

The Windows Concept Journey

version 2.0
October-2023

By Dr. Shlomi Boutnaru



Table of Contents

Table of Contents	2
Introduction	4
Win32 API - Working with Strings	5
Window Messages	6
Recovery Directory	7
COM (Component Object Model)	8
Windows Services	10
What IPC (Inter Process Communication) mechanisms do we have in Windows?	12
Tasks (Windows Scheduler)	13
Objects	14
Object Manager	15
Clipboard	16
Handles	17

Introduction

When starting to learn Windows I believe that they are basic concepts that everyone needs to know about. Because of that I have decided to write a series of short writeups aimed at providing the basic vocabulary and understanding for achieving that.

Overall, I wanted to create something that will improve the overall knowledge of Windows in writeups that can be read in 1-3 mins. I hope you are going to enjoy the ride.

Lastly, you can follow me on twitter - @boutnaru (<https://twitter.com/boutnaru>). Also, you can read my other writeups on medium - <https://medium.com/@boutnaru>.

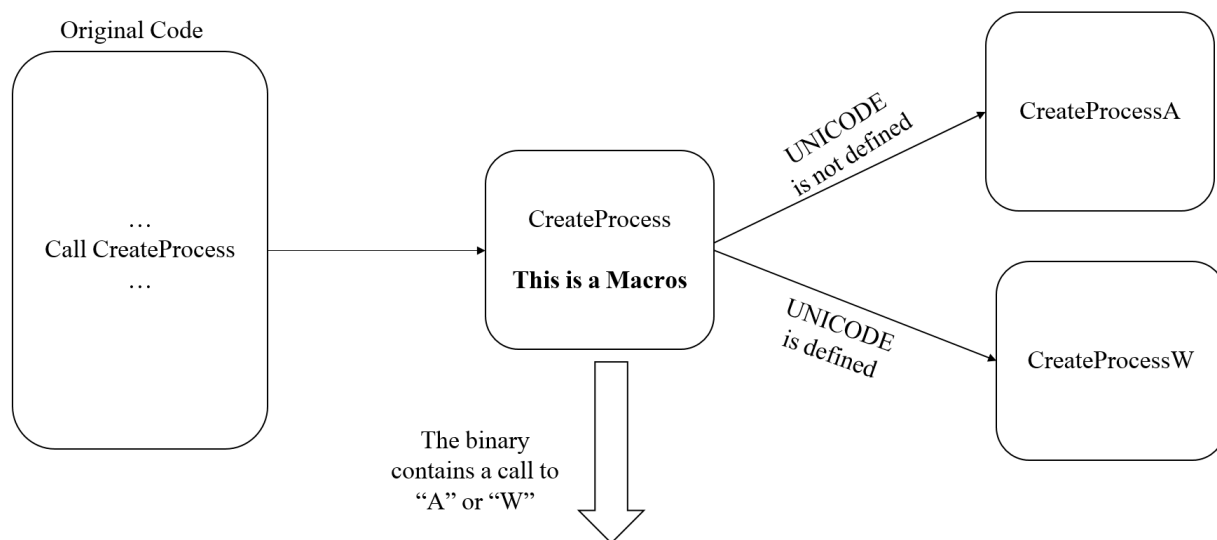
Lets GO!!!!!!

Win32 API - Working with Strings

After Microsoft added support for Unicode as part of Windows, it still needed to support the usage of ANSI strings. The way Microsoft decided to do it is by providing two sets of API (one for ANSI and the second for Unicode). It is important to know that the ANSI version of the API converts the strings to Unicode before calling the relevant syscalls (the kernel is Unicode only).

For example, in order to create a new process we can use “CreateProcessA” or “CreateProcessW” (A for ANSI strings and W for wide char aka Unicode). You might remember that you called CreateProcess and none of the above - so how did it work? The trick is using macros. You called a macro that checked if the UNICODE was defined. If it was defined it made a call to the function ending with “W” else it called the one ending with “A” - see the illustration below for the entire flow.

By the way, in the documentation (like MSDN) the functions names are without the suffix (“A” or “W”) despite the fact it is the name of the macro and not of the functions themselves. After compilation the executable contains a reference/dependency to a specific function (you can see it in the diagram below - I have extracted the strings from different DLLs showing the symbol names).



```
C:\Windows\System32>C:\tools\strings64.exe *.dll | findstr /i CreateProcess
C:\Windows\System32\aadtb.dll: DesktopAPI::CallCreateProcess
C:\Windows\System32\aadtb.dll: CreateProcessW
C:\Windows\System32\accessibilitycp1.dll: CreateProcessW
```

