

# IY5512: Coursework 3: Worked solutions

---

1. *What is the difference between paging and segmentation? Briefly indicate their respective advantages and disadvantages.*

Paging is a virtual memory management scheme which involves dividing memory into fixed-size blocks (**pages**), e.g. of 4 kbytes each. Paging is transparent to an application program. When a program is loaded into memory it is divided into pages, which are loaded into pages of physical memory.

Segmentation is a different memory management scheme which involves dividing memory into blocks (**segments**) of variable size. The use of segmentation is visible to an application program.

The fact that paging is application-transparent is both an advantage and a disadvantage. The advantage of transparency is the lack of overhead, but the disadvantage is that paging cannot be used for policy enforcement.

Similarly, the fact that segmentation is not transparent is both a disadvantage (it imposes a book-keeping burden on an application), and an advantage (it allows different segments be used for different purposes, and treated separately with regards to security policy).

The two schemes both have a fragmentation overhead. Paging gives rise to **internal fragmentation**, where some of the fixed size pages are not fully used. Segmentation gives rise to **external fragmentation**, where the different sizes of segments means that sections of memory address space cannot be allocated to a segment.

A helpful description of segmentation and paging can be found here:

<http://fourier.eng.hmc.edu/e85/lectures/memory/node8.html>

2. *Describe the use made by Windows of the four privilege levels provided by the Intel x86 family of processors. How does this contrast with the intended use of these levels? What problems does this cause?*

The four rings are not used by Windows as was originally intended. Windows only uses two of the four rings – ring 0 for the Operating System (OS) and ring 3 for applications. Rings 1 and 2 were originally intended for use by drivers and system services, respectively.

As a result, all OS activities share the same hardware security level. Every time a single OS component (e.g. a driver) changes, the security of the entire OS is affected. Many attacks result from use of ring 0 for all system activities.

Whilst this could be fixed in principle, in practice it would require the OS, most drivers, and some applications to be rewritten – this is simply not a viable strategy.

3. *What security issues arise from interrupts? Find and briefly describe an example of a security vulnerability arising from misuse of interrupts.*

When certain interrupts occur, the processor will execute a system mode (privileged) program, the address of which is specified in the Interrupt Vector Table. If this table can be modified by an attacker, it might enable the attacker to run code of the attacker's choice with system privileges. This can then be used to do arbitrary damage to the system.

An attacker could use an interrupt to cause the OS to execute a system call, and then use this system call to perform a malicious action. For example, a boot sector virus could issue an interrupt causing data to be written to the boot sector, a specific portion of the hard disk.

A possible attack using interrupts is described in Gollmann's Computer Security (section 5.3.1 of 2nd edition, and section 6.3.4 of the 3rd edition).

4. *Explain the difference between a virtual address and a physical address. Discuss the advantages of the use of virtual memory management.*

A physical address is the fixed address of a location in physical memory in a computer.

A compiled program will refer internally to memory locations using a virtual address, which is translated to a physical memory address by the machine on which it is running at runtime.

There are many advantages of the use of virtual memory management:

- it enables a programme to be loaded anywhere in the physical memory of a machine without recompilation;
- a programme does not need to be loaded into a contiguous block of physical memory – paging enables use of non-contiguous memory – thereby easing memory fragmentation issues;
- if memory is limited, virtual memory management allows only parts of a programme to be loaded into physical memory; indeed, a process can run on a system which has less physical memory than the amount required by the process;
- it enables processes to be security separated using the memory management system.

5. *What security threats can Direct Memory Access (DMA) devices pose?*

DMA by a peripheral bypasses the security protection provided by the memory management system. It enables a physical device to read or write from/to any location in physical memory. This can be used to bypass domain separation.

6. *Briefly describe the mechanism provided within Intel TXT to provide protection against the threats posed by DMA.*

The threat posed by DMA is addressed in TXT by the introduction of the noDMA table, controlled by the VMM. That is, while a DMA access bypasses the VMM and the paging mechanism, DMA is controlled by the noDMA table, and this table is (in turn) controlled by the VMM. The noDMA table is a chipset component. The table controls which physical pages may be accessed using DMA. The VMM is responsible for maintaining synchronisation between the protected pages in the paging mechanism, and entries in the noDMA table.