
NFV ISG PoC Proposal – Virtual Function State Migration and Interoperability

A.1 NFV ISG PoC Proposal Template

A.1.1 PoC Team Members

PoC Project Name: Virtual Function State Migration and Interoperability

Network Operators/ Service Providers:

AT&T

Contact: Steven Wright (sw3588@att.com)

British Telecom

Contact: Don Clarke (don.clarke@bt.com)

Manufacturer A: Broadcom Corporation

Contact: Paul Runcy (pruncy@broadcom.com)

Kumaran Siva (kumarans@broadcom.com)

Manufacturer B: Tieto Corporation

Contact: Mieszko Mularczyk

(mieszko.mularczyk@tieto.com)

A.1.2 PoC Project Goals

- PoC Project Goal #1:

The goal of this Proof-of-concept (PoC) is to demonstrate that a network function that has been appropriately developed using modern open-source development tools, operating systems, and hypervisors can be “transparently” deployed on two different host systems using radically different hardware.

In particular, this PoC will demonstrate the same application source code running on different CPU Instruction Set Architectures (ISA) for the main processor, in this case the x86 architecture and the MIPS-64 architecture. Furthermore, the PoC application has been developed in such a way that the state is separate from the application and can be migrated between server systems running the same code on different CPU ISAs.

Portability of VNFs is identified in several of the ISG documents as an objective of NFV. GS NFV 001 (NFV: Use Cases) identifies the need for portability of VNFs in section 9.2 (virtualized IMS use case) and smooth migration of virtual machines between locations in section 13.5 (virtualization of fixed access network functions). GS NFV 002 (NFV: Architectural framework) in section 7.3.2 describes the scope of NFV to include the ability to guarantee an hardware independent lifecycle, performance and portability requirements of the VNF; and in Section 8.2 describes VNF portability as a study topic in the context of the virtualization layer. GS NFV 004 (NFV: Requirements) in section 5.2 notes that the NFV target is to achieve VNF portability across multiple vendors, hypervisors, and hardware.

This PoC activity investigates VNF migration across Instruction Set Architectures to improve the portability of VNFs in the NFV ecosystem. Software can be run on different ISA as long as the underlying operating platform (operating system, hypervisor, data-plane API) is aligned. This vision is achieved in the PoC using standard off-the-shelf open-source platform components. For the PoC these included: KVM for hypervisor, Linux for the guest operating system, and standard data-plane interface like DPDK or Linaro ODP.

Currently, running applications on different ISA still require recompilation for the specific ISA. In the future, innovations currently in development with open source tool chains could enable ISA independent binary formats (in particular LLVM byte code) that will allow binaries to run unchanged.

A.1.3 PoC Demonstration

- Venue for the demonstration of the PoC:

The PoC will be demonstrated by Broadcom at Mobile World Congress 2014, in Barcelona, Spain on February 24-27.

A.1.4 Publication

The abstract of our PoC will be shown at URL: TBD

A.1.5 PoC Project Timeline

- What is the PoC start date? November 1, 2013 (Already underway)
- (First) Demonstration target date February 24, 2014
- PoC Report target date July 31, 2014
- When is the PoC considered completed? July 31, 2014

A.2 NFV PoC Technical Details

A.2.1 PoC Overview

The Server Platform in this PoC will have identical software platform environments running on both x86 and MIPS-64 – including KVM hypervisor and Linux guest OS (same distribution, but compiled for x86 and MIPS-64 respectively).

As shown in Figure 1, the chosen NFV platform components are open source and available on both x86 and MIPS-64.

