Death of a Salesforce
Whatever happened to anti-virus?

David Harley, ESET Senior Research Fellow, ESET North America
david.harley.ic@eset.com

Larry Bridwell, Independent Researcher
lbridwell@gmail.com

Abstract

Anti-Virus is, it seems, an ex-parrot. We've seen so many announcements of the death of anti-virus we've taken to carrying black ties around with us, ready for the next one. This paper probably won't have much impact on the ludicrously funereal tone of some commentary, but will take an informed look at the reasons most often given for the imminent demise of the AV industry and in the hope of achieving a balanced view of the present role and future evolution of malware analysis. Reports of the (near-) death of static signature detection may not be exaggerated, but anti-malware technology has moved far beyond simple signatures. We consider in depth the accuracy of some of the basic contentions that keep turning up ad infinitum in memoriam...

1. Conclusions based on detection testing and pseudo-testing statistics
2. Anti-virus is ok if you don't have to pay for it
3. Heuristic detection has gone the way of the static signature
4. Spammed out malware is less important than targeted malware
5. New (mobile) platforms require new defensive paradigms

Catching or blocking malware is just part of the security challenge, at home or in the workplace, and malware detection is a very different technology to what it was 20 years ago, but does that mean it's obsolescent? We look at the three primary functions of AV:

- protection in the form of proactive detection and blocking through a range of heuristic, reputational and generic countermeasures
- detection of known malware
- remediation where something is detected and has managed to gain a foothold

We contend and demonstrate that while emphasis has undergone an irreversible shift from detection by signature, to remediation of signature-detected malware, to more generic detection by technologies such as heuristics, behaviour analysis, and reputation, a complete solution addresses all those issues. AV *is* dead, or at best comatose: at any rate, self-replicating malware is a small part of a much larger problem, while signature detection is primarily a fallback technology that helps with remediation rather than a primary layer of protection.

Anti-malware technology moved on long ago. Customer and media perception, though, has lagged way behind. Could it be that when other sectors of the security industry, driven by commercial agendas, engage in inaccurate and at best misinformed anti-AV commentary, that they are also putting their own interests and those of the community at large at risk? Would a world without the mainstream anti-malware industry be such a good place to live?

Introduction

Anti-virus is apparently dead [1], or at best not worth paying for.

Since the anti-malware industry continues to feed elderly security experts who probably should have retired years ago, you might think that the reports of its death have been greatly exaggerated. [2]

Well, if it *isn't* actually an ex-parrot [3], it is at least almost unrecognizable as the descendant of the primitive technologies of the '80s and early '90s. Though we should point out that the first seeds of both the main branches of AV technology – detection of known malware and generic detection – were already present in the form of on-demand scanning and integrity checking (essentially a form of whitelisting). Still, there is no doubt that both the threatscape and the technologies used to counter those threats have changed drastically.

Why AV is Dead (Apparently)

One of those changes, unfortunately, is that despite the range of technological improvements to basic anti-virus technology that have carried it far beyond simple 'known malware' signature detection, security software with core malware detection technology is less effective – in terms of raw malware detection statistics – than it was twenty years ago. Perhaps that's inevitable: 20-25 years ago, most malware was self-replicating, and the faster it spread, the sooner it was likely to be recognized as malicious. While that speed of spreading was restricted by the fact that the Internet was a far smaller place, that restriction also meant that once a malicious program had been identified, an AV customer who diligently updated his anti-virus as soon as signatures were available was likely to see his signatures before he saw the malware (if at all). Even during the heyday of the fast-burning mass-mailer, it was becoming possible for vendors to push out detection of new malware much faster than had been possible when you had to wait for your monthly (or even quarterly) AV update floppies.

Today, viruses are a very much smaller proportion of all known malware, despite the occasional high-profile virus, and an individual malicious sample often has a very short shelf-life before metamorphoses into a different binary, or at any rate one that conceals the same base code behind a layer of obfuscation. And while we're reluctant to lapse into using poorly-defined buzzwords like spear-phishing, APT, AET and so on, there is a significant range of malware that is highly targeted: that, in itself, restricts the opportunities for vendors to catch sight of new malware, especially if it takes advantage of the burgeoning exploit industry.

