Combining IMF with Other Anti-Spam Solutions

February 28, 2006

Many organizations are today adopting the Intelligent Message Filter to block spam at server level. Some rely on IMF exclusively; others are combining it with their current filters to further harden spam filtering.

Since Exchange SP2, IMF is waiting for just one mouse click to kick in. Being so readily available, Administrators can hardly avoid facing some of these questions:

Why shouldn’t IMF be enabled?

What are the benefits of having multiple products doing the same job?

Can IMF coexist with other anti-spam solutions?

How complex is it to manage the different filtering products?

Will the end-user experience be consistent?

Looking at IMF as a second anti-spam filtering layer is indeed an interesting scenario. In this case we are looking at organizations that already invested in an anti-spam solution and want to get the most effective spam filtering setup. Here the fact that IMF is free is certainly of secondary importance.

Picturing IMF

Let’s start by putting the Intelligent Message Filtering technology into perspective and mention some basic facts. This filter is developed by a highly trusted vendor. Since you are reading this, it is probably the vendor most trusted by your organization. This vendor developed most of the software running on your network, including the most business critical applications. Right, you guessed, it is Microsoft.

Secondly this filter is now maturing. The second IMF version improved its email analysis technology. Sender ID verification can be combined with the message rating process. Furthermore IMF is now receiving updates twice a month keeping it in-sync with the latest spamming trends.

To round it off, IMF is easy to configure, readily available on each Exchange 2003 SP2 machine and free.

I intentionally left “free” for last. Too often this tends to overshadow the more important facts about IMF. Being free certainly simplifies the process to adopt

IMF. Nevertheless here we are looking at organizations that already invested in other anti-spam solutions. These are organizations where the filtering result is the critical factor and not pricing.

Layered Filtering

Layering involves combining a number of filtering technologies. Each processes messages trying to separate spam from legitimate emails. Individual layers may immediately filter spam through rejection, deletion, etc. Otherwise a filter may contribute to a shared email rating system. Thus the classification process is distributed over a number of filtering layers.

Layering may at first sound complex. In reality it is the standard way anti-spam filtering is performed. I cannot think of a single anti-spam package that is not composed of a collection of layered filtering technologies. As an example you can look at the various anti-spam filtering layers provided by Exchange 2003 and Outlook 2003 on their own.

So should we simply chain as many filters as possible? Throwing in filters blindly won't give the best results. Effective layering should be based on filtering technologies. A good technology mix should cover all the information within the email delivery process and the email content itself. For a discussion of various filtering technologies from a layering perspective check my article Hardening Anti-Spam Protection.

A chain of filters that covers the broadest spectrum could be one comprising SMTP protocol command filtering, verification of sender reputation, signature based filtering and a self-learning filter. Having filters based on the same technologies is certainly less effective. The filters would in that case end up analyzing the same information potentially missing other valuable data.

Answering the Questions

We are now ready to start addressing some of the introductory questions. Our scenario is quite broad. The answers do depend on the specifically deployed anti-spam filters. Nevertheless we can identify a common approach on how to tackle these questions.

Why shouldn't IMF be enabled?

What are the risks of enabling ANY anti-spam filter? False email classification. The fact that IMF is in use at various organizations and the fact that it comes from a trusted vendor should help us build some confidence. Nevertheless being cautions does not hurt.

I already discussed configuring IMF in the past. IMF SCL Configuration - Getting it Right discusses configuring IMF v1 in a conservative manner. The same approach can be directly applied to IMF v2. Just keep in mind the difference in enablement and deployment between the two IMF versions.

Organizations combining IMF to other filters can take further advantage of their position. Most filters are able to express their email classification in terms of a confidence rating. We could break down this type of classifications into three:

1. Email is most likely Legitimate
2. Email classification is uncertain.
3. Email is most likely Spam

The idea is to operate the filters within the range where email classification is most accurate. Ideally we should only allow filters to block emails when the likeliness of an email being spam is very high. As soon as we fall in the uncertainty range, email classification is considered inconclusive. The decision should then be left to the next filter. Filters built on different technologies look at the same message from a different angle. This gives us a fresh opportunity to classify the email. Hopefully this time the message is classified with a high degree of certainty.

