

Enable new FPGA cards for CAPI2.0

BSP and SNAP

Workgroup Notes

Revision 1.0_pre3 (August 10, 2019)

** OpenPOWER Foundation Work Group Confidential **

REVIEW DRAFT



www.openpowerfoundation.org

System Software Workgroup <syssw-chair@openpowerfoundation.org>
OpenPower Foundation

Copyright © 2018 OpenPOWER Foundation

<http://www.apache.org/licenses/LICENSE-2.0>

The limited permissions granted above are perpetual and will not be revoked by OpenPOWER or its successors or assigns.

OpenPOWER, the OpenPOWER logo, and openpowerfoundation.org are trademarks or registered trademarks of OpenPOWER Foundation, Inc., registered in many jurisdictions worldwide. Other company, product, and service names may be trademarks or service marks of others.

The purpose of this document is to describe how to enable a new customer card to support CAPI SNAP framework. SNAP is a open-source programming framework for FPGA Accelerations. Its homepage is <https://github.com/open-power/snap>. With it, you can develop accelerators with CAPI technology easily.

This document is a Workgroup Note owned by the System Software Workgroup and handled in compliance with the requirements outlined in the *OpenPOWER Foundation Work Group (WG) Process* document. It was created using the *Master Template Guide* version 1.0.0. Comments, questions, etc. can be submitted to the public mailing list for this document at [<capi-snap-doc@mailinglist.openpowerfoundation.org>](mailto:capi-snap-doc@mailinglist.openpowerfoundation.org).

Table of Contents

Preface
1. Conventions
2. Document change history
1. Introduction
1.1. What is CAPI
1.2. Enable PSL IP on FPGA
1.3. SNAP
2. Enable CAPI2.0 BSP
2.1. Structure
2.2. Step-by-step guidance
3. Enable CAPI2.0 SNAP
3.1. Work on github
3.2. SNAP structure
3.3. Modifications to snap git repositories
3.4. Update capi-utils
3.5. Strategy to enable a new card
3.6. Cleanup and submit
A. OpenPOWER Foundation overview
A.1. Foundation documentation
A.2. Technical resources
A.3. Contact the foundation

vi
vi
vii
1
1
2
3
5
5
6
10
10
10
12
14
14
18
20
20
20
21

List of Figures

- 1.1. CAPI basic concept 1
- 1.2. Develop an accelerator in HDK mode 2
- 1.3. Develop an acceleration on SNAP 3
- 2.1. Project hierarchy for HDK mode 5
- 3.1. Project hierarchy for SNAP 10
- 3.2. Repository structure 11
- 3.3. Vivado Lab Edition 15
- 3.4. Add Configuration Memory Device and Program the flash 15

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

Preface

1. Conventions

The OpenPOWER Foundation documentation uses several typesetting conventions.

Notices

Notices take these forms:



Note

A handy tip or reminder.



Important

Something you must be aware of before proceeding.



Warning

Critical information about the risk of data loss or security issues.

Changes

At certain points in the document lifecycle, knowing what changed in a document is important. In these situations, the following conventions will be used.

- *New text will appear like this.* Text marked in this way is completely new.
- ~~Deleted text will appear like this.~~ Text marked in this way was removed from the previous version and will not appear in the final, published document.
- **Changed text will appear like this.** Text marked in this way appeared in previous versions but has been modified.

Command prompts

In general, examples use commands from the Linux operating system. Many of these are also common with Mac OS, but may differ greatly from the Windows operating system equivalents.

For the Linux-based commands referenced, the following conventions will be followed:

\$ prompt Any user, including the root user, can run commands that are prefixed with the \$ prompt.

prompt The root user must run commands that are prefixed with the # prompt. You can also prefix these commands with the **sudo** command, if available, to run them.

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

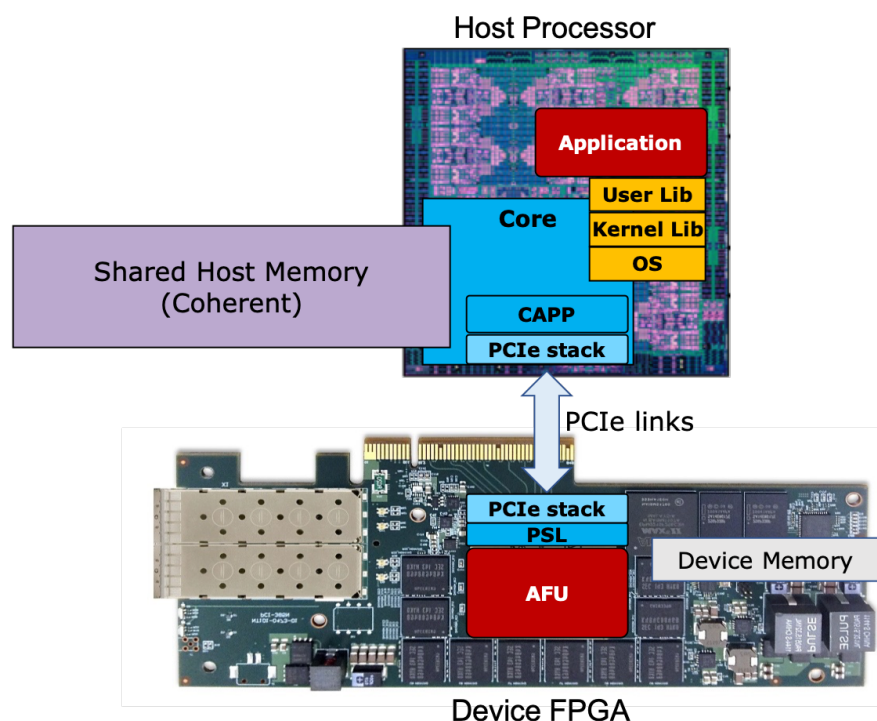
VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

1. Introduction

1.1. What is CAPI

CAPI stands for "Coherent Accelerator Processor Interface" which enables FPGA to access Host memory by virtual address. You can find more introduction about this interface on <https://developer.ibm.com/linuxonpower/capi/>. It is an important feature to develop hardware accelerators in heterogeneous computing. In this document, the "hardware accelerators" are built on FPGA.

