

Vivado Design Suite User Guide

Dynamic Function eXchange

UG909 (v2019.2) January 15, 2020

Revision History

The following table shows the revision history for this document.

Section	Revision Summary
01/15/2020 Version 2019.2	
General	Updated terminology to reflect new Dynamic Function eXchange (DFX) solution name where applicable.
Configuration Time	Added Maximum Bandwidths for Configuration Ports tables for UltraScale and UltraScale+ device families.
06/05/2019 Version 2019.1	
UltraRAM Behavior	Updated information for UltraRAM memory.
Tandem Configuration and Dynamic Function eXchange	New section on ICAP/MCAP ports.

Table of Contents

Revision History	2
Chapter 1: Introduction	
Overview	6
Introduction to Dynamic Function eXchange.....	7
Terminology	8
Design Considerations	11
Dynamic Function eXchange Licensing.....	16
Chapter 2: Common Applications	
Overview	17
Networked Multiport Interface	17
Configuration by Means of Standard Bus Interface.....	19
Dynamically Reconfigurable Packet Processor	21
Asymmetric Key Encryption	22
Summary.....	23
Chapter 3: Vivado Software Flow	
Overview	24
Dynamic Function eXchange Commands	25
Dynamic Function eXchange Constraints and Properties	32
Apply Reset After Reconfiguration.....	41
Software Flow	44
Tcl Scripts	49
Chapter 4: Vivado Project Flow	
Overview	50
Flow Summary	50
Steps for Creating and Using a Dynamic Function eXchange Project.....	51
Supported/Unsupported Features	67
Chapter 5: Design Considerations and Guidelines for All Xilinx Devices	
Overview	68

Partial Reconfiguration IP	68
Design Hierarchy	69
Partition Pin Placement	73
Active-Low Resets and Clock Enables	73
Decoupling Functionality	74
Black Boxes	75
Effective Approaches for Implementation	76
Configuration Analysis Report	78
Managing Constraints for a DFX Design	80
Defining Reconfigurable Partition Boundaries	83
Avoiding Deadlock	83
Design Revision Checks	84
Simulation and Verification	85
Chapter 6: Design Considerations and Guidelines for 7 Series and Zynq Devices	
Overview	86
Design Elements Inside Reconfigurable Modules	86
Global Clocking Rules	87
Creating Pblocks for 7 Series Devices	89
Using High Speed Transceivers	98
Dynamic Function eXchange Design Checklist (7 Series)	98
Chapter 7: Design Considerations and Guidelines for UltraScale and UltraScale+ Devices	
Overview	102
Design Elements Inside Reconfigurable Modules	102
Creating Pblocks for UltraScale and UltraScale+ Devices	103
Global Clocking Rules	111
I/O Rules	112
Using High Speed Transceivers	114
Dynamic Function eXchange Checklist for UltraScale and UltraScale+ Device Designs	114
Chapter 8: Configuring the Device	
Overview	120
Configuration Modes	120
Bitstream Type Definitions	122
Dynamic Function eXchange through ICAP for Zynq Devices	127
Tandem Configuration and Dynamic Function eXchange	127
Formatting BIN Files for Delivery to Internal Configuration Ports	131
Summary of BIT Files for UltraScale Devices	132

System Design for Configuring an FPGA	134
Partial BIT File Integrity	135
Configuration Frames	137
Configuration Time	137
Configuration Debugging.	139
Using Vivado Debug Cores.	141
Chapter 9: Known Issues and Limitations	
Known Issues	145
Known Limitations	146
Appendix A: Hierarchical Design Flows	
Overview	148
Appendix B: Additional Resources and Legal Notices	
Xilinx Resources	149
Solution Centers	149
Documentation Navigator and Design Hubs	149
References	150
Training Resources	151
Please Read: Important Legal Notices	151

Introduction

Overview

Dynamic Function eXchange (DFX) allows for the reconfiguration of modules within an active design. This flow requires the implementation of multiple configurations, which ultimately results in full bitstreams for each configuration and partial bitstreams for each Reconfigurable Module. The number of configurations required varies by the number of modules that need to be implemented. However, all configurations use the same top-level, or static, placement and routing results. These static results are exported from the initial configuration and imported by all subsequent configurations using checkpoints.

DFX is a comprehensive solution that is comprised of many parts. These elements include the Xilinx[®] silicon ability to be dynamically reconfigured, the Vivado[®] software flow for compiling designs from RTL to bitstream, and the complementary features such as IP. In this release, you will see a mix of DFX and Partial Reconfiguration (PR) terminology, with DFX representing the overall solution and PR representing a component technology piece of that solution.

Complementary documentation, such as application notes, white papers, and videos, will not be recaptured with DFX terminology, but all new documentation will show the DFX terms.

The content of this guide includes the following:

- Description of Dynamic Function eXchange as implemented in the Vivado Design Suite[®]
- Assumption of familiarity with FPGA design software, particularly Vivado Design Suite
- Updates specific to the Vivado Design Suite Release 2019.2. This release supports Partial Reconfiguration for the products listed below:
 - 7 Series Devices
 - Nearly all Virtex[®]-7, Kintex[®]-7, Artix[®]-7, and Zynq[®]-7000 SoC devices
 - Note:** Spartan[®]-7 devices, as well as Artix-7 7A12T and 7A25T, are not supported.
 - UltraScale[™] Devices
 - Place and route, as well as bitstream generation is enabled for all production devices.

Note: Access to the VU440 is not restricted in this release. However, expect memory usage to be higher for this device than all others (potentially exceeding 64 MB).

- Bitstream generation is disabled by default for ES2 devices, but place and route can still be performed.
- UltraScale+™ Devices
 - Place and route, as well as bitstream generation, is enabled for all production devices, including all Zynq UltraScale+ RFSoc devices.
 - Place and route is enabled for all Virtex UltraScale+ 58G PAM4 devices, but bitstreams are gated by a parameter as device support is still beta.
 - Place and route is enabled for many engineering silicon (ES1, ES2) versions of UltraScale+ devices. Bitstream generation is disabled by default for these devices.



VIDEO: For an overview of the Vivado Partial Reconfiguration solution in 7 series devices, see the [Vivado Design Suite QuickTake Video: Partial Reconfiguration in Vivado](#).

For an overview of the Vivado Partial Reconfiguration solution in UltraScale devices, see the [Vivado Design Suite QuickTake Video: Partial Reconfiguration for UltraScale](#).

Introduction to Dynamic Function eXchange

FPGA technology provides the flexibility of on-site programming and re-programming without going through re-fabrication with a modified design. Dynamic Function eXchange (DFX) takes this flexibility one step further, allowing the modification of an operating FPGA design by loading a dynamic configuration file, usually a partial BIT file. After a full BIT file configures the FPGA, partial BIT files can be downloaded to modify reconfigurable regions in the FPGA without compromising the integrity of the applications running on those parts of the device that are not being reconfigured.

Figure 1-1 illustrates the premise behind Dynamic Function eXchange.




Figure 1-1: Basic Premise of Dynamic Function eXchange

As shown, the function implemented in Reconfig Block A is modified by downloading one of several partial BIT files, `A1.bit`, `A2.bit`, `A3.bit`, or `A4.bit`. The logic in the FPGA design is divided into two different types, reconfigurable logic and static logic. The gray area of the FPGA block represents static logic and the block portion labeled Reconfig Block "A" represents reconfigurable logic. The static logic remains functioning and is unaffected by the loading of a partial BIT file. The reconfigurable logic is replaced by the contents of the partial BIT file.

There are many reasons why the ability to time multiplex hardware dynamically on a single FPGA is advantageous. These include:

- Reducing the size of the FPGA required to implement a given function, with consequent reductions in cost and power consumption
- Providing flexibility in the choices of algorithms or protocols available to an application
- Enabling new techniques in design security
- Improving FPGA fault tolerance
- Accelerating configurable computing
- Delivering updates (fixes and new features) to deployed systems

In addition to reducing size, weight, power and cost, Dynamic Function eXchange enables new types of FPGA designs that would be otherwise impossible to implement.

Terminology

The following terminology is specific to the Dynamic Function eXchange feature and is used throughout this document.

Bottom-Up Synthesis

Bottom-Up Synthesis is synthesis of the design by modules, whether in one project or multiple projects. In Vivado, bottom-up synthesis is referred to as out-of-context (OOC) synthesis. OOC synthesis generates a separate netlist (or DCP) per OOC module, and is required for Dynamic Function eXchange to ensure no optimization occurs across the module boundary. In OOC synthesis, the top-level (or static) logic is synthesized with `black_box` module definitions for each OOC module.

Configuration

A configuration is a complete design that has one Reconfigurable Module for each Reconfigurable Partition. There might be many configurations in a Dynamic Function eXchange FPGA project. Each configuration generates one full BIT file as well as one partial BIT file for each Reconfigurable Module (RM).

Configuration Frame

Configuration frames are the smallest addressable segments of the FPGA configuration memory space. Reconfigurable frames are built from discrete numbers of these lowest-level elements. In Xilinx devices, the base reconfigurable frames are one element (CLB, block RAM, DSP) wide by one clock region high. The number of resources in these frames vary by device family.

Internal Configuration Access Port (ICAP)

The internal configuration access port (ICAP) is essentially an internal version of the SelectMAP interface. For more information, see the *7 Series FPGAs Configuration User Guide* (UG470) [Ref 9] or the *UltraScale Architecture Configuration User Guide* (UG570) [Ref 10].

Media Configuration Access Port (MCAP)

The MCAP is dedicated link to the configuration engine from one specific PCIe® block per UltraScale device. This entry point can be enabled when configuring the Xilinx PCIe IP.

Partition

A Partition is a logical section of the design, user-defined at a hierarchical boundary, to be considered for design reuse. A Partition is either implemented as new or preserved from a previous implementation. A Partition that is preserved maintains not only identical functionality but also identical implementation.

Partition Definition (PD)

This is a term used within the project flow only. A Partition Definition defines a set of Reconfigurable Modules that are associated with the module instance (or Reconfigurable Partition). A PD is applied to all instances of the module, and cannot be associated with a subset of module instances.

Partition Pin

Partition pins are the logical and physical connection between static logic and reconfigurable logic. The tools automatically create, place, and manage Partition Pins.

Partial Reconfiguration (PR)

Partial Reconfiguration is the Xilinx silicon technology that enables users to modify a subset of logic in an operating FPGA design by downloading a partial bitstream.

Processor Configuration Access Port (PCAP)

The processor configuration access port (PCAP) is similar to the internal configuration access port (ICAP) and is the primary port used for configuring a Zynq-7000 SoC device. For more information, see the *Zynq-7000 All Programmable SoC Technical Reference Manual* (UG585) [Ref 11].

Programmable Unit (PU)

In the UltraScale architecture, this is the minimum required resources for reconfiguration. The size of a PU varies by resource type. Because adjacent sites share a routing resource (or Interconnect tile) in the UltraScale architecture, a PU is defined in terms of pairs.

Reconfigurable Frame

Reconfigurable frames (in all references other than “configuration frames” in this guide) represent the smallest reconfigurable region within an FPGA. Bitstream sizes of reconfigurable frames vary depending on the types of logic contained within the frame.

Reconfigurable Logic

Reconfigurable logic is any logical element that is part of a reconfigurable module. These logical elements are modified when a partial BIT file is loaded. Many types of logical components can be reconfigured such as LUTs, flip-flops, block RAM, and DSP blocks.

Reconfigurable Module (RM)

A Reconfigurable Module (RM) is the netlist or HDL description that is implemented within a reconfigurable partition. Multiple RMs exist for a reconfigurable partition.

Reconfigurable Partition (RP)

Reconfigurable Partition (RP) is an attribute set on an instantiation that defines the instance as reconfigurable. The Reconfigurable Partition is the level of hierarchy within which different Reconfigurable Modules are implemented. Tcl commands such as `opt_design`, `place_design` and `route_design` detect the `HD.RECONFIGURABLE` property on the instance and process it correctly.

Static Logic

Static logic is any logical element that is not part of an RP. The logical element is never partially reconfigured and is always active when RPs are being reconfigured. Static logic is also known as top-level logic.

Static Design

The static design is the part of the design that does not change during partial reconfiguration. The static design includes the top-level and all modules not defined as reconfigurable. The static design is built with static logic and static routing.

Design Considerations

Dynamic Function eXchange is an expert flow within the Vivado Design Suite. The following requirements and expectations need to be understood before embarking on a DFX project.

Design Requirements and Guidelines

- Dynamic Function eXchange requires the use of Vivado 2013.3 or newer.
 - Partial Reconfiguration is supported in the ISE Design Suite as well. Use the ISE Design Suite for PR only with Virtex-6, Virtex-5 and Virtex-4 devices. See the *Partial Reconfiguration User Guide* (UG702) [Ref 12] for more information.
- Floorplanning is required to define reconfigurable regions, per element type.
 - For 7 series, vertically align Pblocks with frame/clock region boundaries. This produces the best QoR and allows RESET_AFTER_RECONFIG to be enabled.
 - For UltraScale, the floorplanning is more flexible. Xilinx recommends stopping the Pblock short of frame/clock region boundaries to allow for expanded routing, which can greatly improve routability and QoR.
 - Horizontal alignment rules also apply. See [Create a Floorplan for the Reconfigurable Region in Chapter 3](#) for more information.
- Bottom-up/OOC synthesis (to create multiple netlist/DCP files) and management of Reconfigurable Module netlist files is the responsibility of the user.
 - For third party synthesis tools, I/O insertion must be disabled.
 - For Vivado OOC synthesis, I/O insertion is automatically disabled in the out_of_context mode.
- Standard timing constraints are supported, and additional timing budgeting capabilities are available if needed.
- A unique set of design rule checks (DRCs) has been established to help ensure successful design completion.
- A DFX design must consider the initiation of partial reconfiguration as well as the delivery of partial BIT files, either within the FPGA or as part of the system design.
- The Vivado Design Suite includes support for the Partial Reconfiguration Controller IP. This customizable IP manages the core tasks for partial reconfiguration in any Xilinx

device. The core receives triggers from hardware or software, manages handshaking and decoupling tasks, fetches partial bitstreams from memory locations, and delivers them to the ICAP. More information on the [PR Controller IP](#) is available on the Xilinx website.

- A Reconfigurable Partition must contain a super set of all pins to be used by the varying Reconfigurable Modules implemented for the partition. If an RM uses different inputs or outputs from another RM, the resulting RM inputs or outputs might not connect inside of the RM. The tools handle this by inserting a LUT1 buffer within the RM for all unused inputs and outputs. The output LUT1 is tied to a constant value and the value of the constant can be controlled by `HD.PARTPIN_TIEOFF` property on the unused output pin. For more information on this property refer to [Black Boxes in Chapter 5](#)
- Black boxes are supported for bitstream generation. See [Black Boxes in Chapter 5](#) for details about how to tie off ports with constant values.
- For user reset signals, determine if the logic inside the RM is level or edge sensitive. If the reset circuit is edge sensitive (as it may be in some IP such as FIFOs), then the RM reset should not be applied until after reconfiguration is complete.

Design Performance

Performance metrics vary from design to design, and the best results are achieved if you follow the Hierarchical Design techniques suggested in [Appendix A, Hierarchical Design Flows](#). This document was created for the ISE Design Suite, but the methodologies contained therein apply for the Vivado Design Suite. You can find additional design recommendations in the *UltraFast Design Methodology Guide for the Vivado Design Suite* (UG949) [\[Ref 13\]](#).

However, the additional restrictions that are required for silicon isolation are expected to have an impact on most designs. The application of partial reconfiguration rules, such as routing containment, exclusive placement, and no optimization across reconfigurable module boundaries, means that the overall density and performance is lower for a DFX design than for the equivalent flat design. The overall design performance for DFX designs varies from design to design, based on factors such as the number of reconfigurable partitions, the number of interface pins to these partitions, and the size and shape of Pblocks.

Any potential Dynamic Function eXchange design must have extra timing slack and resource overhead before considering this solution. See the [Building Up Implementation Requirements, page 76](#) section for more information on evaluating a design for DFX.

Design Criteria

- Some component types can be reconfigured and some cannot.

For 7 series devices, the component rules are as follows:

- Reconfigurable resources include CLB, block RAM, and DSP component types as well as routing resources.
- Clocks and clock modifying logic cannot be reconfigured, and therefore must reside in the static region.
 - Includes BUFG, BUFR, MMCM, PLL, and similar components
- The following components cannot be reconfigured, and therefore must reside in the static region:
 - I/O and I/O related components (ISERDES, OSERDES, IDELAYCTRL)
 - Serial transceivers (MGTs) and related components
 - Individual architecture feature components (such as BSCAN, STARTUP, ICAP, XADC.)

For UltraScale and UltraScale+ devices, the list of reconfigurable component types is more extensive:

- CLB, block RAM, and DSP component types as well as routing resources
- Clocks and clock modifying logic, including BUFG, MMCM, PLL, and similar components
- I/O and I/O related components (ISERDES, OSERDES, IDELAYCTRL)
 - Note:** The types of changes for I/O components is limited. See [I/O Rules in Chapter 7](#) for more information.
- Serial transceivers (MGTs) and related components
- PCIe, CMAC, Interlaken, and SYSMON blocks
- Bitstream granularity of these new components require that certain rules are followed. For example, partial reconfiguration of I/O require that the entire bank, plus all clocking resources in that frame are reconfigured together.
- Only the configuration components (such as BSCAN, STARTUP, ICAP, and FRAME_ECC) must remain in the static portion of the design.
- Global clocking resources to Reconfigurable Partitions are limited, depending on the device and on the clock regions occupied by these Reconfigurable Partitions.
- IP restrictions may occur due to components used to implement the IP or due to connections required by the IP. Examples include:
 - Vivado Debug Cores (See [Using Vivado Debug Cores](#) for more information on using debug cores inside of RMs)
 - IP modules with embedded global buffers or I/O (7 series only)
 - Memory IP controller (MMCM and BSCAN)

- Reconfigurable Modules must be initialized to ensure a predictable starting condition after reconfiguration. For all devices other than 7 series, GSR is automatically applied after DFX completes. For 7 series devices, GSR can be turned on, after meeting Pblock requirements, with the **RESET_AFTER_RECONFIG** Pblock property.
- Decoupling logic is highly recommended to disconnect the reconfigurable region from the static portion of the design during the act of partial reconfiguration.
 - GSR events hold all logic inside the RM in reset until configuration completes. However, RM outputs can be random and all downstream logic should be decoupled. For 7 series, if **RESET_AFTER_RECONFIG** is not used, additional decoupling of clocks and inputs can be required to prevent unintended capture of erroneous data of during reconfiguration (e.g. spurious write to memory).
 - The Vivado Design Suite includes the Partial Reconfiguration Decoupler IP. This IP allows users to easily insert MUXes to efficiently decouple AXI4-Lite, AXI4-Stream, and custom interfaces. More information on the [PR Decoupler IP](#) is available on the Xilinx website.
- A Reconfigurable Partition must be floorplanned with a Pblock, so the module must be a block that can be physically isolated and meet timing. If the module is complete, it is recommended to run this design through a non-DFX flow to get an initial evaluation of placement, routing, and timing results. If the design has issues in a non-DFX flow, these should be resolved before moving on to the DFX flow.
- Optimize an RP's interface as much as possible. An excessive number of interface pins on an RP can cause timing and routing issues. This is especially true if the Partition Pins are densely placed. This can happen for two reasons:
 - a. RP Pblock is relatively small compared to the number of Partition Pins
 - b. All the Partition Pins are placed in a small area due to Static connections.

Consider the RP interface when designing and floorplanning for DFX.

- Virtex-7 SSI devices (7V2000T, 7VX1140T, 7VH870T, 7VH580T) have two fundamental requirements. These requirements are:
 - Reconfigurable regions must be fully contained within a single SLR. This ensures that the global reset events are properly synchronized across all elements in the Reconfigurable Module, and that all super long lines (SLL) are contained within the static portion of the design. SLL are not partially reconfigurable.
 - If the initial configuration of a 7 series SSI device is done through an SPIx1 interface, partial bitstreams must be delivered to the ICAP located on the SLR where the Reconfigurable Partition exists, or to an external port, such as JTAG. If the initial configuration is done through any other configuration port, the master ICAP can be used as the delivery port for partial bitstreams.

- UltraScale devices have a new requirement related to partial reconfiguration events. Before a partial bitstream for a new Reconfigurable Module is loaded, the current Reconfigurable Module must be “cleared” to prepare for reconfiguration. UltraScale+ devices do not have this limitation. For more information, see [Summary of BIT Files for UltraScale Devices in Chapter 8](#).
- Dedicated encryption support for partial bitstreams is available natively. See [Known Limitations, page 146](#) for specific unsupported use cases for UltraScale devices.
- Devices can use a per-frame CRC checking mechanism, enabled by `write_bitstream`, to ensure each frame is valid before loading.
- Optimization across the DFX boundary is prohibited by the implementation tools. Often the WNS paths in a DFX design are high fanout control/reset signals that cross the RP boundary. Avoid high fanout signals crossing the RP boundary because the drivers cannot be replicated. To allow the tools maximum flexibility of optimization/replication, consider the following:
 - For inputs to the RP, make the signal crossing the RP boundary a single fanout net, and register the signal inside the RM before the fanout. This can be replicated as necessary inside the RM (or put on global resources).
 - For outputs, again make the signal crossing the DFX boundary a single fanout net. Register the signal in static before the fanout for replication/optimization.
- For design with multiple RPs, Xilinx recommends not having direct connections between two RPs. This includes connections that go through asynchronous static logic (not registered in static). If direct connections exist between two RPs, all possible configurations must be verified in static timing analysis to ensure timing is met across these interfaces. This can be done for closed systems that are fully owned and maintained by a single user, but can be impossible to verify for designs where different RMs are developed by multiple users. Adding a synchronous endpoint in static ensures timing is always met on any configuration, as long as the configuration where the RM was implemented met timing.

Dynamic Function eXchange is a powerful capability within Xilinx devices, and understanding the capabilities of the silicon and software is instrumental to success. While trade-offs must be recognized and considered during the development process, the overall result is a more flexible implementation of your FPGA design.

Dynamic Function eXchange Licensing

Dynamic Function eXchange is available as a feature within the Vivado Design Suite. Starting with Vivado 2019.1, no specific license code is necessary to use this feature for any edition of Vivado.

For older versions of Vivado, a Partial Reconfiguration license is included with every System Edition and Design Edition seat, and is available for purchase for WebPACK Edition seats.

Common Applications

Overview

The basic premise of Dynamic Function eXchange (DFX) is that the device hardware resources can be time-multiplexed similar to the ability of a microprocessor to switch tasks. Because the device is switching tasks in hardware, it has the benefit of both flexibility of a software implementation and the performance of a hardware implementation. Several different scenarios are presented here to illustrate the power of this technology.

Networked Multiport Interface

Dynamic Function eXchange optimizes traditional FPGA applications by reducing size, weight, power, and cost. Time-independent functions can be identified, isolated, and implemented as Reconfigurable Modules and swapped in and out of a single device as needed. A typical example is a 40G OTN muxponder application. The ports of the client side of the muxponder can support multiple interface protocols. However, it is not possible for the system to predict which protocol will be used before the FPGA is configured. To ensure that the FPGA does not have to be reconfigured and thus disable all ports, every possible interface protocol is implemented for every port, as illustrated in [Figure 2-1, page 18](#).

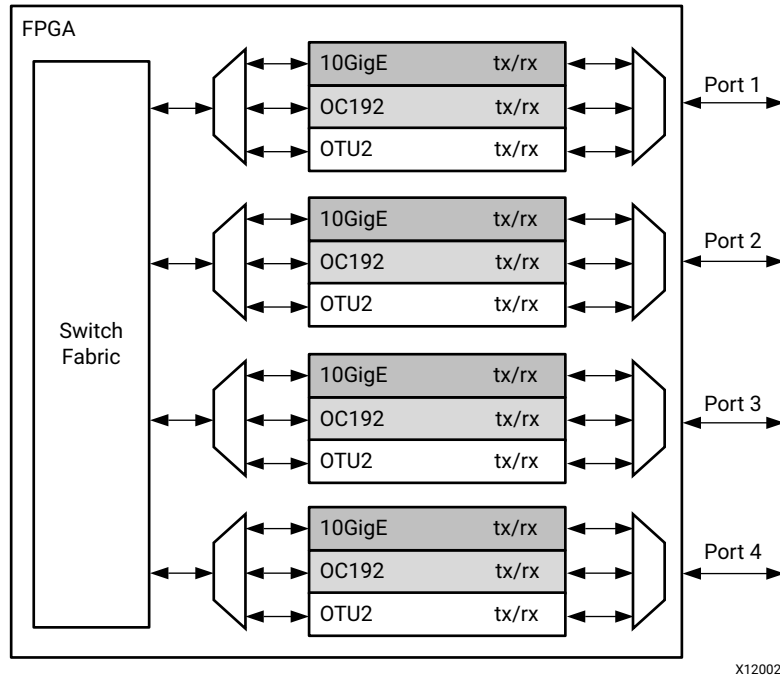


Figure 2-1: Network Switch Without Partial Reconfiguration

This is an inefficient design because only one of the standards for each port is in use at any point in time. Dynamic Function eXchange enables a more efficient design by making each of the port interfaces a Reconfigurable Module, as shown in Figure 2-2. This also eliminates the MUX elements required to connect multiple protocol engines to one port.

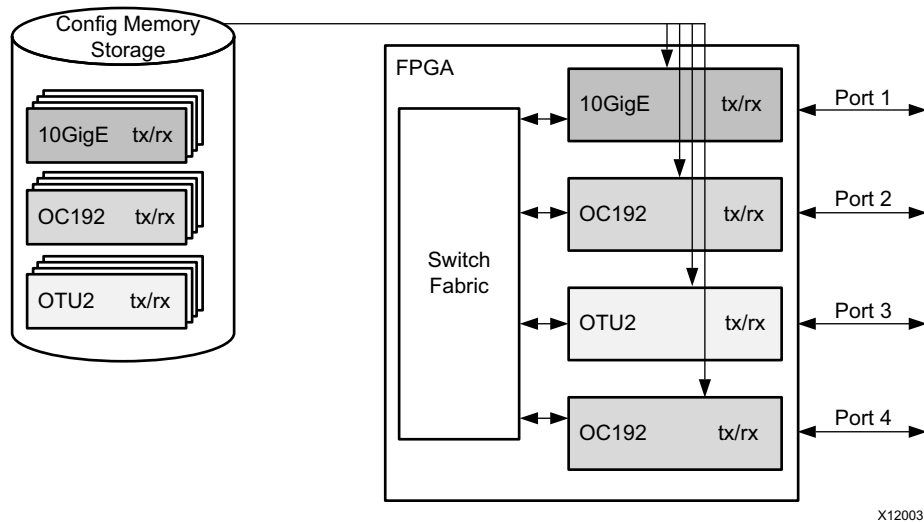


Figure 2-2: Network Switch With Partial Reconfiguration

A wide variety of designs can benefit from this basic premise. Software defined radio (SDR), for example, is one of many applications that has mutually exclusive functionality, and which sees a dramatic improvement in flexibility and resource usage when this functionality is multiplexed.

There are additional advantages with a dynamically reconfigurable design other than efficiency. In the [Figure 2-2, page 18](#) example, a new protocol can be supported at any time without affecting the static logic, the switch fabric in this example. When a new standard is loaded for any port, the other existing ports are not affected in any way. Additional standards can be created and added to the configuration memory library without requiring a complete redesign. This allows greater flexibility and reliability with less down time for the switch fabric and the ports. A debug module could be created so that if a port was experiencing errors, an unused port could be loaded with analysis/correction logic to handle the problem real-time.

In the [Figure 2-2, page 18](#) example, a unique partial BIT file must be generated for each unique physical location that could be targeted by each protocol. Partial BIT files are associated with an explicit region on the device. In this example, sixteen unique partial BIT files to accommodate four protocols for four locations.

Configuration by Means of Standard Bus Interface

Dynamic Function eXchange can create a new configuration port using an interface standard more compatible with the system architecture. For example, the FPGA could be a peripheral on a PCIe® bus and the system host could configure the FPGA through the PCIe connection. After power-on reset the FPGA must be configured with a full BIT file. However, the full BIT file might only contain the PCIe interface and connection to the internal configuration access port (ICAP).

Bitstream compression can be used to reduce the size and therefore configuration time of this initial device load, helping the FPGA configuration meet PCIe enumeration specifications.

The system host could then configure the majority of the FPGA functionality with a partial BIT file downloaded through the PCIe port as shown in [Figure 2-3](#). An example of fast configuration over PCIe is shown in XAPP1338, with an example targeting UltraScale+ included.




Figure 2-3: Configuration by Means of PCIe Interface

The PCIe standard requires the peripheral (the FPGA in this case) to acknowledge any requests even if it cannot service the request. Reconfiguring the entire FPGA would violate this requirement. Because the PCIe interface is part of the static logic, it is always active during the dynamic reconfiguration process, thus ensuring that the FPGA can respond to PCIe commands even during reconfiguration.

Tandem Configuration is a related solution that at first glance appears to be the same as is shown here. However, the solution using Dynamic Function eXchange differs from Tandem Configuration in two regards:

- The configuration process with DFX is a full device configuration, made smaller and faster through compression, followed by a partial bitstream that overwrites the black box region to complete the overall configuration. Tandem Configuration is a two-stage configuration where each configuration frame is programmed exactly once.
- Tandem Configuration for 7 series devices does not permit dynamic reconfiguration of the user application. Using DFX, the dynamic region can be reloaded with different user applications or field updates. Tandem Configuration for UltraScale devices does permit Field Updates and compatibility with DFX in general. The overall flow is Tandem Configuration for a two-stage initial load, followed by partial reconfiguration to dynamically modify the user application.

Tandem Configuration is designed to be a specific solution for a specific goal: fast configuration of a PCIe endpoint to meet enumeration requirements. For more information, see the following manuals:

- *7 Series FPGAs Integrated Block for PCI Express Product Guide* (PG054) [Ref 14]
- *Virtex-7 FPGA Gen3 PCIe Integrated Block for PCI Express Product Guide* (PG023) [Ref 15]
- *LogiCORE IP UltraScale FPGAs Gen3 Integrated Block for PCI Express Product Guide* (PG156) [Ref 16]
- *UltraScale+ Devices Integrated Block for PCI Express Product Guide* (PG213) [Ref 31]

Dynamically Reconfigurable Packet Processor

A packet processor can use Dynamic Function eXchange to change its processing functions quickly, based on the packet types received. In Figure 2-4, a packet has a header that contains the partial BIT file, or a special packet contains the partial BIT file. After the partial BIT file is processed, it is used to reconfigure a co-processor in the FPGA. This is an example of the FPGA reconfiguring itself based on the data packet received instead of relying on a predefined library of partial BIT files.

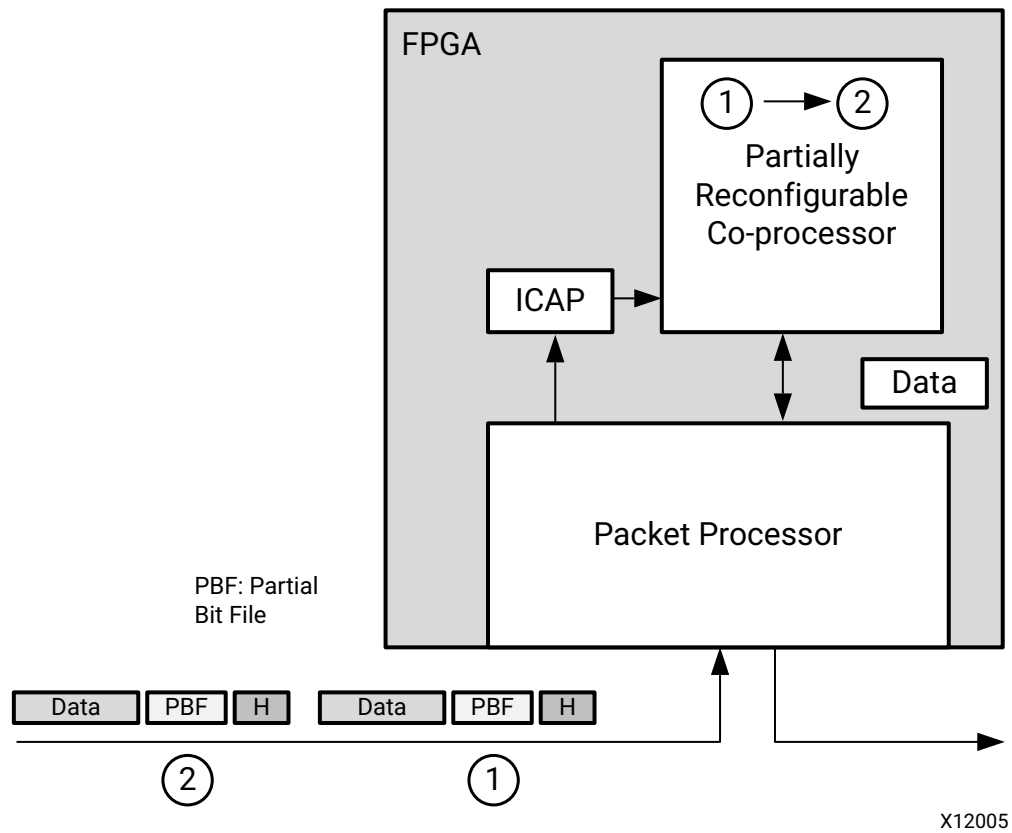


Figure 2-4: Dynamically Reconfigurable Packet Processor

Asymmetric Key Encryption

There are some new applications that are not possible without Dynamic Function eXchange. A very secure method for protecting the FPGA configuration file can be architected when Dynamic Function eXchange and asymmetric cryptography are combined. (See [Public-key cryptography](#) for asymmetric cryptography details.)

