

Distributed Systems Principles and Paradigms

Maarten van Steen

VU Amsterdam, Dept. Computer Science
Room R4.20, steen@cs.vu.nl

Chapter 03: Processes

Version: February 21, 2011



Contents

Chapter
01: Introduction
02: Architectures
03: Processes
04: Communication
05: Naming
06: Synchronization
07: Consistency & Replication
08: Fault Tolerance
09: Security
10: Distributed Object-Based Systems
11: Distributed File Systems
12: Distributed Web-Based Systems
13: Distributed Coordination-Based Systems

Introduction to Threads

Basic idea

We build **virtual processors** in software, on top of physical processors:

Processor: Provides a set of instructions along with the capability of automatically executing a series of those instructions.


Thread: A minimal software processor in whose **context** a series of instructions can be executed. Saving a thread context implies stopping the current execution and saving all the data needed to continue the execution at a later stage.

Process: A software processor in whose context one or more threads may be executed. Executing a thread, means executing a series of instructions in the context of that thread.

Architecture of VMs

Observation


Virtualization can take place at very different levels, strongly depending on the [interfaces](#) as offered by various systems components:



13/36

13/36

Process VMs versus VM Monitors



- **Process VM:** A program is compiled to intermediate (portable) code, which is then executed by a runtime system (Example: Java VM).
- **VM Monitor:** A separate software layer mimics the instruction set of hardware \Rightarrow a complete operating system and its applications can be supported (Example: VMware, VirtualBox).

14/36

14/36

VM Monitors on operating systems

Practice

We're seeing VMMs run on top of existing operating systems.

- Perform **binary translation**: while executing an application or operating system, translate instructions to that of the underlying machine.
- Distinguish **sensitive instructions**: traps to the original kernel (think of **system calls**, or **privileged instructions**).
- Sensitive instructions are replaced with calls to the VMM.

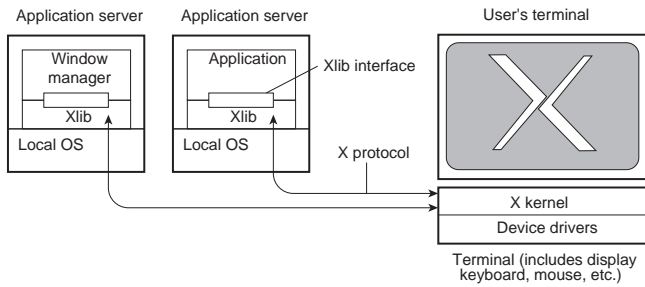
15/36

15/36

Clients: User Interfaces

Essence

A major part of client-side software is focused on (graphical) user interfaces.



16 / 36

16 / 36

Clients: User Interfaces

Compound documents

User interface is application-aware \Rightarrow interapplication communication:

- **drag-and-drop**: move objects across the screen to invoke interaction with other applications
- **in-place editing**: integrate several applications at user-interface level (word processing + drawing facilities)


17 / 36

17 / 36

Client-Side Software

Generally tailored for distribution transparency

- **access transparency**: client-side stubs for RPCs
- **location/migration transparency**: let client-side software keep track of actual location
- **replication transparency**: multiple invocations handled by client stub:



- **failure transparency**: can often be placed only at client (we're trying to mask server and communication failures).

18 / 36

18 / 36

Servers: General organization

Basic model

A server is a process that waits for incoming service requests at a specific transport address. In practice, there is a one-to-one mapping between a port and a service.

ftp-data	20	File Transfer [Default Data]
ftp	21	File Transfer [Control]
telnet	23	Telnet
	24	any private mail system
smtp	25	Simple Mail Transfer
login	49	Login Host Protocol
sunrpc	111	SUN RPC (portmapper)
courier	530	Xerox RPC

19/36

19/36

Servers: General organization

Type of servers

Superservers: Servers that listen to several ports, i.e., provide several independent services. In practice, when a service request comes in, they start a subprocess to handle the request (UNIX *inetd*)

Iterative vs. concurrent servers: Iterative servers can handle only one client at a time, in contrast to concurrent servers

20/36

20/36

Out-of-band communication

Issue

Is it possible to **interrupt** a server once it has accepted (or is in the process of accepting) a service request?

Solution 1

Use a separate port for urgent data:

- Server has a separate thread/process for urgent messages
- Urgent message comes in ⇒ **associated request is put on hold**
- Note: we require **OS supports priority-based scheduling**

Solution 2

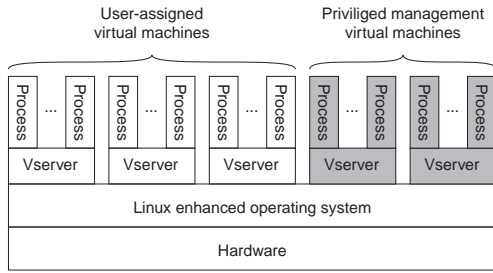
Use out-of-band communication facilities of the transport layer:

- Example: TCP allows for urgent messages in same connection
- Urgent messages can be caught using OS signaling techniques

21/36

21/36

Example: PlanetLab

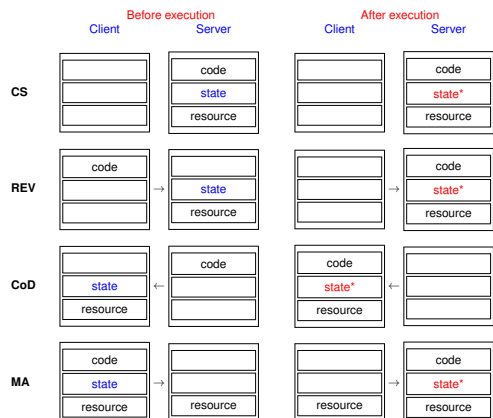


Vserver: Independent and protected environment with its own libraries, server versions, and so on. Distributed applications are assigned a collection of vservers distributed across multiple machines (slice).

Code Migration

- Approaches to code migration
- Migration and local resources
- Migration in heterogeneous systems

Code Migration: Some Context



Managing local resources

Object-to-resource binding

- **By identifier:** the object requires a specific instance of a resource (e.g. a specific database)
- **By value:** the object requires the value of a resource (e.g. the set of cache entries)
- **By type:** the object requires that only a type of resource is available (e.g. a color monitor)

34 / 36

34 / 36

Managing Local Resources (2/2)

	Unattached	Fastened	Fixed
ID	MV (or GR)	GR (or MV)	GR
Value	CP (or MV, GR)	GR (or CP)	GR
Type	RB (or MV, GR)	RB (or GR, CP)	RB (or GR)

GR = Establish global systemwide reference
MV = Move the resource
CP = Copy the value of the resource
RB = Re-bind to a locally available resource

35 / 36

35 / 36

Migration in heterogenous systems

Main problem

- The target machine may not be **suitable to execute the migrated code**
- The definition of process/thread/processor context is **highly dependent on local hardware, operating system and runtime system**

Only solution

Make use of an **abstract machine** that is implemented on different platforms:

- Interpreted languages, effectively having their own VM
- Virtual VM (as discussed previously)

36 / 36

36 / 36