Joel Snyder
Opus One
February 2009

BEST PRACTICES IN USING REPUTATION-BASED ANTI-SPAM SERVICES FOR EMAIL SECURITY

One of the most efficient, least costly and, remarkably, most effective anti-spam techniques is IP reputation filtering. An incredibly inexpensive technique for the anti-spam gateway, IP reputation filtering can be used to identify 80% (or more) of spam without even looking at message content. IP reputation services have become a best practice for any anti-spam gateway.

This white paper discusses the origins of IP reputation services; test results on Cisco IronPort's own reputation service, SenderBase; best practices in using reputation services; and the ROI of reputation services.

IP reputation services come from the observation that you can often identify a message as spam simply by the IP address it comes from. Starting in 1997, reputation services (originally called "blacklists") have become a staple of anti-spam vendors, with nearly 200 open source and commercial reputation services available.

Opus One's testing of Cisco IronPort's SenderBase reputation service has shown that, when used as recommended, an average of 88% of spam can be identified and blocked without regard to content. Opus One's testing of reputation services has shown that, when used as recommended, enterprise-class reputation services can identify and block an average of 88% of spam without regard to message content. The results of our testing also show that IP reputation services have a much lower false positive rate than content filters.

TABLE OF CONTENTS

Executive Summary 1
Introduction 2
What Are IP Reputation Services and Where Did They Come From? 2
IP Reputation Service Performance 4
Best Practices in Using IP Reputation Services 7
- SMTP-time use of IP Reputation Services 7
- Content-filter Use of Reputation Data 8
- Recommendations for SenderBase Deployment 9

ROI for IP Reputation Services 10
Conclusions 12
Vendor Case Study 13

The most efficient use of IP reputation services is during message receipt, sometimes called "SMTP-time." When used properly, IP reputation services can dramatically reduce the amount of spam entering an enterprise, increase spam catch rates, and save considerable money and resources. Using enterprise-class reputation services a typical enterprise could see a 73% reduction in total email flow, saving not only capital expenses, but operational expenses (power, rack space, cooling), administrative expenses (fewer gateways require less time to manage), and overhead such as reporting and database server costs.

Introduction

The most effective anti-spam products are based on an optimized "cocktail" approach to spam detection. An anti-spam product using a cocktail mixes multiple tests and techniques together, both to increase the spam catch rate and to decrease the false positive rate. The idea behind using multiple approaches was best described by a researcher who paraphrased P.T. Barnum: "you can fool some of the tests all of the time, and you can fool all of the tests some of the time, but you can't fool all of the tests all of the time." In other words, mixing techniques together reduces the possibility that a clever spam sender will be able to bypass protections. Opus One's side-by-side testing of anti-spam products has demonstrated the power of the cocktail approach. We have found that products using multiple complementary and even overlapping techniques fare better when confronted with the constantly shifting landscape of spam.

The cocktail approach has another advantage, which is that it can reduce the load on anti-spam gateways. Not every anti-spam test takes the same amount of CPU time or memory. If early tests with lower cost can determine that a message is spam or is not spam, then the message can be released or rejected that much more quickly. This increases throughput through the gateway, reduces latency, and, ultimately, reduces costs.

This white paper discusses the origins of IP reputation services; test results on both commercial and community-supported reputation services; best practices in using reputation services; and the ROI of reputation services.

What Are IP Reputation Services and Where Did They Come From?

IP reputation services come from the observation that you can often identify a message as spam simply from its IP address.

The first IP reputation service, in the form of an IP-based block list, was started in January 1997 by Paul Vixie, already known at that time as the primary maintainer of BIND, the de-facto standard for DNS server software. Vixie envisioned a way of sharing information that he was gathering about IP addresses that were used by known spammers. At that time, anti-spam products had not advanced enough to use this information directly, and Vixie's initial strategy was to distribute the list using the BGP routing protocol. This was an efficient approach that was compatible with the routing infrastructure at ISPs, who could use the list to simply block all traffic (email and otherwise) from the known sources of spam. The first users of the "Realtime Blackhole List" came after a NANOG meeting in San Francisco in February, 1997, where Vixie widely publicized his thoughts in a talk "Controlling Network Abuse Using BGP-4." Amusingly enough, the first retaliation for his anti-spam work came soon after, on February 17th, 1997, when a spammer forged mail

in the name of Vixie to members of the US Senate and House of Representatives in an attempt to discredit him.

