
PathWave Test Sync Executive 2023

User Manual

Notices

Copyright Notice

© Keysight Technologies 2020-2023

No part of this manual may be reproduced in any form or by any means (including electronic storage and retrieval or translation into a foreign language) without prior agreement and written consent from Keysight Technologies, Inc. as governed by United States and international copyright laws.

Manual Part Number

KS2201-90000

Published By

Keysight Technologies
1400 Fountaingrove Parkway
Santa Rosa,
CA 95403-1738

Edition

Edition 2023_U0_00, June, 2023
Keysight Technologies, USA

Regulatory Compliance

This product has been designed and tested in accordance with accepted industry standards, and has been supplied in a safe condition. To review the Declaration of Conformity, go to <http://www.keysight.com/go/conformity>.

Warranty

THE MATERIAL CONTAINED IN THIS DOCUMENT IS PROVIDED "AS IS," AND IS SUBJECT TO BEING CHANGED, WITHOUT NOTICE, IN FUTURE EDITIONS. FURTHER, TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, KEYSIGHT DISCLAIMS ALL WARRANTIES, EITHER EXPRESS OR IMPLIED, WITH REGARD TO THIS MANUAL AND ANY INFORMATION CONTAINED HEREIN, INCLUDING BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF

MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. KEYSIGHT SHALL NOT BE LIABLE FOR ERRORS OR FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, USE, OR PERFORMANCE OF THIS DOCUMENT OR OF ANY INFORMATION CONTAINED HEREIN. SHOULD KEYSIGHT AND THE USER HAVE A SEPARATE WRITTEN AGREEMENT WITH WARRANTY TERMS COVERING THE MATERIAL IN THIS DOCUMENT THAT CONFLICT WITH THESE TERMS, THE WARRANTY TERMS IN THE SEPARATE AGREEMENT SHALL CONTROL.

KEYSIGHT TECHNOLOGIES DOES NOT WARRANT THIRD-PARTY SYSTEM-LEVEL (COMBINATION OF CHASSIS, CONTROLLERS, MODULES, ETC.) PERFORMANCE, SAFETY, OR REGULATORY COMPLIANCE, UNLESS SPECIFICALLY STATED.

Technology Licenses

The hardware and/or software described in this document are furnished under a license and may be used or copied only in accordance with the terms of such license.

U.S. Government Rights

The Software is "commercial computer software," as defined by Federal Acquisition Regulation ("FAR") 2.101. Pursuant to FAR 12.212 and 27.405-3 and Department of Defense FAR Supplement ("DFARS") 227.7202, the U.S. government acquires commercial computer software under the same terms by which the software is customarily provided to the public. Accordingly, Keysight provides the Software to U.S. government customers under its standard commercial license, which is embodied in its End User License Agreement (EULA), a copy of which can be found at <http://www.keysight.com/find/sweula>. The license set forth in

the EULA represents the exclusive authority by which the U.S. government may use, modify, distribute, or disclose the Software. The EULA and the license set forth therein, does not require or permit, among other things, that Keysight: (1) Furnish technical information related to commercial computer software or commercial computer software documentation that is not customarily provided to the public; or (2) Relinquish to, or otherwise provide, the government rights in excess of these rights customarily provided to the public to use, modify, reproduce, release, perform, display, or disclose commercial computer software or commercial computer software documentation. No additional government requirements beyond those set forth in the EULA shall apply, except to the extent that those terms, rights, or licenses are explicitly required from all providers of commercial computer software pursuant to the FAR and the DFARS and are set forth specifically in writing elsewhere in the EULA. Keysight shall be under no obligation to update, revise or otherwise modify the Software. With respect to any technical data as defined by FAR 2.101, pursuant to FAR 12.211 and 27.404.2 and DFARS 227.7102, the U.S. government acquires no greater than Limited Rights as defined in FAR 27.401 or DFAR 227.7103-5 (c), as applicable in any technical data.

Safety Notices

CAUTION

A CAUTION notice denotes a hazard. It calls attention to an operating procedure, practice, or the like that, if not correctly performed or adhered to, could result in damage to the product or loss of important data. Do not proceed beyond a CAUTION notice until the indicated conditions are fully understood and met.

WARNING

A WARNING notice denotes a hazard. It calls attention to an operating procedure, practice, or the like that, if not correctly performed or adhered to, could result in personal injury or death. Do not proceed beyond a WARNING notice until the indicated conditions are fully understood and met.

The following safety precautions should be observed before using this product and any associated instrumentation.

This product is intended for use by qualified personnel who recognize shock hazards and are familiar with the safety precautions required to avoid possible injury. Read and follow all installation, operation, and maintenance information carefully before using the product.

WARNING

If this product is not used as specified, the protection provided by the equipment could be impaired. This product must be used in a normal condition (in which all means for protection are intact) only.

The types of product users are:

- Responsible body is the individual or group responsible for the use and maintenance of equipment, for ensuring that the equipment is operated within its specifications and operating limits, and for ensuring operators are adequately trained.
- Operators use the product for its intended function. They must be trained in electrical safety procedures and proper use of the instrument. They must be protected from electric shock and contact with hazardous live circuits.
- Maintenance personnel perform routine procedures on the product to keep it operating properly (for example, setting the line voltage or replacing consumable materials). Maintenance procedures are

described in the user documentation. The procedures explicitly state if the operator may perform them. Otherwise, they should be performed only by service personnel.

- Service personnel are trained to work on live circuits, perform safe installations, and repair products. Only properly trained service personnel may perform installation and service procedures.

WARNING

Operator is responsible to maintain safe operating conditions. To ensure safe operating conditions, modules should not be operated beyond the full temperature range specified in the Environmental and physical specification. Exceeding safe operating conditions can result in shorter lifespans, improper module performance and user safety issues. When the modules are in use and operation within the specified full temperature range is not maintained, module surface temperatures may exceed safe handling conditions which can cause discomfort or burns if touched. In the event of a module exceeding the full temperature range, always allow the module to cool before touching or removing modules from chassis.

Keysight products are designed for use with electrical signals that are rated Measurement Category I and Measurement Category II, as described in the International Electrotechnical Commission (IEC) Standard IEC 60664. Most measurement, control, and data I/O signals are Measurement Category I and must not be directly connected to mains voltage or to voltage sources with high transient over-voltages. Measurement Category II connections require protection for high transient over-voltages often associated with local AC mains connections.

Assume all measurement, control, and data I/O connections are for connection to

Category I sources unless otherwise marked or described in the user documentation.

Exercise extreme caution when a shock hazard is present. Lethal voltage may be present on cable connector jacks or test fixtures. The American National Standards Institute (ANSI) states that a shock hazard exists when voltage levels greater than 30V RMS, 42.4V peak, or 60VDC are present. A good safety practice is to expect that hazardous voltage is present in any unknown circuit before measuring.

Operators of this product must be protected from electric shock at all times. The responsible body must ensure that operators are prevented access and/or insulated from every connection point. In some cases, connections must be exposed to potential human contact. Product operators in these circumstances must be trained to protect themselves from the risk of electric shock. If the circuit is capable of operating at or above 1000V, no conductive part of the circuit may be exposed.

Do not connect switching cards directly to unlimited power circuits. They are intended to be used with impedance-limited sources. NEVER connect switching cards directly to AC mains. When connecting sources to switching cards, install protective devices to limit fault current and voltage to the card.

Before operating an instrument, ensure that the line cord is connected to a properly-grounded power receptacle. Inspect the connecting cables, test leads, and jumpers for possible wear, cracks, or breaks before each use.

When installing equipment where access to the main power cord is restricted, such as rack mounting, a separate main input power disconnect device must be provided in close proximity to the equipment and within easy reach of the operator.

For maximum safety, do not touch the product, test cables, or any other instruments while power is applied to the circuit under test. ALWAYS remove power from the entire test system and discharge any capacitors before: connecting or disconnecting cables or jumpers, installing or removing switching cards, or making internal changes, such as installing or removing jumpers.

Do not touch any object that could provide a current path to the common side of the circuit under test or power line (earth) ground. Always make measurements with dry hands while standing on a dry, insulated surface capable of withstanding the voltage being measured.

The instrument and accessories must be used in accordance with its specifications and operating instructions, or the safety of the equipment may be impaired.

Do not exceed the maximum signal levels of the instruments and accessories, as defined in the specifications and operating information, and as shown on the instrument or test fixture panels, or switching card.

When fuses are used in a product, replace with the same type and rating for continued protection against fire hazard.

Chassis connections must only be used as shield connections for measuring circuits, NOT as safety earth ground connections.

If you are using a test fixture, keep the lid closed while power is applied to the device under test. Safe operation requires the use of a lid interlock.

Instrumentation and accessories shall not be connected to humans.

Before performing any maintenance, disconnect the line cord and all test cables.

To maintain protection from electric shock and fire, replacement components in mains

circuits – including the power transformer, test leads, and input jacks – must be purchased from Keysight. Standard fuses with applicable national safety approvals may be used if the rating and type are the same.

Other components that are not safety-related may be purchased from other suppliers as long as they are equivalent to the original component (note that selected parts should be purchased only through Keysight to maintain accuracy and functionality of the product). If you are unsure about the applicability of a replacement component, call an Keysight office for information.

WARNING

No operator serviceable parts inside. Refer servicing to qualified personnel. To prevent electrical shock do not remove covers. For continued protection against fire hazard, replace fuse with same type and rating.

PRODUCT MARKINGS:



The CE mark is a registered trademark of the European Community.



Australian Communication and Media Authority mark to indicate regulatory compliance as a registered supplier.



This symbol indicates product compliance with the Canadian Interference-Causing Equipment Standard (ICES-001). It also identifies the product is an Industrial Scientific and Medical Group 1 Class A product (CISPR 11, Clause 4).



South Korean Class A EMC Declaration. This equipment is Class A suitable for professional use and is for use in electromagnetic environments outside of the home. A 급 기기 (업무용 방송통신기자재) 이 기기는 업무용 (A 급) 전자파적합기로서 판매자 또는 사용자는 이 점을 주의하시기 바라며, 가정외의 지역에서 사용하는 것을 목적으로 합니다.



This product complies with the WEEE Directive marketing requirement. The affixed product label (above) indicates that you must not discard this electrical/electronic

product in domestic household waste.

Product Category: With reference to the equipment types in the WEEE directive Annex 1, this product is classified as “Monitoring and Control instrumentation” product. Do not dispose in domestic household waste. To return unwanted products, contact your local Keysight office, or for more information see <http://about.keysight.com/en/companyinfo/environment/takeback.shtml>.



This symbol indicates the instrument is sensitive to electrostatic discharge (ESD). ESD can damage the highly sensitive components in your instrument. ESD damage is most likely to occur as the module is being installed or when cables are connected or disconnected. Protect the circuits from ESD damage by wearing a grounding strap that provides a high resistance path to ground. Alternatively, ground yourself to discharge any built-up static charge by touching the outer shell of any grounded instrument chassis before touching the port connectors.



This symbol on an instrument means caution, risk of danger. You should refer to the operating instructions located in the user documentation in all cases where the symbol is marked on the instrument.



This symbol indicates the time period during which no hazardous or toxic substance elements are expected to leak or deteriorate during normal use. Forty years is the expected useful life of the product.

Contents

KS2201A - PathWave Test Sync Executive User Manual	10
Chapter 1: Introduction	11
Chapter 2: Install PathWave Test Sync Executive	14
System Requirements	15
Install Main Components	17
Install Additional Components	28
Chapter 3: Installing Licenses	31
PathWave Test Sync Executive License Requirements	32
Supported Licensing Modes	35
The Licensing Process	36
Installing Licenses with PathWave License Manager	37
Chapter 4: HVI Technology Elements	40
About Instruments	41
About PathWave Test Sync Executive	42
Language Support	43
HVI Use Model	44
HVI Engines	46
Resources	48
HVI Sequences	50
HVI Statements	52
HVI Diagrams	59
Timing	63
Chapter 5: HVI integration with PathWave FPGA	70
PathWave FPGA and HVI Overview	71
Using FPGA-Sandbox Resources with HVI	76
HVI FPGA-Memory Maps and HVI FPGA-Register Banks in FPGA-Sandboxes	79
Actions, Events and Triggers in an FPGA-Sandbox	83
FPGA Fast Data Sharing	85
FPGA-Instruction	87
HVI Statements for using FPGAs	90
Chapter 6: Multi-Chassis Systems and System Synchronization Modules	92
System Synchronization Modules	93
Configuring a System with SSMs and System Sync Connectivity	99

Clocking	103
Configuring the Reference Clock	109
Chapter 7: The TSE API	123
TSE API Use Model	124
TSE API Common Functionalities	127
System Initialization	143
The SystemDefinition Object	155
HVI Engines and their Resources	158
Chassis and Interconnects (SyncModule)	166
Synchronization Resources	173
Synchronization Signals and Sync Modes	178
Non-HVI Clocks	179
System Initialization in the SystemDefinition Object	180
User-Defined Trigger Routing	182
System Clocking Configuration	188
The Sequencer Object	193
SyncSequence and Sequence objects	196
Sequence Statements	198
Sync Statements	200
Local Statements	210
EngineView Components	227
InstructionSet Object	232
HVI Registers and Scopes	237
HVI Compilation	242
System Initialization in the Sequencer Object	244
Sequence Representation	245
The Hvi Object	254
Engine Runtime Components	256
Load to Hardware and Run	260
System Initialization during Load To Hardware	261
Real-time Hardware Execution Error Handling	262
Chapter 8: Building an Application with the TSE API	264
Planning an HVI with the HVI Use Model	265
1 Create the SystemDefinition	269

2. Program HVI Sequences	280
3. Compile Your Sequences	290
4. Load To Hardware	291
5. Modify Initial Register Values (Optional)	292
6. Execute Sequences	293
7. Release All Resources	295
HVI Logging and Troubleshooting	296
Chapter 9: TSE Service and Multi-Host support	300
About TSE Service and Multi-Host Systems	301
TSE Service Operation Modes	304
TSE Service Configuration	311
Accessing Remote Resources	320
Using TSE Service in an Application	327
Chapter 10: HVI Time Management and Latency	330
Timing Concepts	331
Synchronization Clocks, Signals, and Modes	332
General Timing Concepts	339
Sync and Local Flow-Control Statement Timing Concepts	342
HVI Instruction Timing Concepts	348
Sync Statement Timing	349
Local Flow-Control Statement Timing	383
HVI Instruction Timing	397
Minimum Start Delay Calculation for Local Flow-Control and Sync Statements	416
Errors in Start Delay or Duration specification	425
Appendix A: Supported Instruments	428
Appendix B: Additional Documentation and Examples	430
Appendix C: Timing Tables	433
Sync Statement Timing Tables	434
Local Flow-Control Statement Timing Tables	441
HVI Instruction Timing Tables	448

KS2201A - PathWave Test Sync Executive User Manual

This User Manual describes the PathWave Test Sync Executive programming environment, which is based on Keysight's Hard Virtual Instrument (HVI) technology. HVI enables you to develop and execute synchronous, real-time operations across multiple instruments. The real-time sequencing and synchronization capabilities of PathWave Test Sync Executive make it a powerful tool for Multi-Input Multi-Output (MIMO) applications that require tight synchronization and real-time control and feedback.

NOTE

PathWave Test Sync Executive (KS2201A) is **not compatible** with the older M3601A. You cannot use them together and they cannot run the same HVI Sequences.

Chapter 1: Introduction

This chapter introduces Keysight KS2201A, PathWave *Test Sync Executive* (TSE) and HVI technology.

Keysight PathWave Test Sync Executive Overview

PathWave *Test Sync Executive* (TSE) is a programming environment based on Keysight's *Hard Virtual Instrument* (HVI) technology, that enables you to develop and execute synchronous real-time operations across multiple instruments.

The real-time sequencing and synchronization capabilities of PathWave Test Sync Executive make it a powerful tool for Multi-Input Multi-Output (MIMO) applications that require tight synchronization and real-time control and feedback. For example:

- Radar.
- Bit error testing.
- Communication systems.
- Massive-scale quantum physics experiments.

PathWave Test Sync Executive supports:

- Multi-chassis configuration.
- HVI Sequence design using an Application Programming Interface (API) for Python.
- Programming of multiple instruments.
- Execution of time-deterministic sequences of operations.
- Precision synchronization and execution.

About HVI Technology

HVI technology enables you to program one or more instruments to execute time-deterministic sequences of operations with precise synchronization. It achieves this by deploying a code executable onto the hardware of each instrument. This executes on an HVI Engine, which is an IP block that is integrated into the instrument. The code executes on these Engines in parallel, across multiple instruments.

The user-defined hardware operation sequence for a group of instruments is called a Hard Virtual Instrument or just HVI. The sequences of operations or instructions executed by the HVI Engines are called HVI Sequences. The operations and instructions that make up HVI Sequences are known as HVI Statements.

When creating an HVI, you can include any instrument that integrates the HVI technology. For example, the Keysight M3xxxA, M53xxA are families of PXI instruments with HVI support. This User Manual includes code examples of the HVI Instrument-specific API that complement the code examples that explain the functionality of the *Test Sync Executive* (TSE) API.

TSE Application Programming Interface

The TSE API is the set of programming classes, properties and methods that enable you to create and program a TSE instance. The TSE API supports the Python and C# languages. Unless otherwise noted, this document refers to the Python API in explanations.

Python Help

A complete description of the HVI Python API is provided in the help file provided with the PathWave Test Sync Executive installer. It is found inside the installation directory for PathWave Test Sync Executive inside the *api\python\Help* subdirectory, by default this is:

C:\Program Files\Keysight\PathWave Test Sync Executive 2023\api\python\Help

Alternatively, you can enter *Python API Help* into the Windows Search.

C# Help

The TSE API documentation for C# is located at:

C:\Program Files\Keysight\PathWave Test Sync Executive 2023\api\dotNet\Help

API Use Model: The HVI-native API and the HVI Instrument Specific API

Each instrument extends the TSE API functionality with an instrument specific API. The TSE API is common to all products and only the instrument specific TSE API is different, depending on the instrument. It is important to differentiate between the HVI-native API features and the instrument-specific

extensions. The extensions enable a heterogeneous array of instruments and resources to coexist in a common framework.

The HVI-native API exposes all HVI functions and is a common API for all products. It defines the base interfaces, classes and properties that are used to create an HVI, control the hardware execution flow, and operate with data, HVI Triggers, HVI Events and HVI Actions, but it alone does not include the ability to control instrument-specific operations. The TSE API defines the hard virtual instrumentation framework, and it is the job of the instrument-specific TSE API extensions to enable instrument functions in an HVI. These functions are exposed by the instrument-specific add-on definitions. This is done by an HVI instrument add-on API provided by each instrument that describes the instrument-specific resources and operations that can be executed or used within HVI Sequences.

HVI instrument-specific definitions are listed in your Instrument documentation. For a list of supported instruments see [Appendix A: Supported Instruments](#).

Chapter 2: Install PathWave Test Sync Executive

This chapter explains how to install PathWave Test Sync Executive and related required components.

It contains the following sections:

- [System Requirements](#)
- [Install Main Components](#)
- [Install Additional Components](#)

System Requirements

This section describes the system requirements for PathWave Test Sync Executive.

PathWave Test Sync Executive Installation Requirements

To install PathWave Test Sync Executive you require the following:

- Python 3.7.x or higher, 64-bit.
- Keysight PathWave Test Sync Executive installer.

To install these, see [Install Main Components](#).

Additional Components Required

To run PathWave Test Sync Executive with hardware, you require:

- One or more PXIe chassis.
- One or more PXIe instruments.
- Associated software, libraries, drivers, and firmware.

Chassis

PathWave Test Sync Executive is compatible with any PXIe chassis, however Keysight recommends the following Keysight chassis so you can make use of their capabilities and multi-instrument and multi-chassis scalability:

- M9019A.
- M9018B.
- M9010A.
- M9046A.

These chassis include an enhanced PXI trigger bridge that provides the capabilities required by PathWave Test Sync Executive to provide support for multi-segment/chassis operation. You can use other chassis without limitation for single segment operation, and you can also use other chassis for multi-segment/multi-chassis operations, but these impose limitations on the complexity of the HVI Sequences that you can execute.

For most chassis, the enhanced PXI trigger bridge functionality is delivered by a firmware update, see your chassis user manual for details. The PathWave Test Sync Executive programming examples show how to verify the correct firmware version for specific chassis. The programming examples are described in [Appendix B: Additional Documentation and Examples](#).

NOTE

The Programming Examples are often updated so ensure you check for the latest versions.

Instruments

PathWave Test Sync Executive works with a number of PXIe instruments.

For more information see the *PathWave Test Sync Executive Release Notes* and [Appendix A: Supported Instruments](#).

Older versions of HVI technology

PathWave Test Sync Executive (KS2201A) and the previous version M3601A, are not compatible. You cannot use them together.

If you use M3601A, the additional components required by HVI use different versions, so they must be reinstalled every time you change between running M3601 and KS2201A.

Install Main Components

This section explains how to install the main components of PathWave Test Sync Executive, it contains the following sections:

1. Install Python, the minimum version is 3.7.x, 64-bit.
2. Install PathWave Test Sync Executive.
3. Manual Installation of Python APIs.

NOTE

PathWave License Manager must not be running when you install PathWave Test Sync Executive.

If PathWave License Manager is running, you must close it before installing the main components.

1: Install Python

PathWave Test Sync Executive requires 64-bit Python. Versions 3.7, 3.8, 3.9, and 3.10 are supported along with their sub-versions. Multiple versions can also be supported.

1. Download the Python installer from the Python web site: python.org.
2. Run the installer.
 - a. Add Python 3.x to the PATH system Variable. To do this, ensure the check box **Add python 3.x to PATH** is checked. This is shown in the following screenshot:



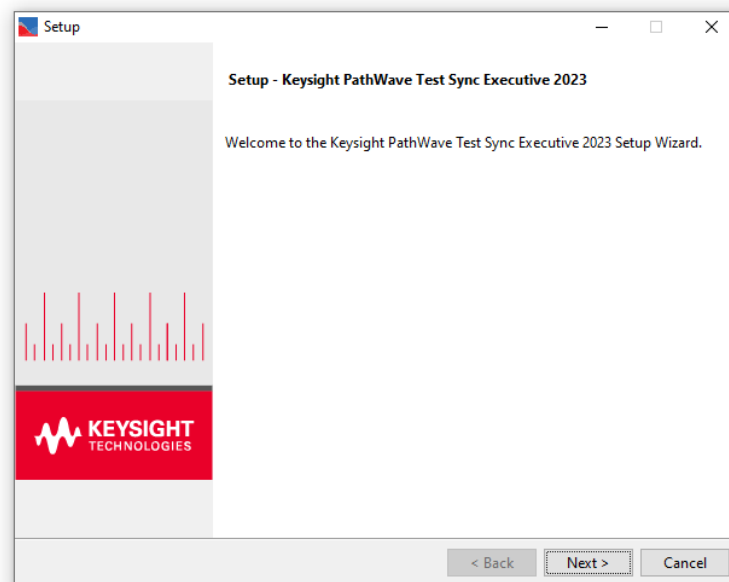
2: Install PathWave Test Sync Executive

Use the following procedure to install PathWave Test Sync Executive:

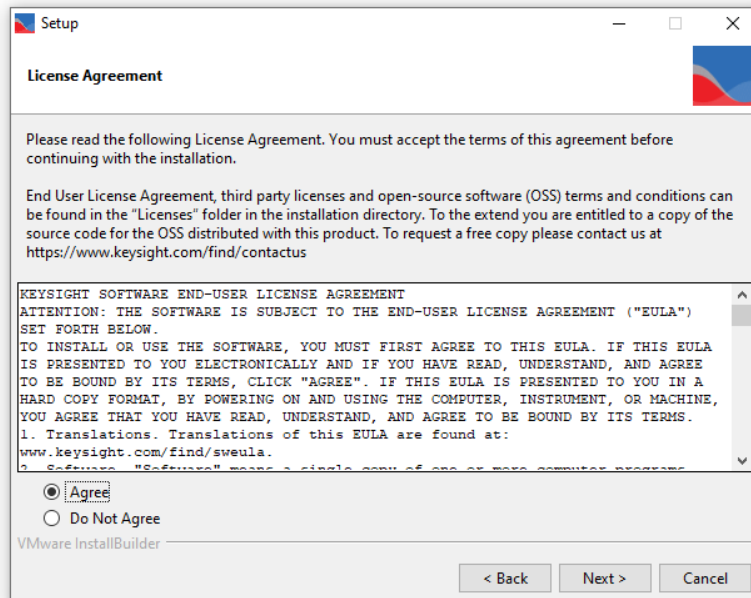
NOTE You must install Python 64-bit before installing PathWave Test Sync Executive. If PathWave License Manager is running, you must close it before installing PathWave Test Sync Executive.

Execute the installer file:

The **Setup** screen is shown:



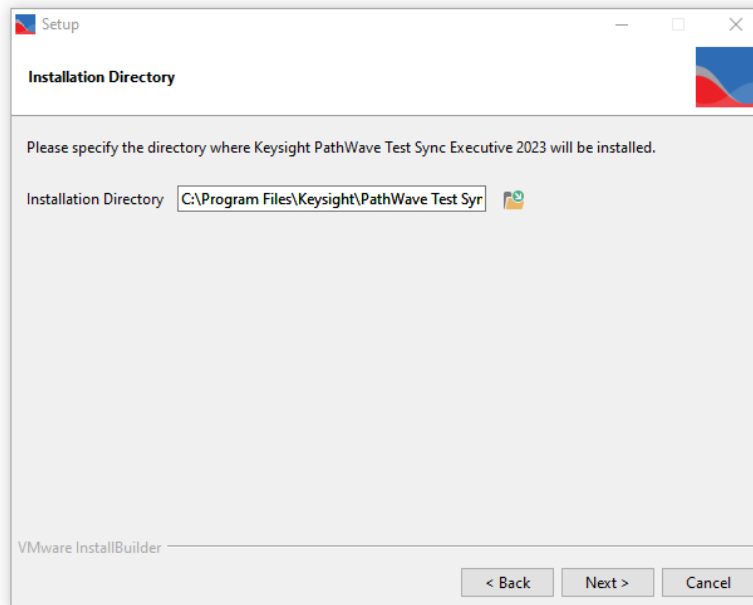
The next screen is the **License Agreement** screen. You must accept the license to continue:



You can change the installation directory on the **Installation Directory** screen.

By default, PathWave Test Sync Executive is installed to:

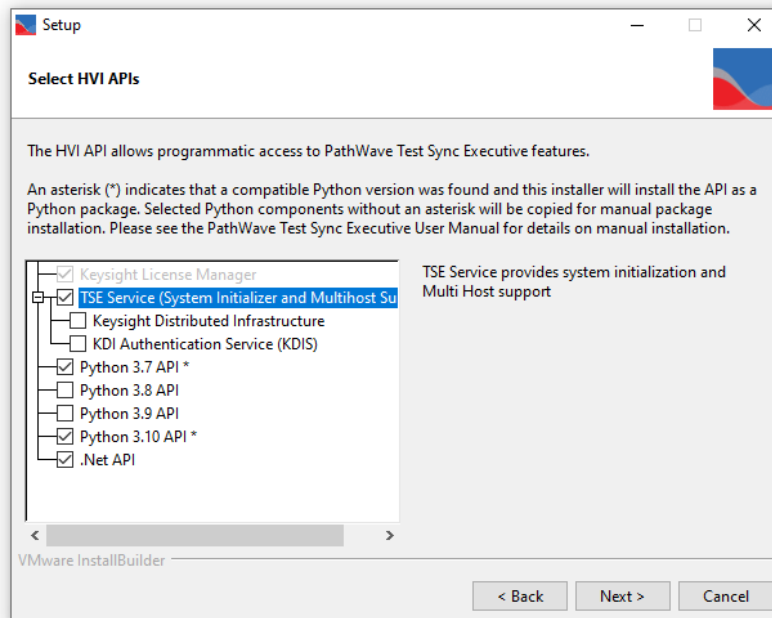
`C:\Program Files\Keysight\PathWave Test Sync Executive 2023`



TSE Service and KDI Options

TSE 2023 introduce extended multi-host and system initialization automation with TSE Service. See [Chapter 9: TSE Service and Multi-Host support](#) for more details on TSE Service capabilities.

To install TSE Service, select the check box labelled TSE Service in the installer. TSE Service is typically installed with KDI, details of the KDI installation options are shown in the section *KDI Installation Overview*.



KDI Installation

KDI can be installed as part of PathWave Test Sync Executive installation when enabling TSE Service.

When Keysight Distributed Infrastructure is selected, TSE Service is configured to be started automatically at windows boot-up. The default configuration after installation will start TSE Service in Free-Running mode configured to autodetect and open all PXI chassis and Instruments supported by TSE Service.

Once TSE Service is selected, to enable KDI installation, you can select:

Keysight Distributed Infrastructure (KDI client).

KDI must be installed on every host,

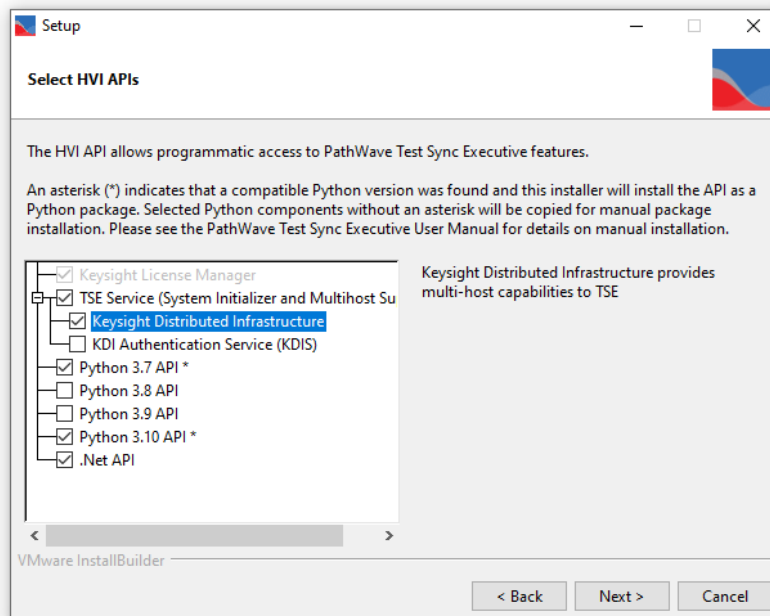
KDI Authentication Service (KDIS).

This is installed on only one host.

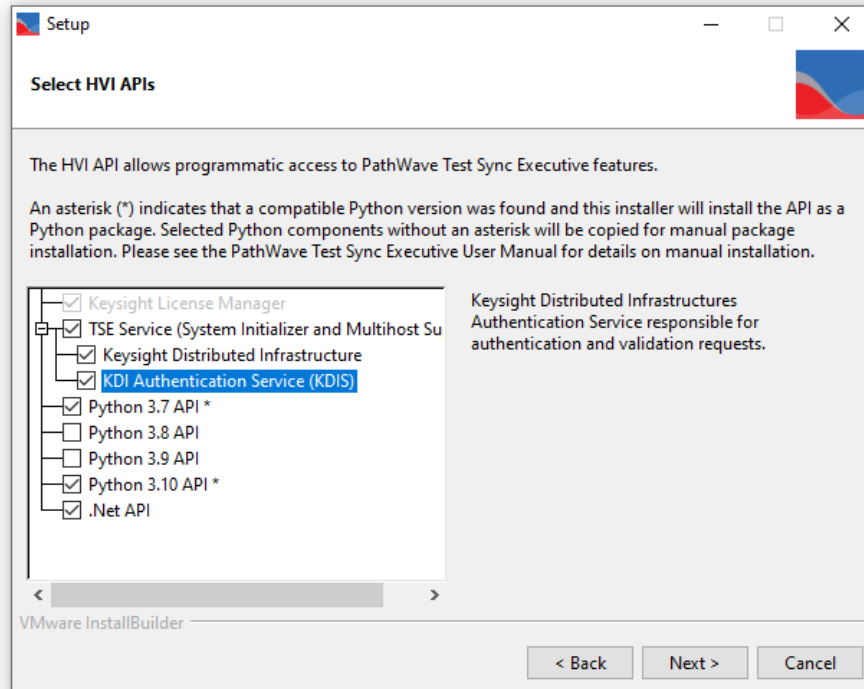
NOTE

Because of a limitation in KDI infrastructure (up to at least release 3.0.187), if you intend to use the TSE Service in Leader-follower mode, then the KDIS must be installed in the Leader TSE Service host.

To install the KDI client, select **Keysight Distributed Infrastructure**. Do this for every host.



To install the KDIS and the KDI client, select **Keysight Distributed Infrastructure** and **KDI Authentication Service (KDIS)**. Do this in one host, in a Leader-Follower system, this must be the leader.

**NOTE**

Keysight recommends that only one KDIS instance is installed in a given network. In this case, any host installing KDI Client will automatically detect the KDIS service for authentication. If for some reason more than one KDIS must be installed, then it is the responsibility of the user to direct each KDI client to the right KDIS host. For more information see the KDI documentation.

KDI and TSE Service Configuration

Once TSE installation is completed in all hosts, you must configure the KDI Authentication Service, KDI clients and TSE Services. For details see section 9. *System initialization with TSE Service and Multi-Host support* in the *PathWave Test Sync Executive Setup Guide*.

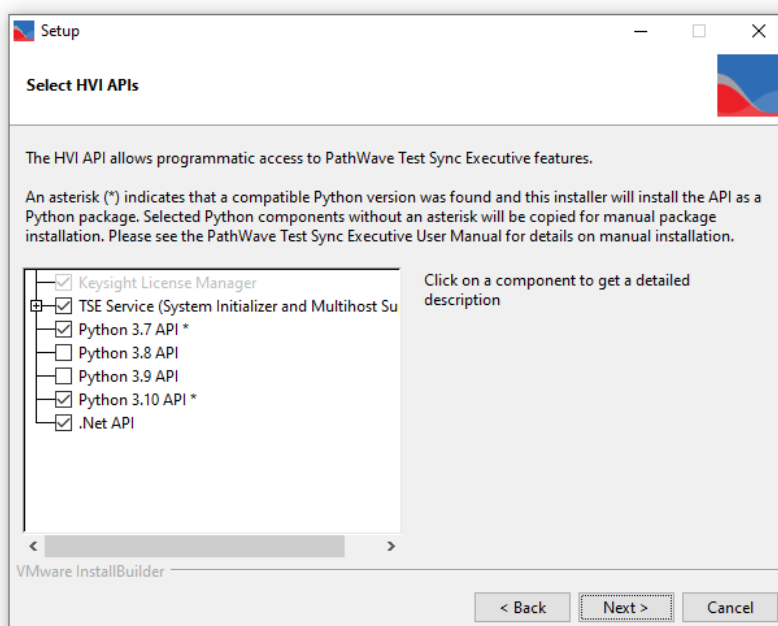
Select Python Versions

You can also select the Python API versions you want to install on the **Select HVI APIs** page.


If a Python version component is marked with an asterisk and selected, the installer will install the Python package.

If the Python version component is *not* marked with an asterisk, but is selected with a check mark, an additional step is required; see Manual Installation of Python APIs below.

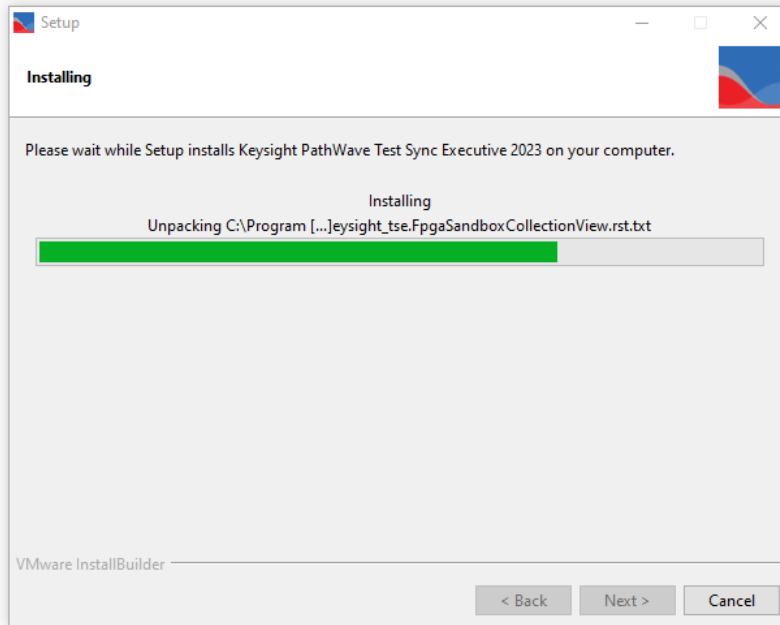
Required components are selected by default and you cannot de-select them.



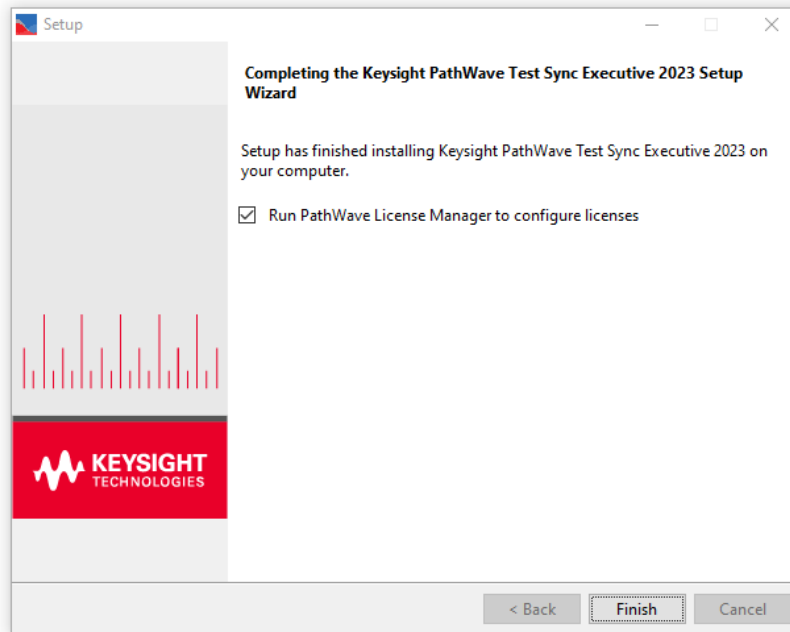
When you have selected the components, the next screen is **Ready to Install**. Select **Next** to install PathWave Test Sync Executive.



The Installer first installs the License manager. It then installs PathWave Test Sync Executive:



The following screen is shown when the installer has completed installing: Select **Finish** to close the installer.



Manual Installation of Python APIs

If you selected Python APIs when installing PathWave Test Sync Executive that were not automatically installed, you can complete the installation process with the `pip` command.

For example, to install Python APIs for Python 3.9, type the following command at a command prompt:

```
py -3.9 -m pip install "C:\Program Files\Keysight\PathWave Test Sync Executive 2023\api\Python\Python39"
```

Install Additional Components

To use PathWave Test Sync Executive, you require both hardware and software.

To work with PathWave Test Sync Executive, instruments and chassis require minimum specific software and firmware versions. These are listed on line at: [Instrument and Chassis Software and Firmware Requirements for KS2201A](#) .

Ensure you have all the following components and they are all up to date:

- Keysight IO Libraries.
- Keysight Instrument Drivers, Libraries, and Software Front Panel.
- Keysight Instrument FPGA Firmware.
- Keysight Chassis Family Driver.
- Keysight Chassis Driver and Firmware.

Install Keysight IO Libraries

Install the IO Libraries. These are available at [Keysight IO Libraries Suite](#).

Install Keysight Instrument Drivers, Libraries, and Software Front Panel

To install the instrument drivers and libraries, install the software for your instruments:

- For the M5302A instrument see: [M5302A Software](#).
- For the M3xxxA instruments see: [Keysight SD1 Software](#).

NOTE

Ensure you check the driver release notes, so that your drivers that are compatible with the version of PathWave Test Sync Executive you have installed.

Update Keysight Instrument FPGA Firmware

You can update the FPGA firmware of your PXI instruments from your Software Front Panel. For information about how to install SW and FPGA firmware for Keysight instruments, see the instrument documentation:

These are available at [Keysight PXI Products](#).

NOTE

Ensure you check the firmware release notes, so that you install firmware that is compatible with the version of PathWave Test Sync Executive you have installed.

Install Keysight Chassis Family Driver

Install the Chassis Family Driver, which is available at [Keysight PXI Chassis](#). When you install the Keysight Chassis Family Driver, PXIe Chassis Software Front Panel software is automatically installed.

Update Keysight Chassis Firmware

In PXIe Chassis Software Front Panel, you can:

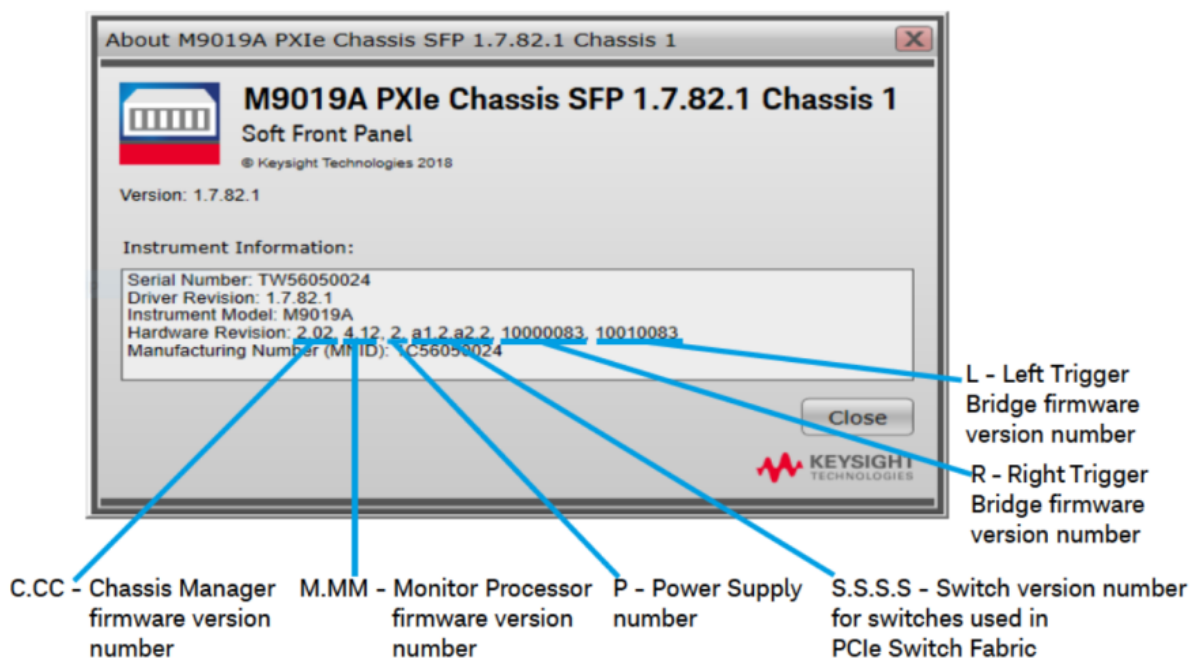
- Check the chassis firmware version in the help window.
- Update the chassis firmware with the Utilities window of PXIe Chassis Software Front Panel (SFP).

You can use the Utilities window of PXIe Chassis SFP to update the chassis firmware. For more information about updating Chassis firmware, see [PXIeChassisFirmwareUpdateGuide.pdf](#) at [Keysight PXI Chassis](#).

NOTE

Ensure you check the firmware release notes, so that you install firmware that is compatible with the version of PathWave Test Sync Executive you have installed.

The following screenshot shows an example of the chassis firmware version shown in the help window of the PXIe Chassis SFP. In this case the chassis is a Keysight Chassis model M9019A.



The following screenshot shows a breakdown of components of different versions of the M9019A chassis firmware:

M9019A Firmware Version Components

Firmware Component	2017	2018	2019StdTrig	2019EnhTrig
Chassis Manager	2.02	2.02	2.02	2.02
Monitor Processor	3.11	3.11	4.12	4.12
Switch version number for switches used in PCIe Switch Fabric	a1.2.a2.2	a1.2.a2.2	a1.2.a2.2	a1.2.a2.2
Right Trigger Bridge	0	10000083	0	10000083
Left Trigger Bridge	0	10010083	0	10010083

Chapter 3: Installing Licenses

This chapter provides a brief introduction to PathWave Test Sync Executive licensing. It contains the following sections:

- [PathWave Test Sync Executive License Requirements](#)
- [Supported Licensing Modes](#)
- [The Licensing Process](#)
- [Installing Licenses with PathWave License Manager](#)

PathWave Test Sync Executive License Requirements

Each instrument used in your HVI implementation must be licensed to be used with PathWave Test Sync Executive.

There are 2 types of licensing for instruments:

1. For instruments with no -HVx option installed, you require 1 license for each instrument (including Sync Modules).
2. For instruments with the -HVx option (-HV1 or -HV2) installed, a single license covers all of the instruments with the -HVx option in the same chassis.

The following table shows an example of the number of licenses required for a single chassis system:

Chassis	Number of Instruments without -HVx option	Number of Instruments with -HVx option	Licenses required
A	1	4	2 <ul style="list-style-type: none"> • 1 license for the instrument without the -HVx option. • 1 license for the 4 instruments with the -HVx option.

The following table shows an example of the number of licenses required for a 3 chassis system:

Chassis	Number of Instruments without -HVx option	Number of Instruments with -HVx option	Licenses required
A	1	4	2 <ul style="list-style-type: none"> • 1 license for the instrument without the -HVx option. • 1 license for the 4 instruments with the -HVx option.
B	4	0	4 <ul style="list-style-type: none"> • 1 license each for the 4 instruments without the -HVx option.
C	2	4	3 <ul style="list-style-type: none"> • 1 license each for the 2 instruments without the -HVx option. • 1 license for all the instruments with the -HVx option.
Total licenses required			9

NOTE

The **-HVx** option was previously required to be purchased for an instrument to be used with PathWave Test Sync Executive.

The **-HVx** option is now deprecated, but existing instruments with the **-HVx** option are still supported.

- Keysight M3xxxA PXI Instruments used the **-HV1** option.
- Keysight M5302A Digital I/O instruments previously used the **-HV2** option.
- Keysight M9415A VXT Vector Transceiver uses the **-HV2** option.

Licenses Requirements per Process

All HVI instances running in the same process share the same licenses, but HVI instances running in different processes require different licenses.

For example, if you have 3 HVI instances running in a single process, the licenses are reused.

The following table shows the number of licenses required for scenarios where these are 1 or 3 processes:

Description	HVI instance 1	HVI instance 2	HVI instance 3	Licenses required
3 HVI instances in the same process	3	6	10	10
3 HVI instances in 3 different processes	3	6	10	19

Licensing Requirements for TSE Service

TSE Service to run in **Leader Mode** has the same licensing requirements as any client application as described above. Following the guidelines above, the number of TSE licenses required depends on the number of instruments, *System Synchronization Modules* (SSMs) and Chassis as defined in the TSE Service configuration files ([tse_config.yml](#) and [system_definition.yml](#)).

Adding licenses to an existing setup

TSE licenses added to an existing set-up must be **co-termed** to match the expiration date of the existing TSE licenses. This is done with the help of your Keysight Field Engineer or Account Manager who will look up your current licenses in order to capture the date-based version. They will then work with internal Keysight systems (KSM) to manually match the end dates of the added licenses to work with your earlier licenses. This process is not customer accessible.

Supported Licensing Modes

The following types of licenses are supported:

Commercial licenses:

- Node-Locked, perpetual and 6, 12, 24, and 36 months, subscription.
- USB Portable, perpetual and 6, 12, 24, and 36 months, subscription.
- Floating/Networked, perpetual and 6, 12, 24, and 36 months, subscription.
- Transportable, perpetual and 6, 12, 24, and 36 months, subscription.

Trial licenses:

- 30 days Node-locked.

NOTE

- To obtain a trial or commercial license, see the product download page.
- As part of the licensing process you will require a Host ID (probably a Mac address) for your workstation. The product license manager might display this, if not, the help or documentation for the license manager shall tell you how to obtain a Host ID.

Transportable Licenses

If you want to reconfigure your systems so a different number of chassis are used, you can use a transportable license. These enable you to move your licenses between systems without any need to contact Keysight, so you don't have to keep buying new licenses.

For example, say you have two systems: one with three chassis and a second system with two chassis. If you want to move the third chassis from the first system to the second, the second system will require a third license. The first system has three licenses, but it shall no longer require all three. A transportable license enables you to move the third license from the first system to the second system. You can then use the new configuration without having to buy a new license.

The Licensing Process

The Keysight licensing process uses the following steps:

1. Purchase and fulfillment

For most Keysight licensed product options, your entitlement certificate is sent to you as a PDF attachment via email immediately after your purchase. In some cases, you receive a paper copy of your certificate with your purchased product. The licensed product options may be software products or upgraded features of an instrument.

2. Getting a license

Using the entitlement certificate you received when you ordered, you can request your licenses on the [Keysight Software Manager](#) web site. To do this, you'll need to choose a host instrument or PC, and provide its identifying information (the Host ID) when you request your licenses. Once you begin the process, Keysight Software Manager will guide you step by step through requesting your licenses and you will receive the license files via email.

You might need to create a *myKeysight* login when you first go to the Keysight Software Manager site, and you will need to log in anytime you go to the site.

3. Installing your license

To enable the licensed software, after you receive a license file from Keysight Software Manager, you must install it on your instrument or computer or on a central licensing server accessible from your instrument or computer. If you are installing node-locked or transportable licenses on the same local PC where you execute KS2201A, ensure you place your license files in a public folder, for example, `C:\Users\public\folder_name`.

To install the license:

1. Install PathWave Test Sync Executive.
2. Use PathWave License Manager to install your license. The installation process is described in the email that comes with your license.

Installing Licenses with PathWave License Manager

You can install licenses from the PathWave License Manager. This is installed when you install Keysight PathWave Test Sync Executive. You can use a local license on your computer or a floating license from a license server.

Full details describing how to install licenses are provided by email when you purchase a license.

If you are upgrading without purchasing a new license, have a more complex setup, or did not get a licensing email, see the [Licensing Quick Start Guide](#), this provides comprehensive information about the licensing process and how to solve problems.

NOTE

If you are upgrading from a previous version of PathWave Test Sync Executive that used a different license manager, Keysight recommends that you keep the old license manager installed.

Potential Conflicts Between License Managers of Different HVI Software

Some previous versions of KS2201A software used a different license manager. Specifically:

- **KS2201A Pathwave Test Sync Executive 2021 Release** and later use **PathWave License Manager (PLM)**.
- **KS2201A Pathwave Test Sync Executive 2020 Update 1 Release** uses **PathWave License Manager (PLM)**.
- **KS2201A PathWave Test Sync Executive 2020 Release** uses **Keysight License Manager 6**.

The license managers described above are compatible with each other and they can detect and show the licenses installed using the other license managers. For node-locked or transportable licenses, conflicts can arise if any licenses were not installed in a public folder, for example, **C:\Users\public\folder_name** . In this case, the license must be reinstalled from scratch using the license manager of the product the license belongs to.

If you are moving from one HVI software to another version that uses a different license manager, to update the floating license installation on your license server, see the instructions provided.

NOTE

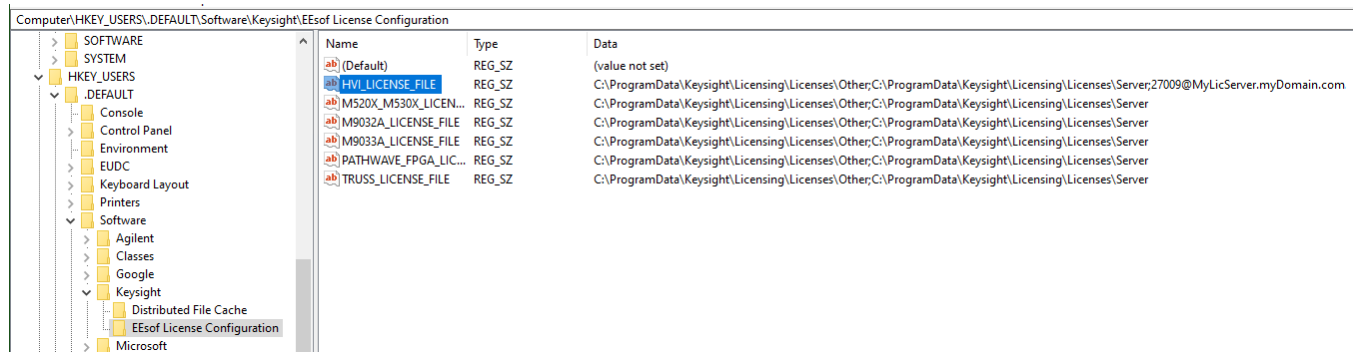
- If you need to uninstall any PathWave Test Sync Executive software, always use the provided software uninstaller. Manually uninstalling a license manager can cause corruption to other license managers.
- If you have licenses located in user-specific locations (such as **C:\Users\Fred\Desktop**), these licenses may not be accessible to the license service created by PathWave License Manager. Using the license manager provided with the appropriate product, remove and reinstall such licenses in a generally accessible location, such as **C:\Users\public**

Licensing Configuration for TSE Service

For TSE Service to run in **Leader-Follower Mode**, it has the same licensing requirements as any client application. The number of TSE licenses required depends on the number of instruments, Sync Modules and Chassis as defined in the TSE Service configuration files (**tse_config.yml** and **system_definition.yml**).

