Debugging Live Migration

About Me

- Alexander Graf
- Freelance developer for SUSE and Freescale
- KVM and Qemu developer
- Server class PowerPC KVM port
- S390x Qemu guest support
- x86 Mac OS X in KVM
- Nested SVM
- ...

What is Live Migration

What is Live Migration

What is Live Migration

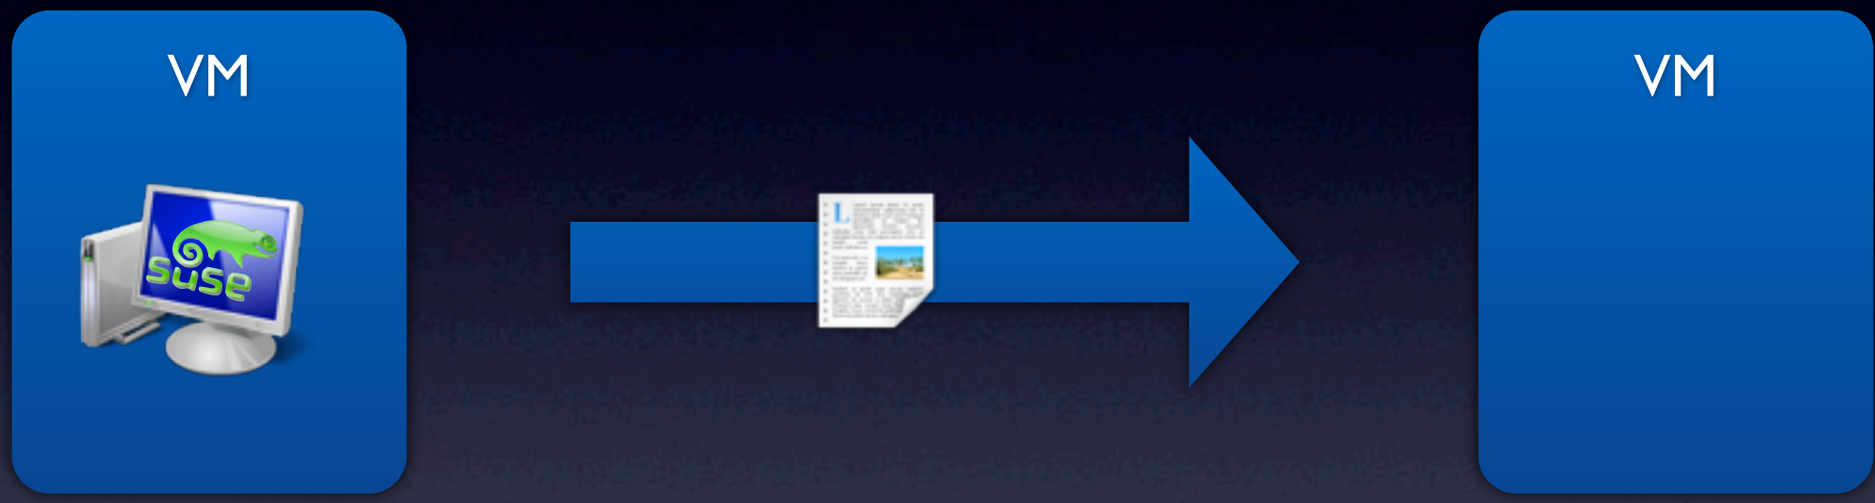

What is Live Migration

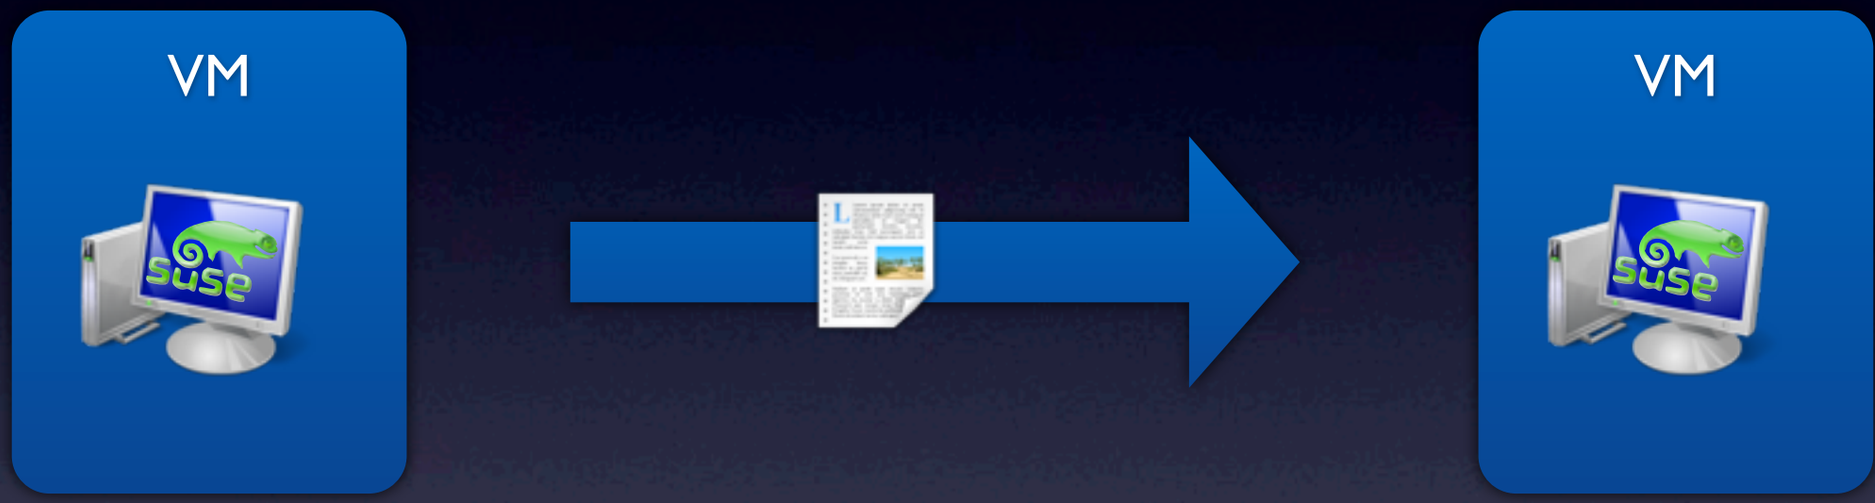

What is Live Migration

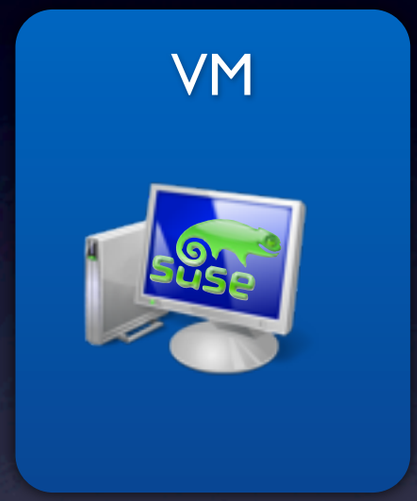

File Format

File Format

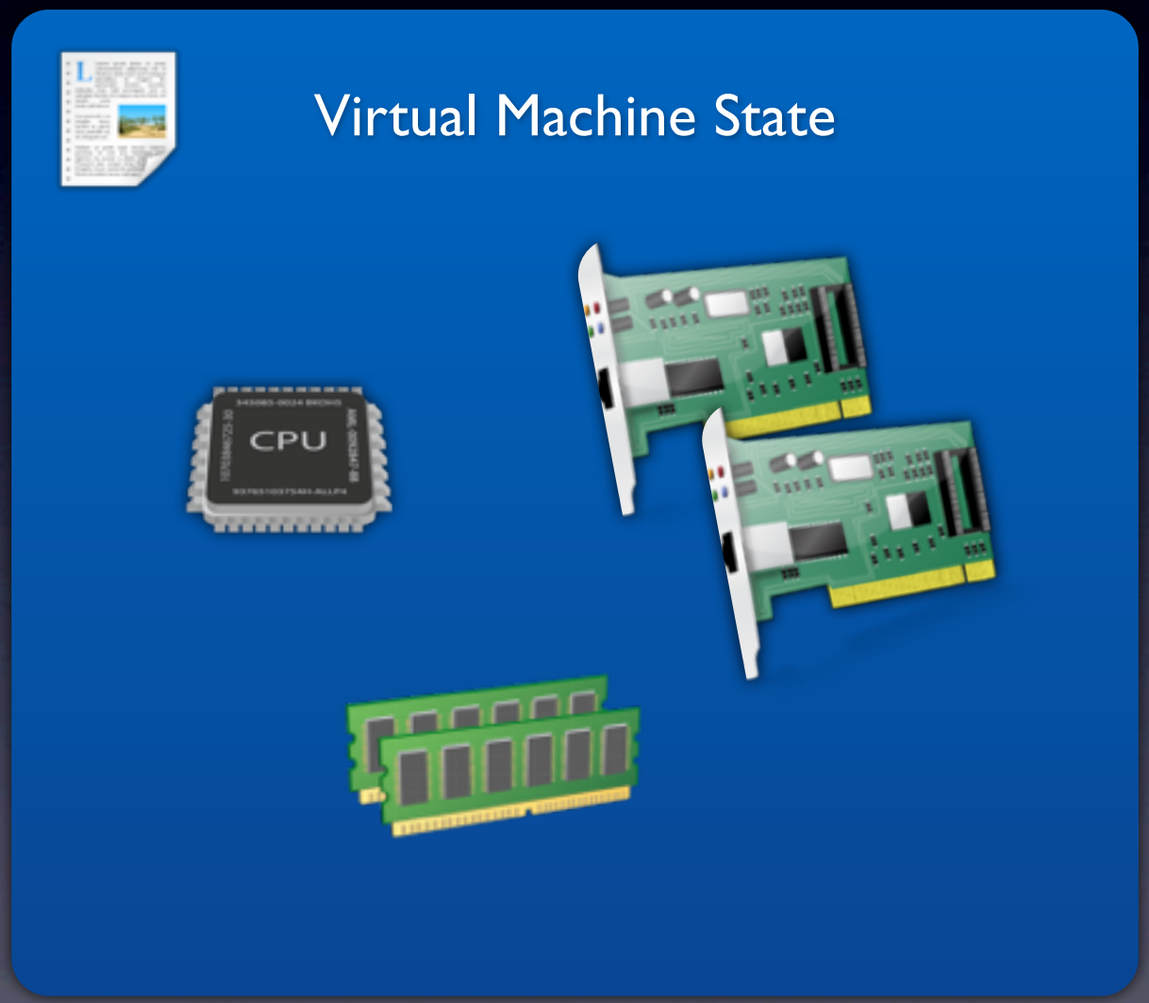

File Format

```
00000000  51 45 56 4d 00 00 00 03  01 00 00 00 03 03 72 61  |QEVM..........ra|
00000010  6d 00 00 00 00 00 00 00  04 00 00 00 00 08 89 40  |m..............@|
00000020  04 06 70 63 2e 72 61 6d  00 00 00 00 08 00 00 00  |..pc.ram........|
00000030  08 76 67 61 2e 76 72 61  6d 00 00 00 00 00 80 00  |.vga.vram.......|
                          [...]
```

File Format

```
00000000  51 45 56 4d 00 00 00 03  01 00 00 00 03 03 72 61  |QEVM..........ra|
00000010  6d 00 00 00 00 00 00 00  04 00 00 00 00 08 89 40  |m..............@|
00000020  04 06 70 63 2e 72 61 6d  00 00 00 00 08 00 00 00  |..pc.ram........|
00000030  08 76 67 61 2e 76 72 61  6d 00 00 00 00 00 80 00  |.vga.vram.......|
                          [...]
```

File Format

```
00000000  51 45 56 4d 00 00 00 03  01 00 00 00 03 03 72 61  |QEVM..........ra|
00000010  6d 00 00 00 00 00 00 00  04 00 00 00 00 08 89 40  |m..............@|
00000020  04 06 70 63 2e 72 61 6d  00 00 00 00 08 00 00 00  |..pc.ram........|
00000030  08 76 67 61 2e 76 72 61  6d 00 00 00 00 00 80 00  |.vga.vram.......|
                                [...]
```