In September, 1998, the BGP-based RBL had spawned an organization, MAPS (the Mail Abuse Protection Service). At MAPS, Eric Ziegast and Dave Rand worked with Vixie to make the RBL available using DNS, and usable in Sendmail, the dominant SMTP MTA at that time. While DNS wasn't as efficient a mechanism for distributing the information as BGP, the simplicity of implementation and ability to target just email traffic made the DNS-based RBL very popular. An estimated 3,000 mail servers were participating in the DNS-based RBL by February, 1998; only 122 networks were using the BGP feed at that time.

MAPS was spun out of Vixie's consulting company as a non-profit and led by Dave Rand, originally of the ISP Abovenet and the first subscriber to the RBL. In 2004, MAPS was turned into Kelkea, which was eventually purchased by Trend Micro in June, 2005.

A major innovation in reputation services came in early 2003, when IronPort introduced its SenderBase reputation service that allowed for the possibility of both positive and negative information about any individual IP address. This type of reputation service gives email managers a better tool to help manage the flow of spam. Rather than only being able to say "this IP address sends mostly spam," now it was possible to also say "this IP address has a long history of **not** sending spam." SenderBase's combination positive/negative reputation system was copied by other anti-spam vendors within 18 months, including Ciphertrust (now McAfee) Trusted Source and Brightmail (now Symantec).

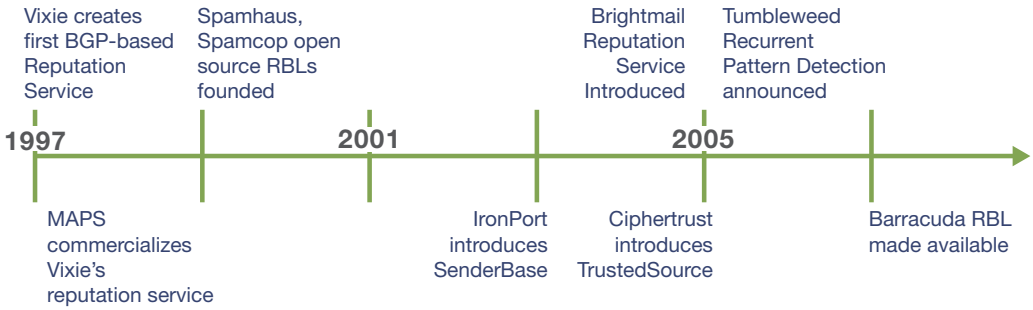

Since their 1997 start, IP reputation services have established themselves as major weapons in the war on spam. Originally, the idea of reputation services was to identify the IP address space used by spammers themselves. To avoid being identified, spammers began to use "open relays," email servers that would accept messages from anyone and take responsibility for final delivery. This evasion technique reduced the effectiveness of IP reputation services, because now some IP addresses were sending both legitimate and spam email. A back-and-forth war has emerged between spammers and the operators of reputation services. In some cases, this has resulted in wholesale change in Internet behavior. For example, no legitimate mail server operator intentionally operates as an open relay anymore.

This set of border skirmishes between spammers and mail managers has also spawned many different reputation services, each with different goals and policies. For example, Spamhaus maintains the Policy Block List (PBL), populated based on input from ISPs who identify IP addresses that should not be operating mail servers, typically broadband residential IP addresses. IP addresses are on the PBL not because they've sent spam in the past, but because their ISP has said that they shouldn't be directly sending mail to the internet at all.

