

# Keyloggers: How they work and how to detect them (Part 1)

In February 2005, Joe Lopez, a businessman from Florida, filed a [suit](#) against Bank of America after unknown hackers stole \$90,000 from his Bank of America account. The money had been transferred to Latvia.

An investigation showed that Mr. Lopez's computer was infected with a malicious program, Backdoor.Coreflood, which records every keystroke and sends this information to malicious users via the Internet. This is how the hackers got hold of Joe Lopez's user name and password, since Mr. Lopez often used the Internet to manage his Bank of America account.

However the court did not rule in favor of the plaintiff, saying that Mr. Lopez had neglected to take basic precautions when managing his bank account on the Internet: a signature for the malicious code that was found on his system had been added to nearly all antivirus product databases back in 2003.

Joe Lopez's losses were caused by a combination of overall carelessness and an ordinary keylogging program.

## About Keyloggers

The term 'keylogger' itself is neutral, and the word describes the program's function. Most sources define a keylogger as a software program designed to secretly monitor and log all keystrokes. This definition is not altogether correct, since a keylogger doesn't have to be software – it can also be a device. Keylogging devices are much rarer than keylogging software, but it is important to keep their existence in mind when thinking about information security.

Legitimate programs may have a keylogging function which can be used to call certain program functions using "hotkeys," or to toggle between keyboard layouts (e.g. Keyboard Ninja). There is a lot of legitimate software which is designed to allow administrators to track what employees do throughout the day, or to allow users to track the activity of third parties on their computers. However, the ethical boundary between justified monitoring and espionage is a fine line. Legitimate software is often used deliberately to steal confidential user information such as passwords.

Most modern keyloggers are considered to be legitimate software or hardware and are sold on the open market. Developers and vendors offer a long list of cases in which it would be legal and appropriate to use keyloggers, including:

- Parental control: parents can track what their children do on the Internet, and can opt to be notified if there are any attempts to access websites containing adult or otherwise inappropriate content;

- Jealous spouses or partners can use a keylogger to track the actions of their better half on the Internet if they suspect them of “virtual cheating”;
- Company security: tracking the use of computers for non-work-related purposes, or the use of workstations after hours;
- Company security: using keyloggers to track the input of key words and phrases associated with commercial information which could damage the company (materially or otherwise) if disclosed;
- Other security (e.g. law enforcement): using keylogger records to analyze and track incidents linked to the use of personal computers;
- Other reasons.

However, the justifications listed above are more subjective than objective; the situations can all be resolved using other methods. Additionally, **any** legitimate keylogging program can still be used with malicious or criminal intent. Today, keyloggers are mainly used to steal user data relating to various online payment systems, and virus writers are constantly writing new keylogger Trojans for this very purpose.

Furthermore, many keyloggers hide themselves in the system (i.e. they have rootkit functionality), which makes them fully-fledged Trojan programs.

As such programs are extensively used by cyber criminals, detecting them is a priority for antivirus companies. Kaspersky Lab’s malware classification system has a dedicated category for malicious programs [with keylogging functionality](#): Trojan-Spy. Trojan-Spy programs, as the name suggests, track user activity, save the information to the user’s hard disk and then forward it to the author or ‘master’ of the Trojan. The information collected includes keystrokes and screen-shots, used in the theft of banking data to support online fraud.

## Why keyloggers are a threat

Unlike other types of malicious program, keyloggers present no threat to the system itself. Nevertheless, they can pose a serious threat to users, as they can be used to intercept passwords and other confidential information entered via the keyboard. As a result, cyber criminals can get PIN codes and account numbers for e-payment systems, passwords to online gaming accounts, email addresses, user names, email passwords etc.

Once a cyber criminal has got hold of confidential user data, s/he can easily transfer money from the user’s account or access the user’s online gaming account. Unfortunately access to confidential data can sometimes have consequences which are far more serious than an individual’s loss of a few dollars. Keyloggers can be used as tools in both industrial and political espionage, accessing data which may include proprietary commercial information and classified government material which could compromise the security of commercial and state-owned organizations (for example, by stealing private encryption keys).

Keyloggers, phishing and social engineering (see [‘Computers, Networks and Theft’](#)) are currently the main methods being used in cyber fraud. Users who are aware of security issues can easily protect themselves against phishing by ignoring phishing emails and by not entering any

personal information on suspicious websites. It is more difficult, however, for users to combat keyloggers; the only possible method is to use an appropriate security solution, as it's usually impossible for a user to tell that a keylogger has been installed on his/ her machine.

[According](#) to Cristine Hoepers, the manager of Brazil's Computer Emergency Response Team, which works under the aegis of the country's Internet Steering Committee, keyloggers have pushed phishing out of first place as the most-used method in the theft of confidential information. What's more, keyloggers are becoming more sophisticated – they track websites visited by the user and only log keystrokes entered on websites of particular interest to the cyber criminal.

In recent years, we have seen a considerable increase in the number of different kinds of malicious programs which have keylogging functionality. No Internet user is immune to cyber criminals, no matter where in the world s/he is located and no matter what organization s/he works for.

## How cyber criminals use keyloggers

One of the most publicized keylogging incidents recently was the [theft](#) of over \$1 million from client accounts at the major Scandinavian bank Nordea. In August 2006 Nordea clients started to receive emails, allegedly from the bank, suggesting that they install an antispyware product, which was supposedly attached to the message. When a user opened the file and downloaded it to his/ her computer, the machine would be infected with a well known Trojan called Haxdoor. This would be activated when the victim registered at Nordea's online service, and the Trojan would display an error notification with a request to re-enter the registration information. The keylogger incorporated in the Trojan would record data entered by the bank's clients, and later send this data to the cyber criminals' server. This was how cyber criminals were able to access client accounts, and transfer money from them. According to Haxdoor's author, the Trojan has also been used in attacks against Australian banks and many others.

On January 24, 2004 the notorious MyDoom worm caused a [major epidemic](#). MyDoom broke the record previously set by Sobig, provoking the largest epidemic in Internet history to date. The worm used social engineering methods and organized a DoS attack on [www.sco.com](#); the site was either unreachable or unstable for several months as a consequence. The worm left a Trojan on infected computers which was subsequently used to infect the victim machines with new modifications of the worm. The fact that MyDoom had a keylogging function to harvest credit card numbers was not widely publicized in the media.

In early 2005 the London police [prevented](#) a serious attempt to steal banking data. After attacking a banking system, the cyber criminals had planned to steal \$423 million from Sumitomo Mitsui's London-based offices. The main component of the Trojan used, which was created by the 32-year-old Yeron Bolondi, was a keylogger that allowed the criminals to track all the keystrokes entered when victims used the bank's client interface.

In May 2005 in London the Israeli police arrested a [married couple](#) who were charged with developing malicious programs that were used by some Israeli companies in **industrial**

**espionage.** The scale of the espionage was shocking: the companies named by the Israeli authorities in investigative reports included cellular providers like Cellcom and Pelephone, and satellite television provider YES. According to reports, the Trojan was used to access information relating to the PR agency Rani Rahav, whose clients included Partner Communications (Israel's second leading cellular services provider) and the HOT cable television group. The Mayer company, which imports Volvo and Honda cars to Israel, was suspected of committing industrial espionage against Champion Motors, which imports Audi and Volkswagen cars to the country. Ruth Brier-Haephtrati, who sold the keylogging Trojan that her husband Michael Haephtrati created, was sentenced to four years in jail, and Michael received a two-year sentence.

In February 2006, the Brazilian police [arrested 55 people](#) involved in spreading malicious programs which were used to steal user information and passwords to banking systems. The keyloggers were activated when the users visited their banks' websites, and secretly tracked and subsequently sent all data entered on these pages to cyber criminals. The total amount of money stolen from 200 client accounts at six of the country's banks totaled \$4.7million.

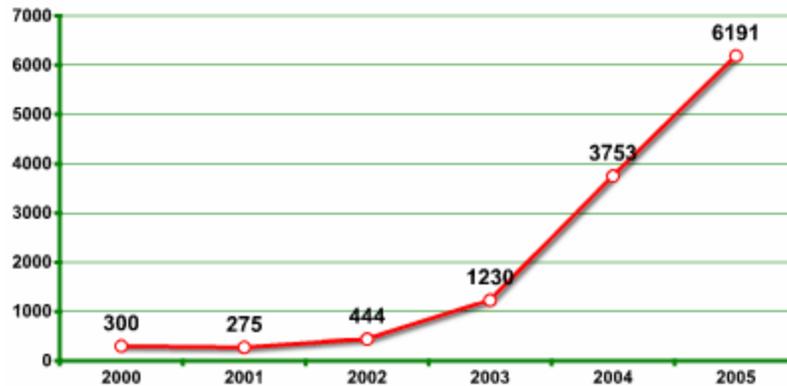
At approximately the same time, a similar criminal grouping made up of young (20 – 30 year old) Russians and Ukrainians was arrested. In late 2004, the group began sending banking clients in France and a number of other countries email messages that contained a malicious program – namely, a keylogger. Furthermore, these spy programs were placed on specially created websites; users were lured to these sites using classic social engineering methods. In the same way as in the cases described above, the program was activated when users visited their banks' websites, and the keylogger harvested all the information entered by the user and sent it to the cyber criminals. In the course of eleven months over one million dollars [was stolen](#).