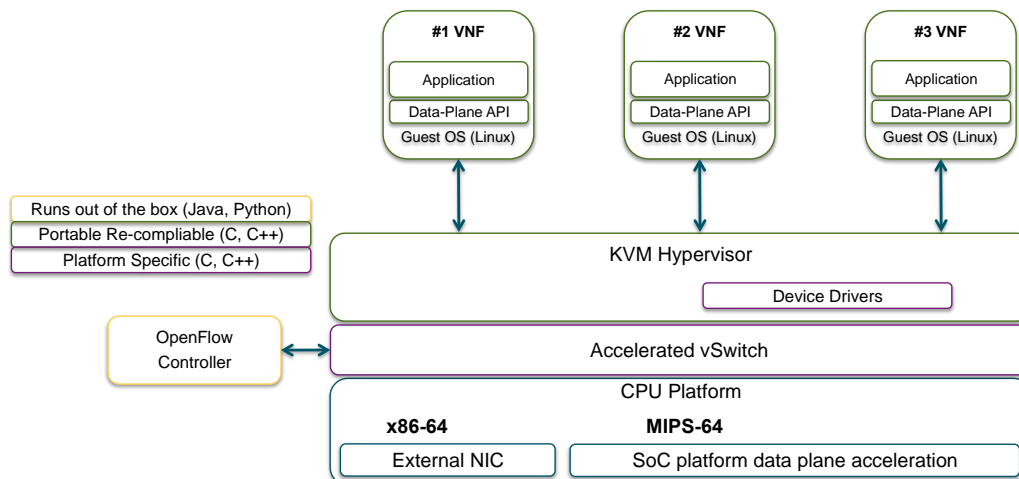


Figure 1 - Server Platform Software Architecture

To enable this PoC, the software platform on both hardware platforms need to be aligned including the development tools (compilers etc.), hypervisor and operating system between the host systems. To be efficient, Independent Software Vendors need to be able to develop network functions using a common tool chain that will compile network functions to different ISAs and deploy on a standard environment.

While not strictly required, open-source and open standards are a natural choice to support

cross-ISA software platforms. By the fact that the source code is made open, silicon and software vendors have the ability to recompile the platform for the different hardware platforms.

For this PoC, the follow open-source components are utilized and cross compiled for both x86 and MIPS-64 Instruction Set Architectures:

1. **Compiler tool chain.** Currently, the best choice is the GNU “C” Compiler or GCC. In the future, Low-Level Virtual Machine compiler (or LLVM) will be the available and provides the benefit of enabling an ISA independent “byte-code”.
2. **Hypervisor.** KVM, Kernel-based Virtual Machine, is an open-source project that leverages heavily from Linux and is optimized to host Linux guest operating systems. XEN is also a possible alternative as a hypervisor.
3. **Operating System.** Linux operating system has gradually become the standard operating system for networking applications. Early on, performance of system calls impaired Linux as a mainstream option for networking applications. But as frameworks, such as DPDK and ODP, enable user-space data-plane applications to be created, the Linux operating system has become a viable platform for networking functions.

The source code for the vEPC VNF will be compiled with identical compiler tool chains for both x86 and MIPS-64. The initial setup will have an IP based video source providing a video stream to the vEPC running on the x86 system. The video will then flow across the vEPC back out to a simulated BTS and finally to a device (User Equipment).

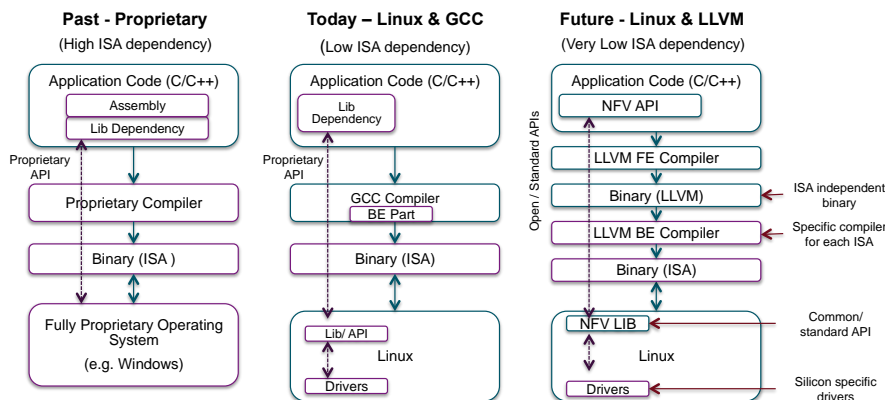


Figure 2 – Tool Chain & Operating System Evolution

In the past tool chains and operating systems were very tightly associated with the hardware platform and in particular the instruction set architecture. Today, the situation has dramatically improved where open-source compilers like GCC are mostly ISA independent. Similarly, open source operating systems like Linux are largely free of ISA dependencies except at for a small (well defined) portion of the Kernel code and hardware specific drivers.

Looking forward, the development tools are evolving towards LLVM, which have a formal separation between front-end compilation which create a non-ISA specific intermediate form binary. This intermediate form (“LLVM byte code”) can be compiled into a hardware specific ISA before the code is executed. The Google Chrome browser utilizes LLVM byte code to enable a multi-platform browser that can run the same LLVM byte code at “near native” speeds on multiple architectures. Currently the Google implementation, known as Portable Native Client is available on ARM, x86, and MIPS.