SenderBase reputation service offered by IronPort Systems, a Cisco business unit, offers another innovation: a ranking system, with a score ranging from -10 to +10 assigned to every IP address in the database. Having a score, rather than a "yes/no" answer, gives the email manager much greater flexibility in using the IP reputation service to control mail flows and traffic. The benefit of having continuous scores, as compared to separate discrete lists of "certain" and "likely" spammers, is that the decision about how much to trust the reputation service is in the hands of the email manager, rather than being left to the operator of the reputation service. This provides more control to the email manager and thus more confidence in the overall security of the system.

In SenderBase, highly positive scores represent systems that are extremely unlikely to be sending spam, while highly negative scores are assigned to systems that are almost certainly sending spam. For example, an email manager might choose to block all mail from systems with scores from -10 to -4.0, throttle mail from systems with scores from -3.9 to -1, and even bypass spam scanning for systems scores above +5.0. Because SenderBase has continuous scores from -10 to +10, the email manager can adjust their blocking, throttling, and bypass thresholds to match the characteristics of their organization's message flows, their own tolerance for false positives, and their desire to increase performance.

IP Reputation Service Performance

Users of IP reputation services must ask two key questions about any reputation service. First, how well does it block spam? Second, what is the false positive rate? We ask the first question to understand whether or not there's any benefit to using a reputation service. If the service doesn't block much spam, then there's no point in using it. We ask the second question to understand the negative impact of a reputation service. If it has a high false positive rate, then the resulting help desk calls and service interruption may outweigh the benefits.

Measuring the spam block rate of a reputation service can be fairly difficult. Many anti-spam products report their spam block rate by counting refused connections and assuming some multiplier based on internal research and testing for the number of messages that would have been sent over that connection. These help the email manager to understand how well reputation-based blocking is working from day to day, but only approximate actual performance.

To provide greater insight into the true behavior of reputation services, Opus One has designed a testing methodology to fully measure the ability of each reputation service to block incoming spam. By manually evaluating each message received in a real corporate email stream, and comparing that message to its reputation service result, a more reliable and comprehensive picture can be constructed. (A fuller description of Opus One's testing methodology is available at http://www.opus1.com/www/whitepapers/spamtestmethodology.pdf)

The graph below highlights SenderBase performance from November 2007 to November 2008. Each of these snapshots was created by looking at 10,000 incoming messages over a period of 7 to 14 days each month, selected from a real corporate mail stream. No artificial mailing list, spam trap or synthetic traffic was part of this stream.

Because SenderBase has a continuous score for each IP from -10.0 to +10.0, we selected four separate cutoff numbers to show how SenderBase would block incoming spam, completely independently of any other content based filtering. In other words, if you operated a spam filter which did nothing but look at the IP reputation service, here is how well it would block spam. All enterprise spam filters combine multiple tests—the "cocktail" mentioned in the introduction to this white paper—but this graph highlights the performance of just one part of the cocktail, the IP reputation service.

An email manager with minimum tolerance for false positives might select a range of -10.0 to -4.0 for spam blocking, as shown in the blue area on the graph. An email manager with more aggressive goals in blocking spam could select a range of -10.0 to -1.0, as shown in the green area. IronPort's out-of-the-box settings recommend blocking senders with a reputation score in the range -10.0 to -3.0, the green area.

