

TCSS 558: APPLIED DISTRIBUTED COMPUTING

**Processes:
Threads & Virtualization,
Clients & Servers**

Wes J. Lloyd
 School of Engineering
 & Technology (SET)
 University of Washington - Tacoma

OBJECTIVES - 1/28

- **Questions from 1/26**
- **Assignment 0: Cloud Computing Infrastructure Tutorial**
 - New testFibService.sh script
- **Assignment 1: Key/Value Store**
- **Chapter 3: Processes**
 - Chapter 3.1: Threads
 - Context Switches
 - Threading Models
 - Multithreaded clients/servers
 - Chapter 3.2: Virtualization
 - Chapter 3.3: Clients
 - Chapter 3.4: Servers

January 28, 2021
TCSS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.2

ONLINE DAILY FEEDBACK SURVEY

- Daily Feedback Quiz in Canvas - Available After Each Class
- Extra credit available for completing surveys **ON TIME**
- Tuesday surveys: due by ~ Wed @ 10p
- Thursday surveys: due ~ Mon @ 10p

TCSS 558 A > Assignments

Winter 2021

Home

Announcements

Assignments

Zoom

Chat

Upcoming Assignments

TCSS 558 - Online Daily Feedback Survey - 1/5
 Not available until Jan 5 at 1:30pm | Due Jan 6 at 10pm | ~75 pts

January 28, 2021
TCSS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.3

TCSS 558 - Online Daily Feedback Survey - 1/5

Due Jan 6 at 10pm Points 1 Questions 4
 Available Jan 5 at 1:30pm - Jan 6 at 11:59pm 1 day Time Limit None

Question 1 0.5 pts

On a scale of 1 to 10, please classify your perspective on material covered in today's class:

1	2	3	4	5	6	7	8	9	10
Mostly Review To Me		Equal New and Review				Mostly Now To Me			

Question 2 0.5 pts

Please rate the pace of today's class:

1	2	3	4	5	6	7	8	9	10
Slow			Just Right				Fast		

January 28, 2021
TCSS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.4

MATERIAL / PACE

- Please classify your perspective on material covered in today's class (22 respondents):
- 1-mostly review, 5-equal new/review, 10-mostly new
- **Average - 7.29** (↑ - previous 6.73)
- Please rate the pace of today's class:
- 1-slow, 5-just right, 10-fast
- **Average - 5.81** (↑ - previous 5.50)

January 28, 2021
TCSS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.5

FEEDBACK FROM 1/26

- **The hybrid architecture remains the least clear to me.**
- A **hybrid** architecture combines **more than one** architecture:
- **EDGE COMPUTING EXAMPLE:**
- **End of network:** Unstructured peer-to-peer
- **Edge Server:** Centralized routes to Internet
- **Cloud data center:** Heterogeneous virtual machines: Structured peer-to-peer

January 28, 2021
TCSS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.6

FEEDBACK - 2

- **Furthermore, I am wondering about the difference between unstructured and structured peer to peer system.**
- **Structured peer-to-peer:**
 Features deterministic message routes and routing times
 More rigid, changes require reconfiguration
 - Dependable messaging
- **Unstructured peer-to-peer:**
 Message routes need to be discovered
 Topology is constantly changing
 Changes are discovered on-the-fly
 - Common with wireless systems where devices have unreliable networks and power sources
 - Best effort messaging

January 28, 2021 TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma L8.7

FEEDBACK - 3

- **For Dynamic Topology - Chord System, how is the shortest path $O(\log N)$? (N is the number of nodes)**
- Chord provides an alternative to implement a DHT but without a fixed size such as with the four-dimensional hypercube
- Each node keeps a finger table containing m entries
 - m is the number of bits in the hash key
- A query is sent to an arbitrary node
- The node will look up the hash k in the finger table
- The finger table identifies the node to send the query to
- Nodes in the chord system are responsible for maintaining up-to-date finger tables

January 28, 2021 TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma L8.8

HOW TO COMPUTE FINGER TABLE (FT)

- i^{th} entry in ft at peer with id n is **first node** $\geq (n+2^i) \pmod{2^m}$
- For our example hash has 4 bits ($m=4$)
- Consider that we have 5 nodes
- Let's compute the finger table for **n3**
- Everytime a node wants to lookup a key it will pass the query to the **first node** which is the closest successor or predecessor (going clockwise) of k in it's finger table
- **N3**

i	ft[i]
4	n6 (3+2 ⁰) mod 2 ⁴ hash l=0
5	n6 (3+2 ¹) mod 2 ⁴ hash l=1
7	n10 (3+2 ²) mod 2 ⁴ hash l=2
11	n13 (3+2 ³) mod 2 ⁴ hash l=3

January 28, 2021 TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma L8.9

5-NODE CHORD SYSTEM

- Consider a 5 node Chord system with a 4-bit hash
- A query is sent to an arbitrary node

January 28, 2021 TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma L8.10

TO FIND THE DATA

- To lookup a item with hash key k , the node will pass the query to the closest successor or predecessor of k in the finger table (the node with the highest ID in the circle whose ID is smaller than k)
- If $k=8$ and the query first goes to node $n3$
- Query is passed to node $n10$
- Data each node is responsible for storing in this 5-node chord:

n0	$k=\{14,15,0\}$
n3	$k=\{1,2,3\}$
n6	$k=\{4,5,6\}$
n10	$k=\{7,8,9,10\}$
n13	$k=\{11,12,13\}$

- Path to data $n3 \rightarrow n10$ (data found) - 1 hop $\approx O(\log n)$

January 28, 2021 TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma L8.11

The Instructor is unavailable for next Tuesday's class due to a community service commitment with the National Science Foundation. What is your preference for next Tuesday's class?