There are many more examples of cyber criminals using keyloggers – most financial cybercrime is committed using keyloggers, since these programs are the most comprehensive and reliable tool for tracking electronic information.

## **Increased use of keyloggers by cyber criminals**

The fact that cyber criminals choose to use keyloggers time and again is confirmed by IT security companies.

One of [VeriSign's recent reports](#) notes that in recent years, the company has seen a rapid growth in the number of malicious programs that have keylogging functionality.



Source: iDefense, a VeriSign Company

One report issued by Symantec shows that almost 50% of malicious programs detected by the company's analysts during the past year do not pose a direct threat to computers, but instead are used by cyber criminals to harvest personal user data.

According to [research](#) conducted by John Bambenek, an analyst at the SANS Institute, approximately 10 million computers in the US alone are currently infected with a malicious program which has a keylogging function. Using these figures, together with the total number of American users of e-payment systems, possible losses are estimated to be \$24.3 million.

Kaspersky Lab is constantly detecting new malicious programs which have a keylogging function. One of the first virus alerts on [www.viruslist.com](http://www.viruslist.com), Kaspersky Lab's dedicated malware information site, was published on 15th June 2001. The warning related to TROJ\_LATINUS.SVR, a Trojan with a keylogging function. Since then, there has been a steady stream of new keyloggers and new modifications. Kaspersky antivirus database currently contain records for more than 300 families of keyloggers. This number does not include keyloggers that are part of complex threats (i.e. in which the spy component provides additional functionality).

Most modern malicious programs are hybrids which implement many different technologies. Due to this, any category of malicious program may include programs with keylogger (sub)functionality. The number of spy programs detected by Kaspersky Lab each month is on the increase, and most of these programs use keylogging technology.

## Keylogger construction

The main idea behind keyloggers is to get in between any two links in the chain of events between when a key is pressed and when information about that keystroke is displayed on the monitor. This can be achieved using video surveillance, a hardware bug in the keyboard, wiring or the computer itself, intercepting input/ output, substituting the keyboard driver, the filter driver in the keyboard stack, intercepting kernel functions by any means possible (substituting addresses in system tables, splicing function code, etc.), intercepting DLL functions in user mode, and, finally, requesting information from the keyboard using standard documented methods.

Experience shows that the more complex the approach, the less likely it is to be used in common Trojan programs and the more likely it is to be used in specially designed Trojan programs which are designed to steal financial data from a specific company.

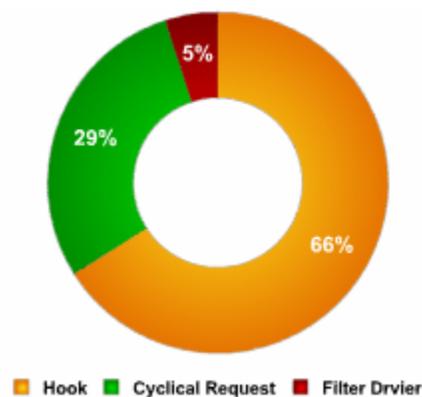
Keyloggers can be divided into two categories: keylogging devices and keylogging software. Keyloggers which fall into the first category are usually small devices that can be fixed to the keyboard, or placed within a cable or the computer itself. The keylogging software category is made up of dedicated programs designed to track and log keystrokes.

The most common methods used to construct keylogging software are as follows:

- a system hook which intercepts notification that a key has been pressed (installed using WinAPI SetWindowsHook for messages sent by the window procedure. It is most often written in C);
- a cyclical information keyboard request from the keyboard (using WinAPI Get(Async)KeyState or GetKeyboardState – most often written in Visual Basic, sometimes in Borland Delphi);
- using a filter driver (requires specialized knowledge and is written in C).

We will provide a detailed explanation of the different ways keyloggers are constructed in the second half of this article (to be published in the near future). But first, here are some statistics.

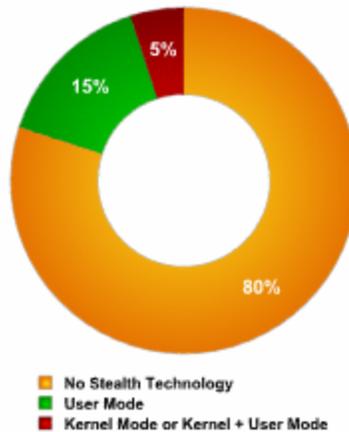
A rough breakdown of the different types of keyloggers is shown in the pie chart below:



Recently, keyloggers that disguise their files to keep them from being found manually or by an antivirus program have become more numerous. These stealth techniques are called rootkit technologies. There are two main rootkit technologies used by keyloggers:

- masking in user mode;
- masking in kernel mode.

A rough breakdown of the techniques used by keyloggers to mask their activity is shown in the pie chart below:



## How keyloggers spread

Keyloggers spread in much the same way that other malicious programs spread. Excluding cases where keyloggers are purchased and installed by a jealous spouse or partner, and the use of keyloggers by security services, keyloggers are mostly spread using the following methods):

- a keylogger can be installed when a user opens a file attached to an email;
- a keylogger can be installed when a file is launched from an open-access directory on a P2P network;
- a keylogger can be installed via a web page script which exploits a browser vulnerability. The program will automatically be launched when a user visits a infected site;
- a keylogger can be installed by another malicious program already present on the victim machine, if the program is capable of downloading and installing other malware to the system.

## How to protect yourself from keyloggers

Most antivirus companies have already added known keyloggers to their databases, making protecting against keyloggers no different from protecting against other types of malicious program: install an antivirus product and keep its database up to date. However, since most antivirus products classify keyloggers as *potentially malicious*, or *potentially undesirable programs*, users should ensure that their antivirus product will, with default settings, detect this type of malware. If not, then the product should be configured accordingly, to ensure protection against most common keyloggers.

Let's take a closer look at the methods that can be used to protect against unknown keyloggers or a keylogger designed to target a specific system.

Since the chief purpose of keyloggers is to get confidential data (bank card numbers, passwords, etc.), the most logical ways to protect against unknown keyloggers are as follows:

1. using one-time passwords or two-step authentication,

2. using a system with proactive protection designed to detect keylogging software,
3. using a virtual keyboard.

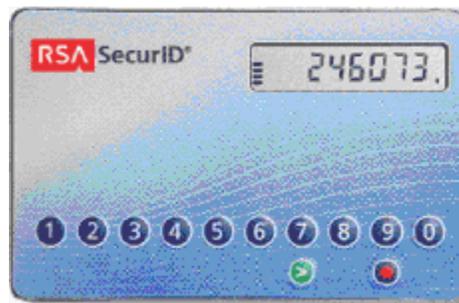
Using a one-time password can help minimize losses if the password you enter is intercepted, as the password generated can be used one time only, and the period of time during which the password can be used is limited. Even if a one-time password is intercepted, a cyber criminal will not be able to use it in order to obtain access to confidential information.

In order to get one-time passwords, you can use a special device such as:

1. a USB key (such as [Aladdin eToken NG OTP](#)):



2. a 'calculator' (such as [RSA SecurID 900 Signing Token](#)):



In order to generate one-time passwords, you can also use mobile phone text messaging systems that are registered with the banking system and receive a PIN-code as a reply. The PIN is then used together with the personal code for authentication.

If either of the above devices is used to generate passwords, the procedure is as described below:

1. the user connects to the Internet and opens a dialogue box where personal data should be entered;
2. the user then presses a button on the device to generate a one-time password, and a password will appear on the device's LCD display for 15 seconds;
3. the user enters his user name, personal PIN code and the generated one-time password in the dialogue box (usually the PIN code and the key are entered one after the other in a single pass code field);
4. the codes that are entered are verified by the server, and a decision is made whether or not the user may access confidential data.

When using a calculator device to generate a password, the user will enter his PIN code on the device 'keyboard' and press the ">" button.

One-time password generators are widely used by banking systems in Europe, Asia, the US and Australia. For example, Lloyds TSB, a leading bank, decided to [use](#) password generators back in November 2005.