TSE Licenses must be properly configured using PathWave License Manager (PLM). Special care must be taken when using KDI to automate TSE Service launch at Windows boot-up, it is required that PLM is properly configured for the Windows' **SYSTEM** user, since KDI runs at Windows boot-up using this user. This is particularly relevant when using floating licenses configured in a different host. You must make sure the "**HKEY_USERS\DEFAULT\Software\Keysight\EEsof License**

`Configuration\HVI_LICENSE_FILE` key list all paths and servers to search for the TSE Licenses, the screenshot below is an example:



Troubleshooting the License Installation

If you have difficulties with installing or using your licenses see [Licensing Quick Start Guide](#). If the problem persists, please contact Keysight Technical Support and share the log files.

Log files are saved by PathWave License Manager in:

C:\ProgramData\Keysight\Licensing\Log

Chapter 4: HVI Technology Elements

This chapter describes the elements that make up an HVI.

It contains the following sections:

- [About Instruments](#)
- [About PathWave Test Sync Executive](#)
- [Language Support](#)
- [HVI Use Model](#)
- [HVI Engines](#)
- [Resources](#)
- [HVI Sequences](#)
- [HVI Statements](#)
- [HVI Diagrams](#)
- [Timing](#)

About Instruments

Instruments are modules or cards that can capture or generate various kinds of electronic signals. Many kinds of instruments are available with different kinds of functions.

Different kinds of instruments can perform various functions with electronic signals:

- Measure signals.
- Record signals.
- Perform signal analysis.
- Perform signal conditioning.

Some types of instruments can generate different kinds of outputs:

- Signals.
- Voltages.
- Pulses.
- Arbitrary waveforms.
- Digital outputs.

Instruments can be supplied as modules or cards that fit into a chassis. The chassis enables you to fit multiple modules together. The instruments in a chassis are synchronized to a common digital clock reference that is shared by all the instruments. The chassis also offers shared triggering and communication resources.

For this User Manual, the specific instruments referred to are PXI modular instruments that are inserted into a PXI chassis.

For a full list of Keysight instruments, see [Keysight.com](https://www.keysight.com).

About PathWave Test Sync Executive

PathWave Test Sync Executive enables you to program multiple instruments together. They operate together, tightly orchestrated with other instruments, so they behave like a single instrument.

PathWave Test Sync Executive enhances individual instruments by enabling them to:

- Execute real-time sequences of operations with full time determinism.
- Precisely synchronize instrument operations.
- Fast, real-time hardware exchange of information and decisions between instruments.

You define a new virtual instrument made up of a combination of instruments. This is known as a *Hard Virtual Instrument* (HVI). Once the HVI resources are defined, you can program multiple instruments to work together as if they were a single instrument.

To program the HVI, you write an application using the TSE API. When you run your application, it generates the HVI instance and the binary code that is executed by the hardware in the instruments.

When creating an HVI, you can include any instrument that supports PathWave Test Sync Executive, such as Keysight's M3xxxA family of PXI instruments.

Each instrument that supports PathWave Test Sync Executive has specific instructions that enable you to use its functionalities within HVI. These instructions are documented in the instrument documentation.

Language Support

The TSE API is the set of programming classes, properties and methods that enable you to create and program an HVI instance. PathWave Test Sync Executive 2021 and above supports the Python and C# languages.

The C# API is similar to the Python API except for the following differences:

- Class and property names are in camel case, that is, the beginning of individual words are capitalized.
- Variable names are also in camel case, except the first letter of the first word is not capitalized.
- There are no spaces, underscores, or dashes between words in class and property names.
- The first letter of methods and functions is capitalized.

The following table shows examples in Python and C#:

Type	Python	C#
Type names	<code>SystemDefinition</code>	<code>SystemDefinition</code>
Variables	<code>multi_seq_block_1</code>	<code>multiSeqBlock1</code>
Methods	<code>add_sync_multi_sequence_block()</code>	<code>AddSyncMultiSequenceBlock()</code>

The following blocks of Python and C# code are equivalent:

Python code:

```
# Add a Sync Multi-Sequence Block:
multi_seq_block_1 = sync_while.sync_sequence.add_sync_multi_sequence_block("multi_seq_block_1",
210)
```

C# code:

```
// Add a Sync Multi-Sequence Block
var multiSeqBlock1 = syncWhile.SyncSequence.AddSyncMultiSequenceBlock("multiSeqBlock1", 210);
```

A complete description of the HVI Python API is provided in the help file installed with the PathWave Test Sync Executive installer.

It is found inside the installation directory for PathWave Test Sync Executive inside the `api\python\Help` subdirectory, by default this is:

C:\Program Files\Keysight\PathWave Test Sync Executive 2023\api\python\Help

Alternatively, you can enter *Python API Help* into the Windows Search.

The TSE API documentation for C# is located at:

C:\Program Files\Keysight\PathWave Test Sync Executive 2023\api\dotNet\Help

HVI Use Model

This section describes the TSE API use model, and the steps it involves.

HVI uses a program-within-a-program model, that is, HVI can be seen as a real-time hardware program that runs within a software program.

HVI Use Model Steps

To use the TSE API, your application must follow a series of steps to define and run an HVI instance. These steps are broadly defined by three different classes within the TSE API:

1. SystemDefinition.
2. Sequencer.
3. Hvi.

1 SystemDefinition

You instantiate a SystemDefinition object from this class. You use this properties of this object to define all the instrument and platform resources that are required to set up the HVI:

- Chassis.
- Interconnects.
- Clocks.
- Synchronous signals.
- Trigger routing.

You also use this object to define the resources in the instruments that are available to your HVI:

- HVI Engines - IP blocks in the FPGA or instrument hardware that executes HVI Sequences.
- HVI Actions - these initiate instrument-specific operations.
- HVI Events - these indicate instrument-specific operations have occurred.
- HVI Triggers - signals used to communicate between instruments.

When you have defined these resources, you must add them to the relevant collections. Collections are special objects that associate resources with individual HVI Engines, so that you can use the resources on those Engines.

2 Sequencer

You instantiate a Sequencer object and use it to program and compile your Sequences:

- You add instructions and operations known as Statements to Sequences. These can be synchronized across instruments or local to a specific instrument.
- You also add and use HVI Registers within this object. Registers are small, fast memories on the HVI Engines that you can use as program Variables.
- Once you have defined all the Sequences that define your HVI, you must compile it. The compilation process returns an Hvi object.

3 Hvi

Hvi is the runtime or executable object. With this object, you load the HVI Sequences into the relevant Engines and execute them.

This object also enables you to interact with the hardware resources assigned to the HVI and initialize all resources before the actual execution happens.

Execution Flow of the HVI

When you run your application, the HVI instance is generated, compiled, and downloaded into the instruments and infrastructure. It is executed across all the instruments and the infrastructure resources, and then the HVI instance takes control of the individual instruments and platform components. The HVI configures the required resources and downloads the hardware programs that, when executed, run on the instruments and platform hardware synchronously.

An application can create multiple HVI instances, but if the resources are shared, only one can be downloaded and executed in hardware at a time. If the HVI instances do not share any resources, they can be executed in parallel.

HVI Engines

For HVI to control an instrument, the instrument requires one or more HVI Engines. An HVI Engine is an Intellectual Property (IP) block that controls the functions of the instrument and the timing of operations. The HVI Engine is included directly in the instrument hardware or it can be programmed into the *Field Programmable Gate Array* (FPGA) in the instrument.

HVI works by deploying a binary executable to each hardware instrument to be executed by the HVI Engine. Different binaries execute on the different HVI Engines in parallel, across multiple instruments.

When you write an application that includes an HVI, you create HVI Sequences. These are sequences of HVI Statements, these are operations that control the instrument. The HVI Sequences are compiled into the binary executables that the HVI Engine executes.

About Instrument FPGAs

An FPGA is an electronic component on the Instrument. The FPGA in an instrument might include pre-programmed IP for the instrument's functionality and this can include HVI IP components and regions you can configure.

In addition to any existing IP and HVI Engines, instrument FPGAs include an FPGA-Sandbox, this is a user-configurable region in the instrument FPGA. You can configure the FPGA-Sandbox to implement your own specific functionality. This can include custom logic and memory. To take advantage of this feature, you must use *PathWave-FPGA* to create your design in the FPGA-Sandbox. For more information see [Chapter 5: HVI integration with PathWave FPGA](#).

For HVI to control an instrument, the instrument requires one or more HVI Engines. An HVI Engine is an Intellectual Property (IP) block that controls the functions of the instrument and the timing of operations. The HVI Engine is included directly in the instrument hardware or it can be programmed into the FPGA in the instrument.

HVI works by deploying a binary executable to each hardware instrument to be executed by the HVI Engine. Different binaries execute on the different HVI Engines in parallel, across multiple instruments.

When you write an application that includes an HVI, you create HVI Sequences. These are sequences of HVI Statements, these are operations that control the instrument. The HVI Sequences are compiled into the binary executables that the HVI Engine executes.

About Instrument FPGAs

An FPGA is an electronic component on the Instrument. The FPGA in an instrument might include pre-programmed IP for the instrument's functionality and this can include HVI IP components and regions you can configure.

In addition to any existing IP and HVI Engines, instrument FPGAs include an FPGA-Sandbox, this is a user-configurable region in the instrument FPGA. You can configure the FPGA-Sandbox to implement your own specific functionality. This can include custom logic and memory. To take advantage of this feature, you must use *PathWave-FPGA* to create your design in the FPGA-Sandbox. For more information see [Chapter 5: HVI integration with PathWave FPGA](#).

Resources

The HVI Engine executes Sequences that are made up of Statements. These statements or instructions can operate on different resources in real-time. HVI can operate on the following resources:

- HVI Actions.
- HVI Events.
- HVI Triggers.
- Clock signals.
- HVI Registers.
- HVI FPGA-Registers and HVI FPGA-Memory Maps.

Actions, Events and Triggers are concepts within HVI. They are used to initiate operations, wait for operations, send signals, and receive signals.

Actions

HVI Actions are digital electronic pulsed or level signals that are sent from the HVI Engine to control instrument operations outside of the HVI Engine.

You use Actions in HVI Sequences to initiate operations. Typically, Actions initiate instrument-specific operations. For example, in a digitizer instrument, a **StartAcquisition** Action sends a digital pulse to start an acquisition operation.

Events

HVI Events are digital electronic pulsed or level signals that are sent to the HVI Engine and used as notifications when instrument operations have occurred outside of the HVI Engine.

You use HVI Events in HVI Sequences as notification events that the execution has to wait for. Typically, events indicate instrument-specific operations have occurred. For example, in an AWG, the AWG will send a digital pulse through the **waveformDone** event when a waveform execution has been completed.

Triggers

HVI Triggers are electronic signals that the HVI Engines can send or receive.

HVI Triggers are used to send signals and share data between instruments. You can use these to initiate operations, communicate states, or share information. There are multiple types of triggers depending on how they are connected, for example:

- Front panel triggers (usually a SMA connector on the instrument front panel).
- PXIe triggers (connected to the PXIe backplane of the chassis).
- General purpose digital IO (LVDS connector in the instrument front panel).

HVI Registers

HVI Registers are similar to Variables in a programming language. They hold values that can be modified at runtime and can be used as parameters for instructions and Statements. Physically, HVI Registers are small hardware memories located in HVI Engines. Their contents can be shared between HVI Engines by using specific instructions.

HVI FPGA-Registers and HVI FPGA-Memory Maps

Some instrument FPGAs provide a user-configurable region in the instrument FPGA known as an FPGA-Sandbox. This enables you to program the instrument with logic that implements your own custom functionality. FPGA-Registers and FPGA-Memory Maps are components in the FPGA-Sandbox that you can use as resources in your HVI Sequences. For more information see [Chapter 5: HVI integration with PathWave FPGA](#).

For the instruments that support an FPGA-Sandbox, HVI can support the sharing of data between the FPGA-Sandbox and the HVI Engine in an instrument or between the FPGA-Sandboxes of different instruments. This functionality depends of the availability of specific interfaces inside the FPGA-Sandbox. To take advantage of these features, you must use [PathWave FPGA](#) to create your design in the sandbox.

NOTE

The exact resources available and how they are configured is instrument dependent. Each instrument defines the actions and events available, how it uses triggers and the number and type of registers available. For the specific definitions and availability of resources in each instrument, see your instrument documentation.

HVI Sequences

You control instruments with HVI Statements. Statements operate on resources such as Actions, Events, and Triggers. There are different types of Statements that perform different types of operations. HVI Statements are the building blocks of HVI Sequences. These Sequences are compiled in your application and are executed in real-time on the HVI Engines.

An HVI Sequence is an ordered list of HVI Statements with associated timing information. A Sequence is executed in a time-deterministic manner by the hardware HVI Engine located within an instrument. An HVI instance is made up of one or more Sequences that run in parallel and synchronously.

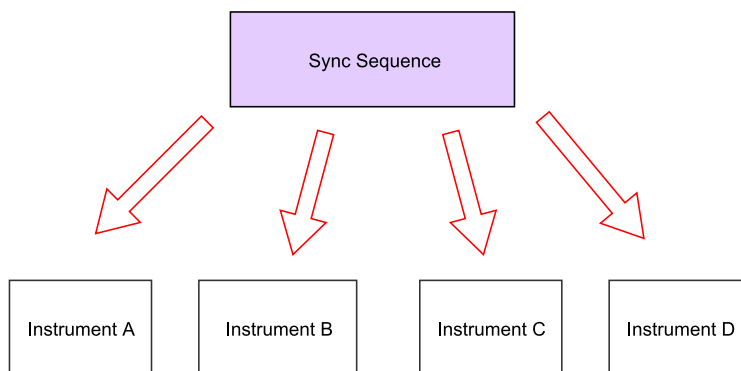
There are two types of sequences:

- **Sync Sequences**
- **Local Sequences**

HVI Sequences are organized in a hierarchy with Sync Sequences at the top.

Sync Sequences

A Synchronized Sequence (called a Sync Sequence) contains commands known as Sync Statements that execute across multiple instruments:

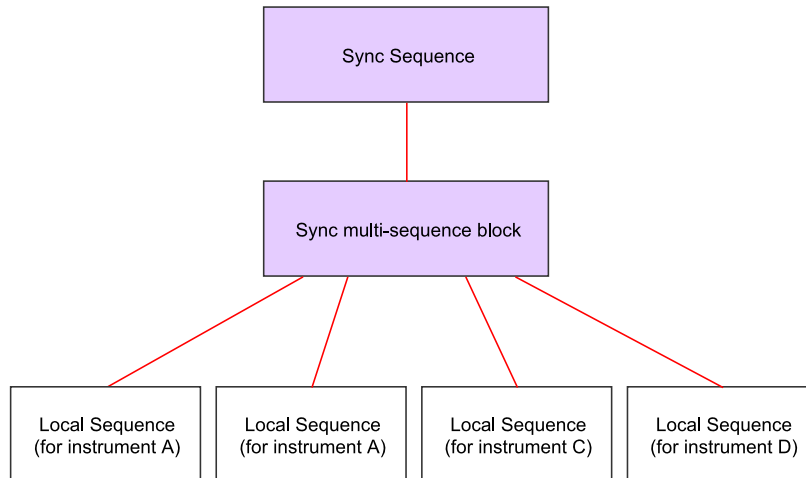


Local Sequences

The Local Sequences are executed by each individual HVI Engine in an instrument.

Local Sequences are contained within *Sync Multi-Sequence Blocks* (SMSB). A Sync Multi-Sequence Block is a type of Sync Statement that is contained in a Sync Sequence.

The following diagram shows the relationship between a Sync Sequence, a Sync Multi-Sequence Block, and Local Sequences:



HVI Statements

This section describes HVI Statements, it contains the following sections:

- **HVI Sync Statements**
 - Sync While
 - Sync Register-Sharing
 - Sync FPGA Data-Sharing
 - Sync Multi-Sequence Block
- **HVI Local Statements**
 - HVI Instructions
 - Local Flow-Control Statements

HVI Statements are the commands or operations that make up an HVI Sequence. HVI Sequences are the ordered lists of HVI Statements that are executed with precise timing. If you think of an HVI Sequence as a poem, the HVI Statements are the possible words you can use to write the poem and the TSE API is the language you use to write it. HVI Statements are FPGA-level operations that are executed by the HVI Engines.

HVI Statements are broadly divided into two groups:


HVI Sync Statements

Synchronized (Sync) Statements are used to execute operations or control the flow of execution across all the hardware HVI Engines. Sync Statements are executed synchronously among all HVI Engines.

HVI Local Statements

These are the commands or operations you put in the Local Sequences to be executed on a specific HVI Engine that is in a specific hardware instrument.

The following diagram shows the different kinds of Statements and how they relate to Sync Sequences and Local Sequences:



HVI Sync Statements

These are used to execute operations or control the flow of execution across all HVI hardware Engines. Sync Statements are executed synchronously among all HVI Engines.

HVI Sync Statements are contained in a Sync Sequence. HVI Sync Statements execute across all instruments.

The Sync Sequence enables multiple Engines to execute statements in lockstep.

The following HVI Sync Statements are available:

- Sync While
- Sync Register-Sharing
- Sync FPGA Data-Sharing
- Sync Multi-Sequence Block

Sync While

Enables a while loop to execute synchronously on all Engines.

The Sync While Statement enables you to execute a Sync Sequence in a loop while a condition is met. The condition is evaluated each time before starting the Sync Sequence execution. When the condition is false and the Sync Sequence reaches the end, the Sync while jumps out of the loop and the Sync Sequence containing the Sync while continues execution with the next Sync Statement.

Sync Register-Sharing

The Sync Register-Sharing statement enables you to share data from a source register to a destination register in any other HVI Engine.

It enables you to share the contents of N adjacent bits from a source register and write it to a destination register in another HVI Engine in your HVI.

Sync FPGA Data-Sharing

Enables you to share data from one FPGA-Sandbox to one or more other FPGA-Sandboxes in different instruments. The data-sharing is orchestrated by the HVI Engines of the different instruments.

Data can be shared between instruments in a single chassis or across instruments in multiple chassis. Sync FPGA Data-Sharing utilizes the *Fast Data Sharing* (FDS) functionality to enable the low-latency transfer of data.

Data is sent 4 bits at a time and can be sent from one to one, or from one to many FPGA-Sandboxes.

Sync Multi-Sequence Block

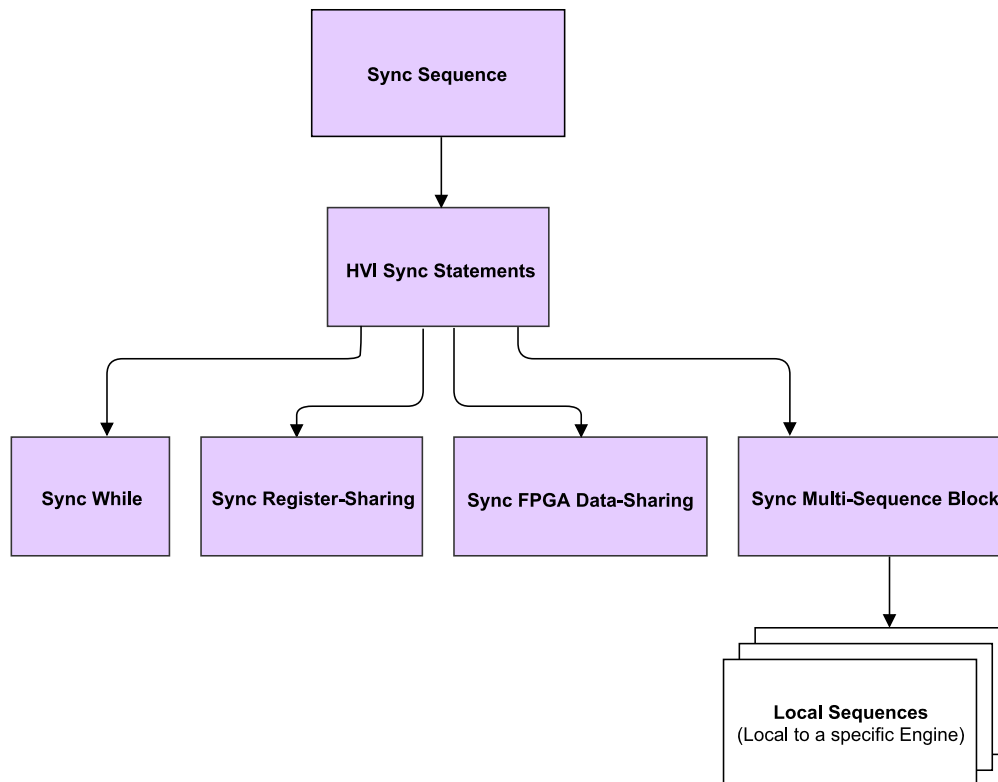
Enables the execution of multiple, simultaneous, Engine-specific Sequences.

Sync Multi-Sequence Blocks are a type of Sync Statement that contain a set of Local Sequences. The Local Sequences execute on individual HVI Engines within the instruments. All Local Sequences contained in a Sync Multi-Sequence Block start and end at the same time.

The Sync Multi-Sequence Block enables you to run different Sequences on each Engine concurrently. It ensures that the execution of all the Local Sequences starts exactly at the same time and that the Sync Sequence remains synchronous afterwards. It serves as a boundary between sections and a container where each Engine operates individually.

All HVI Local Sequences operate within HVI Sync Statements. The HVI Sync Statements determine global or synchronized operations, or synchronization points.

The following diagram shows how the HVI Sync Statements fit in the Sync Sequence:



HVI Local Statements

HVI Local Statements are the commands or operations that make up Local Sequences. These are the commands or operations you put in the Local Sequences to be executed on a specific HVI Engine in a specific hardware instrument. There are two types of Local Statements:

- HVI Instructions.
- HVI Local Flow-Control Statements.

HVI Instructions

These are operations that are executed by the HVI Engine in the instrument hardware and do not impact the execution flow.

There are two types of Local Instructions:

HVI Native Instructions

HVI-native instructions are instrument independent, general-purpose instructions present on all instruments, for example, math operations, writing Triggers and executing Actions. HVI Native Instructions are defined by the TSE API.

Instrument-Specific Instructions

These are instructions that are specific to instruments. You can use these when you program an HVI with those specific instruments.

These instructions can change instrument settings such as amplitude and frequency. They can also trigger instrument functions such as queuing waveforms for playback, outputting a waveform, or triggering a data acquisition.

Instrument-Specific Instructions are defined by the HVI instrument add-on API and are exposed in each instrument driver as instrument-specific HVI definitions.

NOTE

The User Guides for the M320xA PXI AWGs and M310xA PXI Digitizers describe all the HVI Instructions available for each of the M3xxxA PXI instruments.

Local Flow-Control Statements

Local Flow-Control Statements are used to control the execution flow within each Local Sequence. These Statements are depicted with yellow boxes in the HVI diagrams displayed in this User Manual.

These are used to control the execution flow of a specific HVI Engine. They are divided into two types:

- Wait Statements
- Conditional Flow-Control Statements

Wait Statements:

Local Wait-For-Event

Waits for a condition that can be determined by an HVI Event, an HVI Trigger, or any logical combination of any of these types of conditions.

Local Wait-For-Time

Waits for an amount of time specified in a register.

Local Delay

Delays a Sequence for a time you specify.

Conditional Flow-Control Statements:


Local If

This acts as a Local If-Elseif-Else, Local If executes one of a set of possible Local Sequences depending on the value of a defined condition.

Local While

Executes while a condition is true.

The following diagram shows the different types of Local Statements and their relationship to the Local Sequences:



HVI Diagrams


This section shows HVI diagrams. These are used to illustrate HVI Sequences.

In the HVI diagrams, the following colors are used to indicate different kinds of Statements:


Sync Statements	
HVI Instructions	
Local Flow-Control	

Statement diagram color code


The following diagram shows a single Sync Statement with flow and time for the block:



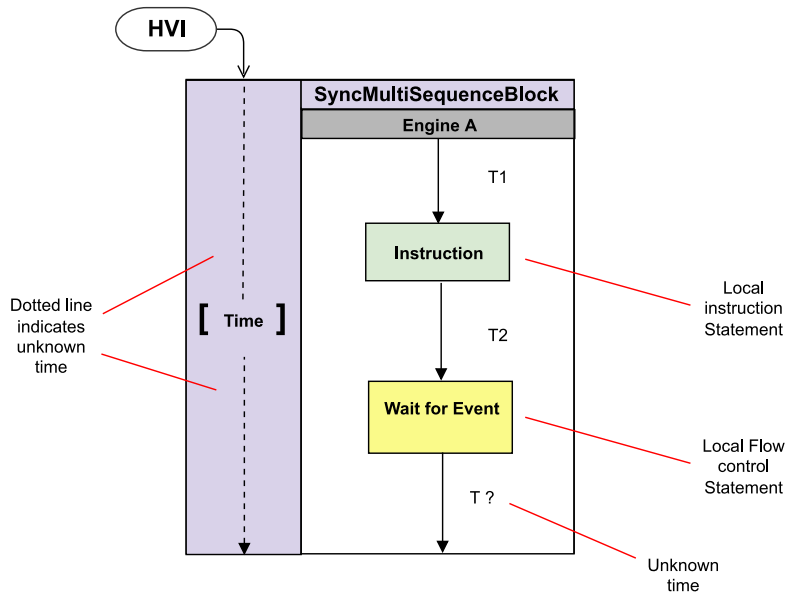
The diagrams can show nesting of Statements within Statements. For example, the following diagram shows a Sync Statement that is within another Sync Statement:



Local Sequences are placed within their HVI Engines in Sync Multi-Sequence Blocks. The following diagram shows a pair of Local Sequences with an HVI Instruction each inside a Sync Multi-Sequence Block:

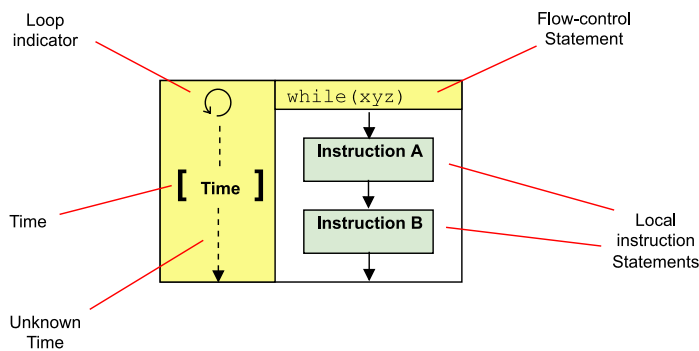


A dotted line indicates that execution time is not known at compile time. This is often the case with flow-control Statements. In this case the Wait-For-Event Statement shall not release until the Event occurs. It is not known at compile time when this is, so the time cannot be calculated at compile time.




The following diagram shows a Local Flow-Control Statement that encloses a pair of HVI Instructions. The color Yellow indicates a Local Flow-Control Statement.

The circular symbol is a loop indicator that shows that the block iterates.



The following diagram shows a more complex example. The Sync Multi-Sequence Block contains two Local Sequences, one per HVI Engine. The Local Sequences execute operations on their associated HVI Engines in parallel.



Timing

This section introduces the basic HVI timing concepts, it contains the following sections:

- [HVI Statement Timing Definitions](#)
- [Timing Descriptions for Different Statement Types](#)
 - [Start Delay operation for different types of Statements](#)
 - [HVI Instruction Timing](#)
 - [Local Flow-Control timing](#)
 - [Sync Statement timing](#)

HVI timing is a complex topic that involves you understanding how to calculate the timing between statements. The calculations required and parameters involved are described in detail in [Chapter 10: HVI Time Management and Latency](#).

HVI Statement Timing Definitions

When you are programming an HVI, you have precise control over the timing of HVI Statement execution. To do this correctly, you must understand the following time definitions:

- Start time.
- End time.
- Execution time.
- Start Delay.

Start time

This is the instant of time when the HVI starts the execution of a Statement.

End time

This is the instant of time when the execution of a Statement is completed and the result is available, or an operation has completed.

Execution time


This is the time interval from the Start time to the End time of the Statement.

Start Delay

The Start Delay defines the period between the execution of consecutive statements. The Start Delay enables you to have full control of the timing of operations and ensures there is enough time for correct execution. Start Delay is a parameter that you set when you add statements to a Sequence.

NOTE If you do not specify a valid Start Delay, the compiler generates an error and indicates the minimum valid minimum value.

The following diagram shows the HVI statement timing definitions:



Timing Descriptions for Different Statement Types

This section describes statement timing and provides a set of examples. It contains the following subsections:

- Start Delay operation for different types of Statements.
- HVI Instruction timing.
- Local Flow-Control Statement timing.
- Sync Statement timing.

Start Delay operation for different types of Statements

Start Delay is always specified between two consecutive Statements, from the previous Statement to the current Statement.

You define a Start Delay in one of 2 different ways:

- From the beginning of the previous Statement.
- From the end of the previous Statement.

The way you define the Start Delay depends on the type of the previous Statement. For example, say you have 2 Statements: A followed by B. The Start Delay for Statement A is already specified and you want to specify the Start Delay for Statement B.

The current Statement is Statement B, so the Start Delay of Statement B depends on the type of the previous Statement A:

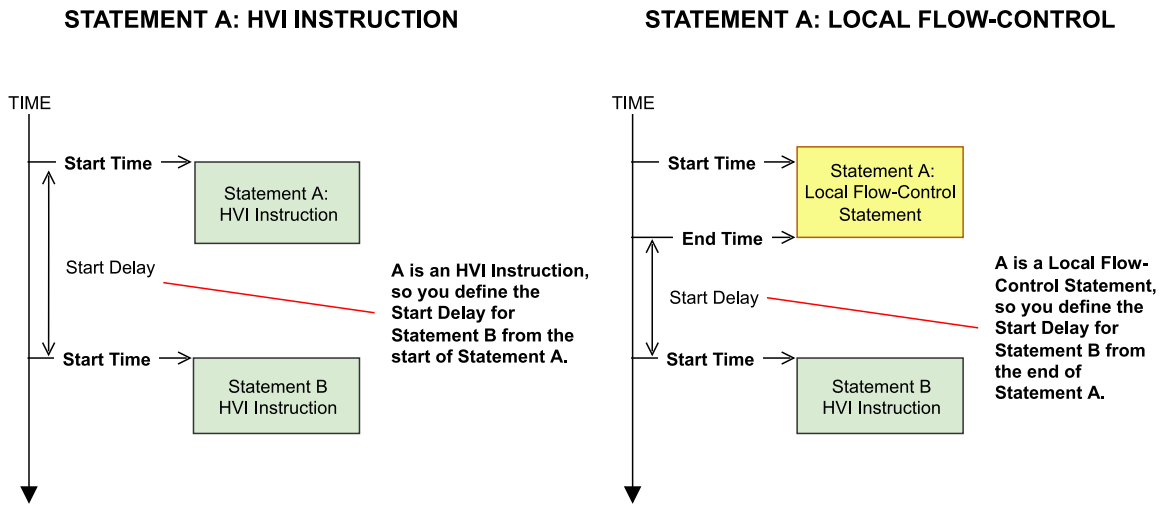
HVI Instructions

If Statement A is an **HVI Instruction**, the Start Delay of Statement B starts at, and is measured from, the **Start time** of the Statement A.

Sync Statements and Local Flow-Control Statements

If Statement A is a **Sync Statement** or a **Local Flow-Control Statement**, the Start Delay of Statement B starts at, and is measured from, the **End time** of Statement A.

The following diagram shows the different Start Delay definitions:

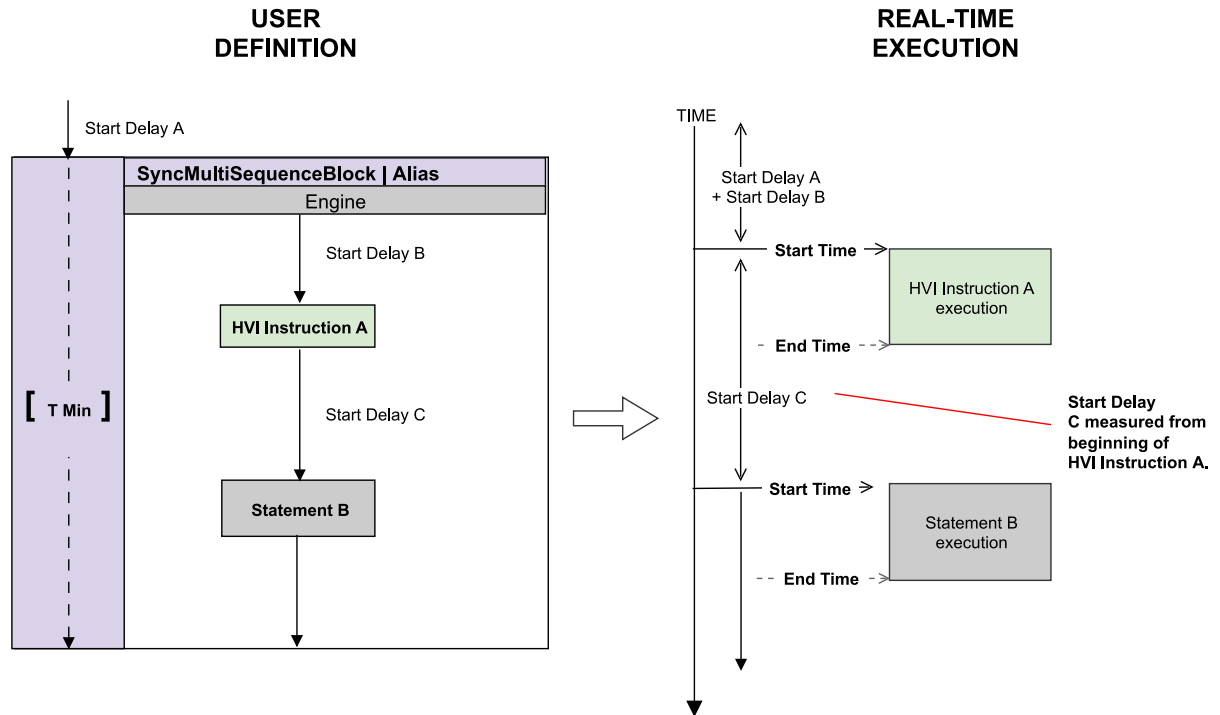


HVI Instruction Timing

The following diagram shows the timing when the **previous** Statement is a HVI Instruction.

For HVI Instructions, the Start Delay of the following instruction *is measured from the start of the previous instruction*.

The following diagram shows two Statements and their timing where the first Statement in a HVI Instruction:



Local Flow-Control timing

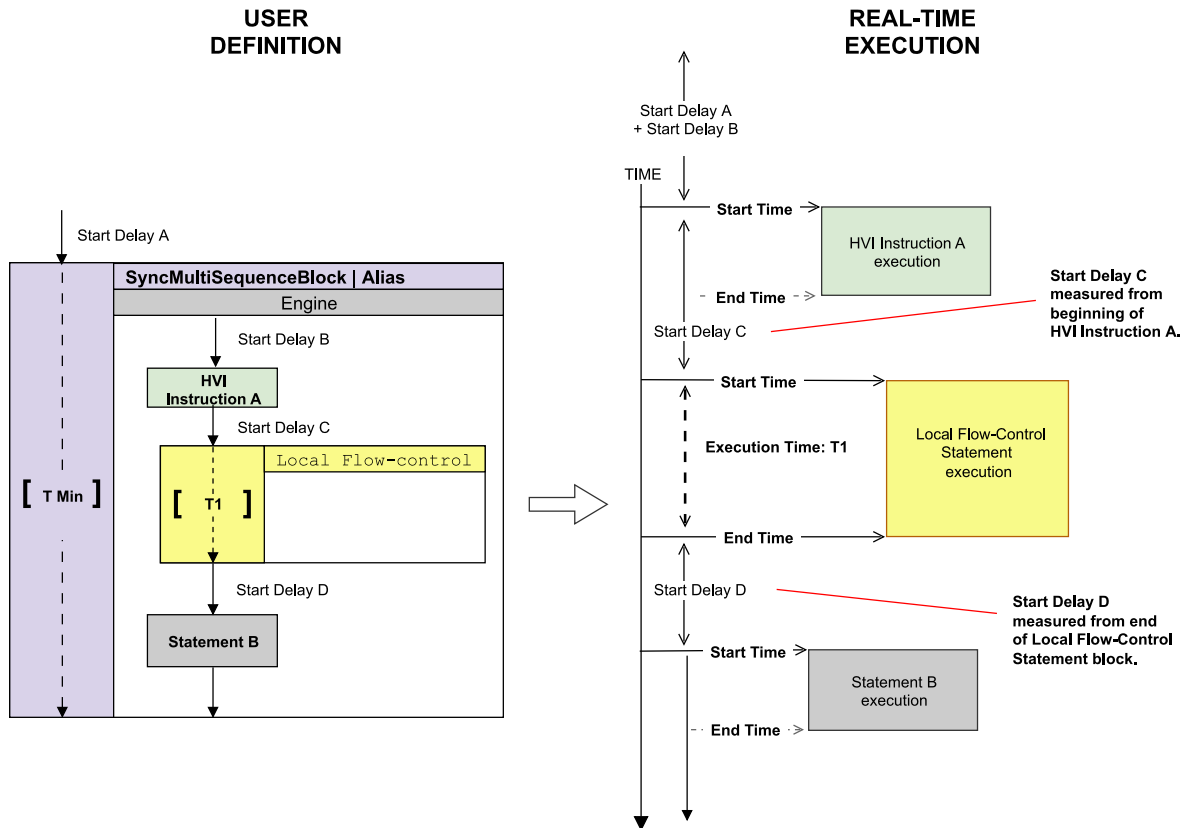
For Local Flow-Control Statements, the Start Delay of the next Statement *is measured from the end of the previous Local Flow-Control Statement*. This is because the HVI Engine is busy during the execution of the flow-control Statement and the execution of a flow-control Statements cannot be overlapped with any following Statements.

For the Local Flow-Control Statement after instruction A, the Start Delay (Start Delay C) is measured from the start of the previous instruction (instruction A).

For instruction B, that follows the Local Flow-Control Statement, the Start Delay (Start Delay D) is measured from the end of the flow-control block.

The execution time of Local Flow-Control Statements can be known at compile time, or might be unknown, the dotted line in the diagram below indicates that the execution time of the Local Flow-Control block T1 is not known at compile time.

The following diagram shows the difference between measuring timing of HVI Instructions and Local Flow-Control Statements.



Sync Statement timing

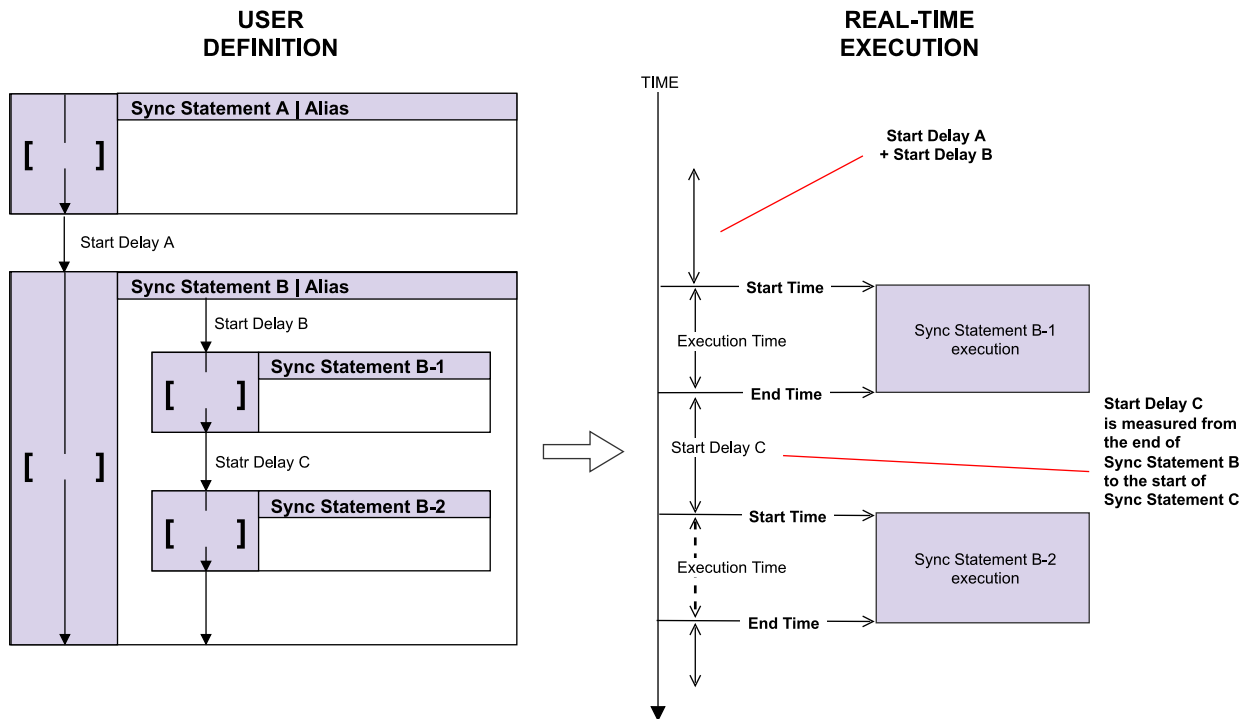
For Sync Statements, the Start Delay is measured from the end of one Sync Statement to the start of the following Sync Statement.

The following diagram shows two Sync Statements, A and B. Sync Statement B is a container for two further Sync Statements, B-1 and B-2. The times indicated are Start Delay A, Start Delay B, Start Delay C, T1, and T2.

The time between the end of Sync Statement A and the start of Sync Statement B-1 is Start Delay A + Start Delay B. The time between the end of Sync Statement B-1 and the start of Sync Statement B-2 is Start Delay C.

The execution time of Sync Statements can be known at compile time, as shown below with a solid line.

The following diagram shows the timing between Sync Statements:



Chapter 5: HVI integration with PathWave FPGA

This chapter describes PathWave Test Sync Executive integration with PathWave FPGA. It contains the following sections:

- [PathWave FPGA and HVI Overview](#)
- [Using FPGA-Sandbox Resources with HVI](#)
- [HVI FPGA-Memory Maps and HVI FPGA-Register Banks in FPGA-Sandboxes](#)
- [Actions, Events and Triggers in an FPGA-Sandbox](#)
- [FPGA Fast Data Sharing](#)
- [FPGA-Instruction](#)
- [HVI Statements for using FPGAs](#)

PathWave FPGA and HVI Overview

What is an FPGA?

A *Field programmable Gate Array* (FPGA) is a digital electronic component on many Keysight instruments, whose behavior can be modified for different use cases.

Keysight instruments use FPGAs to implement complex functionality and data processing. Some instruments also make a region in the FPGA available to enable the addition of custom logic and real-time processing into the instruments. You can customize the FPGA with **PathWave FPGA** software.

PathWave FPGA

PathWave FPGA is a graphical software tool that enables you to rapidly customize logic in the *Sandbox* section of the FPGA in supported Keysight instruments. By doing this you can modify or enhance the default behavior of these instruments.

Instruments that support both PathWave FPGA and PathWave Test Sync Executive enable you to combine your customized logic with the real-time capabilities of PathWave Test Sync Executive. For example, you can have DSP processing in the FPGA-Sandbox, triggered in real time by an HVI Sequence.

FPGA-Sandbox

In addition to any existing *Intellectual Property* (IP) and HVI Engines, a Keysight instrument FPGA can include one or more FPGA-Sandboxes. An FPGA-Sandbox is the region in the FPGA that you can configure using PathWave FPGA.

You can configure the FPGA-Sandbox to implement your own custom IP, signal processing and other functionality. This can include custom logic, HVI FPGA-Registers and memory interfaces. HVI can interact with this custom logic using HVI-specific interfaces.

PathWave FPGA includes an *Intellectual Property* (IP) library that includes Logic/Math, Memory, and DSP blocks that you can place in the FPGA-Sandbox. The *Real-time HVI* design interface catalog enables you to add memories and FPGA-Registers, and also contains specific HVI interfaces for HVI Actions, HVI Events, HVI Triggers, and FPGA-Instruction Statements.

.k7z files

When you have completed your design, PathWave FPGA enables you to easily build a **.k7z file** for your FPGA from your schematic.

The **.k7z** file contains the bitfile that is used to program the FPGA-Sandbox design into the FPGA. The **.k7z** file also contains all the information about the FPGA-Sandbox design, such as names, addresses, ranges of the HVI FPGA-Registers, and memory-mapped locations, etc, including resources that are connected to the HVI Engine.

You add the **.k7z** into your HVI instance in the SystemDefinition. This file is used by the HVI to get all the definitions required so you can utilize your customizations.

HVI Resources in the FPGA-Sandbox

PathWave FPGA enables you to add HVI to your logic design in the FPGA-Sandbox. These resources enable your HVI to interact with the logic in the sandbox.


You can add the following types of components and access them from your HVI Sequences:

- HVI FPGA-Registers.
- HVI FPGA-Memory Maps.
- HviAction interfaces.
- HviEvent interfaces.
- HviTrigger interfaces (HviTriggerIoIn, HviTriggerIoOut, HviTriggerIoT).
- FPGA *Fast Data Sharing* (FDS) ports.
- FPGA-Instruction (HviFPGAInstructions).

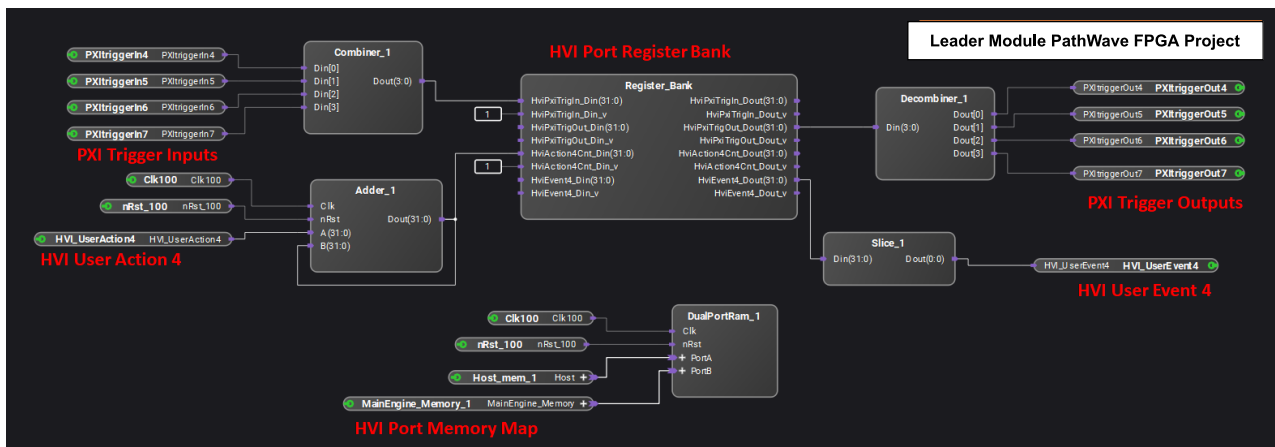
NOTE

The exact resources you can add depends on the capabilities of the instrument you are using. For example, FDS ports are only available on instruments that support them.

The following diagram shows a screenshot of PathWave FPGA with some example resources:



The following image is a screenshot of from PathWave FPGA taken from Programming Example 3. The image shows a set of FPGA blocks, a number of HVI resources, and the connections between them in an FPGA-Sandbox. For more information about Programming Example 3 see: [Appendix B: Additional Documentation and Examples](#). For information about how to use PathWave FPGA see the documentation at [PathWave FPGA](#).




PathWave Test Sync Executive includes a number of HVI instructions that enable your HVI Sequences to interact with the IP blocks in the FPGA-Sandbox. The instructions include access to FPGA-Registers or FPGA-Memory Maps, and send data or commands into the FPGA-Sandbox.

If you want to send an Action into the FPGA-Sandbox, you must add an **HviAction** interface to the FPGA-Sandbox, you just add this to the design and connect them to your customized logic. In a Sequence you are able to interact with the logic using the relevant instructions. Actions in the FPGA-Sandbox are accessed from Sequences in the same way you use any other HVI Actions, no special instructions are required. **HviEvents** and **HviTriggers** work the same way.

HVI Sequences and FPGA-Sandbox resource interaction

When you run your HVI, the HVI Engine reads and executes the individual commands within your HVI Sequences. When the HVI Engine executes an HVI Statement that involves interaction with a resource in the FPGA-Sandbox, the HVI Engine communicates with the FPGA-Sandbox.

The following diagram shows an FPGA-Sandbox that contains custom IP blocks with connections to an HVI Engine and the instrument physical interfaces:



Using PathWave FPGA with PathWave Test Sync Executive

The full flow to customize the logic in an FPGA-Sandbox and then use these customizations in HVI Sequences is:

In PathWave FPGA:

- Open the instrument *Board Support Package* (BSP) using PathWave FPGA.
- Customize the logic by adding logic blocks.
- Add any HVI FPGA-Registers and memories required.
- Add HVI interfaces so HVI can interact with your logic (Actions, Events, Triggers).
- Connect your customized logic to the relevant I/O signals in the HVI interfaces.
- Generate the **.k7z** file.

Once you have configured the FPGA with PathWave FPGA, added the relevant HVI interfaces as required, and generated a **.k7z** file, you must load the definitions into PathWave Test Sync Executive:

- Load the **.k7z** file into your SystemDefinition.
- Write your HVI Sequences and use the FPGA resources.
- Load your HVI instance to hardware (this step loads the **.k7z** as required).

You can use the HVI resources in the FPGA in the same way as you use any other HVI resources.

Using FPGA-Sandbox Resources with HVI

This section describes what to do in the different HVI programming stages, so you can use the FPGA customizations you made in PathWave FPGA in your HVI Sequences.

Load the .k7z file in SystemDefinition

When you have completed and built a design, PathWave FPGA generates a **.k7z** file.

Before you can use any FPGA-Sandbox resources with HVI, you must first load the **.k7z** into your HVI System Definition.

The **.k7z** file contains a *bitfile*, that is used to load the design into the FPGA-Sandbox. The **.k7z** file also contains information about the resources in the design, such as names of ports and interfaces, addresses, ranges of the registers, memory-mapped locations, etc. The **.k7z** file is used to program your customizations into the FPGA and it is also used by HVI to get all the definitions required so you can utilize your customizations.

The following code shows how to load the **.k7z** file:

```
# This must be the name that the instrument has defined for the target FPGA-Sandbox
sandbox_name = "InstrumentSandbox1"
# Get FPGA-Sandbox
sandbox = system_definition.engines[engine_name].fpga_sandboxes[sandbox_name]
# Load the k7z file to HVI
sandbox.load_from_k7z(k7z_file_path)
```

Using FPGA-Sandbox resources in an HVI Sequence

When you load the **.k7z**, file into a specific FPGA-Sandbox, HVI is able to access the resources defined in PathWave FPGA for that specific FPGA-Sandbox, allowing you to use them in your HVI Sequences or at runtime.

The following example shows how to get and write to an HVI FPGA-Memory Map inside a Sequence:

```
# Get FPGA-Memory Map object by name, as this is defined in the PathWave FPGA design
hvi_memory_map = sandbox.hvi_memory_maps["memory_map_name"]
#
# Write FPGA-Memory Map
engine_sequence = multi_sequence_block_statement.sequences[engine_name]
fpga_array_write_instruction = engine_sequence.instruction_set.fpga_array_write
write_mem_map_instruction_statement = engine_sequence.add_instruction("Write FPGA-Memory Map",
10, fpga_array_write_instruction.id)
write_mem_map_instruction_statement.set_parameter(fpga_array_write_instruction.fpga_memory_
map.id, hvi_memory_map)
write_mem_map_instruction_statement.set_parameter(fpga_array_write_instruction.fpga_memory_map_
```



```
offset.id, 0)
write_mem_map_instruction_statement.set_parameter(fpga_array_write_instruction.value.id, 10)
```

Actions, Events, and Triggers are treated in a different way.

The Actions, Events, and Triggers available in the FPGA-Sandbox are provided by the Instrument the FPGA is located in.

You define and use these the same way as you do other instrument Actions, Events and Triggers. For more information, see your instrument documentation.

The following code shows how to use an `action_execute` instruction to execute an Action:

```
# First, the Action that goes to the FPGA-Sandbox must be added to the Engine:
#
# Get the Action ID using the instrument's API, e.g.:
action_id = instrument.hvi.actions.user_sandbox3
#
# Specify a name for the Action to be used in the context of your HVI program
action_name = "MySandboxAction"
#
# Add the Action to the Engine of the instrument
action = system_definition.engines[engine_name].actions.add(action_id, action_name)
#
# Then, use the Action in your Sequence
engine_sequence = multi_sequence_block_statement.sequences[engine_name]
action_execute_instruction = engine_sequence.instruction_set.action_execute
action = engine_sequence.engine.actions[action_name]
instruction = sequence.add_instruction("Execute Action", 10, action_execute_instruction.id)
instruction.set_parameter(action_execute_instruction.action.id, action)
```

NOTE

When you write HVI Sequences, you must use the same names you used in the PathWave FPGA project to access the HVI FPGA resources, FPGA-Memory Maps, FPGA-Registers and FDS ports.

