

```

Sep 19 14:27:41 amd64 syslog-ng[7653]: STATS:
Sep 20 01:00:01 amd64 /usr/sbin/cron[2937]: (root) CMD (/sbin/evlogmgr -c "severity=DEBUG")
Sep 20 01:00:01 amd64 syslog-ng[7653]:
Sep 20 02:00:01 amd64 /usr/sbin/cron[2937]: (root) CMD (/sbin/evlogmgr -c "age > *30d")
Sep 20 02:00:01 amd64 syslog-ng[7653]:
Sep 20 12:46:44 amd64 sshd[6531]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62004
Sep 20 12:48:41 amd64 sshd[6544]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62105
Sep 20 12:54:44 amd64 sshd[6544]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62514
Sep 20 15:27:35 amd64 sshd[6544]: Accepted rsa for esser from ::ffff:87.234.201.207 port 64242
Sep 20 15:27:35 amd64 sshd[6544]: Accepted rsa for esser from ::ffff:87.234.201.207 port 63375
Sep 20 16:11:11 amd64 sshd[6544]: Accepted rsa for esser from ::ffff:87.234.201.207 port 63546
Sep 20 02:00:01 amd64 /usr/sbin/cron[2937]: (root) CMD (/sbin/evlogmgr -c "severity=DEBUG")
Sep 20 02:00:01 amd64 /usr/sbin/cron[2937]: (root) CMD (/sbin/evlogmgr -c "age > *30d")
Sep 23 02:00:01 amd64 /usr/sbin/cron[2937]: (root) CMD (/sbin/evlogmgr -c "age > *30d")
Sep 23 02:00:01 amd64 /usr/sbin/cron[2937]: (root) CMD (/sbin/evlogmgr -c "age > *30d")
Sep 23 18:04:05 amd64 sshd[6554]: Accepted publickey for esser from ::ffff:192.168.1.5 port 59771 ssh2
Sep 23 18:04:05 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 23 18:04:34 amd64 sshd[6606]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62093
Sep 24 01:00:01 amd64 /usr/sbin/cron[12436]: (root) CMD (/sbin/evlogmgr -c "severity=DEBUG")
Sep 24 01:00:01 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 24 02:00:01 amd64 /usr/sbin/cron[12521]: (root) CMD (/sbin/evlogmgr -c "age > *30d")
Sep 24 02:00:01 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 24 11:15:48 amd64 sshd[20998]: Accepted rsa for esser from ::ffff:87.234.201.207 port 64456
Sep 24 13:49:08 amd64 sshd[23197]: Accepted rsa for esser from ::ffff:87.234.201.207 port 61330
Sep 24 13:49:08 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 24 15:42:07 amd64 kernel: snd_seq_midi_event: unsupported module, tainting kernel.
Sep 24 15:42:07 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 24 15:42:07 amd64 kernel: snd_seq_oss: unsupported module, tainting kernel.
Sep 24 20:25:31 amd64 sshd[29391]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62566
Sep 24 20:25:31 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 01:00:02 amd64 /usr/sbin/cron[662]: (root) CMD (/sbin/evlogmgr -c "severity=DEBUG")
Sep 25 02:00:01 amd64 /usr/sbin/cron[1484]: (root) CMD (/sbin/evlogmgr -c "age > *30d")
Sep 25 02:00:02 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 10:59:25 amd64 sshd[8889]: Accepted rsa for esser from ::ffff:87.234.201.207 port 64183
Sep 25 10:59:25 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 10:59:47 amd64 sshd[8921]: Accepted rsa for esser from ::ffff:87.234.201.207 port 64253
Sep 25 11:30:02 amd64 sshd[9372]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62029
Sep 25 11:59:05 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 14:05:37 amd64 sshd[11554]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62822
Sep 25 14:05:37 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 14:06:10 amd64 sshd[11566]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62951
Sep 25 14:07:17 amd64 sshd[11608]: Accepted rsa for esser from ::ffff:87.234.201.207 port 63392
Sep 25 14:08:33 amd64 sshd[11630]: Accepted rsa for esser from ::ffff:87.234.201.207 port 63709
Sep 25 15:25:33 amd64 sshd[12930]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62778

