email in order to communicate with the Internet. It is important that these organizations allow only authorized email servers to relay messages.

Some reasons that sites will want to relay email include getting email through network firewalls, network routability issues, and aggregation and segmentation of electronic mail services within a dispersed organization. These are all legitimate uses for relaying as part of email service architectures that can be considered a part of best practice.

When email relaying is necessary, the right approach is to restrict relaying as tightly as possible. An organization should be able to create a list of host names or individual IP addresses that need to relay off of a given server, and these machines should be the only ones that are allowed to do so. If it can be assured that a given range of IP addresses will *only* be used for email service, and every machine on that network will be maintained by implicitly trusted individuals, then it would be reasonable to add that address range to the list of allowed relays.

If an organization cannot arrive at a discrete list of servers that are allowed to relay, then we view this as a failing within the organization rather than a good reason to adopt a more liberal policy toward relaying. To us, relaying an entire domain, subdomain, or IP address range that includes machines that perform a mix of services demonstrates a lack of IT discipline. Consequently, we strongly discourage this approach.

Given the proliferation of maliciously owned zombie computers that send out malware via email, many networks set up a general block on outbound traffic with a destination of TCP port 25, the SMTP service. Doing so greatly inhibits the damage these zombies can do, which is beneficial for the entire Internet. Service providers should require that outbound email from their dynamic IP space be relayed through their email servers for this reason. Of course, business customers who elect to provide their own email servers should be permitted to send email directly to the Internet, although they may be encouraged to use the relays controlled by the service provider.

In a corporate or governmental organization, outbound email should be restricted to only those servers that are specially configured and carefully monitored for potential abuse or violations of institutional policy. Special servers operated by individuals or departments with special needs may be allowed to send email directly to the Internet, but these should be special exceptions to the general rule that all outbound email goes through the sanctioned relays.

The same prohibition is also appropriate for an academic environment, although the number of exceptions may be significantly larger, depending on the history of computing services at that institution and the email service architecture that they have deployed.

Organizations have a responsibility to provide some measure of protection to the Internet as a whole against the damage that might be caused by compromised local computing resources. To the extent that the quality of email service within a network is not impaired, it is every organization's duty to ensure that locally compromised machines are not used to spam, harass, or attack other networks. Filters and policy should be set up to minimize the damage that compromised computers can do, and local networks should be vigilantly inspected for rogue servers. If a computer is discovered acting maliciously, it is that organization's responsibility to remove it from the network and not reattach it until the offending machine has been sanitized.

## Mailbombs

A mailbomb is one or more large email messages sent for the purpose of consuming resources, typically storage and bandwidth, at the remote end. Mailbombing another site is the email equivalent of an act of war, and a severe response to such an action is warranted.

If a site is suspected of sending mailbombs or perpetrating any sort of deliberate denial-of-service attack against another site, we believe that the proper course of action is, at the very least, to reject all network connections from the attacking site and then inform the postmaster at that site of one's refusal to reestablish connectivity until the situation has been adequately addressed. It's reasonable to keep such a block in place until the attempted connections from that site fall to a civilized level.

If the matter cannot be resolved quickly, we believe it is certainly appropriate to contact that site's upstream network connectivity providers in an attempt to involve them in the situation. In some cases, informing law enforcement agencies will be entirely appropriate.

## Email Rewriting

Many sites rewrite portions of email messages, especially message headers, for various purposes. In some cases they do this to make sender address information more uniform, in other cases they want to remove information so as to not disclose their internal email service architecture. In these cases the purposes are likely to be legitimate. In some cases this goes so far as to be classified as email forgery.

We support the notion of rewriting sender information for the sake of uniformity or to avoid confusion. As long as the sender is still easily and uniquely identifiable, and as long as responses to that message go where they will be read by the expected person or organization, we have no problems with this practice. In fact, we would encourage it at sites where the internal email architecture may be quite complex, in order to present a simpler interface to the rest of the Internet.

It is our opinion that removing or rewriting Received: email header lines in order to hide internal architectural details is misguided at best. The benefits from restricting the release of this information are small, and these headers are important for identifying email problems and breaking out of SMTP loops. RFC 2821, section 3.8.2, states that

"a gateway ... MUST NOT alter in any way a Received: line that is already in the header." This includes deleting such an entry. Altering host names to obscure internal email architectural details in existing Received: headers definitely violates the RFC, and we recommend not doing it, but we also consider it to be a lesser sin than deleting them altogether.

Many bulk messages we receive, legitimate as well as spam, indicate that responses to this message will not be read, or often even accepted. While we think it entirely appropriate to keep responses to widely disseminated email announcements from going back to the whole list, we deplore the use of email to send messages where the original sender refuses to acknowledge responses. We believe it is improper to enjoy the convenience of email as a mechanism for communicating with potential customers while being unwilling to accept responses to that message. Many of these messages suggest that the proper venue for responses will be through some Web form, but often these do not allow us to request or obtain the information we desire. Many times we've found that messages that include an admonition against responding to them do so because the sender is doing something they know will invite dissent, and they don't want to have to handle these responses as a consequence. We believe that using email in this manner is improper in almost every case.

Some sites require disclaimers to be attached to outgoing messages because they have been advised to do so for legal reasons. We cannot evaluate whether these reasons are justified or not, but we recognize that this is common practice and understand that email software will be called upon to support this convention.

From a technical aspect, one has to be careful not to blindly modify the bodies of several types of messages, such as those that are PGP-signed, S/MIME, DKIM-signed, or single-part MIME messages (although we wouldn't object if the use of single-part MIME messages were deprecated and client software modified to not send email using this technique anymore). Instead, if a disclaimer is to be attached to single-part MIME messages, the best option is probably to convert them to multi-part MIME messages first. We believe, in addition, that these disclaimers should be as brief and unobtrusive as possible. If a disclaimer is to be added to messages, it is critical that it be tested against all of the message types mentioned in this paragraph before it is deployed.

Some email modifications can be most properly classified as forgeries. An email message crosses the line when it attempts to deceive the reader, whether it's about the individual sender or the sender's organization. Certainly, it is permissible to contract with other organizations to send email on one's behalf, but in this case the identity of both the real and represented senders should be discernible from an examination of the message. If a reasonably knowledgeable Internet citizen would be deceived by the information found in the header of an email message, sending a message in that form is almost certainly inappropriate.

# 5. Technical Issues

Email service is a precise business. The systems and the environment in which it operates need to be well maintained if the service is to meet the high expectations of its customers. In this chapter we discuss some of the technical aspects to maintaining a robust email service.

## Internal Email Architecture

We won't discuss the merits of various possible internal email service architectures here. Depending on the circumstances, it may be prudent to deploy a centralized or a decentralized service. At some sites the optimum configuration will be monolithic, in others it will be distributed. Some sites will be best served using popular commercial solutions, while some will be better off using open source packages, and still others will need to develop their own software.

Regardless of which architecture is deployed, it is important that an organization base its service on a unified plan. If similar functions will be performed at different locations within an organization, to the extent that it is possible we recommend deploying the same hardware, configuration, and software to fulfill these functions. This will simplify inventory management and maintenance. In a pinch, equipment and personnel can be deployed efficiently to assist other locations. This is true even if the management of these components is decentralized.

It may be appropriate to have different individuals or groups maintaining different components of a large, decentralized email service. If this is the way in which an organization works best, then it is crucial that both the protocol interactions between servers and the delegation of responsibilities be carefully defined. It is important to understand the extents and limitations of each component. It is also important to understand how the components interact in the service of the whole. An organization also needs to determine who owns and is responsible for each of these components.

Good communication is paramount in such a relationship. While routine updates between the parties responsible for maintaining email service can take place over email, it is critical that an out-of-band communication channel be established. Many times when it is most important that email administrators talk, it is precisely because the service itself is unreliable.

## User Information Databases

For an email service dependent on an external function such as NIS, Kerberos, LDAP, or Active Directory, email flow will be no more reliable than access to that external service. Consequently, steps must be taken to make sure that this external database is robust and that access to it is efficient and reliable.

