

Methodology for Process and Image Computation.

Amlal El Mahrouss
amlal@nekernel.org

December 2025
Last Edited: January 2026

Abstract

CoreProcessScheduler governs how the scheduling backend and policy of the system works, It is the common gateway for schedulers inside NeKernel based systems.

1 I: Introduction.

CoreProcessScheduler (now referred as CPS) serves as the foundation between the scheduler backend and system. It takes care of process life-cycle management, team-based process grouping, and affinity-based CPU based allocation to mention the least.

1.1 II: The Affinity System.

Processes are given CPU time affinity hints using an affinity kind type, these hints help the scheduler run or adjust the current process.

1.2 Illustration: The Affinity Kind.

The following sample is C++ code. The smaller the value, the more critical the process.

```
1 enum struct AffinityKind : Int32 {  
2     kRealTime      = 100,  
3     kVeryHigh      = 150,  
4     kHigh          = 200,  
5     kStandard      = 1000,  
6     kLowUsage      = 1500,  
7     kVeryLowUsage  = 2000,  
8 };
```

1.3 III: The Team Domain System.

The team system holds process metadata for the backend scheduler to run on. It holds methods and fields for backend specific operations. One implementation of such team is the UserProcessTeam object inside NeKernel.

1.4 Illustration: Team System.

The following sample is used to hold team metadata. This is part of the NeKernel source tree.

```

1 class UserProcessTeam final {
2 public:
3     explicit UserProcessTeam();
4     ~UserProcessTeam() = default;
5
6     NE_COPY_DEFAULT(UserProcessTeam)
7
8     Array<UserProcess, kSchedProcessLimitPerTeam>& AsArray();
9     Ref<UserProcess>&                               AsRef();
10    ProcessID&                                       Id() noexcept;
11
12 public:
13     UserProcessArray mProcessList;
14     UserProcessRef   mCurrentProcess;
15     ProcessID        mTeamId{0};
16     ProcessID        mProcessCur{0};
17 };

```

1.5 IV: The Process Image System

The ‘ProcessImage’ container is a system designed to contain process data, and metadata, its purpose comes from the need to separate data and code. Such usage is useful for validation and vetability. This approach helps separate concerns and give modularity to the system, as the image and process structure are not mixed together.

1.6 Illustration: Process Image System.

The following sample is a container used to hold process data and metadata. This is part of the NeKernel source tree.

```

1 using ImagePtr = VoidPtr;
2
3 struct ProcessImage final {
4     explicit ProcessImage() = default;
5
6 private:
7     friend UserProcess;
8     friend KernelTask;
9
10    friend class UserProcessScheduler;
11
12    ImagePtr fCode;
13    ImagePtr fBlob;
14
15 public:
16     Bool HasCode() const { return this->fCode != nullptr; }
17
18     Bool HasImage() const { return this->fBlob != nullptr; }
19
20     ErrorOr<ImagePtr> LeakImage() {
21         if (this->fCode) {
22             return ErrorOr<ImagePtr>{this->fCode};
23         }
24     }
25
26     ErrorOr<ImagePtr> LeakBlob() {
27         if (this->fBlob) {
28             return ErrorOr<ImagePtr>{this->fBlob};
29         }
30     }
31 };

```

2 V: Conclusion.

The CPS is a system with provable domains and separation of computation, it governs how a Process Domain shall compute its child processes, it does not provide the algorithms. Which is why, one scheduler backend (such as the UserProcessScheduler) takes care of user process scheduling and fairness.

3 References

1. CoreProcessScheduler.h (2025), github.com
2. NeKernel.org (2025), nekernel.org
3. Scheduling: Introduction (2012), pages.cs.wisc.edu
4. Processor Affinity, Multiple CPU Scheduling (2003), tmurgent.com