

System Observability, Analytics & Insights Platform

Team Kowalski



Members



Jake Borneman
> Team Leader
> Testing
> Sanitization



Erick Salazar
> Data Storage
> Visualization
> Pipeline



Bailey McCauslin
> Data Collection
> Testing



Nick Wiltshire
> Visual Dashboard
Manager

Client, Staff, and Mentor



Old Client

Rajpal Singh

Ex-WD R&D/
Technologist



New Client

Igor Steinmacher

NAU Associate
Professor +
Capstone
Professor



Mentor

Saisri Muttineni

NAU Computer
Science Graduate
Student

**Individually
Collect**

**Manually
Analyze**

**Store
Analysis**

Issues at Hand:

- Silent Error/Failure Detection at Kernel Level
- Limited long-term performance monitoring
- Everyone needs to be an expert

Workflow Inefficiencies:

- Manual Testing Process
- No Data Analysis Automation
- Individual Device Testing

The Problem


Solution Overview

Comprehensive Data Analytics Dashboard:

- Takes user inputs on what to scan, how long, logging, etc.
- Automates data collection and data storage
- Handles data analysis to be displayed on visual dashboard

Raw Data

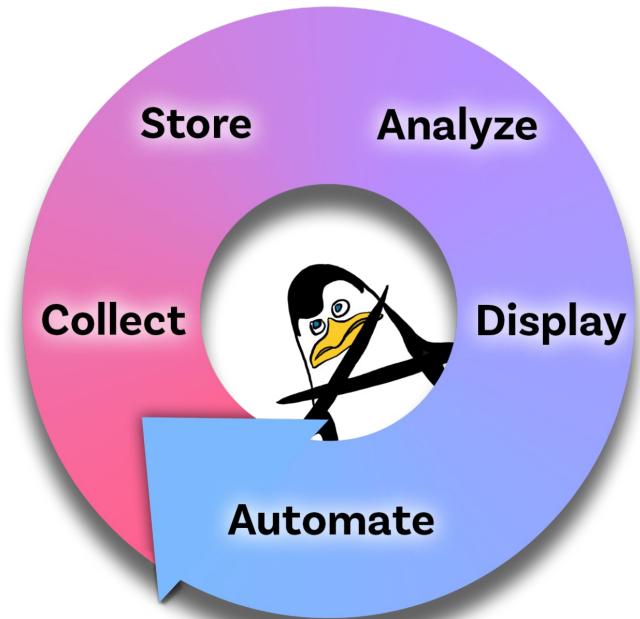
```
325
326 @usecs[3226, URL Classifier]:
327 [256, 512] 16 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@|
328 [512, 1K] 16 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@|
329 [1K, 2K] 12 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@|
330 [2K, 4K] 3 |@@@@@@@@|
331 [4K, 8K] 1 |@|
332
333 @usecs[3226, DOM Worker]:
334 [256, 512] 4 |@@@@|
335 [512, 1K] 35 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@|
336 [1K, 2K] 5 |@@@@@@|
337 [2K, 4K] 4 |@@@@|
338 [4K, 8K] 1 |@|
339
340 @usecs[1508, llvmpipe-1]:
341 [256, 512] 17 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@|
342 [512, 1K] 34 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@|
343 [1K, 2K] 5 |@@@@@@|
344 [2K, 4K] 2 |@@|
345
346 @usecs[58, kworker/u10:2]:
347 [256, 512] 16 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@|
348 [512, 1K] 18 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@|
349 [1K, 2K] 10 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@|
```




Cleaned and
Uploaded Data

Requirements Review

- Show kernel level operations to the end user through a visual dashboard.
- Store all previous kernel data for future analysis.
- Automate kernel level data collection.



Architecture Review



Implementation Review

Transmits input data to adjust
Trace Program based on selection

main.py

Trace.py

sql_manage.py

Uploads data
to DB

Flask Server

Sends Data to Visualization

Svelte
Dashboard

D3 Designs

Sanitization.py

Formatter.py

Data from Trace run through
Sanitizer and Formatter

Prototype Review

GUI Interaction and Startup

The image shows a code editor window with a terminal pane on the left and a process tree on the right.

