CS370 Operating Systems

Colorado State University
Yashwant K Malaiya

Fall 2024 L28
Final Review Part 2

Slides based on

- Text by Silberschatz, Galvin, Gagne
- Various sources

Project Slides/Videos

- Need slides (8-10) and videos (7-8 min) for both Research and Development Projects posted in channels
  - **Research Project Slides and Videos:** See deadlines there
  - **Devp Project Slides and Videos:** See deadlines there
    - Also need to sign-up for 15 min demos (Dec 2,3,4)
- Each student will need to view/evaluate by Dec 05.
  - 2 assigned project reports in Canvas (assigned Nov 23rd)
  - 7 videos/slides for A research projects
  - 3 videos/slides for B Development projects

Needed

.

- Please finish course survey (Available in Canvas) by ASAP, if not already done.
- Special Feedback Quiz for Distance Students available today Due Dec 10.

Final

- Final: Comprehensive but mostly from the second half. 2 Hours.
- Must have laptop with Respondus Lockdown browser installed and tested test quiz available
- Sec 001, 801 local: Th 12/12/2022, 9:40-11:40 AM
  - may not sit next to usual neighbors or fellow team members. May not leave the room without permission.
- Sec 801 non-local: Th 12/12/2022, Two hours, 9:40 AM-11:50 PM window (must start at 9:50 PM, those with accommodation must start earlier to finish by 11:50 PM)

Grading

- Project D1, D2, D3, D4, D5 (raw/adjusted)
- Participation (raw/adjusted)
- Final (raw/adjusted)
- Letter Grades
  - Default: Given on course website
    - ≥ 90 is an A, ≥ 88 is an A-, ≥86 is a B+, ≥80 is a B, ≥78 is a B-, ≥76 is a C+, ≥70 is a C, ≥60 is a D, and <60 is an F.
  - *may* cut lower

Study/Resources

- Terms, concepts, implementations, algorithms, problems
- Lecture slides
  - Also see Midterm Review Slides on website
  - Possible questions not limited to Review Slides
- Quizzes, assignments
- Textbook

HW6

- Discuss after the review.

Effective Access Time

- Hit ratio = α
  - Hit ratio – percentage of times that a page number is found in the TLB
- Associative Lookup = ε time unit
- Memory access time = 100 ns

- **Effective Access Time** (**EAT**)

  EAT = $(100 + \varepsilon)\ \alpha + (200 + \varepsilon)(1 - \alpha)$

Consider α = 80%, ε = 20ns for TLB search, 100ns for memory access

  - EAT = 120 x 0.80 + 220 x 0.20 = 140ns

- Consider higher hit ratio -> α = 99%, ε = 20ns for TLB search, 100ns for memory access
  - EAT = 120 x 0.99 + 220 x 0.01 = 121ns

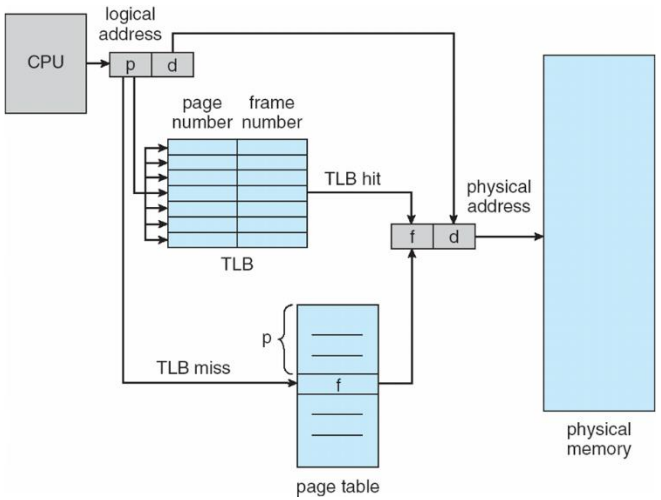

Colorado State University

Least Recently Used (LRU) Algorithm

- Use past knowledge rather than future
- Replace page that has not been used in the most amount of time
- Associate time of last use with each page

