

# ResilientDB Visualization Framework

Yu-Hsuan Lin, Yi-Hsiang Chiu, Tzu-Hsin Yang, Pin-Ting Liu, Yu-Cheng Hwang

Department of Computer Science

University of California, Davis

{zyhlin,yihchiu,zixyang,ptliu,timhwang}@ucdavis.edu

December 11, 2021

**Abstract**—Nowadays, Cryptocurrency, Bitcoins, Ethereum, Dogecoin, and so on, become more and more popular as investments and transactions. The blockchain technology behind these cryptocurrencies has also entered people’s field of sight. Nonetheless, the mechanisms of blockchain are complex and confusing because it is difficult to look into the details of the mechanisms. Therefore, visualizing the view of blockchain is critical work. It can assist people to grab more ideas behind the blockchain. Also, blockchain researchers can make use of visualization systems to monitor the status of blockchain. As we know, blockchain is a distributed system, and each block works on a different machine independently. In this project, our goal is to visualize the status of those individual machines. We will use ResilientDB as the example of blockchain, and monitor the situation of each machine (replica), and then build a web user interface (UI) to show the information such as CPU utilization, throughput, and so on during a transaction.

## I. INTRODUCTION

In order to adopt blockchain technology in a wider range of fields, we need to explore and analyze transaction data to better understand user behaviors and mechanisms in the blockchain system. Therefore, visualization tools can support humans in analyzing transactions and deriving hypotheses and models used in the blockchain.

There are many previous work developing visualization systems. [1]–[6]. However, most of them only provide descriptive statistics or transaction details. In order to help users and new bitcoin learners can understand clearly about the mechanism of bitcoin. We would like to visualize the communication process between replicas.

There were many types of visualization tools proposed for above purposes. Among them, time series data were the main visualization data type for blockchains data [7]–[13]. It is very straightforward to monitor the health condition of blockchains by observing the time series data because blockchains is a complex time dependent system. The metrics to monitor can be CPU, memory, network bandwidth usage, system load and disk R/W data, etc.

## II. RELATED WORK

### A. Logging

The purpose of collecting loggings is to gather the real time data from the blockchain. Most of the data that we want to fetch rely on loggings to export them. Therefore, without a doubt, loggings are the most essential and important part for our visualization framework. Yet, the original loggings in ResilientDB are more complicated than we need. To collect those data, we carefully review the code we need on how ResilientDB logs those data.

### B. Parameters Decision

The parameter decision will impact the way and efficiency of manipulating the dataset which we use to save the data exported from the blockchain. The precision of timestamp, the format of timestamp, the compressing mechanism, and the tag for each type of message will directly affect the chart in the frontend. Although high precision and less compression log can bring us meticulous charts, it consumes more storage and CPU power, or larger