Figure 1.1. CAPI basic concept



A complete accelerator has software part (APP, or Applications) running on CPU Processor and the hardware part (AFU, Acceleration Function Unit) running on FPGA chip. APP and AFU are sharing host memory, that means, they both can read and write the 2^{64} range of virtual memory address. To make it happen, CAPI technology has a CAPP (Coherent Acceleration Processor Proxy) logic unit in Processor chip, and also needs a PSL (Processor Service Layer) logic unit in FPGA chip. For CAPI1.0 and CAPI2.0, the interconnection between processor and FPGA is using PCIe physical links and PCIe form factor.

CAPI1.0 uses PCIe Gen3x8.

CAPI2.0 uses PCIe Gen4x8 or Gen3x16.

OpenCAPI is not covered in this document. Visit <https://opencapi.org> for more information.

1.2. Enable PSL IP on FPGA

This document only applies to the cards using Xilinx FPGA chips.

A customer FPGA card needs to have a PSL module (Processor Service Interface) to become a "CAPI-enabled" card. This PSL module is provided by OpenPower Foundation.

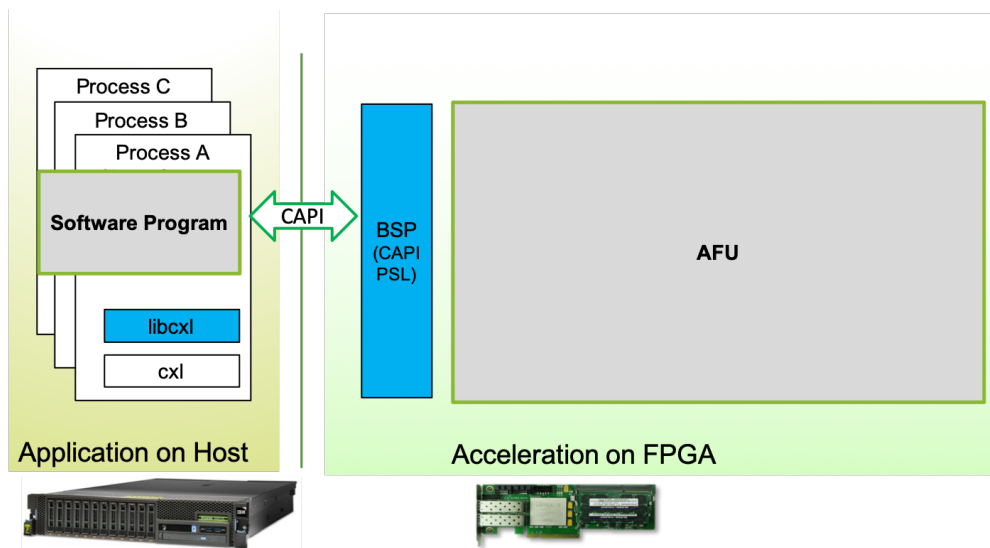
- For CAPI1.0, PSL module and the surrounding board specific modules are provided in the form of a routed dcp file (Xilinx Vivado design checkpoint). It's usually called b_route_design.dcp.
- For CAPI2.0, PSL is an IP package with encrypted source code. It's named like ibm.com_CAPI_PSL9_WRAP_2.00.zip.

They can be downloaded at <https://www.ibm.com/systems/power/openpower>. From the menu, select "CAPI", "Coherent Accelerator Processor Interface (CAPI)" or directly click the "CAPI" icon to go to the CAPI section. Then download the appropriate files depending on your target system being POWER8 (CAPI 1.0) or POWER9 (CAPI 2.0). You need to register an IBM ID to download them.

Users can develop CAPI accelerators in two modes: HDK and SNAP.

HDK is the abbreviation of Hardware Development Kit. As shown in the diagram below, on the FPGA side, the Xilinx Vivado project includes two parts: BSP (Board Supporting Package, containing PSL module) and AFU (Acceleration Function Unit). How to generate BSP will be introduced in Chapter [Enable CAPI2.0 BSP \[5\]](#)

Figure 1.2. Develop an accelerator in HDK mode



AFU is where to implement user-defined functions. The developer working on AFU needs to understand the protocol between AFU and BSP, which is called PSL/AFU interface specification. Please refer to [CAPI1.0 PSL Spec](#) and [CAPI2.0 PSL Spec](#) or search "PSL/AFU interface" in your web browser.

When developing an acceleration, **PSLSE** (PSL Simulation Engine) is also needed to make a software-hardware co-simulation which guarantees the correctness of accelerator design.

When deploying the acceleration to OpenPower servers, it requires user library `libcxl` and kernel module `cx1` to run the application.