File Magic

Always QEVM

File Format

```
00000000  51 45 56 4d 00 00 00 03  01 00 00 00 03 03 72 61  |QEVM..........ra|
00000010  6d 00 00 00 00 00 00 00  04 00 00 00 00 08 89 40  |m..............@|
00000020  04 06 70 63 2e 72 61 6d  00 00 00 00 08 00 00 00  |..pc.ram........|
00000030  08 76 67 61 2e 76 72 61  6d 00 00 00 00 00 80 00  |.vga.vram.......|
                              [...]
```

File Magic — Always QEVM

File Version — Always 3

File Format

```
00000000  51 45 56 4d 00 00 00 03  01 00 00 00 03 03 72 61  |QEVM..........ra|
00000010  6d 00 00 00 00 00 00 00  04 00 00 00 00 08 89 40  |m..............@|
00000020  04 06 70 63 2e 72 61 6d  00 00 00 00 08 00 00 00  |..pc.ram........|
00000030  08 76 67 61 2e 76 72 61  6d 00 00 00 00 00 80 00  |.vga.vram.......|
                                [...]
```

File Magic — Always QEVM

File Version — Always 3

Section Type — QEMU_VM_SECTION_START

File Format

```
00000000  51 45 56 4d 00 00 00 03  01 00 00 00 03 03 72 61  |QEVM..........ra|
00000010  6d 00 00 00 00 00 00 00  04 00 00 00 00 08 89 40  |m..............@|
00000020  04 06 70 63 2e 72 61 6d  00 00 00 00 08 00 00 00  |..pc.ram........|
00000030  08 76 67 61 2e 76 72 61  6d 00 00 00 00 00 80 00  |.vga.vram.......|
                                [...]
```

File Magic — Always QEVM

File Version — Always 3

Section Type — QEMU_VM_SECTION_START

Section ID — Section 3

File Format

```
00000000  51 45 56 4d 00 00 00 03  01 00 00 00 03 03 72 61  |QEVM..........ra|
00000010  6d 00 00 00 00 00 00 00  04 00 00 00 00 08 89 40  |m..............@|
00000020  04 06 70 63 2e 72 61 6d  00 00 00 00 08 00 00 00  |..pc.ram........|
00000030  08 76 67 61 2e 76 72 61  6d 00 00 00 00 00 80 00  |.vga.vram.......|
                                [...]
```

File Magic	Always QEVM
File Version	Always 3
Section Type	QEMU_VM_SECTION_START
Section ID	Section 3
Section Class Name Length	3 bytes

File Format

```
00000000  51 45 56 4d 00 00 00 03  01 00 00 00 03 03 72 61  |QEVM..........ra|
00000010  6d 00 00 00 00 00 00 00  04 00 00 00 00 08 89 40  |m..............@|
00000020  04 06 70 63 2e 72 61 6d  00 00 00 00 08 00 00 00  |..pc.ram........|
00000030  08 76 67 61 2e 76 72 61  6d 00 00 00 00 00 80 00  |.vga.vram.......|
                                [...]
```

File Magic	Always QEVM
File Version	Always 3
Section Type	QEMU_VM_SECTION_START
Section ID	Section 3
Section Class Name Length	3 bytes
Section Class Name	“ram”

File Format

```
00000000  51 45 56 4d 00 00 00 03  01 00 00 00 03 03 72 61  |QEVM..........ra|
00000010  6d 00 00 00 00 00 00 00  04 00 00 00 00 08 89 40  |m..............@|
00000020  04 06 70 63 2e 72 61 6d  00 00 00 00 08 00 00 00  |..pc.ram........|
00000030  08 76 67 61 2e 76 72 61  6d 00 00 00 00 00 80 00  |.vga.vram.......|
                                [...]
```

File Magic	Always QEVM
File Version	Always 3
Section Type	QEMU_VM_SECTION_START
Section ID	Section 3
Section Class Name Length	3 bytes
Section Class Name	“ram”
Section Instance ID	RAM Instance 0

File Format

```
00000000  51 45 56 4d 00 00 00 03  01 00 00 00 03 03 72 61  |QEVM..........ra|
00000010  6d 00 00 00 00 00 00 00  04 00 00 00 00 08 89 40  |m..............@|
00000020  04 06 70 63 2e 72 61 6d  00 00 00 00 08 00 00 00  |..pc.ram........|
00000030  08 76 67 61 2e 76 72 61  6d 00 00 00 00 00 80 00  |.vga.vram.......|
                                [...]
```

File Magic	Always QEVM
File Version	Always 3
Section Type	QEMU_VM_SECTION_START
Section ID	Section 3
Section Class Name Length	3 bytes
Section Class Name	“ram”
Section Instance ID	RAM Instance 0
Section Version ID	RAM Version 4

File Format

```
00000000  51 45 56 4d 00 00 00 03  01 00 00 00 03 03 72 61  |QEVM..........ra|
00000010  6d 00 00 00 00 00 00 00  04 00 00 00 00 08 89 40  |m..............@|
00000020  04 06 70 63 2e 72 61 6d  00 00 00 00 08 00 00 00  |..pc.ram........|
00000030  08 76 67 61 2e 76 72 61  6d 00 00 00 00 00 80 00  |.vga.vram.......|
                                [...]
```

File Magic	Always QEVM
File Version	Always 3
Section Type	QEMU_VM_SECTION_START
Section ID	Section 3
Section Class Name Length	3 bytes
Section Class Name	“ram”
Section Instance ID	RAM Instance 0
Section Version ID	RAM Version 4
Section Data	RAM Bulk Data

File Format

```
                                  [...]
03c78240  00 00 00 00 00 00 04 00  00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00  03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75  73 65 00 00 00 00 00 00  |...vmmouse......|
                                  [...]
```

File Format

```
                                        [...]
03c78240  00 00 00 00 00 00 04 00   00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00   03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75   73 65 00 00 00 00 00 00  |...vmmouse......|
                                        [...]
```

Section Type

QEMU_VM_SECTION_FULL

File Format

```
                                        [...]
03c78240  00 00 00 00 00 00 04 00  00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00  03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75  73 65 00 00 00 00 00 00  |...vmmouse......|
                                        [...]
```

Section Type — QEMU_VM_SECTION_FULL

Section ID — Section 0x17

File Format

```
                                                [...]
03c78240  00 00 00 00 00 00 04 00  00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00  03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75  73 65 00 00 00 00 00 00  |...vmmouse......|
                                                [...]
```

Section Type — QEMU_VM_SECTION_FULL

Section ID — Section 0x17

Section Class Name Length — 5 bytes

File Format

```
                                        [...]
03c78240  00 00 00 00 00 00 04 00   00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00   03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75   73 65 00 00 00 00 00 00  |...vmmouse......|
                                        [...]
```

Section Type	QEMU_VM_SECTION_FULL
Section ID	Section 0x17
Section Class Name Length	5 bytes
Section Class Name	“pckbd”

File Format

```
                                        [...]
03c78240  00 00 00 00 00 00 04 00   00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00   03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75   73 65 00 00 00 00 00 00  |...vmmouse......|
                                        [...]
```

Section Type	QEMU_VM_SECTION_FULL
Section ID	Section 0x17
Section Class Name Length	5 bytes
Section Class Name	“pckbd”
Section Instance ID	pckbd Instance 0

File Format

```
                                                [...]
03c78240  00 00 00 00 00 00 04 00  00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00  03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75  73 65 00 00 00 00 00 00  |...vmmouse......|
                                                [...]
```

Section Type	QEMU_VM_SECTION_FULL
Section ID	Section 0x17
Section Class Name Length	5 bytes
Section Class Name	“pckbd”
Section Instance ID	pckbd Instance 0
Section Version ID	pckbd Version 3

File Format

```
                                                [...]
03c78240  00 00 00 00 00 00 04 00  00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00  03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75  73 65 00 00 00 00 00 00  |...vmmouse......|
                                                [...]
```

Section Type	QEMU_VM_SECTION_FULL
Section ID	Section 0x17
Section Class Name Length	5 bytes
Section Class Name	“pckbd”
Section Instance ID	pckbd Instance 0
Section Version ID	pckbd Version 3
Section Data	pckbd Bulk Data

File Format

```
                                                [...]
03c78240  00 00 00 00 00 00 04 00  00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00  03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75  73 65 00 00 00 00 00 00  |...vmmouse......|
                                                [...]
```

Len	
Section Type	QEMU_VM_SECTION_FULL
Section ID	Section 0x17
Section Class Name Length	5 bytes
Section Class Name	“pckbd”
Section Instance ID	pckbd Instance 0
Section Version ID	pckbd Version 3
Section Data	pckbd Bulk Data

File Format

```
                                                [...]
03c78240  00 00 00 00 00 00 04 00  00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00  03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75  73 65 00 00 00 00 00 00  |...vmmouse......|
                                                [...]
```

Len		
1	Section Type	QEMU_VM_SECTION_FULL
	Section ID	Section 0x17
	Section Class Name Length	5 bytes
	Section Class Name	“pckbd”
	Section Instance ID	pckbd Instance 0
	Section Version ID	pckbd Version 3
	Section Data	pckbd Bulk Data

File Format

```
                                                [...]
03c78240  00 00 00 00 00 00 04 00  00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00  03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75  73 65 00 00 00 00 00 00  |...vmmouse......|
                                                [...]
```

Len		
1	Section Type	QEMU_VM_SECTION_FULL
4	Section ID	Section 0x17
	Section Class Name Length	5 bytes
	Section Class Name	“pckbd”
	Section Instance ID	pckbd Instance 0
	Section Version ID	pckbd Version 3
	Section Data	pckbd Bulk Data

File Format

```
                                        [...]
03c78240  00 00 00 00 00 00 04 00   00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00   03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75   73 65 00 00 00 00 00 00  |...vmmouse......|
                                        [...]
```

Len		
1	Section Type	QEMU_VM_SECTION_FULL
4	Section ID	Section 0x17
1	Section Class Name Length	5 bytes
	Section Class Name	“pckbd”
	Section Instance ID	pckbd Instance 0
	Section Version ID	pckbd Version 3
	Section Data	pckbd Bulk Data

File Format

```
                                                [...]
03c78240  00 00 00 00 00 00 04 00  00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00  03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75  73 65 00 00 00 00 00 00  |...vmmouse......|
                                                [...]
```

Len		
1	Section Type	QEMU_VM_SECTION_FULL
4	Section ID	Section 0x17
1	Section Class Name Length	5 bytes
✔	Section Class Name	“pckbd”
	Section Instance ID	pckbd Instance 0
	Section Version ID	pckbd Version 3
	Section Data	pckbd Bulk Data

File Format

```
                                        [...]
03c78240  00 00 00 00 00 00 04 00   00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00   03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75   73 65 00 00 00 00 00 00  |...vmmouse......|
                                        [...]
```

Len		
1	Section Type	QEMU_VM_SECTION_FULL
4	Section ID	Section 0x17
1	Section Class Name Length	5 bytes
✔	Section Class Name	“pckbd”
4	Section Instance ID	pckbd Instance 0
	Section Version ID	pckbd Version 3
	Section Data	pckbd Bulk Data

File Format

```
                                                [...]
03c78240  00 00 00 00 00 00 04 00  00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00  03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75  73 65 00 00 00 00 00 00  |...vmmouse......|
                                                [...]
```

Len		
1	Section Type	QEMU_VM_SECTION_FULL
4	Section ID	Section 0x17
1	Section Class Name Length	5 bytes
✔	Section Class Name	“pckbd”
4	Section Instance ID	pckbd Instance 0
4	Section Version ID	pckbd Version 3
	Section Data	pckbd Bulk Data

File Format

```
                                                [...]
03c78240  00 00 00 00 00 00 04 00  00 00 17 05 70 63 6b 62  |............pckb|
03c78250  64 00 00 00 00 00 00 00  03 00 1c 43 00 04 00 00  |d..........C....|
03c78260  00 18 07 76 6d 6d 6f 75  73 65 00 00 00 00 00 00  |...vmmouse......|
                                                [...]
```

