



NDN-DPDK File Server Benchmark on FABRIC

Junxiao Shi, Davide Pesavento, Lotfi Benmohamed Communications Technology Laboratory National Institute of Standards and Technology



1

FABRIC: an international research infrastructure

- Testbed with large amounts of compute and storage.
- 33 sites around the world.
- High speed, dedicated optical links (100 Gbps speed).



NDN DPDK

What is NDN-DPDK

- Software implementation of the NDN protocol.
- Aimed at high performance.

Main components:

- ✓ network forwarder
- ✓ traffic generator
- \checkmark file server

File Server Protocol

1. Retrieve metadata.

/files/dir/file/32=metadata

Metadata: result of STATX syscall.

- Versioned name (derived from modify time).
- File type and Unix permissions.
- File size, segment size, number of segments.
- Create time, modify time.

This format extends the RDR protocol.

File does not exist – application Nack.

covered by benchmark

- 2. Fetch file content.
 /files/dir/file/v=7/seg=0
- 3. List directory.
 /files/dir/32=ls/v=7/seg=0

Deploy the Experiment on FABRIC

NDN-DPDK Benchmark Webapp

- Part of NDN-DPDK codebase.
- Producer: file server
 - Access local filesystem via io_uring API provided by the Linux kernel.
- Consumer: congestion-aware fetcher
 - CUBIC-like congestion control, parameters tuned for low-latency networks.



FABRIC Experiment Setup

- 1. Create 3-node slice.
- 2. Download, compile, and install NDN-DPDK on each node.
- 3. Configure CPU isolation in systemd, reboot, enable hugepages.
- 4. Upload .env configuration, start benchmark webapp.



spec	forwarder	traffic gen	
CPU	24 cores	12 cores	
RAM	32 GB	32 GB	
OS image	Ubuntu 22	Ubuntu 22	
OS disk	100 GB	100 GB	
NIC	2 ports	1 port	
NVMe	none	1	

FABRIC NIC and NVMe Components



NVM EXPRESS

NVMe storage device: Intel P4510 or similar

- 1TB capacity, NVMe 1.2 or 1.3
- ext4 filesystem
- populated with large files filled with random data

Location Choices: local, near, far



F GQLSERVER=http://192.168.164.44:3030 F NUMA PRIMARY=0 F CORES PRIMARY=6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23 Setup Script Output F_CORES_SECONDARY=4,5 A_GQLSERVER=http://192.168.164.43:3030 A NUMA PRIMARY=0 A_CORES_PRIMARY=6,7,8,9,10,11 A_CORES_SECONDARY=4,5 A_FILESERVER_PATH=/srv/fileserver B_GQLSERVER=http://192.168.164.42:3030 .env file for benchmark webapp B_NUMA_PRIMARY=0 (already uploaded to the nodes) B CORES PRIMARY=6,7,8,9,10,11 B_CORES_SECONDARY=4,5 B_FILESERVER_PATH=/srv/fileserver F_PORT_A=00:08.0 PCI addresses • F_VLAN_A=100 MAC addresses F_HWADDR_A=04:3f:72:cb:03:8a F_PORT_B=00:09.0 VI AN IDs F VLAN B=100 F_HWADDR_B=04:3f:72:cb:03:8b available CPU cores A PORT F=00:08.0 GraphQL endpoint URIs A VLAN F=100 A_HWADDR_F=b8:ce:f6:5d:45:4e B PORT F=00:09.0 B VLAN F=100 B HWADDR F=08:c0:eb:4e:b2:12 Instructions for: Running as unit: ndndpdk-benchmark.service Running as unit: ndndpdk-status-F.service 1. Establish SSH tunnel Running as unit: ndndpdk-status-A.service Running as unit: ndndpdk-status-B.service 2. Access benchmark webapp via browser NDN-DPDK benchmark is ready. Open an SSH tunnel: ssh -L3333:127.0.0.1:3333 -L8006:127.0.0.1:8006 -L8001:127.0.0.1:8001 -L8002:127.0.0.1:8002 -i ~/.ssh/fabric-sliver-20220908 -J fabric ubuntu@2001:400:a100:3030:f816:3eff: feaa:6f81 Access in browser:

http://localhost:3333 - benchmark webapp

http://localhost:8006 - forwarder F status

http://localhost:8001 - trafficgen A status

http://localhost:8002 - trafficgen B status

Run the Benchmark on FABRIC



Benchmark Webapp Usage

- 1. Input parameters.
- 2. Start file servers and fetchers.
- 3. Gather and show statistics.

