


.NET Memory Management Poster II ¹⁾²⁾

by @konradkokosa <https://prodotnetmemory.com>

- Allocation triggers:**
- AllocSmall** - running out of budget on Gen 0 during object allocation.
 - AllocLarge** - running out of budget on LOH during large object allocation
 - OutOfSpaceSOH** - running out of SOH ephemeral segment space
 - OutOfSpaceLOH** - running out of LOH segments space
 - Induced** - called explicitly from code
 - LowMemory** - operating system has sent *low-memory notification*
 - Internal** - various GC reasons, like: *AppDomain* unload or cleaning up *Thread* objects

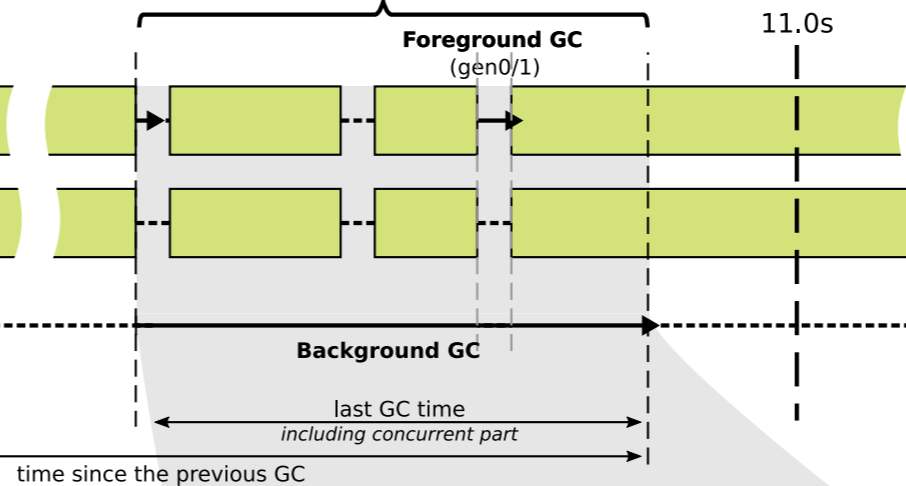
"Stop-the-world" GC

suspends all threads, may compact




Background GC

with optional Foreground GCs
suspends all threads only for some time, may not compact



$$\% \text{ Time in GC} = \frac{\text{last GC time}}{\text{since previous GC time}} [\%]$$

- GC types:**
(as shown in PerfView)
- N** - non-concurrent GC (blocking)
 - B** - Background GC
 - F** - Foreground GC (blocking collection of an ephemeral generations during Background GC)
 - I** - induced (manually triggered) blocking GC
 - i** - induced non-blocking GC



ETW/LTTng events

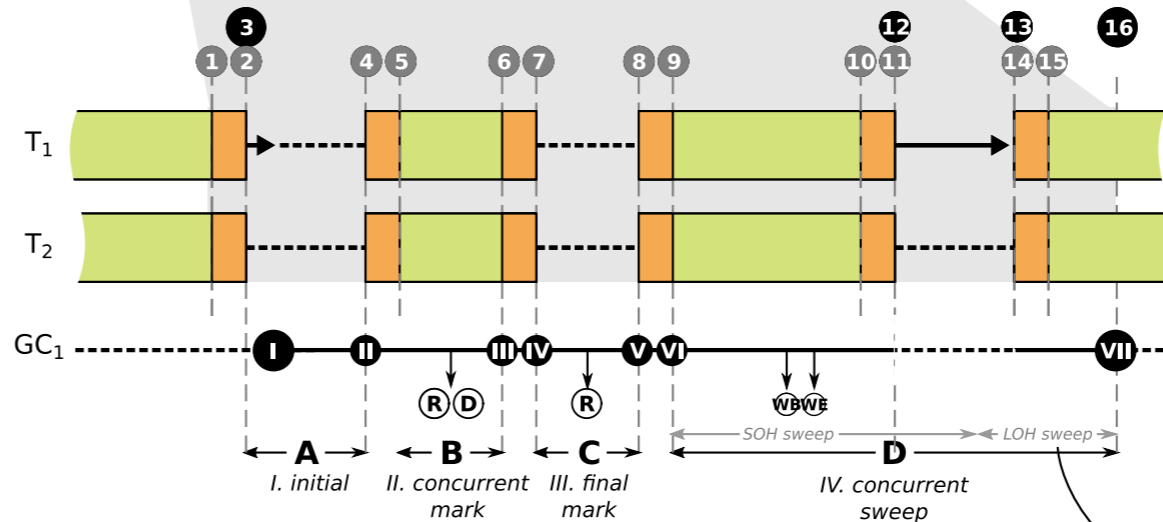
- GC/SuspendEEStart
- GC/SuspendEEStop
- GC/Start
- GC/Stop
- GC/RestartEEStart
- GC/RestartEEStop

