TCSS 558: APPLIED DISTRIBUTED COMPUTING

Introduction

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

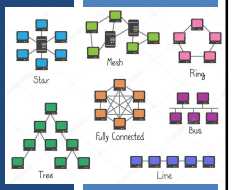

Where are you joining us from? (WORLD VERSION)

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

Where are you joining us from? (PUGET SOUND REGION)

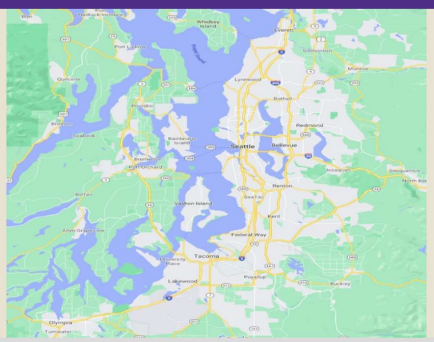

Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app

INTRODUCTIONS: What is your name? nickname / alias? and list one or more areas of interest in Computer Science:

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

OBJECTIVES – 1/5

- **Course Introduction**
- Syllabus
- Demographics Survey
- AWS Cloud Credits Survey
- Chapter 1 - What is a distributed system?
- Design goals of distributed systems:
  - Accessibility: resource sharing & availability
  - Distribution transparency
  - Openness
  - Scalability
- Activity: Design goals of distributed systems (Thursday?)

TCSS 558 A – Winter 2021

- Tuesday / Thursday 1:30-3:30pm ONLINE, Zoom
  - Jan 7: civil twilight - 5:10pm
  - Mar 16: civil twilight - 7:47pm
- 20 class meetings
  - Winter Quarter features 2 Monday holidays: Jan 18, Feb 15

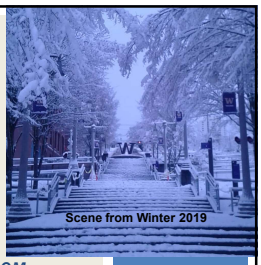

- Will there be snow?
  - *Winter 2021: No commuting w/ ZOOM*
  - *Winter 2019: campus closed 3 full + 2 half days*
  - *Upcoming?* →

- Final exam Tuesday March 16th

TEXT BOOK

- Van Steen, M. and Tanenbaum, A.S., Distributed systems, 3rd edition, Leiden, The Netherlands: Maarten van Steen, 2017.
- FREE online edition at:
- https://www.distributed-systems.net/index.php/books/distributed-systems/ds3-eBook/
- Paperback for $35
- https://www.amazon.com/dp/1543057381/ref=cm_sw_em_r_mt_dp_8Zb9FbH5G7HCK

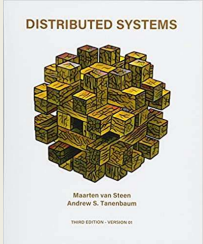

ALTERNATE BOOK

- Distributed Systems: Concepts and Design 5th Edition, 2011
- by George Coulouris, Jean Dollimore, Tim Kindberg, Gordon Blair
- Classic university distributed systems book
- A "denser" read
- Retails for ~$155
- A useful reference
- Not used directly in this course

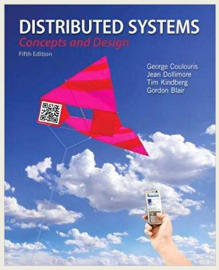

TCS558 COURSE WORK

- **Assignments – 3 [45%]**
  - Most assignments: can work in teams of 2 or 3
  - Assignment 0 is like a tutorial
- **Quizzes / Activities / Tutorials – [15%]**
  - ~ 2-4 total items (??)
  - Variety of formats: in class, online, reading, activity
- **Midterm – [20%]**
  - Open book, note, etc.
- **Final Exam – [20%]**

ONLINE DAILY FEEDBACK SURVEY

- Daily Feedback Quiz in Canvas – Available After Each Class
- Completing 100% of surveys ***ON TIME*** provides 2% extra credit for course grade
- Each survey is worth 1 point
- Tue survey: due Wed @ 10p
- Thur survey: due Mon @ 10p

