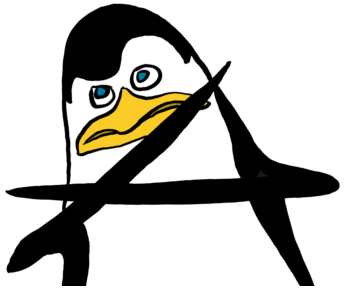

System Observability, Analytics & Insights Platform

Team Kowalski

Members

Jake Borneman
> Team Leader
> Testing
> Sanitization

Erick Salazar
> Data Storage
> Data Pipeline

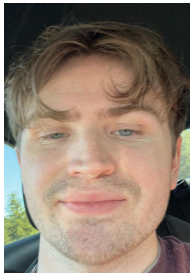

Bailey McCauslin
> Data Collection
> Testing

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

Introduction to Sector:

- Various companies: Adata, Western Digital, Samsung, Crucial, ...
- Market value in the tens of billions annually, with continuous growth.
- Diverse user base from individuals to large corporations and governments.
- Handles exabytes to zettabytes of data, constantly expanding with digital content growth.

Industry Products:

- Hardware: SSDs, HDDs, etc...
- Software: Cloud Storage Services
- Different solutions for different needs

SAMSUNG

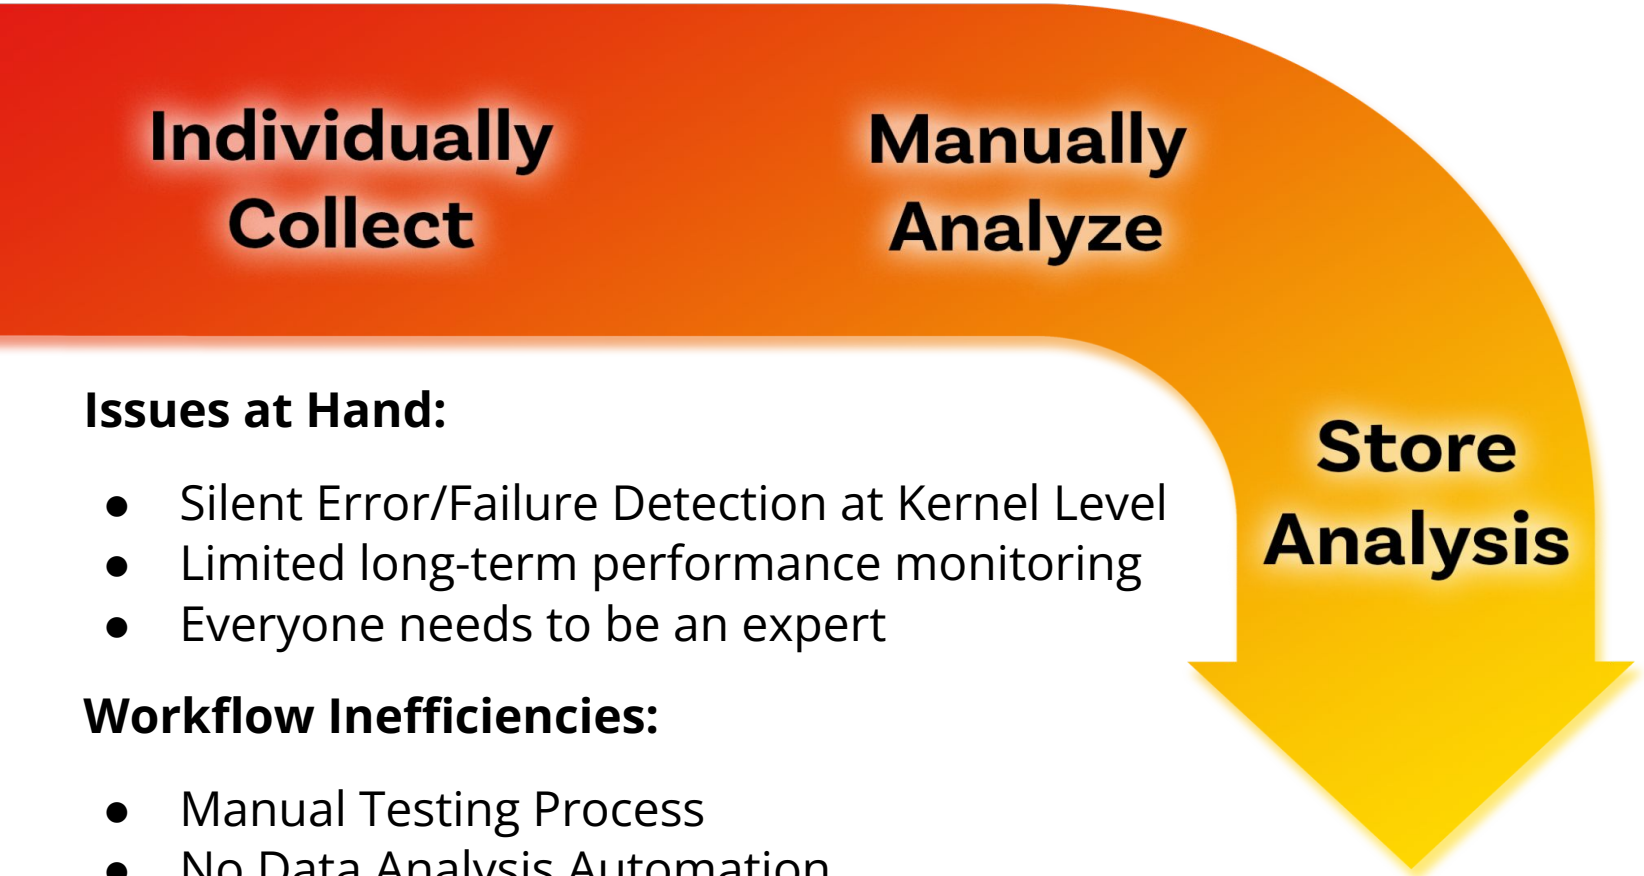

Issues at Hand:

- 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

Silent Errors

Performance Metrics:

- Biolatency
- Biopattern
- Block RQ Complete
- Bioerror
- Block RQ Error

The Black Box

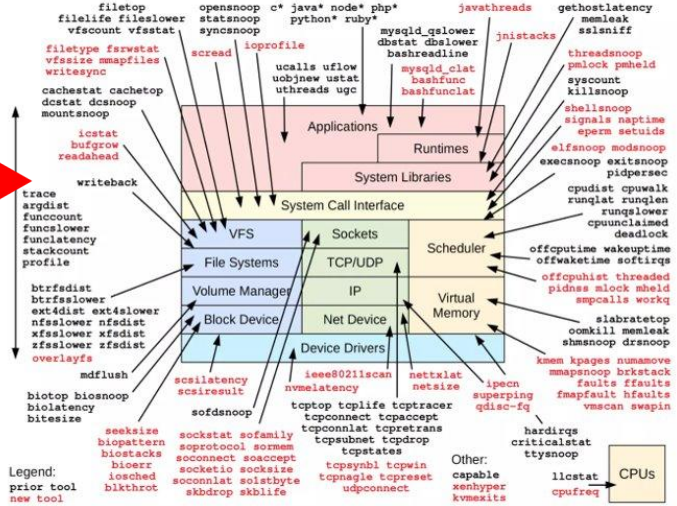

Opening the Black Box

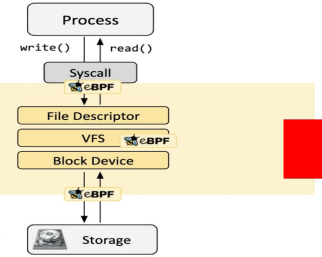

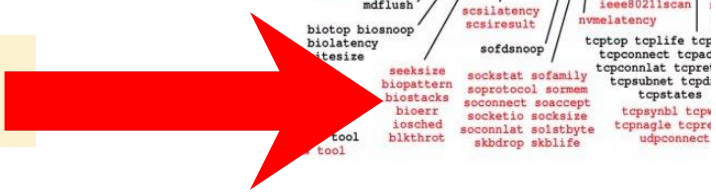

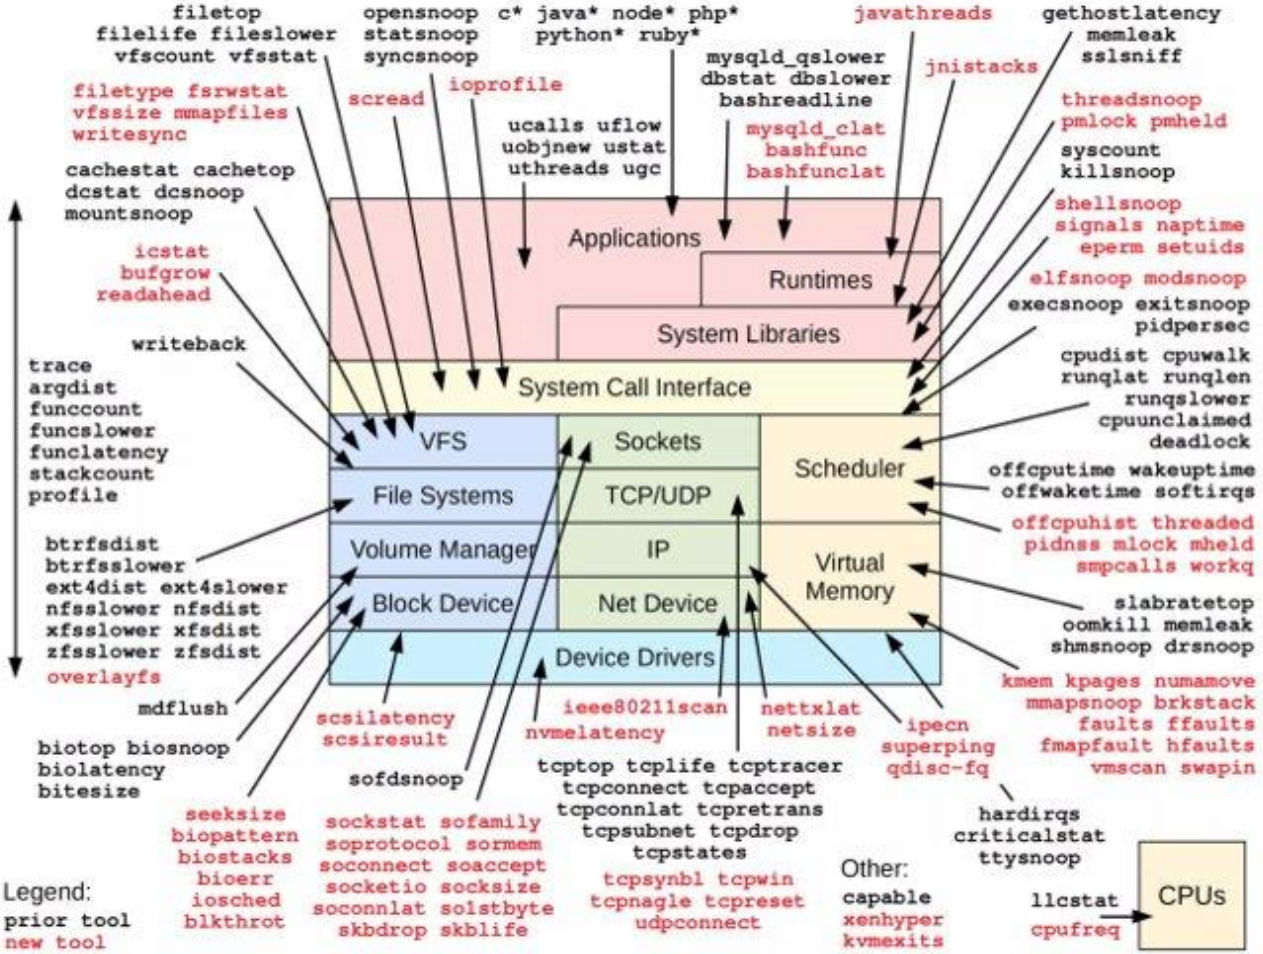
filetop
filelife fileslower
vfscount vfsstat
opensnoop
statsnoop
syncsnoop
c* java* node* php*
python* ruby*
javathreads
gethostlatency
memleak
sslsniff
filetype fsrwstat
vfssize mmapfiles
writesync
scread
ioprofile
mysqld_qslower
dbstat dbslower
bashreadline
jnistacks
threadsnoop
pmlock pmheld
ucalls uflow
uobjnew ustat
uthreads ugc
mysqld_clat
bashfunc
bashfunclat
syscount
killsnoop
cachestat cachetop
dcstat dcsnoop
mountsnoop
shellsnoop
signals naptime
eperm setuids
Applications
Runtimes
System Libraries
System Call Interface
icstat
bufgrow
readahead
elfsnoop modsnoop
execsnoop exitsnoop
pidpersec
writeback
trace
argdist
funccount
funcslower
funclatency
stackcount
profile
VFS
Sockets
Scheduler
File Systems
TCP/UDP
Volume Manager
IP
Virtual Memory
Block Device
Net Device
Device Drivers
cpudist cpuwalk
runqlat runqlen
runqslower
cpuunclaimed
deadlock
offcputime wakeuptime
offwaketime softirqs
offcpuhist threaded
pidnss mlock mheld
smpcalls workq
btrfsdist
btrfsslower
ext4dist ext4slower
nfsslower nfsdist
xfsslower xfsdist
zfsslower zfsdist
overlayfs
slabratetop
oomkill memleak
shmsnoop drsnoop
kmem kpages numamove
mmapsnoop brkstack
faults ffaults
fmapfault hfaults
vmscan swapin
mdflush
scsilatency
scsiresult
ieee80211scan
nvmelatency
nettxlat
netsize
ipecn
superping
qdisc-fq
biotop biosnoop
biolatency
bitesize
sofdsnoop
tcptop tcplife tcptracer
tcpconnect tcpaccept
tcpconnlat tcpretrans
tcpsubnet tcpdrop
tcpstates
hardirqs
criticalstat
ttysnoop
seeksize
biopattern
biostacks
bioerr
iosched
blkthrot
sockstat sofamily
soprotocol sormem
soconnect soaccept
socketio socksize
soconnlat solstbyte
skbdrop skblife
tcpsynbl tcpwin
tcpnagle tcpreset
udpconnect
Other:
capable
xenhyper
kvmexits
llcstat
cpufreq
CPUs
Legend:
prior tool
new tool

