ROARING PENGUIN SOFTWARE INC.

White Paper

Fighting Spam in an ISP Environment:

Challenges, Solutions and Best Practices

April, 2007

Summary

An ISP presents one of the most complex environments for managing spam because of the high volume of e-mail, wide variety of users and high level of service demanded by customers. This white paper reviews the most common approaches to spam management in an ISP environment and details how a suitable solution must address particular challenges for ISP administrators.

The paper provides comprehensive best-practices for spam management in an ISP environment, including guidelines for evaluation, roll out and deployment of an anti-spam solution. Roaring Penguin has developed these best practices through extensive work with ISPs in resolving their spam challenges, and has developed the CanIt-PRO anti-spam solution in conjunction with ISPs across North America.

The Challenge of Spam for ISPs

The challenge to solve the problems of spam e-mail is reaching the desks of ISP CTOs across North America and Europe. A combination of factors has elevated the urgency for better management of spam, including:

- **Pressure from end-users** and the ensuing tax on help-desk staff in dealing with spam-related complaints.
- The **need to protect networks** and end users from malicious threats carried via e-mail.
- The increasing **collusion of commercial spam with fraud schemes**, phishing attempts and malicious viruses.
- A desire for tools to help ISPs **protect customers** from unwanted e-mail.

However, ISPs present one of the most complex environments for managing spam because of the high volume of e-mail, wide variety of users and high level of service demanded by customers. This white paper reviews the most common approaches to spam management for ISPs and details how a suitable solution, such as Roaring Penguin's CanIt-PRO™, must address particular challenges for ISP administrators.

The paper concludes with comprehensive best practices for spam management for ISPs. It includes guidelines for evaluation, roll out and deployment of an anti-spam solution. Roaring Penguin has developed these best practices through extensive work with ISPs in resolving their spam challenges.

Common ISP Solutions to Spam

ISPs generally consider three types of anti-spam solutions:

- **Outsourced filtering services** relay your mail through a third party system housed off site. These systems can sometimes be too costly for the typical ISP. They can also present concerns about a loss of control or security over ISP e-mail.
- **Home grown solutions** are typically based on open-source software such as SpamAssassin™ and Roaring Penguin's MIMEDefang. In most cases, these solutions were sufficient until about 2002, when the volume of spam and spammer's ever-evolving techniques began to render home-grown solutions unmanageable.
- **Third-party in-house solutions** offer a balance between the above two options. While benefiting from the experience of a third party provider, you keep the flow of your email traffic on your own network.

A solution that deals with spam at the mail server is preferable for ISPs. Such a solution is centrally manageable and deals with the problem before spam and viruses reach end-users or consume network resources.

Getting Anti-Spam Right at an ISP

The main challenge for an ISP looking for an anti-spam solution is to find a product for the server that also provides the flexibility ISPs need. For example, ISPs cannot define messages that look like newsletters or product information as spam, as some users may have legitimately requested this information. In addition, some ISPs do not want to be thought of as censoring content, unless explicitly asked to do so.

Other key requirements for an effective ISP anti-spam solution typically include:

- Ability for **administrators to set global policies**, in particular as regards the filtering of viruses and certain types or sizes of mail attachments.
- Ability to let **end-users make the final decision** about what is or is not spam – or even to opt out of spam filtering all together.
- Integration with existing user authentication engines and web-mail interfaces, so that users need only **one password for all mail services**.
- Integration with existing or preferred **anti-virus** solutions.
- **Extensibility**, including the ability to cost-effectively expand the system, to filter outgoing mail, to easily add more users, and to write customized rules and filters if required.
- **Minimizing the performance overhead** that content filtering can put on systems.
- Ability to **roll-out the solution in manageable increments**, reducing strain on help desk staff.
- Ability to **meet privacy concerns** of end users.
- **Cost-effective pricing**, or pricing structures suited to the ISP environment and budget.
- **Automatic updates** of anti-spam rules or “signatures”, and even of the anti-spam software itself.