These external databases may be used by email for username verification, authentication purposes, alias lists or other email routing information, mailing list membership resolution, or to provide mailbox locations. In some cases, email servers may cache this information to improve performance or keep their own local copies of the database, which improves performance and reduces the email server's vulnerability to service outage. If local copies of the data are kept, then even though access to the information is robust, different servers may be acting on different revisions of the data in question. This can cause different servers to simultaneously treat email bound for the same recipient differently. If email servers query a single central source for email routing information, then consistency is assured but the risks that communication between the email server and the data may be disrupted are heightened. These are tradeoffs to consider when it comes to architecting access to data repositories external to the email service.

It is essential that the postmaster be able to perform arbitrary queries against these data sources in order to properly debug potential email problems. It is very convenient if they are permitted to modify these databases, although it's not strictly necessary. Postmasters who are not allowed to fix problems themselves when they encounter them need a direct line of communication to those who can. For postmasters who are allowed to make changes to these data sources, and those databases are controlled by other groups in the organization, it is reasonable to require that the postmasters adhere to whatever procedures the owning group has adopted.

## Supported Client Software

It is important to carefully consider which email clients will be supported by a given email service. It would be nice if everything just worked and all client software implemented the email standards in a uniform and correct way, but that's not the reality of the situation. Many email clients expect features from an email service that not all provide. Some clients and servers interact in unexpected and unpleasant ways. Some software just doesn't work very well.

The providers of an email service may wish to pretend that they're providing a standards-compliant service and if client software doesn't interact well with that service, it's not their problem. However, doing so will lead to customer dissatisfaction, so it is best to take steps to avoid these issues as much as possible in the first place.

Each email service should provide a list of supported client software, including version numbers, to its customers. Each email service should list which email access and ancillary protocols are supported as well. For example, organizations need to determine if the email service will support POP, IMAP, Web mail, and/or direct access to the mes-

sage store. It also should decide explicitly if it will support services commonly associated with email, such as integrated calendar or meeting software.

An organization should decide how it will respond if unsupported client software is used in conjunction with the service. Will this be expressly forbidden? Will it be discouraged but allowed and not supported? Will support be provided on a best effort basis? If so, what is likely to constitute "best effort" support? The customers deserve unequivocal answers to these questions.

In a corporate or government environment, if the range of client platforms is very small it may make sense to greatly constrain the range of permitted email clients. If only one email client is going to be expressly supported, it becomes very tempting to adopt email server software that deviates significantly from the Internet standards. This can be convenient if the additional feature set that this client/server combination embraces is truly useful, but there are reasons to be wary about going down this path.

Rarely can an organization be supremely confident that it will never want or need to expand its client platform list to include those that are not supported by their chosen email client. It would be a shame to become locked in to a desktop operating system solely because one's email service doesn't support clients on otherwise desirable software platforms, especially if the extended features of the proprietary email service are neither necessary nor widely used.

Additionally, this form of vendor lock-in presents other hazards. An organization may decide to replace its email software package down the line but find it extremely difficult to migrate to another system. Migrating customer mailboxes from an undocumented, proprietary storage format can be challenging. Unraveling an organization's dependence on other proprietary features of that system can be even more so.

Consider what happens if one company or governmental bureau acquires or merges with another. If they are using incompatible proprietary email systems, integrating them is likely to be a daunting task. In most cases the only practical way to accomplish a merger of these types of systems is to completely abandon one of them, resulting in considerable loss of investment in time, software, equipment, and familiarity on the part of one organization's employees.

In an educational environment it is usually more difficult, although not impossible, to dictate a uniform client platform to one's customer base. Therefore, it will usually be preferable to deploy a more open, standards-compliant system than would be required in a more structured environment. However, doing so can pose its own set of risks.

Overly permissive customer access methods may make it nearly impossible to improve the system without breaking someone's access to it. For a real example, consider the email service at an educational institution for which one of the authors consulted. This school allowed email access to customer mailboxes via IMAP, POP, Web mail, and direct access to the message store, all on the same system. They explicitly allowed the use of dozens of different email client software packages. As the demands on their system grew, the quality of service started to degrade noticeably. It was determined that

the architecture of the system would need to be redesigned. However, to continue to allow access via all of the currently supported clients would require a Herculean effort to modify code and test everything to ensure compatibility with the new system. As with deploying proprietary systems, it is possible for an organization to paint itself into a corner using open systems as well.

For a service provider, in almost all cases it is essential to strongly support open systems and to make sure that one's email service is compatible with as many of the major client platforms as possible. Enforcing the use of specific client software or particular versions is likely to be impractical, but that doesn't mean that an organization can't strongly suggest that the customers should stick to certain types of software and specific versions. While the most common platforms should receive the greatest amount of testing, other clients should generally still be permitted, even if they're not explicitly supported. This support may be on a best-effort basis, however, and it isn't reasonable to expect a service provider to provide extensive support for the most exotic client platforms.

For email systems that support multiple clients and platforms, it is important that those responsible for maintaining the email service be able to replicate these client environments for testing purposes. Going through the process of creating a test environment also serves as a cautionary exercise on the problems of supporting too many platforms, especially regarding different versions of a given client. Once an organization experiences the challenges in building such a test environment, it is likely that those involved in this task will gain a greater appreciation for the difficulties in supporting a large variety of clients.

As with many other aspects of IT management, the customer-driven desire to support multiple client software packages and versions must be balanced against the costs of supporting those systems. Over time, it will become necessary to add new packages and versions to the supported list. In order to keep the testing process manageable, as new software becomes supported it will be necessary to stop supporting older packages.

When possible, it's good customer relations practice to not force the user community to upgrade or change their software with no warning on short notice. We recommend that a given version of email client software be classified in one of three categories: supported, deprecated, and unsupported. As an organization wants to discourage continued use of formerly supported software, it should be first categorized as deprecated for some period of time before support for it is formally dropped.

Customers should be informed when the list of software that falls into each category is changed and be given every reasonable opportunity to migrate from the deprecated software to another package on their own time. Of course, it is our experience that some will not heed these warnings until the eleventh hour (or beyond), but many will make such a transition smoothly, and every customer who does so eases the support burden on the IT organization. Whenever the supported, deprecated, and unsupported software list changes, the customer base should be explicitly informed. At all other

times the list of supported and deprecated software packages should be easily available to all customers.

By planning ahead, those who support email clients can save themselves a great deal of effort when client software will need to be updated. Instructions to the customer base on how to perform upgrades themselves should be prepared in advance. A list of possible effects and side-effects of the transition should also be made available. People supporting this transition should have already performed the expected upgrades themselves, so support staff should be familiar with these operations. Preparation will pay dividends in reducing the support burden.

## Logging

Busy email servers generate copious amounts of log information. Each organization should have some sort of policy on how to deal with the log data generated by the email service.

Even though the amount of logged information may seem large on an absolute basis, it's generally quite small compared to the volume of email that a server handles or stores. Consequently, keeping the log data around for a considerable amount of time doesn't usually create a large incremental burden on email services. Email logs are valuable sources of data on usage patterns, errors, spammers' tendencies, long-term trends, and other sources of information. Casually deleting these logs may be destroying information that can be used to improve service.

While storage space always seems to be at a premium, storage is generally inexpensive. Consider keeping potentially useful log information around for as long as possible. Some good suggestions discussing logging as a service are available in the SAGE booklet *Building a Logging Infrastructure* by Singer and Bird [SinBir04].

Some jurisdictions require certain industries to maintain email log information, in some cases for long periods of time. On the other hand, email logs may contain information that has legal implications. If it is not a requirement to maintain these logs, and an organization's legal department believes that archiving old logging information may constitute a legal risk, then it may be important to not retain this data for long periods of time. In this case, best practice would be to extract the valuable information from log files as they are rotated out of active use, and then delete the log files themselves.

It is straightforward to write some scripts that can extract email usage patterns for email sending, reception, and access. This information can be used as part of a system baseline for determining whether the service is performing nominally or not. Long-term trends on email usage and message size can be combined with system performance data to predict when a given service may run out of resources. With this information an organization can plan and budget for upgrades before a crisis occurs.

