



CS 423

Operating System Design

<https://cs423-uiuc.github.io>

Tianyin Xu

tyxu@illinois.edu

* Thanks Adam Bates for the slides.



Instructor

- Tianyin Xu

Teaching Assistants

- Cathy Cai (PhD student)
- Peizhe Liu (MS student)

Office Hour

- There is an OH every day.
- Check the website.



I'm working on software and system reliability.

I worked at Facebook on dealing with datacenter level failures before joining UIUC.

- FAQ: academia or industry?
- I gained 20 LB eating free food.

I graduated from UC San Diego in 2017.

- I worked on hardening cloud and datacenter systems against misconfigurations.
- Dream job: a tenured grad student

I applied twice for grad school.

- I failed the first time (as always).

Cathy Cai




- 2nd year PhD student working on system verification (currently)
- Use Arch Linux BTW (since 2019)
- TA for OS & Information Retrieval in my undergrad
 - TBH MPs for that OS should be put in system programming
- Hobbies
 - Japanese anime (just watched for ~40h this summer)
 - Badminton, running & cycling
 - Breaking my homelab
- Got a car for 2,000\$ last year



Peizhe Liu



- MS student working on OS and storage
- Graduated from UIUC ECE
- Hobbies
 - **Retrocomputing**
(founded a club in UIUC)
 - **Blood donation**
(featured on local news)
 - **Teaching**
(CS125, ECE385, ECE391, ECE411 as CA, now CS423 as TA)
 - **Homelab**
(since 2019)



Learning Objectives



Before CS 423:

- Knowledge of C/C++
- Basic knowledge of Linux/POSIX APIs and functions



After CS 423:

- Mastery of Operating Systems concepts
- Comprehensive understanding of virtualization techniques
- Introduction to advanced OS topics: security, energy, redundant storage...
- Become a kernel hacker capable of establishing a kernel development environment and modifying operating system code

Today:

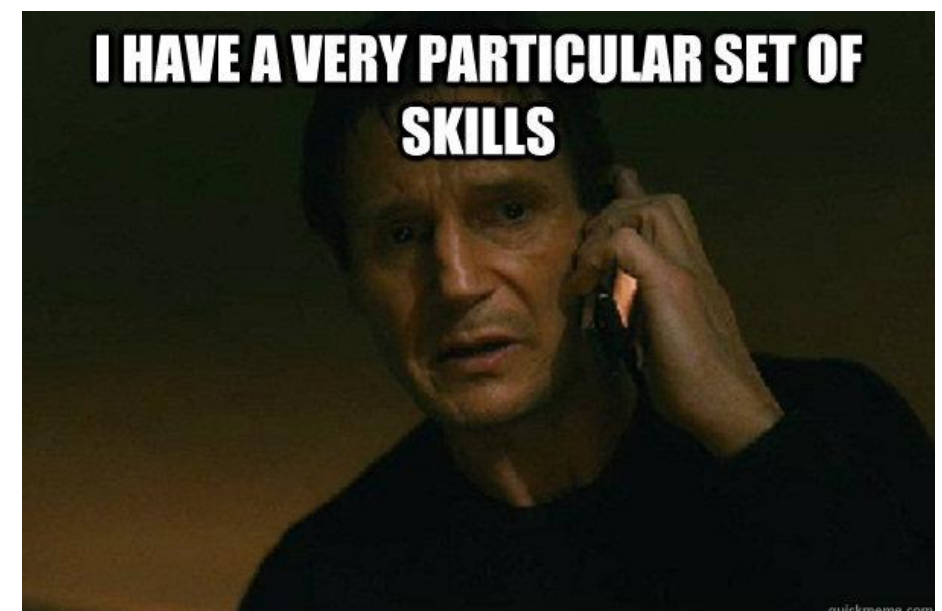
- Introduce the instruction team
- Go over the requirements and expectations for this course

What's in it for you?



- Understand the foundations of all system software
- Apply systems concepts and methodologies to higher layer software systems. Modern browsers, language virtual machines, and IoT devices all run their own forms of operating systems!
- Acquire a very particular (and lucrative) set of skills!

"I attended a Microsoft-organized meeting where the Director of Engineering of (Microsoft in Redmond) talked to me about a great need for engineers who know operating systems/device drivers, and know linux kernel/programming at such lower levels. He bitterly complained that many CS departments are dismanteling their OS programs. I told him that we have actually multiple OS undergraduate classes at UIUC the current instructor to advertise among the students who take these courses that there are many jobs at Microsoft in OS area (more than ever!)."



My Own Version



- Prepare you for the real world!
 - Real-world software is imperfect and buggy
 - Unfortunately, you have to built on top of them;
 - Real-world infrastructure is fragile and vulnerable.
 - Unfortunately, you have to bear with that.
 - Document is obsolete and even misleading
 - You won't have teachers or TAs
- No matter what, have **FUN**!

Prerequisites



- Did you take CS241?
- Did you take ECE391?
- Do you have systems programming experiences from another university?
- If not, you might have a bad time in this course...



piazza

**You are already added on the Piazza.
(if not, find the link on the course website)**

Go here for announcements and to ask questions.

Instruction team will be checking forums regularly!



- “Operating Systems: Three Easy Pieces” OSTEP
Remzi and Andrea Arpaci-Dusseau
 - It is **FREE**.
 - [Why Textbooks Should Be Free](#)
 - The chapters are linked on the website.

Additional Texts



- Alternative Textbooks (Not Free):
 - Operating Systems: Principles & Practice
[Anderson and Dahlin](#), 2018
 - Modern Operating Systems
[Tanenbaum and Bos](#), 2014
 - Operating System Concepts
[Silberschatz, Galvin and Gagne](#), 2012
- Other Recommended Reading:
 - Virtual Machines
[Smith and Nair](#), 2005
 - Linux Kernel Development**
[Love](#), 2010

** Helpful for MPs

CS 423 Requirements



- Attendance/Participation
 - Come to class, Tue/Thu, 2:00-3:15am
 - Participate actively in class and on piazza
- Machine Problems (MPs): 4 major programming assignments
- Periodic Homework: includes “prereqs”, may assign more
- Midterm & Final Exams: Dates TBD
- 4 Credit Class: Read additional assigned literature and submit summaries weekly.

ALL WORK IS TO BE INDEPENDENTLY COMPLETED!

Participation



- Contribute in class — ask questions, respond to questions, share relevant outside knowledge.
- Contribute **good** questions and answers on Piazza!
- *“The kind of answers you get to your technical questions depends as much on the way you ask the questions as on the difficulty of developing the answer.”*
- **How To Ask Questions The Smart Way:**
<http://www.catb.org/esr/faqs/smart-questions.html>
- Other questions (e.g., administrative) on Piazza are also welcome, but won't give you participation credit.

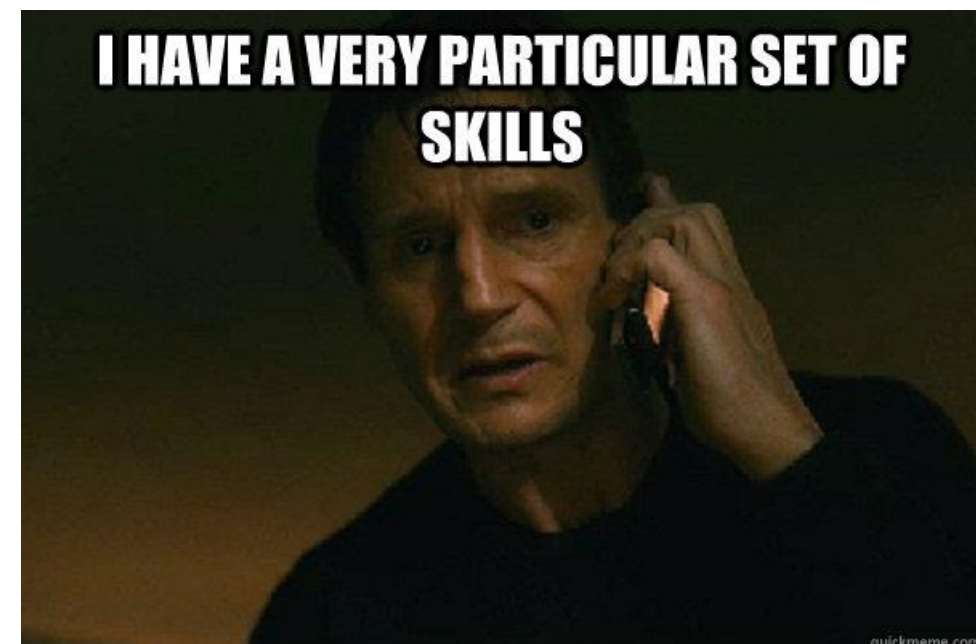
Machine Problems



- Implement and evaluate concepts from class in a commodity operating system
- Kernel Environment: Linux. Not a toy OS, but a real 25 million LoC behemoth.
- Why? Building out a small OS is good experience, but navigating an existing code base is a more practical skill.

- Recall from earlier:

"I attended a Microsoft-organized meeting where the Director of Engineering of (Microsoft in Redmond) talked to me about a great need for engineers who know operating systems/device drivers, and know linux kernel/programming at such lower levels. He bitterly complained that many CS departments are dismantling their OS programs. I told him that we have actually multiple OS undergraduate classes at UIUC the current instructor to advertise among the students who take these courses that there are many jobs at Microsoft in OS area (more than ever!)."





- You will need a kernel development environment.
- How many of you already setup the environment?
- How many of you never worked in any OS stuff?
- **MP 0: Setup a kernel development environment on your own machine**
 - Linux machine?
 - Macbook?
 - Chromebook?