Anti-Spam Best Practices for ISPs

Based on Roaring Penguin Software's work with numerous ISPs, we have assembled best-practices for ISPs looking to evaluate, plan for and deploy an anti-spam solution. Itemized below, these best practices are useful during all phases of an ISP anti-spam project. They will assist in requirements development, evaluation and deployment.

Evaluation Best Practices

ISPs can apply several rules of thumb to evaluate and budget for anti-spam solutions. These rules of thumb are also useful in developing requirements for an anti-spam solution. The shortlisted solution(s) must be tested using real mail, on servers that are connected to the Internet.

Evaluate Realistically

Plan to evaluate the solutions(s) using a broad range of e-mail types. This means that both non-spam and spam messages must be used from mailboxes across the ISP. Testing using just the administrator's mail, for example, won't necessarily provide a realistic sampling of end user mail. Testing using only messages known to be spam is equally problematic, as the solution's statistical filters need both types of input to learn.

Plan to test using all kinds of e-mail from multiple accounts of varying user types. Consider starting with internal accounts only so that you can evaluate the product without disrupting your end users. Be sure to fully evaluate all of the customizations that your end users will take advantage of prior to deploying system-wide.

Set Up Test Mail Streams

Plan how you will setup the evaluation mail streams. Either of these two methods is recommended:

1. Split your e-mail streams so that a portion of e-mail goes thorough the system under evaluation. This is relatively simple if you have multiple e-mail domains. For example, you might choose to set up CanIt-PRO to handle e-mail from two of five domains.
2. Another equally effective method is to select particular e-mail accounts and forward mail addressed to those accounts to the evaluation system. Start with users who have been most vocal about needing a solution, and be sure to include users with unique needs, such as those that receive newsletters or are subscribed to multiple mailing lists.

In either case above, you should also plan how you will inform and train those users whose mail will be used to test the solution so that they can provide you with feedback.

Review the Quarantine

To determine false-positive rates and assess how best to tweak default settings for your particular e-mail environment, you will want to review the quarantined (or “trapped”) e-mail daily during the course of the evaluation period. When evaluating CanIt-PRO, decide whether you will provide the test subjects with access to their own Spam Traps or whether you will appoint a single “spam officer” to review the trapped messages.

Train the Adaptive Filters

The rapidly changing nature of spam demands that an effective anti-spam solution must include adaptive filters. Also known as “Bayesian” or “adaptive” filters, these learn what is or is not spam over time by cataloging words and/or phrases that appear in messages rejected as spam or accepted as legitimate.

To truly evaluation the effectiveness of such a solution, it is critical to train the filters past their threshold level. Plan to train the filters on at least 100 spam and 100 non-spam messages. With CanIt-PRO, the filters are most easily trained by embedding “training links” at the bottom of each e-mail. By clicking on “This is spam” or “This is not spam” (see Figure 1), users can quickly train the filters themselves. Administrators can independently train CanIt-PRO's adaptive filters by taking action (such as “accept” or reject”) on messages in the Spam Trap.