Load to Hardware

The `.k7z` internal *bitfile* is automatically loaded into the hardware at this stage if it is not already loaded. Once it has been loaded, in addition to running the HVI Sequence to control FPGA resources in real-time, you can also access some of the HVI FPGA resources from software, for instance writing to an HVI FPGA-Memory Map:

```
# Load or Deploy Hvi instance to hardware. At this step the k7z is loaded, if it is not already
loaded
hvi_instance.load_to_hw()
#
# Write to HVI FPGA Memory Map, in this example 0 is the offset and 1 is the data.
sandbox.fpga_memory_maps["memory_map_name"].write( 0 , 1 )
#
```

```
# Run the Hvi Sequence to use/control FPGA-Sandbox resources in real-time as described in the  
Sequence  
hvi_instance.Run()
```

HVI FPGA-Memory Maps and HVI FPGA-Register Banks in FPGA-Sandboxes

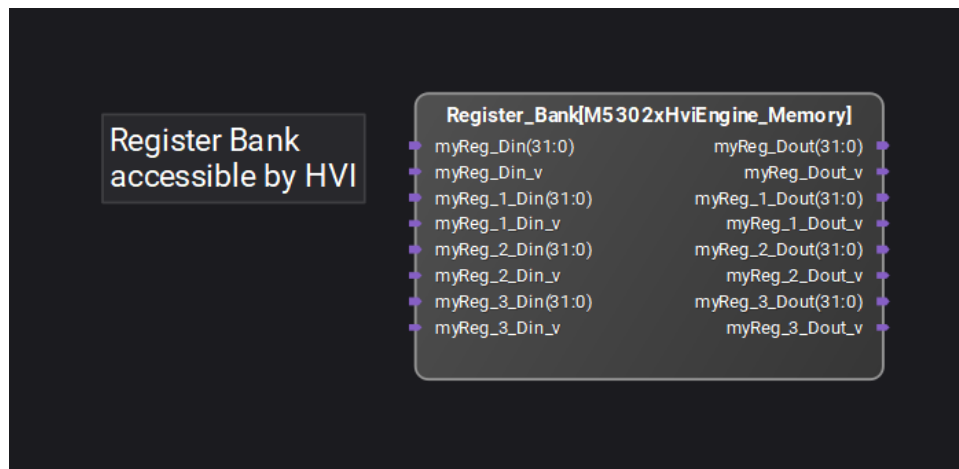
This section describes the HVI FPGA-Registers and HVI FPGA-Memory Maps that you can add to FPGA-Sandboxes in PathWave FPGA.

HVI FPGA-Registers

HVI FPGA-Registers are user-defined hardware registers that are similar to Variables in a programming language. Physically, registers are small hardware memories located in the FPGA-Sandbox. The FPGA-Registers can be accessed and modified by both HVI instructions in real-time during Sequence execution, and can be written in HVI software calls.

The FPGA-Registers can be used for destinations or sources of data, and FPGA-Registers can be treated as signed or unsigned. The size is 32 bits and numerical values must be within the signed or unsigned range. FPGA-Registers are not required to be used for numerical values, you can use the 32 bits however you wish. You can add multiple FPGA-Registers at once as a register bank.

The following image shows a register bank in PathWave FPGA:



In the image, **Din_v** and **Dout_v** indicate signals.

Din_v is used to specify when a value is valid, so the bank will update the internal value of the register for when it is being read.

Dout_v indicates when the **Dout** value is valid, so it can be used by your custom logic.

FPGA-Register Read example

The instruction `fpga_register_read` is an HVI-native instruction that enables you to read from an FPGA-Registers in an FPGA-Sandbox, the destination must be an HVI Register (not an FPGA-Register).

The following code example shows an `fpga_register_read` instruction

```
# Read FPGA-Register into an HVI Register
#
fpga_register = hvi.sync_sequence.engines["engine_name"].fpga_sandboxes["sandbox_name"].hvi_
registers["sandbox_register"]
readFpgaReg0 = sequence.add_instruction("Read FPGA Register_Bank_HviAction4Cnt", 10,
sequence.instruction_set.fpga_register_read.id)
readFpgaReg0.set_parameter(sequence.instruction_set.fpga_register_read.destination.id, hvi_
register)
readFpgaReg0.set_parameter(sequence.instruction_set.fpga_register_read.fpga_register.id, fpga_
register)
```

FPGA-Register Write example

The instruction `fpga_register_write` is an HVI-native instruction that enables you to write to an FPGA-Register in an FPGA-Sandbox. The value to be written to the register is taken from an HVI Register or from a literal.

The following code example shows an `fpga_register_write` instruction:


```
# Write to an HVI Register from an HVI Register used in an HVI Sequence
#
writeFpgaReg0 = seq.add_instruction("Write FPGA Register_Bank_HviPxiTrigOut", 50,
hvi.instruction_set.fpga_register_write.id)
writeFpgaReg0.set_parameter(seq.instruction_set.fpga_register_write.fpga_register.id, fpga_
register)
writeFpgaReg0.set_parameter(seq.instruction_set.fpga_register_write.value.id, hvi_register)
```

HVI FPGA-Memory maps

An HVI FPGA-Memory Map is an interface you add to an FPGA-Sandbox that enables you to connect HVI Sequences to a memory in the FPGA-Sandbox, or to custom logic that includes a memory block. The interface specifies a location and size that you define. The memories are always accessed 32 bits at a time.

To use the interface in HVI Sequences, you must use the same name that you used in PathWave FPGA, otherwise you will not be able to access the memory.

The following image shows an FPGA-Memory Map in PathWave FPGA:



HVI FPGA-Memory Map Read example

The following code example shows an HVI FPGA-Memory Map this is read with an `fpga_array_read` instruction. The destination is always an HVI register:

```
# Read FPGA-Memory Map
#
readMemoryMap = sync_block_1.sequences["engine_name"].add_instruction("Read FPGA-Memory Map",
20, hvi.instruction_set.fpga_array_read.id)
readMemoryMap.set_parameter(seq.instruction_set.fpga_array_read.fpga_memory_map.id, hvi_memory_
map)
readMemoryMap.set_parameter(seq.instruction_set.fpga_array_read.destination.id, register)
readMemoryMap.set_parameter(seq.instruction_set.fpga_array_read.fpga_memory_map_offset.id, 0)
```

HVI FPGA-Memory Map write example

The following code example shows an HVI FPGA-Memory Map that is written by a `fpga_array_write` instruction. The source can be a literal or an HVI register:


```
# Write Memory Map
#
writeMemoryMap = sync_block_1.sequences["engine_name"].add_instruction("Write FPGA-Memory Map",
10, seq.instruction_set.fpga_array_write.id)
writeMemoryMap.set_parameter(seq.instruction_set.fpga_array_write.fpga_memory_map.id, hvi_
memory_map)
writeMemoryMap.set_parameter(seq.instruction_set.fpga_array_write.value.id, register)
writeMemoryMap.set_parameter(seq.instruction_set.fpga_array_write.fpga_memory_map_offset.id, 0)
```

For more information, see [Local Statements](#).

Actions, Events and Triggers in an FPGA-Sandbox

A number of interfaces for HVI Actions, HVI Events, and HVI Triggers are available in PathWave FPGA that you can add to an FPGA-Sandbox. These are provided by the instrument that the FPGA is located in.

The following image shows HviAction, HviEvent, and HviTriggerIoIn interfaces:



Actions

You add Actions to enable you to send signals into the FPGA-Sandbox from your HVI Sequences. For example, you can use an Action to tell the logic in an FPGA-Sandbox to send a signal.

Events

You add Events to inform your HVI Sequences of Events in the FPGA-Sandbox.

For example, you can use an Event to get the FPGA to inform your HVI that an external signal has been received in the FPGA-Sandbox, or a signal has been generated in the FPGA-Sandbox.

You use the Wait-for-Event Statement to command your HVI to wait until an Event occurs. A signal in the FPGA-Sandbox can initiate the Event.

Triggers

You can add Triggers that go into or out of the FPGA-Sandbox.

For example, a Trigger going into the FPGA-Sandbox can initiate an Event in the FPGA-Sandbox.

You can also use an Action to initiate the logic in the FPGA-Sandbox to send a Trigger out of the FPGA-Sandbox, or you can use an Event that is initiated when a Trigger has arrived.

Depending on the Instrument, there may be a number of different Trigger types available, for example:

- HviTriggerIoIn.
- HviTriggerIoOut.
- HviTriggerIoT.
- HviTriggerOutToLvds.
- LvdsToHviTriggerIn.

For a list of Actions, Events and Triggers available for an instrument, see your instrument documentation.

Using FPGA-Sandbox Actions, Events and Triggers in HVI Sequences

You can use the Actions, Events and Triggers that you added to an FPGA-Sandbox in your HVI Sequences. For Actions, Events and Triggers this is the same as you use any other instrument Actions, Events and Triggers, except for Triggers and Events where you must set the source to be `fpga_sandbox`.

For example, for Actions:

- Use the `ActionDefinition` object to define the Action.
- Add the definition to the `ActionDefinitionCollection`.
- Add the Action to an HVI Engine with the add method of the `ActionCollection`.
- In your Sequence, add an Action with `InstructionsActionExecute`.

FPGA Fast Data Sharing

HVI enables FPGAs on instruments to communicate with each other using *Fast Data Sharing* (FDS) technology.

Fast Data Sharing (FDS) is a technology that enables you to share data between FPGA-Sandboxes with a known fixed low latency. You can share data during the execution of Sequences between FPGA-Sandboxes in different instruments in the same, or different chassis. The data sharing is performed using *Sync FPGA Data-Sharing* Statements.

To take advantage of this FDS, you must use **PathWave-FPGA** to create a design in an FPGA-Sandbox that includes FDS ports.

Communication with Fast Data Sharing

FDS enables you to move data such as register values, or values of items such as qubit states. The data can travel between instruments on System Sync cables or on the PXIe DSTARB or DSTARC lines inside a chassis. FDS requires *System Synchronization Modules* (SSM) and PXIe instruments that support FDS technology. An advantage of FDS is that it does not use up additional triggers beyond those PathWave Test Sync Executive requires, so you are not required to reserve any additional triggers to use FDS.

HVI supports the following kinds of FDS transfers:

- Sharing FPGA-Sandbox data with Sync FPGA Data-Sharing Statements.

For FDS enabled instruments, **Pathwave-FPGA** provides the interfaces to use FDS. Timing and routing information is provided by the instruments.

HVI guarantees that the data is sent in the correct order, and that the communication timing and routing is computed automatically by HVI. HVI also automatically calculates the optimal communication timing to avoid collisions when data is transferred.

Accessing FPGA FDS Ports

When a PathWave FPGA project (**.k7z** file) is loaded, the HVI FPGA-Memory Maps, HVI FPGA-Registers, and FDS ports are populated under the sandbox object.

You can access the list of FDS Port locations (**FdsPort** objects) defined in the FPGA-Sandbox by using an FPGA-Sandbox Definition object that is loaded from the **.k7z** file.

The **FdsPort** object enables you to use the FDS port instances placed in the FPGA-Sandbox of a loaded PathWave FPGA project. An **FdsPort** has one property which is the name of the port.

The **FdsPort** can be set as a parameter in a **SyncFpgaDataSharing** Statement.

The following example shows how to use and program the FDS transactions in the `syncFpgaDataSharing` Statement.

The `FdsPortAddress` object enables you to specify both the source and the destination for each FDS transaction. This is done by specifying the name of the FDS port connected in PathWave FPGA to the block transmitting/receiving the data over FDS. The address where the data is read/written is also specified. Once the source and destinations are specified, each transaction can be added to the `syncFpgaDataSharing` Statement by specifying how many bits are shared in each transaction.

Python example:

```
# SyncFpgaDataSharing definition with 3 transactions
#
# Retrieve ports
instrument1_fds_ports = sequencer.sync_sequence.engines[instrument1_engine_name].fpga_sandboxes
[0].fds_ports
instrument2_fds_ports = sequencer.sync_sequence.engines[instrument2_engine_name].fpga_sandboxes
[0].fds_ports
instrument3_fds_ports = sequencer.sync_sequence.engines[instrument3_engine_name].fpga_sandboxes
[0].fds_ports
#
# Sources
instrument1_tx = kthvi.FdsPortAddress(instrument1_fds_ports[instrument1_tx_port_name], src1_
address)
instrument2_tx = kthvi.FdsPortAddress(instrument2_fds_ports[instrument2_tx_port_name], src2_
address)
#
# Destinations
instrument2_rx = kthvi.FdsPortAddress(instrument2_fds_ports[instrument2_rx_port_name], dst2_
address)
instrument3_rx = kthvi.FdsPortAddress(instrument3_fds_ports[instrument3_rx_port_name], dst3_
address)
#
# Adding Sync FPGA Data Sharing Statement
fpga_data_sharing = sequencer.sync_sequence.add_sync_fpga_data_sharing("my statement", start_
delay)
# Transaction 1
fpga_data_sharing.transactions.add(instrument1_tx, instrument2_rx, num_bits_to_share)
#
# Transaction 2
fpga_data_sharing.transactions.add(instrument2_tx, instrument3_rx, num_bits_to_share)
#
# Transaction 3
fpga_data_sharing.transactions.add(instrument1_tx, instrument3_rx, num_bits_to_share)
```

FPGA-Instruction

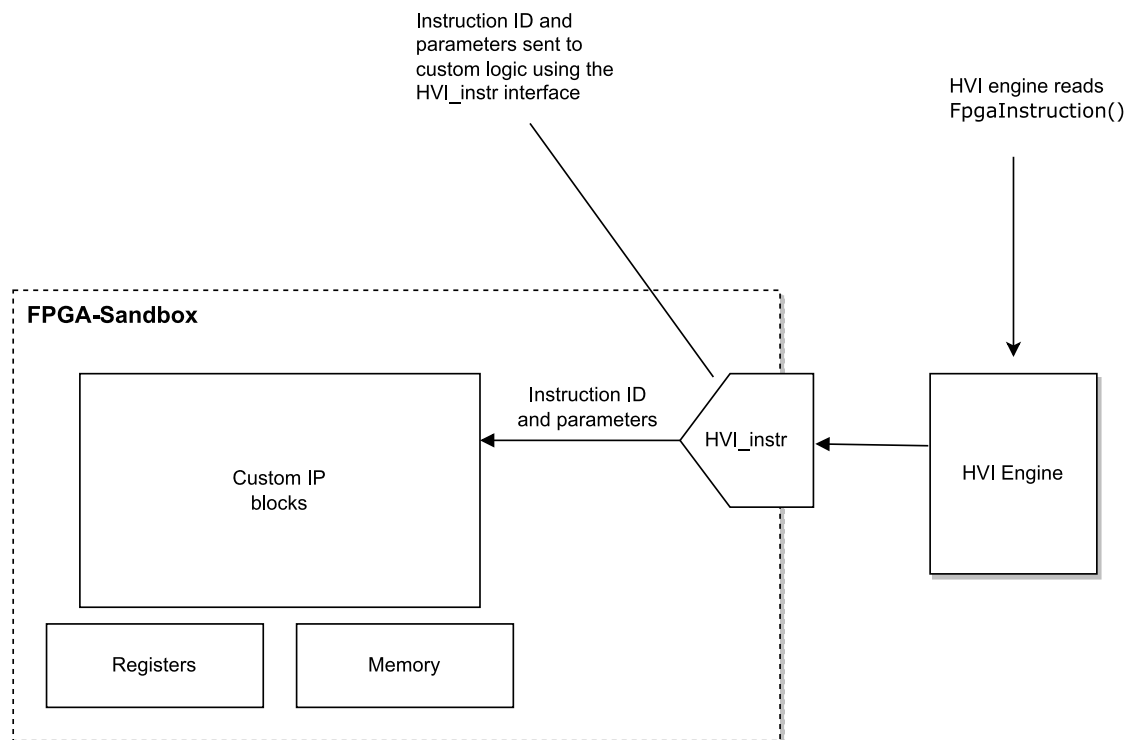
The FPGA-Instruction statement enables you to issue commands to custom FPGA-Sandbox logic from within HVI sequences.

You can customize logic in an FPGA sandbox using PathWave FPGA to create different functions. When you do this, adding an **HVI_Instr** interface enables you to interface with your logic from HVI Sequences.

You use the FPGA-Instruction Statement in your Sequences to issue the commands into the FPGA-Sandbox to utilize the different functionality. This means you can setup custom commands with different functions in the FPGA-Sandbox and utilize them in HVI Sequences.

When an HVI Engine executes an FPGA-Instruction, it also reads the parameters and the instruction ID and passes this data to the **HVI_Instr** interface in the FPGA-Sandbox, this interfaces an instruction parser and your logic.

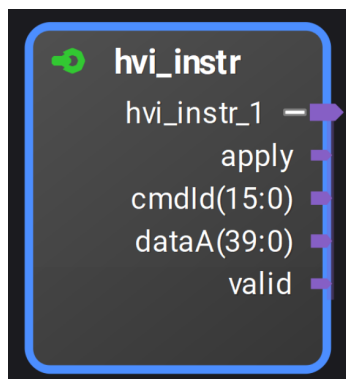
This flow is shown in the following diagram:



Integrating the FPGA-Instruction Statement with FPGA-Sandbox Logic

If you want to issue commands to your logic with the `FpgaInstruction` Statement, you must add an HVI instruction interface to the FPGA-Sandbox. Your logic must receive the parameters provided and then decode and execute the commands. How this is done depends on the instrument you are using, see your instrument documentation for more information.

The following image shows the HVI instruction interface `hvi_instr` as it appears in PathWave FPGA:



The signals are:

apply:

A flag that is typically used to apply stored configuration data in a multi-step setup process.

cmdId:

Command identifier. This is useful when more than one command is supported by the custom logic.

dataA:

General purpose data, 40 bits wide.

valid:

A flag used to identify when the data on the other ports is valid.

The following example shows an FPGA-Instruction Statement:

```
# Set up local Sequence
fpga_inst = local_sequence.instruction_set.fpga_instruction
instruction = local_sequence.add_instruction('fpgaInstruction', 10, fpga_inst.id)
#
port_number = 2
data_a = 1234
command_id = 5
apply = 1
#
instruction.set_parameter(fpga_inst.port_number.id, port_number)
instruction.set_parameter(fpga_inst.data_a.id, data_a)
instruction.set_parameter(fpga_inst.command_id.id, command_id)
instruction.set_parameter(fpga_inst.apply.id, apply)
```

For more information see [Local Statements](#).

HVI Statements for using FPGAs

PathWave Test Sync Executive includes a number of FPGA-specific HVI Statements you can use to interact with the FPGA on an instrument:

Local Statements

FPGA-Register Read

The instruction `fpga_register_read` is an HVI Native Instruction that enables you read from an HVI FPGA-Register to a destination HVI Register.

FPGA-Register Write

The instruction `fpga_register_write` is an HVI Native Instruction that enables you to write an HVI FPGA-Register placed in an FPGA-Sandbox. The value to be written to the HVI FPGA-Register is taken from an HVI Register or from a literal.

FPGA-Memory Map Write

The instruction `fpga_array_write` is an HVI Native Instruction that enables you to write to an HVI FPGA-Memory Map that is located in an FPGA-Sandbox. The value to be written to the HVI FPGA-Memory Map is taken from an HVI Register or from a literal.

FPGA-Memory Map Read

The instruction `fpga_array_read` is an HVI Native Instruction that enables you to read from an HVI FPGA-Memory Map. The value read from the HVI FPGA-Memory Map is written to a destination HVI Register.

FPGA-Instruction Statement

The `fpgaInstruction` Statement enables you to issue commands to your custom FPGA-Sandbox logic from within HVI Sequences. This is an HVI Native Instruction, but it can only be used successfully on instruments that support it.

For more information, see [Local Statements](#).

Sync Statements

Sync FPGA Data-Sharing Statement

The Sync FPGA Data-Sharing Statement enables you to transfer data between FPGA-Sandboxes.

For more information, see [Sync Statements](#).

Chapter 6: Multi-Chassis Systems and System Synchronization Modules

This chapter describes how you use System Synchronization Modules to synchronize a Multi-Chassis System. It contains the following sections:

- [System Synchronization Modules](#)
- [Configuring a System with SSMs and System Sync Connectivity](#)
- [Clocking](#)
- [Configuring the Reference Clock](#)

For information about troubleshooting a multi-chassis system, see the *System Setup Guide*.

System Synchronization Modules

This section describes System Synchronization Modules.

KS2201A PathWave Test Sync Executive includes multi-chassis topologies that use the Keysight M9032A/M9033A PXIe *System Synchronization Modules* (SSMs). The previous means of interconnecting multiple PXI chassis using M9031A modules was discontinued starting from the KS2201A 2022 release. Compared to the discontinued M9031A module, the SSM has a much wider range of functions including:

- Distribution of a precise reference clock.
- Management of *Fast Data Sharing* (FDS).
- Chassis interconnectivity.
- Synchronization of all the PXI instruments in the multi-chassis.

M9032A and M9033A PXIe SSM Overview


The M9032A/M9033A are PXIe *System Synchronization Modules* (SSM). These include an onboard high-quality 10MHz Oven Controlled Crystal Oscillator (OCXO) to achieve a very precise synchronization among various measurement instruments distributed across different chassis. The M9032A/M9033A System Synchronization Module functionalities can only be successfully deployed on chassis compliant with the *PXI-Express* (PXIe) standard. The SSM must be inserted in the timing slot of the PXIe chassis.

Keysight PXIe System Synchronization Module is available in two form factors, which only differ in their connectivity capabilities:

- M9032A is a one-slot PXIe System Synchronization Module with 1 System Sync Upstream and 1 System Sync Downstream ports.
- M9033A is a two-slot PXIe System Synchronization Module with 1 System Sync Upstream and 4 System Sync Downstream ports.

For further information about these SSMs including detailed performance specifications, see the [M9032A/M9033A User's Guide](#), available at [Keysight PXI Products](#).

The following image shows the physical M9032A and M9033A SSMs:



System Sync Module Connectivity

Front Panel

The M9032A and M9033A Front Panel contains various connectors that can be used for both multi-chassis interconnection and configuration of the reference clock source.

Front Panel *Sub Miniature Push-on* (SMP) IOs

Front Panel (FP) SMP connectors are:

SCLK/Ref Out:

Outputs a copy of the system clock or a reference clock signal.

STrig/Trig IO:

Receives an arbitrary trigger signal.

SCLK/Ref In:

Receives the reference clock signal.

PPS/Time Ref:

Receives a Pulse Per Second (PPS) signal.

The front panel SMP connectors can be used to share input and output reference clocks.

System Sync ports

System Sync ports use PCIe *Optical Copper Link* (OCuLink) connectors. System Sync ports are used for chassis interconnection and synchronization in the multi-chassis system. The signals in the System Sync include:

- Clocking (System Sync only).
- Triggering.
- Data.

The different SSM models have the following front panel System Sync ports:

The M9032A has 2 System Sync ports:

- 1 System Sync Upstream.
- 1 System Sync Downstream.

The M9033A has 5 System Sync ports:

- 1 System Sync Upstream.
- 4 System Sync Downstream.

Each System Sync Downstream port can connect to the System Sync Upstream port of another SSM placed in a different chassis. For more information, see the section below about Inter/Intra-chassis Connectivity.

PXIe Backplane DSTAR Connectivity

The M9032A and M9033A are placed in the Timing Slot of a PXIe chassis which enables them to support the DSTAR connectivity built into the chassis.

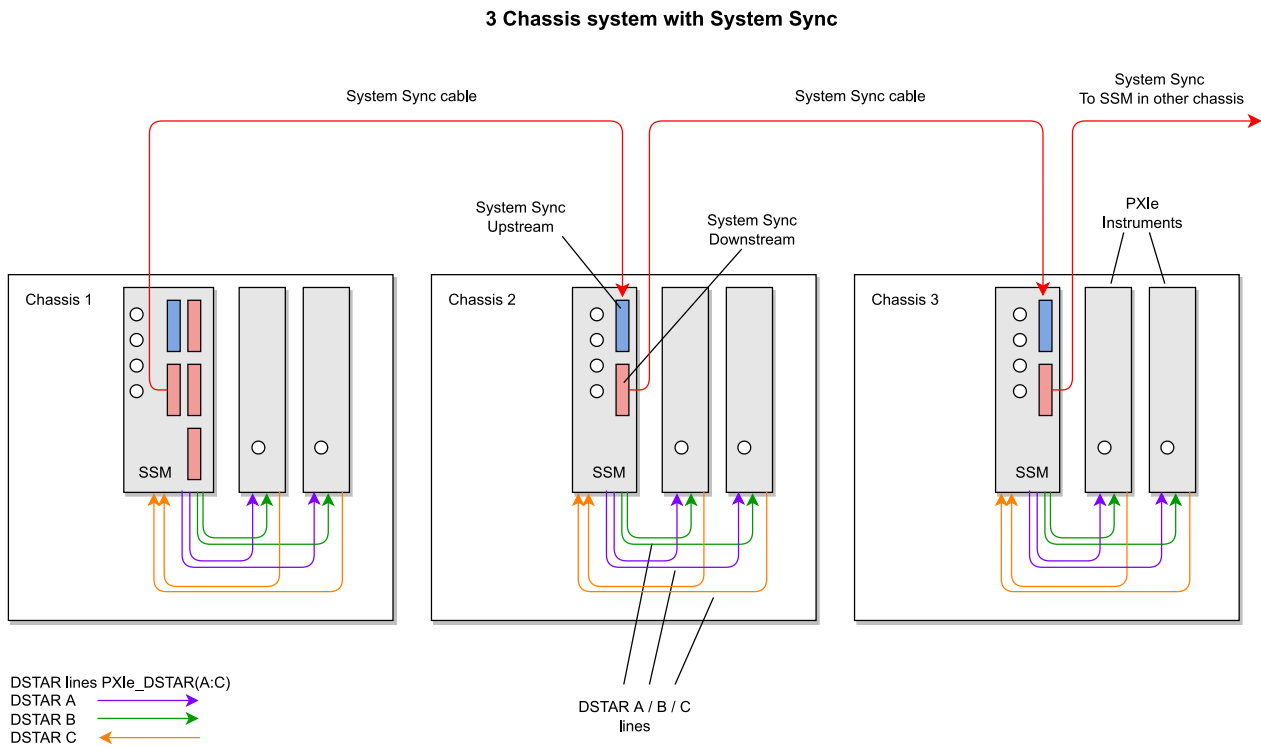
DSTARA/B/C are multi-instrument point to point connections inside a chassis. DSTARA is used to carry the clock signal. DSTARB and DSTARC carry trigger or data signals.

Inter/Intra-chassis connectivity, Synchronization and Data-Sharing Functionality

An SSM can enable both multi-chassis and multi-instrument interconnections. With these connections, SSMs enable synchronization and data sharing across all the instruments in a multi-chassis system.

- Multi-chassis interconnections are made with System Sync connections using their capability to interconnect two SSMs together through their System Sync Downstream/Upstream ports.
- Intra-chassis, multi-instrument interconnections are made with PXIe DSTARA/B/C connections. The SSM can share the precise reference clock over the DSTARA signal.

The following diagram shows a 3 chassis system connected with System Sync cables and DSTARA/B/C signals in each chassis:



Data can be shared across System Sync and DSTAR connections in several different ways:

- The reference clock can be shared between two interconnected SSMs using the System Sync connection between System Sync Downstream/Upstream ports.
- The System Sync connection can share the signals sent over the PXI_TRIG[0:7] trigger buses, from one SSM to the next. This enables the SSMs to share PXI sync resources used by PathWave Test Sync Executive for the *Hard Virtual Instrument* (HVI) across the different chassis.

- System Sync connections can route data shared using *Fast Data Sharing* (FDS) between PXIe instruments.
- The SSM can send the data between two modules located in the same chassis using the DSTARB/C signals.
- Data can be sent through the System Sync connections to route it to instruments located in a different chassis.

Configuring a System with SSMs and System Sync Connectivity

This section describes how you use System Synchronization Modules to synchronize a Multi-Chassis System.

In a multi-chassis system connected with Keysight PXIe System Synchronization Modules, you must include one SSM in each chassis that is part of the system. Each SSM must be inserted in the **timing slot** of your chassis. This is typically slot 10 in Keysight 18-slot chassis, but it can be a different slot number in different chassis models. The SSMs are connected to each other with System Sync cables.


One SSM is automatically chosen as a leader and it is used to synchronize all the instruments in the multi-chassis system. The SSM chosen as leader is the SSM that has no incoming connection to its System Sync Upstream port. The leader SSM distributes a replica of the reference clock signal to the SSMs located in the other chassis. It does this through point-to-point connections between System Sync Downstream/Upstream ports. In the example multi-chassis system shown in the following diagram, the leader SSM is in Chassis 1.

A multi-chassis PXIe system may be configured to use many different reference options. For a list of those options and descriptions of how to configure them, see the section *Clocking* in this document. For one of those reference options, an SSM is chosen as a leader and uses its internal Oven Controlled Crystal Oscillator (OCXO) clock to synchronize all the instruments included in the multi-chassis system.

NOTE

A Multi-chassis system based on the older M9031A modules required an external reference clock generator to distribute the precise common 10 MHz reference clock signal across different chassis. In the multi-chassis topology delivered by PathWave Test Sync Executive, the SSM assumes the function of the **reference clock signal generator/distributor**, by sharing a reference clock generated by an internal PLL. This PLL can be fed by different sources (as explained later in this document) including the OCXO inside the SSM, which generates a 10 MHz sine wave. An external 10 or 100 MHz reference signal can still be connected to the SSM SClk / Ref Input port, to sync it together with other clusters.

The following diagram shows an example of a 6 chassis system connected with SSMs. In this system an M9033A SSM in chassis 1 distributes the reference clock to four M9032A SSMs located in each of the other chassis. The SSM in chassis 5 also forwards the clock to a sixth chassis.



This following code shows how to use the HVI Python API to define and use the SSMs in the multi-chassis system shown in the diagram. Each System Sync Downstream port connects to the System Sync Upstream port of another System Sync Module in a different chassis.

The first step is to define the SSMs placed in each of the chassis during the System Definition phase.

```
# Create SystemDefinition object
my_system = kthvi.SystemDefinition("MySystem")
#
# Define System Sync Modules
resource_id_ssm_1 = 'PXI0::CHASSIS1::SLOT10::INSTR'
resource_id_ssm_2 = 'PXI0::CHASSIS2::SLOT10::INSTR'
resource_id_ssm_3 = 'PXI0::CHASSIS3::SLOT10::INSTR'
resource_id_ssm_4 = 'PXI0::CHASSIS4::SLOT10::INSTR'
resource_id_ssm_5 = 'PXI0::CHASSIS5::SLOT10::INSTR'
resource_id_ssm_6 = 'PXI0::CHASSIS6::SLOT10::INSTR'
#
# In the options, SSMs are set to be simulated with Simulate=true and there are a number of
parameters.
# For the hardware SSM instruments, set options to an empty string.
options1 = "Simulate=true,DriverSetup=Model=M9033A"
options2 = "Simulate=true,DriverSetup=Model=M9032A"
options3 = "Simulate=true,DriverSetup=Model=M9032A"
options4 = "Simulate=true,DriverSetup=Model=M9032A"
options5 = "Simulate=true,DriverSetup=Model=M9032A"
options6 = "Simulate=true,DriverSetup=Model=M9032A"
#
sync_module_1 = my_system.interconnects.add_sync_module(resource_id_ssm_1, options1)
sync_module_2 = my_system.interconnects.add_sync_module(resource_id_ssm_2, options2)
sync_module_3 = my_system.interconnects.add_sync_module(resource_id_ssm_3, options3)
sync_module_4 = my_system.interconnects.add_sync_module(resource_id_ssm_4, options4)
sync_module_5 = my_system.interconnects.add_sync_module(resource_id_ssm_5, options5)
sync_module_6 = my_system.interconnects.add_sync_module(resource_id_ssm_6, options6)
```

NOTE

In the HVI System Definition phase, the SSMs are added to the interconnects collection by using their resource ID and options. Same as for the chassis, it is not necessary to open objects representing the SSMs that are included in the multi-chassis system.

The next step is to define the interconnections among the System Sync Downstream/Upstream ports of each pair of SSMs. The SSM System Sync ports can only be connected Downstream to Upstream.

```
# Define connections among System Sync connectors of the SSMs
#
# Connect SSM 1 to SSM 2
ssm1_downstream_sync1 = sync_module_1.connectivity.systemsync_downstream[0]
ssm2_upstream_sync    = sync_module_2.connectivity.systemsync_upstream[0]
ssm1_downstream_sync1.set_connection(ssm2_upstream_sync)
#
# Connect SSM 1 to SSM 3
ssm1_downstream_sync2 = sync_module_1.connectivity.systemsync_downstream[1]
ssm3_upstream_sync    = sync_module_3.connectivity.systemsync_upstream[0]
ssm1_downstream_sync2.set_connection(ssm3_upstream_sync)
#
# Connect SSM 1 to SSM 4
ssm1_downstream_sync3 = sync_module_1.connectivity.systemsync_downstream[2]
ssm4_upstream_sync    = sync_module_4.connectivity.systemsync_upstream[0]
ssm1_downstream_sync3.set_connection(ssm4_upstream_sync)
#
# Connect SSM 1 to SSM 5
ssm1_downstream_sync4 = sync_module_1.connectivity.systemsync_downstream[3]
ssm5_upstream_sync    = sync_module_5.connectivity.systemsync_upstream[0]
ssm1_downstream_sync4.set_connection(ssm5_upstream_sync)
#
# Connect SSM 5 to SSM 6
ssm5_downstream_sync = sync_module_5.connectivity.systemsync_downstream[0]
ssm6_upstream_sync   = sync_module_6.connectivity.systemsync_upstream[0]
ssm5_downstream_sync.set_connection(ssm6_upstream_sync)
```

Chassis Supported for Multi-Chassis Systems

The following Keysight chassis models are supported:

- M9018B
- M9019A
- M9046A

Software and firmware version requirements are listed on-line here: [Chassis Software and Firmware Requirements for KS2201A](#) .

NOTE

If you mix different chassis models in your multi-chassis setup, you may observe some skew across the different chassis and different performance depending on the different chassis characteristics.

Non Keysight chassis are not supported for multi-chassis systems.

Clocking

This section describes Clocking:

Types of Clock

Clock Types

In a single or multi-chassis system there are 4 types of clocks used for synchronization and instrument-related tasks:

- Reference clock.
- System clocks.
- Analog clocks.
- Sample clocks.

All these clocks are synchronous with one another, but are used for different purposes and can be configured in different ways trading off performance and complexity/cost.

Reference Clock

The Reference clock used as reference for generation of other clocks, it determines the absolute frequency and lowest-frequency offset phase noise performance of the analog instrumentation's inputs and outputs. That is because all of the other clocks are phase-locked to the Reference Clock. A PXIe system can either use its own internal reference clock or phase-lock to an external reference clock. It can also provide external reference clock outputs for other instrumentation to phase-lock to.

System Clocks

The relevant clocks that drive the internal logic of individual instruments. These include clocks automatically reported to HVI and clocks users reports explicitly using the TSE API. The System clocks synchronize all the digital operation of all instruments and the PXIe platform. These clocks are derived from the Reference Clock and are used by, for example, the PathWave FPGAs Sandbox logic, the HVI Engine core clock, Fast Data Sharing and other digital capabilities in the instruments. Basically, a system clock is clock that is neither the reference clock nor an analog clock.

Analog Clocks

Clock responsible for generating the analogue signals in an instrument, the analog signal specifications such as phase noise, depend on the specs of this clock. The Analog Clocks are intermediate frequency clocks from which the instrument's Sample Clocks are derived. Like the Sample clocks, the Analog Clocks affect the overall phase noise performance and skew drift of the instrument analog inputs and outputs. In the simplest clock configurations, each peripheral module generates

its own independent Analog Clock. In the highest fidelity clock configuration, a single common Analog clock is generated by the *High Performance Reference Clock Source* (HPRCS) and is distributed to all the individual peripheral modules through external cables and power dividers.

Sample Clocks

The instrument's ADCs and DACs that digitize analog input signals and generate analog output signals are clocked by their own internal Sample Clocks. The various types of peripheral modules use different sample clock frequencies even though they are ultimately derived from the same Reference clock. These sample clocks determine the overall phase noise performance and skew drift of the analog inputs and outputs because they directly clock the instrument's ADCs and DACs.

System Clock Distribution using SSM and System Sync connectivity

In a multi-chassis system based on the Keysight PXIe SSMs and chassis, the SSM with no other SSM connected to its System Sync Upstream port acts as the leader. This leader SSM forwards a copy of the system clock to other SSMs using System Sync cables. In turn, each SSM shares the forwarded system clock with the instruments located in their respective chassis using the PXIe DSTARA backplane signal.

NOTE

You are not required to set the leader in the TSE API. The leader SSM is determined by the hardware connections. That is, the leader role is automatically taken by the SSM that has no System Sync cable connected to its System Sync Upstream port.

Overview of Supported Clocking Schemes

There are several possible different clocking configurations, the one you should use depends on the hardware and the application requirements. Some of the key aspects to consider when selecting a clocking scheme are:

1. **System and Analog clock sources.** The source for the System and intermediate-frequency analog clocks is a critical element that determines the system synchronization, phase noise and drift performance. The clock sources covered in this section include:
 - a. PXIe chassis.
 - b. System Sync Module.
 - c. PXIe Chassis with *High Performance Reference Clock Source* (HPRCS). This is only available on Keysight PXIe chassis models M904xA.
2. **Internal/external Reference clock.** The clock that serves as reference for the System/Analog clocks can be generated internally by the selected source, or externally provided by the user, generated by a clock source external to the PXIe system. In systems that include the *High Performance Reference Clock Source* (HPRCS), and other external instrumentation that you wish to

share a common Reference Clock, the best overall jitter performance will usually be achieved by phase-locking the other external instrumentation to the HPRCS Reference Clock instead of the other way around. If the overall system needs to be phase-locked to a GPS or atomic standard reference, you should phase-lock the HPRCS to the GPS or atomic standard and phase-lock all the other instrumentation to the HPRCS Reference Clock.

3. **Instruments internal/external Analog Clock.** Most instruments can either use an external Analog Clock or generate their own Analog Clock internally for convenience, however, using a common external Analog Clock will always provide the best performance because all peripheral module sample clocks will jitter and drift together.

The following table shows the different supported/recommended clocking schemes:

Clocking Scheme	Reference Clock Source	Reference Clock Mode	Description	Performance
A.1 Single-chassis, no SSM.	Chassis	Internal 10MHz	An OCXO inside the chassis generates a 10 MHz reference clock. Independent Analog clocks are generated in each peripheral module.	See the chassis datasheet for exact phase noise performance. See the M5xxx PXIe instrument documentation for exact performance of channel to channel skew, jitter, and drift.
	External	External 10MHz	The external reference clock must have a frequency of 10 MHz. As an example, it can come from a <i>Device Under Test</i> (DUT), another instrument that is part of the setup, etc. Independent Analog clocks are generated in each peripheral module.	-
A.2 Single/multiple chassis with at least one M904x, Analog clocks, SSMs.	Chassis	Internal 10MHz	An OCXO inside the chassis generates a 10 MHz reference clock. A common Analog clock is externally distributed to each peripheral module.	See the chassis datasheet for exact phase noise performance. See the M5xxx PXIe instrument documentation for exact performance of channel to channel skew, jitter, and drift.
	External	External 10MHz	The external reference clock must have a frequency of 10 MHz. As an example, it can come from a DUT, another instrument that is part of the setup, etc. A common Analog clock is externally distributed to each peripheral module.	-

Clocking Scheme	Reference Clock Source	Reference Clock Mode	Description	Performance
B No external Analog clocks, SSMs.	SSM	Internal 10MHz	An OCXO inside the SSM generates a 10 MHz reference clock. Independent Analog clocks are generated in each peripheral module.	See the SSM datasheet for exact phase noise performance. See the M5xxx PXIe instrument documentation for exact performance of channel to channel skew, jitter, and drift.
	External	External 10/100MHz	The external reference clock can have a 10 or 100 MHz frequency. As an example, it can come from a DUT, from another instrument that is part of the setup, etc. Independent Analog clocks are generated in each peripheral module.	-
C external Analog clocks, SSMs, HPRCS.	HPRCS	Internal 10MHz	The HPRCS generates a 2.4 GHz sine wave that gets divided in frequency to generate a 100 MHz reference clock signal. A common Analog clock is externally distributed to each peripheral module.	This option provides the best performance in terms of phase noise. For more information, see the <i>Keysight PXIe Chassis M9046A Datasheet</i> , available at Keysight PXI chassis .
	External	External 10/100MHz	The external reference clock for the HPRCS can have a 10 or 100 MHz frequency. As an example, it can come from a DUT or another instrument that is part of the setup, etc. A common Analog clock is externally distributed to each peripheral module.	-

Depending on the chosen clocking scheme, additional connections to the SSM may be required:

- A.1 or B Internal: No extra connections are required.
- B External: Attach the external reference clock source to the SSM **SCLK / Ref In** input.
- A.2 or C Internal: The chassis Ref1 output must be connected to the SSM **SCLK / Ref In** input.
- A.2 or C External: The chassis Ref1 clock output must be connected to the SSM **SCLK / Ref In** input. The external reference clock must be connected to the chassis external **Ref In** input.

The precise connections are described in the configuration sections later in this guide.

You may also phase-lock external instrumentation to the PXI system using either 10 MHz or 100 MHz external reference clock outputs.

Some instruments also require an analog clock. For more information see the later section titled *Configuring Analog Clock Source for Instruments*.

Configuring the Reference Clock

This section describes configuring the reference clock, it contains the following sections:

- [Configuring Clocking Scheme A.1](#)
- [Configuring Clocking Scheme B](#)
- [Configuring Clocking Schemes A.2 and C](#)
 - [Analog clock configuration options](#)
 - [M9046A Front Panel Clocking IO overview](#)
 - [Using the M9046A analog clock source](#)
 - [Using the M9046A's M9546A High Performance Reference Clock Source](#)
- [Enabling Chassis Clock Outputs](#)
 - [Enabling the chassis Analog clock](#)
 - [Configuring the Analog Clock Source in Instruments](#)

Configuring Clocking Scheme A.1

(Single-chassis, no SSM, no external analog clocks)

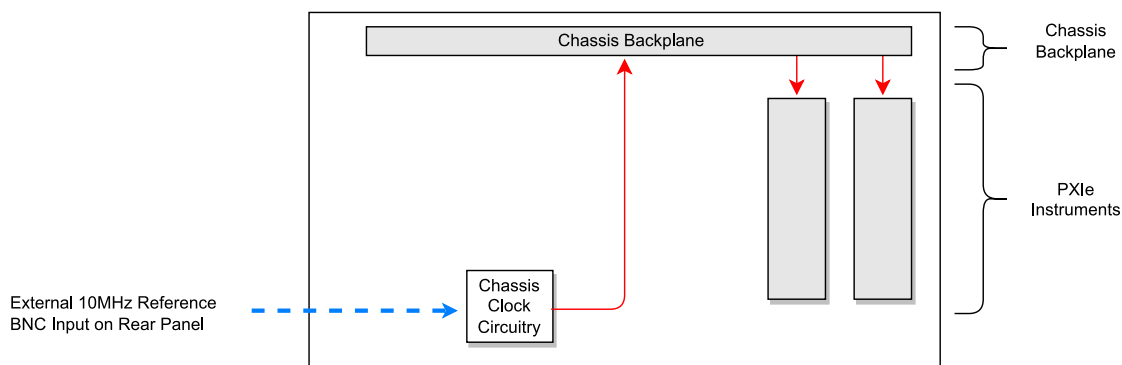
This is the simplest configuration and is the default if you have not specified another.

The chassis is the clock source. For the reference clock, there are two options:

1. **Internal (default):** This is the 10MHz clock built into the chassis (VCXO for the M9019A or OCXO for the M904xA).
2. **External:** A 10MHz signal connected to the 10MHz Ref BNC input located on the chassis rear panel.

The following diagram shows a chassis with an internal clock source (chassis clock) or an external clock source (blue):

Single chassis, no SSM Internal or External Reference clock



All chassis have on their rear panel a 10MHz reference BNC input and a 10MHz reference BNC output. In the case of the M904x chassis, there are two Reference clock SMA outputs on the front panel.

This clocking scheme is rather constrained in terms of features because it only allows for a single chassis and, given that there is no SSM, advanced features like Fast Data Sharing are not available.

The following snippet shows how to configure the chassis as the clock source:

```
# Create SystemDefinition object
my_system = kthvi.SystemDefinition("MySystem")
#
# Define chassis
chassis = my_system.add_chassis(1)
#
# Select the chassis as ref. clock source
clockSource = chassis.clock_source
#
# Set the chassis as clock source
systemDefinition.clocking.reference_source = clockSource
#
# Explicitly set the clock source to use the internal OCXO as the reference clock (this is the
default)
clockSource.set_mode(keysight_tse.ClockingReferenceMode.INTERNAL)
#
# Alternatively you can configure the chassis to use the external clock reference with the 10Mhz
frequency value in Hz
clockSource.set_mode(keysight_tse.ClockingReferenceMode.EXTERNAL, 10e6)
```

Configuring Clocking Scheme B

(Single/multi-chassis system with SSM as clock source and no external analog clocks)

Each SSM is equipped with an onboard high-quality 10MHz Oven Controlled Crystal Oscillator (OCXO) that can be used as the reference clock.

Alternatively, the chassis backplane reference clock output (Scheme A.2) or the optional *High Performance Reference Clock Source* (HPRCS) output (Scheme C) can be used as reference clock. The HPRCS option requires a Keysight PXIe Chassis model M9046A. Clocking scheme B assumes the external reference clock is neither output by the chassis nor by the HPRCS because otherwise further configurations would be required for proper operation.

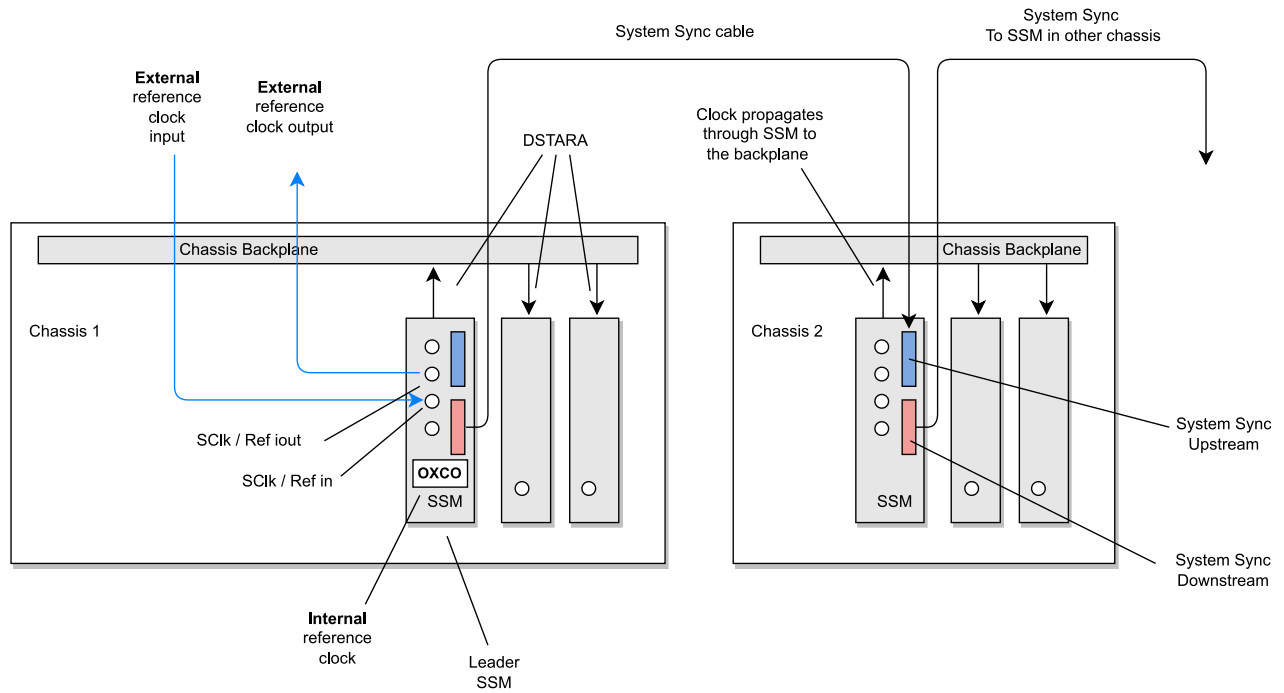
The reference clock can be chosen from two options:

1. **Internal:** This is the default mode. The internal OCXO of the leader SSM is used as the reference clock.
2. **External:** An 10 MHz or 100 MHz external reference clock is connected to the SSM's front-panel **SCLK/Ref In** SMP input.

The reference clock gets propagated to all the PXIe instruments within the same chassis through the DSTARA signal path. It gets propagated to the next SSM through the System Sync cable from the downstream connection on leader SSM to the upstream connection on the follower.

The following diagram shows the operation of the Clocking Scheme B. The clock is generated in the SSM in chassis 1 and is passed to the other instruments in the chassis via the DSTARA signal path in the backplane (red arrows). It is also passed to the next chassis via the System Sync cable (in black) where it propagates via the SSM in that chassis. The internal reference is the SSM's OCXO, and the external reference is shown in blue:

**Multiple chassis with SSMs
SSM as clock source with internal or external reference**



Configuring the SSM as the System Clock source

By default, if you do not specify anything, PathWave Test Sync Executive configures the leader SSM as the reference clock source using its internal OCXO clock. The leader SSM is defined by the hardware connections. In the TSE API, no additional definition other than the connections between SSM is required to identify the leader SSM. You must ensure the connections you define in software match the physical hardware connections between SSMs.

The following code shows how to configure a pair of chassis with SSMs where the OCXO clock is the reference clock source, `options` is set to an empty string:

```
# Create SystemDefinition object
my_system = kthvi.SystemDefinition("MySystem")
#
# Define all necessary follower SSMs depending the number of chassis
leader_ssm = my_system.interconnects.add_sync_module(SSM_1, options)
my_system.interconnects.add_sync_module(SSM_2, options)
#
# Define chassis
my_system.add_chassis(1)
my_system.add_chassis(2)
#
# Select the leader SSM as ref. clock source
clockSource = interconnects[0].clock_source
#
# Set the SSM clock source
systemDefinition.clocking.reference_source = clockSource
#
# Explicitly set the clock source to use the internal OCXO as the reference clock (this is the
default)
clockSource.set_mode(keysight_tse.ClockingReferenceMode.INTERNAL)
```

Configuring the SSM to explicitly use internal OCXO or external reference clock

The SSM leading the synchronization by default with its internal reference clock, can optionally be connected to an external reference clock. The external reference can come from, for example, a DUT or another source such as a PXIe frequency reference.

To use an external reference clock, you must:

- Connect the external reference source to the SSM's **SCLK / Ref in** port.
- In the TSE API you must set the SSM to synchronize to an external reference clock. To do this, set the mode to `EXTERNAL` and set the frequency in Hz.

To use the external reference, change the final line in the previous code snippet to:

```
# Set clock mode to EXTERNAL and set frequency to 10MHz
clockSource.set_mode(keysight_tse.ClockingReferenceMode.EXTERNAL, 10e6)
```

Configuring Clocking Schemes A.2 and C

(Single/multi-chassis with Keysight M904xA chassis, external Analog clocks and chassis or HPRCS as clock source)

Some instruments such as the analog ones in the Keysight M5xxx PXIe family derive their Sample Clocks from an Analog Clock source and perform best when configured to use an externally distributed Analog Clock. This section explains how to distribute the analog clocks to these and similar instruments.

The analog clock can be generated from either the M9546A HPRCS or from the M9046A chassis backplane board. The preferred choice for the analog clock source is the M9546A HPRCS inside a Keysight M9046A PXIe chassis because of its superior phase noise. The HPRCS can generate a sine wave with frequencies of 2.4, 4.8, 9.6, or 19.2 GHz. In this section we assume the analog clock source being set to generate 2.4 GHz, because this is the frequency required by the Keysight M5xxx PXIe family. The frequency can be configured at purchase by choosing the corresponding option for the Keysight M9046A PXIe chassis. For more information, see the *Keysight PXIe Chassis M9046A User Manual* available at [Keysight PXI chassis](#).

Analog clock configuration options

the following table lists the options available:


Source	Analog clock, locked to the reference clock	Performance
M9046A chassis with M9546A HPRCS	2.4, 4.8, 9.6, or 19.2 GHz	Best
M9046A chassis without M9546A HPRCS	2.4 GHz	Medium

M9046A Front Panel Clocking IO overview

The following diagrams show the M9046A chassis front panels and how they are connected in different configurations. The type and number of front panel connectors depend on the purchased hardware option for splitters and HPRCS: (-QS0, -QS1/3, -QS2). In the diagrams a frequency of 2.4 GHz is assumed to have been selected for the analog clock. The analog clock frequency is also chosen as hardware option at purchase time. More info in the *Keysight PXIe Chassis M9046A User Manual*, available at [Keysight PXI chassis](#).


M9046A -QS0 Chassis (no HPRCS)

The following diagram shows the front panel of an M9046A -QS0 chassis.




M9046A -QS2 Chassis with Analog clock splitters

The following diagram shows the front panel of a M9046A chassis with -QS2 option including the front panel analog clock splitters to ease the distribution of the analog clocks to all modules.



M9046A -QS1/3 Chassis with Analog clock splitters and HPRCS

The following diagram shows the front panel of an M9046A -QS1/3 chassis with analog clock splitters and Ref In, Cal In and Cal Out for the M9546A HPRCS.