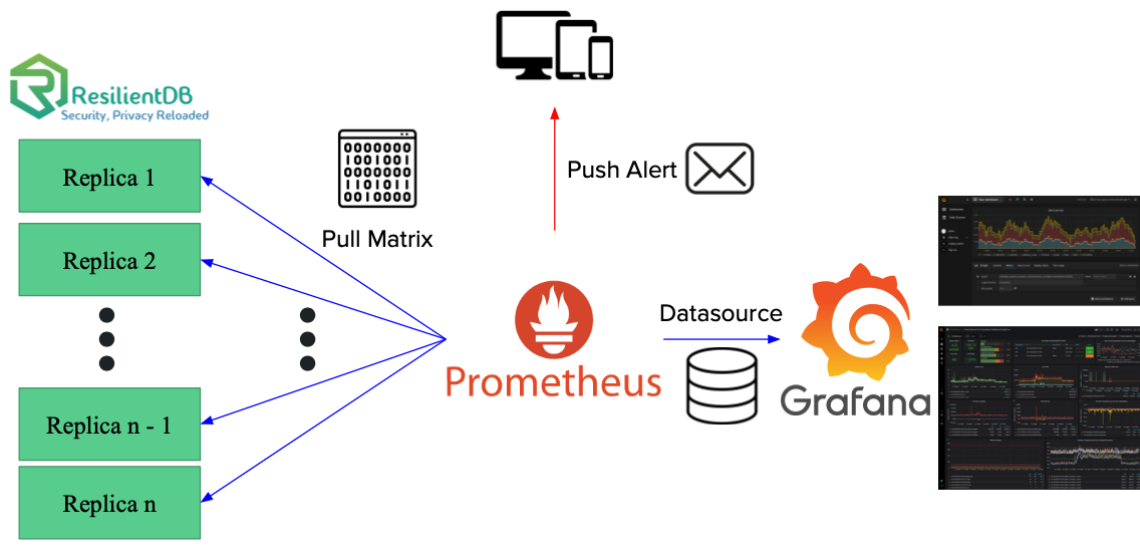


Fig. 1. Flowchart

latency, while rendering the page. Therefore, it will be a trade-off between the user view and server reliability.

### C. Infrastructure Integration

Integrating our library and infrastructure setting to the original codebase is a foreseeable hard work. We will ingrain an exporter on each node, which sends their log to the central unit. The central unit acts as a data collector and a temporary storage. Moreover, we will also set up a web visualization interface as the monitor allowing users to track the status of each component of the blockchain. Both the central unit and the web visualization interface are expected to be built over Docker. As a result, because the intense infrastructure development is involved in this project, the part of bringing up the whole system will be an endeavor.

### D. Original ResilientDB Dashboard

In ResilientDB, we know that it has its own simple visualization implementation, a simple dashboard. After further looking into how the original dashboard works, we realize the current dashboard does not work as we expect. It does not really include the visualization part in the process of installing the ResilientDB. That means if users want to use a dashboard, they have to install a central

unit along with a visualization interface on their own. It requests lots of external work for users to do so. Therefore, in the project, we will build up a solid way for the users who work on ResilientDB in the future to easily install a visualization framework with our project.

### E. ResilientDB Backend

Even though the dashboard would not work right now, we still trace the code to its backend to see how it fetches data from the replica which is an instance in the blockchain. We realized that in the backend of the dashboard it uses some Shell script to get the data it wants. It is not the usual means of backend. Nonetheless, we still take some time to trace down the code and make sure we got some initial data that we needed for our visualization.

### F. Replica IP table

Since we try to fetch data from the replica, it is essential to understand where the data is from. Then, we can ensure the data we get is from the corresponding replica. Yet, the IP table that we get is nothing but a list of IP addresses. Even though we can check the addresses in the container, without any identifier, we can not correspond to the replica, and monitor the status right away. Therefore, while constructing our central unit, we can also connect

the replica through the setting in the central unit, then we don't have to worry about the confusion of the original IP table.

### III. OUR APPROACH

In our approach, we leverage the benefits of Docker [14]. Docker technology has been highly used nowadays. We concatenate the replicas, the central unit, and web visualization interface in the docker network by their domain name. Since the domain name settings are more readable compared to IP settings, using Docker can also address the current IP table problem.

Next, to be able to build new images that integrate all applications we need for ResilientDB, we dive into the source code of ResilientDB, and decide which metrics of replicas are important and necessary for visualization. Then, we capture these features by logging them in the output file. On the other hand, we introduce Prometheus [15] as our central unit, which is an open-source systems monitoring and alerting toolkit. We use the Prometheus server to retrieve the logs from the output file and visualize the statistics of replicas on Grafana [16] Web UI. Also, we configure the alert manager to monitor the status of replicas and notify the administrator by email when errors occur. The Fig. 1 is the flow chart of our system.

#### A. *ResilientDB*

S. Gupta et al. presented the Geo-Scale Byzantine Fault-Tolerant consensus protocol (GeoBFT) [17] to address the problem of geo-scale deployments where many replicas spread across a large area participate in consensus. The main challenge of geo-scale blockchains is to distinguish between local and global communication. Only by distinguishing between these two communication types, can global communication be minimized. GeoBFT performs a topological-aware grouping of replicas into local clusters so that it can minimize global communication. On the other hand, the paper introduces a novel global sharing protocol that optimistically performs minimal inter-cluster communication, still being able to reliably detect communication failures.

The experiment results show that GeoBFT achieves up-to-six times more throughput than existing BFT protocols.