Cancel class - NO CLASS ON TUESDAY Feb 2nd

Reschedule class FOR MONDAY Feb 1st at 6PM for 2 hours

Reduced class ON MONDAY Feb 1st at 6PM for 1 hour

Reduced class ON FRIDAY Feb 5th at 11:30AM for 1 hour

No Preference

Start the presentation to see live content. For screen share software, share the entire screen. Get help at poller.com/app

OBJECTIVES – 1/28

- Questions from 1/26
- **Assignment 0: Cloud Computing Infrastructure Tutorial**
 - New `testFibService.sh` script
- Assignment 1: Key/Value Store
- Chapter 3: Processes
 - Chapter 3.1: Threads
 - Context Switches
 - Threading Models
 - Multithreaded clients/servers
 - Chapter 3.2: Virtualization
 - Chapter 3.3: Clients
 - Chapter 3.4: Servers

January 28, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L8.13
------------------	---	-------

ASSIGNMENT 0

- **Preparing for Assignment 0:**
 - Establish AWS Account
 - Standard account – **** request cloud credits from Instructor ****
 - Specify "AWS CREDIT REQUEST" as subject of email
 - Include email address of AWS account
 - **AWS Educate Starter account** – some account limitations
 - https://awseducate-starter-account-services.s3.amazonaws.com/AWS_Educate_Starter_Account_Services_Supported.pdf
 - Establish local Linux/Ubuntu environment
 - Task 1 – AWS account setup
 - Task 2 – Working w/ Docker, creating Dockerfile for Apache Tomcat
 - Task 3 – Creating a Dockerfile for haproxy
 - Task 4 – Working with Docker-Machine
 - Task 5 – For Submission: Testing Alternate Server Configurations

January 28, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L8.14
------------------	---	-------

TESTING CONNECTIVITY TO SERVER

- `testFibPar.sh` script is a parallel test script
- Orchestrates multiple threads on client to invoke server multiple times in parallel
- To simplify coordinate of parallel service calls in BASH, `testFibPar.sh` script ignores errors !!!
- To help test client-to-server connectivity, have created a new `testFibService.sh` script
- TEST 1: Network layer
 - Ping (ICMP)
- TEST 2: Transport layer
 - TCP: telnet (TCP Port 8080) – security group (firewall) test
- TEST 3: Application layer
 - HTTP REST – web service test

January 28, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L8.15
------------------	---	-------

OBJECTIVES – 1/28

- Questions from 1/26
- Assignment 0: Cloud Computing Infrastructure Tutorial
 - New `testFibService.sh` script
- **Assignment 1: Key/Value Store**
- Chapter 3: Processes
 - Chapter 3.1: Threads
 - Context Switches
 - Threading Models
 - Multithreaded clients/servers
 - Chapter 3.2: Virtualization
 - Chapter 3.3: Clients
 - Chapter 3.4: Servers

January 28, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L8.16
------------------	---	-------

ASSIGNMENT 1

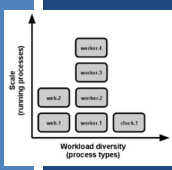
- TCP/UDP/RMI Key Value Store
- Implement a "GenericNode" project which assumes the role of a client or server for a Key/Value Store
- Recommended in Java (11 or 8)
- Client node program interacts with server node to put, get, delete, or list items in a key/value store

January 28, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L8.17
------------------	---	-------

OBJECTIVES – 1/28

- Questions from 1/26
- Assignment 0: Cloud Computing Infrastructure Tutorial
 - New `testFibService.sh` script
- Assignment 1: Key/Value Store
 - **Chapter 3: Processes**
 - Chapter 3.1: Threads
 - Context Switches
 - Threading Models
 - Multithreaded clients/servers
 - Chapter 3.2: Virtualization
 - Chapter 3.3: Clients
 - Chapter 3.4: Servers

January 28, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L8.18
------------------	---	-------



CH. 3: PROCESSES

CH. 3.1: THREADS

L8.19

CHAPTER 3

- Chapter 3 titled “processes”
- Covers variety of distributed system implementation details
- “Grab bag” of topics

- Processes/threads
- Virtualization
- Clients
- Servers
- Code migration


January 28, 2021 TCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma L8.20

OBJECTIVES – 1/28

- Questions from 1/26
- Assignment 0: Cloud Computing Infrastructure Tutorial
 - New testFibService.sh script
- Assignment 1: Key/Value Store
- Chapter 3: Processes
 - **Chapter 3.1: Threads**
 - Context Switches
 - Threading Models
 - Multithreaded clients/servers
 - Chapter 3.2: Virtualization
 - Chapter 3.3: Clients
 - Chapter 3.4: Servers

January 28, 2021 TCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma L8.21


CH. 3.1 - THREADS



- For implementing a server (or client) threads offer many advantages vs. heavy weight processes
- **What is the difference between a process and a thread?**
 - (review?) from Operating Systems
- **Key difference: what do threads share amongst each other that processes do not... ?**
- **What are the segments of a program stored in memory?**
 - Heap segment (dynamic shared memory)
 - Code segment
 - Stack segment
 - Data segment (global variables)

January 28, 2021 TCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma L8.22


THREADS - 2



- **Do several processes on an operating system share...**
 - **Heap segment?**
 - **Stack segment?**
 - **Code segment?**
- **Can we run multiple copies of the same code?**
 - These may be managed as shared pages (across processes) in memory
- Processes are isolated from each other by the OS
 - Each has a separate heap, stack, code segment

January 28, 2021 TCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma L8.23

THREADS - 3



- Threads avoid the overhead of process creation
- No new heap or code segments required
- **What is a context switch?**
 - Context switching among threads is considered to be more efficient than context switching processes
 - Less elements to swap-in and swap-out
- Unikernel: specialized single process OS for the cloud
- Example: Osv, Clive, MirageOS (see: <http://unikernel.org/projects/>)
- Single process operating system with many threads
- Developed for the cloud to run only one application at a time

January 28, 2021 TCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma L8.24

OSV: ONE PROCESS, MANY THREADS

January 28, 2021 TCCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma 18.25

THREADS - 4

- Important implications with threads:
 - (1) multi-threading should lead to performance gains
 - (2) thread programming requires additional effort when threads share memory
 - Known as thread **synchronization**, or enabling **concurrency**
- Access to **critical sections** of code which modify shared variables must be **mutually exclusive**
 - No more than one thread can execute at any given time
 - Critical sections must run **atomically** on the CPU

January 28, 2021 TCCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma 18.26

BLOCKING THREADS

- Example: spreadsheet with formula to compute sum of column
- User modifies values in column
- Multiple threads:
 1. Supports interaction (UI) activity with user
 2. Updates spreadsheet calculations in parallel
 3. Continually backs up spreadsheet changes to disk
- Single core CPU
 - Tasks appear as if they are performed simultaneously
- Multi core CPU
 - Tasks **execute** simultaneously

January 28, 2021 TCCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma 18.27

INTERPROCESS COMMUNICATION

- IPC – mechanism using pipes, message queues, and shared memory segments
- IPC mechanisms incur context switching
 - Process I/O must execute in kernel mode
- **How many context switches are required for process A to send a message to process B using IPC?**
- **#1 C/S:**
Proc A → kernel thread
- **#2 C/S:**
Kernel thread → Proc B

January 28, 2021 TCCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma 18.28