- If you really don't have a machine, we can ask Engr-IT to create a VM for you. But, historically, nobody did well if they are relying on the VM.
- If you brick your machine (happens often), you will need to open a ticket with Engr-IT (≥ 24 hour delay)
- Brick your machine on a weekend? Too bad for you.
- Occasionally, the VM cloud just goes down! That's fun.



- Code repository
 - You will need to submit your source code
 - We will create a private GitHub repo for you.
 - Everything will be based on GitHub.

Four Credit Section



- Intended audience: graduate students, ambitious undergraduate students interested in research.
- Earn your 4th credit by reading and summarizing weekly literature assignments
- **Summaries due on the beginning of each class. The first summaries are due the first class of the 2nd week (8/30).**
- **Print your summary out and bring it to the class.**
- Assigned readings are marked as **C4** in the Assignments section of the class schedule. Other students are not required to read these papers.
- Grading: Summaries will contribute to C4 student's homework and participation credit.

C4 Paper Summaries



- Each summary should be about 1-2 pages in length.
- Structure your summary to cover:
 1. What are the *motivations* for this work?
 2. What is the proposed *solution*?
 3. What is the work's *evaluation* of the proposed solution?
 4. What is *your* analysis of the problem, idea and evaluation?
 5. What are the *contributions*?
 6. What are *future* directions for this research?
 7. What *questions* are you left with?
 8. What is your *take-away message* from this paper?

Grading



- Final Exam: **20%**
- Mid-term Exam: **20%**
- Machine Problems (5 total): **50%**
 - **4%, 12%, 12%, 12%, 10%**
- Participation: **10%**
 - Class/Forum involvement



- **No late homework/MP submissions**
- 1 week window for re-grades from return date
- **Cheating policy: Zero tolerance**
 - 1st offense: get zero
 - 2nd offense: fail class
 - Example: You submitted two MPs in which solutions were not your own. Both were discovered at the same time. You fail class.