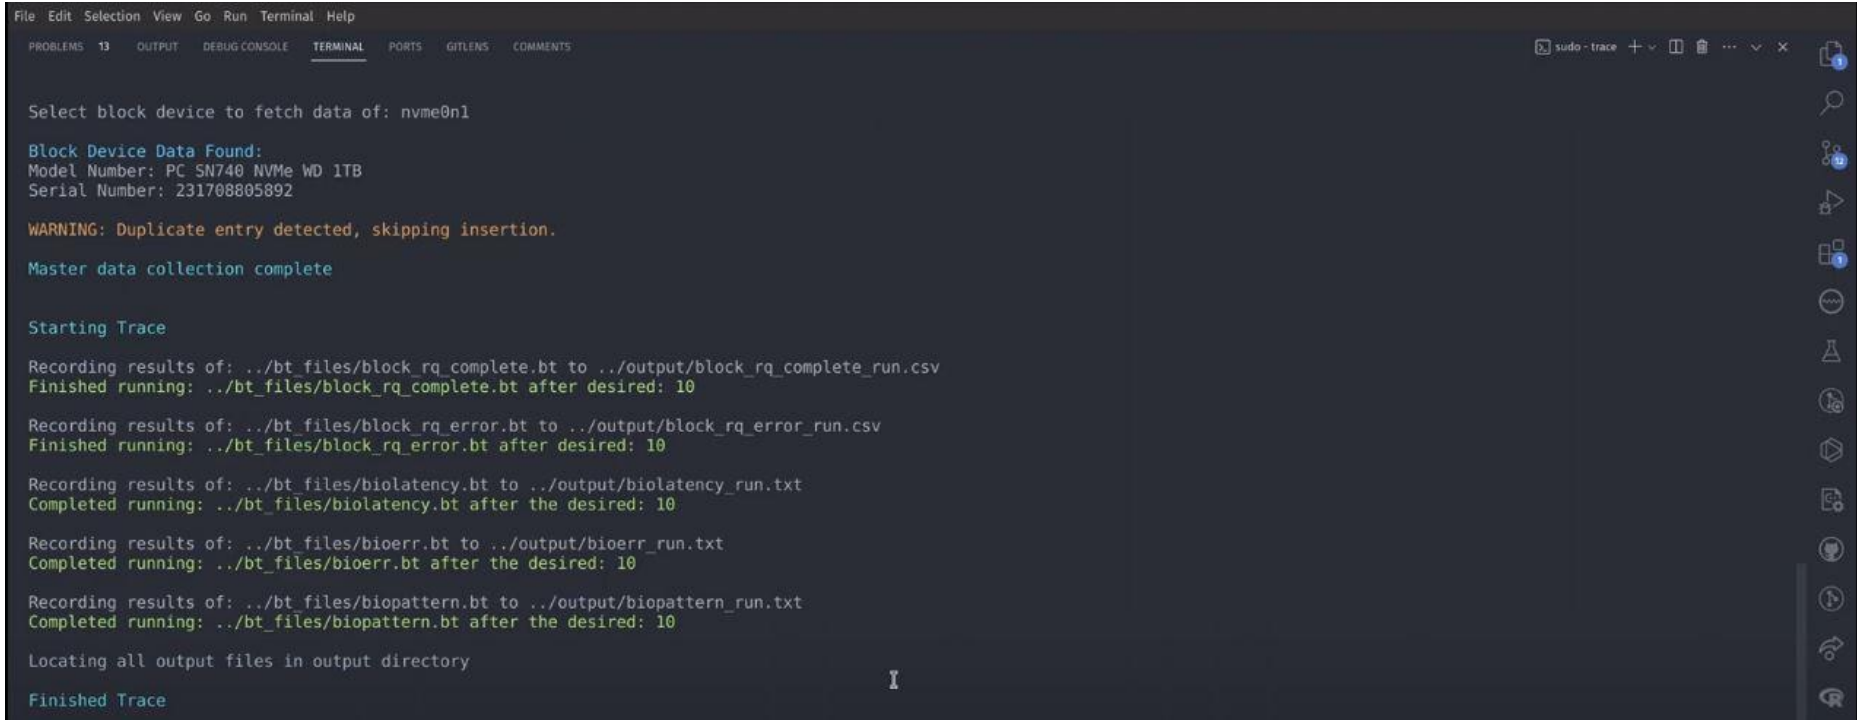
Terminal Output:

```
~/Desktop/School/CS 486C/Capstone-Group-8/soaip/trace | on pipeline !6 ?1  
2:00 PM  
> sudo python3 main.py  
Starting Flask Server and Svelte App  
> svelte-app@1.0.0 dev  
> rollup -c -w  
  
LiveReload enabled on port 35731  
  
> svelte-app@1.0.0 start  
> sirv public --no-clear --dev  
  
Your application is ready~! ✍  
- Port 8080 is taken; using 45473 instead  
- Local: http://localhost:45473  
  
Svelte app running at: http://localhost:45473  
Flask server ready at http://10.21.54.105:5000/  
Flask Server and Svelte App successfully started!  
  
Now collecting user inputs  
[]
```