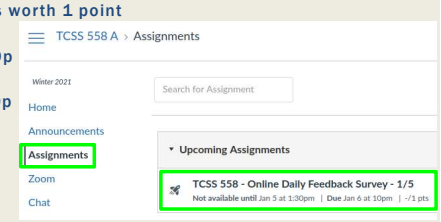

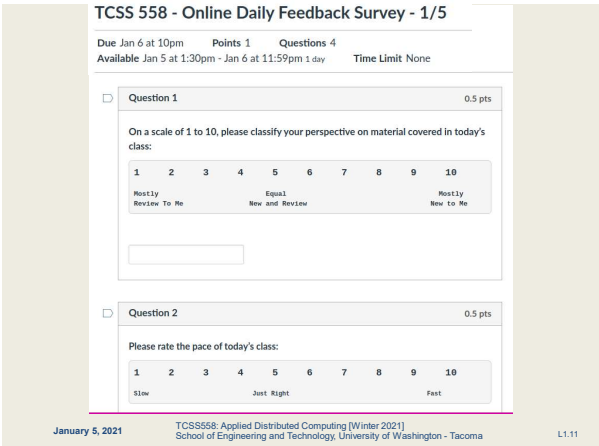

OBJECTIVES – 1/5

- **Course Introduction**
- Syllabus
- Demographics Survey
- AWS Cloud Credits Survey
- Chapter 1 - What is a distributed system?
- Design goals of distributed systems:
  - Accessibility: resource sharing & availability
  - Distribution transparency
  - Openness
  - Scalability
- Activity: Design goals of distributed systems (Thursday?)

OBJECTIVES – 1/5

- **Course Introduction**
- Syllabus
- Demographics Survey
- AWS Cloud Credits Survey
- Chapter 1 - What is a distributed system?
- Design goals of distributed systems:
  - Accessibility: resource sharing & availability
  - Distribution transparency
  - Openness
  - Scalability
- Activity: Design goals of distributed systems (Thursday?)

SURVEY LINKS

AT:

http://faculty.washington.edu/wlloyd/courses/tcss558/announcements.html

TCSS 558: Applied Distributed Computing | UNIVERSITY of WASHINGTON | TACOMA

ANNOUNCEMENTS | Syllabus | Grading | Schedule | Assignments | Home |

Course Announcments

1. Please check the SCHEDULE page for information related to the posting and due dates of the a
2. Please complete the online course demographics survey: [HERE]
3. Please complete the AWS Cloud Credits survey: [HERE]

WE WILL RETURN AT 2:40PM

OBJECTIVES – 1/5

- **Course Introduction**
- Syllabus
- Demographics Survey
- AWS Cloud Credits Survey
- Chapter 1 - What is a distributed system?
- Design goals of distributed systems:
  - Accessibility: resource sharing & availability
  - Distribution transparency
  - Openness
  - Scalability
- Activity: Design goals of distributed systems (Thursday?)

WHAT IS A DISTRIBUTED SYSTEM?

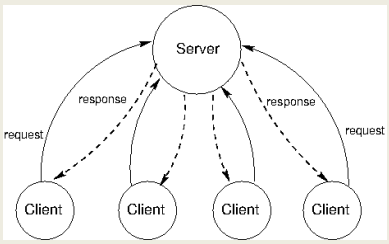

CLIENT/SERVER

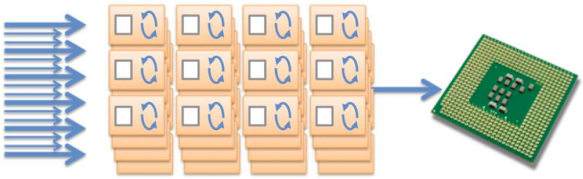

Hundreds of concurrent connections...

require hundreds of heavyweight threads or processes...

competing for limited CPU and memory

SYSTEM ARCHITECTURES

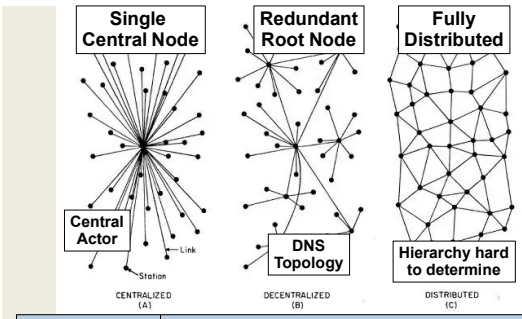

Consider implications for:

- State tracking
- Membership tracking
- Authentication

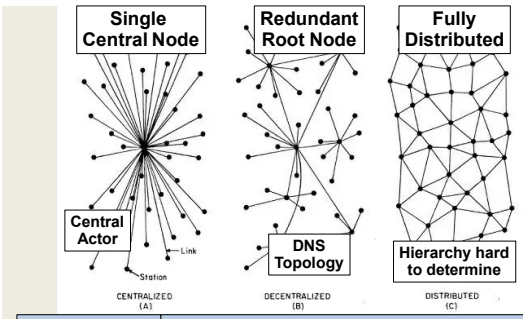

Consider implications for:

- State tracking
- Membership tracking
- Authentication

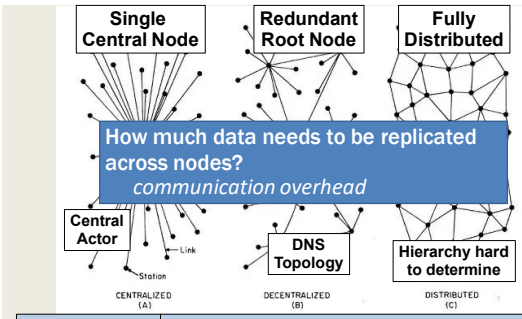

Consider implications for:

- State tracking
- Membership tracking
- Authentication

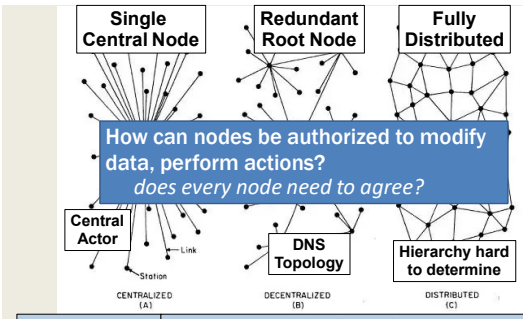

Consider implications for:

- State tracking
- Membership tracking
- Authentication

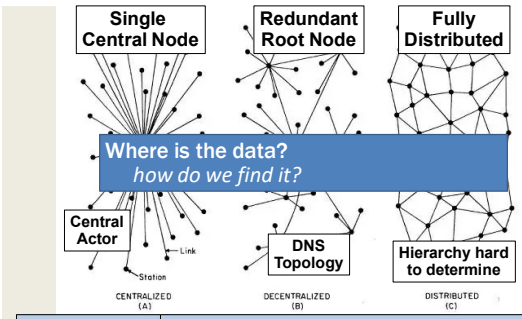

Consider implications for:

- State tracking
- Membership tracking
- Authentication

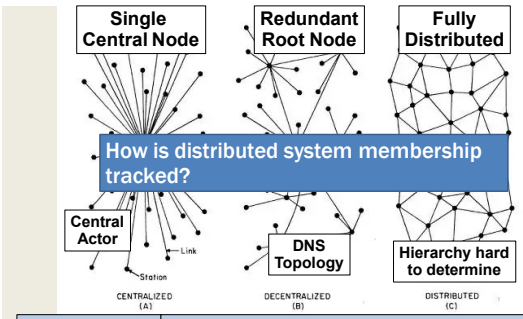

WHAT IS A DISTRIBUTED SYSTEM?

- Definition:
- A ***collection of autonomous computing elements*** that appears to users as a single coherent system.
- How nodes collaborate / communicate is **key**
- Nodes
  - Autonomous computing elements
  - Implemented as hardware or software processes
- Single coherent system
  - Users and applications perceive a single system
  - Nodes collaborate, and provide "abstraction"

CHARACTERISTICS OF DISTRIBUTED SYSTEMS - 1

- **#1: Collection of autonomous computing elements**
  - Node synchronization
  - Node coordination
  - Overlay networks – enable node connectivity
- **#2: Single coherent system**

C1: COLLECTION OF AUTONOMOUS COMPUTING ELEMENTS: NODE SYNCHRONIZATION

- Nodes behave/operate independently
- Maintain separate clocks (notion of time)
  - There is no global clock
- Nodes must address synchronization and coordination

Node synchronization and coordination…
- Subject of chapter 6

C1: COLLECTION OF AUTONOMOUS COMPUTING ELEMENTS: NODE COORDINATION

- Must manage **group membership**
- Nodes can join/leave the group
- Authorized vs. unauthorized nodes
- **Open group**: any node is allowed to join the distributed system
- **Closed group**: communication & membership is restricted
  - Admission control: supports mechanism to enable nodes to join/leave the group