In [Figure 2-5](#), the group of functions in the shaded box can be implemented within the physical package of the FPGA. The cleartext information and the private key never leave a well-protected container.

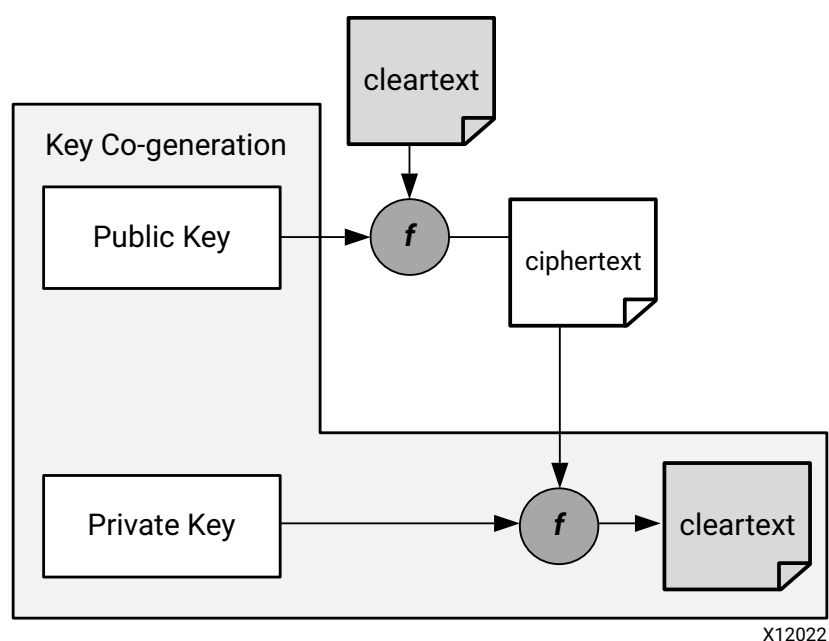


Figure 2-5: Asymmetric Key Encryption

In a real implementation of this design, the initial BIT file is an unencrypted design that does not contain any proprietary information. The initial design only contains the algorithm to generate the public-private key pair and the interface connections between the host, FPGA and ICAP.

After the initial BIT file is loaded, the FPGA generates the public-private key pair. The public key is sent to the host which uses it to encrypt a partial BIT file. The encrypted partial BIT file is downloaded to the FPGA where it is decrypted and sent to the ICAP to partially reconfigure the FPGA, as shown in [Figure 2-6, page 23](#).

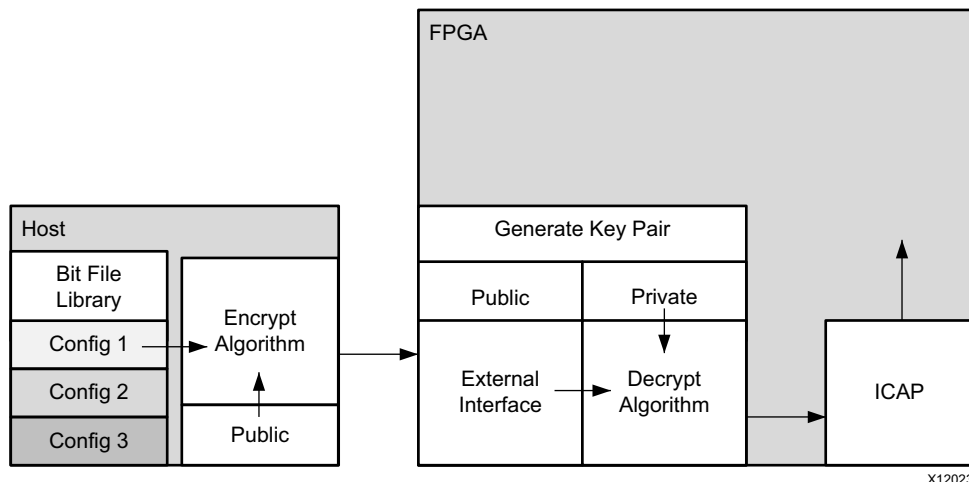


Figure 2-6: Loading an Encrypted Partial Bit File

The partial BIT file could be the vast majority of the FPGA design with the logic in the static design consuming a very small percentage of the overall FPGA resources.

This scheme has several advantages:

- The public-private key pair can be regenerated at any time. If a new configuration is downloaded from the host it can be encrypted with a different public key. If the FPGA is configured with the same partial BIT file, such as after a power-on reset, a different public key pair is used even though it is the same BIT file.
- The private key is stored in SRAM. If the FPGA ever loses power the private key no longer exists.
- Even if the system is stolen and the FPGA remains powered, it is extremely difficult to find the private key because it is stored in the general purpose FPGA programmable logic. It is not stored in a special register. You could manually locate each register bit that stores the private key in physically remote and unrelated regions.

Summary

In addition to reducing size, weight, power and cost, Dynamic Function eXchange enables new types of FPGA designs that would otherwise be impossible to implement.

Vivado Software Flow

Overview

The Vivado[®] Dynamic Function eXchange (DFX) design flow is similar to a standard design flow, with some notable differences. The implementation software automatically manages the low-level details to meet silicon requirements. You must provide guidance to define the design structure and floorplan. The following steps summarize processing a DFX design:

1. Synthesize the static and Reconfigurable Modules separately.
2. Create physical constraints (Pblocks) to define the reconfigurable regions.
3. Set the **HD.RECONFIGURABLE** property on each Reconfigurable Partition.
4. Implement a complete design (static and one Reconfigurable Module per Reconfigurable Partition) in context.
5. Save a design checkpoint for the full routed design.
6. Remove Reconfigurable Modules from this design and save a static-only design checkpoint.
7. Lock the static placement and routing.
8. Add new Reconfigurable Modules to the static design and implement this new configuration, saving a checkpoint for the full routed design.
9. Repeat Step 8 until all Reconfigurable Modules are implemented.
10. Run a verification utility (`pr_verify`) on all configurations.
11. Create bitstreams for each configuration.

Dynamic Function eXchange Commands

The DFX flows are supported through the non-project batch/Tcl interface (no project based commands), as well as within an RTL-based project flow. Example scripts for the non-project flow are provided in the *Vivado Design Suite Tutorial: Dynamic Function eXchange* (UG947) [Ref 1], along with step-by-step instructions for setting up the flows. See that Tutorial for more information.

Even with the introduction of the Dynamic Function eXchange terminology, the underlying design flow remains unchanged. Fundamental Tcl commands remain unchanged so that existing projects and scripts will safely migrate forward. Designs and scripts created prior to Vivado 2019.2 require no modification when updating to this release.

The following sections describe a few specialized commands and options needed for the DFX flows. Examples of how to use these commands to run a DFX flow are given. For more information on individual commands, see the *Vivado Design Suite Tcl Command Reference Guide* (UG835) [Ref 17].

Synthesis

Synthesizing a partially reconfigurable design does not require any special commands, but does require bottom-up synthesis. There are currently no unsupported commands for synthesis, optimization, or implementation.

These synthesis tools are supported:

- XST (supported for 7 series only)
- Synplify
- Vivado Synthesis



IMPORTANT: *Bottom-up synthesis refers to a synthesis flow in which each module has its own synthesis project. This generally involves turning off automatic I/O buffer insertion for the lower level modules.*

This document only covers the Vivado synthesis flow.

Synthesizing the Top Level

You must have a top-level netlist with a black box for each Reconfigurable Partition (RP). This requires the top-level synthesis to have module or entity declarations for the partitioned instances, but no logic; the module is empty.

The top-level synthesis infers or instantiates I/O buffers on all top level ports. For more information on controlling buffer insertion, see this [link](#) in the *Vivado Design Suite User Guide: Synthesis* (UG901) [Ref 18].

```
synth_design -flatten_hierarchy rebuilt -top <top_module_name> -part <part>
```

Synthesizing Reconfigurable Modules

Because each Reconfigurable Module must be instantiated in the same black box in the static design, the different versions must have identical interfaces. The name of the block must be the same in each instance, and all the properties of the interfaces (names, widths, direction) must also be identical. Each configuration of the design is assembled like a flat design.

To synthesize a Reconfigurable Module, turn off all buffer insertions. You can do so in Vivado Synthesis using the `synth_design` command in conjunction with the `-mode out_of_context` switch:

```
synth_design -mode out_of_context -flatten_hierarchy rebuilt -top
<reconfig_module_name> -part <part>
```

Table 3-1: `synth_design` Options

Command Option	Description
<code>-mode out_of_context</code>	Prevents I/O insertion for synthesis and downstream tools. The <code>out_of_context</code> mode is saved in the checkpoint if <code>write_checkpoint</code> is issued.
<code>-flatten_hierarchy rebuilt</code>	There are several values allowed for <code>-flatten_hierarchy</code> , but <code>rebuilt</code> is the recommended setting for DFX flows.
<code>-top</code>	This is the module/entity name of the module being synthesized.
<code>-part</code>	This is the Xilinx® part being targeted (for example, <code>xc7k325tffg900-3</code>)

The `synth_design` command synthesizes the design and stores the results in memory. In order to write the results out to a file, use:

```
write_checkpoint <file_name>.dcp
```

It is recommended to close the design in memory after synthesis, and run implementation separately from synthesis.

Reading Design Modules

If there is currently no design in memory, you must load a design. This can be done in a variety of ways, for either the static design or for Reconfigurable Modules. After the configurations are implemented, checkpoints are exclusively used to read in placed and routed module databases.

Method 1: Add and Link Files

This is the recommended method to load and link all design sources in the most explicit and thorough manner. The following steps pull in all necessary design sources and define the Reconfigurable Partition boundaries.

1. Create a new project in memory. While this allows you to select a target device, the project is not saved.

```
create_project -part <part> -in_memory
```

2. Add all the design sources. This can include multiple checkpoints for static or reconfigurable logic, including lower-level RM sources.

```
add_files <top>.dcp
add_files <rp1_rmA_top>.dcp
add_files <rp1_rmA_lower>.dcp
add_files <rp2_rmA_top>.dcp
```

3. Use the `SCOPED_TO_CELLS` property to define relationships between levels of hierarchy.

```
set_property SCOPED_TO_CELLS {<RP1_module_instance>} [get_files <rp1_rmA_top>.dcp]
set_property SCOPED_TO_CELLS {<RP1_lower_module_instance>} [get_files <rp1_rmA_lower>.dcp]
set_property SCOPED_TO_CELLS {<RP2_module_instance>} [get_files <rp2_rmA_top>.dcp]
```

4. Link the design together, defining all Reconfigurable Partitions.

```
link_design -top <top> -part <part> -reconfig_partitions {<RP1_module_instance> <RP2_module_instance>}
```

Table 3-2: `link_design` Options

Command Option	Description
<code>-part</code>	This is the Xilinx part being targeted (for example, <code>xc7k325tffg900-3</code>)
<code>-top</code>	This is the module/entity name of the module being implemented. This switch can be omitted if <code>set_property top <top_module_name> [current_fileset]</code> is issued prior to <code>link_design</code> .
<code>-reconfig_partitions <args></code>	Specify a list of reconfigurable partitions to load while opening the design. The specified reconfigurable partitions are then marked with the <code>HD.RECONFIGURABLE</code> property for proper handling in the design.
<code>-pr_config <arg></code>	For the project-based design flow only. This option specifies the PR Configuration to apply while opening the design.

Method 2: Read Netlist Design

This approach should be used when modules have been synthesized by tools other than Vivado synthesis.

```
read_edif <top>.edf/edn/ngc
read_edif <rp1_a>.edf/edn/ngc
read_edif <rp2_a>.edf/edn/ngc
link_design -top <top_module_name> -part <part>
```

Method 3: Open/Read Checkpoint

If the static (top-level) design has synthesis or implementation results stored as a checkpoint, then it can be loaded using the `open_checkpoint` command. This command reads in the static design checkpoint and opens it in active memory:

```
open_checkpoint <file>
```

If the checkpoint is for the complete netlist of a reconfigurable module (that is, not for static), then the instance name can be specified using `read_checkpoint -cell`. If the checkpoint is a post-implementation checkpoint, then the additional `-strict` option must be used as well. This option can also be used with a post-synthesis checkpoint to ensure exact port matching has been achieved. To read in a checkpoint in a Reconfigurable Module, the top-level design must already be opened, and must have a black box for the specified cell. Then the following command can be specified:

```
read_checkpoint -cell <cellname> <file> [-strict]
```

Table 3-3: `read_checkpoint` Switches

Switch Name	Description
<code>-cell</code>	Specifies the full hierarchical name of the Reconfigurable Module.
<code>-strict</code>	Requires exact ports match for replacing cell, and checks that part, package, and speed grade values are identical. Should be used when restoring implementation data.
<code><file></code>	Specifies the full or relative path to the checkpoint (DCP) to be read in.



CAUTION! Do not use this method if the synthesized checkpoint has underlying modules that are not included. The `read_checkpoint -cell` approach does not support nesting. Use the `link_design` approach instead in [Method 1: Add and Link Files](#).

Method 4: Open Checkpoint/Update Design

This is useful when the synthesis results are in the form of a netlist (EDF or EDN), but static has already been implemented. The following example shows the commands for the second configuration in which this is true.

```
open_checkpoint <top>.dcp
lock_design -level routing
update_design -cells <rp1> -from_file <rp1_b>.{edf/edn}
update_design -cells <rp2> -from_file <rp2_b>.{edf/edn}
```

Adding Reconfigurable Modules with Sub-Module Netlists

If a Reconfigurable Module has sub-module netlists, it can be difficult for the Vivado tools to process the sub-module netlists. This is because in the DFX flow the RM netlists are added to a design that is already open in memory. This means the `update_design -cells` command must be used, which requires the cell name for every EDIF file, which can be troublesome to get.

There are two ways to make loading RM sub-module netlists easier in the Vivado Design Suite.

Method 1: Create a Single RM Checkpoint (DCP)

Create an RM checkpoint (DCP) that includes all netlists. Use `add_files` to add all of the EDIF (or NGC) files, and use `link_design` to resolve the EDIF files to their respective cells. Here is an example of the commands used in this process:

```
add_files [list rm.edf ip_1.edf ... ip_n.edf]
# Run if RM XDC exists
add_files rm.xdc
link_design -top <rm_module> -part <part>
write_checkpoint rm_v#.dcp
close_project
```



IMPORTANT: *Using this methodology to combine/convert a netlist into a DCP is the recommended way to handle an RM that has one or more NGC source files as well.*

Then this newly-created RM checkpoint can be used in the DFX flow. In the commands below, the single `read_checkpoint -cell` command replaces what could be many `update_design -cell` commands.

```
add_files static.dcp
link_design -top <top> part <part>
lock_design -level routing
read_checkpoint -cell <rm_inst> rm_v#.dcp
```

Method 2: Place the Sub-Module Netlists in the Same Directory as the RM's Top-Level Netlist

When the top-level RM netlist is read into the DFX design using `update_design -cell`, make sure that all sub-module netlists are in the same directory as the RM top-level netlist. In this case, the lower-level netlists do not need to be specified, but they are picked up automatically by the `update_design -cells` command. This is less explicit than Method 1, but requires fewer steps. In this case the commands to load the RM netlist would look like the following:

```
add_files static.dcp
link_design -top <top> part <part>
lock_design -level routing
update_design -cells <rm_inst> -from_file rm_v#.edf
```

In the last (`update_design`) command above, the lower-level netlists are picked up automatically if they are in the same directory as `rm_v#.edf`.

Reading Design Constraints

New constraints can be applied for each configuration at various points in the flow. If an RM is read in as a DCP, then any constraints stored in the DCP are automatically applied. Additionally, the `read_xdc` command can be used to apply constraints scoped to the top-level, or to the specific cell (using `-cell` switch). If constraints are expected to directly or indirectly affect the RM, then the RM must be resolved (not a black box) prior to reading in the new constraints. Otherwise, the constraints may be dropped or not correctly propagated in the constraint system. Because Static is only placed and routed in the initial configuration, all constraints for subsequent configurations (where Static is locked) should be focused strictly on the RP regions being implemented.

Implementation

Because the DFX flow allows for various configurations in hardware, multiple implementation runs are required. Each implementation of a DFX design is referred to as a *configuration*. Each module of the design (static or Reconfigurable Module) can be implemented or imported (if previously implemented). Implementation results for the static design must be consistent for each configuration, so that the design is implemented in one configuration, and then imported in subsequent configurations. Additional configurations can be constructed by importing static, and implementing or importing each Reconfigurable Module.

There are no restrictions to the support of implementation commands or options for DFX, but certain optimizations and sub-routines are not done if they oppose the fundamental requirements of partial reconfiguration. The following list of commands can be run after the logical design is loaded (using `link_design` or `open_checkpoint`):

```
# Run if all constraints are not already loaded
read_xdc

# Optional command
opt_design

place_design

# Optional command
phys_opt_design

route_design
```

Preserving Implementation Data

In the DFX flow, it is a requirement to lock down the placement and routing results of the static logic from the first configuration for all subsequent configurations. The static implementation of the first configuration must be saved as a checkpoint. When the checkpoint is read for subsequent configurations, the placement and routing must be locked, to ensure that the static design remains completely identical from configuration to configuration. To lock the placement and routing of an imported checkpoint (static or reconfigurable), the `lock_design` command is used.

```
lock_design -level [logical|placement|routing] [cell_name]
```

When locking down the static logic with the above command, the optional `[cell_name]` can be omitted.

```
lock_design -level routing
```

To lock the results of an imported RM, the full hierarchical name should be specified within the post-implementation checkpoint:

```
lock_design -level routing u0_RM_instance
```

For Dynamic Function eXchange, the only supported preservation level is `routing`. Other preservation levels are available for this command, but they must only be used for other Hierarchical Design flows.

Dynamic Function eXchange Constraints and Properties

There are properties and constraints unique to the Dynamic Function eXchange flow. These initiate DFX-specific implementation processing and apply specific characteristics in the partial bitstreams. The four areas for constraints and properties for DFX are:

Table 3-4: Constraints and Properties

Constraints and Properties	Necessity
Defining a module as reconfigurable	Required
Creating a floorplan for the reconfigurable region	Required
Applying reset after reconfiguration	Optional, but highly recommended
Turn on visualization scripts	Optional

Define a Module as Reconfigurable

In order to implement a DFX design, it is required to specify each Reconfigurable Module as such. To do this you must set a property on the top level of each hierarchical cell that is going to be reconfigurable. For example, take a design where one Reconfigurable Partition named `inst_count` exists, and it has two Reconfigurable Modules, `count_up` and `count_down`. The following command must be issued prior to implementation of the first configuration.

```
set_property HD.RECONFIGURABLE TRUE [get_cells inst_count]
```

This initiates the Dynamic Function eXchange features in the software that are required to successfully implement a DFX design. The `HD.RECONFIGURABLE` property implies a number of underlying constraints and tasks:

- Sets `DONT_TOUCH` on the specified cell and its interface nets. This prevents optimization across the boundary of the module.
- Sets `EXCLUDE_PLACEMENT` on the cell's Pblock. This prevents static logic from being placed in the reconfigurable region.
- Sets `CONTAIN_ROUTING` on the cell's Pblock. This keeps all the routing for the Reconfigurable Module within the bounding box.
- Enables special code for DRCs, clock routing, etc.

Create a Floorplan for the Reconfigurable Region

Each Reconfigurable Partition is required to have a Pblock to define the physical resources available for the Reconfigurable Module. Because this Pblock is set on a Reconfigurable Partition, these restrictions and requirements apply:

- The Pblock must contain only valid reconfigurable element types. The region may overlap other site types, but these other sites must not be included in the `resize_pblock` commands.
- Multiple Pblock rectangles for each component type may be used to create the Reconfigurable Partition region, but for the greatest routability, they should be contiguous. Gaps to account for non-reconfigurable resources are permitted, but in general, the simpler the overall shape, the easier the design will be to place and route.
- If using the `RESET_AFTER_RECONFIG` property for 7 series devices, the Pblock height must align to clock region boundaries. See [Apply Reset After Reconfiguration](#) for more detail.
- The width and composition of the Pblock must not split interconnect columns for 7 series devices. See [Creating Pblocks for 7 Series Devices in Chapter 6](#) for more detail.
- The resource usage of the largest RM needs to be taken into consideration when defining the Pblock in certain parts. If the largest RM exceeds the documented maximum resource counts of the target device, `write_bitstream` will generate an error.
- The Pblock must not overlap any other Pblock in the design.
- Nesting of Reconfigurable Partitions (a Reconfigurable Partition within another Reconfigurable Partition) is currently not supported. Standard Pblocks for floorplanning logic within a Reconfigurable Partition are supported, as are nested Pblocks.

Table 3-5: Pblock Commands and Properties

Command/Property Name	Description
<code>create_pblock</code>	Command used to create the initial Pblock for each Reconfigurable Partition instance.
<code>add_cells_to_pblock</code>	Command used to specify the instances that belong to the Pblock. This is typically a level of hierarchy as defined by the bottom-up synthesis processing.
<code>resize_pblock</code>	Command used to define the site types (such as SLICE or RAMB36) and site locations that are owned by the Pblock.
<code>RESET_AFTER_RECONFIG</code>	Pblock property used to control the use of the dedicated GSR event on the reconfigurable region. Use of this property is highly recommended and, for 7 series and Zynq devices, requires clock region alignment in the vertical direction.
<code>CONTAIN_ROUTING</code>	Pblock property used to control the routing to prevent usage of routing resources not owned by the Pblock. This property is mandatory for PR and is set to True automatically for Reconfigurable Partitions. Static routing is still allowed to use resources inside of the Pblock.
<code>EXCLUDE_PLACEMENT</code>	Pblock Property used to prevent the placement of any logic, not belonging to the Pblock, inside the defined Pblock RANGE. This property is mandatory for PR and set to true automatically for Reconfigurable Partitions. Static logic can be placed inside of the Reconfigurable Partition with a specific LOC property if <code>RESET_AFTER_RECONFIG</code> is not used.
<code>PARTPIN_SPREADING</code>	Used to control the maximum number of PartPins per INT tile. Default is 5. Setting a lower value (i.e. 3) increases the spreading between partition pin placements. This typically eases routing congestion in areas with dense PartPin placement, but can negatively affect RP interface timing.

The following is an example of a set of constraints for a Reconfigurable Partition:

```
#define a new pblock
create_pblock pblock_count

#add a hierarchical module to the pblock
add_cells_to_pblock [get_pblocks pblock_count] [get_cells [list inst_count]]

#define the size and components within the pblock
resize_pblock [get_pblocks pblock_count] -add {SLICE_X136Y50:SLICE_X145Y99}
resize_pblock [get_pblocks pblock_count] -add {RAMB18_X6Y20:RAMB18_X6Y39}
resize_pblock [get_pblocks pblock_count] -add {RAMB36_X6Y10:RAMB36_X6Y19}
```

Floorplan in the Vivado IDE

The Vivado IDE can be used for planning and visualization tasks. The best example of this is using the Device view to create and modify Pblock constraints for floorplanning.

1. Open the synthesized static design and the largest of each Reconfigurable Module. Here are the commands, using the tutorial design (found in the *Vivado Design Suite Tutorial: Dynamic Function eXchange* (UG947) [Ref 1]) as an example:

```
open_checkpoint synth/Static/top_synth.dcp
set_property HD.RECONFIGURABLE true [get_cells inst_count]
read_checkpoint -cell [get_cells inst_count] synth/count_up/count_synth.dcp
set_property HD.RECONFIGURABLE true [get_cells inst_shift]
read_checkpoint -cell [get_cells inst_shift] synth/shift_right/shift_synth.dcp
```

At this point, a full configuration has been loaded into memory, and the Reconfigurable Partitions have been defined.

2. To create Pblock constraints for the Reconfigurable Partitions, right-click on an instance in the Netlist window (in this case, `inst_count` or `inst_shift`) and select **Draw Pblock**. Create a rectangle in the Device view to select resources for this Reconfigurable Partition.
3. With this Pblock selected, note that the Pblock Properties pane shows the number of available and required resources. The number required is based on the currently loaded Reconfigurable Module, so keep in mind that other modules may have different requirements. If additional rectangles are required to build the appropriate shape (an "L", for example), right-click the Pblock in the Device view and select **Add Pblock Rectangle**.
4. Design rule checks (DRCs) can be issued to validate the floorplan and other design considerations for the in-memory configuration. To run, select **Reports > Report DRC** and ensure the Partial Reconfiguration checks are present (see [Figure 3-1, page 36](#)). Note that if `HD.RECONFIGURABLE` has not been set on a Pblock, only a single DRC is available, instead of the full complement shown below.




Figure 3-1: Partial Reconfiguration DRCs in the Vivado IDE

This set of DRCs can be run from the Tcl Console or within a script, by using the `report_drc` command. To limit the checks to the ones shown here for Partial Reconfiguration, use this syntax:

```
report_drc -checks [get_drc_checks HDPR*]
```

To extend the DRCs to those checked during specific phases of design processing the `-ruledeck` option can be used. For example, the following command can be issued on a placed and routed design:

```
report_drc -ruledeck bitstream_checks
```

To save these floorplanning constraints, enter the following command in the Tcl Console:

```
write_xdc top_fplan.xdc
```

The Pblock constraints stored in this constraints file can be used directly or can be copied to another top-level design constraints file. This XDC file contains all the constraints in the current design in memory not just the constraints recently added.



CAUTION! Do NOT save the overall design from the Vivado IDE using **File > Checkpoint > Save** or the equivalent button. If you save the currently loaded design in this way, you will overwrite your synthesized static design checkpoint with a new version that includes Reconfigurable Modules and additional constraints.

Using Visualization Scripts

For each Reconfigurable Partition, scripts are automatically created to confirm the site ownership for each part of a DFX design. The visualization scripts generated can vary based on architecture and need.

Scripts are automatically created for all RP Pblock in an `hd_visual` directory, which is created in the directory where the run script is launched. To use these scripts, read a routed design checkpoint into the Vivado IDE, then source one of the scripts. These design-specific scripts highlight configuration tiles as you have defined them, show configuration frames used to create the partial bit file, or show sites excluded by the DFX floorplan. Additional scripts are created for other flows, such as Module Analysis or Tandem Configuration, and are not used for DFX.

For 7 series devices, the main script is named `<rp_pblock>_AllTiles.tcl` and shows all the sites owned by the Reconfigurable Partition, for both placement and routing of any implemented Reconfigurable Modules. Other scripts are created for very specific goals and are not needed in most cases.

For UltraScale and UltraScale+, unique scripts named `<rp_pblock>_Placement_AllTiles.tcl` and `<rp_pblock>_Routing_AllTiles.tcl` show the boundaries for the placement and expanded routing for the reconfigurable region. The placement script shows the range available for logic placement after snapping is finished. The routing script shows the expanded routing region and represents the contents of the partial bitstream created for that RP.

For all devices, three additional scripts might be created per design when necessary: `blockedBelsRouteThrus.tcl`, `blockedPins.tcl`, and `blockedSitesInputs.tcl`. When designs encounter higher levels of congestion, these scripts are created to show restricted sites. This information can be used to adjust the size and shape of the RP pblock, and can also be shared with support for troubleshooting purposes.

Timing Constraints

Timing constraints for a partially reconfigurable design are similar to timing constraints for a traditional flat design. The primary clocks and I/Os must be constrained with the corresponding constraints. For more information on these constraints, see this [link](#) (for defining clocks) and this [link](#) (for constraining I/O delays) in the *Vivado Design Suite User Guide: Using Constraints* (UG903) [Ref 19].

After the correct constraints are applied to the design, run static timing analysis to verify the performance of the design. This verification must be run for each reconfigurable module in the overall static design. For more information on how to analyze the design, see the *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [Ref 20].

The Vivado Design Suite includes the capability to run cell level timing reports. Use the `-cell` option for `report_timing` or `report_timing_summary` to focus timing analysis on a specific Reconfigurable Module. This is especially useful on configurations where the static design has been imported and locked from a prior configuration.

There is a **Partition** column added to the timing reports generated by `report_timing` and `report_timing_summary`. It helps identify if failing paths are within static, an RM, or crosses an RP boundary. Both of these commands have a new `-no_pr_attribute` switch to turn this new functionality off. This can be useful if, for example, scripts are being used to parse the timing reports and are negatively affected by this new column.

Partition Pins

Interface points called partition pins are automatically created within the Pblock ranges defined for the Reconfigurable Partition. These virtual I/O are established within interconnect tiles as the anchor points that remain consistent from one module to the next. No physical resources such as LUTs or flip-flops are required to establish these anchor points, and no additional delay is incurred at these points.

The placer chooses locations based on source and loads and timing requirements, but you can specify these locations as well. The following constraints can be applied to influence partition pin placement.

Table 3-6: Context Properties

Command/Property Name	Description
<code>HD.PARTPIN_LOCS</code>	Used to define a specific interconnect tile (INT) for the specified port to be routed. Overrides an <code>HD.PARTPIN_RANGE</code> value. Affects placement and routing of logic on both sides of the Reconfigurable Partition boundary. Do not use this property on clock ports, as this assumes local routing for the clock. Do not use this property on dedicated connections.
<code>HD.PARTPIN_RANGE</code>	Used to define a range of component sites (SLICE, DSP, block RAM) or interconnect tiles (INT) that can be used to route the specified port(s). The value is automatically calculated based on Pblock range if no user-defined <code>HD.PARTPIN_RANGE</code> value exists.

Note: The `PARTPIN_SPREADING` property in Table 3-5, can also be used to affect Partition Pins, but is applied at the Pblock level.

Context Property Examples:

- `set_property HD.PARTPIN_LOCS INT_R_X4Y153 [get_ports <port_name>]`
- `set_property HD.PARTPIN_RANGE SLICE_X4Y153:SLICE_X5Y157 [get_ports <port_name>]`

Instance names for interconnect tile sites can be seen in the Device View with the Routing Resources enabled.

Note: The `HD.PARTPIN_RANGE` is automatically set during `place_design` if no user-defined value is found. Once the value is set, it will not be reset during interactive place and route, such as making experimental changes to the RP Pblocks and running `place_design -unplace`. In this case, the `HD.PARTPIN_RANGE` and `HD.PARTPIN_LOCS` need to be reset manually if Pblock adjustments are made. The properties can be reset like most properties.

The following Tcl proc can be useful when doing this kind of interactive floorplanning on DFX designs:

```
#####
Proc to unroute, uplace, and reset HD.PARTPIN_*
#####
proc pr_unplace {} {
    route_design -unroute
    place_design -unplace
    set cells [get_cells -quiet -hier -filter HD.RECONFIGURABLE]
    foreach cell $cells {
        reset_property HD.PARTPIN_LOCS [get_pins $cell/*]
        reset_property HD.PARTPIN_RANGE [get_pins $cell/*]
    }
}
}
```

Partition pin information can be obtained from placed or routed designs by using the `get_pplocs` command. Use either the `-nets` or `-pins` option to focus the response to a particular Reconfigurable Partition or interface pin.

```
get_pplocs -nets <args> -pins <args> [-count] [-unlocked] [-locked] [-level <arg>]
[-quiet] [-verbose]
```

Table 3-7: `get_pplocs` Options

Name	Description
<code>-nets</code>	List of nets to report its PPLOCs.
<code>-pins</code>	List of pins to report its PPLOCs.
<code>[-count]</code>	Count number of PPLOCs; Do not report PPLOC or node names.
<code>[-unlocked]</code>	Report unlocked/unfixed PPLOCs only.
<code>[-locked]</code>	Report locked/fixed PPLOCs only; use <code>-level</code> to specify locked level.
<code>[-level]</code>	Specify locked level.
<code>[-quiet]</code>	Ignore command errors.
<code>[-verbose]</code>	Suspend message limits during command execution.