Process Tree:

```
Process Time (in seconds):  
10  
Operations (comma separated):  
block.bio  
├── Verbose  
│   ├── Log  
│   └── Ops  
├── merge output  
├── Submit  
└── Upload JSON
```

Data Collection



The image shows a terminal window with a dark background and light text. The window title is "sudo - trace". The terminal output is as follows:

```
File Edit Selection View Go Run Terminal Help  
PROBLEMS 13 OUTPUT DEBUG CONSOLE TERMINAL PORTS GITLENS COMMENTS  
Select block device to fetch data of: nvme0n1  
Block Device Data Found:  
Model Number: PC SN740 NVMe WD 1TB  
Serial Number: Z31708805892  
WARNING: Duplicate entry detected, skipping insertion.  
Master data collection complete  
Starting Trace  
Recording results of: ../bt_files/block_rq_complete.bt to ../output/block_rq_complete_run.csv  
Finished running: ../bt_files/block_rq_complete.bt after desired: 10  
Recording results of: ../bt_files/block_rq_error.bt to ../output/block_rq_error_run.csv  
Finished running: ../bt_files/block_rq_error.bt after desired: 10  
Recording results of: ../bt_files/biolatency.bt to ../output/biolatency_run.txt  
Completed running: ../bt_files/biolatency.bt after the desired: 10  
Recording results of: ../bt_files/bioerr.bt to ../output/bioerr_run.txt  
Completed running: ../bt_files/bioerr.bt after the desired: 10  
Recording results of: ../bt_files/biopattern.bt to ../output/biopattern_run.txt  
Completed running: ../bt_files/biopattern.bt after the desired: 10  
Locating all output files in output directory  
Finished Trace
```