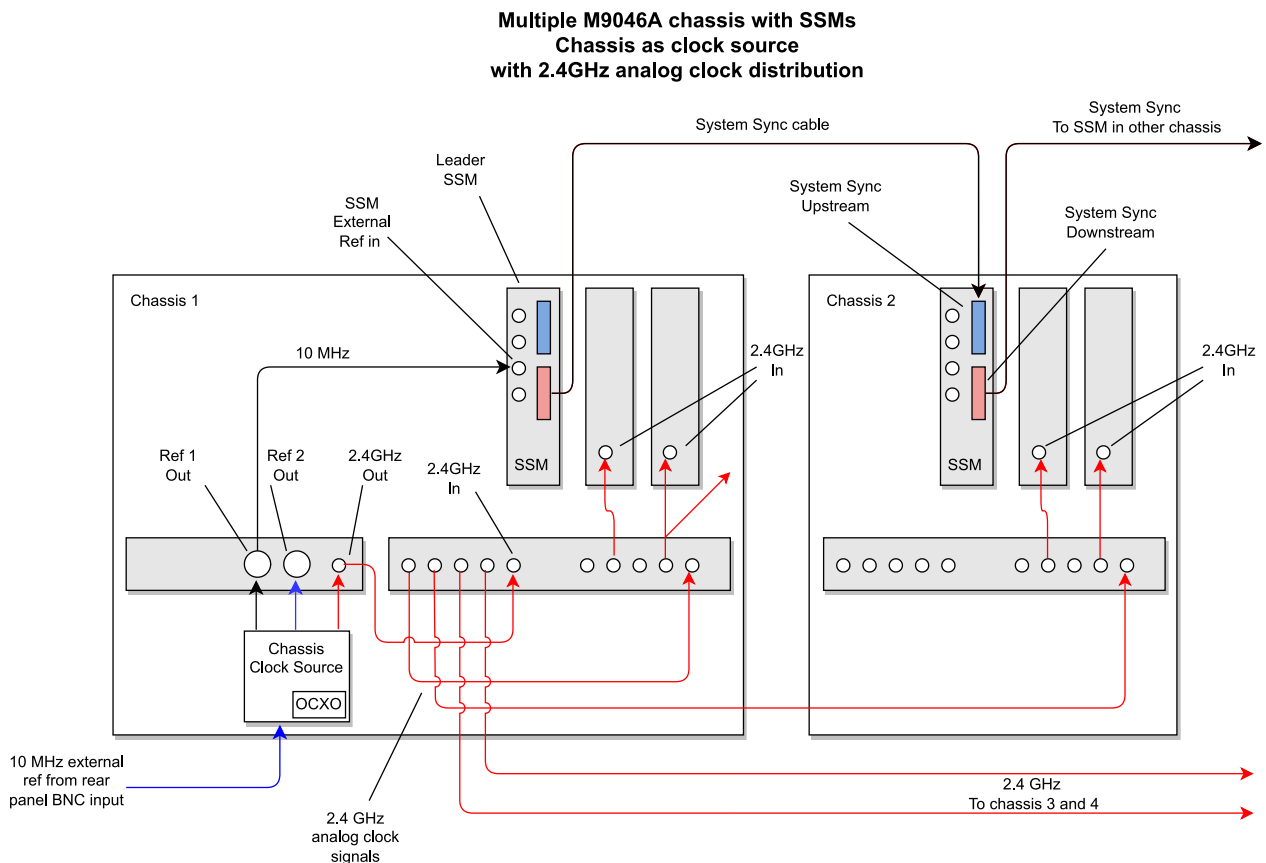
Using the M9046A analog clock source (Clocking Scheme A.2)

This configuration is for single or multiple chassis with SSMs. This configuration is compatible with PXIe Chassis model M9046A with hardware options -QS0, -QS1/3 or -QS2. The internal chassis clock is used as the clock source and this configuration must be defined in the TSE API. The chassis clock must be taken out from the **Ref 1 Out** port on the PXIe M9046 Chassis front panel and must be connected to the **SClk / Ref In** port of the PXIe SSM (see diagram below).

You can use the chassis as the reference clock source with its reference clock set to:

1. **Internal:** Use the chassis internal OCXO.
2. **External:** Using an external reference clock connected to the chassis rear panel 10MHz Ref BNC input.

The following diagram depicts the SSM using the chassis clock (indicated in red) as the clock source. The chassis external reference is indicated by the dotted blue arrow:



Configuring the M9046A as the system and analog clock source

To use the internal chassis clock, you must:

- Connect the Chassis **Ref 1 Out** output to the SSM's **SCLk / Ref In** located in the same chassis.
- In the TSE API you must instruct the SSM to use the chassis clock.

By default, and if it is not specified otherwise, the chassis clock circuitry uses its internal OCXO as the reference clock.

The following code shows how to configure a pair of chassis with SSMs using the chassis clock as the reference clock, options is set to an empty string:

```
# Create SystemDefinition object
ktHvi.SystemDefinition definition("Name")
#
# You must add all necessary follower SSMs depending on the number of chassis
syncModuleLeader = definition.interconnects.add_sync_module(SSM_1, options)
syncModuleFollower = definition.interconnects.add_sync_module(SSM_2, options)
#
# Add chassis
chassis1 = definition.add_chassis(1)
definition.add_chassis(2)
#
# Get the chassis clock
clock_source = chassis1.clock_source
#
# Set as reference
definition.clocking.reference_source = clockSource
#
# Enable the chassis analog clock
clock_output_2_4GHz = ref_chassis.clock_outputs["FP2.4GHzOut"]
clock_output_2_4GHz.set_enabled(True)
```

Configuring the M9046A to use the external reference clock

To use the chassis clock with an external reference clock, you must:

- Connect the external reference clock to the Chassis rear panel's **10 MHz Ref BNC input**.
- Connect the Chassis **Ref 1 Out** to the SSM **SCLk / Ref In** of the SSM in the same chassis.
- In the TSE API you must instruct the chassis to use the external reference clock and set the frequency in Hz.

The following code shows how to set the external reference:

```
# Set the reference mode to use an external reference
# Set clock mode to EXTERNAL and set frequency to 10MHz
clockSource.set_mode(key sight_tse.ClockingReferenceMode.EXTERNAL, 10e6)
```

Using the M9046A's M9546A High Performance Reference Clock Source (Clocking Scheme C)

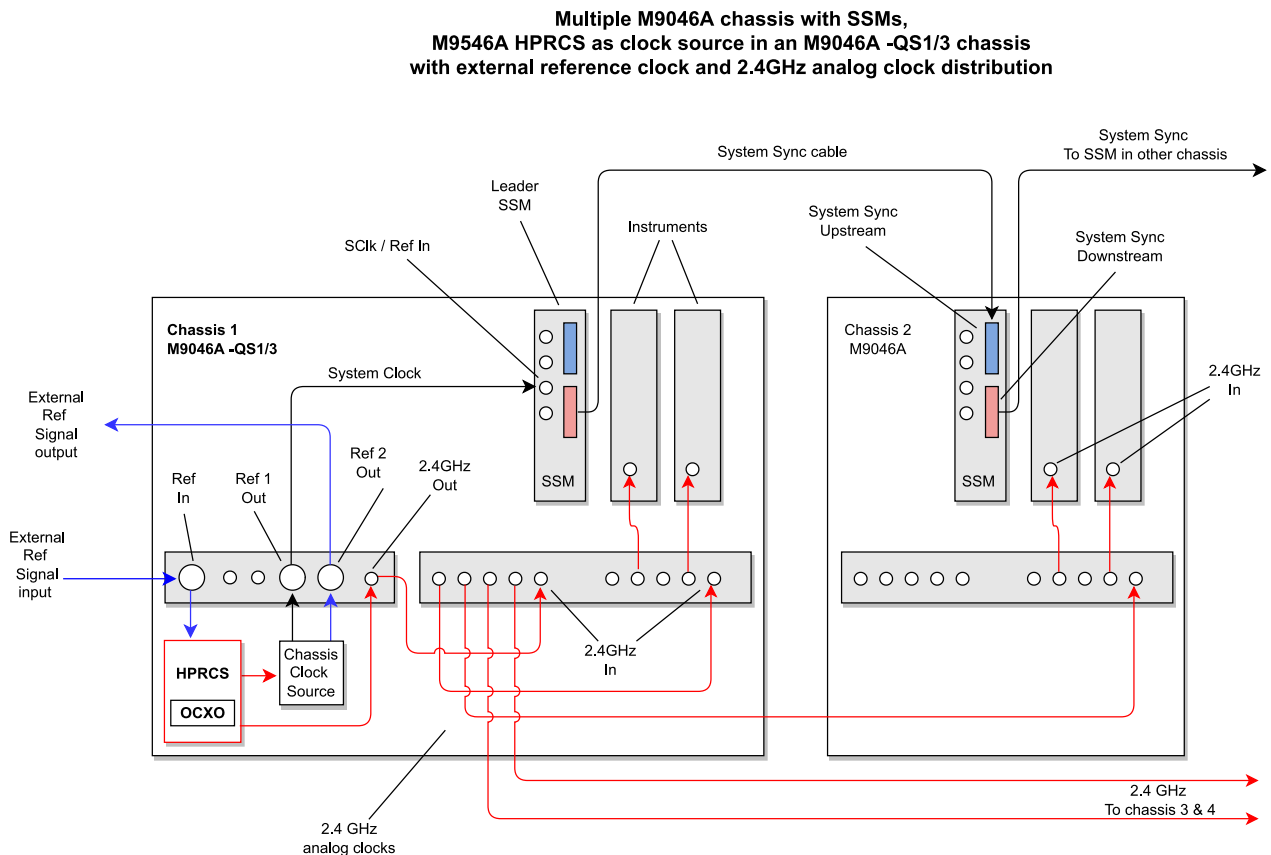
This configuration is for single or multiple chassis with SSMs. For this clocking scheme to be used, the first SSM must be in a Keysight M9046A chassis containing an M9546A High Performance Reference Clock Source (HPRCS).

The HPRCS is used as the clock source and this configuration must be specified in the TSE API.

We can use the HPRCS as a clock source with its reference clock set to:

- **Internal:** Use the HPRCS internal OCXO.
- **External:** Use an external reference clock connected to the chassis front panel **Ref In** input.

The following diagram shows the leader SSM using the M9546A HPRCS (indicated in red) as the clock source in chassis 1 M9046A -QS1/3. The HPRCS external reference is indicated by the blue arrow. The distribution of the 2.4 GHz analog clock to up to 4 chassis is also shown. The connection topology and cables used are critical to achieving the optimal channel skew drift performance. For information about how to connect the analog clock to more than 4 chassis, see *Keysight PXIe Chassis M9046A User Manual* available at [Keysight PXIe chassis](#).



Configuring the M9046A + HPRCS as the system and analog clock source

To use the HPRCS as clock source, you must:

- Connect the chassis **Ref 1 Out** output to the **SCLK / Ref In** of the SSM located in this chassis
- In the TSE API you must:
 - Add the M9046A -QS1/3 chassis with HPRCS to the SystemDefinition.
 - Set the HPRCS to be the clock source.

When no external reference clock for the HPRCS is specified, its internal OCXO is used.

The following code shows how to configure a pair of chassis with SSMs using the HPRCS as clock source, options is set to an empty string:

```
# Create SystemDefinition object
my_system = kthvi.SystemDefinition("MySystem")
#
# Define all necessary SSMs depending on the number of chassis
my_system.interconnects.add_sync_module(SSM_1, options)
my_system.interconnects.add_sync_module(SSM_2, options)
#
# Define chassis
hprcs_chassis = my_system.add_chassis(1)
my_system.add_chassis(2)
#
# Create HPRCS object
clockSource = hprcs_chassis.high_performance_clock_source
#
# Set the HPRCS as the reference clock
my_system.clocking.reference_source = clockSource
#
# Enable the chassis analog clock
clock_output_2_4GHz = ref_chassis.clock_outputs["FP2.4GHzOut"]
clock_output_2_4GHz.set_enabled(True)
```

Configuring the M9046A + HPRCS to use an external reference clock

To use the HPRCS with an external reference clock, you must:

- Connect the external reference clock to the chassis' front panel **Ref In** input.
- Connect the chassis' front panel **Ref 1 Out** output to the **SCLK / Ref In** input of the SSM located in this chassis.
- In the TSE API you must:
 - Add the M9046A -QS1/3 chassis to the SystemDefinition.
 - Set the HPRCS to be the clock source.

- Instruct the HPRCS to use an external reference clock and the desired frequency in Hz.

To use the external reference, set the reference clock mode in the previous code snippet to (defaults to internal):

```
# Set the reference clock mode. Set the HPRCS to use an external reference @10Mhz
clockSource.set_mode(keysight_tse.ClockingReferenceMode.EXTERNAL, 10e6 )
```

Enabling Chassis Clock Outputs

If you are using a clock output from a chassis you can enable it in the TSE API.

The chassis clock outputs are available in the chassis and you can access them by their name as follows:

```
# Get the Clock configuration for the Rear Panel 10MHz output port from the Chassis
ktHvi.SystemDefinition definition("Name")
#
chassis = definition.add_chassis(1)
#
clockOutputRp10Mhz = chassis.clock_outputs["RP10MHzOut"]
clockOutputRP10Mhz.set_enabled(true/false)
```

Some clock outputs support one single frequency and others support multiple frequencies. For the outputs supporting only one frequency, no frequency must be provided when enabling/disabling them. If the clock outputs do support multiple frequencies, you must specify what frequency (in Hz) you want to enable.

When you disable the clock, the frequency argument is ignored.

The following code shows some examples and error cases:

```
clockOutputRp10Mhz = chassis.clock_outputs["RP10MHzOut"]
clockOutputRP10Mhz.set_enabled(true) # Ok
#
clockOutputFpRef20Out = chassis.clock_outputs["FPref20Out"]
clockOutputFpRef20Out.set_enabled(true, 10e6) # Ok
```

Enabling the chassis Analog clock

If you are using an analog clock output from a chassis you must enable it.

The following code shows how to enable a 2.4GHz analog clock output from an M9046A chassis.

```
clock_output_2_4GHz = ref_chassis.clock_outputs["FP2.4GHzOut"]
clock_output_2_4GHz.set_enabled(True)
```

Configuring the Analog Clock Source in Instruments

For instruments that require an analog clock, you must set the source and frequency of the analog clock in your SystemDefinition.

You can set parameters for the analog clock:

- The source as internal or external.
- The frequencies of the sources, in Hz.

For external sources, the source selected depends on the analog clock frequencies that the instrument supports.

- If you indicate multiple frequencies, the first external frequency supported by the instrument is selected.
- If none of the external frequencies are supported, and the instrument has an internal clock, the internal clock is selected.
- If none of the external frequencies are supported, and the instrument does not have an internal clock, an error is generated.

The code is:

```
my_system.clocking.enable_external_analog_clocks(frequencies)
```

For example, if you are using a M9046A chassis with a M9546A HPRCS with the analog clock set to 2.4GHz (Clocking Scheme C), add the following line:

```
my_system.clocking.enable_external_analog_clocks([2400e6])
```

Instruments that support an external analog clock are set to use this clock. Instruments that do not support this external frequency are set to use an internal clock. If the instrument does not support the frequency and does not have an internal clock, an error is generated.

Selecting the best analog clock source for instruments

While it is often convenient for instruments to use their own internally generated analog clocks, the best jitter and drift performance is achieved by using a single common analog clock source generated within the Leader chassis (with or without the HPRCS) and distributing it using the chassis amplified power splitters in a balanced star configuration. This ensures that any low-frequency jitter skew drift is common across the system, minimizing the inter-channel jitter and drift.

In some cases with high channel count configurations, there may not be enough individual copies of the the Analog Clock available from a full balanced star distribution to connect to every instrument. In those cases, a single daisy-chain connection of the Analog Clock between instrument pairs can be used. Noting that the downstream instrument of the daisy-chained pair will have slightly higher skew drift than the non-daisy-chained instrument. Daisy-chained instruments shall have slightly higher

skew drift, so these instruments should be the ones in the system which have the lowest bandwidth. For example, in systems which employ the M5201A Downconverter and the M5200A Digitizer, which are typically used in pairs, it is best practice to route the Analog Clock to the downconverter first and then daisy-chain the downconverter's Analog Clock output to the digitizer's Analog Clock input. This is because the 2 GHz digitizer is less sensitive to the same amount of channel skew than the 16 GHz downconverter.

The Keysight MCX cables are made of a special material that minimizes their propagation delay change with temperature. Matching the total propagation delay from their common clock source to each instrument causes the propagation delay drifts of the clocks to cancel out between instruments.

The balanced-star configuration of external Analog clocks uses custom 4:1 amplified power dividers built into some chassis. These power dividers are designed specifically for minimizing phase noise, temperature drift, and to maintain the Analog clocks amplitude as it is divided many times. Substituting other power dividers to distribute the Analog clocks will degrade jitter and drift performance, so this is not recommended.

Small spurious oscillations can occur within the amplified power divider when any of the outputs are loaded with certain reflective loads. For this reason, terminating unused outputs with 50 ohm loads is recommended. It is only necessary to terminate unused outputs of power dividers that are currently being used to distribute the Analog clocks.

NOTE

If you are using a single M9046A chassis, and you are using an instrument as a synchronization source, ensure this instrument is in the middle segment.

Chapter 7: The TSE API

This chapter describes the TSE API. It describes the main classes and properties required to understand the key programming concepts you must understand to use the TSE API.

The TSE API is used in conjunction with the instrument HVI add-on API:

- The TSE API is the common API that apply to all instruments that support HVI.
- The Instrument HVI add-on API is an instrument-specific set of definition that enable to control instrument-specific capabilities from the TSE API.

NOTE

The TSE API functions alone are not sufficient to fully execute HVI sequences on an instrument. To successfully create an HVI, you must use both APIs.

This chapter contains the following sections:


- [TSE API Use Model](#)
- [TSE API Common Functionalities](#)
- [System Initialization](#)
- [The SystemDefinition Object](#)
- [The Sequencer Object](#)
- [The Hvi Object](#)

TSE API Use Model

PathWave Test Sync Executive has three main classes. You use them in this order:

1. SystemDefinition.
2. Sequencer.
3. Hvi.

Each of the stages in creating an HVI involves instantiating an object instance which is then passed to the next stage. The following diagram shows the stages:



NOTE

Once a step is completed and the SystemDefinition, Sequencer, or Hvi object is created, changes in the previous step instances will not apply on the already created next-step objects. That is:

- Any modification to the **SystemDefinition** instance once at the *Program HVI Sequences* or *Execute HVI* stage do not have any impact on the Sequencer and Hvi objects created/manipulated in these stages.
- Similarly, any modification to the **SystemDefinition** or **Sequencer** instances at the *Execute HVI* stage do not have any impact on the already generated Hvi Instance.

SystemDefinition

This is the first step of building an HVI. You use the `SystemDefinition` object to define the hardware components, configuration and the resources available in your system to be managed in the HVI. You do this by adding each of the resources to the relevant collections.

The SystemDefinition contains properties and methods to define and configure the following elements of an HVI:

- Chassis.
- Interconnects.
- Engines.
- FPGA-Sandboxes.
- HVI system clocks.
- Non-HVI core clocks.
- Sync resources.

Once you have added the resources, you can initialize the system. Ensure you initialize the system after adding all the resources.

NOTE

The default initialization that happens when the Sequencer object is created, initializes all the HVI Engines included in the SystemDefinition object. If you initialize the system using the `initialize()` method in the SystemDefinition instance, ensure that all the HVI Engines are added to the SystemDefinition instance before you call `initialize()`.

NOTE

PathWave Test Sync Executive 2023 release includes TSE Service. If your system uses TSE Service the set-up of a system is different.

The setup of a system that uses TSE Service handles starting up chassis, SSMs, and instruments. It also handles clocking and other infrastructure tasks.

A system definition or equivalent is still required, however this is done by a system administrator.

This arrangement simplifies sequence programming for end users since they are not required to write a full System Definition each time.

Users are still required to add instruments get their engines, and call `initialize()`. For more information, see [Chapter 9: TSE Service and Multi-Host support](#).

Sequencer

Once the `SystemDefinition` object is defined and configured, you define the HVI Sequences with a `Sequencer` object.

The `Sequencer` object is created from a `SystemDefinition` object. The `Sequencer` object contains all the hardware resources/configurations you defined in the `SystemDefinition`, these are available as read-only `view` collections. The view collections enable you to use the hardware resources for the sequence programming, but you cannot add or remove HW from them.

The `Sequencer` object exposes properties and methods for:

- Creating and manipulating the SyncSequence and Local Sequences. The `sync_sequence` property is the entry point for programming the real-time multi-instrument sequence.
- Compilation. Once you have programmed your sequence (Sync and Local sequences), you use the `compile(...)` method to obtain the Hvi object.

Hvi

The `Hvi` object is the actual HVI instance that you load to hardware and execute.

`Hvi` contains runtime versions of the objects that you set up with the `SystemDefinition` and `Sequencer` objects. You use the runtime objects for executing the Sequences on the hardware, but you cannot modify them.

`Hvi` contains the properties and methods to:

- Load the HVI instance to hardware.
- Run the HVI.
- Release the hardware resources used by the HVI.

Further Explanations

Detailed explanations of all the main classes and their functions are provided in the help file provided with the KS2201A PathWave Test Sync Executive installer. This is located at:

C:\Program Files\Keysight\PathWave Test Sync Executive 2023\api\python\Help

TSE API Common Functionalities

This section describes the functionalities that are common across the TSE API. It contains the following sections:

- [Capabilities](#)
- [Collections](#)
- [Error Management](#)
- [Time](#)
- [Component Versions](#)

Capabilities

The TSE API exposed many HVI capabilities or functionalities, including:

- Chassis/PXI backplane resource configuration.
- Multi-chassis/Box Interconnect configuration, for example, with *System Synchronization Modules* (SSMs).
- Access to HVI interfaces and resources in the FPGA-Sandbox.
- Real-time sequencing:
 - Synchronized Flow-Control Statements such as While loops.
 - Synchronized Multi-Sequence Block Statements that provide access to program sequencing operations in each individual instrument that participate in the global sequence.
 - HVI Instructions and operations. These include HVI-Native and Instrument-Specific HVI Instructions.
 - Local Flow-Control Statements such as While loops and If Statements.

Collections

Resources in HVI are grouped into Collections. Collections contain items of the same type, such as:

- Engines.
- Triggers.
- Actions.
- Events.
- Registers.
- FPGA-Sandboxes.

The concept of collections is fundamental in the TSE API use model because every component (Chassis, Engines, Clocks, Triggers etc.) used within the HVI is accessed through its corresponding collection.

For some of the collections, the elements must be registered explicitly by the user when defining an HVI instance, for example, you must add a Trigger to a Trigger collection. Once registered, you can then use the components inside HVI Sequences. There are other collections that are populated automatically by TSE as a result of other resources or components registration.

Collections are particularly useful because the member instances can be accessed by index or string. Collections are exposed at different levels in the TSE API hierarchy and those needed for the real-time sequencing description are exposed in the local and Sync Sequence objects.

The following tables lists the collections:

Name	Description
Collection	Chassis
Description	Contains the chassis added to the SystemDefinition
Access Type	integer: Chassis number identifier (for instance for PXI, as enumerated by PXI manager). Only valid for single host applications string: Chassis resource ID. This can be the VISA resource ID or the TSE resource ID
Exposed in	SystemDefinition::getChassis Hvi::getChassis
<hr/>	
Collection	SystemSyncUpstream/SystemSyncDownstream
Description	Contains the System Sync Upstream/Downstream connectors of a System Synchronization Module (SSM)
Access Type	integer: Index of the System Sync Upstream/Downstream connectors, as this is defined by the Instrument string: Name of the System Sync Upstream/Downstream connectors, as this is defined by the Instrument
Exposed in	SyncConnectivity::getSystemSyncUp SyncConnectivity::getSystemSyncDown
<hr/>	
Collection	Action
Description	Contains the actions that will be used with HVI by a specific engine
Access Type	integer: Position inside the collection. Matches the order of addition by the user. string: Name of the item, when this was added by the user
Exposed in	EngineDefinition::getActions EngineView::getActions EngineRuntime::getActions

Collections (continued)

Collection	Trigger
Description	Contains the triggers that will be used with HVI by a specific engine
Access Type	integer: Position inside the collection. Matches the order of addition by the user. string: Name of the item, when this was added by the user
Exposed in	EngineDefinition::getTriggers EngineView::getTriggers EngineRuntime::getTriggers
Collection	Event
Description	Contains the events that will be used with HVI by a specific engine
Access Type	integer: Position inside the collection. Matches the order of addition by the user. string: Name of the item, when this was added by the user
Exposed in	EngineDefinition::getEvents EngineView::getEvents EngineRuntime::getEvents
Collection	IfBranch
Description	Contains the IF-branches of an IF statement
Access Type	integer: Position inside the collection. Matches the order of addition by the user.
Exposed in	IfStatement::getElselfBranches

Collections (continued)

Name	Description
Collection	RegisterSharingOperation
Description	Contains the RegisterSharing Operations of a SyncRegisterSharing statement
Access Type	integer: Position inside the collection. Matches the order of addition by the user.
Exposed in	SyncRegisterSharingStatement::getRegisterSharingOperations
Collection	ClockOutput
Description	Contains the ClockOutputs supported by a specific chassis
Access Type	string: Name of the clock output on the front panel label
Exposed in	Chassis::getClockOutputs
Collection	FpgaSandbox
Description	Contains the FpgaSandboxes supported by a specific instrument
Access Type	integer: Position inside the collection. Order is defined by the instrument. string: Name of the sandbox as this is defined by the instrument
Exposed in	EngineDefinition::getFpgaSandboxes EngineView::getFpgaSandboxes EngineRuntime::getFpgaSandboxes

Collections (continued)

Name	Description
Collection	FpgaMemoryMap
Description	Contains the FpgaMemoryMaps supported by a specific FPGA sandbox
Access Type	integer: Position inside the collection. Matches the order defined in the k7z file. string: Name of the MemoryMap, as this is defined in the PathWave FPGA project.
Exposed in	FpgaSandboxDefinition::getFpgaMemoryMaps FpgaSandboxView::getFpgaMemoryMaps FpgaSandboxRuntime::getFpgaMemoryMaps
Collection	FdsPort
Description	Contains the FDS ports supported by a specific FPGA sandbox
Access Type	integer: Position inside the collection. Matches the order defined in the k7z file. string: Name of the FDS port, as this is defined by the instrument.
Exposed in	FpgaSandboxDefinition::getFdsPorts FpgaSandboxView::getFdsPorts FpgaSandboxRuntime::getFdsPorts
Collection	Scope
Description	Contains the Scopes per engine added to the SystemDefinition
Access Type	integer: Position inside the collection. Matches the order of engines addition by the user. string: Name of the engine for which this scope applies
Exposed in	SyncSequence::getScopes

Collections (continued)

Name	Description
Collection	Sequence
Description	Contains the Sequences per engine added to the SystemDefinition
Access Type	integer: Position inside the collection. Matches the order of engines addition by the user. string: Name of the engine for which this scope applies
Exposed in	SyncMultiSequenceBlockStatement::getSequences
Collection	Engine
Description	Contains the Engine added to the SystemDefinition
Access Type	integer: Position inside the collection. Matches the order of engines addition by the user. string: Name of the item, when this was added by the user
Exposed in	SystemDefinition::getEngines SyncSequence::getEngines
Collection	SwRecording
Description	Contains the SwRecording added to a specific engine
Access Type	integer: Position inside the collection. Matches the order of addition string: Name of the item, when this was added by the user
Exposed in	EngineDefinition::getRecordings EngineView::getRecordings EngineRuntime::getRecordings

Collections (continued)

Name	Description
Collection	Register
Description	Contains the registers belonging to a specific Scope
Access Type	string: Name of the item, when this was added by the user
Exposed in	Scope::getRegisters
Collection	SyncRegister
Description	Contains the registers belonging to a specific SyncScope
Access Type	string: Name of the item, when this was added by the user
Exposed in	SyncScope::getRegisters
Collection	Interconnect
Description	Contains the Sync Modules added to the SystemDefinition
Access Type	integer: Position inside the collection. Matches the order of addition string: Sync Module resource ID. This can be the VISA resource ID or the TSE resource ID
Exposed in	SystemDefinition::getInterconnects Hvi::getInterconnects

Collections (continued)

Name	Description
Collection	Statement
Description	Contains the statements belonging to a specific Sequence
Access Type	integer: Position inside the collection. Matches the order in the sequence string: Name of the statement, when this was added by the user
Exposed in	Sequence::getStatements
Collection	SyncStatement
Description	Contains the SyncStatements belonging to a specific SyncSequence
Access Type	integer: Position inside the collection. Matches the order in the SyncSequence string: Name of the SyncStatement, when this was added by the user
Exposed in	SyncSequence::getSyncStatements
Collection	Routing
Description	Contains the routing added to the SystemTriggering
Access Type	integer: Position inside the collection. Matches the order of addition string: Name of the item, when this was added by the user
Exposed in	SystemTriggering::getRoutings

Using Collections

Collections have additional access properties beyond those of vectors or lists.

To use the resources in the SystemDefinition you must add them to their respective collections.

For example, If you want to use a specific Engine you must add it to the Engine collection in the SystemDefinition. Once you have added the Engine, you can query it and configure it. To use the resources within each Engine you must add them to their respective collections in that Engine, you can then query them and configure them as required.

Adding a new collection item

You add new collection items by calling the `add()` method.

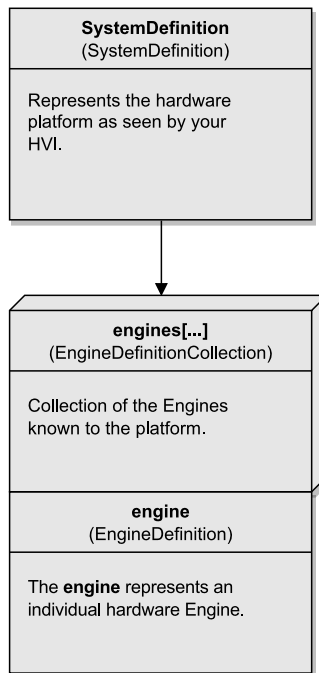
The following code adds a new Engine to the Engines collection property of the SystemDefinition and returns the new Engine. The Engine can be queried with the name `MyEngine_1` :

```
my_engine = my_system_def.engines.add(module.hvi.engines.main_engine, "MyEngine_1")
```

NOTE

You define the name when you add the item to the collection, each name you specify in a collection must be unique in that collection.

The following diagram shows the relationship between the SystemDefinition, the Engine collection and an individual Engine:



Random access by string or by numerical index

You access collection items with the `[]` operator. You can index items with their name, or by a number that indicates their location inside the collection.

For example, you can query any Engine already added to the collection through the `engines` property in the SystemDefinition instance.

The following code returns an EngineDefinition object named `My_Engine_1`:

```
my_system_defintion.engines["My_Engine_1"]
```

To find the number of items in a collection, use either `count` or the built-in `len()` function. For example, the following code returns the number of Engines the instrument has:

```
len(instrument.engines)
```

Managing objects in a collection

The collection is a grouping of elements, but it has no knowledge of the parameters or attributes of each element.

Configuration of the specific element within a collection are managed through the element instances, not the collection itself. For instance, you manage an Engine with the Engine instance obtained from the EngineCollection.

Error Management

Error handling in the TSE API is based on exceptions. If an error occurs during execution, the execution is stopped, and a message is returned that includes an error code and a relevant error message. Error management is done through an Error object.

Time

This section describes the API related to the Time inside HVI.

For time you use a `Duration` object, located in the namespace `time`, to represent a time interval.

You can create a `Duration` object in one of these ways:

- By only providing a single floating point value. In this case, the value is treated as time in nanoseconds.
- By providing a floating point value and the unit of time you want the value to represent.

The signature is:

```
time.Duration(double valueInNanoseconds);  
time.Duration(double value, Time::Unit unit);
```

The units of the duration are defined by `Unit`. The supported units are the following:

- Seconds.
- Milliseconds.
- Microseconds.
- Nanoseconds.
- Picoseconds.

`Minimum` is a specialization of the `Duration` implementation to specify that minimum time instead of a fixed specific time and allow TSE to decide the actual time during sequence compilation or at a later point.

The signature is:

```
time.Minimum();
```

The following is an example of usage:

```
from keysight_tse import time  
a_duration = time.Duration(35.0)  
assert a_duration.type == time.Type.FIXED_DURATION  
assert a_duration.value == 35.0  
assert a_duration.unit == time.Unit.NANOSECONDS  
another_duration = time.Duration(35.78, time.Unit.MICROSECONDS)  
assert another_duration.type == time.Type.FIXED_DURATION
```

```
assert another_duration.value == 35.78
assert another_duration.unit == time.Unit.MICROSECONDS
a_minimum_duration = time.Minimum()
assert a_minimum_duration.type == time.Type.MINIMUM_DURATION
assert a_minimum_duration.value == 0.0
assert a_minimum_duration.unit == time.Unit.NANOSECONDS
```


Component Versions

When using TSE in addition to the PathWave *Test Sync Executive* (TSE) Software Version, you should be aware of the following component versions:

- HVI Core (software) Version.
- HVI Engine Software Version.
- HVI Engine Firmware Version.

	Python	Object	Description
HVI Core (software) Version	<code>keysight_tse.SystemDefinition.hvi_core_version</code>	<code>SystemDefinition</code>	The version of the HVI core component that gets installed by PathWave Test Sync Executive software to deliver the TSE API.
HVI Engine Software Version	<code>Engine.software_version</code> <code>EngineView.software_version</code> <code>EngineRuntime.software_version</code>	<code>EngineDefinition</code> <code>EngineRuntime</code> <code>EngineView</code>	The version of the HVI software component used by the instrument associated with this Engine.
HVI Engine Firmware Version	<code>Engine.firmware_version</code> <code>EngineView.firmware_version</code> <code>EngineRuntime.firmware_version</code>	<code>EngineDefinition</code> <code>EngineRuntime</code> <code>EngineView</code>	The version of the HVI Engine FPGA IP that is programmed into the FPGA of the instrument associated with this Engine.

`hvi_core_version`

This version is exposed as a property of the SystemDefinition.

This is the version of the HVI core component that gets installed by PathWave Test Sync Executive software and also instruments implementing HVI technology. The HVI Core is a shared component in the system, if products with different HVI Core versions are installed in a given host, only newer HVI Core version will be installed and consequently regardless of the order of installation the latest of all HVI Core included in the different products will be available.

software_version

This is a property of each Engine in each product. It can be queried from the [EngineDefintion](#), [EngineRuntime](#) and [EngineView](#) objects.

This is the version of the HVI Core used to build the instrument driver and it is the HVI Core version included in the instrument driver installation. This version does not need to be the same as the HVI Core installed in the system to work, but instruments shipped with older HVI Core versions may not support all the latest features. This software version depends on the version of the instrument drivers that can be queried using the instrument driver API or the instrument Software Front Panel software.

firmware_version

This is a property of each Engine in each product. It can be queried from the [EngineDefintion](#), [EngineRuntime](#) and [EngineView](#) objects.

This is the version of the HVI Engine IP included in the instrument FW. The HVI Engine IP version is associated with the instrument firmware version, different instrument firmwares may include different HVI Engine versions. You can upgrade/downgrade the instrument firmware using the instrument Software Front Panel software. The HVI Engine IP version determine the features that are available for a given instrument.

Major, minor and revisions

The version has the sub-versions: [major](#), [minor](#) and [revision](#).

The following code shows an example of how to get the versions:

```
logging.info("HVI Core : {}".format(sys_def.hvi_core_version.to_string()))
for engine in sys_def.engines:
    logging.info("Firmware Version : {}.{}.{}".format(engine.firmware_version.major,
engine.firmware_version.minor, engine.firmware_version.revision))
    logging.info("Software Version : {}.{}.{}".format(engine.software_version.major,
engine.software_version.minor, engine.software_version.revision))
```

System Initialization

This section describes System initialization, it contains the following sections:

- [The SystemDefinition initialize\(\)](#)
- [Initialization during Sequencer Creation](#)
- [Initialization during Load To Hardware](#)
- [Example of System Initialization](#)

In HVI technology, System Initialization is a process that includes the configuration and alignment of the different systems clocks, including the clocking for each instrument specified as part of the HVI SystemDefinition. This procedure includes the generation and alignment of the instrument internal *Sync* and *SyncBase* signals described in [Chapter 10: HVI Time Management and Latency](#). Initialization also includes the alignment of clocks controlling the *Fast Data Sharing* (FDS) functionality for cases where your HVI definition includes instruments that support this functionality, such as instruments from the Keysight PXIe M5xxx instrument family.

A complete system initialization can be a complex procedure that can take some time. It must be performed when the system is first assembled and powered on, but you are not required to perform this level of initialization every time. However, a basic level of initialization is required every time to ensure that the system is synchronized and the clocks are in alignment.

A number of initialization options are provided to ensure correct setup and minimize initialization time when you want to run operations. You can use the basic options after the system has been fully initialized for the first time, however if you change the hardware setup (instruments, cables, etc.) or clocks (reference clock source, clock connection cables, etc.) you must perform a complete initialization again.

The `SystemDefinition` object includes an `initialize()` method that initializes the hardware included in the SystemDefinition, and performs synchronization and clock alignment. There are 3 cases where system initialization and clock alignment can occur:

- Manually calling `initialize()` in the SystemDefinition.
- When the Sequencer object is created.
- When calling Load to Hardware.

NOTE

The initialization process requires access and control of all of the hardware resources, so it is important that these resources are not already in use by another application or HVI instance already loaded to hardware. An exception is thrown if any of the hardware resources are already in use.

The `SystemDefinition.initialize()`

The `systemDefinition.initialize(...)` method enables you to explicitly trigger a system initialization and alignment. You might want to explicitly control when the initialization process is executed because the initialization process can take some considerable time depending on the parameters and system state.

You can also specify specific alignment modes when calling the `initialize()` method:

Mode	Description
<i>Default</i>	Calling <code>initialize()</code> without parameters performs the default, or minimal-possible initialization. This is the mode intended to be used in normal system operation.
<code>keysight_tse.AlignmentModes.Full</code>	Forces a full system, complete system initialization and alignment.
<code>keysight_tse.AlignmentModes.ResetCalibration</code>	Performs system initialization and alignment resetting and regenerating the stored calibration data.
<code>keysight_tse.AlignmentModes.PreCalibration</code>	Performs system initialization and alignment, ignoring calibration data. This mode is intended for system warm-up or other instances when the use of precise alignment calibration data is not available yet or not required.
<code>keysight_tse.AlignmentModes.Disable_clock_monitoring</code>	Disables the Clock Monitoring on SSMs. By default, Clock Monitoring is disabled unless at least one instrument supports it.
<code>keysight_tse.AlignmentModes.Force_clock_monitoring</code>	Forces the Clock Monitoring on SSMs. By default, Clock Monitoring is disabled unless at least one instrument supports it.

You can combine modes using a bitwise-OR operator, for example:

```
systemDefinition.initialize(AlignmentModes::Full | AlignmentModes::PreCalibration);
```

The API call above will combine the *Full* and *PreCalibration* alignment modes as follows:

The software forces a complete system initialization and clock alignment (*Full* mode). While doing this, if the software detects that any of the instruments included in the HVI SystemDefinition object requires calibration data and this data is missing, it shall allow the system initialization procedure to continue and finalize, instead of throwing an error (*PreCalibration* mode).

```
systemDefinition.initialize(AlignmentModes::PreCalibration);
```

The API call above will perform a default system initialization and clock alignment. While doing this, if the software detects that any of the instruments included in the HVI SystemDefinition object requires calibration data and this is missing, it will allow to the system initialization procedure to continue and finalize, instead of throwing an error.

NOTE

When using the default mode, in order to minimize the initialization time, PathWave Test Sync Executive relies on storing the initialization state of each instrument to decide what initialization steps are required. If hardware or cabling changes have been made in your system, you must make sure the correct initialization modes, *Full* / *PreCalibration* / *ResetCalibration*, is executed as required by your setup. For information about the initialization requirements for cabling or hardware changes, see your instrument documentation.

Default Initialization

This is the mode intended to be used in normal system operation. When calling `systemDefinition.initialize()` without parameters, it performs the default, or minimal-possible initialization. The default initialization tries to minimize the necessary operations to obtain the fastest initialization and synchronization time. The default initialization is automatically executed when a Sequencer object is created from a SystemDefinition, and also when the `LoadToHw()` method is called on the HVI instance.

If you have a system that has been power-cycled, the first call to `systemDefinition.initialize()` with no arguments will actually execute a full initialization and complete clock alignment because the system has not been aligned yet. The full initialization procedure can take several minutes depending on the size and structure of your system, including the number of chassis and instruments, see the *Full Mode* description. Subsequent calls to `systemDefinition.initialize()` after the first call, are very fast.

Some instruments require stored calibration data to initialize correctly, in these cases this calibration data must be available for all instruments in the SystemDefinition for the default initialization to work. If the calibration data is not available, an error is generated. See *ResetCalibration* and *PreCalibration* modes to understand the details on how to manage calibration data. For example, the Keysight PXIe M3xxx family does not require any calibration data, whereas in the Keysight PXIe M5xxx family, the M5300 RF AWG and M5201 Frequency Down-Converter do require calibration data. For information about the calibration requirements, see your instrument documentation.

Full Mode

You can force a full clock alignment by calling: `systemDefinition.initialize(key sight_tse.AlignmentModes.Full)`. The full initialization procedure can take several minutes depending on the size and structure of your system, including the number of chassis and instruments.

NOTE

This mode in general is only needed when some cabling change is done on the system without shutting it down or in cases when we want to force from software a complete reset of the initialization or alignment to recover from an undesired state.

ResetCalibration Mode

You can force the update of alignment calibration data by calling: `systemDefinition.initialize(key sight_tse.AlignmentModes.ResetCalibration)`

This mode erases any existing system calibration data and forces the instruments to re-calculate and store new calibration data. Use this mode if there is no system calibration data available for the current hardware configuration and operating temperature, or if the existing calibration requires recalculation.

This operation must be performed when one or more of the following occur:

- A system setup containing any instrument that requires calibration data is used for the 1st time.
- The hardware in the system has been changed. This includes adding any instrument that requires calibration data, changing the cable connections between System Sync Modules, or changing any of the cables connections (clock, in/out, etc.) of any instrument requiring calibration data.
- The clock configuration has been changed. This includes changing the reference clock source, or any of the cable connections from it to any of the instruments that require calibration data.

NOTE

Ensure you perform a full initialization with ResetCalibration when your system is fully warmed up. For information about the warm-up and temperature stabilization requirements for best performance, see your instrument documentation.

PreCalibration Mode

Calling `systemDefinition.initialize(key sight_tse.AlignmentModes.PreCalibration)` configures the system without the need of having calibration data already stored for the specific temperature condition. Typical use of this mode is for system warm-up to prepare the system for a default initialization or a *ResetCalibration* initialization. Other instances include when calibration data is not available but precise calibration is not required.

NOTE

When using instruments that require precise calibration data, always use this mode to warm-up the system before use. For information about the warm-up and temperature stabilization requirements for best performance, see your instrument documentation.

Initialization during Sequencer Creation

Behind the code used to create a Sequencer object, a system initialization is implicitly executed only if the SystemDefinition contains instruments that include FDS functionality. If the instruments do not include FDS, these operations are skipped. For example, the `initialize()` of a SystemDefinition that only includes M3xxx instruments is not performed at the Sequencer creation because these instruments do not support FDS. If an M5xxx instrument is included in a SystemDefinition, system initialization and clock alignment is performed when creating the Sequencer, unless the `initialize()` had already been explicitly called.

At the Sequencer creation, this only performs a minimal update to the initialization and clock alignment. If the infrastructure, in particular the System Synchronization Modules (SSMs), are not properly initialized and aligned, a full system initialization is triggered, equivalent to calling `SystemDefinition::initialize(AlignmentModes::Full)`. Otherwise, if all of the SSMs are already aligned, it checks if there are any instruments misaligned and force initializes them, skipping any that are already aligned.

NOTE

The initialization process requires exclusive access to the hardware resources involved, so it cannot be executed while another TSE instance is running and has that hardware locked.

Initialization during Load To Hardware

The Load To Hardware operation also contains an implicit default system initialization. In Load to Hardware such system initialization is always performed and it is equivalent to the explicit call `sys_def.initialize()` executed without parameters. This only performs a minimal update to the initialization and clock alignment.

Example of System Initialization

To use the TSE API to initialize and run real-time operation in your system, there are two main procedures that you must follow:

1. **System Warm-up and Calibration**
2. **Normal Operation**

There are also a number of use cases that are variations on these main procedures. The following text describes these procedures along with the use case variations.

NOTE

The initialization process requires access and control of all of the hardware resources, so it is important that these resources are not already in use by another application or HVI instance already loaded to hardware. An exception is thrown if any of the hardware resources are already in use.

System Warm-up and Calibration

The system warm-up must be performed every time the system is turned on or the hardware configuration is changed. This is to enable all of the components to reach a stable and repeatable operating temperature. Once the system is warmed-up, the system can be initialized using the stored System Calibration data.

The System Calibration must be performed in these cases:

1. The very first time that the system is put together and powered-on.
2. When relevant hardware changes are made that require a new system calibration. These hardware changes include:
 - a. Adding/removing a chassis in your SystemDefinition object.
 - b. Adding/removing any instrument that requires clock alignment calibration data, such as an M5300A module or M5201, or changes the operating temperature of the system.
 - c. Changing the cable connections between System Synchronization Modules, even replacing a cable with a similar one with a different serial number.
 - d. Changing any of the external System Clock or Analog Clock cable connections, even replacing a cable with a similar one with a different serial number.
 - e. Making any change to the clock configuration, even if it is only from the TSE API. This is because this triggers the usage of different clock sources or signal paths.
3. Other situations where the system calibration should be updated.
4. On rare occasions, a component in the system can move into an invalid state and a reset of the calibration might be required. For more information, see *System Troubleshooting* in the **System Setup Guide**.

NOTE

Warning : Resetting the system calibration shall in turn require you to recalculate the User Calibration for some instruments. Observe extreme caution when doing this to avoid costly time-consuming recalibration.

Procedure steps:**Procedure steps:**

1. **Power-on the system**
 - a. Power-on all of the chassis. After this is complete, if you are using an external chassis controller, power it on.
2. **Connect to all the instruments**
 - a. For example: `instrument = ktm5300.KtM5300x(resource_id, query, reset, options)`
3. **Create a SystemDefinition** using the TSE API and the instrument drivers:
 - a. Create a SystemDefinition object that we refer to here as *my_system*. Use the *my_system* object to define all the hardware resources in your system: chassis, SSMS, instruments, clocking configuration, reference clock source, etc.
For example: `my_system.chassis.add(1), my_system.clocking.reference_source = chassis.clock_source`
 - b. Add the HVI Engines of each instrument to the SystemDefinition object.
For example: `my_system.engines.add(instrument.hvi.engines.main_engine, "MyEngine")`
4. **System Initialization for Warm-Up**
 - a. Execute `my_system .initialize(keysight_tse.AlignmentModes.FULL | keysight_tse.AlignmentModes.PRE_CALIBRATION)`. The `PRE_CALIBRATION` flag indicates there is no need to apply any previously stored system calibration values because the system is warming-up. This enables the system to execute code without calibration related errors. After this step, instruments may present channel skew errors which are compensated by the next steps.
5. **Wait for System Warm-Up**
 - a. Wait for the required warm-up time, this can range from a few minutes to about 30 minutes. The actual time typically depends on the type and number of instruments in the system, clocking configuration, etc.
 - b. For detailed warm-up time information, see your instrument documentation, for example: *M5300 RF AWG User's Manual* .
6. **System initialization to perform System Calibration**
 - a. Using the SystemDefintion created in step 3, run `my_system.initialize(keysight_tse.AlignmentModes.FULL | keysight_tse.AlignmentModes.RESET_CALIBRATION)` to generate internal system calibration data. At first system turn-on, no previous calibration data is expected to be available.
7. **Calculate User Calibration or channel deskew** (Optional)
 - a. This operation is optional and consists of correcting analog channel skews introduced by cable and signal path delays. Note that in some instruments, the User Calibration must be re-

calculated when a System Calibration is executed. For information about how to do this, see your instrument documentation.

8. Ready for Normal Operation

Use Cases:

Use Case Scenario	Description
First system start-up and calibration	<p>The very first time that the system is put together and powered-on, you must execute a full warm-up and calibration procedure to achieve the best system performance and repeatability:</p> <ul style="list-style-type: none"> • Execute all steps #1 to #7 above.
System start-up using existing calibration	<p>If the system has already been calibrated for the current hardware configuration, then, to reuse the existing calibration to configure the system, wait for the system temperature to stabilize then apply the existing calibration:</p> <ul style="list-style-type: none"> • Execute steps #1 to #5 above. • Skip steps #6 and #7 <i>System initialization to perform System Calibration</i> and <i>Calculate user calibration or channel deskew</i>, and run <code>my_system.initialize(key sight_tse.AlignmentModes.FULL)</code>.
Simplified uncalibrated system start-up	<p>If you want to use the system for test development, or you can tolerate analog channel drift of up to 50ps across reboots/power-cycles:</p> <ul style="list-style-type: none"> • Execute steps #1 to #4 above. • Skip steps #5 to #7 <i>Wait for System Warm-Up</i>, <i>System initialization to perform System Calibration</i> and <i>Calculate user calibration or channel deskew</i>.

NOTE

System hot boot-up: If the system is already warmed-up to the calibration operating conditions, for example after a system restart, you can skip the steps #4 and #5 *System Initialization for Warm Up* and *Wait for System Warm-Up*.

Normal Operation

Once the system is warmed-up and the system calibration has been done, users can use the the TSE API to execute real-time operations:

NOTE

Note that if it is the first system start-up or you have introduced any of the HW changes that require new System/User Calibration you must execute the *First system start-up and calibration* use case described in the *System Warm-up and Calibration* procedure.

Procedure steps:

1. **Connect to all the instruments**, if not already connected.
 - a. For example: `instrument = ktm5300.KtM5300x(resource_id, query, reset, options)`
2. **Apply user calibration to instruments**, You only need to do this if it is required, the user calibration data is available, and it has not been applied already.
 - a. The user calibration is calculated during the *System Warm-up and Calibration* process. For information about how to apply existing calibration, see your instrument documentation, for example: *M5300 RF AWG User's Manual* .
3. **Create a SystemDefinition object**, or reuse an existing one.
4. **Initialize the SystemDefinition object (Optional)**
 - a. Run `my_system.initialize()`. This call executes the minimal or default initialization, provided a Full Initialization has been executed already as described in the *System Warm-up and Calibration* procedure. If the full initialization has not been executed, this step requires calibration data. If the calibration data is not available this operation will fail. To run the system initialization without calibration you can specify the `PRE_CALIBRATION` flag: `my_system.initialize(key sight_tse.AlignmentModes.PRE_CALIBRATION)`
 - b. Note that you can skip the call to `my_system.i nitialize()` because the minimal or default initialization happens implicitly in steps #5 and #7 described below.
5. **Create a Sequencer object**
 - a. For example: `sequencer = keysight_tse.Sequencer("MySequencer", my_system)`
 - b. Note that the sequencer creation operation implicitly executes a default initialization, this is equivalent to calling `SystemDefinition:Initialize()`.
6. **Create an HVI object**
 - a. For example: `hvi = sequencer.compile()`
 - b. The `Hvi` object is created by compiling the Sequencer object after all the HVI Sequences have been programmed._

7. Load HVI to HW

- a. For example: `hvi.load_to_hw()`
- b. Note that the `load_to_hw()` operation implicitly executes a default initialization, this is equivalent to calling `SystemDefinition:Initialize()`.

8. Run HVI

- a. For example: `hvi.run(hvi.no_timeout)`

9. Release HW

- a. For example: `hvi.release_hw()`

NOTE

Forcing a full initialization. You can optionally force a full initialization. Forcing the full initialization can be useful to unblock a system if it is in a bad state, when some temporary hardware changes in the system are done such as reconnecting cabling using the same cables, or in general when it is useful to ensure the system is fully initialized to discard any previous state. To force the full initialization run:

1. `my_system.initialize(key sight_tse.AlignmentModes.FULL)`.
2. Or if you are using the system without calibration, add the `PRE_CALIBRATION` flag:
`my_system.initialize(key sight_tse.AlignmentModes.FULL | key sight_tse.AlignmentModes.PRE_CALIBRATION)`

NOTE

User Calibration not required or already applied: If user calibration is not required or has already been applied to the instruments, you can skip step #2 ***Apply user calibration to instruments***. For more information on how to handle User Calibration in instruments, see your instrument documentation.