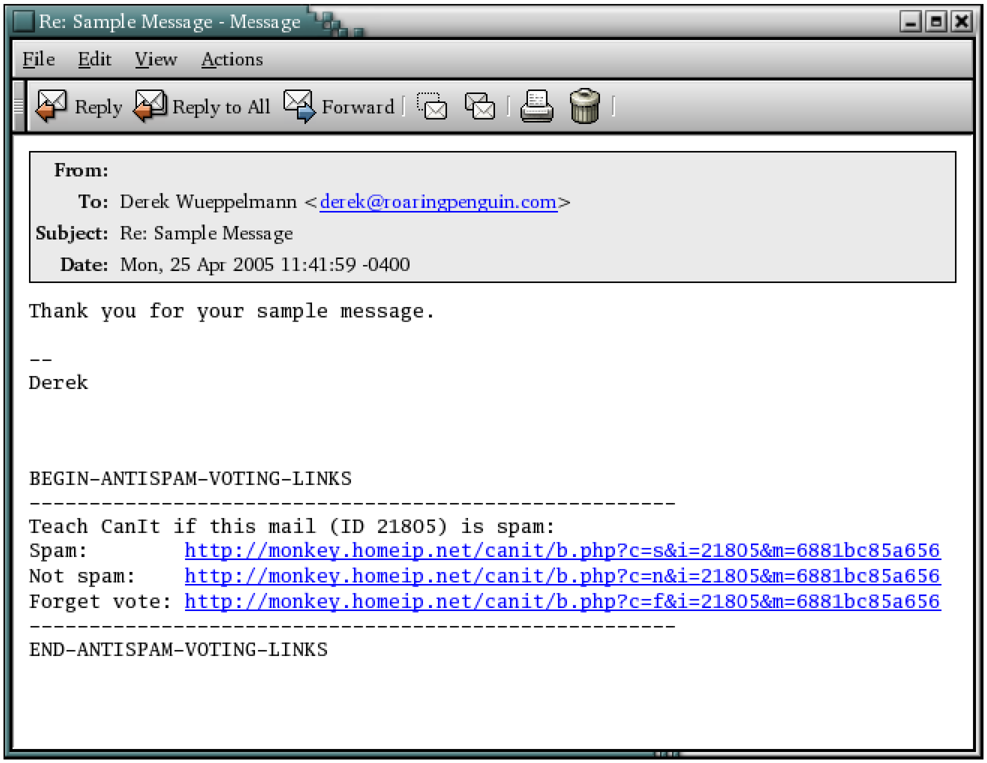

From:
To: Derek Wueppelmann <derek@roaringpenguin.com>
Subject: Re: Sample Message
Date: Mon, 25 Apr 2005 11:41:59 -0400

```
Thank you for your sample message.

--
Derek



BEGIN-ANTISPAM-VOTING-LINKS
------------------------------------------------------
Teach CanIt if this mail (ID 21805) is spam:
Spam:        http://monkey.homeip.net/canit/b.php?c=s&i=21805&m=6881bc85a656
Not spam:    http://monkey.homeip.net/canit/b.php?c=n&i=21805&m=6881bc85a656
Forget vote: http://monkey.homeip.net/canit/b.php?c=f&i=21805&m=6881bc85a656
------------------------------------------------------
END-ANTISPAM-VOTING-LINKS
```

Figure 1: Sample Training Links Embedded by CanIt-PRO

Ballpark the Cost

As the various vendors of anti-spam solutions have different pricing structures, a direct price comparison can be difficult. However, certain rules of thumb can be applied when setting your anti-spam budget. These are:

- **Paying per user will be more cost-effective than paying per-server**. Paying per-server is like paying a tax as the volume of spam increases, whereas paying per-user allows you to scale the solution across as many servers as you need.
- **Expect volume discounts based on number of users**. If your ISP has more than 1000 users or more, you should expect volume pricing. Ask about how pricing scales with the number of licenses purchased.
- **Do not pay extra to stop different or new “types” of spam**. Spam

is spam. From a technology perspective, spammers have a finite number of methods at their disposal, and an effective solution will be able to detect all kinds of spam. Don't be tricked into paying extra for "unique" kinds of spam like *phishing* attempts, *fraud* and the like – because exactly the same anti-spam techniques will catch them all.

- **There may be hidden costs**. When comparing the prices of various solutions, consider what non-advertised costs may be involved. For example:
  - How much hardware will be required, and how powerful does it need to be?
  - Will you need to purchase additional anti-virus software, and if so, at what cost?
  - Will you own the code or will you be dependent upon the provider in perpetuity?
  - How extensive and scalable is the solution?
  - Will you likely need a new solution if spam increases and as spammers change their tactics?
  - What will be the additional costs to add end user functionality or add new features in response to end user feedback?

  All of these considerations are relevant to assessing the cost of an anti-spam solution.

Roll-Out Best Practices

The following best practices have proved to increase the success of ISP anti-spam deployments. Ideally, all of these practices should be followed prior to and during the roll-out of a new solution; however, they are most relevant for a solution that gives end users control, such as CanIt-PRO.

Develop Concise User Documentation

Basic user documentation should be made available in conjunction with the initial roll-out. Look to your anti-spam provider to provide you with the documentation templates as part of the sale. Typically, you will customize the documentation to suit your particular user population and deployment. Plan to provide the documentation via a web page within the e-mail services website.