In this case, however, the company has to spend a considerable amount of money as it had to acquire and distribute password generators to its clients, and develop/ purchase the accompanying software.

A more cost efficient solution is proactive protection on the client side, which can warn a user if an attempt is made to install or activate keylogging software.



Proactive protection against keyloggers in Kaspersky Internet Security

The main drawback of this method is that the user is actively involved and has to decide what action should be taken. If a user is not very technically experienced, s/he might make the wrong decision, resulting in a keylogger being allowed to bypass the antivirus solution. However, if developers minimize user involvement, then keyloggers will be able to evade detection due to an insufficiently rigorous security policy. However, if settings are too stringent, then other, useful programs which contain legitimate keylogging functions might also be blocked.

The final method which can be used to protect against both keylogging software and hardware is using a **virtual keyboard**. A virtual keyboard is a program that shows a keyboard on the screen, and the keys can be 'pressed' by using a mouse.

The idea of an on-screen keyboard is nothing new - the Windows operating system has a built-in on-screen keyboard that can be launched as follows: Start > Programs > Accessories > Accessibility > On-Screen Keyboard.



An example of the Windows on-screen keyboard

However, on-screen keyboards aren't a very popular method of outsmarting keyloggers. They were not designed to protect against cyber threats, but as an accessibility tool for disabled users. Information entered using an on-screen keyboard can easily be intercepted by a malicious program. In order to be used to protect against keyloggers, on-screen keyboards have to be specially designed in order to ensure that information entered or transmitted via the on-screen keyboard cannot be intercepted.

## Conclusions

This article has provided an overview of how keyloggers – both keylogging software and hardware - function and are used.

- Even though keylogger developers market their products as legitimate software, most keyloggers can be used to steal personal user data and in political and industrial espionage.
- At present, keyloggers – together with phishing and social engineering methods – are one of the most commonly used methods of cyber fraud.
- IT security companies have recorded a steady increase in the number of malicious programs that have keylogging functionality.
- Reports show that there is an increased tendency to use rootkit technologies in keylogging software, to help the keylogger evade manual detection and detection by antivirus solutions.
- Only dedicated protection can detect that a keylogger is being used for spy purposes.
- The following measures can be taken to protect against keyloggers:
  - use a standard antivirus that can be adjusted to detect potentially malicious software (default settings for many products);
  - proactive protection will protect the system against new ,modifications of existing keyloggers;
  - use a virtual keyboard or a system to generate one-time passwords to protect against keylogging software and hardware.

## Keyloggers: Implementing keyloggers in Windows. Part Two

This article is a continuation of the previous report on keyloggers. It offers a detailed analysis of the technical aspects and inner workings of keyloggers. As was noted in the first article, keyloggers are essentially designed to be injected between any two links in the chain whereby a signal is transmitted from a key being pressed to symbols appearing on the screen. This article provides both an overview of which links exist in this chain, and how both software and hardware keyloggers work.

This article is written for technical specialists and experienced users. Other users, who are not part of this target group, should simply be aware that Windows offers a multitude of ways in which data entered via the keyboard can be harvested, although the vast majority of keyloggers only use two of these methods (see: [Designing keyloggers, the first part of the article](#)).

It should be stressed that this article does not include any keylogger source code; we do not share the opinion of some researchers who believe it is acceptable to publish such code. Instead, we focus on understanding how keyloggers work, so we can better implement effective protection against them.

- [Processing data entered via the keyboard in Windows](#)
  - [The keyboard as a physical device: how it works](#)
  - [Low-level interaction with the keyboard via the input/ output port](#)
  - [The architecture of “interactive input devices”](#)
  - [Kernel mode drivers for PS/2 keyboards](#)
    - [Driver stack for system input devices](#)
    - [Driver stack for Plug and Play PS/2 keyboards](#)
    - [Device stack for Plug and Play PS/2 keyboards](#)
  - [Processing keyboard input via applications](#)
    - [Raw input thread \(data received from the driver\)](#)
    - [Processing of messages by a specific window](#)
    - [Keyboard key status array](#)
  - [Keyboard hooks](#)
  - [Processing](#)
  - [Raw input model](#)
- [Implementing keyloggers: the variants](#)
  - [1. User mode keyloggers](#)
    - [1.1. Setting hooks for keyboard messages](#)
    - [1.2. Using cyclical querying of the keyboard](#)
    - [1.3. Injection into processes and hooking message processing function](#)
    - [1.4. Using the raw input model](#)
  - [2. Kernel mode keyloggers](#)
    - [2.1. Using the keyboard driver filter Kbdclass](#)
    - [2.2. Using the filter driver of the i8042prt functional driver](#)
    - [2.3. Modifying the dispatch table of the Kbdclass driver](#)
    - [2.4. Modifying the system service table KeServiceDescriptorTableShadow](#)
    - [2.5. Modifying the code of the NtUserGerMessage or NtUserPeekMessage function by splicing](#)
    - [2.6. Substituting a driver in the keyboard stack of drivers](#)

- [2.7 Implementing a handler driver for interrupt 1 \(IRQ 1\)](#)
- [Conclusion](#)

## Processing data entered via the keyboard in Windows

There are several basic technologies which can be used to intercept keystrokes and mouse events, and many keyloggers use these technologies. However, before examining specific types of keylogger, it's necessary to understand how data entered via the keyboard is processed by Windows. To describe the process – from a key being pressed on the keyboard to the keyboard system interrupt controller being activated and an active WM\_KEYDOWN message appearing, three sources have been used:

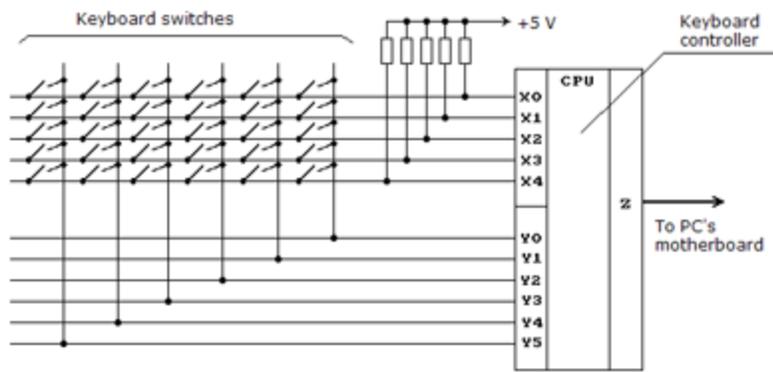
1. “Apparatnoe obeshpechenie IBM PC” by Alexander Frolov and Grigory Frolov, Volume 2, Book 1. Published by Dialog-MIFI, 1992. The second chapter, “The Keyboard” described how a keyboard functions, the ports used, and keyboard hardware interrupts;
2. The section related to “HID / Human Input Devices” of the MSDN library, which describes the low-level (driver) part of the process by which keyboard input is processed;
3. Jeffrey Richter's book ‘Creating effective Win32 applications for 64-bit Windows’. Chapter 27, “A model of hardware input and the local input condition” contains a description of the high level part of the process by which keyboard input is processed (in user mode).

### The keyboard as a physical device: how it works

Today, the majority of keyboards are a separate device connected to the computer via a port – most frequently PS/2 or USB. There are two micro-controllers which support the processing of keyboard input data; one is part of the motherboard, the other is within the keyboard itself. Consequently, it could be said that a PC keyboard is itself a small computer system. It uses an 8042 microcontroller which constantly scans keys being pressed on the keyboard independently of central CPU activity.

Each key on the keyboard has a specific number assigned to it; this is linked to keyboard matrix map and is not directly dependent on the value shown on the surface of the key itself. This number is called a “scan code” (the name highlights the fact that the computer scans the keyboard to search for keystrokes). Scan codes are random values which were selected by IBM back in the days when the company created the first keyboard for PCs. A scan code does not correspond to the ASCII code of a key; actually, a single key may correspond to several ASCII code values. A table of scan codes can be found in the twentieth chapter of “The Art of Assembly Language Programming”.

In actual fact, the keyboard generates two scan codes for each key: one for when the user presses the key, and another for when the user releases the key. The fact that there are two scan codes is important, as some keys only have a function when they are pressed and held (eg. Shift, Control or Alt).



**Fig. 1: The keyboard: a simplified scheme**

When the user presses a key on the keyboard, s/he closes an electrical circuit. As a result, when performing the next scan, the micro-controller detects that a specific key has been pressed, and sends the scan code of the key pressed to the central computer, together with an interrupt request. The same action is performed when the user releases the key that s/he has previously pressed.

The second micro-controller receives the scan code, converts it, makes it accessible on the input/output port 60h and then generates a central processor hardware interrupt. The handler which processes the interrupt can then get the scan code from the designated input/ output port.