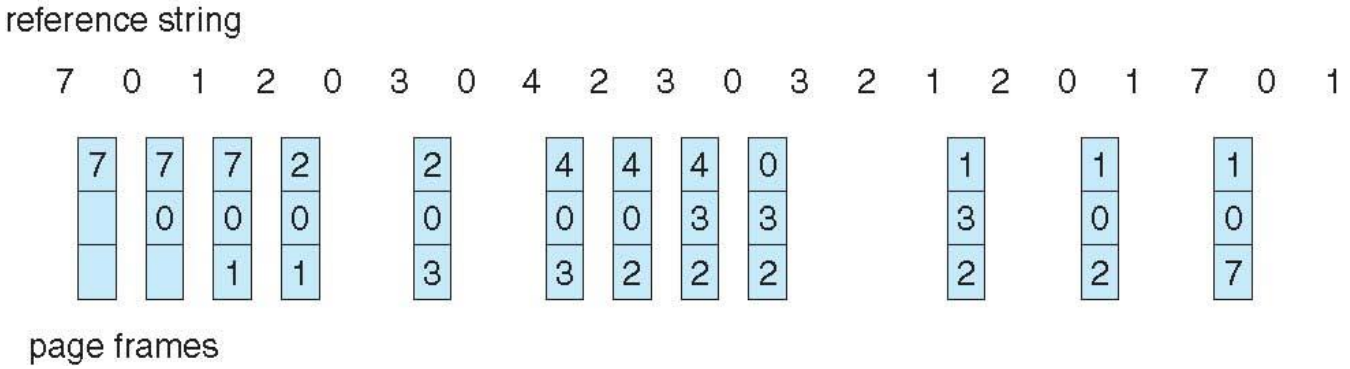

- Blank: implies a Hit with no page fault.
- 12 faults – better than FIFO but worse than OPT
- Generally good algorithm and frequently used
- Approximate Implementations:
  - Counter implementation time of use field
  - Stack implementation
  - Reference bit
  - Second chance

Hadoop: Core components

- Hadoop (originally): MapReduce + HDFS
- For **Big Data** applications.
- **MapReduce**: A programming framework for processing parallelizable problems across huge datasets using a large number of commodity machines.
- **HDFS**: A **d**istributed **f**ile **s**ystem designed to efficiently allocate data across multiple machines, and provide self-healing functions when some of them go down

