Poseidon House
Castle Park
Cambridge CB3 0RD
United Kingdom

TELEPHONE: **Cambridge (01223) 515010**
INTERNATIONAL: **+44 1223 515010**
FAX: **+44 1223 359779**
E-MAIL: **apm@ansa.co.uk**

Training

ANSAwise - Distributed Operating Systems [Eurocontrol]

Mark Madsen

Abstract

This module of the ANSAwise training course programme describes and demarcates the role of networking in a distributed system. The approach is based on the OSI Reference Model, and identifies its uses and limitations.

The thrust is that distributed systems architects build on the network architecture, and so must have a clear grasp of key networking issues.

APM.1613.00.01 **Draft** 12th October 1995

Briefing Note

Distribution:
Supersedes:
Superseded by:

Distributed and Networked Operating Systems

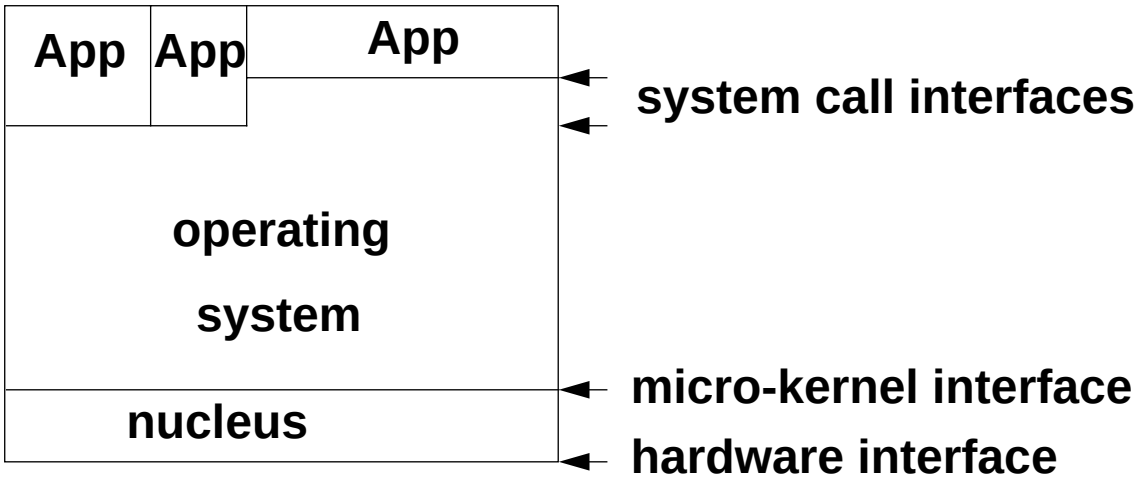

In this session

- *Review the basic ideas of networking*
- *Explain the trends in the new generation of operating systems*
- *Show how distributed systems can exploit these trends*
- *Explore relevant standards*

The OSI Reference Model (est. 1978)