Example:

```
get_pplocs -pins [get_pins u_count/*]
```


In UltraScale or UltraScale+ designs, not all interface ports receive a partition pin. With the routing expansion feature, as explained in [Expansion of CONTAIN_ROUTING Area](#), some interface nets are completely contained within the expanded region. When this happens, no partition pin is inserted; the entire net, including the source and all loads, is contained within the area captured by the partial bit file. Rather than pick an unnecessary intermediate point for the route, the entire net is rerouted, giving the Vivado tools the flexibility to pick an optimal solution.

Apply Reset After Reconfiguration

With the Reset After Reconfiguration feature, the reconfiguring region is held in a steady state during partial reconfiguration, and then all logic in the new Reconfigurable Module is initialized to its starting values. Static routes can still freely pass unaffected through the region, and static logic (and all other dynamic regions) elsewhere in the device continue to operate normally during Partial Reconfiguration. Dynamic Function eXchange with this feature behaves in the same manner as the initial configuration of the FPGA, with synchronous elements being released in a known, initialized state.



IMPORTANT: *Release of global signals such as GSR (Global Set Reset) and GWE (Global Write Enable) are not guaranteed to be synchronized chip-wide. If functionality within a Reconfigurable Module relies on synchronized startup of initialized sequential elements, the clock(s) driving the logic in that module or Clock Enables on these elements can be disabled during reconfiguration, then re-enabled after reconfiguration has been completed. For more details, see the "Design Advisory for techniques on properly synchronizing flip-flops and SRLs" answer record (AR#44174) [\[Ref 36\]](#).*

This is the **RESET_AFTER_RECONFIG** property syntax:

```
set_property RESET_AFTER_RECONFIG true [get_pblocks <reconfig_pblock_name>]
```

If the design uses the DRP interface of the 7 series XADC component, the interface will be blocked (held in reset) during partial reconfiguration when RESET_AFTER_RECONFIG is enabled. The interface will be non-responsive (busy), and there will be no access during the length of the reconfiguration period. The interface will become accessible again after partial reconfiguration is complete.

To apply the Reset After Reconfiguration methodology for 7 series and Zynq-7000 SoC devices, Pblock constraints must align to reconfigurable frames. Because the GSR affects every synchronous element within the region, exclusive use of reconfiguration frames is required; static logic is not permitted within these reconfigurable frames (static routing is permitted). Pblocks must align vertically to clock regions, since that matches the base region for a reconfigurable frame. The width of a Pblock does not matter when using **RESET_AFTER_RECONFIG**.

UltraScale™ and UltraScale+™ devices do not have this clock region alignment requirement, and GSR can be applied at a fine granularity. Because of this, **RESET_AFTER_RECONFIG** is automatically applied for all Reconfigurable Partitions in the UltraScale and UltraScale+ architecture. This capability cannot be disabled.

In [Figure 3-2](#), the Pblock on the left (`pblock_shift`) is frame-aligned because the top and bottom of the Pblock align to the height of clock region X1Y3. The Pblock on the right (`pblock_count`) is not frame-aligned.

- For 7 series devices: Pblocks that are not frame-aligned (such as `pblock_count` in the figure below) cannot have **RESET_AFTER_RECONFIG** set because any static logic placed between it and the clock region boundary above it would be affected by GSR after that module was partially reconfigured.
- For UltraScale and UltraScale+ devices: because of the improved GSR controls, both Pblocks automatically use **RESET_AFTER_RECONFIG**.

Using the **SNAPPING_MODE** constraint automatically creates legal, reconfigurable Pblocks. See [Automatic Adjustments for Reconfigurable Partition Pblocks in Chapter 6](#) (for 7 series devices) or [Automatic Adjustments for PU on Pblocks in Chapter 7](#) (for UltraScale and UltraScale+ devices) for more information.




Figure 3-2: RESET_AFTER_RECONFIG Compatible (Left) and Incompatible (Right) Pblocks

The GSR capabilities are embedded within the partial bitstreams, so nothing extra must be done to include this feature during reconfiguration. However, because this process utilizes the **SHUTDOWN** sequence (masked to the reconfiguring region only), the external **DONE** pin are pulled **LOW** when reconfiguration starts, then pull **HIGH** when it successfully completes. This behavior must be considered when setting up the board. Using the **STARTUP** block **DONE0** is not an option to prevent the **DONE** pin from changing state, since this block is disabled during shutdown. Nor can **STARTUP** be used for other purposes, such as generating a configuration clock for partial reconfiguration if **RESET_AFTER_RECONFIG** is used.

Moreover, in order to open the GSR mask for only the dynamic region when reconfiguration occurs, the mask for the entire design begins as closed after the initial configuration. Each partial bitstream opens the mask for the target region, loads new configuration data, issues a GSR event for this region, then closes the mask. For UltraScale only, this process is split between two bitstreams -- see the Clearing Bitstreams section in Chapter 8 for more information. Because the mask is closed when reconfiguration is not occurring, full-device access to GSR is not permitted.

For 7 Series only, an alternative approach would be to forego this property and apply a local reset to any reconfigured logic that requires initialization to function properly. This approach does not require vertical alignment to clock region boundaries. Without GSR or a local reset, the initial starting value of a synchronous element within a reconfigured module cannot be guaranteed.

Software Flow

This section describes the basic flow, and gives sample commands to execute this flow.

Synthesis

Each module (including Static) needs to be synthesized bottom-up so that a netlist or checkpoint exists for static and each Reconfigurable Module.

1. Synthesize the top level:

```
read_verilog top.v (and other HDL associated with the static design, including  
black box module definitions for Reconfigurable Modules)
```

then:

```
read_xdc top_synth.xdc  
synth_design -top top -part xc7k70tfbg676-2  
write_checkpoint top_synth.dcp
```

2. Synthesize a Reconfigurable Module:

```
read_verilog rp1_a.v  
synth_design -top rp1 -part xc7k70tfbg676-2 -mode out_of_context  
write_checkpoint rp1_a_synth.dcp
```

3. Repeat for each remaining Reconfigurable Module

```
read_verilog rp1_b.v  
synth_design -top rp1 -part xc7k70tfbg676-2 -mode out_of_context  
write_checkpoint rp1_b_synth.dcp
```

Implementation

Create as many configurations as necessary to implement all Reconfigurable Modules at least once. The first configuration loads in synthesis results for top and the first Reconfigurable Module. You must then mark the module as being reconfigurable, then run implementation. Write out a checkpoint for the complete routed configuration, and optionally for the Reconfigurable Module so it can be reused later if desired. Finally, remove the Reconfigurable Module from the design (`update_design -cell -black_box`) and write out a checkpoint for the locked static design alone.

Configuration 1:

```
open_checkpoint top_synth.dcp
read_xdc top_impl.xdc
set_property HD.RECONFIGURABLE true [get_cells rp1]
read_checkpoint -cell rp1 rp1_a_synth.dcp
opt_design
place_design
route_design
write_checkpoint config1_routed.dcp
write_checkpoint -cell rp1 rp1_a_route_design.dcp
update_design -cell rp1 -black_box
lock_design -level routing
write_checkpoint static_routed.dcp
```

For the second configuration, load the placed and routed checkpoint for static (if it was closed), which currently has a black box for the Reconfigurable Module. Then load in the synthesis results for the second Reconfigurable Module and implement the design. Finally write out an implementation checkpoint for the second version of the Reconfigurable Module.

Configuration 2:

```
open_checkpoint static_routed.dcp
read_checkpoint -cell rp1 rp1_b_synth.dcp
opt_design
place_design
route_design
write_checkpoint config2_routed.dcp
write_checkpoint -cell rp1 rp1_b_route_design.dcp
```



TIP: Keep each configuration in a separate folder so that all intermediate checkpoints, log and report files, bit files, and other design outputs are kept unique.

If multiple Reconfigurable Partitions exist, then other configurations may be required. Additional configurations can also be created by importing previously implemented Reconfigurable Modules to create full designs that exist in hardware. This can be useful for creating full bitstreams with a desired combination for power-up, or for performing static timing analysis, power analysis, or simulation.

Full place and route results for each Reconfigurable Module checkpoint is preserved completely, so creating new configurations is easily done by loading a collection of routed checkpoints. However, there are limitations to be aware of when using the flow. Using `write_checkpoint -cell` to save the RM implementation results does not preserve constraints local to this module. For RMs with internal clock constraints or timing exceptions starting and/or ending within the RM, these constraints need to be reapplied for timing analysis after creating the new configuration. RMs with Xilinx or 3rd party IP are good examples of modules that might be exposed to this limitation.

Incremental Compile

Dynamic Function eXchange designs can use the Vivado Incremental Compilation feature for any configuration run. The flow is documented [here](#) in *Vivado Design Suite User Guide: Implementation* (UG904), and best results are seen when 95% of design instances match.

For any Incremental Compile usage, be sure to select a prior checkpoint for that specific configuration to match as much of the design as possible. Note that for the second configuration and beyond, all static logic and routing is locked, so the Incremental Compile effort will be focused within Reconfigurable Modules.

Reporting

Each step of the implementation flow performs design rule checks (DRCs) unique to partial reconfiguration. Keep a close eye on the messages given by the implementation steps to ensure no critical warnings are issued. These messages provide guidance to optimize module interfaces, floorplans, and other key aspects of DFX designs.

Most reports that can be generated do not have DFX-specific sections, but useful information can be extracted nonetheless. For example, utilization information can be obtained by using the `-pblocks` switch for the `report_utilization` command. This shows the used and available resources within a given reconfigurable module. Here is an example using the design from the *Vivado Design Suite Tutorial: Dynamic Function eXchange* (UG947) [\[Ref 1\]](#):

```
report_utilization -pblocks [get_pblocks pblock_count]
```

For clock reporting, however, `report_clock_utilization` shows the clocks reserved for partial reconfiguration implementation.

The Dynamic Function eXchange flow can be used in conjunction with the IEEE-1735 v2 encryption capability available within Vivado. Static design checkpoints can be encrypted and shared with other users without exposing details of the design. Rights management can be set such that details such as LUT contents and schematic details can be hidden, and netlist export and design modification can be disabled. Developers of dynamic regions can still insert their reconfigurable logic and implement within this locked static context. If permission is given, these developers can generate partial bitstreams from within this encrypted context for their dynamic function.

Note that a license is required to use this feature, and any licensed IP within the static region will still require a valid license to open that checkpoint even if it is encrypted.

For more information on creating encrypted design checkpoints and the options available, please consult Chapter 6 of UG1118.

Verifying Configurations

Once all configurations have been completely placed and routed, a final verification check can be done to validate consistency between these configurations using `pr_verify`. This command takes in multiple routed checkpoints (DCPs) as arguments, and outputs a log of any differences found in the static implementation and Partition Pin placement between them. Placement and routing within any RMs is ignored during the comparison.

When just two configurations are to be compared, list the two routed checkpoints as `<file1>` and `<file2>`. The `pr_verify` command loads both in memory and makes the comparison.

When more than two configurations are to be compared, provide a master configuration using the `-initial` switch, then list the remaining configurations by using the `-additional` switch, listing configurations in braces (`{` and `}`). The initial configuration is kept in memory and the remaining configurations are compared against the initial one. Bitstreams should not be generated for any configurations if any pair of configurations do not pass the PR Verify check.

```
pr_verify [-full_check] [-file <arg>] [-initial <arg>] [-additional <arg>]
          [-quiet] [-verbose] [<file1>] [<file2>]
```

Table 3-8: `pr_verify` Options

Command Option	Description
<code>-full_check</code>	Default behavior is to report the first difference only; if this option is selected, <code>pr_verify</code> reports all differences in placement or routing.
<code>-file</code>	Filename to output results to. Send output to console if <code>-file</code> is not used.
<code>-initial</code>	Select one routed design checkpoint against which all others will be compared.
<code>-additional</code>	Select one or more routed design checkpoints to compare against the initial one. List multiple checkpoints within braces, separated by a space, as in this example: <code>{config2.dcp config3.dcp config4.dcp}</code>
<code>-quiet</code>	Ignore command errors.
<code>-verbose</code>	Suspend message limits during command execution.

The following is a sample command line comparing two configurations:

```
pr_verify -full_check config1_routed.dcp config2_routed.dcp -file pr_verify_c1_c2.log
```

The following is a second example verifying three configurations:

```
pr_verify -full_check -initial config1.dcp -additional {config2.dcp config3.dcp} -file
three_config.log
```

The scripts provided with the *Vivado Design Suite Tutorial: Dynamic Function eXchange* (UG947) [Ref 1] have a Tcl Proc called `verify_configs` that automatically runs all existing configurations through `pr_verify`, and reports if the DCPs are compatible or not.

Bitstream Generation

As in a flat flow, bitstreams are created with the `write_bitstream` command. For each design configuration, simply issue `write_bitstream` to create a full standard configuration file plus one partial bit file for each Reconfigurable Module within that configuration.

Xilinx recommends providing the configuration name and Reconfigurable Module names in the `-file` option specified with `write_bitstream`. Only the base bit file name can be modified, so it is important to record which Reconfigurable Modules were selected for each configuration.

Using the previous design, the following is an example of reading routed checkpoints (configurations) and creating bitstreams for all implemented Reconfigurable Modules.

```
open_checkpoint config1_routed.dcp
write_bitstream config1
```

This command generates all possible bitstreams for this particular configuration. It creates a full design bitstream called `config1.bit`. This bitstream should be used to program the device from power-up and includes the functionality of any Reconfigurable Modules contained within. It also creates partial bit files `config1_pblock_rp1_partial.bit` and `config1_pblock_rp2_partial.bit` that can be used to reconfigure these modules while the FPGA continues to operate. For UltraScale devices, it creates clearing bitstreams that pair with each partial bitstream, allowing you to prepare the partition for the next partial image. Repeat these steps for each configuration.



TIP: Rename each partial bit file to match the Reconfigurable Module instance from which it was built to uniquely identify these modules. The current solution names partial bit files only on the configuration base name and Pblock name: `<base_name>_<pblock_name>_partial.bit`

The size of each partial bitstream is reported in the output from `write_bitstream`. As this command is run, these messages will be reported for each partial and clearing bit file.

```
Creating bitmap...
Creating bitstream...
Partial bitstream contains 3441952 bits.
Writing bitstream ./Bitstreams/right_up_pblock_inst_shift_partial.bit...
```

Bitstream compression, encryption, and other advanced features can be used. See [Known Limitations, page 146](#) for specific unsupported use cases for UltraScale devices.

Generating Partial Bitstreams Only

If the full design configuration file is not required, then a single partial bitstream can be created on its own. With a full design configuration checkpoint loaded in memory, use the `-cell` option to identify the instance for which a partial bitstream is needed. The name of this partial bitstream can be given, as it is not automatically derived from the Pblock name.

```
write_bitstream -cell rp1 RM_count_down_partial.bit
```

This creates *only* a partial bitstream for the Reconfigurable Partition identified.



CAUTION! Do not run `write_bitstream` directly on Reconfigurable Module checkpoints; only use full design checkpoints. Reconfigurable module checkpoints, while they are placed and routed submodules, have no information regarding the top level design implementation, and therefore would create unsuitable partial bit files.

Generating Full Configuration Bitstreams Only

If only power-on design bitstream is desired, the `-no_partial_bitfile` option can be used to bypass creation of partial bitstreams.

```
write_bitstream -no_partial_bitfile config3
```

Using this option skips the stage that creates partial and clearing bitstreams. It saves `write_bitstream` runtime for scenarios where either you are looking to test only the full design without DFX, or if the partial bitstreams already exist.

Generating Static-only Bitstreams

If a power-on configuration of the static design only is desired, run `write_bitstream` on the checkpoint that has empty Reconfigurable Partitions (after `update_design -black_box` and `update_design -buffer_ports` have run). This "grey box configuration" can be compressed to reduce the bit file size and configuration time.

Tcl Scripts

Scripts are provided to run this flow in the *Vivado Design Suite Tutorial: Dynamic Function eXchange* (UG947) [Ref 1]. The details of these sample scripts are documented in the tutorial itself and in the `readme.txt` contained in the sample design archive.

Vivado Project Flow

Overview

Dynamic Function eXchange (DFX) in Xilinx® FPGAs and SoCs introduces new design requirements compared to traditional solutions. These requirements include unique approaches to source and runs management, as both bottom-up synthesis and multi-pass implementation are needed. These needs are met with the Vivado® Design Suite DFX Project Flow.

Flow Summary

The Dynamic Function eXchange Project Flow inserts the key requirements of Dynamic Function eXchange into the existing Vivado project solution, accessible within the Vivado IDE as well as via Tcl commands. These key requirements include:

- Defining Reconfigurable Partitions within the design hierarchy
- Populating a set of Reconfigurable Modules for each Reconfigurable Partition
- Creating a set of top-level and module-level synthesis runs
- Creating a set of related implementation runs
- Managing dependencies as sources, constraints or options are modified
- Checking rules and results
- Verifying configurations
- Generating compatible sets of full and partial bitstreams

These fundamental aspects are implemented for this release, featuring support for a front-to-back implementation for RTL-based designs including IP. Partial Reconfiguration (PR) terminology remain for this release but will be transitioned to DFX in a future release. Underlying Tcl commands will remain unchanged so that existing projects and scripts will safely migrate forward.

One expectation of this flow is that all sources (at least the top level RTL and post-synthesis netlists for sub-modules) are managed within a single DFX-enabled project. A project

cannot be broken up or exported as Vivado would no longer be able to track dependencies between runs and sources.

Tcl Commands

Like with most everything within the Vivado IDE, the features and tasks for Dynamic Function eXchange you see are driven behind the scenes by Tcl commands. One of the key goals for DFX project support is to be able to work seamlessly between GUI and script and command line on the same project. You can examine the specific Tcl commands called by examining the Vivado journal file for this project. This can be seen by selecting **File > Project > Open Journal File**. These Tcl commands are not currently documented in this user guide. The full set of commands used to create the entire project up to its current state can be generated by selecting **File > Project > Write Tcl**. Additional information for each command can be found using the `-help` option of each command.

Steps for Creating and Using a Dynamic Function eXchange Project

This section describes the general flow as well as the unique features and capabilities within the Vivado IDE for Dynamic Function eXchange design flows. For a front-to-back tutorial using a specific design that targets a Xilinx Evaluation Platform, please see [Lab 3](#) in the *Vivado Design Suite Tutorial: Dynamic Function eXchange* (UG947) [Ref 1].

Creating a Dynamic Function eXchange Project

The initial creation of a DFX project is no different than for a standard design flow. Step through the New Project wizard to select the target device, design sources and constraints, and set all the main project details. When creating a new project, all source files and constraints for the static portion of the design should be added. You have the option of including the RTL and IP design sources for the first Reconfigurable Module for each Reconfigurable Partition, or you can leave these as black boxes for now.

Note: Only add sources for one RM during the initial project creation. The Partial Reconfiguration wizard is used to add additional RMs to the project. This is discussed in more detail later in this chapter.

Once the project has been created, define it to be a Dynamic Function eXchange project. This is done by selecting **Tools > Enable Partial Reconfiguration**. This prepares the project for the DFX design flow. Once this is set it cannot be undone, so Xilinx recommends archiving your project prior to selecting this option.

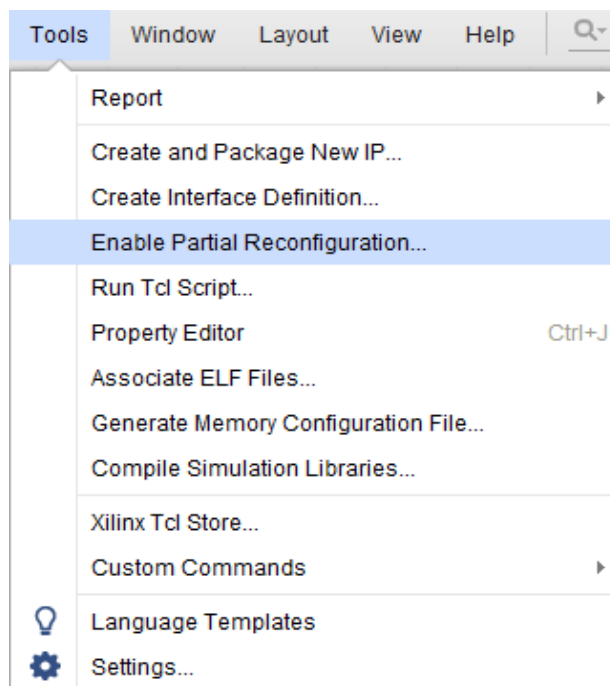


Figure 4-1: Enabling Dynamic Function eXchange

A subsequent dialog box will ask you to confirm this one-way project transition. When this is done, the project will show a few DFX-specific menu options and window tabs. These include:

- A link to the Partial Reconfiguration Wizard in the Flow Navigator
- The Partition Definitions view in the Sources window
- The Configurations window

Defining Reconfigurable Partitions

Once the project has been turned into a DFX project, Reconfigurable Partitions can be defined within the RTL source hierarchy. Suitable instances within the design hierarchy are those that:

- Are defined by RTL, DCP or EDIF sources
- Do not pass parameter and generic values to that level of hierarchy from above. Parameters and generics can exist on the Reconfigurable Partition boundary but must be evaluated locally prior to partition creation.
- Do not contain out-of-context (other than IP) or EDIF modules in the underlying RTL
- Do not have IP as the top level
- Do not contain block diagram (.bd) sources

1. Right-click on the desired module and select **Create Partition Definition** to begin the process of Reconfigurable Partition creation. This module can be a black box if design sources do not yet exist.

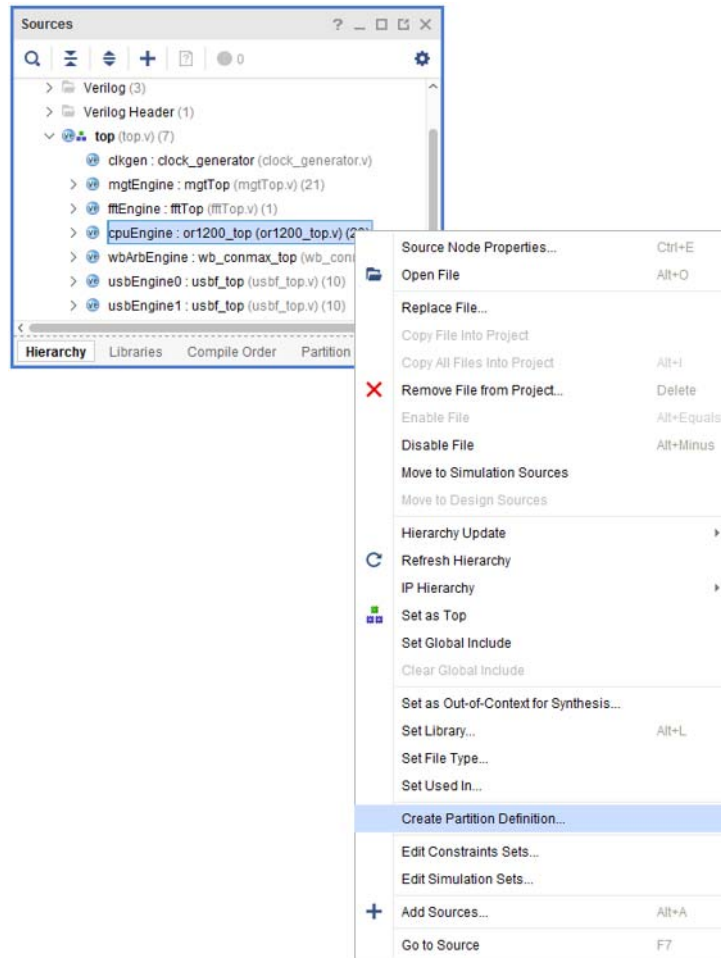


Figure 4-2: Creating a Reconfigurable Partition

2. In the subsequent dialog that appears, give this Partition Definition a unique name. Also define a name for the first Reconfigurable Module (RM). This RM is created from the RTL or netlist sources currently residing in this level of hierarchy; more RM are added or created later in the flow.



IMPORTANT: Every instance of the selected module will be turned into a Reconfigurable Partition. In order for one instance to be defined as reconfigurable and another instance to remain static, the two instances must be given unique module names.




Figure 4-3: Defining the Reconfigurable Partition and the first Reconfigurable Module

- After selecting **OK**, this module displays differently in the Vivado IDE. Each instance of the module is shown in the Hierarchy view with a diamond, indicating that it is a Reconfigurable Partition. The design sources are moved to the Partition Definitions view to be managed separately. Repeat this step for all unique Reconfigurable Partitions required within the design.




Figure 4-4: Partition Definitions as seen in the Sources Window

Completing the Dynamic Function eXchange Project Structure

After defining the Reconfigurable Partitions, enter the full details of the project. This consists of adding more Reconfigurable Modules for each of the Reconfigurable Partitions, defining a full set of Configurations that combine RMs with the static design, and declaring the set of runs that will be used to implement all the Configurations. All of these additions are done within the Partial Reconfiguration Wizard. Any further modifications or additions can be made by returning to the wizard.



TIP: No actions requested in the Partial Reconfiguration Wizard will take effect until the Finish button is clicked. You may step forward and back within the wizard until everything is completed to your liking, and you may cancel and throw away all edits at any time.

Open the Partial Reconfiguration Wizard by selecting the step in the Flow Navigator or from the Tools menu.




Figure 4-5: The Partial Reconfiguration Wizard in the Flow Navigator

When the PR Wizard opens, step through each stage of DFX project management.

Editing Reconfigurable Modules

1. After selecting **Next** to progress past the introduction, the first content page allows you to define new Reconfigurable Modules for any Partition Definitions defined. The first RM for each PD has already been included here, if the RTL/netlist source was present when the PD was created. Click on the blue + to create a new RM and give it a unique name. Be sure to select the correct Partition Definition if more than one exists in the design. If netlist sources are selected, select the **Sources are already synthesized** check box and declare the Top Module within the netlist.




Figure 4-6: Creating a new Reconfigurable Module

- Repeat this process for all existing RMs for every Partition Definition. If a greybox module is desired, no action is required. Note that RMs may be edited by clicking on the pencil icon or removed by clicking the red - icon. When all RM are accounted for click **Next**.

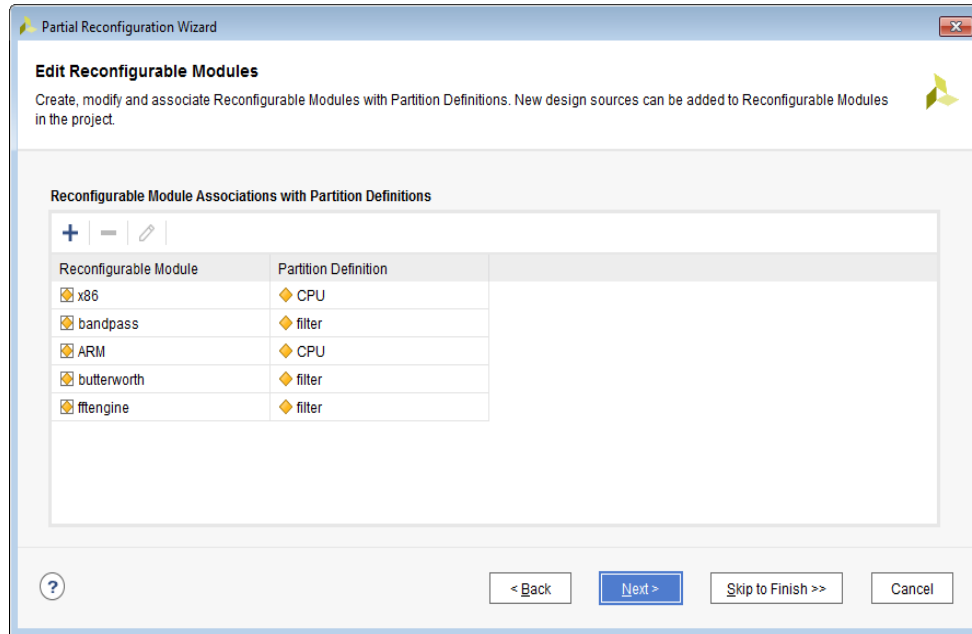


Figure 4-7: Set of Reconfigurable Modules defined

Editing Configurations

With a set of Reconfigurable Modules defined, Configurations can be declared. Each Configuration is a combination of the static logic plus one RM per RP; each Configuration is a full design image.

While each Configuration can be created manually, the simplest path is to let Vivado create the minimum set of Configurations automatically. This is done by selecting the **automatically create configurations** link in the middle of this screen. This will create as many Configurations as necessary to ensure that all RMs are included at least once. This option is only available if no Configurations have been defined yet.

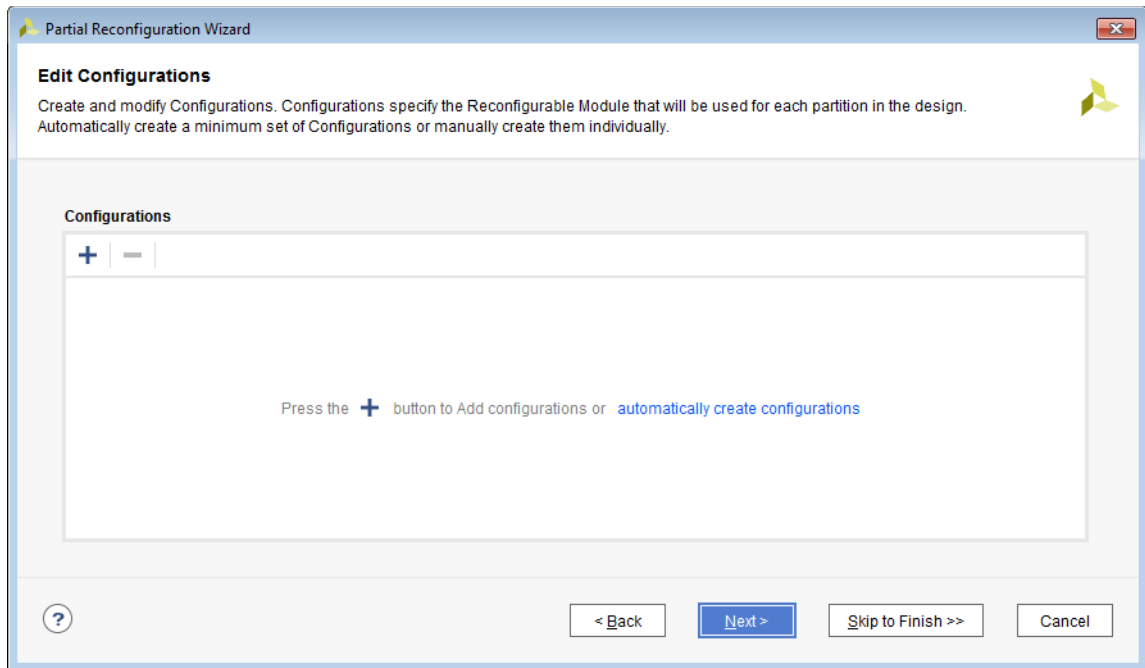


Figure 4-8: Edit Configurations before creating any Configurations

If one Partition Definition has more RMs than another, a greybox RM will automatically be used for any RP that has all its RMs covered by prior Configurations. These default Configurations can be modified or renamed, and additional Configurations may be created if desired.



TIP: Greybox modules are different than black box modules, as they are not truly empty. Greybox RMs have tie-off LUTs inserted to complete legal design connectivity in the absence of an RM and they ensure outputs do not float during operation. Vivado creates these by calling `update_design -buffer_ports` on selected modules.

