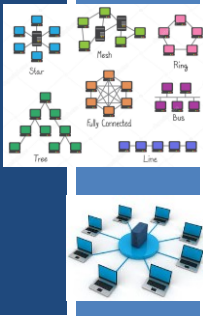


TCSS 558:  
APPLIED DISTRIBUTED COMPUTING

Distributed System Architectures – II, Middleware Organization

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



1

OBJECTIVES – 1/18

Questions from 1/16

Assignment 0: Cloud Computing Infrastructure Tutorial

Chapter 2: Distributed System Architectures:

- Chapter 2.1 – Architectural Styles
  - Resource-centered architectures
    - Representational state transfer (REST)
  - Event-based
    - Publish and subscribe (Rich Site Summary RSS feeds)

Class Activity: Architectural Styles

Chapter 2.2: Middleware Organization

Chapter 2.3: System Architectures

  - Centralized system architectures
  - Decentralized peer-to-peer architectures
  - Hybrid architectures

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

LS.2

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

Search for Assignment

Upcoming Assignments

TCSS 558 - Online Daily Feedback Survey - 1/5

Next available until Jan 5 at 1:30pm | Due Jan 6 at 10pm | -7.5 pts

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

LS.3

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 New 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 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

LS.4

4

MATERIAL / PACE

Please classify your perspective on material covered in today's class (23 respondents):

1-mostly review, 5-equal new/review, 10-mostly new

Average – 7.04 (↑ - previous 5.63)

Please rate the pace of today's class:

1-slow, 5-just right, 10-fast

Average – 6.09 (↑ - previous 5.00)

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

LS.5

5

FEEDBACK FROM 1/16

As someone with no networking background, I'm a bit intimidated by all the protocols and layers

In lecture 4, we presented the Open Systems Interconnection (OSI) model, a conceptual model created by the International Organization for Standardization (ISO).

The OSI model provides a common model that enables diverse communication systems to communicate using standard protocols.

The OSI model provides an excellent example of a layered architecture with 7 layers.

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

LS.6

6

OSI MODEL LAYERS

7

Application Layer

Human-computer interaction layer, where applications can access the network services

6

Presentation Layer

Ensures that data is in a usable format and is where data encryption occurs

5

Session Layer

Maintains connections and is responsible for controlling ports and sessions

4

Transport Layer

Transmits data using transmission protocols including TCP and UDP

3

Network Layer

Decides which physical path the data will take

2

Data Link Layer

Defines the format of data on the network

1

Physical Layer

Transmits raw bit stream over the physical medium

January 18, 2024

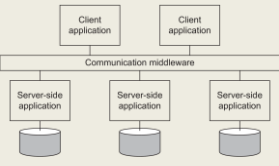
TCSS558: Applied Distributed Computing (Winter 2024)  
School of Engineering and Technology, University of Washington - Tacoma

LS.7

7

FEEDBACK - 2

- What are the key defining traits of a Distributed Information System that differentiates it from other system types?
- Distributed Information Systems are client/server apps found in organizations which over time were integrated to form enterprise-wide information systems
- Key features include:
  - Multi-client/server
  - Use of atomic transactions
  - Separation into components
  - Remote integration
    - RPC, RMI, Message queues
- Some common SaaS apps are managed Dist Info Sys

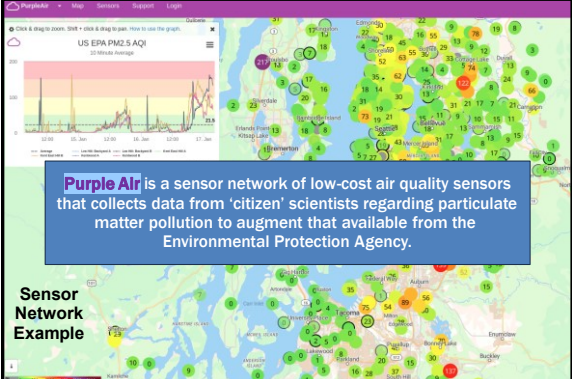


January 18, 2024

TCSS558: Applied Distributed Computing (Winter 2024)  
School of Engineering and Technology, University of Washington - Tacoma

LS.8

8



Purple Air is a sensor network of low-cost air quality sensors that collects data from 'citizen' scientists regarding particulate matter pollution to augment that available from the Environmental Protection Agency.

Sensor Network Example

January 18, 2024

TCSS558: Applied Distributed Computing (Winter 2024)  
School of Engineering and Technology, University of Washington - Tacoma

LS.9

9

OBJECTIVES - 1/18

- Questions from 1/16
- Assignment 1: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
  - Chapter 2.1 - Architectural Styles
    - Resource-centered architectures
      - Representational state transfer (REST)
    - Event-based
      - Publish and subscribe (Rich Site Summary RSS feeds)
  - Class Activity: Architectural Styles
  - Chapter 2.2: Middleware Organization
  - Chapter 2.3: System Architectures
    - Centralized system architectures
    - Decentralized peer-to-peer architectures
    - Hybrid architectures

January 18, 2024

TCSS558: Applied Distributed Computing (Winter 2024)  
School of Engineering and Technology, University of Washington - Tacoma

LS.10

10

ASSIGNMENT 1

- Preparing for Assignment 1: Intro to Cloud Computing Infrastructure and Load Balancing
  - Establish AWS Account - Standard account
- Coming Soon - - PREVIEW:
  - Task 0 - Establish local Linux/Ubuntu environment
  - Task 1 - AWS account setup, obtain user credentials
  - Task 2 - Intro to: Amazon EC2 & Docker: create Dockerfile for Apache Tomcat
  - Task 3 - Create Dockerfile for haproxy (software load balancer)
  - Task 4 - Working with Docker-Machine
  - Task 5 - Submit Results of testing alternate server configs

