



# Storage Systems (StoSys) Course Project Handbook

(XM\_0092)

P1 (September - October), 2023

Version: 4.0

**Animesh Trivedi** ([a.trivedi@vu.nl](mailto:a.trivedi@vu.nl))

**Matthijs Jansen** ([m.s.jansen@vu.nl](mailto:m.s.jansen@vu.nl)), **Zebin Ren** ([z.ren@vu.nl](mailto:z.ren@vu.nl)),  
**Krijn Doekemeijer** ([k.doekemeijer@vu.nl](mailto:k.doekemeijer@vu.nl)), **Nick Tehrany** ([n.a.tehrany@vu.nl](mailto:n.a.tehrany@vu.nl)) and  
**Sudarsan Sivakumar** ([s.sivakumar@student.vu.nl](mailto:s.sivakumar@student.vu.nl))

## Storage-Networking Research Group

<https://stonet-research.github.io/>



Course mascot (courtesy: Gilles Magalhaes via DALL-E)

<b>Storage Systems (StoSys) Course Project Handbook</b>	<b>1</b>
<b>Revision Information</b>	<b>3</b>
<b>Code of Conduct (Plagiarism)</b>	<b>4</b>
<b>1. Overview</b>	<b>5</b>
<b>2. Logistics</b>	<b>5</b>
2.1 What is given to you	6
2.2 How to ask for help when stuck in the project	7
2.3 Late/Slip days	7
2.4 Group Management Logistics	8
<b>3. [Read-carefully] Coding Convention and the Assessment Process</b>	<b>8</b>
3.1 Important Coding Conventions	8
3.2 The grading process and automated scoring system	9
3.3 Grading guidelines during the Interview	10
3.4 Automatic Grading Flow	11
3.5 Notes on partial assessment for Milestones	12
3.6 What are the test parameters for automated testing	13
<b>4. StoSys Project Work</b>	<b>14</b>
4.1 Milestone 1 (week 1) : A new device is in town	15
4.1.1 Task-1 setup QEMU 6.1	16
4.1.2 Prepare the Ubuntu 20 in the QEMU VM	17
4.1.3 Task-3 Familiarize yourself with online resources	20
4.1.4 Task-4 Build and run nvme cli command tool	21
4.1.5 Task-4 Build and prepare the coding environment	24
4.1.6 Task-5 Finish the missing code pieces and take canvas quiz	26
4.1.7 All together	28
4.2 Milestone 2 (week 2) : I can't read, is there a translator here?	29
4.3 Milestone 3 (week 3) : It's 2023, we recycle	32
4.4 Milestone 4 (week 5) : We love Rock(sDB) 'n' Roll!	35
4.5 Milestone 5 (week 7) : Wake up, Neo	42
<b>5. Bonus Ideas and Timeline</b>	<b>44</b>
<b>Appendix</b>	<b>46</b>
A. Step by step guide for setting up Ubuntu installation	46
B. Setup of VNC through Remmina	47
C. Running inside Virtualbox, VirtualMachineManager or VMWare	48

## Revision Information

- August 24th, 2023: A few memory leaks in libnvme fixed thanks to Boyuan Xiao and Krijn Doekemeijer.
- September 7th, 2022: Corrections from Paulo Liedtke on ZNS appends commands and QEMU filename typos.
- November 1st, 2021: version 2.0 is released with NVMe ZNS devices and RocksDB. OC-SSD based experiments are removed.
- October 25th, 2020: version 1.1 with some corrections from Corné Lukken
- October 25th, 2020: version 1.0

## Code of Conduct (Plagiarism)

We have a **ZERO-TOLERANCE** cheating, fraud, and plagiarism policy.

You are encouraged to discuss specific issues with other students, search online, read Linux/Storage code, and share knowledge with other fellow students. After all, the whole field of computer science is built with open-collaboration. ***HOWEVER, in any circumstances you are expected to write your own code. Do not buy or sell code online. Do not copy code. Do not give out code to other students.*** It is quite unlikely that a project of this scope will result in two students/groups writing almost identical code. Any similarities in the code, or style will lead to further investigation, with a possibility of reporting to the examination board - an outcome we both would like to avoid.

What is **not allowed** (including but not limited to):

- 1. Copying code from yesteryear' project codes**
- 2. Sharing code or solution with each other**
- 3. ChatGPT/GitHub Co-Pilot driven development**

Any incident involving a hint of cheating will be passed on to the Examination Board. We do not and will not deal with individual cases on the course level. Allegation of plagiarism is a serious incident that is dealt with at the department and university level. No grades will be assigned to you if you are being investigated for fraud.

Read and understand:

- <https://vu.nl/en/student/examinations/academic-integrity>
- Generative AI and ChatGPT, <https://vu.nl/en/student/examinations/generative-ai-your-use-our-expectations>

## 1. Overview

These are very exciting times to do storage and systems research and the availability of many flexible, easy-to-modify storage stacks gives you an opportunity to think about how a modern, high-performance, storage stack should look like. *Why now?* For the last 50-60 years we have been happy with the hard disk drives (HDDs), with the block layer and associated file system inside the Linux kernel. However, this basic design was constrained by the performance of the rotating magnetic HDDs. With the advent of modern non-volatile memory (NVM) storage technology the storage game has changed completely. These new storage devices (such as NAND flash Solid State Drives, SSDs - a type of NVM storage) offer almost 2-3 orders of magnitude performance gains. At the same time these SSD devices internally are complex, thus offering a variety of tradeoffs in performance, energy, cost, and reliability. There is no one storage stack design that can fit all. In this project, we will build one from scratch without being hindered by the legacy block and file interfaces.

The objective of this course project is to develop a working storage stack in user space. More specifically you will need to

1. Get familiar with Zone Namespace (ZNS) devices and their capabilities
2. Implement a Flash Translation Layer (FTL) with garbage collection (GC)
3. Implement an application-specific file system on top of the FTL and integrate it into a production-quality KV store, RocksDB
4. Make the device FTL, and file system persistent

These concepts are covered in the due time in the lecture. Hence, the idea here is to give you conceptual as well as practical knowledge regarding how an SSD works with an state-of-the-art flash SSD interface, i.e., ZNS. ZNS specification is ratified and merged in the NVMe specification 2.0 in August, 2021.

Furthermore, building such storage stacks in the user-space is not a “hack” just to demonstrate something. As Moore’s Law is slowing down, user-space specialization of storage stacks is one of the most promising and popular ways of delivering performance to modern data-intensive applications such as BigData processing, graph computation, Bioinformatics applications, etc. There are many thesis topics and ideas around the “storage specialization” area of research.

We also have the real ZNS hardware (donated by Western Digital) and it is one of the bonuses to demonstrate that the code you have developed can run on the real hardware. Mind that there are some subtle differences between how QEMU emulates ZNS devices and how the real hardware behaves. For more interesting details, see

- Krijn Doekemeijer, Nick Tehrani, Balakrishnan Chandrasekaran, Matias Bjørling, Animesh Trivedi (2023), Performance Characterization of NVMe Flash Devices with Zoned Namespaces (ZNS), In the IEEE International Conference on Cluster Computing (CLUSTER), 2023.
  - <https://stonet-research.github.io/publications/publications.html>

## 2. Logistics

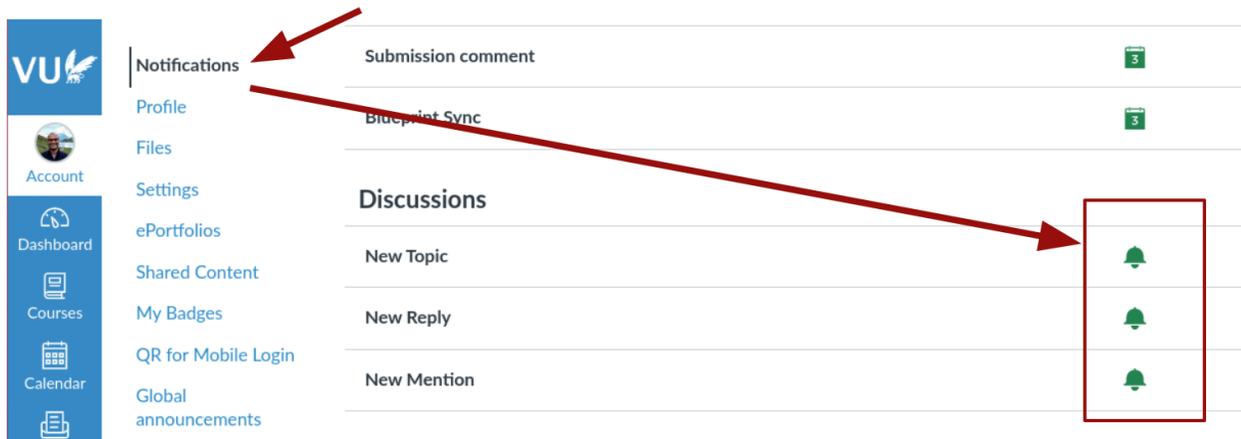
The project is the **team project** and consists of **5 milestones**. The first milestone is an individual milestone. This way you can assess your and your teammate’s capabilities before making teams during the first week. The next 4 milestones are group milestones, to be done in a group of **2 students**.

**When are the deadlines:** check the Canvas page to see when milestones are due.

Some pertinent and practical issues to consider:

- Be aware that you need to form a group in the first week. If you do not know anyone or have not managed to do so, then post it on the Discussion board of the Canvas page.
- You can choose to do the milestones alone, but we do not recommend that. You will not get any lenient marking from us just because you choose to do it alone.

**Important:** please enable notifications (<https://canvas.vu.nl/profile/communication>, you can also do course specific setting) for the Discussion forum as we will be posting many clarifications and information there. In any case do not forget to repeatedly check information there. See this image to enable notifications:



**Weekly lab hours:** There are weekly lab sessions. Please check the Canvas / Rooster. The lab hours are open to any sort of discussion including but not limited to questions about the project (we assume that would be the majority of the questions), clarification about lecture topics, general questions about research, etc.

## 2.1 What is given to you

We have provided a basic framework in which you will develop the project. This framework is a CMake project with example programs that should run at the end of each project milestone. Please make yourself familiar with CMake/C/C++ (and if needed gdb, and friends). We assume that you are familiar with the basic executable concepts in Linux (shared library, compilation, running and installing packages and programs from the source).

In the first milestone you are supposed to prepare a development environment, read the specification, and get yourself familiar with the project environment. This is a relatively simple milestone, and we would **strongly recommend** making use of this time to familiarize yourself with the concepts and source code, and the example programs.

Make sure to check out the due dates on Canvas for projects and ensure that you fulfill all the evaluation criterias. The whole framework and this complete handbook is released to you on day 1, and feel free to code ahead and finish as much as possible. On the demo day you can still just demo us the required

functionality while being ahead in the project time. You can use the additional time for bonuses. See the bonus section later in this handbook.

**Note:** *we do not have any support to do this project on **Windows** or **Mac** or **Windows' Linux subsystem**. So please make sure that you have a working setup in the first week.*

## 2.2 How to ask for help when stuck in the project

Building any system is a complex task. The project that you build uses (or will use) many advanced concepts. So here is the protocol if you get stuck with issues

1. **Read the source code:** Everything given to you is a part of an open-source project and is open and accessible. Hence, the majority of issues or confusions could be clarified by just looking into the source code. We expect that you will spend time reading the framework given to you to understand its semantics. Such setup is very common in day-to-day life as a software developer or researcher where you need to quickly read and understand someone else's code.
2. **Ask Google:** Try to understand if the issue is related to the coding/Linux/C/C++/CMake/Bash environment and look for it appropriately. Often errors which might look strange can simply be resolved by an appropriate google search. There are a very few project specific things that you will not find while googling.
3. **Ask in the Discussion forum:** if the first two steps do not yield any results, then you can ask in the discussion forum. However, you should not ask for a direct solution to the milestones. Try to ask specific issues. We or other fellow students will jump in to answer your question. However, be prepared to hear that "debugging such issues is part of your coding milestones". Building any real system is a complex job, and you need to put in time to solve certain issues. There is no way around it.
4. **Ask in the lab hour:** If the issue is not time critical you can ask in the office hour and we can help you out there. You can share your screen and show us what went wrong with your code.
5. **Reach out to us:** If you believe that you are really stuck and there is something fundamentally wrong (hey, it can happen, no one is perfect ;) ) please reach out to us. However, make sure to include what you have tried to do to solve the problem, what possibilities you have investigated and what you suspect is wrong. Try to include as much information as possible. Again: be prepared to hear "*debugging such issues is part of your coding milestones*".

**Important:** We cannot guarantee any help hours before the deadline. So, please plan ahead, and ask with plenty of due time ahead. You only have one weekly lab session, try to make the most out of it.

## 2.3 Late/Slip days

In the course you have **3 slip days** that you can use to extend the deadline of code/quiz submissions (milestones 1, 2, and 4). There will be no extensions after that, and you will be awarded zero points.

Unfortunately you **cannot** use these days to move the interview dates. They are fixed, and you need to show whatever you have on that day.

## 2.4 Group Management Logistics

Please make sure that you regularly talk and coordinate with your group members and keep in touch. Meet early and plan ahead who does what and when to synchronize. If you are having troubles then contact us early enough. Too many times we have seen students showing up a day before the deadline saying that the team member decided not to work, or dropped out of the course - we cannot help a day before the deadline. We will consider exceptional circumstances like COVID related interrupts, however, beyond that there is very limited flexibility.

**We will not be responsible for resolving disputes in your group.** If you cannot do experiments as a group of 2, then do it alone. Every year students do it alone and it is completely doable. But you will not get a concession from us in grading for that.

## 3. [Read-carefully] Coding Convention and the Assessment Process

Each milestone is worth some points, and in total you can earn 50 points for the complete project. You can earn additional marks from bonuses, but you are restricted to **2 bonuses max** with a maximum of **additional 20 points** (on top of the maximum 50 points from the project). Hence, the total points achievable from the project work is 70.