Implement a Training Program

CanIt-PRO ISP customers use the following methods to train their user population about the system:

- Help desk staff training. These are usually the first groups to obtain hands-on training, typically as part of the evaluation process.
- On-line user documentation, as detailed above. This provides the end user a place to read up on how to use the new anti-spam features.

- Encourage use, as explained below.

Roll Out in Phases

Perhaps the most important aspect of an ISP anti-spam project is a phased roll out. Typically, CanIt-PRO customers deploy the solution in an “opt-out -by-default” state. This means that:

- Global filtering policies, such as filtering for viruses and attachments, are immediately applied to all e-mail boxes.
- Users must voluntarily login and choose to have their e-mail filtered for spam.

This makes the system available to all users while allowing for incremental increases in use of the system. That minimizes burden on the help desk and allows administrators to carefully monitor impact on servers as usage increases.

Having deployed the solution in an opt-out state, the administrators can begin to advertise availability of the anti-spam option. This can be done by sending out a system wide email notification about the new service, or including links on your web site in a section detailing services or in a user's web portal section. Help desk staff who are fielding e-mail related questions can also point out the new features to the end users.

If you plan to provide users with advanced features such as access to individual Spam Traps, whitelists and blacklists and custom settings, do so only once they have become aware of and accustomed to the basic solution. Based on user feedback, you can determine what advanced features are the most important. When the time comes to roll out an advanced user interface with additional options, be sure to revisit these best practices.

Encourage Use

Informing users about the solution once is not generally enough. Users must form new habits to make the system effective, and their use of the system should be encouraged. Two features of CanIt-PRO are particularly useful in this:

- **Embedded training links**. As shown in Figure 1, CanIt-PRO enables administrators to automatically embed “training links” at the end of every incoming e-mail. Users can simply click a link to tell the system that the e-mail is “spam” or “not spam”. This is a constant visible reminder to users and a very easy way to train CanIt-PRO's adaptive filters.
- **E-mail notifications**. Through an automatic e-mail, users can be notified periodically (daily, weekly, etc.) about the system in a way that causes them to act. For example, they may be notified that there are messages pending in their Spam Traps. Or, users may be notified that messages in their Spam Traps will be deleted within a certain time frame if they do not take action on them.

Deployment Best Practices

Although every networked e-mail environment and each anti-spam solution will be unique, the following rules of thumb are applicable to most deployments. For specific, detailed information and custom rules for deploying a CanIt product, please contact Roaring Penguin directly.

Scan First for Viruses

Architect your scanning solution to eliminate e-mail-borne viruses and malicious attachments before you scan messages for spam. This is the safest approach to mail scanning and creates a more efficient system by minimizing the amount of traffic that the anti-spam solution filters.

Blacklist Repeat Offenders

Sender blacklisting is less useful than whitelisting because spammers tend to change their e-mail addresses very frequently. However, blacklisting selected persistent abusers can be very effective.

Use DNS-Base Realtime Blacklists (RBLs)

Realtime blacklists should be used with caution. Messages that are received from a host that is found on a RBL should be held, not rejected. The advantage of RBL lookups is that they are very cheap – you can hold or block a message without doing any content-scanning. The disadvantage is that you are entirely at the mercy of the organization running the real-time blacklist. Some realtime blacklists (such as www.spamhaus.org) are very responsible and have clear criteria for listing and removing machines from their blacklists. Others are less responsible, and still others are positively aggressive. So, be sure you know and trust the managers of an RBL before you reject mail based on an RBL lookup.

Use Realtime URL Lookups

Realtime URL lookups are similar to RBL lookups in that they use DNS. However rather than looking at the IP address of the connection machine, they look at URLs in the message body. This is useful because, while spammers can easily mutate the bodies of their messages to get around signature schemes, it's a lot harder for them to change websites as quickly – registering and setting up a domain name is not something you can do several times a minute.

Use Sender Policy Framework (SPF) if you Agree With its Philosophy

SPF allows the owner of a domain to declare which relays are designated as originators of mail from his domain. For example, www.roaringpenguin.com has a small list of servers that are designated as outgoing mail servers for the domain, and we publish this list using SPF. The drawback of SPF is that it relies on domain owners to publish these lists; you cannot rely on SPF for a comprehensive list of legitimate mail servers.