It's unsurprising, therefore, that today no reputable anti-malware researcher is likely to claim that his company has 100% malware detection.

Testing and Pseudo-Testing

But when does less effective become ineffective? Conclusions based on detection testing and pseudo-testing statistics are all too common, and are frequently accepted uncritically by journalists [4], even when the statistics come from an arguably marketing-tainted source. In many cases, they come from providers of alternative security technologies wanting a bigger slice of customer pie. And that's OK as long as potential customers aren't misled about the capability of *either* technology. [5]

We aren't going to tell you that no anti-malware vendor ever used misleading statistics to make it look better than its competitors (anti-virus companies or other security technology providers) and you wouldn't believe us if we did. We are going to tell you that the same errors we've observed time and time again in comparative tests [6; 7] also turn up regularly in tests (using the term loosely) where the comparison is not between AV products, but between the security technology

represented by the tester (or the company that sponsors the test) and the technology often referred to generically (if not altogether correctly, nowadays) as the anti-virus industry. There have also been tests by security intelligence services that didn't make direct comparisons with other technologies but did, nevertheless, use performance metrics that misled because of faulty methodology. For example, Cyveillance concluded that

> …the average time it takes for the top AV solutions to catch up to new malware threats ranges from 2 to 27 days. Given that this metric does not include malware files still undetected even after thirty days, the AV industry has much room for improvement. [8]

However, its attempt to make its point using a form of Time-to-Update (TtU) testing was impaired [9] by the use of a small sample set and inadequate sample validation by scanning samples with multiple engines and accepting or discarding each sample according to the number of scanners that detected it. That approach causes (at least) two problems: if three or more scanners incorrectly classify a sample as malicious (yes, it can certainly happen), then other products are penalized for not generating false positives. Conversely, if samples are discarded because less than three scanners detected it, the detection rates of those scanners are negatively skewed.

In fact, while some form of longitudinal testing is often considered better than a 'snapshot' test, especially in certification testing, TtU testing isn't generally considered a useful approach, certainly on the Windows platform where it arguably doesn't matter if a product doesn't detect a sample 30 days down the line that had an effective lifetime of five minutes due to server-side polymorphism.

It's not clear exactly what form of (pseudo-)validation Cyveillance used, but according to Kevin Townsend [10], M86/Trustwave, used VirusTotal to evaluate AV performance regarding detection of Asprox and a Skype vulnerability. While apparently agreeing that this is bad test methodology, they argued that 'The value in using VirusTotal is that it reflects what a lot of organizations will be using in live environments rather than a test lab.' [This strikes us as making almost as much sense as not tuning a piano before every concert performance because most pianos aren't used for concerts and are therefore less frequently tuned.]

The report [5] by Imperva and the Technion-Israeli Institute of Technology was certainly founded on the assumption that VirusTotal reports provide a reliable way of ascertaining whether an AV product does or doesn't detect a given sample of malicious software.

On the basis of a set of 80 samples submitted to VirusTotal, Imperva's study asserted that

- The initial detection rate of a newly created virus is less than 5%.
- Some vendors may take up to four weeks to detect a new virus.
- Free packages have the best detection rates despite their high false positive rate.

The quasi-test – implemented despite VirusTotal's own recommendations and commentary [11] – frustrated the AV research community, who were clearly being targeted as being 'proven' ineffective, leading to the conclusion that "easing the need for AV could free up money for more effective security measures." [5]

In fact, it should have frustrated *anyone* who cares about accurate testing and evaluation of security products and strategies. VirusTotal is intended and designed to give some idea of whether a given file is likely to be malicious, not to evaluate the ability of one or more security products to detect it. At best, it tells you whether those products are capable of detecting it using the particular program module and configuration used by VirusTotal. [12]

Gunter Ollman of IOactive, another company with alternative services to sell, put up an assessment of AV effectiveness that makes Imperva's look generous, though he doesn't say how he reaches that conclusion [13]:

> …desktop antivirus detection typically hovers at 1-2% … For newly minted malware that is designed to target corporate victims, the rate is pretty much 0% and can remain that way for hundreds of days after the malware has been released in to the wild.