ResilientDB has originally provided a visualization platform. However, we found that the way the dashboard works is not available now. We then found that in the backend of the dashboard, it applies some shell scripts to get the data it needs as shown in Fig. 3. Moreover, the dashboard acquires a pre-established intra-logging server in the internet. Nevertheless, the log server establishing scripts is not existed in the original codebase. Therefore, we find that an error occurs when people uses logging system. Another problem is that the original logging function posts logging data to the log server, InfluxDB, using different shell scripts. The dependency leads dashboard malfunctioning. To be more specific, this approach highly couples with the original log output format. Once the original code has been changed, the log scripts need to be adjusted accordingly. If we implement new function in ResilientDB and want to see the logs of that function, we have to add several log collectors and post functions. It is an unusual way of backend engineering for monitoring. It humongously hampers the future implementation, prolongs the development cycle, and hinders the research on ResilientDB. Therefore, we aim to improve this system by building a visualization tool that makes the service more visible and easily accessible.

#### B. *Prometheus*

Our Monitoring architecture aims to achieve low dependency, high integration, and open for future expansion. Because the monitoring is related to time series databases, log collectors, and visualization, we take several existing open third party time series databases into considering, including InfluxDB [18], Prometheus, and Elasticsearch [19]. The Prometheus have robust community, which gives the database long-term support and comprehensive document. Therefore, we choose Prometheus as our central unit which can provide more scalability for future implementation. Also, Prometheus is developed by Golang, which is highly




Fig. 2. Framework

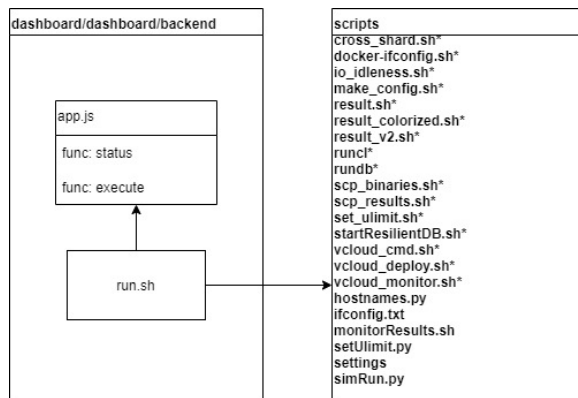


Fig. 3. Component Dependency

efficient language and able to achieve low dependency. More, Prometheus allows huge amount of third party libraries to interact with it. While it is compatible with many UIs, it comes with a default UI which we can first test whether our Prometheus server is working or not. In conclusion, Prometheus is an ideal choice for us as a third-party library to be the central unit. Prometheus is highly scalable as a node exporter and inverts the current dependency of visualization implementation in ResilientDB.

Prometheus is an open source monitoring system developed by engineers at SoundCloud in 2012. In 2016, Prometheus was accepted into the Cloud Native Computing Foundation. The Prometheus monitoring system includes a rich, multidimensional data model, a concise and powerful query language

called PromQL, an efficient embedded timeseries database, and over 150 integrations with third-party systems.

With the help of Prometheus, one can quickly locate the current application's problems, thus handling solutions immediately. The benefit of Prometheus is that we can integrate with any system and extract the part we want to see. For example, when implementing GeoBFT, one can easily observe the network latency with specific metrics data or discover the network throughput of the system. Besides, if there exists an idle node in the network, it would be found out in seconds by observing the inbound/outbound rate of the instance.

### C. Grafana

We use Grafana as our web visualization interface. Grafana was written in Golan and first released in 2014 by Torkel Ödegaard at Orbitz. Grafana is an interactive visualization web application. It demonstrates diagrams, charts, curves and graphs on the web when it connects to data sources. In our work, we use Prometheus as a monitoring solution for storing time series data like metrics, and use Grafana to visualize the data stored in Prometheus, including all the statistics regarding the status of the replicas in ResilientDB.

