

Introduction

More has changed in the past five years of computer virus protection, than in the previous 20 years. Statistics from the global network of Network Box Security Response centers show:

- A dramatic increase in the number of distinct threats seen (from less than 40,000, five years ago, to more than 250,000 today). A few years ago, one new variant of a particular virus would be seen, perhaps, once a month. Now, computer virus researchers are often seeing half a dozen, or more, new variants in a 24-hour period.
- A dramatic decrease in the time taken for a threat to reach peak infection levels. A common measure is the time taken for the infection rate (detections/blocks per minute) to reach half its peak level. Now, variants of the Mydoom, NetSky and Bagle worms (for example), can reach this point within less than 1 hour.
- New techniques for bypassing protection devices have appeared. These include password-protected archives, vulnerability exploitation, social engineering, message fragmentation, and standards bypass.
- The use of spamming techniques for the initial propagation stage of virus seeding (greatly increasing the initial rate of spread).

In the early 1980s, the vast majority of computer viruses were classic file and disk infectors. Now, in 2008, over 99.98% of all viruses are email or network carried worms.

In the early 1980s, it would take weeks for a new virus to propagate around the world. Now, computer worms often take less than 2.5 hours to reach their peak infection rate.

In the early 1980s, there were less than 3,000 viruses known to exist. Now, over 3,000 new virus variants are seen every month.

Statistics for spam are even more impressive – a doubling of the amount of spam sent, each year for the past five years. That has managed to outpace even Moore’s law.

While heuristic technology has improved, protection for the core threats is still primarily signature-based. This is particularly true for anti-virus, anti-spam and intrusion detection systems. If the protection signature is not in place, installed and working, the protection is, most likely, not effective.

CONTENT

Introduction..... 1

Delivery of Protection Signatures and Code Updates.....2

Why PUSH?.....3

Problems with PUSH Technology...5

Conclusion.....6

MARCH 2008

No part of this publication including text, examples, or illustrations may be reproduced, transmitted, or translated in any form or by any means, electronic, mechanical, manual, optical or otherwise, for any purpose, without prior written permission of Network Box Corporation Limited.

Network Box Corporation Limited,
16th Floor, Metro Loft,
38 Kwai Hei Street, Kwai Chung,
Kowloon, Hong Kong
Telephone: +852 2736-2083
Fax: +852 2736-2778
www.network-box.com

Delivery of Protection Signatures and Code Updates

There are two alternatives for delivery of protection signatures and code updates:

- **PULL**

In which the protection device periodically contacts the provider to enquire as to the availability of any 'new' protection signatures/code. If any are available, the protection device downloads the signature/code from the provider, then installs and activates the new protection.

- **PUSH**

In which the provider contacts the protection device whenever new signature/code updates become available; then sends, installs and activates the new protection.

The PULL system has been in use for decades. Originally, it was a manual process – the user would run a program to retrieve, install and activate the protection updates. Then, as that failed to keep up with the threats, systems were enhanced to provide schedulers to allow for automated checks for monthly and even daily updates. Over the past year or so, modern PULL systems have now brought this down to hourly updates. The increased frequency of updates has been in response to the increased number of threats and the higher initial propagation speed.

True PUSH systems have only recently become widespread (although Network Box has been based on this technology for eight years now). Such systems involve the protection provider releasing an 'update' and rapidly contacting each and every protection device to:

1. Send the updated signatures/code to the device.
2. Install the updated signatures/code.
3. Activate the new signatures/code.

Often systems, which are advertised as 'push' are not truly PUSH – they merely send out a signal to initiate a PULL download. These go some way to speeding up the initial download of the signatures (requirement [1], above), but do not provide for verified installation or activation of the new signatures/code (requirements [2] and [3]).

Why PUSH?

Given that the actual download of the updates utilizes similar technology (a file transfer over a network connection, with speed primarily determined by available bandwidth) for both pull and push systems, what is the advantage of PUSH Technology? Why PUSH?

PUSH Technology provides for three primary advantages over PULL:

1. Speed

Reduces latency (the time from when the update is available until when delivery commences) in the delivery of updates to a minimum.

2. Acknowledgment

Allows for the provider to be certain that updates are installed and activated correctly.


3. Optimization

Provides for optimization of the update system, from the provider's point of view (making the most optimum use of the provider's network for delivery of updates, in both resource utilization and source of updates).

SPEED

Let's imagine a provider releases an update at time T (a random time of the day). In the traditional PULL update system, protection devices will be polling the provider at regular intervals for updates. Let's focus on one particular protection device and look at its behavior. Say the protection device is set to poll once a day. In such a case, the best case would be if the device polls just after update T is released (say T + 1 second), the worst case would be if the device polls just before update T is released (say T - 1 second), and the average case would be half a day. Generally speaking, for PULL update systems with a poll period P, the best case is near zero, the worst case is near P and the average case is P/2.