- ***Standardized as ISO 7498, ITU-T X.200***

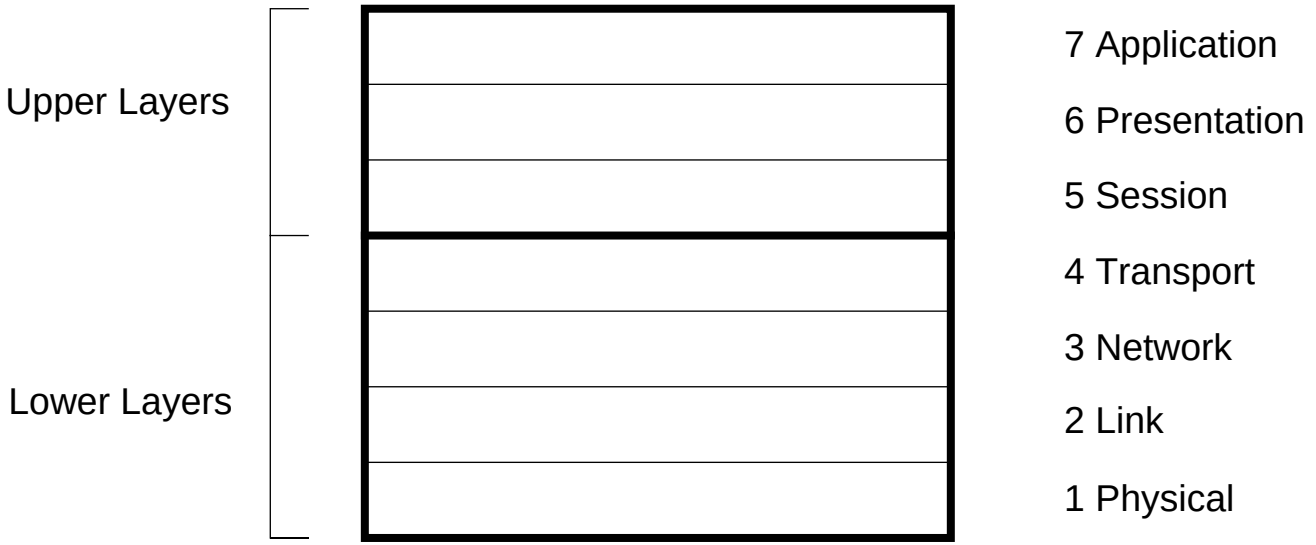

Principles of the OSI Reference Model

- ***Simplicity***
- ***Boundaries***
- ***Layering***
- ***Independence***

OSI Layer Entities

- ***Entities interface via Service Access points and protocols***

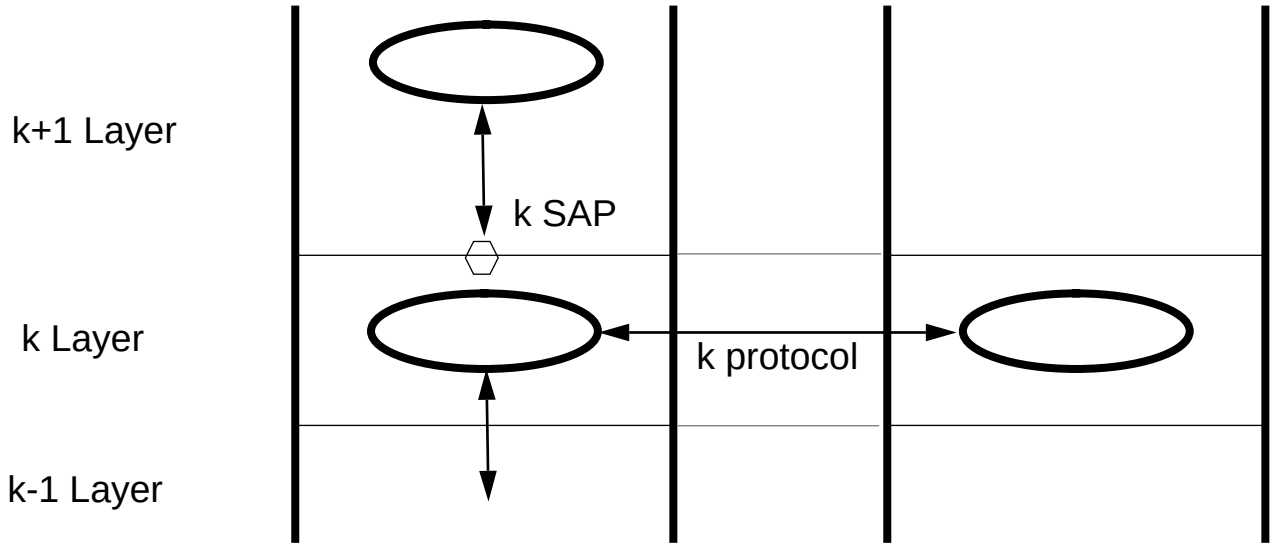

Purposes of the layers

- *Application: common application service elements*
- *Presentation: data syntax*
- *Session: connection management, concurrency management*
- *Transport: fragmentation and reassembly*
- *Network: routing*
- *Data link: reliable point-to-point communications, error correction*
- *Physical: interface to the transmission medium*

Influence of the OSI Reference Model

- *The influence of the OSI Reference Model has been profound*
  - the *principles* have withstood the test of time
  - it has been a framework for shared understanding...
  - ...all network architectures refer to this model
- *You can apply the OSI Reference Model without OSI protocols*
  - the OSI Reference Model does not imply OSI protocols
  - other protocol stacks can match the OSI Reference Model