Grafana provides a handy dashboard to let users organize the metrics data pulled from Prometheus

```

1 FROM docker.io/resilientdb/res-machine
2
3 COPY ./monitor/node_exporter/node_exporter /node_exporter/
4 ENTRYPOINT ["/node_exporter/node_exporter"]
5

```

Fig. 4. Dockerfile Configuration

```

root@localhost:~# docker ps
CONTAINER ID        IMAGE               COMMAND                  CREATED          STATUS          PORTS          NAMES
02748046d1        docker.io/prometheus -- /bin/prometheus --c 7 days ago      Up 3 days      0.0.0.0:9090->9090/tcp    prometheus
0274f023ad        docker.io/resilientdb/res-machine:0.0.4 -- /node_exporter --c 7 days ago      Up 3 days      0.0.0.0:3300->3300/tcp    nrbr
12436c383e        docker.io/resilientdb/res-machine:0.0.4 -- /node_exporter --c 7 days ago      Up 3 days      0.0.0.0:3300->3300/tcp    nrbr
f2080f286a        grafana/grafana    /run.sh                  7 days ago      Up 3 days      0.0.0.0:3000->3000/tcp    grafana
0e013174f4        docker.io/resilientdb/res-machine:0.0.4 -- /node_exporter --c 7 days ago      Up 3 days

```

Fig. 5. Docker Status

```

root@localhost:~# sudo docker images
REPOSITORY          TAG                 IMAGE ID            CREATED            SIZE
docker.io/resilientdb/res-machine 0.0.4              06080c2f7799       7 days ago       410 MB
docker.io/res-machine 0.0.4              06080c2f7799       7 days ago       410 MB
docker.io/resilientdb/res-machine 0.0.3              96e4e26976ff       7 days ago       410 MB
docker.io/resilientdb/res-machine 0.0.2              db938edc21cc       7 days ago       410 MB
docker.io/resilientdb/res-machine 0.0.1              5c2f9e0cfb78       7 days ago       410 MB
docker.io/grafana/grafana latest             ddfae340d088       2 weeks ago      253 MB
docker.io/prom/prometheus latest             227ae28c1b04       2 months ago    193 MB
docker.io/prometheus latest             960100a9344f       2 months ago    4.37 GB
docker.io/ubuntu xenial             b6f507652425       3 months ago    135 MB
docker.io/bash latest             8856e160078        3 months ago    12.9 MB
r_pkg 4.0.3a           f23e93fd5889       11 months ago   614 MB
r_engine 4.0.3a          aab9ead320b1       11 months ago   1.18 GB
r_engine 4.0.3b          c2b9ead42266       11 months ago   1.46 GB
my_r 4              920b13021d98       11 months ago   414 MB
my_r 3              c3d86ea419e5       11 months ago   1.27 GB
my_r 2              347c53233983       11 months ago   290 MB
my_r 1              a5b3077c1de0       11 months ago   1.2 GB
docker.io/r-base latest             964f0d043f8f       12 months ago   204 MB
docker.io/centos 7              8952019fcb4c       12 months ago   204 MB
nbr_sftp 20201019a      7dfcb262935f       13 months ago   154 MB
dockerisk 18.0.2        3a318e37b1be       23 months ago   37.4 MB
nbr_sftp latest             385b6ff904e        2 years ago     394 MB
nbr_analyzer 20180726a      59051bcc38a9       3 years ago     137 MB
nbr_web 20180717b     db11002118c        3 years ago     658 MB
nbr_decoder 20180715a     ba7cc4b22416       3 years ago     381 MB
nbr_finder 20180713b   cba76ff21e07       3 years ago     395 MB
nbr_capturer 20180620c     ca45e08370e3       3 years ago     214 MB
docker.io/mariadb 10.3           389a0b143f41       3 years ago     408 MB

```

Fig. 6. Docker Image

with the highly customizing interface. For example, we can compare the network throughput of each replica to observe the data exchange between them.

## IV. RESULTS

### A. Environment Settings

The new architecture involves three parts, Docker, Prometheus, and Grafana. For the integration part, we create Dockerfile, docker-compose, and modify the original ResilientDB-docker script. In Fig. 4, we ingrain the node\_exporter with original ResilientDB images and modify the Dockerfile. The new image is named docker.io/resilientdb/res-machine with the version number of 0.0.4, as shown in Fig. 6, allows Prometheus to scrape data from each node. We bring up all of the services by using docker-compose, as shown in Fig. 7. We established the ResilientDB essential services: replicas, clients, Prometheus, and Grafana. All services join the same docker network, resilientdb\_default. The network runs in bridge mode and allows services to expose their ports. We can access the docker containers via those exposing ports. In Fig. 8, we adjust the ResilientDB script to prevent it from not creating a new docker-compose file while starting the simulation. Besides, the monotonous docker-compose uses the domain name of the containers, improving readability and simplifying the configuration.