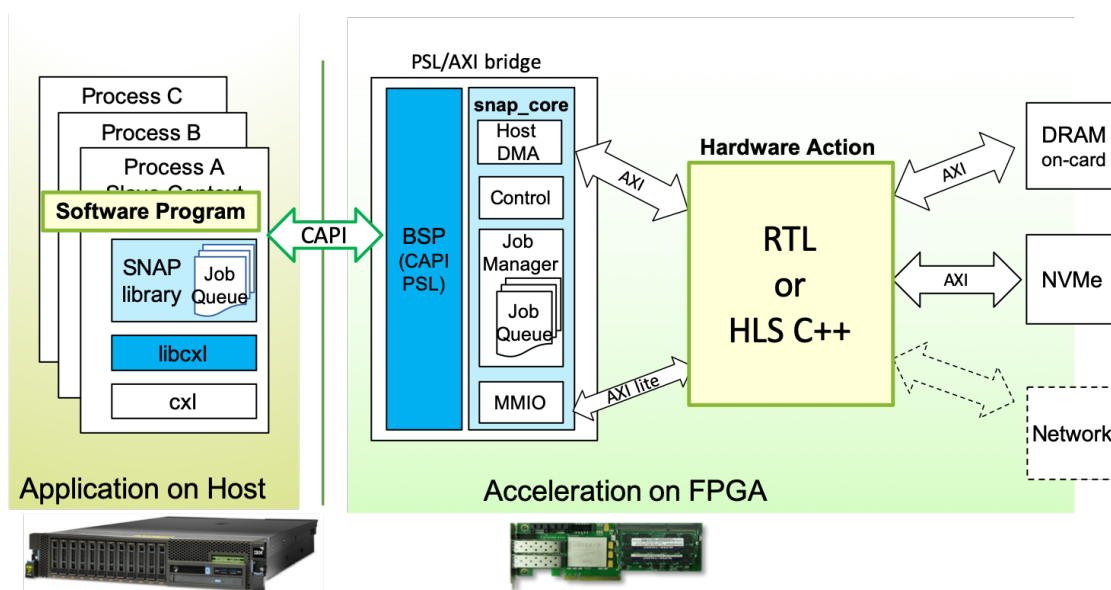
In all, HDK mode will provide the maximum control, utilization of resources and shortest latency. However, SNAP mode simplifies and standardizes the application development significantly and is more recommended.

1.3. SNAP

SNAP is the abbreviation of Storage, Networking and Analytics Programming. It is an open-source acceleration development framework <https://github.com/open-power/snap>. The benefits are:

1. On the FPGA side, SNAP framework adds a bridge to provide AXI interface to developers. So the developer can focus on acceleration function logic design, and doesn't need to study the details of PSL interface specification. AXI is the defacto industry standard for on-chip bus interconnections and is part of [AMBA \(Advanced Microcontroller Bus Architecture\)](#).
2. It also provides DDR SDRAM controller and an optional NVMe controller. The developer can use the card memory or storage directly.
3. SNAP supports using HLS (High Level Synthesis) to develop the acceleration functional unit ("Hardware Action" in yellow box). Developers can write C++ functions and Vivado HLS will compile/convert them to Verilog or VHDL design automatically.
4. A new layer of user library "libsnap" provides more convenient APIs.
5. SNAP is an integrated developing environment that the developer can configure, create Vivado project, run co-simulation or build bitstream with simple commands.
6. Many action examples help new developers to get started.

Figure 1.3. Develop an acceleration on SNAP



Equipping the new FPGA card with SNAP framework needs a few additional steps and is introduced in Chapter [Enable CAPI2.0 SNAP \[10\]](#)



Note

This document focuses on CAPI2.0. For CAPI1.0 enablement, please contact [<capi-snap-doc@mailinglist.openpowerfoundation.org>](mailto:capi-snap-doc@mailinglist.openpowerfoundation.org) for more information.

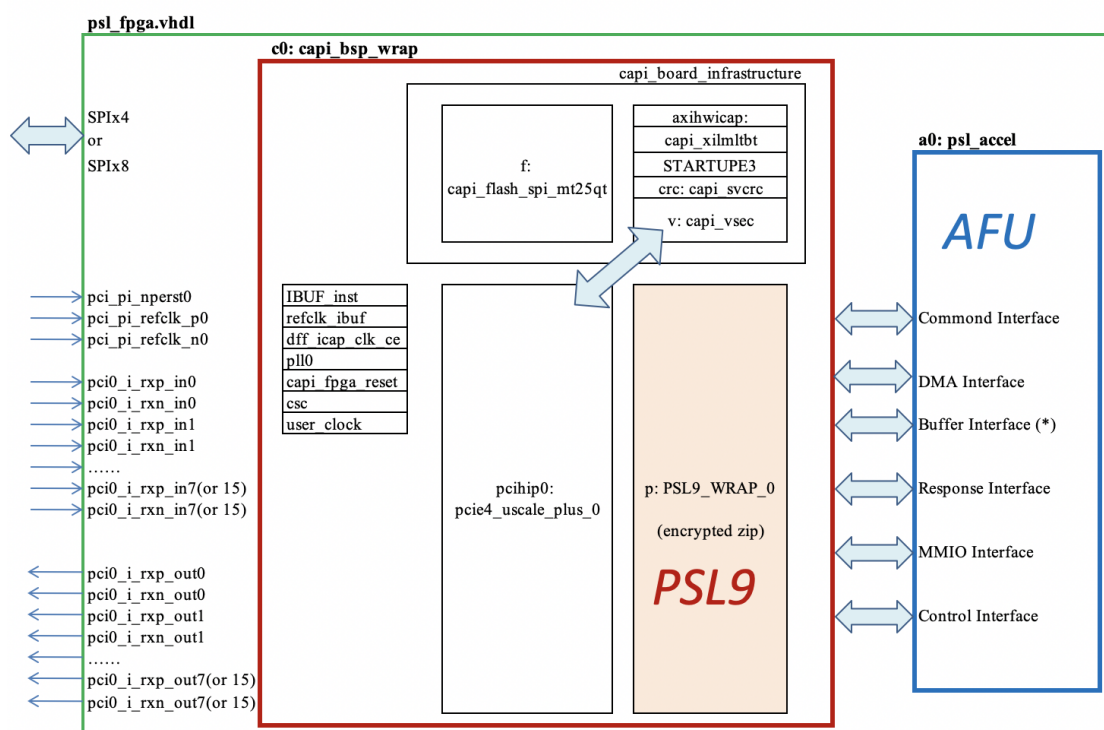
It is assumed the reader knows how to work on Vivado Project and SNAP already. Many materials on how to develop an accelerator with SNAP can be found in "docs" folder of [snap github](#) or other webpages so they are not discussed in this document.