SMB Protocol Suite

- ***The protocols are mainly used on LANs***

NB:
Upper Layer
position

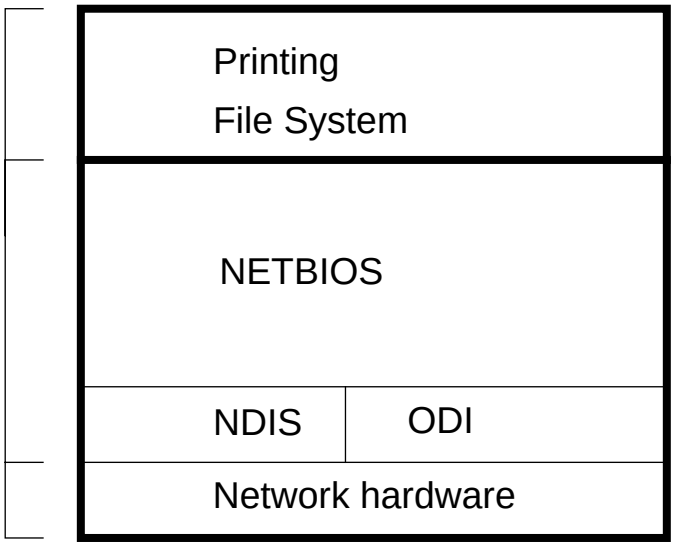

7 Application
6 Presentation
5 Session
4 Transport
3 Network
2 Link
1 Physical

Internet Protocol Suite

- ***The protocols operate over LAN and WAN***

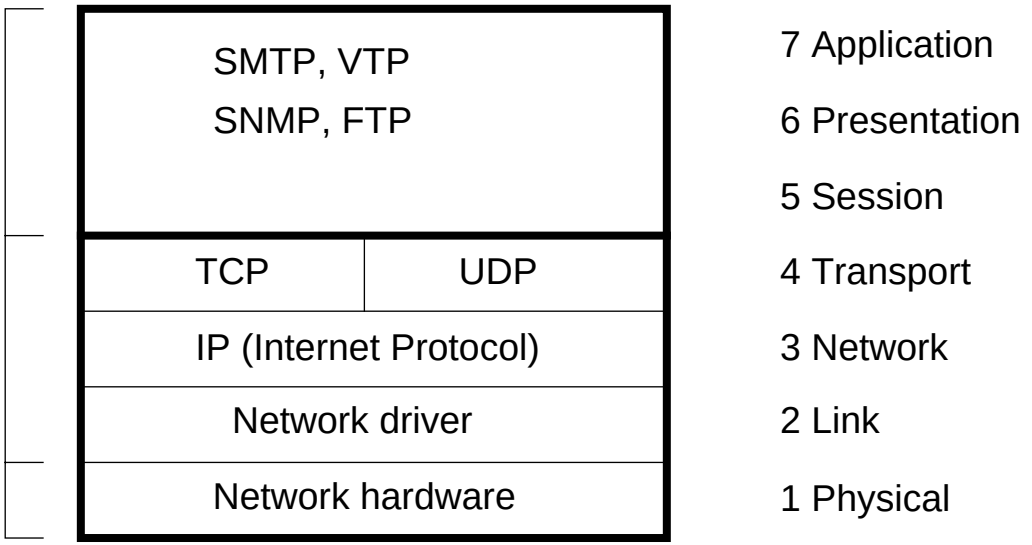

The success of the Internet Protocol Suite

- *It has been driven by*
  - Timely solutions
  - Low-cost (or even free) implementations for many systems
  - Enthusiastic (and tolerant!) technical community
- *It has been hindered by*
  - Complex installation and configuration
  - Expensive implementations for PC LANs
- *These hindrances are now disappearing*
- *The rapid growth of the Internet will continue this success*

Is the OSI Reference Model enough?

- ***Distributed systems need more...***
  - **interaction types more complex than peer-to-peer**
  - **object-based applications rather than simple layering**

OSI and Distributed Systems views of the world

- ***OSI***

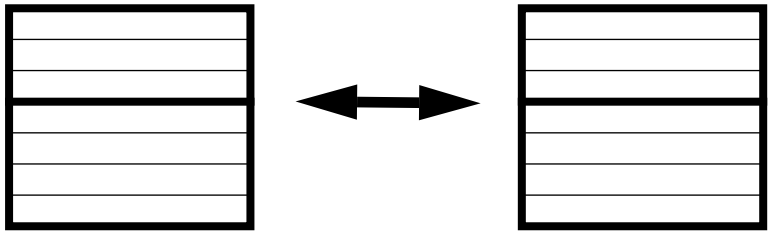

- ***Distributed systems***

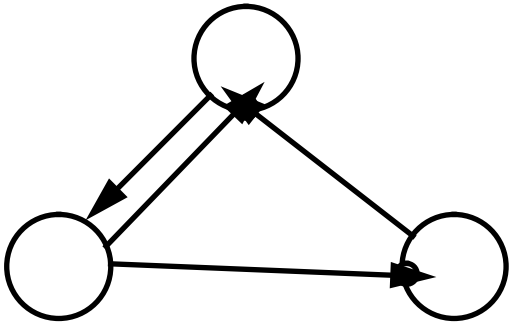

The end-to-end principle

- ***With hindsight, the OSI Reference Model lacked one important principle***
  - **the ‘end-to-end’ principle**
- ***This can be stated as***
  - **“Only add reliability mechanism to a layer if it improves performance”**
- ***Higher layers have to perform error checking anyway***
  - **to provide end-to-end guarantees**