Many people have solved this problem before, and many packages that can assist with log management issues are available on the Internet. It would definitely be worthwhile to examine some of these before starting development on one's own set of tools.

## Backup and Archiving

Email backup and archiving issues should be a part of an organization's overall document retention policies. It is important that the retention of email data be consistent with that of other documents, electronic and paper, as a part of the organization's goals. Consideration of the extent and length of time email needs to be retained will shape an organization's plans for email backup and archiving.

For a general consideration of the issues regarding data backup, we direct the reader to the SAGE booklet *Backups and Recovery* by Preston and Skelly [PreSke02]. In this booklet we will concern ourselves only with issues specific to email.

Email backups are fundamentally different from other types of important data. Most email message stores and all mail queues are extremely transient, making it difficult to capture useful backups of these file systems. After 24 hours, a POP-based service provider's message store may see in excess of 50% turnover in data, making the value of a previous day's backup questionable, and a day-old backup of an email queue almost certainly has little, if any, value. Additionally, performing incremental backups, particularly on transient file systems, often makes very little sense.

Consequently, when considering email backups, as with all data backups, we must consider how quickly the images of the data we're backing up will lose their value. Backups of an IMAP message store may still have some value for months after they have been taken, while an image of an SMTP message queue may have no value in just a few hours.

In order to create a good backup policy for email, we must look at an organization's goals regarding data recovery. What does an organization wish to accomplish when it considers backing up its email service? We need to answer this question in the light of the technical realities of the service itself in order to formulate an appropriate backup strategy.

Backups are designed to allow the recovery of data in case of a mishap. The first thing to consider is what sort of mishaps an organization is trying to protect against. Are we talking about a regional natural disaster, a system failure, or the accidental deletion of a customer's mailbox? Each circumstance calls for a different response. To protect against a natural disaster requires off-site backups. On-site complete backups are appropriate for protecting against hardware failure. File-system snapshots may be sufficient protection against an inadvertent use of the "d" key. Full backups are typically too unwieldy to be used to recover a single accidentally deleted message, much less those that are stored off-site, and local snapshots won't protect against total equipment failure or a large-scale disaster.

Another factor to consider is the time it would take to restore if that were required. Some email services comprise terabytes of data that could literally take days to completely recover, and once restored the data may be very much out of date. For an organization such as a service provider, a multi-day restore of week-old data may not be worth the effort, even if the message store is completely wiped out.

If such a restore were performed under these circumstances, the first priority would have to be to restart the service itself, with restoration of old data as a secondary effort. One would also need a plan to merge the new messages arriving on the new service with the messages coming from backups as they are restored. The costs of performing a full system restore under these circumstances begin to look problematic, making it reasonable to at least ask whether the restoration of the old data would be worthwhile. If performing such a restore is of questionable value, and we believe that this is a fair position to hold under these circumstances, then what would be the point of making such a backup in the first place? This doesn't mean that backups of email should not be performed, but it makes little sense to back up this data without thinking carefully about how it would be restored.

No matter how frequently or how carefully an email service is backed up, there can be no guarantee that any particular message has been saved. In between backups there is always the possibility that a new message will have arrived, been read, and been deleted before the next image of the file system is recorded. This also impacts the value of message store backups.

For an email service built around the IMAP protocol it very likely makes sense to perform backups on the file system on which the mailboxes used for persistent message archiving are stored. Much like the information found in users' home directories, this data is much less transient, less time sensitive, and possibly has more long-term value than that stored in the users' inboxes. For IMAP servers, it likely makes sense to back up this data. For IMAP inboxes or POP-mostly message stores, however, this is less certain and very much depends on the circumstances.

Each organization should actively consider the ramifications of performing email backups, resolve the technical issues, and implement a backup policy. At this point it is important to inform the customers about what is or is not being done to protect their email against data loss. That email inboxes are not being transferred to an offline medium may make plenty of sense for an organization, but the customers should know this before disaster strikes. Customers should also be informed that no matter how diligent a backup system may be, the possibility for lost email will always exist.

Even though both involve the long-term storage of email messages, backup and archiving are quite different processes with distinct purposes and implementations. Backups are concerned with capturing an image of a file system as a whole, while email archival focuses on the preservation of valuable data on a much more granular, typically message-by-message, level. Backups create an image of a data repository at one particular moment in time, while archiving focuses on making sure each piece of data has been saved without regard to how that data may have originally been stored or aggregated.

Because individual messages may be received, stored, read, and discarded between backups, traditional backup mechanisms are not suitable for archival purposes. However, repositories of archived data may be transferred to offline storage using backup technologies.

Email archiving can be performed in two ways: automatically by the email service software or manually by the senders and recipients. The latter method is typically easy to implement and is more discriminating but far less reliable than the former.

Automated email archiving is typically built into the email service software, sometimes with an interface for the user to add information by applying some classification information to each message sent or received. These systems are strongly preferred in those situations in which laws or regulations require the retention of some or all of an organization's email.

There is a downside to email backup and archiving when it is not necessary. Many times this information has been successfully subpoenaed in court cases and has been used as evidence at trial or as ammunition to improve one side's position in obtaining a settlement. For this reason some organizations make it their policy not to create email backups. Evaluating the risks of data loss vs. the risks of legal repercussions is a conversation that should take place between the legal department and information technology of just about any organization, and final approval for an organization's resulting strategy should be given at a very high level. As is usual for email services, a wide range of possibilities can be justified, but what is most important is that the organization's participants aren't surprised by the consequences.

An organization that chooses not to back up email needs to be careful about performing inadvertent backups. For example, on a UNIX host a periodic backup of the system disk for disaster recovery purposes may include the /var directory, and this can catch transient email in /var/spool/mqueue (or the equivalent) or delivered email in /var/mail. Because of this, one may elect to make /var a separate file system and to not back it up, but if a restore becomes necessary, the system won't function very well without creating at least the proper directory structure under /var. This can be a tricky situation.

On email servers, we recommend that directories that contain delivered or transient email should reside under special-purpose mount points, separate from other system functions, so that they can more easily be backed up or excluded from backups as the situation warrants. Computers that are not email servers are unlikely to contain sensitive information in the email spools or queues, even on a temporary basis, so the risk of capturing this data is probably minimal.

# 6. User Issues

Quite a few of postmasters' duties involve direct interaction with the customers they support. This shouldn't be surprising. The entire field of information technology isn't really about computers, it's about enabling technology to improve lives, whether through increasing efficiency, providing entertainment, or enhancing communications. It is paramount for a postmaster to remember that email is a tool to benefit people, and this consideration should drive every aspect of the job.

Customers can be expected to contact the postmaster when they encounter email problems. This is entirely appropriate. The channels by which they are instructed to do so will be dependent on the needs of each organization, but it is desirable for at least one of these channels to be out of band with the primary data network used. It is important to be able to report the fact that primary communications are not available.

There are many ways in which problems people report to the postmaster can be categorized. For the purposes of this section we find it convenient to divide them up by whether they occur on networks controlled by the cognizant organization or external to it. Here we consider email problems to be either local or remote.

## Local Problems

Many of these problems involve either the reception of email or access to messages. Providing general debugging tips and technical advice is beyond the scope of this booklet, but it is appropriate to consider some of the possible consequences of troubleshooting here. It can be difficult to track incoming email while respecting customers' privacy. We provide some recommendations that can assist in this endeavor.

When tracking email problems, check the logs first. There should be no privacy issues regarding any of the information logged by default by any email software of which we're aware. The logs will at least provide some context for the problem in question, if not immediately indicate the problem and resolution. By examining the logs, one will likely be able to determine what the system thought happened to the message in question or at least know where one needs to look next.

Once the first transaction regarding the message in question has been located in the logs, the next step is determining what happened to it. If it has been delivered, then that's it for the logs. If it has been queued or rejected, then it would be prudent to continue looking for subsequent log entries regarding that message. If the initial delivery attempt was unsuccessful, one or more subsequent delivery attempts may have

occurred leading to the message ultimately being bounced or delivered. It is also entirely possible that multiple messages fitting the search criteria (sender, time of arrival, subject line, etc.) have been sent, so it's important to be certain that one is tracking the message in question. Often it's not trivial to match up the email a customer expects with the relevant log entries.

