

Carnegie Mellon University Master of Software Engineering

17-647: Engineering Data Intensive Scalable Systems - section D4 (remote) MW 2:00 pm - 3:20 pm ET Course Zoom Link: TBD Meeting ID: TBD Passcode: TBD [A4, Spring 2024, 6 Units]

Instructor	Email	Office location & hours
Paulo Merson	pmerson@andrew.cmu.edu	Zoom ID: 7741902826 By appointment
Teaching assistant		
Ruchi Thosar	rthosar@andrew.cmu.edu	Virtual (zoom link TBD) (Can help with Python Flask or Django, Node+Express)
Nikita Agarwala	nikitaag@andrew.cmu.edu	Virtual (zoom link TBD) (Can help with Node+Express, Python Flask, or FastAPI, Spring Boot)

**Course Description.** Internet services companies such as Google, Amazon, Netflix, and Meta, have pioneered systems that have achieved unprecedented scale while still providing high availability and high cost-performance. These systems differ from mainstream high performance systems in fundamental ways. They are data intensive rather than compute intensive. They are also cloud-native distributed systems that are often deployed as microservices and require integration with external systems via networked-based APIs. They need to inherently support scalability, typically having high throughput, security, reliability and availability demands as well.

Designing and building these systems require a specialized set of skills. **This course will cover design principles and strategies needed to design and implement data intensive scalable distributed systems.** In this space engineers not only need to know how to design, implement, package and deploy systems to multi-platform cloud environments, but to do so in a way that also supports high scalability, availability, reliability, security, and performance. Given the nature of these systems, basic distributed systems concepts such as time, location transparency, data consistency and synchronization are also important. These systems largely operate around the clock, have components running on different nodes and platforms, so the course also emphasizes deployability and operational concerns.

The course includes a hands-on project where students get to create and discuss design alternatives, and implement the solutions in a public cloud environment. Concepts, design strategies, tools, and patterns for distributed systems will be discussed in lectures and applied in the project. The students will gain exposure to the core concepts needed to design and build such systems as well as current technologies in this space. Class size will be limited.

**Prior Knowledge.** Students are expected to be familiar with programming in at least one language suitable for creating backend services, such as node JS, Java, Python, Kotlin, and C#. Formal training in Software Architecture is helpful, but not required. Basic understanding of http communication and distributed systems is expected. In particular, students should have the ability to implement and run a simple http (REST) service. Students in doubt regarding their experience should obtain instructor's permission.

## **Learning Objectives**

Students will be introduced to concepts, design principles and patterns for distributed systems in general, with a focus on cloud-native http services, asynchronous event-driven communication, and SQL and NoSQL data stores. Students will be expected to demonstrate the application of this knowledge through the construction of systems with increasing expectations.

Students in this class will learn what it takes to engineer systemic properties into data intensive distributed systems. Namely, we will focus on interoperability, scalability, availability and reliability, loose coupling, modifiability, and distributed data. Students will learn about the interplay among these concerns when applying several design patterns and technologies. Managing tradeoffs appropriately is a fundamental aspect of software design and will be a key concern in the lectures and in different stages of the project.

**Learning Resources.** The course and all course materials will be distributed online and accessible with a CMU account via Canvas. We will also use the AWS cloud via <u>AWS Academy</u>.

**Use of Zoom in the Class.** In our class, we will be using Zoom. The link is available on Canvas. Please make sure that your Internet connection and equipment are set up to use Zoom and you are able to share audio and video during class meetings. (See <u>this page</u> for Computing Resources for information on the technology you are likely to need.) Let me know if there is a gap in your technology set-up (pmerson@cmu.edu) as soon as possible, and we can see about finding solutions.

*Sharing Video*: In this course, being able to see one another helps to facilitate a better learning environment and promote more engaging discussions. Therefore, our default will be to expect students to have their cameras on during lectures and discussions. However, I also completely understand there may be reasons students would not want to have their cameras on. If you have any concerns about sharing your video, please email me as soon as possible (pmerson@cmu.edu) and we can discuss possible adjustments. Note: You may use a background image in your video if you wish; just check in advance that this works with your device(s) and internet bandwidth.

*Technical Difficulties*: From time to time we all experience unstable internet connections, unstable computers, etc. In those cases, you may find it necessary to turn your camera off. If you experience technical difficulties during class, please let me know via private chat in Zoom prior to turning your

camera off. If technical difficulties are a recurring issue, please reach out to your <u>HUB liaison</u> who will help you access the appropriate resources.

During our class meetings, please keep your mic muted unless you are sharing with the class or your breakout group. If you have a question or want to answer a question, please unmute your mic and speak up. Alternatively, you may use the chat or the "raise hand" feature (available when the participant list is pulled up).

# Assessments

Students learn more by applying and explaining ideas to others, thus, the course requires the following activities:

- Lectures and reading assignments: students are expected to attend the classes and read all texts indicated in the weekly reading assignments.
- **Class participation:** Students are stimulated to contribute questions, answers and comments during class, and participate in online discussions in Piazza.
- **Project assignments**: Students will work individually on a software project with incremental deliverables to put in practice concepts and technologies discussed in class. We'll use <u>AWS</u> <u>Academy</u> for the project.
- Weekly quizzes: These are short online questions derived from the reading assignments and lectures, to be answered individually. All quizzes will count towards your grade, that is, we are NOT going to drop the lowest grade.