January 18, 2024

TCSS558: Applied Distributed Computing (Winter 2024)  
School of Engineering and Technology, University of Washington - Tacoma

LS.11

11

OBJECTIVES - 1/18

- Questions from 1/16
- Assignment 1: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
  - Chapter 2.1 - Architectural Styles
    - Resource-centered architectures
      - Representational state transfer (REST)
    - Event-based
      - Publish and subscribe (Rich Site Summary RSS feeds)
  - Class Activity: Architectural Styles
  - Chapter 2.2: Middleware Organization
  - Chapter 2.3: System Architectures
    - Centralized system architectures
    - Decentralized peer-to-peer architectures
    - Hybrid architectures

January 18, 2024

TCSS558: Applied Distributed Computing (Winter 2024)  
School of Engineering and Technology, University of Washington - Tacoma

LS.12

12

CH 2.1 - ARCHITECTURAL STYLES

Layered

Lecture 4

Object-based

Service oriented architecture (SOA)

Resource-centered architectures

Representational state transfer (REST)

Event-based

Publish and subscribe (Rich Site Summary RSS feeds)

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.13

13

OBJECT-BASED ARCHITECTURES

Enables loose and flexible component organization

Objects == components

Enable distributed node interaction via function calls over the network

Began with C - Remote Procedure Calls (RPC)

Straightforward: package up function inputs, send over network, transfer results back

Language dependent

In contrast to web services, RPC calls originally were more intimate in nature

Procedures more "coupled", not as independent

The goal was not to decouple and widgetize everything

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.14

14

OBJECT-BASED ARCHITECTURES - 2

A diagram showing four 'Object' nodes. Two objects at the top are connected by a double-headed arrow. Two objects at the bottom are also connected by a double-headed arrow. A single-headed arrow points from the top-left object to the bottom-left object, labeled 'Method call'.

Distributed objects Java- Remote Method Invocation (RMI)

Adds object orientation concepts to remote function calls

Clients bind to proxy objects

Proxy provide an object interface which transfers method invocation over the network to the remote host

How do we replicate objects?

Object marshalling - serialize data, stream it over network

Unmarshalling- create an object from the stream

Unmarshall local object copies on the remote host

JSON, XML are some possible data formats

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.15

15

DISTRIBUTED OBJECTS

A diagram showing a 'Client machine' and a 'Server machine' connected via a 'Network'. The client machine contains a 'Client' box with a 'Proxy' box below it. The server machine contains a 'Server' box with 'Object', 'State', 'Method', and 'Interface' components. A 'Skeleton' box is also shown. Arrows indicate 'Client invokes a method' from the client to the proxy, 'Same interface as object' between proxy and skeleton, 'Skeleton invokes same method at object' from skeleton to server, and 'Marshaled invocation is passed across network' from client to server.

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.16

16

DISTRIBUTED OBJECTS - 2

A counterintuitive feature is that state is not distributed

Each "remote object" maintains its own state

Remote objects may not be replicated

Objects may be "mobile" and move around from node to node

Common for data objects

For distributed (remote) objects consider

Pass by value

Pass by reference .... (does this make sense?)

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.17

17

CH 2.1 - ARCHITECTURAL STYLES

Layered

Object-based

Service oriented architecture (SOA)

Resource-centered architectures

Representational state transfer (REST)

Event-based

Publish and subscribe (Rich Site Summary RSS feeds)

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.18

18

SERVICE ORIENTED ARCHITECTURE

- Services provide always-on encapsulated functions over the internet/web
- Leverage redundant cloud computing infrastructure
- Services may:
  - Aggregate multiple languages, libraries, operating systems
  - Include (wrap) legacy code
- Many software components may be involved in the implementation
  - Application server(s), relational database(s), key-value stores, in memory-cache, queue/messaging services

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

LS.19

19

SERVICE ORIENTED ARCHITECTURE - 2

- Are more easily developed independently and shared vs. systems with distributed object architectures
- Less coupling
- An error while invoking a distributed object may crash the system
- An error calling a service (e.g. mismatching the interface) generally does not result in a system crash

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

LS.20

20

OBJECTIVES – 1/18

- Questions from 1/16
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
  - Chapter 2.1 – Architectural Styles
    - Resource-centered architectures**
      - Representational state transfer (REST)**
    - Event-based
      - Publish and subscribe (Rich Site Summary RSS feeds)
  - Class Activity: Architectural Styles
  - Chapter 2.2: Middleware Organization
  - Chapter 2.3: System Architectures
    - Centralized system architectures
    - Decentralized peer-to-peer architectures
    - Hybrid architectures

January 18, 2024


TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

LS.21

21

RESOURCE BASED ARCHITECTURES

- Motivation:
  - Increasing number of services available online
  - Each with specific protocol(s), methods of interfacing
  - Connecting services w/ different TCP/IP protocols → integration nightmare
    - Need for specialized client for each service that speaks the application protocol "language"...
- Need standardization of interfaces
  - Make services/components more pluggable
  - Easier to adopt and integrate
  - Common architecture



January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

LS.22

22

REST SERVICES

- Representational State Transfer (REST)
- Built on HTTP
- Four key characteristics:
  - Resources identified through single naming scheme
  - Services offer the same interface
    - Four operations: GET PUT POST DELETE
  - Messages to/from a service are fully described
  - After execution server forgets about client
    - Stateless execution

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

LS.23

23

HYPertext TRANSPORT PROTOCOL (HTTP)