Apparently he knows this from his own experience, so there's no need to justify the percentages. Imperva, however, in the face of some fairly high-powered criticism [14] tried to defend its '5%' position by citing the 'similar' results from an AV-Test report to which no link was given, and we were unable to find a public version at the time. However, the screenshot included in Imperva's blog seemed to show the average industry detection results in three scenarios:

- 0-day malware attacks: avg = 87% (n=102)
- malware discovered over last 2-3 months: avg = 98% (n=272,799)
- widespread and prevalent malware: avg=100% (n=5,000)

We don't claim to be statisticians, but are fairly sure that the disparity between 5% and 87% is fairly significant in some statistical sense.

Virus Total was never intended as a test of scanner performance, and a number of factors make it unsuitable for such pseudo-testing. It's a fundamental tenet of testing that testing should not introduce bias [15] but, depending on the test objective, testers can introduce bias according to how they approach configuration of products under test. VirusTotal can't be accused of introducing bias, of course, since it isn't intended to compare performance, but (mis)use of the service for pseudo-testing *does* effectively introduce a bias. Some products will flag Possibly Unwanted Applications as malware: on some products, this is because of default settings, and in other cases, because VirusTotal has been asked to turn on a non-default option. In other words, some products – as configured by VT – may never detect certain samples because they're not unequivocally malicious, yet would detect them as 'possibly unwanted' if the configuration was changed, while others may flag certain samples inappropriately. Some products may be able to detect certain samples on-access, but not on-demand, because not all approaches to behaviour analysis and heuristics can be implemented in static/passive scanning. Of course, it would be a legitimate test target to compare default settings over a range of products as long as it wasn't done in such a way as to generate inaccurate detection results by ignoring the impact of configuration on detection.

In this case, we have no idea what samples Imperva was looking at and whether they were correctly classified as malware, still less about their prevalence or the criteria that governed the choice of those specific samples.

Virus Total itself has spoken out several times against the misuse of the service as a substitute for testing, as discussed in a paper by David Harley and Julio Canto. [16]

Here are some salient extracts.

> *VirusTotal was not designed as a tool to perform AV comparative analyses, but to check suspicious samples with multiple engines, and to help AV labs by forwarding them the malware they failed to detect….*

> *VirusTotal uses a group of very heterogeneous engines. AV products may implement roughly equivalent functionality in enormously different ways, and VT doesn't exercise all the layers of functionality that may be present in a modern security product.*

VirusTotal uses command-line versions: that also affects execution context, which may mean that a product fails to detect something it would detect in a more realistic context.

It uses the parameters that AV vendors indicate: if you think of this as a (pseudo)test, then consider that you're testing vendor philosophy in terms of default configurations, not objective performance.

Some products are targeted for the gateway: gateway products are likely to be configured according to very different presumptions to those that govern desktop product configuration.

Some of the heuristic parameters employed are very sensitive, not to mention paranoid.

VirusTotal is self-described as a TOOL, not a SOLUTION: it's a highly collaborative enterprise, allowing the industry and users to help each other. As with any other tool (especially other public multi-scanner sites), it's better suited to some contexts than others. It can be used for useful research or can be misused for purposes for which it was never intended, and the reader must have a minimum of knowledge and understanding to interpret the results correctly. With tools that are less impartial in origin, and/or less comprehensively documented, the risk of misunderstanding and misuse is even greater.

Free and For-Fee

Imperva also asserted that it was not proposing that AV be dispensed with altogether, but that 'both consumers and enterprises should look into freeware as well as new security models for protection'. A theme taken up by Ollman:

If it's free, never ever bothers me with popups, and I never need to know it's there, then it's not worth the effort uninstalling it and I guess it can stay…

As it happens, the anti-malware industry isn't necessarily averse to free AV [17], even though we do need products that generate a revenue stream: even AV gurus like to eat occasionally. Free versions of commercial products have some benefit to the user community but they're also a very viable marketing tool. They usually take one of the following three forms:

- free-for-personal-use scanners with limited functionality and support
- online scanners that give instant access to an up-to-date engine with limited functionality and support
- fully-featured evaluation copies.

