NAMIBIA UNIVERSITY
OF SCIENCE AND TECHNOLOGY

FACULTY OF COMPUTING AND INFORMATICS

DEPARTMENT OF SOFTWARE ENGINEERING

QUALIFICATION: BACHELOR OF COMPUTER SCIENCE	
QUALIFICATION CODE: 07BCMS, 07BAIT	**LEVEL:** 6
COURSE: DISTRIBUTED SYSTEMS AND APPLICATIONS	**COURSE CODE**: DSA612S
DATE: DECEMBER 2025	**PAPER:** THEORY
DURATION: 3 HRS	**MARKS:** 80

SUPPLEMENTARY/SECOND OPPORTUNITY EXAMINATION QUESTION PAPER	
EXAMINER(S)	**Ms. NDINELAGO NASHANDI**
MODERATOR:	**PROF JOSE QUENUM**

THIS QUESTION PAPER CONSISTS OF 6 PAGES
(Including this front page)

INSTRUCTIONS TO STUDENTS:

1. Read all the questions, passages, scenarios, etc., carefully before answering.
2. **Answer all the questions.**
3. Number each answer clearly and correctly.
4. Write neatly and legibly.
5. Making use of any crib notes may lead to disqualification and disciplinary action.
6. Use the allocated marks as a guideline when answering questions.
7. Looking at other students' work is strictly prohibited.
8. This paper consists of three **(3)** pages, including the cover page.
9. The paper consists of 5 sections: A, B, C, D, E and F.

SECTION A: 5 marks

Multiple Choice Questions (WRITE the correct letter)

1. What is the main objective of Service Engineering?
    A. Developing standalone applications that require minimal external services.
    B. Creating reusable and dependable services that can be integrated into multiple applications.
    C. Ensuring that all services are tightly coupled within a single system.
    D. Eliminating the need for external APIs in software development.

2. Which of the following statements about the Distributed File System characteristics is not true?
    A. It supports location transparency.
    B. It allows the automatic migration of files between machines.
    C. Files are permanently associated with specific disk blocks.
    D. It provides better fault tolerance through data replication.

3. Which of the following best describes the purpose of logical clocks in distributed systems?
    A. To synchronise all system clocks to a universal time reference.
    B. To establish a total ordering of events across all distributed processes.
    C. To ensure causality by allowing a partial ordering of events in a distributed system.
    D. To measure the time elapsed between any two events in a distributed system.

4. A subscription in an Event Notification Service is a _____ that determines which events a subscriber is interested in.
    A. log entry
    B. constraint
    C. message queue
    D. global event broadcast

5. How does Google File System achieve fault tolerance?
    A. By using parity checks to verify the correctness of stored and transmitted data
    B. By maintaining multiple replicas of data chunks across different servers for reliability
    C. By compressing data to reduce storage overhead and improve transmission efficiency
    D. By encrypting data to ensure security and confidentiality during storage and transfer

SECTION B: 14 marks

- *Answer all the questions in the provided booklet.*
- *The section consists of 3 questions.*

1. Define the following terms in distributed systems [4 marks]
    A. Distributed system
    B. Distributed file system
2. Explain three (3) significant characteristics of a distributed system using an example of a distributed web server. [6 marks]
3. State and explain any two (2) problems facing designers of distributed systems [4 marks]

SECTION C: 12 marks

A software company is building a distributed application where clients must communicate efficiently with backend services deployed across multiple servers. The development team has chosen to use gRPC to implement the communication between clients and servers.

1. Discuss Remote Procedure Call (RPC) in distributed systems? [6 marks]
2. Using the above scenario, outline the steps involved in implementing a gRPC-based system. [6 marks]

SECTION D: 14 marks

- *Answer all the questions in the provided booklet.*
- *The section consists of 4 questions.*

Context. An e-commerce platform uses a distributed publish/subscribe overlay (see attached network diagram). Events originate at b0. Each broker forwards an event only if the predicate on the outgoing link is satisfied (see the routing table below). Subscribers apply their own filters at the broker where they are attached.