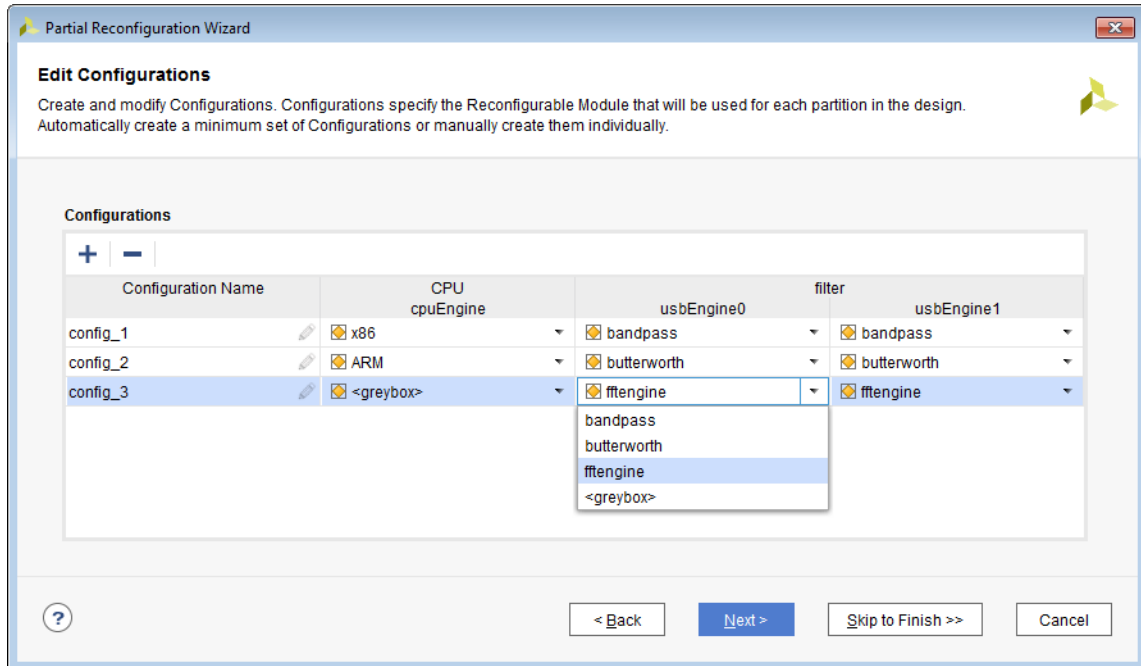


Figure 4-9: Edit Configurations after automatic generation of Configurations

Note that if one RM is used in more than one Configuration, the implementation results may be different, as place and route is performed each time, but only if the RM was initially implemented in a child run. RM implementation results will be reused if they were originally done in the parent configuration. This allows the Vivado project to track dependencies between parent and child.

Editing Configuration Runs

With all the Configurations defined, move to the final screen to manage the Configuration Runs associated with them. Just like the Configurations themselves, Vivado can automatically create a set of Configuration Runs. The first Configuration in the list are defined as the parent, and all remaining Configurations are set as children to that parent.




Figure 4-10: Automatically generated Configuration Runs

This structure assumes that the first configuration is the most critical or challenging. Users are free to change the parent-child relationship by setting that value in the **Parent** column. A Parent of a synthesis run (**synth_1** here) indicates the Configuration (most notably the static part) will be implemented from the synthesized netlist, and a Parent of an implementation run (**impl_1** here) indicates the parent's locked static implementation result will be used as the starting point.

As you explore place and route options, timing closure techniques, and otherwise elaborate on the DFX design, multiple independent parent runs can be used for exploration. Multiple parent runs can be launched in parallel, then child runs can be launched after parent runs complete. Vivado project management handles all the DFX-specific details for creating and storing intermediate checkpoints, including a "static-only" checkpoint for a routed parent run. Ultimately, a single parent run must be selected to establish a golden static implementation result on which all Configurations will be based.



IMPORTANT: *To ensure a safe working environment in silicon, a locked static image must remain consistent across all Configurations so bitstream generation will create compatible full and partial bitstreams. This is managed in the Vivado DFX Project flow by establishing a parent-child relationship for related Configurations*

Add new Configuration Runs by selecting the green + icon. When all Configuration Runs have been created, click **Next**. On the final screen, the number of new elements are listed. Clicking **Finish** will actually perform all the requested changes in the project.

In the Design Runs window, out-of-context synthesis runs are created for each Reconfigurable Module, and all Configuration Runs are generated. Relationships between parent and child runs are shown by the levels of indentation.

Name	Configuration	Constraints	Status	WNS
synth_1 (active)		constrs_2	Not started	
impl_1 (active)	config_1	constrs_2	Not started	
child_0_impl_1	config_2	constrs_2	Not started	
child_1_impl_1	config_3	constrs_2	Not started	
Out-of-Context Module Runs				
x86_synth_1		x86	Not started	
bandpass_synth_1		bandpass	Not started	
ARM_synth_1		ARM	Not started	
butterworth_synth_1		butterworth	Not started	
fftengine_synth_1		fftengine	Not started	

Figure 4-11: Design Runs for synthesis and implementation

In addition, the Configurations window now shows the composition details of each Configuration available in the project.

Name	cpuEngine	usbEngine0	usbEngine1
config_1 (active)	x86	bandpass	bandpass
config_2	ARM	butterworth	butterworth
config_3	<greybox>	fftengine	fftengine

Figure 4-12: Configurations available in the project

Adding or Modifying Reconfigurable Modules or Configurations

The Partial Reconfiguration Wizard is the central mechanism for making any changes to Reconfigurable Modules or Configurations. This includes adding new RMs, modifying source lists for any RM, creating new Configurations or Runs, or removing any of the above. When working within the wizard, nothing is saved or executed until the Finish button is clicked, so you can move forward or back through the screens, making adjustments as needed.

When changes to the RTL sources themselves are needed, they can be seen and opened from the Partition Definitions view in the Source window. This shows each Reconfigurable Module in the same way as the full design is shown in the Hierarchy tab, but scoped to that level of hierarchy and below. This includes all sources and constraints that have been declared for each RM.




Figure 4-13: Sources shown in Partition Definitions view

Adding or Creating IP Sources

IP are permitted within Reconfigurable Modules. They cannot be the top level of an RM, but they can exist within any level below the top. Include existing `.xci` or `.xcix` files along with RTL when adding sources to a Reconfigurable Module.

IP that exist within Reconfigurable Modules can be synthesized either globally or out-of-context. Either value for the Synthesis Options shown below can be selected.

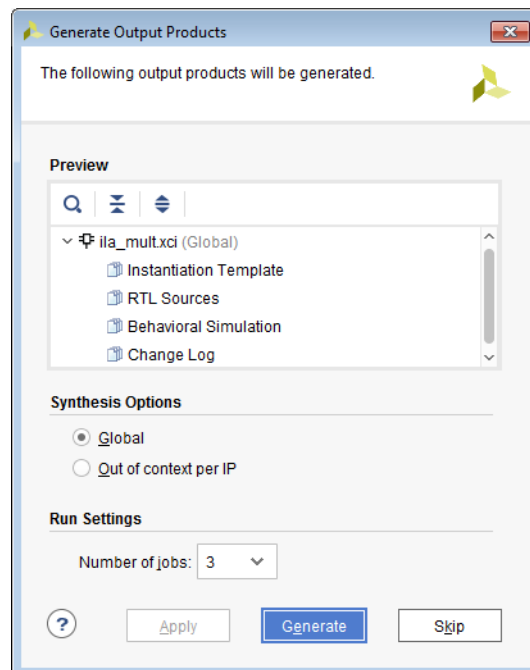


Figure 4-14: Specifying Global when Generating Output Products

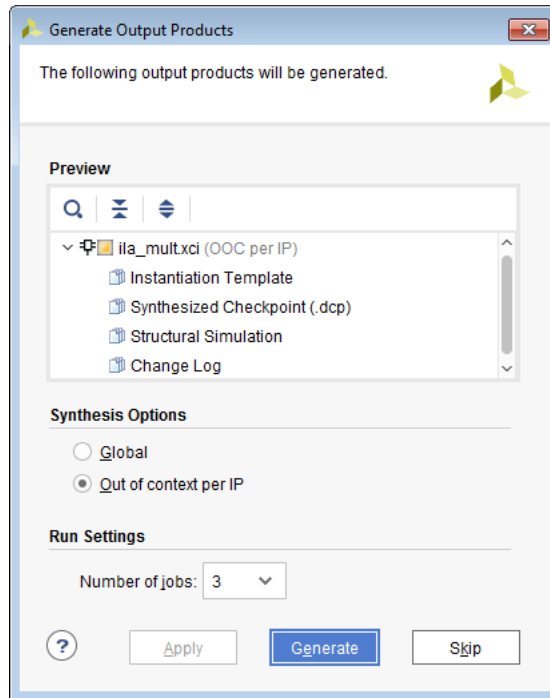


Figure 4-15: Specifying Out of Context per IP when Generating Output Products

IP can be created from the IP Catalog within a DFX project. After the IP has been created, it is added to the primary blockset of the design, as the IP generation flow does not know which RM the IP is for. To assign a new IP to a specific RM, right-click on the IP in the Sources window and select **Move to Reconfigurable Module**. Select the appropriate RM and click **OK**.

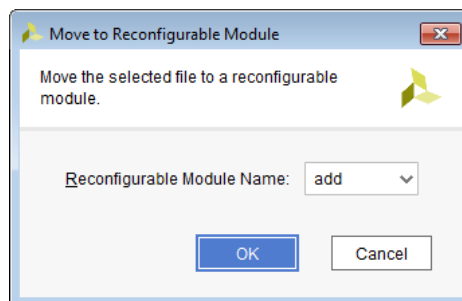


Figure 4-16: Moving IP to a Reconfigurable Module

One final requirement is that IP must be unique per Reconfigurable Module. If two different RMs each contain the same IP function, two unique instances must be created. The best way to do this is to replicate one IP to create a new identical IP. Right-click on an existing IP and select **Copy IP**. Once created, this new IP must be moved to the target Reconfigurable Module as described above.




Figure 4-17: Copying IP

Implementing the DFX Design

With all the necessary Configuration Runs defined, the design can be synthesized and implemented. The Flow Navigator can be used to pull through any steps of synthesis, place and route, and even bitstream generation. The Flow Navigator works on the active run just like a standard flow, but it will launch all child runs in addition to the active parent.

One detail that is required for Dynamic Function eXchange designs is a Pblock for each Reconfigurable Partition. Without a Pblock defined, the following error will be issued in `place_design`:

```
ERROR: [DRC 23-20] Rule violation (HDPR-30) Missing PBLOCK On Reconfigurable Cell
```

If this necessary floorplan is present in a top-level design constraints file, you can pull all the way from synthesis to bitstream generation. If not, the easiest way to create one is to stop and open the design after top-level synthesis. In the Netlist hierarchy view, right-click on the module that corresponds to the Reconfigurable Partition and select **Floorplanning > Draw Pblock**.

After you draw the Pblock for a Reconfigurable Partition, its properties can be seen in the Pblock Properties window under the Properties view. Available here are two options unique to Reconfigurable Partitions: **RESET_AFTER_RECONFIG** (7 series only) and **SNAPPING_MODE**. The Statistics view reports the resources available and used for the currently loaded Reconfigurable Module, so it is important to consider the needs for the other RMs as well.

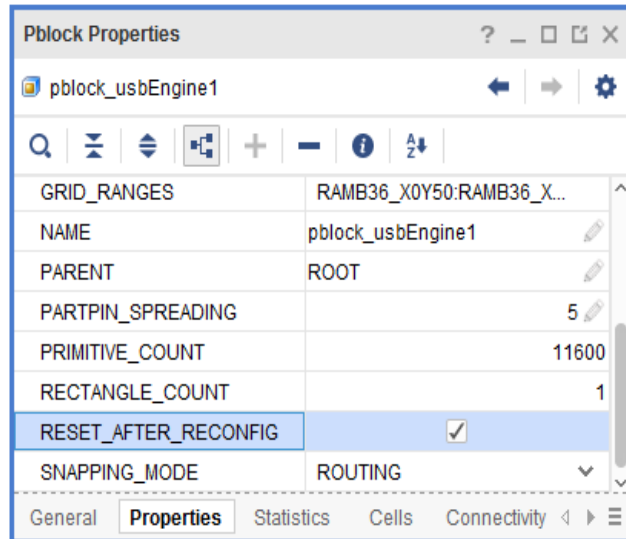


Figure 4-18: Dynamic Function eXchange properties on a Pblock (7 series)

Once a Pblock has been created for each Reconfigurable Partition, each Configuration can be implemented. The **Run Implementation** button in the Flow Navigator launches place and route on the active parent run first. Upon completion, all child runs will be launched in parallel, using the static design results of the parent as a starting point.

The Vivado project takes care of the underlying details of the DFX solution. Database management is one of these details. Upon completion of the parent run, the routed database for the entire design is saved, as well as a cell-level checkpoint for each Reconfigurable Module. Then Vivado calls `update_design -black_box` to carve out each RM, resulting in a static-only design checkpoint, which is the basis for all of its child runs. When child implementations runs are launched, Vivado assembles the configuration from the static-only routed parent checkpoint and post-synthesis checkpoints for each Reconfigurable Module. At this time, only routed module checkpoints from the parent run can be reused in child Configurations; if the same RM is selected for one RP in multiple child runs, the results will be different.




Figure 4-19: Implementing multiple Configurations in parallel

Just as with a standard project, Vivado tracks dependencies between runs. When design sources, constraints, options or settings are modified, any synthesis or implementation run that depends on them are marked out-of-date. One example: if an RTL design source for one RM is updated, that out-of-context module run will be marked out-of-date, and any Configuration Runs that include that RM will also be marked out-of-date. Another example: if any implementation option for a parent Configuration Run is changed, it and all child runs will be marked out-of-date.

Only parent Configuration Runs can be set as active. The Flow Navigator acts upon the active run, but in the DFX flow, all child runs are also included in whatever action is requested. Pop-up messages (completed run, error, etc.) can relate to multiple runs but default to the parent run. Use the pull down selection to choose the desired implementation run.

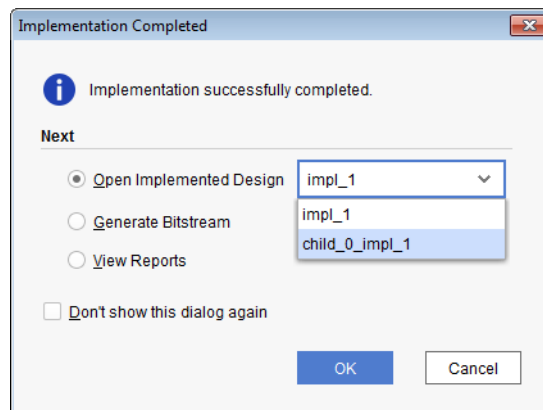


Figure 4-20: Implementation Completed dialog box

Everything shown in the different windows, as shown in Figure 4-19, relate to the active parent run. In order to see details about a child implementation run, select that run and look at the Implementation Run Properties window to see all the inputs (properties, options) and outputs (log, reports, messages) for that specific run.

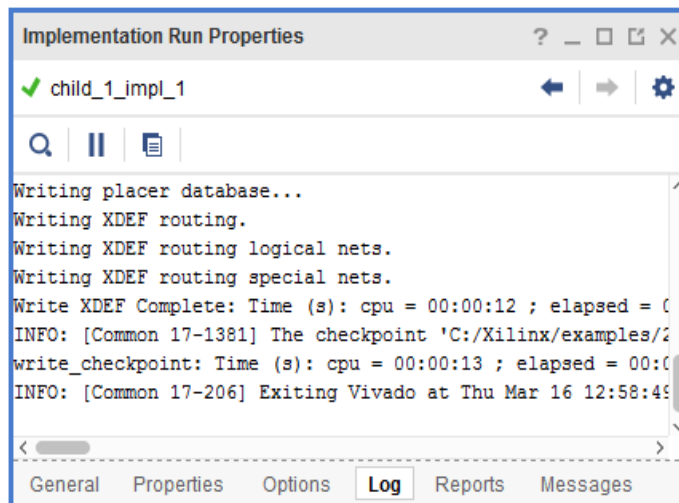


Figure 4-21: Implementation Run Properties window for a child run

Generating Bitstreams

Once all desired Configurations have been placed and routed, bitstreams may be generated. Just as with Implementation, the **Generate Bitstream** button in the Flow Navigator may be utilized. This will launch `write_bitstream` for all child runs as well as the active parent. A local right-click call to `write_bitstream` on any Configuration is also available.

The `pr_verify` utility is automatically called prior to `write_bitstream` on each child Configuration run. This routed database is compared to the parent database to ensure all DFX rules have been met. The results of this check are stored in the run directory under the name `<impl_name>_pr_verify.log`.

By default, a full design bitstream and all partial (and for UltraScale, clearing) bitstreams are generated for all routed configurations. You can request specific bitstreams only by utilizing the `write_bitstreams` options available. Under the Write Bitstream options category in the Options pane for the Implementation Run Properties, use the **More Options** field to select one of these options:

- `-no_partial_bitfile` generates only the full configuration file and no partial bitstreams.
- `-cell <cell>` generates only the partial bitstream for the requested cell.

Supported/Unsupported Features

This section lists the current lists of supported and unsupported features for DFX Projects.

Supported Features

- Device support: All 7 series, Zynq, UltraScale and UltraScale+ devices supported by the Dynamic Function eXchange flow.
- Source types for Reconfigurable Modules: RTL, DCP, EDIF, XDC, XCI, XCIX.
 - XCI or XCIX (Xilinx IP) cannot be the top level.
 - EDIF cannot be a sub-module.
- Module-level constraints must be scoped to the hierarchical instance.
- Greybox (black box module with LUT tie-off) implementation can be done
- An extensive set of Design Rule Checks can be issued from within the project environment.
- All synthesis and implementation design switches can be used.
- PR Verify is automatically called prior to bitstream generation for any child configuration.

Unsupported Features

The following features are not currently implemented:

- IP Integrator support is not in place. Block Diagrams cannot be included as RMs or within RMs. Modules within Block Diagrams cannot be set as RMs.
- Simulation is not supported from within the project.

Known Limitations

- Once Partitions are defined, they cannot be undone. The only way to return to a flat non-PR project is to create a new one.
- Reuse of implemented Reconfigurable Modules from a child run is not supported. Only implementation results from RMs from the parent run can be reused in a child run.
- Child implementation runs cannot be set active. Flow Navigator actions work on just the parent run, or the parent and all child runs, depending on the action.

Design Considerations and Guidelines for All Xilinx Devices

Overview

This chapter explains design requirements that are unique to Dynamic Function eXchange (DFX), and covers specific DFX features within the Xilinx® design software tools.

To take advantage of the dynamic reconfiguration capability of Xilinx FPGAs, you must analyze the design specification thoroughly, and consider the requirements, characteristics, and limitations associated with PR designs. This simplifies both the design and debug processes, and avoids potential future risks of malfunction in the design.

This chapter describes the design requirements that apply to all Xilinx 7 series and UltraScale™ devices. For design requirements specific to the individual FPGA and SoC architectures, see the following chapters in this manual:

- [Chapter 6, Design Considerations and Guidelines for 7 Series and Zynq Devices](#)
- [Chapter 7, Design Considerations and Guidelines for UltraScale and UltraScale+ Devices](#)

Partial Reconfiguration IP

Xilinx has created four pieces of intellectual property specifically for the use within Partial Reconfiguration designs. There is no charge for any of these IP, and DFX designs do not require them. They are available to assist users in quickly and easily implementing key aspects of a reconfigurable design. The IP are all found under the Partial Reconfiguration heading within the IP Catalog, and each have their own landing page on Xilinx.com with a detailed product guide.

These four IP for Dynamic Function eXchange are currently available in Vivado® and can be used for any Xilinx device that supports Dynamic Function eXchange in Vivado. DFX terminology will be introduced in a future Vivado release, and users will be given instructions on how to migrate these IP to use the new names.

- **Partial Reconfiguration Controller.** The PR Controller core provides management functions for self-controlling partially reconfigurable designs. It is intended for enclosed systems where all of the Reconfigurable Modules (RM) are known to the controller. The optional AXI4-Lite register interface allows the core to be reconfigured at run time, so it can also be used in systems where the RMs can change in the field. The core can be customized for many Virtual Sockets, RMs per Virtual Socket, operations and interfaces. Labs 5, 6, and 7 in *Vivado Design Suite Tutorial: Dynamic Function eXchange* (UG947) [Ref 1] show examples of the PR Controller IP in a sample design.
- **Partial Reconfiguration Decoupler.** The PR Decoupler can be used to provide a safe and managed boundary between the static logic and a Reconfigurable Partition during reconfiguration. The core can be customized for the number of interfaces, type of interfaces, decoupling functionality, status and control.
- **Partial Reconfiguration AXI Shutdown Manager.** One or more PR AXI Shutdown Managers can be used to make the AXI interfaces between a Reconfigurable Partition and the static logic safe during reconfiguration. When active, AXI transactions sent to the RM, and AXI transactions emanating from the RM, are terminated because the RM might not be able to complete them. Failure to complete could cause system deadlock. When inactive, transactions pass unaltered.
- **Partial Reconfiguration Bitstream Monitor.** The PR Bitstream Monitor can be used to identify partial bitstreams as they flow through the design. This information can be used for debugging or system applications such as blocking bitstream loads. Identifiers embedded at key places in partial bitstreams are extracted and reported by the core. This information can be passed to Vivado HW Debugger using an ILA core to work out what partial bitstream was fetched, if it was fetched in its entirety, and how far through the datapath it went.

Design Hierarchy

Good hierarchical design practices resolve many complexities and difficulties when implementing a partially reconfigurable FPGA design. A clear design instance hierarchy simplifies physical and timing constraints. Registering signals at the boundary between static and reconfigurable logic eases timing closure. Grouping logic that is packed together in the same hierarchical level is necessary.

These are all well known design practices that are often not followed in general FPGA designs. Following these design rules is not strictly required in a partially reconfigurable design, but the potential negative effects of not following them are more pronounced. The benefits of Dynamic Function eXchange are great, but the extra complexity in design could be more challenging to debug, especially in hardware.

For additional information about design hierarchy, see [Appendix A, Hierarchical Design Flows](#).

Dynamic Reconfiguration Using the DRP

Logic that is in the static region, and therefore is never partially reconfigured, can still be reconfigured dynamically through the Dynamic Reconfiguration Port (DRP). The DRP can be used to configure logic elements such as MMCMs, PLLs, and serial transceivers (MGTs).

Information about the DRP and dynamic reconfiguration, including how to use the DRP for specific design resources, can be found in these documents:

- *7 Series FPGAs Configuration User Guide* (UG470) [\[Ref 9\]](#)
- *7 Series FPGAs GTX/GTH Transceivers User Guide* (UG476) [\[Ref 21\]](#)
- *7 Series FPGAs GTP Transceivers User Guide* (UG482) [\[Ref 22\]](#)
- *MMCM and PLL Dynamic Reconfiguration (7 Series)* (XAPP888) [\[Ref 23\]](#)
- *UltraScale Architecture Configuration User Guide* (UG570) [\[Ref 10\]](#)
- *UltraScale Architecture Clocking Resources User Guide* (UG572) [\[Ref 24\]](#)
- *UltraScale Architecture GTH Transceivers User Guide* (UG576) [\[Ref 25\]](#)
- *UltraScale Architecture GTY Transceivers User Guide* (UG578) [\[Ref 26\]](#)

Packing Logic

Any logic that must be packed together must be placed in the same group, whether it is static or reconfigurable. For example, if a LUT and a flip-flop are expected to be placed within the same slice, they must be within the same partition. Partition boundaries are barriers to optimization.

For Reconfigurable Partitions that include I/O, Clocking, and GT resources, it might be necessary to instantiate any I/O buffers that are automatically inferred by the tools inside that RP level. For example, if `GT_COMMON` is in an RP, the `IBUFDS_GTE` needs to be instantiated. If the associated I/O buffer is in the top-level/static portion, it cannot be packed.

Design Instance Hierarchy

The most simple method is to instantiate the Reconfigurable Partitions in the top-level module, but this is not required because a Reconfigurable Partition may be located in any level of hierarchy. Each Reconfigurable Partition must correspond to exactly one instance—an RP must not have more than one top. The instantiation has multiple modules with which it is associated.

Changes in design hierarchy can be used to merge and/or separate modules and leaf cells into and out of an RP level of hierarchy. There are several reasons to do this:

- To balance device resources between the dynamic region and static region, making the design more efficient. For example, if the target RP takes up most of the device, and there is a module in the static region that requires a high number of Block RAMs unavailable to Static, you can move that module into the dynamic region.
- If you need cells to reside in the same physical area of the device, but they are in a different design hierarchy. For example, if you need `GT_CHANNELS` to be placed in the same UltraScale Clock Region, but the design has GTs in both the Static and RP regions.
- To ensure that dedicated connections, for example from `IBUFDS_GT` to `GT_COMMON`, reside in the same region.

Reconfigurable Partition Interfaces

One of the fundamental requirements of a partially reconfigurable design is consistency between Reconfigurable Modules (RMs). As one module is swapped for another, the connections between the static design and the Reconfigurable Module must be identical, both logically and physically. To achieve this consistency, optimizations across the partition boundary or of the boundary itself are prohibited.



RECOMMENDED: *Keep interface logic connecting to and from RM ports consistent across all RMs. Do not change the levels of logic between RMs, such as using one level of logic in the initial design and five levels of logic in next RM. Also, do not change the driver type, such as using flip-flops in the initial RM and block RAM in next RM. Since the static side of the interface is locked after the initial configuration, the tools are unable to adjust for these changes in later configurations. Xilinx recommends registering all inputs and outputs of the RMs.*

For optimal efficiency, all ports of a Reconfigurable Partition should be actively used on the static design side. For example, if static drivers of the Reconfigurable Partition are driven by constants (0 or 1), they are implemented through the creation of a LUT instance and local tie-off to a constant driver and cannot be trimmed away. Likewise, unconnected outputs remain on Reconfigurable Partition outputs, creating unnecessary waste in the overall design. These measures must be taken by the implementation tools to ensure that all Reconfigurable Modules have the same port map during design assembly.



RECOMMENDED: *Examine the interface of all Reconfigurable Partitions after synthesis to ensure that as few constants or unconnected ports as possible remain. By clearing out dead logic, resource utilization is reduced, and congestion and timing closure challenges easier to address.*

Six different cases are possible for partition interface usage:

1. **Both Static and Reconfigurable Module sides have active logic.** (Applies to partition inputs or outputs)

This is the optimal situation. A partition pin is inserted.



RECOMMENDED: *If partition inputs are driven by VCC or GND, push these constants into the Reconfigurable Modules. This reduces LUT usage and allow the implementation tools to optimize these constants with the RM logic.*

2. **The Static side has an active driver but the Reconfigurable Module does not have active loads.** (Applies to partition inputs)

This is acceptable because it accommodates the situation in which not every Reconfigurable Module has the same I/O requirements. A partition pin is inserted, and the unused input ports are left unconnected.

For example, one module might require CLK_A, while a second might require CLK_B. Clock spines are pre-routed to the Reconfigurable Partition clock regions, but the module only taps into the clock source that is needed. However, if a partition input is not used by any Reconfigurable Module, it should be removed from the partition instantiation.

3. **The Static side has active loads but the Reconfigurable Module does not have an active driver.** (Applies to partition outputs)

This is acceptable and similar to the case above. A partition pin is inserted, and it is driven by ground (logic 0) within the Reconfigurable Module.

4. **The Static side does not have an active driver, but the Reconfigurable Module has active loads.** (Applies to partition inputs)

This results in an error that must be resolved by modifying the partition interface.

The following is an example of an error that may be seen for this scenario:

```
ERROR: [Opt 31-65] LUT input is undriven either due to a missing connection from  
a design error, or a connection removed during opt_design.
```

This error message would be followed by a LUT instance that is within the Reconfigurable Module.

5. **Reconfigurable Module has an active driver, but the Static side has no active loads.** (Applies to partition outputs)

This does not result in an error, but is far from optimal because the RM logic remains. No partition pin is inserted. These partition outputs should be removed.

6. **Neither Static nor Reconfigurable Module sides have driver or loads for a partition port.** (Applies to partition inputs or outputs.)

Nothing is inserted or used, so there is no implementation inefficiency, but it is unnecessary in terms of the instantiation port list.

Partition Pin Placement

Each pin of an RP has a partition pin (PartPin). By default the tools automatically place these PartPins inside of the RP Pblock range (which is required). For many cases, this automatic placement can be sufficient for the design. However, for timing-critical interface signals or designs with high congestion, it might be necessary to help guide the placement of the PartPins. The following is an example of how to achieve this:

- Define user `HD.PARTPIN_RANGE` constraints for some or all of the pins.

```
set_property HD.PARTPIN_RANGE {SLICE_Xx0Yx0:SLICE_Xx1Yy1  
SLICE_XxNYyN:SLICE_XxMYyM} [get_pins <rp_cell_name>/*]
```

By default the `HD.PARTPIN_RANGE` is set to the entire Pblock range. Defining a user range allows the tools to place PartPins in the specified areas, improving timing and/or reducing congestion.



IMPORTANT: *When examining the placement of PartPins, there are limited routing resources available along the edges, and especially in the corners, of the Pblock. The PartPin placer attempts to spread the partition pins, minimizing the number of partition pins per interconnect along the edges, and increasing the PartPin density towards the middle of the Pblock. When defining a custom `HD.PARTPIN_RANGE` constraint, be sure to make the range wide enough to allow for spreading, or you are likely to see congestion around the PartPins.*

Active-Low Resets and Clock Enables

In Xilinx 7 series FPGAs, there are no local inverters on control signals (resets or clock enables). The following description uses a reset as the example, but the same applies for clock enables.

If a design uses an active-Low reset, a LUT must be used to invert the signal. In non-DFX designs that use all active-Low resets multiple LUTs are inferred but can be combined into a single LUT and pushed into the I/O elements (the LUT goes away). In non-DFX designs that use a mix of High and Low, the LUT inverters can be combined into one LUT that remains in the design, but that has minimal effect on routing and the timing of the reset net (output of LUT can still be put on global resources). However, for a design that uses active-Low resets on a partition, it is possible to have inverters inferred inside the partition

that cannot be pulled out and combined. This makes it impossible to put the reset on global resources, and can lead to poor reset timing and to routing issues if the design is already congested.

The best way to avoid this is to avoid using active-Low control signals. However, there are cases where this is not possible (for example, when using an IP core with an Advanced eXtensible Interface (AXI) interface). In these cases the design should assign the active-Low reset to a signal at the top level, and use that new signal everywhere in the design.

As an example:

```
reset_n <= !reset;
```

Use the `reset_n` signal for all cases, and do not use the `!reset` assignments on signals or ports.

This ensures that a LUT is inferred only for the reset net for the whole design and has a minimal effect on design performance.

Decoupling Functionality

Because the reconfigurable logic is modified while the device is operating, the static logic connected to outputs of Reconfigurable Modules must ignore data from Reconfigurable Modules during partial reconfiguration. The Reconfigurable Modules do not output valid data until partial reconfiguration is complete and the reconfigured logic is reset. There is no way to predict or simulate the functionality of the reconfiguring module.

You must decide how the decoupling strategy is solved. A common design practice to mitigate this issue is to register all output signals (on the static side of the interface) from the Reconfigurable Module. An enable signal can be used to isolate the logic until it is completely reconfigured. Other approaches range from a simple 2-to-1 MUX on each output port, to higher level bus controller functions.

The static design should include the logic required for the data and interface management. It can implement mechanisms such as handshaking or disabling interfaces (which might be required for bus structures to avoid invalid transactions). It is also useful to consider the down-time performance effect of a PR module (that is, the unavailability of any shared resources included in a PR module during or after reconfiguration).

A Partial Reconfiguration Decoupler IP is available, allowing users to insert MUXes to efficiently decouple AXI4-Lite, AXI4-Stream, and custom interfaces. More information about the [PR Decoupler IP](#) is available on the Xilinx website.

Black Boxes

You can implement an RP as a *pseudo* black box, referred to in Vivado as a greybox. To do this, the RP must be a black box in the static design (either from bottom-up synthesis results or from running `update_design -black_box`). Then the black box can have LUT1 buffers placed on all inputs and outputs using the command `update_design -buffer_ports` on the black box RP cell:

```
update_design -cell <rp_cellName> -buffer_ports
```

Now you can run this design through implementation to place and route the LUT1 buffers (and static logic, if not already placed and routed).

All the inserted LUT1 output buffers are tied to a logic 0 (ground). If it is necessary to drive a logic 1 (V_{CC}) from the RP outputs, this can be controlled using an RP pin property called `HD.PARTPIN_TIEOFF`. This property can be set at any time (all the way up to pre-write_bitstream), and it controls the LUT equation of the LUT1 buffer connected to the specified port. The default value is '0', which configures the LUT as a route-thru (output is 0). Setting this property to '1' configures the LUT as an inverter (output is 1). You might have to change the output value in some design situations.

```
set_property HD.PARTPIN_TIEOFF 1 [get_pins <RP_cellName>/<output_pinName>]
```

The greybox has no user logic (just the tool-inserted LUT1 buffers). The greybox bitstream contains information for these LUTs, as well as any static logic/routes that use resources inside the RP frames. Static routes that pass through the region, including interface nets up to the partition pin nodes, exist within this region. Programming information for these signals is included in the black box programming bitstream.