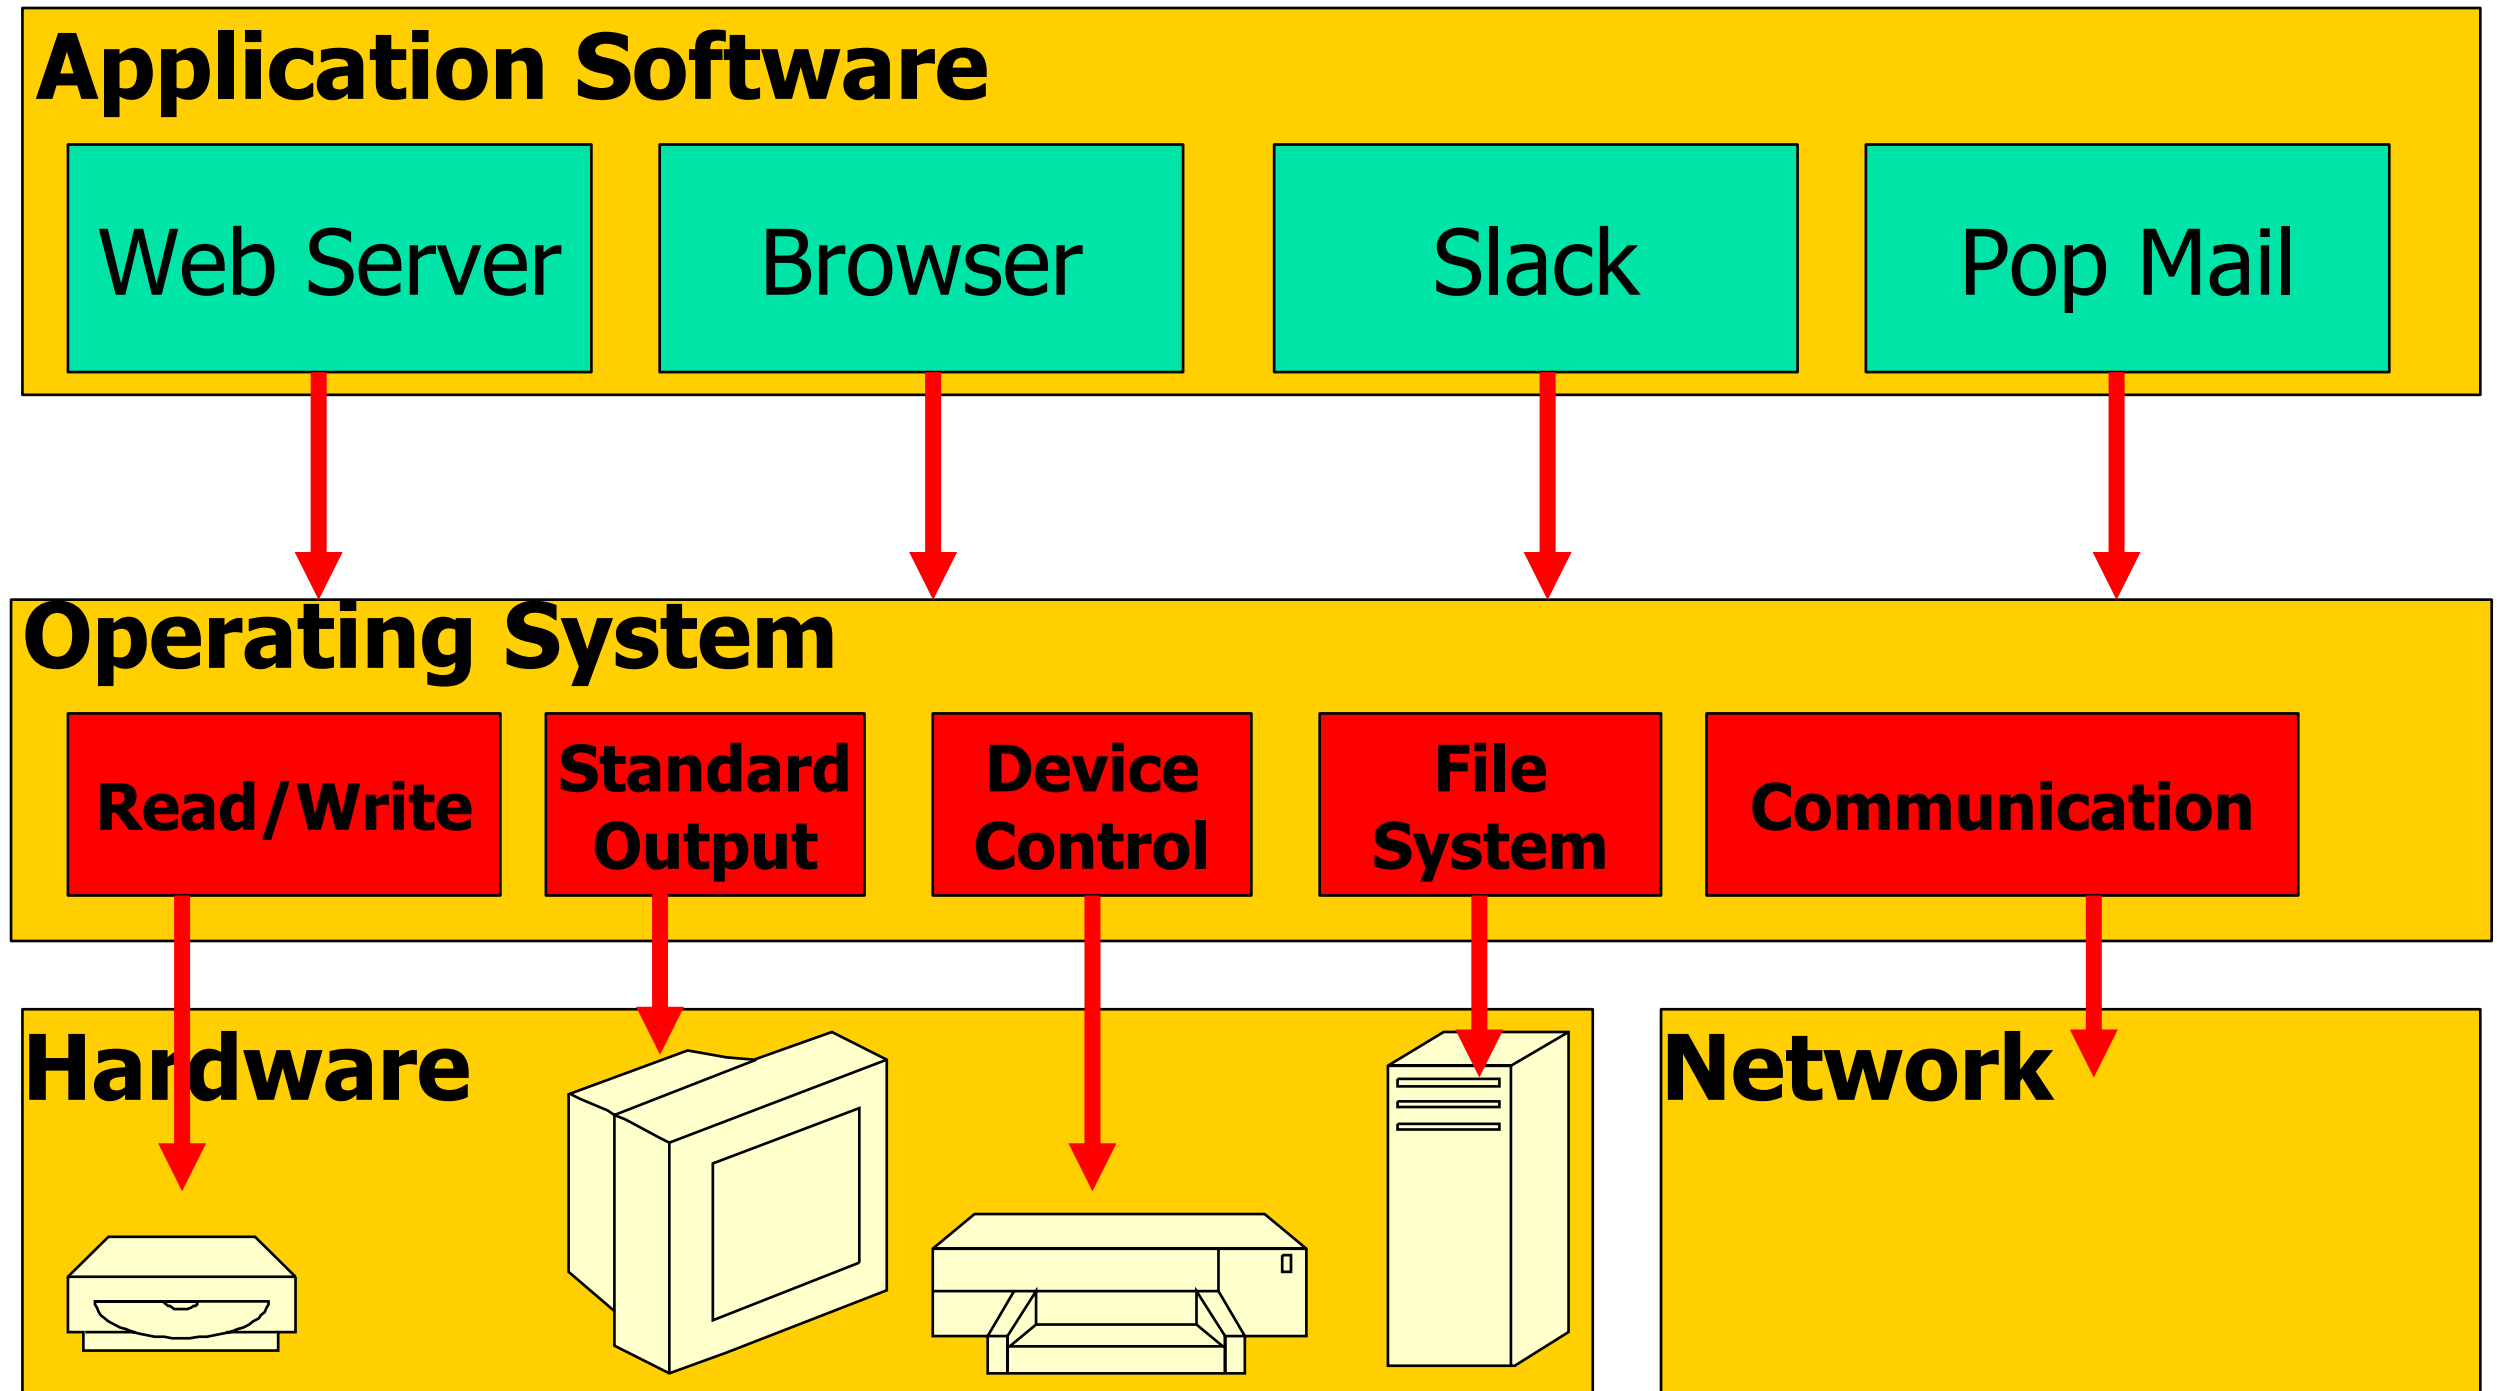


What is an operating system?

Why Operating Systems?