Assessment	Final Grade %	
Weekly quizzes	30%	
Project assignments	60%	
Class participation	10%	

## **Course and Grading Policies**

- Late-work policy: All work is expected to be handed in at the indicated due date and time. For fairness to the whole class, no late submissions will be accepted for the weekly quizzes. The penalty for late submission of project assignments is 10% per day, capped at a maximum penalty of 30%.
- Participation policy. Class participation will be graded by in-class engagement, attendance and punctuality, asking relevant questions based on a critical review of required readings and lectures, and contributions (questions and answers) on Piazza. The lack of attendance, and the use of mobile devices including phones, tablets, and laptops for purposes other than participating in class, will count against your participation grade. Remote learning requires the use of technology. Research has shown that divided attention is detrimental to learning; I encourage you to close any windows not directly related to what we are doing while you are in class. Please turn off your phone

notifications and limit other likely sources of technology disruption, so that you can fully engage with the material, each other, and me. This will create a better learning environment for everyone.

Attendance. Within the first week of our course, please look ahead and determine if you need to miss class for any excusable reason religious observance, job interview, university-sanctioned event, etc.) and notify me as soon as possible. You will be expected to attend all class sessions (unless otherwise discussed with the instructor); the instructor or TA will record attendance. Additionally, you will be expected to participate fully in all in-class discussions, exercises, and case studies. Make meaningful contributions when and where you can. Please note that I expect that you will abide by all behaviors indicated in <u>The Word</u>.

**Recording of Class Sessions.** All synchronous classes will be recorded via Zoom so that students in this course (and only students in the course) can watch or re-watch past class sessions. Recordings become available on Canvas usually a few hours after each class session. Recordings will live on our Canvas website. Please note that you are not allowed to share these recordings. This is to protect your FERPA rights and those of your fellow students.

## **Course Schedule**

The following schedule provides a general overview of topics and assignments.

Week	Торісѕ	Assignments	Notes
1 Mar 9-15	<ul> <li>Course introduction</li> <li>Distributed systems and SOA</li> <li>The Microservice architecture style</li> <li>Microservice in practice</li> <li>Microservice benefits and challenges</li> <li>Microservice security</li> </ul>	<ul> <li>Week 1 reading assignment</li> <li>Week 1 quiz</li> <li>A1 given (dockerized REST service on AWS)</li> </ul>	
2 Mar 16-22	<ul> <li>Types of components and connectors in the Microservice style         <ul> <li>REST vs component technology (gRPC, GraphQL, etc.)</li> <li>REST API design</li> </ul> </li> <li>Asynchronous messaging (producers, consumers, queue/topics)</li> <li>Core principles for microservice design:         <ul> <li>"IDEALS"</li> </ul> </li> </ul>	<ul> <li>Week 2 reading assignment</li> <li>Week 2 quiz</li> <li>A2 given (BFF, SRP, JWT)</li> </ul>	A1 due on Fri
3 Mar 23-29	<ul> <li>Service interceptors, API Gateway, BFF, service mesh</li> <li>Availability strategies         <ul> <li>Monitoring and exception tracking</li> <li>Timeout and Retry</li> <li>Circuit breaker</li> </ul> </li> </ul>	<ul> <li>Week 3 reading assignment</li> <li>Week 3 quiz</li> </ul>	

	<ul> <li>Bulkheads</li> </ul>		
4 Mar 30 - Apr 5	<ul> <li>Messaging patterns used in microservice architectures         <ul> <li>Building blocks: point-to-point, pub-sub</li> <li>For reliability: store-and-forward, transactional outbox, dead-letter channel</li> <li>Other: sync-over-async, push vs pull</li> </ul> </li> <li>Event-Driven Architecture (EDA) and its tradeoffs         <ul> <li>Choreography and orchestration</li> <li>Asynchronous without a message broker</li> <li>NoSQL data stores</li> <li>CAP theorem</li> </ul> </li> </ul>	<ul> <li>Week 4 reading assignment</li> <li>Week 4 quiz</li> <li>A3 Given (K8S, Kafka on AWS MSK, circuit breaker)</li> </ul>	A2 due on Tue
5 Apr 6-12	<ul> <li>Service loose coupling         <ul> <li>Wrapper pattern for services</li> <li>Database per Microservice pattern</li> </ul> </li> <li>Service autonomy and distributed transactions</li> <li>Intro to DDD</li> </ul>	<ul> <li>Week 5 reading assignment</li> <li>Week 5 quiz</li> </ul>	
6 Apr 13-19	<ul> <li>From DDD to microservices</li> <li>Saga pattern (compensation)</li> <li>Data replication and eventual consistency</li> <li>CQRS</li> <li>Event sourcing</li> </ul>	<ul> <li>Week 6 reading assignment</li> <li>Week 6 quiz</li> <li>A4 Given as OPTIONAL (NoSQL, CQRS, data replication with ETL)</li> </ul>	A3 due on Fri
7 Apr 20-26	<ul> <li>(Other) deployability strategies         <ul> <li>Blue-green and Canary</li> <li>IaC, externalized configuration</li> <li>Log consolidation, monitoring, tracing</li> <li>Serverless</li> <li>Continuous delivery</li> </ul> </li> <li>Moving from a monolithic to a microservice architecture (Strangler Fig pattern)</li> </ul>	<ul> <li>Week 7 reading assignment</li> <li>Week 7 quiz</li> </ul>	
8 Apr 27 - May 3	no classes		A4 (optional) due on Tue