2. Enable CAPI2.0 BSP

2.1. Structure

Each card supplier may design their FPGA board with different FPGA chips, circuit components, memory and IOs, so BSP (Board support package) is different from card to card. That's why an open-sourced project is so helpful: It allows card supplier and every developer to explore the functions of the card freely, and get benefits from CAPI technology.

Figure 2.1. Project hierarchy for HDK mode



CAPI2.0 BSP (capi_bsp_wrap) contains following modules:

- PSL9 (PSL9_WRAP)
- PCIe hard IP core (pcie3/4_ultrascale_plus_0)
- Flash Controller (capi_flash_spi_mt25qt)
- VSEC: Vendor Specific Extended Capability (capi_vsec)
- Xilinx MultiBoot control logic (capi_xilmltbt)
- Other miscellaneous logic

In this diagram, psl_fpga.vhdl is the top design. How to connect the ports to FPGA pins is different from card to card, and the information is provided by Card Supplier. They may include:

- Flash interface (usually SPIx4 or dual SPIx4)

- PCIe interface: perst, refclk, TX and RX data lanes (PCIe Gen3.0x16, or Gen4.0x8)
- Peripheral IPs: I2C, LED, DDR, Ethernet, etc.

At least Flash interface pins and PCIe interface pins are required to be assigned in xdc files precisely.

Between capi_bsp_wrap and user AFU (psl_accel), there are 6 groups of signals: Command, DMA, Buffer, Response, MMIO and Control. Please refer to [CAPI2.0 PSL/AFU interface Spec](#) for the details.

2.2. Step-by-step guidance

2.2.1. Work on github

capi2-bsp is a public Github repository. You need to have a Github account first. Then create a "fork" (Click the "fork" button) on <https://github.com/open-power/capi2-bsp>.

```
git clone https://github.com/[YOUR_USERNAME]/capi2-bsp
```



Note

capi2-bsp is also a submodule of snap.

Keep working on your own capi2-bsp fork, when it has been validated to work well, submit a pull request to "open-power/capi2-bsp" and request merging into the public upstream.

2.2.2. Preparations

First, define a FPGACARD name. It can start from the company's name, following with the card name and be short. For example, ADKU3 = Alpha-Data ADM-PCIE-KU3. Get information from the card supplier.

Table 2.1. Information to collect

Item	Description
FPGACARD	Short card name
FPGACHIP	FPGA part name, for example, xcvu9p-fsgd2104-2L-e
Flash Type and IO pins	Flash chip that attached to FPGA, for example mt28gu01gaax1e-bpi-x16. And the related xdc files for FPGA config.
PCIe Config and IO pins	The tcl/xdc information about the PCIe hardware IP for this card.
DDR MC IP	DDR memory controller Vivado IP tcl/xdc file.
Other peripherals	NVMe IP, Ethernet IP and so on (Optional)



Note

DDR MC and other peripheral IP's configurations are not needed at the beginning. They are not in the capi_bsp_wrap diagram. However, when you query information from the FPGA Supplier, it's wise to include them also.

2.2.3. Download PSL Zip package

Download PSL9 Zip package from [OpenPower Portal](#) and put it in "capi2-bsp/psl" directory.

2.2.4. Copy and modify

Choose an existing card as a base. Copy the folder to your new FPGACARD name. Then modify the contents according to the information collected from FPGA supplier.



Important

Make sure the information in xdc/tcl files are permitted to be open-source.

There are some other modifications you should pay attention to:

1. PCIe core IP creation:

- "Vendor ID" and "Device ID" have to be 0x1014 and 0x0477, so kernel module cxl can recognize the card as a CAPI device. (See in [pci.c](#))
- If the card vendor has a code allocated by PCISIG (See in [PCISIG Member companies](#)), use it as "Subsystem Vendor ID". "Subsystem Device ID" can be chosen freely.
- If the card vendor doesn't have a code allocated by PCISIG, or just for testing and evaluation purpose, please use default "Subsystem Vendor ID" = 0x1014, and send email to aclwg-chair@openpowerfoundation.org to get a distinct "Subsystem Device ID" to differentiate this card from others.