PULL TECHNOLOGY - DAILY POLL

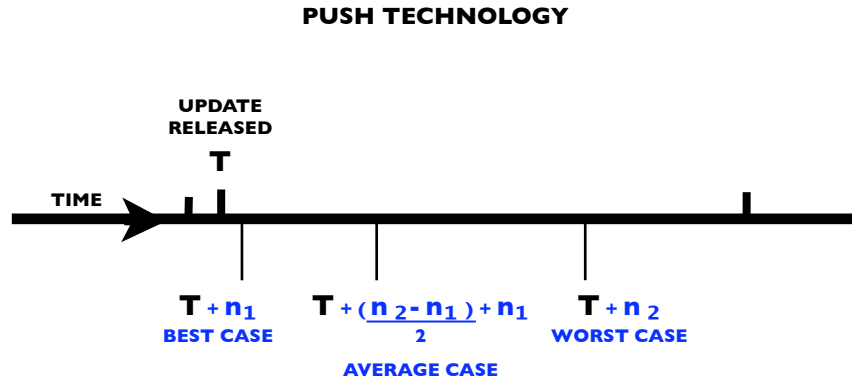


SUMMARY

DAILY POLL: 12 HOURS DELAY

HOURLY POLL: 30 MINUTE DELAY

Now, imagine that provider uses PUSH update technology. The update is released at time T (a random time of the day) and the provider immediately starts to PUSH the update out to protection devices. The protection update will be installed in a time L, dependent on the speed of the provider’s network and push update systems. The time L can be deterministically uniform, as the network is optimized to minimize L - this is typically measured in seconds.



<p>SUMMARY</p> <p>NETWORK BOX</p> <p>PUSH AVERAGE: 23 SECONDS</p> <p>BEST CASE: 1/2 SECOND</p> <p>WORST CASE: 45 SECONDS</p>

Taking the Network Box PUSH Technology as an example, the time for the push update to start (best case; first protection device updated) is half a second, and the time for the last (worst case; last protection device updated) is 45 seconds – leading to an average time of 22.75 seconds. Compare that to 30 minutes (hourly polling) to 12 hours (daily polling) for PULL technology.

ACKNOWLEDGEMENT

In a PULL update system, the provider receives a request to download the update, but has no indication that the update has been installed or activated. There is no way for the provider to know if a particular device has, or hasn’t, activated a particular protection update.

In a PUSH update system, the provider’s equipment sends the update to the protection device and then installs and activates it (receiving an acknowledgement from the device when each of these steps is complete). This gives the provider confirmation that a particular device has the protection signatures/code in place, installed and activated. It also allows the provider to see which devices have the protection in place and control the release of the protection.

OPTIMIZATION

While administrators of protection devices are most likely blissfully unconcerned with the optimization of the provider's update delivery network (except in times of overload and unavailability or slow speed of updates), they should be. PUSH update technology permits the provider complete control over the delivery of updates. The provider can choose to balance the update traffic over links and server systems, to fully utilize all available network capacity. The provider can also choose the source that updates are to be delivered from (to optimize regional delivery, where available).

Compare this to the PULL update system, where the provider is hit with an uncontrollable collection of clients retrieving updates from whichever parts of the provider's network, at whatever time, they like. Such an arrangement leads to an unbalanced update system (with some parts of the provider's systems overloaded, and other parts under-utilized). Overall, the clients suffer and updates take longer to deliver.

Problems with PUSH Technology

So, if PUSH Technology is faster, more deterministic, optimized and allows for acknowledgement of the installation and activation of updates, why use PULL technology at all?

The answer is that PULL Technology has been around for decades, and protection signature providers are used to, and comfortable with, the model. They have an investment in their update delivery networks that is costly to change. Historically, with viruses taking days/weeks to emerge, there was no real requirement for a rapid update delivery system. But, as the threat landscape has changed, most providers have responded by gradually releasing options for increased polling frequency (manual -> monthly -> daily -> hourly) or pseudo-PUSH systems (where a signal is sent out to indicate the availability of a protection update).

You might ask what are the disadvantages of PUSH Technology? Most technology experts will be able to suggest just one – in the event of the release of a faulty signature, it will be sent, installed and activated on all devices rapidly. The problem with this logic is who informs the provider of the faulty signature? The answer is that it is (at least, in the vast majority of cases) the customer, and real-world figures show that customers take an average of between 2 and 48 hours to inform the provider of a problem with signatures. In such a case, 100% of push update, 100% of pull update (hourly polls), 100% of pull update (bi-hourly polls), and 16.7% of pull update (daily polls) customers would be affected in the minimum case of a 2 hour notification (assuming the provider immediately corrected the update).

Looking at the flip side of the problem, assuming the provider fixed the signature and released a new protection set to resolve the problem, how long would it take to install and activate the revised protection? The answer is the same, although in this case favouring PUSH technology. The disadvantage balances the advantage.

Note that in both cases (PULL and PUSH), extensive validation checks are undertaken before the new signatures/code is released to the update system. The difference between the systems is purely from the point it is released to the public.

Conclusion

In this paper, we have presented a background to the technology behind releasing signature/code updates to protection devices. We have shown how the increase in frequency and volume of malicious threats has required providers to improve their technology for delivering protection updates to their customers. We believe that we have also clearly demonstrated the superiority of PUSH technology over the older PULL technology.

Real world experience with PULL vs PUSH also backs up the clear statistical and mathematical analysis. PUSH technology, quite simply, offers the best way to deliver updates to the signatures and code on protection devices.