Percent of Spam Blocked at Selected Level

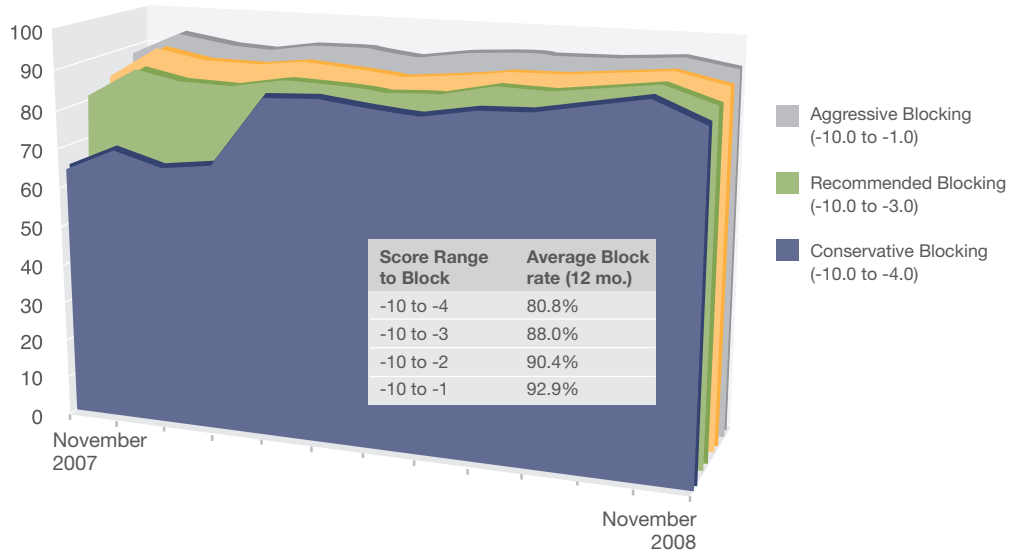

Score Range to Block	Average Block rate (12 mo.)
-10 to -4	80.8%
-10 to -3	88.0%
-10 to -2	90.4%
-10 to -1	92.9%

The results show that IP reputation services can identify an enormous percentage of incoming spam without even looking at the message content. This has important implications for performance, because the IP address of the sender is known before the message is accepted, which means that spam can be blocked before it is even received. Using a recommended threshold of -10.0 to -3.0 for blocking spam, an average of 88% of spam can be identified, and blocked, without regard to content.

This brings us to the second key question: what is the false positive rate? When an IP reputation service is used properly, false positives have a minimum impact (compared to typical content filter false positives). Regardless, a high false positive rate is obviously undesirable.

The results of our testing show that IP reputation services have a much lower false positive rate than content filters. As a typical example, in May, 2008, Opus One tested 46 different anti-spam engine scenarios. In that test, the average content-filter false positive rate was 0.21% (median 0.17%, standard deviation 0.2%). Over the past 12 months, the average false positive rate for SenderBase when blocking in a range of -10.0 to -3.0 was 0.018%--about 1/10th the rate of content filters. Even with a very aggressive blocking range of -10.0 to -1.0, the false positive rate for SenderBase is 0.12%, still less than the rate for content filters.[1]

The graph below, scaled to show 1% as a maximum, shows the false positive rate for SenderBase at various blocking levels. The graph has to be scaled up to this level in order to discern differences in the blocking strategies.