Len		
1	Section Type	QEMU_VM_SECTION_FULL
4	Section ID	Section 0x17
1	Section Class Name Length	5 bytes
✔	Section Class Name	“pckbd”
4	Section Instance ID	pckbd Instance 0
4	Section Version ID	pckbd Version 3
?	Section Data	pckbd Bulk Data

File Format

```
                                                [...]
03c78240   00 00 00 00 00 00 04 00   00 00 17 05 70 63 6b 62   |............pckb|
03c78250   64 00 00 00 00 00 00 00   03 00 1c 43 00 04 00 00   |d..........C....|
03c78260   00 18 07 76 6d 6d 6f 75   73 65 00 00 00 00 00 00   |...vmmouse......|
                                                [...]
```

Len		
1	Section Type	QEMU_VM_SECTION_FULL
4	Section ID	Section 0x17
1	Section Class Name Length	5 bytes
✔	Section Class Name	“pckbd”
4	Section Instance ID	pckbd Instance 0
4	Section Version ID	pckbd Version 3
?	Section Data	pckbd Bulk Data

File Format

```
00 1c 43 00
```

File Format

00 1c 43 00

```
static const VMStateDescription vmstate_kbd = {
    .name = "pckbd",
    .version_id = 3,
    .minimum_version_id = 3,
    .minimum_version_id_old = 3,
    .fields      = (VMStateField []) {
        VMSTATE_UINT8(write_cmd, KBDState),
        VMSTATE_UINT8(status, KBDState),
        VMSTATE_UINT8(mode, KBDState),
        VMSTATE_UINT8(pending, KBDState),
        VMSTATE_END_OF_LIST()
    }
};
```

File Format

00 1c 43 00

```
static const VMStateDescription vmstate_kbd = {
    .name = "pckbd",
    .version_id = 3,
    .minimum_version_id = 3,
    .minimum_version_id_old = 3,
    .fields      = (VMStateField []) {
        VMSTATE_UINT8(write_cmd, KBDState),
        VMSTATE_UINT8(status, KBDState),
        VMSTATE_UINT8(mode, KBDState),
        VMSTATE_UINT8(pending, KBDState),
        VMSTATE_END_OF_LIST()
    }
};
```

File Format

```
static const VMStateDescription vmstate_kbd = {
    .name = "pckbd",
    .version_id = 3,
    .minimum_version_id = 3,
    .minimum_version_id_old = 3,
    .fields      = (VMStateField []) {
        VMSTATE_UINT8(write_cmd, KBDState),
        VMSTATE_UINT8(status, KBDState),
        VMSTATE_UINT8(mode, KBDState),
        VMSTATE_UINT8(pending, KBDState),
        VMSTATE_END_OF_LIST()
    }
};
```

File Format

VMSD

File Format

VMSD

File Format

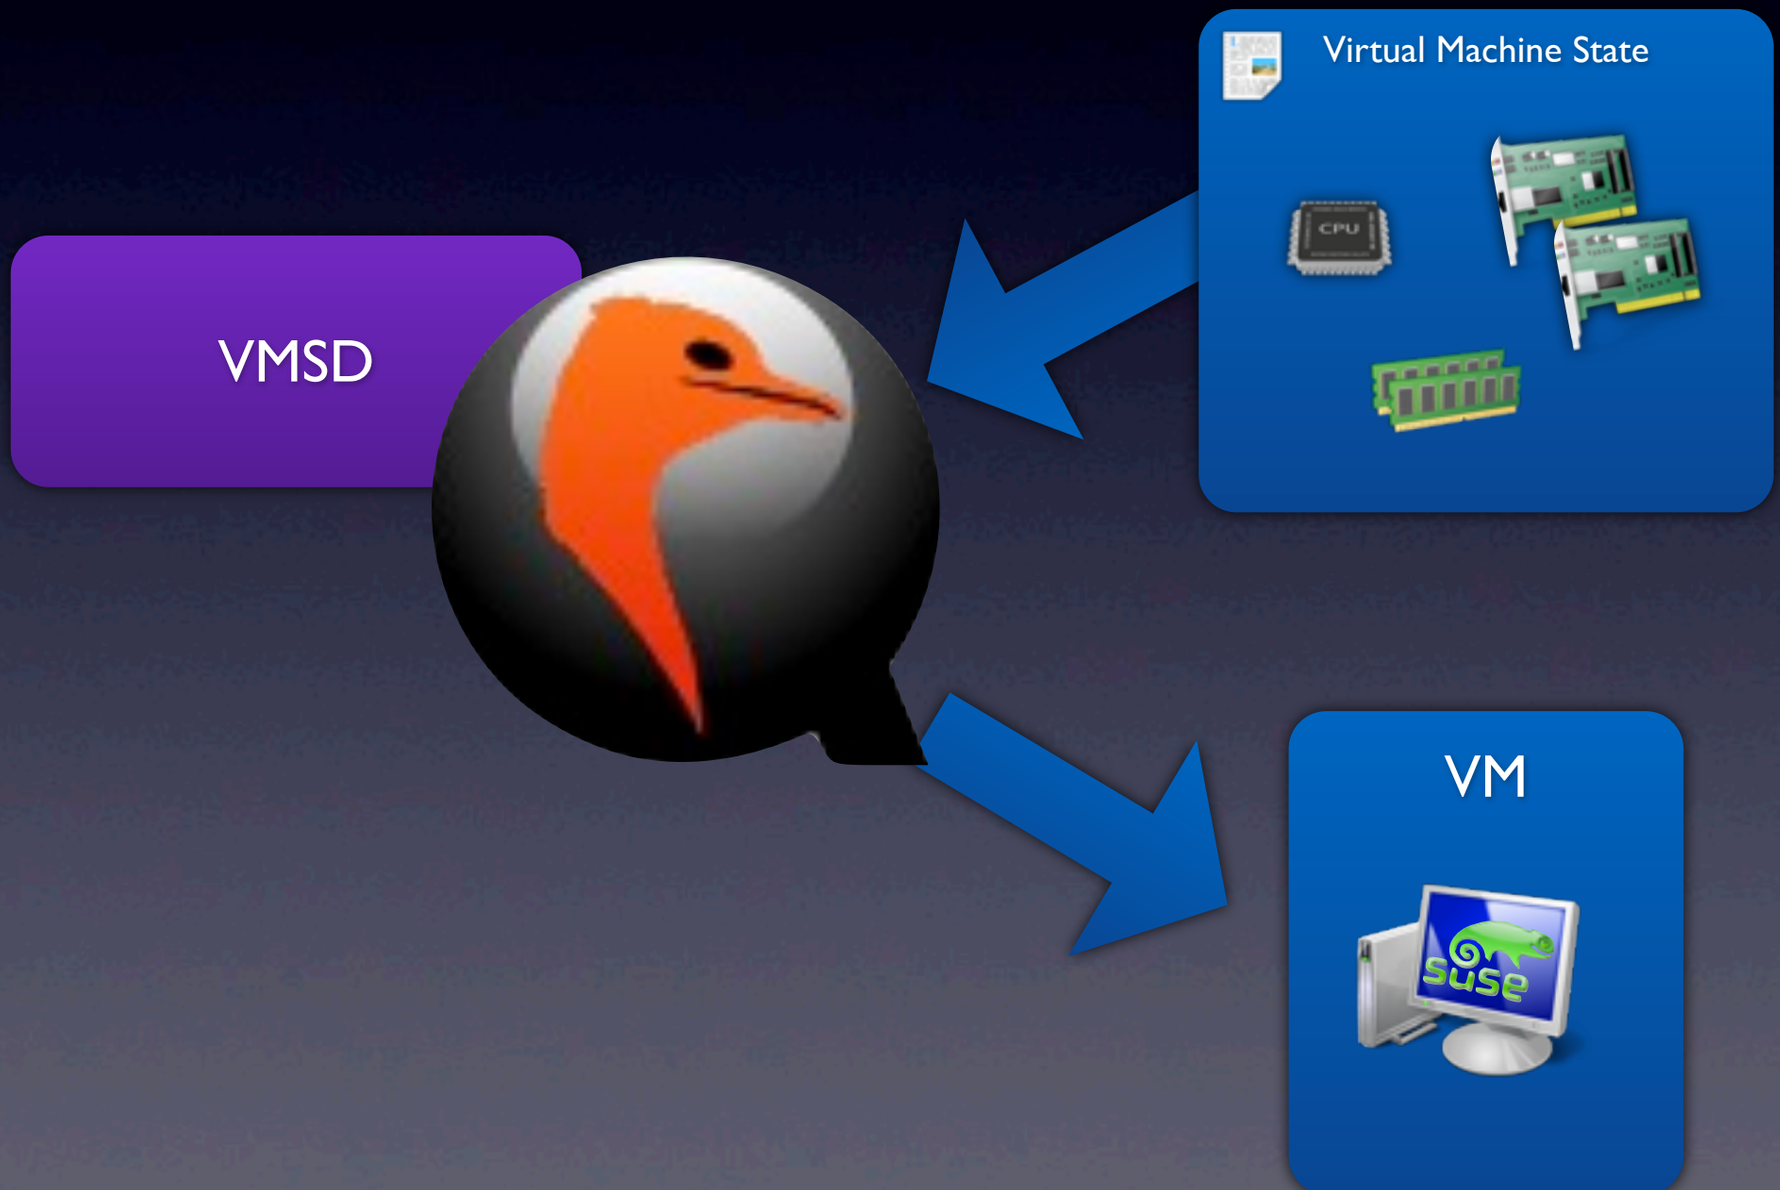

File Parsing

File Parsing

00 1c 43 00

```
static const VMStateDescription vmstate_kbd = {
    .name = "pckbd",
    .version_id = 3,
    .minimum_version_id = 3,
    .minimum_version_id_old = 3,
    .fields      = (VMStateField []) {
        VMSTATE_UINT8(write_cmd, KBDState),
        VMSTATE_UINT8(status, KBDState),
        VMSTATE_UINT8(mode, KBDState),
        VMSTATE_UINT8(pending, KBDState),
        VMSTATE_END_OF_LIST()
    }
};
```

```
VMSTATE_UINT8(write_cmd, KBDState),
VMSTATE_UINT8(status, KBDState),
VMSTATE_UINT8(mode, KBDState),
VMSTATE_UINT8(pending, KBDState),
```

```
{
    "name": "pckbd",
    "instance_id": 0,
    "vmsd_name": "pckbd",
    "versions": {
        "3": {
            "fields": [
                {
                    "name": "kbd", "size": 56, "type": "struct",
                    "struct": {
                        "version_id": 3,
                        "fields": [
                            { "name": "write_cmd", "size": 1, "type": "uint8" },
                            { "name": "status", "size": 1, "type": "uint8" },
                            { "name": "mode", "size": 1, "type": "uint8" },
                            { "name": "pending", "size": 1, "type": "uint8" }
                        ]
                    }
                }
            ]
        }
    }
},
```

```
VMSTATE_UINT8(write_cmd, KBDState),
VMSTATE_UINT8(status, KBDState),
VMSTATE_UINT8(mode, KBDState),
VMSTATE_UINT8(pending, KBDState),
```

```
{
    "name": "pckbd",
    "instance_id": 0,
    "vmsd_name": "pckbd",
    "versions": {
        "3": {
            "fields": [
                {
                    "name": "kbd", "size": 56, "type": "struct",
                    "struct": {
                        "version_id": 3,
                        "fields": [
                            { "name": "write_cmd", "size": 1, "type": "uint8" },
                            { "name": "status", "size": 1, "type": "uint8" },
                            { "name": "mode", "size": 1, "type": "uint8" },
                            { "name": "pending", "size": 1, "type": "uint8" }
                        ]
                    }
                }
            ]
        }
    }
},
```