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

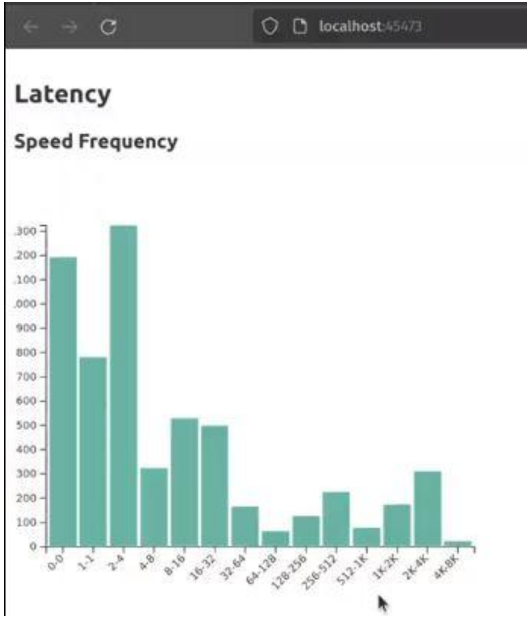

Cleaned and Uploaded Data

Raw Data

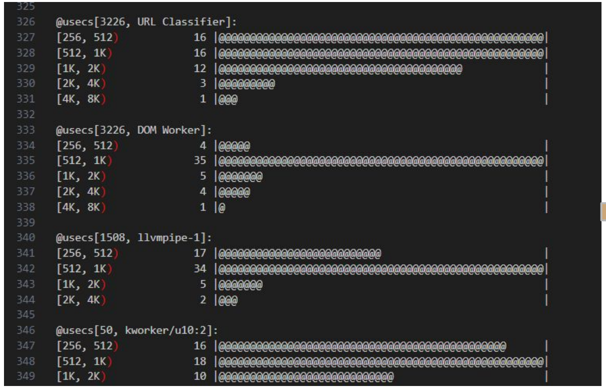

Requirements Review

- 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.