Window Messages

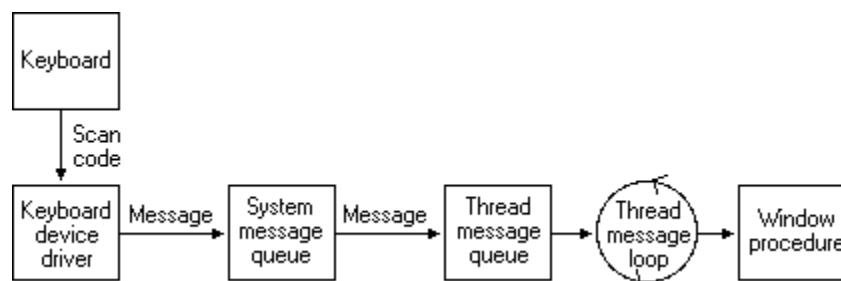
GUI applications under Windows react to different events from the user and the operating system itself. In case of user events, think about: keys pressed, touch-screen gestures, mouse clicks and more. In the case of OS events, think about plugging a new hardware device or changing a power-state¹ (like hibernation or sleep).

Overall, Windows communicates with windows created by applications using messages (aka “Window Messages”). In essence, a message is a number that defines a specific event sent to a window. Every window has a “Window Procedure”, which is a function that processes all messages sent to windows of the same class².

By the way, from the perspective of the Windows kernel the following classes can be created using the “CreateWindow” API call: BUTTON, COMBOBOX, EDIT, LISTBOX, MDICLIENT, SCROLLBAR and STATIC³.

Moreover, due to the fact an application can receive many messages and it can have several windows (each with its own window procedure) a loop to retrieve the message is needed. It is called “The Message Loop” which dispatches the message to the correct window⁴.

To summarize, in a Windows GUI application, “The Message Loop” is a continuous loop that retrieves messages from the operating system and dispatches them to the appropriate “Window Procedure”. It is responsible for handling all of the user input (and other events) that occur in the application - as shown in the diagram below⁵. In order to pull a message from the queue we can call “GetMessage”⁶.



¹ <https://learn.microsoft.com/en-us/windows/win32/learnwin32/window-messages>

² <https://learn.microsoft.com/en-us/windows/win32/winmsg/window-procedures>

³ <http://winapi.freetechnet.com/win32/WIN32CreateWindow.htm>

⁴ <https://learn.microsoft.com/en-us/windows/win32/learnwin32/window-messages>

⁵ <https://learn.microsoft.com/en-us/windows/win32/inputdev/about-keyboard-input>

⁶ <https://learn.microsoft.com/en-us/windows/win32/api/winuser/nf-winuser-getmessage>

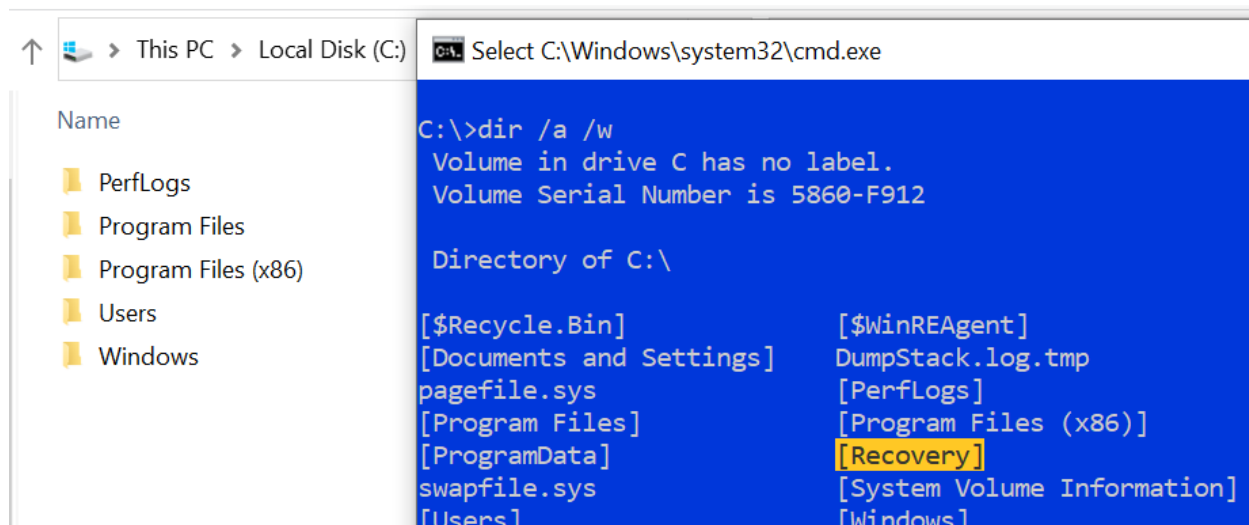
Recovery Directory

A bunch of folks have asked me about what is the goal of different directories in a Windows filesystem hierarchy. So I have decided to write a short series about that. In this writeup we are going to talk about the “Recovery” directory.

It could be that you have never seen this directory before on your root drive (“C:\Recovery”). The reason for that is that the directory is marked “hidden” - as shown in the screenshot below.