It should be noted that the keyboard contains an internal 16 byte buffer which it uses to exchange data with the computer.

### **Low-level interaction with the keyboard via the input/ output port**

Interaction with the keyboard system controller takes place via the input/ output port 64h. Reading a byte from this port makes it possible to determine the status of the keyboard controller; writing a byte makes it possible to send a command to the controller.

Interaction with the micro-controller within the keyboard itself takes places via the input/ output ports 60h and 64h. The 0 and 1 bits in the status byte (port 64h in read mode) make it possible to control the interaction: before writing data to these ports, bit 1 of port 64h should be 0. When data is read-accessible from port 60h, bit 1 of port 64h is equal to 1. The keyboard on/ off bits in the command byte (port 64h in write mode) determine whether or not the keyboard is active, and whether the keyboard controller will call a system interrupt when the user presses a key.

Bytes written to port 60h are sent to the keyboard controller, while bytes written to port 64h are sent to the keyboard system controller. A list of permissible commands which can be sent to the keyboard controller can be found in, for example, “8042 Keyboard Controller IBM Technical Reference Manual” or in the twentieth chapter of “The Art of Assembly Language Programming”.

Bytes read from port 60h come from the keyboard. When reading, port 60h contains the scan code of the last key pressed, and in write mode this is used for extended management by the keyboard. When using port 60h in write mode, a program has the following additional options:

- Establishing a wait period before the keyboard goes into auto repeat mode
- Establishing the interval for generating scan codes in auto repeat mode
- Management of LEDs located on the outer surface of the keyboard – Scroll Lock, Num Lock, Caps Lock.

To summarize, in order to read data entered via the keyboard, one only has to be able to read the values of input/ output port 60h and 64h. However, user-level applications in Windows are unable to work with ports; this function is fulfilled by operating system drivers.

## **The architecture of “interactive input devices”**

So, what processes the hardware interrupts which are generated when data sent by the keyboard appears on port 60h? Clearly, it's done through the handler of the keyboard hardware interrupt IRQ 1. In Windows, this processing is conducted by the system driver i8042prt.sys. In contrast to MS DOS, when each system component was a law unto itself, as it were, in Windows all components are constructed in accordance with a clear architecture and work in accordance with strictly defined rules assigned by the program interface. So before examining i8042prt, let's take a look at the architecture of interactive input devices, within the confines of which all program components connected with the processing of keyboard (and 'mouse') input function.

In Windows, devices which are used to manage operations on the computer are called 'interactive input devices'. A keyboard is one such device, together with the mouse, joysticks, trackballs, game controllers, wheels, virtual reality helmets etc.

The architecture of interactive input devices is based on the USB Human Interface Device standard put forward by the USB Implementers Forum. However, this architecture is not limited to USB devices, and supports other input devices, such as Bluetooth keyboards, PS/2 keyboards and mice and devices connected to I/O port 201 (the gaming port).

Later in this article we will examine how the driver stack and the device stack for a PS/2 keyboard are constructed (given that historically the PS/2 keyboard was the first type of device described above). USB keyboards use a range of elements which were introduced during the development of program support for PS/2 keyboards).

## **Kernel mode drivers for PS/2 keyboards**

### **Driver stack for system input devices**

Regardless of how the keyboard is physically connected, keyboard drivers use keyboard class system drivers to process data. This happens regardless of the hardware used. Actually, these drivers are called class drivers because they support system requirements independent of the hardware requirements of a specific class of device.

The corresponding functional driver (port driver) supports the execution of input/ output operations in correlation with the device being used. In x86 Windows this is implemented in a single system keyboard driver (i8042) and mouse driver.

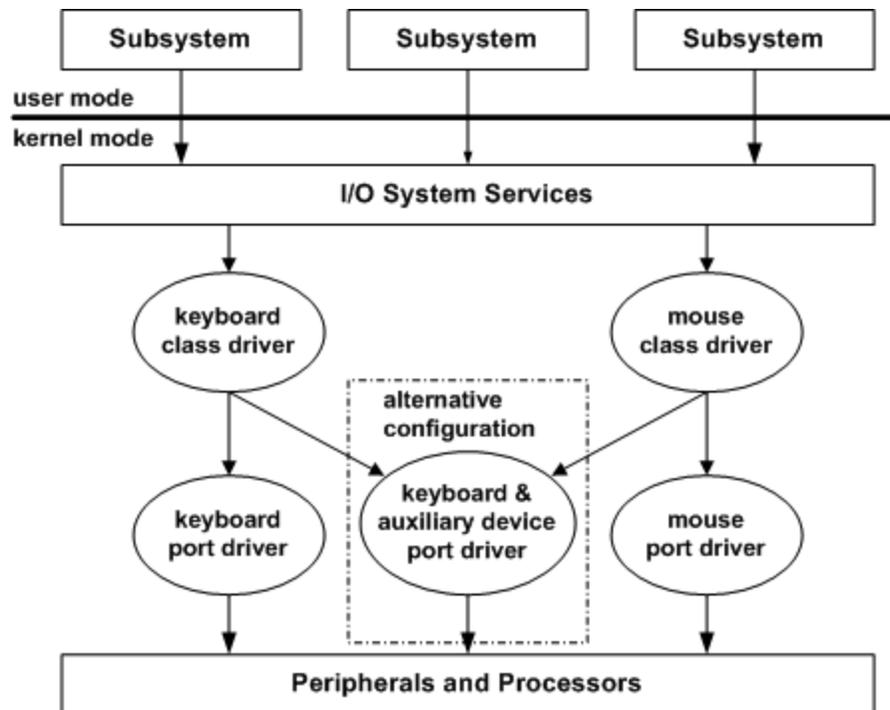


Fig 2.: Driver stack for system entry devices: keyboard and mouse

### Driver stack for Plug and Play PS/2 keyboards

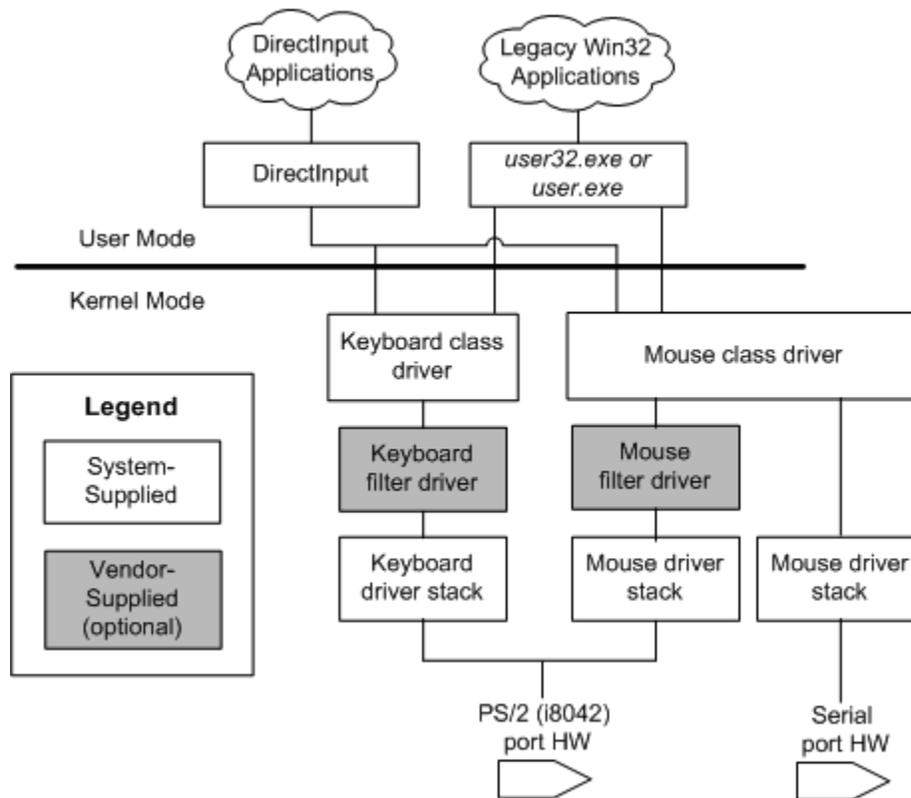


Fig 3. Driver stack for PS/2 keyboards

The driver stack (from top to bottom):

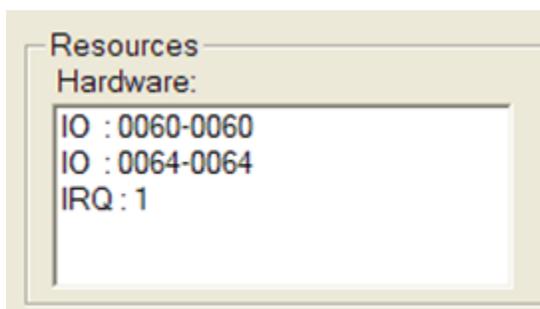
1. Kbdclass – high level filter driver, keyboard class;
2. optional high level filter driver, keyboard class;
3. i8042prt – functional keyboard driver
4. root bus driver

In Windows 2000 and older versions of Windows, the keyboard class driver is Kbdclass. The main tasks of this driver are:

- To support general and hardware-dependent operations of the device class
- To support Plug and Play, support power management and Windows Management Instrumentation (WMI)
- To support operations for legacy devices
- Simultaneous execution of operations from more than one device
- To implement the class service callback routine, which is called by the functional driver to transmit data from the device input buffer to the device driver data buffer.