### B. Prometheus

In Figure 9 we set the basic information of replicas in the YAML file for Prometheus to connect to the instances. And the status of the Prometheus

```

1 version: '3.1'
2 services:
3   s1:
4     image: docker.io/resilientdb/res-machine:0.0.4
5     container_name: s1
6     volumes:
7       - ./:/home/expo/resilientdb
8     restart: always
9   s2:
10    image: docker.io/resilientdb/res-machine:0.0.4
11    container_name: s2
12    volumes:
13      - ./:/home/expo/resilientdb
14    restart: always
15  s3:
16    image: docker.io/resilientdb/res-machine:0.0.4
17    container_name: s3
18    volumes:
19      - ./:/home/expo/resilientdb
20    restart: always
21  s4:
22    image: docker.io/resilientdb/res-machine:0.0.4
23    container_name: s4
24    volumes:
25      - ./:/home/expo/resilientdb
26    restart: always
27  c1:
28    image: docker.io/resilientdb/res-machine:0.0.4
29    container_name: c1
30    volumes:
31      - ./:/home/expo/resilientdb
32    restart: always
33
34  monitor_db:
35    image: docker.io/prom/prometheus
36    container_name: prometheus
37    ports:
38      - 9090:9090
39    volumes:
40      - ./monitor/prometheus:/etc/prometheus
41    command:
42      - --config.file=/etc/prometheus/prometheus.yml
43    restart: always
44
45  monitor_ui_grafana_alpine:
46    image: grafana/grafana
47    container_name: grafana
48    ports:
49      - 3000:3000
50    restart: always
51

```

Fig. 7. Docker Compose

after executing the YAML is shown as Fig. 10. We have already pulled the metrics data from each replica, and we only take several data as examples in our project demo. The data we retrieve is enough for monitoring the health status of the whole system. The Prometheus provides a basic UI dashboard as shown in Fig. 11, but later we will further integrate with Grafana with a more user-friendly interface.

```

64 clients=1
65 fi
66 echo -e "Number of Replicas:\t${replicas}"
67 echo -e "Number of Clients:\t${clients}"
68
69 if [ -f "docker-compose.yml" ]; then
70   echo "Stopping previous containers..."
71   command -v docker-compose >/dev/null 2>&1 || {
72     printf "${RED}ResilientDB requires docker and docker-compose\nAborting...${NC}\n"
73     exit 1
74   }
75   $(docker-compose down)
76   printf "${GREEN}Successfully stopped${NC}\n"
77 fi
78
79 #echo "Creating docker compose file ..."
80 #echo "version: '3'" >>docker-compose.yml
81 #echo "services:" >>docker-compose.yml
82
83 #t- "
84 #for i in $(seq 1 $replicas); do
85 #  echo -e "${t}${i}:" >>docker-compose.yml
86 #  echo -e "${t}${t}image: resilientdb/res-machine" >>docker-compose.yml
87 #  echo -e "${t}${t}container_name: s${i}" >>docker-compose.yml
88 #  echo -e "${t}${t}volumes:" >>docker-compose.yml
89 #  echo -e "${t}${t}${t} - ./:/home/expo/resilientdb" >>docker-compose.yml
90 #  echo -e "${t}${t}restart: always" >>docker-compose.yml
91 #done
92 #for i in $(seq 1 $clients); do
93 #  echo -e "${t}c${i}:" >>docker-compose.yml
94 #  echo -e "${t}${t}image: resilientdb/res-machine" >>docker-compose.yml
95 #  echo -e "${t}${t}container_name: c${i}" >>docker-compose.yml
96 #  echo -e "${t}${t}volumes:" >>docker-compose.yml
97 #  echo -e "${t}${t}${t} - ./:/home/expo/resilientdb" >>docker-compose.yml
98 #  echo -e "${t}${t}restart: always" >>docker-compose.yml
99 #done
100 #printf "${GREEN}Docker compose file created -> docker-compose.yml${NC}\n"
101
102 echo "Starting the containers..."
103 $(docker-compose up -d)
104
105 scripts/docker-ifconfig.sh
106
107 printf "\nChecking Dependencies...\n"
108 if [ -d "deps/crypto" ]; then
109   printf "Installing dependencies...\n"
110   cd deps && $(ls | xargs -n 1 tar -xvf 2>/dev/null)
111   cd ..
112 fi
113 pwd
114 printf "${GREEN}Dependencies has been installed${NC}\n"