Use of greyboxes is an effective way to reduce the size of a full configuration BIT file, and therefore reduce the initial configuration time. The compression feature might also be enabled to reduce the size of BIT files. This option looks for repeated configuration frame structures to reduce the amount of configuration data that must be stored in the BIT file. The compression results in reduced configuration and reconfiguration time. When the compression option is applied to a routed DFX design, all of the BIT files (full and partial) are created as compressed BIT files. To enable compression, set this property prior to running `write_bitstream`:

```
set_property BITSTREAM.GENERAL.COMPRESS TRUE [current_design]
```

Effective Approaches for Implementation

There are trade-offs associated with optimizing any FPGA design. Dynamic Function eXchange is no different. Partitions are barriers to optimization, and reconfigurable frames require specific layout constraints. These are the additional costs to building a reconfigurable design. The additional overhead for timing and area needs vary from design to design. To minimize the impact, follow the design considerations stated in this guide.

When building Configurations of a reconfigurable design, the first Configuration to be chosen for implementation should be the most challenging one. Be sure that the physical region selected has adequate resources (especially elements such as block RAM and DSP48) for each Reconfigurable Module in each Reconfigurable Partition, then select the most demanding (in terms of either timing or area) RM for each RP. If all of the RMs in the subsequent Configurations are smaller or slower, it is easier to meet their demands. Timing budgets should be established to meet the needs of all Reconfigurable Modules.

If it is not clear which reconfigurable module is the most challenging, each can be implemented in parallel in context with static, allowing static to be placed and routed for each. Examine resource utilization statistics and timing reports to see which configuration met design criteria most easily and which had the tightest tolerances, or which missed by the widest margins.



IMPORTANT: *Focus attention on the configuration that is the furthest from meeting its goals, iterating on design sources, constraints, and strategies until needs are met. At some point, one configuration must be established as the golden result for the static design, and that implementation of the static logic will be used for all other configurations.*

Building Up Implementation Requirements

Implementation of Dynamic Function eXchange designs requires that certain fundamental rules are followed. These rules have been established to ensure that a partial bitstream can be accurately created and safely delivered to an active device. As noted throughout this document, these rules include these basic premises:

- The logical and physical interface of a Reconfigurable Partition remains consistent as each Reconfigurable Module is implemented.
- The logic and routing of a Reconfigurable Module is fully contained within a physical region which is then translated into a partial bitstream.
- The logic of the static design must be kept out of the reconfigurable region if the dedicated initialization feature is used.

These requirements necessitate specific implementation rules for optimization, placement and routing. Application of these rules might make it more difficult to meet design goals, including timing closure. A recommended strategy is to build up this set of requirements one at a time, allowing you to analyze the results at each step. Starting with the most challenging configuration and the full set of timing constraints, implement the design through place and route and examine the results, making sure you have sufficient timing slack and resources available to continue to the next step.

1. Implement the design with no Pblocks. Use bottom-up synthesis and follow general Hierarchical Design recommendations, such as registered boundaries, to achieve a baseline result.
2. Add Pblocks for the design partitions that will later be marked reconfigurable. This floorplan can be based on the results established in the bottom-up synthesis run from Step 1. Logic from the Reconfigurable Modules must be placed in the Pblocks, but static logic may be included there as well.

While creating these Pblocks, the **HD.RECONFIGURABLE** property (and optionally, the **RESET_AFTER_RECONFIG** property) can be added temporarily to run PR-specific Design Rule Checks. This ensures that the floorplan created meets PR size and alignment requirements.

3. With the floorplan established, separate the placement of static design resources from those to be reconfigurable by adding the **EXCLUDE_PLACEMENT** property to the Pblocks. This keeps static logic placed outside the defined Pblocks.
4. Keep the routing for Reconfigurable Modules bound within the Pblocks by applying the **CONTAIN_ROUTING** property to the Pblocks. With the properties from this and the previous step, the only remaining rules relate to boundary optimization procedures as well as PR-specific Design Rule Checks.
5. Finally, mark the Reconfigurable Partition Pblocks as **HD.RECONFIGURABLE**. The **EXCLUDE_PLACEMENT** and **CONTAIN_ROUTING** properties are now redundant and can be removed.

If design requirements are not met at any of these steps, you have to opportunity to review the design structure and constraints in light of the newly applied implementation condition.

Configuration Analysis Report

The Dynamic Function eXchange design flow uses multiple versions of the design that must be implemented through place and route. These different configurations have common static design results, but differing modules within each Reconfigurable Partition. Designers must set up timing constraints and floorplans that account for these different modules that will be swapped on the fly.

The DFX Configuration Analysis report compares each Reconfigurable Module that you select to give you information on your DFX design. It examines resource usage, floorplanning, clocking, and timing metrics to help you manage the overall PR design.

The DFX Configuration Analysis report is currently run in the Tcl Console or within a Tcl script. The top level design must be open before issuing this command:

```
report_pr_configuration_analysis -cells <RP_name> -dcps {<list_of_RM_checkpoints>}
```

Either select a single cell (RP) and multiple DCPs (each representing an RM) that can be inserted into that cell for a comprehensive analysis of that RP, or select multiple cells with no subsequent DCPs for a top-level analysis of the static design and interfaces into each RP.

By default, three aspects of the PR design are analyzed. You can select one or more of these switches to narrow the focus of the report.

- The **-complexity** switch focuses the report on resource usage, including the maximum number of each resource type required for the RP.
- The **-clocking** switch focuses the report on clock usage and loads for each RM, helping you plan the overall clocking distribution of the design.
- The **-timing** switch focuses the report on boundary interface timing details, allowing you analyze bottlenecks in and out of RMs.

Additionally, the **-rent** switch adds Rent metrics to the report. The Rent exponent calculates the routing complexity and can be an indication of how much congestion is likely to be seen. For more information on Rent, see this [link](#) in *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* (UG906) [Ref 20]. Note that this option can take a long time to run on large designs.

When this analysis is done, each RM is examined based on information in the checkpoints provided. While post-synthesis checkpoints can be supplied, if the RM contains IP that have been synthesized out-of-context, or if debug cores are to be inserted, information will be missing from these checkpoints. The most complete information is not available until after **opt_design** when all the linking and expansion has been done. We advise you to create fully assembled RM checkpoints after **opt_design** by calling **write_checkpoint -cell** for each configuration, then run the configuration analysis report using these files.

Here are some example sections from a report for a design with a single Reconfigurable Partition that has three Reconfigurable Modules.

Complexity

First is the resource usage table for the `-complexity` switch:

Categories	Grid Type	Current	RM1	RM2	RM3	Max
Slice Logic						
Slice LUTs	SLICE	936 (23.40%)	936 (23.40%)	927 (23.17%)	1091 (27.28%)	1091 (27.28%)
LUT as Logic	SLICE	836 (20.90%)	836 (20.90%)	827 (20.67%)	977 (24.43%)	977 (24.43%)
LUT as Memory	SLICE	100 (5.00%)	100 (5.00%)	100 (5.00%)	114 (5.70%)	114 (5.70%)
LUT as Distributed RAM	SLICE	32 (1.60%)	32 (1.60%)	32 (1.60%)	42 (2.10%)	42 (2.10%)
LUT as Shift Register	SLICE	68 (3.40%)	68 (3.40%)	68 (3.40%)	72 (3.60%)	72 (3.60%)
Slice Registers	SLICE	1775 (22.19%)	1775 (22.19%)	1613 (20.16%)	1654 (20.68%)	1775 (22.19%)
Register as Flip Flop	SLICE	1775 (22.19%)	1775 (22.19%)	1613 (20.16%)	1654 (20.68%)	1775 (22.19%)
CARRY8	SLICE	14 (2.80%)	14 (2.80%)	16 (3.20%)	16 (3.20%)	16 (3.20%)
F7 Muxes	SLICE	6 (0.30%)	6 (0.30%)	6 (0.30%)	6 (0.30%)	6 (0.30%)
Unique Control Sets	SLICE	105 (21.00%)	105 (21.00%)	100 (20.00%)	102 (20.40%)	105 (21.00%)
Memory						
RAMB18	RAMB18	1 (5.00%)	1 (5.00%)	3 (15.00%)	2 (10.00%)	3 (15.00%)
PPLOCs (INT Tile Ratio)	-	28 (0.09)	28 (0.09)	28 (0.09)	28 (0.09)	28 (0.09)

Notice that RM1 requires the most resources for Slice Registers, RM2 requires the most BlockRAM, and RM3 requires the most Slice LUTs. These maximum values for each resource type are summarized in the Max column—this column should be used to plan Pblock resource sizes. Remember that additional overhead is advised—packing densities for a given Reconfigurable Partition is similar to a complete design.

Clocking

The `-clocking` switch summarizes the full set of clocks in the design, then breaks down the clock loads in each Reconfigurable Module. It also reports the number of RM clock loads in each clock region (not shown here).

Static Clock Summary						
Clock Name	Static Loads	RP1	Max	Loads		
mb_prc_wrapper/mb_prc_i/ddr4/inst/u_ddr4_infrastructure/dbg_clk	2889			0		
mb_prc_wrapper/mb_prc_i/ddr4/inst/u_ddr4_infrastructure/CLK	1639			0		
dbg_hub/inst/BSCANID.u_xsdbm_id/itck_i	496			365		
mb_prc_wrapper/mb_prc_i/ddr4/inst/u_ddr4_infrastructure/addn_clkout1	8534			1403		
mb_prc_wrapper/mb_prc_i/ddr4/inst/u_ddr4_infrastructure/c0_ddr4_ui_clk	15098			0		
mb_prc_wrapper/mb_prc_i/ddr4/inst/u_ddr4_infrastructure/addn_ui_clkout2	11791			0		
Reconfigurable Module Clocking RP1						
Clock Name	Current	RM1 Loads	RM2 Loads	RM3 Loads	Max Loads	
Static Clocks						
dbg_hub/inst/BSCANID.u_xsdbm_id/itck_i	365	365	360	350	365	
mb_prc_wrapper/mb_prc_i/ddr4/inst/u_ddr4_infrastructure/addn_clkout1	1403	1403	1385	1344	1403	

Timing

The `-timing` switch analyzes the worst interface paths on the RP boundary based on logic levels. The default is to examine the 10 worst paths but this can be changed using the `-nworst` option. The Logic Path field shows the levels of logic and defines if each level is in the static (S) or reconfigurable (RM) partition. Here is a sample of a single boundary path:

```
Reconfigurable Module Boundary Timing RP1
+-----+-----+
| Characteristics | Paths |
+-----+-----+
| Path #1 | ----- |
| RP Boundary Pin | S_BSCAN_shift |
| RM With Worst Path | RP1 1st Configuration |
| Static Logic Levels | 3 |
| RM Logic Levels | 2 |
| Logic Path | FDRE (S) LUT3 (S) LUT6 (S) LUT3 (S) LUT4 (RM) LUT6 (RM) FDRE (RM) |
| Start Point Clock | itck_i |
| End Point Clock | itck_i |
| High Fanout | 45 |
| Boundary Fanout | 1 |
+-----+-----+
```

This information can help you optimize boundary paths. Insertion of pipeline registers can break up these timing challenges and even create a decoupling point between reconfigurable and static logic.

Summary

Run the `report_pr_configuration_analysis` command early in your design flow, after you have established the logic in each Reconfigurable Module but before you finalize the floorplan of the design. This report will help you optimize each Pblock for the Reconfigurable Partitions in your design, give you guidance on the clocking usage throughout the design, and provide insight as you close timing on each configuration in the overall project.

Managing Constraints for a DFX Design

Dynamic Function eXchange requires all the same constraints as any other design, and requires additional physical constraints (Pblocks) and may require various sets of constraints that define the various configurations defined by a set of RMs. Constraints be defined as either global or scoped.

- Global - These are constraints that are applied to the entire design, and all object references (cell/pin/net/port) are with respect to the full design hierarchy.
- Scoped - These are constraints that are written with respect to a module other than the top module and are scoped to a single or all instances of that module. For the following discussion on Dynamic Function eXchange, it is assumed scoped constraints are scoped to one or more instances of a reconfigurable module. For more information on XDC

scoping, see this [link](#) in *Vivado Design Suite User Guide: Using Constraints* (UG903) [Ref 19].

Constraint Creation

Constraints should be broken up into multiple files based on the type of constraint. For Dynamic Function eXchange, Xilinx recommends breaking up the design constraints into the three following types:

- **Static** - Constraints (physical or timing) that reference only static objects (cell/pins/nets/ports), and do not reference any objects within an RP. These are global constraints and are applied to static synthesis, or at implementation of the initial configuration. It is not recommended to reapply these constraints for subsequent configurations where static is imported, as these constraints should already exist within the static DCP and reapplying them can cause unintended constraint interactions.
- **Boundary** - These are timing constraints (such as `set_false_path`) that reference objects in both static and the RP. Write these constraints referencing the RP pin object, and not objects internal to the RM. For example, take the following two constraints:

```
set_false_path -from [get_pins static_reset/C] -through [get_pins rp_inst/rst]
set_false_path -from [get_pins static_reset/C] -to [get_pins rp_inst/*foo*/D]
```

The first constraint is preferred, as this constraint will not be lost when the RMs are converted to a black box after the initial configuration. Even when the RMs are carved out, the RP interface still exists in the static design, and the `rp_inst/rst` pin referenced in the constraint still exists in the design. If a constraint is defined using objects inside the RM, as shown in the second constraint, the constraint becomes invalid and is dropped from the design after carving. These constraints would have to be reapplied for every configuration if written using this syntax.

- **RM** - Constraints (physical or timing) that reference only RM logic. While these constraints should all be scoped to the RM, how they get applied will depend on what type of constraints they are:
 - **General RM constraints** - These constraints apply to every instance of the RM. An example is a timing constraint, like a false path, multi-cycle exception, or a `create_clock` for a local RM clock. These apply to the RM regardless of the physical location.
 - **RP specific RM constraints** - These constraints are specific to the location (Pblock) in which the RM is being implemented. An example is physical constraints like specific block RAM placement, or `PACKAGE_PIN` constraints for embedded IO.



IMPORTANT: *If an RP has embedded IO, the IO (`PACKAGE_PIN`, `IOSTANDARD`, `direction`) must be reapplied for every configuration, even if the IO pins are identical between every RM. In the Vivado database, a port is a top-level object. However, the IO constraint information associated with that port is intentionally cleared out during carving of the RM if the associated IO buffer is part of an RP.*

In addition to this, it is also recommended to separate physical (Pblock, LOC, PACKAGE_PIN, etc) and timing constraints into separate XDC files. This allows for better control of when constraints are applied (synthesis vs implementation), and easier management of constraint files if constraints need to be modified.

Constraint Application

Once all of the constraint have been broken up, then the method of how and when to apply these constraints become much easier.

- Static Timing Constraints - Apply these at static synthesis and mark them for use in synthesis as well as implementation so that they are passed into the static synthesis DCP. As long as these constraints are properly applied, and exist in the post synthesis DCP, there is no need to reapply them at implementation time.
- Static Physical Constraints - Apply these at implementation of the initial configuration. There is no need to reapply these for subsequent configurations, as the routed static DCP contains all of these, as well as the static timing constraints.
- Boundary Timing Constraints - Apply these at implementation of the initial configuration. If the constraints were written without referencing any objects internal to the RM, then these constraints do not need to be reapplied for subsequent configurations.
- General RM Constraints - Apply these at RM synthesis time and mark them for use in synthesis and implementation. These constraints exist in the RM DCP and are applied when the DCP is linked into the full design.



IMPORTANT: *If there are additional OOC specific timing constraints that are used for OOC synthesis only (such as a `create_clock` for a module port), mark these for use in `out_of_context` so that they do not get applied to the full design. Failure to do this can result in unwanted constraint interaction when the top-level constraints are applied and conflict with these lower level OOC constraints.*

- RP Specific RM Constraints - Apply these at implementation of any configuration involving the specific RM at the specific RP location. All physical constraints within the RM (including IOB) are cleared out during carving, and must be reapplied.

If a post-route_design RM DCP is being reused, the physical information (block RAM, IOB, etc) is all included within the RM DCP and there is no need to reapply these constraints. However, the RM DCP (created with `write_checkpoint -cell`) does not contain any RM specific timing constraints, and all general RM constraints need to be applied as well as any boundary timing constraint that reference RM objects.

Defining Reconfigurable Partition Boundaries

Partial reconfiguration is done on a frame-by-frame basis. As such, when partial BIT files are created, they are built with a discrete number of configuration frames. The size of a partial bit file depends on the number and type of frames included. You can see this size in the header of a raw bit file (.rbit) created by `write_bitstream -rawbitfile`.

Partition boundaries do not have to align to reconfigurable frame boundaries, but the most efficient place and route results are achieved when this is done. Static logic is permitted to exist in a frame that will be reconfigured, as long as:

- It is outside the area group defined by the Pblock
- It does not contain dynamic elements such as Block RAM, Distributed (LUT) RAM, or SRLs (7 series only).

When static logic is placed in a reconfigured frame, the exact functionality of the static logic is rewritten, and is guaranteed not to glitch.

Irregular shaped Partitions (such as a T or L shapes) are permitted but discouraged. Placement and routing in such regions can become challenging, because routing resources must be entirely contained within these regions. Boundaries of Partitions can touch, but this is not recommended, as some separation helps mitigate potential routing restriction issues as these partitions connect to the static design. Nested or overlapping Reconfigurable Partitions (partitions within partitions) are not permitted. Design rule checks (**Reports > Report DRC**) validate the Partitions and settings in a PR design.

Only one Reconfigurable Partition can exist per physical Reconfigurable Frame.

A Reconfigurable Frame is the smallest size physical region that can be reconfigured, and its height aligns with clock region or I/O bank boundaries. A Reconfigurable Frame cannot contain logic from more than one Reconfigurable Partition. If it were to contain logic from more than one Reconfigurable Partition, it would be very easy to reconfigure the region with information from an incorrect Reconfigurable Module, thus creating contention. The software tools are designed to avoid that potentially dangerous occurrence.

Avoiding Deadlock

Some transactions across an RM boundary can take multiple cycles to complete. Removing an RM after a transaction has started but before it completes causes the system to deadlock (for example, the master, which initiated the transaction, waits for a response from a slave which no longer exists).

Additionally, the RM itself can cause deadlock. For example, assume some software is polling an RM register for a particular value. If the RM is removed, the software might stall as it continues to wait. It could also stall while waiting on a large block transfer to complete.

Any Dynamic Function eXchange design should be built with some sort of handshaking, ensuring that the removal of a Reconfigurable Module occurs when it is safe to do so. This request or acknowledgment pairing is part of the user design and can be built in any fashion you deem appropriate.

Design Revision Checks

A partial bitstream contains programming information and little else, as described in [Chapter 8, Configuring the Device](#). While you do not need to identify the target location of the bitstream (the die location is determined by the addressing that is part of the BIT file), there are no checks in the hardware to ensure the partial bitstream is compatible with the currently operating design. Loading a partial bitstream into a static design that was not implemented with that Reconfigurable Module revision can lead to unpredictable behavior.

Xilinx suggests that you prefix a partial bitstream with a unique identifier indicating the particular design, revision and module that follows. This identifier can be interpreted by your configuration controller to verify that the partial bitstream is compatible with the resident design. A mismatch can be detected, and the incompatible bitstream can be rejected, before being loaded into configuration memory. This functionality must be part of your design, and would be similar to or in conjunction with decryption and/or CRC checks, as described in *PRC/EPRC: Data Integrity and Security Controller for Partial Reconfiguration* (XAPP887) [\[Ref 33\]](#).

A bitstream feature provides a simple mechanism for tagging a design revision. The **BITSTREAM.CONFIG.USR_ACCESS** property allows you to enter a revision ID directly into the bitstream. This ID is placed in the **USR_ACCESS** register, accessible from the FPGA programmable logic through a library primitive of the same name. Partial Reconfiguration designs can read this value and compare it to information a user can add to a header of a partial bitstream to confirm the revisions of the design match. More information on this switch can be found in the application note *Bitstream Identification with USR_ACCESS using the Vivado Design Suite* (XAPP1232) [\[Ref 34\]](#).



CAUTION! Do not use the *TIMESTAMP* feature because this value is not consistent for each call to *write_bitstream*. Only select a consistent, explicit ID to be used for all *write_bitstream* runs.

Simulation and Verification

Configurations of Dynamic Function eXchange designs are complete designs in and of themselves. All standard simulation, timing analysis, and verification techniques are supported for DFX designs. Partial reconfiguration itself cannot be simulated. Specifically, the delivery of a partial bitstream to a configuration port like the ICAP to see the resulting change (including intermediate states) in a Reconfigurable Partition.

Design Considerations and Guidelines for 7 Series and Zynq Devices

Overview

This chapter explains design requirements that are unique to Dynamic Function eXchange (DFX), and are specific to 7 series and Zynq®-7000 SoC devices.

To take advantage of the Dynamic Function eXchange capability of Xilinx® devices, you must analyze the design specification thoroughly, and consider the requirements, characteristics, and limitations associated with DFX designs. This simplifies both the design and debug processes, and avoids potential future risks of malfunction in the design.

Design Elements Inside Reconfigurable Modules

Not all logic is permitted to be actively reconfigured. Global logic and clocking resources must be placed in the static region to not only remain operational during reconfiguration, but to benefit from the initialization sequence that occurs at the end of a full device configuration.

Logic that can be placed in a Reconfigurable Module includes:

- All logic components that are mapped to a CLB slice in the device. This includes LUTs (look-up tables), FFs (flip-flops), SRLs (shift registers), RAMs, and ROMs.
 - Block RAM and FIFO:
 - RAMB18E1, RAMB36E1, BRAM_SDP_MACRO, BRAM_SINGLE_MACRO, BRAM_TDP_MACRO
 - FIFO18E1, FIFO36E1, FIFO_DUALCLOCK_MACRO, FIFO_SYNC_MACRO
- Note:** The IN_FIFO and OUT_FIFO design elements cannot be placed in an RM. These design elements must remain in static logic.
- DSP blocks: DSP48E1
 - PCIe® (PCI Express): Entered using PCIe IP

All other logic must remain in static logic and must not be placed in an RM, including:

- Clocks and Clock Modifying Logic - Includes BUFG, BUFR, MMCM, PLL, and similar components
- I/O and I/O related components (ISERDES, OSERDES, IDELAYCTRL, etc.)
- Serial transceivers (MGTs) and related components
- Individual architecture feature components (such as BSCAN, STARTUP, XADC, etc.)

Global Clocking Rules

Because the clocking information for every Reconfigurable Module for a particular Reconfigurable Partition is not known at the time of the first implementation, the DFX tools pre-route each BUFG output driving a partition pin on that RP to all clock regions that the Pblock encompasses. This means that clock spines in those clock regions might not be available for static logic to use, regardless of whether the RP has loads in that region.

In 7 series devices, up to 12 clock spines can be pre-routed into each clock region. This limit must account for both static and reconfigurable logic. For example, if 3 global clocks route to a clock region for static needs, any RP that covers that clock region can use the 9 global clocks available, collectively, in addition to those three top-level clocks.

In the example shown in [Figure 6-1, page 88](#), `icap_clk` is routed to clock regions X0Y1, X0Y2, and X0Y3 prior to placement, and static logic is able to use the other clock spines in that region.




Figure 6-1: Pre-routing Global Clock to Reconfigurable Partition



RECOMMENDED: If there are a large number of global clocks driving an RP, create area groups that encompass complete clock regions to ease placement and routing of static logic. Global clocks can be downgraded to regional clocks (for example, BUFR, BUFH) for clocks with fewer loads or less demanding requirements. Shifting clocks from global to local resources allows for more flexibility in floorplanning when the RP requires many unique clocks.

Creating Pblocks for 7 Series Devices

As noted in [Apply Reset After Reconfiguration in Chapter 3](#), the height of the Reconfigurable Partition must align to clock region boundaries if `RESET_AFTER_RECONFIG` is to be used. Otherwise, any height can be selected for the Reconfigurable Partition.

The width of the Reconfigurable Partition must be set appropriately to make most efficient usage of interconnect and clocking resources. The left and right edges of Pblock rectangles should be placed between two resource columns (for example, CLB-CLB, CLB-block RAM or CLB-DSP) and not between two interconnect columns (INT-INT). This allows the placer and router tools the full use of all resources for both static and reconfigurable logic. Implementation tool DRCs provide guidance if this approach is not followed.

Automatic Adjustments for Reconfigurable Partition Pblocks

The Pblock `SNAPPING_MODE` property automatically resizes Pblocks to ensure no back-to-back violations occur for 7 series designs. When `SNAPPING_MODE` is set to a value of `ON` or `ROUTING`, it creates a new set of derived Pblock ranges that are used for implementation. The new ranges are stored in memory, and are not written out to the XDC. Only the `SNAPPING_MODE` property is written out, in addition to the normal Pblock constraints.

In 7 series devices the structure is such that the routing resources, called interconnect tiles, are placed adjacent, or back-to-back. When floorplanning for partial reconfiguration, it is important to understand where these back-to-back boundaries exist. If a Pblock splits these paired interconnect tiles, it is called a back-to-back violation. For more information on back-to-back interconnect please refer to [Creating Reconfigurable Partition Pblocks Manually, page 92](#).

The original Pblock rectangle(s) are not modified when using `SNAPPING_MODE` and can still be resized, moved, or extended with additional rectangles. Whenever the original Pblock rectangle is modified, the derived ranges are automatically recalculated. The `SNAPPING_MODE` property is supported in batch mode, so there is no requirement to open the current Pblock in the Vivado® IDE to set the `SNAPPING_MODE` value, although this option is available when performing interactive floorplanning, as shown in [Figure 6-2, page 90](#).




Figure 6-2: Enabling the SNAPPING_MODE Property in the Vivado IDE

When you set the **SNAPPING_MODE** property using the following syntax (or by selecting the Pblock Property as shown above), the implementation tools automatically see the corrected Pblock ranges.

```
set_property SNAPPING_MODE ON [get_pblocks <pblock_name>]
```

The table below shows **SNAPPING_MODE** property values for 7 series devices.

Table 6-1: SNAPPING_MODE Property Values for 7 Series Devices

Property	Value	Description
SNAPPING_MODE	OFF	Default for 7 series. No adjustments are made and DERIVED_RANGES == GRID_RANGES
	ON	Fixes all back-to-back violations.
	ROUTING	Same behavior as ON, except for the following exceptions: <ul style="list-style-type: none"> Does not fix back-to-back violations across the center clock column to improve routing. Grabs unbonded I/O and GT sites that are within or adjacent to the RP Pblock to improve routing. It can only use these resources for PR routing if the sites are unbonded and if the entire column (Clock Region in height) is included in the Pblock rectangle. This is the recommended value for 7 series and Zynq designs.

The **SNAPPING_MODE** property also works in conjunction with **RESET_AFTER_RECONFIG**. Using **RESET_AFTER_RECONFIG** requires Pblocks to be vertically frame (or clock region) aligned. When **SNAPPING_MODE** is set to **ON** or to **ROUTING** and **RESET_AFTER_RECONFIG** is set to **TRUE**, the derived ranges automatically include all sites necessary to meet this requirement.

Figure 6-3 shows the original user-created Pblock in purple. **RESET_AFTER_RECONFIG** has been enabled, and both left and right edges split interconnect columns. By applying **SNAPPING_MODE**, the resulting derived Pblock (shown in yellow) is *narrower* to avoid INT-INT boundaries, and *taller* to snap to the height of a clock region.





Figure 6-3: Original and Derived Pblocks using SNAPPING_MODE

Creating Reconfigurable Partition Pblocks Manually

If automatic modification to the Reconfigurable Partition Pblock is not desired to fix back-to-back issues, you can create Pblock ranges manually to meet your needs. This is most useful when explicit control is needed for Pblocks that must span non-reconfigurable sites, such as configuration blocks or the center column, which contains clock buffer resources.

In [Figure 6-4](#), note that the left and right edges are drawn between CLB columns for the Pblock highlighted in white. Visualization of the interconnect tiles as shown in this image requires that the routing resources are turned on, using this symbol in the Device View :




Figure 6-4: Optimal - Reconfigurable Partition Pblock Splitting CLB-CLB on Both Left and Right Edges

The Reconfigurable Partition Pblock must include all reconfigurable element types within the shape drawn. In other words, if the rectangle selected encompasses CLB (Slice), block RAM, and DSP elements, all three types must be included in the Pblock constraints. If one of these is omitted, a DRC is triggered with an alert that a split interconnect situation has been detected.

Other considerations must be taken if the Reconfigurable Partition spans non-reconfigurable sites, such as the center-column clocking resources or configuration components (ICAP, BSCAN, etc.), or abuts non-reconfigurable components such as I/O. If a Pblock edge splits interconnect columns for different resource types, implementation tools accept this layout, but restrict placement in the columns on each side of the boundary. If this prohibits sites that are needed for the design (such as the ICAP or BSCAN, for example), the Pblock must be broken into multiple rectangles to clearly define reconfigurable logic usage, or **SNAPPING_MODE** must be used.

The implementation tools automatically prevent placement on both sides of the back-to-back interconnect by creating **PROHIBIT** constraints. If the sites that are prohibited due to a back-to-back violation are not needed in the design, it is acceptable to leave the back-to-back violation in the design. Doing so allows an extra column of routing tiles to be included in the dynamic region, and can reduce congestion in a dynamic region that spans non-reconfigurable sites. In this case, a Critical Warning is issued by DRCs, but the warning can be safely ignored if you understand the trade-offs of placement versus routing resources.

The one exception to this behavior is around the clock column. If a violation occurs at the clock column boundary, **PROHIBIT** constraints are generated for the RM side of the violation (typically SLICE prohibits), but the clocking resources do not get prohibit constraints and are still available to the static logic. The **SNAPPING_MODE** property has a value of **ROUTING**, which takes advantage of this special exception. For example, the initial floorplan shown in [Figure 6-5](#) spans the center column, which contains clock buffer resources (BUFHCE/BUFGCTRL). These resources have not been included in the Pblock, as they are not highlighted in [Figure 6-5](#). There is violation caused by spanning this clock column but the resources can still be used by the static logic.




Figure 6-5: Pblock Spanning Non-Reconfigurable Sites

Prohibited sites appear in placed or routed checkpoints as sites with a red circle with a slash, as shown in Figure 6-6. With this automatic prohibit feature, the routing interconnect associated with reconfigurable sites (CLBs) can still be used for the reconfigurable module even though the CLBs themselves are not used. In Figure 6-6, the column of INT on the left is available for the RM, but the column of INT on the right is only available for static logic because these are part of the clock tile, which is not reconfigurable for 7 series devices.




Figure 6-6: Prohibited Sites in a Checkpoint

If a back-to-back violation prohibits sites that are needed for the design (that is, ICAP or BSCAN sites), a placement error is issued, stating that not enough sites are available in the device.

ERROR: [Common 17-69] Command failed: Placer could not place all instances

To avoid this restriction, create multiple Pblock rectangles that avoid splitting interconnect columns, as shown in Figure 6-7, or use the Pblock `SNAPPING_MODE` property.



RECOMMENDED: In general, spanning non-reconfigurable site types (such as IOB, configuration, or clocking columns) should be avoided whenever possible. If the Pblock must span one of these, the clocking column is the least risky choice, owing to its special nature (described previously). Use `SNAPPING_MODE ROUTING` to cross this boundary as efficiently as possible.




Figure 6-7: Multiple Pblock Rectangles that Avoid Non-Reconfigurable Resources

Figure 6-8 is a close-up of this split, showing Slice (CLB) and Interconnect (INT) resource types. The gap between the two Pblock rectangles gives full access to the BUFHCE components to route completely using static resources. This also leaves one column of CLBs available for the static design to use. Although routing resources exist that can cross these gaps, the overall routability of such structures is notably reduced. This approach is more challenging and should be avoided if possible. When spanning other static boundaries, such as IOB or configuration tiles, the routing gap for the dynamic region becomes two INT resources, and routing becomes difficult.




Figure 6-8: Close-up Showing Columns Reserved for Clock Routing Usage

Irregular shaped Partitions (such as a T or L shapes) are permitted, but you are encouraged to keep overall shapes as simple as possible. Placement and routing in such regions can become challenging because routing resources must be entirely contained within these regions. Boundaries of Partitions can touch, but this is not recommended, as some separation helps mitigate potential routing restriction issues. Nested or overlapping Reconfigurable Partitions (partitions within partitions) are not permitted.