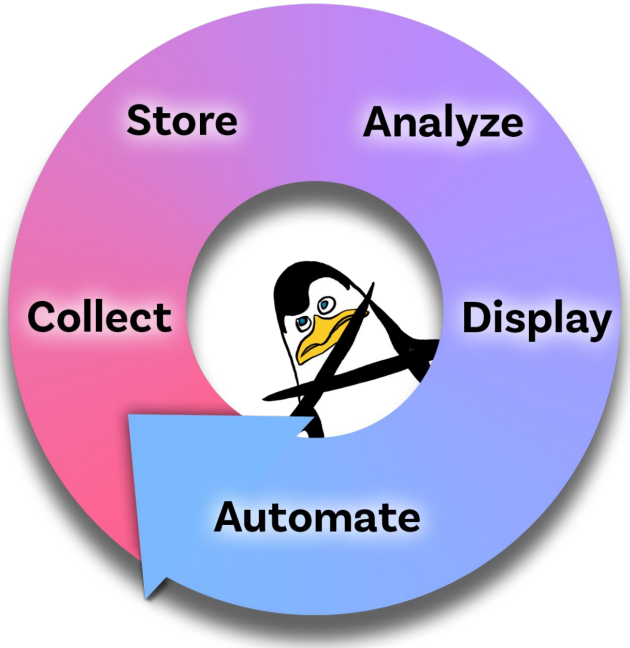

Architecture Review

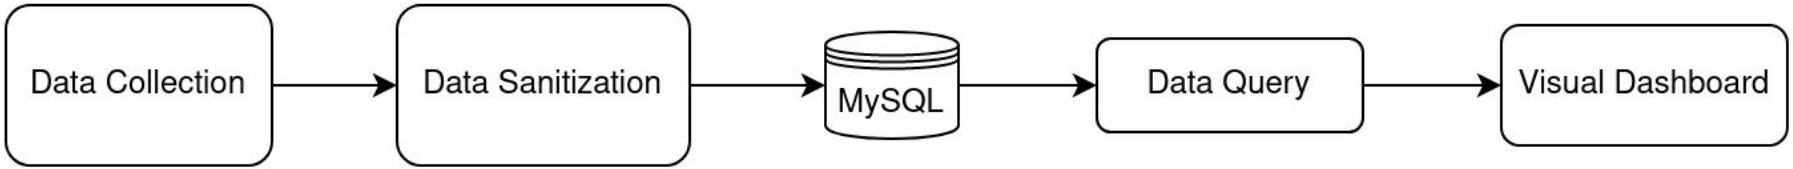

Implementation Review

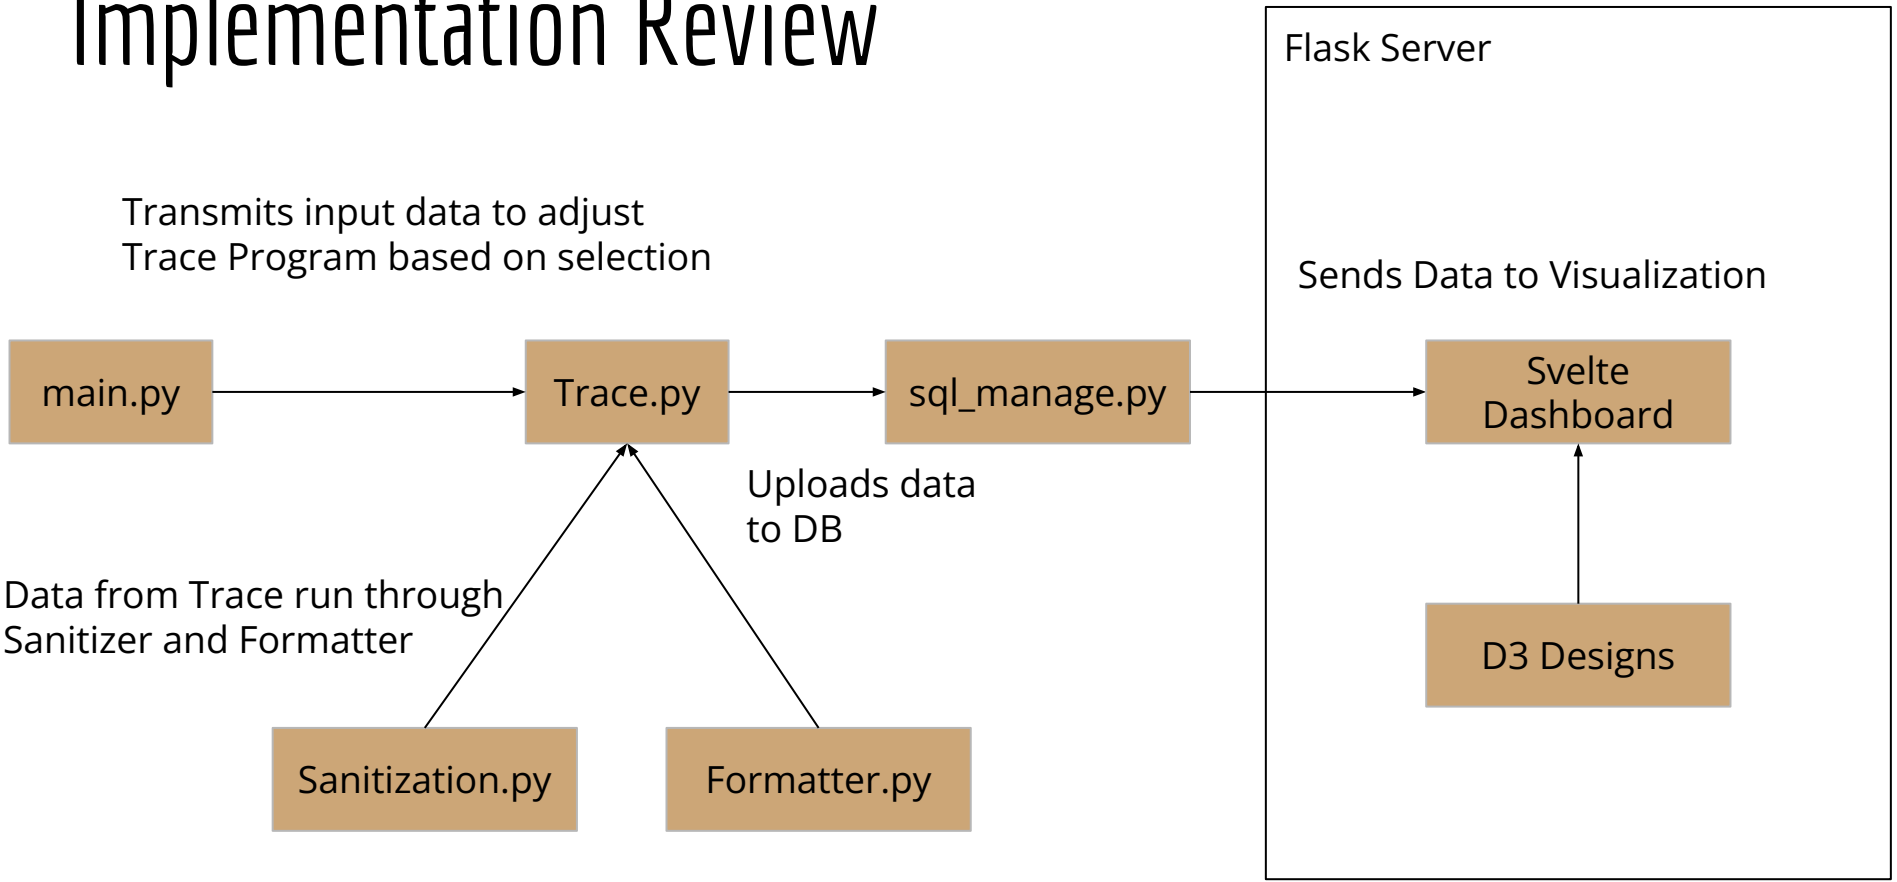

Prototype Review

Initialization of Services

```
Starting Flask Server and Svelte App

> svelte-app@1.0.0 dev
> rollup -c -w

LiveReload enabled

> svelte-app@1.0.0 start
> sirv public --no-clear --dev


  Your application is ready~! 🚀

  - Local:      http://localhost:8080

Svelte app running at: http://localhost:8080

Flask server ready at  http://127.0.0.1:5000/

Flask Server and Svelte App successfully started!
```

User Input Collection

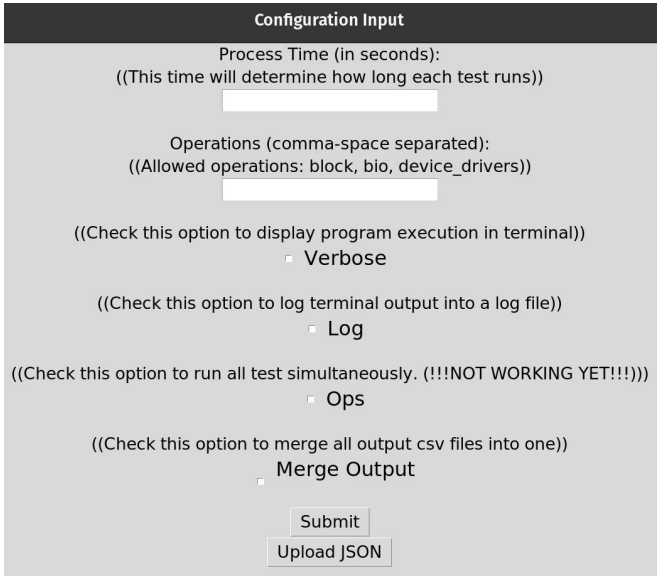

Given user input displayed in terminal

```
Now collecting user inputs

----------------------USER INPUTS----------------------

Process Time: 15, Operations: ['block', 'bio'], Verbose: True, Log: False, Ops: False, Merge: False

-------------------------------------------------------

User Inputs Collected
```

Master Data Collection

```
Beginning master data collection

NAME              SIZE MODEL
zram0            15.3G
nvme0n1         953.9G PC SN740 NVMe WD 1TB
├─nvme0n1p1      1022M
├─nvme0n1p2         4G
├─nvme0n1p3     944.9G
│ └─cryptdata   944.9G
│   └─data-root 944.8G
└─nvme0n1p4         4G
  └─cryptswap       4G

Select block device to fetch data of: nvme0n1

Block Device Data Found:
Model Number: PC SN740 NVMe WD 1TB
Serial Number: 231708805892

Data inserted successfully!
Master data collection complete
```

Block Devices connected to PC

Transactional Data Collection

```
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

```
Begun Sanitizing Files

Sanatizing File: biopattern_run.txt
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
```

```
@usecs[29, ksoftirqd/2]:
[4, 8)                 5 |@                                                   |
[8, 16)                6 |@                                                   |
[16, 32)              14 |@@@@                                                |
[32, 64)              15 |@@@@                                                |
[64, 128)            177 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@|
[128, 256)           130 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@              |
[256, 512)            41 |@@@@@@@@@@@@                                        |
[512, 1K)             53 |@@@@@@@@@@@@@@@                                     |
[1K, 2K)              28 |@@@@@@@@                                            |
[2K, 4K)              52 |@@@@@@@@@@@@@@@                                     |
[4K, 8K)              20 |@@@@@                                               |
[8K, 16K)              6 |@                                                   |
[16K, 32K)             3 |                                                    |