```

Fig. 8. ResilientDB Shell Script

```

1 global:
2 # How frequently to scrape targets by default.
3 scrape_interval: 1m
4
5 # How long until a scrape request times out.
6 scrape_timeout: 10s
7
8 # How frequently to evaluate rules.
9 evaluation_interval: 1m
10
11 # The labels to add to any time series or alerts when communicating with
12 # external systems (federated, remote storage, Alertmanager).
13 external_labels:
14   monitor: 'replicas'
15
16 # A list of scrape configurations.
17 scrape_configs:
18   - job_name: 'prometheus'
19     scrape_interval: 5s
20
21     static_configs:
22       - targets: ['localhost:9090']
23
24   - job_name: 'docker'
25     static_configs:
26       - targets: ['localhost:9323']
27
28   - job_name: 's1'
29     static_configs:
30       - targets: ["s1:9100"]
31     scrape_interval: 5s
32
33   - job_name: 's2'
34     static_configs:
35       - targets: ["s2:9100"]
36     scrape_interval: 5s
37
38   - job_name: 's3'
39     static_configs:
40       - targets: ["s3:9100"]
41     scrape_interval: 5s
42
43   - job_name: 's4'
44     static_configs:
45       - targets: ["s4:9100"]
46     scrape_interval: 5s
47

```

Fig. 9. Prometheus Configuration

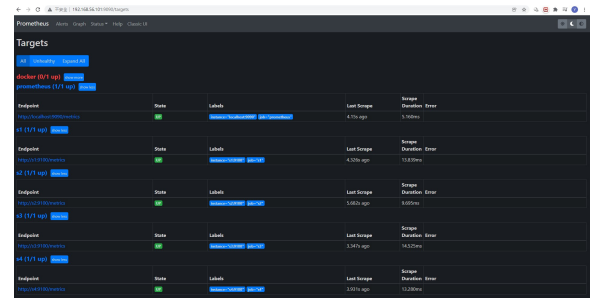


Fig. 10. Prometheus UI

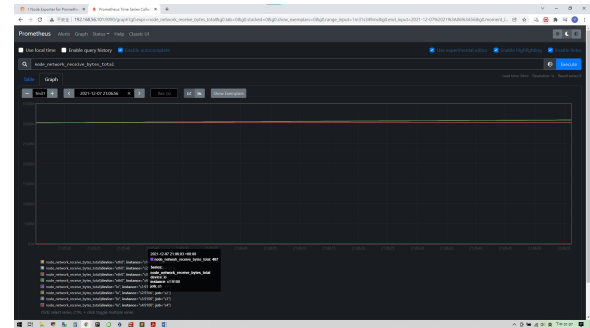


Fig. 11. Prometheus UI Dashboard

### C. Grafana

On Grafana UI web, we can choose which data source we would like to connect. In our case, obviously, we choose Prometheus server as our data sources as shown Fig. 15. We can show the status of all replicas on the Grafana web UI as shown in Fig. 12, including CPU, memory, partition usage. Also, we can see the transmit and receive status among replicas. If we are interested in one specific replica, we can pull out its status, and monitor its statistics in detail as shown in Fig. 13 and Fig. 14. From these figures, we can see the time series values of all monitored metrics. This feature is significant especially when the network is unusual, and we want to find which replica is abnormal.

## V. FUTURE WORK

The current progress of this project has already implemented the basic structure of the system, such as retrieving metrics data from replicas and the customizable visualized dashboard. Prometheus provides more than just extracting the data.