running	trial	3
---------	-------	---

timestamp	duration	Data packets throughput	goodput	
4:54:07 PM	387.209 s	1.065 Mpps	68.150 Gbps	
5:01:53 PM	379.460 s	1.087 Mpps	69.542 Gbps	
current	in progress			
average	383.335 s	1.076 Mpps	68.846 Gbps	
stdev	5.479 s	0.015 Mpps	0.984 Gbps	
			4.0	

Preliminary Results

unique Data packets only excludes Interests & retx content payload only excludes name & signature

	sites A-F-B	NIC	duration	Data throughput	goodput
local-	FIU-FIU-FIU	NIC_Basic	342.833 s	1.203 Mpps	76.981 Gbps
	STAR-STAR-STAR	ConnectX-5 25Gbps	554.743 s	0.743 Mpps	47.569 Gbps
	DALL-DALL-TACC	NIC_Basic	380.541 s	1.084 Mpps	69.384 Gbps
near	MAX-MAX-WASH	ConnectX-5 25Gbps	635.654 s	0.649 Mpps	41.516 Gbps
	DALL-DALL-TACC	ConnectX-6 100Gbps	381.705 s	1.082 Mpps	69.237 Gbps
far -	MAX-GPN-UCSD	NIC_Basic	1573.780 s	0.262 Mpps	16.770 Gbps
	MAX-GPN-UCSD	ConnectX-5 25Gbps	1552.970 s	0.266 Mpps	17.001 Gbps
	MAX-GPN-UCSD	ConnectX-6 100Gbps	1396.210 s	0.295 Mpps	18.911 Gbps

average over 5 runs

NDN-DPDK commit eae4b295 tested on 2023-04-11 ~ 2023-04-13

Preliminary Results: Goodput (Gbps)



File Server Implementations and Deployments

File Server Implementations

Client implementations

- ndncatchunks
- NDN-DPDK fetcher
- NDNgo ndn6file package
- NDNts @ndn/fileserver
- NDNC (from N-DISE project)

Server implementations

- ndn6-file-server (ndn-cxx)
- NDN-DPDK fileserver
- proposed hackathon project to implement in NDNd





File Server Deployments

- Testbed Status Page <u>https://testbed-status.named-data.net</u>
 - Status data is retrieved from ndn6-file-server running on each testbed router.
- NDNts adaptive video https://pushups.ndn.today
 - Video clips are served by ndn6-file-server.
- N-DISE (published at ACM-ICN 2022).
 - Genomics data is served by NDN-DPDK file server.
- NDN testbed traffic traces (published at ACM-ICN 2023)
 - https://www.tntech-ngin.net/datasets/
 - Currently distributed through GitHub LFS (Large File Storage).
 - May transition to HTTPS + NDN file server.

NDN-DPDK Codebase

https://github.com/usnistgov/ndn-dpdk

- forwarder
 - expanded NIC offloads
 - enhanced test coverage
- traffic generator
- file server

- GraphQL management tools
- interactive benchmark webapp
 - with file server support
- NDNgo library
- more user documentation



Information + Q&A

- NDN-DPDK related papers
 - NDN-DPDK: NDN Forwarding at 100 Gbps on Commodity Hardware
 - <u>N-DISE: NDN-based data distribution for large-scale data-intensive science</u>
- NDN-DPDK repository: <u>https://github.com/usnistgov/ndn-dpdk</u>
- Junxiao's FABRIC scripts: https://github.com/yoursunny/fabric

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.