@usecs[35, ksoftirqd/3]:
[4, 8)                23 |@@@@@@                                              |
[8, 16)               19 |@@@@@                                               |
[16, 32)              13 |@@@                                                 |
[32, 64)              11 |@@@                                                 |
[64, 128)            144 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@          |
[128, 256)           178 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@|
[256, 512)            57 |@@@@@@@@@@@@@@@@                                    |
[512, 1K)             74 |@@@@@@@@@@@@@@@@@@@@@                               |
[1K, 2K)              40 |@@@@@@@@@@@                                         |
[2K, 4K)              56 |@@@@@@@@@@@@@@@@                                    |
[4K, 8K)              12 |@@@                                                 |
[8K, 16K)              7 |@@                                                  |
[16K, 32K)             1 |                                                    |

@usecs[41, ksoftirqd/4]:
[4, 8)                 7 |@@                                                  |
[8, 16)               12 |@@@@                                                |
[16, 32)              21 |@@@@@@@@                                            |
[32, 64)               8 |@@@                                                 |
[64, 128)            128 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@   |
[128, 256)           121 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@      |
[256, 512)            77 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@                       |
[512, 1K)            134 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@|
[1K, 2K)              91 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@                 |
[2K, 4K)              83 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@                    |
[4K, 8K)              62 |@@@@@@@@@@@@@@@@@@@@@@@@                            |
[8K, 16K)             20 |@@@@@@@                                             |
[16K, 32K)             0 |                                                    |
[32K, 64K)             2 |                                                    |

@usecs[0, swapper/4]:
[256, 512)           586 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@                   |
[512, 1K)            905 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@|
[1K, 2K)             261 |@@@@@@@@@@@@@@                                      |
[2K, 4K)             171 |@@@@@@@@@                                           |
[4K, 8K)              64 |@@@                                                 |
[8K, 16K)             32 |@                                                   |
[16K, 32K)             3 |                                                    |
[32K, 64K)             2 |                                                    |
```


```
usecs,3226,BgIOThr~Pool #1,"[['256', '512', 1], ['512', '1K', 1]]"
usecs,3226,Backgro~Pool #2,"[['8K', '16K', 2]]"
usecs,3224,bpftrace,"[['4K', '8K', 2]]"
usecs,3564,StreamTrans #13,"[['2K', '4K', 2]]"
usecs,3226,DNS Resolver #1,"[['256', '512', 1], ['512', '1K', 1]]"
usecs,3183,Xwayland,"[['4K', '8K', 1], ['8K', '16K', 1]]"
usecs,3226,BgIOThr~Pool #2,"[['256', '512', 2], ['512', '1K', 0], ['1K', '2K', 1]]"
usecs,2024,gjs,"[['512', '1K', 2], ['1K', '2K', 0], ['2K', '4K', 0], ['4K', '8K', 1]]"
usecs,3226,StreamTrans #15,"[['256', '512', 1], ['512', '1K', 1], ['1K', '2K', 1]]"
usecs,3226,QuotaManager IO,"[['256', '512', 1], ['512', '1K', 0], ['1K', '2K', 2]]"
usecs,3226,sqldb:p~lite #1,"[['2K', '4K', 1], ['4K', '8K', 1], ['8K', '16K', 1]]"
usecs,3410,Socket Process,"[['512', '1K', 2], ['1K', '2K', 1]]"
usecs,3226,RemoteLzyStream,"[['256', '512', 1], ['512', '1K', 3]]"
usecs,3226,glean.init,"[['256', '512', 3], ['512', '1K', 1]]"
usecs,3226,TaskCon~ller #2,"[['512', '1K', 1], ['1K', '2K', 3]]"
usecs,3226,IPC Launch,"[['256', '512', 1], ['512', '1K', 2], ['1K', '2K', 1]]"
usecs,3226,TaskCon~ller #4,"[['512', '1K', 4]]"
usecs,3564,WebExtensions,"[['512', '1K', 3], ['1K', '2K', 2]]"
usecs,3226,SSL Cert #2,"[['256', '512', 1], ['512', '1K', 1], ['1K', '2K', 3]]"
usecs,3226,StreamTrans #3,"[['256', '512', 2], ['512', '1K', 3]]"
usecs,3226,StreamTrans #11,"[['512', '1K', 2], ['1K', '2K', 3]]"
usecs,1508,pool-gnome-shel,"[['512', '1K', 4], ['1K', '2K', 2]]"
usecs,3226,IPDL Background,"[['256', '512', 1], ['512', '1K', 1], ['1K', '2K', 1], ['2K', '4K', 0], ['4K', '8K', 3]]
usecs,182,jbd2/sda3-8,"[['128', '256', 1], ['256', '512', 2], ['512', '1K', 1], ['1K', '2K', 0], ['2K', '4K', 2]]"
usecs,3226,SSL Cert #1,"[['256', '512', 1], ['512', '1K', 7]]"
usecs,3226,Backgro~Pool #1,"[['256', '512', 3], ['512', '1K', 6], ['1K', '2K', 0], ['2K', '4K', 1]]"
usecs,3226,[pango] FcInit,"[['512', '1K', 1], ['1K', '2K', 2], ['2K', '4K', 6], ['4K', '8K', 2]]"
usecs,778,kworker/4:2H,"[['512', '1K', 8], ['1K', '2K', 1], ['2K', '4K', 1], ['4K', '8K', 1]]"
usecs,3434,Privileged Cont,"[['256', '512', 1], ['512', '1K', 3], ['1K', '2K', 4], ['2K', '4K', 2], ['4K', '8K', 1]]
usecs,348,kworker/u10:8,"[['256', '512', 2], ['512', '1K', 7], ['1K', '2K', 3], ['2K', '4K', 1]]"
usecs,3226,Renderer,"[['256', '512', 6], ['512', '1K', 8], ['1K', '2K', 1], ['2K', '4K', 1]]"
usecs,3226,IndexedDB #1,"[['256', '512', 1], ['512', '1K', 8], ['1K', '2K', 3], ['2K', '4K', 6], ['4K', '8K', 0], ['
usecs,79,kworker/u10:4,"[['128', '256', 1], ['256', '512', 8], ['512', '1K', 12], ['1K', '2K', 3], ['2K', '4K', 1]]"
usecs,3226,Cache2 I/O,"[['256', '512', 4], ['512', '1K', 13], ['1K', '2K', 6], ['2K', '4K', 0], ['4K', '8K', 1], ['8
usecs,10,kworker/u10:0,"[['256', '512', 9], ['512', '1K', 11], ['1K', '2K', 4], ['2K', '4K', 4], ['4K', '8K', 1]]"
usecs,3226,Socket Thread,"[['256', '512', 10], ['512', '1K', 16], ['1K', '2K', 5], ['2K', '4K', 1], ['4K', '8K', 3]]
usecs,1508,llvmpipe-2,"[['256', '512', 4], ['512', '1K', 24], ['1K', '2K', 3], ['2K', '4K', 3], ['4K', '8K', 1]]"
usecs,1508,llvmpipe-3,"[['128', '256', 1], ['256', '512', 16], ['512', '1K', 22], ['1K', '2K', 4]]"
usecs,1508,llvmpipe-4,"[['128', '256', 4], ['256', '512', 19], ['512', '1K', 14], ['1K', '2K', 4], ['2K', '4K', 1],
usecs,3226,URL Classifier,"[['256', '512', 16], ['512', '1K', 16], ['1K', '2K', 12], ['2K', '4K', 3], ['4K', '8K', 1
usecs,3226,DOM Worker,"[['256', '512', 4], ['512', '1K', 35], ['1K', '2K', 5], ['2K', '4K', 4], ['4K', '8K', 1]]"
usecs,1508,llvmpipe-1,"[['256', '512', 17], ['512', '1K', 34], ['1K', '2K', 5], ['2K', '4K', 2]]"
usecs,50,kworker/u10:2,"[['256', '512', 16], ['512', '1K', 18], ['1K', '2K', 10], ['2K', '4K', 11], ['4K', '8K', 7],
usecs,1508,llvmpipe-0,"[['256', '512', 24], ['512', '1K', 41], ['1K', '2K', 5]]"
usecs,3347,glxtest,"[['256', '512', 32], ['512', '1K', 55], ['1K', '2K', 20], ['2K', '4K', 4], ['4K', '8K', 3]]"
usecs,3226,Classif~date #1,"[['256', '512', 22], ['512', '1K', 59], ['1K', '2K', 26], ['2K', '4K', 6], ['4K', '8K',
usecs,1508,gnome-shell,"[['256', '512', 32], ['512', '1K', 111], ['1K', '2K', 37], ['2K', '4K', 8], ['4K', '8K', 11]
usecs,3226,firefox,"[['256', '512', 141], ['512', '1K', 160], ['1K', '2K', 50], ['2K', '4K', 21], ['4K', '8K', 8]]"
usecs,15,ksoftirqd/0,"[['0', '0', '1'], ['1', '1', '0'], ['2', '4', 1], ['4', '8', 11], ['8', '16', 18], ['16', '32'
usecs,23,ksoftirqd/1,"[['2', '4', 1], ['4', '8', 22], ['8', '16', 13], ['16', '32', 33], ['32', '64', 19], ['64', '1
usecs,29,ksoftirqd/2,"[['4', '8', 5], ['8', '16', 6], ['16', '32', 14], ['32', '64', 15], ['64', '128', 177], ['128'
usecs,35,ksoftirqd/3,"[['4', '8', 23], ['8', '16', 19], ['16', '32', 13], ['32', '64', 11], ['64', '128', 144], ['12
usecs,41,ksoftirqd/4,"[['4', '8', 7], ['8', '16', 12], ['16', '32', 21], ['32', '64', 8], ['64', '128', 128], ['128'
usecs,0,swapper/4,"[['256', '512', 586], ['512', '1K', 905], ['1K', '2K', 261], ['2K', '4K', 171], ['4K', '8K', 64],
```

Data Upload to Database