By the way it is not enough to display hidden items in explorer to see it. In order to show it we need to unmark “Hide Protected Operating system files (recommended)”⁷ - you can see the entire flow in the following link.

Overall, the directory is a leftover from a previous version of Windows (the version before an upgrade that was made). It is used in cases where there are issues after an upgrade and the user wants to revert back. Thus, after a successful upgrade you can probably delete it⁸.



The screenshot shows a Windows Explorer window with the address bar set to 'This PC > Local Disk (C:)'. The left sidebar shows the 'Name' column with folders like PerfLogs, Program Files, Program Files (x86), Users, and Windows. The main pane shows a command prompt window with the following output:

```
C:\>dir /a /w
Volume in drive C has no label.
Volume Serial Number is 5860-F912

Directory of C:\

[$Recycle.Bin]           [$WinREAgent]
[Documents and Settings]  DumpStack.log.tmp
pagefile.sys           [PerfLogs]
[Program Files]         [Program Files (x86)]
[ProgramData]          [Recovery]
swapfile.sys           [System Volume Information]
[Users]                [Windows]
```

⁷ <https://www.techbout.com/hidden-system-files-windows-10-51145/>

⁸ <https://answers.microsoft.com/en-us/windows/forum/all/can-you-remove-the-large-crecovery-folder-in/e2185d3b-930c-41c0-a1d5-62c753b8a085>

COM (Component Object Model)

COM (Component Object Model) is a platform-independent, distributed, object-oriented system for creating binary software components that can interact with each other. COM is the foundation technology for Microsoft's OLE (compound documents) and ActiveX (Internet-enabled components) technologies. These objects can be within a single process, in other processes, even on remote computers⁹.

COM was introduced by Microsoft in 1993. It is used for IPC (Inter Process Communication) in a variety of programming languages. Also, COM allows the reuse of objects without any knowledge of their internal implementation, it forces component implementers to provide well-defined interfaces that are separated from the implementation¹⁰.

Let us go over a small example of using COM. Excel uses COM to enable users to create/modify/save/share excel files. By using COM we don't need to understand the binary format of excel files in order to perform the different operations. You can see a demonstration for that in the screenshot below.

Moreover, COM objects are registered with the operating system so they could be loaded in the future. The magic behind that is CLSID (Class ID). A CLSID is a globally unique identifier that identifies a COM class object. If your server or container allows linking to its embedded objects, you need to register a CLSID for each supported class of objects¹¹.

CLSID is stored in the registry under HKEY_CLASSES_ROOT\CLSID\{CLSID value}¹². It is used by the operating system to locate the appropriate code for loading. For examples of CLSIDs I suggest going over the following link

<https://www.elevenforum.com/t/list-of-windows-11-clsid-key-guid-shortcuts.1075/>.

They are several related technologies that we are going to talk about in future writeups: COM+, DCOM, Windows Runtime (aka WinRT), XPCOM (aka nano-COM), .NET framework, DEC/RPC, OLE, ActiveX, MSRPC and DDE¹³

⁹ <https://learn.microsoft.com/en-us/windows/win32/com/component-object-model--com--portal>

¹⁰ https://en.wikipedia.org/wiki/Component_Object_Model

¹¹ <https://learn.microsoft.com/en-us/windows/win32/com/clsid-key-hklm>

¹² <https://www.trendmicro.com/vinfo/us/security/definition/clsid>

¹³ <https://learn.microsoft.com/en-us/windows/win32/com/component-object-model--com--portal>

Windows PowerShell

PS C:\tmp> dir

Directory: C:\tmp

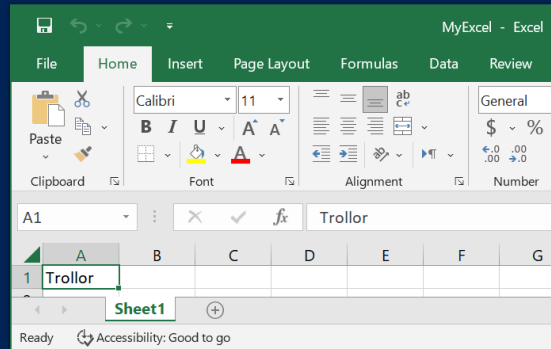
Mode	LastWriteTime	Length	Name
-a----	1/6/2023 10:47 AM	0	empty.txt

```
PS C:\tmp> $MyExcelFile = New-Object -ComObject Excel.Application
PS C:\tmp> $MyWorkbook = $MyExcelFile.Workbooks.Add()
PS C:\tmp> $MySheet = $MyWorkbook.Sheets.Item(1)
PS C:\tmp> $MySheet.Range("A1").Value = "Trollor"
PS C:\tmp> $MyWorkbook.SaveAs("C:\tmp\MyExcel.xlsx")
PS C:\tmp> $MyWorkbook.Close()
PS C:\tmp> $MyExcelFile.Quit()
PS C:\tmp> dir
```

Directory: C:\tmp

Mode	LastwriteTime	Length	Name
-a----	1/6/2023 10:47 AM	0	empty.txt
-a----	1/6/2023 10:48 AM	8668	MyExcel.xlsx

```
PS C:\tmp> start .\MyExcel.xlsx
PS C:\tmp>
```



Windows Services

Before we are going to talk about processes which are related to services handled under Windows (like services.exe and svchost.exe) we have to explain what a service is.