Besides, in the next step, we will implement an alert manager to notify users if the system

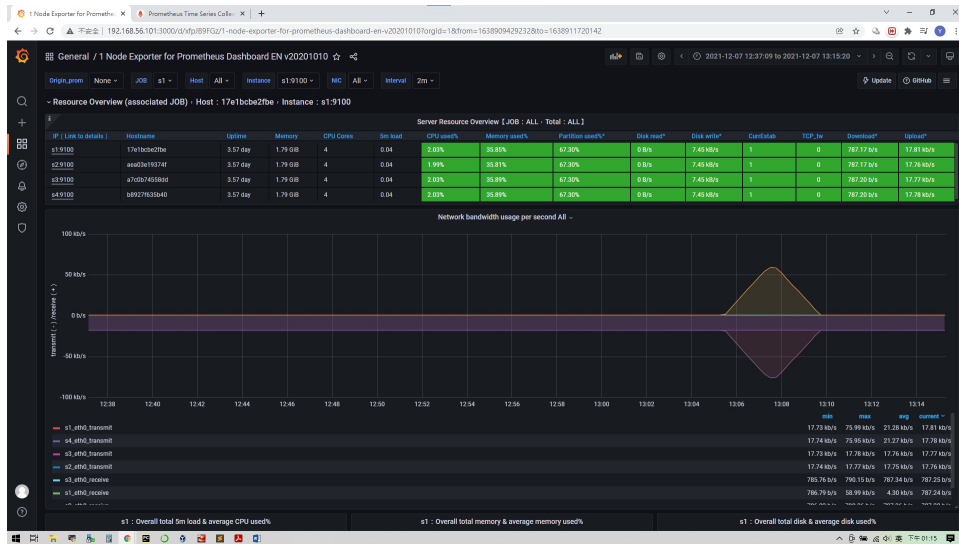


Fig. 12. Grafana UI

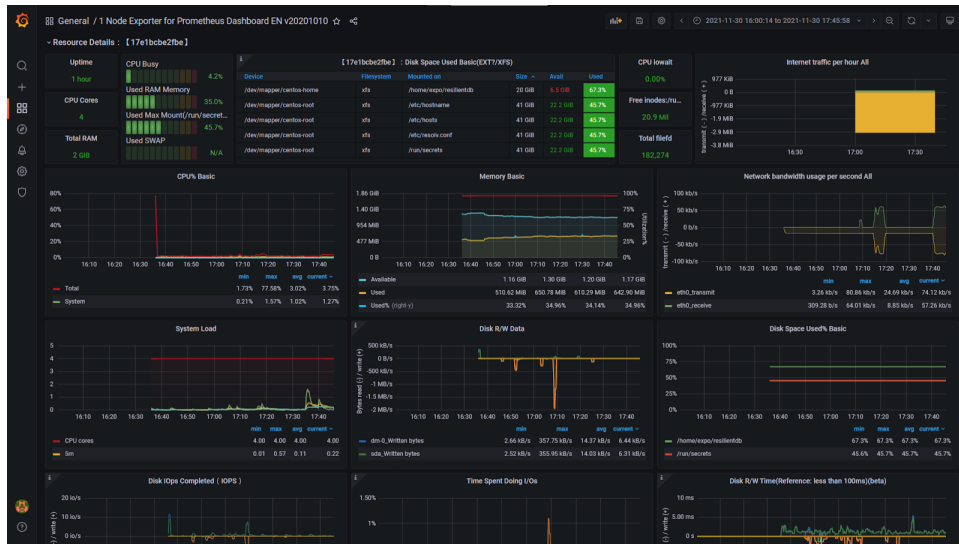


Fig. 13. Grafana UI

encounters some issue in the future, thus resolving the problems in a short time. Also, Grafana can integrate with other related data sources with the whole blockchain networks, such as the data from different chains.

On the other hand, there is more work we can do on ResilientDB. Now, the logs we extract from ResilientDB are only information of each replicas such as CPU and memory usage. In the next step, we expect to store more information about the state of the entire system, especially at which stage it is in, including pre-commit, commit, and view changes. If we can show the stages of ResilientDB,

it can help users learn more about the process and health condition of each transaction.

Moreover, in our work, we manually write the docker compose file and Prometheus YAML file. However, it is not ideal for system integration. Therefore, in the future, we would like to design a shell script which is able to automatically generate docker compose file and Prometheus YAML file.