The cost of producing mainstream AV has to be offset somewhere [18], and it's usually underwritten by income from a for-fee, expanded-functionality version. It's true (and unfortunate) that consumer magazines cater for an audience that doesn't always understand the need for AV, doesn't want to pay for it if it can be helped, and is, not unlike the business sector, far more forgiving towards what it doesn't pay for [19]. We're not convinced that price (let alone the absence thereof) is the best primary determining factor for selecting security software, though you might feasibly do as well with a suitable combination of free products as with a commercial single-vendor suite, if you have the time and expertise to select, configure and maintain those layers of security. (That's a *big* If…)

But marketers actually do potential customers a disservice by suggesting that companies should use free anti-virus so as to be able to afford their own solutions. Not only because that advice ignores the licensing stipulations that usually govern the legitimate use of free versions of commercial products, and not only because free AV has restricted functionality and support, but also because security suites – even if they're the only security software in use – offer more comprehensive, multi-

layered protection than a product - free or otherwise - that only offers one layer of protection. (Though even products that only offer 'anti-virus' functionality actually detect a lot more than just viruses.)

Detection Technology

Heuristic detection has gone the way of the static signature, apparently? Well, when people take this position, we presume they have in mind the relatively unsophisticated heuristics of the early 1990s, rather than the complex multi-layering of a modern security program with core malware detection technology, making use of passive and active heuristic analysis, behaviour blocking, sandboxing, behaviour analysis, whitelisting, integrity checking, traffic analysis and emulation.

There is, in fact, a rational debate to be held on whether AV – certainly raw AV with no multi-layering bells and whistles – *should* be on the point of extinction. The rate of detection for specialized, targeted malware like Stuxnet is indeed very low, with all-too-well-known instances of low-distribution but high-profile malware lying around undetected for years. (If such malware is aimed at parts of the world like Iran where most commercial AV cannot legally reach, that obviously doesn't help.)

Pierre Vandevenne observed [20] in response to an article by David Harley at Securiteam [21]:

> Traditional stand-alone A-V (essentially the scan-detect-protect-clean paradigm) should definitely be dead. Multi-layered protections with web browsing protection, DNS monitoring, in the cloud file checks and heuristics, real time analysis of new and/or infrequent or unique executables (of all kinds) etc… are definitely needed but won't ever reach the near-perfect protection levels the A-V industry offered at very specific and short lived moments in the history of malware.

> But the public's mind remains stuck in the "scan-detect-protect-clean" era thanks to some 20 years of repetitive dumb marketing by A-V companies. Just look at the promotional material and white papers offered by any anti-virus company? Can you find one that doesn't refer to some kind of award won in some "scan-etc…" test? Can you find one that doesn't claim to offer "best" or often "near perfect" detection or protection percentages?

No reputable researcher will claim that an anti-malware product is capable of 100% detection levels or anything close to it. [22]

 But then, there are no absolute solutions. A totally generic solution may get close to blocking 100 per cent of threats, but will discard some 'true positive' objects. Unfortunately, the kind of service that is often found carrying out pseudo-testing in order to prove the ineffectiveness of AV is far less frequently exposed to testing – sound or otherwise – of its own claimed effectiveness. While a defensive anti-malware industry continues to promote the raising of testing standards that affect its own product ranges, it hasn't yet turned its attention to other security sectors, and companies within those sectors are probably quite comfortable with the situation as it stands.

However, the use of detection statistics and test performance as a promotional tool is a permanently contentious point, one that isn't far removed from a point made by Kevin Townsend [23] about WildList testing [24]. WildList testing is a pretty limited measure of anti-malware effectiveness, if not yet completely valueless [25]. But if companies use marketing that suggests that everything in the wild (whatever you may understand by that) is on the WildList, implying that a

100% detection of WildCore = 100% protection, that sets an expectation just as unrealistic as the 0-5% figures bandied by AV's critics. Vandevenne continued:

> Fundamentally, what is attacked is not how a modern A-V works but how it is perceived by the public. And that perception was created by the A-V vendors themselves… We've had the example of a positive reality distortion field with Apple. We're experiencing a negative one on the A-V industry as a whole right now. If I was launching a competing product today, I would probably build its internals quite like those of a modern A-V, but would market it as "Definitely Not an A-V".