*The deadline for getting a bonus greenlight is milestone 3 interview (IMPORTANT). NO EXCEPTIONS. See section 5.*

Here is the evaluation plan for the milestones

- **Milestone 1** will be evaluated as a graded Canvas quiz, plus code upload.
- **Milestone 2** is a code upload with the example target programming running successfully.
- **Milestone 3** will be evaluated with a group interview. All members of the group need to be presented in the interview session to answer questions.
- **Milestone 4** is a code upload with the example target programming running successfully.
- **Milestone 5** will be evaluated with a group interview. All members of the group need to be presented in the interview session to answer questions.

### 3.1 Important Coding Conventions



#### Important Coding Conventions

1. We will use GitHub for development. We will provide you the repo where you will do the code development from milestone 2 onward. There is an assignment on week-1 where you need to

tell us your GitHub id. Please make sure to do this. Also sign up with your partner into one of the self-serving groups under the people tab on Canvas.

2. Make sure to use this convention in the git repo:
  - **main/master**: for your own development. *We will not check this branch.*
  - **test**: submissions to this branch will be tested by the automated grading framework (see later in this section). Only submit code that you want to get tested, do not do active development here. For each milestone (from m2 onward) you will submit a commitID with your final solution to Canvas - this commitID should point to a commit on the test branch. If that's not the case, we can't grade your assignment. *We will checkout/read for this branch.*
  - **result**: result branch where we will commit the results of the unit tests back to your repo. *Do not commit on this - we will commit the results from the unit test.*

**Warning:** Keep the repo private. If we find your repo as public, you are disqualified for the milestone. No arguments here.

### 3.2 The grading process and automated scoring system

Each milestone is worth 10 points and is evaluated with a mix of methods.

- Milestone 1 - learning and testing the code and all 10 points are earned in a canvas quiz where we will ask you to run the code. You still have to upload the code commit hash in the canvas. Take the quiz and you get the marks.
- Milestone 2, 3, 4, 5: we expect you to commit the code you wish to evaluate in the "test" branch. The final commit we find there (before the assignment deadline) is the one that will be evaluated.

**Ranked grading:** If you pass all functional tests you get a grade of 6/10 (actually 7, see below) for that assignment. The further 4 points depend on the performance quartile you are in, which is calculated as follows

**Primary objective: performance of the code on our servers.**

Total performance range (runtime in milliseconds) = fastest code - slowest code.

Total range is then divided in the 4 ranges. And you get an additional +1 point for each step. So even if you wrote the slowest code, you will get a +1 on top of the 6 marks that you will get. Hence, a total of 7 points. For any higher score you must try harder and check the global leader board that is published periodically in the <https://github.com/StoSysBot/> account.

ALL functional tests passed	6 marks
Bottom 25 percentiles	+1
Middle 25-50 percentiles	+2

Middle 50-75 percentiles	+3
Top 25 percentiles	+4

**The global ranked leaderboard will be updated here:**

<https://github.com/StoSysBot/StoSys2023-Leaderboard>

**Milestones 3 and 5 also have interviews:** if we believe that you coded these milestones but were unable to properly explain your design choices we will deduct points from the base 6 points.

***Join the first lab session to get more details.***

### 3.3 Grading guidelines during the Interview

How the grading is done during the interview :

- We will first ask if your group have to project working
  - If yes, show us demo and explain what is happening
  - If not, then demonstrate to us to the best of your abilities (explain, code) where you are in the milestone.
- Explain various pieces of the code that you have written and what design and logic you have considered to implement and why
- The code will be judged on the generality (if you are making assumptions which are reasonable or not), error handling, hardcoding numbers, unexpected sleep/wait conditions to avoid race conditions, etc. We can not give a comprehensive list of criterias on which your code will be evaluated, however, we are reasonable people and we will explain what is being evaluated and why.
- Follow a coherent coding style, something like <https://www.kernel.org/doc/html/latest/process/coding-style.html>. Ugly code with unmanageable code segments will incur points penalties.
- You are graded on (1) if the code is working; (2) how well do you understand what is implemented - all team members, even if you split the work; (3) how well do you answer questions during the interview; and (4) how well do you understand the design space and choices present in your implementation of the code.

More interview logistics will be published as we approach the interview week.

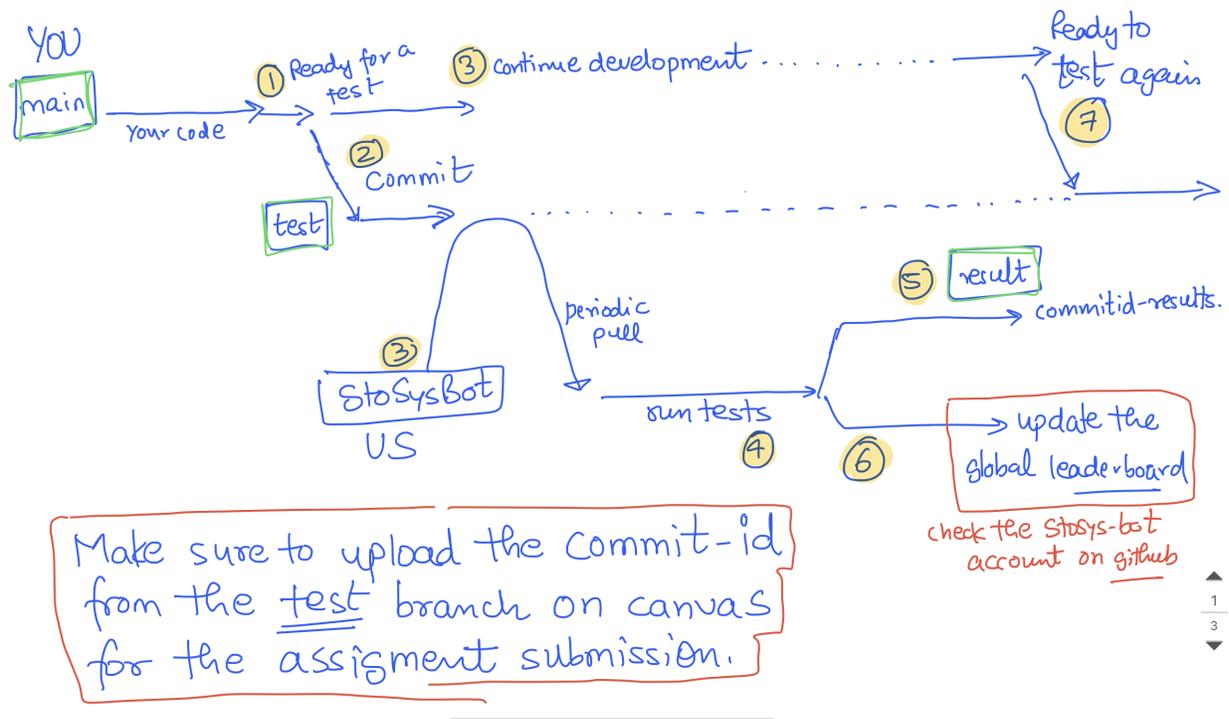
**Note on assessments and marking:** *Often students are under the impression that since they wrote a lot of code, and spent too much time in front of the computer they should get marks.* Unfortunately this is not how programming assignments work. The project that is given to you has been done by us carefully and we know how much it should take as a team. This is a 6 ECTS course, and you need to put time and effort to get it done. A working code is the best proof of your time, energy, and effort put in this course.

**Also** - make sure to start early enough. I am sympathetic to the cause that this is very coding heavy project. You have been warned. If you start one or two days before the deadline there is not much you will get done. The project is designed with a group effort worth 6 ECTS. Keep this in mind. I can not stress this enough: **Start early and make use of the lab session to ask questions and get feedback.**

### 3.4 Automatic Grading Flow

Every time you have a new commit in the “test” branch we will periodically pull the code and commit the unit test results back in the “result” branch.

**Note: This is the second year when this framework is used. We are aware that there are rough edges. Let us know immediately when something breaks. We will make necessary adjustments to make sure no one is penalized unnecessarily.**



This grading framework only runs unit tests for the current milestone. It will automatically shut down once the deadline for that milestone has passed, and will be restarted with the unit tests for the next milestone soon after.

When you push to the test branch, and you do not get any feedback on the result branch in the next hour, push again to the test branch - the framework still has some rough edges. When you still get no feedback on the result branch after that, notify us and we will look into it.

#### Helpful git commands:

- Development on the main branch
  - `git clone repo local_folder`
  - `cd local_folder`
  - `git checkout main`
  - `< code development >`

- git add a, b, c (if there are new files developed)
- git commit a, b, c -m "commit message"
  - OR git commit -am "commit message" (-a says commit everything)
- #pushing upstream
  - git push origin main
- How to push code for testing on the "test" branch
  - # lets say you have code in the main branch that you want to test
  - git checkout test (or use git checkout -b test, if the test branch does not exists)
  - git status
  - \* test (<--- should show that you are on the test branch)
  - # now get all the code from the main and pull and merge in the test
  - git merge main
  - #push to test for testing
  - git push origin test (or git push --set-upstream origin test if it is the first push)
- How to check the result branch?
  - git fetch -a
  - git checkout result
- For submission of M2-M5, get most recent commit hash that is your code submission with: git log --oneline

**IMPORTANT:** Let me emphasize one thing. It happens every year. People do not have all code committed, or forget to add missing files, or all changes are not pushed or what not. It is your responsibility to understand and make sure this does not happen. The commit that you give us to test and upload is THE commit. If it fails during compilation - we are not going to test it and you get a ZERO for the assignment. That is it. If you do not know if the code is complete and ready to test, then do a checkout of main in another location like /tmp folder and see if everything compiles.

### 3.5 Notes on partial assessment for Milestones

It might be the case that you do not have a working code for milestones 2, 3, 4 and 5 at the time of the interview or code upload. So here is the protocol for the assessment

1. You will first be interviewed on your general understanding of the milestone
2. If you demonstrated sufficient expertise then you can demonstrate your partial working code
3. In case you did not demonstrate sufficient understanding of the milestone, we will not review the code and maximum points awarded for this milestone would be whatever you achieve in the interview

**What happens if M3 is not working?** You have a file-backed implementation of M3 given to you.

**A non-working code on our machines will result in maximum marks restricted to the 50% of the maximum marks possible.**

### 3.6 What are the test parameters for automated testing

For milestones 2 - 5, we use the following VM specifications (**these are neither comprehensive or the final version, we can test further configurations if needed. These are shared here to give you an idea of the scope of our testing**):

- CPU: 4 cores, Memory: 8 GB
- NVMe ZNS-1: 64 MB, 4 KB block size, 1 MB zone size and capacity
- NVMe ZNS-2: 64 MB, 8 KB block size, 2 MB zone size and capacity
- NVMe ZNS-3: 96 MB, 4 KB block size, 1 MB zone size and capacity
- NVMe ZNS-4: 96 MB, 8 KB block size, 2 MB zone size and capacity

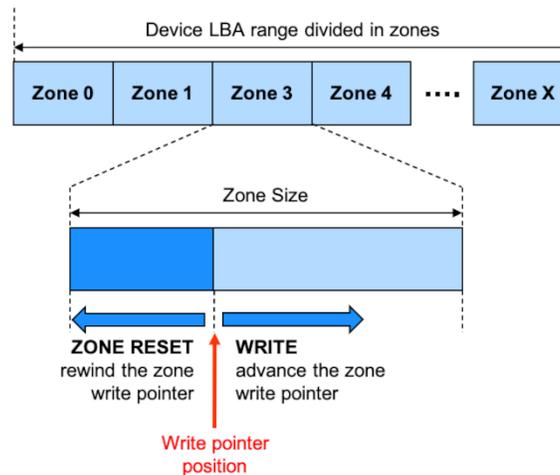
We will use some or all of these disks to test and grade the milestones.

Furthermore, we use the following per-milestone parameters:

- Milestone 2:
  - `-l = {3, 8, 16}`
  - `-s = {1024*1024,1024*1024*2}`
- Milestone 3:
  - `-l 16 -w 1 -o 30000`
  - `-l 32 -w 1 -o 30000`
  - `-l 32 -w 32 -o 30000`
- Milestone 4:
  - `-e 100 -k 10 -v 10`
  - `-e 10000 -k 100 -v 100`
  - `-e 10000 -k 1000 -v 1000`
  - `-e 100000 -k 100 -v 100`
- Milestone 5:
  - `[ -e 100 -k 10 -v 10 ]` for 3 repeats in total
  - `[ -e 5000 -k 10 -v 100 ] [ -e 10000 -k 10 -v 100 ] [ -e 15000 -v 10 -k 100 ]`
  - `[ -e 50000 -k 10 -v 10 ]` for as many repeats as possible to test how many key value pairs we can write to your disk, to make it better.

## 4. StoSys Project Work

**Background:** NVMe is a command specification regarding how one should interact with NVMe devices which typically are high-performance non-volatile memory (NVM) storage devices such as Flash or Optane. The NVMe command set is designed to be lean, light-weight to simplify how many times system software needs to interact with the NVMe storage devices to complete an I/O request.



**Figure 1:** The principle of Zoned Device operations  
(taken from <https://zonedstorage.io/docs/introduction/zoned-storage>)

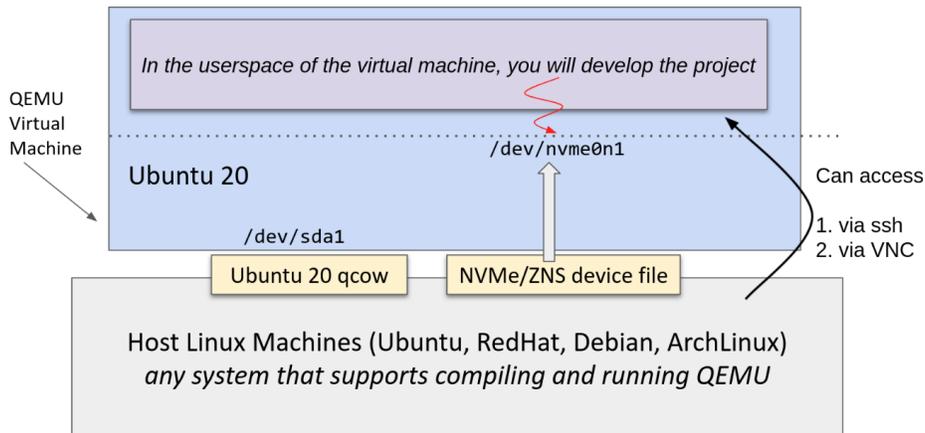
The idea of Zoned Devices, i.e. a device interface with append-only interface with segmented zones can be implemented with many device technologies including hard-disk drives (HDDs). With HDD it makes even more sense as HDDs have high penalty for random, rotational accesses. Hence, there are shingling HDDs in the market that you can buy today that follow the (almost) same Zone interface. Consequently, there are SCSI extensions for Zoned HDDs also (see <https://www.t10.org/ftp/zbcr01.pdf>). These specifications are implemented in a few userspace libraries such as libzbc, see <https://zonedstorage.io/docs/tools/libzbc>.