Whitelist Trusted Senders, but Never Whitelist your own Domain or E-Mail Address.

This is a best practice because spammers often fake mail from their own victims.

Make your Anti-Spam Server the only Mail Gateway to the Internet

The anti-spam server should be the initial MX host to receive e-mail. This is the best deployment scenario to take advantage of advanced features such as hit-and-run detection and realtime blacklists. Figure 2 below illustrates two recommended network configurations for an anti-spam system.

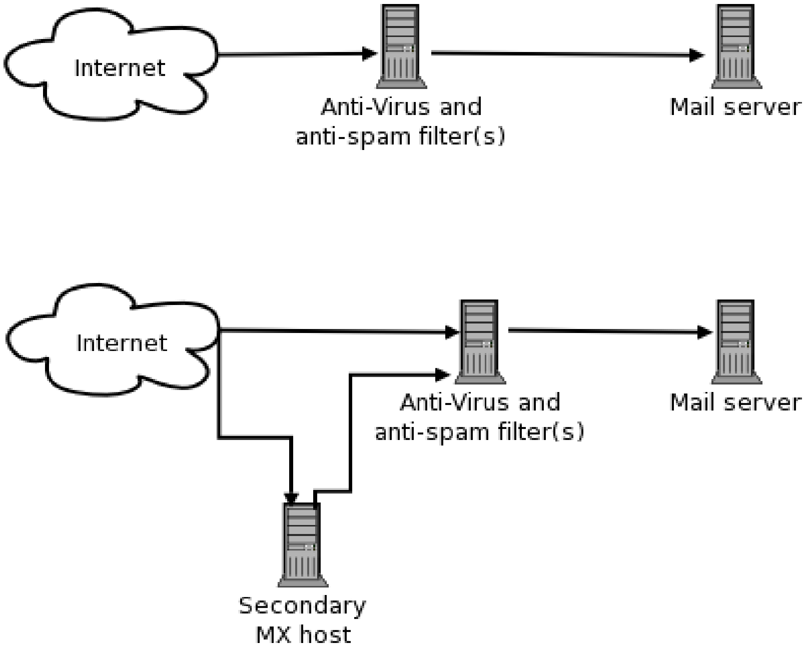
Internet
Anti-Virus and anti-spam filter(s)
Mail server
Internet
Anti-Virus and anti-spam filter(s)
Mail server
Secondary MX host

Figure 2: Two Recommended Network Configurations

Never Publish an Internal Mail Server as a Secondary MX Host

Any unprotected path to your internal mail server is a spam conduit and a security risk. In fact, some spammers deliberately try higher-cost MX servers first, on the assumption that they may be less protected than primary MX servers. You should never have a path from the outside world to your internal mail server that is not protected by the anti-spam solution. Figure 3 below illustrates an example network configuration that should not be used for spam filtering.

Roaring Penguin Software's CanIt-PRO anti-spam solution has been designed specifically to address these challenges for ISPs, large enterprises and educational institutions. CanIt-PRO balances global policies with end-user control, and provides the most flexible and user-friendly solution for campuses that is based on open-source software. Here's how CanIt-PRO works for ISPs.