- An ASCII-based request/reply protocol for transferring information on the web
- HTTP request includes:
  - request method (GET, POST, etc.)
  - Uniform Resource Identifier (URI)
  - HTTP protocol version understood by the client
  - headers—extra info regarding transfer request
- HTTP response from server
  - Protocol version & status code →
  - Response headers
  - Response body

HTTP status codes:  
2xx — all is well  
3xx — resource moved  
4xx — access problem  
5xx — server error

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

LS.24

24

REST-FUL OPERATIONS

Operation	Description	
PUT	Create a new resource	(C)reate
GET	Retrieve state of a resource in some format	(R)ead
POST	Modify a resource by transferring a new state	(U)pdate
DELETE	Delete a resource	(D)elete

- Resources often implemented as objects in OO languages
- REST is weak for tracking state
- Generic REST interfaces enable ubiquitous “so many” clients

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.25

25

EXAMPLE: AMAZON S3

- Amazon S3 offers a REST-based interface
- Requires signing HTTP authorization header or passing authentication parameters in the URL query string
- REST: GET/PUT/POST/DELETE
- SOAP: 16 operations, moving toward deprecation
- Python boto ~50 operations (SDK for Python)
- SDKs for other languages

- AWS SDKs and Explorers
  - Set Up the AWS CLI
  - Using the AWS SDK for Java
  - Using the AWS SDK for .NET
  - Using the AWS SDK for PHP and Running PHP Examples
  - Using the AWS SDK for Ruby - Version 3
  - Using the AWS SDK for Python (Boto)

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.26

26

REST - 2

- Defacto web services protocol
- Requests made to a URI – uniform resource identifier
- Supersedes SOAP – Simple Object Access Protocol
  - SOAP – application protocol specific to web services
- Access and manipulate web resources with a predefined set of stateless operations (known as web services)
- Responses most often in JSON, also HTML, ASCII text, XML, no real limits as long as text-based
- curl – generic command-line REST client:  
<https://curl.haxx.se/>

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.27

27

// SOAP REQUEST - Book Store - Query Price

POST /InStock HTTP/1.1  
Host: www.bookshop.org  
Content-Type: application/soap+xml; charset=utf-8  
Content-Length: nnn

<?xml version="1.0"?>  
<soap:Envelope  
  xmlns:soap="http://www.w3.org/2001/12/soap-envelope"  
  soap:encodingStyle="http://www.w3.org/2001/12/soap-encoding">  
  <soap:Body xmlns:m="http://www.bookshop.org/prices">  
    <m:GetBookPrice>  
      <m:BookName>The Fleamarket</m:BookName>  
    </m:GetBookPrice>  
  </soap:Body>  
</soap:Envelope>

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L11.28

28

// SOAP RESPONSE - Book Store - Query Price

POST /InStock HTTP/1.1  
Host: www.bookshop.org  
Content-Type: application/soap+xml; charset=utf-8  
Content-Length: nnn

<?xml version="1.0"?>  
<soap:Envelope  
  xmlns:soap="http://www.w3.org/2001/12/soap-envelope"  
  soap:encodingStyle="http://www.w3.org/2001/12/soap-encoding">  
  <soap:Body xmlns:m="http://www.bookshop.org/prices">  
    <m:GetBookPriceResponse>  
      <m:Price>10.95</m:Price>  
    </m:GetBookPriceResponse>  
  </soap:Body>  
</soap:Envelope>

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L11.29

29

// Web Service Definition Language (WSDL)  
// Service Definition - Day of Week Service

<?xml version="1.0" encoding="UTF-8"?>  
<definitions name="DayOfWeek"  
  targetNamespace="http://www.roquewave.com/soapwsc/examples/DayOfWeek.wsdl"  
  xmlns:tns="http://www.roquewave.com/soapwsc/examples/DayOfWeek.wsdl"  
  xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/"  
  xmlns:xsd="http://www.w3.org/2001/XMLSchema" xmlns="http://schemas.xmlsoap.org/wsdl/">  
  <message name="DayOfWeekInput">  
    <part name="date" type="xsd:date"/>  
  </message>  
  <message name="DayOfWeekResponse">  
    <part name="dayOfWeek" type="xsd:string"/>  
  </message>  
  <portType name="DayOfWeekPortType">  
    <operation name="GetDayOfWeek">  
      <input message="tns:DayOfWeekInput"/>  
      <output message="tns:DayOfWeekResponse"/>  
    </operation>  
  </portType>  
  <binding name="DayOfWeekBinding" type="tns:DayOfWeekPortType">  
    <soap:binding style="document" transport="http://schemas.xmlsoap.org/soap/http"/>  
    <operation name="GetDayOfWeek">  
      <soap:operation soapAction="getDayOfWeek"/>  
      <input>  
        <soap:body use="encoded" namespace="http://www.roquewave.com/soapwsc/examples" encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>  
      </input>  
      <output>  
        <soap:body use="encoded" namespace="http://www.roquewave.com/soapwsc/examples" encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>  
      </output>  
    </operation>  
  </binding>  
  <service name="DayOfWeekService">  
    <documentation>  
      Returns the day-of-week name for a given date  
    </documentation>  
    <port name="DayOfWeekPort" binding="tns:DayOfWeekBinding">  
      <soap:address location="http://localhost:8090/dayOfWeek/DayOfWeek"/>  
    </port>  
  </service>  
</definitions>

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.30

30

REST CLIMATE SERVICES EXAMPLE

- USDA Lat/Long Climate Service Demo
- Just provide a Lat/Long

```
// REST/JSON
// Request climate data for Washington

{
  "parameter": [
    {
      "name": "latitude",
      "value": 47.2529
    },
    {
      "name": "longitude",
      "value": -122.4443
    }
  ]
}
```

