

## zTidBit#26 Cheatsheet

# ENCLAVES & Speciality Processors

Generally, a dispatchable unit (DU), as generally in z/OS is a serial execution of computer logic code to achieve a business need.

There are **four types** of DUs in z/OS:

- \* Preemptible Task Control Block (TCB)
- \* Nonpreemptible Service Request Block (SRB)
- \* Preemptible Client Service Request Bolck (client SRB)
- \* Preemptible Enclave Service Request Block (enclave SRB)

(NOTE - TCBs are user execution threads – SRBs are System Threads)

Each Dispatchable Unit (DU) is associated with:

1. State, 2. Resources, 3. Priorities, 4. Accounting, 5. Security, 6. Addressing Mode

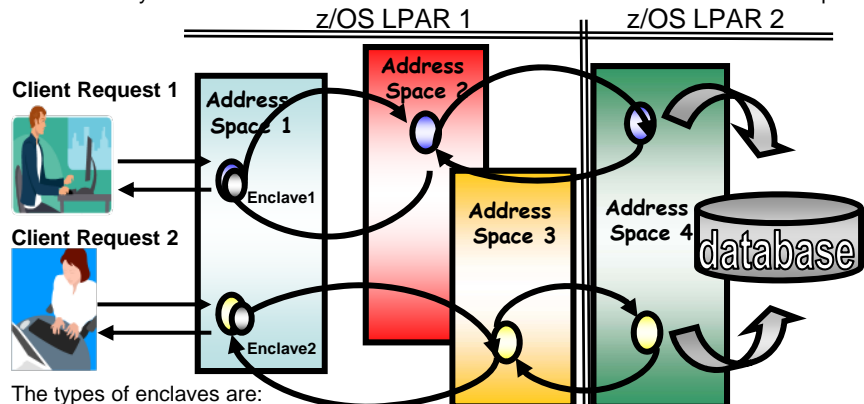
A program and a DU are different entities. In z/OS, you have control blocks describing the DU (TCB and SRB) with a description of their state and held resources. A *program* is a sequence of instructions that are executed on behalf of a DU on a z/OS Central Processor Unit (CPU).

A **preemptive** operating system follows strict CPU dispatching priority (DP). When a dispatchable unit (DU) is active and another DU with a higher DP becomes ready, the active DU is instantaneously suspended (its state is saved) and the higher priority ready DU becomes active.

A **non-preemptible** dispatchable units are traditional SRBs. Once dispatched, they continue to run until they incur a voluntary suspension (such as lock held or page fault) or they complete. External and I/O interrupts are serviced (usually SRBs are interrupt enabled), but the SRB is redispached after each interrupt is processed.

WLM transaction management employs an Enclave which is an extension of dispatchable units. Enclaves provide a much more direct way of managing DUs that span multiple address spaces such as a Websphere (or CICS-MRO) transaction. An *enclave* is an entity that encapsulates the execution units (TCBs and SRBs) that execute programs on behalf of the same work request. You can also think of an enclave as a "Business Unit of Work (BUoW)" without address space boundaries. It is close to the user view of a transaction.

Enclaves are created by WLM on behalf of a client request and therefore one address space can own many enclaves. Enclaves maintain account information of the associated dispatchable unit.



The types of enclaves are:

- \* **Independent enclaves:** These are true WLM transactions separately classified and managed in a WLM service class (SC). SC is a grouping of transaction attributes reported an entity.
- \* **Dependent enclaves:** These are a logical extension of an existing address space and inherit WLM service class from the owner address space. They continue an existing transaction that is running under dispatchable units not associated with the originating address space.

**NOTE:** When an enclave spans a system boundary in a sysplex (more than one LPAR), it is called a *multisystem enclave*.

With z/OS Subsystems composed of multiple address spaces as well as today's current transaction complexity if not for the enclave - accounting for performance, resource and user usage would be a very difficult challenge and error prone even for the most astute z/OS administrator.