### References

Distributed Systems and Software Architecture references

Accommodations for Students Disabilities. If you have a disability and have an accommodations letter from the Disability Resources office, I encourage you to discuss your accommodations and needs with me as early in the semester as possible. I will work with you to ensure that accommodations are provided as appropriate. If you suspect that you may have a disability and would benefit from accommodations but are not yet registered with the Office of Disability Resources, I encourage you to contact them at access@andrew.cmu.edu.

Academic Integrity. Honesty and transparency are important to good scholarship. Plagiarism and cheating, however, are serious academic offenses with serious consequences. If you are discovered engaging in either behavior in this course, you will earn a failing grade on the assignment in question, and further disciplinary action may be taken.

For a clear description of what counts as plagiarism, cheating, and/or the use of unauthorized sources, please see the <u>University's Policy on Academic Integrity</u>.

If you have any questions regarding plagiarism or cheating, please ask me as soon as possible to avoid any misunderstandings. For more information about Carnegie Mellon's standards with respect to academic integrity, you can also check out the <u>Office of Community Standards & Integrity</u> website.

**Generative Artificial Intelligence (AI) Tools and Academic Integrity.** Certain assignments in this course will permit or even encourage the use of generative artificial intelligence (AI) tools, such as ChatGPT. When AI use is permissible, it will be clearly stated in the assignment. Otherwise, the default is that use of generative AI is disallowed. It is each student's responsibility to assess the validity and applicability of any AI output that is submitted. You may not earn full credit if inaccurate on invalid information is found in your work. Deviations from the guidelines above will be considered violations of CMU's academic integrity policy. Note that expectations for "plagiarism, cheating, and acceptable assistance" on student work may vary across your courses and instructors. Please email me if you have questions regarding what is permissible and not for a particular graded activity.

**Student Well-Being.** The last few years have been challenging. We are all under a lot of stress and uncertainty at this time. I encourage you to find ways to move regularly, eat well, and reach out to your support system or me (<u>pmerson@cmu.edu</u>) if you need to. We can all benefit from support in times of stress, and this semester is no exception.

As a student, you may experience a range of challenges that can interfere with learning, such as strained relationships, increased anxiety, substance use, feeling down, difficulty concentrating and/or lack of motivation. These mental health concerns or stressful events may diminish your academic performance and/or reduce your ability to participate in daily activities. CMU services are available, and treatment does work. You can learn more about confidential mental health services available on campus at the <u>Counseling and Psychological Services</u> website. Support is always available (24/7) from Counseling and Psychological Services: 412-268-2922.

If you are worried about affording food or feeling insecure about food, there are resources on campus who can help. Email (<u>cmu-pantry@andrew.cmu.edu</u>) or call (412-268-8704) the CMU Food Pantry Coordinator to schedule an appointment.

We must treat every individual with respect. We are diverse in many ways, and this diversity is fundamental to building and maintaining an equitable and inclusive campus community. Diversity can refer to multiple ways that we identify ourselves, including but not limited to race, color, national origin, language, sex, disability, age, sexual orientation, gender identity, religion, creed, ancestry, belief, veteran status, or genetic information. Each of these diverse identities, along with many others not mentioned here, shape the perspectives our students, faculty, and staff bring to our campus. We, at CMU, will work to promote diversity, equity, and inclusion not only because diversity fuels excellence and innovation, but because we want to pursue justice. We acknowledge our imperfections while we also fully commit to the work, inside and outside of our classrooms, of building and sustaining a campus community that increasingly embraces these core values.

Each of us is responsible for creating a safer, more inclusive environment.

Unfortunately, incidents of bias or discrimination do occur, whether intentional or unintentional. They contribute to creating an unwelcoming environment for individuals and groups at the university. Therefore, the university encourages anyone who experiences or observes unfair or hostile treatment on the basis of identity to speak out for justice and support, within the moment of the incident or after the incident has passed. Anyone can share these experiences using the following resources:

- Center for Student Diversity and Inclusion: csdi@andrew.cmu.edu, (412) 268-2150
- Ethics Reporting Hotline. Students, faculty, and staff can anonymously file a report by calling 844-587-0793 or visiting cmu.ethicspoint.com.

All reports will be documented and deliberated to determine if there should be any following actions. Regardless of incident type, the university will use all shared experiences to transform our campus climate to be more equitable and just.