Services are processes which are managed by the operating system, it resembles demons in Linux (but there are a couple of differences that I am going to talk about in a different writeup).

Due to security reasons services can be executed at least under 3 different entities: System, Network Service and Local Service (each of them with different permissions and privileges - we will cover them in more details in the future). Of course, we can also run a service using a local user or a domain user with any access rights that we want.

There are different ways in which we can administer services, however I am going to focus on probably the 4 well known interfaces. First, Win32 API (it is used also for 64 bit despite its name) such as StartService¹⁴. Second, the mmc snap-in “services.msc”¹⁵. Third, PowerShell by leveraging cmdlets such as: New-Service, Get-Service, Restart-Service, Stat-Service, Stop-Service and more. Fourth, the command line tool “sc.exe”¹⁶.

A service can be in one of the three major states: started, stopped or paused. Also, each service has a startup state which defines what should happen with the service when the OS starts - it could be one of the following: Automatic, Automatic (Delayed Start), Manual or Disabled. Let us go over each and one of them.

Automatic, in this configuration the SCM starts the service as part of the system boot process. In the case of the delayed start, it's an optimization feature to reduce the time it takes the system to boot-up. “Automatic Start” is still run by the SCM but not during the boot process (they are started automatically shortly after the boot process has finished). Manual, in this configuration the SCM does not start the service and it needs to be from some other administrative interface (as we explained above), we can also script it if we want. Disabled, in this configuration even an administrator can't start the service. In order to start the service we first need to enable it by setting it to any setting which is not disabled.

Over the years different security enhancements were added for service hardening (Examples are session isolation, least privileges, restricted network access and service isolation). We are going to speak about it more when talking about security and the process of hardening the OS.

¹⁴ <https://docs.microsoft.com/en-us/windows/win32/services/service-functions>

¹⁵ <https://www.thewindowsclub.com/open-windows-services>

¹⁶ <https://ss64.com/nt/sc.html>

In the screenshot below you can see examples for the things we have talked about regarding the DHCP Client service. On the left we can see the status and the startup type and on the right the user which the service is logging on behalf of.

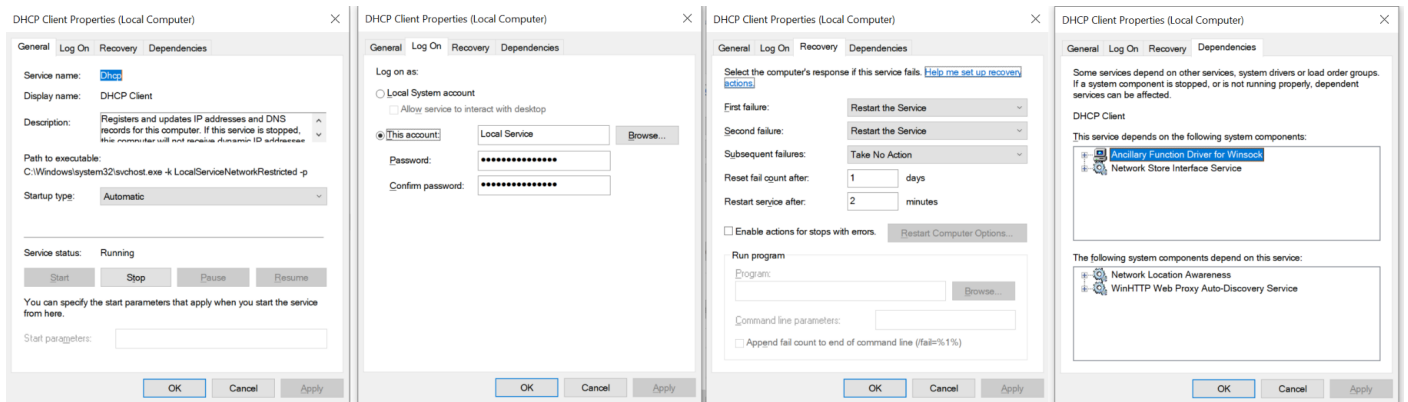
In addition, I am going to talk about: dependency management and recovery handling. In the screenshot below an example of those configurations are shown regarding the “DHCP Client” service (from services.msc). It is time to deep dive into it, so let’s go.

In the case of dependency management in the configuration of each service there is a list of dependent services and the list of the services that depend on it. Those lists are checked by the SCM when starting and stopping services.