Data Sanitization/Formatting

```
Locating all output files in output directory
```

```
Finished Trace
```

```
Begun Sanitizing Files
```

```
Sanatizing File: biopattern_run.txt
```

```
Error: list index out of range
```

```
Error: list index out of range
```

```
Error: list index out of range
```

```
Error: list index out of range
```

```
Error: list index out of range
```

```
Error: list index out of range
```

```
Error: list index out of range
```

```
Attempting to remove past file version
```

```
Successfully removed the past file version
```

```
DONE Sanatizing File: biopattern_run.txt
```

```
Sanatizing File: biolatency_run.txt
```

```
Data has been written to ./output/biolatency_run.csv
```

```
Attempting to remove past file version
```

```
Successfully removed the past file version
```

```
DONE Sanatizing File: biolatency_run.txt
```

```
Sanatizing File: block_rq_complete_run.csv
```

```
DONE Sanatizing File: block_rq_complete_run.csv
```

```
Sanatizing File: bioerr_run.txt
```

```
Attempting to remove past file version
```

```
Successfully removed the past file version
```

```
DONE Sanatizing File: bioerr_run.txt
```

```
Sanatizing File: block_rq_error_run.csv
```

```
DONE Sanatizing File: block_rq_error_run.csv
```

```
Finished Sanitizing Files
```

```
442 @usecs[29, ksfortind/2]:
443 [4, 8] 5 |e
444 [8, 16] 6 |e
445 [16, 32] 14 |eeee
446 [32, 64] 15 |eeee
447 [64, 128] 17 |eeee
448 [128, 256] 138 |eeee
449 [256, 512] 41 |eeee
450 [512, 1K] 53 |eeee
451 [1K, 2K] 28 |eeee
452 [2K, 4K] 52 |eeee
453 [4K, 8K] 20 |eeee
454 [8K, 16K] 6 |e
455 [16K, 32K] 3 |
456
457 @usecs[35, ksfortind/3]:
458 [4, 8] 23 |eeee
459 [8, 16] 19 |eeee
460 [16, 32] 13 |eee
461 [32, 64] 11 |eee
462 [64, 128] 144 |eeee
463 [128, 256] 178 |eeee
464 [256, 512] 57 |eeee
465 [512, 1K] 74 |eeee
466 [1K, 2K] 40 |eeee
467 [2K, 4K] 56 |eeee
468 [4K, 8K] 12 |eee
469 [8K, 16K] 7 |ee
470 [16K, 32K] 1 |
471
472 @usecs[41, ksfortind/4]:
473 [4, 8] 7 |ee
474 [8, 16] 12 |eeee
475 [16, 32] 21 |eeee
476 [32, 64] 8 |eee
477 [64, 128] 121 |eeee
478 [128, 256] 128 |eeee
479 [256, 512] 77 |eeee
480 [512, 1K] 134 |eeee
481 [1K, 2K] 51 |eeee
482 [2K, 4K] 83 |eeee
483 [4K, 8K] 62 |eeee
484 [8K, 16K] 20 |eeee
485 [16K, 32K] 0 |
486 [32K, 64K] 2 |
487
488 @usecs[0, swapper/4]:
489 [256, 512] 586 |eeee
490 [512, 1K] 985 |eeee
491 [1K, 2K] 261 |eeee
492 [2K, 4K] 171 |eeee
493 [4K, 8K] 64 |eee
494 [8K, 16K] 32 |e
495 [16K, 32K] 3 |
496 [32K, 64K] 2 |
```

```
36 usecs,3226,DgUthr-Pool #1,[[['256', '512', 1], ['512', '1K', 1]]]
37 usecs,3226,Backgr-Pool #2,[[['8K', '16K', 2]]]
38 usecs,3224,bpfftrace,[[['4K', '8K', 2]]]
39 usecs,3564,StreamTrans #13,[[['2K', '4K', 2]]]
40 usecs,3226,DNS Resolver #1,[[['256', '512', 1], ['512', '1K', 1]]]
41 usecs,3183,Xwayland,[[['4K', '8K', 1], ['8K', '16K', 1]]]
42 usecs,3226,DgUthr-Pool #2,[[['256', '512', 2], ['512', '1K', 0], ['1K', '2K', 1]]]
43 usecs,3204,13,[[['512', '1K', 2], ['1K', '2K', 0], ['2K', '4K', 0], ['4K', '8K', 1]]]
44 usecs,3226,StreamTrans #15,[[['256', '512', 1], ['512', '1K', 1], ['1K', '2K', 1]]]
45 usecs,3226,QuotaManager 10,[[['256', '512', 1], ['512', '1K', 0], ['1K', '2K', 2]]]
46 usecs,3226,saldp-lite #1,[[['2K', '4K', 1], ['4K', '8K', 1], ['8K', '16K', 1]]]
47 usecs,3410,Socket Process,[[['512', '1K', 2], ['1K', '2K', 1]]]
48 usecs,3226,RemotelStream,[[['256', '512', 1], ['512', '1K', 3]]]
49 usecs,3226,gleam.init,[[['256', '512', 3], ['512', '1K', 1]]]
50 usecs,3226,TaskController #2,[[['512', '1K', 1], ['1K', '2K', 3]]]
51 usecs,3226,IPC launch,[[['256', '512', 1], ['512', '1K', 2], ['1K', '2K', 1]]]
52 usecs,3226,TaskController #4,[[['512', '1K', 4]]]
53 usecs,3564,WebExtensions,[[['512', '1K', 3], ['1K', '2K', 2]]]
54 usecs,3226,SSL Cert #2,[[['256', '512', 1], ['512', '1K', 1], ['1K', '2K', 3]]]