JSON

File Parsing

File Parsing

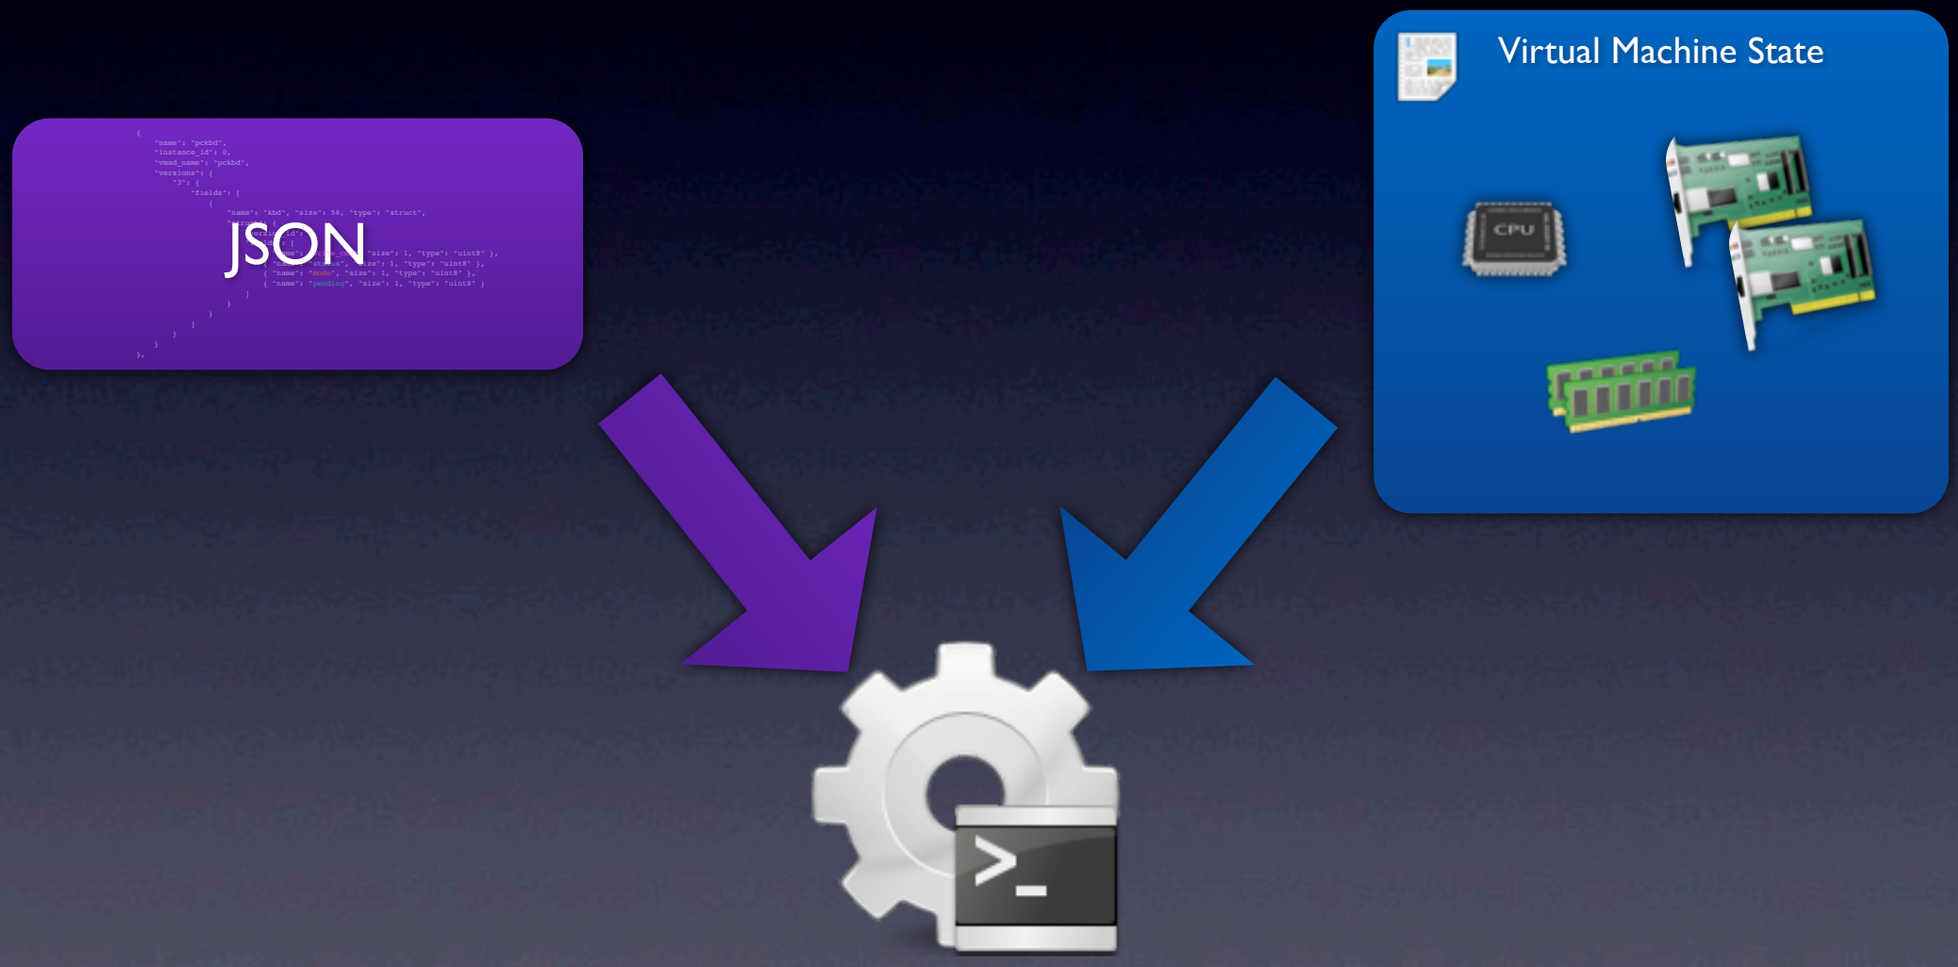

File Parsing

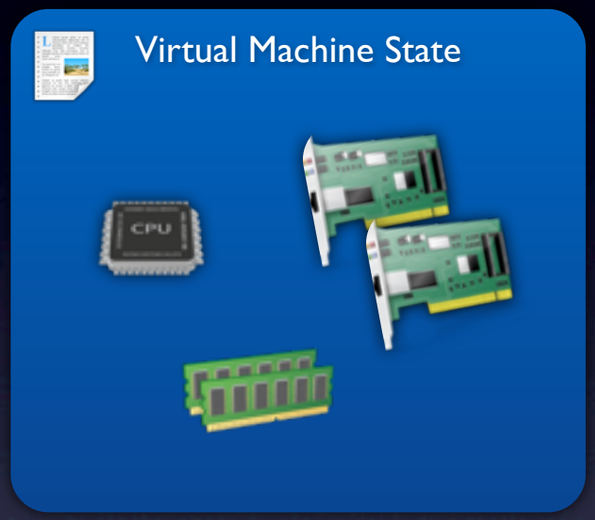

File Parsing

```
"pckbd (23)": {
    "kbd": {
        "write cmd": "0x00",
        "status": "0x1c",
        "mode": "0x43",
        "pending": "0x00"
    }
},
```