Just like that, NVMe compliant devices also have Zone Complaint Specification (originally prepared as Technical Proposal TR 4053) which is now merged in the NVMe Specification 2.0 in August 2021. The TP-4053 and NVMe 1.4b specification PDFs are given to you in the given framework (see the ./doc folder)<sup>1</sup>.

In this project work we will use NVMe Zoned Namespace (ZNS) devices to do the practical work. ZNS devices are NVMe devices. They use NVMe read, write commands (among other management commands) together with controller and namespace identifiers. Apart from that, they have additional ZNS specific operations like ZNS controller and namespace identifiers, append, reset zone commands. All zone related efforts are very nicely documented at <https://zonedstorage.io/>

Here is a high-level development setup in which you will be doing the development during this project.

<sup>1</sup> It might be confusing to follow all these details, but try to read through. We will only be using NVMe/ZNS devices in the project which are properly introduced.



**Figure 2:** Stosys project development setup with QEMU

A note on the terminology used:

- Block, Page, Sector - often (and sometimes confusingly) used to refer to the minimum I/O size for read and write operations. For example, HDDs have 512 bytes sectors. Flash devices can have 512 bytes of 4kB page size. Sometimes a page is also referred to as a block.
- Logical Block Address (LBA) - the address of the logical minimum read/write units.
- Zone address - the starting address of a zone. With ZNS the zones are clear. However, before ZNS sometimes the unit of garbage collection (which is the zone) was also referred to as a block. So just pay attention if we are talking about I/O operations or GC operations. Typically GC operates on larger and different units (like a couple of MegaBytes) than the basic I/O to amortize the cost of the GC operation.

When in doubt, just ask.

**Important:** Do not touch the **m1.cpp**, **m2.cpp**, **m3.cpp**, and **m45.cpp** source files. You can debug whatever you like but for the test commit and for the unit test, you must revert all the changes in those files. All your source code must be in other files.

**Do not forget to add new files in CMakeLists and git.**

## 4.1 Milestone 1 (week 1) : A new device is in town

**Type:** individual

**Maximum points:** 10

**What to do:** prepare coding environment, (i) code zip upload, (ii) tell us your GitHub id and GroupId, and (iii) take canvas quiz for milestone 1

**Where:** in the given framework, and canvas

**Deadline:** see the course canvas page

The goal of milestone 1 is to set up the development environment, and get familiar with the setup, framework, and ZNS and NVMe devices (and specification).

We are going to do the project development on top of Zoned Namespace (ZNS) devices. ZNS are NVMe devices and support all the basic NVMe commands (read, write, controller and namespace querying), plus ZNS specific details.

**Why build from scratch:** We will be building a lot of tools from the source for a particular specific version as this is what we have tested over the years. Even though the newest version of these tools might have all the features that we want, there is no guarantee. Plus many of these frameworks and tools are part of the cutting edge research, hence they evolve very fast. For the stability and support in this course, we require you to run a very specific version of the code.

#### 4.1.1 Task-1 setup QEMU 6.1

QEMU is a whole system emulator (with CPU, devices, BIOS, etc. everything) which is very popular to do real systems research with and is being deployed in the cloud. See <https://www.qemu.org/docs/master/> for more details. We will do the project work in QEMU where the ZNS device will be emulated. The support for these devices was only added recently (and for the sake of debugging), we need a specific version that we compile from the source:

**Step 1:** Download **QEMU 6.1** source code from <https://download.qemu.org/qemu-6.1.0.tar.xz> (or <https://www.qemu.org/download/>)

**Step 2:** Untar the code and then (you have to install some missing packages based on the Linux distribution that you are using as the host machine where QEMU is being compiled at)

```
$ tar xvf qemu-6.1.0.tar.xz
$ cd qemu-6.1.0
$ mkdir build && cd ./build
$ ../configure --target-list=x86_64-softmmu --enable-kvm --enable-linux-aio
--enable-trace-backends=log --disable-werror --disable-gtk
$ make -j
```

**Step 3:** Once the compilation is done successfully, add the QEMU binary path in your `.bashrc` file (no need to install QEMU) as:

```
# find where the PATH variable is defined in your ~/.bashrc file
# put export PATH=<your_qemu_build_directory>:$PATH
# example of PATH line my .bashrc file:
export PATH=/home/atr/vu/qemu-6.1.0/build:$PATH

# do not forget to reload the new PATH variable
$ source ~/.bashrc
```

Assume here that the final path of your binary is `/home/atr/vu/qemu-6.1.0/build`.

**Step 4:** At this point you should have QEMU built and ready to use, you can check it as (see the correct version number as 6.1.0). It is **important** that you run this as `sudo`!

```
$ atr@atr-xps-13:~$ sudo qemu-system-x86_64 --version
QEMU emulator version 6.1.0
Copyright (c) 2003-2021 Fabrice Bellard and the QEMU Project developers
```

At this point we are ready to start the VM and install Ubuntu 20 in it.

### 4.1.2 Prepare the Ubuntu 20 in the QEMU VM

**Pre-built Image:** Maybe you do not want to see how sausages are made, and just want to get on with the project work. In that case, we have provided a cleanly pre-installed image of Ubuntu 20.04 LTS with 5.12 kernel at :

<https://drive.google.com/drive/folders/19J-avjdbY0MQgPFzOWOdNC3ID1R5LN1w?usp=sharing>

Alternate link: <https://surfdive.surf.nl/files/index.php/s/bFU81kTvuxSwYcg>

However, you still have to build QEMU 6.1 from source and make two NVMe and ZNS files as shown below. You can skip the Ubuntu 20.04 installation and v5.12 kernel building. For this image:

**Username:** user

**Password:** stosys!!

This user has sudo privileges. The base root filesystem has 50GB. Feel free to modify this image and install whatever you like to facilitate your code development. I would not recommend a GUI and/or window-manager though. Learn how to do CLI-based code development and use git.

Just like a normal system QEMU needs a hard-disk image where the OS will be installed and booted from. We will also create a second image file that will be used as NVMe and ZNS devices (two different devices for testing and experiments).

**Step 1 creating images:** You can create a Ubuntu base hard-disk where the Ubuntu will be installed as :

```
# (RECOMMENDED) with qcow2 mode, does not consume 50G immediately
$ qemu-img create -f qcow2 ubuntu-2004.qcow2 50G

# otherwise, raw files will consume 50G immediately
$ qemu-img create -f raw ubuntu-2004.img 50G
```

Now create two files which will be the NVMe and ZNS devices that we will experiment with:

```
# For ZNS and NVMe drives, we just create two raw files
$ qemu-img create -f raw znsssd-32M.img 32M
$ qemu-img create -f raw nvme32M.img 32M
```

**Step 2 installing the ubuntu 20.04 server system (no GUI):** Download an ubuntu 20.4 server install iso from (checked on Aug 23d, 2023): <https://old-releases.ubuntu.com/releases/20.04.5/> and download the image: <https://old-releases.ubuntu.com/releases/20.04.5/ubuntu-20.04-beta-live-server-amd64.iso>

Then start the QEMU VM using the full path of the binary (update the path according your local building directory):

```
$ sudo /home/atr/vu/qemu-6.1.0/build/qemu-system-x86_64 -hda ubuntu-2004.qcow2
--enable-kvm -m 4G -smp 2 -boot d -cdrom ubuntu-20.04-live-server-amd64.iso -net
user,hostfwd=tcp::7777-:22 -net nic -cpu host
```

For QEMU related parameters see (checked on Aug 23d, 2023): <https://www.qemu.org/docs/master/system/qemu-manpage.html> (please mind the version that you are looking at)

If with the start of the QEMU machine you get an error like:

```
Could not access KVM kernel module: No such file or directory
```

Try to probe the kvm module. If it is not found you might have to enable virtualization in the BIOS of your machine.

```
sudo modprobe kvm-intel
```

Once the system is running you can connect to the system using VNC using the output from the previous command run. You should see something like this in the terminal

```
VNC server running on 127.0.0.1:5900
```

You can connect to the virtual machine by putting this address in the VNC client. If your host system is also Ubuntu, you should have **Remmina** already installed (see Appendix A and B how to do remote desktop). You can paste the address "127.0.0.1:5900" in the address bar and change the protocol to VNC (not RDP) and connect (see Appendix B for more details). There is no special setup for the Ubuntu 20.04 QEMU installation. Follow the defaults to install the Ubuntu 20.04 image inside the QEMU virtual machine.

**MacOS Bug:** The apple remote desktop doesn't work. Use Royal TSX when connecting to the QEMU VM.

**Remember to enable SSH during** installation to be able to access the VM without VNC! At the end of your installation step, the system will ask if you want to install any new packages then pick **OpenSSH** server. If you missed this step, you can still login through VNC and then install the OpenSSH server. Once the installation is complete, you can stop the VM using Ctrl+C.

If you are having trouble installing the OS, see Appendix A for a detailed step by step guide.

**Step 3 starting the system and ssh:** Once the installation of the QEMU Ubuntu 20.04 is done, you can start the complete virtual machine with all the settings as (you can change -m flag to the largest memory you can afford on your machine like, 4, 8, 16 or 32 Gigabytes)

```
$ sudo /home/atr/vu/qemu-6.1.0/build/qemu-system-x86_64 \
-name qemuzns -m 4G --enable-kvm -cpu host -smp 4 \
-hda ubuntu-2004.qcow2 \
-net user,hostfwd=tcp::7777-:22 -net nic \
-drive file=znsssd-32M.img,id=zns-device,format=raw,if=none \
-drive file=nvmessd-32M.img,id=nvme-device,format=raw,if=none \
-device nvme,serial=zns-dev,id=nvme2,zoned.zasl=5 \
-device
nvme,drive=nvme-device,serial=nvme-dev,physical_block_size=4096,logical_block_size=
4096 \
-device
nvme-ns,drive=zns-device,bus=nvme2,nsid=1,logical_block_size=4096,physical_block_si
ze=4096,zoned=true,zoned.zone_size=1M,zoned.zone_capacity=1M,zoned.max_open=0,zoned
.max_active=0,uuid=5e40ec5f-eeb6-4317-bc5e-c919796a5f79
```

In case you encounter the following error

```
'qemu-system-x86_64: -device nvme,serial=zns-dev,id=nvme2,zoned.zasl=5: Property 'nvme.zoned.zasl'
not found'
```

make sure you are using the correct QEMU version and binary, which you can check with `'sudo which qemu-system-x86_64'`.

Make sure to use the right file location for the files which are highlighted in yellow above. We are not going into the details of what each of these QEMU parameters do. Please see here (August 23rd, 2023) :

- QEMU documentation <https://www.qemu.org/docs/master/system/index.html>.
- NVMe documentation <https://www.qemu.org/docs/master/system/devices/nvme.html>
- Further information: <https://zonedstorage.io/docs/getting-started/zbd-emulation>

After a few seconds you should be able to ssh into the VM as (see the QEMU command above, hostfwd, which forwards traffic from local 7777 port to VM's 22 port which is the port used by the ssh server)

```
$ ssh -X user@localhost -p 7777
```

Otherwise use the username and password you set up during the Ubuntu 20.04 installation. For any subsequent ssh instruction we will use the provided one. If you get too many authentication failures, then disable the public key check and only use the password:

```
$ ssh -X user@localhost -p 7777
Received disconnect from 127.0.0.1 port 7777:2: Too many authentication failures
Disconnected from 127.0.0.1 port 7777
$ ssh -X user@localhost -o PubkeyAuthentication=no -o
PreferredAuthentications=password -p 7777
```

**Step 4 upgrading the Linux kernel:** Since NVMe and ZNS development is very cutting edge, the base kernel that comes with Ubuntu 20.04 (v5.4) is too old. See this: <https://zonedstorage.io/docs/linux/overview#kernel-versions> (only v5.9 and above has support for ZNS devices).

Now we are going to upgrade the kernel. You can choose the latest kernel version, but I strongly recommend to use the **5.12 version** - this is what I tested extensively. The following compiling instructions are based on: <https://wiki.ubuntu.com/KernelTeam/GitKernelBuild>. To compile the v5.12 kernel (you do not have to do this inside the VM, just find the fastest Linux machine you can which can compile a kernel. Any distro will do, but x86-64 architecture is recommended unless you know how to cross compile):

```
$ git clone https://git.kernel.org/
$ cd kernel (or whatever folder it made)
$ git checkout v5.12

# make sure all packages are installed to build a kernel
$ make olddefconfig
# or you can also do
$ make menuconfig

# now we build
$ make clean
$ make -j $(getconf _NPROCESSORS_ONLN) deb-pkg LOCALVERSION=-stosys
```

After a successful make (can take 20-30mins), you should have a bunch of deb files in your parent folder where you compiled the kernel. Then copy the final deb files inside the VM and install them. I did this like this (make sure you change the file names to what you have):

```
# on the machine you compiled the kernel (make sure that the VM is running)
$ scp -P 7777 linux-image-5.12.0+_5.12.0-stosys-6_amd64.deb user@localhost:~/
$ scp -P 7777 linux-headers-5.12.0+_5.12.0-stosys-6_amd64.deb user@localhost:~/

# inside the VM
$ sudo dpkg -i linux-image-5.12.0+_5.12.0-stosys-6_amd64.deb
$ sudo dpkg -i linux-headers-5.12.0+_5.12.0-stosys-6_amd64.deb
$ sudo upgrade-grub2
$ sudo sync
$ sudo reboot
```

At this point, your VM should boot into the 5.12 kernel. You can skip all these steps and can just use the provided 5.12 image directly and move to the next subsection.