Regarding recovery, the configuration allows setting actions for first/second/subsequent failures. The action could be to restart the service/the computer or execute an arbitrary command. We could also reset the fail count after N (we can set it) days and enable actions for cases the service stops with errors.

All the configurations of services are saved as part of the Registry (HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services). It is important to know that kernel drivers configuration is also stored there (we will go over it as part of the SCM discussion).

Next we are going to check out the SCM (Service Control Manager) which is the OS part responsible for managing all the services and their configuration.



What IPC (Inter Process Communication) mechanisms do we have in Windows?

Due to the fact that each process in Windows has its memory address space¹⁷ we can't pass pointers between threads in different processes and expect to see the same data in the same virtual address. It could be that the virtual address is not valid in one of the address spaces or we have a different data stored there (they are of other cases also).

In order to allow different threads (in different processes) to pass data between them we need to use an IPC (Inter Process Communication) mechanism - an illustration of that is seen in the diagram below¹⁸. Moreover, each OS has its own IPC mechanisms. On Windows we have the following mechanisms: clipboard, Windows Messages, COM (Component Object Model), DDE (Dynamic Data Exchange), Shared Memory, File Mapping, Mailslots, Pipes, RPC (Remote Procedure Call), ALPC (Advance Local Procedure Call) and sockets¹⁹.

There are also folks that say we can use files on the filesystem for IPC but it is not a specific mechanism of Windows so I won't speak about it for now. Lastly, when talking about IPC we should also talk about synchronization objects, however I will leave it for a future discussion.

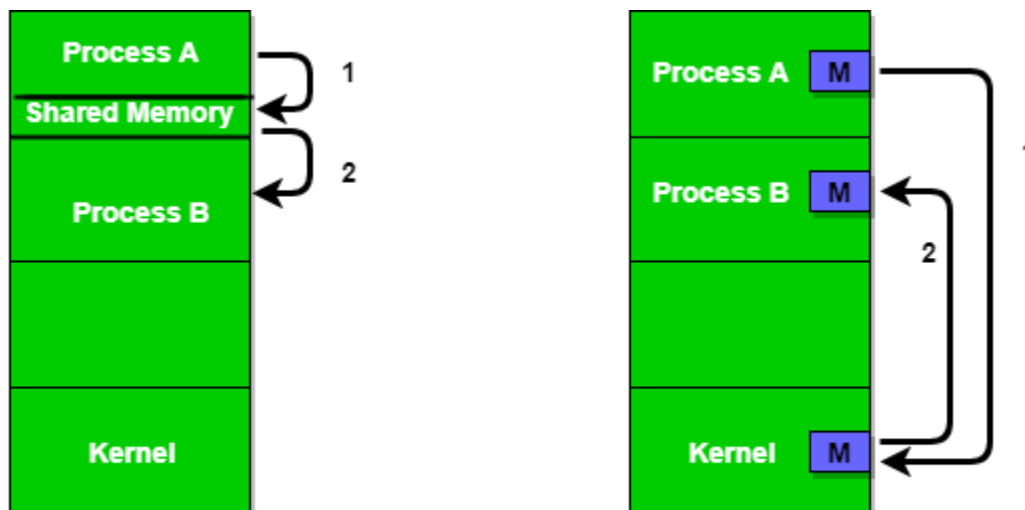


Figure 1 - Shared Memory and Message Passing

¹⁷ <https://medium.com/@boutnaru/linux-memory-management-part-1-introduction-896f376d3713>

¹⁸ <https://www.geeksforgeeks.org/inter-process-communication-ipc/>

¹⁹ https://www.slideshare.net/mfsj_vinothr/ipc-mechanisms-in-windows


Tasks (Windows Scheduler)

Overall, a task is a scheduled work that is performed by the “Task Scheduler” service. Each task has several components: triggers, actions, principals, settings, registration information and data - as shown in the diagram below²⁰.

Triggers are events/time-based conditions which are used as a criteria for starting an execution of a task. A task can have multiple triggers up to a maximum of 48²¹. Also, actions are the actual work performed by a task. A task can have a single/multiple actions up to a maximum of 32 actions. We can have different types of actions: “ComHandler” (COM), “Exec Action”, “Email Action” (sending an email notification) and “Show Message Action”²².

Moreover, principals is the definition of the security context in which the task is executing on behalf of, including UAC settings and more²³. Settings, that is the configuration used by the “Task Scheduler” while running the task. Think about if we can run multiple instances of the task, or what to do with the task if the system is in idle state and more. By default a task will stop after 72 hours, unless we change the “ExecutionTimeLimit”²⁴.

In addition, registration information is the data collected when the task is created/registered. Data elements that can be included (but not limited to) are: author, date, description, task version, security descriptor and more²⁵. We can also have additional documentation for the tasks (this is the data portion in the diagram shown below). Lastly, “Task Scheduler” has two versions (“1.0” and “2.0”) which have differences in the API they support and the configuration that can be made.