In Windows 2000 and older versions of Windows, the functional driver for input devices using the PS/2 port (keyboard and mouse) is the i8042prt driver. The main functions are as follows:

- To support hardware dependent simultaneous operations of PS/2 input devices (the keyboard and mouse share the input and output ports, but use different interrupts, Interrupt Service Routines (ISR) and procedures for terminating interrupt processing;
- To supporting Plug and Play, power management and Windows Management Instrumentation (WMI);
- To supporting operations for legacy devices;
- To call the class service callback routine for classes of keyboards and mice in order to transmit data from the input data buffer i8042prt to the device driver data buffer;
- To call a range of callback functions which can be implemented in high level driver filters for flexible management by a device



**Fig. 4: Hardware resources which use the i8042prt driver**

Figure 4 shows a list of the hardware resources which use the i8042prt driver. These can be viewed, for example, using DeviceTree (<http://www.osronline.com/article.cfm?article=97>), a utility developed by Open Systems Resources. If you've read the sections on 'The keyboard as a physical entity – how it works' and 'Low-level interaction with the keyboard via the input/ output

ports" the values of the input/ output ports of 60h and 64h, and the hardware interrupt (IRQ) 1 will not come as any surprise.

A new **driver filter** can be added above the keyboard class driver in the driver stack shown above in order to, for instance, perform specific processing of data entered via the keyboard. This driver should support the same processing of all types of input/ output requests and management commands (IOCTL) as the keyboard class driver. In such cases, before data is transmitted to the user mode subsystem, the data is passed for processing to this driver filter.

### Device stack for Plug and Play PS/2 keyboards

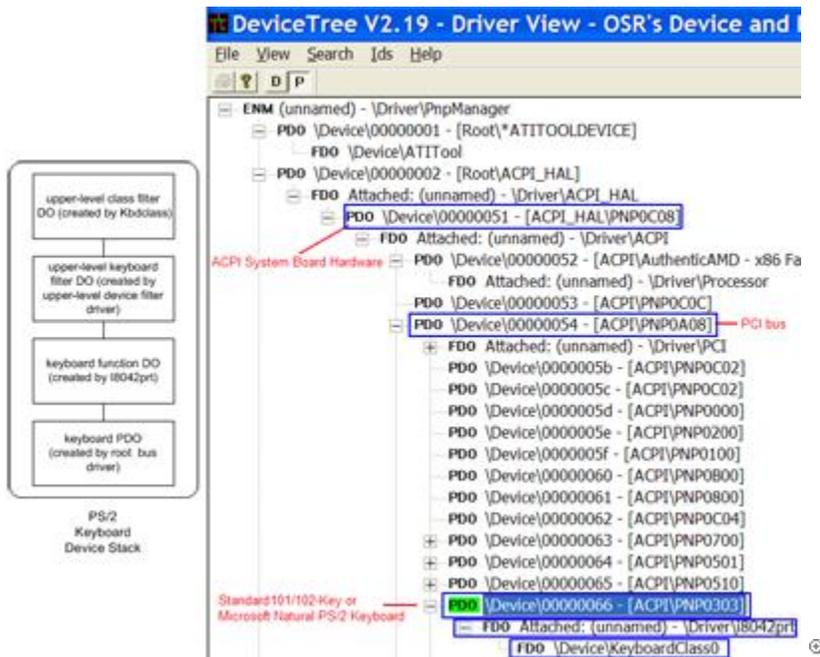


Fig. 5: Configuration of device objects for Plug and Play PS/2 keyboards.

Overall, the device stack (which more correctly should be called the device object stack) for a P/S2 keyboard is made up of:

1. The physical device object (PDO), created by the driver bus (in this case, the PCI bus) – \Device\00000066;
2. The functional device object (FDO), created and connected to the PDO by the i8042prt port – an unnamed object;
3. Optional filter objects for the keyboard device, created by the keyboard driver filters, which are developed by third party developers;
4. High level filter objects for the keyboard device class which are created by the Kbdclass class driver - \Device\KeyboardClass0.

### Processing keyboard input via applications

#### Raw input thread (data received from the driver)

The previous section gives examples of how keyboard stacks are constructed in kernel mode. This section takes a look at how data about keystrokes is transmitted by applications in user mode.

The subsystem of Microsoft Win32 gets access to the keyboard by using the Raw Input Thread (RIT), part of the csrss.exe system process. On boot, the system creates the IRT and the system hardware input queue (SHIQ).

The RIT opens the keyboard class device driver for exclusive use and uses the ZwReadFile function to send it an input/ output request (IRP) of the type IRP\_MJ\_READ. Having received the request, the Kbdclass driver flags it as pending, places it in the queue and returns a STATUS\_PENDING code. The RIT has to wait until the IRP terminates, and in order to determine this it uses the Asynchronous Procedure Call or ACP.

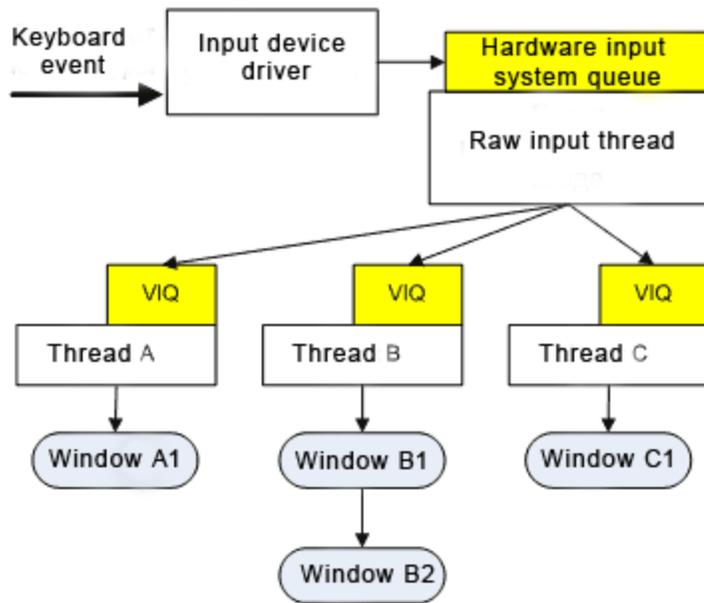
When the user presses or releases one of the keys, the keyboard system controller yields a hardware interrupt. The hardware interrupt processor calls a special procedure to process the IRQ 1 interrupt (the interrupt service routine, or ISR), which is registered in the system by the i8042prt driver. This procedure reads the data which has appeared from the internal keyboard controller queue. The processing of the hardware interrupt should be as quick as possible; because of this, the IRC places a Deferred Procedure Call (or DPC), i8042KeyboardIsrDpc and then terminates. As soon as is possible (the IRQL reverts to DISPATCH\_LEVEL), the DPC is called by the system. At this moment the callback procedure KeyboardClassServiceCallback will be called, which is registered by the i8042 driver Kbdclass driver.

KeyboardClassServiceCallback extracts the pending termination request (IRP) from its queue, completes the maximum amount of KEYBOARD\_INPUT\_DATA which provide all the information required about keys pressed and released and terminates the IRP. The raw input thread is activated again, processes the information received, and sends another IRP to the class driver, which will again be queued until the next key press/ release. This means that the keyboard stack always contains at least one pending termination request in the Kbdclass driver queue.