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.31

31

WE WILL RETURN AT 4:50PM



January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.32

32

OBJECTIVES – 1/18

- Questions from 1/16
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
  - Chapter 2.1 – Architectural Styles
    - Resource-centered architectures
      - Representational state transfer (REST)
    - Event-based
      - Publish and subscribe (Rich Site Summary RSS feeds)****
- Class Activity: Architectural Styles
- Chapter 2.2: Middleware Organization
- Chapter 2.3: System Architectures
  - Centralized system architectures
  - Decentralized peer-to-peer architectures
  - Hybrid architectures

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.33

33

PUBLISH-SUBSCRIBE ARCHITECTURES: EVENT-BASED

- Enables separation between processing and coordination
- Types of coordination:

	Temporally coupled (at the same time)	Temporally decoupled (at different times)
Referentially coupled (dependent on name)	<b>Direct</b> Explicit synchronous service call	<b>Mallbox</b> Asynchronous by name (address)
Referentially decoupled (name not required)	<b>Event-based</b> Event notices published to shared bus, w/o addressing	<b>Shared data space</b> Processes write tuples to a shared data space

Publish and subscribe architectures

January 18, 2024

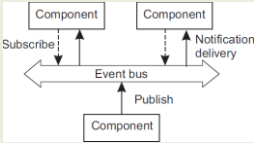
TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.34

34

PUBLISH-SUBSCRIBE ARCHITECTURES - 2

- Event-based coordination**
- Processes do not know about each other explicitly
- Processes:
  - Publish:** a notification describing an event
  - Subscribe:** to receive notification of specific kinds of events
- Assumes subscriber is presently up (*temporally coupled*)
- Subscribers must actively **MONITOR** event bus



January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.35

35

PUBLISH-SUBSCRIBE ARCHITECTURES: SHARED DATA SPACE

- Enables separation between processing and coordination
- Types of coordination:

	Temporally coupled (at the same time)	Temporally decoupled (at different times)
Referentially coupled (dependent on name)	<b>Direct</b> Explicit synchronous service call	<b>Mallbox</b> Asynchronous by name (address)
Referentially decoupled (name not required)	<b>Event-based</b> Event notices published to shared bus, w/o addressing	<b>Shared data space</b> Processes write tuples to a shared data space

Publish and subscribe architectures

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.36

36

### PUBLISH SUBSCRIBE ARCHITECTURES - 3

- **Shared data space**
- Full decoupling (name and time)
- Processes publish "tuples" to shared dataspace (publish)
- Processes provide search pattern to find tuples (subscribe)
- Subscribers are notified of matches (both existing and newly published tuples)
- **Key characteristic:** Processes have no explicit reference to each other

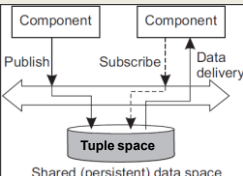


Diagram illustrating Publish-Subscribe Architecture. Two components are shown at the top. The left component has a 'Publish' arrow pointing to a 'Tuple space' (a cylinder) in the center. The right component has a 'Subscribe' arrow pointing to the same 'Tuple space'. Below the 'Tuple space' is the label 'Shared (persistent) data space'. A 'Data delivery' arrow points from the 'Tuple space' to the right component.

January 18, 2024 TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma LS.37

37

### PUBLISH SUBSCRIBE ARCHITECTURES - 4

- Subscriber describes events interested in
- Complex descriptions are intensive to evaluate and fulfil
- **Middleware will:**
  - Publish matching notification and data to subscribers
  - Common if middleware lacks storage
- Publish only matching notification
  - Common if middleware provides storage facility
  - Client must explicitly fetch data on their own
- Publish and subscribe systems are generally scalable
- **What would reduce the scalability of a publish-and-subscribe system?**

January 18, 2024 TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma LS.38

38

### OBJECTIVES – 1/18

- Questions from 1/16
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
  - Chapter 2.1 – Architectural Styles
    - Resource-centered architectures
    - Representational state transfer (REST)
    - Event-based
    - Publish and subscribe (Rich Site Summary RSS feeds)
  - **Class Activity: Architectural Styles**
  - Chapter 2.2: Middleware Organization
  - Chapter 2.3: System Architectures
    - Centralized system architectures
    - Decentralized peer-to-peer architectures
    - Hybrid architectures

January 18, 2024 TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma LS.39

39

### IN-CLASS ACTIVITY: ARCHITECTURAL STYLES



LS.40

40

### CLASS ACTIVITY 2

- We will form groups of ~2-3
  - On Zoom breakout rooms will be created
- Each group will complete a MS Doc worksheet
- Add names to the Doc as they appear in Canvas
- Once completed, **one person** submits a PDF to Canvas
- Instructor will score all group members based on the uploaded PDF file
- To get started – link is under Class Activity 2 in Canvas:
  - Log into your \*\*\* UW NET ID \*\*\*
  - Link to shared doc file on Canvas
  - Follow link:  
<https://canvas.uw.edu/files/114972397/>

January 18, 2024 TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma LS.41

41

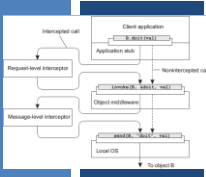
### DISTRIBUTED SYSTEM GOALS TO CONSIDER

- **Consider how the architectural change may impact:**
- Availability
- Accessibility
- Responsiveness
- Scalability
- Openness
- Distribution transparency
- Supporting resource sharing
- Other factors...