```

Memory Management (4)

```

Sep 20 01:00:01 amd64 /usr/sbin/cron[2937]: (root) CMD (/sbin/evlogmgr -c "severity=DEBUG")
Sep 20 02:00:01 amd64 /usr/sbin/cron[2937]: (root) CMD (/sbin/evlogmgr -c "age > *30d")
Sep 20 12:46:44 amd64 sshd[6531]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62004
Sep 20 12:48:41 amd64 sshd[6544]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62105
Sep 20 12:54:44 amd64 sshd[6544]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62514
Sep 20 15:27:35 amd64 sshd[6544]: Accepted rsa for esser from ::ffff:87.234.201.207 port 64242
Sep 20 15:27:35 amd64 sshd[6544]: Accepted rsa for esser from ::ffff:87.234.201.207 port 63375
Sep 20 16:11:11 amd64 sshd[6544]: Accepted rsa for esser from ::ffff:87.234.201.207 port 63546
Sep 20 02:00:01 amd64 /usr/sbin/cron[2937]: (root) CMD (/sbin/evlogmgr -c "severity=DEBUG")
Sep 20 02:00:01 amd64 /usr/sbin/cron[2937]: (root) CMD (/sbin/evlogmgr -c "age > *30d")
Sep 23 02:00:01 amd64 /usr/sbin/cron[2937]: (root) CMD (/sbin/evlogmgr -c "age > *30d")
Sep 23 02:00:01 amd64 /usr/sbin/cron[2937]: (root) CMD (/sbin/evlogmgr -c "age > *30d")
Sep 23 18:04:05 amd64 sshd[6554]: Accepted publickey for esser from ::ffff:192.168.1.5 port 59771 ssh2
Sep 23 18:04:05 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 23 18:04:34 amd64 sshd[6606]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62093
Sep 24 01:00:01 amd64 /usr/sbin/cron[12436]: (root) CMD (/sbin/evlogmgr -c "severity=DEBUG")
Sep 24 01:00:01 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 24 02:00:01 amd64 /usr/sbin/cron[12521]: (root) CMD (/sbin/evlogmgr -c "age > *30d")
Sep 24 02:00:01 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 24 11:15:48 amd64 sshd[20998]: Accepted rsa for esser from ::ffff:87.234.201.207 port 64456
Sep 24 13:49:08 amd64 sshd[23197]: Accepted rsa for esser from ::ffff:87.234.201.207 port 61330
Sep 24 13:49:08 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 24 15:42:07 amd64 kernel: snd_seq_midi_event: unsupported module, tainting kernel.
Sep 24 15:42:07 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 24 15:42:07 amd64 kernel: snd_seq_oss: unsupported module, tainting kernel.
Sep 24 20:25:31 amd64 sshd[29391]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62566
Sep 24 20:25:31 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 01:00:02 amd64 /usr/sbin/cron[662]: (root) CMD (/sbin/evlogmgr -c "severity=DEBUG")
Sep 25 02:00:01 amd64 /usr/sbin/cron[1484]: (root) CMD (/sbin/evlogmgr -c "age > *30d")
Sep 25 02:00:02 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 10:59:25 amd64 sshd[8889]: Accepted rsa for esser from ::ffff:87.234.201.207 port 64183
Sep 25 10:59:25 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 10:59:47 amd64 sshd[8921]: Accepted rsa for esser from ::ffff:87.234.201.207 port 64253
Sep 25 11:30:02 amd64 sshd[9372]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62029
Sep 25 11:59:05 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 14:05:37 amd64 sshd[11554]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62822
Sep 25 14:05:37 amd64 syslog-ng[7653]: STATS: dropped 0
Sep 25 14:06:10 amd64 sshd[11566]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62951
Sep 25 14:07:17 amd64 sshd[11608]: Accepted rsa for esser from ::ffff:87.234.201.207 port 63392
Sep 25 14:08:33 amd64 sshd[11630]: Accepted rsa for esser from ::ffff:87.234.201.207 port 63709
Sep 25 15:25:33 amd64 sshd[12930]: Accepted rsa for esser from ::ffff:87.234.201.207 port 62778