Time	Call/Comp	IRP Addr	Seq Number	Originating Device	Target Device	Major Function	Minor Function	Completion Status
01:58:49.515	Comp	0x89627060	2189	csrss.exe	Device\KeyboardClass	READ	NORMAL	SUCCESS, Info = Dec
01:58:49.515	NTAPI	0x89627060	0	csrss.exe	Device\KeyboardClass	READ	NORMAL	SUCCESS, Info = Dec
01:58:49.515	Call	0x89627060	0	csrss.exe	Device\KeyboardClass	READ	NORMAL	PENDING, Info = Dec
01:58:49.515	NTAPI	0x89627060	0	csrss.exe	Device\KeyboardClass	READ	NORMAL	PENDING, Info = Dec
01:58:49.515	Call	0x89627060	0	csrss.exe	Device\KeyboardClass	READ	NORMAL	PENDING, Info = Dec
01:58:49.515	NTAPI	0x89627060	0	csrss.exe	Device\KeyboardClass	READ	NORMAL	PENDING, Info = Dec
01:58:49.515	Call	0x89627060	0	csrss.exe	Device\KeyboardClass	READ	NORMAL	PENDING, Info = Dec
01:58:49.515	NTAPI	0x89627060	0	csrss.exe	Device\KeyboardClass	READ	NORMAL	PENDING, Info = Dec

**Fig. 6: Sequence of requests from RIT to the keyboard driver**

Using a utility called IrpTracker, developed by the previously mentioned Open Systems Resources, it's possible to track the sequence of calls which takes place when keyboard input is processed.



**Fig. 7: Processing keyboard input in user mode**

How does RIT process incoming data? All incoming keyboard events are placed in the hardware input system queue, and are in turn transformed into Windows messages (e.g. WM\_KEY\*, WM\_?BUTTON\* or WM\_MOUSEMOVE) and are then placed at the end of the virtualized input queue, or VIQ of the active thread. The key scan codes in the Windows messages are replaced by virtual key codes which correspond not to the location of keys on the keyboard but the action that this key performs (function that it fulfils). The mechanism for transforming codes depends on the current (active) keyboard layout, simultaneous pressing of keys (e.g. SHIFT) and other factors.

When the user enters the system, the Windows Explorer process launches a thread which creates the task panel and the desktop (WinSta0\_RIT). This thread binds to the RIT. If the user launches MS Word, then the MS WORD thread, having created a window, will immediately connect to the RIT. The Explorer process will then unhook from the RIT, as only one thread can be connected to RIT at any one time. When a key is pressed, the relevant element will appear in the SHIQ; this leads to the RIT becoming active, transforming the hardware input event into a message from the keyboard which will then be placed in the MS Word VIQ thread.

### Processing of messages by a specific window

How does a thread process messages from the keyboard which have entered the thread's message queue?

The standard message processing cycle usually looks like this:

```

while(GetMessage(&msg, 0, 0, 0))
{
  TranslateMessage(&msg);
}
  
```

```
DispatchMessage(&msg);  
}
```

Using the GetMessage function, keyboard events are extracted from the queue and redirected using the DispatchMessage function to the window procedure which processes messages for the window where input is currently focussed. The input focus is an attribute which can be assigned to a window created by an application or by Windows. If the window has an input focus, all keyboard messages from the system queue will reach the appropriate function of this window. An application can pass the input focus from one window to another, e.g. when another application is switched to using Alt+Tab.

The TranslateMessage function is usually called prior to the DispatchMessage function. This function creates the 'symbolic' messages WM\_CHAR, WM\_SYSCHAR, WM\_DEADCHAR and WM\_SYSDEADCHAR using the WM\_KEYDOWN, WM\_KEYUP, WM\_SYSKEYDOWN, WM\_SYSKEYUP messages as a base. These 'symbolic' messages are placed in the application message queue; however, it should be noted that the original keyboard messages are not deleted from this queue.

### **Keyboard key status array**

One of the aims when developing the Windows hardware input model was to ensure resilience. Resilience is ensured by independent processing of input by threads; this prevents conflicts between threads. However, this is not enough to isolate threads from each other, and the system therefore supports an additional concept: local input status. Each thread has its own input condition, and information about this is stored in THREADINFO. The information includes data about the virtual queue thread, and a group of variables. This last contains management information about the input thread status. The following notifications are supported for the keyboard: which window is currently in the focus of the keyboard, which window is currently active, which keys are pressed and the status of the input cursor.

Information about which keys are being pressed is saved to the synchronous status array of keys. This array is connected to the variables for the local input status of each thread. The array of asynchronous key status, which contains similar information, is shared by all threads. The arrays reflect the status of all keys at a given moment, and the GetAsyncKeyState function makes it possible to determine whether or not a specific key is being pressed at a given time. GetAsyncKeyState always returns 0 (i.e. not pressed) if it is called by a different thread (i.e. not the thread which created the window which is currently the focus of input status).

The GetKeyState function differs from GetAsyncKeyState in that it returns the status of the keyboard at the moment when the most recent keyboard message is extracted from the thread queue. This function can be called at any time regardless of which window is currently in focus.

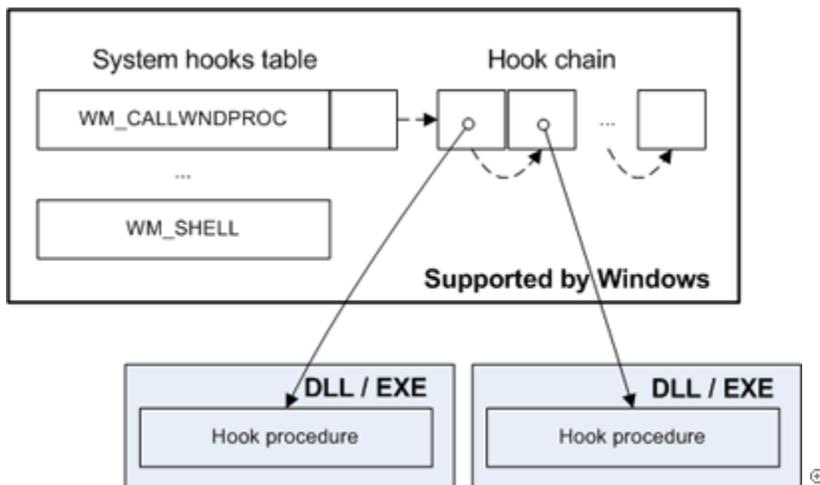
### **Keyboard hooks**

The mechanism used to intercept events using specific functions (e.g. sending Windows messages, data input via the mouse or keyboard) in Microsoft Windows is called 'hooking'. This function can react to an event and, in certain cases, modify or delete events.

Functions which receive notification of events are called filter functions; they differ from each other by which events they can intercept. In order for Windows to call a filter function, the function must be bound to a hook (for instance, to a keyboard hook). Binding one or more filter functions to a hook is called “setting a hook”. An application can use the Win32 API SetWindowsHookEx and UnhookWindowsHookEx functions to set or remove filter functions. Some hooks can either be set across the system as a whole, or for a specific thread.

To avoid conflicts if several filter functions are bound to a single hook, Windows implements a function queue; in such cases, the function most recently bound to the hook will be at the start of the queue, with the first function bound being at the end of the queue. The function filter queue (see figure 8) is managed by Windows itself, which simplifies the writing of filter functions and optimizes operating system productivity.

The system also supports separate chains for each type of hook. A hook chain is a list of pointers to filter functions (specific callback functions determined by the application.) When an event linked to a particular type of hook takes place, the system consecutively sends the message for each type of filter function to the hook chain. The actions which filter functions may perform depend on the type of hook: some function can only track the appearance of an event, while others may modify message parameters or initiate message processing, by preventing the next filter function in the hook chain from being called, or the message processing function for the relevant window from being called.



**Figure 8. Filter function chain in Windows**

When one or more filter functions are bound to a hook and an event takes place which leads to the hook being activated, Windows calls the first function from the filter function queue. With this, its responsibility is over. The filter function is then responsible for calling the next function in the chain, and the Win32 API CallNextHookEx function is used to do this.

The operating system supports several types of hooks, each of which provides access to one aspect of the Windows messaging process mechanism.

Nearly all types of hooks are of potential interest to the creators of keyloggers:

WH\_KEYBOARD, WH\_KEYBOARD\_LL (hooking keyboard events when they are added to the thread event queue), WH\_JOURNALRECORD and WH\_JOURNALPLAYBACK (writing and producing keyboard and mouse events), WH\_CBT (intercepting multiple events, including remote keyboard events from the system hardware input queue), WH\_GETMESSAGE (intercepting an event from the thread event queue.)

