

# SMB3 Offload Update

## Zambezi Project

Chris Hertel

Samba Team

SambaXP-2021






# Introductions

Me:

- Long-haul Samba Team Member (22 years)
- jCIFS project co-founder
- Book: Implementing CIFS
- Lead author of the SMB1 Spec.S: [MS-CIFS], [MS-SMB]

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





# Introductions

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.





# Zambezi Project Overview





# ZAMBEZI SMB3 OFFLOAD



- TOE Cards: TCP Offload Engines
- iSCSI NICs: Networked SCSI
- SMB3 Offload: API Undefined



# ZAMBEZI SMB3 OFFLOAD

- SMB2 and SMB3 — *no SMB1*
- Encryption / Compression
- Syntax Layer
  - Message Packing / Parsing
- Host-provided State
  - Authentication and Keys
  - Feature Support
- Fast I/O Processing







# ZAMBEZI SMB3 OFFLOAD

---

The host handles the Semantics:

- Local Filesystem Interface
  - E.g.: POSIX Layer
  - Sync'd Access (Local, NFS, Object...)
- Metadata Management
  - ACLs, Identity, Attributes
- File System Behaviors
  - Locking, Change Notify, etc.





# ZAMBEZI SMB3 OFFLOAD

## The Semantic Layer

...runs above the offload engine.

- The SMBOE API needs to be useful, stable, and **open**.
- SMB2/3 implementations should all be able to use the same API.
- It should be possible to run different SMBOE implementations in parallel.








# ZAMBEZI SMB3 OFFLOAD

## Building a Rational, Well-Documented API

- Stackable for adding new dialects, features, and capabilities.
- Device Driver / Library / Toolkit / All?
- It needs a community.

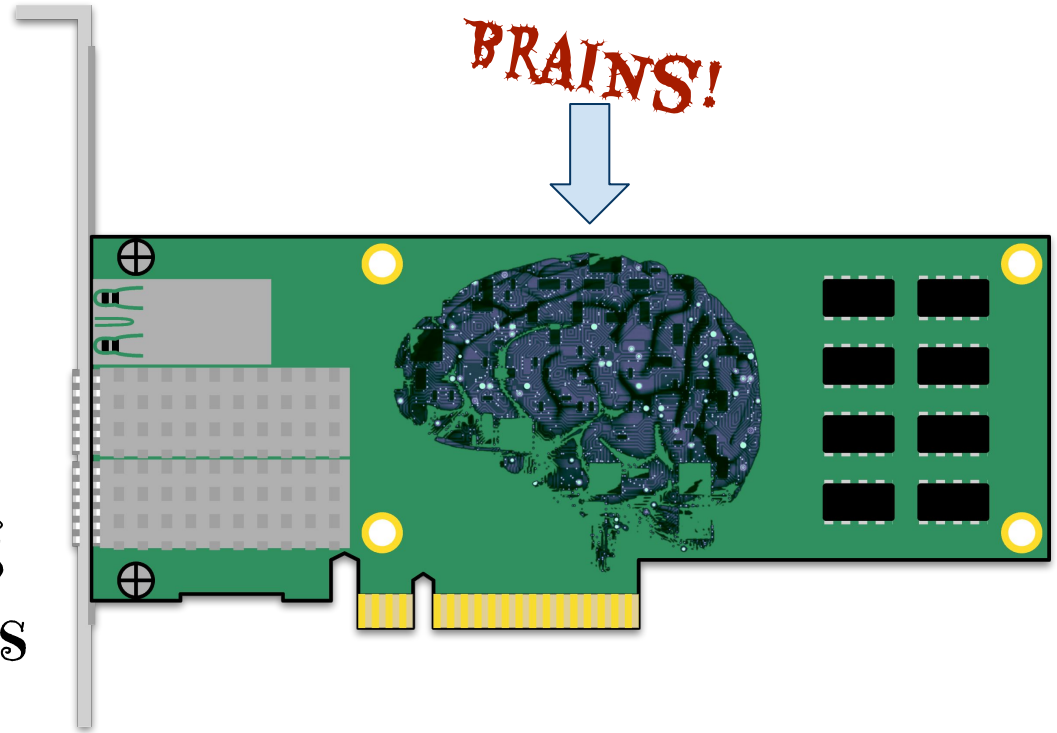
So...

# SmartNICs



# SmartNICs

Network  
Interface  
Cards  
with  
powerful  
processing  
capabilities







# SmartNICs

**Network  
Interface  
Cards**  
→ with  
powerful  
processing  
capabilities

**Programmable:**

- FPGA-based,
- Multi-core ASICs, or
- A Combo of the above





# SmartNICs

## Network Interface Cards

→ with powerful processing capabilities

A generic networking offload device:

- Multi-interface
- Multi-physical layer
- Multi-core processor
- PCIe Interconnect


Ethernet, FC, NVMe-oF, RDMA...