OBJECTIVES - 1/28

- Questions from 1/26
- Assignment 0: Cloud Computing Infrastructure Tutorial
 - New testFibService.sh script
- Assignment 1: Key/Value Store
- Chapter 3: Processes
 - Chapter 3.1: Threads
 - **Context Switches**
 - Threading Models
 - Multithreaded clients/servers
 - Chapter 3.2: Virtualization
 - Chapter 3.3: Clients
 - Chapter 3.4: Servers

January 28, 2021 TCCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma 18.29

CONTEXT SWITCHING

- **Direct overhead**
 - Time spent not executing program code (user or kernel)
 - Time spent executing interrupt routines to swap memory segments of different processes (or threads) in the CPU
 - Stack, code, heap, registers, code pointers, stack pointers
 - Memory page cache invalidation
- **Indirect overhead**
 - Overhead not directly attributed to the physical actions of the context switch
 - Captures performance degradation related to the side effects of context switching (e.g. rewriting of memory caches, etc.)
 - **Primarily cache perturbation**

January 28, 2021 TCCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma 18.30

CONTEXT SWITCH – CACHE PERTURBATION

- Refers to cache reorganization that occurs as a result of a context switch
- Cache is not clear, but elements from cache are removed as a result of another program running in the CPU
- 80% performance overhead from context switching results from this **"cache perturbation"**

The diagram shows a 4x4 grid representing a cache. The top row is labeled 'MRU' and the bottom row is labeled 'LRU'. The grid contains the following elements: Row 1: A, empty, D, empty; Row 2: B, A, empty, empty; Row 3: C, B, A, empty; Row 4: D, C, B, empty.

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.31

OBJECTIVES – 1/28

- Questions from 1/26**
- Assignment 0: Cloud Computing Infrastructure Tutorial
 - New testFibService.sh script
- Assignment 1: Key/Value Store
- Chapter 3: Processes
 - Chapter 3.1: Threads
 - Context Switches
 - Threading Models**
 - Multithreaded clients/servers
 - Chapter 3.2: Virtualization
 - Chapter 3.3: Clients
 - Chapter 3.4: Servers

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.32

THREADING MODELS

- Many-to-one threading:** multiple user-level threads per process
- Thread operations (create, delete, locks) run in user mode
- Multithreaded process mapped to single schedulable entity
- Only one thread per process runs at any given time
- Key take-away: thread management handled by user processes
- What are some advantages of many-to-one threading?**
- What are some disadvantages?**

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.33

THREADING MODELS - 2

- One-to-one threading:** use of separate kernel threads for each user process - also called **kernel-level threads**
- The kernel API calls (e.g. I/O, locking) are farmed out to an existing kernel level thread
- Thread operations (create, delete, locks) run in kernel mode
- Threads scheduled individually by the OS
- System calls required, context switches as expensive as process context switching
- Idea is to have preinitialized kernel threads for user processes
- Linux uses this model...
- What are some advantages of one-to-one threading?**
- What are some disadvantages?**

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.34

APPLICATION EXAMPLES

- Google chrome: processes
- Apache tomcat webservice: threads
- Multiprocess programming avoids synchronization of concurrent access to shared data, by providing coordination and data sharing via interprocess communication (IPC)
- Each process maintains its own private memory
- While this approach avoids synchronizing concurrent access to shared memory, what is the tradeoff(s)??**
 - Replication instead of synchronization – must synchronize multiple copies of the data
- Do distributed objects share memory?**

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.35

WE WILL RETURN AT 2:46PM

OBJECTIVES – 1/28

- **Questions from 1/26**
- **Assignment 0: Cloud Computing Infrastructure Tutorial**
 - New testFibService.sh script
- **Assignment 1: Key/Value Store**
- **Chapter 3: Processes**
 - Chapter 3.1: Threads
 - Context Switches
 - Threading Models
 - **Multithreaded clients/servers**
 - Chapter 3.2: Virtualization
 - Chapter 3.3: Clients
 - Chapter 3.4: Servers

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.37

MULTITHREADED CLIENTS

- **Web browser**
- Uses threads to load and render portions of a web page to the user in parallel
- A client could have dozens of concurrent connections all loading in parallel
- **testFibPar.sh**
- Assignment 0 client script (GNU parallel)
- **Important benefits:**
 - Several connections can be opened simultaneously
 - Client: dozens of concurrent connections to the webserver all loading data in parallel

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.38

MULTIPLE THREADS

- In Linux, threads also receive a process ID (PID)
- To display threads of a process in Linux:
- Identify parent process explicitly:
- `top -H -p <pid>`
- `htop -p <pid>`
- `ps -iT <pid>`
- Virtualbox process ~ 44 threads
- No mapping to guest # of processes/threads

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.39

PROCESS METRICS

CPU

- `cpuUsr`: CPU time in user mode
- `cpuKrn`: CPU time in kernel mode
- `cpuIdle`: CPU idle time
- `cpuIoWait`: CPU time waiting for I/O
- `cpuIntSrcv`: CPU time serving interrupts
- `cpuSftIntSrcv`: CPU time serving soft interrupts
- `cpuNice`: CPU time executing prioritized processes
- `cpuSteal`: CPU ticks lost to virtualized guests
- `contextsw`: # of context switches
- `loadavg`: (avg # proc / 60 secs)

Disk

- `dsk`: disk sector reads
- `dsreads`: disk sector reads completed
- `dskm`: merged adjacent disk reads
- `readtime`: time spent reading from disk
- `dskw`: disk sector writes
- `dskwrites`: disk sector writes completed
- `dskwm`: merged adjacent disk writes
- `writetime`: time spent writing to disk

Network

- `nbs`: network bytes sent
- `nbr`: network bytes received

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.40

LOAD AVERAGE

- Reported by: `top`, `htop`, `w`, `uptime`, and `/proc/loadavg`
- Updated every 5 seconds
- Average number of processes using or waiting for the CPU
- Three numbers show exponentially decaying usage for 1 minute, 5 minutes, and 15 minutes
- One minute average: exponentially decaying average
- Load average = $1 \cdot (\text{avg last minute load}) - 1/e \cdot (\text{avg load since boot})$
- 1.0 = 1-CPU core fully loaded
- 2.0 = 2-CPU cores
- 3.0 = 3-CPU cores . . .

