



Virtual Conference June 8, 2021

"The Network is the Computer" - Revisited

Christopher R. Hertel Samba Team



Generic Disclaimer

The opinions expressed are my own and not necessarily those of:

- my spouse,
- my dog,
- my spirit familiar,
- my colleagues,
- the monster under the bed,
- or the basement bugs.







Hello

About Me:

- Storage Developer,
- Long-haul <u>Samba Team</u> Member,
- jCIFS project co-founder,
- Book: <u>Implementing CIFS</u>,
- Lead author of the SMB1 specs: [MS-CIFS], [MS-SMB].

A ruminant mammal (Geekus geekus) With humped shoulders, spindly legs, rollerblades, and broadly palmated antlers.









The Plan

- Introduce the Zambezi Project
- SmartNICs and DPUs
 - ...with an End-to-End Arguments sidebar
- A Quick Discussion of Software Defined Networks/NAS
 - ...and we'll touch on SONiC
- Distributed and Shared Filesystems in a SmartNetwork
 - The Network Really Is the Computer
- Zambezi Revisited



Zambezi

An SMB3 Offload Engine





Zambezi Project Overview

- An SMB2/3 Developer's Toolkit
- Similar to jCIFS, but differn't...
- Originally aimed at:
 - Embedded systems
 - Application development
- Low-level code:
 - SMB2/3 Message [un]packing
 - Transform headers; Encryption, Compression
- A Syntax-Level SMB3 Toolkit
 - ...with a well-defined API.





Zambezi Project Overview

- An SMB2/3 Developer's Toolkit
- Separate Semantic Layer
 - Filesystem Behaviors
 - Locking and Leasing
 - Ownership & Access Controls
 - Integration with the underlying OS and Filesystem
- Temporarily Out of Scope









Zambezi Project Overview

An SMB2/3 Offload Engine

- Encryption, Compression
- Basic message packing/unpacking
- Syntax and Transport error handling
- SMB Connection Negotiation
- Still with a well-defined API

Hey! A SmartNIC! ...already used for protocol offload.









SmartNICs

...and End-to-End Arguments







STORAGE DEVELOPER CONFERENCE

SmartNICs and End-to-End Arguments

"Choosing the proper boundaries between functions is perhaps the primary activity of the computer system designer."

- End-to-End Arguments in System Design, 1981





SmartNICs and End-to-End Arguments

Network offload requires "proper boundaries between functions".

- What are the logical endpoints...
 - For Compression?
 - For Encryption?
 - For the underlying transport?
 - For connection negotiation & maintenance?







SmartNICs and End-to-End Arguments

Function Boundaries:

- For TOE cards, we have Sockets (or WinSock).
- For iSCSI cards, we have the SCSI interface.
- No such thing exists for SMB, or for NFS.

A Zambezi Project Goal: Produce a draft SMB3 Offload API Specification.







SmartNICs with DPUs



Intelligence Within the SmartNICs Data Processing Units: DPUs

- Multi-core RISC processors
- High I/O capacity
- Multiple I/O channels
- Built for programmability and fast message processing

DPUs are typically found on high-end SmartNICs, but may also be included on computer mainboards.



SmartNICs and DPUs

DPU-based SmartNICs:

- Established technology
- Development boards and toolchains
- Multiple vendors
- Ambitious roadmaps



STORAGE DEVELOPER CONFERENCE

Software Defined NAS

...and SONiC (among others)



STORAGE DEVELOPER CONFERENCE

Network Interior Intelligence

Processing Power on the SmartNICs... Processing Power on the Storage (Computational Storage)...

As processing power increases at the edges, what about the middles?

Smart Switches for building Software Defined Networks



STORAGE DEVELOPER CONFERENCE

Network Interior Intelligence



SONIC

- A Linux-based
 - **Network Operating System (NOS)**
- Developed by Microsoft and
 - the Open Compute Project
- Multiple switches × vendors × ASICs
- P4 Switch Abstraction Interface

STORAGE DEVELOPER CONFERENCE



Distributed and Shared Filesystems

The Network Is the Computer





- Ubiquitous Computing Power
- RISC at the Disk
 - Computational Storage
- Switched On ASICs
 - Software Defined Network Storage
 - In-network packet processing
- DPU meets CPU
 - SmartNICs and Mainboard DPUs



SD@EMEA

With all of this computing power, where do we put stuff?

- The network is a distributed computing environment.
- Pay attention to function boundaries.
- Modularize the code.

Beyond Zambezi

- If SMB3, why not NFS?
- If File Protocols, why not Distributed File systems?
 - Lustre, Ceph, etc.
- If all of the above, plus block storage (obviously),
 why not the complete storage subsystem?
- How would a client access the storage?
 - A complete filesystem stack on the SmartNIC?





Zambezi Project

- A Voice in the Wilderness
- Past Experience with jCIFS is Promising
- It's Your Friendly, Neighborhood
 Open Source Project

<u>
 https://gitlab.com/ubiqx/zambezi</u>



References

- End-to-End Arguments in System Design
 - https://web.mit.edu/Saltzer/www/publications/endtoend/endtoend.pdf
 - Wikipedia: <u>End-to-end principle</u>
- SONIC: <u>https://azure.github.io/SONiC/</u>
- What Is a DPU?
 - <u>NVidia Blog Post</u>
 - Fungible <u>DPU Intro Video</u>
 - Wikipedia: <u>Data Processing Unit</u>
- Some sounds sourced from:
 - FreeSound: <u>https://freesound.org/</u>







Please take a moment to rate this session.

Your feedback is important to us.