Example: (in create_ip.tcl)

```
create_ip -name pcie4_uscale_plus -vendor xilinx.com -library ip -module_name
pcie4_uscale_plus_0 -dir $ip_dir >> $log_file
set_property -dict [list
    CONFIG.PF0_CLASS_CODE {1200ff}
    CONFIG.PF0_REVISION_ID {02}
    CONFIG.VENDOR_ID {1014}
    CONFIG.PF0_DEVICE_ID {0477}
    CONFIG.PF0_SUBSYSTEM_VENDOR_ID {1014}
    CONFIG.PF0_SUBSYSTEM_ID {0661}
    .....
    .....
] [get_ips pcie4_uscale_plus_0] >> $log_file
```

The corresponding "Subsystem Vendor ID" and "Subsystem Device ID" need to be added into [capi-utils](#), file "psl-devices".

2. Product Family support:

If the FPGA chip types are Xilinx VU33P or VU37P who have HBM, this is actually a new FPGA family **virtexuplushbm**. For a new FPGA Production family, additional steps need to take:

- "capi2-bsp/psl/create_ip.tcl": **set_property supported_families ...**, add new family name like "virtexuplushbm Production"
- "capi2-bsp/common/tcl/create_capi_bsp.tcl": **set_property supported_families ...**, do the same as above.
- Add family support to PSL9 ZIP package: unzip the package, do the modifications, and zip them back again. Commands:

3. VSEC starting address:

"capi2-bsp/[FPGACARD]/src/capi_vsec.vhdl": **vsec_addr[21:32]** defines the address for VSEC. It should be matched with PCIe core value **PF0_SECONDARY_PCIE_CAP_NEXTPTR**. Take card U200 for example, its **vsec_addr[21:32]** starts from 12'h400 (12'b0100 0000 0000), and "tcl/patch_ip.tcl" modifies it from default value 12'h480 to 12'h400."

Xilinx PCIe code information for extended configuration space can be found on [PG156 \(for Ultrascale Device\)](#) or [PG213 \(for Ultrascale+ Device\)](#).

For Ultrascale+ HBM device's pcie4c core, the VSEC starts from 12'hE80. At this time **vsec_addr[21:32]** must be changed in "capi_vsec.vhdl". And the above two lines in "patch_ip.tcl" are not needed anymore.

"capi2-bsp/[FPGACARD]/src/capi_vsec.vhdl": Edit the hardcoded **vpd44data_be** to add VPD (Vital Product Data) information. Ideally this information should be read from an I2C EEPROM. The FPGA supplier wrote the content of EEPROM before shipping. Today the simplest way is taken -- writing some hard coded value. "capi2-bsp/common/script" has a script "gen_vsec.sh" to do this.

```
"capi2-bsp/[FPGACARD]/src/capi_xilmltbt.vhdl": Edit the User image starting address
wbstart addr.
```

"capi_xilmltbt.vhdl" has a Xilinx multi-boot core. That means you can create two kinds of images: Factory image and User image. Factory images will be placed at address 0 of FPGA Flash, and User image will be placed at "User_image_address" on the flash. When power-on or the FPGA card is reset, the multiboot core knows where to load the image. Usually a Golden factory image is put on address 0 and is never changed. Multiboot core always tries to load user image first. If the user image has something wrong, multiboot logic will tell the FPGA to "fallback" to factory image. You still see the card in the system and you can just program a new user image to try again.

OpenPOWER Foundation
Work Group Confidential

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

3.1. Work on github

3.1. Work on github

```
git clone https://github.com/[YOUR USERNAME]/snap
```



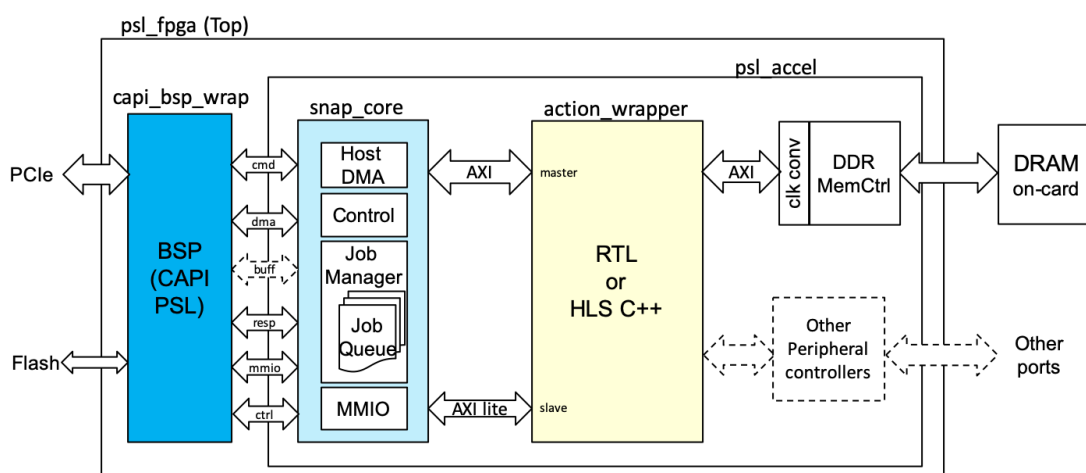
Note

```
git submodule init
git submodule update
```

Anyway, make sure that "hardware/capi2-bsp" is the one just generated in last chapter.