This approach indeed combines the strengths of individual filters in order to minimize false classification. Applying the concept to IMF, one could operate the filter at the higher SCL threshold levels. Once the system is running we can start lowering the thresholds a little to fine tune the system. Having multiple filters brings extra flexibility. If one filter is giving better results than another, then we can lower its filtering benchmark, whilst retaining the thresholds for the other.

What are the benefits of having multiple products doing the same job?

Having filters doing the "same" job is useful if these base their classification process on different technologies and/or information. A filter analyzing the SMTP Protocol command data is very different from a filter analyzing the email body content. Here both the filtering technology and the information analyzed are different.

Filters applying different technologies to the same information may also be complementary and appropriate for layering. Consider a signature based filter and a self-learning filter. It is true that spammers try all kind of tricks to give their emails a legitimate look. Nevertheless a large proportion of spam exhibits common patterns where signature based filters are very effective. Self-learning filters employ a very different process. This enables them to be more adaptable to the individual characteristics of organizations.

IMF is based on the SmartScreen technology. This is proprietary and little information is available on its internals. The spam delivered to hotmail mailboxes is used to construct its filtering intelligence. Thus we can say that IMF primarily operates on the email content information. This makes it an excellent candidate for layering in combination with filters operating on SMTP protocol command data.

Layering IMF with other content based filters introduces some overlap. Since there is little information on the IMF internals, classifying the extent in technology overlap is not possible. Thus here the effectiveness of the filtering combination is best determined through testing.

Can IMF coexist with other anti-spam solutions?

The short answer here should be yes. IMF does not break other filters and I am unaware of other filters breaking IMF. If you run into one such filter, you should certainly question its quality.

Of course basic coexistence is of little value in itself. As already discussed we need a set of complimentary filters. Filters should be complimentary in terms of technology, but ideally also in terms of administration and end-user experience.

How complex is it to manage the different filtering products? Will the end-user experience be consistent?

The basic IMF configuration is certainly not complex. Indeed its lack of configurability is often considered to be one of its biggest limitations. Nevertheless products like IMF Tune overcame this aspect transforming IMF into a feature rich product.

Here we should look beyond the basic IMF configuration. Other more important manageability factors include reporting, archiving and Junk Email repositories. Lack of consistency in these areas can be a true hurdle to adopting multiple filtering layers from different vendors.

Exchange 2003 laid the foundation for consistent spam filtering. It standardized the Junk E-mail functionality and provided client side Safe Senders/Recipient lists. This takes us a long way in achieving a consistent end-user experience. With a standard Junk Email repository users may remain happily unaware of the day-to-day battle taking place at the servers.

Anti-spam filters claiming to be Exchange integrated should today be supporting the Junk Email folder. As for other filters lacking Exchange integration, hope is not lost. Indeed with a little help these may also be transformed into first class Exchange citizens. This subject is discussed in IMF Tune Opens Exchange to Any Anti-Spam Filter.

This covers the end-user experience fairly well. As for the administrative side, there is certainly less consistency. The IMF configuration is available from the Exchange System Manager and the performance monitor provides its monitoring interface. Other products have their own interfaces and reporting tools.

Indeed some more consistency in this area would be welcomed. But achieving consistency from competing parties is clearly not that simple. You will probably have to accept a little more administrative overhead. Nevertheless I am sure the results obtained from the hardened spam filtering setup will completely out weigh this aspect.

Final Tips

The Intelligent Message Filter offers an excellent opportunity to harden spam filtering. Many organizations understood this, have greatly improved spam blocking levels and minimized false detection. The case studies at WinDeveloper confirm this fact.

Layering is most effective when the right technology mix is in place. Deploying the Intelligent Message Filter only requires a few minutes. Thus testing the new environment requires little effort.

Exchange integration is the key to effectively combine different anti-spam solutions. The Exchange 2003 Junk Email folder ensures a consistent end-user experience.

Filters lacking Exchange integration may also join this multi-layered setup. IMF Tune transforms anti-spam solutions running on any platforms, firewall appliances, or external service providers into first class Exchange citizens. Emails identified as spam by these filters are routed to the Junk Email folder just like any other natively integrated solution.