January 18, 2024 TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma LS.42

42

CH 2.2: MIDDLEWARE ORGANIZATION



January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.43

43

OBJECTIVES - 1/18

- Questions from 1/16
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
  - Chapter 2.1 – Architectural Styles
    - Resource-centered architectures
      - Representational state transfer (REST)
    - Event-based
      - Publish and subscribe (Rich Site Summary RSS feeds)
  - Class Activity: Architectural Styles
  - Chapter 2.2: Middleware Organization**
  - Chapter 2.3: System Architectures
    - Centralized system architectures
    - Decentralized peer-to-peer architectures
  - Hybrid architectures

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.44

44

MIDDLEWARE ORGANIZATION

- Relies on two important design patterns:
  - Wrappers
  - Interceptors
- Both help achieve the goal of openness

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.45

45

MIDDLEWARE: WRAPPERS

- Wrappers (also called adapters)**
  - WHY?** Interfaces available from legacy software may not be sufficient for all new applications to use
  - WHAT:** Special “frontend” components that provide interfaces for clients
    - Interface wrappers transform client requests to “implementation” (i.e. legacy software) at the component-level
    - Can then provide modern service interfaces for legacy code/systems
  - Components encapsulate (i.e. abstract) dependencies to meet all preconditions to operate and host legacy code
  - Interfaces parameterize legacy functions, abstract environment configuration (i.e. make into black box)
- Contributes towards system **OPENNESS**
- Example: Amazon S3:** S3 HTTP REST interface
- GET/PUT/DELETE/POST: requests handed off for fulfillment

January 18, 2024

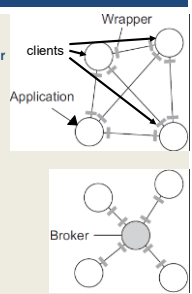
TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.46

46

MIDDLEWARE: WRAPPERS - 2

- Inter-application communication
  - Applications may provide unique interface for every client application
- Scalability suffers
  - $N \text{ applications} \rightarrow O(N^2) \text{ wrappers}$
- ALTERNATE: Use a Broker**
  - Provide a common intermediary
  - Broker knows how to communicate with every application
  - Applications only know how to communicate with the broker



January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.47

47

MIDDLEWARE: INTERCEPTORS

- Interceptor**
- Software construct, breaks flow of control, allows other application code to be executed
- Interceptors send calls to other servers, or to ALL servers that replicate an object while abstracting the **distribution** and/or **replication**
  - Used to enable remote procedure calls (RPC), remote method invocation (RMI)
- Object A calls method belonging to object B
  - Interceptors route calls to object B regardless of location

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.48

48



MIDDLEWARE: INTERCEPTORS - 2

Request-level interceptor transforms:  
`B.doit(val)`  
into generic call:  
`invoke(B, doIt, val)`

Message-level interceptor in middleware sends message through OS (TCP/IP socket) to transfer data:  
`send(B, "doIt", val)`

Non-intercepted:

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.49

49

MIDDLEWARE INTERCEPTION - METHOD

- **MIDDLEWARE:** Provides local interface matching Object B to Object A
- Object A calls Object B's method provided by local interface
- A's call is transformed into a "generic object invocation" by request-level Interceptor
- "Generic object invocation" is transformed into a **message** by message-level Interceptor and sent over Object A's network to Object B
- Interception automatically routes calls to all object replicas

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.50

50

MODIFIABLE MIDDLEWARE

- **GOAL:** It should be possible to modify middleware without loss of availability
  - Software components can be replaced at runtime
- Component-based design
  - Modifiability through composition
  - Systems may have static or dynamic configuration of components
  - Dynamic configuration requires **late binding**
  - Components can be changed at runtime
- Component based software supports modifiability at runtime by enabling components to be swapped out.
- Does a microservices architecture (e.g. AWS Lambda) support modifiability at runtime ?

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.51

51

CH 2.3: SYSTEM ARCHITECTURES

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.52

52

OBJECTIVES - 1/18

- Questions from 1/16
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
  - Chapter 2.1 - Architectural Styles
  - Resource-centered architectures
    - Representational state transfer (REST)
  - Event-based
    - Publish and subscribe (Rich Site Summary RSS feeds)
- Class Activity: Architectural Styles
- Chapter 2.2: Middleware Organization
- **Chapter 2.3: System Architectures**
  - Centralized system architectures
  - Decentralized peer-to-peer architectures
  - Hybrid architectures

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.53

53

SYSTEM ARCHITECTURES

- Architectural styles (or patterns)
- General, reusable solutions to commonly occurring system design problems
- Expressed as a logical organization of **components** and **connectors**
- Deciding on the system components, their interactions, and placement is a "realization" of an **architectural style**
- System architectures represent designs used in practice

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.54

54

OBJECTIVES – 1/18

- Questions from 1/16
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
  - Chapter 2.1 – Architectural Styles
    - Resource-centered architectures
      - Representational state transfer (REST)
    - Event-based
      - Publish and subscribe (Rich Site Summary RSS feeds)
  - Class Activity: Architectural Styles
  - Chapter 2.2: Middleware Organization
  - Chapter 2.3: System Architectures
    - Centralized system architectures
    - Decentralized peer-to-peer architectures
    - Hybrid architectures

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.55

55

TYPES OF SYSTEM ARCHITECTURES

- Centralized system architectures
  - Client-server
  - Multitiered
- Decentralized peer-to-peer architectures
  - Structured
  - Unstructured
  - Hierarchically organized
- Hybrid architectures