Finally, only one Reconfigurable Partition can exist per physical Reconfigurable Frame. A Reconfigurable Frame is the smallest size physical region that can be reconfigured, and aligns with clock region boundaries. A Reconfigurable Frame cannot contain logic from more than one Reconfigurable Partition. If it were to contain logic from more than one Reconfigurable Partition, it would be very easy to reconfigure the region with information from an incorrect Reconfigurable Module, thus creating contention. The Vivado tools are designed to avoid that potentially dangerous occurrence.

Using High Speed Transceivers

Xilinx high speed transceivers (GTP, GTX, GTH,GTZ) are not reconfigurable in 7 series devices, and must remain in static logic. However, settings for the transceivers can be updated during operation using the DRP ports. For more information on the transceiver settings and DRP access, see *7 Series FPGAs GTX/GTH Transceivers User Guide* (UG476) [Ref 21], or *7 Series FPGAs GTP Transceivers User Guide* (UG482) [Ref 22].

Dynamic Function eXchange Design Checklist (7 Series)

Xilinx highly encourages the following items for a 7 series FPGA design using Dynamic Function eXchange:

Recommended Clocking Networks

Are you using Global Clock Buffers, Regional Clock Buffers, or Clock Modifying Blocks (MMCM, PLL)?

These blocks must be in static logic.

See [Design Elements Inside Reconfigurable Modules](#) for more information, and [Global Clocking Rules](#) for complete details on global clock implementation.

Configuration Feature Blocks

Are you using device feature blocks (BSCAN, CAPTURE, DCIRESET, FRAME_ECC, ICAP, STARTUP, USR_ACCESS)?

These featured blocks must be in static logic.

See [Design Elements Inside Reconfigurable Modules](#) for more information.

High Speed Transceiver Blocks

Do you have high speed transceivers in your design?

High speed transceivers must remain in the static partition.

See [Using High Speed Transceivers](#) for specific requirements.

System Generator DSP Cores, HLS cores, or IP Integrator Block Diagrams

Are you using System Generator DSP cores, HLS cores, or IP Integrator block diagrams in your Dynamic Function eXchange design?

Any type of source can be used as long as it follows the fundamental requirements for Dynamic Function eXchange. Any code processed by SysGen, HLS, or Vivado IP Integrator (or other tools) is eventually synthesized. The resulting design checkpoint or netlist must be made up entirely of reconfigurable elements (CLB, block RAM, DSP) for it to be legally included in an RP.

Packing I/Os into Reconfigurable Partitions

Do you have I/Os in reconfigurable modules?

All I/Os must reside in static logic.

See [Design Elements Inside Reconfigurable Modules](#) for more information.

Packing Logic into Reconfigurable Partitions

Is all logic that must be packed together in the same Reconfigurable Partition?

Any logic that must be packed together must be in the same RP and RM.

See [Packing Logic](#) for more information.

Packing Critical Paths into Reconfigurable Partitions

Are critical paths contained within the same partition?

Reconfigurable partition boundaries limit some optimization and packing, so critical paths should be contained within the same partition.

See [Packing Logic](#) for more information.

Floorplanning

Can your Reconfigurable Partitions be floorplanned efficiently?

See [Creating Pblocks for 7 Series Devices](#) for more information.

Recommended Decoupling Logic

Have you created decoupling logic on the outputs of your RMs?

During reconfiguration the outputs of RPs are in an indeterminate state, so decoupling logic must be used to prevent static data corruption.

See [Decoupling Functionality](#) for more information.

Recommended Reset after Reconfiguration

Are you resetting the logic in an RM after reconfiguration?

After reconfiguration, new logic might not start at its initial value. If the Reset After Reconfiguration property is not used, a local reset must be used to ensure it comes up as expected when decoupling is released. Clock and other inputs to the reconfigurable partition can also be disabled during reconfiguration to prevent initialization issues. Alternatively, the Reset After Reconfiguration property can be applied. This option holds internal signals steady during reconfiguration, then issues a masked global reset to the reconfigured logic.

See [Apply Reset After Reconfiguration](#) for more information.

Debugging with Logic Analyzer Blocks

Are you using the Vivado Logic Analyzer with your Dynamic Function eXchange design?

Vivado Logic Analyzer (ILA/VIO debug cores) can be used in your Dynamic Function eXchange design, but care must be taken when connecting these cores to debug hubs. Use the automatic inference solution shown in [Using Vivado Debug Cores](#).

Efficient Reconfigurable Partition Pblocks

Have you created efficient Reconfigurable Partition Pblock(s) for your design?

The height of the Reconfigurable Partition Pblock must align with the top and bottom of a clock region boundary, if the `RESET_AFTER_RECONFIG` property is to be used. Otherwise, any height can be selected for the Reconfigurable Partition Pblock.

See [Creating Pblocks for 7 Series Devices](#) for more information.

Validating Configurations

How do you validate consistency between configurations?

The `pr_verify` command is used to make sure all configurations have matching imported resources.

See [Verifying Configurations](#) for more information.

Configuration Requirements

Are you aware of the particular configuration requirements for Dynamic Function eXchange for your design and device?

Each device family has specific configuration requirements and considerations.

See [Configuring the Device](#) for more information.

Effective Pblock recommendations

Does an RP Pblock extend over the center clock column or the configuration column in the device?

Due to the back-to-back INT tile requirement for 7 series devices, coupled with the **CONTAIN_ROUTING** requirement, extending a Pblock over these specialized blocks in the device can make routing very difficult or impossible. Avoid extending an RP Pblock across these areas whenever possible.

See [Automatic Adjustments for Reconfigurable Partition Pblocks](#) and [Creating Reconfigurable Partition Pblocks Manually](#) for more information on back-to-back requirements.

Design Considerations and Guidelines for UltraScale and UltraScale+ Devices

Overview

This chapter explains design requirements that are unique to Dynamic Function eXchange (DFX), and are specific to UltraScale™ and UltraScale+™ devices.

To take advantage of the Dynamic Function eXchange capability of Xilinx® devices, you must analyze the design specification thoroughly, and consider the requirements, characteristics, and limitations associated with DFX designs. This simplifies both the design and debug processes, and avoids potential future risks of malfunction in the design.

Design Elements Inside Reconfigurable Modules

In UltraScale and UltraScale+ devices, nearly all component types may be partially reconfigured.

Logic that can be placed in a Reconfigurable Module includes:

- All logic components that are mapped to a CLB slice in the FPGA. This includes LUTs (look-up tables), FFs (flip-flops), SRLs (shift registers), RAMs, and ROMs.
- Block RAM and FIFO: RAMB18E2, RAMB36E2, FIFO18E2, FIFO36E2
- DSP blocks: DSP48E2
- PCIe[®] (PCI Express), CMAC (100G MAC), and ILKN (Interlaken MAC) blocks
- UltraRAM blocks: URAM288
- SYSMON (XADC and System Monitor)
- Clocks and Clock Modifying Logic: Includes BUFG, BUFGCE, BUFGMUX, MMCM, PLL, and similar components
- I/O and I/O related components (ISERDES, OSERDES, IDELAYCTRL, etc.)
- Serial transceivers (MGTs) and related components

Note: DNA_PORT -- the Device DNA Access Port (DNA_PORTE2) is the only configuration element in the CONFIG_SITE that is reconfigurable. Any other usage of CONFIG_SITE elements is not permitted.

Only configuration components must remain in the static part of the design. These components are:

- BSCAN
- CFG_IO_ACCESS
- DCIRESET
- EFUSE_USR
- FRAME_ECC
- ICAP
- MASTER_JTAG
- STARTUP
- USR_ACCESS

Creating Pblocks for UltraScale and UltraScale+ Devices

As part of improvements to the UltraScale architecture, the smallest unit that can be reconfigured is much smaller than in previous architectures. The minimum required resources for reconfiguration varies based on the resource type, and are referred to as a Programmable Unit (PU). Because adjacent sites share a routing resource (or Interconnect Tile) in UltraScale, a PU is defined in terms of pairs.

Examples of some of the minimum PU that can be reconfigured based on the site types:

- CLB PU: 2 adjacent CLBs, and the shared interconnect
- Block RAM PU: 1 block RAM/FIFO, the 5 adjacent CLBs, and the shared interconnect
- DSP PU: 1 DSP, the 5 adjacent CLBs, and the shared interconnect
- IOB PU: The IO of the full height of the clock_region and includes BITSlice_CONTROL, BITSlice_RX_TX, BITSlice_TX, BUFGCE, BUFGCE_DIV, BUFGCTRL, IOB, MMCME3_ADV, PLLE3_ADV, PLL_SELECT_SITE, RIU_OR, etc. The adjacent 60 CLBs and the shared interconnect.

Automatic Adjustments for PU on Pblocks

In UltraScale and UltraScale+ devices, there is no `RESET_AFTER_RECONFIG` option. Instead, GSR is always issued at the end of a partial reconfiguration, and there are no Pblock size/shape requirements to enable this like there are in 7 series devices. However, to ensure that the Pblock does not violate any rules for minimum PU sizes, the `SNAPPING_MODE` property is also always on by default, and automatically adjusts the Pblock to make sure it is valid for PR.

Figure 7-1 and Figure 7-2 below give an example of how `SNAPPING_MODE` adjusts the Pblock for PU alignment. In Figure 7-1, despite the larger outer rectangle, only the selected tiles belong to the RP Pblock. The upper block RAM and DSP sites are not included because they are not fully contained in the Pblock, and the associated CLB sites are not included either, based on the PU rules. There are also CLB sites on both the left and right edge that are not included in the Pblock because the adjacent CLBs are not owned by the original rectangle.




Figure 7-1: `SNAPPING_MODE` Example - UltraScale

While **SNAPPING_MODE** made the above Pblock legal for the RP, it is possible that the intent was to include all of these sites. By making a small adjustment to the original Pblock rectangle, you can prevent **SNAPPING_MODE** from removing sites that are intended for the dynamic region. In [Figure 7-2](#) the Pblock has been expanded by one CLB on the left, right, and top edges. The shaded tiles that are owned by the RP Pblock now match the outer rectangle.




Figure 7-2: PU Aligned Pblock

While shading shows what is included in a Reconfigurable Partition, you can best visualize the sites owned by a RP by using highlighting scripts that the Vivado Design Suite tools create automatically for Pblocks of Reconfigurable Partitions. The following steps can be used for debugging/verifying Pblocks:

1. Create or make an adjustment to an RP. The cell assigned to the Pblock must have the **HD.RECONFIGURABLE** property set.
2. Source the highlighting script that was generated by the Vivado tools.

```
source ./hd_visual/<pblock_name>_AllTiles.tcl
```

Note: The scripts in the `hd_visual` directory are updated any time the Pblock constraints are processed. This includes opening a design that contains Pblocks and creating or modifying Pblocks in an open design.

Sharing Configuration Frames between RP and Static Logic

Even though UltraScale and UltraScale+ devices Pblocks are not required to be frame aligned (that is, the height of the clock region), Dynamic Function eXchange still programs the entire configuration frame. This means that logic outside of the RP is overwritten. This does not cause any issues in DFX, but in previous architectures there were some limitations about what kind of static logic could be in the same frame as reconfigurable logic.

For UltraScale and UltraScale+ devices, any static logic can be placed in the same configuration frame as the RM, in any sites not owned by the RP Pblock. This includes block RAM, DSP, and LUT RAM. There is still a restriction however, that there can be only one RP per configuration frame. That means you cannot vertically stack two RPs in the same clock region.

Expansion of CONTAIN_ROUTING Area

The contained routing requirement of RP Pblocks for UltraScale and UltraScale+ devices has been relaxed to allow for improved routing and timing results. Instead of routing being confined strictly to the resources owned by the Pblock, the routing footprint is expanded. This includes resources that are within the Pblock boundary, but not necessarily owned by the Pblock, as well as resources beyond the Pblock rectangle. This means there might be RM nets and Partition Pins outside of the Pblock boundary. However, any partition pin or contained net will still be within the expanded routing footprint.

The expanded routing footprint can be visualized by sourcing one of the `hd_visual` Tcl scripts. These are scripts that are generated automatically during the Dynamic Function eXchange flow, and can be found in the `./hd_visual` subdirectory within the current working directory. The visualization script that shows the expanded routing footprint is named `./hd_visual/<pblock_name>_Routing_AllTiles.tcl`. The expanded routing footprint is actually determined during routing, so this file is not available until `route_design` completes. To see the expanded routing footprint, source this file from the Tcl Console after opening a routed design. This selects all tiles available to the router, and then selected tiles can be highlighted or marked as desired. In the figure below, the user-defined Pblock that bounds placement is shown in blue, and the expanded routing zone is shown in yellow.

```
source ./hd_visual/pblock_inst_shift_low_Routing_AllTiles.tcl
```




Figure 7-3: Highlighted Pblock and expanded routing footprint

Obviously any frames that are used by the RM must be contained with the partial bit files, so one effect of expanding the routing footprint will be larger partial bit files. The increase in size depends on the original Pblock size and shape. Pblocks that are already rectangular can still expand. However, the expansion cannot go beyond a clock region boundary in the vertical direction; it can extend into a new clock region to the left or right. It may help the routability of the RP if the Pblock boundaries stop short of internal clock region edges, especially in the vertical direction. Pblock edges that align to the device edges, such as left or bottom edges, should not be pulled in just to allow for expanded routing. This causes placement issues if the static region now has access to small pockets of resources along the edges. Xilinx recommends keeping this routing expansion enabled, but if the partial bitstream size is more critical than the performance of the design, then this feature can be disabled by setting the following parameter:

```
set_param hd.routingContainmentAreaExpansion false
```



IMPORTANT: *The expanded routing footprint is not supported for 7 series devices.*

UltraRAM Behavior

Just as with a full device configuration, UltraRAM memory is initialized to all 0's during partial reconfiguration. There is no user defined INIT attribute and therefore the content of the SRAM array cannot be initialized to user defined values.

Floorplanning Rules for Clocks inside an RP

UltraScale and UltraScale+ devices support clocking resources within the RP such as BUFG_*, PLL, and MMCM. Designs incorporating this feature should follow the general design restrictions described in [Global Clocking Rules](#) as well as the additional floorplanning rules below. These rules are required to ensure that the clocks internal to the RP can reach the necessary routing resources within the frames owned by the RP Pblock.

1. Create rectangular Pblocks whenever possible. If the Pblock is made up of multiple rectangles, the tallest column of the Pblock must be `clock_region` aligned.
2. The `CLOCK_ROOT` property of the internal RM clock should be set as one of the tallest columns in the Pblock. The tools attempt to pick the correct columns for the `CLOCK_ROOT` automatically, but in some cases this cannot be done.
 - a. If a `USER_CLOCK_ROOT` property exists on the clock net, then the tools will not automatically select the `CLOCK_ROOT`. If the `USER_CLOCK_ROOT` property is set to a column that is not the full height of the Pblock, unroutable connections might occur.
 - b. Certain configurations of `BUFG_GT` require that the `CLOCK_ROOT` be in the same region as the `BUFG_GT`. If this is not the tallest column of the Pblock, unroutable connection might occur. To resolve this, consider splitting the clock net into two `BUFG_GT` (one for user logic, and the other for the direct GT connections). This way each clock can have its own `CLOCK_ROOT`.

As shown in [Figure 7-4](#), a `CLOCK_ROOT` defined in region `x2y2` (top-left of the Pblock) would prevent routing to any loads in region `x3y1` (bottom-right of the Pblock), because the region `x2y1` is not available to the clock. Conversely, if the `CLOCK_ROOT` were defined in either `x3y2` or `x3y1`, no clock routing restrictions would apply.




Figure 7-4: CLOCK_ROOT Restrictions on L-Shaped Pblocks

3. If a CLOCK_ROOT cannot be set to the tallest column, the loads of the clock can be contained to regions accessible by the clock using nested Pblocks within the RP region. The nested Pblock will prevent the placer from putting a load in a region that is not accessible by the clock due to irregularly shaped Pblocks.
4. Do not create U or H shaped Pblocks with large gaps that span an entire `clock_region`, as shown in Figure 7-5 below.




Figure 7-5: **Unsupported Pblock with clock_region Gap**

As shown in Figure 7-6 below, small static gaps, such as an IOB column, are permitted in the row of an RP Pblock. However, Xilinx recommends avoiding these gaps when possible, as they are a potential source of routing congestion because RM routes need to route over these gaps.




Figure 7-6: **Supported Pblock with Small Static Gap**

Small stair-step shaped Pblocks, as shown in Figure 7-7, are sometimes necessary. While they are supported, they can also lead to routing congestion around the corners.




Figure 7-7: Supported Stair-Step Shaped Pblock

5. RP Pblocks with clocking resources cannot share any part of a clock region with any other RP. It can share a clock region with static logic. This is true regardless of where the logic driven by these clock resources exist—inside or outside of the given RP.

Global Clocking Rules

As with architectures previous to UltraScale, all unique clocks driving the Reconfigurable Partition are pre-routed to every clock region in which the RP owns sites. Effectively, this means that the total number of global clocks driving the Reconfigurable Partition (regardless of size) is a maximum of 24. Higher clock utilization is possible when the clock source is in the RM, since these do not need to be pre-routed to every clock region. For this reason it is always important to carefully consider the RP Pblock size and shape. However, one difference in the UltraScale architecture is that there are now 24 global clocks available per clock region instead of the 12 available in 7 series devices.

Note: For BUFGCTRL components, the `PRESELECT_I0` and `PRESELECT_I1` properties are ignored during partial reconfiguration, even with `RESET_AFTER_RECONFIG` enabled. The clock source selected depends only on the select and clock enable inputs of the BUFGCTRL instance.

Clock sources that exist within Reconfigurable Modules can drive logic to the static design or other Reconfigurable Modules. Vivado handles many of the low-level details, such as

consistent clock spine usage. Design rule checks ensure that fundamental rules are applied. There are, however, a few considerations to be aware of:

- Clock resource must be consistent from one RM to the next for a given RP. While you may change characteristics of the clock such as MMCM parameters, the clock driver type (BUFG, etc.) must remain fixed so routing resources can remain consistent in the locked static design.
- Clock behavior is indeterminate during reconfiguration, so consider decoupling. Because clocks will be using high-fanout routing resources, a basic 2-to-1 MUX or register will not be appropriate like it would be for standard interfaces. Instead, a BUFGMUX can be used to block the clock source during reconfiguration. Alternately, synchronous elements in static or other RMs can be disabled or held in reset while the clock source is reconfigured.

I/O Rules

In UltraScale and UltraScale+ devices, I/O logic and buffers can be included in an RP. While the I/O can be modified from one RM to another, there are some rules that must be followed.

The following checks are done between all configurations that use the I/O sites. If an I/O site changes from being used to unused, or vice versa, then these checks are not done for those configurations. If an I/O is unused in a particular configuration, make sure the appropriate property for the design is set on these ports via the PULLTYPE attribute. For more information on setting PULLTYPE, see this [link](#) in *Vivado Design Suite Properties Reference Guide* (UG912) [Ref 31].

- The I/O direction and I/O standard must be the same between all RMs whenever the I/O is used.
- For **DCI_CASCADE**, the member bank assignments between RMs cannot overlap.
 - Legal example: In Configuration 1, **DCI_CASCADE** has banks 12, 13. In Configuration 2, **DCI_CASCADE** has banks 14, 15 and 16. They do not have overlapped banks.
 - Illegal example: In Configuration 1, **DCI_CASCADE** has banks 12 and 13. In Configuration 2, **DCI_CASCADE** has banks 13, 14, 15 and 16. In this case bank 13 overlaps.
- For **DCI_CASCADE**, member banks must be fully contained within the reconfigurable region. All of the member banks for the same **DCI_CASCADE** must be in either the same RP Pblock, or completely in static. **DCI_CASCADE** usage must remain consistent between different Reconfigurable Modules.

Changes to the IOB from one configuration to another are limited by the rules above. This means that the following I/O characteristics may be modified through Dynamic Function eXchange:

- Usage (used vs. unused, per I/O)
- Drive Strength (12mA, 8mA, etc.)
- Driver Output Impedance (40Ω, 48Ω, etc.)
- Driver Input Impedance (40Ω, 48Ω, etc.)
- Driver Slew Rate (slow, fast, etc.)
- ODT Termination (40, 60, etc.)

Adding the I/O sites into the RP requires that the entire PU (encompassing the I/O bank, BITSlice, MMCM, PLL, and one column of CLBs plus shared interconnect) be added. All components in this fundamental region are reconfigured and reinitialized, and adding these other site types to the reconfigurable region can be beneficial in some cases for these reasons:

- Adding I/O sites allows use of the routing resources of the I/O, which reduces congestion (instead of increasing congestion, as it could if the I/O sites were in Static, and caused a gap in the reconfigurable region).
- Allows reconfiguration of other clocking resources like the MMCM and PLL.
- Allows reconfiguration of other I/O logic sites such as BITSlice and BITSlice_CONTROL.

Regardless of whether or not the I/O usage or characteristics change during reconfiguration, the entire bank is reconfigured. During reconfiguration, all I/O in the banks defined by the RP Pblock is held with the dedicated global tri-state (GTS) signal, which is released at the end of reconfiguration.

If the Reconfigurable Module contains an MMCM or PLL component, the size of the partial bitstream will be at its smallest when the lock cycle of these components are set to “no wait”. Similarly, IO with DCI matching requirements will have minimal bitstream sizes when the lock cycle is set to “no wait”. Set these options using this commands:

```
set_property BITSTREAM.STARTUP.LCK_CYCLE NoWait [current_design]
set_property BITSTREAM.STARTUP.MATCH_CYCLE NoWait [current_design]
```

During the Dynamic Function eXchange flow, RMs are carved out using `update_design -black_box`. During this command any embedded IO buffers, and the associated constraints, such as `PACKAGE_PIN` and `IOSTANDARD`, are removed. When the black box RP is filled in with a new RM, these IOB constraints need to be reapplied to the design.

Using High Speed Transceivers

Xilinx high speed transceivers (GTH, GTY) are supported within a Reconfigurable Partition. As with other reconfigurable site types, the entire PU must be included. For the UltraScale GT transceivers, the PU includes

- 4 GT_CHANNEL sites (GT Quad)
- Associated GT_COMMON site
- Associated BUFG_GT_SYNC sites
- Associated BUFG_GT sites
- Associated Interconnect and CLB sites

The required GT PU is the entire height of a clock region. As with previous architectures, it is also possible to leave the GT components in static logic and change the functionality through the DRP. For more information on using UltraScale and UltraScale+ transceivers, see the *UltraScale Architecture GTH Transceivers User Guide* (UG576) [Ref 25] or the *UltraScale Architecture GTY Transceivers User Guide* (UG578) [Ref 26].

Virtex UltraScale+ High Bandwidth Memory (HBM) devices support Dynamic Function eXchange just like any other UltraScale+ device. Users have the choice of where the HBM Controllers are placed: in the static region or in the dynamic region. If the AXI High Bandwidth Memory Controller IP is kept in the static region with an active HBM reference clock, the memory interface will remain active and all memory contents are retained. Self-refresh mode can be used for power savings during reconfiguration or when the memory is not need; memory contents will be maintained in this case as well. If this IP is placed in a Reconfigurable Module, it will be reconfigured and memory contents will be reinitialized per its Vivado customization. For more information on the HBM Controller IP, consult the documentation here:

<https://www.xilinx.com/products/intellectual-property/hbm.html>

Dynamic Function eXchange Checklist for UltraScale and UltraScale+ Device Designs

Xilinx highly encourages the following for an UltraScale and UltraScale+ device design using Dynamic Function eXchange:

Recommended Clocking Networks

Are you using Global Clock Buffers or Clock Modifying Blocks (MMCM, PLL)?

These blocks can be reconfigured, but all elements in this frame type must be reconfigured. This includes an entire I/O bank and all clocking elements in that shared region, plus one column of CLBs that share the interconnect.

See [Design Elements Inside Reconfigurable Modules](#) for more information, and [Global Clocking Rules](#) for complete details on global clock implementation.

In addition, the following restrictions are currently enforced by Vivado Design Suite DRC rules. The use of clocking resources `BUFGCTRL`, `BUFG_CE` and `BUFG_GT` is supported with the following restrictions:

- Xilinx recommends using rectangular Pblock shapes. Non-rectangular shapes are also supported for RPs with clocking logic, as long as the tallest column of the Pblock is aligned vertically and horizontally with the clock region. The tallest column of the RP Pblock must also range the IOB, and this range must cover the full height of all the rectangles that define the RP Pblock, as shown in [Figure 7-8](#). In other words, this vertical column of IOB ranges must be able to access all rows of the Pblock. Pblock shapes like a sideways “L” are not supported unless the vertical section of the shaped includes the IOB range.




Figure 7-8: Tallest Column of Pblock Clock Region Aligned

- A gap is defined as an unranged site type with ranged sites on both sides of it. The following gaps are not allowed:
 - Gaps in the IOB/XIPHY ranges, such as the gap in the IOB column shown in [Figure 7-9](#) below.




Figure 7-9: Unsupported Gap in the IOB Column

- Gaps in the DSP ranges.
- o A clock region cannot be shared by two RP Pblocks if:
 - At least one of them has a global clock source.
 - The other has ranged a global clock source.

Configuration Feature Blocks

Are you using device feature blocks (BSCAN, DCIRESET, FRAME_ECC, ICAP, STARTUP, USR_ACCESS)?

These featured blocks must be in static logic.

See [Design Elements Inside Reconfigurable Modules](#) for more information.

Pblock Boundaries

Have you set the Pblock boundaries?

For UltraScale and UltraScale+ devices, the X-axis boundary of a dynamic region can be set by a PU, including CLB, Block RAM, DSP, and others. The tool adjusts the Pblock automatically for a valid placement. The Y-axis boundary of a PR region can be a clock region and IO bank. However, if BUFGCTRL/BUFG_CE/BUFG_GT are used in the RP, a full clock region must be used.

SSI Technology

Does the Pblock span an SLR of an SSI device?

If using an SSI device it is recommended to keep a dynamic region within a single SLR. However, for UltraScale and UltraScale+ devices, if a RP Pblock must span an SLR, the

necessary Laguna sites must be included to allow for routing across this boundary. This requires that at least one full clock region belongs to the dynamic region on both sides of the SLR boundary.

For more information on SSI Technology devices and Laguna, see Devices using Stacked Silicon Interconnect (SSI) Technology in the *UltraScale Architecture Configurable Logic Block User Guide* (UG574) [Ref 28].

High Speed Transceiver Blocks

Do you have high speed transceivers in your design?

High speed transceivers can be reconfigured. An entire quad, including all component types (GT_CHANNEL, GT_COMMON, BUFG_GT) must be reconfigured together.

See [Using High Speed Transceivers](#) for specific requirements.

System Generator DSP Cores, HLS cores, or IP Integrator Block Diagrams

Are you using System Generator DSP cores, HLS cores, or IP Integrator block diagrams in your Dynamic Function eXchange design?

Any type of source can be used as long as it follows the fundamental requirements for Dynamic Function eXchange. Any code processed by SysGen, HLS, or IP Integrator (or other tools) is eventually synthesized. The resulting design checkpoint or netlist must be comprised entirely of reconfigurable elements in order for it to be legally included in an RP.

Packing I/Os into Reconfigurable Partitions

Do you have I/Os in reconfigurable modules?

I/Os can be partially reconfigured. An entire I/O bank, along with all I/O logic (XiPhy) and clocking resources, must be reconfigured at once. IOSTANDARD and direction cannot change and DCI Cascade rules must be followed. But other I/O characteristics may change from one RM to the next.

See [Design Elements Inside Reconfigurable Modules](#) for more information.

Packing Logic into Reconfigurable Partitions

Is all logic that must be packed together in the same Reconfigurable Partition?

Any logic that must be packed together must be in the same RP and RM.

See [Packing Logic](#) for more information.

Packing Critical Paths into Reconfigurable Partitions

Are critical paths contained within the same partition?

Reconfigurable partition boundaries limit some optimization and packing, so critical paths should be contained within the same partition.

See [Packing Logic](#) for more information.

Floorplanning

Can your Reconfigurable Partitions be floorplanned efficiently?

See [Creating Pblocks for UltraScale and UltraScale+ Devices](#) for more information.

Recommended Decoupling Logic

Have you created decoupling logic on the outputs of your RMs?

During reconfiguration the outputs of RPs are in an indeterminate state, so decoupling logic must be used to prevent static data corruption.

See [Decoupling Functionality](#) for more information.

Recommended Reset After Reconfiguration

Are you resetting the logic in an RM after reconfiguration?

Reset After Reconfiguration is always enabled for UltraScale and UltraScale+ devices. This capability cannot be disabled.

See [Apply Reset After Reconfiguration](#) for more information.

Debugging with Logic Analyzer Blocks

Are you using the Vivado Logic Analyzer with your Dynamic Function eXchange design?

Vivado logic analyzer (ILA/VIO debug cores) can be used in your Dynamic Function eXchange design, but care must be taken when connecting these cores to debug hubs. Use the automatic inference solution shown in [Using Vivado Debug Cores](#).

Efficient Reconfigurable Partition Pblocks

Have you created efficient Reconfigurable Partition Pblock(s) for your design?

A Reconfigurable Partition Pblock can be any height, but multiple Reconfigurable Partitions cannot be stacked vertically within a single clock region.

See [Creating Pblocks for UltraScale and UltraScale+ Devices](#) for more information.

Validating Configurations

How do you validate consistency between configurations?

The `pr_verify` command is used to make sure all configurations have matching imported resources.

See [Verifying Configurations](#) for more information.

Configuration Requirements

Are you aware of the particular configuration requirements for Dynamic Function eXchange for your design and device?

Each device family has specific configuration requirements and considerations.

See [Configuring the Device](#) for more information.

Configuring the Device

Overview

This chapter describes the system design considerations when configuring your device with a partial BIT file, as well as architectural features in the FPGA that facilitate Dynamic Function eXchange (DFX). Because most aspects of Dynamic Function eXchange are no different than standard full configuration, this section concentrates on the details that are unique to DFX.

Configuration Modes

Dynamic Function eXchange is supported using the following configuration modes:

- **ICAP:** A good choice for user configuration solutions. Requires the creation of an ICAP controller as well as logic to drive the ICAP interface.
- **MCAP:** (UltraScale™ and UltraScale+™ devices only) Provides a dedicated connection to the configuration engine from one specific PCIe® block per device.
- **PCAP:** The primary configuration mechanism for Zynq-7000 SoC and Zynq UltraScale+ designs.
- **JTAG:** A good interface for quick testing or debug. Can be driven with the Vivado Logic Analyzer.
- **Slave SelectMAP** or **Slave Serial:** A good choice to perform full configuration and dynamic reconfiguration over the same interface.

Master modes are not directly supported because IPROG housecleaning clears the configuration memory.

Table 8-1: Supported Configuration Ports

Configuration Mode	7 Series	Zynq	UltraScale	UltraScale+	Zynq UltraScale+ MPSoC
JTAG	Yes	Yes	Yes	Yes	Yes
ICAP	Yes	Yes	Yes	Yes	Yes
PCAP	N/A	Yes	N/A	N/A	Yes
MCAP	N/A	N/A	Yes	Yes	Yes
Slave Serial	Yes	N/A	Yes	Yes	N/A
Slave SelectMap	Yes	N/A	Yes	Yes	N/A
SPI (any width)*	No	N/A	No	Yes	N/A
BPI sync mode*	No	N/A	No	Yes	N/A
BPI async mode	Yes	N/A	Yes	Yes	N/A
Master modes	No	N/A	No	No	N/A

*. SPI and BPI flash can be used to store partial bitstreams, but the STARTUP primitive cannot be used to deliver partial bitstreams to the configuration engine for devices prior to UltraScale+. The static design would need to be connected to the flash via user IO, and a controller would be used to fetch bitstreams for delivery to the ICAP.

To use external configuration modes (other than JTAG) for loading a partial BIT file, these pins must be reserved for use after the initial device configuration. This is achieved by using the `BITSTREAM.CONFIG.PERSIST` property to keep the dual-purpose I/O for configuration usage and to set the configuration width. Refer to this [link](#) in the *Vivado Design Suite User Guide: Programming and Debugging* (UG908) [Ref 27]. The Tcl command syntax to set this property is:

```
set_property BITSTREAM.CONFIG.PERSIST <value> [current_design]
```

where `<value>` is either `No` or `Yes`.