```
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: biolatency_run.csv to the DB
Data from CSV '../output/biolatency_run.csv' uploaded to table 'biolatency_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
```

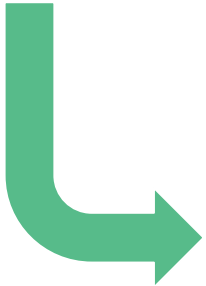

biolatency_run *[all]*

identifier | equals | Enter Value

	identifier varchar(255)	action_id int	process varchar(255)	process_speeds json	serial_number
1	usecs	0	swapper/9	[["128","256",3],["256","512",2]]	231708805892
2	usecs	11002	firefox-bin	[["64","128",3],["128","256",2],["256","512",0],["512","1K",37],["1K","2K",10]]	231708805892
3	usecs	2250	InputThread	[["512","1K",5]]	231708805892
4	usecs	0	swapper/1	[["512","1K",3]]	231708805892
5	usecs	0	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	0	swapper/1	[["128","256",42],["256","512",18],["512","1K",48],["1K","2K",74],["2K","4K",11]]	231708805892
7	usecs	0	swapper/2	[["128","256",1],["256","512",42],["512","1K",1]]	231708805892
8	usecs	22488	kworker/u40:1	[["512","1K",33]]	231708805892

Dashboard/Visualization

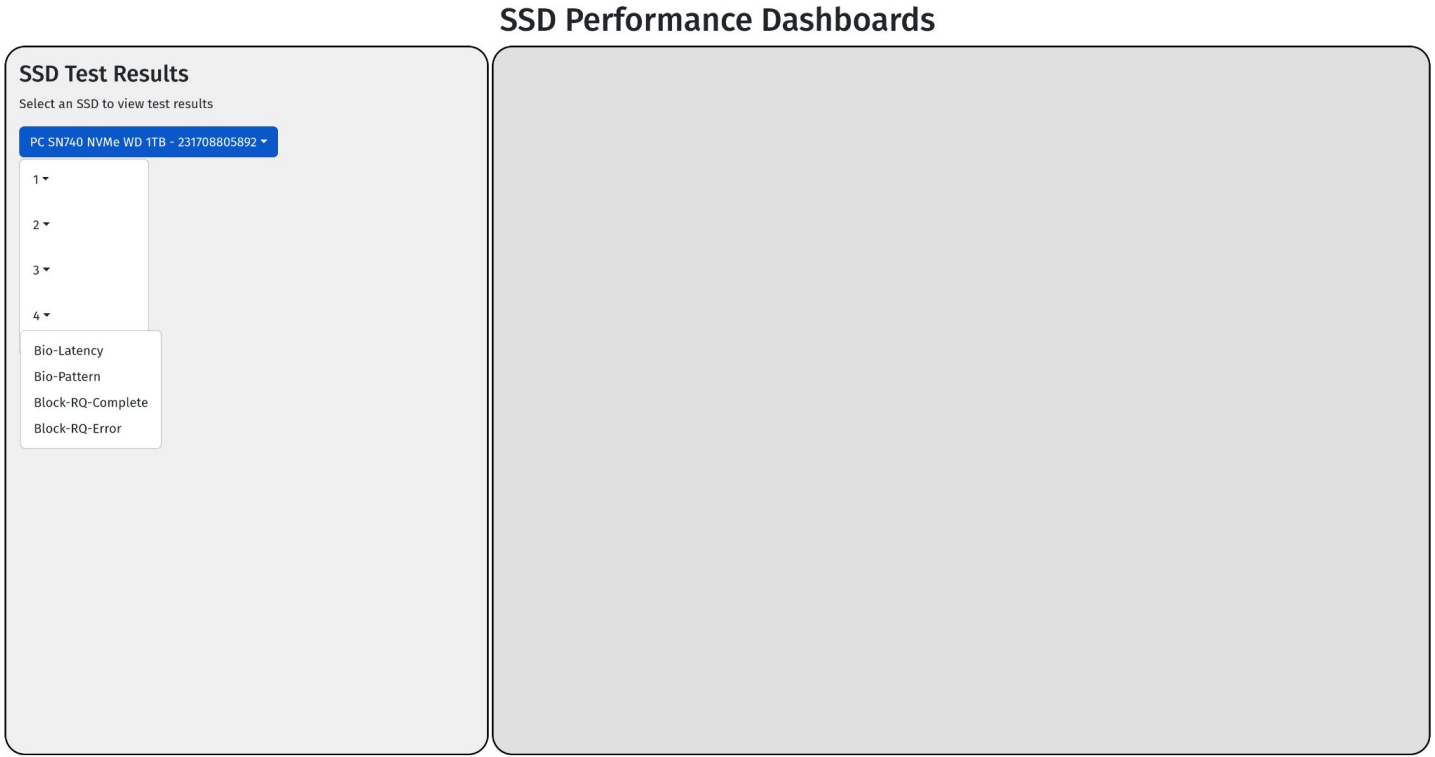

Biopattern

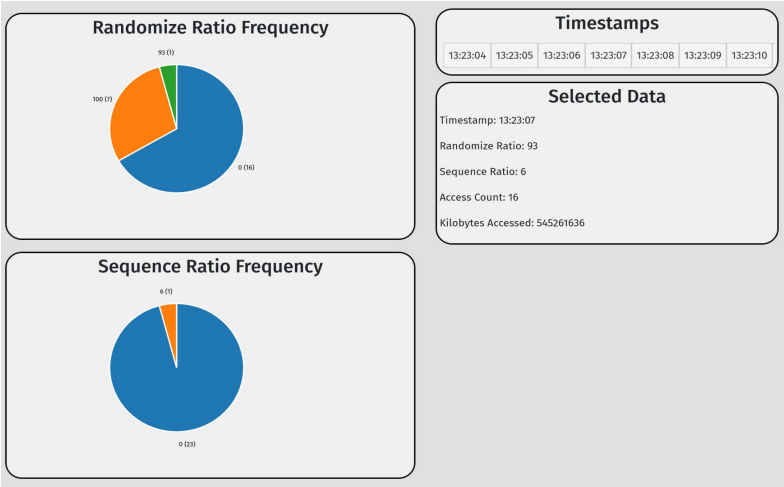

Biolatency

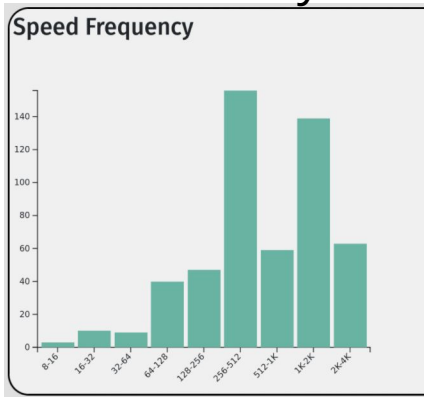

Select individual process

Block RQ Complete & Error

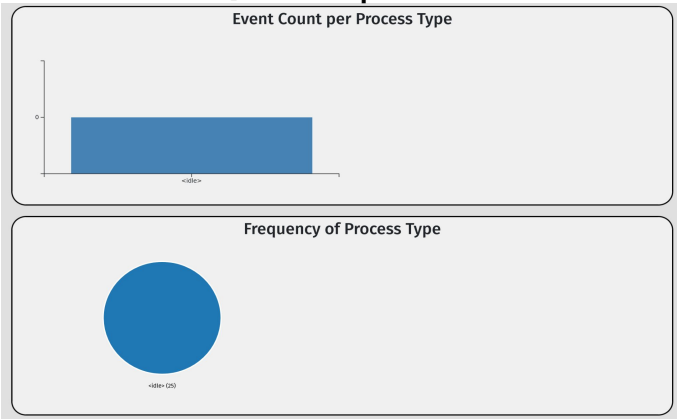

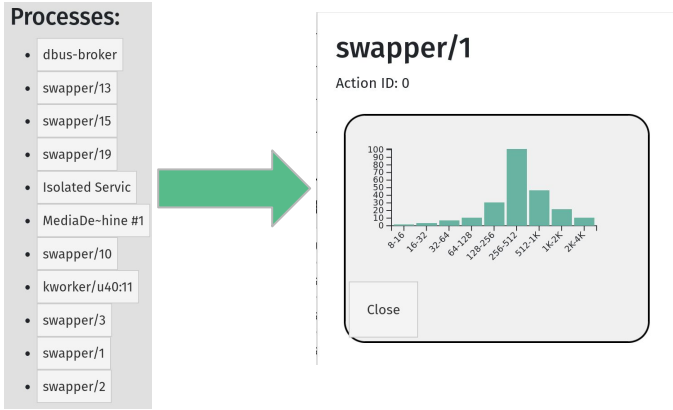

Silent Failure Identification

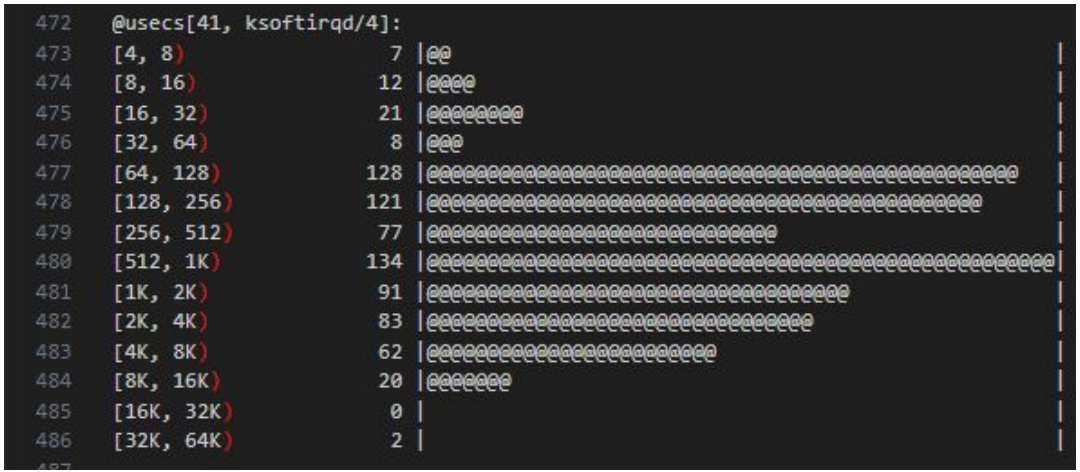

Nominal Data

Client Interest for Silent Failure?

Need more collections to understand