1. Trace the full broker path for m1 and m3 [4 marks]

   m1 = {category: "electronics", amount: 2500}

   m2 = {category: "toys", amount: 800}

2. b) Identify which subscribers (if any) receive each message. [4 marks]

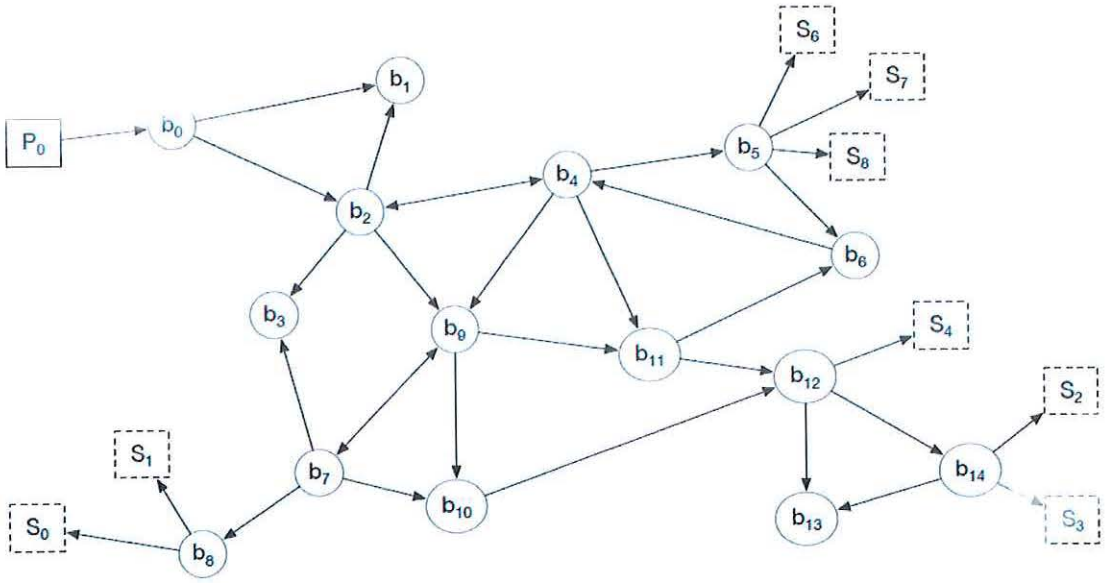

From	Go to	When
b0	b2	Any
b2	b4	category = "elect*" AND amount > 2000
b2	b9	category = "furniture"
b4	b5	category = "electronics" AND amount > 2000
b9	b11	priority = "high"
b11	b12	amount ≥ 2500
b12	b13	Any
B13	b14	category = "furniture"

Subscriptions

- **S7 @ b5:** category = "electronics" AND amount > 2000
- **S4 @ b12:** category = "furniture" AND priority = "high" AND amount ≥ 2500
- **S2 @ b14:** category = "furniture"

3. Using this context, differentiate between Flooding-based and Filter-based routing models in Publish–Subscribe (Pub-Sub) systems. In your answer, explain how each routing model would operate in such a network. [4 marks]

4. Discuss the limitations of each model. [2 marks]

SECTION E: 20 marks

- *Answer all the questions in the provided booklet.*
- *The section consists of 3 questions.*

1. Discuss the Remote Service Model and the Data Caching Model as methods of accessing files in Distributed File Systems (DFS).

   In your answer, explain how each model works, outline their benefits and limitations. [6 marks].

2. Identify one major challenge of caching in distributed file systems and explain how caching update policies can help address this challenge. [5 marks]
3. Provide and explain the architecture of the Hadoop File System. (9 marks)

SECTION F: 15 marks

- *Answer all the questions in the provided booklet.*
- *The section consists of 3 questions.*

1. Explain briefly the difference between logical and physical clocks. [4 marks]
2. Describe the Christians' algorithm for achieving synchronisation in clocks within Distributed Systems. [7 marks]
3. A client's clock reads 3:20:00. The server's clock reads 3:10:00 when they synchronise using the Berkeley algorithm. Assume message delays are negligible. What is the time at the client after synchronisation? [4 marks]

********** End of the Paper **********