FEATURE 2

Rescue Me: Updating Anti-Virus Rescue Systems

Andreas Marx, AV-Test.org
University of Magdeburg, Germany

The problem is an old one: if the PC of a home user becomes infected by a virus, the user is advised to boot the machine from a 'known good' virus-free disk to scan and disinfect the local hard drives.

In times of DOS, *Windows 95*, *98* and *ME* this presented only a minor problem – most virus scanners included bootable disks in their retail package, otherwise the user was able to create them during or after installation. Another solution was a bootable installation CD-ROM, which has the benefit of being write-protected and therefore safe from viruses. The FAT16 and FAT32 platforms were well known and caused no great problems, if file or even boot viruses had to be disinfected.

The Problem: Windows XP

The situation has changed with the arrival of *Windows XP*: NTFS drives have become common as primary file systems for home users. And that's the problem – the majority of virus scanners are unable to scan NTFS drives, if started from their rescue disks or CD-ROMs (despite the fact that the manufacturers of these scanners claimed that their products were completely ready for and compatible with *Windows XP*).

Currently I know of only a few working solutions that can be started directly from the bootable CD. The first is *AntiVirusKit* by *G Data*, which uses a *Linux* version of *Kaspersky Anti-Virus* and is included in the retail product. *Kaspersky Labs* has its own solution, however this is no longer included in the retail product; it is only sold separately. The third is *AntiVir* by *H+BEDV Datentechnik*, which uses *Linux*. The rescue system is included in the retail product as well as on all demonstration CDs of *AntiVir Personal Edition*. The same product is also distributed under the name *Vexira Antivirus* by *Central Command* in the US. According to Igor Muttik, *Network Associates Inc.* has developed a similar rescue system on floppy disks which is based on DOS. However, it is not available due to the high licensing cost of the third-party drivers and tools used.

As far as I am aware, all other anti-virus programs (regardless of whether they are started from the CD or from included or creatable boot disks) will completely ignore and skip NTFS drives. Some programs indicate that NTFS drives have been found, but not scanned, while other programs simply report after a few seconds that they have scanned all available drives and they are confirmed clean – although, in fact, nothing has been scanned.

Under *Windows XP*, there is a feature called Recovery Console, which (if installed) can be started at boot-up by pressing F8. After a log-in using the administrator account and password, the user can access all data on NTFS drives. Command-line operations like 'copy' or 'ren' will work, but no external programs can be started. Therefore, it can be a great help, but only if the problems are known.

Computer Magazine Solutions

In response to a large number of requests from their readers, two German computer magazines have published articles on work-arounds for this problem.

An article in the technical *c't* magazine (*c't* 25/2001, p.250) included a manual instruction guide on how a rescue CD-ROM can be created, not only for use in the case of virus infections, but also for the recovery of accidentally deleted data etc.

Their solution was largely *Linux*-based, because *Linux* has built-in NTFS drivers (besides drivers for various other systems, like FAT16/32, HPFS, its own ext2/3, ReiserFS and so on), which are reliable for read operations. Write operations, however, are dangerous according to the author of the NTFS drivers and according to our own tests. The main reason for this is that very few of the facts about how NTFS works are public and the authors had to reverse-engineer a lot of these complex internal structures.

However, *Linux* is not needed at all – a DOS NTFS driver is available from *Sysinternals* (http://www.sysinternals.com/), which works quite well, even if it consumes a lot of memory. Write operations are not permitted in the freeware version; the registered version costs US$49.

The more end-user-focused *PC-WELT* has published a completely ready solution on their bootable cover CD-ROM, which is based on a *Linux* beta version of *F-Prot* (*PC-WELT* 4/2002, p.154). According to the authors, it took only a few hours to prepare the solution, because it is based mainly on the *Linux* rescue system Rip-45 (see http://www.ibiblio.org/pub/Linux/system/recovery/) and the Live System Knopper (http://www.knopper.net/). The bulk of the work went into creating a menu system, from which the user can select what he or she wants to do.

If two magazines have already published a working solution, it should not be a great deal of trouble for AV companies to do the same in order to provide their customers with better protection. A few other points should be taken into consideration, such as updates of the signature files (CDs are usually old) or the ability to create and save log files.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form without the prior written permission of the publishers.