False Positive Rate

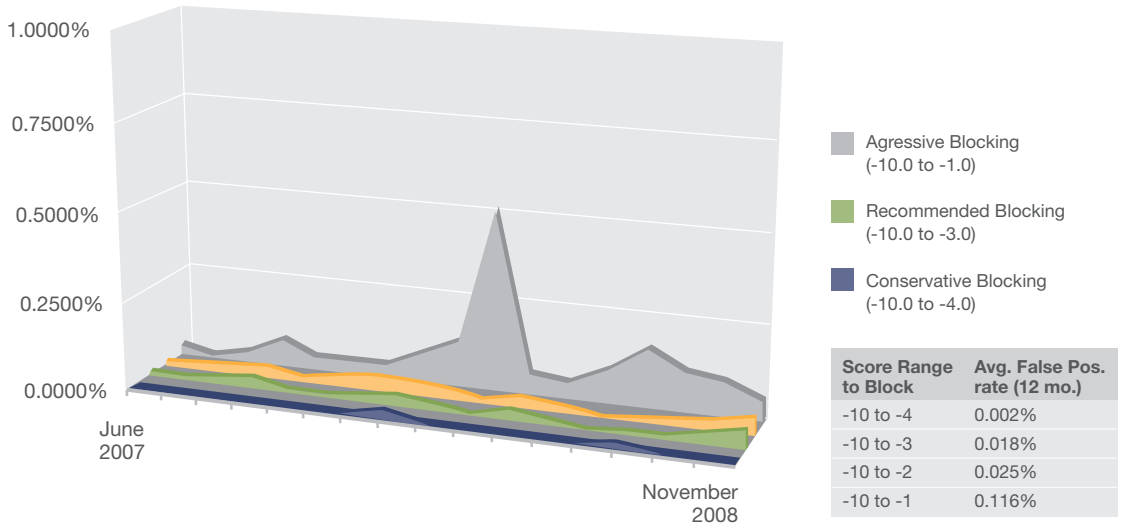

Score Range to Block	Avg. False Pos. rate (12 mo.)
-10 to -4	0.002%
-10 to -3	0.018%
-10 to -2	0.025%
-10 to -1	0.116%

1 - In Opus One's testing, we use the Positive Predictive Value (PPV), a measure of the accuracy of a test, to compute the false positive rate for content filters. We have used the same statistic for reputation filters for consistency with earlier reports. Another common measure used by researchers is specificity. In either case, the comparison between content filter and reputation service accuracy is similar.

Best Practices in Using IP Reputation Services

IP reputation services can be used in many ways as part of the cocktail, but most anti-spam products have settled on one of two main techniques: either IP reputation services are used at message receipt time (at the SMTP connection time), or the reputation data is simply one more input into a single "spam/not-spam" decision made by the content filter's cocktail of tests.

SMTP-TIME USE OF IP REPUTATION SERVICES

The most efficient use of IP reputation services is during message receipt, sometimes called "SMTP-time." When reputation services are used this way, the reputation of a sending system is checked when the sender attempts to connect to the anti-spam gateway. Based on the reputation, the receiving anti-spam gateway could decide not to accept the TCP/IP connection at all, or it could accept the connection, but refuse to accept the message. Some products don't offer a choice of these two strategies, but for those that do the choice should be made based on logging requirements and feedback strategies. If the anti-spam gateway doesn't accept the connection at all, then very little debugging information is available to the email manager—only the IP address of the system trying to connect. If the connection is accepted, but the message refused, then the email manager can at least see the RFC2821 envelope information: who the message is (purportedly) from, and who it was sent to.

ATTRIBUTES OF A PREMIUM REPUTATION SERVICE

- Includes multiple diverse sources of data as part of scoring process to reduce errors and false positives
- Provides reputation scores beyond simply "bad"/"not-bad" to give email manager flexibility in deployment and determining risk tolerance
- Links to web service making it easy to see why an IP address scores the way it does
- Operates using globally distributed service for maximum uptime
- Builds on open policies and procedures to provide maximum transparency and give email managers maximum confidence

While, at first glance, it may seem like the best approach is to simply refuse the connection, since that uses the least resources, a more effective technique is to accept the connection, but refuse the email message. (This is often referred to as a "5xx" SMTP response, the error code that actively refuses a message.) The main reason is that this technique provides strong feedback to the sender of the message that their message will not be accepted. If the TCP/IP connection is not accepted, the sender has no idea why, or whether this is a permanent or transient error. Thus, the sender has no choice but to keep retrying the connection in the hopes that their message will get through. If the TCP/IP connection is accepted, but the message is explicitly refused, then the sender knows immediately what happened.

The reason we recommend refusing the message, rather than the connection, is to eliminate undetected false positives. In the case of a false positive, the actual sender of a message will get an immediate—perhaps within seconds—notification that their message was not accepted. This gives legitimate senders an opportunity to communicate using some other channel, such as via telephone or a different email service. By refusing the message, the impact of false positives on business-critical communications is minimized.

Compare this to choosing to refuse the TCP/IP connections of senders with "bad" IP reputations. In that case, the sender of a message will only receive notification of the problem after their own email MTA has made multiple attempts to deliver, usually after several hours or even several days. This means that detection of false positives is delayed.