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.41

THREAD-LEVEL PARALLELISM

- Metric – measures degree of parallelism realized by running system, by calculating average utilization:

$$TLP = \frac{\sum_{i=1}^N i \cdot c_i}{1 - c_0}$$

- c_i – fraction of time that exactly i threads are executed
- N – maximum threads that can execute at any one time
- Web browsers found to have TLP from 1.5 to 2.5
- Clients for web browsing can utilize from 2 to 3 CPU cores
- Any more cores are redundant, and potentially wasteful
- **Measure TLP to understand how many CPUs to provision**

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.42

MULTITHREADED SERVERS

- Multiple threads essential for servers in distributed systems
- Even on single-core machines greatly improves performance
- Take advantage of idle/blocking time
- Two designs:
 - Generate new thread for every request
 - Thread pool – pre-initialize set of threads to service requests

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.43

SINGLE THREAD & FSM SERVERS

- Single thread server
 - A single thread handles all client requests
 - BLOCKS for I/O
 - All waiting requests are queued until thread is available
- Finite state machine
 - Server has a single thread of execution
 - I/O performing asynchronously (non-BLOCKing)
 - Server handles other requests while waiting for I/O
 - Interrupt fired with I/O completes
 - Single thread “jumps” back into context to finish request

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.44

SERVER DESIGN ALTERNATIVES

- A blocking system call implies that a thread servicing a request synchronously performs I/O
- The thread BLOCKS to wait on disk/network I/O before proceeding with request processing
- Consider the implications of these designs for responsiveness, availability, scalability. . .

Model	Characteristics
Multithreading	Parallelism, blocking I/O
Single-thread	No parallelism, blocking I/O
Finite-state machine	Parallelism, non-blocking I/O

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.45

OBJECTIVES – 1/28

- Questions from 1/26
- Assignment 0: Cloud Computing Infrastructure Tutorial
 - New testFibService.sh script
- Assignment 1: Key/Value Store
- Chapter 3: Processes
 - Chapter 3.1: Threads
 - Context Switches
 - Threading Models
 - Multithreaded clients/servers
 - Chapter 3.2: Virtualization
 - Chapter 3.3: Clients
 - Chapter 3.4: Servers

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.46

CH. 3.2: VIRTUALIZATION

L8.47

VIRTUALIZATION

- Initially introduced in the 1970s on IBM mainframe computers
- Legacy operating systems run in mainframe-based VMs
- Legacy software could be sustained by virtualizing legacy OSes
- 1970s virtualization went away as desktop/rack-based hardware became inexpensive
- Virtualization reappears in 2000s to leverage multi-core, multi-CPU processor systems
- VM-Ware virtual machines enable companies to host many virtual servers with mixed OSes on private clusters
- Cloud computing: Amazon offers VMs as-a-service (IaaS)

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.48

TYPES OF VIRTUALIZATION

- Levels of Instructions:**
 - Library functions
 - Application
 - Library
 - Operating system
 - Hardware
- Hardware:** CPU
 - Privileged instructions
KERNEL MODE
 - General instructions
USER MODE
- Operating system:** system calls
- Library:** programming APIs: e.g. C/C++, C#, Java libraries
- Application:**
- Goal of virtualization:** mimic these interface to provide a virtual computer

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.49

TYPES OF VIRTUALIZATION - 2

- Process virtual machine**
 - Interpret instructions: (interpreters) (JavaVM) byte code → HW instructions
 - Emulate instructions: (emulators) (Wine) windows code → Linux code
- Native virtual machine monitor (VMM)**
 - Hypervisor (XEN): small OS with its own kernel
 - Provides an interface for multiple guest OSes
 - Facilitates sharing/scheduling of CPU, device I/O among many guests
 - Guest OSes require special kernel to interface w/ VMM
 - Supports **Paravirtualization** for performance boost to run code directly on the CPU
 - Type 1 hypervisor

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.50

TYPES OF VIRTUALIZATION - 3

- Hosted virtual machine monitor (VMM)**
 - Runs atop of hosted operating system
 - Uses host OS facilities for CPU scheduling, I/O
 - Full virtualization
 - Type 2 hypervisor
 - Virtualbox**
- Textbook: note 3.5-good explanation of full vs. paravirtualization**
- GOAL:** run all user mode instructions directly on the CPU
- x86 instruction set has ~17 privileged user mode instructions
- Full virtualization:** scan the EXE, insert code around privileged instructions to divert control to the VMM
- Paravirtualization:** special OS kernel eliminates side effects of privileged instructions

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.51

EVOLUTION OF AWS VIRTUALIZATION

From <http://www.brendangregg.com/blog/2017-11-29/aws-ec2-virtualization-2017.html>

AWS EC2 Virtualization Types

VS: Virtualization in software

P: Paravirtual

VH: Virtualization in Hardware

H: Hardware

#	Tech	Type	With	Importance → Least							
				CPU	Memory	Network I/O	Local Storage I/O	Remote Storage I/O	Interrupts, Timers	Multirooted Boot	
1	VM	Fully Emulated		VS	VS	VS	VS	VS	VS	VS	VS
2	VM	Xen PV 3.0	PV drivers	P	P	P	P	P	P	VS	VS
3	VM	Xen HVM 3.0	PVHVM drivers	VH	P	P	P	P	P	VS	VS
4	VM	Xen HVM 4.0.1	PVHVM drivers	VH	P	P	P	P	P	P	VS
5	VM	Xen AWS 2013	PVHVM + SR-IOV(net)	VH	VH	P	P	P	P	P	VS
6	VM	Xen AWS 2017	PVHVM + SR-IOV(net, stor)	VH	VH	VH	P	P	P	P	VS
7	VM	AWS Nitro 2017		VH	VH	VH	VH	VH	VH	VH	VS
8	HW	AWS Bare Metal 2017		H	H	H	H	H	H	H	H
		Bare Metal		H	H	H	H	H	H	H	H

* VM: Virtual Machine, HW: Hardware.
 VS: Virt. in software, VH: Virt. in hardware, P: Paravirt. Not all combinations shown.
 SR-IOV(net): sgbeena driver, SR-IOV(storage): nvme driver.

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.52

AWS VIRTUALIZATION - 2

- Full Virtualization - Fully Emulated**
 - Never used on EC2, before CPU extensions for virtualization
 - Can boot any unmodified OS
 - Support via slow emulation, performance 2x-10x slower
- Paravirtualization: Xen PV 3.0**
 - Software: Interrupts, timers
 - Paravirtual: CPU, Network I/O, Local+Network Storage
 - Requires special OS kernels, interfaces with hypervisor for I/O
 - Performance 1.1x - 1.5x slower than "bare metal"
 - Instance store instances: 1st & 2nd generation- m1.large, m2.xlarge
- Xen HVM 3.0**
 - Hardware virtualization: **CPU, memory (CPU VT-x required)**
 - Paravirtual: network, storage
 - Software: interrupts, timers
 - EBS backed instances
 - m1, c1 instances