55 usecs,3226,StreamTrans #3,[[['256', '512', 2], ['512', '1K', 3]]]
56 usecs,3226,StreamTrans #1,[[['512', '1K', 2], ['1K', '2K', 3]]]
57 usecs,1508,pool-gnome-shell,[[['512', '1K', 4], ['1K', '2K', 2]]]
58 usecs,3226,IPDL Background,[[['256', '512', 1], ['512', '1K', 1], ['1K', '2K', 1], ['2K', '4K', 0], ['4K', '8K', 3]]]
59 usecs,182,libd2-sda3-8,[[['128', '256', 1], ['256', '512', 2], ['512', '1K', 1], ['1K', '2K', 0], ['2K', '4K', 2]]]
60 usecs,3226,SSL Cert #1,[[['256', '512', 1], ['512', '1K', 7]]]
61 usecs,3226,Backgr-Pool #1,[[['256', '512', 3], ['512', '1K', 6], ['1K', '2K', 0], ['2K', '4K', 1]]]
62 usecs,3226,gnome-bin,[[['512', '1K', 4], ['1K', '2K', 2], ['2K', '4K', 6], ['4K', '8K', 2]]]
63 usecs,3434,Privileged Cont,[[['256', '512', 1], ['512', '1K', 3], ['1K', '2K', 4], ['2K', '4K', 1], ['4K', '8K', 1]]]
64 usecs,348,kworker/u10:8,[[['256', '512', 2], ['512', '1K', 7], ['1K', '2K', 3], ['2K', '4K', 1]]]
65 usecs,3226,Renderer,[[['256', '512', 6], ['512', '1K', 8], ['1K', '2K', 1], ['2K', '4K', 1]]]
66 usecs,3226,IndexedDB #1,[[['256', '512', 1], ['512', '1K', 0], ['1K', '2K', 3], ['2K', '4K', 6], ['4K', '8K', 0], ['8K', '16K', 1]]]
67 usecs,79,worker/u10:4,[[['128', '256', 3], ['256', '512', 8], ['512', '1K', 12], ['1K', '2K', 3], ['2K', '4K', 1]]]
68 usecs,3226,Cached I/O,[[['256', '512', 4], ['512', '1K', 13], ['1K', '2K', 6], ['2K', '4K', 0], ['4K', '8K', 1]]]
69 usecs,18,worker/u10:8,[[['256', '512', 8], ['512', '1K', 11], ['1K', '2K', 4], ['2K', '4K', 4], ['4K', '8K', 1]]]
70 usecs,3226,Socket Thread,[[['256', '512', 10], ['512', '1K', 16], ['1K', '2K', 5], ['2K', '4K', 1], ['4K', '8K', 3]]]
71 usecs,1508,llvmpipe-0,[[['256', '512', 4], ['512', '1K', 24], ['1K', '2K', 3], ['2K', '4K', 3], ['4K', '8K', 1]]]
72 usecs,1508,llvmpipe-3,[[['128', '256', 4], ['256', '512', 16], ['512', '1K', 22], ['1K', '2K', 4]]]
73 usecs,1508,llvmpipe-4,[[['128', '256', 4], ['256', '512', 19], ['512', '1K', 14], ['1K', '2K', 4], ['2K', '4K', 1]]]
74 usecs,3226,URL Classifier,[[['256', '512', 16], ['512', '1K', 16], ['1K', '2K', 12], ['2K', '4K', 5], ['4K', '8K', 1]]]
75 usecs,3226,DOM Worker,[[['256', '512', 4], ['512', '1K', 35], ['1K', '2K', 5], ['2K', '4K', 4], ['4K', '8K', 1]]]
76 usecs,1508,llvmpipe-1,[[['256', '512', 17], ['512', '1K', 34], ['1K', '2K', 5], ['2K', '4K', 2]]]
77 usecs,50,worker/u10:2,[[['256', '512', 16], ['512', '1K', 18], ['1K', '2K', 18], ['2K', '4K', 11], ['4K', '8K', 7]]]
78 usecs,1508,llvmpipe-0,[[['256', '512', 24], ['512', '1K', 41], ['1K', '2K', 5]]]
79 usecs,3347,glxtest,[[['256', '512', 32], ['512', '1K', 55], ['1K', '2K', 20], ['2K', '4K', 4], ['4K', '8K', 3]]]
80 usecs,3226,Classif-date #1,[[['256', '512', 2], ['512', '1K', 50], ['1K', '2K', 26], ['2K', '4K', 6], ['4K', '8K', 1]]]
81 usecs,1508,gnome-shell,[[['256', '512', 32], ['512', '1K', 31], ['1K', '2K', 17], ['2K', '4K', 8], ['4K', '8K', 11]]]
82 usecs,3226,firefox,[[['256', '512', 143], ['512', '1K', 160], ['1K', '2K', 50], ['2K', '4K', 21], ['4K', '8K', 1]]]
83 usecs,13,ksfortind/0,[[['0', '1'], ['1', '1'], ['0'], ['4', '8', 1], ['8', '16', 1], ['16', '32'], ['32', '64', 18], ['64', '128'], ['128', '256', 1]]]
84 usecs,23,ksfortind/1,[[['2', '4', 1], ['4', '8', 22], ['8', '16', 13], ['16', '32', 33], ['32', '64', 19], ['64', '128'], ['128', '256', 1]]]
85 usecs,29,ksfortind/2,[[['4', '8', 5], ['8', '16', 6], ['16', '32', 14], ['32', '64', 15], ['64', '128', 17], ['128', '256', 1]]]
86 usecs,39,ksfortind/3,[[['4', '8', 23], ['8', '16', 19], ['16', '32', 13], ['32', '64', 11], ['64', '128', 14], ['128', '256', 1]]]
87 usecs,41,ksfortind/4,[[['4', '8', 7], ['8', '16', 12], ['16', '32', 21], ['32', '64', 8], ['64', '128', 12], ['128', '256', 1]]]
88 usecs,8,swapper,4,[[['256', '512', 586], ['512', '1K', 985], ['1K', '2K', 261], ['2K', '4K', 171], ['4K', '8K', 64],
```