Enclave reporting is captured in System Management Facility (SMF) record type 30 ( and 72 ).

In the past, all workload executed on general processors and the capturing throughput and resource usage for enclaves was fairly simple. BUT, the introduction of speciality engines such as zAAP and zIIP introduced complexity since customers wanted to know how their workload made use of these new processors. SMF records 30 and 72 were modified in order to account for speciality processors.

**NOTE:** The zAAP processor originally was named Integrated Facility for Applications (IFA). SMF reports zAAP under an IFA construct.

The SMF30CPT field in SMF30 record reports on all standard CPU step time in hundredths of a second including enclave time, preemptible class SRB time, client SRB time. It also includes time consumed by zAAP/IIP eligible work running on standard processor. This situation may occur is a zAAP/zIIP is DISABLED ( varied offline ) from the Console or the zAAP/zIIP is busy processing instructions and not available to take on more work. For time spent on a zAAP, use SMF30\_TIME\_ON\_IFA. For time spent on a zIIP, use SMF30\_TIME\_ON\_ZIIP. Because the times used for this field are calculated from SRM service units (MSUs), all types of processors are included in CPU time or SRB time. The CPU time value contained in this field is comparable to the sum of fields SMF30CPT, SMF30\_TIME\_ON\_IFA and SMF30\_TIME\_ON\_ZIIP. If the zAAP or zIIP processors are faster than the general CP (under this situation the customer may be running subcapacity general CPs for cost licensing and zAAPs and zIIPs run at full speed) , the CPU time is then normalized to the time expected on the slower CP. Field SMF30ZNF contains the normalization factor used.

Note that a workload may generate different values for SMF30CPT if some zAAP eligible work runs on a standard processor. If a repeatable value is more desirable than the possible performance benefits of letting zAAP eligible work run on both zAAP and standard processors - specify IFAHONORPRIORITY=NO in the IEAOPTxx parm lib member. Interval records may show this number to be hundredth (1/100) of a second less than other related SMF30 fields (such as SMF30\_TIME\_IFA\_ON\_CP). This difference is due to rounding differentials while calculating delta values, and will not occur for jobstep end and/or job end metrics.

### Examples - SMF RECORD Types and associated fields

#### T30

SMF30\_TIME\_ON\_IFA..... CPU time spent on IFA in hundreds of a second includes Enclaves  
 SMF30\_ENCLAVE\_TIME\_ON\_IFA..... ENCLAVE time spent on IFA in hundreds of a second  
 SMF30\_DEP\_ENCLAVE\_TIME\_ON\_IFA.....Dependent enclave time of IFA in hundreds of a second  
 SMF30\_ENCLAVE\_TIME\_IFA\_ON\_CP..... IFA Enclave time spent on a standard CP in hundreds of second  
 SMF30\_DEP\_ENCLAVE\_TIME\_IFA\_ON\_CP... IFA dependent enclave time spent on a standard CP in 100s of a second.

SMF30\_TIME\_ON\_ZIIP.....Time spent on zIIP in hundreds of a second (including enclave time)  
 SMF30\_ENCLAVE\_TIME\_ON\_ZIIP .....ENCLAVE Time spent on zIIP in hundreds of a second  
 SMF\_Dependc\_TIME\_ON\_ZIIP.....Dependent enclave time spent on zIIP in hundreds of a second  
 SMF30\_TIME\_ZIIP\_ON\_CP.....CP time spent on CP in hundreds of a second (including enclave time)  
 SMF30\_ENCLAVE\_TIME\_ZIIP\_ON\_CP...Enclave time spent on CP in hundreds of a second

#### T72

SMF72PRF.....ZAAPs/ZIIPs Available  
 SMF72CIE.....Independent enclave total transaction active time for enclaves that originated on this system  
 SMF72CXE.....Exported enclave total transaction time ( in 1024-microseconds)  
 SMF72ECT.....CPU time consumed for an address space or enclave while dispatching priority was temporarily raised because the work held a resource that other work needed (in 1024 microseonds units)