### 4.1.3 Task-3 Familiarize yourself with online resources

The following are the key resources that you should read through to get a conceptual understanding of ZNS devices and the Linux development environment and integration.

1. In general, I would strongly recommend flipping through the whole <https://zonedstorage.io> website. You may find yourself looking for details here over and over again.
  - a. Understand what zoned storage is (broader concept): <https://zonedstorage.io/docs/introduction/>
  - b. Introduction to NVMe ZNS concepts : <https://zonedstorage.io/docs/introduction/zns/>
  - c. libnvme user library <https://zonedstorage.io/docs/tools/libnvme> (you will be using this library in the project - but we have some modifications done to it. So use the one what we provide in the framework.)
  - d. Understand the broad Linux and ZNS ecosystem: <https://zonedstorage.io/docs/linux/>
2. You *should* flip through the **NVMe 1.4 specification and ZNS TR 4053** (given in the framework, see the ./docs/ folder).
  - a. In the NVMe 1.4 specification you should read through section: **chapter 1**, sections **3.1** (has details about controllers capabilities and metadata), section **4.1, 4.2, 4.6** (explains the basic I/O data structures. **4.6.1.2** contain error codes), section **5.15** (identify namespace **5.15.2.1** and the NVMe controller **5.15.2.2** are important), **6.15** (write command), and **6.9** (read command). Many of the details of these will be hidden behind the library that you would use but you still need to fill out function parameters which are defined and explained in these sections.
  - b. In the Technical Proposal, TP 4053 you should read: chapter **2**, section **3.1** (ZNS namespace and controller identification structures, section **4.1** (ZNS errors), **4.3** (how to interact with zones, including resetting and closing them), **4.4** (explains how interact with zones and extract their current state), **4.5** (special zone append command). Notice figure 20 for an overview.

The idea of reading these specifications is not to memorize them, but to make sure you get an intuition how the information is organized, and how you can quickly look for specific commands or details. Also, you do not have to read them all at once, instead read them gradually as you move with the coding.

The TP 4053 is now merged in the latest NVMe 2.0 specification, <https://nvmexpress.org/specifications/>. The aspects which are covered in this lab are the same across the 1.4 and 2.0 specification. We have not upgraded the project to this latest specification.

#### 4.1.4 Task-4 Build and run nvme cli command tool

Inside the Ubuntu 20.04 (with 5.12 kernel) VM we are now going to interact with NVMe and ZNS devices. The primary way a user can interact with an NVMe device is via using nvme command line interface tools (nvme-cli). This command/package is available on most of the popular Linux distributions. However, since we need support for the latest ZNS devices, we need to clone and build from source.

nvme-cli <https://github.com/linux-nvme/nvme-cli>

I have tested the code with **v1.14**. Use the specific commit that we show below. See the README.md file in the repo on how to compile the code.

```
#install the dependency
$ sudo apt-get update
$ sudo apt install uuid-dev

$ git clone https://github.com/linux-nvme/nvme-cli
$ cd nvme-cli
$ git checkout f38d0c1c20db1c8e870ba828be1bd960511dba98
$ git submodule init
$ git submodule update
$ make -j
...

$ sudo ./nvme --version
nvme version 1.14.121.gf38d
$ sudo ./nvme list
```

Node	SN	Model	Namespace	Usage	Format	FW Rev
nvme1n1	nvme-dev	QEMU NVMe Ctrl	1	8.39 MB / 8.39 MB	4 KiB + 0 B	1.0
nvme0n1	zns-dev	QEMU NVMe Ctrl	1	8.39 MB / 8.39 MB	4 KiB + 0 B	1.0

If you get an error due to broken NL link like:

E: Failed to fetch

[http://nl.archive.ubuntu.com/ubuntu/pool/main/u/util-linux/libuuid1\\_2.34-0.1ubuntu9.1\\_amd64.deb](http://nl.archive.ubuntu.com/ubuntu/pool/main/u/util-linux/libuuid1_2.34-0.1ubuntu9.1_amd64.deb)

Then change the mirror from the NL to the USA:

- Step 1: open the `sudo vim /etc/apt/sources.list`
- Step 2: change all the links from <http://nl.archive.ubuntu.com> to <http://us.archive.ubuntu.com>

At this point we are ready to interact with the NVMe devices using the command line interface using our newly compiled nvme command. In the example above you can see the two NVMe devices (one normal and one ZNS is shown). *I used 8MB devices for testing, hence 8MB capacity.* You can also put the location of the nvme command binary in your \$PATH variable in the .bashrc file if you do not want to always run the command from the location.

For more examples of how you can use nvme cli tool to interact with the ZNS and/or NVMe devices see:

<https://zonedstorage.io/docs/tools/zns>

**Important:** the way the setup is done, /dev/nvme0n1 is your ZNS device and /dev/nvme1n1 is your ordinary NVMe device.

Your particular tasks is to run these command, fill in the details (*find any valid value to fill out the yellow parts and see if you understand it*), understand their output and match that with the specification (make yourself familiar with how to read the NVMe and ZNS specifications):

1. Extract and identify NVMe controller and Namespace structures for the NVMe device and match the details from the NVMe specifications (5.15.2.1 and 5.15.2.2). Hint use -H flag and compare the output with specification.

```
# read the help in the commands
$ nvme help
$ nvme id-ctrl help
$ nvme id-ns help

# now fill in the command
$ sudo nvme id-ctrl [dev] -H (fill the yellow part)
$ sudo nvme id-ns [dev] -H (fill the yellow part)
```

2. Extract and identify ZNS controller and Namespace structures for the ZNS device and match them with the ZNS specification (3.1.1 and 3.1.2)

```
$ nvme zns help
$ nvme zns id-ctrl help
$ nvme zns id-ns help

# now fill in the command
$ sudo nvme zns id-ctrl [dev] (fill the yellow part)
$ sudo nvme zns id-ns [dev] -H (fill the yellow part)
```

3. Execute NVMe LBA read and write commands and test them using -d flag. With -d (or --data) you can write and read data from files, and then compare their checksums. Keep in mind the ZNS device uses the same NVMe read and write commands.

```
# read the help in the commands
$ nvme read help
$ nvme write help

# now fill in the command to check if you can read, write and compare
outputs
# make a random 4kb file
$ dd if=/dev/random of=./4kb bs=4096 count=1
# now we write and read this file, and compare the content

$ sudo nvme write [dev] -s [ ] -z [ ] -d ./4kb (fill the yellow part)
$ sudo nvme read [dev] -s [ ] -z [ ] -d ./4kb-w (fill the yellow part)

#compare the data, the checksum should match
$ md5sum ./4kb
075ca305497df6cfb9500abb98993438 ./4kb
$ md5sum ./4kb-w
075ca305497df6cfb9500abb98993438 ./4kb-w

# now let's try to rewrite the ZNS again with the same command and same values!
$ sudo nvme write [dev] -s [ ] -z [ ] -d ./4kb (fill the yellow part)
???
```

*← What output do you get? Do you get the same output with NVMe?*

4. Execute ZNS zone report and Zone reset commands:

```
# read the help in the commands
$ nvme zns report-zones help
$ sudo nvme zns report-zones [dev] -s [ ] -H (fill the yellow part)

# reset zones
$ sudo nvme zns reset-zone [dev] -s [ ] (also see -a)
```

5. Execute ZNS zone append command

```
# read the help in the commands
$ sudo nvme zns zone-append help

# fill in the command:
$ sudo nvme zns zone-append [dev] -s [ ] -z [ ] -d ./4kb

# now run this command 3 times,
$ sudo nvme zns zone-append [dev] -s [ ] -z [ ] -d ./4kb
$ sudo nvme zns zone-append [dev] -s [ ] -z [ ] -d ./4kb
$ sudo nvme zns zone-append [dev] -s [ ] -z [ ] -d ./4kb

# Where is the zone write pointer?
$ sudo nvme zns report-zones [dev]
```

**Why are we doing so many practice commands?** When you run into issues with your project code, you can compare the behavior of the ZNS/NVMe devices with these commands (with a similar I/O pattern) and see if you are making any mistake the way you are interpreting commands, their results in your code. You can dump the state of the ZNS device (zone state, write pointers) and diagnose problems. Plus, how these commands are implemented is giving you a reference code to browse. Feel free to read the nvme cli code to get a better understanding.

**How to debug nvme error codes?** The NVMe error codes are 16 bit values and are returned as Status Field in the Completion Queue Entry structure (see section 4.6 in the NVMe 1.4 specification, figure Figure 123). Please read that section to figure out how to decode such values.

For Command Specific Status specific errors (like with ZNS), you might find yourself looking for in the range marked as "80h to BFh - I/O Command Set Specific". Then you should look for them in the ZNS specification.

If you are curious, then you can use the NVMe 2.0 specification (which has everything in one single PDF with the merged specifications) in the section 3.3.3.2 (Common Completion Queue Entry). The 2.0 has ZNS codes merged in the NVMe error codes. NVMe 2.0 (check Aug 23rd, 2023) - <https://nvmexpress.org/wp-content/uploads/NVMe-NVM-Express-2.0a-2021.07.26-Ratified.pdf>

#### 4.1.5 Task-4 Build and prepare the coding environment

The programmatic way to interact with the NVMe/ZNS devices is through libnvme userspace shared library (among other ways, see bonuses for that). In this task we are going to compile the libnvme library from source (with some patches from me) to make sure things work. At the end of this step you should be able to compile the given framework to you for milestone 1.

Copy the given framework with in the virtual machine

```
# from the host machine copy it inside the VM
$ scp -P 7777 msc-stosys-framework-skeleton.zip user@localhost:~/

# ssh into the VM
$ ssh -X user@localhost -p 7777

# unzip and setup
$ sudo apt install unzip (if needed)

$ unzip msc-stosys-framework-skeleton.zip
```

The patched version of the library is given to you in the framework (the actual library is located at <https://github.com/linux-nvme/libnvme>). The patch is nothing special, except that libnvme did not export ZNS specific functions and a few memory leaks were fixed. More recently they have moved their build to meson, which I have not tested. Furthermore, when I prepared the project, the libnvme scan was broken in Linux, so I had to work around it. In order to avoid any unforeseeable issues with the library (a fast moving coding repo). I have a fork that is given to you in the framework. Here is how you can compile and install it:

```
# install the package config dependency
$ sudo apt install pkg-config

# then build and compile
$ cd msc-stosys-framework/libnvme
$ ./configure
$ make -j
$ sudo make install

# if you want to install this in a non-standard path, you can
$ ./configure --prefix=/home/$USER/local/
[...] ← see the output here with the appropriate path names
$ make -j
# if you install in some non-standard path, then you probably do not need sudo
$ make install

# make sure to put your non-standard install (if you did that) path in your -I
(gcc) and LD_LIBRARY_PATH
```

Once you have installed the library, you can prepare the framework:

```

$ cd msc-stosys-framework
$ cmake .
-- Building using CMake version: 3.16.3
compiler is GNU and name is g++
[info] compiler is gcc/g++
[info] ASAN is on, adding additional ASAN fields to the compiler
-- Configuring done
-- Generating done
-- Build files have been written to: /home/atr/msc-stosys-framework
$ make -j

```

At this point you are ready to finish the code for Milestone 1.

### 4.1.6 Task-5 Finish the missing code pieces and take canvas quiz

Milestone 1 (in [m1.cpp](#)) does the following things:

1. Does a full system scan to find all devices (similar to the `nvme list` command as shown above), return a count of all the devices.
2. Does a second scan, but this time splits the found NVMe devices into ZNS-capable or not devices, and returns an array of `struct ss_nvme_ns` devices (size of the array is the number of devices found).
3. Pick the last ZNS capable device and enumerate its properties ([show\\_zns\\_zone\\_status](#))
4. Then it performs the following I/O tests:
  - a. Test 1 [test1\\_lba\\_io\\_test](#): In this function, a series of LBA size tests are conducted. This includes: read, write, zone reset, and zone append.
  - b. Test 2 [test2\\_zone0\\_full\\_io\\_test](#): In this function, a full zone size I/O test (reset, write, and read) are tested.
  - c. Test 3 [test3\\_threaded\\_parallel](#): it starts the “p” number of parallel threads to do I/O for “t” seconds.
5. Close everything and finish.

M1 takes the following arguments:

- -d : which device index to use.
- -p : number of parallel threads.
- -t : test time in seconds.
- -h : shows help, and exits with success. No argument needed

Your goal is to fill up functions in the [m1\\_assignment.cpp](#) so that the aforementioned tests pass. More specifically, you need to provide a complete function implementations for:

- [ss\\_nvme\\_device\\_read\(...\)](#): Write a function to read one or more LBA addresses in a user provided buffer from an NVMe device (the same can be used for ZNS devices).

- `ss_nvme_device_write(...)`: Write a function to write one or more LBA addresses from an NVMe device (the same can be used for ZNS devices).
- `ss_zns_device_zone_reset(...)`: Write a function to reset a given zone on the ZNS device.
- `ss_zns_device_zone_append(...)`: Write a function to append data in a given zone on the ZNS device. Keep in mind, appending also has a maximum append size (ZASL, see Figure 10 in section 3.1.2 in the ZNS TR 4053).
- `update_lba(...)`: Update LBA function takes an LBA address and returns the next LBA address that can be used to write data with a certain offset. For example, if you start writing from LBA 0 with a data size of 2 x LBA size (2 LBA blocks), then the function should return the next LBA address to write after writing the 2x LBA data. It is important to understand this concept because all zones have an associated write pointer in them that is internally incremented with each write in zones, and you must understand how this is done. If you try to write an LBA in the zone which does not match the zone-internal write pointer then you will get an error.
- `get_mdts_size(...)`: Return the maximum data transfer size in bytes. See the QEMU MDTS bug below. This function specifically takes a variable number of arguments. Please see how to set and parse variable number of arguments here: <https://stackoverflow.com/questions/1579719/variable-number-of-parameters-in-function-in-c>
- `ss_nvme_device_io_with_mdts(...)`: Write a function around read/write functions (defined above) where a large I/O request can be split into multiple MDTS size I/O requests. You can pass MDTS size as an argument in the function. For example, if an NVMe device has MDTS size of 512KB, then a complete zone read/write of 1MB should be split into 2x 512 KB I/O requests. This is a performance optimization. Naturally you can always split any I/O request into the smallest unit, i.e, the page size.

