

# SmartNIC all the Things



Andy Gospodarek  
Netconf 2017 Part 2

**Note1:**

**This is not supposed to be an Advertisement  
for Broadcom's SmartNIC**

**Note2:**

**Deviation from Netconf 2017 Part2 trend:**

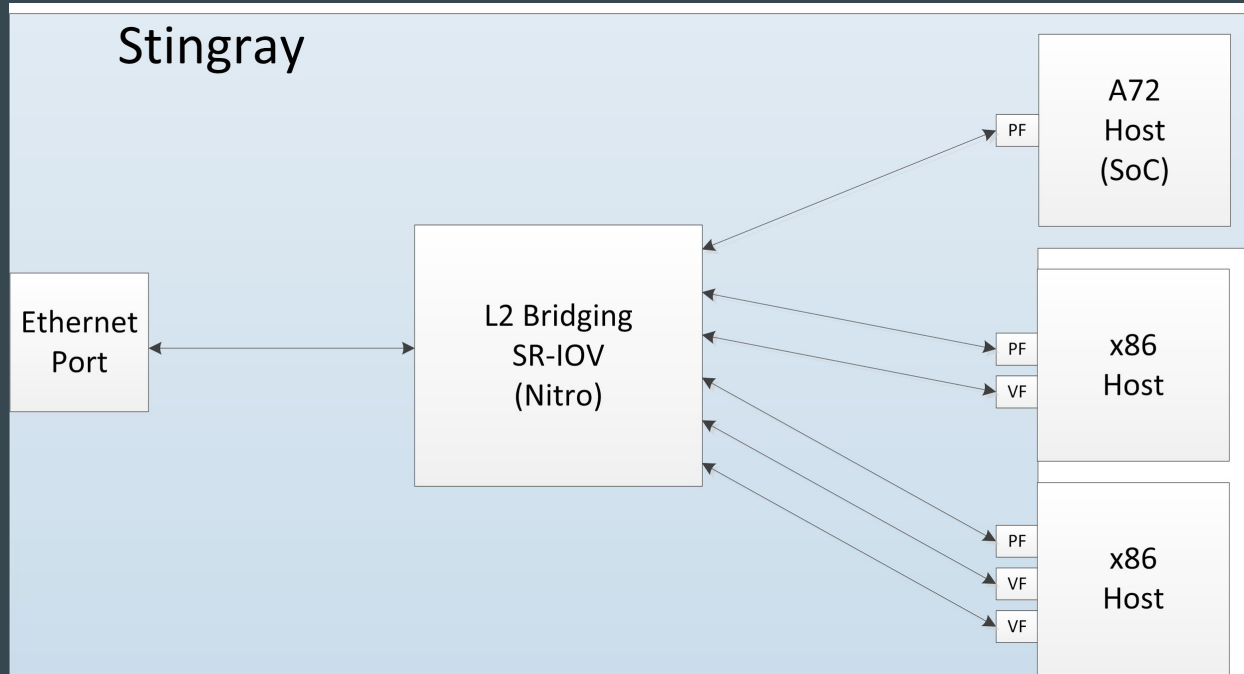
**“You cannot just do this with XDP”**

**But of course, it can run XDP!**

Just another ARMv8 server inside your  
[ARMv8] server

**What does this thing actually look like?**

# Block Diagram Cold boot (+ VFs)

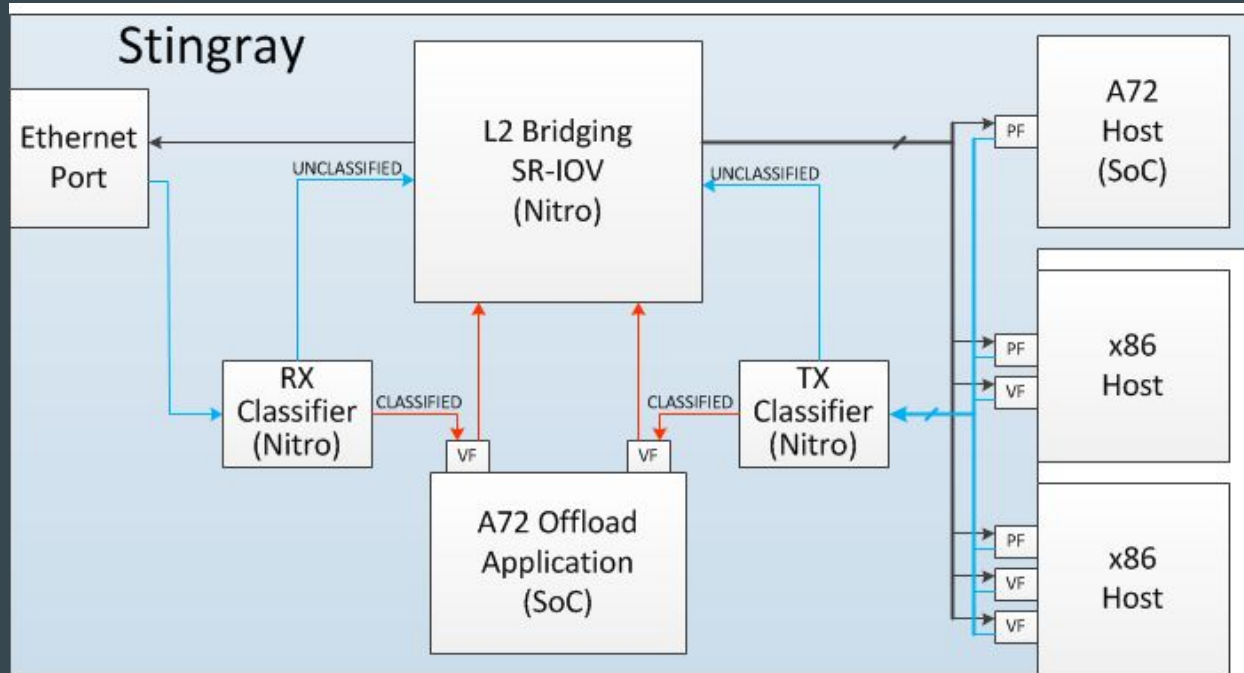


# Cold-boot hardware state

- All PFs have direct access to “Ethernet Port” to get out of the system
- All PFs can communicate with each other via hardware L2 Bridge (much like default PF<->VF communication on SR-IOV NICs)



# Offload All Packet Processing to SmartNIC



# Why would someone want this?

- Free up server cores by offloading processing to the NIC (Standard NIC offload case)
- Network isolation from server
  - Might not be able to trust your baremetal server
  - Network admin != server admin
  - Network policy/forwarding on the server without installing apps on the server OS

# What is the 'A72 Offload Application'

- Standard Kernel forwarding (bridging/routing/etc)
- DPDK
- XDP
- Open vSwitch
- VxLAN termination

Slideware looks *awesome*, but does this  
actually work?

Yes\*

## \*Gospo's To do list

- Transition away from current forwarding setup tool
- Optimize for packet handling via XDP

# Transition away from current forwarding setup tool

- Current workflow:
  - VFs are created on server
  - `bnxt-ctl` created to setup hardware to connect VF to representor and isolate VF from parent PF broadcast domain
  - Great PoC and usable by some customers, but not generic enough to be usable on multiple SmartNICs

# Hardware Forwarding Setup

- Automate creation of netdev representors on A72 based on VF creation on x86 (manual process right now)
- Extension of what is done for VF Representers (Sathya Perla and Michael Chan)
- Use ARM cores to ‘Connect the Wires’ between x86 PFs and VFs in a more abstract/familiar way:
  - switchdev
  - dpipe (maybe)



# Optimize for packet handling via XDP

- Finalize `XDP_REDIRECT` for `bnxt_en`
- Known Wish List:
  - Manipulate TX and RX Buffer Descriptors from XDP program
    - `XDP_TX` probably always going to be higher performance than `XDP_REDIRECT`
    - DPDK can do this already -- not really an original idea
- Unknown Wish List:
  - Desires of [potential] customers

감사합니다