²⁰ <https://learn.microsoft.com/en-us/windows/win32/taskschd/tasks>

²¹ <https://learn.microsoft.com/en-us/windows/win32/taskschd/task-triggers>

²² <https://learn.microsoft.com/en-us/windows/win32/taskschd/task-actions>

²³ <https://learn.microsoft.com/en-us/windows/win32/taskschd/security-contexts-for-running-tasks>

²⁴ https://learn.microsoft.com/en-us/windows/win32/api/taskschd/nf-taskschd-itasksettings-get_executiontimelimit

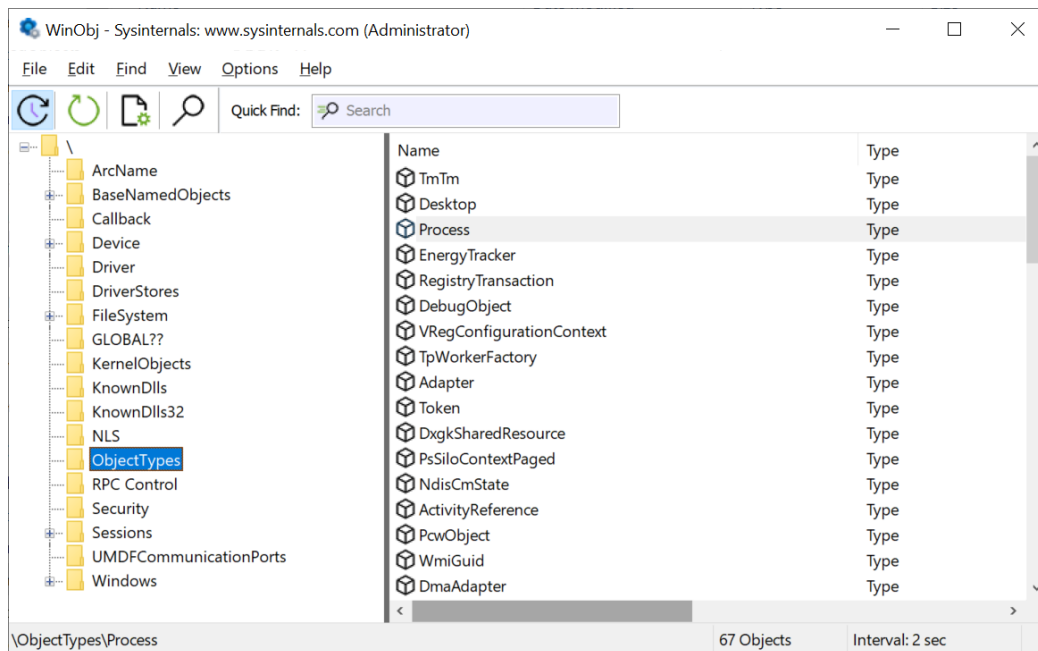
²⁵ <https://learn.microsoft.com/en-us/windows/win32/taskschd/task-registration-information>

Objects

Basically, objects are data structures that represent system resources. Think about things like processes and threads²⁶. Those object are divided to two main parts: the object header “struct _OBJECT_HEADER”²⁷ and the body which holds the specific information regarding a system resource. Examples are: _EPROCESS²⁸, _FILE_OBJECT²⁹ and more.

Moreover, we can see a list of available object types using the “WinObj” tool from Sysinternals - as shown in the screenshot below. By the way, the subsystem that manages the Windows resources is the “Object Manager”, which is part of the “Windows Executive”. The “Windows Executive” is contained as part of “ntoskrnl.exe”³⁰.

Overall, they are three different categories of objects in Windows: user, graphics (GDI objects) and kernel³¹. In User we have objects like: “Hook”³² and “Menu”³³. In GDI we have objects like: “Font”³⁴ and “Region”³⁵. Lastly, in the case of kernel objects we have examples like: “Desktop”³⁶ and “Job”³⁷.



²⁶ <https://learn.microsoft.com/en-us/windows/win32/sysinfo/handles-and-objects>

²⁷ https://www.nirsoft.net/kernel_struct/vista/OBJECT_HEADER.html

²⁸ <https://www.geoffchappell.com/studies/windows/km/ntoskrnl/inc/ntos/ps/eprocess/index.htm>

²⁹ https://learn.microsoft.com/en-us/windows-hardware/drivers/ddi/wdm/ns-wdm-_file_object

³⁰ [https://learn.microsoft.com/en-us/previous-versions/cc768129\(v=technet.10\)](https://learn.microsoft.com/en-us/previous-versions/cc768129(v=technet.10))