HDFS Architecture

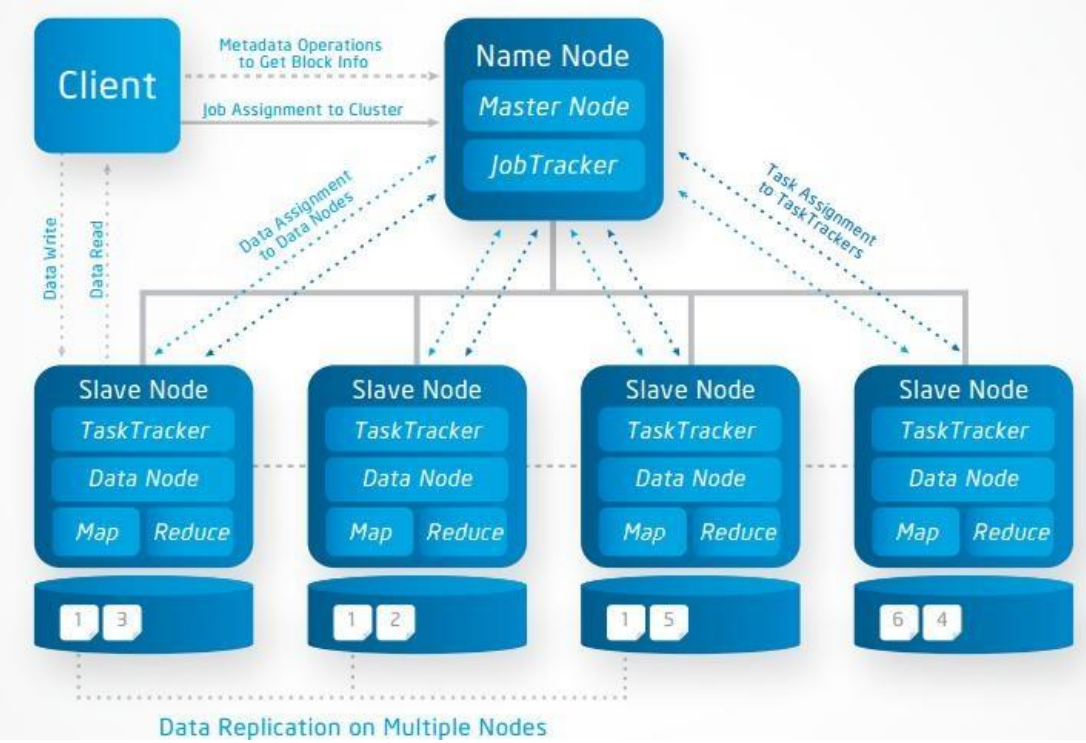

HDFS Block size: 64-128 MB
ext4: 4KB
HDFS is on top of a local file system.

Name Node: metadata, where blocks are physically located
Data Nodes: hold blocks of files (files are distributed)

http://a4academics.com/images/hadoop/Hadoop-Architecture-Read-Write.jpg

HDFS Fault-tolerance

- Individual node/rack may fail.
  - Disks use error detecting codes to detect corruption.
- Data Nodes (on slave nodes):
  - data is replicated. Default is 3 times. Keep a copy far away.
  - Send periodic heartbeat (I'm OK) to Name Nodes. Perhaps once every 10 minutes.
  - Name node creates another copy if no heartbeat.
- Name Node (on master node) Protection:
  - Transaction log for file deletes/adds, etc (only metadata recorded).
  - Creation of more replica blocks when necessary after a DataNode failure
- Standby name node: namespace backup
  - In the event of a failover, the Standby will ensure that it has read all of the edits from the Journal Nodes and then promotes itself to the Active state

Implementation of VMMs

- **Type 1 hypervisors** - Operating-system-like software built to provide virtualization. Runs on 'bare metal".
  - Including VMware ESX, Joyent SmartOS, and Citrix XenServer
- Also includes general-purpose operating systems that provide standard functions as well as VMM functions
  - Including Microsoft Windows Server with HyperV and RedHat Linux with KVM
- **Type 2 hypervisors** - Applications that run on standard operating systems but provide VMM features to guest operating systems
  - Includiing VMware Workstation and Fusion, Parallels Desktop, and Oracle VirtualBox

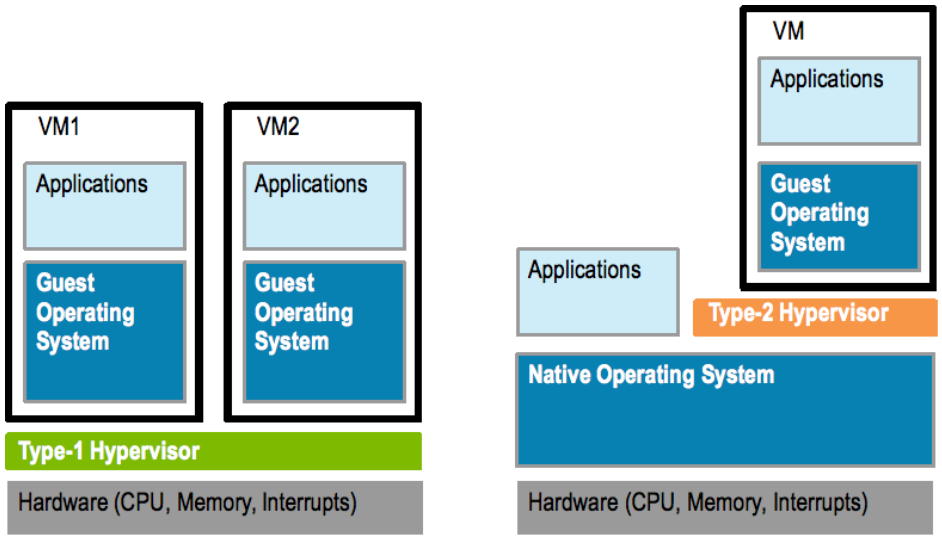

Colorado State University

Memory Management

Memory mapping:

- On a bare metal machine:
  - VPN -> PPN
- VMM: Real physical memory (*machine memory*) is shared by the OSs. Need to map PPN of each VM to MPN (Shadow page table)

  PPN ->MPN
- Where is this done?
  - In Full virtualization?

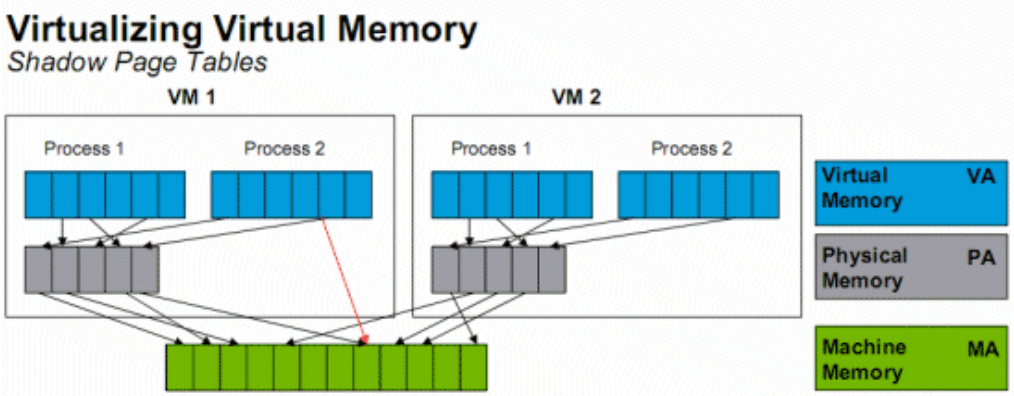

Live Migration

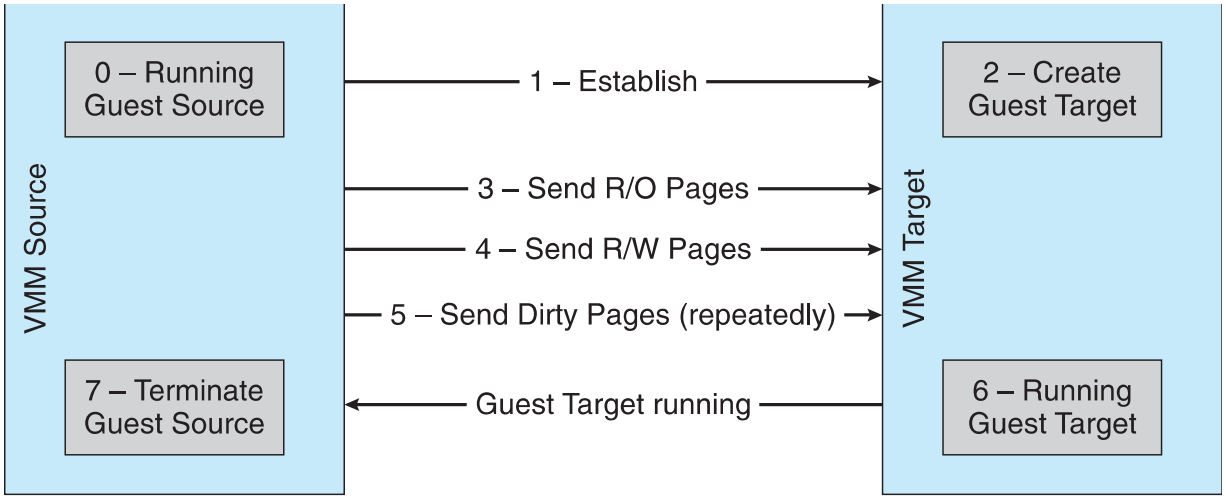

- Migration from source VMM to target VMM
  - Source establishes a connection with the target
  - Target creates a new guest
  - Source sends all read-only memory pages to target
  - Source starts sending all read-write pages
  - Source VMM freezes guest, sends final stuff,
  - Once target acknowledge