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.56

56

CENTRALIZED:  
SIMPLE CLIENT-SERVER ARCHITECTURE

- Clients request services
- Servers provide services
- Request-reply behavior
- Connectionless protocols (UDP)
  - Assume stable network communication with no failures
  - Best effort communication: No guarantee of message arrival without errors, duplication, delays, or in sequence. No acknowledgment of arrival or retransmission
  - Problem: How to detect whether the client request message is lost, or the server reply transmission has failed
  - Clients can resend the request when no reply is received
  - But what is the server doing?

```
sequenceDiagram
    participant Client
    participant Server
    Client->>Server: Request
    Note over Client: Wait
    Server-->>Client: Reply
```

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.57

57

CLIENT-SERVER PROTOCOLS

- Connectionless cont'd
  - Is resending the client request a good idea?
  - Examples:
    - Client message: "transfer \$10,000 from my bank account"
    - Client message: "tell me how much money I have left"
  - Idempotent – repeating requests is safe
- Connection-oriented (TCP)
  - Client/server communication over wide-area networks (WANs)
  - When communication is inherently reliable
  - Leverage "reliable" TCP/IP connections

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.58

58

CLIENT-SERVER PROTOCOLS - 2

- Connection-oriented cont'd
  - Set up and tear down of connections is relatively expensive
  - Overhead can be amortized with longer lived connections
    - Example: database connections often retained
- Ongoing debate:
  - How do you differentiate between a client and server?
  - Roles are blurred
- Blurred Roles Example: Distributed databases
  - DB nodes both service client requests, \*and\* submit new requests to other DB nodes for replication, synchronization, etc.

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.59

59

TCP/UDP

TCP	UDP
Reliable	Unreliable
Connection-oriented	Connectionless
Segment retransmission and flow control through windowing	No windowing or retransmission
Segment sequencing	No sequencing
Acknowledge segments	No acknowledgement

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.60

60

CONNECTIONLESS VS CONNECTION ORIENTED

	Connectionless (UDP) stateless	Connection-oriented (TCP) stateful
Advantages		
Disadvantages		

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.61

61

CONNECTIONLESS VS CONNECTION ORIENTED

	Connectionless (UDP) stateless	Connection-oriented (TCP) stateful
Advantages	<ul style="list-style-type: none"><li>Fast to communicate (no connection overhead)</li><li>Broadcast to an audience</li><li>Network bandwidth savings</li></ul>	<ul style="list-style-type: none"><li>Message delivery confirmation</li><li>Idempotence not required</li><li>Messages automatically resent - if client (or network) is temporarily unavailable</li><li>Message sequences guaranteed</li></ul>
Disadvantages	<ul style="list-style-type: none"><li>Cannot tell difference of request vs. response failure</li><li>Requires idempotence</li><li>Clients must be online and ready to receive messages</li></ul>	<ul style="list-style-type: none"><li>Connection setup is time-consuming</li><li>More bandwidth is required (protocol, retries, multinode-communication)</li></ul>

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.62

62

MULTITIERED ARCHITECTURES

Where should functionality be distributed?

- At the client?
- At the server?

Why should we consider component composition?

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.63

63

**Bell's Number:**

k: number of ways n components can be distributed across containers

n	k
4	15
5	52
6	203
7	877
8	4,140
9	21,147
n	...

SC14: M, D, F, L; SC15: M, D, F, L

M: Tomcat Application Server  
D: PostgreSQL DB  
F: nginx file server  
L: Logging server (high O/H)

64

Resource utilization profile changes from component composition

**M-bound RUSLE2 – Soil Erosion Model Webservice**

- Box size shows absolute deviation (+/-) from mean
- Shows relative magnitude of performance variance

**Two application variants tested**

- M-bound: Standard service, M is compute bound
- D-bound: Modified service, D is compute bound

	M-bound	D-bound
Disk sector reads:	21.8%	111.1%
Disk sector writes:	144.9%	145%
Network bytes received:	143.7%	143.9%
Network bytes sent:		

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.64

65

PERFORMANCE IMPLICATIONS OF COMPONENT DEPLOYMENTS

**Δ Performance Change:**  
Min to max performance

M-bound: 14%  
D-bound: 25.7%

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.65

66

MULTITIERED ARCHITECTURES - 2

- **M D F L** architecture
- **M** – is the application server
- **M** – is also a client to the database (**D**),  
fileserver (**F**), and logging server (**L**)

```
graph TD
    client --> M
    M --> D
    M --> F
    M --> L
```

```
sequenceDiagram
    participant Client
    participant Application server
    participant Database server
    Client->>Application server: Request operation
    Application server->>Database server: Request data
    Database server-->>Application server: Return data
    Application server-->>Client: Return reply
    Note over Client: Wait for reply
    Note over Application server: Wait for data
```

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.67

67

MULTITIERED RESOURCE SCALING

- **Vertical distribution**
- The distribution of “M D F L”
- Application is scaled by placing “tiers” on separate servers
  - **M** – The application server
  - **D** – The database server
- Vertical distribution impacts “network footprint” of application
- Service isolation: each component is isolated on its own HW
- **Horizontal distribution**
- Scaling an individual tier
- Add multiple machines and distribute load
- Load balancing

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.68

68

MULTITIERED RESOURCE SCALING - 2

- **Horizontal distribution cont'd**
- Sharding: portions of a database map” to a specific server
- Distributed hash table
- Or replica servers

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.69

69

OBJECTIVES – 1/18

- Questions from 1/16
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
  - Chapter 2.1 – Architectural Styles
    - Resource-centered architectures
      - Representational state transfer (REST)
    - Event-based
      - Publish and subscribe (Rich Site Summary RSS feeds)
  - Class Activity: Architectural Styles
  - Chapter 2.2: Middleware Organization
  - Chapter 2.3: System Architectures
    - Centralized system architectures
    - **Decentralized peer-to-peer architectures**
    - Hybrid architectures