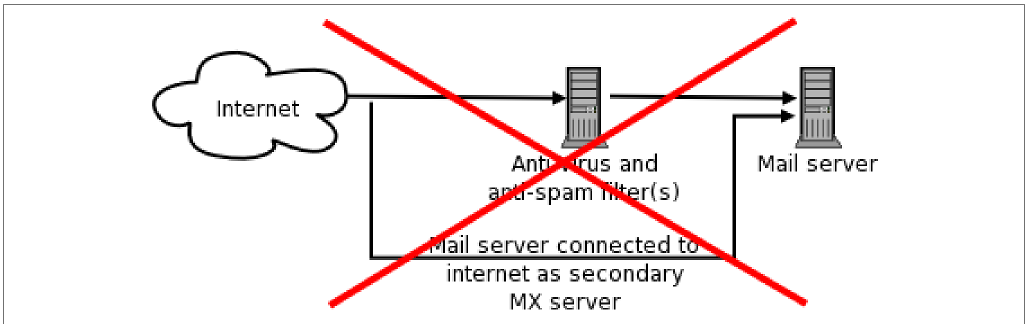

Figure 3: Improper Network Configuration for Anti-Spam

CanIt-PRO Addresses ISP Challenges

Roaring Penguin Software has developed the CanIt-PRO anti-spam solution with the input and ongoing feedback of ISPs ranging in size from 500 users to 100,000 users. As such, CanIt-PRO addresses the challenges listed above.

Specifically, CanIt-PRO:

- **Balances global policy management with end-user control**. Administrators can establish global policies and filter settings, then push these out to end users for the day-to-day management of spam. In this way, ISPs can ensure that all mail is scanned for viruses or particular types of attachments while still enabling users to determine how aggressively they want their mail filtered for spam. ISPs can also choose to give users access to their individual Spam Traps to view and take action on trapped messages, as well as adjusting their filtering preferences. While few users want this much control, more technical users appreciate the option.
- **Integrates seamlessly with existing authentication engines and databases** such as LDAP, POP3, Kerberose, etc. Users can then use one user name and password to access all of their email services. CanIt-PRO's end-user interface is also theme-able for easy integration into an existing web portal system.
- **Integrates with anti-virus software**. CanIt-PRO comes bundled with Clam AntiVirus software and can also be integrated with most commercial anti-virus solutions of an ISP's choice. In fact, many CanIt-PRO customers use both Clam AV and their existing anti-virus solution for double the protection.
- **Provides unparalleled flexibility and extensibility** for investment protection. Since CanIt-PRO is priced per-user, the solution can be scaled across as many servers as an ISP requires, at no extra cost. As an open-source based solution, CanIt-PRO is the most flexible solution of its kind. Administrators can write custom filters using Perl scripts and can integrate the solution into nearly any third-party system on the network.
- **Offers a daily-updated Bayesian database** to catch new types of

spam. This database contains tokens from hundreds of thousands of messages classified by CanIt-PRO customers, and enables your CanIt-PRO installation to adapt automatically to new types of spam.

- **Offers outgoing filtering options** so that ISPs can protect their domain integrity. ISPs are often targeted by spammers that rely on zombie machines to distribute their work, and outgoing filtering allows an ISP to implement checks to detect likely spam activity within its own domains.
- **Is an elegant and fast architecture with low performance overhead**. In fact, CanIt-PRO has been found to require far less hardware than alternative server-based solutions. This is partly due to the fact that CanIt-PRO employs numerous spam filtering techniques that eliminate spam messages even before they reach the mail server.

  An example of such a technique is greylisting, also know as “hit-and-run detection”. Greylisting causes a temporary delivery failure for mail from an unknown sender. This technique identifies spam without allowing it onto the server, because, unlike legitimate mail servers, most spamware does not bother to retry if mail was not successfully delivered. CanIt-PRO also holds suspected spam on the sender's server until it is accepted. These and other techniques keep CanIt-PRO's overhead low.
- **Enables manageable roll-out of the solution**. Since CanIt-PRO operates by separating mail into filtering streams, it can be deployed user-by-user or domain-by-domain. This helps minimize the help desk burden of a new service offering, and supports incremental testing. For example, CanIt-PRO may be rolled out first to users who need it most, or to internal users only, or to a specific domain.
- **Is cost-effective** because it is based on open-source technologies and requires no client software licenses. In addition, all of CanIt-PRO's technical prerequisites are freely available. CanIt-PRO's inherent efficiency also means that you'll likely need comparatively less hardware to run the solution. Roaring Penguin also offers special discounts to ISPs.
- **Comes with exceptional support**. Roaring Penguin's outstanding customer support is well known to the MIMEDefang user community, and continues with all CanIt, CanIt-PRO and CanIt Appliance customers.