3.2. SNAP structure

Figure 3.1. Project hierarchy for SNAP

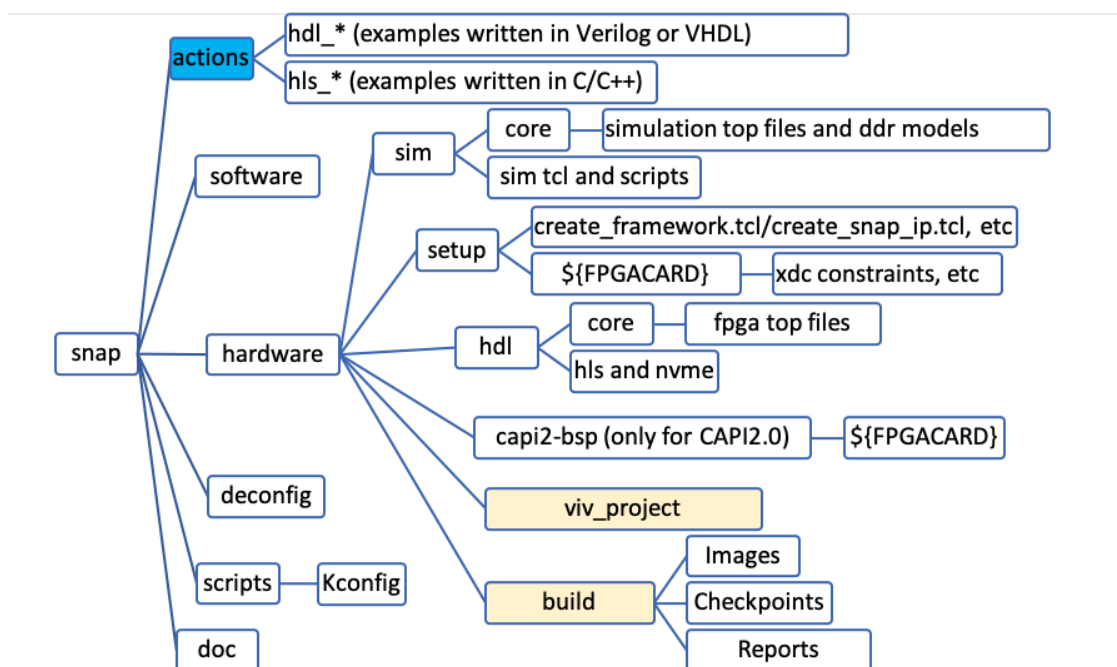


Note

Module **snap_core** on CAPI2.0 implemented the data path with DMA interface. Buffer interface is not used.

The following picture shows the SNAP github repository folders and files.

Figure 3.2. Repository structure



All of the user-developed accelerators are in "**actions**" directory. There are already some examples there. Each "action" has its "sw", "hw", "tests", and other sub-directories. The hardware part uses "action_wrapper" as its top.

"**software**" directory includes libsnap, header files and some tools. "**hardware**" directory is the main focus. "**deconfig**" has the config files for silent testing purpose, and "**scripts**" has the menu settings and other scripts.

How does SNAP work and what are the files used in each step?

- **make snap_config**: The menu to select cards and other options is controlled by "script/Kconfig"
- **make model**: This step creates a Vivado project. It firstly calls "hardware/setup/create_snap_ip.tcl" to generate the IP files in use, then calls "hardware/setup/create_framework.tcl" to build the project. About "create_framework.tcl":
 - It adds BSP (board support package). In CAPI1.0, it is also called PSL Checkpoint file (b_route_design.dcp) or base_image. It uses the path pointed to b_route_design.dcp and adds it into the design. In CAPI2.0, it will call the make process in capi2-bsp submodule to generate "capi_bsp_wrap" if it doesn't exist. This step is skipped if "capi_bsp_wrap" is already generated. Then "create_framework.tcl" adds the capi_bsp_wrap (xcix or xci file) into the design.
 - It adds FPGA top files and snap_core files (in hardware/hdl/core).
 - It adds constrain files: in "hardware/setup/[FPGACARD]" or in "hardware/capi2-bsp/[FPGACARD]"

- It adds user files (in "actions/[ACTION_NAME]/hw"). User's action hardware uses top file named "action_wrapper.vhd"
- It adds simulation files (in "hardware/sim/core") including simulation top files and simulation models. (If **no_sim** is selected in snap_config menu, this step is skipped.)

After above steps, "**hardware/viv_project**" is created. Open it with Vivado GUI, and check the design hierarchy. And it will call the selected simulator to compile the simulation model.

- **make image:** This step runs synthesis, implementation and bitstream generation. It calls "hardware/setup/snap_build.tcl" and also uses some related tcl scripts to work together. In this step, "**hardware/build**" will be created and the output products like bit images, checkpoints (middle products for debugging) and reports (reports of timing, clock, IO, utilization, etc.) If everything runs well and timing passes, user will get the bitstream files (in "build/Images" sub directory) to program the FPGA card.

3.3. Modifications to snap git repositories

For a new FPGA card, the detailed items to update are listed as below.

- Hardware RTL, setup, simulation
- Software and tools
- Testing
- Publishing

The best way is to grep some keywords like "S241" or "AD8K5" under the directories and look for the locations that need modifications.



Note

A file ending with "_source", like "psl_fpga.vhd_source", means this file will be pre-processed to generate the output file without "_source" suffix, like "psl_fpga.vhd". There are **#ifdef** macros or comments like **-- only for NVME_USED=TRUE**. They help to create a target VHDL/Verilog file with different configurations.

