

Architectural ideas behind OS/2 personality

Compatibility with OS/2 (Intel)

osFree combines OS/2 (Intel) ideas with those from OS/2 (PowerPC). It must run programs using OS/2 API, both console and PM ones. So, on Intel architecture, it must be binary compatible with IBM's OS/2. But on other architectures it will have different ABI ¹⁾, but the same portable API ²⁾.

Also, like it was in OS/2 (PowerPC), the binary compatibility with OS/2 (Intel) is possible (even on non-Intel architectures). It includes the translator of PowerPC instructions to Intel ones. The Instruction Set Translator (IST) is a special DLL. It could be made on the base of QEMU binary translation.

LX executable format is unlikely to be used on other architectures, because it is tied to the Intel one. So, most possibly, the ELF format will be used. It is cross-platform, extensible and simple. It supports both a 32-bit and a 64-bit variant. IBM Microkernel used ELF too. And more, IBM used an extended ELF format with special binary sections semantics. It supported inline resources, imports and exports, for example. So, it was possible to use UNIX-style shared objects, as well as OS/2-style DLL's.

So, the userland must be compatible with OS/2 (Intel). This includes some support of 16-bit applications. Although there is a very little of pure 16-bits apps (like cmd.exe and hiew.exe), but some 16-bit API's are widely used today. For example, Kbd/Mou/Vio API's (i.e., Console API) are still 16-bits in OS/2 (Intel), though on OS/2 (PowerPC) they are made 32-bits, because binary compatibility makes no sense on other arch's (only source one makes sense).

TODO: Thinking. Getting rid of thunks in 32-bits executables calling the 16-bit API's.

Kernel-related parts of OS/2

The most of 'kernel' area was decided to be moved to userland of L4 microkernel. So, some services, commonly meant as kernel parts, like filesystems, drivers, memory management, thread scheduling etc. are moved outside the kernel. So, they are not significantly differ from usual applications, and confined inside their address spaces.

The OS/2 personality main server (or simply "OS/2 Server")

Filesystem server

Installable file systems (IFS)

Filesystem router

IFS helpers

Exec server

Virtual Memory Arenas management

Instalable eXecutable Formats (IXF)

Shared memory management

Multiple Virtual Machines (MVM) server

A historical note

The term “MVM” goes from OS/2 (PowerPC). It replaces the term “MVDM” (Multiple Virtual **DOS** Machines). The PowerPC platform could not so easily emulate the Intel processor real mode, so it required more emulation. Though, we must say that it worked just fine, fast, and theoretically may be used not for DOS only. (IBM had plans to provide a binary compatibility with OS/2 (Intel), on the base of the same component which was used to run DOS Apps).

MVM server

The MVM server is a central server of the MVM personality – the infrastructure for running multiple virtual machines on top of L4 microkernel. It is almost separate from OS/2 personality, but can be controlled by OS/2 programs via DosOpenVDD/DosCloseVDD/DosRequestVDD API's.

So, the MVM server exposes some interfaces to other OS personalities to be controlled by them. Also, it starts VM's, which are executed in the context of a VM monitor, running a guest OS (DOS, for example). The VM environment is defined by VDD's³⁾ loaded. The MVM server loads VDD's, which are a kind of plugins. Also, it exports some helper API's for VDD's via MVDM.DLL.

Virtual Device Drivers (VDD's)

A VDD is like a plugin for MVM Server, and it can communicate with outside programs via request API (DosRequestVDD for OS/2 programs). On the other side, from the VM point of view, the VDD emulates some BIOS services/Option ROM's, a hardware devices of guest platform. It installs interrupt handlers, catches I/O ports or MMIO registers accesses from inside the guest OS. A VDD implements its services based on VDD helpers API, served by the MVM server.

Originally, the MV(D)M were used to emulate a 8086 machine with BIOS and DOS/Win 3.1. But now it is a trend for many OS'es to have the kernel virtual machines, like kvm in Linux or VirtualPC in WinNT. Our MVM personality is our solution of the same kind, but it not so monolithic like qemu/kvm – it decomposed to several parts, in best IBM solutions design traditions. The VDD's are like plugins for MVM server, allowing to extend the MVM environment. They're loadable modules (DLL's).

Also, as we're told previously, to share the screen with OS/2 apps and apps of other personalities, IBM

created the solution of “Seamless WinOS/2”. This is the possibility for Win 3.1 programs to share the same screen with OS/2 PM ones.

Virtual Machine Monitor

The DOS emulation kernel (doskrnl)

Instruction Set Translator (IST)

VM86 on Intel, and Hardware-assisted virtualization

Microkernels as Hypervisors

Other OS's personalities

"The one Ring to Bind Them all" (OS/2 as an integration platform for different types of applications)

The current OS/2 personality prototype

osFree PM

The osFree PM is an osFree version of FreePM, which was began by Evgeny Kotsuba, and then abandoned. After that, we did some changes to it, and called it “osFree PM”, to avoid name collision in case Evgeny will continue his initial version.

- [Docs about FreePM, by Evgeny](#), now inaccessible
- [FreePM on SourceForge](#) Forum and code repository
- [Documentation](#) on PM implementation based on FreePM
- [Some parts](#) of FreePM docs written by Evgeny.

Graphical Program Interface (GPI)

The GPI ⁴⁾ is the graphical engine of Presentation Manager. It is based on PM GRE ⁵⁾. The GPI/GRE is the counterparts of Windows GDI. The GPI/GRE pair is designed as an enhanced version of Windows Graphics Engine. Contrary to Windows, they are decomposed to two layers. The GPI is the high-level layer.

The Windows programs API's operate directly on DC ⁶⁾. OS/2 PM is redesigned, so Programs operate on PS ⁷⁾, not the DC. The DC is something related to the instance of graphics device (the video screen, a window or printer).

The PS is a higher level abstraction. It can be treated as a canvas in other graphics libraries. It

maintain such things as current background/foreground color, a palette, a brush shape, a current font etc. Also, the PS can be associated/deassociated to/from the DC, or migrate from one DC to another. This gives a flexibility to graphics API.

The GPI functions have the "Gpi" prefix and reside in PMGPI.DLL. They are built on top of GRE functions.

Graphics Runtime Environment (OS/2 PM GRE) and Presentation Drivers

The GRE is a low-level graphics API which operates DC's. It handles all kinds of graphics devices like video cards and printers.

The GPI contains an array of pointers called the Dispatch Table. This table stores pointers to the most lowlevel GRE functions.

The graphics device drivers are called also "presentation drivers". The presentation drivers get an instance of Dispatch Table, and hook in GRE by installing their own functions pointers in the Dispatch Table, and preserving the old ones.

The GRADD model

Video Protected Mode Interface (PMI)

VIO/KBD/MOU (Console API)

VIO and BVH's (Base Video Handlers)

Framebuffer Interface

Virtual Framebuffer over GRADD driver

BVH and PM GRE on top of a Framebuffer interface

The current osFree PM prototype

1)

Application Binary Interface

2)

Application Program Interface

3)

Virtual Device Drivers

4)

Graphical Program Interface

5)

Graphics Runtime Engine

[6\)](#)

Device Context

[7\)](#)

Presentation Space

From:

<https://ftp.osfree.org/doku/> - **osFree wiki**

Permanent link:

<https://ftp.osfree.org/doku/doku.php?id=en:docs:os2:architecture&rev=1402076709>Last update: **2014/06/06 17:45**