Protocol diversity

- ***We want to exploit maximum diversity in the applications, and in the physical technologies***

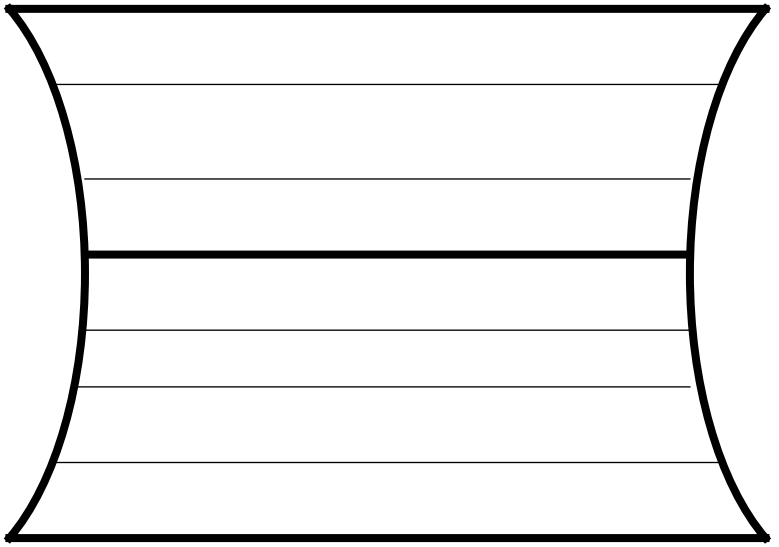

Upper and Lower Layers

- ***As a gross generalization (treat with caution!)***
  - **Distributed systems is concerned with the upper layers**

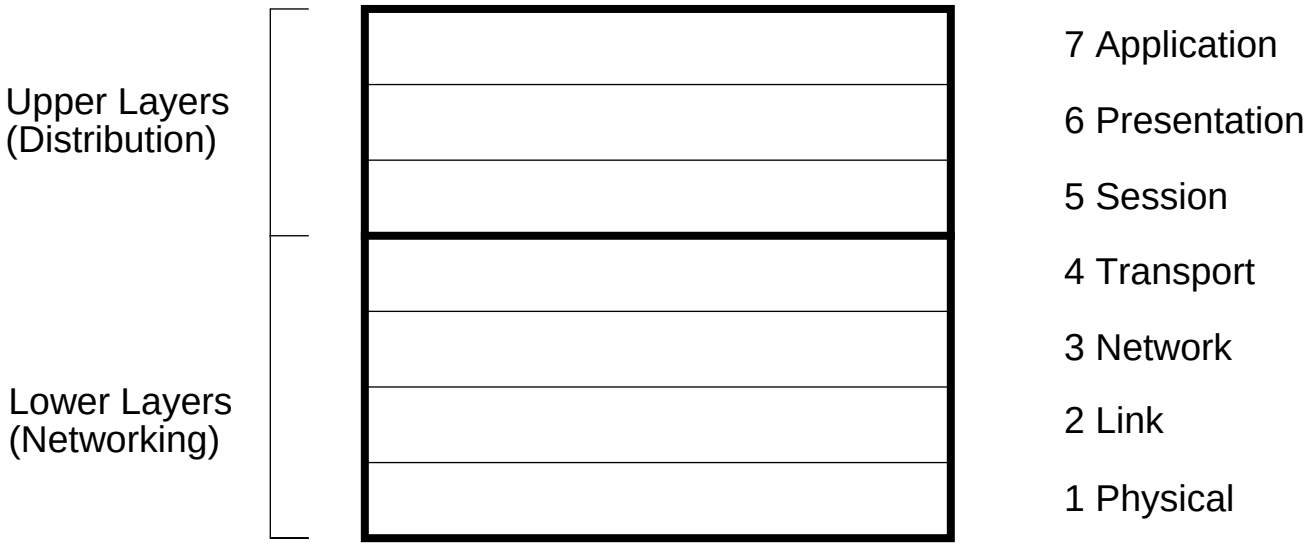

  - **... the lower layers are the concern of networking**

Distributed Systems and Networking

- *It is important to understand the lower layers, because*
  - Distributed systems will make specific demands on the lower layers...
  - ... for example, quality of service/performance characteristics
  - ... you must understand how these demands will be met
- *Many of the lower layers concepts are common with the upper layers*
- *Practical understanding also helps*
  - to assemble a complete system
  - to test and debug a complete system
  - when working with networking specialists

What else is missing?