The SystemDefinition Object

This section describes the `SystemDefinition` object, it contains the following sections:

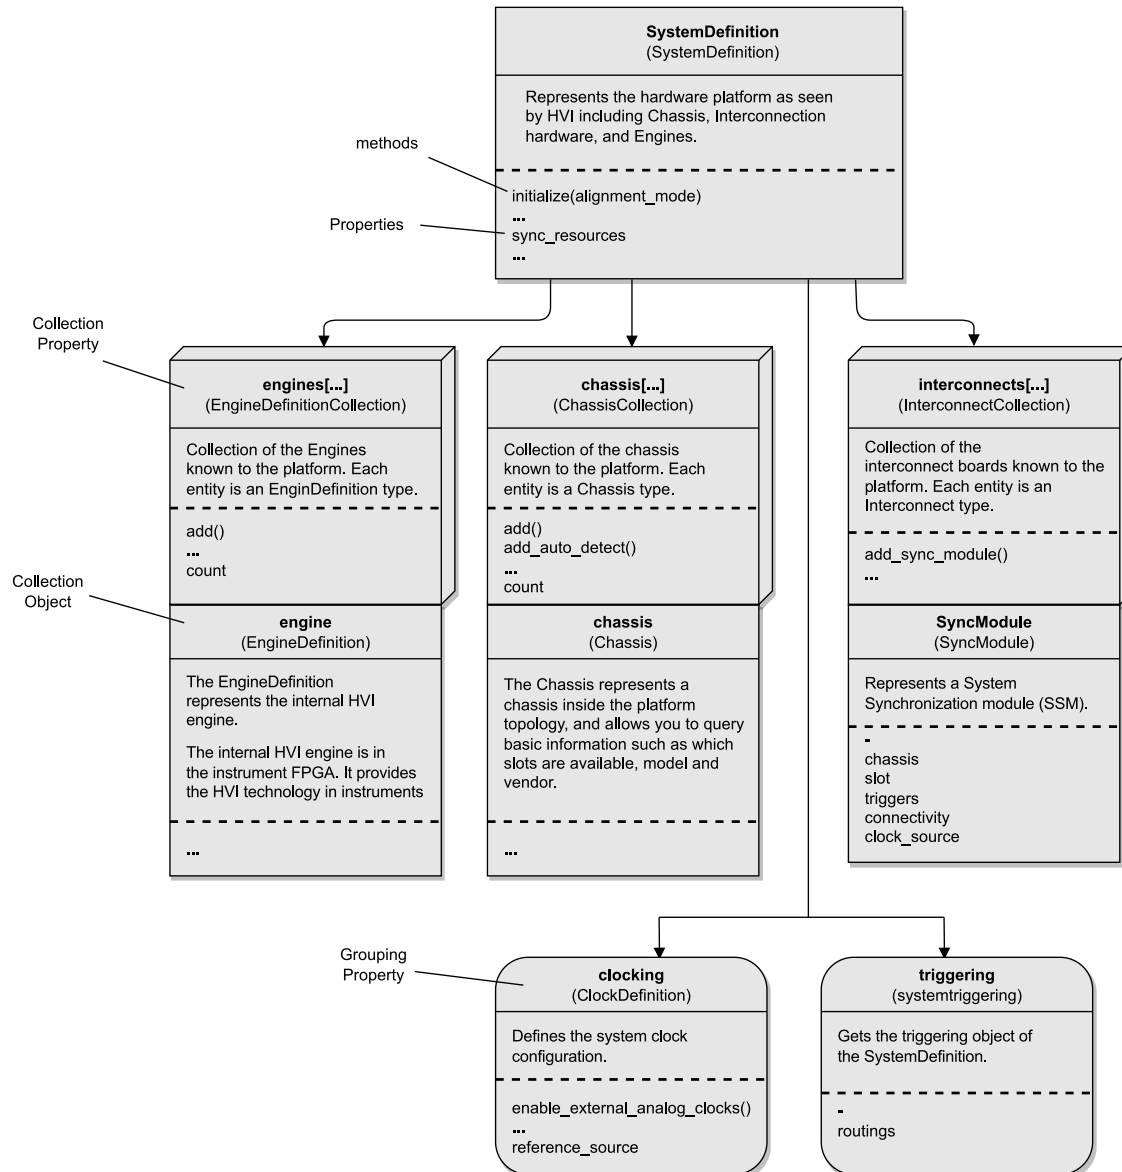
- [HVI Engines and their Resources](#)
- [Chassis and Interconnects \(SyncModule\)](#)
- [Synchronization Resources](#)
- [Synchronization Signals and Sync Modes](#)
- [Non-HVI Clocks](#)
- [System Initialization in the SystemDefinition Object](#)
- [User-Defined Trigger Routing](#)
- [System Clocking Configuration](#)

You must tell your HVI about the components in your system, how they are connected, and the features available. To enable this, the `SystemDefinition` object contains a set of properties (collections) that enable you to define the hardware configuration that your HVI is going to use. When you define your system, you must add the different components you intend to use to the relevant collections (properties) in your `SystemDefinition`. There are collections for Engines, Chassis, and Interconnects.

You use a `SystemDefinition` object to configure the physical hardware resources available to the HVI. You create a `SystemDefinition` with the following code:

```
my_system_definition = keysight_tse.SystemDefinition("My Setup")
```

The following diagram shows the `SystemDefinition` with some of its methods and properties. To simplify the diagram, it does not show all methods or properties:



The properties shown in the diagram above can be classified in three kinds:

Collection Properties

engines[...]

Enables you to add and access engines.

chassis[...]

Enables you to add and access chassis.

interconnects[...]

Enables you to add and access SyncModules.

Grouping Properties (interfaces that group methods and properties)

clocking

A group of methods and properties used for representing clocks.

triggering

A group of properties used for representing trigger routings.

Direct access Properties

sync_resources

Provides a list of trigger IDs you can reserve for use in your HVI.

HVI Engines and their Resources

Engine Collection in SystemDefinition

To include the HVI Engine of an instrument into an HVI instance, you must first add the instrument HVI Engine to the Engine collection in the SystemDefinition instance using the `add()` method. You get the Engine collection in the SystemDefinition through the `engines` property:

```
my_engine_def = my_system_def.engines.add(module.hvi.engines.main_engine, "MyEngine_1")
```


Where the `add(...)` parameters are:

1. The instrument HVI Engine identifier obtained from the instrument driver.
2. The name of the engine assigned within the HVI instances. This is a string used when referring to this HVI Engine in any following TSE API calls.

When you add the instrument HVI engine to the system definition you obtain an *EngineDefinition* object. You can also query any Engine already added to the Engines collection using the `[]` operator on the `engines` property, like in any other collection:

```
my_engine_def = my_system_defintion.engines["My_Engine_1"]
```

The following diagram shows the EngineDefinition with its methods and properties, and how it relates to the SystemDefinition:



Engine (Definition) Properties

The Engine objects returned from the Engine collection in the SystemDefinition are of the type *EngineDefinition*. The EngineDefinition instance contains properties and methods that you use to configure the Engine for use later in the Sequencer and Hvi instances. The EngineDefinition object includes Triggers, Actions, Events, and FPGA-Sandboxes collections that enable you to define and configure the Engine resources to be used within the SystemDefinition instance, and later in the Sequencer and Hvi instances created from this SystemDefinition.

Triggers (Definition)

The Trigger collection in the EngineDefinition object is used to add and manage all the Triggers that you use later throughout the TSE API. There are multiple types of triggers depending on their hardware implementation, for example:

- Front panel triggers (usually a SMA connector on the module's front panel),
- PXIe triggers (connected to the PXIe backplane of the chassis),
- General purpose digital IO (LVDS connector in the module's front panel),
- Any other trigger lines enabled within the instrument.

To use a Trigger from an HVI Engine, you must add it to the Trigger collection using the `add()` method. The Trigger collection in the Engine object is accessed through the `triggers` property:

```
fp_trigger_def = my_engine_def.triggers.add(instrument.hvi.triggers.front_panel_1, "FP Trigger")
```

Where the `add(...)` parameters are:

1. The instrument Trigger identifier, obtained from the instrument driver.
2. The name of the Trigger assigned within the HVI instance. This is a string used when referring to this HVI Trigger in any following TSE API calls.

The object returned when you add a Trigger to the EngineDefinition or when querying the Trigger from the Trigger collection is of the type *TriggerDefinition*. This TriggerDefinition object provides an interface to query the Trigger properties such as Id, Name (user defined name), Hardware Name and Type. It also expose a `config` property to configure the behavior of the Trigger:

```
fp_trigger_def_config = my_engine_def.triggers["FP Trigger"].config
```

Trigger Configuration

The `config` property holds the current configuration of the Trigger hardware and enables you to query and configure the required Trigger behavior. The Trigger `config` property includes the following properties:

Parameter	Description	Possible values	Default value
<code>direction</code>	Get or set the direction of the Trigger	Direction enum: INPUT, OUTPUT	INPUT
<code>polarity</code>	Get or set the polarity of the output Trigger	TriggerPolarity enum: ACTIVE_HIGH, ACTIVE_LOW	ACTIVE_HIGH
<code>trigger_mode</code>	Get or set the Trigger mode	TriggerMode enum: LEVEL, PULSE	LEVEL
<code>sync_mode</code>	Get or set the synchronization mode of the Trigger	SyncMode enum: IMMEDIATE, SYNC, SYNC_BASE	IMMEDIATE
<code>hw_routing_delay</code>	Get or set the delay of the Trigger in nanoseconds	Int	0
<code>pulse_length</code>	Get or set the pulse length of the Trigger in nanoseconds	Int	100ns
<code>source</code>	Get or set an input signal (trigger or event) to route instrument internal signals out of the instrument	Event	Event.SOFTWARE_HVI
<code>sampling_mode</code>	Get or Set the sampling mode	SamplingMode enum: FULL_RATE, CLK_100Mhz	FULL_RATE
<code>input_threshold</code>	Get or Set the analog input threshold to determine the logic 0 and 1 logic values	double	See Instrument documentation

NOTE

Not all triggers support all possible operation modes or configurations. Please refer to the instrument documentation for details on the capabilities supported by the different triggers available in an instrument.

In the following example, the Trigger `fp_trigger` is added to the Trigger collection and configured to be used as input.

```
# Defines the FP Trigger to be used as a wait condition by the digitizer
# Add to the HVI Trigger Collection of each HVI Engine the FP Trigger object of that instrument
#
fp_trigger_id = instrument.hvi.triggers.front_panel_1
fp_trigger_def = my_engine_def.triggers.add(fp_trigger_id, "FP Trigger")
```

```
#  
# Trigger configuration  
fp_trigger_def.config.direction = kthvi.Direction.INPUT  
fp_trigger_def.config.polarity = kthvi.Polarity.ACTIVE_HIGH  
fp_trigger_def.config.hw_routing_delay = 0  
fp_trigger_def.config.trigger_mode = kthvi.TriggerMode.LEVEL
```

NOTE

In the SystemDefinition you only configure the Triggers / Actions etc. to be used later in the Sequencer and Hvi instances created from this SystemDefinition.

Actions

The Action collection in the EngineDefinition object is used to add and manage all the Actions that you use later throughout the TSE API.

To use an HVI Action in an HVI Engine, you must add it to the Actions collection using the `add()` method. The Action collection in the Engine object is accessed through the `actions` property:

```
my_action_def = my_engine_def.actions.add(action_id, action_name)
```

Where the `add(...)` parameters are:

1. The Action identifier, obtained from the instrument driver.
2. The name of the Action assigned within the HVI instance. This is a string that is used when referring to this Action in any following TSE API calls.

The object returned when you add an Action to the EngineDefinition or when querying the Action from the Action collection is of the type *ActionDefinition*. This ActionDefinition object provides an interface to query the Action properties such as Id, Name (user defined name), Hardware Name and type. It also expose a config property to configure the behavior of the Action:

```
my_action_def_config = my_engine_def.actions[action_name].config
```

Actions are used in HVI Sequences with Action-execute instructions.

Events

The Events collection in the EngineDefinition object is used to add and manage all the Events that you use later throughout the TSE API.

To use an HVI Event in an HVI Engine, you must add it to the Events collection using the `add()` method. The Event collection in the Engine is accessed through the `events` property:

```
my_event_def = my_engine_def.events.add(event_id, event_name)
```

Where the `add(...)` parameters are:

1. The Event identifier, obtained from the instrument driver.
2. The name of the Event assigned within the HVI instance. This is a string that is used when referring to this Event in any following TSE API calls.

The object returned when you add an Event to the EngineDefinition or when querying the Event from the Event collection is of the type *EventDefinition*. This EventDefinition object provides an interface to query the Event properties such as Id, Name (user defined name), Hardware Name and Type.

FPGA-Sandboxes

The FPGA-Sandbox collection in the EngineDefinition object is used to manage the FPGA-Sandboxes that you use later throughout the TSE API.

An FPGA-Sandbox is a user-configurable region in the FPGA. You use **PathWave-FPGA** to create your design for the FPGA-Sandbox. An HVI interface is provided for the FPGA-Sandbox for the instruments that support it. Through this interface, HVI can access HVI FPGA-Registers and HVI FPGA-Memory Maps inside the FPGA-Sandbox.

When the design is completed and built, PathWave-FPGA generates a **.k7z** file. This file stores all the information needed about the names, addresses, ranges of the registers and memory-mapped locations that are connected to the HVI interface.

To use an FPGA-Sandbox design in your HVI Engine, you must first get an object that represents the FPGA-Sandbox from the FPGA-Sandbox collection in the Engine. You do this by querying the FPGA-Sandbox collection. The FPGA-Sandbox collection in the Engine object is accessed through the `fpga_sandboxes` property:

```
sandbox = engine.fpga_sandboxes["sandbox0"]
```

The sandboxes available and their names depend on each instrument, so you should refer to the instrument documentation to know the available sandboxes. The object returned when you query the FPGA-Sandboxes collection is of the type *FpgaSandboxDefinition*. To use the *FpgaSandboxDefinition* object you must load the relevant data from the **.k7z** file into it:

```
sandbox.load_from_k7z("c:/fpga/Hvi2SandboxTest.k7z")
```

Once the data from the **.k7z** file is loaded, you can use the interfaces in the *FpgaSandboxDefinition* object to access the contents of the FPGA-Sandbox, including the FPGA-Registers, FPGA-Memory Map, and FDS Ports. These are accessed through the following collections:

FPGA-Registers

Using the *FpgaSandboxDefinition* object, you can access the list of FPGA-Registers defined in the FPGA-Sandbox. The FPGA-Registers have one property, the name of the register.

```
fpga_register = engine.fpga_sandboxes["sandbox0"].fpga_registers[0]  
fpga_register.name
```

When you write sequences, you can set the FPGA-Registers as parameters in certain instructions.

FPGA-Memory Maps

Using the *FpgaSandboxDefinition* object, you can access the list of FPGA-Memory Maps defined in the FPGA-Sandbox. The FPGA-Memory Maps have two properties, the `name` and the `size` of the memory-mapped location.

```
fpga_memory_map = engine.fpga_sandboxes[SANDBOX_0_NAME].fpga_memory_maps[0]
```


The FPGA-Memory Maps can be set as a parameters in certain instructions.

FDS Ports

Using the FpgaSandboxDefinition object, you can access the list of FDS Ports in the FPGA-Sandbox. The FdsPort enables you to use the FDS Port instances placed in your FPGA-Sandbox design. An FdsPort has one property which is the name of the port.

The following example code shows how to get the FDS Ports from the FpgaSandboxDefinitions objects in two different Engines.

```
# get FDS Ports for each Engine
source_port_name = 'fds_tx_output_1'
dst_port_name = 'fds_rx_input_1'
#
module_1_fds_ports = sequencer.sync_sequence.engines["Module_1"].fpga_sandboxes[0].fds_ports
module_2_fds_ports = sequencer.sync_sequence.engines["Module_2"].fpga_sandboxes[0].fds_ports
#
source_address = 10
source_port = module_1_fds_ports[source_port_name]
source = keysight_tse.FdsPortAddress(source_port, source_address)
#
dst1_address = 20
dst1_port = module_2_fds_ports[dst_port_name]
dst1 = keysight_tse.FdsPortAddress(dst1_port, dst1_address)
```

FdsPorts can be set as parameters in Sync-FPGA Data-Sharing Statements.

Chassis and Interconnects (SyncModule)

This section describes the Chassis, and Interconnects (SyncModule) objects and how to use them. It contains the following sections:

- [The Chassis Object](#)
- [The Interconnect \(SyncModule\) Collection in the System Definition](#)
- [Opening Hardware or Simulated Devices](#)

The Chassis Object

The Chassis collection in the SystemDefinition object is used to add and manage all the Chassis that you use later throughout the TSE API.

To use a Chassis in the SystemDefinition, you must add it to the Chassis collection using the `add()` method. The Chassis collection in the SystemDefinition object is accessed through the `chassis` property. You can add a chassis with just a number or you can include additional options:

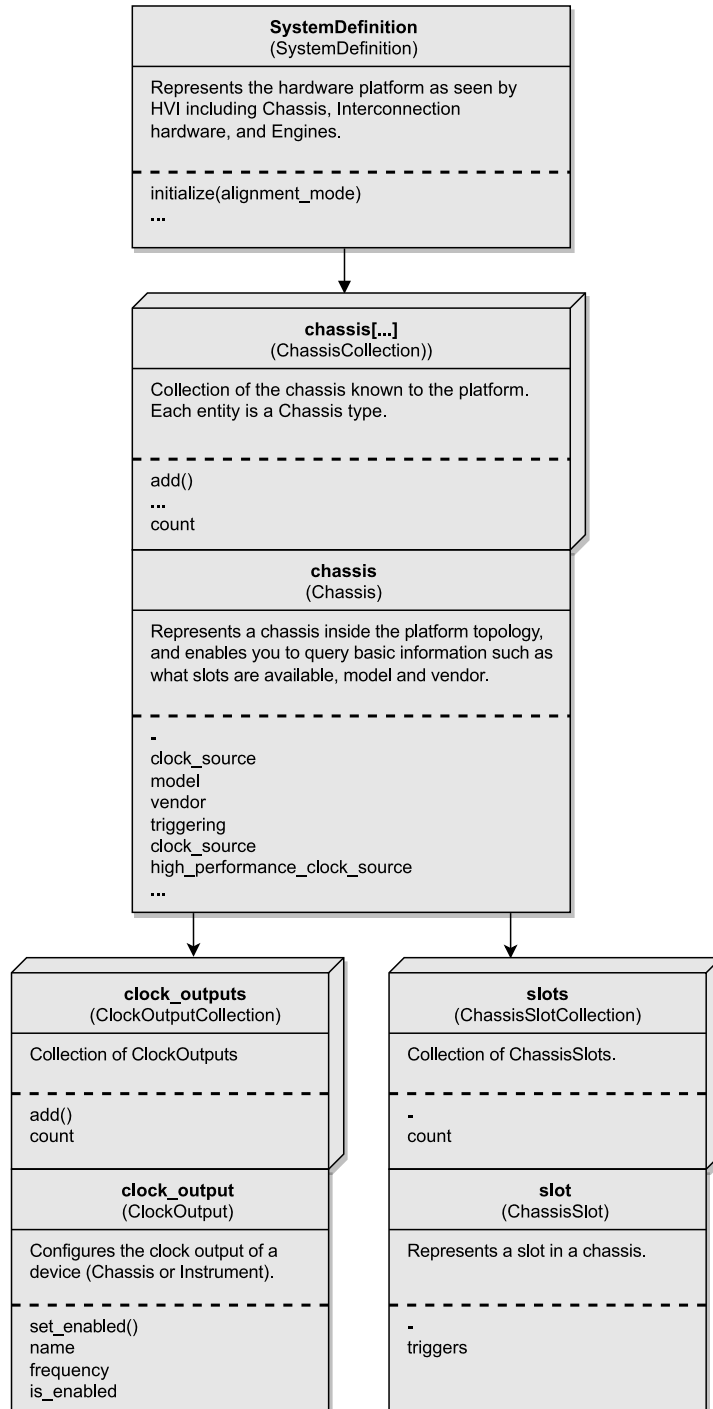
```
# Add chassis with number or options
my_system.chassis.add(chassis_number)
or
my_system.chassis.add(chassis_number, "DriverSetup=model=GenericPcieChassis" )
```

Where the `add(...)` parameters are:

1. The Chassis number, that you define.
2. Optionally, a string with parameters, See *Options for opening a Chassis*.

The object returned when you add a Chassis to the SystemDefinition object or when querying the Chassis from the Chassis collection is of the type `Chassis`. This `Chassis` object provides an interface to query the Chassis properties such as which slots are available, the Chassis model, and Chassis vendor.

The following diagram shows the Chassis, its methods, properties and relationship to the SystemDefinition. To simplify the diagram, some methods and properties are not shown:



A `chassis` has the following properties:

Property	Description
<code>number</code>	The chassis number
<code>model</code>	The chassis model
<code>vendor</code>	The chassis vendor
<code>first_slot</code>	The first slot number in the chassis
<code>last_slot</code>	The last slot number in the chassis
<code>slots</code>	Collection of slots
<code>triggering</code>	The triggering object of the chassis
<code>clock_outputs</code>	Collection of clock outputs in the chassis
<code>high_performance_clock_source</code>	Controls the High Precision Reference Clock Source

The Interconnect (SyncModule) Collection in the System Definition

The interconnects collection represents physical hardware modules or instruments that are used to connect different chassis together. A *System Synchronization Module* (SSM) is a PXIe instrument that can be used for connecting multiple chassis together. It can synchronize the multiple chassis and the instruments within them, share a high performance clock reference across the multi-chassis system, and managing *Fast Data Sharing* (FDS) between PXIe instruments.

To include an SSM into an HVI instance, the first step is to add the SSM to the Interconnect collection in the SystemDefinition instance. This is slightly different from other collections because there is no `add()` method, instead you use the `add_sync_module()` method. You get the Interconnects collection in the SystemDefinition through the `interconnects` property:

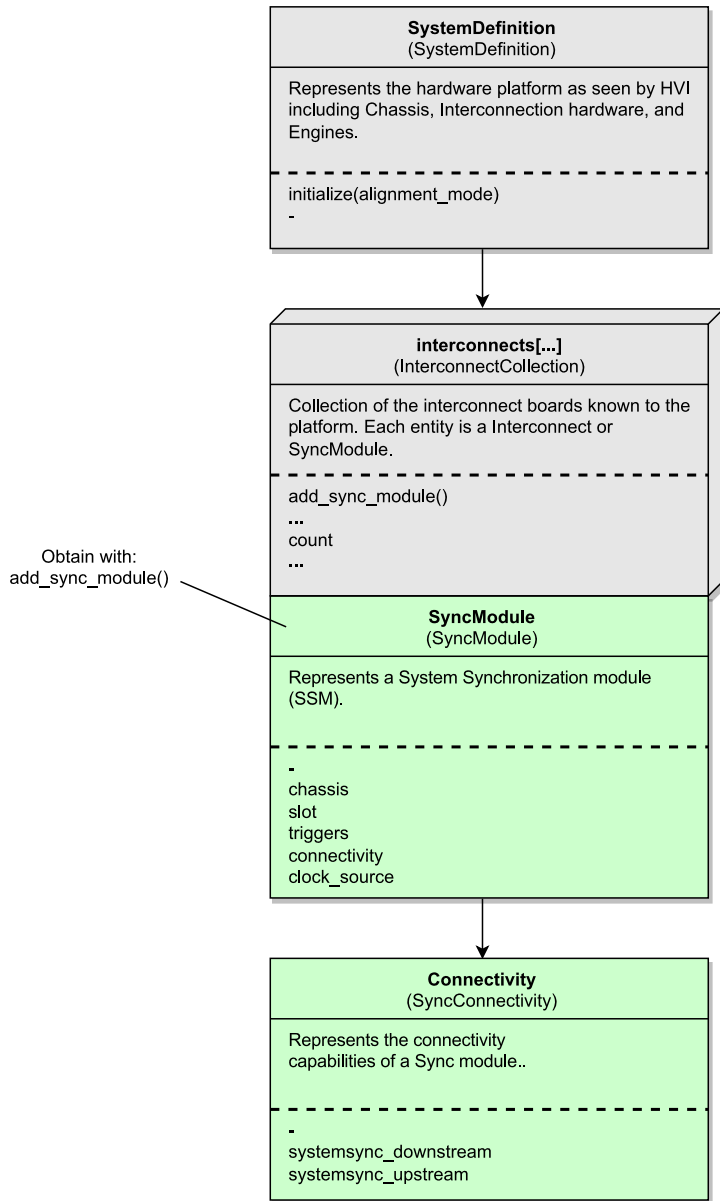
```
# Define System Sync Module
resource_id_ssm_1 = 'PXI0::CHASSIS1::SLOT10::INSTR'
options1 = ""
#
sync_module_1 = my_system_definition.interconnects.add_sync_module(resource_id_ssm_1, options1)
```

Where the `add_sync_module(...)` parameters are:

1. The instrument resource identifier that you define.
2. An options string, see *Options for opening SSMs*.

When you add the SSM to the Interconnects collection you obtain an *SyncModule* object (shown in green in the following diagram). You can also query any SyncModule already added to the Interconnects collection using the `[]` operator on the `Interconnects` property, like in other collections.

The following diagram shows the Interconnects and its methods and properties, the SyncModule, and how these relate to the SystemDefinition:



The SyncModule object has the following properties:

Property	Description
<code>chassis</code>	The chassis number where the SSM is located.
<code>slot</code>	The slot number where the SSM is located.
<code>clock_source</code>	Gets an object to control the SSM clock source.
<code>connectivity</code>	Gets the Connectivity object describing the connectivity capabilities of the SSM. You must use this to specify connections in software that reflect what is connected in your hardware setup.
<code>triggers</code>	Gets the Triggers collection for the SSM.

Connectivity

This describes the connectivity capabilities of an SSM. This gives your access to the physical System Sync ports on the SSM.

Opening Hardware or Simulated Devices

You can use PathWave Test Sync Executive with hardware or simulated devices and instruments. The simulation mode enables you to test your Sequences before running them on actual hardware.

When you are opening a device such as a SSM or a Chassis, you can specify an options string. This is a string that contains a list of comma separated options. The options you specify are specific to the device you are opening, and can change depending on if you are opening real device or using a simulation.

NOTE

In some cases a generic simulation built-in to HVI is provided, this is to enable you to get things up and running. A driver based simulation provides a more accurate simulation of the real hardware, so it is better for testing.

Options for opening SSMs

Hardware SSM

If you are using SSM hardware, the options sting is typically empty. If you want to specify hardware options when using the SSM, see the SSM user manual for available options.

```
# Add SSM to Interconnects Collection
interconnects.add_sync_module(resource_id,"")
```

Simulated SSM

You can simulate a specific SSM with the driver for that SSM. When simulating several options should be specified:

- `Simulate=True`.
- The following option must go after `DriverSetup=`
 - `Model` specifies the model of SSM you want to simulate.

You can add a simulated SSM in the following way:

```
interconnects.add_sync_module(resource_id, 'Simulate=true,DriverSetup=Model=M9033A')
```

Options for opening a Chassis

Hardware Chassis

To add a real chassis do the following:

```
# Add chassis with number  
my_system.chassis.add(chassis_number)
```

or you can include options with:

```
my_system.chassis.add(chassis_number, options)
```

currently, the only options supported are the ones described in Simulated Chassis part below. Any other options you provide are ignored.

Simulated Chassis

You can simulate a chassis using a generic chassis simulation that is built in to HVI.

To enable chassis simulation, use the method: `add(chassis_number, options)`.

- `chassis_number` is the number of the chassis you want to simulate.
- `options` is a string that contains a list of comma separated options. You use these options to enable simulation mode and the chassis simulation, any other options are ignored.

The following code shows how to add a chassis in simulation mode using the built-in generic chassis simulation:

```
sys_def.chassis.add(chassis_number, 'Simulate=True,DriverSetup=Model=GenericPxieChassis')
```

The *GenericPxieChassis* also simulates the the *High Performance Reference Cock Source* (HPRCS) if required. Provided you have the drivers for a commercial chassis installed, you can specify that chassis Model to use the specific chassis driver simulation capabilities.

Synchronization Resources

This section describes Synchronization Resources, it contains the following sections:

- [HVI synchronization resources](#)
- [Where Sync Resources are used](#)
- [Calculating the number of Sync Resources required](#)

HVI provides transparent multi-instrument synchronization and synchronized conditional execution, for example, the Sync While Statement does synchronized conditional execution. To use these capabilities, for a *Device Under Test* (DUT) or instruments that do not integrate HVI technology, you must assign HVI synchronization resources and specify clock frequencies.

HVI synchronization resources

When you set up your system, you must allocate sufficient synchronization resources for your system and Sequences to work correctly. Sync resources in the PXIe platform consist of the PXI Trigger lines. These are a limited resource, so you must be careful when you are allocating them.

The sync resources are used internally by the HVI to implement the following cross-instrument operations, transparently to the user:

- Alignment and Synchronization initialization.
- Real-time Sequencing multi-instrument operations, such as:
 - Sync While.
 - Sync Register-Sharing.
 - Triggered synchronization in a SyncMultiSequenceBlock.

The HVI optimizes the use of Sync resources as much as possible and reuses the same Sync resources when possible for different operations, providing they are executed with sufficient time separation. You can estimate the number of Sync resources you require by working out how many are required at the different stages of your application.

The Sync resources consist of PXI triggers and are defined by the enumeration `keysight_tse.TriggerResourceId`. The resources must be specified in the `SyncResources` property of the `SystemDefintion` object. For example:

```
# Add Sync resources
sys_def.sync_resources = [keysight_tse.TriggerResourceId.PXI_TRIGGER0,
                          keysight_tse.TriggerResourceId.PXI_TRIGGER1,
                          keysight_tse.TriggerResourceId.PXI_TRIGGER2]
```

Where Sync Resources are used

There are 3 areas where you require Sync resources, these are the same 3 stages you set up your HVI in:

System Initialization

Initialization (`SystemDefinition.initialize()` call) requires one Sync resource for instrument synchronization. At this step the Sync resources is configured in HW and used to synchronize all hardware in the `SystemDefinition`, it is important that at this point the Sync resource is available and not in use by any other HVI instance or application. This Sync resource is reused later for the Sequence execution, for example, if you use `PXI_TRIGGER0` for synchronization, it is later reused for the Sequence execution.

Sequence Compilation

The HVI Sequence requires Sync resources to execute specific multi-instrument real-time operations. Some operations that require Sync resources include:

- Sync While.
- Sync Register-Sharing.
- Triggered synchronization in a Sync Multi-Sequence Block.

During the Sequence compilation, HVI allocates the Sync resources assigned in the `SystemDefinition` as required. So it is important that sufficient Sync resources are assigned for the Sequence to compile, if this is not the case, a compilation error will be generated. At the compile stage, Sync resources are not used in hardware, they are just allocated to specific real-time operations in the code resulting of the Sequence compilation. These resources will be configured and used in hardware when the HVI instance is loaded to hardware.

Sequencer Creation

Your Sequence shall require Sync resources to operate, but it can reuse the Sync resources previously used in the `SystemDefinition` for initialization.

If you have not called initialize or it is otherwise required, the initialization still occurs at the beginning of the Sequence. Sync resources are required for this, however these resource are reused by the Sequence.

The numbers of Sync resources required in a Sequence depends on:

- The use of certain Sync Statements such as Sync While require 1 sync resource.
- The use of Sync Register-Sharing Statements requires 1 Sync resource per bit.

- Triggered synchronization requires 1 additional Sync resource.
- The arrangement of your system also affects the number of Sync resources required.

HVI Load to Hardware

The Sync resources required to initialize the system (synchronize all hardware) and those allocated to the HVI Sequence during compilation, are configured into hardware at this step (`Hvi.load_to_hw()` call). The same Sync resources used to initialize the system are also used to run the HVI Sequence. It is important that at the time of the `Hvi.load_to_hw()` call, to ensure the allocated Sync resources are not already in use in hardware by any other HVI instance or application.

Calculating the number of Sync Resources required

Different functionalities require different amounts of Sync resources, this can also depend on the system configuration, in particular if it is a small setup such as a single PXIe-chassis & single segment, or a large system with multiple chassis.

Sync resource usage per functionality

The following table summarizes the Sync resources required by the different functionalities.

Functionality		Sync resources required (for recommended Keysight chassis)	
#	Description	Single PXIe chassis & Segment	Others
1	SystemDefinition::Initialize() and Sequence start in Hvi::Run()	1	
2	Sync While Statement	1	
3	Sync Multi-Sequence Blocks with Triggered-Sync (those with unknown execution time during compilation)	1	2
4	Sync Register-Sharing of N bits	N	

For information about recommended chassis, see [Configuring a System with SSMs and System Sync Connectivity](#).

Sync resource reuse across functionalities

HVI reuses the same Sync resources for different functionalities and also for the same functionality if executed multiple times. The criteria to reuse Sync resources is:

- Functionalities #1, #2 and #3 reuse the Sync resources.
- Functionality #4 (Sync Register-Sharing) reuse Sync resources ONLY when sender module are in the same Chassis and Segment.

Calculating the total Sync resources required

To calculate the total amount of Sync Resources required, use the following formula:

- Total Sync Resources = $\text{Max}(\#1, \#2, \#3) + \text{Sum}(\text{Max}(\#4 \text{ for each segment}))$.
- If a functionality is not used, use 0 in the equation above.

The following table shows examples with the number of sync resources required:

Scenario Description	Functionality				Sync Resource Total
	#1	#2	#3	#4	
System initialization only, SystemDefinition::Initialize() (any number of chassis)	1	-	-	-	1
SyncSequence (1x chassis, 1x segment) No Triggered-Sync SyncMultiSequenceBlocks + Sync-While	1	1	-	-	1
SyncSequence (1x chassis, 1x segment) + Triggered-Sync SyncMultiSequenceBlocks No Sync-While + RegSharing (chassis1, segment 1) (n bits)	1	-	1	n	n + 1
SyncSequence (2+ chassis) + Triggered-Sync SyncMultiSequenceBlocks No Sync-While + RegSharing (chassis1, segment 1) (n bits)	1	-	2	n	n + 2
SyncSequence (2+ chassis) + Triggered-Sync SyncMultiSequenceBlocks + Sync-While + RegSharing (chassis 1, segment 2) (n bits) + RegSharing (chassis 1, segment 2) (m bits)	1	1	2	Max (n,m)	Max(n,m) + 2
SyncSequence (2+ chassis) + Triggered-Sync SyncMultiSequenceBlocks + Sync-While + RegSharing (chassis 1, segment 1) (n bits) + RegSharing (chassis 2, segment 3) (m bits)	1	1	2	n + m	n + m + 2

Synchronization Signals and Sync Modes

HVI uses different periodic digital signals for synchronization purposes. The definition of those digital signals depends on platform and instruments signals. Platform signals are the CLK100 and CLK10 signals in a PXI platform such as a PXI chassis. Instruments have different clock signals inside that are classified as core clocks or system clocks. Platform and instrument clock signals contribute to define the following HVI synchronization signals:

- SYNC.
- SYNC_BASE.

Synchronization modes

You can configure the synchronization mode. This is used, for example, for generating a Trigger value or waiting for an Event.

The following modes are supported:

IMMEDIATE

The Trigger or Action is issued immediately, with no need to wait for any common synchronization clock. For the Wait-For-Event, the HVI execution continues immediately, as soon as the Event is received.

SYNC

The Trigger or Action is issued at the first edge of the SYNC signal. For the Wait-For-Event, the HVI execution continues at the first edge of the SYNC signal, following the Event arrival time.

SYNC_BASE

The Trigger or Action is issued at the first edge of the SYNC_BASE signal. For the Wait-For-Event, the HVI execution continues at the first edge of the SYNC_BASE signal, following the Event arrival time.

For more information about synchronization, see [Synchronization Clocks, Signals, and Modes](#).

Non-HVI Clocks

Non-HVI clocks are clocks on instruments that are present in a system, but do not have HVI Engines. In order to account for these additional clock frequencies on the entire system timing, TSE API allows to specify two types of additional clocks:

Non-HVI system clocks

Non-HVI system clocks are those clocks used by the instrument that do not directly impact the operation of the specific feature that the HVI must trigger. System clocks are used by the HVI to determine the SYNC_BASE period.

You must set these with the property `SystemDefinition.non_hvi_system_clocks`

Non-HVI core clocks

Non-HVI core clocks are instrument clocks that directly impact the operation of the specific feature that the HVI must trigger. Core clocks are used by the HVI to determine both the SYNC and the SYNC_BASE period.

You must set these with the property `SystemDefinition.non_hvi_core_clocks`

For more information see [Synchronization Clocks, Signals, and Modes](#).

System Initialization in the SystemDefinition Object

System Initialization is a process that includes the configuration and alignment of the different systems clocks, including the clocking for each instrument specified as part of the HVI SystemDefinition.

The SystemDefinition object includes an `initialize()` method that initializes the hardware included in the SystemDefinition, and performs synchronization and clock alignment. There are 3 cases where system initialization and clock alignment can occur:

- Manually calling `initialize()` in the SystemDefinition.
- When the Sequencer object is created.
- When calling Load to Hardware.

The SystemDefinition initialize() method

The `systemDefinition.initialize(...)` method enables you to explicitly trigger a system initialization and alignment. You might want to explicitly control when the initialization process is executed because the initialization process can take some considerable time depending on the parameters and system state.

You can also specify specific alignment modes when calling the `initialize()` method.

The following table shows the alignment mode parameters:

Mode	Description
<i>Default</i>	Calling <code>initialize()</code> without parameters performs the default, or minimal-possible initialization. This is the mode intended to be used in normal system operation.
<code>keysight_tse.AlignmentModes.Full</code>	Forces a full system, complete system initialization and alignment.
<code>keysight_tse.AlignmentModes.ResetCalibration</code>	Performs system initialization and alignment resetting and regenerating the stored calibration data.
<code>keysight_tse.AlignmentModes.PreCalibration</code>	Performs system initialization and alignment, ignoring any missing calibration data. This mode is intended for system warm-up or other instances when the use of precise alignment calibration data is not available yet or not required.

Alignment mode parameters (continued).

Mode	Description
<code>keysight_ tse.AlignmentModes.ForceClockMonitoring</code>	During system operation, TSE can run a clock monitoring routing in the background and use the information for optimal system alignment. The clock monitoring is enabled automatically if any of the instruments with HVI Engines added into the SystemDefinition benefits from the clock monitoring information. This flag force enables the clock monitoring, regardless of the HVI engines added to the SystemDefinition.
<code>keysight_ tse.AlignmentModes.DisableClockMonitoring</code>	This flag force disables the clock monitoring regardless of the instrument's requirements. Disabling clock monitoring speeds up the system initialization and load-to-hardware operations but may result in non-optimal system alignment.

For a full explanation of initialization and all the parameters, see [System Initialization](#).

User-Defined Trigger Routing

About Triggering

Triggers are logic signals that enable you to communicate events or states between instruments. In a PXI system, instruments can be located in different slots of a PXI chassis, and in different PXI chassis connected through a SystemSync connection using System Synchronization Modules (SSMs). TSE includes API methods that enable users to configure custom Trigger routings across different instruments and chassis.

Trigger Resources in TSE API

The TSE API supports the following types of Trigger:

Triggers

Instrument Triggers

Instruments with HVI engines expose triggers that can be accessed and controlled with TSE API. These triggers are registered and obtained through the HVI Engines instances when you setup your SystemDefinition, see [HVI Engines and their Resources](#). When using PXI triggers, you must be careful that if they are assigned as sync resources, they cannot be used for general purpose triggering.

Trigger resources available depend on each instrument. Instruments drivers exposes a Triggers property in the Hvi interface that includes all available triggers in that instrument.

The following snippet demonstrates how to register and access Instrument Triggers from TSE API:

```
# Retrieve Engine instance
instr1_engine = my_system_definition.engines["myInstr1Engine"]
#
# Create/Add Trigger object for instrument trigger through the Engine instance
my_instr_trigger = instr1_engine.triggers.add(my_instr1.hvi.triggers.PxiTrigger0,
"DestinationTrig")
#
# Triggers once added can be retrieved by means of the label assigned when adding it
my_instr_trigger = instr1_engine.triggers["DestinationTrig"]
```

PXIe DSTAR Triggers

PXIe DSTAR Triggers are only available in a PXIe chassis when a System Timing Module is present. TSE only supports PXIe DSTAR triggers when M9032A/33A PXIe System Synchronization Modules are present in the chassis. The PXIe DSTAR triggers are accessed through the Chassis Slot instance:

```
system_definition.chassis[...].slot[...].Triggers[string]
```

Where the string identifies the specific trigger, supported valued are:

- "PXIe_DSTARB"
- "PXIe_DSTARC"

The following snippet demonstrates how to retrieve PXIE_DSTARB/C:

```
## Retrieve Trigger object for PXIe_DSTARB in Slot 5
dstar_trig = system_definition.chassis[1].slot[5].Triggers["PXIe_DSTARB"]
```

System Synchronization Module (SSM) Triggers

System Synchronization Modules expose trigger resources that can be accessed in the TSE API through the interconnect instance. TSE 2023 only supports STRIG_IO for M9032A/33A SSMs. The snippet below demonstrates how to access STRIG_IO:

```
# Retrieve Trigger Object for STRIG_IO in SSM 1
SSM1_triggerInOut = my_system_definition.Interconnects["SSM_1"].Triggers["STRIG_IO"]
```

Trigger Configuration

Triggers include a configuration property that enables you to control several different capabilities in the triggers. Not all triggers support all capabilities, see [HVI Engines and their Resources](#) for more information about Trigger configuration. For instruments triggers, refer to the instrument documentation for details about the supported features for the different triggers available. For PXIe DSTAR and SSM triggers, the following table summarizes the properties that are not supported:

Trigger	Properties not supported
PXIe DSTAR	source, hw_routing_delay, input_threshold
SSM STRIG_IO	source, hw_routing_delay

The following snippet demonstrates the use of trigger configuration properties:

```
# Trigger configuration
my_trigger.config.direction = keysight_tse.Direction.INPUT
my_trigger.config.polarity = keysight_tse.Polarity.ACTIVE_HIGH
my_trigger.config.hw_routing_delay = 0
my_trigger.config.trigger_mode = keysight_tse.TriggerMode.LEVEL
```

Trigger Runtime configuration and operation

Once the HVI instance is loaded to hardware, for triggers configured as outputs, it is possible to change the trigger configuration, read trigger state, and change state. See [Engine Runtime Components](#) for more details.

Platform Triggers

Platform Triggers are defined by the platform and are typically used when working with instruments or resources that are not HVI-capable (do not contain an HVI Engine). Platform resources are identified by a TriggerResourceId, supported triggers are:

- PxiTrigger0, PxiTrigger1, PxiTrigger2, PxiTrigger3, PxiTrigger4, PxiTrigger5, PxiTrigger6, PxiTrigger7

Platform triggers are accessed through the triggering interface in the chassis instance:

```
system_definition.chassis[...].triggering.get_platform_trigger(keysight_tse.TriggerResourceId
trigger, int slot)
```

The following snippet demonstrates how to retrieve PXI_TRIGGERS:

```
# Retrieve PXI_TRIGGER[0] in Slot 11
platform_trigger = system_definition.chassis[1].triggering.get_platform_trigger(keysight_
tse.TriggerResourceId.PXI_TRIGGER0, 11)
```

Platform Trigger Configuration

Platform triggers do not have any configuration accessible from the TSE API. The Platform Triggers configuration is determined automatically by TSE depending on the SystemDefinition and Sequencer definitions. For instance, the section below describes how to define trigger routings, but, when defining routings for the platform triggers, their configuration is set automatically depending on the routing definition.

Routing Triggers Across Instruments and Chassis

Defining custom trigger routing

In addition to the automatic triggering configuration performed by TSE on the triggers assigned as Sync Resources (see [Synchronization Resources](#)), TSE exposes methods for users to configure custom trigger routings.

NOTE Triggers assigned as Sync Resources cannot be used for user-defined routings or configured by the user. These triggers are managed and configured automatically by TSE.

Trigger routings have some constraints depending on the type of trigger resource used:

- Routings cannot include together PXI_TRIGGER with PXIE_DSTARB/C or SSM STRIG_IO
- For PXI_TRIGGER routings, source and destinations must all correspond to the same PXI_TRIGGER index.

The trigger routings are configured through the triggering interface in the SystemDefinition:

- `SystemDefinition.triggering.routings.Add(label: str, source: TriggeringSignal, destinations: list[TriggeringSignal])`
- `SystemDefinition.triggering.routings.Add(source: TriggeringSignal, destinations: list[TriggeringSignal])`

Applying custom trigger routing to hardware

Trigger routings are applied to hardware when the HVI instance is loaded to hardware. The triggering interface in the HVI instance enables you to query the trigger routings defined in the SystemDefinition:

- `Hvi.triggering.routings[label: str]`
- `Hvi.triggering.routings[index: int]`

Keysight recommends you use the `add()` method in the SystemDefinition with a label, so this label can be used later in the HVI instance to retrieve the specific routing. When using the integer index to access the routing collection, the index is associated to the order when adding the routing.

The routing returned by the HVI instance enables you to query the following properties:

Property	Type	Description	Note
label	String	Label assigned when added in the systemDefinition instance	-
source_signals	list [TriggeringSignal]	List of source triggers as defined in the systemDefinition instance	-
dest_signals	list [TriggeringSignal]	List of destination triggers as defined in the systemDefinition instance	-
latency	float	Routing latency in nanoseconds	Only supported for PXIE_DSTARB/C and SSM STRIG_IO

PXI_TRIGGER routing example

The following example shows how to route PXI_TRIGGERS using both Platform Triggers and Instrument Triggers. Platform triggers are used in general when the PXI instrument source or destination of the trigger do not support HVI technology, in other cases it is recommended to use the trigger from the HVI Engine:

```
### Create PlatformTrigger through the Chassis object for PXI_TRIGGER[0] in slot 11
slot_trigger = my_system_definition.chassis[1].triggering.get_platform_trigger(keysight_
```

```

tse.TriggerResourceId.PXI_TRIGGER0, 11)
#
### Create Instrument Trigger object through the HVI Engine object
instr_engine = my_system_definition.engines["Instr1"]
instr_trigger = instr_engine.triggers.add(Instr1.hvi.triggers.PxiTrigger0, "Destination")
#
# Configure Instrument trigger
instr_trigger.config.direction = keysight_tse.Direction.INPUT
#... other properties
#
### Add a user routing from PXI_TRIGGER[0] in slot 11 (platform trigger) to Instr1
routing = my_system_definition.triggering.routings.add("MyPxi0Routing", slot_trigger, [instr_
trigger])
#
### Create Sequencer and Hvi instances
my_sequencer = keysight_tse.Sequencer("sequencer", my_system_definition)
#
#... Optionally define a real-time sequence
#
# Compile Sequencer to obtain "executable" HVI instance
my_hvi = my_sequencer.compile()
#
### The trigger routing is applied to HW in the load_to_hw()
my_hvi.load_to_hw()
#
# ... Optionally run the real-time sequence
#
# The trigger routing is released from the hardware
my_hvi.release_hw()

```

SSM STRIG_IO and PXIE_DSTARAR routing example

The following example illustrates how to route the STRIG_IO from an SSM to the DSTARB/C in different slots in different chassis.

```

### Create PXIE_DSTARB destination triggers objects (DSTARB in the slot is always an output)
DStarB_ch1_slot3 = my_system_definition.chassis["PXI0::1::BACKPLANE"].slot[3].Triggers["PXIE_
DSTARB"]
DStarB_ch1_slot12 = my_system_definition.chassis["PXI0::1::BACKPLANE"].slot[12].Triggers["PXIE_
DSTARB"]
DStarB_ch2_slot5 = my_system_definition.chassis["PXI0::2::BACKPLANE"].slot[5].Triggers["PXIE_
DSTARB"]
#
### Create SSM STRIG_IO triggers objects and configure it as input
SSM1_triggerIn = my_system_definition.Interconnects["MySSM_1"].Triggers["STRIG_IO"]
#
# Configure Trigger (use same configuration interface as all triggers)
SSM1_triggerIn.config.direction = keysight_tse.Direction.Input (throw exception if not supported
specified direction)
#
### Define STRIG_IO routing to PXIE_DSTARBs

```

```
my_strig_routing = my_system_definition.triggering.routings.add("StrigRouting", SSM1_triggerIn,
[DStarB_ch1_slot3, DStarB_ch1_slot12, DStarB_ch1_slot5])
#
### Create Sequencer and Hvi instances
my_sequencer = keysight_tse.Sequencer("sequencer", my_system_definition)
#
#... Optionally define a real-time sequence
#
# Compile Sequencer to obtain "executable" HVI instance
my_hvi = my_sequencer.compile()
#
### The trigger routing is applied to HW in the load_to_hw()
my_hvi.load_to_hw()
#
# after the load to hardware can retrieve trigger objects and change configuration
SSM1_triggerIn_runtime = my_hvi.interconnects["MySSM_1"].Triggers["STRIG_IO"]
SSM1_triggerIn_runtime.config.polarity= keysight_tse.Polarity.ACTIVE_LOW
SSM1_triggerIn_runtime.apply_config()
#
# after the load to hardware can query Latency property
my_routing = my_hvi.triggering.routings["StrigRouting"]
my_routing.latency
#
# ... Optionally run the real-time sequence
#
# The trigger routing is released from the hardware
my_hvi.release_hw()
```

System Clocking Configuration

In a hardware system, there are a number of different options for the system clocking. The clocking property in the `SystemDefinition` object enables you to define the source of the system wide clocking reference along with a mode and frequency.

Setting the Source of the Reference Clock

You can set a *System Sync Module* (SSM) as a reference clock source.

```
# Select the SSM as the source
clockSource = interconnects[0].clock_source
#
# Set the SSM clock source
systemDefinition.clocking.reference_source = clockSource
```

Alternatively, you set the chassis as a reference clock source with the following code:

```
# Select the chassis as the source
clockSource = chassis.clock_source
#
# Set the clock reference source
systemDefinition.clocking.reference_source = clockSource
```

Setting clock reference Mode and Frequency

You can set the mode as `INTERNAL` or `EXTERNAL`.

INTERNAL

The reference clock source is internal. This is the default value.

Do not set the frequency, this raises an error.

EXTERNAL

The reference clock source is synchronized to an external clock.

You must set the frequency (in Hz) of the external sources.

To set the *High Performance Reference Clock Source* (HPRCS) as the reference clock do the following:

```
# This will be added only in the main chassis with the leader SSM.
ktHvi.SystemDefinition definition("Name")
#
chassis = definition.add_chassis(1)
clockSource = chassis.high_performance_clock_source
#
# Configuring HPRCS to use its internal clock
clockSource.set_mode(keysight_tse.ClockingReferenceMode.INTERNAL)
#
# Configuring HPRCS to use an external reference @10Mhz
clockSource.set_mode(keysight_tse.ClockingReferenceMode.EXTERNAL, 10e6)
#
# As SSM is required because the HPRCS output is connected to REF_IN input of the SSM.
syncModule = definition.interconnects.add_sync_module(resourceIdStm1, options)
#
definition.clocking.reference_source = clockSource
```

In some cases you may want to use the clock source device (Chassis, SSM or HPRCS) internal clock. This also enables you to use an external clock source to drive it, such as an atomic clock or a device under test. The following code shows how to configure the chassis as the clock source and take the clock reference from an external 10MHz source:

```
clockSource = chassis.clock_source
#
# Set clock mode to EXTERNAL and set frequency to 10MHz
clockSource.set_mode(keysight_tse.ClockingReferenceMode.EXTERNAL, 10e6)
#
systemDefinition.clocking.reference_source = clockSource
```

However, if you want to explicitly configure the clock source to use the chassis internal OCXO clock source:

```
clockSource = chassis.clock_source
#
# Set clock mode to INTERNAL
clockSource.set_mode(keysight_tse.ClockingReferenceMode.INTERNAL)
#
systemDefinition.clocking.reference_source = clockSource
```

You can query the clock reference mode and frequency with the following code:

```
# Get mode and frequency (in Hz)
#
mode = clockSource.mode
frequency = clockSource.frequency
```

Chassis Clock Outputs

The chassis have internal clocks and outputs for them. For instance, the clock output on the rear panel of an M9019A or the clock outputs on the front panel of an M9046A. These clock outputs can be used as a reference clock for instruments in the system and for devices external to the system. The TSE API enables you to enable or disable the chassis clocks.

For the clock options and outputs available on your chassis, see your chassis documentation.

NOTE

Chassis Clock Outputs do not have any default behavior. If the user does not specify any configuration for a clock output (see below) the clock output is left untouched.

Enabling chassis clock outputs

The chassis clock outputs are available in the chassis and you can access them by their name as follows:

```
# Get the Clock configuration for the Rear Panel 10MHz output port from the Chassis
ktHvi.SystemDefinition definition("Name")
#
chassis = definition.add_chassis(1)
#
clockOutputRp10Mhz = chassis.clock_outputs["RP10MHzOut"]
clockOutputRP10Mhz.set_enabled(true/false)
```

Some clock outputs support one single frequency and others support multiple frequencies. For the outputs supporting only one frequency, no frequency must be provided when enabling/disabling them. If the clock outputs do support multiple frequencies, you must specify what frequency (in Hz) you want to enable.

When you disable the clock, the frequency argument is ignored.

The following code shows some examples and error cases:

```
clockOutputRp10Mhz = chassis.clock_outputs["RP10MHzOut"]
clockOutputRP10Mhz.set_enabled(true) # Ok
clockOutputRP10Mhz.set_enabled(true, 10e6) # Throws error, no frequency expected
clockOutputRP10Mhz.set_enabled(false, 10e6) # Ok, frequency is ignored
#
clockOutputFpRef20out = chassis.clock_outputs["FPref20Out"]
clockOutputFpRef20out.set_enabled(true) # Throws error, frequency expected
clockOutputFpRef20out.set_enabled(true, 10e6) # Ok
clockOutputFpRef20out.set_enabled(false, 10e6) # Ok, frequency is ignored
```

Enabling the chassis Analog Clock Output

If you are using an analog clock output from a chassis you must enable it manually.

The following code shows how to enable a 2.4GHz analog clock output from an M9046A chassis.

```
clock_output_2_4GHz = ref_chassis.clock_outputs["FP2.4GHzOut"]
clock_output_2_4GHz.set_enabled(True)
```

Automatic clock output enable

If you defines the HPRCS clock from the M9046A chassis as the clock source, you must connect the Front panel **Ref 1 Out** port (FPRef1Out) from the leader M9046A to the REF_IN of the leader SSM. HVI automatically enables the **Ref 1 Out** clock output port from the leader chassis to let the leader SSM take the clock reference from it.

The following code shows an example:

```
# This is only added in the main chassis with the leader SSM.
ktHvi.SystemDefinition definition("Name")
#
chassis = definition.add_chassis(1)
clockSource = chassis.high_performance_clock_source
#
# Configure the HPRCS to use its internal clock
clockSource.set_mode(key sight_tse.ClockingReferenceMode.INTERNAL)
#
# As SSM is required because the HPRCS output is connected to REF_IN input of the SSM.
syncModule = definition.interconnects.add_sync_module(resourceIdStm1, options)
#
definition.clocking.reference_source = clockSource
#
# The following lines are not required, HVI does this automatically
#clockOutputFpRef1Out = chassis.clock_outputs["FPRef1Out"]
#clockOutputFpRef1Out.set_enabled(true)
```

Enabling the Analog Clock Source in Instruments

For instruments that require an analog clock, you must set the source and frequency of the analog clock in your SystemDefinition.