Debug Device

Debug Device

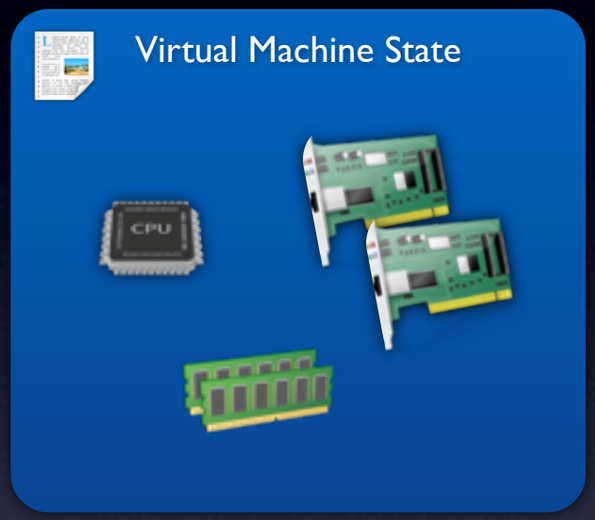

Debug Device

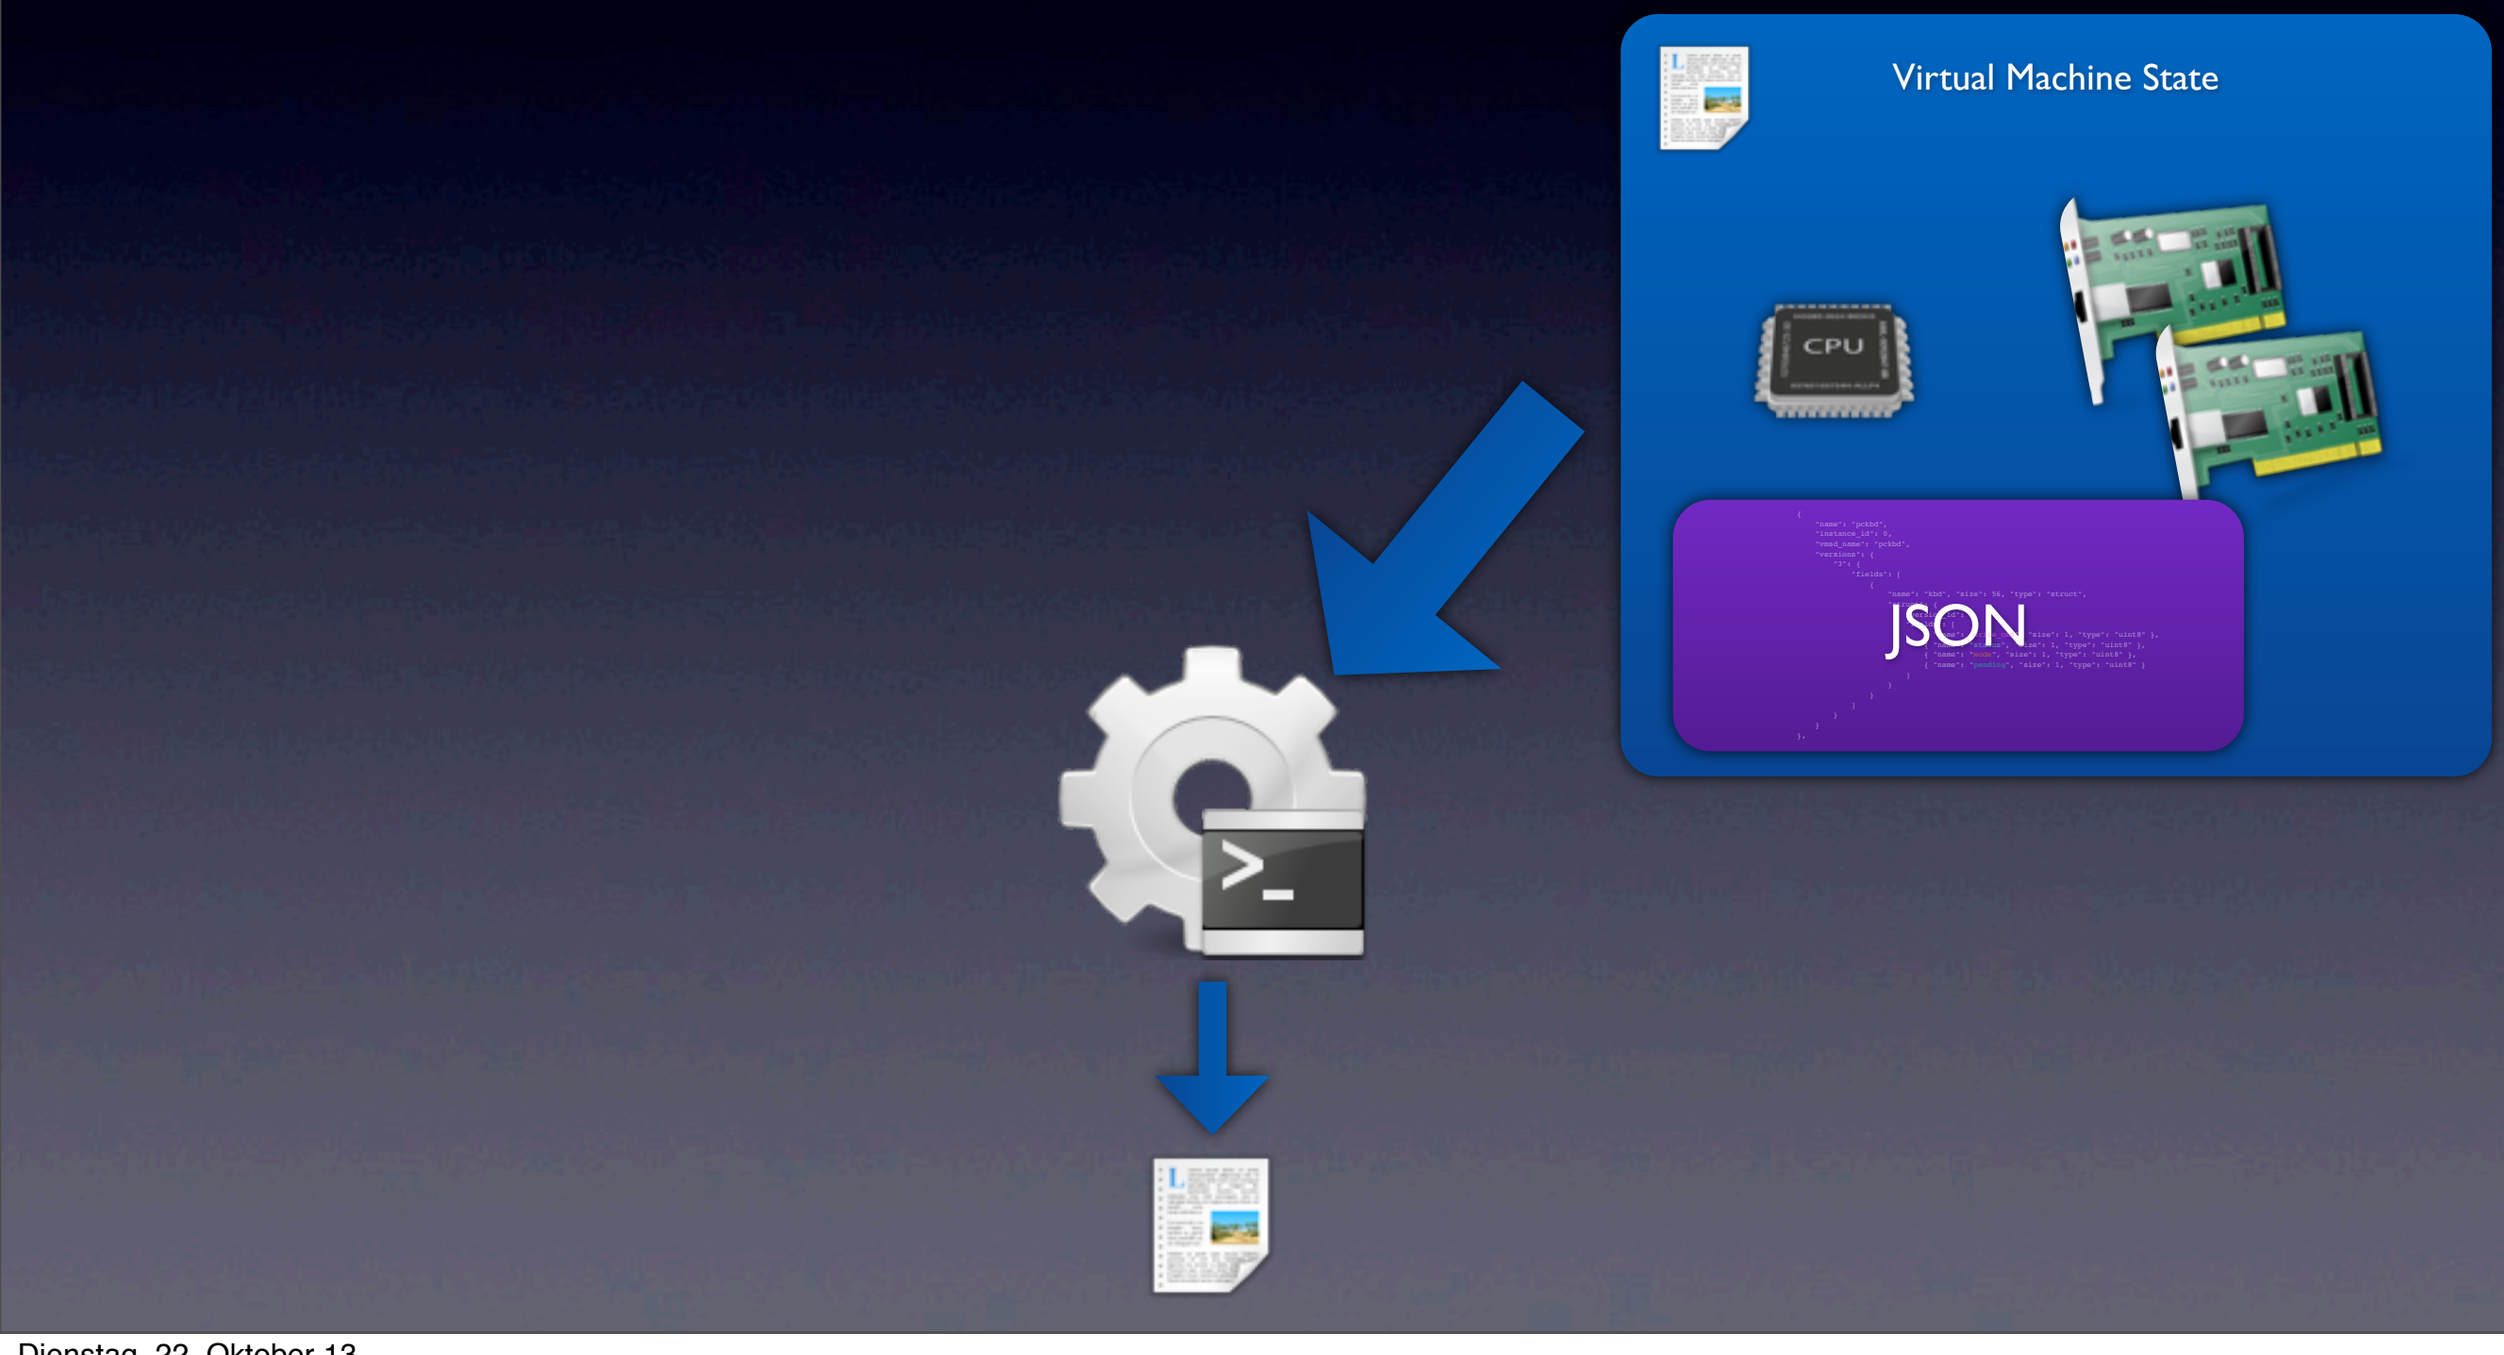

Debug Device

$ qemu-system-x86_64 -device debug-migration

Debug Device

$ qemu-system-x86_64 -device debug-migration

```
[...]
00131860  00 00 00 00 00 00 08 00  00 00 00 04 00 00 00 21  |...............!|
00131870  0f 64 65 62 75 67 2d 6d  69 67 72 61 74 69 6f 6e  |.debug-migration|
00131880  00 00 00 00 00 00 00 01  00 06 c7 da 44 65 62 75  |............Debu|
00131890  67 20 4d 69 67 72 61 74  69 6f 6e 00 7b 20 22 64  |g Migration.{ "d|
001318a0  65 76 69 63 65 73 22 20  3a 20 5b 20 7b 20 22 6e  |evices" : [ { "n|
001318b0  61 6d 65 22 20 3a 20 22  74 69 6d 65 72 22 2c 20  |ame" : "timer", |
001318c0  22 69 6e 73 74 61 6e 63  65 5f 69 64 22 20 3a 20  |"instance_id" : |
001318d0  30 2c 20 22 76 6d 73 64  5f 6e 61 6d 65 22 20 3a  |0, "vmsd_name" :|
[...]
```

Debug Device

$ qemu-system-x86_64 -device debug-migration

```
[...]
00131860  00 00 00 00 00 00 08 00  00 00 00 04 00 00 00 21  |...............!|
00131870  0f 64 65 62 75 67 2d 6d  69 67 72 61 74 69 6f 6e  |.debug-migration|
00131880  00 00 00 00 00 00 00 01  00 06 c7 da 44 65 62 75  |............Debu|
00131890  67 20 4d 69 67 72 61 74  69 6f 6e 00 7b 20 22 64  |g Migration.{ "d|
001318a0  65 76 69 63 65 73 22 20  3a 20 5b 20 7b 20 22 6e  |evices" : [ { "n|
001318b0  61 6d 65 22 20 3a 20 22  74 69 6d 65 72 22 2c 20  |ame" : "timer", |
001318c0  22 69 6e 73 74 61 6e 63  65 5f 69 64 22 20 3a 20  |"instance_id" : |
001318d0  30 2c 20 22 76 6d 73 64  5f 6e 61 6d 65 22 20 3a  |0, "vmsd_name" :|
[...]
```

Magic String

Debug Device

$ qemu-system-x86_64 -device debug-migration

```
[...]
00131860  00 00 00 00 00 00 08 00  00 00 00 04 00 00 00 21  |...............!|
00131870  0f 64 65 62 75 67 2d 6d  69 67 72 61 74 69 6f 6e  |.debug-migration|
00131880  00 00 00 00 00 00 00 01  00 06 c7 da 44 65 62 75  |............Debu|
00131890  67 20 4d 69 67 72 61 74  69 6f 6e 00 7b 20 22 64  |g Migration.{ "d|
001318a0  65 76 69 63 65 73 22 20  3a 20 5b 20 7b 20 22 6e  |evices" : [ { "n|
001318b0  61 6d 65 22 20 3a 20 22  74 69 6d 65 72 22 2c 20  |ame" : "timer", |
001318c0  22 69 6e 73 74 61 6e 63  65 5f 69 64 22 20 3a 20  |"instance_id" : |
001318d0  30 2c 20 22 76 6d 73 64  5f 6e 61 6d 65 22 20 3a  |0, "vmsd_name" :|
[...]
```

