




国家高能物理科学数据中心

National HEP Science Data Center



高能所计算中心


IHEP Computing Center

perfSONAR@IHEP Data Center

Mengyao Qi, Shan Zeng, Tao Cui

IHEP Computing Center

2024-10-11




1 Introduction and Challenges

2 Current status

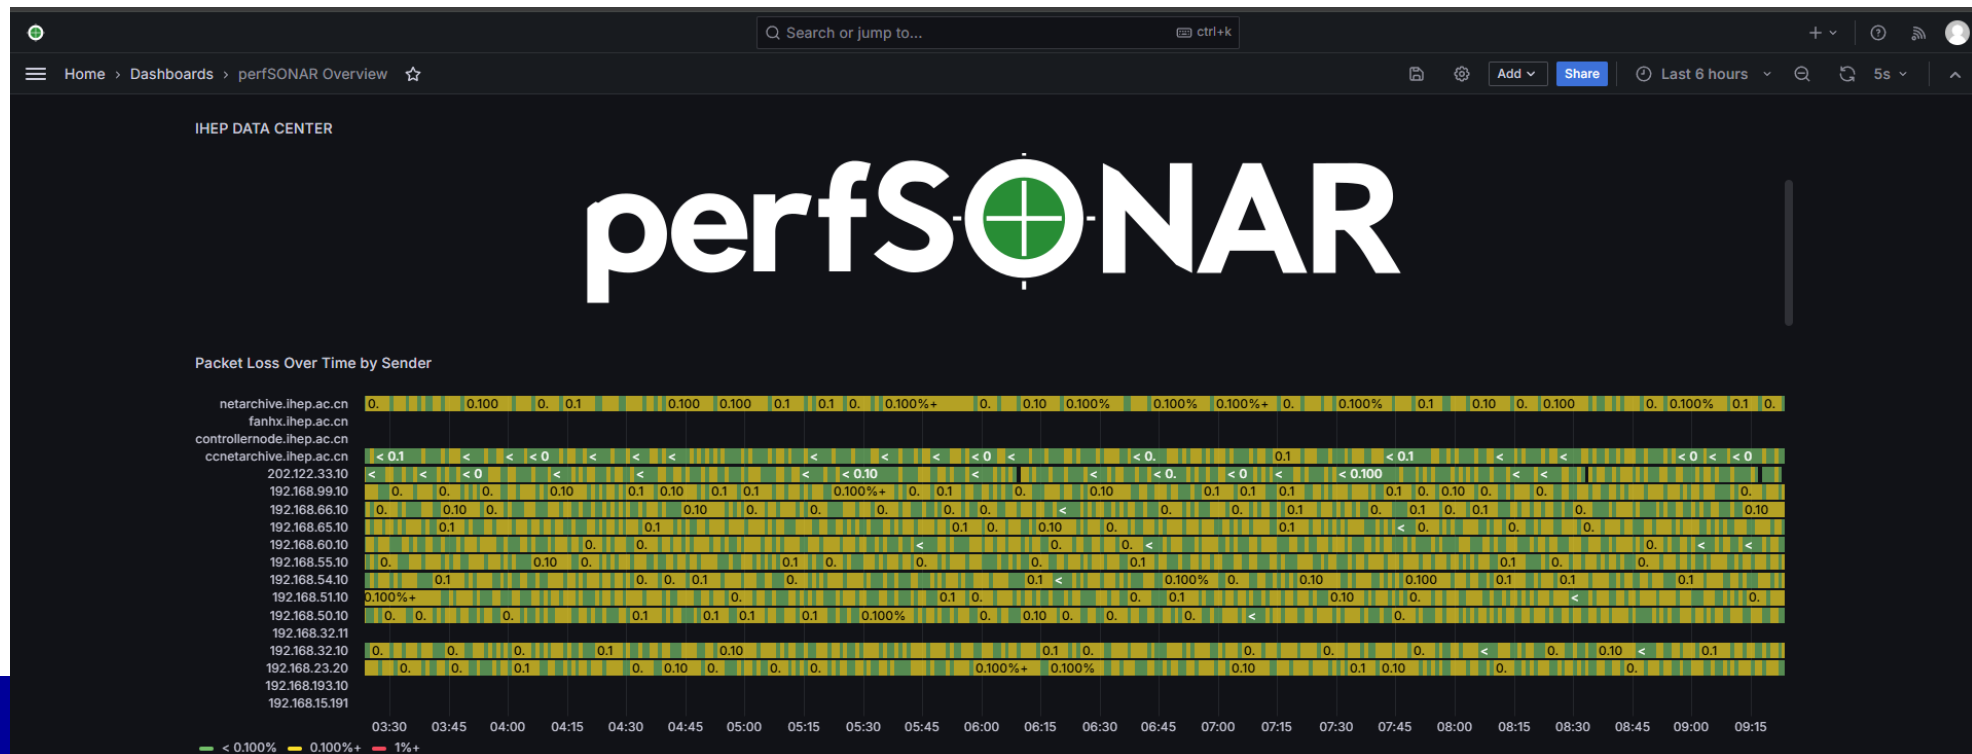
3 Future Plan

4 Summary


Introduction



- Importance of network performance measurement in data centers
 - High-energy physics research relies heavily on data transfer and network stability
- Key metrics: latency, jitter, bandwidth, packet loss, throughput
- To ensure optimal performance and reliability of services



Challenges



❖ Technical :

- Compatibility issues, unexpected test failures

❖ Learning Curve:

- New features and configurations in the latest version

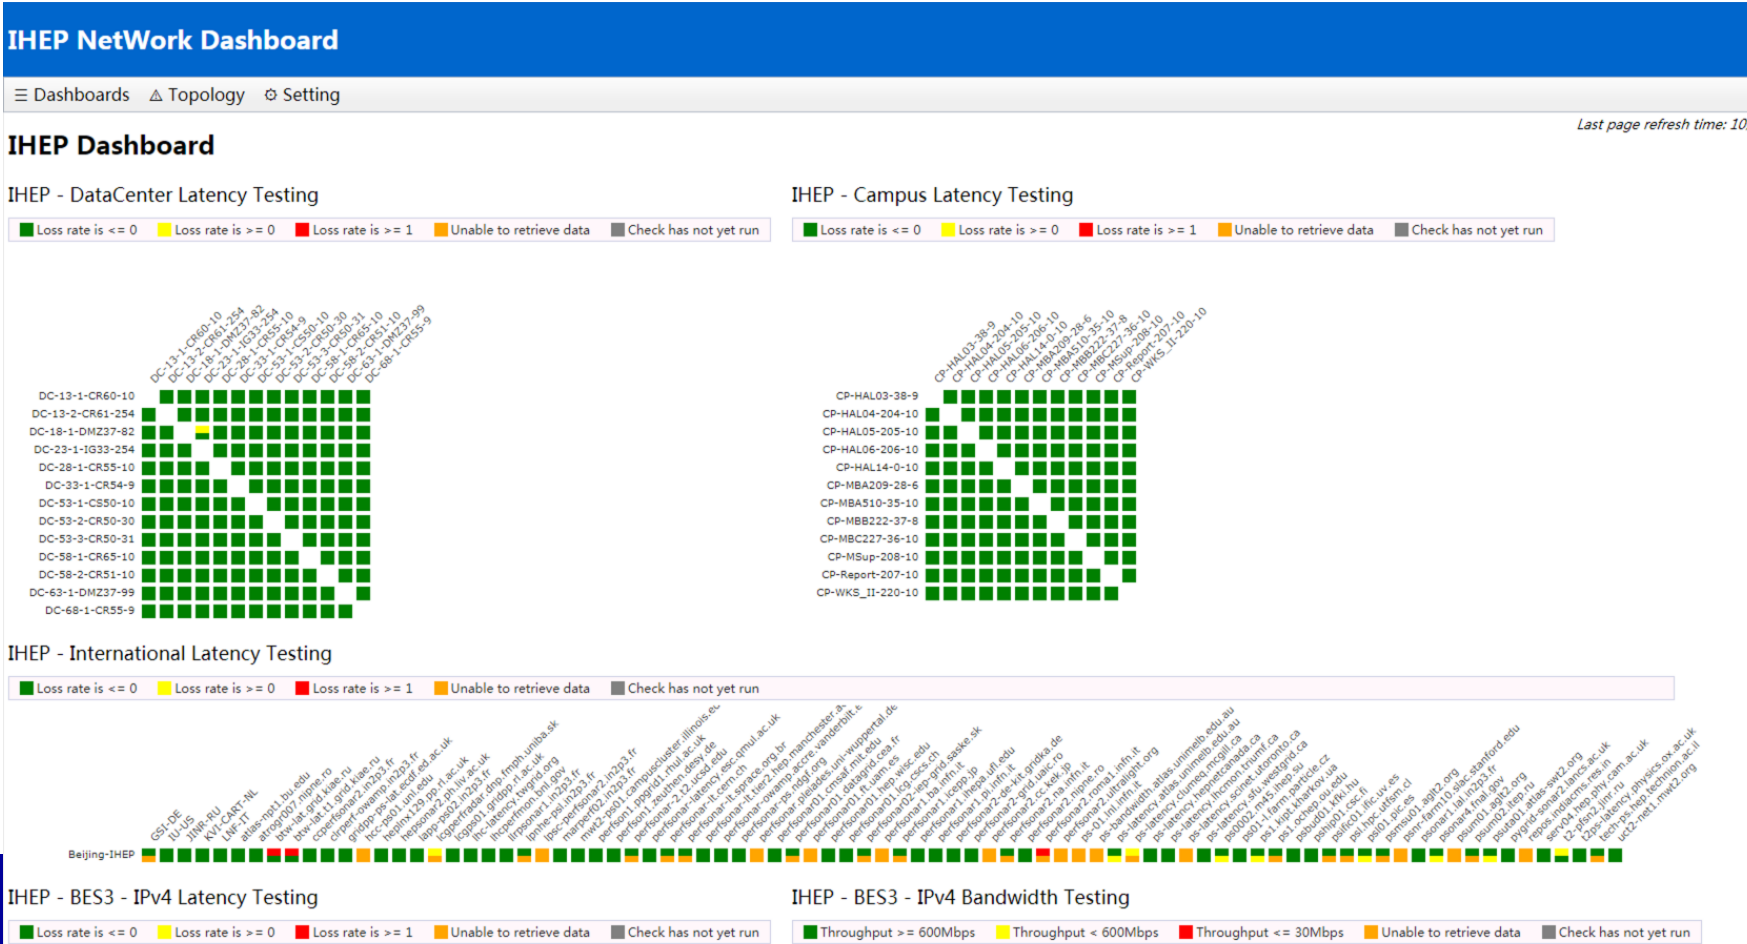


❖ Challenges:

- Limitations of older versions, scalability issues
- A vast number of data center switches and IP subnets.


Current status

❖ Initial Deployment:

- Overview of the starting point, including version and scale



Upgrading



❖ Upgrade Process:

- Centos 6.x to Almalinux9.4
- The version of perfSONAR from 4.x to 5.1.3
- Use a high-performance server as the core node(central opensearch instance).


```
[root@perf-99-10 ~]# df -hl
Filesystem                Size      Used Avail Use% Mounted on
devtmpfs                   4.0M         0  4.0M   0% /dev
tmpfs                      1.8G    440K  1.8G   1% /dev/shm
tmpfs                      712M     66M   647M  10% /run
/dev/mapper/almalinux-root 26G     9.6G   17G   38% /
/dev/sda1                  960M    412M   549M  43% /boot
tmpfs                      356M         0  356M   0% /run/user/0
```

❖ Benefits of Upgrading:

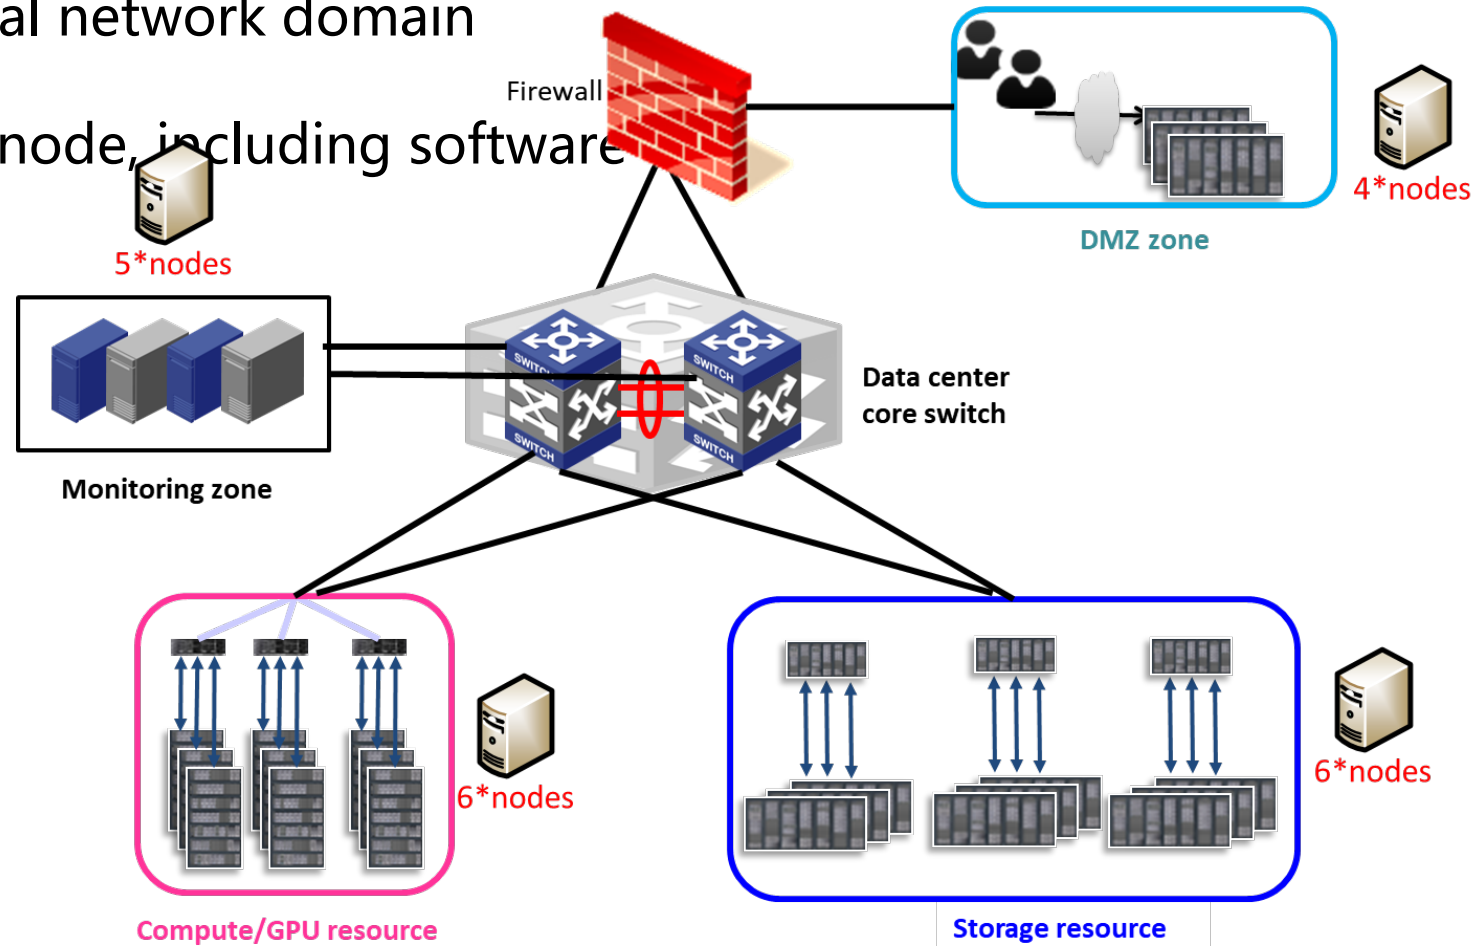
- Enhanced features, better performance, improved user interface, and security patches

```
[root@perf-99-10 ~]# top
top - 09:24:29 up 20 days, 17:17,  1 user,  load average: 0.43, 0.37, 0.41
Tasks: 326 total,  1 running, 322 sleeping,  0 stopped,  3 zombie
%Cpu(s):  2.9 us,  5.9 sy,  0.0 ni, 91.2 id,  0.0 wa,  0.0 hi,  0.0 si,  0.0 st
MiB Mem :  3558.2 total,  800.0 free,  2567.3 used,  568.6 buff/cache
MiB Swap:  3056.0 total,  1324.1 free,  1731.9 used.  990.9 avail Mem
```