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.53

AWS VIRTUALIZATION - 3

- XEN HVM 4.0.1**
 - Hardware virtualization: CPU, memory (**CPU VT-x required**)
 - Paravirtual: network, storage, **Interrupts, timers**
- XEN AWS 2013 (diverges from open source XEN)**
 - Provides hardware virtualization for CPU, memory, **network**
 - Paravirtual: storage, **Interrupts, timers**
 - Called Single root I/O Virtualization (SR-IOV)
 - Allows sharing single physical PCI Express device (i.e. network adapter) with multiple VMs
 - Improves VM network performance
 - 3rd & 4th generation instances (c3 family)
 - Network speeds up to 10 Gbps and 25 Gbps
- XEN AWS 2017**
 - Provides hardware virtualization for CPU, memory, network, **local disk**
 - Paravirtual: remote storage, **Interrupts, timers**
 - Introduces hardware virtualization for EBS volumes (c4 instances)
 - Instance storage hardware virtualization (x1.32xlarge, i3 family)

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.54

AWS VIRTUALIZATION - 4

- **AWS Nitro 2017**
 - Provides hardware virtualization for CPU, memory, network, **local disk, remote disk, Interrupts, timers**
 - All aspects of virtualization enhanced with HW-level support
 - November 2017
 - Goal: provide performance indistinguishable from "bare metal"
 - 5th generation instances - c5 instances (also c5d, c5n)
 - Based on KVM hypervisor
 - Overhead around ~1%


January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.55

OBJECTIVES - 1/28

- **Questions from 1/26**
- **Assignment 0: Cloud Computing Infrastructure Tutorial**
 - New testFibService.sh script
- **Assignment 1: Key/Value Store**
- **Chapter 3: Processes**
 - Chapter 3.1: Threads
 - Context Switches
 - Threading Models
 - Multithreaded clients/servers
 - Chapter 3.2: Virtualization
 - **Chapter 3.3: Clients**
 - Chapter 3.4: Servers

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.56

CH. 3.3: CLIENTS



L8.57

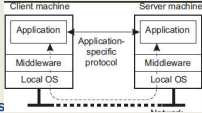
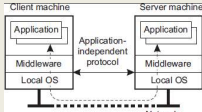
TYPES OF CLIENTS

- **Thick clients**
 - Web browsers
 - Client-side scripting
 - Mobile apps
 - Multi-tier MVC apps
- **Thin clients**
 - Remote desktops/GUIs (very thin)

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.58

CLIENTS

- **Application specific protocol**
 - Thick clients
 - Clients maintain local data
 - Middleware (APIs)
 - Clients synchronize data with remote nodes
 - Example: shared calendar application
- **Application independent**
 - Thin clients
 - Client acts as a remote terminal
 - Provides interface to user (GUI / UI)
 - Server houses entire application stack

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.59

X WINDOWS

- Layered architecture to transport UI over network
- Remote desktop functionality for Linux/Unix systems
- X kernel acts as a server
 - Provides the **X protocol**: application level protocol
 - Xlib instances (client applications) exchange data and events with X kernels (servers)
 - Clients and servers on single machine → Linux GUI
 - Client and server communication transported over the network → remote Linux GUI

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.60

X WINDOWS - 2

- Window manager:**
 - Application running atop of X-windows which provides flair
 - Many variants
 - Without X windows is quite bland

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.61

- Layered architecture**
- X-kernel: low level interface/APIs for controlling screen, capturing keyboard and mouse events (X window Server)**
- Provided on Linux as Xlib**
- Provides network enabled GUI**
- Layering allows for use for custom window managers**

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.62

EXAMPLE: VNC SERVER

- How to Install VNC server on Ubuntu EC2 Instance VM:**
- `sudo apt-get update`
- `# ubuntu 16.04`
- `sudo apt-get install ubuntu-desktop`
- `sudo apt-get install gnome-panel gnome-settings-daemon metacity nautilus gnome-terminal`
- `# on ubuntu 18.04`
- `sudo apt install xfce4 xfce4-goodies`
- `sudo apt-get install tightvncserver # both`
- Start VNC server to create initial config file**
- `vncserver :1`

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.63

EXAMPLE: VNC SERVER - UBUNTU 16.04

- On the VM: edit config file: nano ~/.vnc/xstartup**
- Replace contents as below (Ubuntu 16.04):**

```
#!/bin/sh
export XKL_XMODMAP_DISABLE=1
unset SESSION_MANAGER
unset DBUS_SESSION_BUS_ADDRESS

[ -x /etc/vnc/xstartup ] && exec /etc/vnc/xstartup
[ -x $HOME/.Xresources ] && xrdb $HOME/.Xresources
xsetroot -solid grey

vncconfig -iconic &
gnome-panel &
gnome-settings-daemon &
metacity &
nautilus &
gnome-terminal &
```

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.64

EXAMPLE: VNC SERVER - UBUNTU 18.04

- On the VM:**
- Edit config file: nano ~/.vnc/xstartup**
- Replace contents as below (Ubuntu 18.04):**

```
#!/bin/bash
xrdb $HOME/.Xresources
startxfce4 &
```

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.65

EXAMPLE: VNC SERVER - 3

- On the VM: reload config by restarting server**
- `vncserver -kill :1`
- `vncserver :1`
- Open port 22 & 5901 in EC2 security group:**

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.66

EXAMPLE: VNC CLIENT

- On the client (e.g. laptop):
- Create SSH connection to securely forward port 5901 on the EC2 instance to your localhost port 5901
- This way your VNC client doesn't need an SSH key

```
ssh -i <ssh-keyfile> -L 5901:127.0.0.1:5901 -N -f -l <username> <EC2-instance ip_address>
```

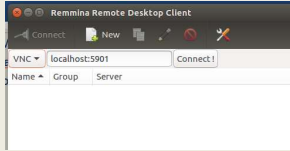
- For example:

```
ssh -i mykey.pem -L 5901:127.0.0.1:5901 -N -f -l ubuntu 52.111.202.44
```

January 28, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L8.67
------------------	---	-------

EXAMPLE: VNC CLIENT - 2

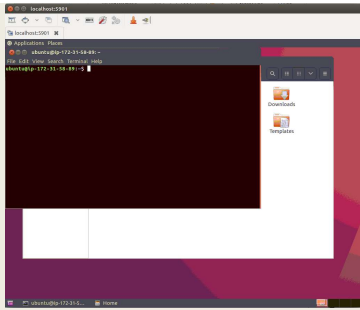
- On the client (e.g. laptop):
- Use a VNC Client to connect
- Remmina is provided by default on Ubuntu 16.04
- Can "google" for many others
- Remmina login:
- Chose "VNC" protocol
- Log into "localhost:5901"



January 28, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L8.68
------------------	---	-------

REMOTE COMPUTER IN THE CLOUD

- EC2 instance with a GUI. . . !!!