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

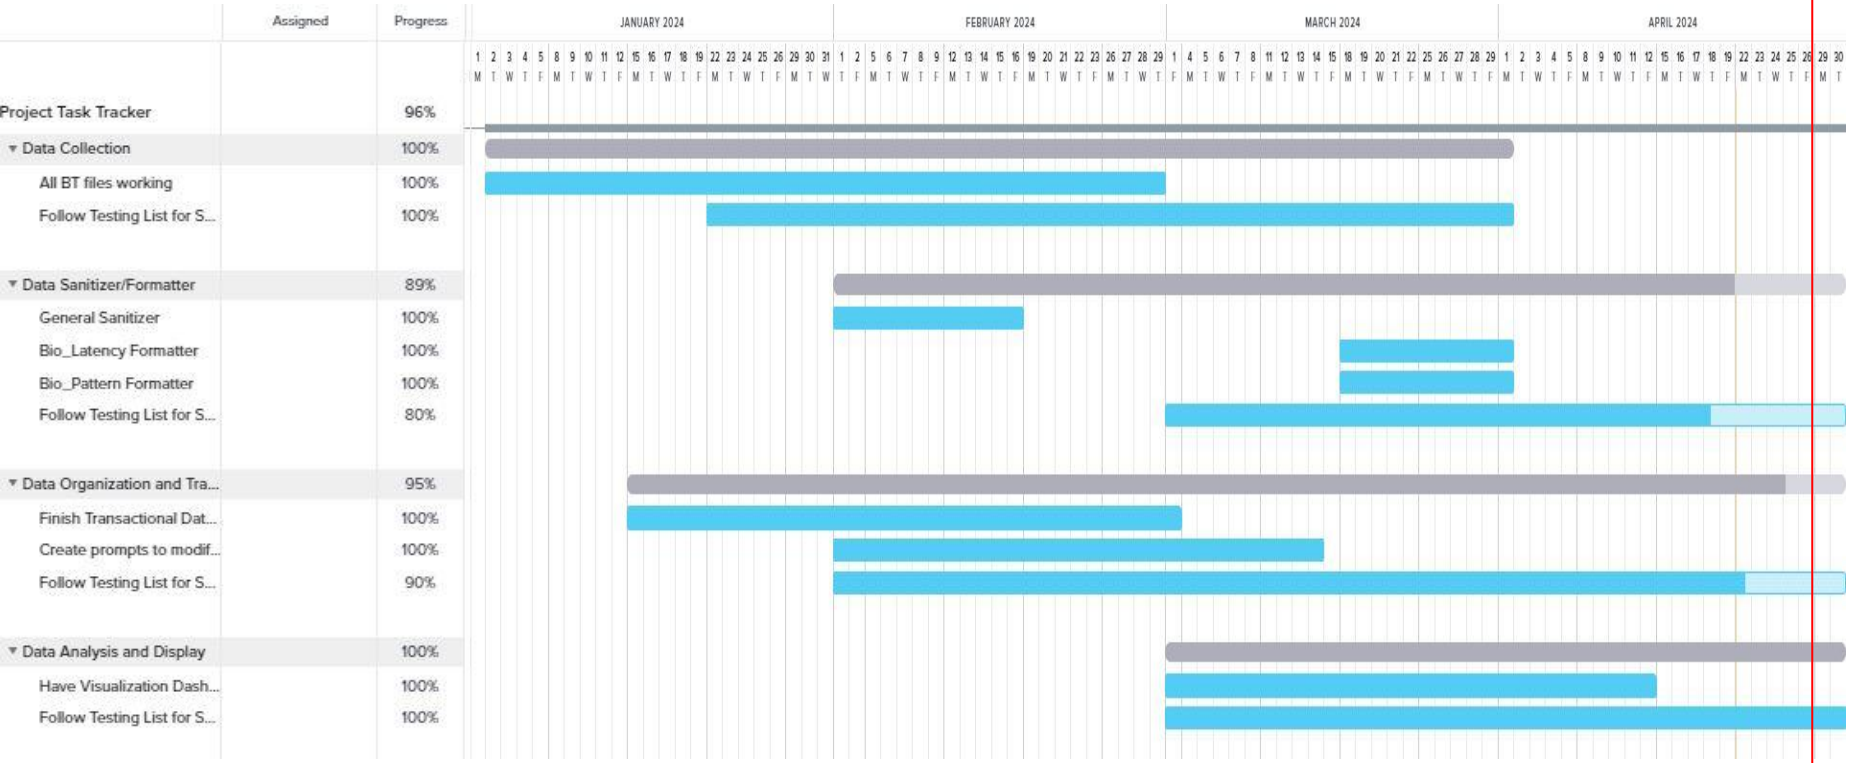
Assigned
Progress
JANUARY 2024
FEBRUARY 2024
MARCH 2024
APRIL 2024
Project Task Tracker 96%
Data Collection 100%
All BT files working 100%
Follow Testing List for S... 100%
Data Sanitizer/Formatter 89%
General Sanitizer 100%
Bio_Latency Formatter 100%
Bio_Pattern Formatter 100%
Follow Testing List for S... 80%
Data Organization and Tra... 95%
Finish Transactional Dat... 100%
Create prompts to modif... 100%
Follow Testing List for S... 90%
Data Analysis and Display 100%
Have Visualization Dash... 100%
Follow Testing List for S... 100%

Testing Plan

Unit Testing
(Uses PyTest,
Unittest library)

Current:
Trace.py
Sanitizer.py
Formatter.py

Integration Testing

Current:
System Pipeline
Repeated Collections

Usability Testing

Plan:
Survey
Ticket submission

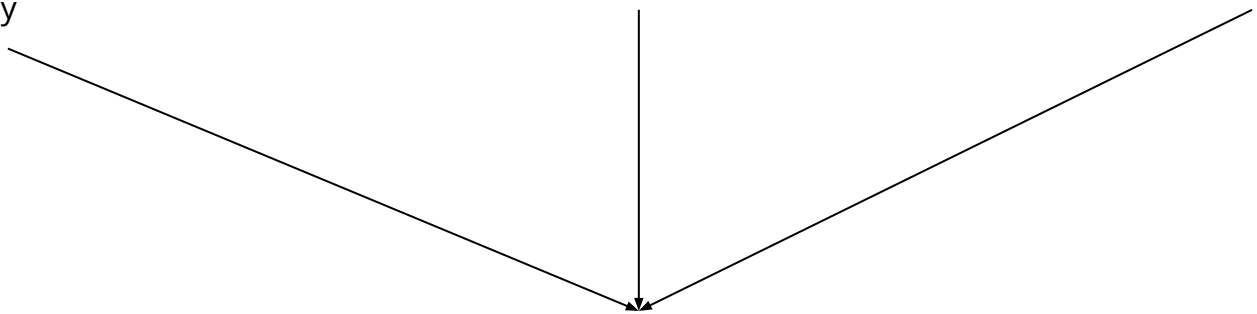

Refined Product

Future Work

- Dockerize the program
- Implement cloud storage to gather information from multiple sources
- Create a prediction program to predict silent failures ahead of time

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.

Value: Reduced expenses in man hours and avoiding recalls/refunds from broken devices

Outcomes: Old and new client were happy with the product

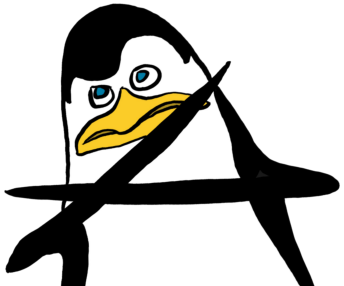

THANK YOU

Old Client

Rajpal Singh

New Client

Igor Steinmacher

Mentor

Saisri Muttineni

QUESTIONS?

QR Codes

Our Website

Walkthrough