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.70

70

TYPES OF SYSTEM ARCHITECTURES

- Centralized system architectures
  - Client-server
  - Multitiered
- Decentralized peer-to-peer architectures
  - Structured
  - Unstructured
  - Hierarchically organized
- Hybrid architectures

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.71

71

DECENTRALIZED PEER-TO-PEER ARCHITECTURES

- Client/server:
  - Nodes have specific roles
- Peer-to-peer:
  - Nodes are seen as **all equal...**
- **How should nodes be organized for communication?**

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.72

72

STRUCTURED PEER-TO-PEER

- Nodes organized using specific **topology** (e.g. ring, binary-tree, grid, etc.)
  - Organization assists in data lookups
- Data indexed using "semantic-free" indexing
  - Key / value storage systems
  - Key used to look-up data
- Nodes store data associated with a subset of keys

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.73

73

DISTRIBUTED HASH TABLE (DHT)

- Distributed hash table (DHT) (ch. 5)
- Hash function
  - $key(data\ item) = hash(data\ item's\ value)$
  - Hash function "generates" a unique key based on the data
  - No two data elements will have the same key (hash)
  - System supports data lookup via key
  - Any** node can receive and resolve the request
  - Lookup function determines which node stores the key
    - $existing\ node = lookup(key)$
- Node forwards request to node with the data

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.74

74

FIXED HYPERCUBE EXAMPLE

- Example where topology helps **route** data lookup request
- Statically sized 4-D hypercube, every node has 4 connectors
- 2 x 3-D cubes, 8 vertices, 12 edges
- Node IDs represented as 4-bit code (0000 to 1111)
- Hash data items to 4-bit key (1 of 16 slots)
- Distance (number of hops) determined by identifying number of varying bits between neighboring nodes and destination

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.75

75

FIXED HYPERCUBE EXAMPLE - 2

- Example: fixed hypercube**  
node 0111 (7) retrieves data from node 1110 (14)
- Node 1110 is not a neighbor to 0111
- Which connector leads to the shortest path?**

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.76

76

WHICH CONNECTOR LEADS TO THE SHORTEST PATH?

- Example:** node 0111 (7) retrieves data from node 1110 (14)
- Node 1110 is not a neighbor to 0111

**[0111] Neighbors:**  
1111 (1 bit different than 1110) 0011 (3 bits different - bad path)  
0110 (1 bit different than 1110) 0101 (3 bits different - bad path)

- Does It matter which node is selected for the first hop?**

January 18, 2024

TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.77

77

DYNAMIC TOPOLOGY

- Fixed hypercube requires static topology
  - Nodes cannot join or leave
- Relies on symmetry of number of nodes
- Can force the DHT to a certain size
- Chord system - DHT (again in ch.5)
  - Dynamic topology
  - Nodes organized in ring
  - Every node has unique ID
  - Each node connected with other nodes (shortcuts)
  - Shortest path between any pair of nodes is ~ order  $O(\log N)$
  - N is the total number of nodes

January 18, 2024


TCSS558: Applied Distributed Computing [Winter 2024]  
School of Engineering and Technology, University of Washington - Tacoma

L5.78

78

## CHORD SYSTEM

- Data items have m-bit key
- Data item is stored at closest "successor" node with ID  $\geq$  key k
- Each node maintains finger table of successor nodes
- Client sends key/value lookup to **any** node
- Node forwards client request to node with m-bit ID closest to, but not greater than key k
- Nodes must continually refresh finger tables by communicating with adjacent nodes to incorporate node joins/departures



January 18, 2024

TCCS558: Applied Distributed Computing [Winter 2024]  
 School of Engineering and Technology, University of Washington - Tacoma

L5.79

79

## UNSTRUCTURED PEER-TO-PEER

- No topology: How do nodes find out about each other?**
- Each node maintains adhoc list of neighbors
- Facilitates nodes frequently joining, leaving, adhoc systems
- Neighbor:** node reachable from another via a network path
- Neighbor lists constantly refreshed
  - Nodes query each other, remove unresponsive neighbors
- Forms a "random graph"
- Predetermining network routes not possible
  - How would you calculate the route algorithmically?
- Routes must be discovered

January 18, 2024

TCCS558: Applied Distributed Computing [Winter 2024]  
 School of Engineering and Technology, University of Washington - Tacoma

L5.80

80

## SEARCHING FOR DATA: UNSTRUCTURED PEER-TO-PEER SYSTEMS

- Flooding**
- [Node u] sends request for data item to **all neighbors**
- [Node v]
  - Searches locally, responds to u (or forwarder) if having data
  - Forwards request to **ALL neighbors**
  - Ignores repeated requests
- Features
  - High network traffic
  - Fast search results by saturating the network with requests
  - Variable # of hops
  - Max number of hops or time-to-live (TTL) often specified
  - Requests can "retry" by gradually increasing TTL/max hops until data is found

January 18, 2024

TCCS558: Applied Distributed Computing [Winter 2024]  
 School of Engineering and Technology, University of Washington - Tacoma

L5.81

81

## SEARCHING FOR DATA - 2

- Random walks**
- [Node u] asks a randomly chosen neighbor [node v]
- If [node v] does not have data, forwards request to a random neighbor
- Features
  - Low network traffic
  - Akin to sequential search
  - Longer search time
  - [node u] can start "n" random walks simultaneously to reduce search time
  - As few as  $n=16..64$  random walks sufficient to reduce search time (LV et al. 2002)
  - Timeout required - need to coordinate stopping network-wide walk when data is found...