January 28, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L8.69
------------------	---	-------

THIN CLIENTS

- Thin clients
- X windows protocol
- A variety of other remote desktop protocols exist:

Remote desktop protocols include the following:

- Apple Remote Desktop Protocol (ARD) – Original protocol for Apple Remote Desktop on macOS machines.
- Appliance Link Protocol (ALP) – a Sun Microsystems-specific protocol featuring audio (play and record), remote printing, remote USB, accelerated video
- HP Remote Graphics Software (RGS) – a proprietary protocol designed by Hewlett-Packard specifically for high end workstation remoting and collaboration.
- Independent Computing Architecture (ICA) – a proprietary protocol designed by Citrix Systems
- NX technology (NoMachine NX) – Cross platform protocol featuring audio, video, remote printing, remote USB, H264-enabled.
- PC-over-IP (PCoIP) – a proprietary protocol used by VMware (licensed from Teradici)^[2]
- Remote Desktop Protocol (RDP) – a Windows-specific protocol featuring audio and remote printing
- Remote Frame Buffer Protocol (RFB) – A framebuffer level cross-platform protocol that VNC is based on.
- SPICE (Simple Protocol for Independent Computing Environments) – remote-display system built for virtual environments by Qumranet, now Red Hat
- Splashtop – a high performance remote desktop protocol developed by Splashtop, fully optimized for hardware (H.264) including Intel / AMD chipsets, NVIDIA media codecs. Splashtop can deliver high frame rates with low latency, and also low power consumption.
- X Window System (X11) – a well-established cross-platform protocol mainly used for displaying local applications; X11 is network transparent

January 28, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L8.70
------------------	---	-------

THIN CLIENTS - 2

- Applications should separate application logic from UI
- When application logic and UI interaction are tightly coupled many requests get sent to X kernel
- Client must wait for response
- Synchronous behavior and app-to-UI coupling adversely affects performance of WAN / Internet
- Protocol optimizations:** reduce bandwidth by shrinking size of X protocol messages
- Send only differences between messages with same identifier
- Optimizations enable connections with 9600 kbps

January 28, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L8.71
------------------	---	-------

THIN CLIENTS - 3

- Virtual network computing (VNC)
- Send display over the network at the pixel level (instead of X lib events)
- Reduce pixel encodings to save bandwidth – fewer colors
- Pixel-based approaches lose application semantics
- Can transport any GUI this way
- THINC**- hybrid approach
- Send video device driver commands over network
- More powerful than pixel based operations
- Less powerful compared to protocols such as X

January 28, 2021	TCCS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma	L8.72
------------------	---	-------

TRADEOFFS: ABSTRACTION OF REMOTE DISPLAY PROTOCOLS

- Tradeoff space: abstraction level of remote display protocols

Pixel-level VNC

Graphics IIB X.1.1

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.73

TRADEOFFS: ABSTRACTION OF REMOTE DISPLAY PROTOCOLS

- Tradeoff space: abstraction level of remote display protocols

Pixel-level VNC

Graphics IIB X.1.1

- Generic – no app context
- Graphics data
- Higher network bandwidth
- Fewer colors
- Utilize graphics compression
- More network traffic

- Application context is available
- UI data/operations
- Lower network bandwidth
- More colors

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.74

CLIENT ROLES IN PROVIDING DISTRIBUTION TRANSPARENCY

- Clients help enable distribution transparency of servers
- Replication transparency
 - Client aggregates responses from multiple servers
 - Only the client knows of replicas

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.75

CLIENT ROLES IN PROVIDING DISTRIBUTION TRANSPARENCY - 2

- Location/relocation/migration transparency
 - Harness convenient naming system to allow client to infer new locations
 - Server inform client of moves / Client reconnects to new endpoint
 - Client hides network address of server, and reconnects as needed
 - May involve temporary loss in performance
- Replication transparency
 - Client aggregates responses from multiple servers
- Failure transparency
 - Client retries, or maps to another server, or uses cached data
- Concurrency transparency
 - Transaction servers abstract coordination of multithreading

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.76

OBJECTIVES - 1/28

- Questions from 1/26
- Assignment 0: Cloud Computing Infrastructure Tutorial
 - New testFibService.sh script
- Assignment 1: Key/Value Store
- Chapter 3: Processes
 - Chapter 3.1: Threads
 - Context Switches
 - Threading Models
 - Multithreaded clients/servers
 - Chapter 3.2: Virtualization
 - Chapter 3.3: Clients
 - Chapter 3.4: Servers**

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.77

CH. 3.4: SERVERS

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.78

SERVERS

- Cloud & Distributed Systems – rely on **Linux**
- <http://www.zdnet.com/article/it-runs-on-the-cloud-and-the-cloud-runs-on-linux-any-questions/>
- IT is moving to the cloud. And, what powers the cloud?
 - **Linux**
- Uptime Institute survey - 1,000 IT executives (2016)
 - 50% of IT executives – plan to migrate majority of IT workloads to off-premise to cloud or colocation sites
 - 23% expect the shift in 2017, 70% by 2020...
- Docker on Windows / Mac OS X
 - Based on **Linux**
 - Mac: Hyperkit Linux VM
 - Windows: Hyper-V Linux VM

January 28, 2021
TCCSS58: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.79

SERVERS - 2

- Servers implement a specific service for a collection of clients
- Servers wait for incoming requests, and respond accordingly
- **Server types**
- **Iterative:** immediately handle client requests
- **Concurrent:** Pass client request to separate thread
- Multithreaded servers are concurrent servers
 - E.g. Apache Tomcat
- **Alternative:** fork a new process for each incoming request
- **Hybrid:** mix the use of multiple processes with thread pools

January 28, 2021
TCCSS58: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.80

END POINTS

- Clients connect to servers via:
IP Address and Port Number
- How do ports get assigned?
 - Many protocols support “default” port numbers
 - Client must find IP address(es) of servers
 - A single server often hosts multiple end points (servers/services)
 - When designing new TCP client/servers must be careful not to repurpose ports already commonly used by others

January 28, 2021
TCCSS58: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.81