You can set parameters for the analog clock:

- The source as internal or external.
- The frequencies of the sources, in Hz.

For external sources, the source selected depends on the analog clock frequencies that the instrument supports.

- If you indicate multiple frequencies, the first external frequency supported by the instrument is selected.
- If none of the external frequencies are supported, and the instrument has an internal clock, the internal clock is selected.

- If none of the external frequencies are supported, and the instrument does not have an internal clock, an error is generated.

```
my_system.clocking.enable_external_analog_clocks(frequencies)
```

If the instrument does not support the frequency and does not have an internal clock, an error is generated.

The Sequencer Object

This section describes the Sequencer object, it contains the following sections:

- [SyncSequence and Sequence objects](#)
- [Sequence Statements](#)
- [Sync Statements](#)
- [Local Statements](#)
- [EngineView Components](#)
- [InstructionSet Object](#)
- [HVI Registers and Scopes](#)
- [HVI Compilation](#)
- [System Initialization in the Sequencer Object](#)
- [Sequence Representation](#)

You use the Sequencer object to define, program, and compile your HVI Sequences.

The Sequencer object includes a top level Sync Sequence, known as the Global Sync Sequence. Within this:

- You can add Sync Statements within the top level Sync Sequence using the methods provided in the `sync_sequence` object.
- You can create local sequences to control each instrument (or HVI engine) individually using the `SyncMultiSequenceBlockStatement` which exposes a `sequences` collection with Local Sequence objects for all HVI engines.
- You add Local Statements to Local Sequences using the methods of each `sequence` object.

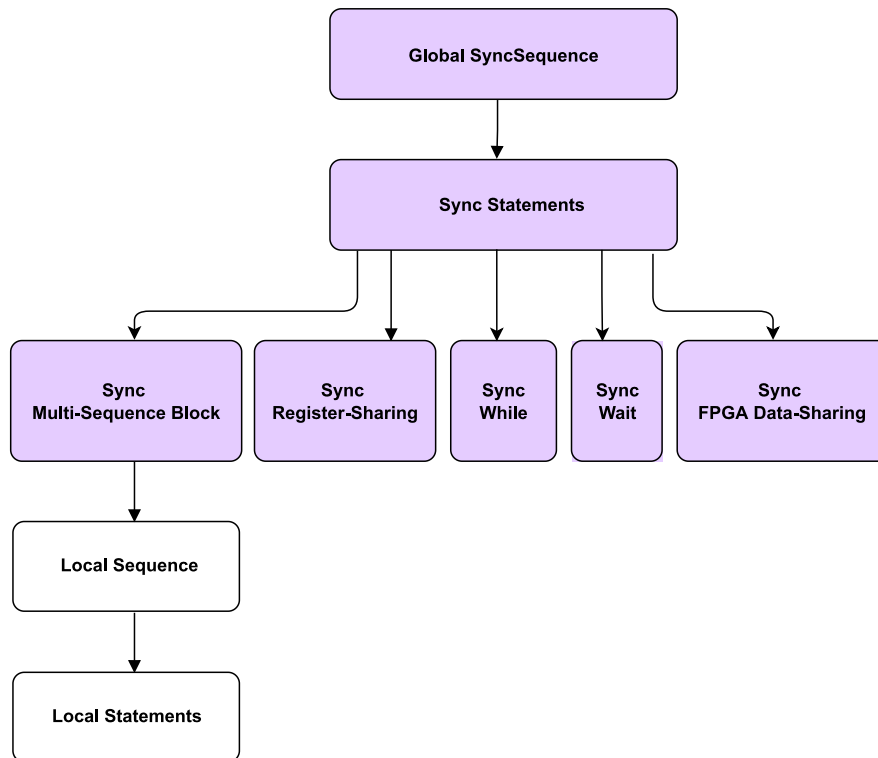
The Sequences and Statements you add can access the resources previously defined using the `EngineCollectionView` property of the `sync_sequence`. The view properties and objects enable you to see the definitions you have set up, but you cannot modify them.

NOTE


WARNING : Once a Sequencer instance is created, any change to the SystemDefinition will not affect the Sequencer object or any objects or definitions inside the Sequencer (Engines, Triggers, Actions, Events,...). All API calls on the Sequencer object *must* use the properties in the Sequencer and its hierarchy. Do not use properties objects from the SystemDefinition or other Sequencer instances under any circumstances.

Once you have defined all the Sequences that define your HVI, you must compile it. The HVI instance `Hvi`, is generated when you compile the Sequencer object successfully.

The following diagram shows the hierarchy of Sequences and Statements:



The following diagram shows the Sequencer methods and properties:



SyncSequence and Sequence objects

There are two types of HVI Sequence objects that enable HVI Sequence programming and usage:

- `SyncSequence`.
- `Sequence`.

HVI uses a `sync_sequence` object to manage all of the HVI Engine Sequences simultaneously. It exposes the add statement methods such as `add_sync_while()`. All of the Statements added are collected in the `sync_sequence` object.

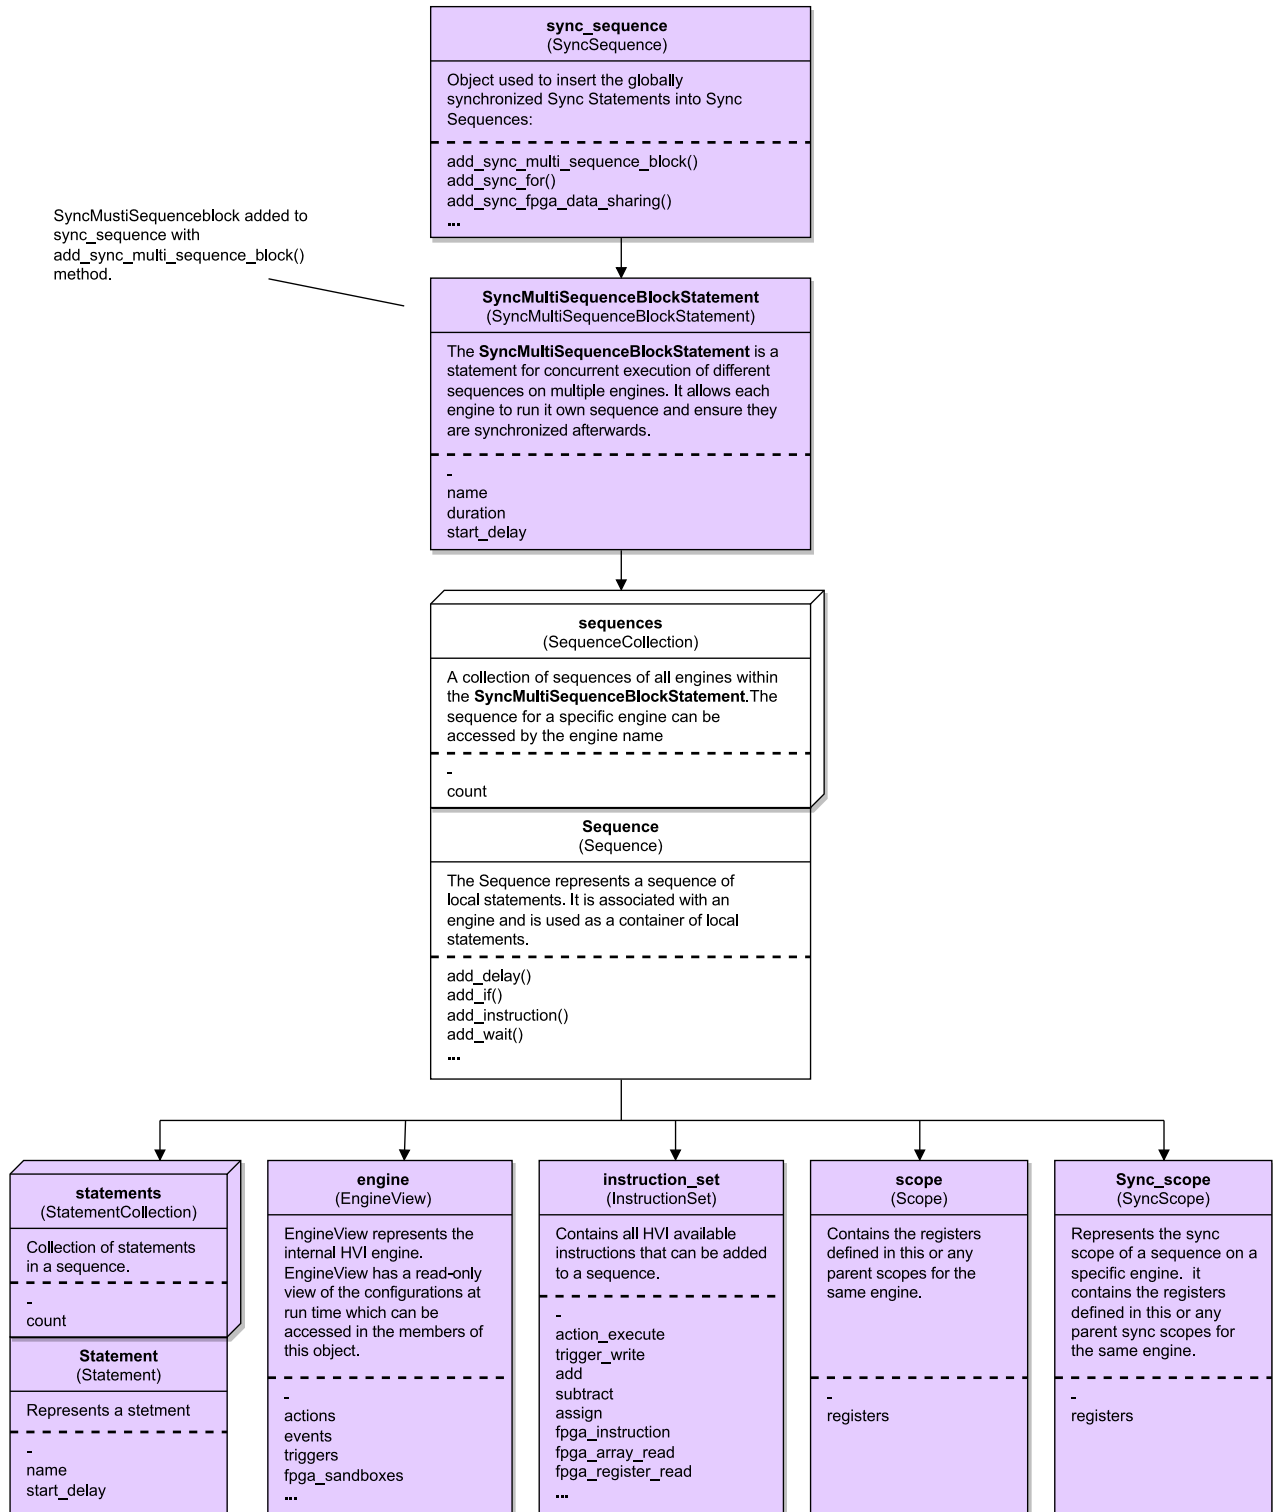
Synchronization and timing information are added within each `sync_sequence` so that all Sequences across the HVI are coordinated precisely. The `SyncMultiSequenceBlockStatement` exposes HVI Instructions and Local Flow-Control Statements that are sent by the `Sequence` object. The other Sync Statements are all synchronized across all the Sequences in the HVI.

An HVI Sequence contains the list of Local Statements to be executed by the HVI Engine.

The `Sequence` object exposes the add Statement methods such as `add_while()`. You add Local Flow-Control Statements such as If or While directly into the Sequence. All HVI Instructions are added using `add_instruction()`. The list of available Statements for the `add_instruction()` Statement is shown in [Local Statements](#).

The `Sequence` object stores a collection of all the Statements added to it, along with the scope Variables and registers needed for this Sequence. These are sent to a `SyncMultiSequenceBlockStatement`. This object exposes access and execution of Local Sequences.

The following diagram shows the `SyncMultiSequenceBlockStatement` and its relationship to `sync_sequences` and `sequences`:



Sequence Statements

Sequencing statements exposed in the TSE API are divided into two types:

Sync Statements

Sync Statements are the building blocks used to program Sync Sequences. The following types of Sync Statement are available:

- Sync While.
- Sync Multi-Sequence Block.
- Sync Register-Sharing.
- Sync FPGA Data-Sharing.

For a description of each Sync Statement with examples and a description of the Statement execution, see [Sync Statements](#).

Local Statements

Local Statements are programmed on Engines in individual instruments. They are always programmed within a Sync MultiSequenceBlock Statement. Local Statements are in the form of Local-Flow-control, and HVI-Native or Instrument-Specific instructions:

Local Flow-Control Statements

- Local If
- Local If with matched branches
- Local While
- Local Wait-for-Event
- Local Wait-for-time
- Local Delay

HVI-native Instructions

- Action Execute
- Register Increment
- Trigger Write
- Register Assign
- FPGA-Register Read
- FPGA-Register Write

- FPGA-Memory Map Read
- FPGA-Memory Map Write
- FPGA-Instruction Statement

Instrument-specific Instructions

- For Instrument-Specific HVI Instructions, see your instrument documentation.

For a description of each HVI-Native Instruction with examples and a description of the Statement execution, see [Local Statements](#).

Sync Statements

This section describes the HVI Statements in the TSE API that you use to program HVI Sequences. The functions of each Statement are explained in detail along with a corresponding HVI diagram. Python code examples are provided showing how to program the Statements with the HVI Python API.

Sync Statements are the building blocks used to program Sync Sequences. The following types of Sync Statement are available:

- [Sync While](#)
- [Sync For](#)
- [Sync Multi-Sequence Block](#)
- [Sync Register-Sharing](#)
- [Sync FPGA Data-Sharing](#)

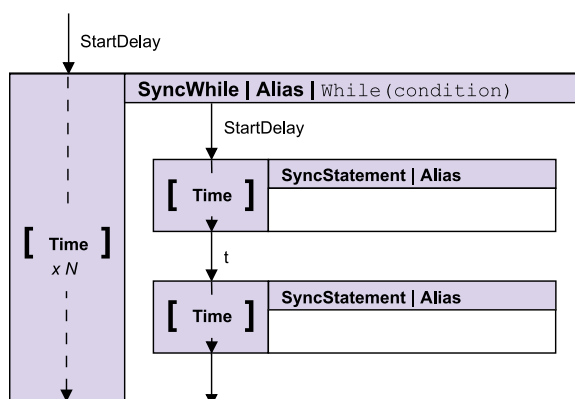
Sync While

A Sync While enables you to synchronously execute multiple Local Sequences while a condition you specify is met. The Sync While condition is evaluated each time at the beginning of the Statement execution. If the condition is true, an iteration of the Sync While Statement is executed. If the condition is false, the HVI execution jumps to the Statement following the Sync While.

You can add other Sync Statements inside a Sync while. To define Local Sequences within the Sync While, you must use a Sync Multi-Sequence Block.

You can set a duration for this Statement. For more information see [Chapter 10: HVI Time Management and Latency](#).

A Sync while that contains a pair of Sync Statements is shown in the following diagram:



If you are using a Sync while Statement across multiple Engines, during its execution, one of the Engines is set to the role of *Leader* and the remaining Engines have the role of *Follower*:

Leader

The condition of the Sync While Statement is evaluated in this Engine and the result is propagated to the other Engines through hardware resources, for example, PXI triggers in a PXI platform.

Follower

A *Follower* Engine monitors the result of the condition and acts on it, following the *Leader*.


The condition expression assigned to the Sync While must use resources that belong to the same HVI Engine. The *Leader* Engine of the Sync While is selected automatically by the HVI compiler from the condition expression.

The following code example shows how to add a Sync While Statement and access the Sync Sequence in the Sync while.

```
# Configure Sync While Condition
sync_while_condition = keysight_tse.Condition.register_comparison(reg, keysight_
tse.ComparisonOperator.GREATER_THAN, 10)
#
# Add Sync While to a sync-Sequence
sync_while = my_sync_seq.add_sync_while("sync_while", 10, sync_while_condition)
#
# Access the sync Sequence in the Sync-While and add Sync-Statements inside
sync_block = sync_while.sync_sequence.add_sync_multi_sequence_block("exec_block",10)
```

Sync For

The Sync-For statement enables you to execute synchronized real-time looping, across multiple instruments over a range of values you pre-define. The Sync-For runs with a Sync-Register that enables significantly faster real-time looping than using Sync-While with Local HVI registers. By using Sync Registers, the Sync-For avoids having to use triggered-synchronization points. This results in a multi-instrument execution that is as fast a single engine Local-For. See [HVI Registers and Scopes](#) for more information about Sync-Registers.



The Sync-For statement is created with a Sync-Iterator, the Sync-Iterator takes a Sync Register and the range and step for the sweep:

Sync-Iterator Parameter	Type	data range
syncRegister	SyncRegister	NA
beginValue	std::int64_t	signed range $[-2^{\text{bits}-1}, 2^{\text{bits}-1} - 1]$, where bits is the number of bits the register can hold
endValue	std::int64_t	signed range $[-2^{\text{bits}-1}, 2^{\text{bits}-1} - 1]$, where bits is the number of bits the register can hold
stepValue	std::int64_t	signed range $[-2^{\text{bits}-1}, 2^{\text{bits}-1} - 1]$, where bits is the number of bits the register can hold

The following snippet demonstrates how to use the Sync-For statement:

Example API usage

```
# Get the Sync registers collection
sync_register_collection = my_sequencer.sync_sequence.sync_scope.registers
# Create Sync register
my_sync_register = sync_register_collection.add("my_sync_register", keysight_
```

```

hvi.RegisterSize.SHORT)
# Create Sync iterator with start value 1, end value 10, step 1.
sync_iterator = keysight_hvi.sync_iterator( my_sync_register , 1, 10, 1)
# Create Sync-For statement
sync_for_statement = sequencer.sync_sequence.add_sync_for("sync_for", 100, sync_iterator)
# Program Sync Sequence inside Sync-For as any other Sync Sequence
sync_for_sequence = sync_for_statement.sync_sequence
    
```

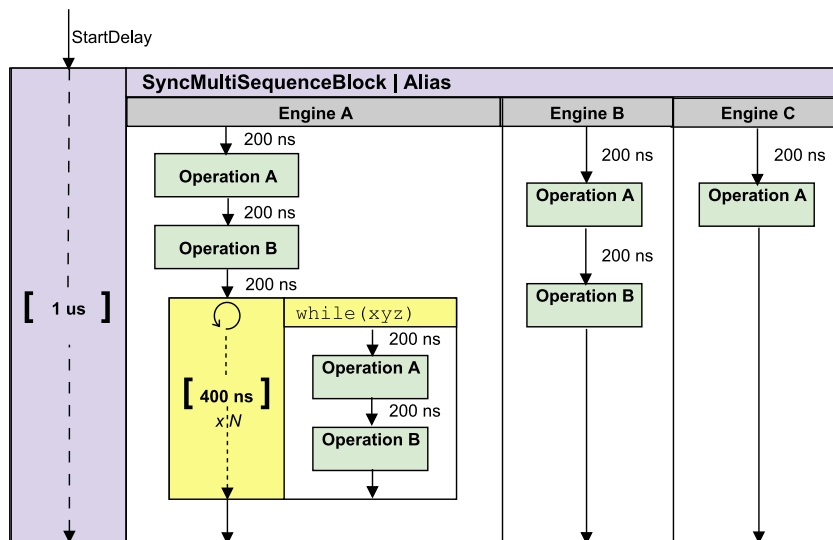
Sync Multi-Sequence Block

Sync Multi-Sequence Blocks are a type of Sync Statement that contains a set of local Sequences. It serves as a container and boundary between sections, where each local Sequence executes on an individual Engine within a specific instrument.

The Sync Multi-Sequence Block enables you to program each Engine to do specific operations and run them on each Engine concurrently. The Sync Multi-Sequence Block synchronizes all the Engines so that all of the contained Local Sequences start at exactly the same time. You can define which Local Statements each Engine is going to execute, and the exact time each Local Statement starts to execute compared to the previous one. The Sync Multi-Sequence Block ensures that all the Engines complete their Sequences at the same time so the sync Sequence remains synchronous afterwards.

You can set a duration for this Statement. For more information see [Chapter 10: HVI Time Management and Latency](#).

The following diagram shows a Sync Multi-Sequence Block that contains three Local Sequences:



The following code snippet shows a Sync Multi-Sequence Block being added with the call `add_sync_multi_sequence_block()`, a Local Sequence is then obtained and an instruction added to it:


```
# Add Sync Multi-Sequence Block
sync_block = keysight_tse.sync_sequence.add_sync_multi_sequence_block("TriggerAWGs")
#
# Add instruction to a Local Sequence in the block
sequence = sync_block.sequences["Main Engine"]
inst = sequence.add_instruction("Add Instruction", 10, seq.instruction_set.add_instruction.id)
```

Sync Register-Sharing

Sync Register-Sharing enables you to share data from a source register to a destination register in any Engine in your HVI. Specifically, you share the contents of *n* adjacent bits from a source register to a destination register.

Sync Register-Sharing is defined with a `SyncRegisterSharingStatement`.

You can set a duration for this Statement. For more information see [Chapter 10: HVI Time Management and Latency](#).

In the following code example, Sync Register-Sharing is used to share the content of the digitizer register `feedback` and write into the AWG register `wfm_id`:

```
# Digitizer registers
feedback = keysight_tse.sync_sequence.scopes["Dig Engine"].registers.add("Feedback Reg",
keysight_tse.RegisterSize.SHORT)
feedback.initial_value = 0
#
# AWG registers
wfm_id = keysight_tse.sync_sequence.scopes["AWG Engine"].registers.add("Wfm ID", keysight_
tse.RegisterSize.SHORT)
wfm_id.initial_value = 0
#
# Add Sync Register Sharing
bits_to_share = 3
sync_while_2.sync_sequence.add_sync_register_sharing("Share feedback->wfm_id", 10, steps, wfm_
id, bits_to_share)
```

Sync FPGA Data-Sharing

This Statement enables you to share data between FPGA-Sandboxes in different instruments. It uses Fast Data Sharing technology to share data.

You set up each transfer with an `FpgaDataSharingTransaction` which you add to a collection in the `SyncFPGADataSharing` Statement. Each `FpgaDataSharingTransaction` enables the sharing of a specified number of bits. The number of bits to be shared must be a multiple of 4.

The start and the end of the Statement are time synchronized. The `SyncFPGADataSharing` Statement blocks the execution in all instruments until the last instrument has received the last 4 bits of data. You can also set a duration for this Statement. For more information see [Chapter 10: HVI Time Management and Latency](#).

The following table shows related objects :

Name	Description
<code>FpgaDataSharingTransaction</code>	A Sync FPGA Data-Sharing transaction enables you to share a specified number of bits from one FPGA-Sandbox to one or more FPGA-Sandboxes in different instruments. You add a transaction to a <code>SyncFpgaDataSharingStatement</code> using one of the <code>FpgaDataSharingTransactionCollection.add()</code> methods.
<code>FpgaDataSharingTransactionCollection</code>	A list of Sync FPGA Data-Sharing transactions.
<code>FdsPortAddress</code>	Represents the source or destination address of a Fast Data Sharing transaction. It consists of the FDS port of an instrument that shall be used to transmit or receive data, as well as the address of the IP connected to that FDS port, where the data will be read from / written to.

Use the following method to add a `SyncFpgaDataSharingStatement` Statement into a Sequence:

New Method	Description
<code>SyncSequence.add_sync_fpga_data_sharing</code>	Adds a <code>SyncFpgaDataSharingStatement</code>

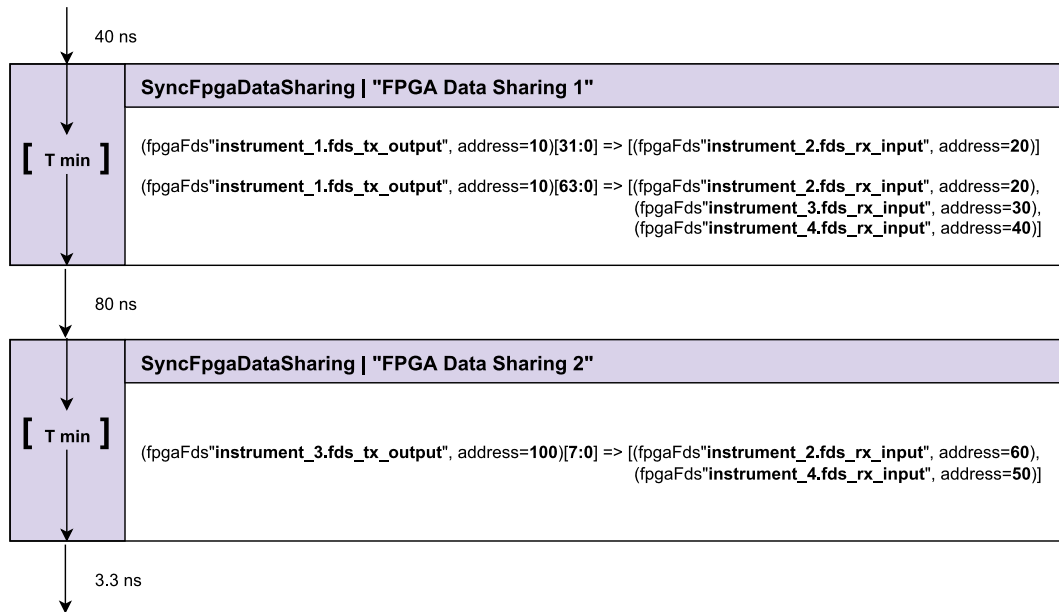
The following code snippet shows examples of `SyncFpgaDataSharingStatement` Statements:

```
# The name for each sandbox provided below should match exactly the name that each instrument
has given to each FPGA-Sandbox
# For simplicity, we assume that all instrument use the same name for their sandboxes
instrument_1_sandbox_name = "SampleSandbox"
instrument_2_sandbox_name = "SampleSandbox"
instrument_3_sandbox_name = "SampleSandbox"
instrument_4_sandbox_name = "SampleSandbox"
#
# The name for each FDS port provided below should match exactly the name that was given by each
user to the FDS port in the FPGA-Sandbox
# For simplicity, we assume that all Tx ports share a common name
source_port_name = 'fds_tx_output'
# For simplicity, we assume that all Rx ports share a common name
dst_port_name = 'fds_rx_input'
#
instrument_1_fds_ports = sequencer.sync_sequence.engines["instrument_1"].fpga_sandboxes
[instrument_1_sandbox_name].fds_ports
instrument_2_fds_ports = sequencer.sync_sequence.engines["instrument_2"].fpga_sandboxes
[instrument_2_sandbox_name].fds_ports
instrument_3_fds_ports = sequencer.sync_sequence.engines["instrument_3"].fpga_sandboxes
[instrument_3_sandbox_name].fds_ports
instrument_4_fds_ports = sequencer.sync_sequence.engines["instrument_4"].fpga_sandboxes
[instrument_4_sandbox_name].fds_ports

# Add an FPGA Data-Sharing Statement
name = 'Sync FPGA Data-Sharing 1'
start_delay = 40
fpga_data_sharing_st = sequencer.sync_sequence.add_sync_fpga_data_sharing(name, start_delay)
```

```
# Add transaction 1: single destination
source_address = 10
source_port = instrument_1_fds_ports[source_port_name]
source = keysight_tse.FdsPortAddress(source_port, source_address)
dst1_address = 20
dst1_port = instrument_2_fds_ports[dst_port_name]
dst1 = keysight_tse.FdsPortAddress(dst1_port, dst1_address)
num_bits_to_share = 32
fpga_data_sharing_st.transactions.add(source, dst1, num_bits_to_share)
#
# Add transaction 2: multiple destinations
dst2 = keysight_tse.FdsPortAddress(instrument_3_fds_ports[dst_port_name], 30)
dst3 = keysight_tse.FdsPortAddress(instrument_4_fds_ports[dst_port_name], 40)
dests = [dst1, dst2, dst3]
num_bits_to_share = 64
fpga_data_sharing_st.transactions.add(source, dests, num_bits_to_share)
#
# Add another FPGA Data-Sharing Statement
name = 'Sync FPGA Data-Sharing 2'
start_delay = 80
fpga_data_sharing_st = sequencer.sync_sequence.add_sync_fpga_data_sharing(name, start_delay)
source = keysight_tse.FdsPortAddress(instrument_3_fds_ports[source_port_name], 100)
dst1 = keysight_tse.FdsPortAddress(instrument_2_fds_ports[dst_port_name], 60)
dst2 = keysight_tse.FdsPortAddress(instrument_4_fds_ports[dst_port_name], 50)
dests = [dst1, dst2]
num_bits_to_share = 8
fpga_data_sharing_st.transactions.add(source, dests, num_bits_to_share)
```

The following diagram shows a pair of `SyncFpgaDataSharing` Statements:



Sync FPGA Data-Sharing Errors

The `SyncFpgaDataSharing` Statement uses FDS for transfers. It has the following error conditions:

1. The number of bits to share is not a positive multiple of 4.
2. A transaction has multiple FDS ports from the same Engine.
3. A transaction contains a singular FDS port more than once. The same FDS port cannot be a source and a destination in the same transaction.

Local Statements

This section describes the Local Statements available in the TSE API, that you use to program HVI Sequences. It contains the following sections:

- [Programming Local Sequences](#)
- [HVI Instructions](#)
- [Instrument-Specific HVI Instructions](#)
- [HVI Instructions](#)
 - [Action Execute](#)
 - [Register Increment](#)
 - [Trigger Write](#)
 - [Register Assign](#)
 - [FPGA-Register Read](#)
 - [FPGA-Register Write](#)
 - [FPGA-Memory Map Read](#)
 - [FPGA-Memory Map Write](#)
 - [FPGA-Instruction](#)
- [Local Flow-Control Statements](#)
 - [Local If Statement](#)
 - [Local If with matched branches](#)
 - [Local While](#)
 - [Local Wait-For-Event](#)
 - [Local Wait-For-Time](#)
 - [Local Delay](#)

The functions of each Statement are explained in detail together with Python code examples showing how to program the Statements with the HVI Python API.

Programming Local Sequences

You program Local Sequences within a Sync Multi-Sequence Block or a Local Flow-Control Statement (Local While or Local If). The following code shows an example of a Local Sequence programmed within a Sync Multi-Sequence Block.

```
# Add Statements to each Local Sequence within the Sync Multi-Sequence Block
# HVI Local Sequence collection is automatically created from the
# user-defined HVI Engine Collection
# Each HVI Local Sequence can be retrieved using the name of the corresponding HVI Engine
sequence = sync_block.sequences[engine_name]
#
# Add FP Trigger ON to all instruments
instr_trigger_write = sequence.instruction_set.trigger_write
instr_trigger_ON = sequence.add_instruction("FP Trigger ON", 10, instr_trigger_write.id)
instr_trigger_ON.set_parameter(instr_trigger_write.trigger, fp_trigger)
instr_trigger_ON.set_parameter(instr_trigger_write.sync_mode.id, instr_trigger_write.sync_
mode.immediate)
instr_trigger_ON.set_parameter(instr_trigger_write.value.id, instr_trigger_write.value.on)
```

HVI Instructions

HVI Instructions are operations that can be executed by the instrument hardware within an HVI Sequence. There are two types of HVI instructions:

- Instrument-Specific HVI Instructions.
- HVI Native instructions.

Instrument-Specific HVI Instructions

Instrument-Specific HVI Instructions are specific to individual instruments. They are defined by the instrument add-on API and exposed in each instrument driver as instrument specific HVI definitions. Instrument-Specific HVI Instructions can change instrument settings such as amplitude, frequency, or trigger an instrument function such as output a waveform or trigger a data acquisition. For example, the M3xxxA documentation describes all the Instrument-Specific HVI instructions available for each of the M3xxxA PXI instruments.

The following code is an example of using the `awgQueueWaveform` custom instruction that is part of the HVI instruction set of the Keysight M320xA AWG instrument. This example shows how to add an Instrument-Specific HVI Instruction to a Local Sequence using the `add_instruction()` API method and also how to set the instruction parameters using the `set_parameter()` method:

```
# Retrieve engine sequence:
seq = sync_block.sequences["engine_0"]
#
# Add and program AWG Queue Waveform instruction:
instr_queue_wfm = instrument.hvi.instruction_set.queue_waveform
instruction0 = seq.add_instruction("awgQueueWaveform", 10, .id)
#
# Set instruction parameters:
instruction0.set_parameter(instr_queue_wfm.waveform_number.id, seq.registers
[waveformNumberRegisterName])
instruction0.set_parameter(instr_queue_wfm.channel.id, nAWG)
instruction0.set_parameter(instr_queue_wfm.trigger_mode.id, keysightSD1.SD_
TriggerModes.SWHVITRIG)
instruction0.set_parameter(instr_queue_wfm.start_delay.id, startDelay)
instruction0.set_parameter(instr_queue_wfm.cycles.id, nCycles)
instruction0.set_parameter(instr_queue_wfm.prescaler.id, prescaler)
```


HVI Instructions

HVI Instructions are HVI Native Instructions that are available on all Keysight instruments. They are general purpose and instrument independent. They include HVI Instructions and Local Flow-Control Statements. The HVI Native Instructions and parameters are defined in the interface `hvi.instruction_set`.

The set of HVI Instructions include:

- Action Execute.
- Register Increment.
- Trigger Write.
- Register Assign.
- FPGA-Register Read.
- FPGA-Register Write.
- FPGA-Memory Map Read.
- FPGA-Memory Map Write.
- FPGA-Instruction.

Action Execute

To add Actions to an HVI Sequence, you must add them with the `add()` method of the `ActionCollection` property of the HVI Engine.

Once the required Actions are added to the list of the HVI Engine Actions for the instruments, an instruction to execute them can be added to the instrument's Sequence using the property `action_execute`. One or multiple Actions can be executed at the same time within the same Action Execute instruction.

The following code example shows an Action Execute instruction being used to initiate an AWG Trigger:

```
# Previously defined Actions to be executed within the experiment
awg_trigger_12 = [hvi.sync_sequence.engines["engine_name"].actions["previously_defined_action_1"], hvi.sync_sequence.engines["engine_name"].actions["previously_defined_action_2"]]
#
# AWG Trigger CH1, CH2 - Generates first pulse
sequence = sync_block_2.sequences["engine_name"]
inst_awg_trigger = sequence.add_instruction("AwgTrigger(CH1, CH2)", 10, sequence.instruction_set.action_execute.id)
inst_awg_trigger.set_parameter(hvi.instruction_set.action_execute.action.id, awg_trigger_12)
```

Register Increment

You can implement a register increment within a Sequence with the property `add`. The same instruction can be used to add registers and constant values (operands) and put the result in another register (result). To increment the register, it must have been added previously to the scope of the corresponding HVI Engine.

The following code shows an example of a register increment:

```
# Previously defined
counter = sync_sequence.scopes["AWG Engine"].registers.add("Counter Reg", keysight_
tse.RegisterSize.SHORT)
#
# Increment counter register
instruction = awg_sequence.add_instruction("Increment counter", 10, awg_sequence.instruction_
set.add.id)
instruction.set_parameter(awg_sequence.instruction_set.add.destination.id, counter)
instruction.set_parameter(awg_sequence.instruction_set.add.left_operand.id, counter)
instruction.set_parameter(awg_sequence.instruction_set.add.right_operand.id, 1)
```

Trigger Write

The following code example shows a Trigger write instruction used to implement a front panel trigger ON/OFF instruction. The instruction is added to the Sequence with the method `add_instruction()`. Instruction parameters are set using the API method `set_parameter()`. All HVI Native Instructions and parameters are defined in the `hvi.instruction_set` interface.

```
# Add FP Trigger ON to all instruments
sequence = sync_block.sequences[engine_name]
instr_trigger_write = sequence.instruction_set.trigger_write
instr_trigger_ON = sequence.add_instruction("FP Trigger ON", 10, instr_trigger_write.id)
instr_trigger_ON.set_parameter(instr_trigger_write.trigger, fp_trigger)
instr_trigger_ON.set_parameter(instr_trigger_write.sync_mode.id, instr_trigger_write.sync_
mode.immediate)
instr_trigger_ON.set_parameter(instr_trigger_write.value.id, instr_trigger_write.value.on)
```

Register Assign

A register assign Statement can be used to initialize a register to an initial value using the `instruction_set` property `assign`. The same instruction can be used to assign a register value (source) to another register (destination). Each register can also be initialized outside an HVI Sequence, before its execution, by using the property `register.initial_value`.

The following code shows an example of Register Assign:

```
# Previously defined registers
wfm_id = hvi.sync_sequence.scopes["AWG Engine"].registers.add("Wfm ID", keysight_
tse.RegisterSize.SHORT)
#
# Initialize Waveform ID
seq = sync_block_1.sequences["AWG Engine"]
instruction = seq.add_instruction("Initialize Wfm ID", 10, seq.instruction_set.assign.id)
instruction.set_parameter(seq.instruction_set.assign.destination.id, wfm_id)
instruction.set_parameter(seq.instruction_set.assign.source.id, 0)
```

FPGA-Register Read

The instruction `fpga_register_read` is an HVI Native instruction that enables you read from an HVI FPGA-Register. The value read from the HVI FPGA-Register is written to a destination HVI FPGA-Register.

The following code example shows an FPGA-Register Read instruction:

```
# Read FPGA Register into an HVI Register used in the HVI Sequence
sequence = sync_block_1.sequences["engine_name"]
hvi_register = hvi.sync_sequence.scopes["engine_name"].registers["my_register"]
fpga_register = hvi.sync_sequence.engines["engine_name"].fpga_sandboxes["sandbox_name"].hvi_
registers["sandbox_register"]
readFpgaReg0 = sequence.add_instruction("Read FPGA Register_Bank_TseAction4Cnt", 10,
sequence.instruction_set.fpga_register_read.id)
readFpgaReg0.set_parameter(sequence.instruction_set.fpga_register_read.destination.id, hvi_
register)
readFpgaReg0.set_parameter(sequence.instruction_set.fpga_register_read.fpga_register.id, fpga_
register)
```

FPGA-Register Write

The instruction `fpga_register_write` is an HVI Native instruction that enables you to write an HVI FPGA-Register placed in an FPGA-Sandbox. The value to be written to the HVI FPGA-Register is taken from an HVI FPGA-Register or from a literal.

The following code example shows an FPGA-Register Write instruction:

```
# Write FPGA Register from an an HVI Register used in the HVI Sequence
hvi_register = hvi.sync_sequence.scopes["engine_name"].registers["my_register"]
fpga_register = hvi.sync_sequence.engines["engine_name"].fpga_sandboxes["sandbox_name"].hvi_registers["sandbox_register"]
seq = sync_block_1.sequences["engine_name"]
writeFpgaReg0 = seq.add_instruction("Write FPGA Register_Bank_TsePxiTrigOut", 50,
hvi.instruction_set.fpga_register_write.id)
writeFpgaReg0.set_parameter(seq.instruction_set.fpga_register_write.fpga_register.id, fpga_register)
writeFpgaReg0.set_parameter(seq.instruction_set.fpga_register_write.value.id, hvi_register)
```

FPGA-Memory Map Read

The instruction `fpga_array_read` is an HVI Native Instruction that enables you to read from an HVI FPGA-Memory Map. The value read from the HVI FPGA-Memory Map is written to a destination HVI register.

The following code example shows an FPGA-Memory Map read instruction:

```
# Register and Memory map
register = sync_sequence.scopes["engine_name"].registers["my_register"]
hvi_memory_map = sync_sequence.engines["engine_name"].fpga_sandboxes["sandbox_name"].hvi_memory_maps["memory_map_name"]
# Read Memory Map
seq = sync_block_1.sequences["engine_name"]
readMemoryMap = sync_block_1.sequences["engine_name"].add_instruction("Read FPGA-Memory Map",
20, hvi.instruction_set.fpga_array_read.id)
readMemoryMap.set_parameter(seq.instruction_set.fpga_array_read.fpga_memory_map.id, hvi_memory_map)
readMemoryMap.set_parameter(seq.instruction_set.fpga_array_read.destination.id, register)
readMemoryMap.set_parameter(seq.instruction_set.fpga_array_read.fpga_memory_map_offset.id, 0)
```

FPGA-Memory Map Write

The instruction `fpga_array_write` is an HVI Native Instruction that enables you to write to an HVI FPGA-Memory Map that is located in an FPGA-Sandbox. The value to be written to the HVI FPGA-Memory Map is taken from an HVI register or from a literal.

The following code example shows an FPGA-Memory Map write instruction:

```
# Register and Memory map
register = sync_sequence.scopes["engine_name"].registers["my_register"]
hvi_memory_map = sync_sequence.engines["engine_name"].fpga_sandboxes["sandbox_name"].hvi_memory_
maps["memory_map_name"]
# Write Memory Map
seq = sync_block_1.sequences["engine_name"]
writeMemoryMap = sync_block_1.sequences["engine_name"].add_instruction("Write FPGA-Memory Map",
10, seq.instruction_set.fpga_array_write.id)
writeMemoryMap.set_parameter(seq.instruction_set.fpga_array_write.fpga_memory_map.id, hvi_
memory_map)
writeMemoryMap.set_parameter(seq.instruction_set.fpga_array_write.value.id, register)
writeMemoryMap.set_parameter(seq.instruction_set.fpga_array_write.fpga_memory_map_offset.id, 0)
```

FPGA-Instruction

The FPGA-Instruction Statement enables you to issue commands to custom FPGA-Sandbox logic from within HVI Sequences.

You can add the FPGA-Instruction Statement in an HVI Sequence. When the HVI Sequence is running, the HVI Engine executes the FPGA-Instruction Statement and sends the command and data parameters to a parser and your logic in the FPGA-Sandbox. Your logic reads the parameter data and executes the command as indicated.

This instruction can only be used successfully on instruments that support it. For more information, see your instrument documentation.

FPGA-Instruction Statement Parameters

The `FpgaInstruction` Statement has the following parameters:

Parameter	Description	Size	Notes
Port Number	Selects the port in the FPGA-Sandbox	-	Number of available ports defined by the instrument
Command ID	An identifier for a command you have implemented in custom logic	16 bits	
Data A	The data to send to the IP in the sandbox	40 bits	Supports registers <ul style="list-style-type: none"> • If the source register is a short (32 bits), the top 8 (most significant) bits are set to 0. • If the source register is a long (48 bits), the top 8 (most significant) bits are truncated.
Apply ¹	A 1 bit field which determines if the command is applied immediately or is set up for later execution	1 bit	<ul style="list-style-type: none"> • 0 = Set up now, apply the instruction later. • 1 = Apply instruction immediately (this is the default).

¹ To apply the instruction after setup, you initiate it with the next instruction with Apply set to 1. You can set a number of instructions each with Apply = 0, then the next instruction with Apply = 1 shall trigger all of them.

The Apply=0 followed by an Apply=1 provides a set up now, and apply later option. This enables you to set up the command and then delay the execution so it happens at a specifically timed interval. This also enables you to set up a number of commands and then have them execute simultaneously.

The following code shows an example of an FPGA-Instruction:

```
# Set up local Sequence
fpga_inst = local_sequence.instruction_set.fpga_instruction
instruction = local_sequence.add_instruction('fpgaInstruction', 10, fpga_inst.id)
#
port_number = 2
data_a = 1234
command_id = 5
apply = 1
#
instruction.set_parameter(fpga_inst.port_number.id, port_number)
instruction.set_parameter(fpga_inst.data_a.id, data_a)
instruction.set_parameter(fpga_inst.command_id.id, command_id)
instruction.set_parameter(fpga_inst.apply.id, apply)
```

Local Flow-Control Statements

Local Flow-Control Statements are HVI Native Instructions that are available on all Keysight instruments. Local Flow-Control Statements execute within Local Sequences. These include wait Statements, loops such as Local While, and conditional execution like Local If. Local Flow-Control Statements are depicted with a yellow box in the HVI diagrams in this User Manual.

Local Flow-Control Statements include:

Local If

Selects and executes a different possible Local Sequence according to the value of a defined condition (if / elseif / else).

Local While

Executes the same Sequence in a loop while the condition is met.

Local Wait-For-Event

Causes the Sequence to stop and wait for a condition to evaluate true. Once the condition is true, for example, when the selected Event occurs, the next instruction is executed. In future releases, this will be extended to more complex conditions.

Local Wait-For-Time

Causes the Sequence to wait for a certain time specified in an HVI register. Once the time has elapsed, the Sequence will continue.

Local Delay

Delays the Sequence for a time you specify. The delay is specified in nanoseconds.

All Local Flow-Control Statements except wait Statements, include one or more Local Sequences. For instance, Local While Statements have a single Sequence and the Local If Statement can have multiple Sequences. These Statements have the following common characteristics:

- Sequences in Local Flow-Control Statements can contain any Local Statement including Local Flow-Control Statements.
- Only Local Statements can be added inside Local Sequences and consequently inside Local Flow-Control Statements. You cannot add Sync Statements inside Local Flow-Control Statements.


Local If Statement

The Local If Statement is a flow-control Statement that conditionally executes one of a set of different possible Local Sequences (if/ elseif / else) depending on the value of predefined conditions.

The conditions are evaluated in the order they are inserted. The possible sequences are:

- At least one Sequence that is conditionally executed. This is the main **If** branch.
- Optional conditional Sequences where their conditions are evaluated in order. The first Sequence with a true condition is executed if the conditions in previous branches evaluated false. These are the **Elseif** branches.
- If more than one Elseif condition evaluates to true, only the first is executed.
- One optional Else Sequence, which is executed if all above previous conditions evaluate to false. This is the **Else** branch.

The following diagram shows a Local If Flow-Control Statement:



The `add_if()` method of a Sequence enables you to add an If-Elseif-Else construct within the main HVI Sequence of any HVI Engine. The Local If Statement contains one If branch, zero or more Elseif branches and one Else branch. The instructions and Statements contained in each If or Else branch are executed if the condition of each branch is met.

You can program the branch Sequences with the same methods and properties used to program the main HVI Sequence, using the properties `if_branch`, `else_branch` and `else_if_branches`. You define the condition of each branch with a `ConditionalExpression` object, The conditions are stored in registers.

You can set a duration for this Statement. For more information see [Chapter 10: HVI Time Management and Latency](#).

The following code is an example of a Local If Statement:

```
# Configure IF condition
if_condition = keysight_tse.Condition.register_comparison(reg, keysight_
tse.ComparisonOperator.SMALLER_THAN, 10);
#
# Set flag that enables to match the execution time of all the IF branches
enable_matching_branches = True
if_statement = my_sync_multi_seq_block.add_if("MyIfBlock", 10, if_condition, enable_matching_
branches)
#
# Program IF branch
if_sequence = if_statement.if_branch.sequence
#
# Add Statements in if-sequence
instruction = ifSequence.add_instruction("ExecuteAction0", 10, if_sequence.instruction_
set.action_execute.id)
instruction.set_parameter(...) ...
#
# Program Else-If branches
# Else-If Condition
else_if_condition_1 = keysight_tse.Condition.register_comparison(reg, keysight_
tse.ComparisonOperator.SMALLER_THAN, 15)
else_if_branch_1 = if_statement.else_if_branches.add(else_if_condition_1)
#
# Program Else-If branch
else_if_sequence_1 = else_if_branch_1.sequence
#
# Add Statements in Else-If-sequence
instruction = else_if_sequence_1.add_instruction("SetFrequency", 10, module.HVI.instruction_
set.set_frequency.id)
instruction.set_parameter(...) ...
#
# Eventually add more Else-If-branches
else_if_condition_2 = ... else_if_branch_1 = ... ...
#
# Else-branch
# Program Else branch
else_sequence = else_branch.sequence
#
# Add Statements in Else-Sequence
instruction = else_sequence.add_instruction(...) ...
```

Local If with matched branches

Unlike other flow-control options, the Local If Statements can have different execution paths, each with different times. The matched branches option enables you to control how the HVI deals with them.

Enabling matched branches ensures the HVI synchronizes the times of the branches, so they are the same. The shorter branches get an additional delay added when they are finished so that the durations of all the branches are equal. If the matched branches option is not enabled, the branches can end at different times, that is, they are *de-synchronized*.

The following code shows how to enable matched branches:


```
# Set flag that enables to match the execution time of all the IF branches
#
enable_matching_branches = True
if_statement = my_sync_multi_seq_block.add_if("MyIfBlock", 10, if_condition, enable_matching_
branches)
```

Local While

The Local While Statement executes a same Sequence in a loop while a condition is met. The value for the condition is stored in a register.

You can set a duration for this Statement. For more information see [Chapter 10: HVI Time Management and Latency](#).

The following diagram shows a Local While:



The following code is an example of a Local While Statement:

```
# Configure while condition
while_condition = keysight_tse.Condition.register_comparison(reg, keysight_
tse.ComparisonOperator.NOT_EQUAL, 1)
#
# Add WHILE Sequence within the Sequence of "engine_0"
seq = sync_block.sequences["engine_0"]
while_loop = seq.add_while("While Loop", 10, while_condition)
#
# Program local while Sequence
instruction = while_loop.sequence.add_instruction("Initialize Pulse Counter", 10,
seq.instruction_set.assign.id)
instruction.set_parameter(seq.instruction_set.add.destination.id, pulse_counter)
instruction.set_parameter(seq.instruction_set.add.source.id, 0)
```

Local Wait-For-Event

The Local Wait-For-Event Statement causes the HVI Sequence to stop and wait for a condition to evaluate true. Once the condition is true, for example the selected Event occurs, the next instruction is executed.

The Local Wait Statement is implemented with the Sequence method `add_wait()`. This Sequence blocking Statement sets an instrument to wait for a condition. The condition can be defined by a Trigger, an Event, or a combination of them using logical operators. You can only use one Event in the condition.

In the following example, the Local Wait is used to set a digitizer instrument to wait for an external front panel trigger. The Local Wait Statement is set to wait for a Trigger falling edge using the `.wait` mode `keysight_tse.WaitMode.TRANSITION` combined with a Trigger configuration as `ACTIVE_LOW`. The sync mode `keysight_tse.SyncMode.IMMEDIATE` sets the wait Event to let the execution continue immediately, that is, as soon as the trigger Event is received:

```
# Trigger resource to be used as a wait condition
fp_trigger_id = module_list[0].hvi.triggers.front_panel_1
fp_trigger = sync_sequence.engines[digitizer_engine_name].triggers.add(fp_trigger_id, "FP
Trigger")
#
# Trigger configuration
# NOTE: Trigger to be used as WaitEvent conditions must be configured
# as keysight_tse.Direction.INPUT
fp_trigger.configuration.direction = keysight_tse.Direction.INPUT
fp_trigger.configuration.drive_mode = keysight_tse.DriveMode.PUSH_PULL
fp_trigger.configuration.polarity = keysight_tse.TriggerPolarity.ACTIVE_LOW
fp_trigger.configuration.hw_routing_delay = 0
fp_trigger.configuration.trigger_mode = keysight_tse.TriggerMode.LEVEL
#
# Define the condition for the Local Wait Statement
wait_condition = keysight_tse.Condition.trigger(hvi.sync_sequence.engines[digitizer_engine_
name].triggers["FP Trigger"])
#
# Add a Wait For Event
wait_event = sync_block_1.sequences[digitizer_engine_name].add_wait("Wait for FP Trigger", 100,
wait_condition)
wait_event.set_mode(keysight_tse.WaitMode.TRANSITION, keysight_tse.SyncMode.IMMEDIATE)
```

For information about timing implications for wait for Event Statements, *Synchronization Points and Sync Sequence Start* in [Chapter 10: HVI Time Management and Latency](#).

Local Wait-For-Time

The Local Wait-For-Time Statement causes the Sequence to wait for a certain time specified in an HVI register. Once the time is elapsed, the Sequence continues.

The following code is an example of a Local Wait-For-Time Statement:

```
# Wait Time makes the HVI Sequence wait for an amount of time specified by
# a register (register 'tau' in this example)
#
waitTau = sync_block.sequences["digitizer_engine"].add_wait_time("WaitTau", 10, tau)
```

Local Delay

The Local Delay Statement delays the execution of a Local Sequence for a time you specify. The default unit is nanoseconds but the delay is specified in any unit of seconds. The delay is fixed and cannot be changed during HVI execution, so the delay value must be known at the time of creating the HVI Sequence.

The Local Delay Statement works in a similar way as the Start Delay Statement parameter. The difference is that the Start Delay can only be specified before the other Statements in a Sequence. The Local Delay Statement enables you to place a fixed delay at the end of Sync Multi-Sequence Block or a Flow-Control Statement.