```

Segmentation with Paging (2)

Two ways of implementing this:

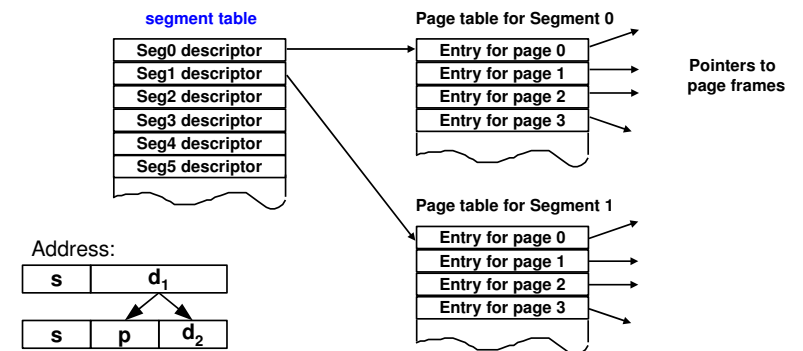
- One segment table and for each segment (per process) one page table. Entries in segment table point to the page table of the segment (e.g. Unix).
- One segment table and one single page table (per process) (e.g. Intel CPUs).
 - Entries in the segment table contain the base address of the segment in a linear (virtual) address space.
 - Add the relative address within the segment to this base address. The resulting linear (virtual) address is translated into a physical address by using the page table.

Segmentation with Paging (1)

- Partition programs into segments, which are placed in the physical memory using Paging (non-contiguous).
- Programmer uses segment numbers and relative addresses within each segment. Paging is handled by the hardware – transparent for the program.

Segmentation with Paging (3)

- One page table per segment:



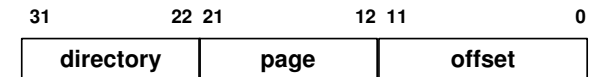
Segmentation with Paging (4)

- Addresses always consist of
 - a segment number and
 - a linear address within the segment.

Splitting of the linear address into page number and offset (for paging) is transparent for the program.

Segmentation with Paging: Intel 80386

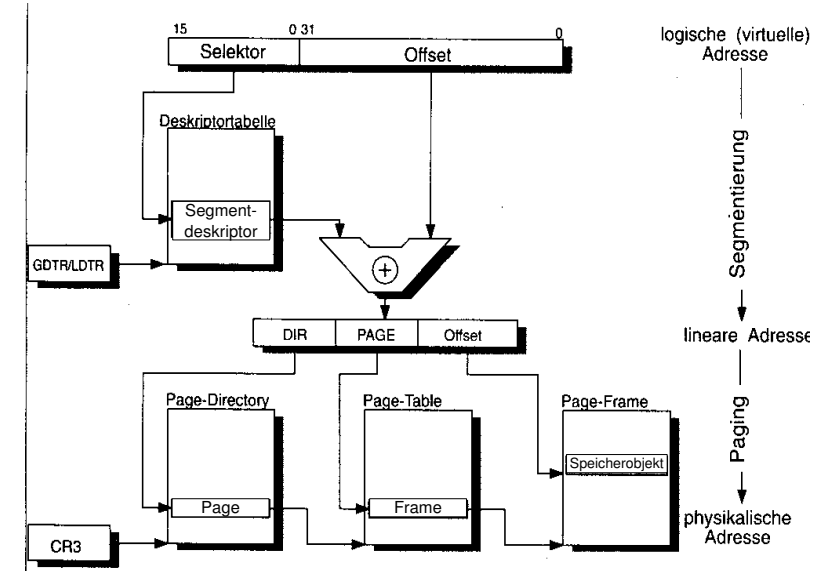
- Six internal registers which contain the corresponding segment descriptors.
- An **address** is a tuple (segment selector, offset).
- The segment descriptor is the base address of the segment in a 32 bit address space plus the offset, delivering the so-called linear address.
- The linear address looks like this (**two-level paging**):