That's a lesson that even established companies with a core anti-malware product functionality are learning: one vendor, evidently keen to disassociate itself from the tainted term 'anti(-)virus' announced that

> "Antivirus only stopped 49% of malware in 2012…The security of your data depends on your ability to move from reactive to proactive defense that involve both intelligence and policy based protection. [26]

We don't doubt the effectiveness of whatever functionality is incorporated in the product range addressed in the webcast that this announcement heralded, but 'a proactive approach to endpoint protection' is hardly incompatible with modern anti-malware products and security suites.

While products continue to improve, even as we move decades away from static detection as a primary detection technique, we need to make it clearer to people who write and read reviews and security articles, both influencers and customers, what our products *really* do and what they can *realistically* expect from us.

APTitude Adjustment

Is spammed out malware less important than targeted malware? It has long been acknowledged by the AV research community [27] that heavily obfuscated, stealthy and highly targeted attacks are an area in which AV is weak. Mikko Hyppönen [28] described the fact that for a very lengthy period AV vendors had samples of Flame yet didn't realize its significance as a 'spectacular failure' on the part of the anti-malware industry. As, in PR terms, it undoubtedly was, allowing our critics to immediately jump on other examples of 'incompetence'. [27]

Mikko has probably forgotten more about malware and anti-malware technology than we ever knew between the two of us, but in this instance he has it back to front, or has at least allowed our critics to misrepresent the situation. In an era where anti-malware labs process hundreds of thousands of samples a day, failure to realize the significance of a vanishingly small set of stealthy, low-prevalence samples – however great their subsequent impact – while hardly describable as a success, is hardly a spectacular failure in statistical terms.

It can, of course, be described as a failure of automated preliminary analysis – clearly, manual analysis of several hundred thousand samples per day is beyond the resources of any lab we know of – to flag those samples as requiring further investigation. Yes, once again security technology failed to provide 100% protection. Perhaps we should just give up and go back to 19th century technologies, since in general those security vendors now fighting for their own share of the budgets currently allocated to AV didn't do any better.

David Harley addressed the assertion by Schneier [29] and others that "anti-virus technology does not bother to block non-generic, targeted attacks" in an email discussion subsequently reported by Dan Raywood. [22]

> The sheer number of malicious attacks does mean that anti-virus labs have to prioritize to some extent, but that prioritization is rather more complex than that and it is far from the only factor in detection. The relationship between a given binary and other malware families, for instance, is a big factor in determining whether that binary is detected at a time when there is no malware-specific detection for it.

> It's quite possible that a single, targeted attack won't be detected initially by many or any anti-virus solutions, especially if it involves the combination of a zero-day and the use of multi-scanning to tweak the binary until no common engine detects it, but to dismiss anti-virus on those grounds is to throw out a whole generation of babies with a very small quantity of bathwater.

> The fact that anti-virus is focused on malicious binaries does make it less effective in attack scenarios that are more generic in nature, but that's why you need multi-layering. Horses for courses.

Stuck Inside of Mobile (with the smartphone blues again) [30]

It has been asserted that new (mobile) platforms require new defensive paradigms. Is the mobile device (smartphone, tablet, smartwatch, wearable computer) the future of computing?

We're as fond of a gadget as anyone else, but we don't anticipate parting with our full-featured desktop and laptop machines (ok, laptops are certainly somewhat mobile, but that's not what the evangelists of hypermobility are usually discussing) until technologists magically marry extreme miniaturization of footprint with the ergonomic and functional convenience of type-able keyboards and apps that have the same rich functionality as desktop versions. Nevertheless, there's no denying that the restrictions imposed on app building and approval by such platforms as iOS, Windows RT 8, and Android has had a decided impact on the feasibility, methodology and testing of security programs, even if that impact has not always translated to an absence of malware within a given environment [31]. While we can debate how you measure the impact and prevalence of malware in a semi-liberal environment like Android [32; 33] and the nature of the primary threats on these platforms is not the same as it is on older desktop operating systems, the challenge is not to find any use at all for security software, but to adapt security technology to substantially different threats and operating environments. Given the evidence of substantial adaptation and evolution in anti-malware over the past few decades, we aren't writing malware detection technology off just yet.