If the message in question has been queued, it is reasonable to examine the status of the queued messages. For email software that splits queued messages into separate envelope-plus-header and message body files, it would be appropriate to examine the envelope-plus-header file for clues about why delivery did not occur as expected. The message envelope and header information should be considered fair game for troubleshooting without overstepping any expectation of privacy the message sender or recipient might expect.

Sometimes the problem will be in one of a customer's mailboxes. Debugging these issues while still respecting user privacy can be tricky. Each organization should have a part of their email administration policy that covers this situation.

Generally, we believe it should be mandatory to receive explicit permission from the owner before searching or editing a given mailbox in order to resolve problems with that mailbox. Of course, if permission isn't granted, options for solving the problem may be severely limited, but this is the mailbox owner's option. It is also recommended that if such authorization is given, the email administrator take special measures to ensure that the authorization is genuine.

It may be a good idea to have a colleague present at an authorized examination of a user's mailbox to concur that authorization was given and to witness what actually happens during access. This isn't exactly a "two-key" system, but working in pairs should reduce the chance of a major mistake or of someone exceeding the acceptable bounds while trying to address the problem.

If a customer is having problems accessing the contents of a mailbox and doesn't want to grant anyone access to it, generally the best solution is to move the mailbox away, package it up, and give it to the customer to deal with via an appropriate mechanism (which could be emailing it back to them as an attachment or moving the mailbox to a private directory in that user's file space.) These are entirely reasonable approaches to this situation.

## Remote Problems

Diagnosing problems with remote email services is obviously much more difficult than handling local problems. Again, the first place to look will always be the email logs. By knowing the sender, the recipient, and the approximate time the message was sent, one should be able to determine what the local service thinks happened to the message in question, whether it was accepted by the remote service or not. Beyond that, unless a DSN is received from the recipient email server there is very little else that can be done to diagnose the situation.

A postmaster may be called upon to assist a customer in interpreting a DSN and to

help them to understand what went wrong with the message they were trying to send. A service provider may not have the bandwidth to handle these sorts of requests, but they're reasonable at most other types of organizations.

If sent messages don't seem to be reaching the recipient, but everything looks like it is functioning correctly at the local service, then the last resort is to contact the postmaster of the remote service and ask for assistance. It may be better for the local postmaster to make this contact on behalf of the customer in question. A query coming from a fellow postmaster is likely to get more respect than one from another user, and the message is more likely to be phrased using language that will expedite a useful response.

## Handling Complaints

If the complaints are internally generated due to received email, these should be handled in the same manner as any other internally generated request. Of course, it may be necessary to explain to the customer that it's not appropriate to respond to every undesired or unpleasant piece of email that a person receives, but these sorts of issues arise in every aspect of system administration, and email issues are no different in this regard. In any case, for most IT organizations user education on technical issues is part of their charter.

External complaints are a different matter. These should be taken seriously and acted upon promptly. Generally, acting on them within one work shift is a reasonable response time. If they can't be handled within that time, it is reasonable for the complainant to expect a response indicating when action will be taken.

In a service provider setting, so many complaints may be received on a daily basis that providing personal responses to submissions simply isn't practical. In this case, an automated response may be appropriate. If so, the automated responses should be tagged with an internally significant incident tracking number so that updates or questions regarding that response can be linked to the event in question.

Insisting that one who files a complaint switch to another method of communication or resubmit their complaint to another email address is inappropriate. It is appropriate to ask that future correspondence on similar issues be addressed to a particular location, but these should all be RFC 2142 addresses. Outgoing correspondence on these matters should have the Reply-To: email header set to wherever an organization desires responses to go.

Sometimes a service provider will insist that complaints of a certain type be submitted to a given non-RFC 2142 address, and that complaints sent to other addresses will be ignored. This behavior is inappropriate. A service provider cannot expect the Internet at large to remember the way one organization wants its email routed when that differs from standard practice.

Even though an organization may not be able to respond personally to every complaint made against its customer base, it is important to realize that to an outside person an organization that is addressing the problem often looks remarkably similar to

one that is ignoring the complaint. Some effort should be expended on ways of minimizing the frustration of those who file legitimate grievances against the customers of one's own email service. By alerting the postmaster to inappropriate behavior by one's customers, those who complain are doing that organization a favor and should be treated accordingly.

## Internal Distribution Lists

Many organizations use email as a broadcast medium to communicate with a large subset or even everyone within an organization. Often, sending a message to these large distribution lists can consume significant resources, in terms both of computational effort and of employee time. Consequently, it is often appropriate to restrict the circumstances in which email should be sent to these lists.

The mechanism for restricting access to these lists is a technical issue that is beyond the scope of this booklet, but an organization that supports large internal distribution email address lists should have email use policies that specify when email should be sent to these lists and who is allowed to do so.

A small organization typically doesn't need these sorts of restrictions, and larger organizations usually only put such policies in place to deal with the distribution lists with large numbers of recipients. As organizations grow, however, they often find it worthwhile to curtail open use of lists at some point. This is entirely appropriate, as long as the membership is informed as policy changes.

## Large Internal Messages

Email client software has made it trivially easy to add various types of files to email messages, including some that are larger than most email services can handle gracefully. The world would be a better place if all email software would notify users of the sizes of the attachments they're trying to send before dropping the message onto the network. There aren't a lot of email administrators out there who haven't had one of their customers at least try to unintentionally mailbomb their own organization.

User education will only go so far here, although if it prevents one significant email service outage, it will repay the effort. Basically, people should check file sizes before emailing attachments. Moreover, even if it is appropriate to send a large file by email, a wise person will keep the distribution list as small as practical. If it's necessary to distribute an especially large document or to send large documents to many people, it's usually much better to put the file in some shared file space and email its location to the distribution list.

Needless to say, education isn't foolproof, and no matter how robust an email service might be, there exist a file size and number of recipients that will lead to a significant service disruption. Consequently, it's quite reasonable to configure an email service to reject messages above some size threshold. This should obviously be larger than the maximum message size that an organization should reasonably expect to send, but smaller than a size that would cause problems for that service.

## Access to External Email Services

Few Internet-connected people use only one active email address these days. Since so many people use email as a primary means of keeping in contact, it's natural that they will want to be able to access their personal email from work, much as they expect to be able to carry their personal mobile phone.

Given the fact that email is a serious channel for malware to access an organization's computing infrastructure, connections to external email services can be both a necessity and a risk and should be viewed as such. Policies regarding access to external email services is something that should be carefully considered by every organization.

Of course, for a service provider or a loose organization, allowing customers access to other email sources will almost certainly be a necessity. Customer satisfaction is paramount, and a policy that prohibited customers from reading email stored on someone else's service is unreasonable. Creating a prohibition on accessing external email services would be much more plausible in a corporate or governmental environment. Let's examine the pros and cons of allowing access to external email services from an office network.

One reason to allow this connectivity is that many employees will want it. Having access to external email sources is convenient and people are used to being reachable by email. Consequently, allowing access to external email services is good for morale.

If an organization's official email service isn't as reliable as one might like, people will find that having an alternative email address can be quite convenient. Suppose a customer wants to send important information by email but the corporate email service isn't operational. An employee could direct that customer to send their query to another email provider where that request could be retrieved. This flexibility isn't available if access to external email services is restricted.

One issue organizations have to deal with is whether and how to allow access to internal email services from outside the organization's network. If external access to official email services is not available from outside the network, and employees often work from remote locations, such as working from home or when traveling, then they will most likely need to use service providers for official email. Consequently, they are also going to want to be able to access these external email services when they're within the organization's network perimeter.

Of course, using both internal and external email services may cause a great deal of confusion for the customers, and employers may not want to have official company correspondence handled by services outside of their control. This is especially true for those businesses that require archiving of correspondence for compliance purposes. So there certainly may be sound business reasons for why an organization may prohibit the use of external email services for official purposes.

Providing access to external email services also comes with technical risks. It's entirely possible that service providers will not provide as much protection against malware and other email-based attacks as the local email service. This would increase the chance

that these services may be a mechanism by which malware may penetrate an organization's network perimeter.