Almost all of these functions can be implemented in *10s of lines of code* by directly calling appropriate functions in libnvme. Browse libnvme code and header files. `get_mdts_size()` may be more complicated, see below.

**Libnvme does not export all functions by default.** Hence if you find yourself wanting to include and build with any function which is not exported from the libnvme then include the function name in the `libnvme/src/libnvme.map` file and rebuild and install the library. Please also let us know.

**Calculating MDTS is challenging.** MDTS is defined in the identify controller data structure (5.15.2.2 Identify Controller data structure (CNS 01h)), but the unit is (CAP.MPSMIN) page size. The value MPSMIN (indicates the minimum host memory page size that the controller supports, see section 3.1.1 Figure 69: Offset 0h: CAP – Controller Capabilities) is defined in a Read-Only Capability register (as PCIe BAR). You need to figure out how to read this register (**Hint**: see `nvme show-regs` command, and corresponding code in the `nvme-cli` repository). See Issue-2 below with the QEMU MDTS bug.

Please pay attention to compiler warnings and memory leak errors. We will look at them properly and deduct marks, if appropriate. See the CMakeLists to see which flags are we using.

**Example of a successful run:**

```

user@stosys-qemu-vm:~$ sudo ./bin/m1
[sudo] password for user:
=====
This is M1. The goal of this milestone is to explore the framework
=====
settings, timeout : 10 seconds, max_parallel: 4 threads, and dev_index: -1
nvme-subsys1 - NQN=nqn.2019-08.org.qemu:nvme-dev
  ` - nvme1 pcie 0000:00:05.0 live
    ` - nvme1n1lba size:4096 lba max:8192
nvme-subsys0 - NQN=nqn.2019-08.org.qemu:zns-dev
  ` - nvme0 pcie 0000:00:04.0 live
    ` - nvme0n1lba size:4096 lba max:8192

<...>

=====
Milestone 1 results
Test 1 (read, write, append, reset) : Passed
Test 2 (Large zone read, write)      : Passed
Test 3 (multi-threaded I/O)         : Passed
=====
user@stosys-qemu-vm:~$

```

After finishing - upload the code and take the canvas quiz. **Make sure your code does not have any memory leaks!**

### Issue: QEMU MDTS bug

There is some issue (unofficially tested, but not confirmed) that whatever MDTS QEMU NVMe reports, use one less than that. For example, if the MDTS reported is 6, then use 5. In my testing the NVMe (and ZNS) device fails non-deterministically to complete the I/O with the reported MDTS size. However, this does not happen on real NVMe hardware (like NVMe Optane or Flash), only in QEMU NVMe devices.

## 4.1.7 All together

**Specific goals for M1:** To recap this milestone (as there are many small things to do):

1. Make sure you have a working development setup with QEMU and ZNS.
2. Read the online documentation (<https://zonedstorage.io/>), TR 4053, and NVMe 1.4b specifications (only the sections prescribed above. The whole 1.4b specification is 400+ pages).
3. Make sure nvme-cli is working and you can browse the code.
4. Complete the Milestone 1 missing functions and make sure the milestone 1 tests passes.
5. Upload the full framework code on Canvas as a zip file with the format : milestone1.zip (or tar or whatever) Check the upload deadline on canvas.
6. Take the canvas quiz.
7. Finish the assignment where you need to tell us your Github id and your group number.
8. **[Not graded - IMPORTANT]** Find yourself a team-mate to work as a team for Milestone 2 onwards. When choosing a teammate ask them about the details of this milestone to get a good skill match.

**Important:** In the quiz we will ask you to run certain commands and find the right answer from the options given. So it is very important that you keep the VM running when you take the quiz.

## 4.2 Milestone 2 (week 2) : I can't read, is there a translator here?

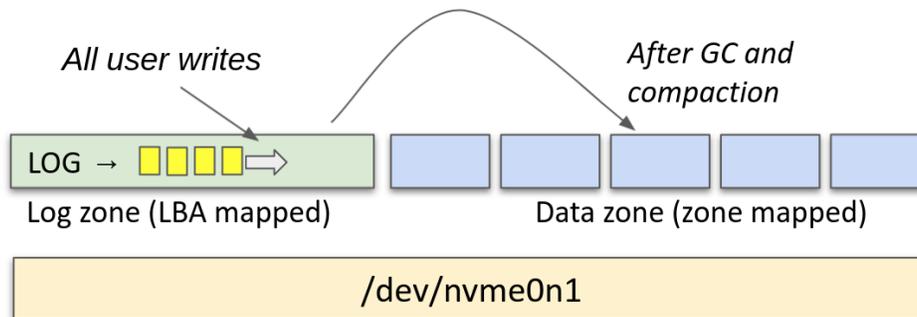
**Type:** group project

**Maximum points:** 10

**What to do:** implement a hybrid log-structured FTL to hide ZNS device complexity (*without GC*) and upload the final code commit git hash from the test branch in the canvas assignment

**Where:** Where: code development in the given framework

**Deadline:** see the course canvas page



**Figure 2 : A high-level layout of the hybrid log-structured FTL design**

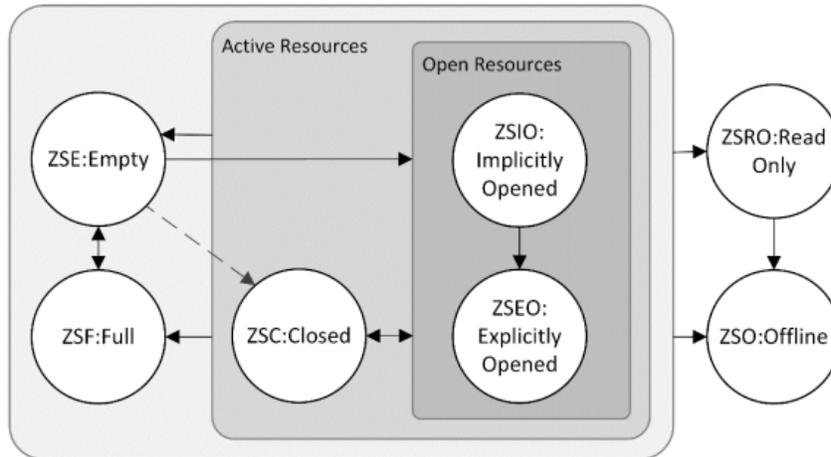
The purpose of this milestone is to design and implement a **hybrid log-structured flash translation layer (FTL)** on top of the ZNS device. With this FTL in place, the ZNS device should appear as a normal block device where a user can read and write any block offset without any restriction.

The design of the hybrid log-structured FTL is discussed in the class. See the class slides.

The key design decisions that you have to make is

1. What data structure to use to keep track of page and zone -level maps
2. Are these data structures scalable (can scale to TBs) and thread safe? (M4-M5 are multi-threaded).
3. How do you plan to support more than page-level I/O in a single call? (if you want to do this by leveraging MDTS).
4. [Plan] what will GC do and when a log content needs to be moved from the log to data zone or vice-versa

When designing and implementing the FTL, make sure you understand how ZNS devices are expected to be read, written and reset.



**Figure 3:** NVMe ZNS zone state machine

**Practical information:** Milestone 2 (M2) is decoupled from milestone 1. M2 is built as a shared library (`libstosys.so`) which is then eventually linked against the `m2.cpp` code (the main file which runs the tests). You should aim to implement high-level functions (your FTL device open, read, write, and close) as given in the `zns_device.h` file. In that file there is also a `struct user_zns_device`. This structure is the primary device structure that should be used everywhere in your code. There are following fields:

1. `uint32_t lba_size_bytes`: Logical block size that the user should adhere to when doing I/O in this device. This can be the same as the underlying LBA size of the NVMe/ZNS device or some multiple of it.
2. `uint64_t capacity_bytes`: Total user-usable capacity in bytes. This would be the actual ZNS device capacity, minus the FTL metadata and log zones. A user should be able to read/write this area `[0, capacity_bytes]` in `lba_size_bytes` pages any number of times (for this milestone without GC, only once).
3. `void *_private`: A private pointer area that you can use to save any information (like your own private structure with information). The user is not supposed to touch or change this.

Make sure you respect the arguments passed (initialized in the struct `zdev_init_params`) to the M2 program. It takes:

1. `-d` : `/dev/nvmeXpY` - in this format with the full path.
2. `-l` : the number of zones to use for log/metadata (default, minimum = 3).
3. `-r` : resume if the FTL can. (useful for M5).

You need to implement the following functions (see the :

1. `int init_ss_zns_device(struct zdev_init_params *params, struct user_zns_device **my_dev)` : this function creates a user-device with appropriate values.
2. `int zns_udevice_read(struct user_zns_device *my_dev, uint64_t address, void *buffer, uint32_t size)` : a user-facing read function. A user should be able to read any arbitrary LBA address (any size, which is multiple of the page size).
3. `int zns_udevice_write(struct user_zns_device *my_dev, uint64_t address, void *buffer, uint32_t size)` : a user-facing write function. A user should be able to write to any arbitrary LBA address (any size, which is multiple of the page size).

4. `int deinit_ss_zns_device(struct user_zns_device *my_dev)` : a shutdown close function for a device.

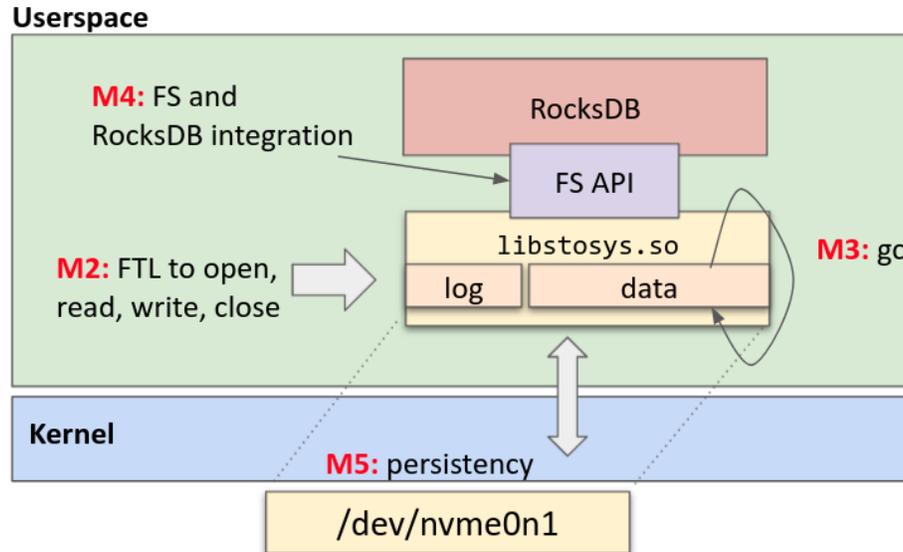
The m2.cpp file has following tests that your FTL implementation should pass (at this stage)

1. Test 1 - a simple LBA zero write and read (and pattern matching)
2. Test 2 - the whole log zone worth of pages are written randomly (the log zone capacity is known, hence, writing the whole log zone does not and should not create a GC). A pattern is matched.
3. Test 3 - A same test as Test-2, but with a larger buffer. Here, if your implementation is optimized with MDTS sizes then you will see performance gains.

As you may have noticed the mix of C and C++ languages in the project. Feel free to do so, but be careful with the linkage between C and C++ (when in doubt use the appropriate extern "C" { } blocks)

**IMPORTANT: Caching or buffering of data is not allowed:** Any kind of data caching or buffering is not allowed. What does this mean? It means when an I/O operation is issued (write or read) you must go to the device to retrieve or write the value. The FTL and mapping related "metadata" (mostly mapping structures) can be maintained in memory buffers (as they are not persistent) but all data operations must read/write data directly from the NVMe ZNS device.

**Plan ahead:** You do not have to implement GC and support for overwriting blocks as in the current testing there are none. This will come in milestone 3 with the garbage collection support. But we advise to plan for it already in the choices of data structures you have. Also, persistence (shutdown and resume) in the milestone 5 (so plan a bit ahead).



**Figure 4:** Various milestones and their location in the overall architecture.

#### Specific goals for M2:

1. Design and implement a hybrid log-structured FTL which
  - Gives user a block device to read and write at any LBA address (captured in `struct user_zns_device`)
  - Keeps track of log and data mappings

2. Make sure the m2 tests pass (execute `./bin/m2` with various arguments) and try out different parameters that the test takes (execute the binary with `-h` flag and read the source code).
3. Commit the code in **the “test” branch of your repo** (this is where we will pull the code, not the master/main) and upload the git commit hash in the canvas. Check the upload deadline on canvas.

**Don’t do stupid designs:** Every year we have discussions with students regarding what kind of design is allowed or not. We do not believe in telling you very specifically what is allowed and what is not allowed. This is the beauty of doing systems work. However, we also believe that you are sensible enough to come up with a decent system. There are still stupid designs that we have seen - so please avoid them

- Allocating a giant array of a few GBs to store all the data in memory and not writing anything to the log. This called buffering is not allowed. Each I/O call must go to the device to read and write data.
- You are only allowed to use the standard C++/STL library code. We will not install anything additionally on our server than what is mentioned in the preparation of milestone 2.
- Do not change the gcc compiler flags.
- Do not touch the files containing the `main()` function.
- Does your design have some fundamental limitations? How many bits are sufficient? How much memory would you need? Etc. Systems work is about well thought through, but reasonably running practical code.

When in doubt - ask us!