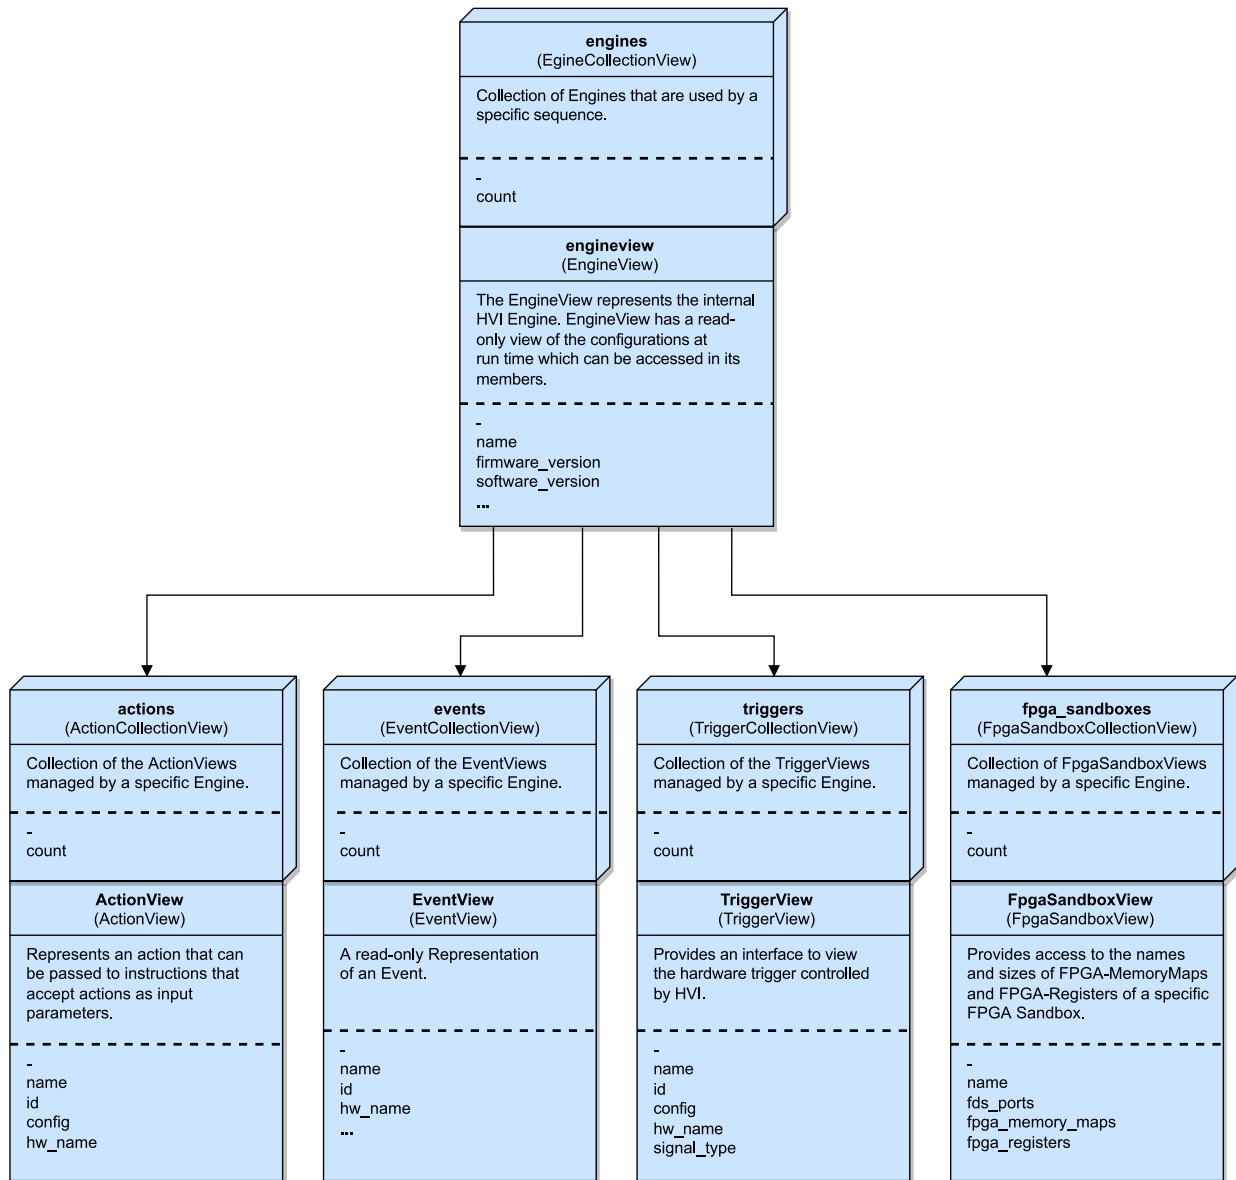
If you require a Variable delay that can be changed during HVI execution, use the Local Wait-For-Time Statement.

The following code shows an example of a Local Delay Statement:

```
# Delay makes the HVI Sequence wait for an amount of time specified by a constant
#
wait = sync_block.sequences["digitizer_engine"].add_delay("Delay", 30)
```

EngineView Components

A number of runtime components are under the **EngineView**.



TriggerCollectionView

Collection of the TriggerViews managed by the Engine.

TriggerView

TriggerView provides an interface to view the hardware Trigger controlled by HVI.

ActionCollectionView

Collection of the Actions managed by the Engine.

ActionView

Represents an Action that can be passed to instructions that accept Actions as input parameters.

EventCollectionView

Collection of the Events managed by the Engine.

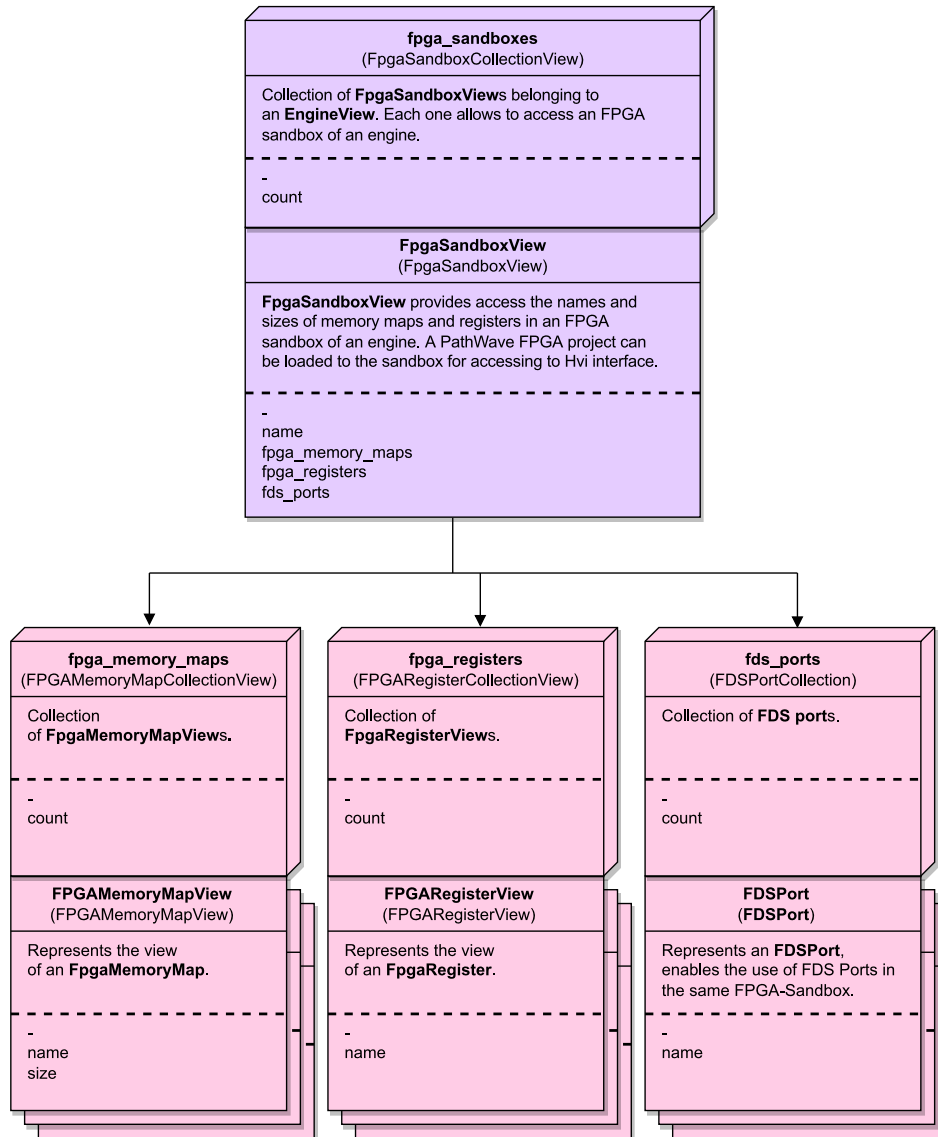
EventView

EventView represents a read-only configuration of an Event.

FPGA-Sandbox View

This section describes the FPGA-SandboxView properties and methods().

The following diagram shows the `FPGASandboxCollectionView`, its methods and properties:



FPGA-Sandboxes and FPGA-Memory Maps

`FpgaSandboxView` provides access to FPGA-Memory Maps by providing handles to FPGA-Registers and FPGA-Memory Maps that are defined in the FPGA memory. You can use `FpgaRegisterView` and `FpgaMemoryMapView` as parameters for HVI instructions for reading or writing FPGA memory. You must load the PathWave FPGA project as part of the SystemDefinition and then you can use the `FpgaSandboxView` in the Sequencer.

FpgaRegisterView

Once the sandbox project is loaded, you can access the contents of the FPGA-Sandbox and use them as parameters for HVI instructions. The FPGA write operation can accept registers and literal values as parameters. The following example shows writing FPGA-Registers:

```
# Retrieve FPGA-Register object from FPGA-Registers collection
# All sandbox object collections are populated when loading a bit file generated by PathWave
FPGA
fpga_register_view = sequencer.sync_sequence.engines[0].fpga_sandboxes[SANDBOX_0_NAME].fpga_
registers[FGPA_REGISTER_NAME]
# Write FPGA-Register
fpga_regw_instruction = sequence.instruction_set.fpga_register_write
fpga_register_write = sequence.add_instruction("my_fpga_register_write", 10, fpga_regw_
instruction.id)
fpga_register_write.set_parameter(fpga_regw_instruction.fpga_register.id, fpga_register_view )
fpga_register_write.set_parameter(fpga_regw_instruction.value.id, value_register)
```

FpgaMemoryMapView

Like FPGA-Registers, the `FpgaMemoryMapView` can be used after the PathWaveFPGA project has been loaded. The destination of FPGA read operation must be a register. The following example shows how you use it to read from an FPGA-Memory Map:

```
# Retrieve memory map object from memory maps collection
# All sandbox object collections are populated when loading a bit file generated by PathWave
FPGA
memory_map = sequencer.sync_sequence.engines[0].fpga_sandboxes[SANDBOX_0_NAME].fpga_memory_maps
[FGPA_MEMORY_MAP_NAME]
# Read Memory Map
fpga_arrayr_instr = sequence.instruction_set.fpga_array_read
fpga_array_read = sequence.add_instruction("my_fpga_array_read", time_ns, fpga_arrayr_instr.id)
fpga_array_read.set_parameter(fpga_arrayr_instr.fpga_memory_map.id, memory_map)
fpga_array_read.set_parameter(fpga_arrayr_instr.fpga_memory_map_offset.id, 1)
fpga_array_read.set_parameter(fpga_arrayr_instr.value.id, value_register))
```

FdsPort

FdsPort enables you to access the FDS ports you defined in the SystemDefinition. FdsPort is a object in an [fds_ports](#) collection which provides the name of a sandbox FDS port, it enables you to use an FDS port instance that you placed in the sandbox of a loaded PathWave FPGA project.

As with the other FPGA software definitions, FdsPort can only be used after the PathWaveFPGA project has been loaded.

InstructionSet Object

HVI instructions can be one of two types, HVI Native Instructions or Instrument-Specific HVI Instructions:

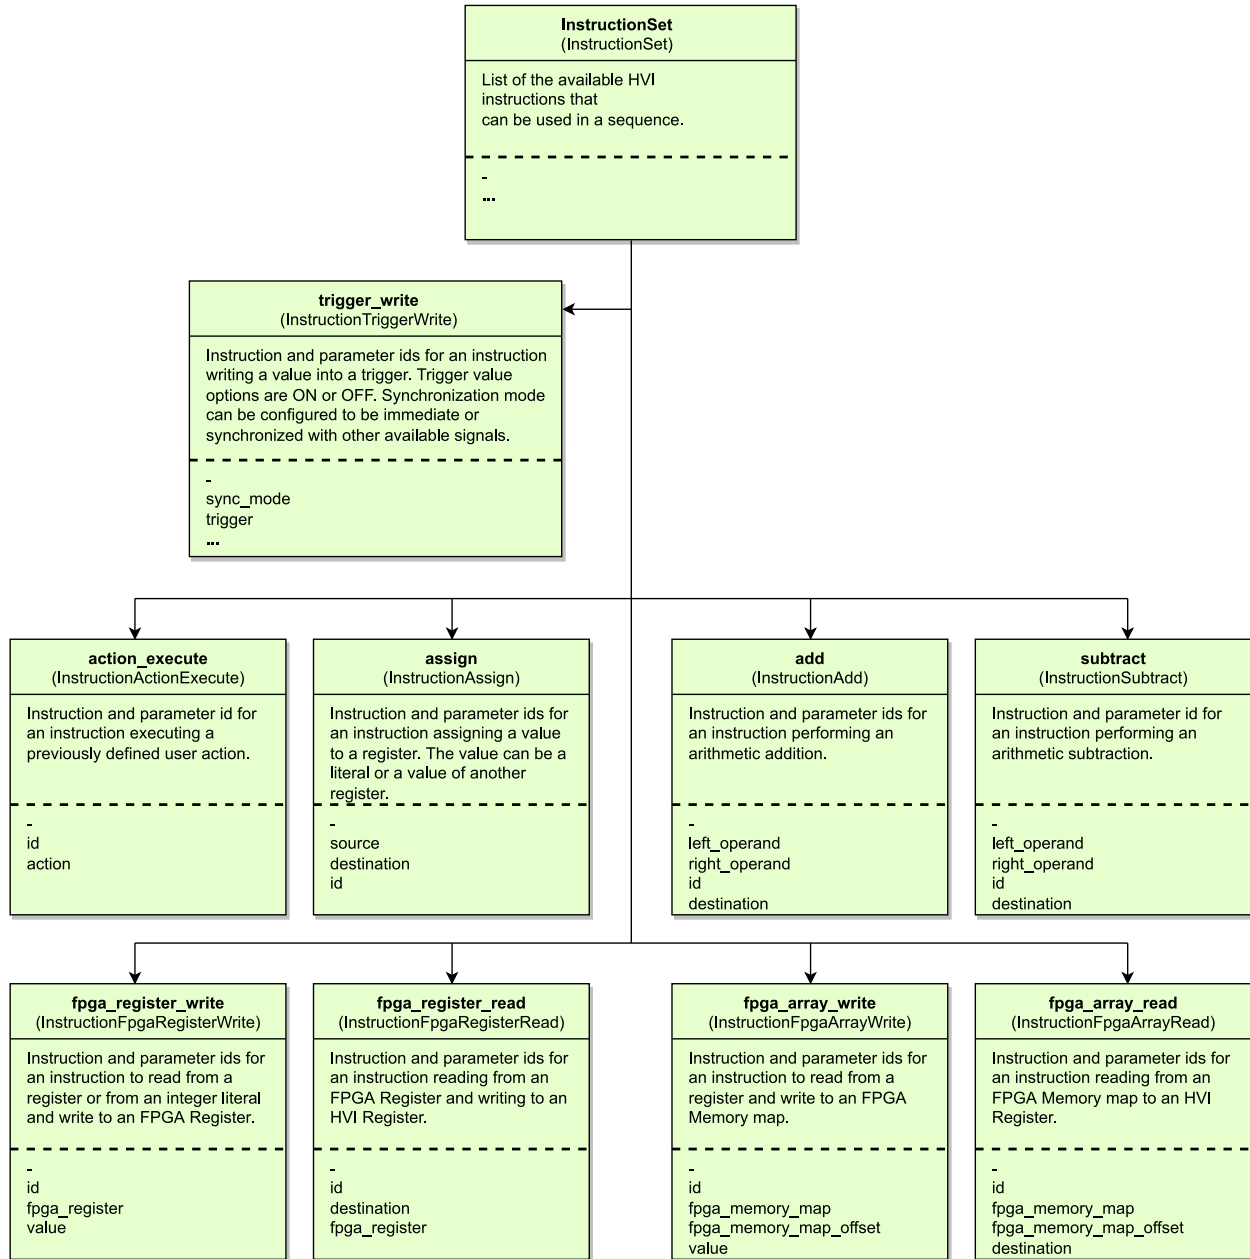
- HVI Native Instructions are part of the `instruction_set` property of a Sequence.
- Instrument-Specific HVI Instructions are documented in instrument user guides.

The InstructionSet contains the set of available HVI Native Instructions that can be executed within an HVI Statement. These include instructions for:

- Register arithmetic.
- Reading and writing I/O Trigger ports.
- Executing Actions.
- Communicating with the instrument sandbox using an HVI Host Interface, previously called an HVI Port.

HVI Native Instructions are executed within an instruction execute Statement, this is, the same way the Instrument-Specific HVI Instructions are executed.

The following diagram shows the InstructionSet and its properties:



Using the instruction set

You program HVI Instructions into Local Sequences with the `add_instruction()` method. You can set HVI Instruction parameters with the `set_parameter()` method and set each parameter with its `parameter.id` property. Some HVI Instruction parameters must be set to literal values or to an HVI Register, for example, the source and destination parameters in the `InstructionAssign` from the native `InstructionSet`.

You can set other HVI Instruction parameters such as the `SyncMode` and `TriggerValue` of the `trigger_write` instruction to one value of a pre-defined set of possible values. In this case, the possible values available are stored in properties contained within the parameter.

```
# Pseudo-code explaining the HVI Instruction programming concept
hvi_instr = sequence.instruction_set.hvi_instruction_X
instr = sequence.add_instruction("My HVI Instruction", 10, hvi_instr.id)
instr.set_parameter(hvi_instr.parameter_A.id, hvi_instr.parameter_A.VALUE_1)
instr.set_parameter(hvi_instr.parameter_B.id, hvi_instr.parameter_B.VALUE_XY)
```

Trigger write HVI Instruction example

The following example shows an example of the HVI Native Instruction `trigger_write`. For the meaning of each parameter value, see the TSE API help that is installed with PathWave Test Sync Executive. It is located at:

`C:\Program Files\Keysight\PathWave Test Sync Executive 2023\api\python\Help`

`C:\Program Files\Keysight\PathWave Test Sync Executive 2023\api\dotNet\Help`

The following table show the parameters for the HVI Native Instruction: `trigger_write`

Parameter ID	Parameter Values
<code>trigger.id</code>	<code>Trigger</code> object taken from the <code>TriggerCollection</code> object of the Engine in the Sequence.
<code>sync_mode.id</code>	<code>sync_mode.immediate</code>
	<code>sync_mode.sync</code>
<code>value.id</code>	<code>value.on</code>
	<code>value.off</code>

The following example code shows a `trigger_write` HVI Instruction.:

```
# Write FP Trigger to ON value
fp_trigger = sequence.engine.triggers["FP Trigger"]
trigger_write_instr = sequence.instruction_set.trigger_write
instr_trigger_ON = sequence.add_instruction("FP Trigger ON", 10, trigger_write_instr.id)
instr_trigger_ON.set_parameter(trigger_write_instr.trigger.id, fp_trigger)
instr_trigger_ON.set_parameter(trigger_write_instr.sync_mode.id, trigger_write_instr.sync_mode.IMMEDIATE)
instr_trigger_ON.set_parameter(trigger_write_instr.value.id, trigger_write_instr.value.ON)
```

WARNING You must take the Trigger from the Engine inside the Sequence. Taking the Trigger from the Engine via `SystemDefinition` raises an error when `set_parameter(trigger_write_instr.trigger.id, fp_trigger)` is called.

This also applies to Actions.

Instrument-Specific HVI Instructions

You program Instrument-Specific HVI Instructions into your HVI Sequences using the same methods as HVI Native Instructions, that is, you add Instrument-Specific HVI Instructions to Local Sequences with the `add_instruction()` API method. Parameters of Instrument-Specific HVI Instructions are also set with the `set_parameter()` API method. For documentation on Instrument-Specific HVI Instructions and their parameters, see your instrument documentation. For M3xxxA PXI instruments, the information is located in the *SD1 3.x Software for M320xA / M330xA Arbitrary Waveform Generators User's Guide* available at [M3201A PXIe Arbitrary Waveform Generator](#).

HVI Registers and Scopes

HVI Registers

HVI Registers are similar to Variables in a programming language. They hold values that can be modified at runtime and can be used as parameters for instructions and Statements. Physically, HVI Registers are small hardware memories located in HVI Engines. The number of registers available depends on the instrument, for more information see the HVI Engine setting [HviRegCount](#).

Registers are specific to individual HVI Engines and cannot be accessed by other HVI Engines. To transfer data between HVI registers in different instruments you must use register sharing instructions.

You define HVI Registers by adding them to the HVI Register collection in the scope corresponding to each specific HVI Engine.

HVI Sync Registers

While HVI Registers are HVI Engine specific, HVI Sync Registers are global registers present in all HVI Engines in a sequence that holds the same value. HVI Sync Registers are internally implemented by the "local" HVI Registers, and the HVI Statements and the HVI compiler are responsible for maintaining the consistency of the Sync Registers across instruments. To guarantee consistency, Sync Registers can only be modified by Sync Statements, such as Sync-For. Sync Registers can be used in any statement for reading purposes.

The use of Sync Registers in Sync Statements enables the HVI Compiler to optimize the compiled code to avoid using triggered-resynchronizations, this results in significantly faster real-time execution. For instance, the Sync-For Statement using Sync-Registers can loop all instruments, fully synchronized, with the same performance of a Local-For statement in a single instrument.

Sync Registers are defined and accessed through the Sync-Scope exposed in the Sync Sequence instances.

HVI Scope

HVI Sync Sequences and HVI Local Sequences both include the concept of *the scope of registers*, this is similar to the concept of *the scope of Variables* in programming languages. The scope defines what registers or memory resources can be used within each HVI Sequence, and when they can be used.

Each scope is associated with a specific Sequence and HVI Engine. Registers can only be defined within the Global Sync Sequence scope, but they can be retrieved from any child Sequence scope providing it is on the same HVI Engine. Registers are always defined with a clear connection to a specific Engine and their visibility only propagates to child Sequences that execute on the same Engine.

HVI Engines do not have visibility of, and cannot access registers that are in the scopes of other HVI Engines.

NOTE

Registers can only be added to the HVI top Sync Sequence scopes. This means that you can only add global registers that are visible in all child Sequences.

NOTE

Registers are created using the Sequencer, but to read/write Registers during HVI execution, you must use the `RegisterRunTime` within the `Hvi` object. For more information, see [The Hvi Object](#).

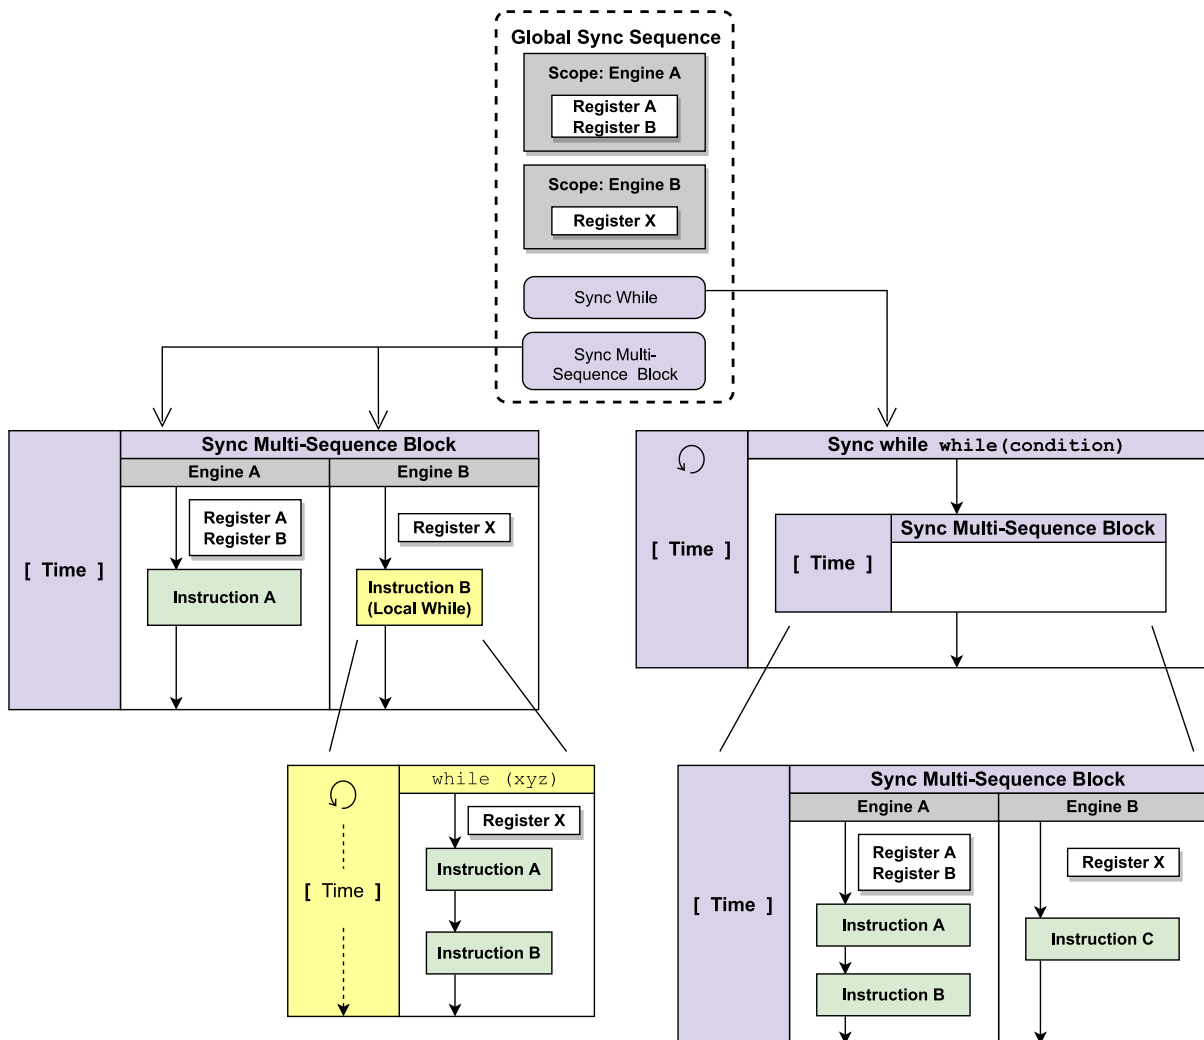
The following diagram shows the scope concepts. The registers available are shown in the Sequences and child Sequences.

In the Global Sync Sequence a scope is defined for each of **Engine A** and **Engine B**.

- Engine A scope contains Register A and Register B.
- Engine B scope contains Register X.

The Global Sync Sequence contains Sync Statements including a Sync while and a Sync Multi-Sequence Block. These are expanded as HVI diagrams. The Sync Multi-Sequence Block contains Sequences for both Engines. These are shown with the registers available in blue for Engine A, red for Engine B.. The Sequence for Engine B contains a Local While. This is expanded below with the available Register X shown in red.

The Sync while in the Global Sync Sequence is also shown, it contains another Sync Multi-Sequence Block which is shown expanded.



Scope, SyncScope and ScopeCollection

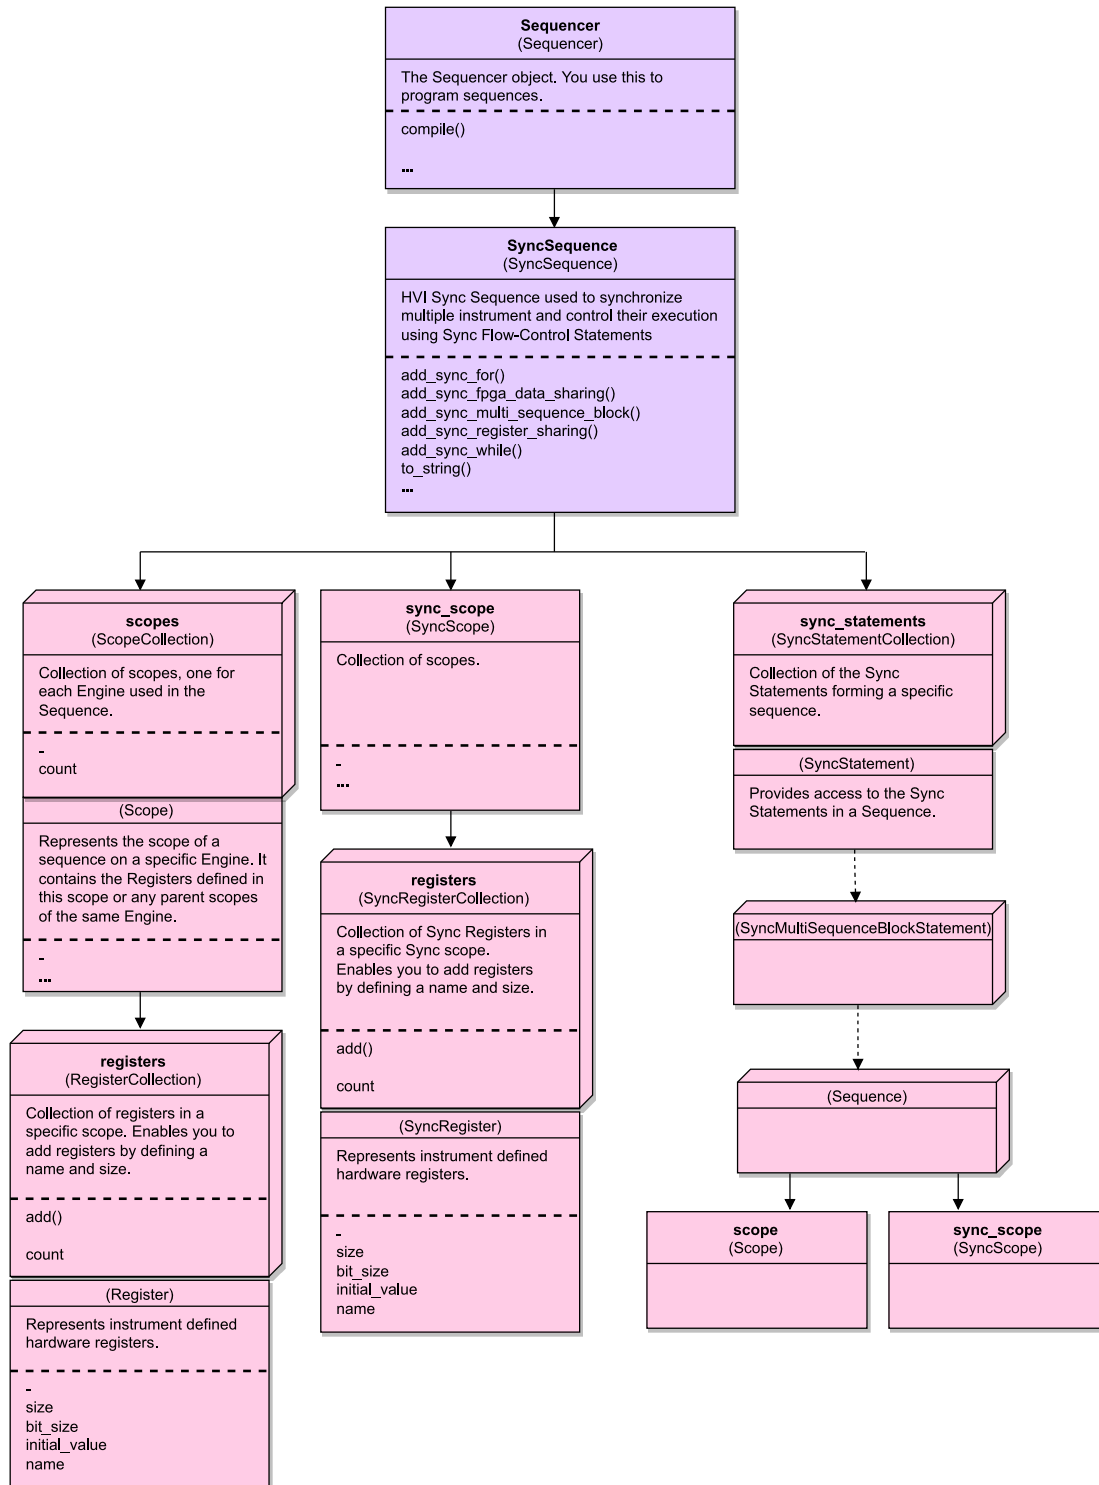
The register scopes of HVI Sequences are managed through a Scope. Each Local Sequence instance is associated to a specific HVI Engine and has its own Scope object. Sync Sequences are associated to multiple HVI Engines and consequently have an HVI Scope collection that contains a Scope for each associated HVI Engine.

Each HVI Scope can be accessed from the Scope collection using the same name as the corresponding HVI Engine. HVI Scopes are used to define the registers within a sequence.

Similarly, the scope of Sync Registers is managed through the SyncScope instance. Each SyncScope instance is owned by the Sync Sequence, and can be accessed through the SyncScope property available in the Sync Sequence and also in Local Sequences inside that specific Sync Sequence.

To use Registers or Sync Registers in HVI Sequences, you must add them to the Register collection within the scope or SyncScope of the corresponding HVI Sequence or Sync Sequence.

The following diagram shows the Scopes and their relationship to the Sequencer and Sync Sequence:



HVI Compilation

Once you have programmed all of your HVI Sequences, the next step is to compile them. The compilation process returns the `Hvi` object that is used to run the created Sequences on hardware.

Call the `compile()` method in the `Sequencer` object to perform the compilation operation. If successful, this method returns an `Hvi` object, if the compilation fails, it throws an exception.

The compilation process translates the programmed Sequence into binary instructions to be loaded into the hardware. During this process, the compiler applies the compilation rules, evaluates the specified constraints, and determines if the number of resources required (PXI Triggers, Actions, Events, HVI Registers) is available in hardware and can be acquired. The compiler returns an error if any of the HVI Statements was not programmed properly inside the HVI Sequence or if any of the HVI resources are missing or not registered properly.

NOTE At this point you can no longer modify Sequences, Actions, Events or Triggers.

Information returned

The value returned from the compilation procedure is an `Hvi` object. This object can be used to:

- Load and execute the binary instructions by each Engine.
- Retrieve the `CompileStatus` object.

Errors returned

If the compilation fails, the object `keysight_tse.CompilationFailed` is thrown. This contains a `CompileStatus` object.

In the following Python snippet, the `CompileStatus` object is retrieved from the exception object thrown:


```
try:
    hvi = sequencer.compile()
    print('Compilation completed successfully!')
except kthvi.CompilationFailed as err:
    print('Compilation failed!')
    compile_status = err.compile_status
    print(compile_status.to_string()) # This line will print all the errors and warnings
    collected during compilation raise err
```

Compile status

The `CompileStatus` object contains the following information:

- The elapsed time of the compilation process.
- The warning and error messages generated by the compilation.
- Information about the PXI sync resources that must be reserved.

The following diagram shows `CompileStatus` and its properties:



System Initialization in the Sequencer Object

When a Sequencer object is created, a system initialization is implicitly executed, but only if it is required. This system initialization is only intended to ensure that the hardware included in the System-Definition can work together and that the minimal information required to build real-time sequences is available.

In the case that the infrastructure, in particular the System Sync Modules (SSMs), are not properly initialized and aligned, a full system initialization is triggered, equivalent to calling `SystemDefinition::initialize(AlignmentModes::Full)`. If all of the System Sync Modules are already aligned, it checks if there are any instruments misaligned and force initializes them, skipping any others that are already aligned.

NOTE

The initialization process requires exclusive access to the hardware resources involved, so it cannot be executed while another TSE instance is running and has that hardware locked.

For a full explanation of initialization, see [System Initialization](#).

Sequence Representation

This section describes Sequence Representation, it contains the following sections:

- [Using Sequence representation](#)
- [Format of the Sequence Representation Output](#)
- [Format Variations](#)
- [Sequence Representation Error Messages](#)

PathWave Test Sync Executive enables you to troubleshoot your Sequences with Sequence Representation.

The Sequence representation displays Statements, timing values, and Statement parameters. The output is designed so you can read it and see what your Sequences are doing.

NOTE

- Sequence representation is only available for Sync Sequences in this release.
- There is also an HVI Logger that provides more detailed information, this is primarily designed for support. See [HVI Logging and Troubleshooting](#).

Using Sequence representation

To activate the output, In Python use the Sequence method `to_string()`:

```
output = sequencer.sync_sequence.to_string(kthvi.OutputFormat.DEBUG)
print(output)
```

If you are programming with C#, use the method `ToString`:

```
var output = GlobalSequence.ToString(OutputFormat.Debug);
System.Console.WriteLine(output);
```

Format of the Sequence Representation Output

Sequence representation has a basic structure with variations for different types of Statements.

The representation out format has one Statement per line and uses curly braces to begin and end any inner or Local Statements.

The basic format is:

```
Time-related information => "User-assigned Label" : Statement Name(Parameter List) {  
Optional statements  
}
```

For example:

```
+10ns => "FP Trigger ON": TriggerWrite(Trigger = [trigger"TriggerIO"], Mode = IMMEDIATE, Value = ON)
```

For Arithmetic-like and FPGA Statements the format is:

```
Time-related information => "User-assigned Name" : ASSIGNEE = EXPRESSION
```

where:

- **ASSIGNEE** is a named reference, such as `event`, `trigger`, `action`, `reg`, or `fpgaReg` followed by the label in quotes.
- **EXPRESSION** is a mathematical expression with binary operators, such as addition, subtraction, and assignment.

For example:

```
+10ns => "Increment counter register": reg"LeaderEngine.Loop Counter" = reg"LeaderEngine.Loop Counter" + 1
```

Time related information

The time information section of the representation output is in the following format:

```
+Start_delay <Duration> Absolute_time =>
```

NOTE

There are a number of limitations in this release:

- Duration is shown as **Min** or **?**.
- Absolute time is not shown in this release.

Indicators

The representation output uses the following characters to indicate different pieces of information:

Category	Indicators	Description
Timing-related information	+	Appears at the start, the number with this indicates the Start delay.
	<>	Encloses a Duration if it is set. If the Duration is not set, this defaults to <i>min</i> , which is the minimum time possible.
		Absolute time (not supported in this release).
Separator	=>	Separator. The time information for the Statement is on the left of this and information about the Statement is on the right.
Command label and name	" "	Encloses labels
	:	Divides the label and the command description.
Blocks and parameters	{ ... }	Encloses blocks of Statements: <ul style="list-style-type: none"> • Sync multi-Sequence block. • Engine instructions. • Sync Flow-Control. • Local Flow-Control.
	(...)	Enclose parameters. These can be optional.
	[...]	Enclose lists. For example [element, ...], or for named element lists [name"username", ...]
Register indicators	reg	Indicates a register.
	fpgaReg	Indicates an FPGA-Register

Code blocks

Code blocks are indented and shown within curly braces. Code blocks include code in Sync multi-Sequence blocks, Engines, and flow-control Statements.

The following example from *Programming Example 1* shows a Sync multi-Sequence block `TriggerAWGs` that contains a pair of Engines `AwgEngine0` and `AwgEngine1`.

```
+30ns<Min> => "TriggerAWGs": SyncMultiSequenceBlock {
  Engine "AwgEngine0" {
    +10ns => "FP Trigger ON": TriggerWrite(Trigger = [trigger"TriggerIO"],
      Mode = IMMEDIATE, Value = ON)
    +100ns => "FP Trigger OFF": TriggerWrite(Trigger = [trigger"TriggerIO"],
      Mode = IMMEDIATE, Value = OFF)
    +10ns => "AWG trigger": ActionExecute([action"GenMarker1", action"GenMarker2"])
  }
  Engine "AwgEngine1" {
    +10ns => "FP Trigger ON": TriggerWrite(Trigger = [trigger"TriggerIO"],
      Mode = IMMEDIATE, Value = ON)
    +100ns => "FP Trigger OFF": TriggerWrite(Trigger = [trigger"TriggerIO"],
      Mode = IMMEDIATE, Value = OFF)
    +10ns => "AWG trigger": ActionExecute([action"GenMarker1", action"GenMarker2"])
  }
}
```

If an Engine does not execute any Statements, the Engine is shown with empty braces. For example, in the previous example, if the Engine `AwgEngine1` didn't have any instructions, it would be shown as:

```
Engine "AwgEngine1" {}
```

Format Variations

There are variations of the Sequence representation output format for different types of Statement.

Sync Statements

The following example shows a Sync register-sharing command that copies the contents of the `steps` register in the Digitizer Engine to the `wavefrom ID` register in the AWG Engine:

```
+190ns<Min> => "Share steps->wfm_id": SyncRegisterSharing {
    reg"Digitizer Engine.Steps"[1:0] => [reg"AWG Engine.Waveform ID"]
}
```

Sync multi-Sequence blocks

The output for a Sync multi-Sequence block indicates any Engines it contains. The Sequences and the Statements they contain are shown within each Engine.

The following example shows the output for a Sync multi-Sequence block that contains 2 Engines. The first Engine is labelled Digitizer Engine and contains a Sequence with a pair of Local Statements. A second Engine labelled AWG Engine does not contain any Sequences. This is indicated with empty braces.

Representation output for a Sync multi-Sequence block:

```
+260ns<Min> => "Loop Delay": SyncMultiSequenceBlock {
    Engine "Digitizer Engine" {
        +10ns => "loops++": reg"Digitizer Engine.Loops" = reg"Digitizer Engine.Loops" + 1
        +30ns<?> => "WaitTime: loop_delay": WaitTime(reg"Digitizer Engine.Loop Delay")
    }
    Engine "AWG Engine" {}
}
```

Sync flow-control and Local flow-control Statements

Flow control Statements show the flow-control condition and the Statements executed if the condition is met.

The following example shows a Local If. The condition is indicated along with the matching branches parameter and the Statement executed is also shown inside braces.

Representation output for a Local If Statement:

```
+70ns<?> => "Check wfm_id": If(condition = (reg"AWG Engine.Waveform ID" > = 1), MatchingBranches = TRUE) {
    +30ns => "wfm_id = 0": reg"AWG Engine.Waveform ID" = 0
}
```

If a flow control instruction contains multiple branches, these are also listed.

The following example contains a Local If with a condition and an Else branch that is executed when the If condition is not met.

```
+70ns<?> => "Queue Wfm AWG": If(condition = (reg"AWG Engine0.Queue Reg" == 0), MatchingBranches = TRUE) {
    +100ns => "Queue Waveform A CH1": M30xxA.AwgQueue(Channel = 1, WaveformId = reg"AWG Engine0.Wfm A", Cycles = 3, StartDelay = 0, Prescaler = 0, TriggerMode = AUTOTRIG)
}
Else {
    +100ns => "Queue Waveform B CH1": M30xxA.AwgQueue(Channel = 1, WaveformId = reg"AWG Engine0.Wfm B", Cycles = 2, StartDelay = 0, Prescaler = 0, TriggerMode = AUTOTRIG)
}
```

HVI instructions

HVI Instructions show the Start delay, the label, instruction and any parameters. For example:

```
+10ns => "Increment counter register": reg"LeaderEngine.Loop Counter" = reg"LeaderEngine.Loop Counter" + 1
```

Custom instructions

Custom instructions indicate the product family before the instruction in the form:

```
ProductFamily.CustomInstructionName
```

In the following example, the product family `kM30xxA` is indicated before the custom instruction `QueueWaveform`:

```
+100ns => "QueueWaveform(CH1, wfm_id)": M30xxA.AwgQueue(Channel = 1, WaveformId = reg"AWG Engine.Waveform ID", Cycles = 1, StartDelay = 0, Prescaler = 0, TriggerMode = AUTOTRIG)
```

Examples

The following example is an excerpt from Programming Example 1. It shows the Python code for setting up the TriggerWrite and ActionExecute instructions and the resulting Sequence representation output that is generated.

Python Code:

```
# Write FP Trigger ON to all instruments
fp_trigger = sequence.engine.triggers[config.fp_trigger_name]
trigger_write = sequence.instruction_set.trigger_write
instr_trigger_ON = sequence.add_instruction("FP Trigger ON", 10, trigger_write.id)
instr_trigger_ON.set_parameter(trigger_write.trigger.id, fp_trigger)
instr_trigger_ON.set_parameter(trigger_write.sync_mode.id, trigger_write.sync_mode.immediate)
instr_trigger_ON.set_parameter(trigger_write.value.id, trigger_write.value.on)
# Write FP Trigger OFF to all instruments
instr_trigger_OFF = sequence.add_instruction("FP Trigger OFF", 100, trigger_write.id)
instr_trigger_OFF.set_parameter(trigger_write.trigger.id, fp_trigger)
instr_trigger_OFF.set_parameter(trigger_write.sync_mode.id, trigger_write.sync_mode.immediate)
instr_trigger_OFF.set_parameter(trigger_write.value.id, trigger_write.value.Off)
# Execute AWG trigger from the HVI sequence of each module
# "Action Execute" instruction executes the AWG trigger from HVI
action_list = sequence.engine.actions
instruction1 = sequence.add_instruction("AWG trigger", 10, sequence.instruction_set.action_
execute.id)
instruction1.set_parameter(sequence.instruction_set.action_execute.action.id, action_list)
```

The Sequence representation output from the previous code:

```
+10ns => "FP Trigger ON": TriggerWrite(Trigger = [trigger"TriggerIO"],
Mode = IMMEDIATE, Value = ON)
+100ns => "FP Trigger OFF": TriggerWrite(Trigger = [trigger"TriggerIO"],
Mode = IMMEDIATE, Value = OFF)
+10ns => "AWG trigger": ActionExecute([action"GenMarker1", action"GenMarker2"])
```


Sequence Representation Error Messages

The Sequence representation system can detect and report errors.

Errors can be part of an assignment expression, destination, or parameter value.

For example, if a parameter has not been set in an instruction. In the case of register parameters this can result in values completely missing from the output or exceptions being thrown.

Error formats

Errors are indicated in the following formats:

Errors with a message

An error is indicated in the following manner, the messages provided do not contain @ symbols:

```
@ERROR: <message>@
```

Errors with no message

In some cases, an error is indicated without a message:

```
@ERROR@
```

The following example output shows some example errors:

```
+90ns<Min> => "Sync MIMO Trigger": SyncWhile(reg"AwgEngine0.Loops" < 3) {
  +250ns<Min> => "TriggerAWGs":SyncMultiSequenceBlock {
    Engine "AwgEngine0" {
      +10ns => "assign":@ERROR: register is not set@ = @ERROR@
      +100ns => "QueueWaveform(CH1, wfm_id)": M30xA.AwgQueue(Channel = @ERROR@,
        WaveformId = @ERROR: invalid id@, Cycles = 1, StartDelay = 0, Prescaler = 0,
        TriggerMode = AUTOTRIG)
    }
  }
}
```

The Hvi Object

This section describes the Hvi object, it contains the following sections:

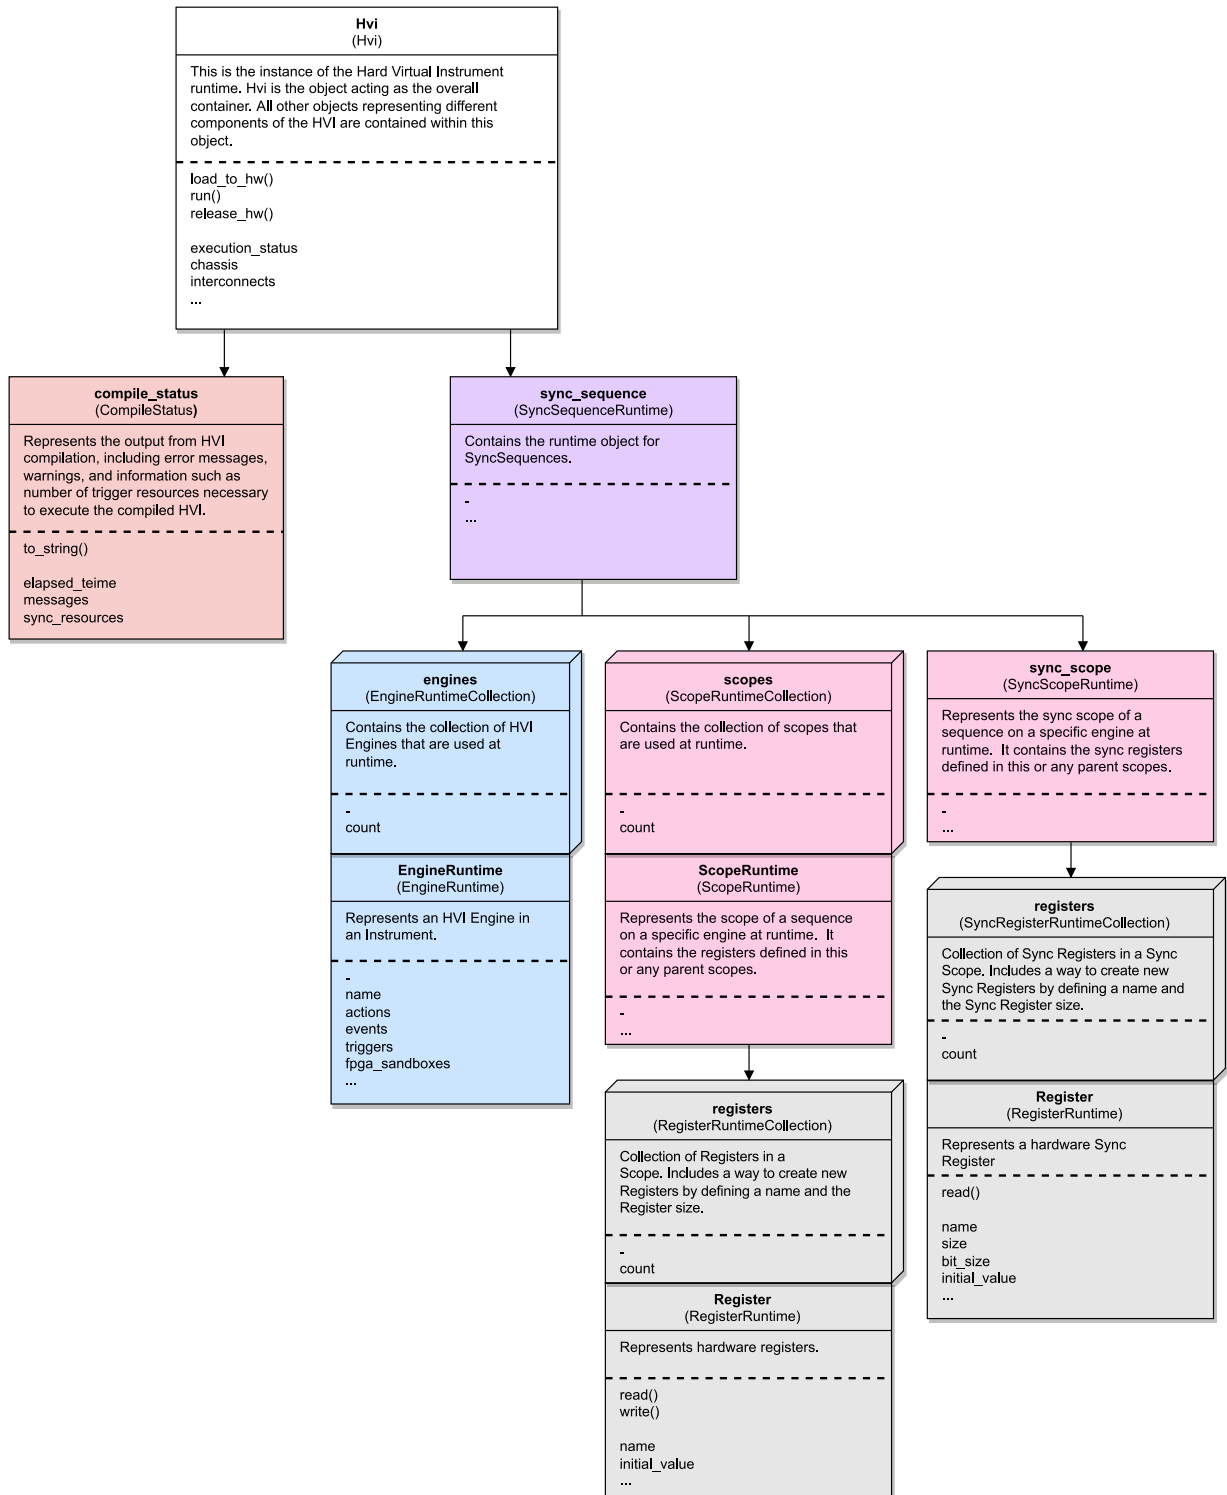
- [Engine Runtime Components](#)
- [Load to Hardware and Run](#)
- [System Initialization during Load To Hardware](#)
- [Real-time Hardware Execution Error Handling](#)

The `Hvi` object is the actual HVI instance that is loaded to hardware and executed. It contains the runtime versions of the objects you set up with the `SystemDefinition` and `Sequencer` instances. The runtime objects are the instances of the objects that operate while the HVI is running. For most cases, you cannot change the configuration of these objects at runtime, but you can access and use these resources such as HVI registers or an FPGA-Memory Map.

NOTE

The `Hvi` object is the runtime object. once you have compiled it, you can no longer change resources or Sequences.