Another downside to allowing access to non-work email is that this access can tempt an employee to waste work time on non-work issues. Of course, this is a problem with other forms of communication and access to other information. If restrictions against external email access are viewed as necessary for this reason, it is important that the organization treats all similar distractions in a uniform manner. Blocking external email in the name of workplace efficiency but allowing employees to watch television on the job would seem to be an inconsistent set of policies for most work environments.

An organization may encounter significant resistance from employees to its decision to restrict access to external email services. Sometimes this will be due to personal preference, but it might indicate that the official email service is inadequate in some important way.

If the objections are based on legitimate issues—for example, lack of external access or reliability problems—these should be addressed before putting stringent restrictions in place regarding access to external email services. Once technical issues have been ironed out there may still be objections, but these are more likely to be based on criteria that don't impact the organization's bottom line.

It may be the case that the internal email service does not meet all the needs of the organization and that solving these deficiencies isn't practical in the near term. In such a case, an organization may want to consider an arrangement with a particular email service provider and allowing external access only to accounts created with that service. The organization can work with this particular service provider to make sure that its service satisfies the technical, information security, and perhaps regulatory needs of the organization.

It is strongly recommended that a corporation or government organization that explicitly sanctions an external email service account for official business specially create and use this account only for that organization's business. The username and access password for that account should be escrowed somewhere within that organization so that information sent to the account may be retrieved even if the employee who typically accesses it becomes unavailable. If the account is used for official business, then the organization ought to be able to access that account.

We're not advocating that employers should necessarily restrict access to external email services. However, we believe that organizations should be aware of the risks posed by allowing this type of access and should make plans accordingly. The decision to use external email services should be made explicitly. Doing so will have an impact on an organization's network security policy, and this should be well understood by any site that allows such access.

Those familiar with the issues involved in email security are certain to realize that allowing employees to ssh to their home computers and run a text-based email client, and allowing employees to access external Webmail accounts create very different levels of risk for a network from which these connections originate. Nonetheless, trying to

explain the distinction between these two types of external email access to employees will be a challenge in many environments. It will be difficult to explain why one form of access might be allowed and the other would not to those without a nuanced understanding of network issues. It's important that a policy regarding external email access be universal and clear. If an exception is going to be made by access method, then this exception and the reasoning behind it should be spelled out in the organization's email usage policy.

## Account Management

Creating and deleting email accounts is one of those postmaster activities that on the surface seems to be very simple but turns out to be quite complex. Organizations should carefully consider how they want to manage email accounts, especially in light of their particular needs and organizational structure.

In most regards creating new accounts is the simplest part of this issue. Some thought should be given to how one will securely communicate a temporary password for that account (as well as for non-email accounts) from the person who created it to the owner. Some mechanism should be in place for ensuring that the temporary password is quickly changed and that the new password is reasonably strong.

What happens when someone leaves an organization is more complex. One issue that should be covered in policy documents is whether email for departing people should be forwarded, and if so, for how long.

Consider four stages of email forwarding:

1. Silent forwarding
2. Forwarding with a warning
3 Failed delivery
4. Email address reallocation

All email account deactivations can be considered as a specific case of moving through these four steps. At some sites and in some cases, the amount of time spent in any particular stage will be zero.

Silent forwarding is the typical email forwarding function. Email directed to one address is forwarded to another without the sender being aware of this process.

Some email software provides a configuration option to forward email on to another email address while at the same time sending a DSN to the sender indicating the recipient's new email address.

Email sent to an invalid address will bounce. With some email server software the DSN may contain the intended recipient's new email address.

Finally, the email address may be reallocated for use by another individual.

Determining which of these stages to adopt and how long an email address should be maintained in each of these stages depends on the nature of the organization doing the forwarding.

Within academic institutions wholesale changes to most of the information infrastructure, including email, occur on an annual basis. Providing forwarding services for one year after departure seems like a reasonable criterion.

For service providers, email forwarding is a courtesy for those who switch services. In some sense, providing this service removes an incentive for subscribers to maintain their present accounts, but doing so is likely to engender goodwill. Providing forwarding on the time frame of a few months seems reasonable here.

For a corporate or governmental organization, an employee's email is likely to be viewed as the property of the company. Therefore, forwarding a person's email to a new address external to that organization may be inappropriate. However, when an employee leaves the service of an organization, it may be reasonable to configure the system to provide an automated response informing the sender that the intended recipient has moved on, who the new point(s) of contact might be, and, if desired and appropriate, new contact information for the intended recipient.

In any environment, email account issues should be addressed before a person leaves the organization. In all cases it is appropriate to remove that email address from all internal email lists, regardless of whether the account will be forwarded, will be deactivated, or will remain accessible. If the account will not be accessed by the original owner going forward, then they should be encouraged to follow the appropriate procedure to unsubscribe to all external mailing lists.

It may be necessary to forward a departing person's email to someone else for processing for a period of time. If so, some mechanism should be granted to allow the new recipient to be able to send email on behalf of that user in order to unsubscribe to mailing lists that may require some sort of authentication. Because of this, it may be expedient to give the new recipient of this email access to the old account as a whole, or the new recipient may just forward requests to unsubscribe from various lists to the local organization's postmaster for processing. In any case, it's not appropriate for anyone to send out email posing as someone who has left an organization for anything other than the purpose of unsubscribing from a mailing list.

A related issue is whether or when a given email address should be made available for reuse by an email service. After someone has left an organization and an email address has been forwarded somewhere else for a while, it should go unused for some additional amount of time before being pressed back into service. For a loose organization, company, or government organization, it may make sense to permanently retire any email address after its only owner has left the organization, to avoid any possible confusion. For a large service provider or a large academic institution this may be impractical. In any case, at the very least several months should pass before a "fallow" email address is put back in service, and we would recommend that waiting for a year would be entirely appropriate in most circumstances.

The amount of time before email addresses can be reused should be spelled out in the email administration policy document, and some mechanism besides personal memory should be used to ensure that this occurs. This method could be as simple as

replacing the account's password field in the passwd file with an "account may be reused on <date>" for small organizations, to much more complex automated procedures for large service providers.

In corporate or governmental realms, archiving old email accounts is likely to be at least a good idea if not a legal requirement. Once either the email account has been turned off or forwarding has been implemented and tested, existing email should be moved to archival storage. This might be a special file system or some pre-archival staging area used as temporary storage until the data is transferred to some other medium. In most cases, email archiving should be performed in conjunction with archiving other data that needs to be preserved from the same account.

## Username Assignment

A great deal of heat can be generated on the subject of the assignment of the user portion of email addresses. This booklet weighs in on these issues, but at the very least we hope there can be agreement from the technical community that this issue is more subtle than it might at first appear. For every username scheme anyone can come up with, we can think of several ways in which that scheme is inferior to another. What is most important is that each organization consider the consequences of possible email addressing schemes before implementation.

We believe that email addresses should primarily be considered to be opaque identifiers. That is, exclusivity and clarity should be the foremost considerations in adopting a naming convention. While ease of recollection, brevity, and likelihood of association between the email account and the real person are all very desirable traits for email addresses, these should be secondary considerations. Not everyone agrees with this prioritization, but there are good technical reasons to adopt this philosophy. Even if technical merit does not ultimately carry the debate on this topic, an organization depends on its technical staff to present the technical side of an argument, so we advocate this point of view here.

This document isn't the first to recommend against adopting a firstname.lastname email address convention or similar theme, but we join the chorus in believing that while doing so initially may appear attractive, it is quite likely to lead to considerable confusion and misdirected email. For a more thorough dissection of this issue, we direct the readers to question 3.5 of the Sendmail FAQ [Sendma97] or to section 19.1.2 of *The Practice of System and Network Administration* by Limoncelli and Hogan [LimHog02]. Suffice it to say that we recommend against this naming scheme for the same reasons as these other sources.

We would like to recommend a single, simple, universal email address naming scheme that worked perfectly in all cases. Unfortunately, every scheme that we have seen adopted, and the authors have probably seen every or nearly every possible method in practice, has disadvantages as well as advantages.