Below lists the files to change:

- snap_config and environmental files
- Hardware: psl_accel and psl_fpga (top) RTL files
- Hardware: tcl files for the workflow
- Hardware: xdc files for IO, floorplan, clock and bitstream settings
- Hardware: create DDR Memory controller IP (mig) in create_snap_ip.tcl, create DDR memory sim model, and other xdc files
- Hardware: create_ip, sim model and xdc files for other IPs
- Software: New card type, register definition

- VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

- VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -



VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

powered on, after Vivado Lab sees the FPGA, right click the device, "program device..." and select the bit file **immediately**. This action only takes about 10 seconds and can be done before skiboot on the server starts to scan PCIe devices.

Be aware of the fact that now only FPGA chip is programmed, (the flash memory is still empty or holding old data), so when the server is powered off or reboot the recent programming to FPGA chip will be lost.

Figure 3.3. Vivado Lab Edition

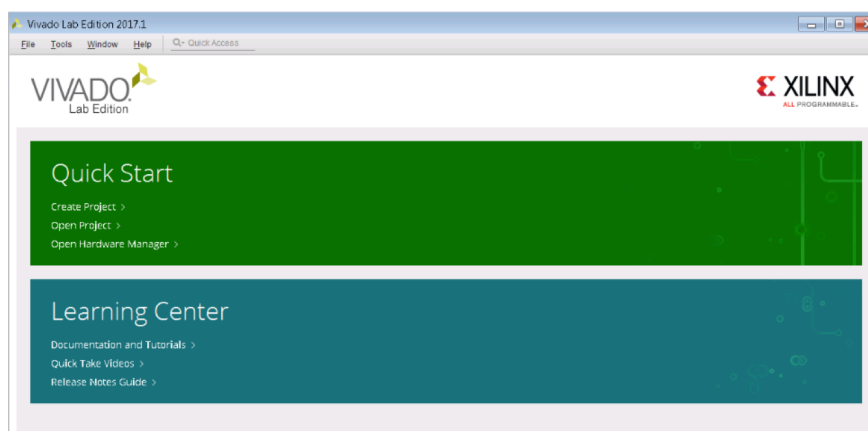
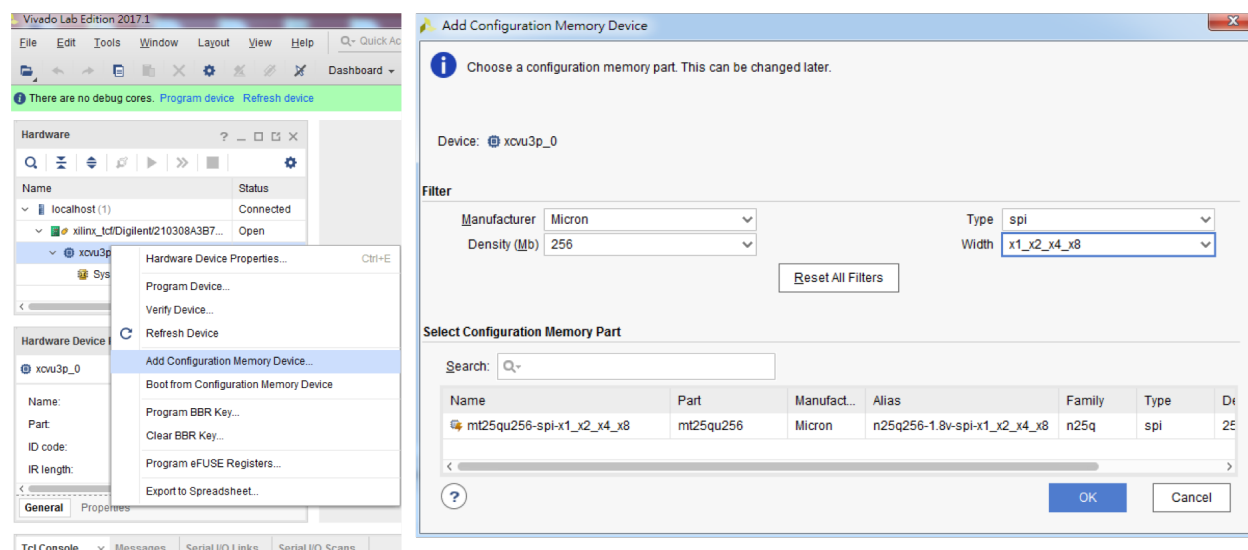


Figure 3.4. Add Configuration Memory Device and Program the flash



Important

When installing **Vivado Lab**, please choose as same version as the Vivado tool which was used to build images.

Tips to debugging:

1. Seeing 0477 by **lspci** is the most important milestone. If not, please check file **"/sys/firmware/opal/msglog"** to see whether there are link training failed messages. A successful message

```
[ 63.403485191,5] PHB#0000:00:00.0 [ROOT] 1014 04c1 R:00 C:060400 B:01..01 SLOT=CPU1 Slot2 (16x)
[ 63.403572553,5] PHB#0000:01:00.0 [EP ] 1014 0477 R:02 C:1200ff (
device) LOC_CODE=CPU1 Slot2 (16x)
```

- ```
[9.301403] cxl-pci 0000:01:00.0: Device uses a PSL9
[9.301523] cxl-pci 0000:01:00.0: enabling device (0140 -> 0142)
[9.303327] cxl-pci 0000:01:00.0: PCI host bridge to bus 0006:00
[9.306749] cxl afu0.0: Activating AFU directed mode
```

- ```
modinfo cxl
```

- ```
ls /dev/cxl
afu0.0m afu0.0s
```

```
sudo lspci -s 0000:01:00.0 -vvv
```

```
0000:01:00.0 Processing accelerators: IBM Device 0477 (rev 02) (prog-if ff)
Subsystem: IBM Device 0660
```