Segmentation with Paging: Intel 80386

- Segments:
 - Up to 8 K process-private segments, described by entries in the **local descriptor table (LDT)**.
 - Up to 8 K **shared** segments (used by all processes), described by entries in the **global descriptor table (GDT)**.
 - Each segment can be up to 4 GByte in size.
- **Six segment registers**, at each time a process can use up to six segments.

Address translation with Intel 80386



PAE and Linux 2.6 (3)

- Memory partitioned:
 - ZONE_DMA: First 16 MByte
for old ISA devices
 - ZONE_NORMAL: 16 MByte – 896 MByte
Standard (kernel and user mode)
 - ZONE_HIGHMEM: 896 MByte – Rest
- Kernel has 1 GByte of virtual memory, of which the first 896 MByte (ZONE_DMA und ZONE_NORMAL) are statically mapped.
- 128 MByte remain, which can be dynamically mapped to access memory > 1 GByte.

PAE and Windows 2000 (2)

- The Windows PAE kernel supports the following physical memory sizes:
 - W2K Professional: 4 GByte
 - W2K Server: 4 GByte
 - W2K Advanced Server: 8 GByte
 - W2K Datacenter 64 GByte
- Microsoft offers „Address Windowing Extensions (AWE) API“ which let applications allocate physical memory for their exclusive usage and let them map parts of their address spaces to this memory.

PAE and Windows 2000 (1)

- In order to use PAE, the operating system must be modified.
- For Windows 2000, Microsoft has developed a special PAE kernel (ntkrnlpa.exe). If the CPU is PAE-capable and the system has more than 4 GByte memory, this kernel is loaded.

Demand Paging (1)

- The address space of a process need not reside completely in main memory.
 - Locality principle says that a process will only access few, close (local) addresses in any given time span.
 - Parts of the program code possibly never execute during regular processing (e.g. special cases; error handling routines).

Demand Paging (2)

- **Demand Paging** means that
 - a page is only loaded into memory when a process accesses it,
 - a page can be removed from memory.
- Advantages of Demand Paging:
 - Address space of a process can be larger than the total physical memory.
 - Processes need less space in main memory, thus more processes can be (in memory and) ready at the same time.

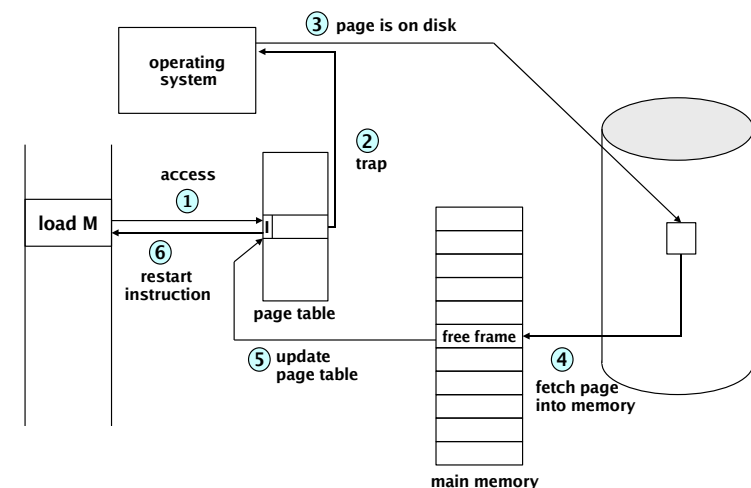
Requirements for Demand Paging (2)

- If there is **no free page frame** in the main memory, some other page has to be replaced. For selecting the page that is to be replaced, a replacement strategy must be implemented.
- The instruction that was interrupted by the page fault, must be re-executed (and it must be possible to do that).