#### Successful example run for the m2 milestone

```
$ make -j
# pass the ZNS device
$ sudo ./bin/m2 -d /dev/nvme0n1 -l 8 -s 2097152
=====
This is M2. The goal of this milestone is to implement a hybrid log-structure ZTL
(Zone Translation Layer) on top of the ZNS (no GC)
=====
parameter settings are: device-name nvme0n1 log_zones 8 gc-watermark 1 force-reset
yes
<...>
=====
Milestone 2 results
[stosys-result] Test 1 (write, read, and match on LBA0)      : Passed
[stosys-result] Test 2 (1792 LBA write, read, match)        : Passed
[stosys-result] Test 3 (large writes, read, match)          : Passed
=====
[stosys-stats] The elapsed time is 538 milliseconds
=====
```

### 4.3 Milestone 3 (week 3) : It’s 2023, we recycle

**Type:** group project

**Maximum points:** 10

**What to do:** implement the GC logic in the FTL; (ii) upload the final code commit git hash from the test branch in the canvas assignment; and (iii) show up for the interview

**Where:** code development in the given framework, and on the campus interview

**Deadline:** see the course canvas page

In the previous milestone, we wrote the LBA address 0 and a random offset of some LBAs. We try to make sure we do not write the same address again or overflow the log zone. In this milestone you will build a GC capable FTL. With the GC, the device will become a completely general purpose read/write-able block device. The GC logic is completely FTL's internal business and no API changes are necessary to the `zns_device.h` file that is developed in the Milestone 2.

When does a GC occur? (1) when you run out of the log blocks; (2) during active wear leveling. In this Milestone you can focus on the case (1) only.

**GC Design choices:** Consider what GC algorithm you want to implement, how you want to ensure ordering, what kind of information you want to track about hot-cold pages or zones. Even though you are given a minimum `-w` watermark threshold, this is the absolute minimum number of zones you need to keep free. You are always free to start GC a bit early. Is it a good design? Or you rather wait for the minimum number of free zones. How many log zones do you want to clean, 1 or more?

**Important:** You are NOT allowed to implement an in-place GC, where GC only erases a block/zone whenever a new write comes in the same user-thread. This way a new write could result in an erase-merge-write cycle. The GC MUST be implemented as a separate thread in your FTL, and run independently from the rest of the user-driven I/O requests. Of course, that means that the GC thread must take the right locks and synchronize with other threads not to erase log zones when they are being concurrently written or read. This overall thread-safe architecture will keep in handy in Milestone 4 and 5 as well.

**Testing:** The testing framework is similar to milestone 2. Specifically, the tests are given in the `m3.cpp` file. Please read through the file to understand the intent of the tests. In the file, we have following tests:

1. Write, read, and match pattern over the whole device sequentially once (from LBA 0 to max LBA)
2. Write, read, and match pattern over the whole device randomly once
3. Write, read, and match pattern over the whole device by randomly hammering certain LBA locations for "x" number of times (x is passed as the parameter `-o`)

While writing the data to the ZNS device, the test also writes data in parallel to a local shadow file and then compares the results from what was read from the device and what is in the file. The file is deleted at the end of the test. If you want to hold on to the file, feel free to modify the test program and remove the `remove_file` call.

Feel free to add your own tests.

**Don't do stupid designs:** Every year we have discussions with students regarding what kind of design is allowed or not. We do not believe in telling you very specifically what is allowed and what is not allowed. This is the beauty of doing systems work. However, we also believe that you are sensible enough to come up with a decent system. There are still stupid designs that we have seen - so please avoid them

- Allocating a giant array of a few GBs to store all the data in memory and not writing anything to the log

- Not doing any gc until the log is full. Then draining all the log to the data zones. Then writing again the whole log. This kind of start-stop world is not a good design. Clean “a few” zones in a one go, and continue to let a user write.

When in doubt - ask us! You are given all the configuration parameters (-l, -w, and -o test with them thoroughly).

### Specific goals for M3:

1. Enhance your FTL design with a GC, with the GC the ZNS device can be read from or written to any number of times
2. Make sure the m3 tests in the framework pass (execute `./bin/m3 -d /dev/nvme0n1`)
3. Commit the code in **the “test” branch of your repo** (this is where we will pull the code, not the master/main) and upload the git commit hash in the canvas. Check the upload deadline on canvas.
4. **Show up for the interview**, give demo, and explain you code

### Example successful execution of m3

```

user@stosys-qemu-vm:~$ sudo ./bin/m3 -d /dev/nvme0n1
=====
This is M3. The goal of this milestone is to implement a hybrid log-structure ZTL
(Zone Translation Layer) on top of the ZNS WITH a GC
                ^^^^^^^^^^
=====
parameter settings are: device-name nvme0n1 log_zones 3 gc-watermark 1 force-reset
yes hammer-time 10000
Opening the namespace device: /dev/nvme0n1
device nvme0n1 opened successfully 3
<...>
fallocate OK with ./tmp-output-fulld and size 0x1d00000
the ZNS user device has been written (ONCE) completely OK
verifying the content of the ZNS device ....
Verification passed on the while device
fallocate OK with ./tmp-output-fulld and size 0x1d00000
the ZNS user device has been written (ONCE) completely OK
verifying the content of the ZNS device ....
Verification passed on the while device
fallocate OK with ./tmp-output-fulld and size 0x1d00000
the ZNS user device has been written (ONCE) completely OK
Hammering some random LBAs 10000 times
Hammering done, OK for 10000 times
verifying the content of the ZNS device ....
Verification passed on the while device
the number of GC events 1391
=====
Milestone 3 results
[stosys-result] Test 1 sequential write, read, and match (full device)
: Passed

```

```
[stosys-result] Test 2 randomized write, read, and match (full device)
: Passed
[stosys-result] Test 3 randomized write, read, and match (full device, hammer 10000
) : Passed
=====
[stosys-stats] The elapsed time is 85264 milliseconds
=====
user@stosys-qemu-vm:~$
```

## 4.4 Milestone 4 (week 5) : We love Rock(sDB) 'n' Roll!

**Type:** group project

**Maximum points:** 10

**What to do:** integrate your FTL with a file system design into RocksDB key-value store, and upload the final code commit git hash from the test branch in the canvas assignment

**Where:** in the framework

**Deadline:** see the course canvas page

**What happens if I could not complete M3?** There is a protocol in place where you can develop M4 and M5 on top of your local host file system (ext4, or F2F2). However, your maximum points will be restricted to the 50% of the maximum attainable grades. Please talk to us and inform us if you would like to exercise this option. This decision must be made around the M3 completion deadline.

**IMPORTANT:** Since RocksDB FS API is in C++, use C++ as the choice of language. Be aware, *use the CC file extension, not CPP here*. I need to look through it systematically but I had issues with cpp file extension not compiling inside RocksDB. To be on the safer side, just use the extension cc for this milestone.

At this point, you (should) have a working FTL implementation with GC. It is time to put it to a good use and integrate into a widely used key-value store called RocksDB. RocksDB is an embedded key-value store, i.e., it provides an efficient KV interface for users to store data locally. RocksDB uses LSM trees to store and read key-value data. More about it can be found at <https://github.com/facebook/rocksdb/wiki>

It is a very large project with many advanced topics. Naturally, we cannot cover all these details in a short 5 week class project. In this milestone, we are going to focus on writing a file system engine for it on top of the ZNS device with FTL that you developed in the previous milestone.

If you enjoy such project building see our work-in-progress research key-value store, **TropoDB**, <https://github.com/stonet-research/TropoDB> (ask Krijn ask about it)

### Step 1: Get and compile RocksDB with a file system hook for your filesystem

Install all RocksDB dependencies:

<https://github.com/facebook/rocksdb/blob/main/INSTALL.md#dependencies>

```

$ sudo apt install autoconf automake libgflags-dev libtool autoconf-archive m4
$ sudo apt install libsnpappy-dev zlib1g-dev libbz2-dev liblz4-dev libzstd-dev

$ git clone https://github.com/westerndigitalcorporation/libzbd.git
$ cd libzbd

# I tested the code with this
$ git checkout e9c593685e1c957bc903d3b937dffc7f15710324

$ sh ./autogen.sh
$ ./configure
$ make
$ sudo make install

```

Then we need to get RocksDB and make it aware of our FTL/file system code.

```

$ git clone https://github.com/facebook/rocksdb.git
$ cd rocksdb

# may be the master would also work, but I tested the code with this
$ git checkout c0ec58ecb94d276f69ccdfa956c5152d4c91a2ec

# now we need to copy the folder content from the framework into rocksdb plugin
# folder (update the paths accordingly what you have on your machine)
$ mkdir -p plugin/s2fs
$ cp -vr ~/msc-stosys-framework/src/rocksdb-plugin/* ./plugin/s2fs/

# now we compile (below is one line command)
$ DEBUG_LEVEL=0 ROCKSDB_PLUGINS="s2fs" USE_RTTI=1 PREFIX=/home/$USER/local/
DISABLE_JEMALLOC=1 DISABLE_WARNING_AS_ERROR=1 make -j4 db_bench install
...
[at the end]
CC      utilities/transactions/lock/range/range_tree/range_tree_lock_tracker.o
echo 'prefix=/home/atr/local/' >> rocksdb.pc
echo 'exec_prefix=${prefix}' >> rocksdb.pc
echo 'includedir=${prefix}/include' >> rocksdb.pc
echo 'libdir=/home/atr/local//lib' >> rocksdb.pc
echo '' >> rocksdb.pc
echo 'Name: rocksdb' >> rocksdb.pc
echo 'Description: An embeddable persistent key-value store for fast storage' >>
rocksdb.pc
echo Version: 6.25.0 >> rocksdb.pc
echo 'Libs: -L${libdir} -ldl -lrocksdb' >> rocksdb.pc
echo 'Libs.private: -lpthread -lrt -ldl -lsnappy -lgflags -lz -libz2 -llz4 -lzstd
-lnuma -u stosys_s2fs_reg' >> rocksdb.pc
echo 'Cflags: -I${includedir} -std=c++11 -faligned-new -DHAVE_ALIGNED_NEW
-DROCKSDB_PLATFORM_POSIX -DROCKSDB_LIB_IO_POSIX -DOS_LINUX -fno-builtin-memcmp
-DROCKSDB_FALLOCATE_PRESENT -DSNAPPY -DGFLAGS=1 -DZLIB -DBZIP2 -DLZ4 -DZSTD -DNUMA
-DROCKSDB_MALLOC_USABLE_SIZE -DROCKSDB_PTHREAD_ADAPTIVE_MUTEX -DROCKSDB_BACKTRACE
-DROCKSDB_RANGESYNC_PRESENT -DROCKSDB_SCHED_GETCPU_PRESENT
-DROCKSDB_AUXV_GETAUXVAL_PRESENT -march=native -DHAVE_SSE42 -DHAVE_PCLMUL
-DHAVE_AVX2 -DHAVE_BMI -DHAVE_LZCNT -DHAVE_UINT128_EXTENSION

```

```

-DROCKSDB_SUPPORT_THREAD_LOCAL -isystem third-party/gtest-1.8.1/fused-src' >>
rocksdb.pc
install -d /home/atr/local/lib
install -d /home/atr/local/lib/pkgconfig
for header_dir in `find "include/rocksdb" -type d`; do \
    install -d //home/atr/local/$header_dir; \
done
for header in `find "include/rocksdb" -type f -name *.h`; do \
    install -C -m 644 $header //home/atr/local/$header; \
done
for header in ; do \
    install -d //home/atr/local/include/rocksdb/`dirname $header`; \
    install -C -m 644 $header //home/atr/local/include/rocksdb/$header; \
done
install -C -m 644 rocksdb.pc /home/atr/local/lib/pkgconfig/rocksdb.pc
GEN      util/build_version.cc
CC       util/build_version.o
AR       librocksdb.a
/usr/bin/ar: creating librocksdb.a
install -d /home/atr/local/lib
CCLD    db_bench
install -C -m 755 librocksdb.a /home/atr/local/lib
[ -e librocksdb.so.6.25.0 ] && make install-shared || :

```

**Important:** Do not give RocksDB compilation more than 2 or 4 threads, it will freeze your VM.

**What is RTTI: Run time type information.** gcc flag, `-fno-rtti`, disables generation of information about every class with virtual functions for use by the C++ runtime type identification features (`dynamic_cast` and `typeid`). If you don't use those parts of the language, you can save some space by using this flag. Note that exception handling uses the same information, but it will generate it as needed. The `dynamic_cast` operator can still be used for casts that do not require runtime type information, i.e. casts to `void*` or to unambiguous base classes.

Also we are disabling the use of `jemalloc`, a multicore scalable malloc implementation useful in multithreaded high performance implementations. Typically it would be helpful to debug some memory leaks or segmentation faults. Disabling it is not strictly needed, but I did have to debug some hard memory errors in my code so I am recommending to disable it for now.

**What is happening in the code:** We integrated a skeleton code in the RocksDB compilation that registers a function with an URI API into RocksDB library. If you browse the copied (single) file [rocks\\_s2fs.cc](#):

```

extern "C" FactoryFunc<FileSystem> stosys_s2fs_reg;
FactoryFunc<FileSystem> stosys_s2fs_reg =
    ObjectLibrary::Default()->Register<FileSystem>(
        "s2fs:.*://.*", [](const std::string &uri, std::unique_ptr<FileSystem>
            *ret_fs,
            std::string *errmsg) {
                ret_fs->reset(FileSystem::Default().get());
                return ret_fs->get();
            });

```

```
});
```

The code only does two useful things (1) it tells the RocksDB library that the `stosys_s2fs_reg` factory function exists (it is compiled and linked with `--as-needed` and with `-u` flags. From the gcc man page, `-u` symbol: “Pretend the symbol symbol is undefined, to force linking of library modules to define it. You can use `-u` multiple times with different symbols to force loading of additional library modules.” (<https://man7.org/linux/man-pages/man1/gcc.1.html>, search for “`-u` symbol”). This forces the compiler/linker to link our `libstosys.so` library. (2) tell RocksDB what is the URI of our new file system for which the calls will be forwarded to our code. In this case, it is `"s2fs:.*://.*"`. It follows the RegEX syntax, hence, `s2fs:.*` (one or more parameters, here you can pass the device name), then `“://”` the separator, and the rest of the path name. In the current code, it returns the `FileSystem::Default()` (which is your normal POSIX backed FS). In the given framework to you, there is a `S2FileSystem` class defined with empty function. Please implement your functions there.