LnkSta: Speed 8GT/s, Width x16, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt-

```
Capabilities: [b0] Vital Product Data
Product Name: U200 PCIe CAPI2 Adapter
Read-only fields:
 [PN] Part number: Xilinx.U200
 [V1] Vendor specific: 0000000000000000
 [V2] Vendor specific: 0000000000000000
 [V3] Vendor specific: 0000000000000000
 [V4] Vendor specific: 0000000000000000
 [RV] Reserved: checksum good, 3 byte(s) reserved
End
```

```
Capabilities: [400 v1] Vendor Specific Information: ID=1280 Rev=0 Len=080 <?>
Kernel driver in use: cxl-pci
Kernel modules: cxl
```

- 
- Workgroup Notes  
Non-standard Track

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -

- VIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT - REVIEW DRAFT -



- Flash Data Port

The details are described in [Coherent Accelerator Interface Architecture](#), Chapter 12.3, "CAIA Vendor-Specific Extended Capability Structure". So the C file in [capi-utils](#) reads FPGA bitstream "bin" file, and writes the data to VSEC "Flash Data Port" register. So the bytes are sent through PCIe, to Flash controller and finally arrive to flash memory on the card.

### 3.5.3. Stage 3: Verify DDR interface

1. Select another action example (hdl\_example with DDR) or hls\_memcpy.
2. **make model** and **make sim**. Make sure the DDR simulation model works well.
3. **make image** to generate the bitstream files.
4. Use capi-utils to program the bitstream "bin" file to the card.
5. Run the application to see whether it works.

Basically SNAP only implemented one DDR Bank (or channel) while most cards have two to four banks. (N250S+ is one of the rare card which has only one DDR bank). To implement more DDR channels, depending on user's needs, there are two options: the first is to just extend the size of the first bank by adding this second bank on the same DDR memory controller. The other option is to use two (or more) memory controllers in parallel to have a higher throughput. This later option means duplicating the DDR memory controller in place and this will take twice the place in the design. In this case, the `action_wrapper` also requires change to add the additional DDR ports. For HLS design, another HLS DDR port should be added into `"actions/[YOUR_ACTION]/hw/XXX.CPP"`. As for an opensource project, everyone is welcomed to add your contribution by implementing it and add it to the SNAP design.

### 3.5.4. Stage 4: Verify Other IO interface

This step is decided by the card's capabilities and the specific IOs that the card can provide. Like the second or more DDR channels, user can add their code for other IO interface freely.

### 3.5.5. Stage 5: Performance Validation

Check the result of "snap/actions/hls\_memcopy/tests/test\*\_throughput.sh" for bandwidth and "snap/actions/hls\_latency\_eval/test/test\*.sh" for latency.

### 3.5.6. Stage 6: Pressure Test

Prepare bitstream files for basic tests, throughput tests, latency tests, max-power tests. Adding image flashing tests, card reset tests and others. Run them intensively.

### 3.6. Cleanup and submit

Now a new FPGA card has been enabled to CAPI2.0 SNAP. Cleanup your workspace, check files and submit your work!

- Submit the pull request of your "capi2-bsp fork" before "snap fork". Assign the reviewer and wait capi2-bsp to be merged into <https://github.com/open-power/capi2-bsp> master branch

- Update the submodule pointer to the latest "open-power/capi2-bsp" master and then submit the pull request of your forked snap.
- Capi-utils is independent. Just create a pull request and assign a reviewer. It can only be merged into master branch after having been reviewed.

The OpenPOWER Foundation was founded in 2013 as an open technical membership organization that will enable data centers to rethink their approach to technology. Member companies are enabled to customize POWER CPU processors and system platforms for optimization and innovation for their business needs. These innovations include custom systems for large or warehouse scale data centers, workload acceleration through GPU, FPGA or advanced I/O, platform optimization for SW appliances, or advanced hardware technology exploitation. OpenPOWER members are actively pursuing all of these innovations and more and welcome all parties to join in moving the state of the art of OpenPOWER systems design forward.

## A.1. Foundation documentation

- [\*Bylaws of OpenPOWER Foundation\*](#)
- [\*OpenPOWER Foundation Intellectual Property Rights \(IPR\) Policy\*](#)
- [\*OpenPOWER Foundation Membership Agreement\*](#)
- [\*OpenPOWER Anti-Trust Guidelines\*](#)

## A.2. Technical resources

- Foundation work groups
- Remote development environments (VMs)
- Development systems
- Technical specifications
- Software
- Developer tools

Workgroup Notes  
Non-standard Track

## A.3. Contact the foundation

To learn more about the OpenPOWER Foundation, please use the following contact points:

- General information -- <info@openpowerfoundation.org>
- Membership -- <membership@openpowerfoundation.org>
- Technical Work Groups and projects -- <tsc-chair@openpowerfoundation.org>
- Events and other activities -- <admin@openpowerfoundation.org>
- Press/Analysts -- <press@openpowerfoundation.org>

More contact information can be found at [openpowerfoundation.org/get-involved/contact-us](https://openpowerfoundation.org/get-involved/contact-us).