C1: COLLECTION OF AUTONOMOUS COMPUTING ELEMENTS: OVERLAY NETWORKS

- Simply implies "on top of" another network
- Typically the internet
- Nodes in a collection communicate only with other nodes in the system
- The set of neighbors may be dynamic, or may even be known only implicitly (i.e., requires a lookup).
- **Structured:** each node has a well-defined set of neighbors with whom it can communicate (tree, ring).
- **Unstructured:** each node has references to randomly selected other nodes from the system.
- Always connected, communication paths are available

C1: COLLECTION OF AUTONOMOUS COMPUTING ELEMENTS: OSI MODEL

- Open systems interconnect:
- Standardization of the functionalities in a communication system via abstract layers

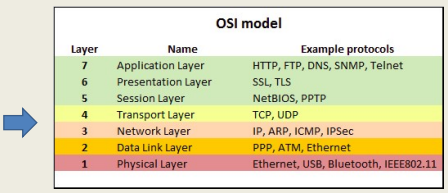

OSI model

Layer	Name	Example protocols
7	Application Layer	HTTP, FTP, DNS, SNMP, Telnet
6	Presentation Layer	SSL, TLS
5	Session Layer	NetBIOS, PPTP
4	Transport Layer	TCP, UDP
3	Network Layer	IP, ARP, ICMP, IPSec
2	Data Link Layer	PPP, ATM, Ethernet
1	Physical Layer	Ethernet, USB, Bluetooth, IEEE802.11

C1: COLLECTION OF AUTONOMOUS COMPUTING ELEMENTS: PEER-TO-PEER NETWORK

- Distributed systems leverage the ability of computing systems to collaborate and aggregate resources across many nodes providing potential for great scale and robustness than centralized client-server models
- How can ***fault tolerance*** be provided in the client/server model?
- How can ***fault tolerance*** be provided by the peer-to-peer model?

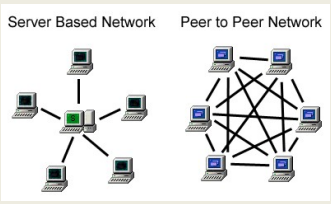

CHARACTERISTIC 2: SINGLE COHERENT SYSTEM

- Collection of nodes operates the same, regardless of where, when, and how interaction between a user and the system takes place
- **Distribution transparency:**
- From the user's perspective, they can't discern how the distributed system is implemented
- What are some examples of transparent distributed systems that you frequently use?

C2: SINGLE COHERENT SYSTEM DISTRIBUTION TRANSPARENCY

- An end user cannot tell where a computation takes place
- Where data is stored is abstracted (hidden)
- State of data replication is abstracted (hidden)
  - Is data consistent?
- Devices accessing services deployed on "The Cloud" is one example of distributed transparency

C2: SINGLE COHERENT SYSTEM DISTRIBUTION TRANSPARENCY - 2

- Partial failures: when part of a distributed system fails
- Hiding partial failures and their recovery is challenging
- Leslie Lamport, a distributed system is:
  "... one in which the failure of a computer you didn't even know existed can render your own computer unusable"

C2: SINGLE COHERENT SYSTEM MIDDLEWARE

- The OS of distributed systems:
- Provide facilities for inter-node communication
- Security, user account services (authentication, access control)
- Reliability: masking of and recovery from failures
- Service protocols
- Transaction support: "atomic" transactions all-or-nothing

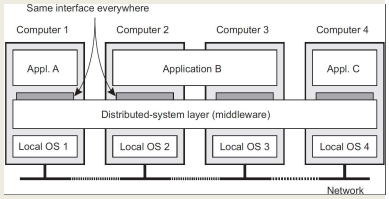

OBJECTIVES – 1/5

- **Course Introduction**
- Syllabus
- Demographics Survey
- AWS Cloud Credits Survey
- Chapter 1 - What is a distributed system?
- Design goals of distributed systems:
  - Accessibility: resource sharing & availability
  - Distribution transparency
  - Openness
  - Scalability
- Activity: Design goals of distributed systems (Thursday?)