## Processing

Let's sum up all the information above on the procedure of keyboard input in a single algorithm: the algorithm of the passing of a signal from a key being pressed by the user to the appearance of symbols on the screen can be presented as follows:

1. When starting, the operating system creates a raw input thread and a system hardware input queue in the csrss.exe process.
2. The raw input thread cyclically sends read requests to the keyboard driver, which remains in a waiting condition until an event from the keyboard appears.
3. When the user presses or releases a key on the keyboard, the keyboard micro-controller detects that a key has been pressed or released and sends both the scan code of the pressed/ released key to the central computer and an interrupt request.
4. The keyboard system controller gets the scan code, processes it then makes it accessible on input/output port 60h and generates a central processor hardware interrupt.
5. The interrupt controller signals the CPU to call the interrupt processing procedure for IRQ1 – ISR, which is registered in the system by the functional keyboard driver i8042prt.
6. The ISR reads the data which has appeared from the internal keyboard controller queue, transforms the scan codes to virtual key codes (independent values which are determined by the system) and queues “I804KeyboardlsrDPC”, a delayed procedure call.
7. As soon as possible, the system calls the DPC which in turn executes the callback procedure KeyboardClassServiceCallback registered by the Kbdclass keyboard driver.
8. The KeyboardClassServiceCallback procedure extracts a pending termination request from the raw input thread from its queue and returns it with information about the key pressed.
9. The raw input thread saves the information to the system hardware input queue and uses it to create the basic Windows keyboard messages WM\_KEYDOWN, WM\_KEYUP, which are placed at the end of the VIQ virtual input queue of the active thread.
10. The message processing cycle thread deletes the message from the queue and sends the corresponding window procedure for processing. When this happens, the system function TranslateMessage may be called, which uses basic keyboard messages to create the additional “symbol” messages WM\_CHAR, WM\_SYSCHAR, WM\_DEADCHAR and WM\_SYSDEADCHAR.

## Raw input model

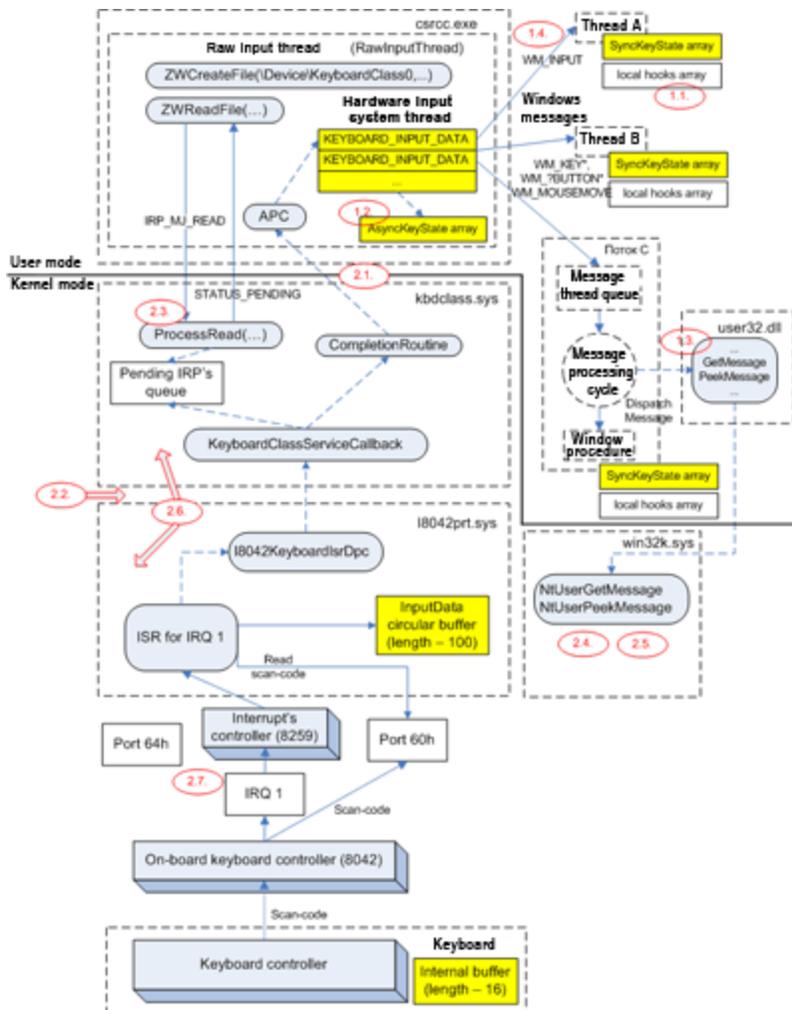
The keyboard input model described above has a number of shortcomings from the point of view of those who develop applications. In order to get input from a non-standard input device, the application has to perform a significant number of operations: mount the device, periodically query the device, consider the possibility of parallel use of the device by other applications etc. For this reason, later versions of Windows offer an alternative input model, called the raw input model, which simplifies the development of applications which use non-standard input devices (see Fig 3, DirectInput applications)

The raw input model differs from the original Windows input model. In the latter, the application gets device independent input in the form of a message (such as (WM\_CHAR), which is sent to application windows. In the raw input model, the application has to register the device it wants to receive data from. The application then receives user input via WM\_INPUT messages. Two methods of data transmission are supported: a standard method and a buffering method. In order to interpret entered data, the application has to get information about the input device, which can be done using the GetRawInputDeviceInfo function.

## **Implementing keyloggers: the variants**

Now we take a look at the main methods used by malware authors for implementing keyloggers, taking the model of processing keyboard input in Windows described above as a basis. Knowing how the model is structured makes it easier to understand which mechanisms can be used by keyloggers.

Keyloggers can be injected at any stage in the processing sequence, intercepting data about keys pressed which is transmitted by one processing subsystem to another subsystem. The methods examined below for creating software keyloggers are divided into user mode methods and kernel mode methods. Figure 9 shows all the subsystems processing keyboard input and their interdependencies. Next to some of the subsystems are numbers in red circles, and these indicate the section of the article which describes how keyloggers which use or substitute the corresponding subsystem can be created.



**Fig 9: Overview of how Windows processes keyboard input**

This section will not look at creating hardware keyloggers. It is enough to note that there are three types of hardware keylogger: keyloggers incorporated into the keyboard itself; keyloggers incorporated into the cable connecting the keyboard to the system block; keyloggers incorporated into the computer's system block itself. The most common of these is the second type of hardware keylogger, and one of the best known examples of this type is the "KeyGhost USB Keylogger" (<http://www.keyghost.com>).

Unfortunately, at the moment only a few antivirus solutions offer adequate protection against the potential threat posed by keyloggers. Such products are for instance Kaspersky Lab 6.0 products and newer.

## 1. User mode keyloggers

User mode keyloggers are the easiest to create, but also the easiest to detect, as they use well known and well documented Win32 API functions to intercept data.

### 1.1. Setting hooks for keyboard messages

This is the most common method used when creating keyloggers. Using SetWindowsHookEx, the keylogger sets a global hook for all keyboard events for all threads in the system (see the section on 'Keyboard hooks'). The hook filter function is located in a separate dynamic library which will be injected into all processes in the system. When any keyboard message thread is extracted from the message queue, the system calls the filter function installed.

One of the advantages of this method of interception is its simplicity and the guaranteed interception of all pressed keys. As for disadvantages, it should be noted that it's necessary to have a separate dynamic library file, and it is also relatively easy to detect the keylogger due to the fact that it has been injected into all processes.

Some keyloggers which use this approach are AdvancedKeyLogger, KeyGhost, Absolute Keylogger, Actual Keylogger, Actual Spy, Family Key Logger, GHOST SPY, Haxdoor, MyDoom and others.

Kaspersky Internet Security proactively detects this type of Keylogger as 'invader (loader)'; the option 'Windows hooks' in the 'Application activity analyzer' subsystem in the PDM module of KIS should be enabled.

## **1.2. Using cyclical querying of the keyboard**

This is the second most common method used when implementing keyloggers. The GetAsyncKeyState and GetKeyState are used to periodically query the status of all keys at rapid intervals. This function returns arrays of asynchronous or synchronous key status (see section Keyboard key status array); by analysing these it's possible to understand which keys have been pressed or released since the last query was carried out.

The advantages of this method are that it is extremely easy to implement, there is no additional module (in contrast to the use of hooks); the disadvantages are that there is no guarantee that all keystrokes will be intercepted – some may be missed; and this method can easily be detected by looking for processes which query the keyboard with high frequency.

This method is used in keyloggers such as Computer Monitor, Keyboard Guardian, PC Activity Monitor Pro, Power Spy, Powered Keylogger, Spytecor, Stealth KeyLogger, Total Spy.

Kaspersky Internet Security proactively detects this type of Keylogger as 'Keylogger'; the option 'Keylogger detection' in the 'Application activity analyzer' subsystem in the PDM module should be enabled.

## **1.3. Injection into processes and hooking message processing functions**

This method is used infrequently. The keylogger is injected into all processes and intercepts the GetMessage or PeekMessage functions from the user32.dll library. A range of methods can be used to do this: splicing (a method used to intercept API function calls; essentially, the method consists of substituting a few - normally five - of the first bytes of the JMP instruction function which passes control to the intercepting code), modifying the address table of IAT import

functions, intercepting the GetProcAddress function which returns a function address from the library which has been loaded. A keylogger can be created either in the form of a DLL or by injecting the code directly into the target process.