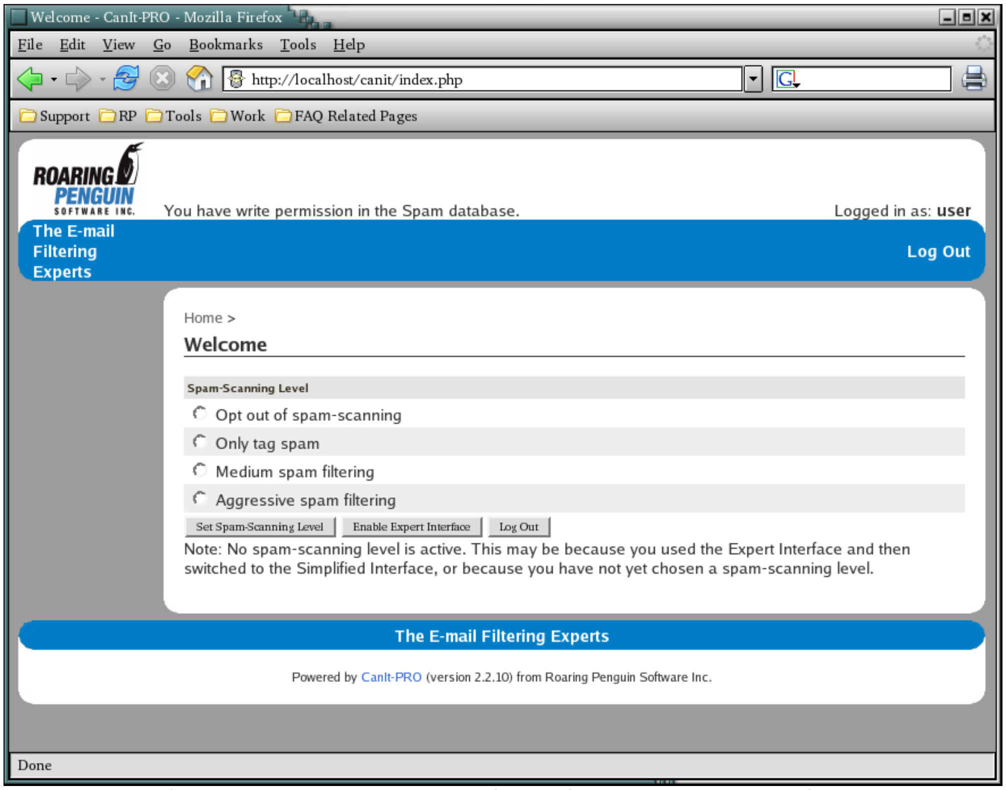
Welcome - CanIt-PRO - Mozilla Firefox
File Edit View Go Bookmarks Tools Help
http://localhost/canit/index.php
Support RP Tools Work FAQ Related Pages
ROARING PENGUIN SOFTWARE INC.
The E-mail Filtering Experts
Log Out
Home >
Set Spam-Scanning Level
Enable Expert Interface
Log Out
The E-mail Filtering Experts
Done

You have write permission in the Spam database.

Logged in as: **user**

Welcome

Spam-Scanning Level

○ Opt out of spam-scanning
○ Only tag spam
○ Medium spam filtering
○ Aggressive spam filtering

Note: No spam-scanning level is active. This may be because you used the Expert Interface and then switched to the Simplified Interface, or because you have not yet chosen a spam-scanning level.

Powered by CanIt-PRO (version 2.2.10) from Roaring Penguin Software Inc.

Figure 4: CanIt-PRO Sample End-User Opt-In Interface

Conclusion

This white paper has reviewed the main challenges, solutions and best practices for ISPs seeking or deploying an anti-spam solution. Roaring Penguin's CanIt-PRO anti-spam solution has been developed in conjunction with ISPs, and as such addresses ISP e-mail challenges with the ideal combination of central and distributed end-user control.

Please contact Roaring Penguin for a free CanIt-PRO evaluation, and ask about our ISP Discount Program.

Contact Information:

Roaring Penguin Software
17 Grenfell Cres., Suite 209C
Ottawa, Ontario
Canada K2G 0G3

E-mail: info@roaringpenguin.com
Tel: +1 (613) 231-6599
Toll-Free: 1-800-210-0984
Web: www.roaringpenguin.com