DESIGN GOALS OF DISTRIBUTED SYSTEMS

- **Accessibility**: support for sharing resources
- **Distribution transparency**: the idea that how a system is distributed is hidden from users
- **Openness**: avoid vendor lock-in
- **Scalability**: ability to adapt and perform well with an increased or expanding workload or scope

OBJECTIVES – 1/5

- **Course Introduction**
- Syllabus
- Demographics Survey
- AWS Cloud Credits Survey
- Chapter 1 - What is a distributed system?
- Design goals of distributed systems:
  - Accessibility: resource sharing & availability
  - Distribution transparency
  - Openness
  - Scalability
- Activity: Design goals of distributed systems (Thursday?)

ACCESSIBILITY: RESOURCE SHARING

- Easy for users (and applications) to ***SHARE*** remote resources
  - Storage, compute, networks, services, peripherals, …
  - **Example:** Field programmable arrays (FPGAs) "as a service":

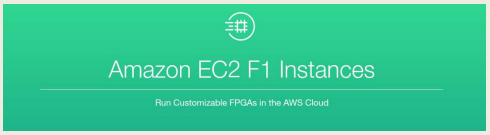

  - https://aws.amazon.com/ec2/instance-types/f1/
  - Make resources more ***AVAILABLE*** to end users
  - Nearly any resource can be shared

OBJECTIVES – 1/5

- **Course Introduction**
- Syllabus
- Demographics Survey
- AWS Cloud Credits Survey
- Chapter 1 - What is a distributed system?
- Design goals of distributed systems:
  - Accessibility: resource sharing & availability
  - Distribution transparency
  - Openness
  - Scalability
- Activity: Design goals of distributed systems (Thursday?)

DISTRIBUTION TRANSPARENCY

- In distributed systems, aspects of the implementation are hidden from users
- End users can simply use / consume the resource (or system) without worrying about the implementation details
- Technology aspects required to implement the distribution are abstracted from end users
- **The distribution is transparent to end users.**
- End users are not aware of certain mechanisms that do not appear in the distributed system because transparency confines details into layer(s) below the one users interact with. *(abstraction through layered architectures)*
- Users perceive the system as a single entity even though it's implementation is spread across a collection of devices.

DISTRIBUTION TRANSPARENCY - 2

- Types of distribution transparency
- Object is a resource or a process

Transparency	Description
Access	Hide differences in data representation and how an object is accessed.
Location	Hide where an object is located
Relocation	Hide that an object may be moved to another location while in use
Migration	Hide that an object may move to another location
Replication	Hide that an object is replicated
Concurrency	Hide than an object may be shared by several independent users
Failure	Hide the failure and recovery of an object

DISTRIBUTION TRANSPARENCY - 3

- **Location transparency**:
- Provided with Uniform resource locator (URLs) …
- Location is abstract: no client reconfiguration needed for relocation
- Users can't tell where an object physically is
- **Example:** during covid-19 students have location transparency from instructor enabled by Zoom
- **Relocation transparency:**
- Resource(s) can migrate from one server to another
- Initiated by the distributed system, possibly for maintenance
- Should a resource move while in use, users are unable to notice
- **Example:** Student changes Zoom client from laptop to cell phone - instructor does not notice
- **Migration transparency:**
- Feature offered by distributed systems
- Users are unaware if a resource possesses the ability to move to a different location

DISTRIBUTION TRANSPARENCY - 4

- **Replication transparency:**
- Hide the fact that several copies of a resource exist
- What if a user is aware of, or has to interact with the copies?
- **Reasons for replication:**
- Increase availability
- Improve performance
- Fault tolerance: a replica can take over when another fails

DISTRIBUTION TRANSPARENCY - 5

- **Concurrency transparency:**
- Concurrent use of resources requires synchronization w/ locks
- Transactions are often used
- Having concurrency transparency implies the client is unaware of locking mechanisms, etc.
- No special knowledge is needed
- **Failure transparency:**
- Masking failures is one of the hardest issues in dist. systems
- How do we tell the difference between a failed process and a very slow one?
- When do we need to "fail over" to a replica?
- Subject of chapter 8…

DEGREES OF DISTRIBUTION TRANSPARENCY

- Full distribution transparency may be impractical
- Communication latencies cannot be hidden
- Completely hiding failures of networks and nodes is impossible
  - Difference between slow computer and failing one
  - Transactions: did operation complete before crash?
- Full transparency will lead to slower performance:
  - Performance vs. transparency tradeoff
- Synchronizing replicas with a master requires time
- Immediately commit writes in fear of device failure