Of course, either of these techniques are preferable to the false positives that come from content filters. In most anti-spam gateways, spam detected by a content filter simply disappears into a black hole, which means that neither the receiver nor the sender have any information that an important message was lost. Systems that provide quarantines ameliorate this problem somewhat. It is clear, though, that the impact on business of a false positive from a properly configured IP reputation service, is much lower and presents much less risk than improperly configured IP reputation services or most content filters.

An additional benefit comes from using reputation services, which can provide a mixed reputation for a particular IP address. The mixed reputation is neutral since it contains a balance of both positive and negative information about the IP. Examples of this include Cisco's SenderBase reputation service, Commtouch IP reputation service and the Trend Micro Email Reputation Service. When a mixed reputation is available, email managers should choose to either throttle connections or respond with "temporary failure" codes to senders that have bad—but not too bad—reputations.

For example, throttling can be used to lower the likelihood of a business impact on an improperly classified sender. If a non-spamming sender has a "slightly bad" reputation and is then throttled, for example, to 10 messages per day, then a single user would be unlikely to be affected by the misclassification. On the other hand, if the sender is actually generating large volumes of spam, then the throttle would reduce the amount of spam accepted, and thus reduce both load on the email systems and the likelihood of spam getting through.

Another strategy, similar to throttling,is to return a "temporary failure" (sometimes called a 4xx SMTP code) in the presence of negative reputation data. This can be used with throttling, or in place of throttling. This is appropriate when reputations change quickly up and down and there is a reasonable expectation that someone who has a "slightly bad" reputation will either be given a clean bill of health or move to the "always bad" category quickly.

CONTENT-FILTER USE OF REPUTATION DATA

Reputation data can also be used after a message is accepted and during content filter analysis of the message. This might be done for any number of reasons. The most common is that the gateway that actually received the message does not have the ability to properly use IP reputation data. Typically, such a gateway will simply receive

mail and then pass it on to the anti-spam gateway for inspection and content filtering.[2] In that case, the content filtering gateway has no choice but to accept the message because the reputation of the sender is not available during the SMTP conversation.

Content filtering decisions use IP reputation services as one factor in the decision on whether to mark a message as spam or not. The IP reputation is carried with the message and when the content filter starts looking at the message, it is one component in the total score or verdict.

Content filtering is significantly more expensive from a resource point of view than simply refusing to accept a message, so this approach should be avoided if at all possible.

In the case of IP reputation services like SenderBase, which have a broad range of rankings possible, reputation data can be used during content filtering when the reputation data alone is not sufficient reason to reject a message—or when reputation data may help to identify trusted senders and bypass anti-spam scanning. For example, SenderBase scores in the range of -0.9 to 0.0 are highly correlated with spam senders, but also have a high false positive rate. Rather than simply rejecting messages from a sender with a -0.1 score, it might be better to run additional tests on the message to see if other findings correlate with the slightly negative score.

RECOMMENDATIONS FOR SENDERBASE DEPLOYMENT

Cisco has generally recommended that its customers block incoming messages with SenderBase scores in the range of -10.0 to -3.0.

Our research confirms this recommendation as the “sweet spot” for SenderBase that balances out a high spam block rate with a low false positive rate. The table below combines 12-month averages for block rate and false positive rate of SenderBase.

SenderBase Blocking Threshold	Incoming Spam Blocked (percentage)	False Positive Rate (percentage)
-10.0 to -4.0	80.8%	0.0023%
-10.0 to -3.0	88.0%	0.0185%
-10.0 to -2.0	90.4%	0.0246%
-10.0 to -1.0	92.9%	0.1162%

With a spam detection setting of -10.0 to -3.0, SenderBase-enabled systems will block approximately 88% of incoming spam, and have a false positive rate of 0.0185%. Blocking from -10.0 to -2.0 will add only 2% to the block rate, but will increase the false positive rate by 33%. Thus, -10.0 to -3.0 seems to be an appropriate blocking rate for a normal organization with average risk tolerances and a desire to reduce false positive rates.

2 This can also occur when the content filtering gateway is co-resident with the receiving SMTP MTA, but the SMTP receiver is not properly integrated with the content filtering anti-spam software. This is common in open source deployments, especially those using multiple reputation services to achieve a “consensus” reputation for any sender.