Note: When configuration pins are persisted, the ICAP is disabled; the two features are mutually exclusive. For more information on the ICAP, see *7 Series FPGAs Configuration User Guide* (UG470) [Ref 9] or *UltraScale Architecture Configuration User Guide* (UG570) [Ref 10], depending on your device.

Partial bitstreams contain all the configuration commands and data necessary for Dynamic Function eXchange. The task of loading a partial bitstream into an FPGA does not require knowledge of the physical location of the RM because configuration frame addressing information is included in the partial bitstream. A valid partial bitstream cannot be sent to the wrong part of the FPGA.

A DFX controller retrieves the partial bitstream from memory, then delivers it to a configuration port. The DFX control logic can either reside in an external device (for example, a processor) or in the programmable logic of the FPGA to be reconfigured. A user-designed internal DFX controller loads partial bitstreams through the ICAP interface. As with any other logic in the static design, the internal DFX control circuitry operates without interruption throughout the reconfiguration process.

Internal configuration can consist of either a custom state machine, or an embedded processor such as MicroBlaze™. For a Zynq-7000 SoC and MPSoC, the Processor Subsystem (PS) can be used to manage partial reconfiguration events.

Note: For Zynq-7000 SoC devices, the Programmable Logic (PL) can be partially reconfigured, but the Processing System cannot.

As an aid in debugging Dynamic Function eXchange designs and DFX control logic, the Vivado® Logic Analyzer can be used to load full and partial bitstreams into an FPGA by means of the JTAG port.

For more information on loading a bitstream into the configuration ports, see the Configuration Interfaces chapter in these documents:

- *7 Series FPGAs Configuration User Guide (UG470)* [Ref 9]
- *UltraScale Architecture Configuration User Guide (UG570)* [Ref 10]
- *Zynq-7000 All Programmable SoC Technical Reference Manual (UG585)* [Ref 11]

Bitstream Type Definitions

When designs are compiled for Dynamic Function eXchange in Xilinx devices, different types of bitstreams are created. This section defines terminology and explains the details for each type of bitstream for 7 series and UltraScale devices. The types of bitstreams are:

- [Full Configuration Bitstreams](#)
- [Partial Bitstreams](#)
- [Blanking Bitstreams](#)
- [Clearing Bitstreams](#)

Full Configuration Bitstreams

All DFX designs start with standard configuration of the full device using a full configuration bitstream. The format and structure is no different than for a flat design solution (with one exception), and there is no difference in how this bitstream can be used to initially program the FPGA. The one exception is that the global signal mask for a DFX design is closed; it is opened with each partial (or clearing) bitstream to affect only the reconfigurable region. Because of this exception, chip-wide GSR events (after the initial configuration) cannot be issued. However, note that the design itself has been processed in preparation for partial reconfiguration of the device after the full programming has been done. Standard features, such as encryption and compression, are supported.

Reconfigurable Partitions (RP) set as black boxes are supported, so Reconfigurable Modules (RM) with no functionality can be delivered as part of the initial configuration, to be replaced later with a desired Reconfigurable Module. Bitstream compression can be effective in this case, reducing bitstream size and initial configuration time.

Downloading a Full BIT File

The FPGA in a digital system is configured after power on reset by downloading a full BIT file, either directly from a PROM or from a general purpose memory space by a microprocessor. A full BIT file contains all the information necessary to reset the FPGA, configure it with a complete design, and verify that the BIT file is not corrupt. The figure below illustrates this process.




Figure 8-1: Configuring with a Full BIT File

After the initial configuration is completed and verified, the FPGA enters user mode, and the downloaded design begins functioning. If a corrupt BIT file is detected, the DONE signal is never asserted, the FPGA never enters user mode, and the corrupt design never starts functioning.

Partial Bitstreams

Partial bitstreams are delivered during normal device operation to replace functionality in a pre-defined device region. These bitstreams have the same structure as full bitstreams but are limited to specific address sets to program a specific portion of the device. Dedicated DFX features such as per-frame CRC checks (to ensure bitstream integrity) and automatic initialization (so the region starts in a known state) are available, as well as full bitstream features such as encryption and compression.

The size of a partial bitstream is directly proportional to the size of the region it is reconfiguring. For example, if the Reconfigurable Partition is composed of 20% of the device resources, the partial bitstream is roughly 20% the size of the full design bitstream.

Partial bitstreams are fully self-contained, so they are delivered to an appropriate configuration port. All addressing, header, and footer details are contained within these bitstreams, just as they would be for full configuration bitstreams. You deliver partial bitstreams to the FPGA through any external non-master configuration mode, such as JTAG, Slave Serial, or Slave SelectMap. Internal configuration access includes the ICAP (all devices), PCAP (Zynq-7000 SoC devices), and MCAP (UltraScale and UltraScale+ devices through PCIe).

Partial bitstreams are automatically created when `write_bitstream` is run on a DFX configuration. Each partial bitstream file name references your top-level design name, plus the pblock name for the Reconfigurable Partition, plus `_partial`. For example, for a full design bit file `top_first.bit`, a partial bit file could be named `top_first_pblock_red_partial.bit`.



RECOMMENDED: *The pblock instance is always the same, regardless of the RM contained within, so it is recommended that you use a descriptive base configuration name or rename the partial bit files to clarify which module it represents.*

Downloading a Partial BIT File

A partially reconfigured FPGA is in user mode while the partial BIT file is loaded. This allows the portion of the FPGA logic not being reconfigured to continue functioning while the reconfigurable portion is modified. Figure 8-2 illustrates this process.

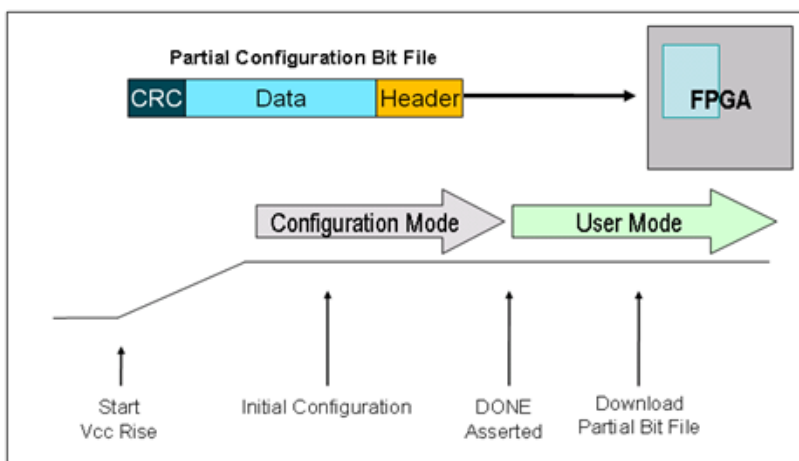


Figure 8-2: Configuring with a Partial BIT File

The partial BIT file has a simplified header, and there is no startup sequence that brings the FPGA into user mode. The BIT file contains (essentially, and with default settings) only frame address and configuration data, plus a final checksum value. Additional CRC checks can be inserted, if desired, to perform bitstream integrity checking.

If Reset After Reconfiguration is used, the DONE pin pulls LOW when reconfiguration begins and pulls HIGH again when partial reconfiguration successfully completes, although the partial bitstream can still be monitored internally as well.

Note: With UltraScale devices, the DONE pin pulls LOW at the beginning of the clearing bitstream and remains low until the end of the partial bitstream because the two bitstreams together constitute a complete partial reconfiguration sequence. The DONE pin does NOT return high at the end of the clearing bitstream.

If Reset After Reconfiguration is not selected, you must monitor the data being sent to know when configuration has completed. The end of a partial BIT file has a DESYNCH word (0000000D) that informs the configuration engine that the BIT file has been completely delivered. This word is given after a series of padding NO OP commands, ensuring that once the DESYNCH has been reached, all the configuration data has already been sent to the target frames throughout the device. As soon as the complete partial BIT file has been sent to the configuration port, it is safe to release the reconfiguration region for active use.

Blanking Bitstreams

A blanking bitstream is a specific type of partial bitstream, one that represents a logical black box. In Vivado this is referred to as a greybox, as it is not completely empty. It removes the functionality of an existing Reconfigurable Module by replacing it with new functionality, which consists simply of tie-off LUTs on all appropriate module I/O.

To create a greybox Reconfigurable Module, you remove the logical and physical representation of a fully placed and routed design configuration and replace it with tie-off LUTs. Starting with a routed configuration (with the static design locked) in active memory, run these steps:

```
update_design -cell <foo> -black_box
update_design -cell <foo> -buffer_ports
place_design
route_design
```

The design must be placed and routed to implement the LUTs that have been inserted into the design. Outputs of the greybox RM are tied to ground by default, but can be set to Vcc by setting the HD.PARTPIN_TIEOFF on desired ports.

Compression can be used to greatly reduce the size of blanking bitstreams. Note that these bitstreams still contain, not only the tie-off LUTs, but also any static routing that happens to pass through this region of the FPGA. Blanking bitstreams are generated and named in the same way as standard partial bitstreams, as the greybox variation is saved as another configuration checkpoint.

Advisories had been given for prior versions of Vivado software recommending the use of blanking bitstreams for 7 series and Zynq devices to avoid potential glitching conditions. Starting with Vivado 2016.1, these rare glitching scenarios are automatically avoided by embedding specific blanking events in each partial bitstream. Blanking bitstreams, while still available to remove logic from a Reconfigurable Partition, are no longer required to avoid any potential glitch events. Automatically embedding blanking events results in an increased size of the partial bit files; compression can be used to reduce these effects.

Clearing Bitstreams

Unlike the bitstream types noted above, this type is for UltraScale devices only (UltraScale+ does not have this requirement). A new requirement for this architecture is to clear an existing module before loading a new module. This clearing bitstream prepares the device for the delivery of any subsequent partial bitstream for that Reconfigurable Partition by establishing the global signal mask for the region to be reconfigured. Although the existing module is technically not removed (the current logical module remains), it is easiest to think of it this way. If a clearing bitstream is not delivered, the subsequent reconfigurable module will not be initialized.

Clearing bitstreams are *not* partial bitstreams. They comprise less than 10% of the frames for the target region and are therefore less than 10% the size of the corresponding partial bitstreams. They do not change the functionality but shut down clocks driving logic in the region. They must be delivered between partial bitstreams and should always be followed as soon as possible by the next partial bitstream.

Each clearing bitstream is built for a specific Reconfigurable Module and must be applied after that module has been used, and must be sent to the configuration engine immediately before the next partial bitstream is delivered. For example, to transition from module A to module B, the clearing bitstream for A must be delivered just before the partial bitstream for B is delivered. To transition from module B back to module A, the clearing bitstream for B must be delivered just before the partial bitstream for A is delivered. This is the case even if any partial bitstream in question is a blanking bitstream.

Clearing bitstreams are automatically generated and have the same name as partial bitstreams with `_clear` at the end. Looking at the example above, if `top_first` is an UltraScale device design, the clearing bit file name would be `top_first_pblock_red_partial_clear.bit`.

Dynamic Function eXchange through ICAP for Zynq Devices

The primary configuration mechanism for the programmable logic (PL) of Zynq-7000 and Zynq MPSoC devices is through the processing system (PS), which delivers bitstreams to the PCAP. The most common mechanism for partial reconfiguration is also through this path. However, to manage partial reconfiguration completely within the PL (either through the PR Controller IP or through a custom-designed controller module), partial bitstreams can also be delivered to the ICAP, just as they can be for FPGA devices.

The PCAP and ICAP interfaces are mutually exclusive and cannot be used simultaneously. Switching between ICAP and PCAP is possible, but you must ensure that no commands or data are being transmitted or received before changing interfaces. Failure to do this could lead to unexpected behavior.

To enable the ICAP for Zynq-7000 devices, set bit 27 (**PCAP_PR**) of the Control Register (**devc.CTRL**). This bit selects between ICAP and PCAP for PL reconfiguration. The default is **PCAP (1)**, but that can be changed to **ICAP (0)** to enable this configuration port. Note that bit 28 (**PCAP_MODE**) must also be set to 1, which is the default. For more details, see the *Zynq-7000 All Programmable SoC Technical Reference Manual* (UG585) [Ref 11].

To enable the ICAP for Zynq MPSoC devices, set the **PCAP_PR** field of the **pcap_ctrl1** (CSU) register. This bit selects between ICAP (or MCAP) and PCAP for PL reconfiguration. The default is **PCAP (1)**, but that can be changed to **ICAP / MCAP (0)** to enable this configuration port. For more details, see the *Zynq UltraScale+ MPSoC Technical Reference Manual* (UG1085) [Ref 32] and the *Zynq UltraScale+ MPSoC Register Reference* (UG1087) [Ref 33].

The Zynq UltraScale+ MPSoC Xilfpga library supports the delivery of partial bit files for Linux and bare-metal applications. In the current release, only non-secure bitstreams (without encryption or authentication) are supported. For more information and examples, visit the Xilfpga wiki page (www.wiki.xilinx.com/xilfpga).

Tandem Configuration and Dynamic Function eXchange

UltraScale devices introduced the MCAP, a dedicated connection from one specific PCIe block on a device to the configuration engine, providing an efficient mechanism for delivering partial bitstreams. No explicit routes are required to connect the PCIe block to the configuration engine, saving considerable resources.

The MCAP is enabled by customizing a Xilinx PCIe IP with Dynamic Function eXchange or Tandem Configuration features. These features are available for three IP cores that support PCI Express:

- *UltraScale Architecture Gen3 Integrated Block for PCI Express* (PG156) [Ref 16]
- *AXI Bridge for PCI Express Gen3 Subsystem* (PG194) [Ref 28]
- *DMA Subsystem for PCI Express* (PG195) [Ref 30]
- *UltraScale+ Devices Integrated Block for PCI Express* (PG213) [Ref 31]

Tandem Configuration utilizes a two-stage methodology that enables the IP to meet the configuration time requirements indicated in the PCI Express Specification. The following use cases are supported with this technology:

- **Tandem PROM:** Load the single two-stage bitstream from the flash.
- **Tandem PCIe:** Load the first stage bitstream from flash, and deliver the second stage bitstream over the PCIe link to the MCAP.
- **Tandem with Field Updates:** After a Tandem PROM (UltraScale only) or Tandem PCIe (UltraScale or UltraScale+) initial configuration, update the entire user design while the PCIe link remains active. The update region (floorplan) and design structure are pre-defined, and a Tcl scripts are provided.
- **Tandem + Dynamic Function eXchange:** This is a more general case of Tandem Configuration followed by Dynamic Function eXchange of any size or number of dynamic regions.
- **DFX over PCIe:** This is a standard configuration followed by DFX, using the PCIe / MCAP as the delivery path of partial bitstreams.

The Tandem and DFX combined solutions has few additional requirements. This approach requires that the Pblocks for the `HD.TANDEM_IP_PBLOCK` and `HD.RECONFIGURABLE` parts of the design do not overlap. Otherwise, like a standard DFX design, any number or size Reconfigurable Partitions can be defined.

To enable any of these capabilities, select the appropriate option when customizing the core. In the **Basic** tab:

1. Change the **Mode** to **Advanced**.
2. Change the **Tandem Configuration or Partial Reconfiguration** option according to your desired case in UltraScale:
 - **Tandem** for Tandem PROM, Tandem PCIe or Tandem + Dynamic Function eXchange use cases
 - **Tandem with Field Updates ONLY** for the pre-defined Field Updates use case
 - **PR over PCIe** to enable the MCAP link for Dynamic Function eXchange, without enabling Tandem Configuration

3. Change the **Tandem Configuration or Partial Reconfiguration** option according to your desired case in UltraScale+:
 - **Tandem PROM** for Tandem PROM or Tandem + Dynamic Function eXchange use cases
 - **Tandem PCIe** for Tandem PCIe or Tandem + Dynamic Function eXchange use cases
 - **Tandem PCIe with Field Updates ONLY** for the pre-defined Field Updates use case; Tandem PROM does not support Field Updates in UltraScale+
 - **PR over PCIe** to enable the MCAP link for Dynamic Function eXchange, without enabling Tandem Configuration

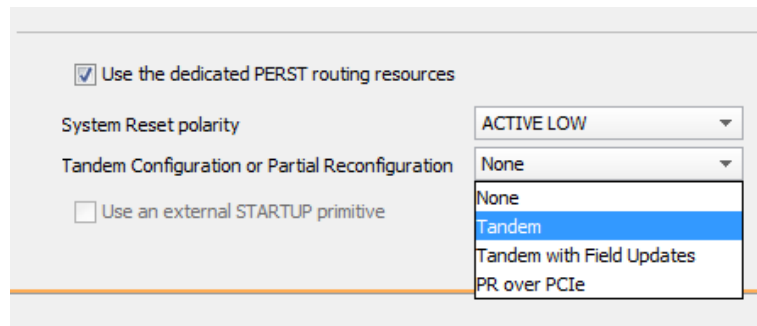


Figure 8-4: Customizing the Core in the Basic tab

The PCIe block that must be selected in most cases is the lowest instance in the device, except for SSI devices with three super logic regions (SLRs), in which case it is the lowest PCIe instance in the center SLR. A complete listing of the specific supported blocks is shown below in [Table 8-2](#). All other PCIe blocks do not have the dedicated MCAP feature.

For complete information about Tandem Configuration, including required PCIe block locations, design flow examples, requirements, restrictions and other considerations, see this [link](#) in *UltraScale Architecture Gen3 Integrated Block for PCI Express* (PG156) [Ref 16] for UltraScale devices. For UltraScale+ devices, see *UltraScale+ Devices Integrated Block for PCI Express* (PG213) [Ref 31].

Table 8-2: UltraScale: PCIe Block and Reset Locations Supporting DFX, by Device

Device	Package PCIe Block	PCIe Reset Location	Status
Kintex® UltraScale			
XCKU025	PCIE_3_1_X0Y0	IOB_X1Y103	Production
XCKU035	PCIE_3_1_X0Y0	IOB_X1Y103	Production
XCKU040	PCIE_3_1_X0Y0	IOB_X1Y103	Production
XCKU060	PCIE_3_1_X0Y0	IOB_X2Y103	Production
XCKU085	PCIE_3_1_X0Y0	IOB_X2Y103	Production
XCKU095	PCIE_3_1_X0Y0	IOB_X1Y103	Production

Table 8-2: UltraScale: PCIe Block and Reset Locations Supporting DFX, by Device (Cont'd)

Device	Package PCIe Block	PCIe Reset Location	Status
XCKU115	PCIE_3_1_X0Y0	IOB_X2Y103	Production
Virtex® UltraScale			
XCVU065	PCIE_3_1_X0Y0	IOB_X1Y103	Production
XCVU080	PCIE_3_1_X0Y0	IOB_X1Y103	Production
XCVU095	PCIE_3_1_X0Y0	IOB_X1Y103	Production
XCVU125	PCIE_3_1_X0Y0	IOB_X1Y103	Production
XCVU160	PCIE_3_1_X0Y1	IOB_X1Y363	Production
XCVU190	PCIE_3_1_X0Y2	IOB_X1Y363	Production
XCVU440	PCIE_3_1_X0Y2	IOB_X1Y363	Production

Table 8-3: UltraScale+: PCIe Block Locations Supporting DFX, by Device

Device	Package PCIe Block	Status*
Kintex® UltraScale+		
KU3P	PCIE40E4_X0Y0	Production
KU5P	PCIE40E4_X0Y0	Production
KU11P	PCIE40E4_X1Y0	Production
KU15P	PCIE40E4_X1Y0	Production
Virtex® UltraScale+		
VU3P	PCIE40E4_X1Y0	Production
VU5P	PCIE40E4_X1Y0	Production
VU7P	PCIE40E4_X1Y0	Production
VU9P	PCIE40E4_X1Y2	Production
VU11P	PCIE40E4_X0Y0	Production
VU13P	PCIE40E4_X0Y1	Production
VU27P	PCIE40E4_X0Y0	Production
VU29P	PCIE40E4_X0Y0	Production
VU31P	PCIE4CE4_X1Y0	Production
VU33P	PCIE4CE4_X1Y0	Production
VU35P	PCIE4CE4_X1Y0	Production
VU37P	PCIE4CE4_X1Y0	Production
Zynq MPSoC		
ZU4EV/EG/CC	PCIE40E4_X0Y1	Production
ZU5EV/EG/CC	PCIE40E4_X0Y1	Production
ZU7EV/EG/CC	PCIE40E4_X0Y1	Production

Table 8-3: UltraScale+: PCIe Block Locations Supporting DFX, by Device

Device	Package PCIe Block	Status*
ZU11EG	PCIE40E4_X1Y0	Production
ZU17EG	PCIE40E4_X1Y0	Production
ZU19EG	PCIE40E4_X1Y0	Production

*. For the AXI streaming core without Field Updates. For the most up-to-date information on core and device support status, consult the product guide for the specific version of the IP you wish to use.

Note: Any device not listed in this table does not have a PCIe site in the programmable logic portion of the device, or, like Zynq RFSoc devices, does not have an MCAP-enabled PCIe site in the programmable logic. Unlike UltraScale, UltraScale+ does not have a dedicated connection to a PCIe Reset pin, but Xilinx recommends using a pin in Bank 65.

The MCAP is capable of operating at 200 MHz with a 32-bit data path. Traditionally bitstreams are loaded into the MCAP from a host PC through PCI Express configuration packets. In these systems the host PC and host PC software are the main factors which limit MCAP performance and bitstream throughput. Because PCIe performance of specific host PC and host PC software can vary widely, overall MCAP performance throughput might vary.

For more information and sample drivers, see the answer record, *Bitstream Loading across the PCI Express Link in UltraScale Devices for Tandem PCIe and Partial Reconfiguration* (AR# 64761) [Ref 7].

If the performance of partial bitstream delivery via the MCAP port is insufficient, the ICAP can be used instead. While this approach does require additional logic to funnel configuration data from the PCIe end point to this internal configuration port, the ICAP can be saturated with 32-bit configuration data at the maximum clock rate (200MHz for monolithic devices, 125MHz for SSI devices). See XAPP1338 *Fast Partial Reconfiguration over PCI Express* for more information and an example design [Ref 38].

Formatting BIN Files for Delivery to Internal Configuration Ports

Partial bit files have the same basic format as full bit files, but they are reduced to the set of configuration frames for the target region and restricted to the set of events that make sense for active devices. Partial bit files can be:

- Delivered to external interfaces, such as JTAG or slave configuration ports.
- Reformatted as BIN files to be delivered to the internal configuration ports: ICAP (7 series or UltraScale devices), PCAP (Zynq devices only) or MCAP (UltraScale devices only).

Generate BIN files using the `write_cfgmem` utility. Three options are critical:

- Set `-format` as BIN to generate that file type.
- Use `-interface` to select the SelectMap width, and use **SMAPx32** for PCAP or **MCAP** for UltraScale ICAP.
 - **SMAPx16** and **SMAPx8** (default) can also be used for the 7 series ICAP.
 - **SMAPx8** is required for 7 series encrypted partial bitstreams.
- You must use `-disablebitswap` to target the PCAP or MCAP.

Examples

ICAP (for 7 series devices)

```
write_cfgmem -format BIN -interface SMAPx8 -loadbit "up 0x0 <partial_bitfile>"
```

ICAP (for UltraScale devices)

```
write_cfgmem -format BIN -interface SMAPx32 -loadbit "up 0x0 <partial_bitfile>"
```

PCAP (for Zynq-7000 SoC devices) or MCAP (for one specific PCIe block per UltraScale device)

```
write_cfgmem -format BIN -interface SMAPx32 -disablebitswap -loadbit "up 0x0  
<partial_bitfile>"
```

Summary of BIT Files for UltraScale Devices

This section applies specifically to UltraScale devices and does not apply to 7 series, UltraScale+, Zynq or Zynq MPSoC devices. With the finer granularity of global signals (that is, GSR) and the ability to reconfigure new element types, a new configuration process is necessary. Prior to loading in a partial bitstream for a new Reconfigurable Module, the existing Reconfigurable Module must be cleared. This clearing bitstream prepares the device for delivery of any subsequent partial bitstream for that Reconfigurable Partition by establishing the global signal mask for the region to be reconfigured. Although the existing module technically is not removed, it is easiest to think of it this way.

When running `write_bitstream` on a design configuration with Reconfigurable Partitions, a clearing BIT file per RP is created. For example, take a design in which two Reconfigurable Partitions (RP1 and RP2), with two Reconfigurable Modules each, A1 and B1 into RP1, and A2 and B2 into RP2, have been implemented. Two configurations (configA and configB) have been run through place and route, and `pr_verify` has passed. When bitstreams are generated, each configuration produces five bitstreams. For configA, these could be named:

- `configA.bit` - This is the full design bitstream that is used to configure the device from power-up. This contains the static design plus functions A1 and A2.

- **configA_RP1_A1_partial.bit** - This is the partial BIT file for function A1. This is loaded after another RM has been cleared from this Reconfigurable Partition.
- **configA_RP1_A1_partial_clear.bit** - This is the clearing BIT file for function A1. Before loading in any other partial BIT file into RP1 *after function A1*, this file must be loaded.
- **configA_RP2_A2_partial.bit** - This is the partial BIT file for function A2. This is loaded after another RM has been cleared from this Reconfigurable Partition.
- **configA_RP2_A2_partial_clear.bit** - This is the clearing BIT file for function A2. Before loading in any other partial BIT file into RP2 *after function A2*, this file must be loaded.

Likewise, configB produces five similar bitstreams:

- **configB.bit** - This is the full design bitstream that is used to configure the device from power-up. This contains the static design plus functions B1 and B2.
- **configB_RP1_B1_partial.bit** - This is the partial BIT file for function B1. This is loaded after another RM has been cleared from this Reconfigurable Partition.
- **configB_RP1_B1_partial_clear.bit** - This is the clearing BIT file for function B1. Before loading in any other partial BIT file into RP1 *after function B1*, this file must be loaded.
- **configB_RP2_B2_partial.bit** - This is the partial BIT file for function B2. This is loaded after another RM has been cleared from this Reconfigurable Partition.
- **configB_RP2_B2_partial_clear.bit** - This is the clearing BIT file for function B2. Before loading in any other partial BIT file into RP2 *after function B2*, this file must be loaded.

The sequence for any reconfiguration is to first load a clearing BIT file for a current Reconfigurable Module, immediately followed by a new Reconfigurable Module. For example, to transition Reconfigurable Partition RP1 from function A1 to function B1, first load the BIT file **configA_RP1_A1_partial_clear.bit**, then load **configB_RP1_B1_partial.bit**. The first bitstream prepares the region by opening the mask, and the second bitstream loads the new function, initializes only that region, then closes the mask.


If a clearing bit file is not loaded, initialization routines (GSR) have no effect. If a clearing file for a different Reconfigurable Partition is loaded, then that RP is initialized instead of the one that has been just reconfigured. If the incorrect clearing file for the proper RP is used, the current RM or possibly even the static design could be disrupted until the following partial bit file has been loaded.

System Design for Configuring an FPGA

A partial BIT file can be downloaded to the FPGA in the same manner as a full BIT file. An external microprocessor determines which partial BIT file should be downloaded, where it exists in an external memory space, and directs the partial BIT file to a standard FPGA configuration port such as JTAG, Select MAP or serial interface. The FPGA processes the partial BIT file correctly without any special instruction that it is receiving a partial BIT file.

It is common to assert the INIT or PROG signals on the FPGA configuration interface before downloading a full BIT file. This must not be done before downloading a partial BIT file, as that would indicate the delivery of a full BIT file, not a partial one.

Any indication to the working design that a partial BIT file will be sent (such as holding enable signals and disabling clocks) must be done in the design—and not by means of dedicated FPGA configuration pins. Figure 8-5 shows the process of configuring through a microprocessor.



X12033

Figure 8-5: Configuring Through a Microprocessor

In addition to the standard configuration interfaces, Dynamic Function eXchange supports configuration by means of the Internal Configuration Access Port (ICAP). The ICAP protocol is identical to SelectMAP and is described in the Configuration User Guide for the target device. The ICAP library primitive can be instantiated in the HDL description of the FPGA design, thus enabling analysis and control of the partial BIT file before it is sent to the configuration port. The partial BIT file can be downloaded to the FPGA through general

purpose I/O or gigabit transceivers and then routed to the ICAP in the FPGA programmable logic.

Rules for DFX ports and formats:

- Encrypted partial bitstreams can be delivered to any port, but only when the initial configuration was also encrypted. The same key must be used for all bitstreams.
- If the initial configuration of a device is encrypted, unencrypted partial bitstreams can be used only if they are delivered to the ICAP.
- Bitstream authentication for partial bitstreams is not supported for Virtex or Kintex devices, only for Zynq devices.
- The ICAP must be used with an 8-bit bus only for Dynamic Function eXchange for encrypted 7 series BIT files.
- Reconfiguration through external configuration ports is permitted only when bitstream readback security is not set to Level 2.

Partial BIT File Integrity

Error detection and recovery of partial BIT files have unique requirements compared to loading a full BIT file. If an error is detected in a full BIT file when it is being loaded into an FPGA, the FPGA never enters user mode. The error is detected after the corrupt design has been loaded into configuration memory, and specific signals are asserted to indicate an error condition. Because the FPGA never enters user mode, the corrupt design never becomes active. You must determine the system behavior for recovering from a configuration error such as downloading a different BIT file if the error condition is detected.

When you download partial BIT files, you cannot use this methodology for error detection and recovery. The FPGA is by definition already in user mode when the partial BIT file is loaded. Because the configuration circuitry supports error detection only after a BIT file has been loaded, a corrupt partial BIT file can become active, potentially damaging the FPGA if left operating for an extended period of time.

If a CRC error is detected during a partial reconfiguration, it asserts the INIT_B pin of the FPGA (INIT_B goes Low to indicate a CRC error). In UltraScale devices, this behavior is echoed on the PRERROR output pin of the ICAP. It is important to note that if a system monitors INIT_B for CRC errors during the initial configuration, a CRC error during a partial reconfiguration might trigger the same response. To detect the presence of a CRC error from within the FPGA, the CRC status can be monitored through the ICAP block. The Status Register (STAT) indicates that the partial BIT file has a CRC error, by asserting the CRC_ERROR flag (bit 0).

There are two types of partial BIT file errors to consider: data errors and address errors (the partial BIT file is essentially address and data information). Given that static routes are free to pass through reconfigurable regions, both types of errors can corrupt the static design, although the likelihood is very small. The only method for completely safe recovery is to download a new full BIT file to ensure the state of the static logic, which requires the entire FPGA to be reset.

Many systems do not need a complex recovery mechanism because resetting the entire FPGA is not critical, or the partial BIT file is stored locally. In that case, the chance of BIT file corruption is not appreciable. Systems in which the BIT files are at risk of becoming corrupted (such as sending the partial BIT file over a radio link) should use a dedicated silicon feature that avoids the problem.

The configuration engines of 7 series, UltraScale, and UltraScale+ FPGAs, as well as Zynq-7000 SoC devices, have the ability to perform a frame-by-frame CRC check and do not load a frame into the configuration memory if that CRC check fails. A failure is reported on the INIT_B pin (it is pulled Low) and gives you the opportunity to take the next steps: retry the partial bit file, fall back to a golden partial bit file, etc. The partially loaded reconfiguration region does not have valid programming in it, but the CRC check ensures the remainder of the device (static region and any other reconfigurable modules) stays operational while the system recovers from the error.

To enable this feature for these devices, set the **PerFrameCRC** property prior to running **write_bitstream**. The default is No, and Yes inserts the extra CRC checks. The size of an uncompressed bit file increases four to five percent with this option enabled. Note that this feature is not compatible with bitstream compression. No other specific design considerations are necessary to select this option, but your partial reconfiguration controller solution should be designed to choose the course of action should the INIT_B pin indicate a failure has occurred.

The syntax for setting the **PerFrameCRC** property is:

```
set_property bitstream.general.perFrameCRC yes [current_design]
```

This property inserts per-frame CRC checks in all bitstreams created from the current checkpoint, not just partial bitstreams. Full device bitstreams for the initial configuration of the device would also contain the extra CRC checks.

After a partial bit file has been loaded (with or without the per-frame CRC checks), the overall configuration of the device has changed. If the POST_CRC feature for SEU mitigation is enabled, the SEU mitigation engine automatically recalculates the embedded SEU CRC value after the partial bitstream has been loaded and after you have de-synced the configuration interface. Upon completion of the CRC recalibration, the FRAME_ECCE2 FRAME_VALID output toggles again to indicate that SEU detection has resumed.

Configuration Frames

All user-programmable features inside Xilinx FPGA and SoC devices are controlled by volatile memory cells that must be configured at power-up. These memory cells are collectively known as configuration memory. They define the LUT equations, signal routing, IOB voltage standards, and all other aspects of the design.