DEGREES OF DISTRIBUTION TRANSPARENCY - 2

- Abstracting location when user desires to interact intentionally with local resources / systems
- **Exposing** the distribution may be good:
  - Location-based-services (find nearby friends)
  - Help a user understand what's going on
  - When a server doesn't respond for a long time – is it far away?
  - Users in different times zones?
- Can you think of examples where distribution is not hidden?
  - Eventual consistency
  - Many online systems no longer update instantaneously
  - Users are getting accustomed to delays

OBJECTIVES – 1/5

- **Course Introduction**
- Syllabus
- Demographics Survey
- AWS Cloud Credits Survey
- Chapter 1 - What is a distributed system?
- Design goals of distributed systems:
  - Accessibility: resource sharing & availability
  - Distribution transparency
  - Openness
  - Scalability
- Activity: Design goals of distributed systems (Thursday?)

OPENNESS

- Capability of a system consisting of components that are easily used by, or integrated into other systems
- **Key aspects of openness:**
- Interoperability, portability, extensibility
- **Interoperability**: ability for components from separate systems to work together (different vendors?)
- Though implementation of a common interface
- How could we measure interoperability of components?

OPENNESS - 2

- **Portability**: degree that an application developed for distributed system A can be executed without modification on distributed system B
- How could we evaluate portability of a component?
- What percentage of portability is expected?
- The degree of portability will also reflect the ***reusability*** of the software

OPENNESS - 3

- **Extensibility**: easy to reconfigure, add, remove, replace components from different developers
- Example: replace the underlying file system of a distributed system
- To be open, we would like to ***separate policy from mechanism***
- Policy may change
- Mechanism is the technological implementation
- Avoid coupling policy and mechanism
- Enables flexibility
- Similar to separation of concerns, modular/OO design principle

ENABLING OPENNESS

- **Interfaces**: provide general syntax and semantics to interact with distributed components
- Services expose interfaces: functions, parameters, return values
- Semantics: describe what the services do
  - Often informally specified (via documentation)
- General interfaces enable alternate component implementations

SEPARATING POLICY FROM MECHANISM

- Example: **web browser caching**
- **Mechanism:** browser provides facility for storing documents
- **Policy:** Users decide which documents, for how long, …
- Goal: Enable users to set policies dynamically
- For example: browser may allow separate component plugin to specify policies
- **Tradeoff**: management complexity vs. policy flexibility
- Static policies are inflexible, but are easy to manage as features are barely revealed.
- AWS Lambda (Function-as-a-Service) abstracts configuration polices from the user resulting in management simplicity

OPENNESS EXAMPLE

- **Which of the following designs is more open?**
- Acme software corporation hosts a set of public weather web services (e.g. web service API)
- **DESIGN A:** API is implemented using MS .NET Remoting
- .NET Remoting is a mechanism for communicating between objects which are not in the same process. It is a generic system for different applications to communicate with one another. .NET objects are exposed to remote processes, thus allowing inter process communication. The applications can be located on the same computer, different computers on the same network, or on computers across separate networks.

OPENNESS EXAMPLE - 2

- **DESIGN B**: API is implemented using Java RMI
- The Java Remote Method Invocation (RMI) is a Java API that performs remote method invocation to allow Java objects to be distributed across different Java program instances on the same or different computers. RMI is the Java equivalent of C remote procedure calls, which includes support for transfer of serialized Java classes and distributed garbage-collection.
- **DESIGN C**: API is implemented as HTTP/RESTful web interface
- A RESTful API is an API that uses HTTP requests to GET, PUT, POST and DELETE data. RESTful APIs are referred to as a RESTful web services

Which of the following designs is more open?

- Design A: API is implemented using MS .NET Remoting
- Design B: API is implemented using Java RMI
- Design C: API is implemented with HTTP/RESTful web interface
- None of the above
- None of the above

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

OBJECTIVES – 1/5

- **Course Introduction**
- Syllabus
- Demographics Survey
- AWS Cloud Credits Survey
- Chapter 1 - What is a distributed system?
- Design goals of distributed systems:
  - Accessibility: resource sharing & availability
  - Distribution transparency
  - Openness
  - Scalability
- Activity: Design goals of distributed systems (Thursday?)