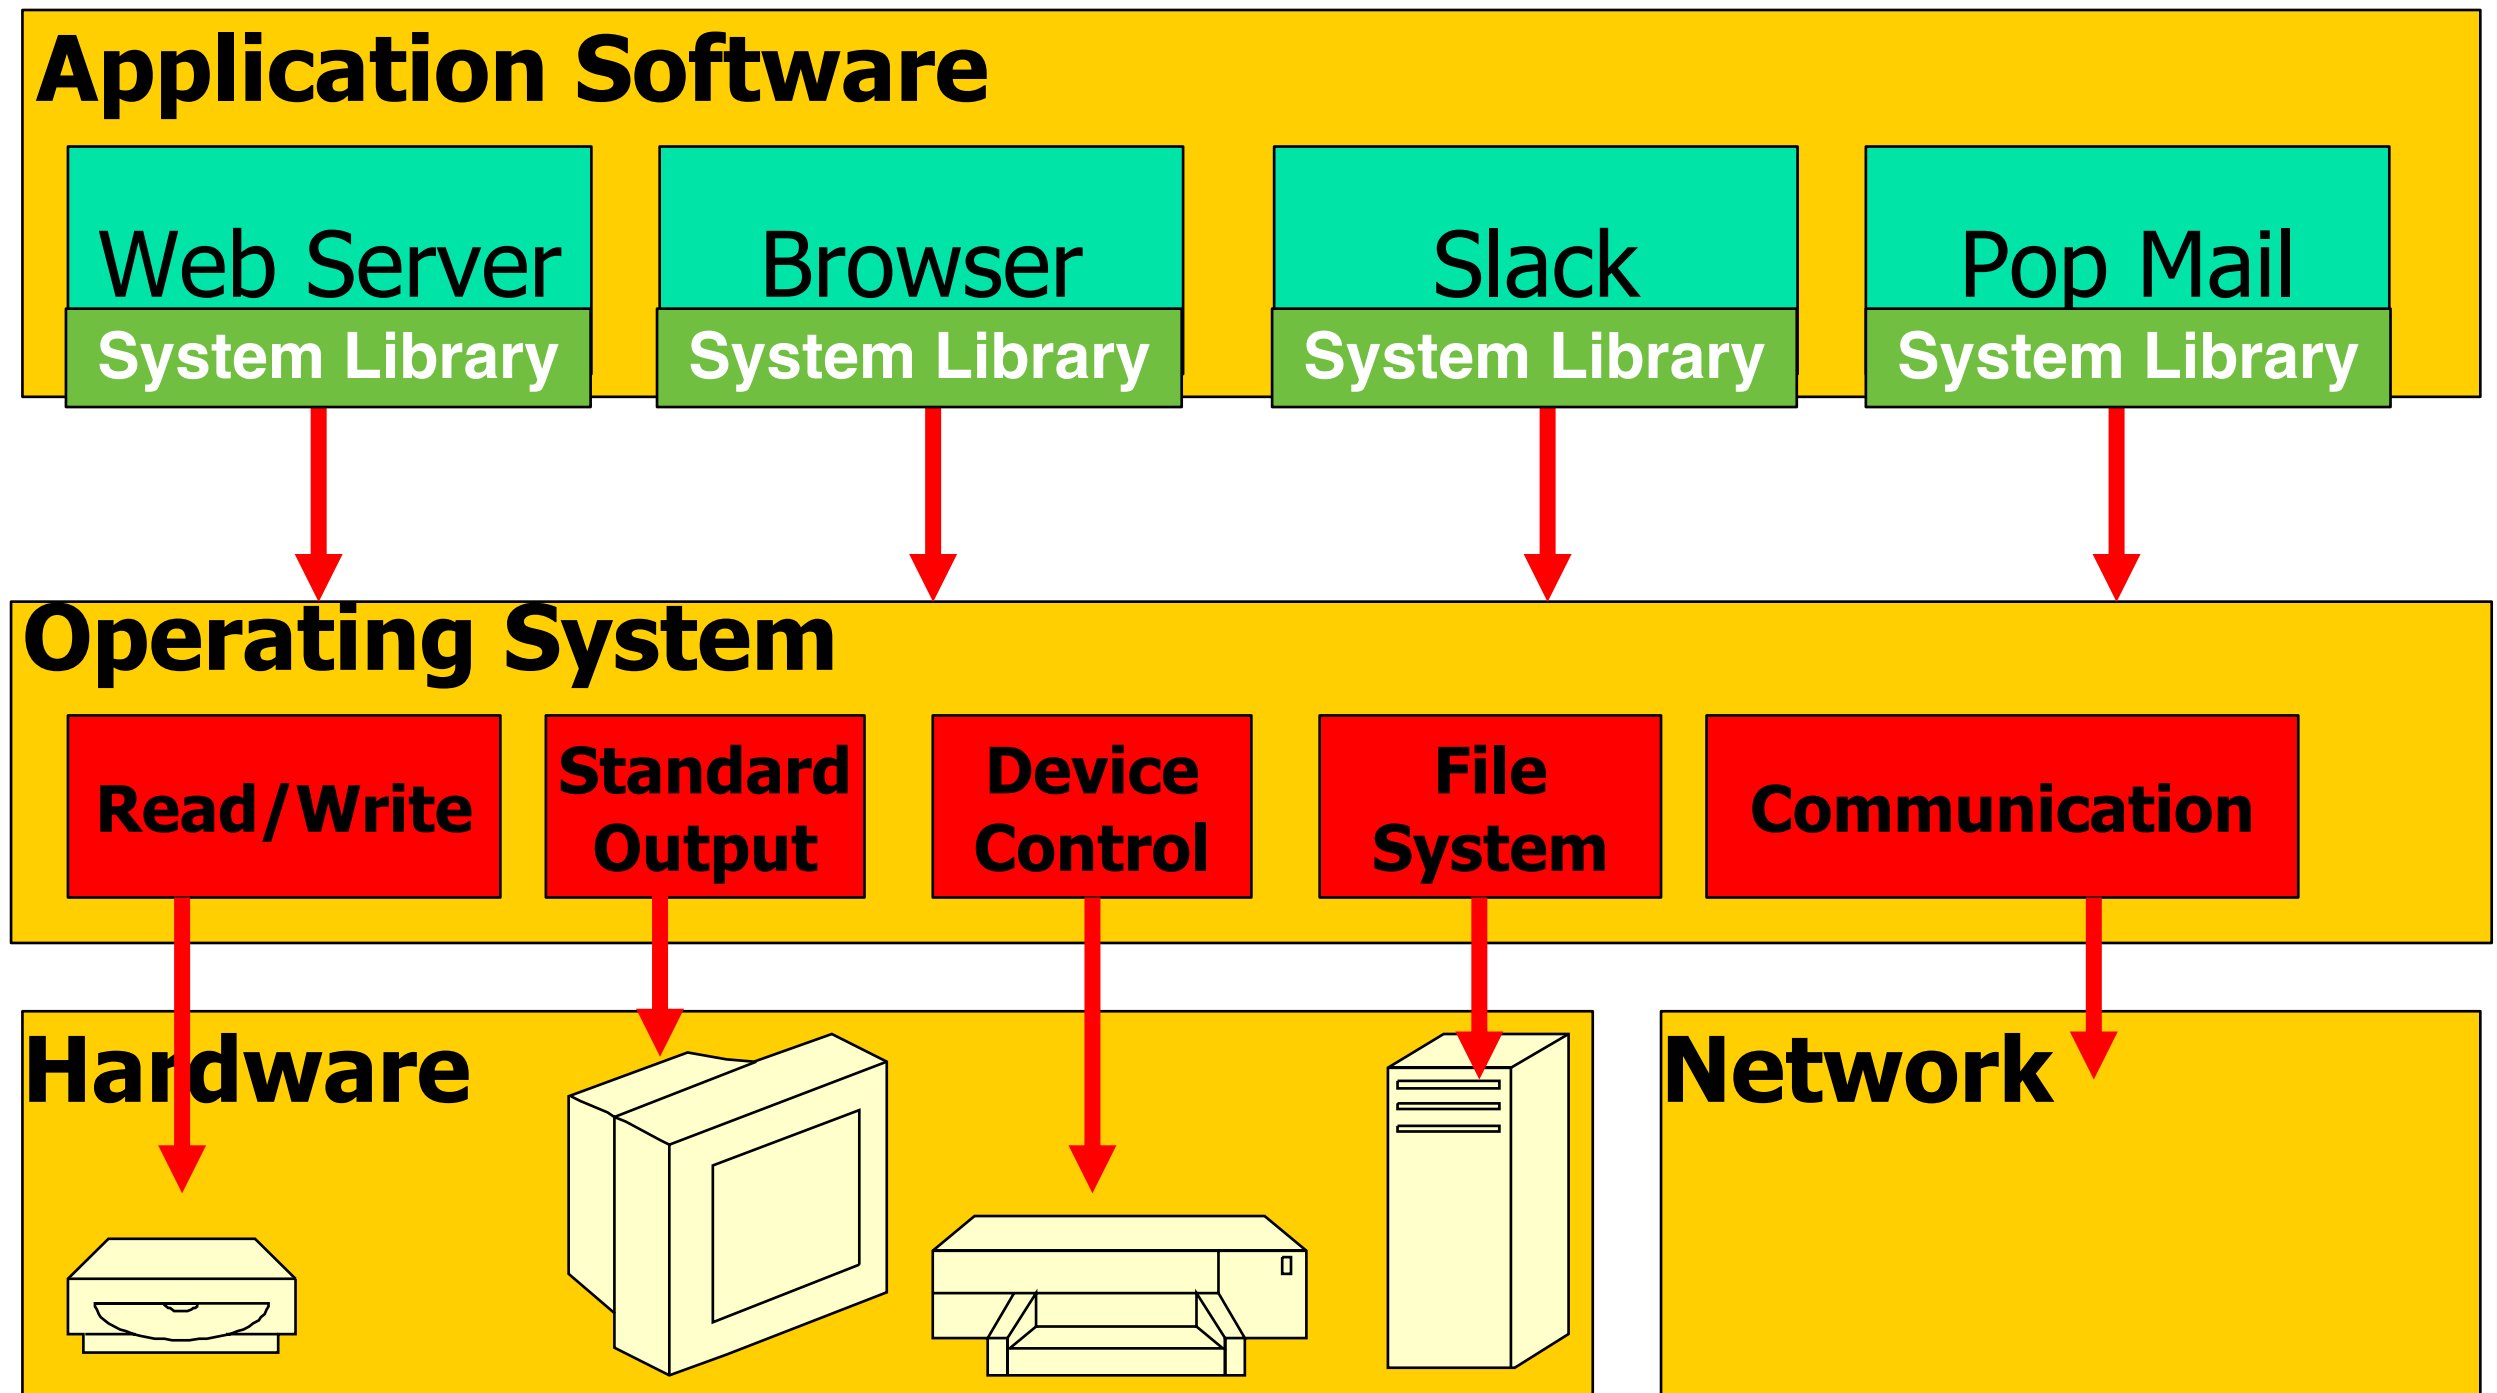
Software to manage a computer's resources for its user