- *No control over quality-of-service (QoS)*
- *No synchronization capability*
- *No multicast capability*
- *Ease of management*

Group communications mechanisms

- ***Unicast***

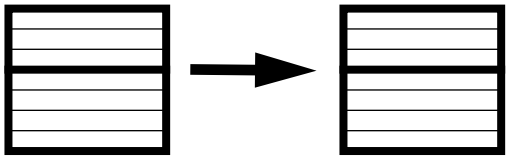

- ***Broadcast***

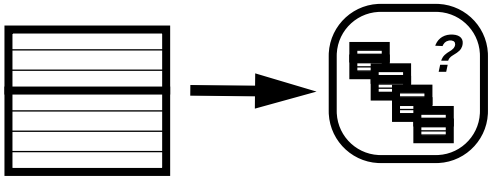

- ***Multicast***

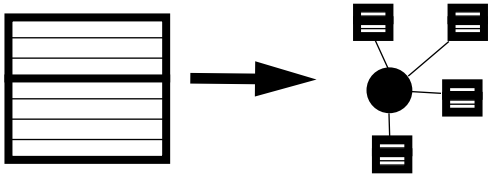

Group communications mechanisms

- ***Unicast (point-to-point)***
  - **simple, but only suitable for 1-1 communications**
- ***Broadcast***
  - **Broadcast protocols don't scale to large networks**
  - **The scope of a broadcast protocol is usually physical (e.g. an Ethernet LAN segment; this is unrelated to logical distributed systems needs**
- ***Multicast (point-to-multipoint)***
  - **promising, but requires a sophisticated mechanism to manage the membership of the multicast group**

Connectionless and connection-orientated

- *Connectionless protocols are suitable for call-response interactions*
- *Connection-orientated protocols are needed for continuous media*
  - for example, video, audio
- *A connectionless k-protocol can be layered over a connection-orientated k-1 protocol (and vice-versa)*
- *Newer networks have faster call setup*
  - encouraging connection-orientated use