Thread suspend/resume

Every managed method may be:

- Fully interruptible** - thread running it may be suspended at any time
- Partially interruptible** - thread running it may be suspended at specified *safe points* (mostly, method calls)

GCInfo data knows where *GC roots* live (what registers, where on stack) for the current method instruction



ETW/LTTng events

- GC/SuspendEEStart
- GC/SuspendEEStop
- GC/Start
- GC/RestartEEStart
- GC/RestartEEStop
- GC/SuspendEEStart
- GC/SuspendEEStop
- GC/RestartEEStart
- GC/RestartEEStop
- GC/RestartEEStart
- GC/SuspendEEStart
- GC/SuspendEEStop
- GC/Start
- GC/Stop
- GC/RestartEEStart
- GC/RestartEEStop
- GC/Stop
- BGCStart
- BGC1stNonConStop
- BGCRevisit
- BGCDrainMark
- BGC1stConStop
- BGC2ndNonConStart
- BGCRevisit
- BGC2ndNonConStop
- BGC2ndConStart
- BGC2ndConStop

LOH allocations not allowed
(produces *LOH Allocation Pause* (due to *background GC*) > 200 Msec report in PerfView)

- VB BGCAllocWaitBegin
- VE BGCAllocWaitEnd


- Choose **condemned generation**
- Marking** of reachable objects in the condemned and younger generations
- Plan** - the GC decides whether compacting is worth doing or maybe sweeping is just enough
- Sweep** or...
- Relocate** and **Compact** - to update all addresses to the new ones

- Choose **condemned generation**
- Initial marking** - finding GC roots
- Concurrent Marking** of reachable objects in the condemned and younger generations, starting from the initial GC roots
- Final Mark** - revisit reachability from pages that has been modified during concurrent marking
- Concurrent **Sweep**

Workstation GC


minimize pauses for better **interactivity**

Workstation Non-concurrent



- No GC threads
- Only "stop-the-world" GC (may compact)

Background Workstation




- Single GC thread
- Gen 0/1 (and sometimes 2) - "stop-the-world" GC (may compact)
- Mostly - Background GC (concurrent sweep)

Server GC

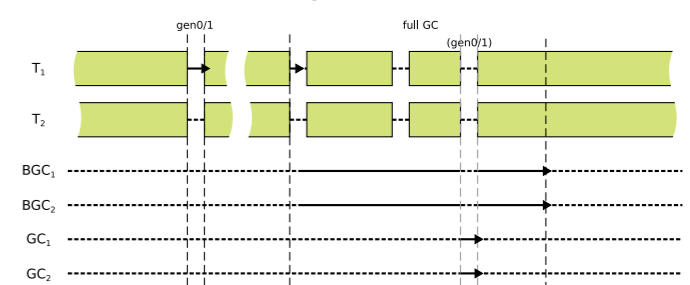
maximize resources usage for better **throughput**

Server Non-concurrent



- By default, N GC threads (N - # of cores)
- Only "stop-the-world" GC (may compact)

Background Server



- By default, N GC threads (N - # of cores)
- Gen 0/1 (and sometimes 2) - "stop-the-world" GC (may compact)
- Mostly - Background GC (concurrent sweep)

Legend: Working user thread (green bar), Suspending thread (orange bar), Working GC thread (solid line), Suspended thread (dashed line)

1) Based on .NET Core and 2 CPU cores with 2 user threads, 64-bit Windows
2) Many identifiers (like **AllocSmall** or **% Time in GC**) are used as defined in *PerfView* tool