perfSONAR Topology@IHEP DC



- ❖ Creation of over 20+ test nodes across the data center network
- ❖ Strategic distribution to cover all critical network domain
- ❖ Configuration: Detailed setup of each node, including software installation and test configurations



Network Performance Measurement Tests



- ❖ **Types of Tests:** Ping, traceroute, iperf3, owamp, bwctl, throughput
- ❖ **Test Frequency:** Continuous monitoring with periodic in-depth tests
- ❖ **Data Collection:** Real-time and historical data storage for analysis
- ❖ **Data Sources:** OpenSearch


Configure tests between this host and other hosts.

+ Host

+ Test

View by: Test | [Host](#)

TEST NAME	TYPE	INTERVAL	TEST MEMBERS	ENABLED	ACTIONS
one-way latency	One-way latency		15 hosts	☑	⚙️ 🗑️
perfSONAR Toolkit Default Traceroute Test	Traceroute	10 minutes	0 hosts	☑	⚙️ 🗑️
round-trip latency	Round-trip latency	30 minutes	15 hosts	☑	⚙️ 🗑️
throughput	Throughput - TCP	1 hour	15 hosts	☑	⚙️ 🗑️
traceroute	Traceroute	1 hour	15 hosts	☑	⚙️ 🗑️



perf5ONAR on ccnetarchive.ihep.ac.cn


Home Configuration ? Help

Throughput (Max)
948 Mb/s

Packet Loss (Max)
0.833%

Jitter (Max) **119 ms**

Number of Tests Run



Interfaces

IP	Speed
	n/a
	n/a
	n/a

Test Results by Type

Type	#
trace	162
throughput	127
rtt	12

Host Info

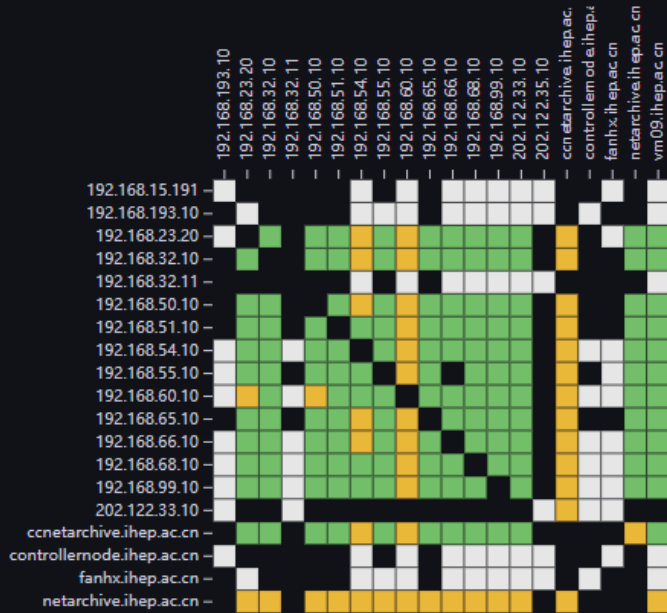
My Dashboards

Test Result Summary

Source	Destination	IP Version	Packet Loss (Max)	Throughput (Min)	Latency (Min)	RTT (Min)
netarchive.ihep.ac.cn	ccnetarchive.ihep.ac.cn	4	0.667%	154 Mb/s	-346 ms	
netarchive.ihep.ac.cn	202.122.33.10	4	0.500%	872 Mb/s	0.0200 ms	
netarchive.ihep.ac.cn	192.168.99.10	4	0.500%	831 Mb/s	-0.0900 ms	
netarchive.ihep.ac.cn	192.168.66.10	4	0.833%	821 Mb/s	-0.01000 ms	
netarchive.ihep.ac.cn	192.168.65.10	4	0.333%	866 Mb/s	-0.01000 ms	
netarchive.ihep.ac.cn	192.168.60.10	4	0.667%	833 Mb/s	0.0300 ms	

