

1. The tasks described in this worksheet are part of the formative assessment. They serve the purpose to prepare you for the examination. We will discuss the solutions during the next **interactive session** after they are handed out – while they fit to the lecture of the week they are handed out, they might be discussed in two weeks time due to the bi-weekly exercise schedule.
2. Make sure to plan your time for the whole sheet carefully. The complete exercise should represent approximately three hours of independent study. The time limit indicates how much time you should spend on each task, and not how much time you may actually need; it is important that you engage with the material and not that you complete all tasks perfectly. Feel free to collaborate and team up.
3. The exercises are designed to challenge you and train you further as guided self-study. The time limit might be too ambitious for you; you may team up with colleagues. It is not an issue as long as you manage to at least partially resolve each task within the time budget. If you (and team) are struggling, reach out for help in Teams! You may also share your thoughts on the channel.
4. We recommend that you create a (private) GIT repository where you store your findings and outcomes while processing the exercises. This portfolio of work could be useful in the future.

## Contents

<b>Task 1: Introduction: Understanding Architectures and Performance (60 min)</b>	<b>1</b>
<b>Task 2: Data Schemas for Wikipedia Articles (90 min)</b>	<b>2</b>
<b>Task 3: Word Frequencies in Moby-Dick (90 min)</b>	<b>3</b>
3.1 Visualization of Individual Word Frequencies . . . . .	3
3.2 Aggregation of Word Occurrences . . . . .	3
3.3 Most Common Words . . . . .	3
3.4 Hints and Code-Skeleton . . . . .	3

## Task 1: Introduction: Understanding Architectures and Performance (60 min)

Performance is an important aspect of **High-Performance** Data Analytics (HPDA) and of parallel and distributed systems. Imagine you are 20% more efficiently using server infrastructure than your competition - and you have 1000s of servers. That is a business advantage...

In this task, you will research relevant performance characteristics of a single computer system and a distributed system using literature.

### Tasks

1. Research what the term “performance” could mean for an HPDA application. Note that there is not the single performance characteristics but multiple characteristics that are meaningful. Document why this use case is high-performance. Your notes should cover about 1/4-page.

2. Research alternatives for “supercomputer architecture” in the literature, e.g., using Google. Research two architectures and select diagrams that describe them; the architecture could be a specific supercomputer. Your notes should cover about 1/4-page.
3. Research what the term “performance” could mean for a single computer system. Document relevant characteristics single computer system, i.e., a server or desktop system. Find typical (performance) values for a decent (current) computer system in the document. Your notes should cover about 1/4-page.
4. Think about performance of multiple computers that are interconnected to a distributed system. How would you expand upon your performance characteristics for a single computer system? Look at the TOP500 list<sup>1</sup> and document which performance characteristics are provided in the list.

## Portfolio (directory: 2/performance)

<code>2/performance/architectures/</code>	Store supercomputer architecture documents here.
<code>2/performance/performance-hpda.txt</code>	Performance characteristics of the HPDA application.
<code>2/performance/performance-1node.txt</code>	Performance characteristics of a single node.
<code>2/performance/performance-distributed.txt</code>	Performance characteristics of a distributed system and performance metrics used in the TOP500.

## Task 2: Data Schemas for Wikipedia Articles (90 min)

In this task, we will create various data schema for typical Wikipedia data.

Firstly, we need to consider the data we want to store. Each article should store the following properties:

- Title
- Text
- Category
- Links to related articles

The following operations should be possible:

- Access article details based on the article’s “title”
- Finding related articles (those that link to one another) from a given article
- Retrieving all articles for one category

You can view some Wikipedia articles to understand this better.

Please document how you would create the following data models:

- Relational
- Columnar
- Key-Value
- Document
- Graph

Think about what a fact-based model implementation could add to any of these data models.

---

<sup>1</sup><https://www.top500.org>

## Portfolio (directory: 2/data-model)

2/data-model/data-model.txt	Store a description of the respective data model, describe how the operations can be executed, and give at least one example.
2/data-model/fact-based-discussion.txt	Discuss here what the fact-based data model would add.

## Task 3: Word Frequencies in Moby-Dick (90 min)

In this task, you will write a Python program to compute word frequencies for a text. We will work with the book "Moby-Dick" by Herman Melville.

You can find it as a text file `moby-dick.txt` in Google.

### 3.1 Visualization of Individual Word Frequencies

Your program should take a word as input via command line argument, then split the book into its chapters and for each chapter calculate the relative frequency of the word given by the command line.

Your program will compute the occurrences of any given term across the books chapters. Visualize the frequency in a graph (see the hints).

### 3.2 Aggregation of Word Occurrences

Now we want to output the absolute occurrences of *each* word in each chapter. Program output should be a CSV file with the following header:

chapter-name, chapter-length, word-occurrences

The word-occurrences should be stored as JSON.

### 3.3 Most Common Words

Write a program to print the 10 most common words in the whole book.

### 3.4 Hints and Code-Skeleton

Visualization of the word frequency can be done modifying the code provided below:

```
1 #!/usr/bin/env python3
2 import matplotlib.pyplot as plt
3
4 if __name__ == "__main__":
5     numbers = [1,22,33,5,44] # Save your samples here
6     plt.bar(range(len(numbers)), numbers, 0.35)
7     plt.ylabel("selected word")
8     plt.xlabel("occurrence over all chapters")
9     plt.show()
```

For this task the module `matplotlib` is required, it is already installed on the cluster. In case you want to test your work locally, install the package via the command `pip` or your distribution's package manager.

You may use dictionaries and the `json` module.

---

## Portfolio (directory: 2/word-freq)

<code>2/word-freq/graph.py</code>	Your source code to visualize a word's occurrence.
<code>2/word-freq/csv.py</code>	Your source code to extract the word occurrences.
<code>2/word-freq/most-common.py</code>	Your source code to find the most common words.