Magic String

Debug Device

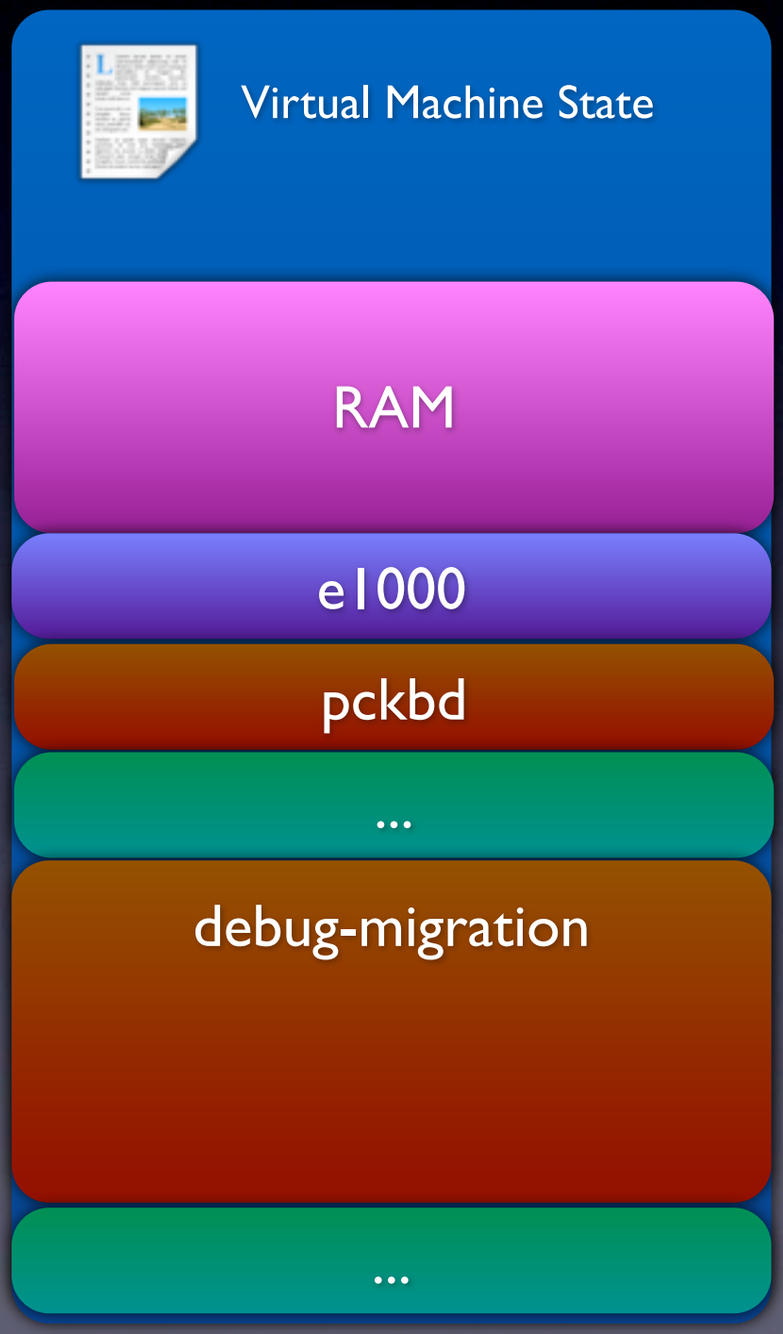

Debug Device

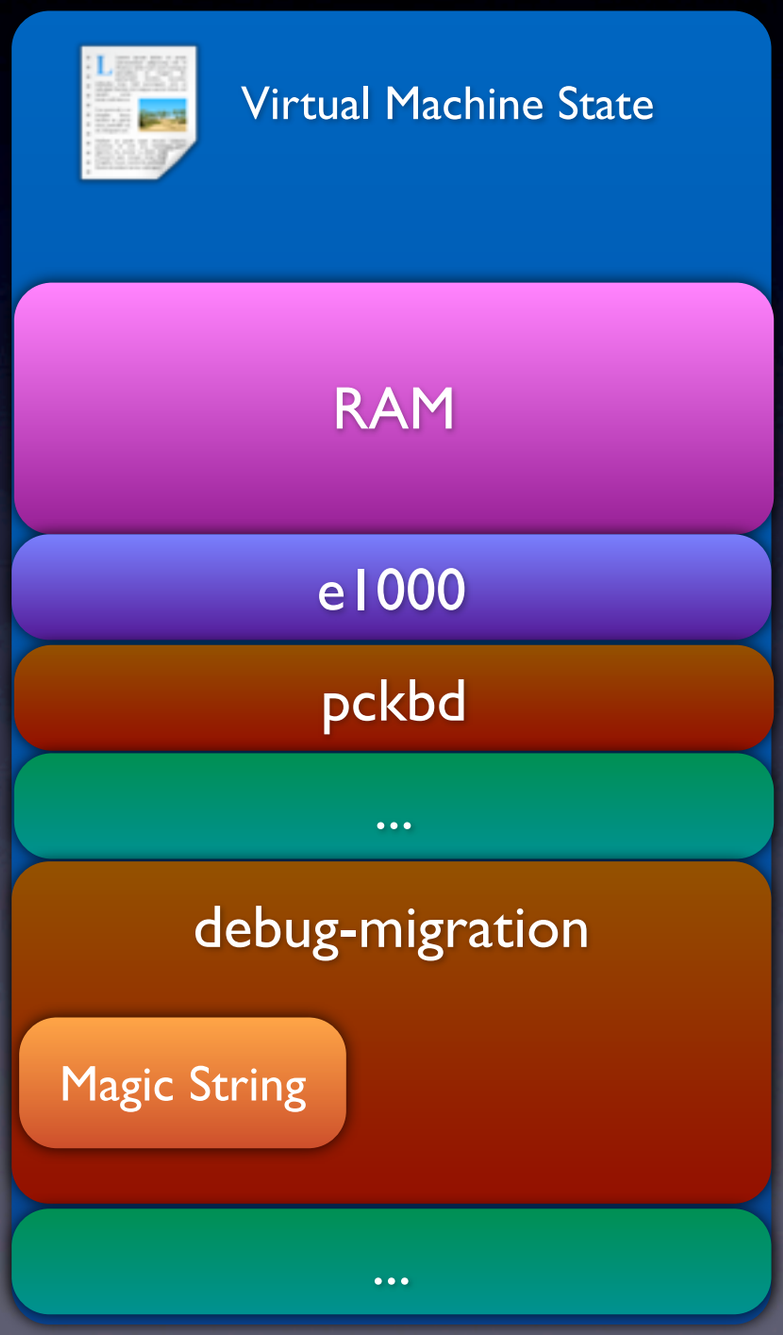

Debug Device

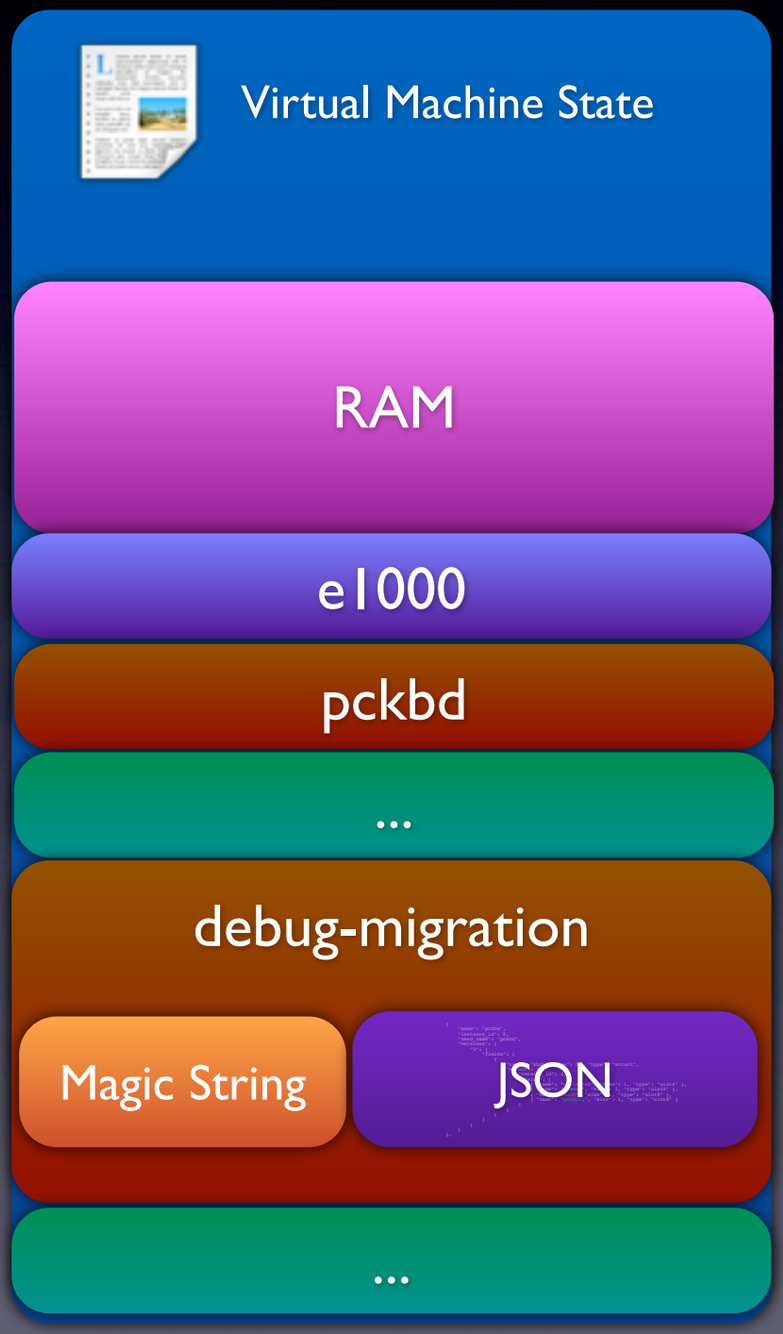

Debug Device

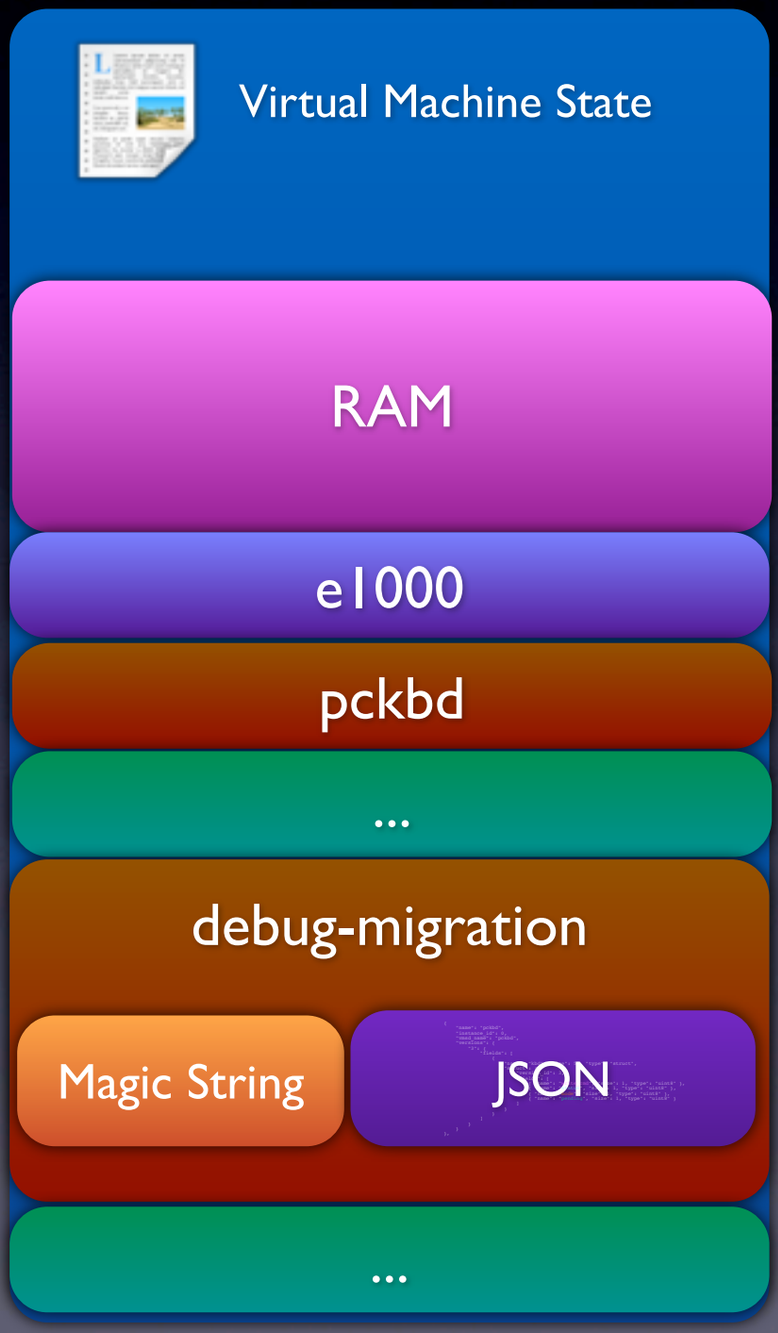

Debug Device

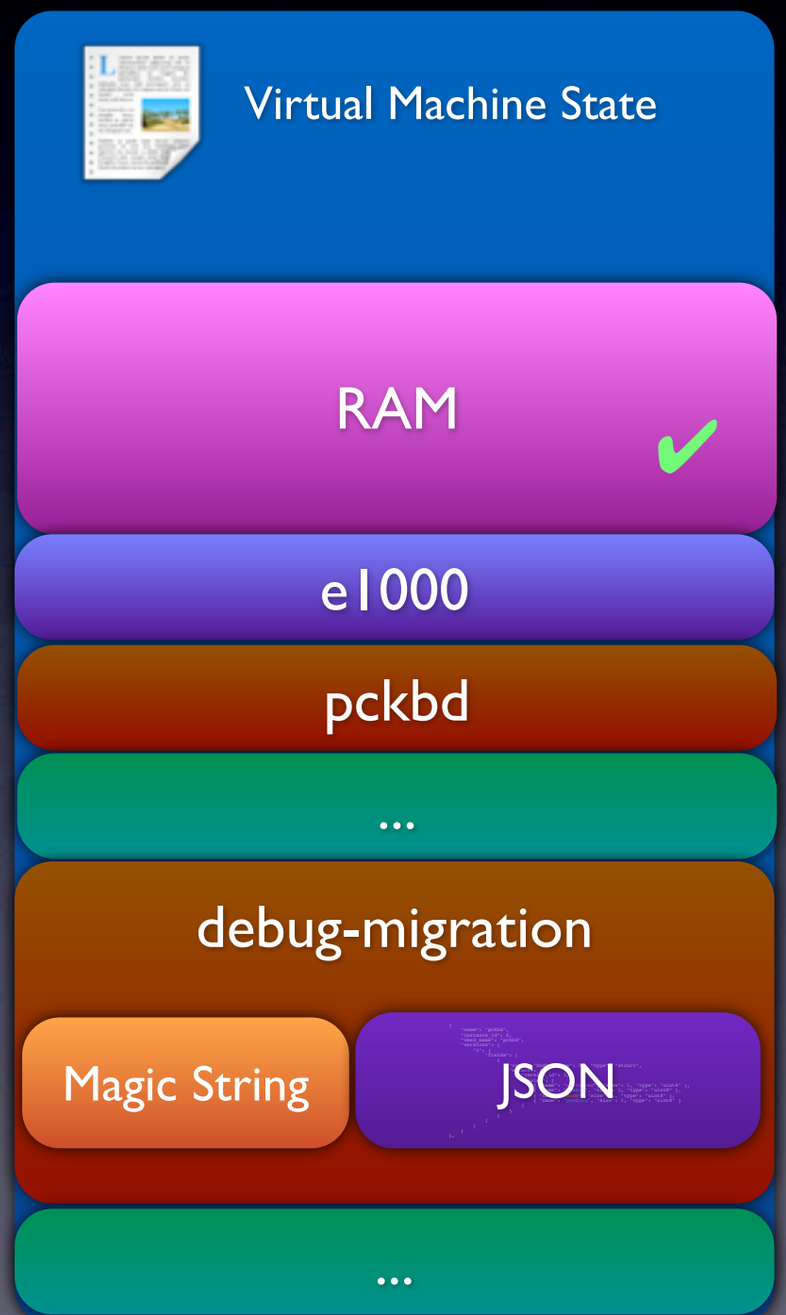

Debug Device

Debug Device

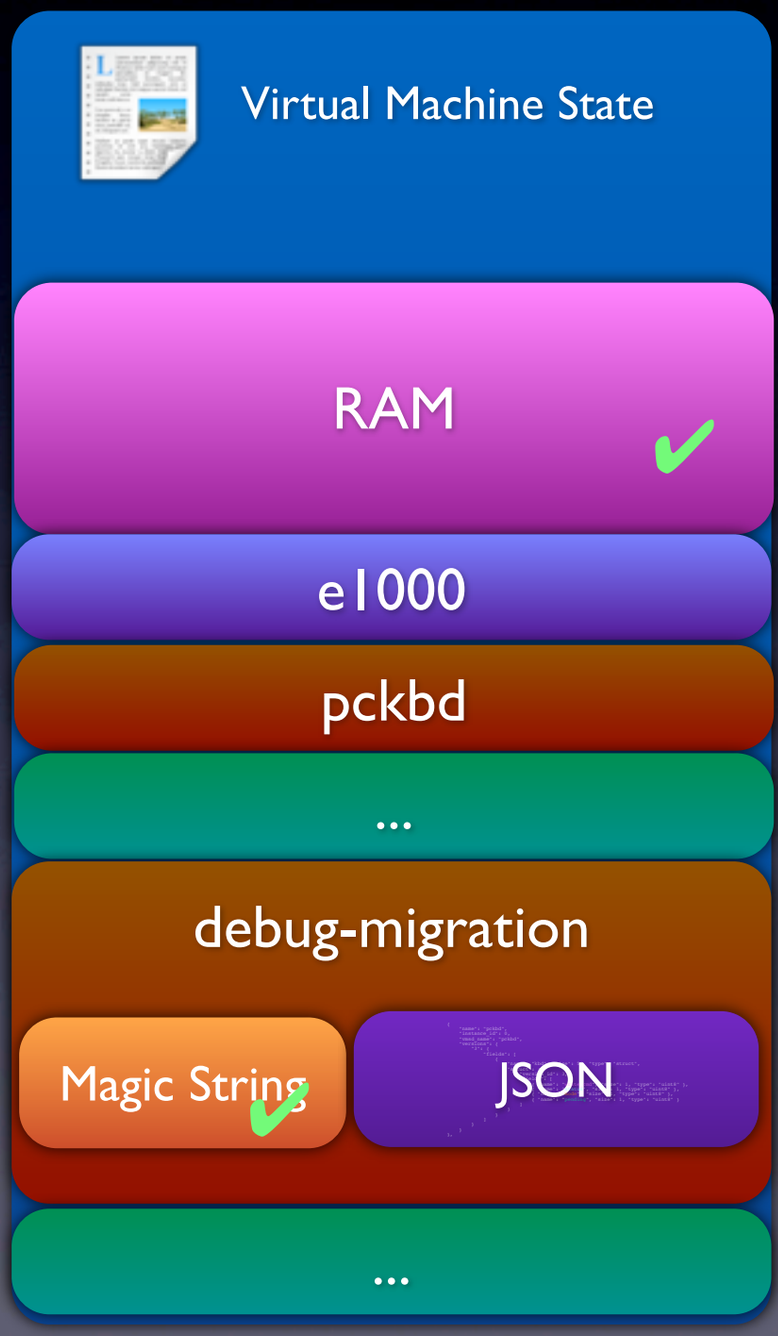

Debug Device

Debug Device

Debug Device

Use Cases

Use Cases

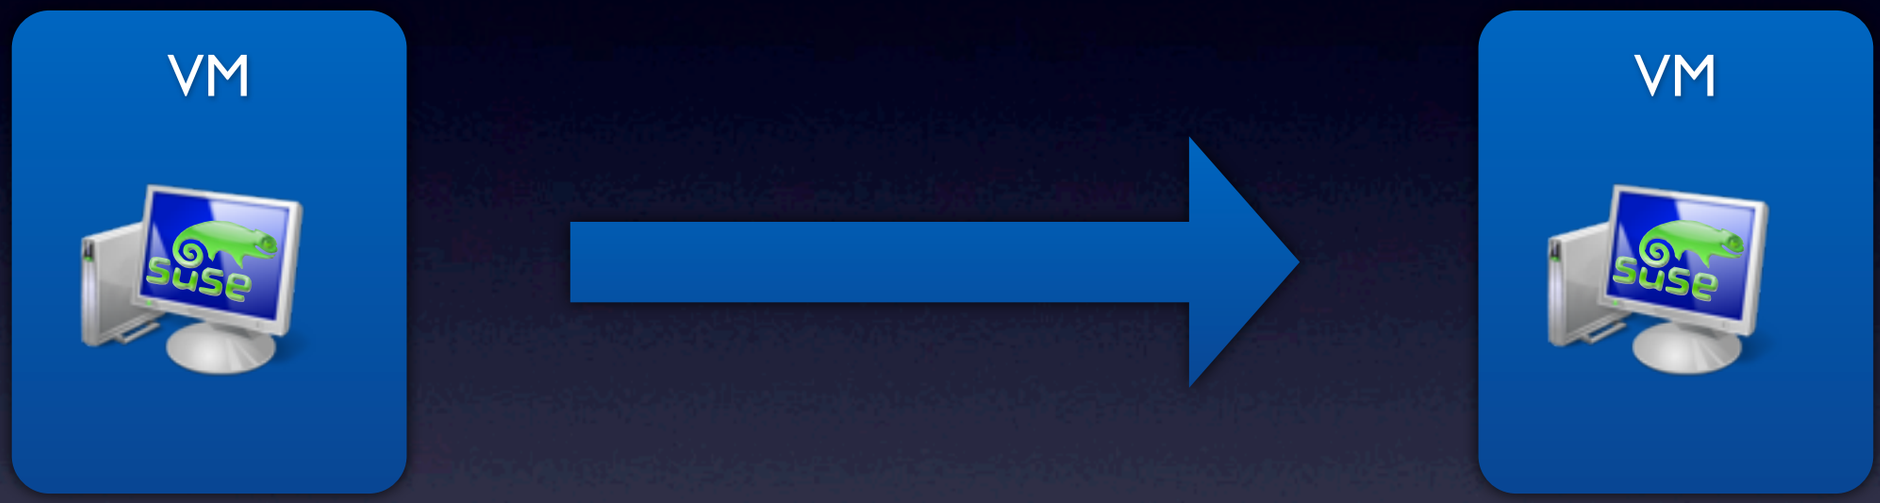

AMD new

Use Cases

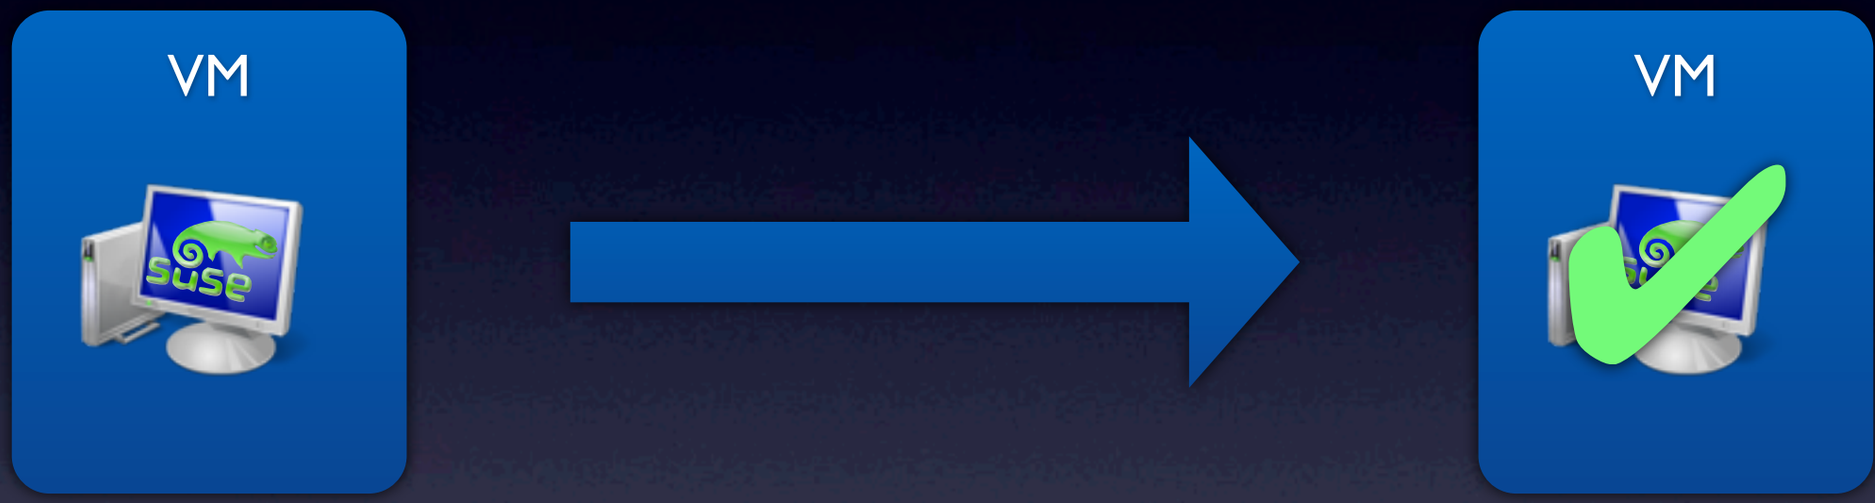

AMD new

AMD new

Use Cases

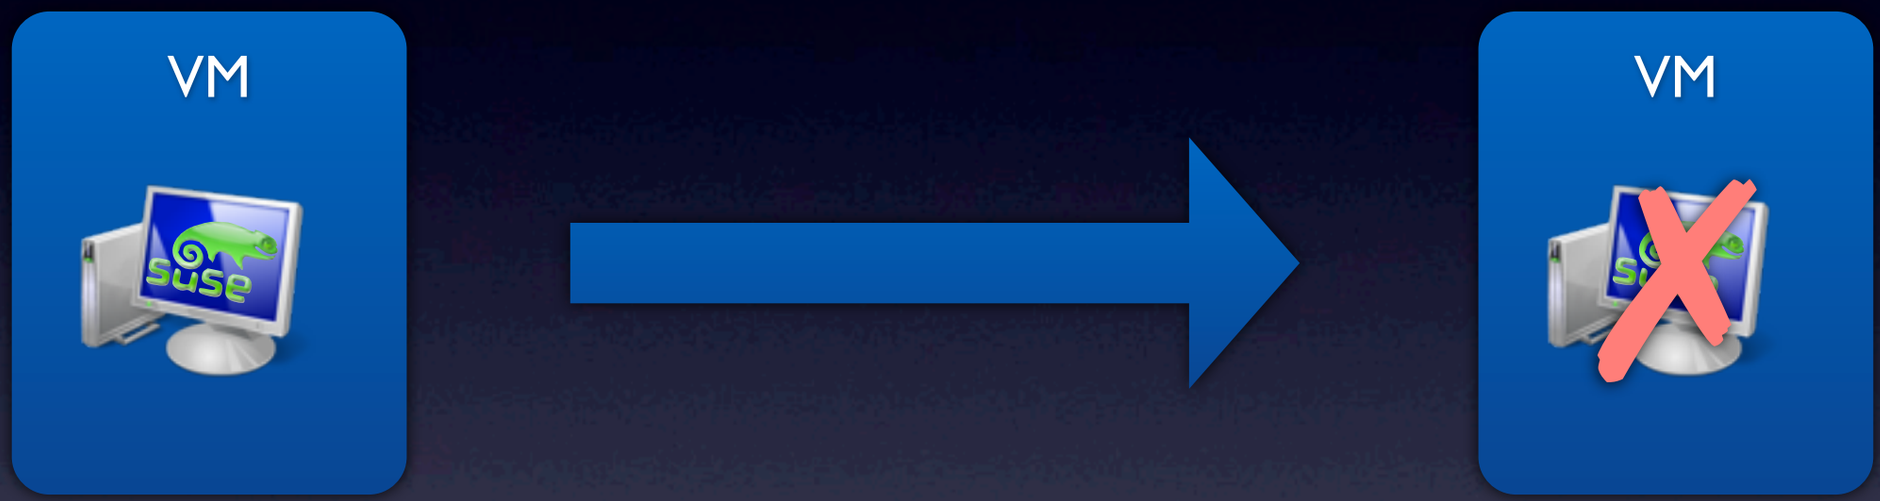

AMD new

AMD old

Use Cases

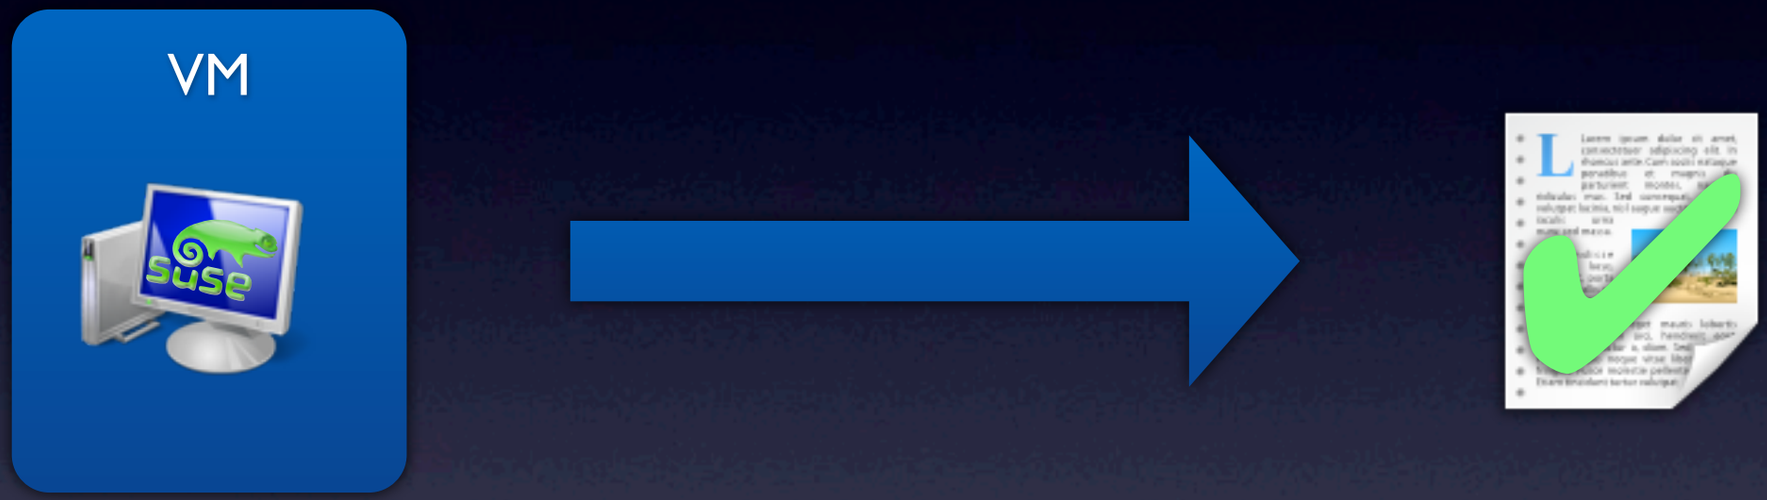

AMD new

Use Cases

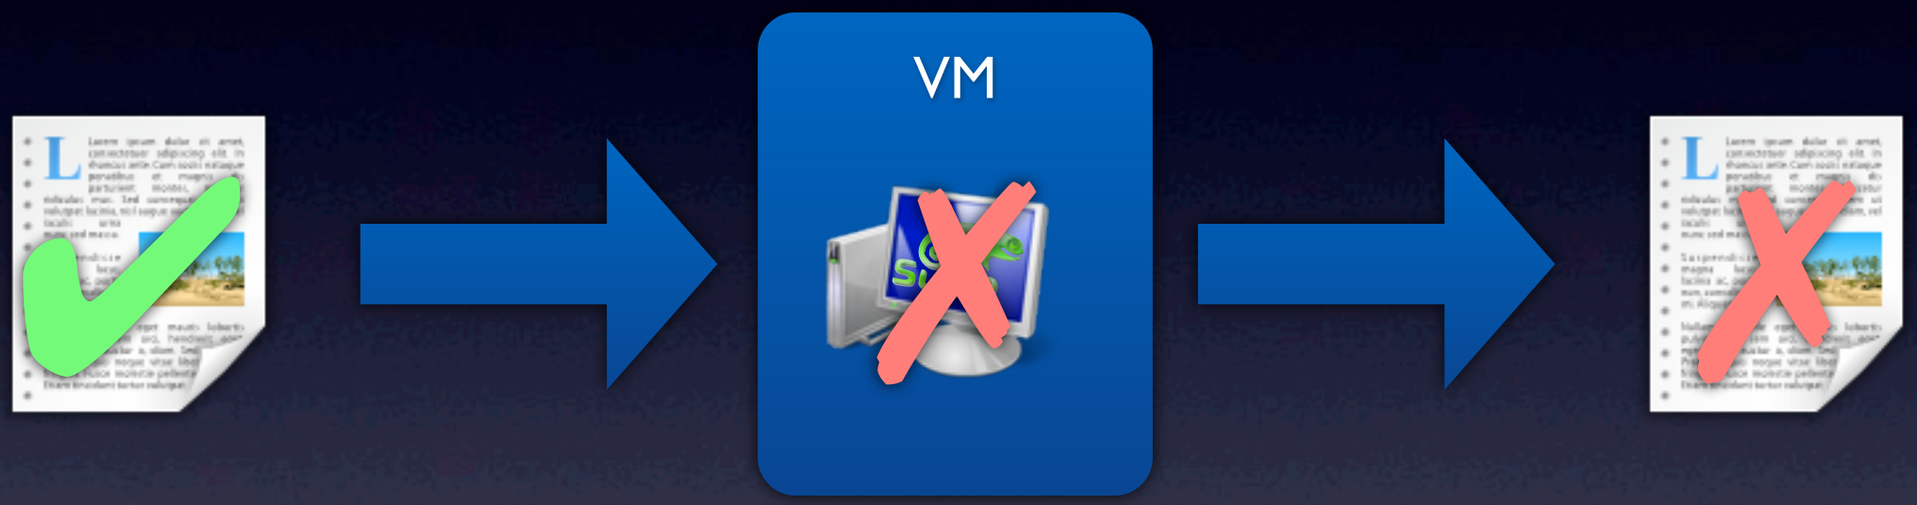

AMD old

Use Cases

Use Cases

Use Cases

diff

```
@@ -33,90 +33,90 @@
          ],
          "env.eip": "0xffffffff810300a2",
          "env.eflags": "0x0000000000000246",
-         "env.hflags": "0x0040cab4",
+         "env.hflags": "0x00000044",
          "env.fpuc": "0x037f",
          "env.fpus_vmstate": "0x0000",
          "env.fptag_vmstate": "0x0000",
          "env.fpregs_format_vmstate": "0x0000",
          "env.fpregs[0]": [
-             "00 00 00 00 00 00 00 34 ff ff",