### Step 2: enable compiling milestone 4 in CMakeList.txt file

After you have installed the RocksDB, you should enable compiling Milestone 4 and 5 in your `CMakeLists.txt` file by

```
# Project configuration specific parameters
set(STOSYS_M45 OFF)
set(STOSYS_CMAKE_DEBUG OFF)
set(STOSYS_ASAN ON)
```

Set the `STOSYS_M45` to ON

```
set(STOSYS_M45 ON)
```

Then make sure to

```
# in the msc-stosys-framework directory
$ make clean
$ cmake . ← redo the CMake after modifying the CMakeFile
$ make -j
```

In case making or linking problems occur, try to clean the make and cmake build files and redo the build steps.

```
# delete all build files
$ make clean
$ rm -rf cmake-build-debug/ CMakeFiles/ CMakeCache.txt cmake_install.cmake
Makefile
```

### Step 3: Enable your file system implementation (S2FileSystem) and complete the file system code

You are given a code base that compiles. It compiles a testing framework ([m45\\_main.cc](#)) and two file system implementations (1) a DummyFSForward file system and (2) S2FileSystem.cc file system. The DummyFS file system just forwards all the calls to the default file system implementation which is your normal POSIX file system wherever you are running the RocksDB. You can use this file system to see what is being passed, what semantics are expected, and how the calls should behave, in case it is not immediately clear from the description.

In case you want to browse the RocksDB POSIX file system implementation it is at: [https://github.com/facebook/rocksdb/blob/c0ec58ecb94d276f69ccdfa956c5152d4c91a2ec/env/fs\\_posix.cc](https://github.com/facebook/rocksdb/blob/c0ec58ecb94d276f69ccdfa956c5152d4c91a2ec/env/fs_posix.cc) (see the class PosixFileSystem : public FileSystem implementation and calls).

When you run the given code after compiling from the last step, you will see something like

```
$ sudo ./bin/m45 -p s2fs:nvme0n1:///tmp/testdb/ -e 1 -k 10 -v 10
=====
Welcome to milestone 4/5, which is about integration into RocksDB (also:
congratulations for bearing with us so far!)
=====
=====
test_uri s2fs:nvme0n1:///tmp/testdb/ , shadow uri posix:///tmp/shadowdb
entries 1, each entry ksize 10 bytes, vsize 10 bytes, readonly verify : 0 deleteall
: 0 single : 0
=====
Opening database at posix:///tmp/shadowdb with uri posix and the db_path as
/tmp/shadowdb
Environment from URI posix:///tmp/shadowdb for FS Posix File System
## Database opened at posix:///tmp/shadowdb db name is /tmp/shadowdb , attached FS
is --> Posix File System<--
Opening database at s2fs:nvme0n1:///tmp/testdb/ with uri s2fs:nvme0n1 and the
db_path as /tmp/testdb/
Initialization uri is s2fs:nvme0n1:///tmp/testdb/ and errmsg:
Environment from URI s2fs:nvme0n1:///tmp/testdb/ for FS DummyFSForward to Posix
File System
  call_seq: 0 tid: 14034391480468988153  func: GetAbsolutePath line: 350
  call_seq: 1 tid: 14034391480468988153  func: CreateDirIfMissing line: 236
<...>
  call_seq: 47 tid: 14034391480468988153  func: DeleteFile line: 212
  call_seq: 48 tid: 14034391480468988153  func: GetChildren line: 187
## Database opened at s2fs:nvme0n1:///tmp/testdb/ db name is /tmp/testdb/ ,
attached FS is --> DummyFSForward to Posix File System<--
Preparing the map to insert values for 1 entries, max size 20
a testmap is filled with 1 of max_size 20 each
All values inserted, number of entries 1, expected size stores would be 20 bytes
Starting data reading from the shadow db at posix:///tmp/shadowdb | single 0
*****
OK: all 1 values matched successfully
*****
  call_seq: 49 tid: 14034391480468988153  func: GetChildren line: 187
  call_seq: 50 tid: 14034391480468988153  func: UnlockFile line: 321
database(s) closed, test is done OK
=====
[stosys-stats] The elapsed time is 98 milliseconds
=====
```

You can see that the dummy FS just displays a line with the call sequence and thread id, and just forwards the calls to the underlying POSIX file system.

Now you should enable your own S2FileSystem implementation by changing `rocks_s2fs.cc`:

```
if(false) {
    S2FileSystem *z = new S2FileSystem(uri, true);
    ret_fs->reset(z);
} else {
    //DummyFSForward is forwarding implementation - I should be left in peace
    class DummyFSForward *m = new DummyFSForward();
    ret_fs->reset(m);
}
```

Change **false** to **true** in the if condition to get calls to your file system (sorry I did not have a much smarter way to do it for now). The constructor of the file system is filled out for you to give you an example of how parameters could be passed to it from URI. The constructor currently initializes the ZNS device using the libstosys API. The majority of the file system calls are left unimplemented. These are the ones that you have to fill out.

What choices you have to make:

1. What is the file system layout
2. How do you save the file system metadata and data
  - a. Feel free to choose a log-structured file system or a simple FAT/inode style. See the design of ext2, which is quite simple and doable (and a popular choice among the past students).
  - b. What is the largest file size or file name you can support
  - c. How do you keep track of blocks allocated
  - d. How do you keep track of file structure
3. How do you plan to shutdown and resume
4. **[There is a separate special bonus, above the 10 points which are awarded for this work]** We will ask you what is the largest key-value configuration that you can run on your file system implementation and points will be awarded in a ranked manner. More details to follow later.

**Hints:**

1. You do not have to implement all the calls. Populate calls on demand when they are called.
2. RocksDB is multithreaded, so keep appropriate locks.

**Don't do stupid designs:** Every year we have discussions with students regarding what kind of design is allowed or not. We do not believe in telling you very specifically what is allowed and what is not allowed. This is the beauty of doing systems work. However, we also believe that you are sensible enough to come up with a decent system. There are still stupid designs that we have seen - so please avoid them

- Don't do a toy file system with 4kB max file sizes. Even though the development device is small we expect that you will be able to support 100MB-1GB+ files.
- **[Update Sep 5th 2023]** Do not design a file system where your file system reads all the file system metadata in at the start and only does data I/O to the device. **This is considering buffering and caching and those are not allowed.** At the completion of each I/O call, all user data and file system metadata should be stored in the device.

**Step 4: Run the test program**

Here is an example of how you can run the test program:

```
$ sudo ./bin/m45 -p s2fs:nvme0n1:///tmp/atr1/
Welcome to milestone 4/5, which is integration into RocksDB
=====
test_uri s2fs:nvme0n1:///tmp/atr1/ , shadow uri posix:///tmp/shadowdb
entries 1, each entry ksize 10 bytes, vsize 90 bytes, readonly verify : 0 deleteall
: 0
=====
Opening database at posix:///tmp/shadowdb with uri posix and the db_path as
/tmp/shadowdb
Environment from URI posix:///tmp/shadowdb for FS Posix File System
## Database opened at posix:///tmp/shadowdb db name is /tmp/shadowdb , attached FS
is --> Posix File System<--
Opening database at s2fs:nvme0n1:///tmp/atr1/ with uri s2fs:nvme0n1 and the db_path
as /tmp/atr1/
[...]

Preparing the map to insert values for 1 entries, max size 100
a testmap is filled with 1 of max_size 100 each
All values inserted, number of entries 1, expected size stores would be 100 bytes
Starting data reading from the shadow db at posix:///tmp/shadowdb | single 0
*****
OK: all 1 values matched successfully
*****
database(s) closed, test is done OK
```

**What does the test program (m45) do and what parameters does the test program takes:**

```
$ sudo ./bin/m45 -h
Welcome to milestone 4/5, which is integration into RocksDB
options: -p val -t val -r -s val -e val -d -h

-p test_db_uri, example: -p posix:///tmp/testdb or -p s2fs:nvme0n1:///tmp/atr1/ ,
delimited is :// (3 char)

-t is test/shadow DB which will be populated in parallel with the same keys as the
testdb and is used when using roverify in ro mode

-r says do a readonly verification from the shadowdb, it is assumed that the shadow
db is primed before by running this program once

-k integer: max size of the key

-v integer: max size of the value

-e integer: total number of entries to test for

-D delete the databases at the test and shadow locations
```

```
-S do a single DB test on the test_db with -p, no comparison
-d show debugging information
-h show this help
```

Just like in milestone 3, the test program here creates a shadowDB on the side using the default POSIX file system that you are running. Any key-value that is inserted in the KV store running on top of the S2FS file system, is also inserted in the POSIX file system. In the end they are scanned and matched. Some options that you want to know about

- -S : does not do shadow comparison, just simply inserts and reads the value back. Can be useful if you want to debug your S2FS code with a simple test program. You can pass -e 1 and just insert and test one kv tuple.
- -k, -v and -e control the amount of data that is inserted. These are in bytes, and randomly generated strings are inserted as key-value values.
- -r does a read-only verification of the test db against a shadow db. This is used for persistence tests. It assumes that there have been previous runs through which data has been inserted in the two databases.

#### Specific goals for M4:

1. Design and implement a file system API implementation on top of ZNS device in RocksDB KV store using your FTL-enabled device
2. Pass the test conditions. Please test what is the maximum number of keys you can test with 10+10 bytes K+V tests. Naturally the larger key or value sizes means more space on the ZNS device.
3. Commit the code in **the “test” branch of your repo** (this is where we will pull the code, not the master/main) and upload the git commit hash in the canvas. Check the upload deadline on canvas.

In case you run into some strange key-value pairs, try to delete the POSIX database path by hand (rm -rf).

## 4.5 Milestone 5 (week 7) : Wake up, Neo

**Type:** group project

**Maximum points:** 10

**What to do:** Make everything persistent (FTL and FS), upload the final code commit git hash from the test branch in the canvas assignment

**Where:** code development in the given framework, and on the campus interview

**Deadline:** see the course canvas page

The last milestone is simple to explain. Make everything persistent. That means the value and data that I have stored in my previous RockDB test runs, I should be able to read and verify their content. For this milestone, you need to reserve zones to write out metadata at the FTL and file system level, and then resume next time when the system starts.

**How do I test if things are working?** We are going to use the same testing setup that we used before. For example, we can insert multiple 10,000 key-values in the test and shadow db by executing the following command, and then verify them as

```
$ sudo ./bin/m45 -p s2fs:nvme0n1:///tmp/atr1/ -e 10000 -k 10 -v 10
...
Starting data reading from the shadow db at posix:///tmp/shadowdb | single 0
*****
OK: all 10000 values matched successfully
*****
database(s) closed, test is done OK

$ sudo ./bin/m45 -p s2fs:nvme0n1:///tmp/atr1/ -e 10000 -k 10 -v 10
...
Starting data reading from the shadow db at posix:///tmp/shadowdb | single 0
*****
OK: all 20000 values matched successfully
*****
database(s) closed, test is done OK

$ sudo ./bin/m45 -p s2fs:nvme0n1:///tmp/atr1/ -e 10000 -k 10 -v 10
...
Starting data reading from the shadow db at posix:///tmp/shadowdb | single 0
*****
OK: all 30000 values matched successfully
*****
database(s) closed, test is done OK

# now we run the benchmark in verify only mode as
$ sudo ./bin/m45 -p s2fs:nvme0n1:///tmp/atr1/ -r
...
Starting data reading from the shadow db at posix:///tmp/shadowdb | single 0
*****
OK: all 30000 values matched successfully
*****
database(s) closed, test is done OK
```

**Hint:** make sure you understand what is the current state of the DB and the ZNS device between the last set of writes and the new one. Whenever unsure, delete all the old databases and reset the ZNS device to start from scratch. For example, you may want to do something like:

```
# shadow db
$ rm -rf /tmp/shadowdb

# another db location used in the code, this is the default testdb
$ rm -rf /tmp/testdb

# reset the ZNS device
```

```
$ sudo nvme zns reset-zone -a /dev/nvme0n1
```

Sometimes even after a reset the ZNS device can return the old content. After the rest the zones are in EMPTY state, and read on an EMPTY state is not defined. Hence, any garbage content can be returned. In this case it happened to be the last value. So when this happens, just use nvme-cli command to write out on the desired zones, or LBA to zero them out. Or if you want to make a sensible design, when you open a zone for I/O, always check the write pointer location and never read beyond it.

You may also want to browse the file content and sizes in the POSIX /tmp location to see how it should look on your file system on ZNS for comparison.

Also check the code in S2FileSystem.cc and make sure to pass force\_reset as false when you are testing for restoring from the device.

```
//make sure to setup these parameters properly and check the forced reset flag for M5
params.name = strdup(device.c_str());
params.log_zones = 3;
params.gc_wmark = 1;
params.force_reset = true;
```

### Specific goals for M5:

1. Implement persistence so that data and context can be saved across multiple runs.
2. Pass multiple run setup as shown above.
3. You are not expected to support power-failure cases, hence you can design your implementation around an ordered shutdown.
4. Make sure that you pass the test code with some reasonably high value of key-values.
5. Upload the final code commit hash from the test branch in the canvas assignment. Check the upload deadline on canvas.
6. **Show up for the interview**, give a demo, and explain your code.

## 5. Bonus Ideas and Timeline

For the brave (or Masochists) among us we are not yet satisfied with the level of pain offered in this project, we also do offer bonuses that you can do within the project framework. Each of these bonus ideas is worth 10 points. There are no partial marks for the bonus. **The maximum points that you can achieve in bonuses is limited to 20.** You need to implement the bonus, and demonstrate that it works by implementing additional code that exercises your bonus feature implementation, and show conclusively that your implementation works. The burden of proof of your bonus code is on you, unfortunately.

**Timeline and Logistics:** send us (via email) the bonus plan and the number of experiments, and results that you intend to show. This should not be more than 1-2 pages. The idea of this write up is to tell us up

front what you will be demonstrating and how you plan to show that your bonus works. We can either give a go ahead, or ask you to delete some or include some more stuff. If it looks sufficient, we will green light and give you a timeslot for demo and explanation. **The deadline for getting a bonus greenlight is milestone 3 interview (IMPORTANT). NO EXCEPTIONS**. The bonus demo date will be after the milestone 5 interview.