For small organizations, the use of first names is popular. As email addresses these have the advantages of tending to be short and easy to remember. On the downside,

only in the smallest organizations are they unique. Once a second person named "Greg," for example, joins an organization, email WILL get sent to the wrong person. As the organization matures, the consequences of this may not be as benign as they were at its outset.

Unfortunately, the only way to avoid ambiguity problems in their entirety is to use indecipherable identifiers, such as random strings of letters and numbers, such that they could never be mistakenly associated with any particular person. Of course, this scheme produces email addresses that cannot be remembered, which is also undesirable. Any reasonable naming scheme will have to compromise human factors (brevity, ease of recollection) with the potential for confusion.

Another problem with the first-name convention is that these addresses are popular with spammers performing dictionary attacks against domains. Someone with the email address of "jim" at a given domain will receive more spam than will "jsl2645." This is a good reason to make usernames at least a little more complex than just using first names. It is possible that the desired level of complexity will need to be increased as spamming techniques evolve.

Other mechanisms that are short, reasonably memorable, and usually provide at least some resistance to dictionary attacks are initials. First name plus the last initial, first initial plus the last name, or first two initials plus the last name (if the names in question are reasonably short) work well for moderately sized organizations, although collisions are still possible.

In any case, we believe the ideal lengths for the user portion of email addresses be between three and eight characters, especially if they will also be used as usernames. Longer strings of characters are appropriate if they are especially easy to remember (such as postmaster), if they are not used as login accounts, or if the number of email addresses supported is so large that shorter names would inordinately restrict the name space. These conditions certainly hold for the larger service providers.

A question each organization needs to address is whether users will be allowed to select their own email addresses. In the case of a service provider, the answer is almost certainly "yes," although some necessary restrictions on length and potentially on content may be appropriate. For large service providers, certain portions of the name space may be so densely packed that it would be prudent for the system to make some suggestions during the registration process if a user's preferences aren't available.

Many business and governmental organizations prefer to assign email addresses in order to present a uniform and businesslike image. The authors appreciate the desire to present a professional demeanor, but we generally believe that under most circumstances users should be permitted to select their preference as long as it is appropriate for the organization. Requiring the email address to be based on the employee's name is certainly a reasonable restriction, but we don't see the need to be much more restrictive than that in almost every case. Of course, we do not mean to suggest that preferences should overrule technical concerns, such as name lengths if email addresses are also used as usernames.

On occasion someone will want to change their email address. There may be good reasons for this, or the desire to do so may be based on a whim. Legal name changes and abandoning addresses that are inundated with spam are two good reasons to want to change an email address, and changing them for these reasons should be possible and allowable.

For service providers and loose organizations, it should be possible for customers to make such a change, although it's reasonable to restrict the frequency with which these changes can be made, say, to some small number of times per year. In a more formal organization email address changes can be more costly, as it likely requires the update of mailing lists, the reprinting of business cards, the potential for confusion, etc. Because of these costs it's reasonable to restrict the circumstances under which these changes can occur, but the email service should be set up to handle such changes as gracefully as possible.

The old email address should remain as an alias for the new one for a reasonable period of time. The process of deprecating and expiring old usernames should follow the same sort of path that was presented for forwarding and email address deactivation earlier in this chapter. That is, first the address should be transparently forwarded, possibly followed by a period in which email is forwarded but the sender is warned about the address change, followed by a period when the address is guaranteed to bounce, potentially followed by putting the address back in the available pool for reuse.

## Misdirected Email

It's guaranteed that a person will send email to someone (often several someones) they did not intend to be recipients. Everyone who has used email as a communications method for a significant period of time has both sent and received email of this sort. Sometimes such email can trigger quite a panic, so it's probably a good idea to consider how this should be handled before it happens.

If someone receives email by accident it is appropriate to reply to the message to let the sender know that their email was missent. If the email was obviously sent by accident to a wide distribution list, then it's probably not necessary to respond to it: the sender will be made aware of this fact by one, if not many, of the recipients. Many times spammers will disguise the messages they send in order to make them seem like they are misdirected when they are not. It can be difficult to tell the differences between these and legitimately misdirected email messages on occasion, but email recipients should be aware of this possibility.

A question arises as to what to do if the misdirected message content is especially sensitive. If it contains indications that a crime is being or will be committed, the message ought to be revealed to the proper authorities. Demonstration of intent to act contrary to the mission of the organization (e.g., to inappropriately divulge sensitive information, defraud the organization, threaten or harass a colleague, etc.) should also be shown to the appropriate authorities.

It is difficult to know how to respond to a message that is inappropriate but that

does not rise to the level of the offenses mentioned in the previous paragraph. In a corporate, governmental, or academic environment, there should be someone to whom such an issue can be brought in confidentiality so that the recipient can decide what would be the appropriate response. In the absence of specific advice to the contrary, our recommendation for handling email that discusses sensitive topics that are neither illegal nor contrary to the mission of the institution is to inform the sender that the message was misrouted and then to ignore its contents.

## Email Encryption

Encryption of email messages is done to protect the message against inappropriate disclosure, either through accident or malice. If the concern is that email should be protected against eavesdroppers who may be passively recording email traffic, then it is appropriate to employ an opportunistic host-to-host encryption system, such as the standards-based STARTTLS [Hoffma99] mechanism. An application-layer encryption system such as S/MIME [Ramsde04] or PGP [Lucas06] should be used if it is necessary to provide protection against disclosure at the end points as well as in transit.

A nice feature of host-to-host encryption is that once it is set up the encryption requires no additional effort from the user in order to realize its benefits. Whether the encryption is employed from email server to email server via STARTTLS, or from email client to IMAP server via TLS, or network to network via an SSH tunnel, the encryption occurs transparently. Additionally, since the messages themselves are not stored in an encrypted format, there are no issues regarding key escrow to ensure the ability to be able to recover official encrypted email.

The downside is that these messages receive no protection at the end points. Sending email encrypted via STARTTLS won't protect against the message being read by the email administrator on the recipient's server. Basically, these methods are useful for protecting communication against prying eyes on the wire, but don't provide much protection beyond that.

With an application-layer encryption mechanism the message is encrypted either by an email client or on the local file system before being imported to the email client, and then the encrypted message is sent as an opaque payload via the email protocols to the recipient. It is presumed that the recipient has the means to decode the message. In this manner the message remains protected at every intermediate step along the way.

Where laws allow the unrestricted use of encryption, for a service provider, a loose organization, and, usually, an academic institution, it's fine if only the recipient can decrypt a given message. A problem arises, however, if the message contains information valuable to a company or governmental organization and the recipient is not available to decrypt the message. We would expect those organizations who use application-layer encryption as a supported part of their email service to set up a form of key escrow system in order to ensure that important messages will be institutionally recoverable. Key escrow may also be provided by academic institutions and service providers to allow the use of encryption by their email customers without running afoul of local laws.

With a key escrow system in operation, each message is, of course, no more secure than access to the key escrow system. These systems require considerable thought to set up and maintain properly. As a technical matter, doing so is beyond the scope of this booklet, except to say that if a key escrow system is in use at an organization, information on its use should be a part of the appropriate email usage policy document.

## Email Quotas

It seems that regardless of the amount of storage space that is deployed, usage quickly expands to fill whatever is available. Consequently, in many organizations it is necessary to manage how much space each user occupies for various purposes. Sometimes it is necessary to manage email storage by limiting how much space is consumed by each customer.

For a message store that is accessed directly or by the POP protocol, applying quotas to email storage is fairly straightforward. Allow arbitrary email delivery until the user's mailbox exceeds some threshold, after which delivery of additional messages will not be permitted.

When an email delivery fails due to the recipient's mailbox being over quota, some email server software will treat such an event as a temporary failure and will queue the message up for redelivery. Other software will treat this event as a permanent failure and bounce the message. On servers that are lightly loaded and have few concerns about being mailbombed, as is often the case for loose organizations, treating an over-quota delivery as a temporary failure can make sense. For most email servers, however, especially for those that handle a high volume, it is almost always a better idea to bounce the message. In some cases, implementation-specific issues may add weight to the argument for just queueing these messages.