Silent Failure Identification

```
472 @usecs[41, ksoftirqd/4]:
473 [4, 8) 7 |@@
474 [8, 16) 12 |@@@
475 [16, 32) 21 |@@@@@@@
476 [32, 64) 8 |@@@
477 [64, 128) 128 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
478 [128, 256) 121 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
479 [256, 512) 77 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
480 [512, 1K) 134 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
481 [1K, 2K) 91 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
482 [2K, 4K) 83 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
483 [4K, 8K) 62 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
484 [8K, 16K) 20 |@@@@@@@
485 [16K, 32K) 0 |
486 [32K, 64K) 2 |
```

Nominal Data

Client Interest for Silent Failure?

Need more collections to understand

Data Upload to Database

Finished Sanitizing Files

Starting File Upload

Uploading: biopattern_run.csv to the DB

Data from CSV '../output/biopattern_run.csv' uploaded to table 'biopattern_run' successfully.

Uploading: block_rq_complete_run.csv to the DB

Data from CSV '../output/block_rq_complete_run.csv' uploaded to table 'block_rq_complete_run' successfully.

Uploading: biolateness_run.csv to the DB

Data from CSV '../output/biolateness_run.csv' uploaded to table 'biolateness_run' successfully.

Uploading: block_rq_error_run.csv to the DB

Data from CSV '../output/block_rq_error_run.csv' uploaded to table 'block_rq_error_run' successfully.

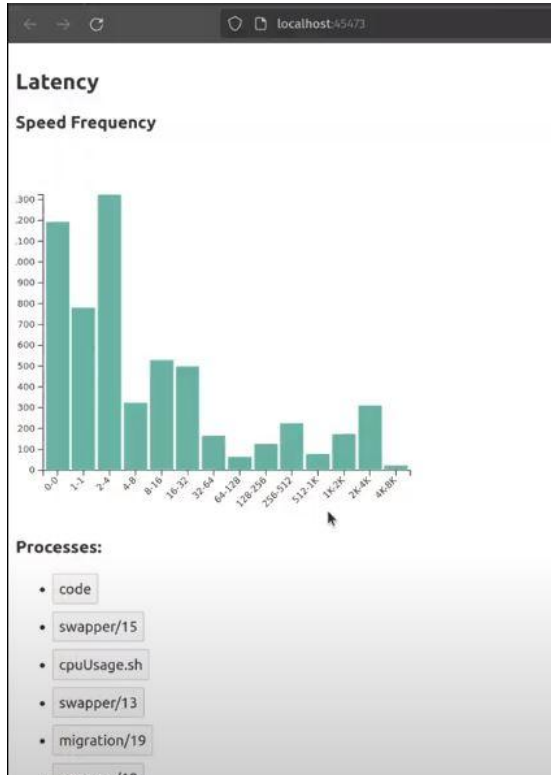
Finished File Upload

DATA COLLECTION TO DATA UPLOAD SUCCESSFULLY RAN



identifier	action_id	process	process_speeds	serial_number
1	usecs	swapper/9	[["128","256",3],["256","512",2]]	231708805892
2	usecs	firefox-bin	[["64","128",3],["128","256",2],["256","512",0],["512","1K",37],["1K","2K",10]]	231708805892
3	usecs	InputThread	[["512","1K",5]]	231708805892
4	usecs	swapper/1	[["512","1K",3]]	231708805892
5	usecs	swapper/10	[["4","8",1],["8","16",2],["16","32",4],["32","64",11],["64","128",19],["128","256",44],["256","512",2],["512","1K",26],["1K","2K",94],["2K",4...	231708805892
6	usecs	swapper/1	[["128","256",42],["256","512",18],["512","1K",48],["1K","2K",74],["2K","4K",11]]	231708805892
7	usecs	swapper/2	[["128","256",1],["256","512",42],["512","1K",1]]	231708805892
8	usecs	kworker/u40:1	[["512","1K",33]]	231708805892

Dashboard/Visualization



When selecting a certain process

- swapper/18
- Timer
- kworker/u40:8
- swapper/1
- swapper/3
- swapper/19
- swapper/2
- code
- swapper/4
- swapper/3
- gdbus
- swapper/17
- kworker/u40:30
- swapper/18
- code
- swapper/15
- tokio-runtime-w
- swapper/19
- swapper/3
- code
- swapper/10