The Changing Role of Malware Detection Technology

In our opinion, malware detection as implemented by what some still refer to as the anti-virus industry has three main components. Although each component has been regarded as its 'primary' function by different commentators at different times, all three still have a part to play in a modern anti-malware product.

Proactive Detection and Blocking

The holy grail of security software is protection in the form of proactive blocking through a range of heuristic, reputational and generic countermeasures. In other words, stop badware (and other forms of attack) gaining a foothold on a protected system in the first place.

Detection of Known Malware

First there were viruses (in the broad sense of self-replicating malware in many guises). And yes, there were trojans too, but in much smaller quantities, except in the limited sense in which viruses can also be described as trojans – or at any rate, virus-infected code can be described as trojanized. As the balance between self-replicating and non-replicating malware slowly shifted, detection technology also changed, from exact identification to near-exact, to passive heuristics, to active heuristics and sandboxing, to reputational analysis and so on. Unfortunately, malware technology also evolved in ways that reduced the effectiveness of these enhancements. Nonetheless, a high proportion of threats and threat variants continue to be detected either specifically or using more generic detections. We don't advocate that you discard mainstream technology because it can't achieve 100% detection unless you have an acceptably-convenient alternative that *does* achieve (near-) perfection.

Remediation

Remediation where something is detected after it has gained a foothold (i.e. infected and made some undesirable modification to the system.

There is no way back: Elvis has left the building and the genie has broken the bottle. For a while detection got the lion's share of the developer's attention. As the glut problem began to bite and detection by static signature declined in effectiveness, infection became more sophisticated and harder to reverse, and remediation needed more attention, though we've rarely agreed with those who've said that once you're infected, there's nothing to do but re-image. Then things began to change with heuristics, behaviour analysis, reputation and the rest. Today AV *is* dead – in the sense signature detection is either dead or at least taking a nap until it's needed. Now it is anti-malware and protection is achieved through reputation, behaviour, advanced heuristics, and signatures are primarily used for remediation where proactive methods have failed.

Conclusion

As Kurt Wismer pointed out recently [34], the anti-malware research community has been urging people not to rely on anti-malware alone for decades.

Don't worry about offending the AV industry– no-one else cares, and we've had to grow thick skins – but consider whether you want to base your security strategy (at home or at work) on a PR exercise based on statistical misrepresentation and misunderstanding. Don't be too optimistic about finding The One True (probably generic) Solution: look for combinations of solution that give you the best coverage at a price you can afford. The principle applies to home users too: the right free antivirus is a lot better than no protection, but the relatively low outlay for a competent security suite is well worth it for the extra layers of protection. As one of us stated elsewhere [22]:

> "Personally (and in principle) I'd rather advocate a sound combination of defensive layers than advocate the substitution of one non-panacea for another, as vendors in other security spaces sometimes seem to. Actually, a modern anti-virus solution is already a compromise between malware-specific and generic detection, but I still wouldn't advocate anti-virus as a sole solution, any more than I would IPS, or whitelisting, or a firewall."

Since, by definition, we can't say what '100%' of known and unknown malware means at any moment in time, we can't say what percentage of that totality is *detected* at any one time by any one AV product, let alone all products. We can say, though, that a very significant proportion of new threats are immediately detected by some form of code and/or behaviour analysis. It's nothing like 100%, and no AV researcher worth listening to would claim that it is, but it's a lot more than 0%.

More to the point, if there's *any* single security solution (not just AV) that offers 100% detection and/or blocking of all malware and is still easy and convenient to use, totally transparent to all business processes, and never, ever generates some form of false positive, perhaps someone would tell us what it is so we can go and buy a copy.