Why Operating Systems?



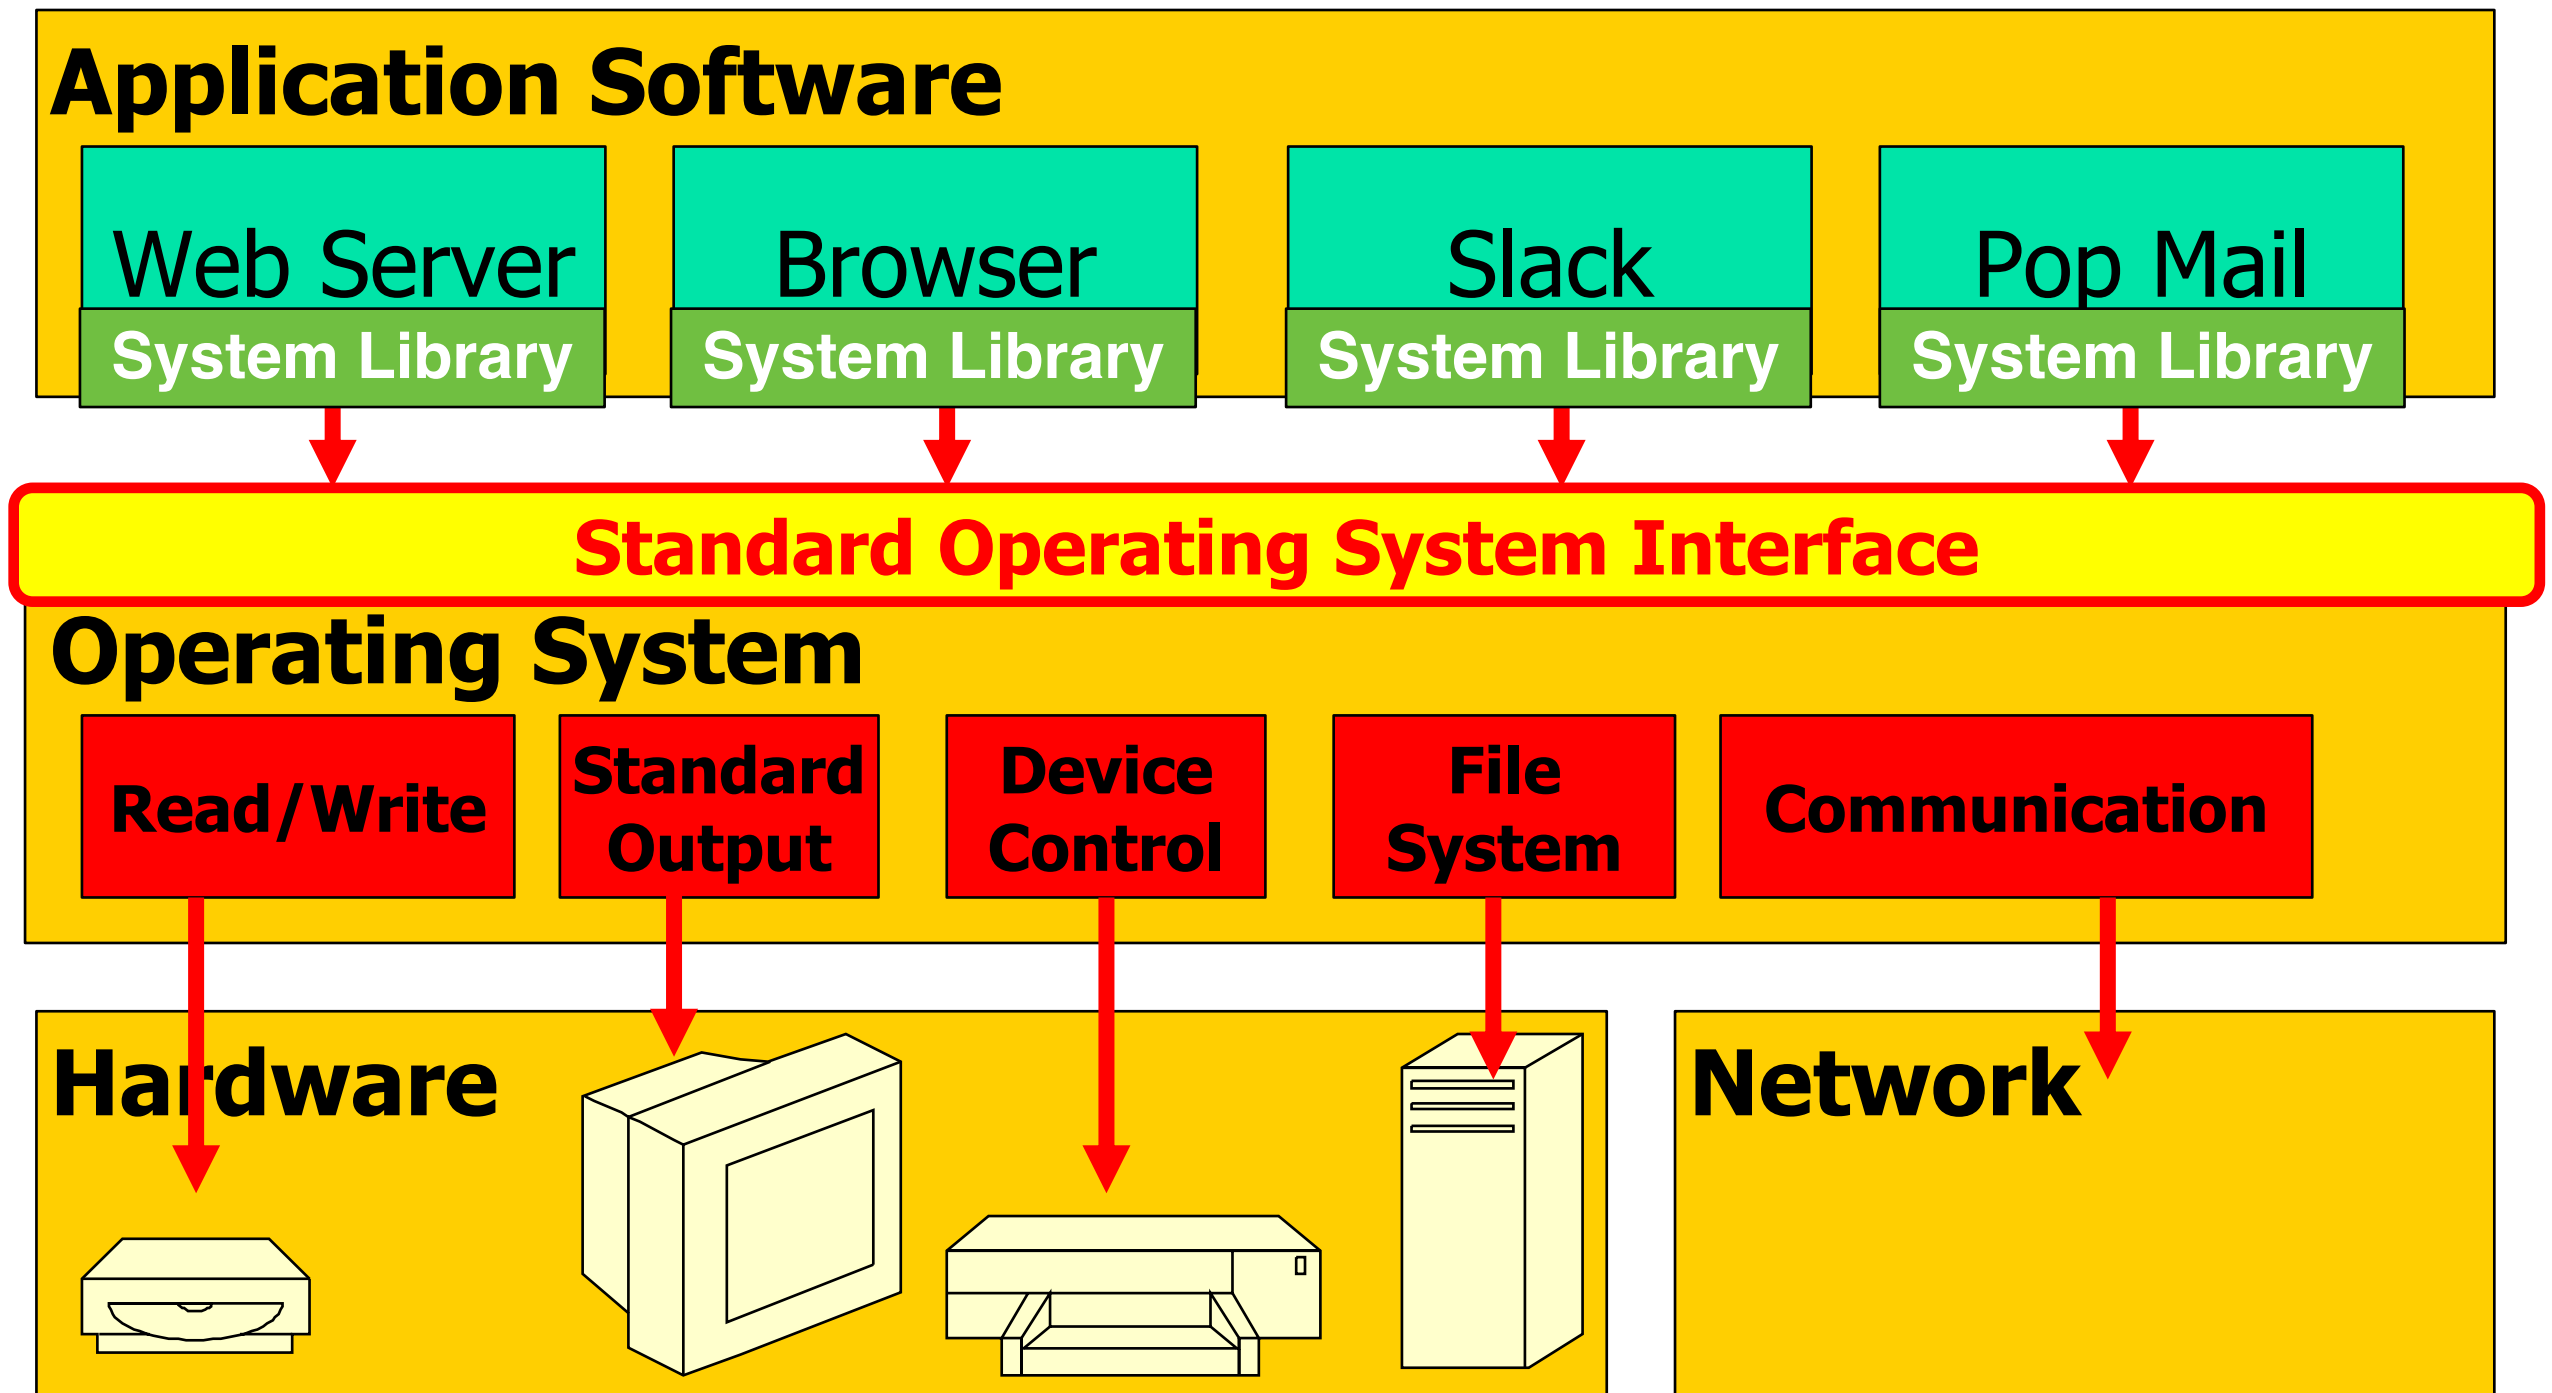
Software to manage a computer's resources for its user



Why Operating Systems?