A.2.2 PoC Scenarios

Scenario 1 - VNF state migration from x86 blade server to Broadcom XLP MIPS-64 based blade server

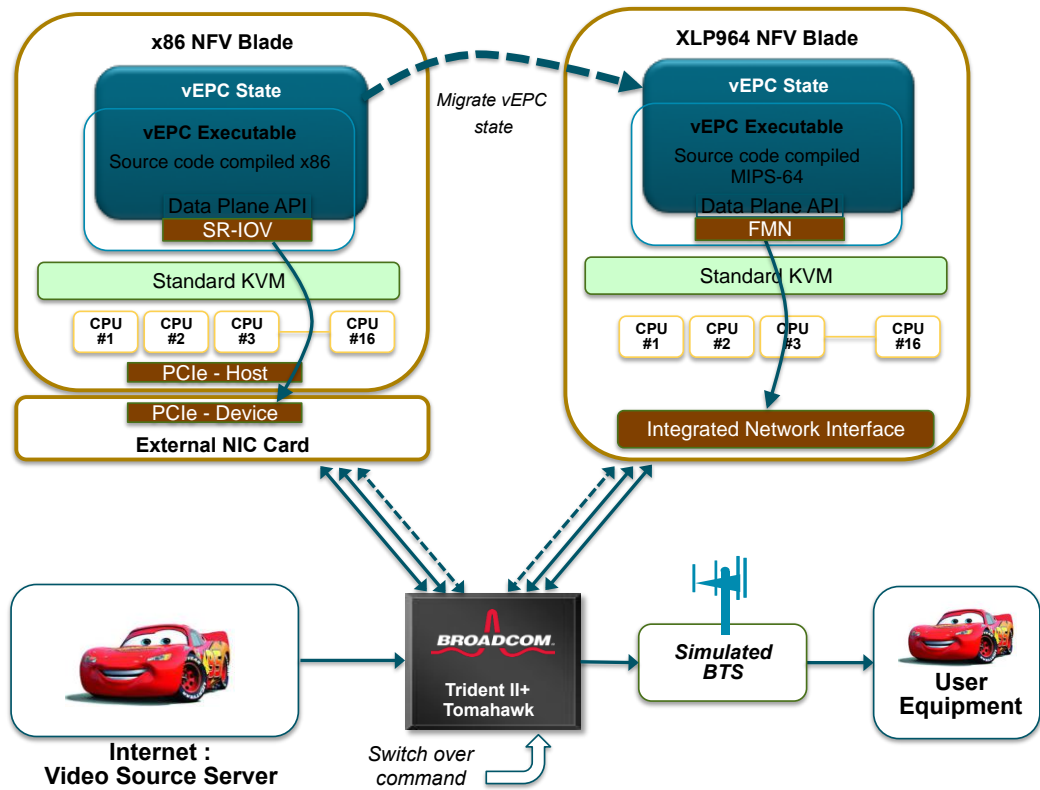


Figure 3 - VNF Migration PoC Detail

1. Show vEPC with video traffic flowing through to the simulated BTS running on x86 system.
2. VNF on the x86 will be placed in a standby state, and the state information will be frozen.
3. VNF state will be moved from the x86 blade to the XLP MIPS-64 blade.
4. VNF will be restarted on the MIPS-64 blade with the transitioned state.
5. Traffic will flow will be rerouted to the vEPC running on the XLP MIPS-64

A.2.3 Mapping to NFV ISG Work

Mapping of our work to documents of NFV ISG is shown below:

	Use Case	Requirement	E2E Arch	Comments
Scenario 1	UC#1	-		<i>This scenario demonstrates IaaS across two different administration domains that are potentially using different hardware.</i>
...				

- 1) (Optional) If this PoC intends to solve or validate any challenge or ongoing work in NFV ISG working groups, complete the table below:

	INF	SWA	MAN	REL	PER	Comments
Scenario 1					x	<i>This scenario extends the definition of portability to encompass applications and environments that are ISA independent.</i>
...						

A.2.4 PoC Success Criteria

The success criteria is straight forward for this PoC:

- VNF should successfully migrate from the x86 blade server to the XLP MIPS-64 based blade server.
- Bearers that are setup on the x86 should be transferred to the XLP MIPS-64 blade server with the delay being the only impact and will be transparent to the user equipment.
- This PoC does not attempt to demonstrate performance of either the x86 or MIPS-64 platforms. Only bearers that have been setup are transitioned over to the XLP platform.

A.2.5 Expected PoC Contribution

List of contributions towards specific NFV ISG Groups expected to result from the PoC Project:

1. PoC Project Contribution #1: VNF Software Development Best Practices to achieve ISA Independence, to NFV Group SWA (NFV-SWA001)
 - This contribution will include recommendations for a fully ISA Independent Open-Source Server Software Architecture including hypervisor, operating system, data-plane interface and cloud management pieces.
 - vEPC will be shown as an example of what can be accomplished using ISA independent software architecture.
2. PoC Project Contribution #2: vEPC active migration use case, to NFV Group INF (NFV-INF008)