In the meantime, as Blaze reminds us [35], the real bad boys here are the cybercriminals, not AV, the security industry, the media or even government agencies, though the latter have stretched out patience and credulity pretty thin lately. However, it's also worth reminding ourselves that while the lion's share of the blame belongs to the criminal, marketing also plays a part if it encourages the customer to be totally dependent on panacea-du-jour technical controls without ever using his own common sense, and that it's not only vendors who are in the business of self-promotion in order to justify their budget allocation.

A major justification for malware-oriented security software is that it provides OS vendors and the vendors behind competing technologies with an incentive to try to keep ahead of malware (and anti-malware). If all the researchers in AV labs retired or went to social media startups, the long-term impact on the overall detection (and therefore blocking) of malware would be considerable. Those free products are effectively subsidized by commercial products – though they constitute a loss leader that may help to sell those same commercial products – and considerable resources and expertise are needed to maintain a quality anti-malware product. We don't see how [36] effective but free anti-malware technology – as opposed to less effective products maintained by enthusiastic amateurs or as a 'value-add' to a different kind of security product – could survive. As one of the authors has previously [36] remarked:

> "…would the same companies currently dissing AV while piggybacking its research be able to match the expertise of the people currently working in anti-malware labs?"

References

[1] Chirgwin R., *Anti-virus products are rubbish, says Imperva*, The Register, 2013: http://www.theregister.co.uk/2013/01/01/anti_virus_is_rubbish/

[2] Clements, S. (Mark Twain): http://www.brainyquote.com/quotes/quotes/m/marktwain141773.html

[3] Monty Python: http://www.davidpbrown.co.uk/jokes/monty-python-parrot.html

[4] Old Mac Bloggit, *Journalism's Dirty Little Secret*, Anti-Malware Testing, 2013: http://antimalwaretesting.wordpress.com/2013/01/02/journalisms-dirty-little-secret/

[5] Imperva, *Assessing the Effectiveness of Antivirus Solutions*, 2012: http://www.imperva.com/docs/HII_Assessing_the_Effectiveness_of_Antivirus_Solutions.pdf

[6] Kosinár, P.; Malcho, J.; Marko, R.; Harley, D., *AV Testing Exposed*, Proceedings of the 20th Virus Bulletin International Conference, 2010: http://go.eset.com/us/resources/white-papers/Kosinar-etal-VB2010.pdf

[7] Harley, D., *How to Screw Up Testing*, Anti-MalwareTesting, 2010: http://antimalwaretesting.wordpress.com/2010/06/21/how-to-screw-up-testing/

[8] Cyveillance, *Malware Detection Rates for Leading AV Solutions*, Cyveillance, 2010:
https://www.cyveillance.com/web/docs/WP_MalwareDetectionRates.pdf

[9] Abrams, R., *How to Screw Up and Skew a Test*, ESET, 2010:
http://www.welivesecurity.com/2010/08/06/how-to-screw-up-and-skew-a-test/

[10] Townsend, K., *Anti-Malware Testing Standards Organization: a dissenting view*, 2010
http://kevtownsend.wordpress.com/2010/06/27/anti-malware-testing-standards-organization-a-dissenting-view/

[11] VirusTotal: https://www.virustotal.com/about

[12] Harley, D., *Imperva, VirusTotal and whether AV is useful*, ESET, 2013
http://www.welivesecurity.com/2013/01/03/imperva-virustotal-and-whether-av-is-useful/

[13] Ollman, G., *The Demise of Desktop Antivirus*, 2013: http://blog.ioactive.com/2013/01/the-demise-of-desktop-antivirus.html

[14] Eddy, M., *Experts Slam Imperva Antivirus Study*, PC Magazine, 2013:
http://securitywatch.pcmag.com/none/306552-experts-slam-imperva-antivirus-study

[15] Anti-Malware Testing Standards Organization, 2008:
http://www.amtso.org/released/20081031_AMTSO_Fundamental_Principles_of_Testing.pdf

[16] Harley, D. & Canto, J., *Man, Myth, Malware and Multi-scanning*, ESET, 2011:
http://go.eset.com/us/resources/white-papers/cfet2011_multiscanning_paper.pdf