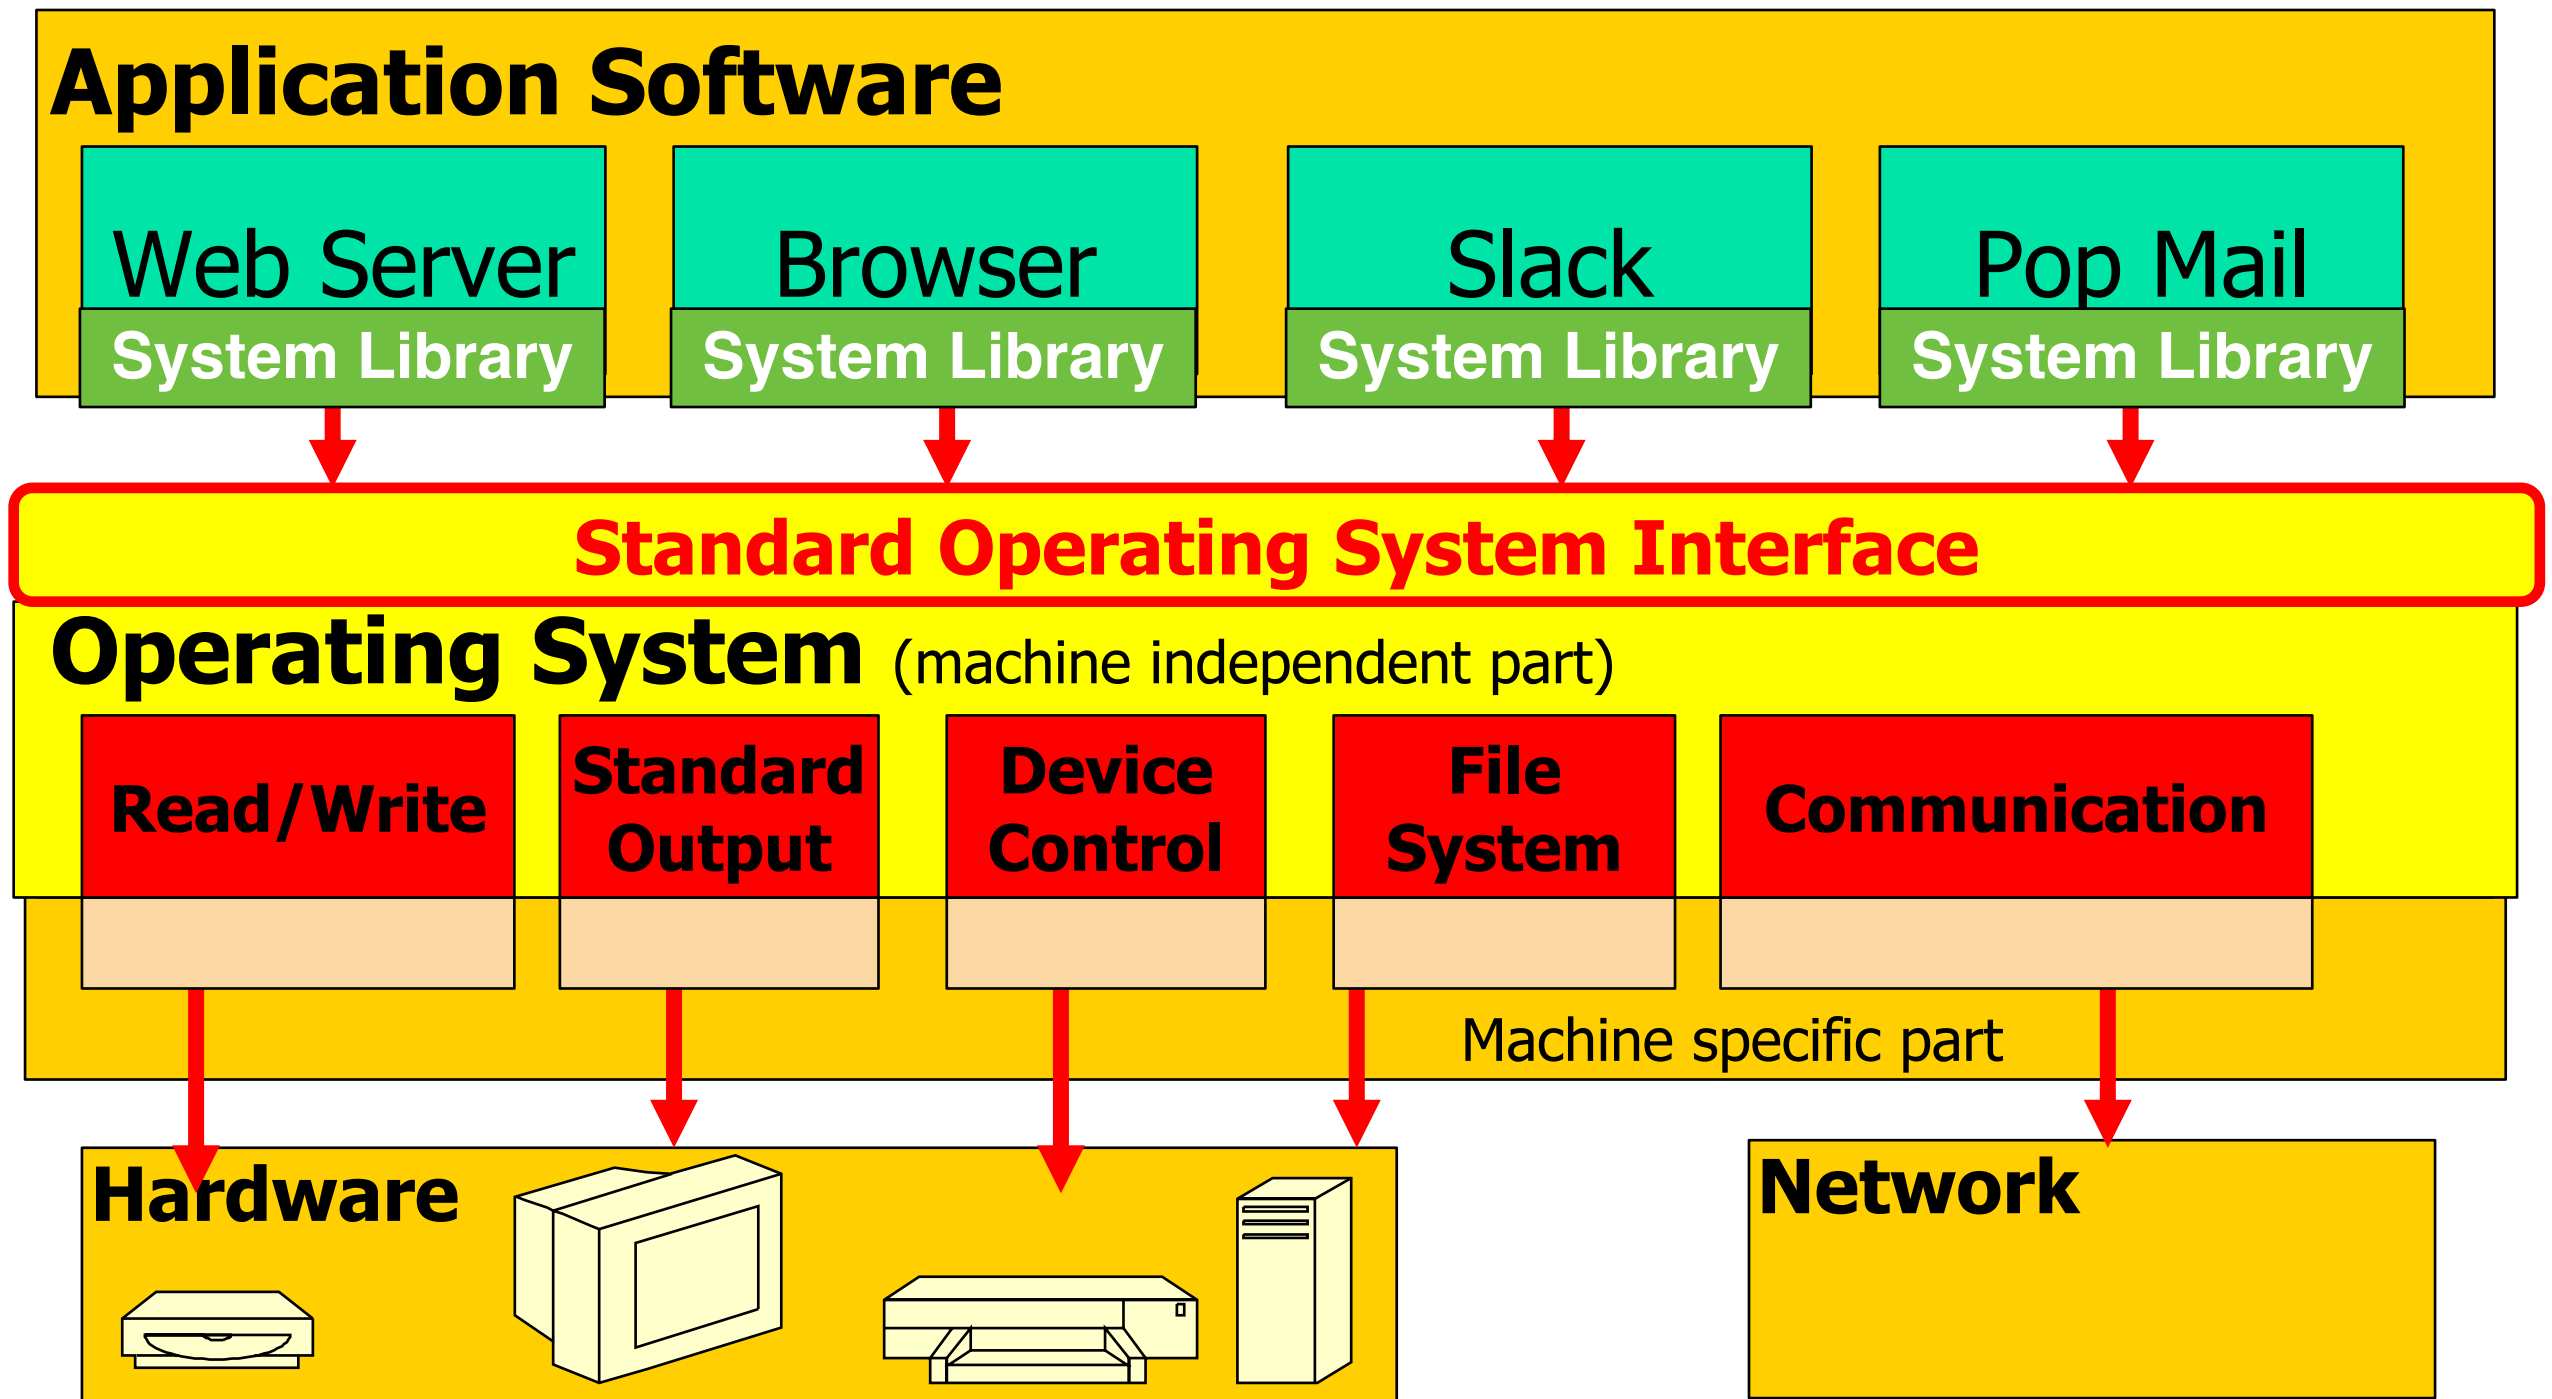
The OS exports a user interface. Why?



Why Operating Systems?




Standard interface increases portability and reduces the need for machine-specific code.



Why Operating Systems?



OS Runs on Multiple Platforms while presenting the same Interface:





What are the responsibilities of an operating system?



Role #1: Referee

- Manage resource allocation between users and applications
- Isolate different users and applications from one another
- Facilitate and mediate communication between different users and applications



Role #2: Illusionist

- Allow each application to believe it has the entire machine to itself
- Create the appearance of an Infinite number of processors, (near) infinite memory
- Abstract away complexity of reliability, storage, network communication...



Role #3: Glue

- Manage hardware so applications can be machine-agnostic
- Provide a set of common services that facilitate sharing among applications
- **Examples of “Glue” OS Services?**



Role #3: Glue

- Manage hardware so applications can be machine-agnostic
- Provide a set of common services that facilitate sharing among applications
- **Examples of “Glue” OS Services?**
 - Cut-and-paste, File I/O, User Interfaces...

Consider file systems and storage devices...