Challenges/Resolution

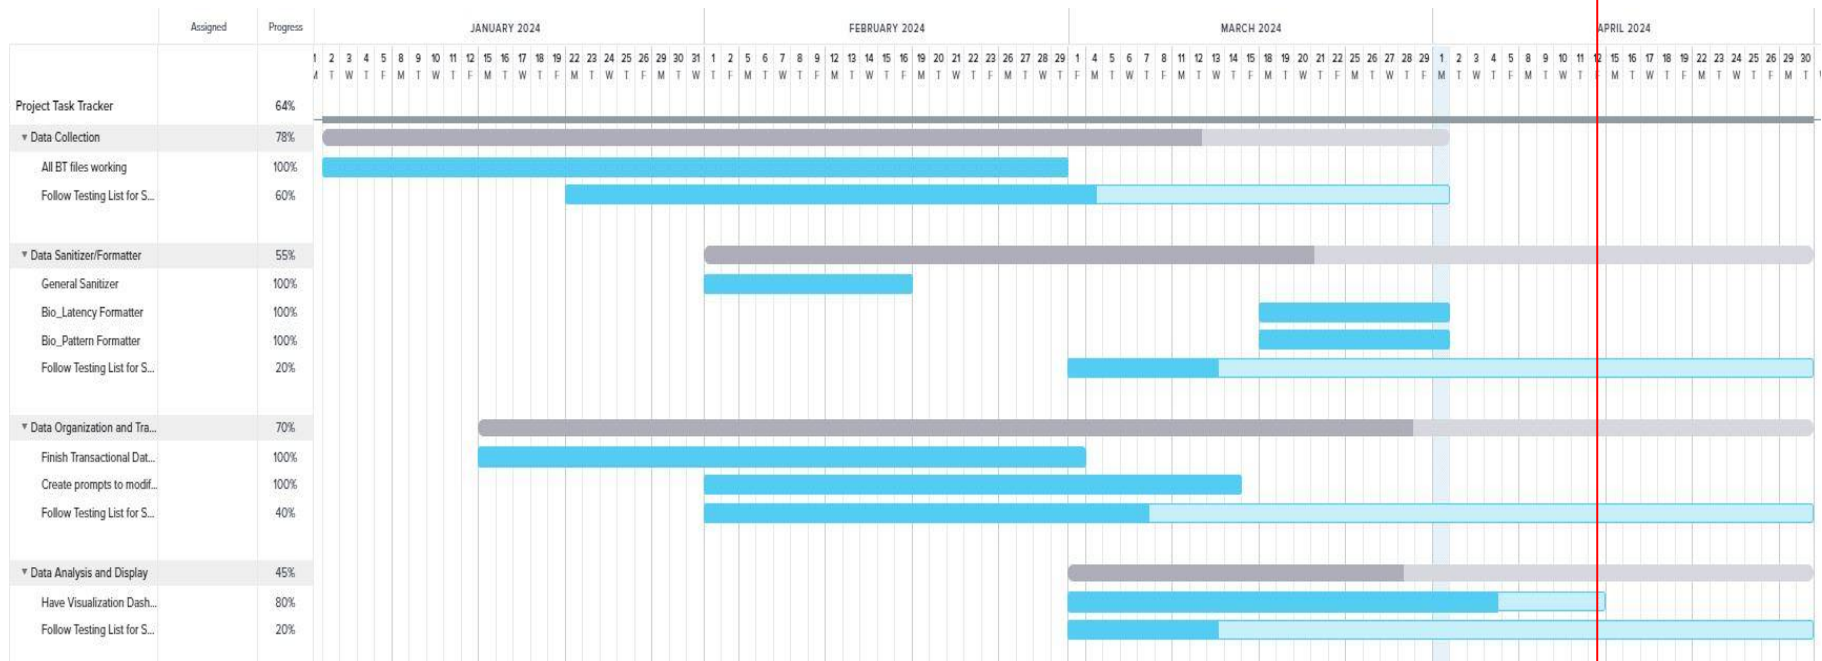
Challenges

- NVME Latency - Unable to work on system
- Visualization - Low documentation and compatibility for previous dashboard (Open search and Prometheus)