The following are the most important issues that must be dealt with in order to achieve a working solution. The facts are not limited to NTFS drives or *Windows XP*, but generic and useable for every platform. Furthermore, these features are useful not only for home user systems, but also for system administrators as well as computer retail stores, for example, if they are given a computer for further analysis.

Linux, BSD or DOS?

One of the first questions that has to be answered is which platform should be used.

Linux is available free of charge, but solutions that are based on *Linux* must be published as source code, according to the GNU General Public License (see http://www.fsf.org/licenses/gpl.txt). Of course, the source code of the virus scanner does not need to be published, but all kernel and script modifications that are based on GPL-protected program code must be.

An advantage of using *Linux* is its easy implementation, because the system can be configured as needed and the 32-bit program code of the scanner and helper programs can be run without any memory management problems.

FreeBSD, *OpenBSD* and *NetBSD* include essentially the same (optional) NTFS driver as *Linux*, but the *Linux* version is updated more frequently than the BSD port. However, these operating systems do not require any source code to be published, even after changes have been made (see http://www.freebsd.org/copyright/). [In the rest of this article, I shall not differentiate between *Linux* and *BSD*, but use *Linux* as a synonym for all of the open source operating systems mentioned above.]

There are two main possibilities for using a DOS platform. MS-DOS is shipped with *Windows 98*, where licence fees must be paid, but is the most compatible platform for other programs. FreeDOS could be used as an alternative to MS-DOS (http://www.freedos.org). This can be used free of charge under GPL terms, just like *Linux*. However, FreeDOS is only 99 per cent compatible with other DOS applications.

DOS will require a few additional drivers for the CD-ROM (if needed) and SCSI drives, as well as for the memory management. The scanner itself is likely to require a DOS extender – as well, of course, as a tool such as NTFSDOS, to make NTFS drives accessible. It might be a good idea to include a few tools such as ‘Fdisk’, ‘Sys’ and ‘Format’. Free file managers, like the Midnight Commander for *Linux* or the Volkov Commander for DOS can also be useful for later manual rescue operations.

Start-up Process

The boot process is fairly similar for these platforms. For *Linux* systems, the compiled monolithic compressed kernel must be started, which includes all the necessary drivers and, for DOS, the usual configuration and start-up files config.sys and autoexec.bat will run. This can be done from a simple 1.44MB boot disk or a bootable CD-ROM, which uses almost the same method: the BIOS will simply load and start a disk image which is stored in a special area at the very beginning of the CD-ROM.

Of course, the user has to change the boot order first, so that the A: drive or the CD-ROM is used before the hard disk in the boot sequence. The rescue system should notify the user that these changes have to be undone after a successful scan or repair session to prevent boot virus infections.

After loading all the required drivers, the scan process should not start automatically. Instead the user should be prompted with a simple menu, with options such as ‘scan all hard drives’, ‘scan selected drives’, ‘scan floppy disk’, ‘exit to operating system’, ‘run a special command or program’, and so on.

The ability to test the complete hard disk for read errors (simply try to read everything, sector by sector) would be a useful feature. In many cases the root of the problem is not a virus, but hardware – and in particular hard disk – errors. Of course, a help window or help function with short instructions should be included as well.

First, however, the user should be able to update the virus scanner, because the CD-ROM or the rescue disks that have been created will usually be quite old. For this, it should be possible to update the signature files from an external drive, such as a floppy disk – alternatively the signature databases stored on the local hard disk should be used, if they are valid.

The process is more complex if engine or scanner updates are needed, because such files are more likely to be infected by a virus and are highly platform-dependent (a *Windows* DLL won’t run easily under DOS or *Linux*), therefore only trustworthy sources should be used. However, this should not be a problem if a scanner uses an ‘integrated’ solution, where all data (both engine and signatures) are stored in one or more encrypted, digitally signed file(s).

There is one main limitation: only 1.44 MB of data can be stored on a disk. All scanner files should be able to fit on a disk, which means that none of the files should be larger than 1.44 MB. Due to the size of most scanner databases, it’s likely that more than one disk will be needed.

It may be possible to download the required updates from the local network or even the Internet or other kinds of dial-up or DSL connections, after loading all the necessary network and TCP/IP drivers etc., but that is likely to be a prohibitively complex task.

