Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Operating Systems and Networks Spring 2015
Assignment 6 Due 2 April 2015

1. **I/O**
   (a) **General Questions**
      i. State three advantages and disadvantages of placing functionality in a device controller instead of in the kernel.

> **Solution:**
> Advantages:
> - Bugs are less likely to cause the operating system to crash.
> - Performance can be improved by utilizing dedicated hardware.
> - Less functionality has to be implemented in the kernel.
>
> Disadvantages:
> - Bugs are harder to fix; a new firmware version or new hardware is needed.
> - There is less flexibility since functionality is hardcoded in the firmware.
> - The firmware could harm performance by interfering with the kernel or user processes.

      ii. The Unix-like operating systems differentiate between block and character devices. What is the difference between them?

> **Solution:** Block devices allow random access while character devices don't. In addition, block devices may be buffered.

      iii. What is the purpose of an IOMMU?

> **Solution:** Using an IOMMU has the following benefits:
> - Memory is protected from malicious devices; a device cannot read or write memory that has not been mapped.
> - Virtualized guest operating systems can use devices that do not explicitly support virtualization.
> - Devices that do not support memory addresses long enough to address the entire physical memory can still address the entire memory through the IOMMU, avoiding overhead associated with bounce buffers.

   (b) **DMA**
      i. Although DMA does not use the CPU, the maximum transfer rate is still limited. Consider reading a block from disk. Name three factors that might ultimately limit the file transfer.

> **Solution:**
> - The Speed of the I/O device, i.e., the throughput of disk reads.
> - The bus bandwidth.
> - The size of buffers within disk controllers.

      ii. A DMA controller has 4 channels. The controller is capable of requesting a 32-bit word every 100 ns. A response takes equally long. How fast does the bus have to be to avoid being a bottleneck?

Solution: Assuming each bus transaction consists of a request and a response each taking 100 ns, each transaction requires 200 ns. In other words, we can have at most 5 million transactions per second. If each one is 4 bytes, the bus should be able to handle 20 MB/s. The distribution of the transactions over the 4 channels is irrelevant.

2. **Virtual Machines**

(a) List all machine resources that must be virtualized and discuss why. List some uses of virtual machines.

Solution: To a hypervisor, the different guest operating systems are like different untrusted user applications to a normal operating system. Therefore, any resource which requires privileges to access must be virtualized by the hypervisor, e.g., CPUs, memory, I/O devices, etc.

Virtualization has been used for dynamic allocation of resources in data centers, running legacy code, and debugging operating systems.

(b) What is the difference between native (type 1) and hosted (type 2) virtual machines? Which one is more suitable for a data center (cloud provider) and which one is more suitable for your laptop?

Solution: The native type virtual machines run on bare metal. They are more suited for data centers as there are fewer layers between the hardware and guests. Note that on these types of virtual machines, one guest OS is privileged, so it can be used to manage other guests.

For instance, Xen is a hypervisor originally developed at the University of Cambridge. It is widely deployed and used for Amazon EC2, Linode, and Rackspace. In Xen, the privileged guest runs in dom0 (domain zero) while the remaining guests run in domU (unpriviledged domain).

The hosted type runs on top of an existing operating system. This is more suitable for your laptop. Your native OS is probably the one you use most regularly while the guest runs in an isolated environment. One such example is VirtualBox.

Popular virtualization solutions include Xen, QEMU/KVM, VMware, VirtualBox, and Hyper-V.

(c) Discuss the differences between full virtualization and paravirtualization, i.e., state the benefits each one has over the other.

Solution: In the former, the actual hardware is (nearly) completely simulated, so guests can run unmodified. On the other hand, hardware is not simulated in the latter, but they still run in isolated environments. While this requires modifications to guests, sometimes they can run more efficiently.

(d) Discuss the functionality of ballooning. What problem is it trying to solve and how does it solve it?

Solution: Ballooning is used reclaim memory from guest OSes. Classical operating systems use the MMU to reclaim memory from applications. When the system runs low on memory, it reclaims paged memory from applications. When applications try to access it, a page fault is triggered and the memory is returned. In the absence of such hardware support, virtual machines use ballooning.

A balloon is installed on every guest OS. To reclaim memory from the guest, the balloon expands taking up memory from the guest and returning it to the hypervisor. To return memory to the

guest, the balloon contracts.