³¹ <https://learn.microsoft.com/en-us/windows/win32/sysinfo/object-categories>

³² <https://learn.microsoft.com/en-us/windows/win32/winmsg/hooks>

³³ <https://learn.microsoft.com/en-us/windows/win32/menurc/menus>

³⁴ <https://learn.microsoft.com/en-us/windows/win32/gdi/fonts-and-text>

³⁵ <https://learn.microsoft.com/en-us/windows/win32/gdi/regions>

³⁶ <https://learn.microsoft.com/en-us/windows/win32/winstation/desktops>

³⁷ <https://learn.microsoft.com/en-us/windows/win32/procthread/job-objects>

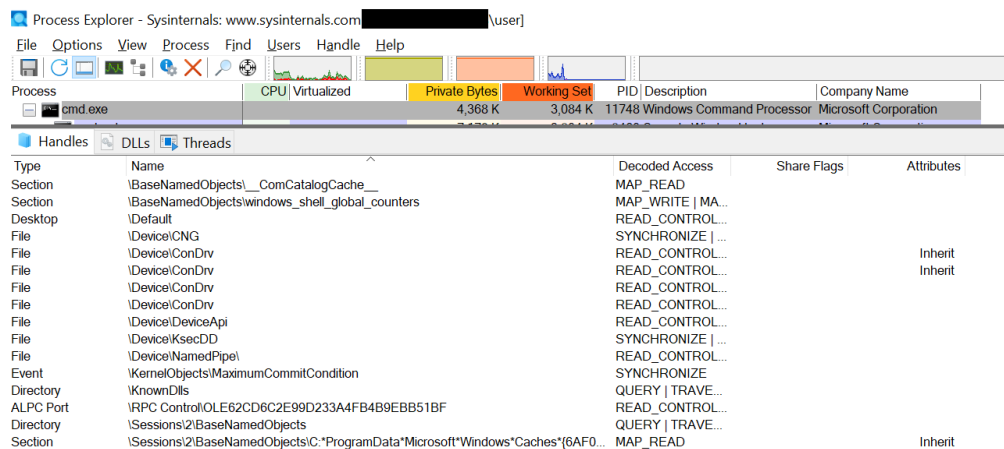
Object Manager

As part of the Windows operating system there is a single “Object Manager” that maintains all the objects³⁸. Among the tasks the “Object Manager” does are: creating objects, verifying that a process has a right to use an object, creating object handles³⁹ and returning them to the caller, maintaining resource quotas, duplicating handles and closing handles⁴⁰.

Overall, Windows has more than 25 types of objects. Some examples are: files, devices, threads, processes, events, mutexes, jobs, registry keys, sections, access tokens - as shown in the screenshot below taken from “Process Explorer”. The kernel routines which provide a direct interface with the “Object Manager” are prefixed with “Ob”⁴¹. Some examples are: “ObGetObjectSecurity”⁴² and “ObReferenceObjectByHandle”⁴³.

Moreover, we can go over the reference implementation of the “Object Manager” as part of ReactOS⁴⁴. There is also the internal header file⁴⁵.

Lastly, we can summarize the “Object Manager” as being responsible for keeping track of all the resources in Windows. It also provides a way for applications to access and manage those resources in a secure and efficient way⁴⁶.



³⁸ <https://medium.com/@boutnaru/windows-objects-2c289da600bf>

³⁹ <https://medium.com/@boutnaru/windows-handles-594b36c39d2f>

⁴⁰ <https://learn.microsoft.com/en-us/windows/win32/sysinfo/object-manager>

⁴¹ <https://learn.microsoft.com/en-us/windows-hardware/drivers/kernel/windows-kernel-mode-object-manager>

⁴² <https://learn.microsoft.com/en-us/windows-hardware/drivers/ddi/wdm/nf-wdm-obgetobjectsecurity>

⁴³ <https://learn.microsoft.com/en-us/windows-hardware/drivers/ddi/wdm/nf-wdm-obreferenceobjectbyhandle>

⁴⁴ <https://github.com/reactos/reactos/tree/master/ntoskrnl/ob>

⁴⁵ <https://github.com/reactos/reactos/blob/master/ntoskrnl/include/internal/ob.h>

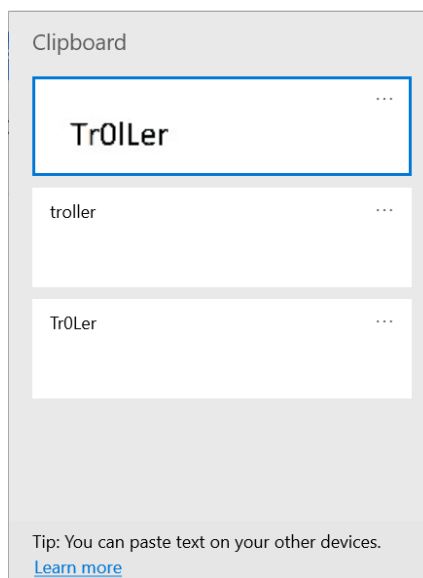
⁴⁶ <https://www.i.u-tokyo.ac.jp/edu/training/ss/lecture/new-documents/Lectures/01-ObjectManager/ObjectManager.pdf>

Clipboard

Overall, the clipboard is a temporary storage in Windows that can store different data types. Among those data types are images and text. The operation which stores data in the clipboard is called “copy” (Ctrl+C), also we can use “cut” using “Ctrl+X”. In order to paste data from the clipboard we use “Ctrl+V”⁴⁷.