# SmartNICs

Some SmartNICs  
can run Linux.








# SmartNICs & DPUs

DPU: Data Processing Unit

...the **BRAINS** of the SmartNIC.



DPUs can also be placed on motherboards.



# SmartNICs & DPUs

DPUs are typically  
RISC-based:

- ♣ ARM
- ♣ MIPS
- ♣ RISC-V?

DPUs have *lots*  
of I/O capacity.



# Zambezi Update



May 2021





# Zambezi Update

For convenience, I arranged the SMB3 commands into six debatably semi-logical categories:

<b>Managing Connections</b> <ul style="list-style-type: none"><li>• NEGOTIATE (0x0000)</li><li>• SESSION_SETUP (0x0001)</li><li>• LOGOFF (0x0002)</li><li>• ECHO (0x000D)</li></ul>	<b>Share Access</b> <ul style="list-style-type: none"><li>• TREE_CONNECT (0x0003)</li><li>• TREE_DISCONNECT (0x0004)</li></ul>	<b>Open/Close, Lock/Unlock</b> <ul style="list-style-type: none"><li>• CREATE (0x0005)</li><li>• CLOSE (0x0006)</li><li>• LOCK (0x000A)</li><li>• OPLOCK_BREAK (0x0012)</li></ul>
<b>Fundamental I/O</b> <ul style="list-style-type: none"><li>• READ (0x0008)</li><li>• WRITE (0x0009)</li><li>• FLUSH (0x0007)</li></ul>	<b>Metadata Query and Set</b> <ul style="list-style-type: none"><li>• QUERY_DIRECTORY (0x000E)</li><li>• CHANGE_NOTIFY (0x000F)</li><li>• IOCTL (0x000B)</li><li>• QUERY_INFO (0x0010)</li><li>• SET_INFO (0x0011)</li></ul>	<b>Odds and Ends</b> <ul style="list-style-type: none"><li>• CANCEL (0x000C)</li><li>• SMB2 Error Response</li></ul>



# Zambezi Update

---

Several messages have the same format:

```
typedef struct
```

```
{
```

```
    uint16_t StructureSize;
```

```
    uint8_t  Reserved[2];
```

```
} smb2_BaseMsg;
```



- |  |   |
|--|---|
| <ul style="list-style-type: none"><li>• LOGOFF Request/Response</li><li>• TREE_DISCONNECT Request/Response</li><li>• ECHO Request/Response</li></ul> | <ul style="list-style-type: none"><li>• CANCEL Request</li><li>• LOCK Response</li><li>• FLUSH Response</li></ul> |
|--|---|



# Zambezi Update



## Consider Create

- In SMB2/3, Create is a Monster.
  - Lots of flags fields
  - Access & Share Modes
  - Leases/OpLocks
  - Contexts

It's practically a sub-protocol.  
How should it be presented?








# Zambezi Update

---

## Short Term Goals:

- ★ Proxy Dæmon
- ★ Transform Headers
- ★ Unit Tests
- ★ Finish Messages  
(including Create)
- ★ Study other code





# Zambezi Update

---

Where else might this code be useful?


- Software Defined Storage Switches
- Proxy and Cache Servers
- WAN Accelerators
- Hypervisors
- Remote Access Portals





# Zambezi Update

- LGPLv3 & AGPLv3
- Published modules are not well tested, but are fairly complete.
- ...and excessively well documented.




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





# Distractions



In the '67 Doctor Dolittle movie...

The **Giant Luna Moth** is attracted by the light of the moon, and flies there. Once on the moon the Moth is attracted by the light of the earth, and so flies back again.

The cycle repeats.



# Distractions

---

- SmartNIC and DPU Docs and Videos
  - I enjoy learning about all this new stuff
- Embedded Tinkering Platforms (RPi, etc.)
  - [Ultibo](#) for Raspberry Pi:
    - Development toolkit based on FreePascal
    - Built-in support for NTFS



## Block Storage

- AoE is small and simple, so...



# Summery [sic]






# Project Zambezi

---

## Goals:

- ★ **Git 'er done.**
- ★ **Work with the SNIA**
  - Standardize the API
  - Fork a reference implementation under an additional license
- ★ **Partner with others to implement on SmartNICs**
- ★ **Find new & interesting use cases**







[crh@ubiqx.org](mailto:crh@ubiqx.org)



[https://gitlab.com/ubiqx/  
zambezi](https://gitlab.com/ubiqx/zambezi)

# The End