## VI. CONCLUSIONS

In this work, we build a visualization system based on ResilientDB, which is the Geo-Scale Byzantine Fault-Tolerant consensus proto-



Fig. 14. Grafana UI

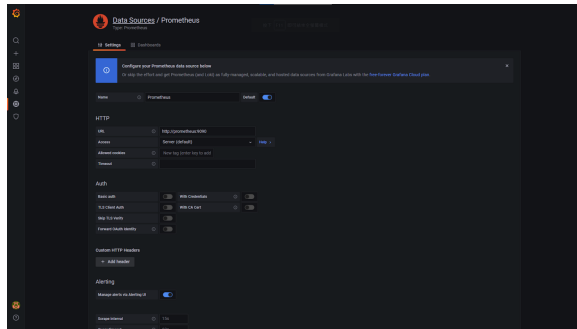


Fig. 15. Data Source

col (GeoBFT) aiming to address the problem of geo-scale deployments where many replicas spread across a large area participate in consensus. Since the original visualization platform of ResilientDB has issues of components dependency, we construct a new visualization system to provide users and learners to understand the mechanisms and status of replicas more easily. We built the Prometheus server to pull all the statistics of replicas in ResilientDB. Then, we build the Grafana user interface, and let users decide which metrics they would like to monitor. When users query the metrics, the Grafana will pull the specific metrics from the Prometheus server, and demonstrate the charts and curves on the web UI. This work is very valuable for providing an independent service to monitor the communication process within a protocol. Moreover, our work has

high scalability. For those users and developers on ResilientDB, it is costly-less for integrating our project into the current ResilientDB. Our work can not only monitor ResilientDB, but it can be used on a variety of protocols.

## REFERENCES

- [1] “Bitbonkers.com,” <https://bitbonkers.com/>.
- [2] “Bitcoin transaction visualization.com,” <http://bitcoin.interaqt.nl/>.
- [3] “Bitnodes.earn.com,” <https://bitnodes.earn.com/>.
- [4] “Realtime bitcoin globe,” <https://blocks.wizb.it/>.
- [5] “Blockchain.info,” <https://blockchain.info/tree/114688199>.
- [6] “Elliptic.co,” <https://www.elliptic.co/>.
- [7] A. Bogner, “Seeing is understanding: Anomaly detection in blockchains with visualized features,” p. 5–8, 2017.
- [8] C. Kinkeldey, J. Fekete, T. Blascheck, and P. Isenberg, “Visualizing and analyzing entity activity on the bitcoin network,” *CoRR*, vol. abs/1912.08101, 2019.
- [9] H. Kuzuno and C. Karam, “Blockchain explorer: An analytical process and investigation environment for bitcoin,” pp. 9–16, 2017.
- [10] D. McGinn, D. Birch, D. Akroyd, M. Molina-Solana, Y. Guo, and W. Knottenbelt, “Visualizing dynamic bitcoin transaction patterns,” *Big Data*, vol. 4, pp. 109–119, 06 2016.
- [11] D. McGinn, D. McIlwraith, and Y. Guo, “Toward open data blockchain analytics: A bitcoin perspective,” *CoRR*, vol. abs/1802.07523, 2018.
- [12] S. Meiklejohn, M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. M. Voelker, and S. Savage, “A fistful of bitcoins: Characterizing payments among men with no names,” *Commun. ACM*, vol. 59, no. 4, p. 86–93, mar 2016.
- [13] L. Norbutas, “Offline constraints in online drug marketplaces: An exploratory analysis of a cryptomarket trade network.” *The International journal on drug policy*, vol. 56, pp. 92–100, 2018.



- [14] “Docker,” <https://www.docker.com/>.
- [15] “Prometheus,” <https://prometheus.io>.
- [16] “Grafana,” <https://github.com/grafana/grafana>.
- [17] S. Gupta, S. Rahnama, J. Hellings, and M. Sadoghi, “Resilientdb: Global scale resilient blockchain fabric,” *CoRR*, vol. abs/2002.00160, 2020.
- [18] “Influxdb,” <https://www.influxdata.com/>.
- [19] “Elasticsearch,” <https://www.elastic.co/>.