In Windows 10 we can store in the clipboard up to 25 items. We can access the “Clipboard History” using the keyboard shortcut “Winkey + V”. While pressing this shortcut a menu showing the current clipboard history - as shown in the screenshot below⁴⁸. We can see that the first item is an image and the other two are strings.

Lastly, in Windows 10/11 we can also copy text/images from one PC to another using a cloud based clipboard. We can also pin items that we would like to use. In order to do that we need to sign with a Microsoft account/work account⁴⁹.



⁴⁷ <https://www.howtogeek.com/671222/how-to-enable-and-use-clipboard-history-on-windows-10/>

⁴⁸ <https://www.emailoverloadsolutions.com/blog/windows-clipboard-history-feature>

⁴⁹ <https://support.microsoft.com/en-us/windows/clipboard-in-windows-c436501e-985d-1c8d-97ea-fe46ddf338c6>

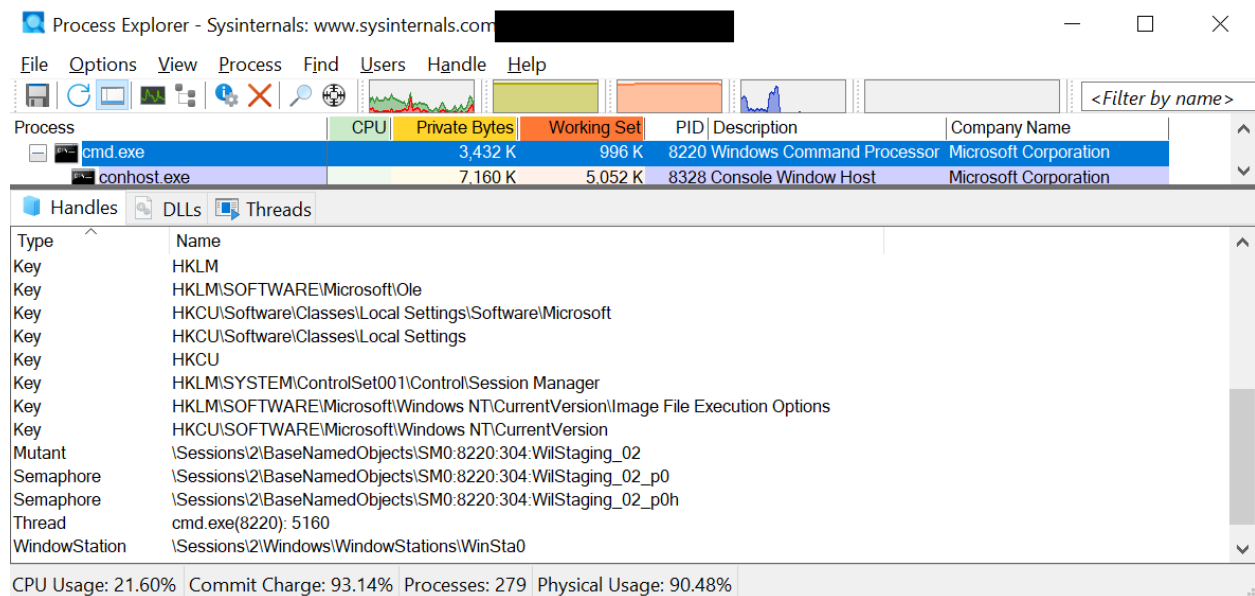
Handles

Overall, applications running in user mode can't directly access the Windows objects⁵⁰ or object data held by the "Object Manager". Due to that, an application needs to obtain a handle to the specific object. For each handle there is an entry in a handle table which resides in kernel space, there is one for each process⁵¹.

Thus, an object in Windows is accessed by the user using a per process "handle table". Opening an object results in adding a pointer to the object to the specific "handle table". The return value is "an index" to that table⁵².

We can see when using different Win32 API function like when opening/creating a file using "CreateFileW" the return value is of type HANDLE⁵³.

Lastly, we can use Sysinternals' tools like "Process Explorer" and handle.exe⁵⁴ in order to see which open handles each process has. An example showing the handles for "cmd.exe" using "Process Explorer" is shown in the screenshot below.



⁵⁰ <https://medium.com/@boutnaru/windows-objects-2c289da600bf>

⁵¹ <https://learn.microsoft.com/en-us/windows/win32/sysinfo/handles-and-objects>

⁵² https://www.cs.miami.edu/home/burt/journal/NT/handle_table.html

⁵³ <https://learn.microsoft.com/en-us/windows/win32/api/fileapi/nf-fileapi-createfilew>

⁵⁴ <https://learn.microsoft.com/en-us/sysinternals/downloads/handle>