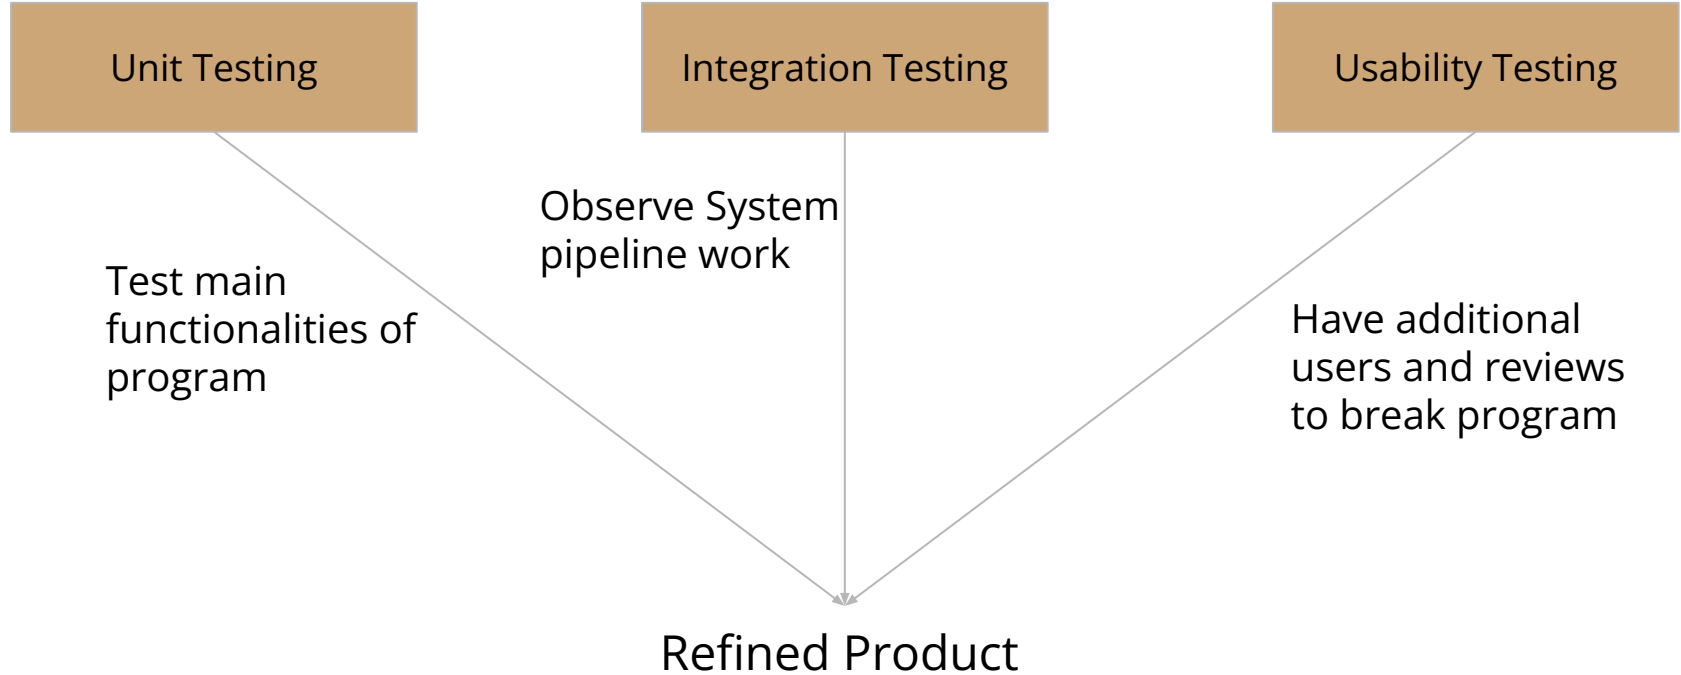
Resolution

- NVME Latency - Most recent Ubuntu version removed tracepoints entirely. Use current NVME Latency file on deprecated Version of Ubuntu with tracepoints existing
- Visualization - Research completed and will remain on flask, svelte and D3

Schedule



Testing Plan



Unit Test Example

```
def test_regex_extractor(self):
    # Mocking the behavior of pd.read_csv to read the dummy CSV file
    with patch('Sanitizer.pd.read_csv') as mock_read_csv:
        # Load dummy CSV data
        dummy_csv_data = pd.read_csv('block_rq_complete_run.csv')
        mock_read_csv.return_value = dummy_csv_data
        # Assuming 'process' column contains data like "b'<process_name>'"
        data_list = dummy_csv_data['process'].tolist()
        result = self.sanitizer.regex_extractor(data_list)
        # Assertions
        self.assertNotEqual(result, ["b'kworker/3:1'", "b'<idle>'"])
```


Conclusion

Problem:

- Silent Failures, bad for companies who rely on SSD's in their infrastructure.
- For SSD Manufacturers:
 - Long data collection process.
 - Money Loss due to time usage and manual actions.

Solution:

- Make R&D Process more efficient.
- System observability, insights, and analytics platform.



QUESTIONS?