Linux Containers and Docker

- Linux containers (LXC) are “lightweight” VMs
- Comparison between LXC/docker and VM

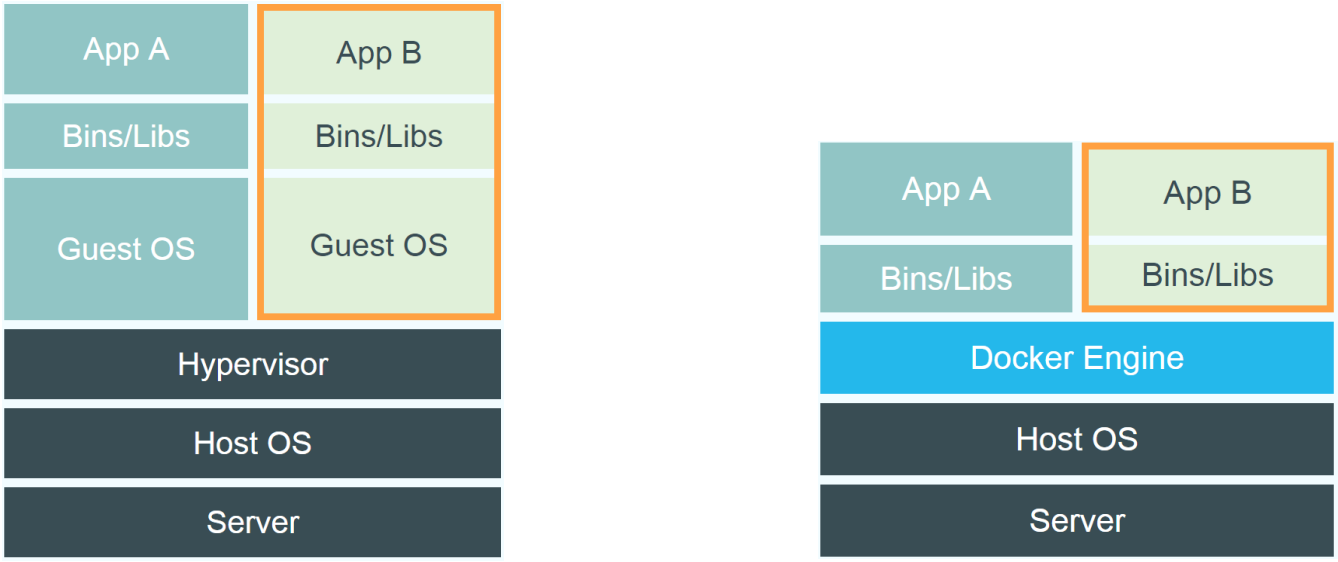

- Containers provide “OS-level Virtualization” vs “hardware level”.
- Containers can be deployed in seconds.
- Very little overhead during execution, just like Type 1.

Microservices Characteristics

- Many smaller (fine grained), clearly scoped services
  - Single Responsibility Principle
  - Independently Managed
- Clear ownership for each service
  - Typically need/adopt the “DevOps” model
- 100s of MicroServices
  - Need a Service Metadata Registry (Discovery Service)
- May be replicated as needed
- A microservice can be updated without interruption

Cloud Capacity provisioning

User has a variable need for capacity. User can choose among

Fixed resources: Private data center

- Under-provisioning when demand is too high, or
- Provisioning for peak

Variable resources:

- Use more or less depending on demand
- Public Cloud has elastic capacity (i.e. way more than what the user needs)
- User can get exactly the capacity from the Cloud that is actually needed

Why does this work for the provider?