PC LAN protocols in practice

- ***PC networks in Times Top 1000 Companies***

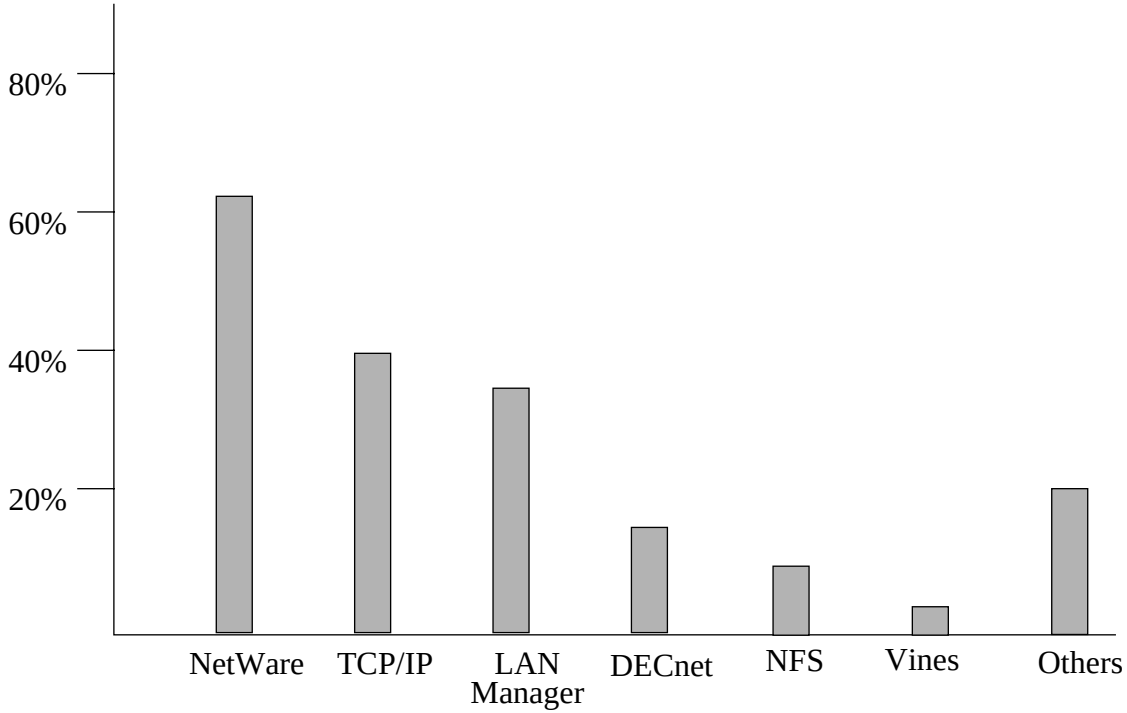

LAN protocol diversity - an example

- *Hughes Aircraft Corp. have 51,000 staff*
- *Hughes Enterprise Network Architecture has to cope with*
  - 50,000 ports
  - 860 hubs
  - 9 protocols (NetWare IPX, TCP/IP, Vines IP, Appletalk, DECnet, LAT, NETBIOS, SNA, OSI TP)
- *They aim to get down to 6 protocols*
- *They'd like to get down to 2 protocols*
- *They think it will take until the year 2000*

LAN protocol summary

- ***This is a legacy problem...***
  - **it makes technical support difficult (particularly routing and bridging)**
  - **it slows down the use of new technologies**
- ***...But it is one that the network architects must solve***

Networking challenges - protocols

- *Provide new and updated protocols for*
  - scaleable performance
  - dependability (fault tolerance)
  - predictability (real-time communications
  - mobility
  - ubiquity
  - security (highly controversial!)

Networking challenges - new technology

- ***Develop new hardware technologies***
  - **wireless (radio and infra-red)**
  - **very low cost networks**
  - **very high speed networks (broadband)**

Corresponding challenges for distributed systems

- *Distributed systems depend on these for new applications*
  - multimedia
  - telepresence
  - UPT (universal personal telecommunications)
- *The challenges will affect*
  - the engineering of distributed systems
  - the design of distributed applications
  - the management of both distributed systems and distributed applications

LAN/Network Operating Systems

- ***PC centralized networks***
  - **Microsoft/IBM (NT Advanced Server, LAN Server)**
  - **Novell NetWare**
  - **Banyan Vines**
- ***Unix decentralized networks***
- ***Mainframe networking***

Peer-to-peer network operating systems

- *In a peer-to-peer network, any machine can run a server process*
- *For example:*
  - Mac Appletalk
  - Unix NFS
  - Microsoft Windows for Workgroups
- *Again, the network protocol varies (Appletalk, NFS, NETBIOS)*

Distributed Operating Systems

- ***The research community has been studying distributed operating systems since the late 1970's***
  - **The ANSA project started looking at this in 1983**
- ***Attempts have been made to produce fully-transparent distributed operating systems***
  - **location transparency for files, processes, memory, resources**

One distributed operating system?