COMMON PORTS

packetlife.net

TCP/UDP Port Numbers			
7	Echo	554	RTSP
19	Chargen	546-547	DHCPv6
20-21	FTP	560	monitor
22	SFTP	563	SFTP over SSL
23	Telnet	587	SMTTP
25	SMTTP	591	FileMaker
42	WINS Replication	593	Microsoft DCOM
43	WHOIS	631	Internet Printing
49	TACACS	636	LDAP over SSL
53	DNS	639	MSDP (PIM)
67-68	DHCP/BOOTP	646	LDP (MPLS)
69	TFTP	691	MS Exchange
70	Gopher	860	ISCSI
79	Finger	873	rsync
80	HTTP	902	VMware Server
88	Kerberos	988-990	FTP over SSL
102	MS Exchange	993	IMAP over SSL
110	POP3	995	POP3 over SSL
113	Ident	1025	Microsoft RPC
119	NTP (Usenet)	1026-1029	Windows Messenger
123	NTP	1080	SOCKS Proxy
135	Microsoft RPC	1080	MyDoom
137-139	NetBIOS	1194	OpenVPN
143	IMAP4	1214	Kazaa
161-162	SNMP	1241	Nessus
177	XDMCP	1311	Dell OpenManage
179	BGP	1317	WACFT
2745	Segin.H	2967	Symantec AV
3050	Interbase DB	3074	MSRPC
3124	HTTP Proxy	3127	MSRPC
3128	HTTP Proxy	3222	GLBP
3260	ISCSI Target	3306	MySQL
3389	Terminal Server	3689	iTunes
3689	iTunes	3690	Subversion
3724	World of Warcraft	3784-3785	Venrilo
4333	mSQL	4444	Skype
4664	Google Desktop	4672	Utorrent
4899	Radmin	5000	UPnP
5001	Slingbox	5001	iperf
5004-5005	RTSP	5050	Yahoo! Messenger
5060	SIP	5060	SIP
5130	XMPP	5190	XMPP
6891-6901	Windows Live	6970	QuickTime
7212	GhostSurf	7648-7649	CSRSERVE
8000	Internet Radio	8080	HTTP Proxy
8086-8087	Kaspersky AV	8118	Privoxy
8200	VMware Server	8500	Adobe ColdFusion
8767	TeamSpeak	8767	TeamSpeak
8866	Bagle.B	9100	HP JetDirect
9101-9103	Bacula	9119	RMK
9800	WebDAV	9898	Skype
9988	ibmi/Saybot	9999	Urchin
10000	Webmin	10000	BackupExec
10113-10116	NetIQ	10205-10206	Second Life
11371	OpenPGP	12345	NetBus
13720-13721	NetBackup	14567	RealPlayer

January 28, 2021
TCCSS58: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.84

TYPES OF SERVERS

- **Daemon server**
 - Example: NTP server
- **Superserver**
- **Stateless server**
 - Example: Apache server
- **Stateful server**
- **Object servers**
- **EJB servers**

January 28, 2021
TCCSS58: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.83

NTP EXAMPLE

- **Daemon servers**
 - Run locally on Linux
 - Track current server end points (outside servers)
 - Example: network time protocol (ntp) daemon
 - Listen locally on specific port (ntp is 123)
 - Daemons routes local client traffic to the configured endpoint servers
 - University of Washington: time.u.washington.edu
 - Example “ntpd -p”
 - Queries local ntp daemon, routes traffic to configured server(s)

January 28, 2021
TCCSS58: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.84

SUPERSERVER

- Linux inetd / xinetd
 - Single superserver
 - Extended internet service daemon
 - Not installed by default on Ubuntu
 - Intended for use on server machines
 - Used to configure box as a server for multiple internet services
 - E.g. ftp, pop, telnet
 - inetd daemon responds to multiple endpoints for multiple services
 - Requests fork a process to run required executable program
- Check what ports you're listening on:
 - `sudo netstat -tap | grep LISTEN`

January 28, 2021

TCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma

L8.85

INTERRUPTING A SERVER

- Server design issue:
 - Active client/server communication is taking place over a port
 - How can the server / data transfer protocol support interruption?
- Consider transferring a 1 GB image, how do you pass a unrelated message in this stream?
 1. **Out-of-band** data: special messages sent in-stream to support interrupting the server (*TCP urgent data*)
 2. Use a separate connection (different port) for admin control info
- Example: sftp secure file transfer protocol
 - Once a file transfer is started, can't be stopped easily
 - Must kill the client and/or server

January 28, 2021

TCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma

L8.86

STATELESS SERVERS

- Data about state of clients is not stored
- Example: web application servers are typically stateless
 - Also function-as-a-service (FaaS) platforms
- Many servers maintain information on clients (e.g. log files)
- Loss of stateless data doesn't disrupt server availability
 - Losing log files typically has minimal consequences
- **Soft state**: server maintains state on the client for a limited time (*to support sessions*)
- Soft state information expires and is deleted

January 28, 2021

TCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma

L8.87

STATEFUL SERVERS

- Maintain persistent information about clients
- Information must be explicitly deleted by the server
- Example:
 - File server - allows clients to keep local file copies for RW
 - Server tracks client file permissions and most recent versions
 - Table of (client, file) entries
- If server crashes data must be recovered
- Entire state before a crash must be restored
- Fault tolerance - Ch. 8

January 28, 2021

TCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma

L8.88

STATEFUL SERVERS - 2

- Session state
 - Tracks series of operations by a single user
 - Maintained temporarily, not indefinitely
 - Often retained for multi-tier client server applications
 - Minimal consequence if session state is lost
 - Clients must start over, reinitialize sessions
- Permanent state
 - Customer information, software keys
- Client-side cookies
 - When servers don't maintain client state, clients can store state locally in "cookies"
 - Cookies are not executable, simply client-side data

January 28, 2021

TCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma

L8.89

OBJECT SERVERS

- **OBJECTIVE**: Host objects and enable remote client access
- Do not provide a specific service
 - Do nothing if there are no objects to host
- Support adding/removing hosted objects
- Provide a home where objects live
- Objects, *themselves*, provide "services"
- Object parts
 - State data
 - Code (methods, etc.)
- **Transient object(s)**
 - Objects with limited lifetime (< server)
 - Created at first invocation, destroyed when no longer used (i.e. no clients remain "bound").
 - Disadvantage: initialization may be expensive
 - Alternative: preinitialize and retain objects on server start-up

January 28, 2021

TCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma

L8.90

OBJECT SERVERS - 2

- **Should object servers isolate memory for object instances?**
 - Share neither code nor data
 - May be necessary if objects couple data and implementation
- Object server threading designs:
 - Single thread of control for object server
 - One thread for each object
 - Servers use separate thread for client requests
- Threads created on demand **vs.** Server maintains pool of threads
- **What are the tradeoffs for creating server threads on demand vs. using a thread pool?**

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.91

EJB – ENTERPRISE JAVA BEANS

- EJB- specialized Java object hosted by a EJB web container
- 4 types: stateless, stateful, entity, and message-driven beans
- Provides "middleware" standard (framework) for implementing back-ends of enterprise applications
- EJB web application containers integrate support for:
 - Transaction processing
 - Persistence
 - Concurrency
 - Event-driven programming
 - Asynchronous method invocation
 - Job scheduling
 - Naming and discovery services (JNDI)
 - Interprocess communication
 - Security
- Software component deployment to an application server