Requirements for Demand Paging (1)

- Each entry in the page table has a **valid bit** that tells whether the page is in memory or not.
- When a process accesses a page that is not in memory, this causes a special exception, the so-called **page fault**.
- An operating system routine, the **page fault handler**, loads the requested page from disk.

Page Fault Handling



Page Replacement (1)

- When a Page Fault occurs and there is **no free page frame** available, the operating system must free one.
- An algorithm chooses this page frame according to a specific strategy.

Page Replacement (3)

- An unmodified page can later (if needed) be loaded from the old, known location on disk.
- In the page table entry for the replaced page
 - the **valid bit** will be deleted,
 - information about the location, from where the page can be retrieved, will be saved.

Page Replacement (2)

- If the page to be replaced has been changed since it was last fetched from the disk, its content must be saved:
 - A **modify bit** (or **dirty bit**) in the page table entry tells whether the page has been modified.
 - A modified page will be written to disk (to the **page or swap area**).

Page Replacement Strategies (1)

- Goal: As few Page Faults as possible.
- Two principal types of page replacement strategy:
 - local replacement
 - global replacement

Page Replacement Strategies (2)

local replacement:

Always replace a page that belongs to the process which requests a new page.

- The number of pages which can be allocated by a process has an upper bound. The maximum number is assigned per process (e.g. by the system administrator) and it can strongly influence runtime behavior of a process.
- A process which causes many Page Faults (e.g. because it does not conform to the locality principle) affects only itself, not the whole system.

Optimal Strategy

Replace the page which will not be accessed for the longest time (in the future).

- **Advantage:** This strategy causes the smallest number of page faults.
- **Disadvantage:** This strategy is not implementable.

However, the optimal strategy can serve as a model in order to evaluate other strategies.

Page Replacement Strategies (3)

global replacement:

Replace an arbitrary page in memory.

- processes take away each other's page frames.
- a process that causes many page faults automatically receives more page frames. (This can be an advantage or disadvantage.)

First In First Out (FIFO)

Replace the page which has been in memory for the longest time.

- **Advantage:** very easy to implement:
 - Manage a chained list of all pages in memory (global strategy) respectively of a process' pages (local strategy).
 - When a page fault occurs, replace the first page in the list and append the new page to the end of the list.
- **Disadvantage:** The replaced page can be in constant use and might be re-requested immediately.

Least Recently Used (LRU) (1)

Replace the page, that **has not been used for the longest time**.

- **Advantage:** Typically less Page Faults than with FIFO strategy.
- **Nachteil:** complex implementation.

two approaches for implementation:

- with counter
- with chained list

Least Recently Used (LRU) (3)

- Implementation with **chained list**:
 - a chained list contains all pages.
 - at each memory access: bring the accessed page to the head of the list. (walk through list and change order, i.e. change pointers)
 - replace page at the end of the list.

Least Recently Used (LRU) (2)

- Implementation with **counter**:
 - system-wide counter that is incremented at each memory access.
 - write recent value of counter into a field in the data structure which describes the accessed page.
 - replace page with lowest value.

Using a Reference Bit (1)

- each page table entry can contains a **reference bit**
 - that is set at accessing this page (by the hardware),
 - that will be unset according to specific criteria (by the software).
- a reference bit
 - holds the information, whether a page was accessed since last unsetting of the bit,
 - says nothing about when the page was accessed,
 - says nothing about the order of accesses of several pages.

Using a Reference Bit (2)

- reference bits let you implement further page replacement strategies, e.g.
 - modifications of LRU which are less complex.
 - Second Chance Algorithm, an improved version of the FIFO strategy.

Modification of LRU (2)

- replace one of the pages having the minimal counter value.
 - If several pages have the same counter value, it is unknown which one was accessed last.
 - Accesses which happened longer ago, will first be weighted less and finally be "forgotten" (shifted out of the counter).

Modification of LRU (1)

- create a **binary counter** for each page.
- in regular intervals:
 - shift each counter one position to the right ("**aging**"),
 - copy the reference bit into the highest bit of the counter,
 - unset the reference bit.