**How is the OS a referee?
An illusionist?
Glue?**

Ex: File System Support



Referee

- Prevent users from accessing each other's files without permission
- Even after a file is deleting and its space re-used

Illusionist

- Files can grow (nearly) arbitrarily large
- Files persist even when the machine crashes in the middle of a save

Glue

- Named directories, printf, other system calls for File I/O

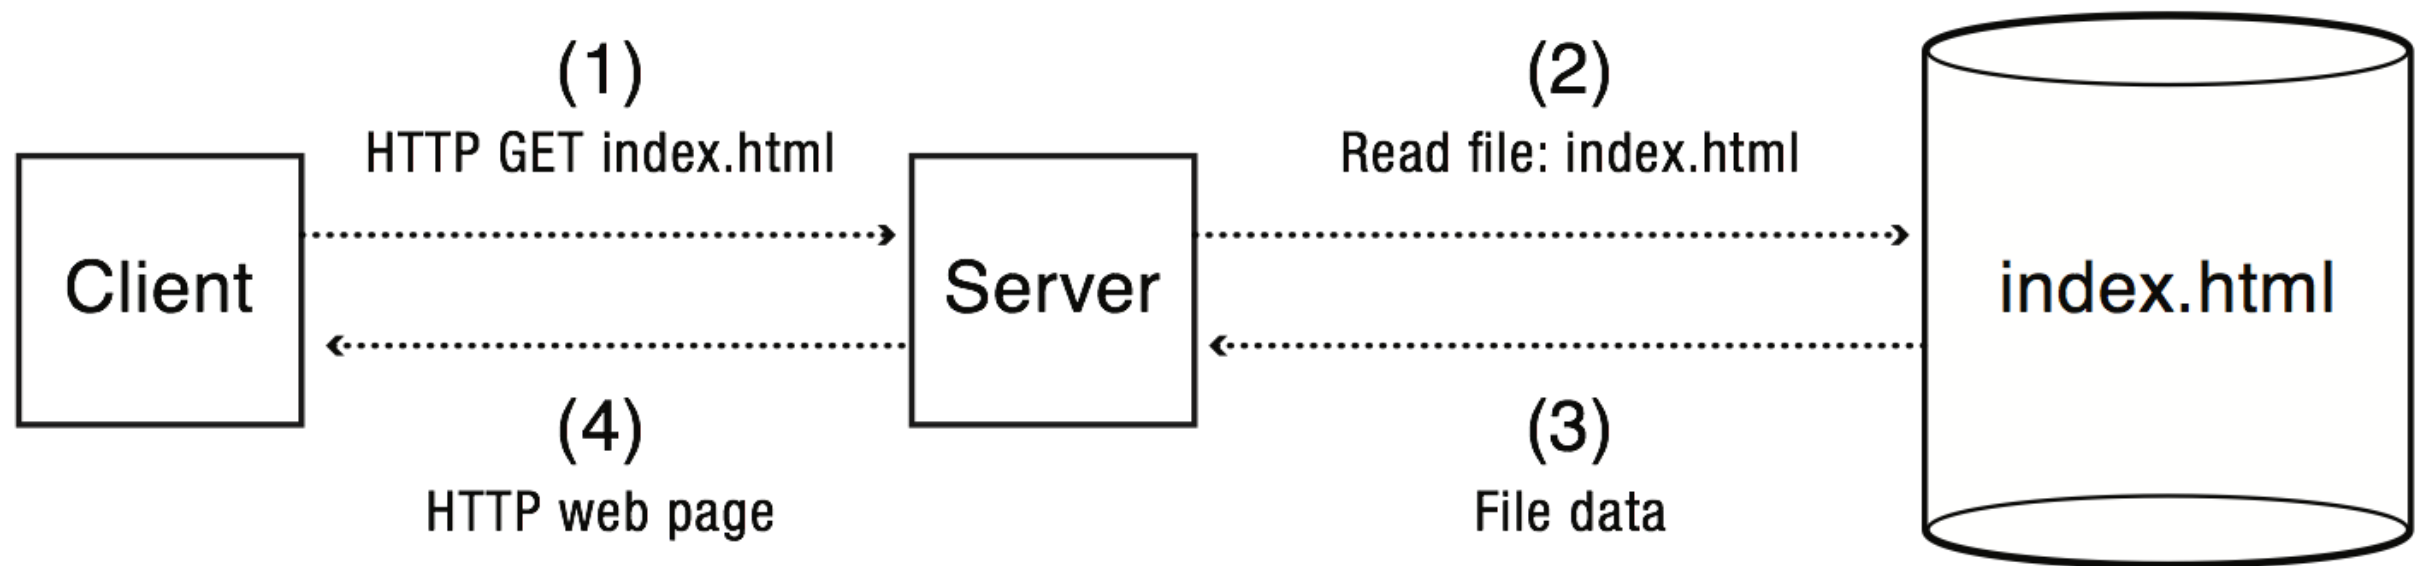


What does an OS need to do in order safely run an untrustworthy application?



How should an operating system allocate processing time between competing uses?

Example: Web Service



- How does the server manage many simultaneous client requests?
- How do we keep the client safe from spyware embedded in scripts on a web site?
- How do handles updates to the web site such that clients always see a consistent view?



Reliability

- Does the system do what it was designed to do?

Availability

- What portion of the time is the system working?
- Mean Time To Failure, Mean Time to Repair

Security

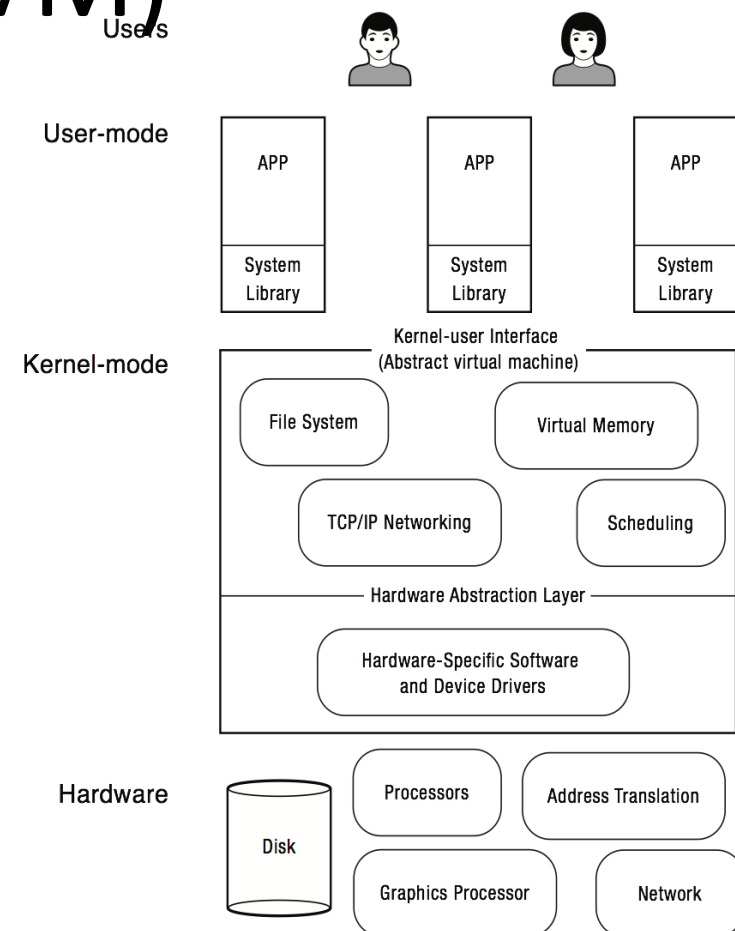
- Can the system be compromised by an attacker?

Privacy

- Data is accessible only to authorized users

Portability

- For programs:
Application programming interface (API)
Abstract virtual machine (AVM)
- For hardware
Hardware abstraction layer





Performance

Latency/response time

How long does an operation take to complete?

Throughput

How many operations can be done per unit of time?

Overhead

How much extra work is done by the OS?