Xilinx FPGA and SoC architectures have configuration memory arranged in frames that are tiled about the device. These frames are the smallest addressable segments of the device configuration memory space, and all operations must therefore act upon whole configuration frames.

Reconfigurable Frames are built upon these configuration frames, and these are the minimum building blocks for performing dynamic reconfiguration.

- Base Regions in 7 series FPGAs are:
 - **CLB:** 50 high by 1 wide
 - **DSP48:** 10 high by 1 wide
 - **Block RAM:** 10 high by 1 wide
- Base Regions in UltraScale and UltraScale+ FPGAs are:
 - **CLB:** 60 high by 1 wide
 - **DSP48:** 24 high by 1 wide
 - **Block RAM:** 12 high by 1 wide
 - **I/O and Clocking:** 52 I/O (one bank), plus related XiPhy, MMCM, and PLL resources
 - **Gigabit Transceivers:** 4 high (one quad, plus related clocking resources)

Configuration Time

The speed of configuration is directly related to the size of the partial BIT file and the bandwidth of the configuration port. The different configuration ports in each device family have the maximum bandwidths shown in [Table 8-4](#), [Table 8-7](#), and [Table 8-8](#).

Table 8-4: Maximum Bandwidths for Configuration Ports in 7 Series Devices

Configuration Mode	Max Clock Rate	Data Width	Maximum Bandwidth
ICAP	100 MHz	32 bit	3.2 Gb/s
SelectMAP	100 MHz	32 bit	3.2 Gb/s

Table 8-4: Maximum Bandwidths for Configuration Ports in 7 Series Devices

Configuration Mode	Max Clock Rate	Data Width	Maximum Bandwidth
Serial Mode	100 MHz	1 bit	100 Mb/s
JTAG	66 MHz	1 bit	66 Mb/s

Table 8-5: Maximum Bandwidths for Configuration Ports in UltraScale Devices

Configuration Mode	Max Clock Rate	Data Width	Maximum Bandwidth
ICAP/MCAP	200 MHz	32 bit	6.4 Gb/s
ICAP/MCAP (SSI)*	125 MHz	32 bit	4.0 Gb/s
SelectMAP	125 MHz	32 bit	4.0 Gb/s
Serial Mode	150 MHz	1 bit	150 Mb/s
Serial (SSI Devices)	100 MHz	1 bit	100 Mb/s
JTAG	50 MHz	1 bit	50 Mb/s
JTAG (SSI Devices)	20 MHz	1 bit	20 Mb/s

*. When delivered to the master SLR to program any SLR; when programming to the same SLR as the ICAP, the faster rate can be used. Note that configuration clock frequency maximums may be less than these values depending on speed grade or operating voltage. Consult the Data Sheet for your target device for more information.

Table 8-6: Maximum Bandwidths for Configuration Ports in UltraScale+ Devices

Configuration Mode	Max Clock Rate	Data Width	Maximum Bandwidth
ICAP/MCAP	200 MHz	32 bit	6.4 Gb/s
ICAP/MCAP (SSI*)	125 MHz	32 bit	4.0 Gb/s
BPI	125 MHz	16 bit	2.0 Gb/s
QSPI	125 MHz	4 bit	500 Mb/s
SelectMAP	125 MHz	32 bit	4.0 Gb/s
Serial Mode	125 MHz	1 bit	125 Mb/s
JTAG	66 MHz	1 bit	66 Mb/s
JTAG (SSI Devices)	20 MHz	1 bit	20 Mb/s

*. When delivered to the master SLR to program any SLR; when programming to the same SLR as the ICAP, the faster rate can be used. Note that configuration clock frequency maximums may be less than these values depending on speed grade or operating voltage. Consult the Data Sheet for your target device for more information.

In addition to being reported by `write_bitstream`, The exact bitstream length is available in the created `.rbt` file by using the `-raw_bitfile` option for `write_bitstream`. Use this number along with the bandwidth to calculate the total configuration time. Here is an example of the header in a raw bit file:

```
Xilinx ASCII Bitstream
Created by Bitstream 2019.2
Design name: led_shift_count;UserID=0XFFFFFFFF
Architecture:kintex7
Part:          7k325tffg900
```

Date: Mon Mar 16 16:42:05 2015
 Bits: 1211072
 111111111111111111111111111111111111

Configuration Debugging

The ICAP interface can be used to monitor the configuration process when it is used as the port for delivering bitstreams. The "O" port of the ICAP block is a 32-bit bus, but only the lowest byte is used. The mapping of the lower byte is as follows:

Table 8-7: ICAP "O" Port Bits

ICAP "O" Port Bits	Status Bit	Meaning
O[7]	CFGERR_B	Configuration error (active-Low) 0 = A configuration error has occurred. 1 = No configuration error.
O[6]	DALIGN	Sync word received (active-High) 0 = No sync word received. 1 = Sync word received by interface logic.
O[5]	RIP	Readback in progress (active-High) 0 = No readback in progress. 1 = A readback is in progress.
O[4]	IN_ABORT_B	ABORT in progress (active-Low) 0 = Abort is in progress. 1 = No abort in progress.
O[3:0]	1	Reserved

The most significant nibble of this byte reports the status. These Status bits indicate whether the Sync word been received and whether a configuration error has occurred. The following table displays the values for these conditions.

Table 8-8: ICAP Sync Bits

O[7:4]	Sync Word?	CFGERR?
9	No Sync	No CFGERR
D	Sync	No CFGERR
5	Sync	CFGERR
1	No Sync	CFGERR

Note: In the above table, the entries in the first column are different depending on the board. For 7 series devices, they end with "F", so 9F, DF, etc. For UltraScale+ devices, they end with "B", so 9B, DB, etc.

Figure 8-6 shows a completed full configuration, followed by a partial reconfiguration with a CRC error, and finally a successful partial reconfiguration. Using the table above, and the description below, you can see how the “O” port of the ICAP can be used to monitor the configuration process. If a CRC error occurs, these signals can be used by a configuration state machine to recover from the error. These signals can also be used by Vivado Logic Analyzer to capture a configuration failure for debug purposes. With this information Vivado Logic Analyzer can also be used to capture the various points of a partial reconfiguration.

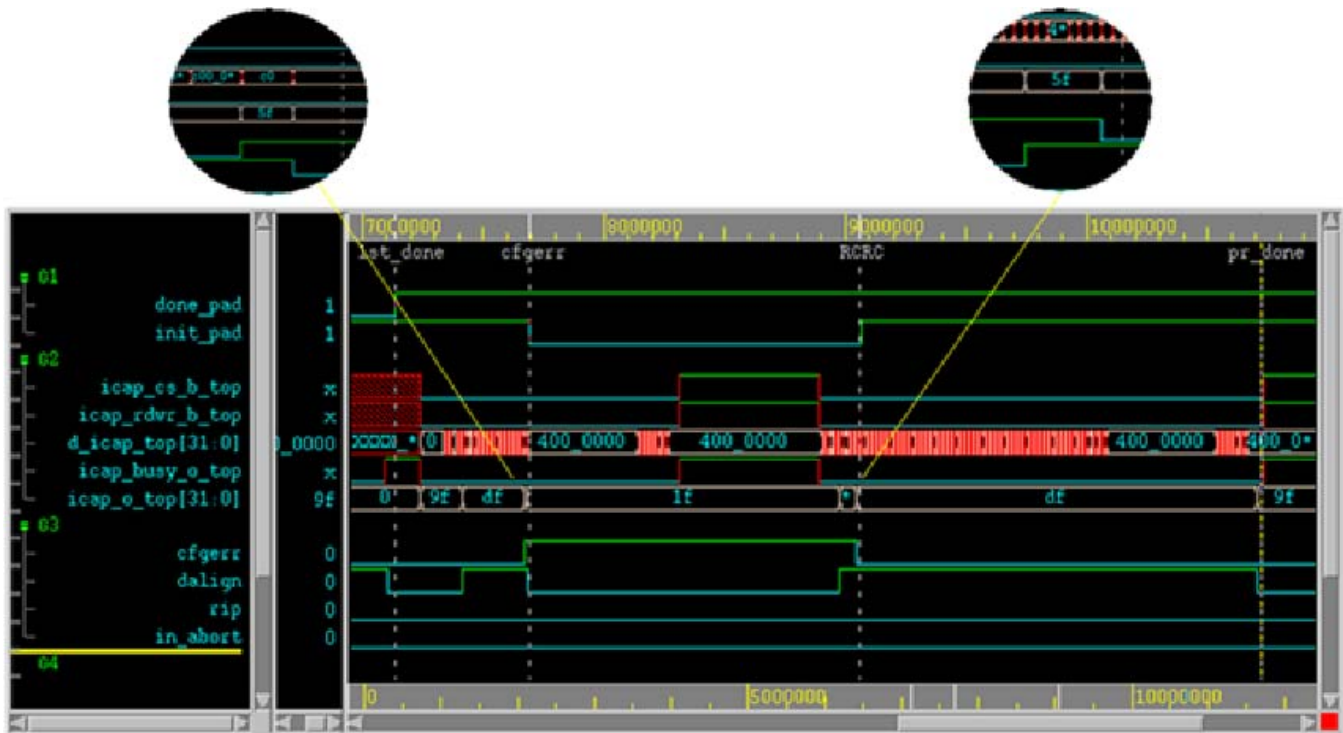


Figure 8-6: Vivado Logic Analyzer Display for Dynamic Function eXchange

The markers in the Vivado Logic Analyzer display indicate the following:

- **1st_done**

This marker indicates the completion of the initial full bitstream configuration. The DONE pin (`done_pad` in this waveform) goes HIGH.

- **cfgerr**

This marker indicates a CRC error is detected while loading partial bitstream. The status can be observed through O[31:0] (`icap_o_top [31:0]` in the waveform).

- `Icap_o_top [31:0]` starts at 0x9F
- After seen SYNC word, `Icap_o_top [31:0]` change to 0xDF

- After detect CRC error, `Icap_o_top[31:0]` change to 0x5F for one cycle, and then switches to 0x1F
- INIT_B pin is pulled Low (`init_pad` in the waveform)
- **RCRC**

This marker indicates when the partial bitstream is loaded again. The RCRC command resets the `cfgerr` status, and removes the pull-down on the INIT_B pin (`init_pad` in this waveform).

 - `Icap_o_top[31:0]` change from 0x1F to 0x5F when the SYNC word is seen
 - `Icap_o_top[31:0]` change from 0x5F to 0xDF when RCRC command is received
- **pr_done**

This marker indicates a successful partial reconfiguration.

 - `Icap_o_top[31:0]` change from 0xDF to 0x9F when the DESYNC command is received and no configuration error is detected.

In addition to the techniques described above, the UltraScale architecture introduced two new dedicated ports on the ICAP to aid in Dynamic Function eXchange:

- The PRDONE pin is intended to echo the external DONE pin. However, it should only be used for the following:
 - Monolithic devices
 - SSI devices when the RP is on the master SLR
 - When the ICAP used is on the same SLR as the RP.

Any partial bitstream that configures a slave SLR from the master ICAP will see multiple PRDONE events due to the construction of these partial bitstreams. Instead, use the EOS pin on the STARTUP block on the master SLR as a reliable indication of the completion of partial reconfiguration.

- The PRERROR pin echoes the external INIT_B pin. It drops LOW when a CRC error occurs, either with the standard full CRC value at the end of the bit file, or with any per-frame CRC value.

Using Vivado Debug Cores

Vivado debug cores (ILA, VIO, etc.) can be placed in any part of a Dynamic Function eXchange design, including within Reconfigurable Modules. A specific design methodology is necessary to connect these cores to a central Debug Hub for communication throughout the device.

The connectivity between the central Debug Hub in static can be set up automatically, and this is triggered by a specific naming convention for the port names on the Reconfigurable Partition. These twelve pins are required, and the hub will be inferred if these exact names are used. Examples in Verilog and VHDL are below.

Verilog instantiation in the static design:

```
my_count counter_inst (
    .clk(my_clk),
    .dout(dout),
    .S_BSCAN_drck(),
    .S_BSCAN_shift(),
    .S_BSCAN_tdi(),
    .S_BSCAN_update(),
    .S_BSCAN_sel(),
    .S_BSCAN_tdo(),
    .S_BSCAN_tms(),
    .S_BSCAN_tck(),
    .S_BSCAN_runtest(),
    .S_BSCAN_reset(),
    .S_BSCAN_capture(),
    .S_BSCAN_bscanid_en()
);
```

VHDL component declaration and instantiation in the static design:

```
component my_count is
    Port ( clk : in STD_LOGIC;
          dout : out STD_LOGIC;
          S_BSCAN_drck: IN std_logic := '0';
          S_BSCAN_shift: IN std_logic := '0';
          S_BSCAN_tdi: IN std_logic := '0';
          S_BSCAN_update: IN std_logic := '0';
          S_BSCAN_sel: IN std_logic := '0';
          S_BSCAN_tdo: OUT std_logic;
          S_BSCAN_tms: IN std_logic := '0';
          S_BSCAN_tck: IN std_logic := '0';
          S_BSCAN_runtest: IN std_logic := '0';
          S_BSCAN_reset: IN std_logic := '0';
          S_BSCAN_capture: IN std_logic := '0';
          S_BSCAN_bscanid_en: IN std_logic := '0'
        );
end component;

...
counter_inst: my_count
    port map (clk => my_clk,
             dout => dout,
             S_BSCAN_drck => open,
             S_BSCAN_shift => open,
             S_BSCAN_tdi => open,
             S_BSCAN_update => open,
             S_BSCAN_sel => open,
             S_BSCAN_tdo => open,
             S_BSCAN_tms => open,
             S_BSCAN_tck => open,
             S_BSCAN_runtest => open,
             S_BSCAN_reset => open,
```

```

        S_BSCAN_capture => open,
        S_BSCAN_bscanid_en => open
    );

```

Note: These input ports must receive an initial value to use the `open` keyword, and that initial value must be 0. Ports tied to 1 will not connect to the local debug hub.

Within the Reconfigurable Module top-level RTL, leave these twelve ports unconnected. Debug Hubs are inserted as black boxes, one in static and one in each Reconfigurable Module during synthesis. These inserted IP are then expanded during `opt_design`. This is done for each RM, even if there are no debug cores within that RM (including greybox RMs).

If these exact port names cannot be used for any reason, such as the need to explicitly connect to multiple BSCAN instances, attributes can be used to drive the Debug Hub insertion. This approach must also be used if the first configuration processed does not have any debug cores present; inference of Debug Hubs is not done if no debug cores within the RM can be found by Vivado implementation tools.

In the syntax shown below, do not change anything other than the port names in order to assign Debug ports. These attributes are to be used in all RM top level source files.

Verilog attributes in the RM top level:

```

(* X_INTERFACE_INFO = "xilinx.com:interface:bscan:1.0 S_BSCAN drck" *) (* DEBUG="true" *)
input my_drck;
(* X_INTERFACE_INFO = "xilinx.com:interface:bscan:1.0 S_BSCAN shift" *) (* DEBUG="true" *)
input my_shift;
(* X_INTERFACE_INFO = "xilinx.com:interface:bscan:1.0 S_BSCAN tdi" *) (* DEBUG="true" *)
input my_tdi;
(* X_INTERFACE_INFO = "xilinx.com:interface:bscan:1.0 S_BSCAN update" *) (* DEBUG="true" *)
input my_update;
(* X_INTERFACE_INFO = "xilinx.com:interface:bscan:1.0 S_BSCAN sel" *) (* DEBUG="true" *)
input my_sel;
(* X_INTERFACE_INFO = "xilinx.com:interface:bscan:1.0 S_BSCAN tdo" *) (* DEBUG="true" *)
output my_tdo;
(* X_INTERFACE_INFO = "xilinx.com:interface:bscan:1.0 S_BSCAN tms" *) (* DEBUG="true" *)
input my_tms;
(* X_INTERFACE_INFO = "xilinx.com:interface:bscan:1.0 S_BSCAN tck" *) (* DEBUG="true" *)
input my_tck;
(* X_INTERFACE_INFO = "xilinx.com:interface:bscan:1.0 S_BSCAN runttest" *) (* DEBUG="true" *)
input my_runttest;
(* X_INTERFACE_INFO = "xilinx.com:interface:bscan:1.0 S_BSCAN reset" *) (* DEBUG="true" *)
input my_reset;
(* X_INTERFACE_INFO = "xilinx.com:interface:bscan:1.0 S_BSCAN capture" *) (* DEBUG="true" *)
input my_capture;
(* X_INTERFACE_INFO = "xilinx.com:interface:bscan:1.0 S_BSCAN bscanid_en" *)
(* DEBUG="true" *) input my_bscanid_en;

```

VHDL attributes in the RM top level:

```
attribute X_INTERFACE_INFO : string;
attribute DEBUG : string;
attribute X_INTERFACE_INFO of my_drck: signal is "xilinx.com:interface:bscan:1.0 S_BSCAN
drck";
attribute DEBUG of my_drck: signal is "true";
attribute X_INTERFACE_INFO of my_shift: signal is "xilinx.com:interface:bscan:1.0 S_BSCAN
shift";
attribute DEBUG of my_shift: signal is "true";
attribute X_INTERFACE_INFO of my_tdi: signal is "xilinx.com:interface:bscan:1.0 S_BSCAN
tdi";
attribute DEBUG of my_tdi: signal is "true";
attribute X_INTERFACE_INFO of my_update: signal is "xilinx.com:interface:bscan:1.0 S_BSCAN
update";
attribute DEBUG of my_update: signal is "true";
attribute X_INTERFACE_INFO of my_sel: signal is "xilinx.com:interface:bscan:1.0 S_BSCAN
sel";
attribute DEBUG of my_sel: signal is "true";
attribute X_INTERFACE_INFO of my_tdo: signal is "xilinx.com:interface:bscan:1.0 S_BSCAN
tdo";
attribute DEBUG of my_tdo: signal is "true";
attribute X_INTERFACE_INFO of my_tms: signal is "xilinx.com:interface:bscan:1.0 S_BSCAN
tms";
attribute DEBUG of my_tms: signal is "true";
attribute X_INTERFACE_INFO of my_tck: signal is "xilinx.com:interface:bscan:1.0 S_BSCAN
tck";
attribute DEBUG of my_tck: signal is "true";
attribute X_INTERFACE_INFO of my_runttest: signal is "xilinx.com:interface:bscan:1.0 S_BSCAN
runttest";
attribute DEBUG of my_runttest: signal is "true";
attribute X_INTERFACE_INFO of my_reset: signal is "xilinx.com:interface:bscan:1.0 S_BSCAN
reset";
attribute DEBUG of my_reset: signal is "true";
attribute X_INTERFACE_INFO of my_capture: signal is "xilinx.com:interface:bscan:1.0 S_BSCAN
capture";
attribute DEBUG of my_capture: signal is "true";
attribute X_INTERFACE_INFO of my_bscanid_en: signal is "xilinx.com:interface:bscan:1.0
S_BSCAN bscanid_en";
attribute DEBUG of my_bscanid_en: signal is "true";
```

There are currently two limitations to this solution:

1. All debug cores within Reconfigurable Modules must be instantiated. The MARK_DEBUG insertion flow is not yet supported.
2. A greybox configuration cannot be the first one processed. The debug bridge must be established within a Reconfigurable Module with debug cores to establish connectivity with the debug hub before moving to versions that do not contain debug cores.

For an example of this core insertion as well as functionality within the Vivado Hardware Manager, see this [link](#) in *Vivado Design Suite Tutorial: Dynamic Function eXchange* (UG947) [Ref 1].

Known Issues and Limitations

Known Issues

This is a list of issues that might be encountered when using Dynamic Function eXchange (DFX) in the current Vivado® Design Suite release. If you encounter any of these issues, or discover any others, contact Xilinx® Support and send an example design that shows the issue. These test cases are very helpful to improve the overall solution.

Report to Xilinx all cases of fatal or internal errors, incomplete routing (partial antennas), or other rule violations that prevent place and route, `pr_verify`, and `write_bitstream` from succeeding. Including a design showing the failure is critical for proper analysis and implementation of fixes.

- If the initial configuration of a 7 series SSI device (7V2000T, 7VX1140T) is done through an SPI interface, partial bitstreams cannot be delivered to the master (or any) ICAP; they must be delivered to an external port, such as JTAG. If the initial configuration is done through any other configuration port, the master ICAP can be used as the delivery port for partial bitstreams. Contact Xilinx Support for a workaround.
- Do not drive multiple outputs of a single reconfigurable module with the same source. Each output of an RM must have a unique driver.
- When using Virtex UltraScale+ VU29P devices, connections between the IBUFDS_GTM and GTM_DUAL sites might be unroutable if the placer does not place them on the same SLR and the same side of the device. You might encounter route_design Route 35-7 in this case. If this occurs, you must LOC both the IBUFDS_GTM and GTM_DUAL instances to appropriate locations in the same SLR on the same side of the device.
- Engineering Silicon (ES) for UltraScale or UltraScale+ devices does not officially support Dynamic Function eXchange. To investigate the capabilities of DFX on ES devices, contact Xilinx Support for advice.

Known Limitations

Certain features are not yet developed or supported in the current release. These include:

- When selecting Pblock ranges to define the size and shape of the Reconfigurable Partition, do not use the CLOCKREGION resource type for 7 series or Zynq-7000 designs. Pblock ranges must only include types SLICE, RAMB18, RAMB36, and DSP48 resource types.
- Do not use Vivado Debug core insertion features within Reconfigurable Partitions. This flow inserts the debug hub, which includes BSCAN primitive, which is not permitted inside reconfigurable bitstreams. Vivado Debug cores must be instantiated or included within IP, then the Debug Hub can be inferred as described in [Using Vivado Debug Cores](#).
- UpdateMEM does not support partial or Tandem bitstreams. To associate memory files with DFX designs, the ELF Association flow must be applied. See this [link](#) in the *Vivado Design Suite User Guide: Designing IP Subsystems Using IP Integrator* (UG994) [Ref 29] for details.
- The Soft Error Mitigation (SEM) IP core is supported in conjunction with DFX in monolithic devices. For UltraScale devices, the SEM IP core is not supported when using Dynamic Function eXchange on SSI devices. For UltraScale+ devices, the SEM IP core is supported when using DFX on SSI devices. For more information on using the SEM IP in DFX designs, see *Demonstration of Soft Error Mitigation IP and Partial Reconfiguration Capability on Monolithic Devices* (XAPP1261) [Ref 6].
- The STARTUP primitive does not support loading of partial bitstreams for 7 series and UltraScale devices, as its clock will stop once a partial bitstream enters the configuration engine. IP, such as the AXI SPI IP or the AXI EMC IP, should not be configured to use the STARTUP primitive to clock or deliver partial bitstreams from external flash. For these architectures, partial bitstreams may be stored in BPI or SPI flash, but they must be moved to DDR or another location before being shifted into the ICAP.
- Two use cases regarding encryption will not be supported when using new features within UltraScale and UltraScale+ devices:
 - a. If RSA authentication is selected for the initial configuration, then encrypted partial reconfiguration is not supported. RSA authentication is not supported on FPGAs for partial bitstreams.
 - b. If the initial configuration bitstream uses an obfuscated AES-256 key stored in either the eFUSE or BBRAM, then any encrypted partial bitstreams must use the same obfuscated key. Encrypted partial bitstreams using a different key than the initial bitstream is not supported.

In either of these two cases, an unencrypted partial bitstream may be delivered to the ICAP to partially reconfigure the device.

- Bitstream compression and per-frame CRC checks cannot be enabled at the same time for a partial bitstream.
- The `update_design` command in general permits multiple targets for the `-cells` switch. However, when using this command for PR designs (for use with `-black_box` or `-buffer_ports`), specify one cell (Reconfigurable Partition) at a time. Performing these actions on more than RP requires multiple calls to `update_design`.
- Cascaded global clocking buffers across RM boundary is not a supported use case and is not guaranteed to be successfully routed. If cascaded BUFS are unavoidable in the design, it is recommended to keep them both, either in static or reconfigurable partition.

Hierarchical Design Flows

Overview

The Dynamic Function eXchange (DFX) flow is an ideal hierarchal design tool for a top-down in-context use case. The ability to implement the static portion of the design, meet timing, and then reuse those results meets the needs of other hierarchical design (HD) flows.

Platform Reuse takes advantage of the tool's ability to separate out the top-level (Platform) from the lower level partitions. Platform Reuse can be used for a number of use cases including:

- Reduced design cycle and timing closure. By closing timing on the top-level, interface logic that changes once (such as memory controllers, networking IP, high speed interconnect) and then locking down and reusing the placed/routed result for the Platform which partition logic is being developed.
- Platform delivery to end users. In this case a Platform can be developed by one user, and then fully placed, routed, and locked DCP (with black boxes for partitions) and be delivered to another user for development of customer partition logic.

Both of these flows use the Dynamic Function eXchange flow, to implement the initial design, carve out the partitions, and lock down the Static/Platform logic. All features of DFX (like greybox support) can be used, and all requirements of the DFX flow must be followed. From the tool's perspective, there is no difference between the flows, and all DFX DRCs will be applied.

Since partial bit files are not required for this flow, Xilinx recommends using the `-no_partial_bitfile` switch of the `write_bitstream` command to avoid producing partial bitstreams.

Additional Resources and Legal Notices

Xilinx Resources

For support resources such as Answers, Documentation, Downloads, and Forums, see [Xilinx Support](#).

Solution Centers

See the [Xilinx Solution Centers](#) for support on devices, software tools, and intellectual property at all stages of the design cycle. Topics include design assistance, advisories, and troubleshooting tips.

Documentation Navigator and Design Hubs

Xilinx Documentation Navigator provides access to Xilinx documents, videos, and support resources, which you can filter and search to find information. To open the Xilinx Documentation Navigator (DocNav):

- From the Vivado IDE, select **Help > Documentation and Tutorials**.
- On Windows, select **Start > All Programs > Xilinx Design Tools > DocNav**.
- At the Linux command prompt, enter `docnav`.

Xilinx Design Hubs provide links to documentation organized by design tasks and other topics, which you can use to learn key concepts and address frequently asked questions. To access the Design Hubs:

- In the Xilinx Documentation Navigator, click the **Design Hubs View** tab.
- On the Xilinx website, see the [Design Hubs](#) page.

Note: For more information on Documentation Navigator, see the [Documentation Navigator](#) page on the Xilinx website.

References

1. *Vivado Design Suite Tutorial: Dynamic Function eXchange* ([UG947](#))
2. *Partial Reconfiguration Controller LogiCORE IP Product Guide* ([PG193](#))
3. *Partial Reconfiguration Decoupler LogiCORE IP Product Guide* ([PG227](#))
4. *Partial Reconfiguration Bitstream Monitor Product Guide* ([PG304](#))
5. *Partial Reconfiguration AXI Shutdown Manager* ([PG305](#))
6. *Demonstration of Soft Error Mitigation IP and Partial Reconfiguration Capability on Monolithic Devices* ([XAPP1261](#))
7. *Bitstream Loading across the PCI Express Link in UltraScale Devices for Tandem PCIe and Partial Reconfiguration* ([AR# 64761](#))
8. *Partial Reconfiguration of a Hardware Accelerator with Vivado Design Suite for Zynq-7000 SoC Processor* ([XAPP1231](#))
9. *7 Series FPGAs Configuration User Guide* ([UG470](#))
10. *UltraScale Architecture Configuration User Guide* ([UG570](#))
11. *Zynq-7000 SoC Technical Reference Manual* ([UG585](#))
12. *Partial Reconfiguration User Guide* ([UG702](#)) - For ISE Design Tools
13. *UltraFast Design Methodology Guide for the Vivado Design Suite* ([UG949](#))
14. *7 Series FPGAs Integrated Block for PCI Express LogiCORE IP Product Guide* ([PG054](#))
15. *Virtex-7 FPGA Gen3 Integrated Block for PCI Express Product Guide* ([PG023](#))
16. *UltraScale FPGAs Gen3 Integrated Block for PCI Express LogiCORE IP Product Guide* ([PG156](#))
17. *Vivado Design Suite Tcl Command Reference Guide* ([UG835](#))
18. *Vivado Design Suite User Guide: Synthesis* ([UG901](#))
19. *Vivado Design Suite User Guide: Using Constraints* ([UG903](#))
20. *Vivado Design Suite User Guide: Design Analysis and Closure Techniques* ([UG906](#))
21. *7 Series FPGAs GTX/GTH Transceivers User Guide* ([UG476](#))
22. *7 Series FPGAs GTP Transceivers User Guide* ([UG482](#))
23. *MMCM and PLL Dynamic Reconfiguration (7 Series)* ([XAPP888](#))
24. *UltraScale Architecture Clocking Resources User Guide* ([UG572](#))
25. *UltraScale Architecture GTH Transceivers User Guide* ([UG576](#))
26. *UltraScale Architecture GTY Transceivers User Guide* ([UG578](#))

27. *Vivado Design Suite User Guide: Programming and Debugging* ([UG908](#))
 28. *AXI Bridge for PCI Express Gen3 Subsystem Product Guide* ([PG194](#))
 29. *Vivado Design Suite User Guide: Designing IP Subsystems Using IP Integrator* ([UG994](#))
 30. *DMA/Bridge Subsystem for PCI Express Product Guide* ([PG195](#))
 31. *UltraScale+ Devices Integrated Block for PCI Express LogiCORE IP Product Guide* ([PG213](#))
 32. *Zynq UltraScale+ MPSoC Technical Reference Manual* ([UG1085](#))
 33. *Zynq UltraScale+ MPSoC Register Reference* ([UG1087](#))
 34. *Bitstream Identification with `USR_ACCESS` using the Vivado Design Suite* ([XAPP1232](#))
 35. *Local Partial Reconfiguration Using Embedded Processing for 3D ICs* ([XAPP1099](#))
 36. *Design Advisory for Techniques on Properly Synchronizing Flip-Flops and SRLs* ([AR#44174](#))
 37. [Vivado Design Suite Documentation](#)
 38. *Fast Partial Reconfiguration PCI Express* ([XAPP1338](#))
-

Training Resources

Xilinx provides a variety of training courses and QuickTake videos to help you learn more about the concepts presented in this document. Use these links to explore related training resources:

1. [Vivado Design Suite QuickTake Video Tutorials](#)
 2. [Vivado Design Suite QuickTake Video: Partial Reconfiguration in Vivado](#)
 3. [Vivado Design Suite QuickTake Video: Partial Reconfiguration for UltraScale](#)
 4. [Vivado Design Suite QuickTake Video: Partial Reconfiguration for UltraScale+](#)
 5. [Partial Reconfiguration Flow on Zynq using Vivado](#)
 6. [Xilinx Partial Reconfiguration Tools and Techniques](#)
-

Please Read: Important Legal Notices

The information disclosed to you hereunder (the "Materials") is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available "AS IS" and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special,

incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of Xilinx's limited warranty, please refer to Xilinx's Terms of Sale which can be viewed at <https://www.xilinx.com/legal.htm#tos>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in such critical applications, please refer to Xilinx's Terms of Sale which can be viewed at <https://www.xilinx.com/legal.htm#tos>.

AUTOMOTIVE APPLICATIONS DISCLAIMER

AUTOMOTIVE PRODUCTS (IDENTIFIED AS "XA" IN THE PART NUMBER) ARE NOT WARRANTED FOR USE IN THE DEPLOYMENT OF AIRBAGS OR FOR USE IN APPLICATIONS THAT AFFECT CONTROL OF A VEHICLE ("SAFETY APPLICATION") UNLESS THERE IS A SAFETY CONCEPT OR REDUNDANCY FEATURE CONSISTENT WITH THE ISO 26262 AUTOMOTIVE SAFETY STANDARD ("SAFETY DESIGN"). CUSTOMER SHALL, PRIOR TO USING OR DISTRIBUTING ANY SYSTEMS THAT INCORPORATE PRODUCTS, THOROUGHLY TEST SUCH SYSTEMS FOR SAFETY PURPOSES. USE OF PRODUCTS IN A SAFETY APPLICATION WITHOUT A SAFETY DESIGN IS FULLY AT THE RISK OF CUSTOMER, SUBJECT ONLY TO APPLICABLE LAWS AND REGULATIONS GOVERNING LIMITATIONS ON PRODUCT LIABILITY.

© Copyright 2012–2020 Xilinx, Inc. Xilinx, the Xilinx logo, Artix, ISE, Kintex, Spartan, Virtex, Vivado, Zynq, and other designated brands included herein are trademarks of Xilinx in the United States and other countries. PCI, PCIe and PCI Express are trademarks of PCI-SIG and used under license. All other trademarks are the property of their respective owners.