pSConfig Dashboards

Maximum Packet Loss



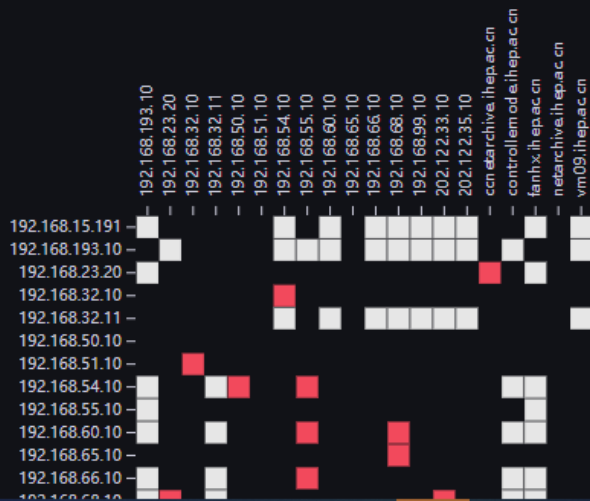
Highest Average Packet Loss by Sender

Sender Host	Max Packet Loss
192.168.55.10	0.500%
192.168.66.10	0.500%
192.168.23.20	0.333%
192.168.32.10	0.333%
192.168.50.10	0.333%
192.168.51.10	0.333%
192.168.54.10	0.333%
192.168.60.10	0.333%
192.168.65.10	0.333%
192.168.68.10	0.333%

Highest Average Packet Loss by Receiver

Receiver Host	Max Packet Loss
192.168.60.10	0.500%
ccnetarchive.ihep.ac.cn	0.500%
192.168.54.10	0.333%
192.168.55.10	0.333%
vm09.ihep.ac.cn	0.333%
192.168.23.20	0.167%
192.168.32.10	0.167%
192.168.50.10	0.167%
192.168.51.10	0.167%
192.168.65.10	0.167%

Minimum Throughput



Lowest Minimum Throughput by Sender


Sender Host	Receiver Throughput
192.168.23.20	287 Mb/s
192.168.68.10	932 Mb/s
192.168.60.10	932 Mb/s
192.168.54.10	934 Mb/s
192.168.66.10	935 Mb/s
ccnetarchive.ihep.ac.cn	935 Mb/s
192.168.65.10	937 Mb/s
192.168.99.10	937 Mb/s
192.168.51.10	937 Mb/s

Lowest Minimum Throughput by Receiver


Receiver Host	Min Throughput
ccnetarchive.ihep.ac.cn	287 Mb/s
192.168.23.20	932 Mb/s
192.168.68.10	932 Mb/s
192.168.50.10	934 Mb/s
192.168.55.10	935 Mb/s
netarchive.ihep.ac.cn	935 Mb/s
202.122.33.10	936 Mb/s
192.168.32.10	937 Mb/s
192.168.54.10	941 Mb/s

Source ccnetarchive.ihep.ac.cn Destination 192.168.32.10 Filters +


Throughput




Retransmits





Latency



Latency (Reverse)