-             "00 00 00 00 00 00 00 0c ff ff",
-             "00 00 00 00 00 00 00 00 ff ff",
-             "00 00 00 00 00 00 00 00 ff ff",
              "00 00 00 00 00 00 00 00 00 00",
              "00 00 00 00 00 00 00 00 00 00",
-             "80 00 00 00 00 00 00 00 ff ff",
-             "00 00 00 00 00 00 00 00 ff ff"
+             "00 00 00 00 00 00 00 00 00 00",
+             "00 00 00 00 00 00 00 00 00 00",
+             "00 00 00 00 00 00 00 00 00 00",
+             "00 00 00 00 00 00 00 00 00 00",
+             "00 00 00 00 00 00 00 00 00 00",
+             "00 00 00 00 00 00 00 00 00 00"
          ],
          "env.segs": [
              {
                  "selector": "0x00000000",
                  "base": "0x0000000000000000",
-                 "limit": "0xffffffff",
-                 "flags": "0x00000000"
+                 "limit": "0x0000ffff",
+                 "flags": "0x00009300"
              },
```

Use Cases

- Xsave on AVX capable system always saves AVX state
- Non-AVX capable system can not restore AVX state

```
int kvm_arch_put_registers(CPUState *cpu, int level)
{
    X86CPU *x86_cpu = X86_CPU(cpu);
    int ret;

    assert(cpu_is_stopped(cpu) || qemu_cpu_is_self(cpu));

    ret = kvm_getput_regs(x86_cpu, 1);