  - Varying demand is statistically smoothed out over many users, their peaks may occur at different times
  - Prices set low for low overall demand periods

Cloud Instance types/Service/Management models

Instance types

- On-Demand instances
- Spot Instances
- Reserved Instances
- Dedicated Hosts

Service models

- IaaS: Infrastructure as a Service
- PaaS: Platform as a Service
- SaaS: Software as a Service

Cloud Management models

- Public clouds
- Private clouds
- Hybrid clouds:

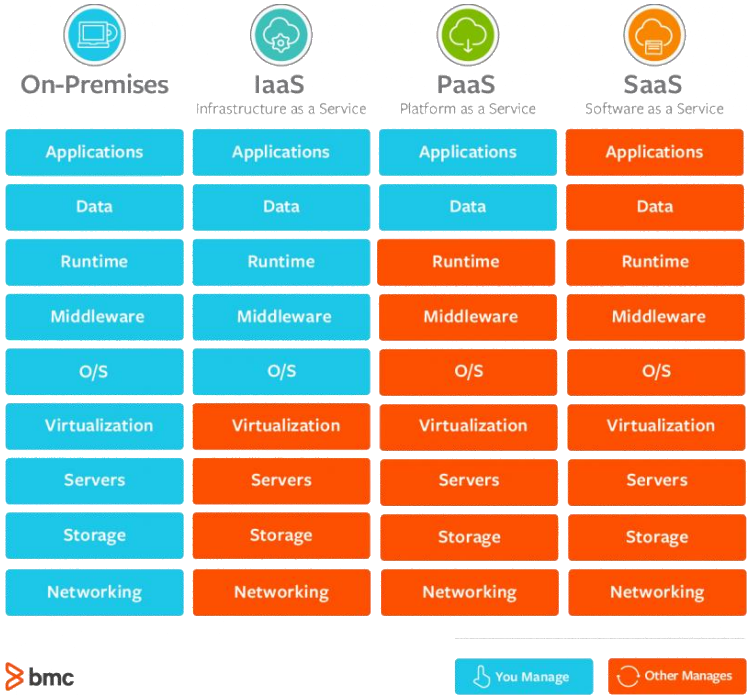

Colorado State University

Assets, Risk, Threat, Vulnerability

System Resource (Asset): what needs protection by the defenders.

Risk: A measure of the adverse impacts and the likelihood of occurrence.

Threat: potential attempts by an adversary.

Vulnerability: Weakness in an information system that may be exploited.

Note of caution: In pre-cyber-security days, classical risk literature used the term vulnerability with a different meaning.