January 18, 2024

TCCS558: Applied Distributed Computing [Winter 2024]  
 School of Engineering and Technology, University of Washington - Tacoma

L5.82

82

## SEARCHING FOR DATA - 3

- Policy-based search methods**
- Incorporate history and knowledge about the adhoc network **at the node-level** to enhance effectiveness of queries
- Nodes maintain lists of preferred neighbors which often succeed at resolving queries
- Favor neighbors having highest number of neighbors
  - Can help minimize hops

January 18, 2024

TCCS558: Applied Distributed Computing [Winter 2024]  
 School of Engineering and Technology, University of Washington - Tacoma

L5.83

83

## HIERARCHICAL PEER-TO-PEER NETWORKS

- Problem:** Adhoc system search performance does not scale well as system grows
- Allow nodes to assume **ROLES** to improve search
- Content delivery networks (CDNs) (*video streaming*)
  - Store (cache) data at nodes local to the requester (client)
  - Broker node - tracks resource usage and node availability
    - Track where data is needed
    - Track which nodes have capacity (disk/CPU resources) to host data
- Node roles
  - Super peer** - Broker node, routes client requests to storage nodes
  - Weak peer** - Store data

January 18, 2024

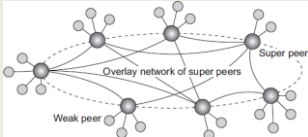
TCCS558: Applied Distributed Computing [Winter 2024]  
 School of Engineering and Technology, University of Washington - Tacoma

L5.84

84

## HIERARCHICAL PEER-TO-PEER NETWORKS - 2

- Super peers
  - Head node of local centralized network
  - Interconnected via overlay network with other super peers
  - May have replicas for fault tolerance
- Weak peers
  - Rely on super peers to find data
- Leader-election problem:
  - Who can become a super peer?
  - What requirements must be met to become a super peer?



The diagram illustrates a hierarchical peer-to-peer network. It shows a central 'Overlay network of super peers' where several 'Super peer' nodes are interconnected. These super peers are then connected to a larger group of 'Weak peer' nodes. The weak peers rely on the super peers to find data.

January 18, 2024	TCSS558: Applied Distributed Computing [Winter 2024] School of Engineering and Technology, University of Washington - Tacoma	LS.85
------------------	---	-------

85

## OBJECTIVES - 1/18

- Questions from 1/16
- Assignment 0: Cloud Computing Infrastructure Tutorial
- Chapter 2: Distributed System Architectures:
  - Chapter 2.1 – Architectural Styles
    - Resource-centered architectures
      - Representational state transfer (REST)
    - Event-based
      - Publish and subscribe (Rich Site Summary RSS feeds)
  - Class Activity: Architectural Styles
  - Chapter 2.2: Middleware Organization
  - Chapter 2.3: System Architectures
    - Centralized system architectures
    - Decentralized peer-to-peer architectures
    - **Hybrid architectures**

January 18, 2024	TCSS558: Applied Distributed Computing [Winter 2024] School of Engineering and Technology, University of Washington - Tacoma	LS.86
------------------	---	-------

86

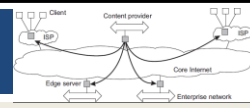
## TYPES OF SYSTEM ARCHITECTURES

- Centralized system architectures
  - Client-server
  - Multitiered
- Decentralized peer-to-peer architectures
  - Structured
  - Unstructured
  - Hierarchically organized
- Hybrid architectures

January 18, 2024	TCSS558: Applied Distributed Computing [Winter 2024] School of Engineering and Technology, University of Washington - Tacoma	LS.87
------------------	---	-------

87

## HYBRID ARCHITECTURES

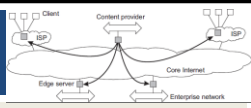


- Combine centralized server concepts with decentralized peer-to-peer models
- **Edge-server systems:**
  - Adhoc peer-to-peer devices connect to the internet through an edge server (origin server)
- Edge servers (provided by an ISP) can optimize content and application distribution by storing assets near the edge
- **Example:**
  - AWS Lambda@Edge: Enables Node.js Lambda Functions to execute "at the edge" harnessing existing CloudFront Content Delivery Network (CDN) servers
  - <https://www.infoq.com/news/2017/07/aws-lambda-at-edge>

January 18, 2024	TCSS558: Applied Distributed Computing [Winter 2024] School of Engineering and Technology, University of Washington - Tacoma	LS.88
------------------	---	-------

88

## HYBRID ARCHITECTURES - 2



- **Fog computing:**
  - Extend the scope of managed resources beyond the cloud to leverage compute and storage capacity of end-user devices
- End-user devices become part of the overall system
- Middleware extended to incorporate managing edge devices as participants in the distributed system
- Cloud → in the sky
  - compute/resource capacity is huge, but far away...
- Fog → (devices) on the ground
  - compute/resource capacity is constrained and local...

January 18, 2024	TCSS558: Applied Distributed Computing [Winter 2024] School of Engineering and Technology, University of Washington - Tacoma	LS.89
------------------	---	-------

89

## COLLABORATIVE DISTRIBUTED SYSTEM EXAMPLE

- **BitTorrent Example:**
  - File sharing system – users must contribute as a file host to be eligible to download file resources
- Original implementation features hybrid architecture
- Leverages idle client network capacity in the background
- User joins the system by interacting with a central server
- Client accesses global directory from a **tracker** server at well known address to access torrent file
- Torrent file tracks nodes having chunks of requested file
- Client begins downloading file chunks and immediately then participates to reserve downloaded content **or network bandwidth is reduced!!**
- Chunks can be downloaded in parallel from distributed nodes

January 18, 2024	TCSS558: Applied Distributed Computing [Winter 2024] School of Engineering and Technology, University of Washington - Tacoma	LS.90
------------------	---	-------

90