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.92

APACHE WEB SERVER

- Highly configurable, extensible, platform independent
- Supports TCP HTTP protocol communication
- Uses hooks – placeholders for group of functions
- Requests processed in phases by hooks
- Many hooks:
 - Translate a URL
 - Write info to log
 - Check client ID
 - Check access rights
- Hooks processed in order enforcing flow-of-control
- Functions in replaceable modules

Module Module Function Module
 Hook Hook Hook Hook Link between function and hook
 Functions called per hook Apache core Request Response

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.93

SERVER CLUSTERS

- Hosted across an LAN or WAN
- Collection of interconnected machines
- Can be organized in tiers:
 - Web server → app server → DB server
 - App and DB server sometimes integrated

Logical switch (possibly multiple) Application/compute servers Distributed file/database system
 Client requests Dispatched request First tier Second tier Third tier

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.94

LAN REQUEST DISPATCHING

- Front end of three tier architecture (logical switch) provides distribution transparency – hides multiple servers
- Transport-layer switches: switch accepts TCP connection requests, hands off to a server
 - Example: hardware load balancer (F5 networks – Seattle)
 - HW Load balancer - OSI layers 4-7
- Network-address-translation (NAT) approach:
 - All requests pass through switch
 - Switch sits in the middle of the client/server TCP connection
 - Maps (rewrites) source and destination addresses
- Connection hand-off approach:
 - **TCP Handoff:** switch hands of connection to a selected server

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.95

LAN REQUEST DISPATCHING - 2

- Who is the best server to handle the request?
- Switch plays important role in distributing requests
- Implements load balancing
- **Round-robin** – routes client requests to servers in a looping fashion
- **Transport-level** – route client requests based on TCP port number
- **Content-aware request distribution** – route requests based on inspecting data payload and determining which server node should process the request

Logically a single TCP connection Response Server
 Client Request Switch Request (handed off) Server

January 28, 2021
TCCS558: Applied Distributed Computing [Winter 2021]
School of Engineering and Technology, University of Washington - Tacoma
L8.96

WIDE AREA CLUSTERS

- Deployed across the internet
- Leverage resource/infrastructure from Internet Service Providers (ISPs)
- Cloud computing simplifies building WAN clusters
- Resource from a single cloud provider can be combined to form a cluster
- For deploying a cloud-based cluster (WAN), what are the implications of deploying nodes to:**
 - (1) a single availability zone (e.g. us-east-1e)?
 - (2) across multiple availability zones?

January 28, 2021

TCCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma

L8.97

WAN REQUEST DISPATCHING

- Goal:** minimize network latency using WANs (e.g. Internet)
- Send requests to nearby servers
- Request dispatcher:** routes requests to nearby server
- Example:** Domain Name System
 - Hierarchical decentralized naming system
- Linux:** find your DNS servers:


```
# Find you device name of interest
nmcli dev
# Show device configuration
nmcli device show <device name>
```

January 28, 2021

TCCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma

L8.98

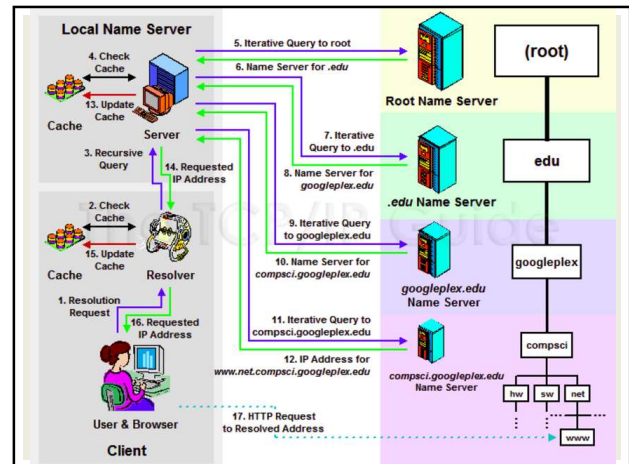
DNS LOOKUP

- First query local server(s) for address
- Typically there are (2) local DNS servers
 - One is backup
- Hostname may be cached at local DNS server
 - E.g. www.google.com
- If not found, local DNS server routes to other servers
- Routing based on components of the hostname
- DNS servers down the chain mask the client IP, and use the originating DNS server IP to identify a local host
- Weakness:** client may be far from DNS server used. Resolved hostname is close to DNS server, but not necessarily close to the client

January 28, 2021

TCCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma

L8.99



DNS: LINUX COMMANDS

- `nslookup <ip addr / hostname>`
- Name server lookup - translates hostname or IP to the inverse
- `traceroute <ip addr / hostname>`
- Traces network path to destination
- By default, output is limited to 30 hops, can be increased

January 28, 2021

TCCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma

L8.101

DNS EXAMPLE - WAN DISPATCHING

- Ping www.google.com in WA from wireless network:
 - `nslookup`: 6 alternate addresses returned, choose (74.125.28.147)
 - Ping 74.125.28.147: Average RTT = **22.458 ms (11 attempts, 22 hops)**
- Ping www.google.com in VA (us-east-1) from EC2 instance:
 - `nslookup`: 1 address returned, choose 172.217.9.196
 - Ping 172.217.9.196: Average RTT = 1.278 ms (11 attempts, 13 hops)
- From VA EC2 instance, ping WA www.google.com server
 - Ping 74.125.28.147: Average RTT 62.349ms (11 attempts, 27 hops)
 - Pinging the WA-local server is ~60x slower from VA
- From local wireless network, ping VA us-east-1 google :
 - Ping 172.217.9.196: Average RTT=81.637ms (11 attempts, 15 hops)

January 28, 2021

TCCS558: Applied Distributed Computing [Winter 2021]
 School of Engineering and Technology, University of Washington - Tacoma

L8.102

DNS EXAMPLE – WAN DISPATCHING

- Ping www.google.com in WA from wireless network:
 - nslookup: 6 alternate addresses returned, choose (74.125.28.147)


Latency to ping VA server in WA: ~3.63x
WA client: local-google 22.458ms to VA-google 81.637ms

Latency to ping WA server in VA: ~48.7x
VA client: local-google 1.278ms to WA-google 62.349!

- From local wireless network, ping VA us-east-1 google :
- Ping 172.217.9.196: Average RTT=81.637ms (11 attempts, 15 hops)

January 28, 2021 TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma L8.103

QUESTIONS



January 28, 2021 TCSS558: Applied Distributed Computing [Winter 2021] School of Engineering and Technology, University of Washington - Tacoma L8.104