While it is possible to have a more aggressive approach to blocking spam based on reputation services, the tradeoff between a blocking level of (for example) -10.0 to -3.0 and a blocking level of -10.0 to -1.0 is fairly expensive. For an additional 5% of spam blocked at the incoming gateway, the aggressive email manager must tolerate a six-fold increase in the false positive rate. Our advice is that scores in the range of -2.9 to 0.0 may be more appropriately handled via throttling or at the content filtering and message analysis part of the anti-spam gateway.

Choosing an aggressive scanning level is really only a good idea when message gateways are strained for capacity. For example, if you would receive 1,000,000 messages (spam and non-spam) a day, choosing a block level of -10.0 to -3.0 means that you will have to receive and search for spam in 270,000 messages each day. On the other hand, selecting a block level of -10.0 to -1.0 would reduce that number to 229,000, a savings of about 15%.

ROI for IP Reputation Services

The Return on Investment for IP reputation services comes in several ways. The simplest way to realize value is by understanding that reputation services improve the spam catch rate of message scanning content filters. In Opus One's testing, we find that adding a reputation service can raise the catch rate for a content filter by a significant percentage. For example, in a recent 2008 test of 54 different anti-spam scenarios, the average increase in catch rate we observed was 3.9%—a readily discernible difference in spam volume.

A second value from IP reputation services is their ability to reduce the number and size of email gateways an enterprise needs. By simply refusing to accept spam email, the load on message content filters is lower and quarantines can be smaller. A message that isn't received doesn't have to be logged and archived. And, for enterprises that prefer to use tag-and-deliver or special "spam" folders, loads on message servers (such as Microsoft Exchange or Lotus Notes) are dramatically lower.

Calculating the reduction in anti-spam gateway costs is illustrative of the savings. Consider the example of an organization that has a load of 1,000,000 messages/day, including spam, moving through their anti-spam gateways. In Opus One's testing, the amount of spam varies from month to month, but 2008's average is 83%, meaning that of the 1,000,000 messages, 830,000 will, on average, be spam. Using SenderBase's IP reputation service as an example, an enterprise blocking using the recommended range of -10.0 to -3.0 will see an average block rate of spam of 88% (measured across the past 12 months). In other words, SenderBase—used according to the best practices described in this white paper—would have blocked about 730,000 of those messages from being received at the enterprise.

Performance Advantage of Reputation Services

Using enterprise-class reputation services to refuse spam can reduce total gateway load by 63% to 73%.

	SenderBase Example	Open Source Product Example
Number of messages received a day Percent spam for average enterprise (measured for last 12 months)	1,000,000 83%	1,000,000 83%
Actual spam messages you received (# messages * % spam)	830,000	830,000
Reputation Service Block Rate	SenderBase -10 to -3 blocks 88% of spam	Open Source list blocks 76.9% of spam
Number of Spam Messages Blocked Before They Enter the Enterprise (Actual Spam * Block Rate%)	730,000	631,000
Messages Left that your Anti-Spam Gateway Will Have to Handle (Total Messages – Blocked Spam)	1,000,000 – 730,000 = 270,000	1,000,000 – 631,000 = 369,000

This blocking of 73% of the incoming mail means that the enterprise could get by with 2/3 to 3/4 fewer anti-spam gateways. They don't need to handle 1,000,000 messages a day, but only 270,000 messages a day. Even a conservative approach to this would allow the number of gateways to be reduced by half.

This savings is significant, because reducing gateways saves not only capital expenses, but operational expenses (power, rack space, cooling), administrative expenses (fewer gateways require less time to manage), and overhead such as reporting and database server costs.