Here are some ideas, however, we are open to ideas from your side as well (many of these ideas can be expanded into a MSc thesis work as well)

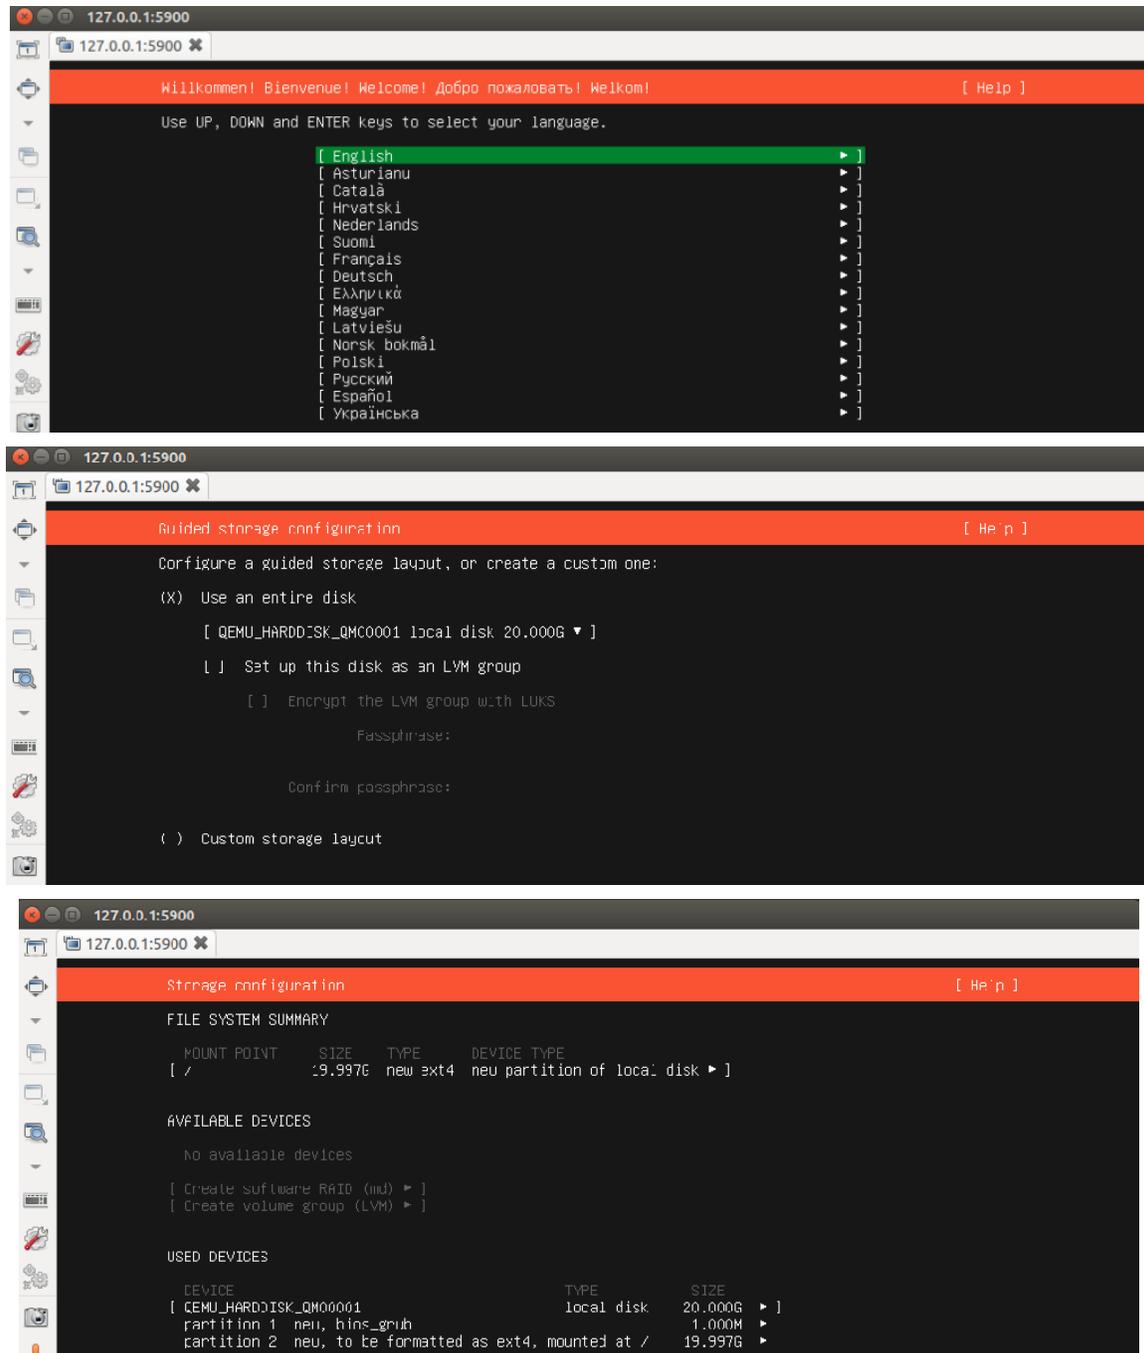
1. Deploy your code on real NVMe ZNS devices and benchmark. We have multiple NVMe ZNS devices (4TB in size). First come first serve here. We might have more but not yet confirmed.
2. Integrate another NVMe access backend (other than libnvme), for example, SPDK or xNVMe libraries. (PS~ these are asynchronous libraries, not blocking synchronous like libnvme).
3. ZNS emulation in QEMU is not persistent. If you reboot your VM then the ZNS devices come back all clean. Make it persistent.
4. Showcase multi-tenancy of your implementation by running two (or multiple) parallel RocksDB threads with the expectation that the performance should scale linearly with the number of cores.
5. Run db\_bench on your implementation, and compare and explain results with a run from a file system like F2FS or ext4 and your implementation
6. Integrate and run YCSB benchmark (<https://github.com/brianfrankcooper/YCSB>) on RocksDB on top of your Milestone 5. YCSB already has a RocksDB plugin, you need to integrate the whole thing and show that it works, and produce experimental results.
7. Integrate your FTL-device into another KV or SQL engine (MySQL, FoundationDB, ScyllaDB).
8. Implement any ideas from any papers discussed in the class, any advanced FTL design, GC algorithm, etc.
9. Design and implement a crash-consistent FTL and file system design, and show that it works by crash stopping and resuming it.
10. **Your idea?** I am curious to hear it.

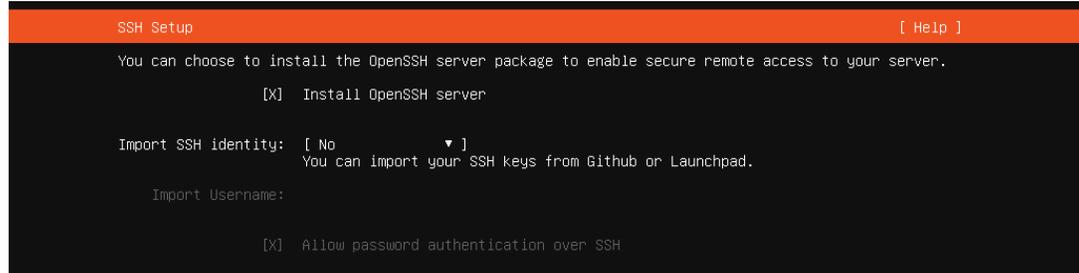
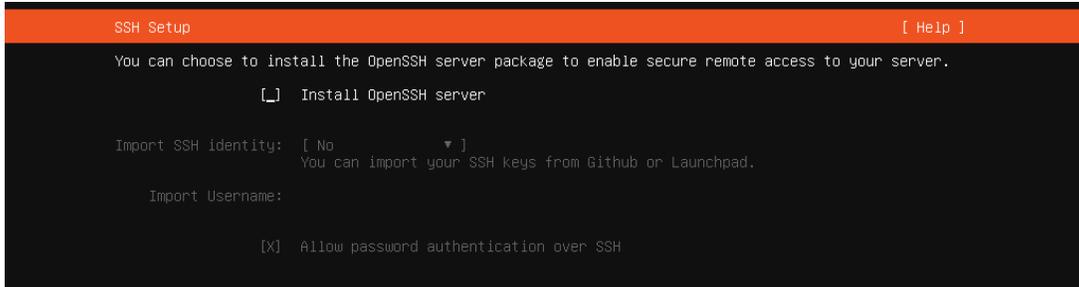
For more information please ask us in the lab hours.

## Appendix

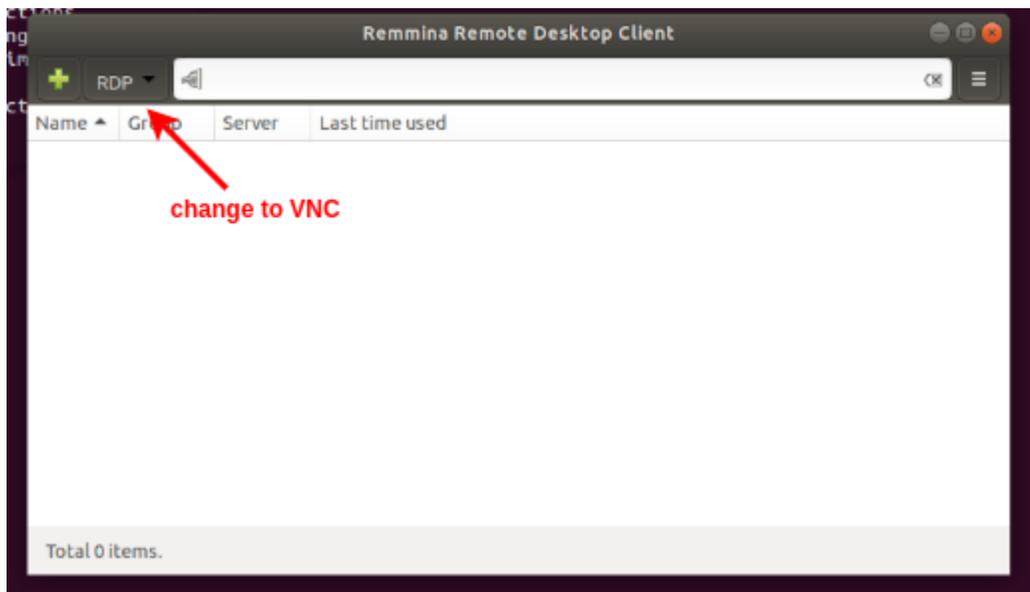
### A. Step by step guide for setting up Ubuntu installation

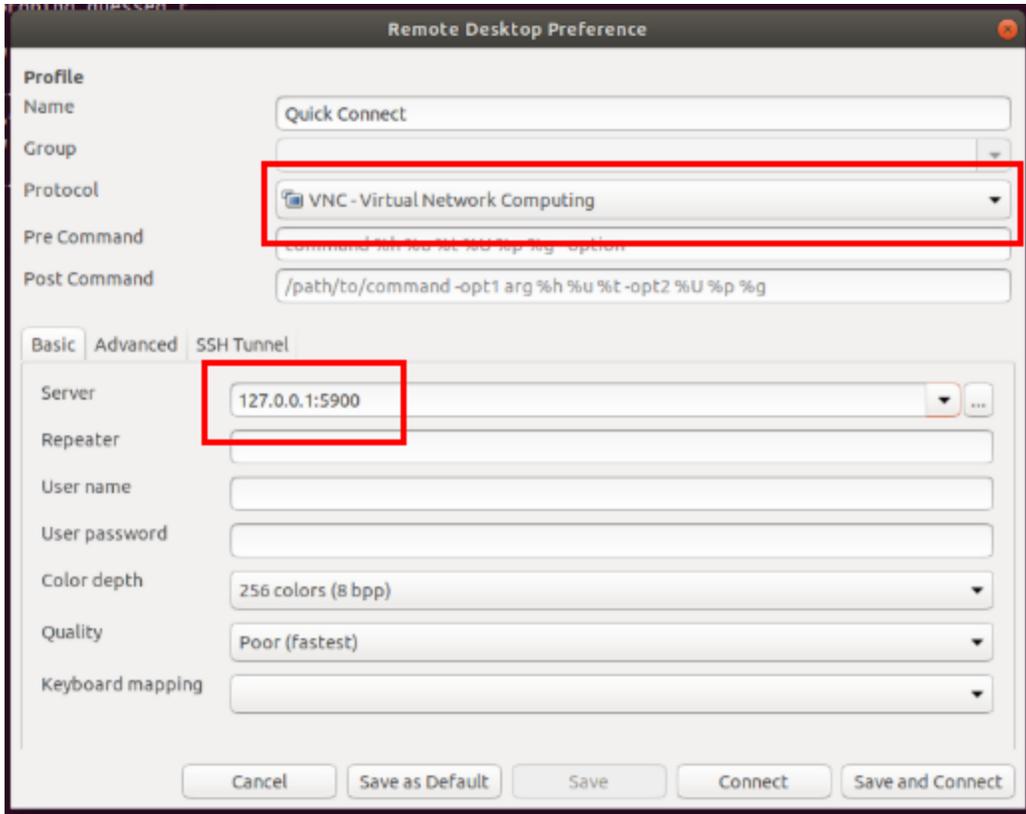
Here are “some” sequences of print screens that you should see when installing the Ubuntu OS on your image.





## B. Setup of VNC through Remmina





### C. Running inside Virtualbox, VirtualMachineManager or VMWare

If you run Virtualbox, VirtualMachineManager or VMWare the command to launch qemu-system will fail with

```
ioctl(KVM_CREATE_VM) failed: 16 Device or resource busy
qemu-system-x86_64: failed to initialize KVM: Device or resource busy.
```

This is due to a limitation in CPU virtualization extensions Intel VT-x and AMD-V. You can either disable all other virtualization software or remove `-enable-kvm` However, this will be significantly slower and you will need to specify a cpu model.

Also see: <https://ostechnix.com/how-to-enable-nested-virtualization-in-virtualbox/>