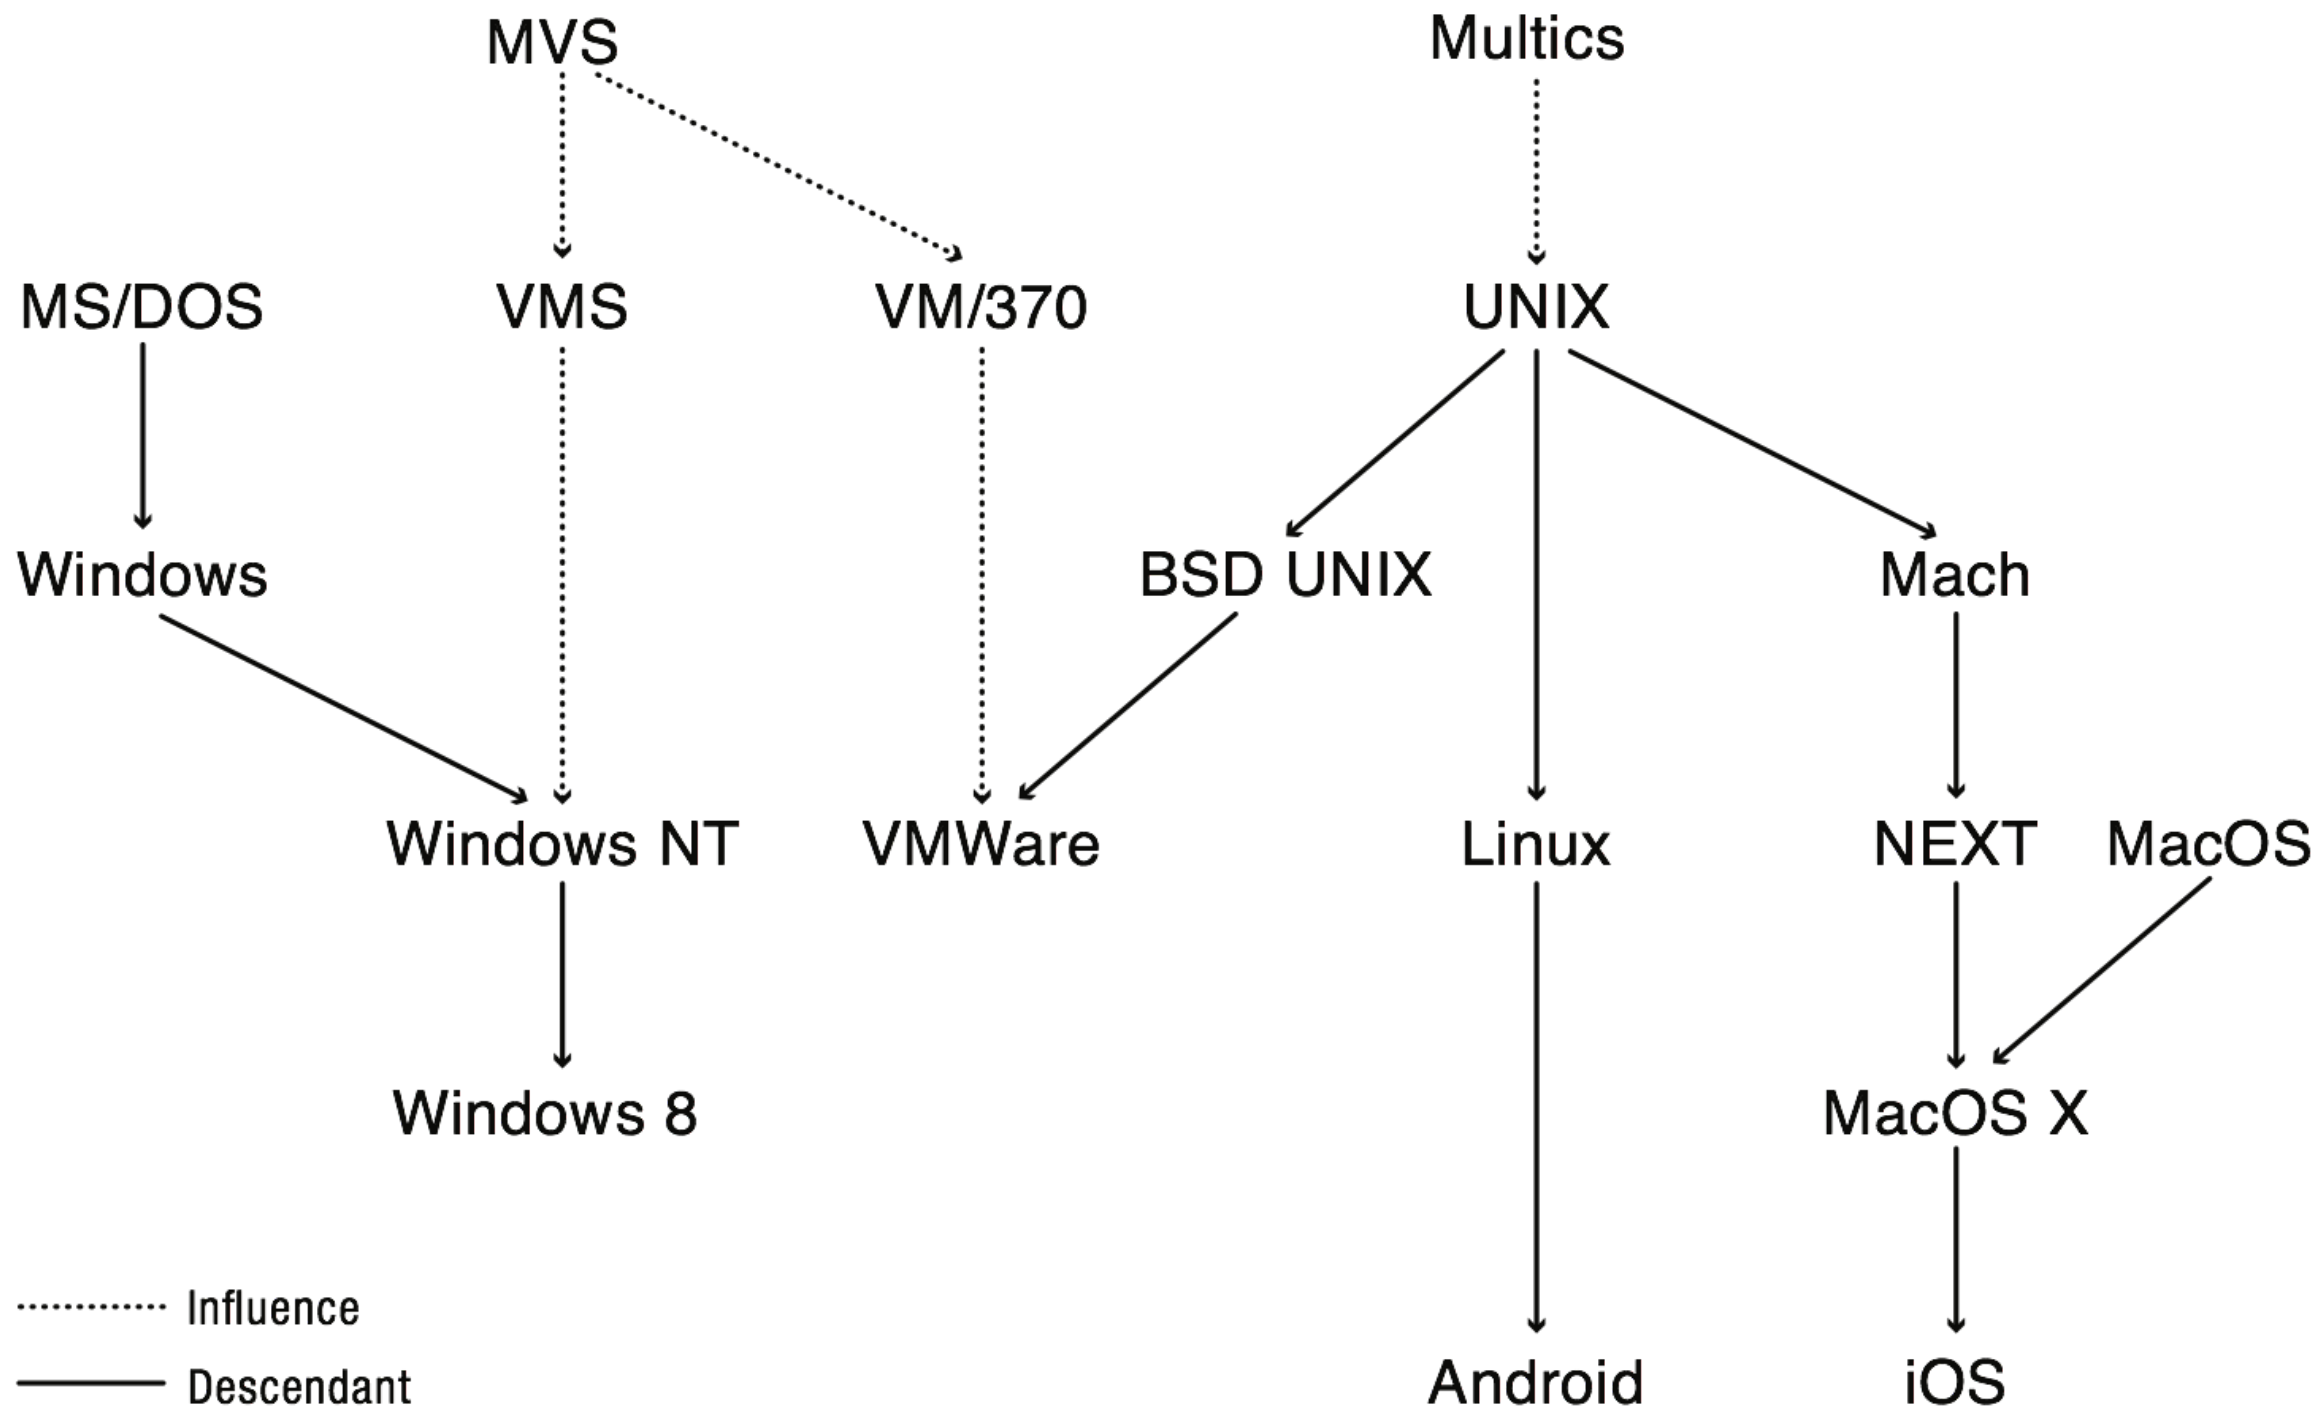
Fairness

How equal is the performance received by different users?

Predictability

How consistent is the performance over time?

OS Family Tree




Performance / Time



	1981	1997	2014	Factor (2014/1981)
Uniprocessor speed (MIPS)	1	200	2500	2.5K
CPUs per computer	1	1	10+	10+
Processor MIPS/\$	\$100K	\$25	\$0.20	500K
DRAM Capacity (MiB)/\$	0.002	2	1K	500K
Disk Capacity (GiB)/\$	0.003	7	25K	10M
Home Internet	300 bps	256 Kbps	20 Mbps	100K
Machine room network	10 Mbps (shared)	100 Mbps (switched)	10 Gbps (switched)	1000
Ratio of users to computers	100:1	1:1	1:several	100+

Early Operating Systems



One application at a time

- Had complete control of hardware
- OS was runtime library
- Users would stand in line to use the computer

Batch systems

- Keep CPU busy by having a queue of jobs
- OS would load next job while current one runs
- Users would submit jobs, and wait, and wait, and

Time-Sharing OSs



Multiple users on computer at same time

- Multiprogramming: run multiple programs at same time
- Interactive performance: try to complete everyone's tasks quickly
- As computers became cheaper, more important to optimize for user time, not computer time

Today's OSs



- Smartphones
- Embedded systems
- Laptops
- Tablets
- Virtual machines
- Data center servers

Tomorrow's OSes