[17] Harley, D., *Security Software & Rogue Economics: New Technology or New Marketing?*, EICAR 2011 Conference Proceedings: http://smallbluegreenblog.files.wordpress.com/2011/05/eicar-2011-paper.pdf

[18] Schrott, U., *How Free is Free Antivirus*, ESET, 2010:
http://www.welivesecurity.com/2010/04/14/guest-blog-how-free-is-free-antivirus/

[19] Harley, D., *Untangling the Wheat from the Chaff in Comparative Anti-Virus Reviews*, Small Blue-Green World, 2006-2008: http://go.eset.com/us/resources/white-papers/AV_comparative_guide.pdf

[20] Harley, D., *The Posthumous Role of AV*, Anti-Malware Testing, 2013:
http://antimalwaretesting.wordpress.com/2013/01/09/the-posthumous-role-of-av/

[21] Harley, D., *The death of AV. Yet again*, Securiteam, 2013:
http://blogs.securiteam.com/index.php/archives/2037

[22] Raywood, D. *Palo Alto Networks CTO: anti-virus technology can't stop targeted attacks*, SC Magazine, 2011: http://www.scmagazineuk.com/palo-alto-networks-cto-anti-virus-technology-cant-stop-targeted-attacks/article/211543/

[23] Townsend, K., *Old Mac Bloggit isn't really a grumpy old man*, 2013:
http://kevtownsend.wordpress.com/2013/01/02/old-mac-bloggit-isnt-really-a-grumpy-old-man/

[24] Harley, D., *Going beyond Imperva and VirusTotal*, 2013:
http://antimalwaretesting.wordpress.com/2013/01/03/going-beyond-imperva-and-virustotal/

[25] Harley, D. and Lee, A., *Call of the WildList: Last Orders for WildCore-Based Testing?*, Virus Bulletin Conference Proceedings, 2010: http://www.welivesecurity.com/media_files/white-papers/Harley-Lee-VB2010.pdf

[26] Symantec, 2013: https://symantecevents.verite.com/?action=event.dsp_event&event_id=29201&view_reg=1&acode=119104

[27] Tung, L., *AV just doesn't work for targeted attacks: Schneier*, CSO Online (Australia), 2012: http://www.cso.com.au/article/428143/av_just_doesn_t_work_targeted_attacks_schneier/

[28] Hyppönen, M., *Why Antivirus Companies Like Mine Failed to Catch Flame and Stuxnet*, Wired, 2012: http://www.wired.com/threatlevel/2012/06/internet-security-fail/

[29] Schneier, B., *The Failure of Anti-Virus Companies to Catch Military Malware*, 2012: https://www.schneier.com/blog/archives/2012/06/the_failure_of_3.html

[30] Wikipedia: Stuck Inside of Mobile with the Memphis Blues Again
http://en.wikipedia.org/wiki/Stuck_Inside_of_Mobile_with_the_Memphis_Blues_Again

[31] Cobb, S., *Android security issues: does a Microsoft Windows analogy make sense?*, ESET, 2013: http://www.welivesecurity.com/2013/03/11/android-security-issues-does-a-microsoft-windows-analogy-make-sense/

[32] Harley, D. & Myers, L., *Mac Hacking: the Way to Better Testing?*, Virus Bulletin Conference Proceedings, 2013: http://www.virusbtn.com/conference/vb2013/abstracts/HarleyMyers.xml

[33] Harley, D., *Memoirs of a Charlatan Scammer*, Mac Virus, 2011: http://macviruscom.wordpress.com/2011/11/21/memoirs-of-a-charlatan-scammer/

[34] Wismer, K., Twitter, 2013: https://twitter.com/imaguid/status/377067760620408832

[35] Blaze, *The Malware Blame Game*, Blaze's Security Blog, 2013: http://bartblaze.blogspot.co.uk/2013/09/malware-blame-game.html

[36] Harley D., *Anti-Virus: Last Rites, or Rites of Passage?*, Virus Bulletin, February 2013: http://antimalwaretesting.files.wordpress.com/2013/05/dharley-feb2013.pdf