Scan Selected Directories

The scan process itself can raise a few problems, starting with the drive selection, especially under *Linux*. Here, the common drive letters such as ‘A:’ or ‘C:’ are not available,

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form without the prior written permission of the publishers.

instead names like '/dev/fd0' or '/dev/hda1' are used. For user convenience, it should be possible to display the *Windows* drives convention as well. This should be quite an easy task provided these mappings have not been changed manually using the *Windows* Drive Management tools. It may also be useful to display a few further details of the drive, such as the label, file system type and its size. The rest of the scan process is a simple run of the DOS or *Linux* virus scanner with a few parameters, like the selected drives or folders.

Currently, there is no known ability to scan NTFS online encrypted files and folders in this situation, even if the password of the user account is known. These files can only be accessed in the recovery console or if the user is logged in under *Windows*. If a virus scanner attempts to access such a file or folder, it should display a brief warning message rather than skipping the file silently. Therefore, it's important not to store programs in such encrypted areas of the hard disk, which will start automatically, for example from the Autostart folder or 'Run' registry key.

Disinfection Trouble

Boot viruses are one of the oldest forms of these digital parasites, but even ten-year-old boot viruses are still found in the wild. They do not cause very much trouble for DOS-based *Windows* versions, such as *95*, *98* and *ME* – often, these can still be started after infection.

However, *NT*-based *Windows* versions, such as *2000* and *XP*, will not start after a boot virus infection, making a rescue system very important. While DOS-based scanners can identify and repair boot viruses quickly, in *Linux*-based programs the detection routines to scan the MBR and all bootsectors for this kind of malware are often not yet implemented, making boot viruses invisible to the scanner.

Macro viruses should not cause any greater trouble, but file viruses, Trojan horses, backdoors and worms can cause a few problems, if they change INI files or Registry values. For example, if they start automatically at boot-up time ('Run', 'RunService' keys etc.) or if they change the properties of a file type, as Pretty_Park and SubSeven will do for an EXE file. In this case, it would be a good idea to replace the malicious files by a helper program, which can undo the changes at the next clean *Windows* start. Direct writes to the complex, not completely documented Registry file structure should not be made.

Other kinds of change, such as the deletion of additional program files (for example the Badtrans.B keylogger DLL), or the renaming of files, as would be necessary in order to clean the QAZ worm (delete the worm file notepad.exe and rename the backup copy note.com to notepad.exe), should be easy, too.

However, all these write operations will be problematic in the case of the complex internal structure of NTFS drives. Both *Linux* and the registered version of NTFSDOS have some problems while creating, copying and changing a large number of files on an NTFS drive. Usually, the file structure will be a little corrupted, causing *Windows* to display an error message or record something in its Event Log at the next startup, but the good news is that they are often recoverable.

Furthermore, there are usually only minor changes needed in the file for the cleaning process, meaning that the size of the file is unchanged, or only decreased. Therefore, it should be easy to heal the file without greater, possibly corruption-causing changes in the NTFS file structure. However, such an NTFS clean functionality should be tested carefully and disabled by default.

Another issue is related to the backup of the infected files before the clean process is carried out. This should not cause any problems on FAT drives (unless there is not sufficient space available on disk), but on NTFS such write operations can corrupt the drive. If (infected) backup copies of files need to be kept, it's a better idea to save them on a non-NTFS drive, such as a floppy or ZIP disk. Under *Linux*, it would also be possible to burn a CD using a simple command-line tool, but that is probably too complex a task.

Report Files

Every scan and disinfection process should produce a readable log file, in which all the relevant information is stored. This should include the date and time of the scan as well as the last update of the signature files and the main virus scan engine. One important task that is often forgotten is to record all the activities of the program during the cleaning process. Usually, the report includes information such as the name of the infected file and whether this file has been cleaned or deleted, but not what kind of Registry and other file changes have been made, which would make the clean process much more transparent.

The use of helper files should also be documented – for example the fact that a second restart of the computer will be necessary after the helper program is executed and has undone the malware changes. Finally, it should be possible to save the report file as ASCII text to disk or to print it out. Don't forget that *Linux* will use a simple line break only, and not two characters, like DOS or *Windows*.

Conclusion

Today's anti-virus rescue systems are too limited to be useful against today's file system and complex malware. It seems that these routines have been written once and not updated for a very long time.

However, with a little research, it should be a relatively easy task to help infected customers with a powerful, menu-driven rescue system which will scan and clean the local hard disks. This could be extended to include a few more emergency rescue programs, such as an undelete utility or a disk editor.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form without the prior written permission of the publishers.