The result is as follows: when the application calls, for example, the GetMessage function in order to get the next message from the message queue, this call will be passed to the hook code. The hook calls the outgoing GetMessage function from user32.dll and analyses the results returned for message type objects. When a keyboard message is received, all information about it is extracted from the message parameters and logged by the keylogger.

The advantages of this method are its effectiveness: due to the method being not very common, only a few programs are able to detect keyloggers of this type. Additionally, standard onscreen keyboards are useless against this type of keylogger, as the messages sent by them will also be intercepted.

The disadvantages include the fact that modifying the IAT table does not guarantee interception, as the function address of user32.dll may have been saved prior to the keylogger being injected; splicing has certain difficulties connected with, for example, the need to rewrite the body of the function on the fly.

Kaspersky Internet Security proactively detects this type of Keylogger as 'invader'; the option 'Intrusion into process (invaders)' in the 'Application activity analyzer' subsystem in the PDM module should be enabled.

#### **1.4. Using the raw input model**

At the time of writing, there is only one known example of this method: a utility used to test how well protected the system is against keyloggers. Essentially the method uses the new Raw Input module (see the section "Raw Input model"), where the keylogger registers the keyboard as a device from which it wants to receive input. The keylogger then starts to get data about keys which have been pressed and released using the WM\_INPUT message.

The advantage is that this method is easy and effective to implement; in view of the fact that this method became well known only recently, almost none of today's protection software currently has the ability to combat such keyloggers.

The disadvantage is that fact that such keyloggers are easy to detect due to the need for special registration in order for the application to get raw input (by default, no system processes do this).

Kaspersky Internet Security 7.0 will proactively detect such keyloggers as 'Keylogger'; the option 'Keylogger detection' in the 'Application activity analyzer' subsystem in the PDM module should be enabled.

## **2. Kernel mode keyloggers**

In terms of both implementation and detection, kernel mode keyloggers are significantly more complex than user mode keyloggers. In order to write such keyloggers, higher knowledge is required, but this makes it possible to create keyloggers which will be completely unnoticed by all user mode applications. Both methods documented in the Microsoft Driver Development Kit and undocumented methods may be used for interception.

### **2.1. Using the keyboard driver filter Kbdclass**

The hook method is documented in the DDK (see the section "Driver stack for Plug and Play PS/2 keyboard"). Keyloggers which use this approach intercept requests to the keyboard by installing a filter above the device "\Device\KeyboardClass0", created by the Kbdclass driver (see Driver stack for Plug and Play PS/2 keyboards). Only IRP\_MJ\_READ requests need to be filtered as it is these which make it possible to get the codes of keys which have been pressed and released.

The advantage is the guarantee that all keystrokes will be intercepted; such keyloggers cannot be detected without using a driver, and due to this, not all anti-keyloggers will detect them. The disadvantage is that the driver itself has to be installed.

The best known keyloggers which use this approach are ELITE Keylogger and Invisible KeyLogger Stealth.

Kaspersky Internet Security proactively detects such keyloggers as Keylogger by monitoring the keyboard device stack (the "Detect keyboard intercepts" option in the "Application activity analysis" option in the proactive detection module (PDM) should be enabled).

### **2.2. Using the filter driver of the i8042prt functional driver**

This method of interception is documented in the DDK. Keyloggers based on this method intercept requests to the keyboard by installing a filter on top of an unnamed device, created by the i8042prt driver for the Device\KeyboardClass0 (see section "Device stack for Plug and Play PS/2 keyboards"). The i8042prt driver is a program interface for adding an additional IRQ1 (IsrRoutine) interrupt processing function, within which data received from the keyboard can be analysed.

The advantages and disadvantages of this method are the same as those detailed in the previous point. However, there is an additional disadvantage – it is dependent on the type of keyboard. The i8042prt driver processes requests only from PS/2 keyboards, and because of this, this method is not suitable for intercepting data entered via a USB or wireless keyboard.

Kaspersky Internet Security proactively detects such keyloggers as Keylogger by monitoring the keyboard device stack (the "Detect keyboard intercepts" option in the "Application activity analysis" option in the proactive detection module (PDM) should be enabled).

### **2.3. Modifying the dispatch table of the Kbdclass driver**

Keyloggers based on this principle intercept requests to the keyboard by changing the IRP\_MJ\_READ entry point in the dispatch table for the Kbdclass driver. This is functionally similar to the Kbdclass driver filter (see section “2.1 Driver filters for kbdclass driver”). The specifics are the same. Another variant hooks a different function for processing requests: IRP\_MJ\_DEVICECONTROL. In this case, the keylogger becomes a driver filter similar to the i8042 driver (see section “2.2 Using the filter driver of the i8042prt functional driver”).

Kaspersky Internet Security proactively detects such keyloggers as Keylogger by monitoring the keyboard stack devices (the “Detect keyboard interception” in the “Application activity analysis” option in the proactive detection module (PDM) should be enabled.)

#### **2.4. Modifying the system service table KeServiceDescriptorTableShadow**

This is a relatively common method for implementing keyloggers, which is similar to the user mode method described in section “1.3 Injection into processes and hooking message processing functions”. Keyloggers which use this method intercept requests to the keyboard by patching the entry point for NtUserGetMessage/PeekMessage in the second table of system services (KeServiceDescriptorTableShadow of the win32k.sys driver. Information about keys pressed is transmitted to the keylogger when a call to the GetMessage or PeekMessage function is terminated within a thread.

The advantage of this method is that it is difficult to detect. The disadvantage is that it is difficult to implement (searching for KeServiceDescriptorTableShadow is not an easy task; in addition to this, other drivers may already have patched the entry point in this table) and the need to install a separate driver.

Kaspersky Internet security proactively detects such keyloggers as Keylogger by monitoring the keyboard stack devices (the “Detect keyboard interception” in the “Application activity analysis” option in the proactive detection module (PDM) should be enabled.)

#### **2.5. Modifying the code of the NtUserGerMessage or NtUserPeekMessage function by splicing**

This is an extremely rare type of keylogger. Keyloggers which use this method intercept requests to the keyboard by modifying the code of the NtUserGetMessage or NtUserPeekMessage using splicing. These functions are implemented in the system kernel through the win32k.sys driver and are called by corresponding functions in the user32.dll library. As shown in section 1.3 “Injection into process and hooking message processing functions”, these functions make it possible to filter all messages received by applications and get data about the pressing/ releasing of keys from keyboard messages.

The advantage of this method is that it is difficult to detect. The disadvantage is that it is difficult to implement (it’s necessary to rewrite the body of the function on the fly, and is dependent on the operating system version and software installed) and it’s also necessary to install a driver.

#### **2.6. Substituting a driver in the keyboard stack of drivers**

This method involves substituting a Kbdclass driver or a low-level keyboard driver with a driver expressly developed for the purpose. The clear disadvantage of this method is that it is difficult to implement, as it is impossible to know in advance which type of keyboard is being used. It is therefore relatively easy to detect driver substitution. As a consequence, this method is almost never used.

## **2.7. Implementing a handler driver for interrupt 1 (IRQ 1).**

This involves writing a kernel mode driver which will hook the keyboard interrupt (IRQ1) and which directly contacts the keyboard input/ output ports (60h, 64h). As this is difficult to implement, and due to the fact that it is not entirely clear how it interacts with the system processor for interrupt IRQ1 (i8042prt.sys), this method remains, at the moment, purely theoretical.

## **Conclusion**

This article surveys the progression of data from a key being pressed by the user to the appearance of symbols on the screen; the links in the chain of processing the signal; and methods for implementing keyloggers which intercept keyboard input at specific stages of the process.

1. The process of keyboard input in Windows is relatively complex, and it's possible to install a hook at any stage. Keyloggers have already been created for some of the stages, while for some they do not yet exist.
2. There is a connection between how common a keylogger is and the difficulty of creating such a keylogger. For instance, the more common methods of interception – such as hooking input events and cyclical querying of the keyboard – are the simplest to implement. Even a person who has only been programming for a week would be able to write a keylogger which uses these methods.
3. The majority of keyloggers currently in existence are relatively simple and can easily be used with malicious intent, with the main aim being to steal confidential information entered via the keyboard.
4. Most important of all: Antivirus companies should protect their users against the threat posed by malicious keyloggers.

As protection mechanisms become more sophisticated, the cybercriminals who create keyloggers will be forced to implement more complex methods using Windows kernel drivers – there are still many unexploited possibilities in this area.