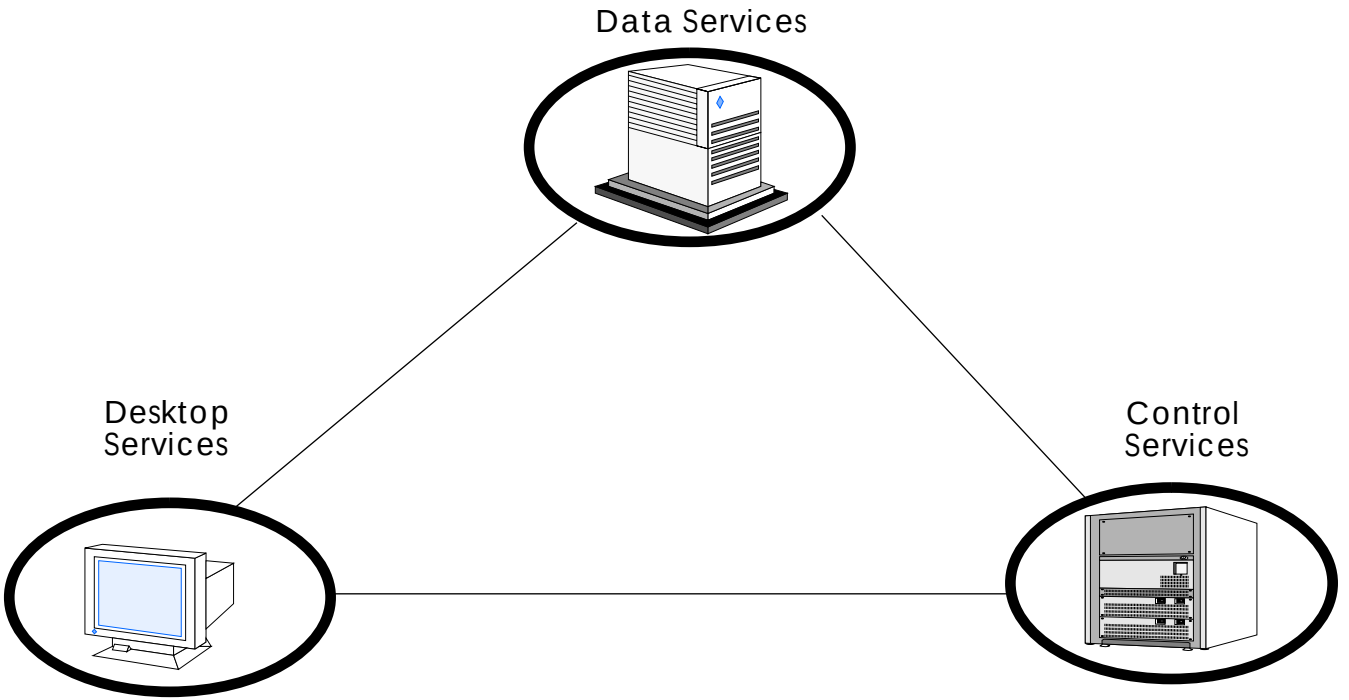

The problem with true distributed operating systems

- *Complete location transparency can be achieved...*
- *... but even this is impractical*
  - it places far too many overheads on the system (at several levels)
  - it won't scale beyond a small group of machines
  - it won't scale over a wide-area network
  - it makes it difficult to support a diverse collection of hardware
  - it doesn't allow for interworking with other operating systems
- *Research continues, but full transparency is no longer a goal*

Three kinds of systems

- ***Centralized***
  - **single location**
- ***Decentralized (peer-to-peer)***
  - **multiple location, un-coordinated components**
- ***Distributed***
  - **multiple location, coordinated components**

New generation operating systems

- *They have common aims...*
  - Dependability, to support mission-critical applications
  - Portability, to support new and old hardware platforms
  - Modularity, a component-based architecture
  - Scalability, from low-cost to high-performance
  - Ease of use, having a familiar graphical user interface
- *Applications must inherit these benefits*
  - but they do not come automatically