Because the percentage of received mail that is spam is so high at typical enterprises, small changes in the amount of spam blocked will have large effects on the performance of the email network as a whole. If spam were only 20% of the mail received, then a few percentage points would only add up to a small amount of difference. But because spam represents more than 80% of the email in an enterprise—and as much as 90% in some enterprises and at some times of the year—a few percentage points difference have a disproportionate effect on the number of messages that must be handled. For example, in the reputation service example above, the two services have effectiveness rates of 76% and 88%. This 12% difference in effectiveness is magnified as the amount of mail to be handled increases from 270,000 (88% effective reputation service) to 369,000 (76% effective reputation service): a 37% increase in load!

Enterprises that deal with those messages using quarantine servers or, even worse, their own enterprise mail servers, will see an even more dramatic ROI. Using the example above, a quarantine server without a reputation service would have to handle 830,000 spam a day; with a reputation service, only 100,000 spam a day: an eight-fold decrease

in capacity requirements. Enterprises that actually tag-and-deliver messages or send them onto a special "spam" mailbox in their main mailbox servers will see even more dramatic differences in costs, because services such as Exchange and Notes are much more expensive per message. Without reputation services, an enterprise would have to triple the capacity of their enterprise messaging system to handle the same amount of Internet-incoming email.

These values were calculated using an average of per-message costs published by various industry analyst firms.

If your enterprise receives 1,000,000 messages a day, and …	Using reputation services could save you …
… quarantines spam using a dedicated spam quarantine server	$25,550.00 each year
… archives spam using an enterprise email archiving product	$153,300.00 each year
… delivers spam to a per-user "spam" folder in Exchange or Notes	$511,000.00 each year

Conclusions

IP reputation services are a valuable tool for any anti-spam gateway. With a low rate of false positives, IP reputation services can dramatically reduce the amount of spam entering an enterprise, increase spam catch rates, and save considerable money and resources. Our testing has proven that an industry-leading reputation service such as Cisco's SenderBase can block up to 88% of the spam before it hits the network, providing significant future proofing as the volume of email—and spam—continues to increase. When IP reputation services are used during "SMTP time," the amount of spam entering an enterprise can be dramatically reduced, while spam catch rates increase saving considerable money and resources.

Vendor Case Study: Cisco and SenderBase Reputation Service

Effective use of reputation services can go beyond simply refusing email at SMTP time. As a case study, we looked at Cisco IronPort's mail gateways to investigate how reputation services could be leveraged to provide other services within the gateways.

Since SenderBase is a scaled reputation service, rather than a go/no-go service, IronPort mail gateways can both block and throttle mail based on the reputation service. For example, scores of -3.0 might be blocked entirely, while scores up to -2.0 might be throttled to 10 messages per hour, and so on. These settings are available to the network manager who can select as granular and detailed a throttling policy as needed.

Recipient throttling also helps reduce false positives in the slightly negative, "suspect" range by allowing the valid senders in the "suspect" SBRS range to continue sending while slowing down the egregious senders in the same "suspect" SBRS range. Given the recent prevalence of reputation hijacking, due to botnet activity and compromised webmail accounts, this can provide significant incremental savings.

SenderBase reputation services is also used by IronPort as part of its directory harvest protection capabilities. A common feature on high-end mail gateways, directory harvest protection is used to keep malicious senders from "walking through" all possible email addresses at an organization in search of valid ones. Most directory harvest protection is based on invalid recipient counts: once a sender goes past a pre-determined threshold of invalid recipients, the mail gateway assumes they are trying to harvest the directory and takes a proactive action, such as closing the connection. IronPort ties their directory harvest protection (called DHAP, Directory Harvest Attack Protection, by IronPort) to reputation. In particular, users can set DHAP thresholds in the "suspect" range lower than that for positive reputation ranges. As organizations are more likely to be victims of DHAs from IP's with suspect reputation, this feature provides higher protection. For senders with good reputations, such as business partners, this threshold can be much higher. This enables these business partners who may not have the most updated email distribution lists to continue sending important business communications (industry newsletters, benefits information, etc.) without the risk of getting blocked by a DHAP policy. This is a representative example of both the security integration and holistic approach of SenderBase into Cisco IronPort's Email Security service.reputation services could be leveraged to provide other services within the gateways.