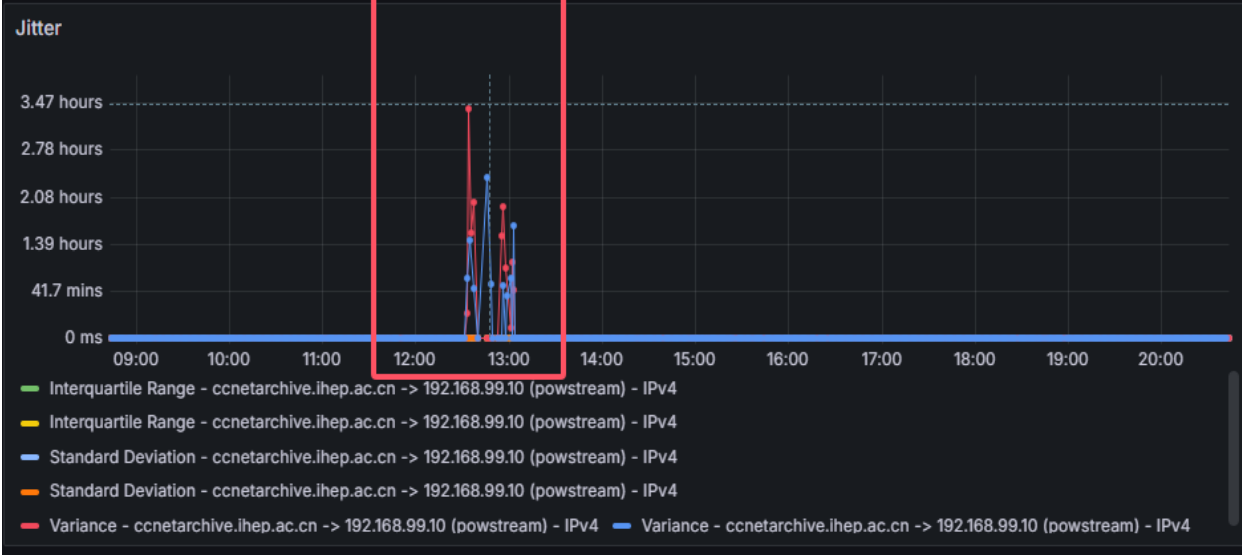







```
[root@perf-32-10 ~]# ifconfig
enp2s0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 192.168.32.10 netmask 255.255.255.0 broadcast 192.168.32.255
inet6 2401:de00:1:2032::9669 prefixlen 128 scopeid 0x0<global>
inet6 fe80::2e0:4cff:fe69:3cfa prefixlen 64 scopeid 0x20<link>
ether 00:e0:4c:69:3c:fa txqueuelen 1000 (Ethernet)
RX packets 3626365612 bytes 4785341488435 (4.3 TiB)
RX errors 0 dropped 414224 overruns 0 frame 0
TX packets 3843783958 bytes 5161542930266 (4.6 TiB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

net_ipv4_tcp_rmem_default	87380	204800
net_ipv4_tcp_rmem_max	33554432	33554432
net_ipv4_tcp_rmem_min	4096	102400
net_ipv4_tcp_wmem_default	65536	102400
net_ipv4_tcp_wmem_max	33554432	33554432
net_ipv4_tcp_wmem_min	4096	51200








Average CPU Utilization



Lessons Learned



❖ 1: Latency Reduction

- **Scenario:** Elevated latency between data servers and analysis workstations
- **Analysis:** Used perfSONAR data to isolate the problematic segment
- **Solution:** Upgraded network hardware


❖ 2: Packet Loss Mitigation

- **Scenario:** Frequent packet loss during high-data-rate transfers
- **Root Cause:** Faulty interface on a tor switch
- **Resolution:** Replaced the faulty hardware and implemented redundancy

❖ **Importance of Regular Updates:** Keeps the system secure and functional

❖ **Comprehensive Testing:** Ensures reliability before full-scale deployment

❖ **Collaborative Effort:** Engaging with the perfSONAR community for support and best practices



❖ Expanding Network Coverage

- **Network Segment Expansion:** Plan to include more network segments in the monitoring system
- **Enhanced Visibility:** Fuller figure of network health and performance


❖ Scalability Considerations

- **Horizontal Scaling:** Adding more test nodes as the network grows
- **Vertical Scaling:** Upgrading existing nodes with more powerful hardware

❖ Enhancements to Monitoring Capabilities

- New features to be integrated into the monitoring system
- Plans for integrating AI and machine learning for predictive analysis
- Enhancing visualization capabilities

Summary



- ❖ Overview of the perfSONAR deployment at IHEP data center
- ❖ Network Performance Analysis System works well
- ❖ Positive impact on network performance and research productivity
- ❖ Optimistic future with ongoing improvements and expansion