The structure of a micro-kernel operating system

- *From a complete micro-kernel Unix system...*

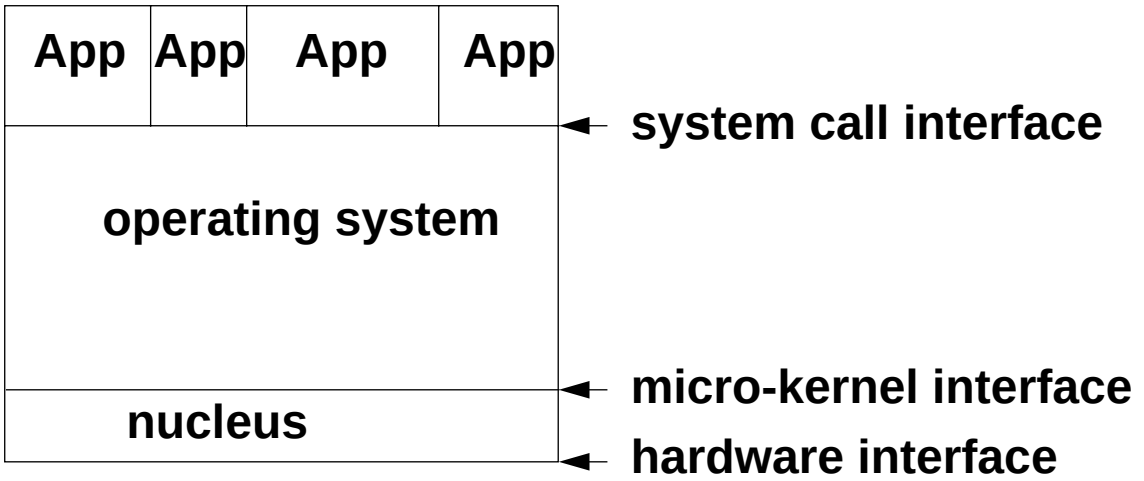

Micro-kernel for modular operating systems

- *.... unnecessary services can be removed...*

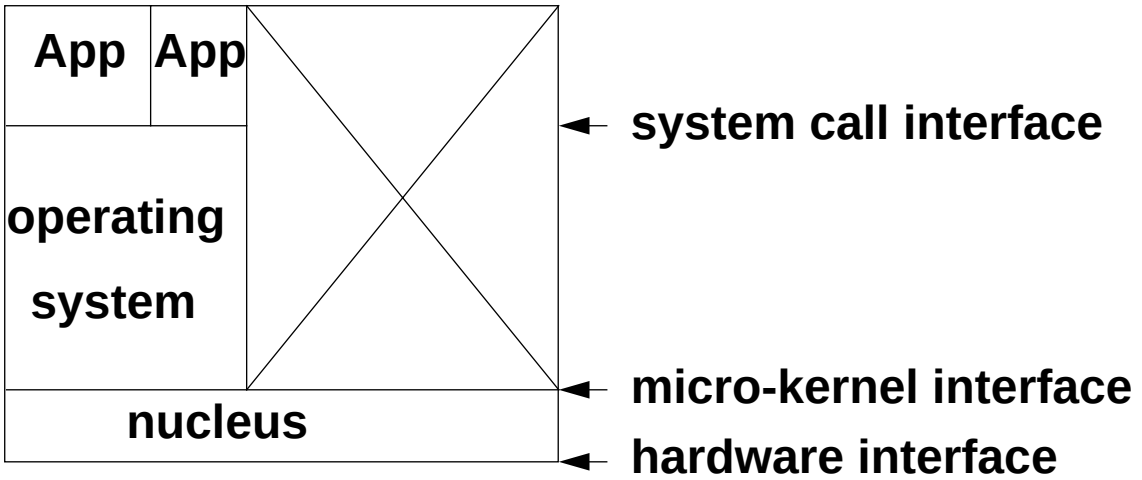

Micro-kernel for modular operating systems

- *.... or different operating systems emulated simultaneously*

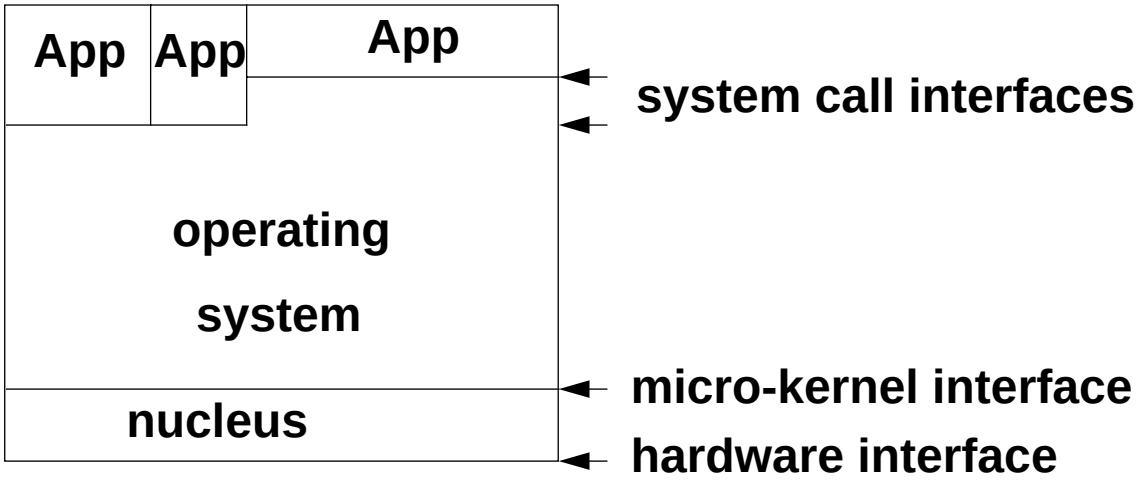

- *for example, different variants of Unix*

The importance of micro-kernel

- *A micro-kernel provides the minimum necessary facilities*
  - interprocess communication
  - low-level scheduling and process management
  - basic memory management
  - low-level I/O
- *Once again, the kernel can be small (less than 100 Kbytes)*
  - generic low-cost embedded systems can built
  - the micro-kernel is isolated; it can be developed for reliability

Operating System Dependability

- ***Micro-kernels help here***
  - **they help contain the effects of faults**
- ***Operating system design and engineering has improved***
- ***Product maturity helps too***
  - **...every application exercises the operating system**
  - **...bugs get fixed eventually (“time cures all wounds”)**
- ***Fault tolerance is still an issue***
- ***Application dependability is still an issue***

Operating System Portability

- *Operating system suppliers gain a wider market from portability...*
  - ... but what do users gain?
- *Basically, investment protection*
  - so you can move your application if you must
- *Unix is portable, but...*
  -
  -
  -

Summary

- ***The new generation of operating systems and hardware platforms provides useful technology***
  - **we must be able to exploit them...**
  - **... but they don't provide any assistance in building distributed systems**
- ***The applications developer's job is no simpler***
  - **Operating systems developers have solved *their* problems, not yours**
- ***To build distributed systems, we need something extra***