

MVM/DOS personality

This is an infrastructure for running VM's with unmodified OS's

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 the VDD's¹⁾ 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.

This is done with a special VDD's, like VVIDEO (VVGGA, VSVGA, etc) implementing a video mode support in a DOS window (or fullscreen). The DOS window is implemented as a special DLL (pmvdmp.dll) which was loaded by PM. It communicates with a VM via DosRequestVDD and implement the video functions via GPI calls (for windowed mode, so it is a PM application, based on VIO Shield (pmviop.dll), working via BVH drivers, like bvhwndw.dll for windowed OS/2 and DOS sessions, or bvhvga.dll+bvhsvga for fullscreen OS/2 or DOS sessions).

For windowed WinOS/2 sessions, it existed the solution of using a so-called PM shield (seamless.dll) and WinOS/2 shield (winsheld.exe). The 1st one is an “avatar” of Windows application in OS/2 PM world. And vice versa, the WinOS/2 shield is a representative of a PM app in WinOS/2 world.

Also, the second solution exists, based on GRADD video driver model. It works via VVMI (vmanwin.sys in the Intel OS/2). It is the VDD related to the communication of windows video driver (ifgdi2vm.driv for fullscreen, isgdi2vm.driv for seamless mode) with GRADD's VMAN²⁾. The special thread in VMAN polls the VVMI driver to communicate with Windows driver. So, the Windows driver is a generic one,

but it communicates with a “real” driver. This results in WinOS/2 and OS/2 PM shared the same screen using access to a common video driver. See the GRADD-related section for more details about multiple graphics engines sharing the same screen/video driver.

Virtual Machine Monitor

The VMM ³⁾ is a program implementing the environment of guest hardware platform, with help of VDD's. It maintains the contexts of all VM's, and handles the traps redirecting them to the needed VDD, loads the IST ⁴⁾ for different processor instruction sets, utilizes the hardware emulation features of the CPU, like VM86, AMD SVM, Intel VT-x etc.

It maintains the address space layout of a VM application, loads a firmware (for BIOS, the SeaBIOS can be used) and DOS emulation kernel.

The DOS emulation kernel (doskrnl)

The DOS emulation kernel is a special rehosted DOS kernel working via OS/2 (or PN ⁵⁾) services. For example, file system API's of int 21h are implemented via OS/2 (or PN) file API's. Access to OS/2 API's done via SVC API (supervisor call?) which is trap of HLT instruction followed by call number and inverted call number.

Instruction Set Translator (IST)

The IST is a DLL, emulating the instructions of Guest hardware via Host CPU instructions. It exports a set of entry points, each corresponding the emulated instruction.

The similar component exist in QEMU – but it is linked statically with the emulator binary.

VM86 on Intel, and Hardware-assisted virtualization

Some processors implement special compatibility modes (VM86 allows creation of special task in protected mode, which emulates a virtual i8086 processor) or special instruction to assist the Virtual Machine Monitors creation. (Like “hypercall” to change context to a hypervisor, to execute its service and exit “hypervisor” mode). Also, the very new processors implement the IOMMU (a hardware support for sharing a hardware).

These extensions can be used to run unmodified OS'es on top of a hypervisor (it is supported in newer versions of Xen, VBox, VMWare, VPC).

Microkernels as Hypervisors

Microkernels and Hypervisors are very similar things. Microkernels implement similar features. For example, the Fiasco.OC microkernel supports SVM and VT-x and allows to run unmodified Linux in very thin VM's. This feature can be utilized in our MVM personality too.

SVC API

Supervisor calls API aimed to call host services from VM.

Ordinal	Description
0	
1	
2	Terminate VDM and return code to parent process from AX
3	
4	
5	
6	
7	
8	
9	
10	
11	
12	
13	
14	
15	
16	
17	
18	
19	
20	
21	
22	
23	
24	
25	
26	
27	
28	
29	
30	
31	
32	
33	
34	
35	
36	
37	
38	
39	
40	

Ordinal	Description
41	
42	
43	
44	
45	
46	
47	
48	
49	
50	
51	
52	
53	
54	
55	
56	
57	
58	
59	
60	
61	
62	
63	
64	
65	
66	
67	
68	
69	
70	
71	
72	
73	
74	
75	
76	
77	
78	
79	
80	
81	
82	
83	
84	
85	

1)

Virtual Device Drivers

2)

Video Manager

3)

Virtual Machine Monitor

4)

Instruction Set Translator

5)

Personality Neutral

From:

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

Permanent link:

<http://osfree.org/doku/doku.php?id=en:docs:general:mvm&rev=1700615850>

Last update: **2023/11/22 01:17**