SCALABILITY

- The capability of a system to handle a growing amount of work by adding resources to the system
- Scalability is measured over multiple dimensions
- Two types: ***horizontal*** (scale out by adding more nodes) and ***vertical*** (scale up by adding resource to a single node)

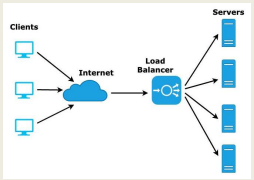

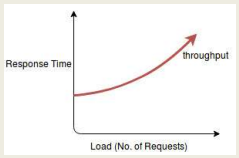

SCALABILITY DIMENSIONS

- **Size scalability**: distributed system can grow easily ***without*** impacting performance
  - Supports adding new users, processes, resources
- **Geographical scalability**: users and resources may be dispersed, but communication delays are negligible
- **Administrative scalability:** Policies are scalable as the distributed system grows to support more users… (security, configuration management policies are agile enough to deal with growth) ***Goal: have administratively scalable systems !***
- Most systems only account for size scalability
- One solution is to operate multiple parallel independent nodes

SIZE SCALABILITY

- Centralized architectures have limitations
- At some point a single central coordinator/arbitrator node can't keep up
  - Centralized server: limited CPU, disk, network capacity
- Scaling requires surmounting bottlenecks

Lloyd W, Pallickara S, David O, Lyon J, Arabi M, Rojas K. Migration of multi-tier applications to infrastructure-as-a-service clouds: An investigation using kernel-based virtual machines. InGrid Computing (GRID), 2011 12th IEEE/ACM International Conference on 2011 Sep 21 (pp. 137-144). IEEE.

GEOGRAPHIC SCALABILITY

- Nodes dispersed by great distances
  - Communication is slower, less reliable
  - Bandwidth may be constrained
- How do you support synchronous communication?
  - Latencies may be higher
  - Synchronous communication may be too slow and timeout
  - WAN links can be unreliable

ADMINISTRATIVE SCALABILITY

- Conflicting policies regarding usage (payment), management, and security
- How do you manage security for multiple, discrete data centers?
- Grid computing: how can resources be shared across disparate systems at different domains, etc. ?

APPROACHES TO SCALING

- **Hide communication latencies**
  - Use asynchronous communication to do other work and hide latency
  - Remote server runs in parallel in the background – client not locked
  - Separate event handler captures return response from server
- Hide latency by moving key press validation to client:

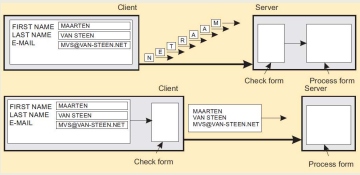

APPROACHES TO SCALING - 2

- Partitioning data and computations across machines
- Just one copy
  - Where is the copy?
- Move computations to the client
  - Thin client → thick client
  - Edge, fog, cloud….
- Decentralized naming services (DNS)
- Decentralized information services (WWW)

APPROACHES TO SCALING - 3

- Replication and caching – make copies of data available at different machines
- Replicated file servers and databases
- Mirrored web sites
- Web caches (in browsers and proxies)
- File caches (at server and client)
- **LOAD BALANCER** (or proxy server)
  - Commonly used to distribute user requests to nodes of a distributed system

PROBLEMS WITH REPLICATION

- Having multiple copies leads to inconsistency (cached or replicated)
- Modifying one copy invalidates all of the others
- Keeping copies consistent requires global synchronization
- Global-synchronization prohibits large-scale up
  - Best to synchronize just a few copies or synchronization latency becomes too long, entire system slows down!
  - ***Consider how synchronization time increases with system size***
- Can these inconsistencies be tolerated?

1. Current temperature and wind speed from weather.com
2. Bank account balance – for a read only statement
3. Bank account balance – for a transfer/withdrawal transaction

DEVELOPING DISTRIBUTED SYSTEMS

- Developing a distributed system is a formidable task
- Many issues to consider:
- Reliable networks do not exist
- Networked communication is inherently insecure

FALSE ASSUMPTIONS ABOUT DISTRIBUTED SYSTEMS

- The network is reliable
- The network is secure
- The network is homogeneous
- The topology does not change
- Latency is zero
- Bandwidth is infinite
- Transport cost is zero
- There is one administrator

QUESTIONS