    if (ret < 0) {
        return ret;
    }
    ret = kvm_put_xsave(x86_cpu);
    if (ret < 0) {
        return ret;
    }
[...]
    ret = kvm_put_sregs(x86_cpu);
    if (ret < 0) {
        return ret;
    }
[..]

void kvm_cpu_synchronize_post_init(CPUState *cpu)
{
    kvm_arch_put_registers(cpu, KVM_PUT_FULL_STATE);
    cpu->kvm_vcpu_dirty = false;
}
```

```
int kvm_arch_put_registers(CPUState *cpu, int level)
{
    X86CPU *x86_cpu = X86_CPU(cpu);
    int ret;

    assert(cpu_is_stopped(cpu) || qemu_cpu_is_self(cpu));

    ret = kvm_getput_regs(x86_cpu, 1);
    if (ret < 0) {
        return ret;
    }
    ret = kvm_put_xsave(x86_cpu);
    if (ret < 0) {
        return ret;
    }
[...]
    ret = kvm_put_sregs(x86_cpu);
    if (ret < 0) {
        return ret;
    }
[..]

void kvm_cpu_synchronize_post_init(CPUState *cpu)
{
    kvm_arch_put_registers(cpu, KVM_PUT_FULL_STATE);
    cpu->kvm_vcpu_dirty = false;
}
```

Use Cases

- Bug in kvm cpu register sync
- Failures don’t get propagated, just skips syncing the rest
- End up with incomplete register sync

Use Cases

- -M compatibility debugging
- Human readable device introspection
- Send broken system snapshots around the world for analyzation

Thank You