If a message bound for an over-quota mailbox is queued rather than bounced, space on the email server really isn't being saved, it's just being consumed somewhere else. For performance reasons, it is more efficient to deliver a message than have it take up space in the queue. Moreover, if someone mailbombs a particular email address, it may become necessary to just stop accepting messages for that account at some point, and this doesn't do any good if the messages are queued.

## Delegated Mailing List Management

If an organization supports mailing lists, usually at least some of them are managed by the postmaster group. In some cases, it may be decided to turn over day-to-day management of some lists to people who aren't responsible for email service maintenance. If this is done, those mailing list managers (here we refer to the people, not the software) need to abide by the rules set down by the postmaster group.

In essence, it is important that everyone understands that even if the list was created, is owned, and is being managed by someone outside of email administration, email list administration is still a responsibility that is being delegated from those who have

authority over the email service as a whole. Email administration still has final say over how the email service performs, and their rules regarding the operation of the mailing lists they ultimately support must be followed.

As mailing lists grow, it's possible that their use can interfere with the operation of the greater email service. In this case, special measures may need to be taken in order to reconcile the demands both uses place on the system. It is important to spell out what the email service's priorities will be before such a conflict arises.

## Email Use in Marketing

Email is a powerful, low-cost medium by which an organization can keep in contact with its customers or constituents. As we all know, the reasons it is favored as a legitimate communications medium are the same reasons that make it prone to abuse through the practice that has come to be called spamming.

Not everyone is in complete agreement as to what constitutes spam and what is legitimate marketing. Some well-known brands and organizations operate within this gray area. Our recommendation is that this is a realm in which it is especially important to behave with the utmost propriety. We'd advise any organization to do whatever they need to do to minimize the chance that current and potential customers would be annoyed or offended.

Whether or not a given email service is subject to laws requiring that email lists be opt-in only, we believe that this is the only proper way to proceed. Under no circumstances should email be sent to an address unless the owner of the address has explicitly taken action to opt-in to that specific list. Moreover, the list owner should take some action to verify that the recipient is the one that is subscribing to the list. This can be accomplished with a confirmation message of some sort.

All bulk email messages should include clear instructions on how a recipient can unsubscribe from that particular list. One of the methods available should include responding to the message in question and requesting that the recipient be unsubscribed. The sending message should clearly identify both the organization standing behind the message itself and the organization that is doing the actual sending, if these are different. One of the main differences between spammers and legitimate email marketing is that folks with whom the recipient has a beneficial relationship don't have anything to hide.

Occasionally, someone signs up for an email list, confirms this fact, and then complains that the list content is spam. In this situation the recipient should immediately be removed from the list in question. It does no good to argue with the complainant about whether they actually signed up or not. It's clear from their response that they don't want such email, and that's sufficient reason to remove them from the list.

If an organization, such as an ISP, site postmaster, or anti-spam database maintainer inquires based on such complains from their own customers, then producing evidence that the person in question actively signed up may be appropriate. It's important to be

polite, responsive, understanding, and open with these sorts of inquiries. Inquirers don't see this sort of response from spammers, so it's usually the most effective way to indicate subscription to a different philosophy.

Some bulk email senders provide notice that it may take some large number of days to remove someone from their address lists, sometimes several months. Most recipients will take this as an indication that customer service is not a high priority for these organizations. Having run very large email services ourselves, the authors cannot think of a good reason why removing an address from a distribution list should take longer than 48 hours at the absolute outside, unless customer satisfaction really isn't important.

More than with other forms of marketing, a poorly run email campaign has the potential to quickly and severely damage an organization's reputation with customers and potential customers. It is important that each organization decide how it will handle bulk electronic interactions with its customers, that these policies be carefully considered and explicitly documented, and that high ethical standards be maintained by the organization and the agents that may be communicating on its behalf.

# Appendix 1: Email Policy Checklists

Every organization will have its own requirements regarding what issues need to be included in its email policy documents. We can't hope to create an exhaustive list here, nor will it be necessary to include every consideration listed here in every organization's documents.

Our purpose is to list as many issues as we can that may be important to some organizations. When creating policy documents, these lists can be consulted as a collection of suggestions. In each document we expect some of the issues we provide to be included and some not. However, we hope that the information here will serve as a useful starting point for formulating the list topics each organization's policy documents will cover.

Presented here are two checklists. The first, an Email Usage Policy Checklist, is, as the name suggests, a list of potential topics for inclusion in email usage policy documents. These documents focus on the email provider's expectations for customer behavior with regard to email use and customer expectations of the email service.

The second, an Email Administration Policy Checklist, is a list of topics for policies regarding the administration of the email service. Such policies would spell out what professional behavior is required of all postmasters, email administrators, system administrators, and other personnel involved in the support of electronic mail services.

## Email Usage Policy Checklist

**General Usage**

[ ] For what purposes can electronic mail be used?
- [ ] Business
- [ ] Personal
- [ ] Confidential communication

[ ] What activities are forbidden?
- [ ] Spam
- [ ] Malware dissemination
- [ ] Illegal activities
- [ ] Harassment

[ ] Who is the primary owner of the email account?
- [ ] Customer
- [ ] Employer

**New Accounts**

[ ] How are usernames selected?

[ ] What are policies regarding passwords?
- [ ] Temporary password dissemination
- [ ] Time frame for changing passwords
- [ ] Recovery of forgotten passwords
- [ ] Password generation policies (strength of passwords)

**Changing Accounts**

[ ] Under what circumstances may account names be changed?

[ ] How are changes requested?

[ ] What are the policies for forwarding after a name change?

**Removing Accounts**

[ ] How is email forwarding handled?
- [ ] Is it permitted at all?
- [ ] For what time frame?

**Email Privacy**

[ ] What is the expectation of privacy for email users?

[ ] Under what circumstances can email be read by another?

[ ] There is always the possibility of accidental disclosure.

**Content Filters**

[ ] What content filters are in use?
- [ ] Anti-spam
- [ ] Malware

[ ] How much control is the user allowed to have in their email filtering?

[ ] Is suspicious email quarantined?
    [ ] User access to quarantined email

**Access and Storage**

[ ] How much space is allocated to each account?
[ ] Are there expectations on how frequently email is checked?
[ ] Are there expectations on how email is stored?
[ ] How are users responsible for email archiving?
[ ] Is there a maximum permitted message size?

**Remote Access to Local Email**

[ ] Can local email services be accessed from outside the LAN?
[ ] What precautions must be taken to limit vulnerabilities caused by this access?

**Local Access to Remote Email**

[ ] Are connections to remote email services allowed from the local network?
[ ] What precautions must be taken to limit vulnerabilities caused by this access?
[ ] Can outside email addresses be used for conducting official business?

**Encryption**

[ ] When is it appropriate/legal/required to encrypt email?
[ ] Is key escrow performed?

**Misdirected Email**

[ ] How should misdirected email be handled?
    [ ] When sent
    [ ] When received

**Personal Messages and Liability**

[ ] Who is responsible for the contents of email sent?

**Consequences of Inappropriate Use**

[ ] Under what circumstances will accounts be deactivated?
[ ] Under what circumstances will account be reactivated?
[ ] What behaviors will result in account/employee suspension?
[ ] What behaviors will result in employee dismissal or service termination?
[ ] What is the channel for appeals?

**Updates to Policy**

[ ] How are updates to this policy handled?
[ ] How is information about these updates disseminated?

## Email Administration Policy Checklist

**Account Changes**

[ ] Who authorizes new account creation?
[ ] Who authorizes account renaming?
[ ] Who authorizes account suspension/termination?
[ ] By what mechanism are accounts temporarily suspended?
[ ] Will email backups be performed?
[ ] If so, how are email backups to be performed?
[ ] How are unintentional backups avoided?
[ ] How are backups tested?

**External Resources**

[ ] What connections to external data sources will the email service make (DNS, NIS, LDAP, Active Directory, etc.)?
[ ] How will these connections and services be secured?

**Monitoring**

[ ] In what ways will the email system be monitored?
[ ] How might that affect the confidentiality of email?