RFC 2828, Internet Security Glossary

Cyber attack types

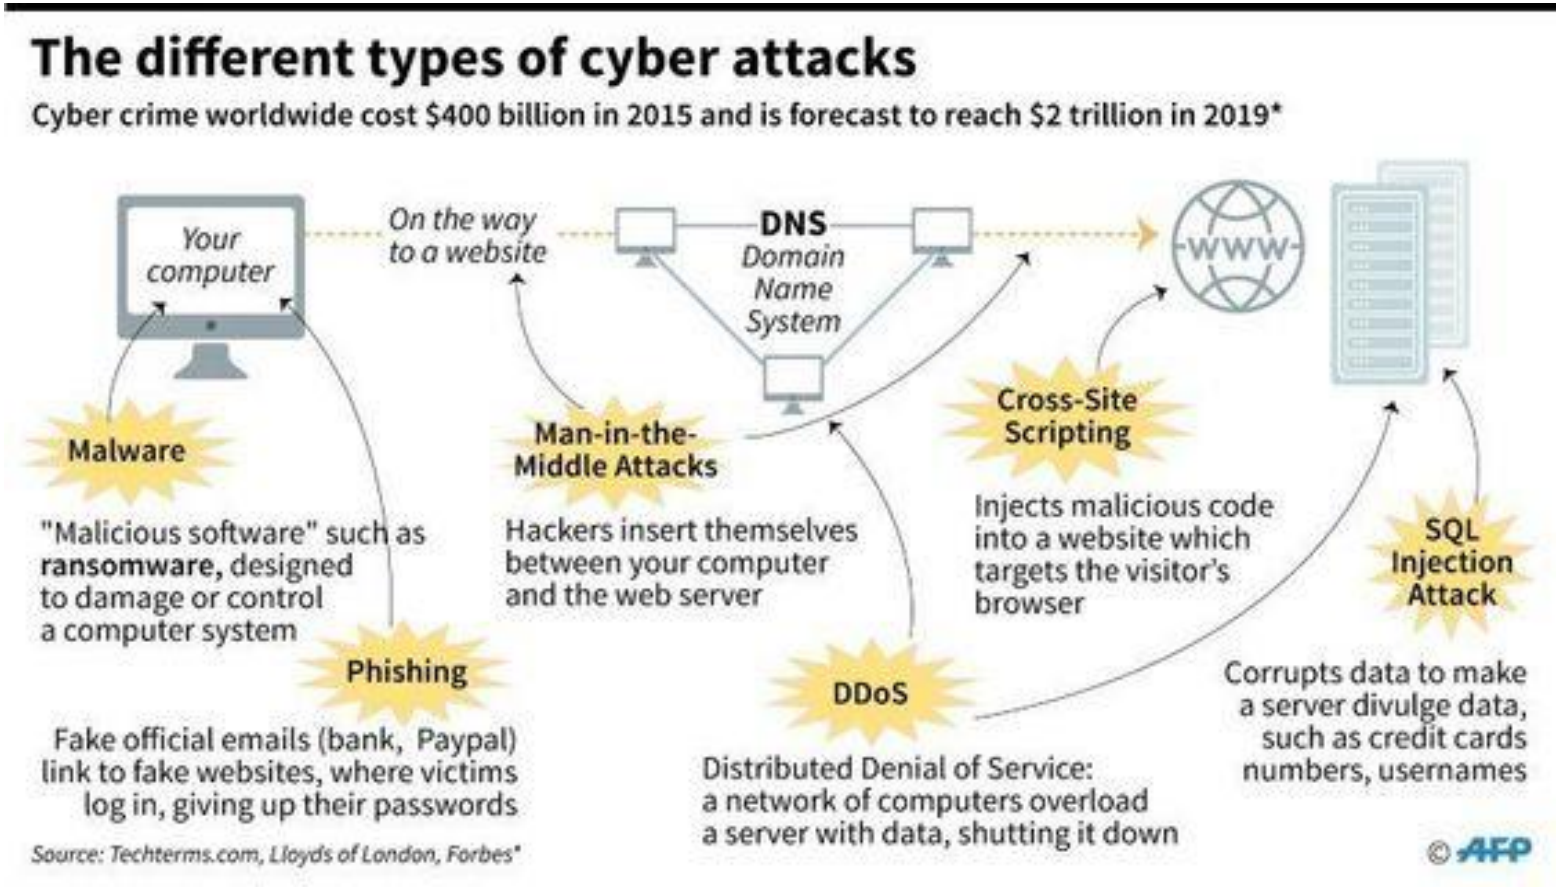

Example: Access Control Matrix

		OBJECTS			
		File 1	File 2	File 3	File 4
SUBJECTS	User A	Own Read Write		Own Read Write	
	User B	Read	Own Read Write	Write	Read
	User C	Read Write	Read		Own Read Write

(a) Access matrix

Access Control List (ACL): Every object has an ACL that identifies what operations subjects can perform. Each access to object is checked against object's ACL.

May be kept in a relational database. Access recorded in file metadata (inode).

Authentication Methods

Three existing and two new.

- Something a user knows
  - Password, answers to questions
- Something a user has
  - Ex. Id card, Phone
- Something a user is
  - Biometric (face, iris, fingerprint)
- Somewhere you are geographically
- Something you do based on recognizable patterns of behavior

- Can be multifactor to reduce false positives
- After-access confirmation

See you in the Final

Colorado State University