**Looking at Others' Email**

[ ] Who can authorize access to someone's email?
[ ] What is the procedure for doing this?
[ ] What steps are taken to reduce the chances of accidental disclosure?
[ ] What should be done if email is disclosed without authorization?
[ ] What can employees look at without authorization when troubleshooting email services?

**Updating Content Filters**

[ ] Under what circumstances is this done?
[ ] How is this tested?

**Upgrading Email Hardware/Software**

[ ] For typical upgrades, how is this performed?
    [ ] How much warning?
    [ ] When are upgrades performed?
    [ ] Target for frequency of update
    [ ] Testing procedure
[ ] What is the procedure for emergency upgrades (e.g., security patches)?
[ ] What is done to ensure that no email is lost during upgrades?

**Dealing with Legal Issues**

[ ] How are subpoenas handled?

[ ] Under what circumstances is law enforcement contacted?

    [ ] Contact list

[ ] What legal requirements are in place for email services?

**Policy Violations**

[ ] What behaviors will result in account/employee suspension?

[ ] What behaviors will result in employee dismissal or service termination?

[ ] What is the channel for appeals?

# Appendix 2: Other Resources

There are many sources of good information on providing email service, far too many to list here. We expect that Internet postmasters may find some sources of information especially valuable, however, and these we have listed below.

## RFCs

The Internet RFCs (Request for Comments) form the basis for the interoperability of the Internet and, therefore, are invaluable to those responsible for supporting Internet-based services such as Internet email. Many Internet email protocols, including SMTP, POP, IMAP, MIME, STARTTLS, SMTP AUTH, among others, are standardized in RFCs. Many other protocols useful or necessary for Internet email operation are also standardized in RFC documents, including DNS, TCP/IP, LDAP, and many more.

Aside from Internet standards, many RFCs provide useful advice or important historical context for understanding how the Internet works on a day-to-day basis. These may be very valuable, especially to those who are new to electronic mail administration or who have not had a detailed education in the operation of the Internet.

Many RFCs relevant to email use have been cited in this booklet and are listed in the References section. Many more that may be useful are available through the public RFC repositories. One source for RFC documents is http://www.rfc-editor.org/rfc.html.

A current list of the Internet RFCs is at http://www.isi.edu/in-notes/rfc-index.txt.

## Books

A fair number of books deal with the nuts and bolts of system administration, but very few discuss the "big picture" issues that IT professionals face. One of these that we strongly recommend has already been referenced in this booklet: *The Practice of System and Network Administration,* by Limoncelli and Hogan. We recommend it to all system administrators, especially those in decision-making positions.

*The ePolicy Handbook* (AMACOM, 2000), by Nancy Flynn, discusses corporate policies governing information technology issues in some depth. This book focuses on policies for more structured environments, such as companies and governmental organizations, but some of its suggestions can be adapted to other settings. A few of the author's recommendations are presented less flexibly than we'd prefer, and the book is aimed more at corporate management and human resources folks than technical people, but we believe it can be valuable to those looking for additional resources on this topic. In many ways it provides a less technical counterpart to this booklet.

With Randolph Kahn, Flynn has written another book relevant to some of the topics discussed in this booklet: *E-Mail Rules* (AMACOM, 2003). This book is essentially an expansion of the discussion of email topics found in *The ePolicy Handbook*. Again, it is aimed at a less technical audience than this booklet, and it discusses only those organizational types where a more authoritarian approach toward email will make sense. It does, though, reiterate and expand upon many of the issues we bring up in this booklet.

## Booklets

This booklet is part of the ongoing Short Topics in System Administration series published by SAGE, the USENIX Special Interest Group for Sysadmins. Many booklets in this series are on topics that may be useful to an Internet postmaster.

A complete list of these booklets and links to ordering information can be found at http://www.sage.org/pubs/short_topics.html.

## Internet Services

Several services available on the Internet likely will prove valuable to postmasters. We list several of these here.

**rfc-ignorant.org**

This Web site hosts a database of those domains that do not follow testable Internet standards on electronic mail. Self-respecting postmasters do not want the domains they manage to be included on this list.

Submissions to this list are manual, and some of the criteria for inclusion are borderline standards violations, but on the whole we expect those sites who follow the guidelines in this booklet and other Internet email best practices are unlikely to be added to this database.

**dnsreport.com**

This is a wonderful free service that both aids in setting up DNS and provides a great deal of useful DNS information. Just type in the name of the domain one wishes to test and the site will report a large amount of data on the DNS records for that domain, including information on email routing. This Web site provides good advice that is worthy of consideration.

**whois**

The public WHOIS databases contain information about domains and IP address ranges. This is the primary method by which contact information is obtained for Internet domains. Most operating systems come with a WHOIS software client. Web-based WHOIS interfaces are available as well, including the public WHOIS service hosted by Internic: http://www.internic.net/whois.html.

## SANS Policy Documents

The SANS (SysAdmin, Audit, Network, Security) Institute is an excellent source for information on information security and policy. One of the resources they provide is a set of examples and templates for various types of IT policies. An index of these documents is available at http://www.sans.org/resources/policies/.

PDF-formatted documents about specific policies that are especially relevant to the topics of this booklet include the following:

Email Use Policy: http://www.sans.org/resources/policies/Email_Policy.pdf.

Email Retention Policy: http://www.sans.org/resources/policies/email_retention.pdf.

# References


[Allman05] Allman, E. "DomainKeys Identified Mail (DKIM) Introduction and Overview." 2005, http://www.dkim.org/info/DKIM-Intro-Allman.html.

[Beerte93] Beertema, P. "Common DNS Data File Configuration Errors." RFC 1537, October 1993.

[Crocke82] Crocker, D. "Standard for the Format of ARPA Internet Text Messages." RFC 822, August 1982.

[Crocke97] Crocker, D. "Mailbox Names for Common Services, Roles and Functions." RFC 2142, May 1997.

[Dijker96] Dijker, B. *A Guide to Developing Computer Policy Documents.* Short Topics in System Administration #2, SAGE, 1996.

[Hoffma99] Hoffman, P. "SMTP Service Extension for Secure SMTP over TLS." RFC 2487, January 1999.

[Klensi01] Klensin, J. "Simple Mail Transfer Protocol." RFC 2821, April 2001.

[Law03] Law, G. "Sobig.F Breaks Speed Records." *PC World,* August 21, 2003. http://www.pcworld.com/news/article/0,aid,112108,00.asp.

[LimHog02] Limoncelli, T., and Hogan, C. *The Practice of System and Network Administration.* Addison-Wesley, 2002.

[Lucas06] Lucas, M. *PGP & GPG: Email for the Practical Paranoid.* No Starch Press, 2006.

[PreSke02] Preston, W.C., and Skelly, H. *Backups and Recovery.* Short Topics in System Administration #9, SAGE, 2002.

[Ramsde04] Ramsdell, B. "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.1 Message Specification." RFC 3851, July 2004.

[Sendma97] Sendmail Consortium, "Sendmail FAQ." http://www.sendmail.org/faq/index.html.

[SinBir04] Singer, A., and Bird, T. *Building a Logging Infrastructure.* Short Topics in System Administration #12, SAGE, 2004.

[VanBok90] Van Bokkelen, J. "Responsibilities of Host and Network Managers: A Summary of the 'Oral Tradition' of the Internet." RFC 1173, August 1990.

[Vaudre96] Vaudreuil, G. "Enhanced Mail System Status Codes." RFC 1893, January 1996.

[WonSch06] Wong, M., and Schlitt, W. "Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1." RFC 4408, April 2006.

## About the Authors

Nick Christenson (http://www.jetcafe.org/~npc/) has designed, implemented, and maintained reliable, high-performance Internet services based on open systems while at a wide variety of organizations, including the Jet Propulsion Laboratory, EarthLink, Sendmail Inc., and Sistina Software. Nick is currently working as a consultant based in Las Vegas, Nevada.

Brad Knowles (http://www.shub-internet.org/brad/) has specialized in Internet email and DNS administration for more than a decade, and has provided the benefit of his experience to the U.S. Department of Defense, America Online, and Collective Technologies, among others. Among other things, he is currently setting up his own consulting company in Austin, TX, and is writing his first book.