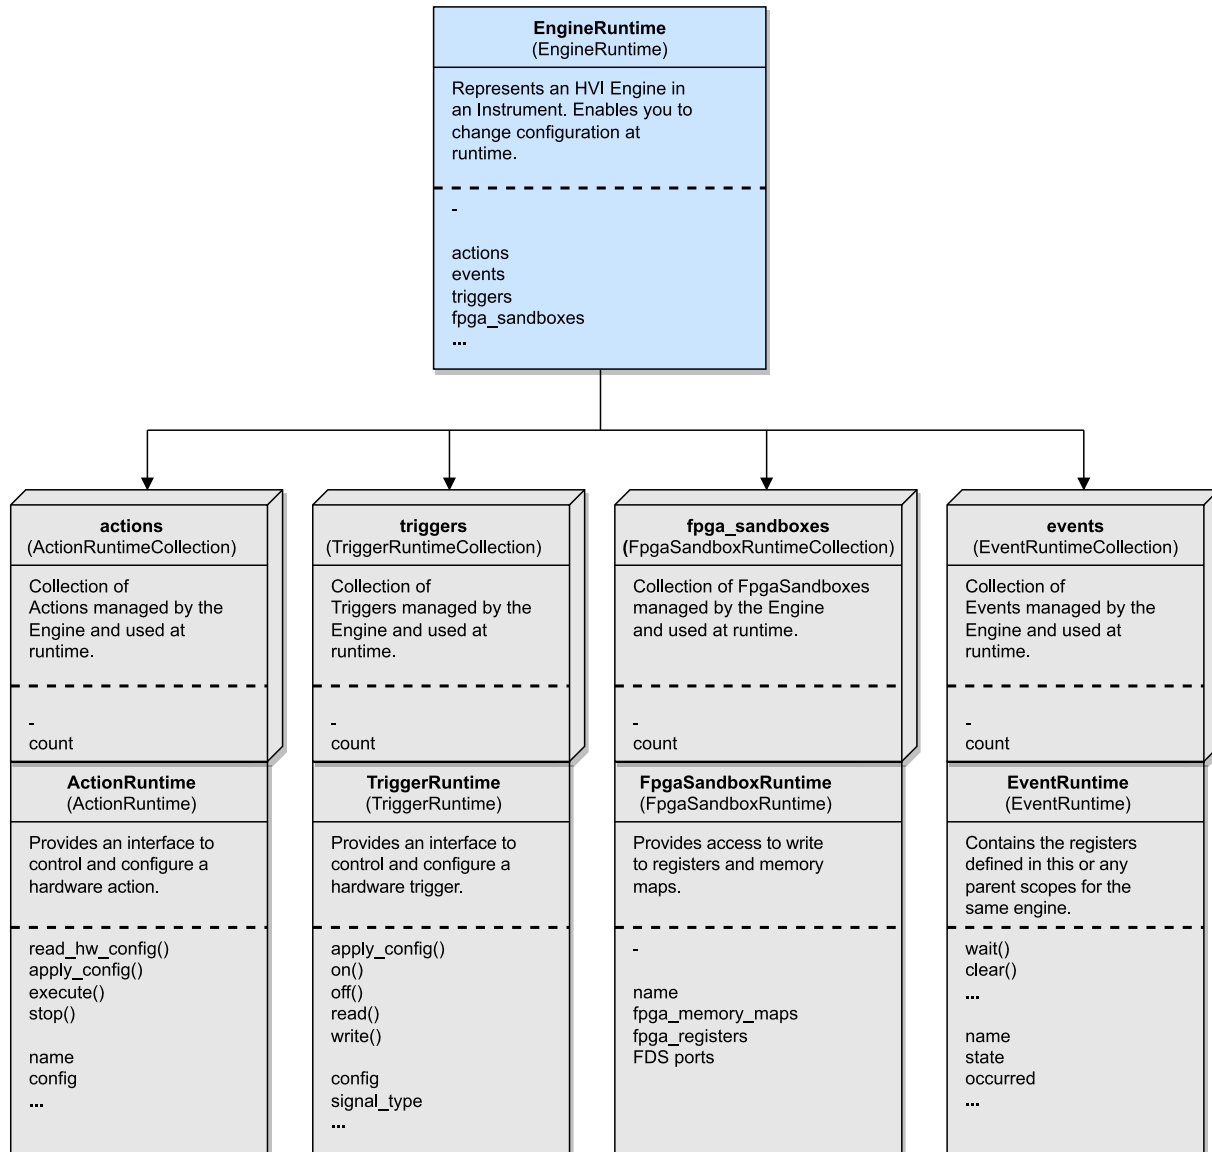
The following diagram shows the Hvi object:



Engine Runtime Components

A number of runtime components are under the `EngineRuntime`.

The following diagram shows the `EngineRuntime` object and properties:



ActionRuntime

Once the HVI instance is loaded to hardware, it is possible to change the actions configuration and trigger the action execution or state change for actions previously configured in the systemDefinition instance.

TriggerRuntime

Once the HVI instance is loaded to hardware, it is possible to change the trigger configuration, read trigger state and change state in the case of trigger configured as outputs for triggers previously configured in the systemDefinition instance.

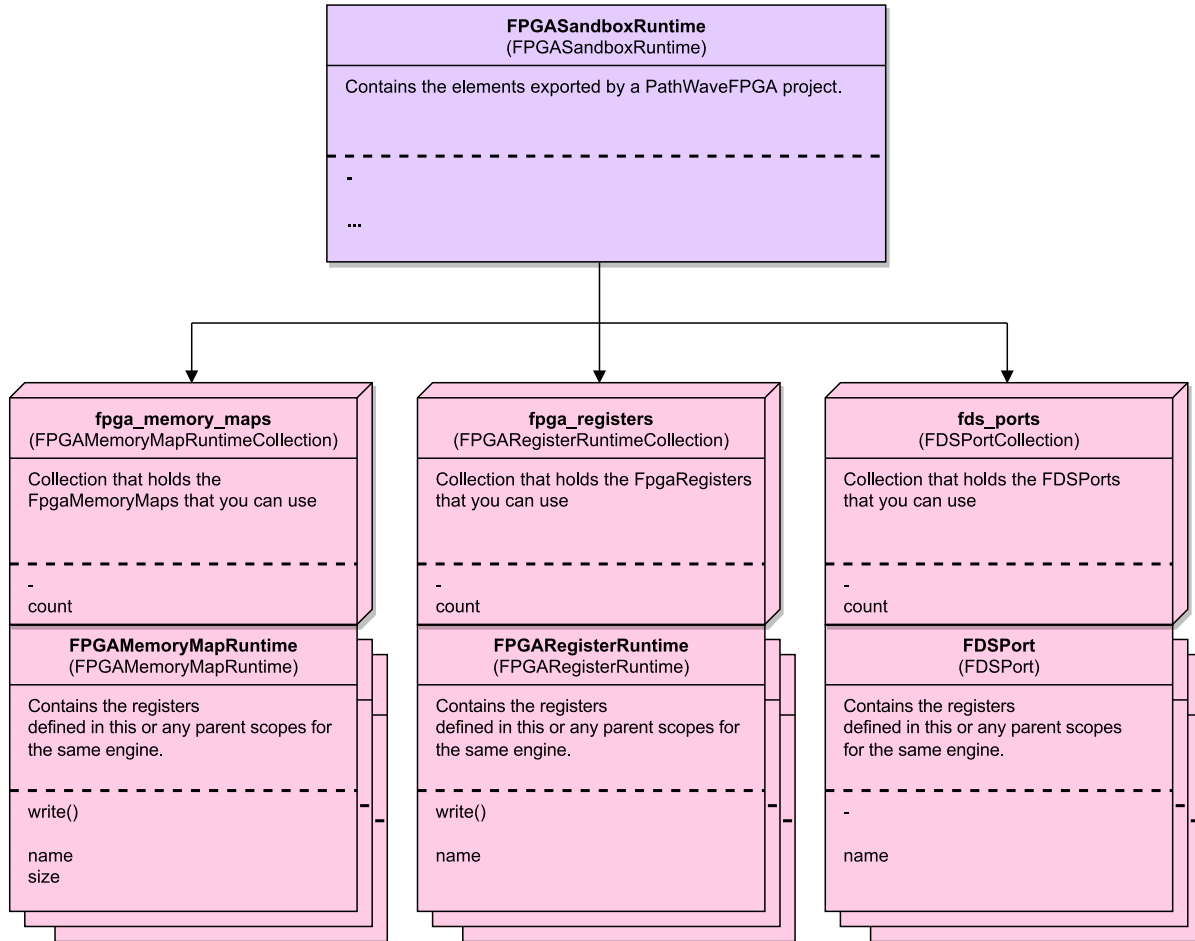
EventRuntime

Once the HVI instance is loaded to hardware, it is possible to query and wait for events previously configured in the systemDefinition instance.

FpgaSandboxRuntime

This section describes the FPGA-Sandbox runtime.

`FPGASandboxRuntime` contains all the FPGA-Registers and FPGA-Memory Maps available at runtime. The following diagram shows the objects:



The `FPGASandboxRuntime` object can be obtained from the `Hvi` object:

```
SANDBOX_0_NAME = "sandbox0" sandbox = hvi.sync_sequence.engines[0].fpga_sandboxes[SANDBOX_0_NAME]
```

NOTE Hvi resources can only be read or written loaded, that is, between the `load_to_hw()` and `release_hw()` calls. Any attempt to read or write resources without having the instrument loaded to hardware results in an exception being thrown.

FPGA-Registers

Once the Sequencer has been compiled and the HVI has been loaded to hardware, the register can be read and written. If the HVI is not loaded, an exception is thrown.

```
fpga_register = hvi.sync_sequence.engines[0].fpga_sandboxes[SANDBOX_0_NAME].fpga_registers[0]
hvi.load_to_hw()
fpga_register.write(1) # ok
hvi.release_hw()
fpga_register.write(1) # exception is thrown
```

FPGA-Memory Maps

As with registers, FPGA-Memory Maps can be used after HVI has been loaded to hardware. They can only be accessed, read, or written while the HVI is loaded to hardware.

```
fpga_memory_map = hvi.sync_sequence.engines[0].fpga_sandboxes[SANDBOX_0_NAME].fpga_memory_maps
[0]
hvi.load_to_hw()
fpga_memory_map.write(0x10, 0x1245) # ok
hvi.release_hw()
fpga_memory_map.write(0x10, 0x1245) # exception is thrown
```

FDS Ports

As with registers, FPGA-Memory Maps can only be accessed after HVI has been loaded to hardware.

Load to Hardware and Run

A successful compilation of the Sequencer instance returns the `hvi` instance. To execute it, you must load it to hardware and run it.

These operations are performed using the following methods of the Hvi object.

To load the HVI to hardware call the method `hvi.load_to_hw()`.

The `hvi.load_to_hw()` method deploys HVI to hardware and does all of the resource configuration including:

- Synchronization resources.
- Trigger resources.
- Clocks.

The `hvi.load_to_hw()` method also loads the binaries containing information to execute the HVI Sequences, to the relevant HVI Engines.

Once the HVI has been loaded to hardware, you can execute your Sequences by calling `hvi.run()`. The HVI execution in Hardware finishes when the HVI Sequence reaches the end. The `Stop()` method can be used to stop or cancel the HVI execution.

When the HVI has finished execution and it is not needed to run the HVI again, call the method `ReleaseHw()` to release or free all resources used by the HVI.

System Initialization during Load To Hardware

During Load to Hardware, a system initialization is always performed. The initialization is equivalent to the explicit call `sys_def.initialize()` executed without parameters. This only performs a minimal update to the initialization and clock alignment.

For a full explanation of initialization and all the parameters, see [System Initialization](#).

Real-time Hardware Execution Error Handling

This section describes real-time hardware error handling during the HVI execution.

If a hardware error is detected while a Sequence is running, it is possible the execution results are invalid or unreliable. HVI captures some critical hardware errors and reports them to the user. Errors are indicated in a different way, depending on if you are running your HVI Sequence in blocking mode or non-blocking mode:

Errors in an HVI Sequence Running in Blocking Mode

If an error occurs when the HVI is executed in blocking mode, the Sequence execution is halted, an exception is thrown, and a list of the errors can be queried.

The error reporting Sequence goes in the following order:

1. Call `loadtoHw()`
2. Call `run()`. When an error Event occurs while the HVI is running in hardware:
 - a. The HVI Sequence is halted.
 - b. The error list is updated.
 - c. An exception is thrown.
3. Call `getExecutionErrors()` to obtain details of the errors that has occurred during the execution.

Errors in an HVI Sequence Running in Non-Blocking Mode

If an error occurs when the HVI is in non-blocking mode, Sequence execution is not halted automatically and no exception is thrown, since in non-blocking mode the `run()` method returns immediately as soon as the Sequence execution starts. In order to query the status of the HVI execution you must call `getExecutionStatus()`. This method returns the enum `ExecutionStatus` with values for:

- Not started.
- Running.
- Sequence completed successfully.
- Sequence stopped due to timeout.
- Sequence stopped due to error.
- Sequence stopped by user.

Retrieving the Errors List

If a Sequence stops because of an error, a list is populated with all the errors detected during execution.

To retrieve the errors call the method `getExecutionErrors()`, it returns a list of any errors that occurred during the last run.

If you call `getExecutionErrors()` more than once, it returns the same list. Calling `run()`, `loadToHw()` or `ReleaseHw()` clears the list.

The `ExecutionError` object provides the following information:

- A complete string-based representation for easy printing.
- The name of the HVI Engine that reported the error.
- The product code, this is an integer defined by the instrument that identifies the specific error.
- The product message, this is a string provided by the instrument with any relevant details.

For instrument specific errors, see your instrument documentation for a description of the errors a specific instrument returns.

Chapter 8: Building an Application with the TSE API

This chapter describes the steps you must follow to use the TSE API. If you do not follow these steps your application shall not work correctly.

HVI uses *program-within-a-program* model. That is, the HVI enables you to define a program that runs on the instrument's hardware while the software programs run in parallel and can interact with the instruments. HVI is also responsible for all the setup, compilation, and hardware execution management. When you run your application, it generates an HVI instance and the Sequences within it are executed on the instruments.

This chapter contains the following sections:

- [Planning an HVI with the HVI Use Model](#)
- [1 Create the SystemDefinition](#)
- [2. Program HVI Sequences](#)
- [3. Compile Your Sequences](#)
- [4. Load To Hardware](#)
- [5. Modify Initial Register Values \(Optional\)](#)
- [6. Execute Sequences](#)
- [7. Release All Resources](#)
- [HVI Logging and Troubleshooting](#)

NOTE

The code examples provided in this chapter are in both Python and C#.

Planning an HVI with the HVI Use Model

Programming and executing an HVI requires you to follow a precise use model. You must write your code in the correct order and be aware of the requirements of each stage, or your application shall not work correctly.

The HVI Use Model consists of 3 broad stages:

1. System Definition

You must define hardware resources before you can use them in HVI. The resources you can use depends on your hardware set up, what instruments you have, what capabilities they have, and how they are arranged. You set these up first and then you can use their functionalities in your HVI Sequences. This operation is called *System Definition* and it can be done by using an instance of the `SystemDefinition`.

The initialization of the system you have defined is also important to understand. By default, the defined system is initialized at the code line that is creating the `Sequencer` object from the `SystemDefinition` object. If you use the default initialization, this ensures that the complete system is correctly initialized.

There are some use cases when you might need to use the `initialize()` API method to perform a custom initialization, for example, a full realignment of the HVI Engine clocks. For more information, see the description of *AlignmentMode* list in [System Initialization](#) and the Python API Help. In this case it is important to make sure that the `SystemDefinition` object is not modified after calling the `initialize()` API method.

NOTE

Ensure you initialize your system after all the resources have been added and defined. If you call the `initialize()` API method before the system is fully defined, the system shall not be initialized properly. Consequences of an improper initialization might be that some instruments included in the HVI might be out of sync or that their HVI Sequence execution will misbehave by for example, missing a Trigger or not playing a waveform.

For example, in the following code `initialize()` is called incorrectly before all the Engines are added to the `SystemDefinition`.

```
# call initialize()
#
sys_def.initialize()
#
# Incorrect usage, Engines added after the initialize() call are not initialized.
sys_def.engines.add(...)
```

NOTE

PathWave Test Sync Executive 2023 release includes TSE Service. If your system uses TSE Service the set-up of a system is different.

The setup of a system that uses TSE Service handles starting up chassis, SSMs, and instruments. It also handles clocking and other infrastructure tasks.

A system definition or equivalent is still required, however this is done by a system administrator.

This arrangement simplifies sequence programming for end users since they are not required to write a full System Definition each time.

Users are still required to add instruments get their engines, and call `initialize()`. For more information, see [Chapter 9: TSE Service and Multi-Host support](#).

2. Program HVI Sequences

As a next step, the `SystemDefinition` object is passed into the `Sequencer` object at the Sequencer creation.

Once the `Sequencer` object is created, the `SystemDefinition` instance is fixed. All resources added and defined using the `SystemDefinition` object must be modified before this step. You cannot make any changes to the `SystemDefinition` instance after this. Any changes made in the `SystemDefinition` after this point are not passed into the `Sequencer` object and therefore are not included in the HVI.

Once the hardware is set up and resources assigned, you can write your Sequences and set initialization values. You create Sync Sequences for globally synchronized operations, and you create Local Sequences for operations in the HVI Engines in individual instruments.

3. Execute the HVI

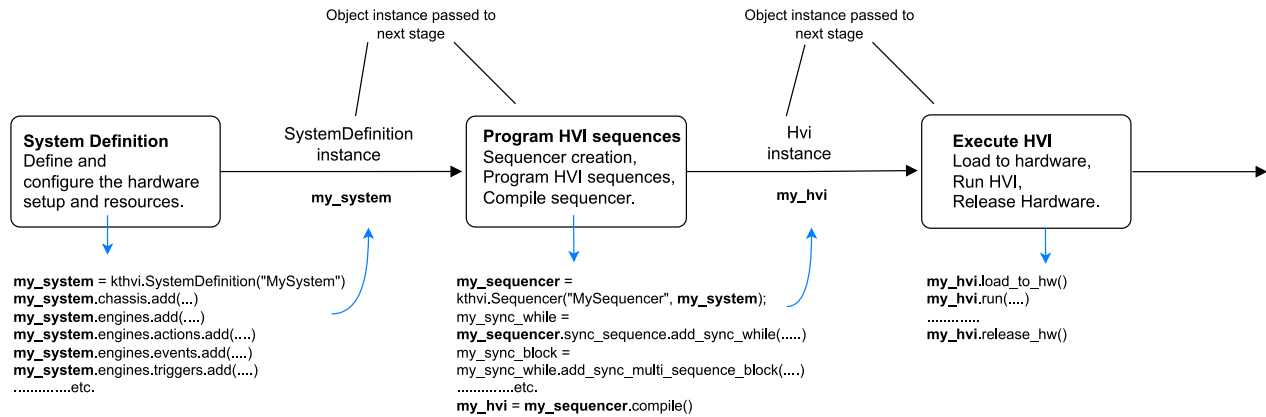
When you have programmed your Sequences, you call the compile software method to create an instance of the `Hvi` class from the `Sequencer` object instance containing the information about the programmed Sequences.

After a successful Sequencer compilation, the Sequencer configuration is passed into the `Hvi` object when it is created. Once the `Hvi` object is created, the Sequencer instance is fixed. You cannot make any changes to the Sequencer or `SystemDefinition` instances after this point.

The compilation generates binary files that can be loaded to hardware and execute your HVI. Before running the HVI, you can redefine the initial values of some of the resources that are included in the HVI, such as HVI registers for different Engines.

Object Instances in the HVI Use Model

The following diagram shows the stages and highlights how each of the stages in the HVI use model creates and uses an object instance which is then passed to the next stage:



NOTE

Once an instance of SystemDefinition, Sequencer, or Hvi classes is created and configured, you cannot modify it in the next HVI step. If you attempt to modify one of these instances at a later stage, the modifications will not apply to your HVI. That is:

- You shall not modify the **SystemDefinition** object at the "Program HVI Sequences" or "Execute HVI" stage.
- You also shall not modify the **SystemDefinition** or **Sequencer** instances at the "Execute HVI" stage.

Correct Example

In the following example the value `non_hvi_core_clocks` in `SystemDefintion` is set.

This is set before the Sequencer is created so this is the correct place to do this.

```

# Define SystemDefinition
my_system = kthvi.SystemDefinition("MySystem");
#
# Set value of non_hvi_core_clocks (in Hz)
sys_def.non_hvi_core_clocks = [10e6]
#
# Create the sequencer
sequencer = kthvi.Sequencer("MySequencer", my_system); .....
#
# Get the Hvi
hvi = sequencer.compile().
  
```

Incorrect Example

In the following example the value `non_hvi_core_clocks` in `SystemDefintion` is set. In this case the value is set *after* the Sequencer is defined.

This example will not work because you cannot change a value in `SystemDefintion` after you have created the Sequencer.

```
# Define SystemDefinition
my_system = kthvi.SystemDefinition("MySystem");
#
# Create the Sequencer
sequencer = kthvi.Sequencer("MySequencer", my_system); .....
#
# Set value of non_hvi_core_clocks (in Hz)
sys_def.non_hvi_core_clocks =[10e6] # THIS FAILS
#
# Get the Hvi
hvi = sequencer.compile().
```

NOTE

If you need to make a change to `SystemDefintion` object after creating the Sequencer, you must create a new Sequencer for the change to have an effect.

1 Create the SystemDefinition

Setting up the HVI requires several steps:

- Include the HVI Library in your Application
- Define the Hardware in your HVI
 - Define the chassis
 - Define the chassis interconnects
 - Define the synchronization resources
 - Define the clocks
- Define and Configure HVI Resources
 - Define HVI Engines
 - Define HVI Actions
 - Define HVI Events
 - Define HVI Triggers
 - Define FPGA-Sandbox resources

Include the HVI Library in your Application

Include the HVI library in your application:

Python code:

```
import keysight_tse as kthvi
```

C# code:

```
using Keysight.Tse;
```

You must first create an instance of a `SystemDefinition` object.

Python code:

```
# Create SystemDefinition instance  
my_system = keysight_tse.SystemDefinition("Multi-chassis Setup")
```

C# code:

```
// Create SystemDefinition instance  
var sysDef = new SystemDefinition("My System");
```

When you have done this, specify the hardware and hardware resources that you require in your HVI:

- Define the hardware in your HVI.
- Define the HVI resources.
- Add the resources to the relevant collections.
- Initialize HVI hardware resources for the HVI.

Define the Hardware in your HVI

Add the hardware resources in your system to the `SystemDefinition` object, including:

- Chassis.
- Chassis interconnections.
- PXI Trigger synchronization resources.
- Synchronization clocks.

Define the chassis

Python code:

```
# Add chassis with number or options
my_system.chassis.add(chassis_number)
my_system.chassis.add(chassis_number, "DriverSetup=model=GenericPxieChassis")
```

C# code:

```
// Add chassis with number or options
sysDef.Chassis.Add(1);
sysDef.Chassis.Add(1, "Simulate=True,DriverSetup=model=GenericPxieChassis");
```

Define the chassis interconnects

You must first define the SystemSync modules. The options specify a number of parameters about each module:

Python code:

```
# Define SystemSync Modules
ssm_m9032_resource_id_ssm_1 = 'PXI0::CHASSIS1::SLOT10::INSTR'
ssm_m9033_resource_id_ssm_2 = 'PXI0::CHASSIS2::SLOT10::INSTR'
ssm_m9032_options = "Simulate=true,DriverSetup=Model=M9032A, HviEngineIpVersion=1.1.0,
HviGatewayFeatureVersion=2,model=M9032"
ssm_m9033_options = "Simulate=true,DriverSetup=Model=M9033A, HviEngineIpVersion=1.1.0,
HviGatewayFeatureVersion=2,model=M9033"
system_sync_modules_descriptors = [SystemSyncModuleDescriptor
('PXI0::CHASSIS1::SLOT10::INDEX0::INSTR', ssm_options)]
```

C# code:

```
// Define SystemSync Modules
public static string Ssm9032Options { get; set; } = "Simulate=true,DriverSetup=Model=M9032A,
HviEngineIpVersion=1.1.0, HviGatewayFeatureVersion=2,model=M9032"
public static string Ssm9033Options { get; set; } = "Simulate=true,DriverSetup=Model=M9033A,
HviEngineIpVersion=1.1.0, HviGatewayFeatureVersion=2,model=M9033"
public List<SystemSyncModuleDescriptor> SystemSyncModulesDescriptors { get; set; } = new
List<SystemSyncModuleDescriptor>
{
    new SystemSyncModuleDescriptor("PXI0::CHASSIS1::SLOT10::INDEX0::INSTR", Ssm9032Options),
    new SystemSyncModuleDescriptor("PXI0::CHASSIS2::SLOT10::INDEX0::INSTR", Ssm9033Options),
};
```

You must add the modules to the interconnects collection within the SystemDefinition:

Python code:

```
# Add SystemSync Modules to chassis
ssm_m9032 = my_system.interconnects.add_sync_module(ssm_m9032_resource_id, ssm_m9032_options)
ssm_m9033 = my_system.interconnects.add_sync_module(ssm_m9033_resource_id, ssm_m9033_options)
```

C# code:

```
// Add SystemSync Modules to chassis
ssmList.Add(interconnects.AddSyncModule(descriptor.ResourceId, descriptor.Options));
```

Once done, get the interface objects for each of the SSM connectors:

NOTE

The items in the collection `systemsync_downstream` are indexed from 0.

Python code:

```
# Get the 8x SystemSync downstream connector on first SSM
ssm_m9032_down = ssm_m9032.connectivity.systemsync_downstream[0]

# Get the 8x SystemSync upstream connector on second SSM
ssm_m9033_up = ssm_m9033.connectivity.systemsync_upstream[0]
```

C# code:

```
// Get the 8x SystemSync downstream connector on first SSM
ssm1Down = ssm1.Connectivity.SystemsyncDownstream[0]

// Get the 8x SystemSync upstream connector on second SSM
ssm2Up = ssm2.Connectivity.SystemsyncUpstream[0]
```

Set the connection between the connectors. This tells the HVI that these connections are connected together.

Python code:

```
# Set the connection
ssm_m9032_down.set_connection(ssm_m9033_up)
```

C# code:

```
// Set the connection.
ssm1.Connectivity.SystemSyncDownstream[0].SetConnection(ssm2.Connectivity.SystemSyncUpstream
[0]);
```

Define the synchronization resources

Python code:

```
# Define sync resources
my_system.sync_resources = [keysight_tse.TriggerResourceId.PXI_TRIGGER0,
                             keysight_tse.TriggerResourceId.PXI_TRIGGER1,
                             keysight_tse.TriggerResourceId.PXI_TRIGGER2]
```

C# code:

```
// Define sync resources
TriggerResourceId[] resources = {
    TriggerResourceId.PxiTrigger0,
    TriggerResourceId.PxiTrigger1,
    TriggerResourceId.PxiTrigger2};
```

Define the clocks

In simple setups this is only required when dealing with instruments that do not support HVI technology, or *Devices Under Test* that have specific clocking requirements.

For more complex setups see the *System Setup Guide*.

Python code:

```
# clocks configuration
my_system.non_hvi_core_clocks = [100MHz]
my_system.non_hvi_system_clocks = [500MHz]
```

C# code:

```
// clocks configuration
sysDef.NonHviCoreClocks = {100};
sysDef.NonHviCoreClocks = {500};
```

Define and Configure HVI Resources

Triggers, Actions, and Events are all HVI resources that can be used by the HVI Engine and the HVI Sequence to interact with the outside world, that is, with other instruments, the instrument sandbox, or any other external entities.

You must first define the Engines, then add the resources you are going to use to the relevant collections in the Engines you want to use them with. You must do this for the following types of resources:

- HVI Engines.
- Actions.
- Events.
- Triggers.
- FPGA-Sandbox resources.

Define HVI Engines

First, you must add the Engines you want to use to an Engine collection. The method `add_engine()` returns an Engine.

Python code:

```
# Add Engines
engine0 = my_system.engines.add(module.hvi.engines.main_engine, "Receiver")
engine1 = my_system.engines.add(module.hvi.engines.main_engine, "Transmitter")
```

C# code:

```
// Add Engines
sysDef.Engines.Add(module.Hvi.Engines.MainEngine, "Receiver");
sysDef.Engines.Add(module.Hvi.Engines.MainEngine, "Transmitter");
```

The procedure for adding the other HVI resources follows the same pattern. As a first step, the resource must be added to the corresponding collection using the method `add()` within the properties `TriggerCollection`, `ActionCollection`, `EventCollection`, etc.

For example, for an Event, do the following:

There is an Event collection for each Engine. Get the Event collection with the property `engine.events`. This returns the `EventCollection` object. Add the Events you want to use to the Event collection with the `add()` method of `EventCollection`. To add each Event you must specify both an `event id` and an `event name`:

Python code:

```
my_event = engine.events.add(module.hvi.events.PXI0, "My Event")
```

C# code:

```
myEvent = Engine.Events.Add(module.Hvi.Events.Pxi0, "My Event")
```

Actions, Triggers, and FPGA-Sandboxes all have their own collection properties, for example `ActionCollection` is for Actions.

Use the same procedure to get collections and add Actions, Triggers, and FPGA-Sandboxes to their respective collections. The ID of Engines, Actions, Events, and Triggers related to a specific instrument are defined by the instrument API, typically within the `instrument.hvi` interface of an `instrument` object.

Define HVI Actions

The following code example defines all HVI Actions necessary to perform AWG (*Arbitrary Waveform Generator*) Trigger operations. The AWG Trigger Actions for each AWG channel is added to the `ActionCollection` of the AWG Engine that needs to execute them in its Local Sequence.

Python code:

```
# Define AWG Trigger Actions for all AWG channels
for ch_index in range(1, num_channels + 1):
#
# Actions need to be added to the Engine's Action list so that they can be executed
action_name = "AWG Trigger CH" + str(ch_index)          # arbitrary user-defined name
instrument_action = "awg{}_trigger".format(ch_index)    # name decided by instrument API
action_id = getattr(instrument.hvi.Actions, instrument_action)
my_system.engines[awg_engine_name].actions.add(action_id, action_name)
```

C# code:

```
// Define AWG Trigger Actions for 4 AWG channels
// Actions must be added to the Engine's Action list so that they can be executed
//
mySystem.Engines[engineName].Actions.Add(module.Hvi.Actions.Awg0Trigger, "awg0trigger")
mySystem.Engines[engineName].Actions.Add(module.Hvi.Actions.Awg1Trigger, "awg1trigger")
mySystem.Engines[engineName].Actions.Add(module.Hvi.Actions.Awg2Trigger, "awg2trigger")
mySystem.Engines[engineName].Actions.Add(module.Hvi.Actions.Awg3Trigger, "awg3trigger")
```

Define HVI Events

The code example below adds the AWG CH1 Waveform Start Event to the Event collection of an M320xA AWG's HVI Engine object called `awg_engine`. For further information on M320xA Events see SD1 3.x Software for M320xA / M330xA Arbitrary Waveform Generators User's Guide available at [M3201A PXIe Arbitrary Waveform Generator](#).

Python code:

```
wfm_start_event = awg_engine.events.add(instrument.hvi.events.awg1_waveform_start, "AWG CH1 Wfm Start Event")
```

C# code:

```
// adding wait for Trigger Event
wfmStartEvent = awgEngine.Events.Add(instrument.Hvi.Events.Awg1WaveformStart, "AWG CH1 Wfm Start Event")
```

Define HVI Triggers

The code example below defines a *Front Panel* (FP) Trigger to be used by a digitizer instrument. The `TriggerCollection` is accessed through the `dig_engine.triggers` interface, where `dig_engine` is an HVI Engine object.

Python code:

```
# Defines the FP Trigger to be used as a wait condition by the digitizer
# Add to the HVI Trigger Collection of each HVI Engine the FP Trigger object of that same
instrument
#
fp_trigger_id = instrument.hvi.triggers.front_panel_1
fp_trigger = dig_engine.triggers.add(fp_trigger_id, "FP Trigger")
#
# Trigger configuration
# NOTE: Trigger to be used as WaitEvent conditions must be configured as kthvi.Direction.INPUT
# DriveMode (e.g. PushPull/OpenDrain) is defined by the instrument and cannot be changed by the
user
fp_trigger.config.direction = kthvi.Direction.INPUT
fp_trigger.config.polarity = kthvi.Polarity.ACTIVE_HIGH
fp_trigger.config.hw_routing_delay = 0
fp_trigger.config.trigger_mode = kthvi.TriggerMode.LEVEL
```

C# code:

```
// Defines the FP Trigger to be used as a wait condition by the digitizer
// Add to the HVI Trigger Collection of each HVI Engine the FP Trigger object of that same
instrument
//
fpTriggerId = instrument.Hvi.Triggers.frontPanel1;
fpTrigger = digEngine.Triggers.Add(fpTriggerId, "FP Trigger");
//
// Trigger configuration
// NOTE: Trigger to be used as WaitEvent conditions must be configured as Direction.Input
// DriveMode (e.g. PushPull/OpenDrain) is defined by the instrument and cannot be changed by the
user
fpTrigger.Config.Direction = Direction.Input;
fpTrigger.Config.Polarity = Polarity.ActiveHigh;
fpTrigger.Config.HwRoutingDelay = 0;
fpTrigger.Config.TriggerMode = TriggerMode.Level;
```

Define FPGA-Sandbox resources

The `SandboxCollection` is accessible through the `engine.fpga_sandboxes` interface of an Engine object. Unlike other HVI collections, this collection is already populated by a number of FPGA-Sandboxes where the number of FPGA-Sandboxes depends on the instrument being used. Most instruments have a single sandbox region in their FPGA, but some instruments might have multiple sandbox regions. FPGA-Sandbox objects do not need to be added to the collection, you only need to access them.

Python code:

```
# NOTE: The M3xxxA_sandbox name is not arbitrary and cannot be changed.
# The sandbox name is defined by each instrument. See SD1 3.x M3xxxA documentation for further
info
sandbox_name = 'sandbox0'
awg_sandbox = awg_engine.fpga_sandboxes[sandbox_name]
```

C# code:

```
// NOTE: The M3xxxA_sandbox name is not arbitrary and cannot be changed.
// The sandbox name is defined by each instrument. See SD1 3.x M3xxxA documentation for further
info
sandboxName = "sandbox0";
awgSandbox = AwgEngine.FpgaSandboxes[sandboxName];
```

2. Program HVI Sequences

Programming HVI Sequences requires a number of steps:

- [Create a Sequencer Object](#)
- [Create Sequences](#)
- [Start with the Global Sync Sequence](#)
- [Adding Sync Statements and Sync Sequences](#)
- [Adding Local Statements](#)
- [Adding HVI Instructions](#)
- [Using Triggers, Actions, and Events](#)
- [Using FPGA-Sandbox Resources](#)

Create a Sequencer Object

Before you can begin writing Sequences, you must create a [Sequencer](#) object and pass the [SystemDefinition](#) to the [Sequencer](#) object:

Python code:

```
sequencer = keysight_tse.Sequencer("sequencer", my_system)
```

C# code:

```
Sequencer seq = new Sequencer("sequencer", sysDef);
```

Define HVI Registers and Initialize Register Values

Define the HVI Registers resource you require in each Engine and use the `add()` method to add them to the register collection for that Engine. Then define their initial values:

Python code:

```
loop_register = sequencer.sync_sequence.scopes["Engine 1"].registers.add("Loop Register",  
keysight_tse.RegisterSize.SHORT)  
loop_register.initial_value = 0
```

C# code:

```
var loopRegister = sequencer.SyncSequence.Scopes["Engine 1"].Registers.Add("Loop Register",  
RegisterSize.SHORT);  
loopRegister.InitialValue = 0;
```

The registers that you to use in the HVI Sequences must be added to the register collection within the scope of the corresponding HVI Sequence. This can be done using the `RegisterCollection` property that is within the `Scope` object corresponding to each Sequence. HVI Registers belong to a specific HVI Engine because they refer to hardware registers of that specific instrument. Registers from one HVI Engine cannot be used by other Engines or outside of their scope. Registers can only be added to the HVI top Sync Sequence scopes. This means that you can only add global registers that are visible in all child Sequences. The number and size of registers is defined by each instrument.

To reserve a register resource:

1. Get the register collection from the Engine you want to reserve the register on.
2. Add the registers you require. Use the `add()` method to the register collection for that Engine

NOTE Register size is defined by the following:

- SHORT = 32 bit
- LONG = 48 bit

Create Sequences

After you have got the `Sequencer` object and set up the registers you require, you can write the program the HVI executes, this is composed of:

- Sequences.
- Statements.
- Instructions.
- Time restrictions.

To define your program, you must:

- Create Sequences.
- Add Statements and instructions.

Start with the Global Sync Sequence

When HVI starts execution, it starts in a Global Sync Sequence, this is defined by the `Sequencer` object. This is used in the previous example when the registers were reserved:

Python code:

```
engine_1_registers = sequencer.sync_sequence.scopes["Engine 1"].registers
```

C# code:

```
var engine1Registers = seq.SyncSequence.Scopes[engine1Name].Registers;
```

Adding Sync Statements and Sync Sequences

You add Sync Statements to the `SyncSequence` property with `add_statement()` methods such as `SyncSequence.add_sync_while()`:

Python code:

```
# Create Sync While Statement (loop_register < SYNC_WHILE_LOOP_ITERATIONS):
SYNC_WHILE_LOOP_ITERATIONS = 5
sync_while_condition = keysight_tse.Condition.register_comparison( engine_1_registers["loop_
register"], keysight_tse.ComparisonOperator.LESS_THAN, SYNC_WHILE_LOOP_ITERATIONS)
sync_while = sequencer.sync_sequence.add_sync_while("sync_while", 100, sync_while_condition)
```

C# code:

```
// create Sync While Statement (loop_register < SYNC_WHILE_LOOP_ITERATIONS)
var syncWhileCondition = Condition.RegisterComparison(
engine1Registers["loop_register"], ComparisonOperator.LessThan, SYNC_WHILE_LOOP_ITERATIONS);
var syncWhile = seq.SyncSequence.AddSyncWhile("sync_while", 100, syncWhileCondition);
```

You can also add Sync Sequences within the global Sync Sequence and add Sync Statements within the Sync Sequences.

Adding Local Statements

To add HVI Instructions or Local Flow-Control Statements, you must add them within a Sync Multi-Sequence Block. You must add this Sync Multi-Sequence Block within a Sync Sequence by using the `add_sync_multi_sequence_block()` method:

Python code:

```
# Add a Sync Multi-Sequence Block:
multi_seq_block_1 = sync_while.sync_sequence.add_sync_multi_sequence_block("multi_seq_block_1",
210)
```

C# code:

```
// Add a Sync Multi-Sequence Block
var multiSeqBlock1 = syncWhile.SyncSequence.AddSyncMultiSequenceBlock("multiSeqBlock1", 220);
```

To add the Local Statements, you must get a `Sequence` object for each Engine in the Sync Multi-Sequence Block and add them using the corresponding `add_XXX()` method. HVI Instructions can be added to a Sync Multi-Sequence Block using the `add_instruction()` method. For each instruction parameter, use the `set_parameter()` method to set it.

By adding Local Statements to the Sequences, you define the Local Sequence that each local Engine executes in parallel with the other Engines.

Adding HVI Instructions

There are two types of HVI instructions:

- HVI Native Instructions.
- Instrument-Specific HVI Instructions.

HVI Native Instructions

The `InstructionSet` property contains the set of native instructions that can be executed within an HVI Statement, including:

- Register arithmetic.
 - Add / Subtract.
 - Assign.
- Read/write I/O Trigger ports.
- Communications operations with the instrument sandbox using an HVI Host Interface.
 - FPGA-Register read/write.
 - FPGA-Memory Map read/write (`fpga_array_read` and `fpga_array_write`).
- Action Execute.
- Trigger Write.

To use the HVI Native Instructions, you must use the `InstructionSet` property. You get this from the Local `Sequence` object:

Python code:

```
# Initialize loop_register
loop_reg = multi_seq_block.scope.registers["loop_register"]
awg_sequence = multi_seq_block.sequences["AWG Engine"]
instruction_a = multi_seq_block.add_instruction("loop_register = 0", 10, awg_
sequence.instruction_set.assign.id)
instruction_a.set_parameter(awg_sequence.instruction_set.assign.destination.id, loop_reg)
instruction_a.set_parameter(awg_sequence.instruction_set.assign.source.id, 0)
#
# Increment pulse_counter
pulse_counter = multi_seq_block_1.scope.registers["pulse_counter"]
instruction = multi_seq_block_1.add_instruction("Increment Pulse Counter", 10, awg_
sequence.instruction_set.add.id)
instruction.set_parameter(awg_sequence.instruction_set.add.destination.id, pulse_counter)
instruction.set_parameter(awg_sequence.instruction_set.add.left_operand.id, pulse_counter)
instruction.set_parameter(awg_sequence.instruction_set.add.right_operand.id, 1)
```

C# code:


```
// Initialize loop_register
var reg = sequence.Scope.Registers[registerName];
var instructionA = sequence.AddInstruction(registerName + "_assign", startDelay,
sequence.InstructionSet.Assign.Id);
instructionA.SetParameter(sequence.InstructionSet.Assign.Value.Id, value);
instructionA.SetParameter(sequence.InstructionSet.Assign.Destination.Id, reg);
//
// Increment register by 1
private void incrementRegisterBy1(ISequence sequence, string registerName, int startDelay)
{
    var reg = sequence.Scope.Registers[registerName];
    var instructionA = sequence.AddInstruction("Increment Pulse Counter",
startDelay, sequence.InstructionSet.Add.Id);
    instructionA.SetParameter(sequence.InstructionSet.Add.LeftOperand.Id, reg);
    instructionA.SetParameter(sequence.InstructionSet.Add.RightOperand.Id, 1);
    instructionA.SetParameter(sequence.InstructionSet.Add.Destination.Id, reg);
}
```

Instrument-Specific HVI Instructions

Instrument-Specific HVI Instructions are described in the documentation for the instrument. For example, the following code shows how to set a channel amplitude value:

Python code:

```
# Set CH1 amplitude to 1.0 V:
instruction = multi_seq_block_1.add_instruction("Set CH1 amplitude to 1.0 V", 10,
instrument.hvi.instruction_set.set_amplitude.id)
instruction.set_parameter(instrument.hvi.instruction_set.set_amplitude.channel.id, ch1)
instruction.set_parameter(instrument.hvi.instruction_set.set_amplitude.value.id, 1.0)
```

C# code:

```
// Set CH1 amplitude to 1.0 V
instruction = multiSeqBlock1.AddInstruction("Set CH1 amplitude to 1.0 V", 10,
instrument.Hvi.InstructionSet.SetAmplitude.id);
instruction.SetParameter(instrument.Hvi.InstructionSet.SetAmplitude.Channel.id, ch1);
instruction.SetParameter(instrument.Hvi.InstructionSet.SetAmplitude.Value.id, 1.0);
```

Using Triggers, Actions, and Events

The examples below provide an overview of how to use Triggers, Actions and Events within an HVI Sequence.

Using Triggers

There are two typical use cases of Trigger objects (previously defined by the user during system definition). The first one is the usage of the Trigger object as a wait condition inside a Wait Statement:

Python code:

```
# Add a wait Statement that has a FP Trigger as a condition
fp_trigger = awg_engine.triggers["fp_trigger"]
wait_condition = keysight_tse.Condition.trigger(fp_trigger)
wait_event = awg_sequence.add_wait("wait for fp trigger", 10, wait_condition)
wait_event.set_mode(kthvi.WaitMode.TRANSITION, kthvi.SyncMode.IMMEDIATE)
```

C# code:

```
// Add a wait Statement that has a FP Trigger as a condition
fpTrigger = awgEngine.Triggers["fpTrigger"];
waitCondition = Condition.Trigger(fpTrigger);
waitEvent = awgSequence.AddWait("wait for trigger", 10, waitCondition);
waitEvent.SetMode(WaitMode.Transition, SyncMode.Immediate);
```

The second use case involves the `TriggerWrite` HVI Native instruction, where the Trigger object can be used to specify which electrical Trigger line can be written from the HVI Sequence:

Python code:

```
# Write FP Trigger to ON value
fp_trigger = awg_engine.triggers["fp_trigger"]
trigger_write_instr = sequence.instruction_set.trigger_write
instr_trigger_ON = sequence.add_instruction("FP Trigger ON", 10, trigger_write_instr.id)
instr_trigger_ON.set_parameter(trigger_write_instr.sync_mode.id, trigger_write_instr.sync_mode.immediate)
instr_trigger_ON.set_parameter(trigger_write_instr.trigger.id, fp_trigger)
instr_trigger_ON.set_parameter(trigger_write_instr.value.id, trigger_write_instr.value.on)
```

C# code:

```
// Write FP Trigger to ON value
var tw = sequence.InstructionSet.TriggerWrite;
var instOn = sequence.AddInstruction("Trigger On", 20, tw.Id);
instOn.SetParameter(tw.Trigger.Id, trigger);
instOn.SetParameter(tw.SyncMode.Id, tw.SyncMode.Immediate);
instOn.SetParameter(tw.Value.Id, tw.Value.On);
```

Using Actions

User-defined Actions can be executed using the HVI native instruction `ActionExecute`. A list of Actions `action_list`, can be executed simultaneously within the same instruction. The `action_list` object must have been previously defined.

Python code:

```
# "Action Execute" instruction executes the AWG Trigger from HVI
instruction = sequence.add_instruction("AWG trigger", 10, sequence.instruction_set.action_
execute.id)
instruction.set_parameter(sequence.instruction_set.action_execute.action.id, action_list)
```

C# code:

```
// "ActionExecute" instruction executes the AWG Trigger from HVI
var actionArray = sequence.Engine.Actions.ToArray();
instruction = sequence.AddInstruction("AWG trigger", 10,
sequence.InstructionSet.ActionExecute.id);
instruction.SetParameter(sequence.InstructionSet.ActionExecute.Action.id, actionArray);
```

Using Events

The typical use case of Events within HVI Sequences is as a condition for a Wait Statement:

Python code:

```
# Add a wait Statement that waits for AWG CH1 queue to be empty
awg_queue_empty = awg_engine.events["Awg1QueueIsEmpty"]
wait_condition = keysight_tse.Condition.event(awg_queue_empty)
wait_event = awg_sequence.add_wait("Wait for AWG Queue to be Empty", 10, wait_condition)
wait_event.set_mode(kthvi.WaitMode.TRANSITION, kthvi.SyncMode.IMMEDIATE)
```

C# code:

```
// adding wait for Trigger
var waitTrigger = sequence.Engine.Triggers["wait_trigger"];
var waitEvent = sequence.AddWait("wait for trigger", 10, Condition.Trigger(waitTrigger));
waitEvent.SetMode(WaitMode.Transition, SyncMode.Immediate);
```

Using FPGA-Sandbox Resources

To use FPGA-Sandbox resources, the sandbox must be loaded using the `load_from_k7z()` method specifying the path containing the `.k7z` file produced compiling a project designed using PathWave FPGA. For more information see the PathWave FPGA User Manual at [PathWave FPGA](#) . Once the FPGA-Sandbox is loaded, all the HVI FPGA-Registers and HVI FPGA-Memory Maps that were inserted in the specified PathWave FPGA project file can be accessed to be used in the HVI Sequence. Please note that the same names used in the PathWave FPGA project must be used to access the FPGA resources. In the following example, the FPGA-Register name `Register_Bank_MyCounter` is not arbitrary but assumed to be taken from the PathWave FPGA project that generated the file `MySandboxProject.k7z` :

Python code:

```
sandbox = engine.fpga_sandboxes["sandbox0"]
sandbox.load_from_k7z("MySandboxProject.k7z")
counter_register = sandbox.fpga_registers["Register_Bank_MyCounter"]
```

C# code:

```
sandbox = Engine.FpgaSandboxes["sandbox0"];
sandbox.LoadFromk7z("MySandboxProject.k7z");
counterRegister = sandbox.FpgaRegisters["registerBankMyCounter"];
```

Write to FPGA resources

The following example shows how to write to an FPGA-Register and read an FPGA array. The process in both cases is very similar:

Python code:

```
# Write FPGA-Register
fpga_register = engine.fpga_sandboxes[sandbox_name].fpga_registers[register_name]
fpga_regw_instruction = sequence.instruction_set.fpga_register_write
fpga_register_write = sequence.add_instruction("my_fpga_register_write", 10, fpga_regw_
instruction.id)
fpga_register_write.set_parameter(fpga_regw_instruction.fpga_register.id, fpga_register)
fpga_register_write.set_parameter(fpga_regw_instruction.value.id, value_register)
#
# Read FPGA array
memory_map = engine.fpga_sandboxes[sandbox_name].fpga_memory_maps[0]
fpga_arrayr_instr = sequence.instruction_set.fpga_array_read
fpga_array_read = sequence.add_instruction("my_fpga_array_read", time_ns, fpga_arrayr_instr.id)
fpga_array_read.set_parameter(fpga_arrayr_instr.fpga_memory_map.id, memory_map)
fpga_array_read.set_parameter(fpga_arrayr_instr.fpga_memory_map_offset.id, loop_reg)
fpga_array_read.set_parameter(fpga_arrayr_instr.value.id, value_register))
```

C# code:

```
// Write FPGA-Register
fpga_register = engine.fpga_sandboxes[sandbox_name].fpga_registers[register_name];
fpga_regw_instruction = sequence.instruction_set.fpga_register_write;
fpga_register_write = sequence.add_instruction("my_fpga_register_write", 10, fpga_regw_
instruction.id);
fpga_register_write.set_parameter(fpga_regw_instruction.fpga_register.id, fpga_register);
fpga_register_write.set_parameter(fpga_regw_instruction.value.id, value_register);
//
// Read FPGA array
memoryMap = Engine.fpgaSandboxes[sandbox_name].fpgaMemoryMaps[0];
fpgaArrayrInstr = sequence.InstructionSet.FpgaArrayRead;
fpgaArrayRead = sequence.AddInstruction("myFpgaArrayRead", timeNs, fpgaArrayrInstr.id);
fpgaArrayRead.SetParameter(fpgaArrayrInstr.FpgaMemoryMap.id, memoryMap);
fpgaArrayRead.SetParameter(fpgaArrayrInstr.FpgaMemoryMapOffset.id, loopReg);
fpgaArrayRead.SetParameter(fpgaArrayrInstr.Value.id, valueRegister));
```

3. Compile Your Sequences

After writing the Sequences, you must add the command that compiles the HVI. Call the `compile()` method in the `Sequencer` object to perform the compilation operation. The `compile()` method returns the HVI instance `Hvi`.

Python code:

```
# Compile HVI Sequences:
try:
    hvi = sequencer.compile()
    print('HVI Compiled')
except keysight_tse.CompilationFailed as err:
    print(err.compile_status.to_string())
    raise err
```

C# code:

```
// Compile HVI Sequences:
try
{
    hvi = sequencer.Compile();
    Console.WriteLine("compile DONE");
}
catch (CompilationFailed err)
{
    Console.WriteLine(err.CompileStatus.ToString());
    throw err;
}
```

NOTE

At this point you can no longer modify Sequences, Actions, Events or Triggers.

The property `hvi.sync_resources` provides information about the PXI sync resources you must reserve.

Python code:

```
print("This needs to reserve {} PXI trigger resources to execute".format(len(hvi.sync_
resources)))
```

C# code:

```
Console.WriteLine("This needs to reserve {} PXI trigger resources to execute".Format(len
(Hvi.SyncResources)));
```

If the compilation fails, the object `keysight_tse.CompilationFailed` is thrown. This contains compilation error messages that you can print.

4. Load To Hardware

Before your compiled Sequences can be executed, they must be uploaded into the HVI Engines in the instrument hardware. To upload the compiled Sequences, you must use the `Hvi` method `load_to_hw()`.

Python code:

```
# Load HVI to hardware:
Hvi.load_to_hw()
print("HVI Loaded to hardware")
```

C# code:

```
// Load HVI to hardware:
Hvi.LoadToHw();
Console.WriteLine("load DONE");
```

5. Modify Initial Register Values (Optional)

The HVI execution can be parameterized using Registers, the initial values of all Registers are updated when the `run()` method in `Hvi` is called. To modify the initial value of the Registers in the `Hvi` object, use:

Python code:

```
# Modify register initial value
value = 10
register_runtime = hvi.sync_sequence.scopes[0].registers[loop_register.name]
register_runtime.initial_value = value
```

C# code:

```
// Modify register initial value
var value = 10;
registerRuntime = Hvi.SyncSequence.Scopes[0].registers[loopRegister.name];
registerRuntime.initialValue = value;
```

Once the instrument has been loaded to hardware, you can write to the FPGA-Memory Map.

Python code:

```
memory_map.write(0, 1)
memory_map.write(1, 2)
memory_map.write(2, 3)
```

C# code:

```
memoryMap.Write(0, 1);
memoryMap.Write(1, 2);
memoryMap.Write(2, 3);
```


6. Execute Sequences

To execute the binaries, call the `run()` method in `Hvi`. The HVI can be run in a blocking or non-blocking mode:

Blocking mode

In blocking mode, the execution is blocked at the HVI execution code line for a fixed amount of time specified by the timeout input parameter. If `timeout = hvi.no_timeout` is used as an input parameter, the execution can be blocked until the HVI Sequences finish their execution.

Python code:

```
hvi.run(hvi.no_timeout)
```

C# code:

```
hvi.Run(System.TimeSpan.FromSeconds(10));
```

Non-blocking mode

In non-blocking mode, the execution is not blocked. This enables you to initiate a second HVI instance to run in parallel.

Python code:

```
# Execute HVI in non-blocking mode
# This mode allows SW execution to interact with HVI execution:
hvi.run(hvi.no_wait)
print("HVI Running...")
```

C# code:

```
// Execute HVI in non-blocking mode
// This mode allows SW execution to interact with HVI execution:
hvi.Run(IHvi.no_wait);
Console.WriteLine("HVI Running...");
```

While and after execution is finished, you can read or write registers and execute the binaries again.

Python code:

```
# Modify Register initial value
value = 20
register_runtime = hvi.sync_sequence.scopes[0].registers[loop_register.name]
register_runtime.initial_value = value
hvi.run(hvi.no_timeout)
```

C# code:

```
// Modify Register initial value
var value = 20;
registerRuntime = hvi.SyncSequence.Scopes[0].Registers[loopRegister.name];
registerRuntime.initialValue = value;
hvi.Run(IHvi.NoTimeout);
```

7. Release All Resources

To release all HVI resources and enable other applications or HVI instances to use the hardware, you must release the hardware. Your application cannot perform any operation with the hardware resources in the HVI after this point.

Python code:

```
# Unlock and release hardware resources:  
hvi.release_hw()  
print("Releasing Hardware...")
```

C# code:

```
// Unlock and release hardware resources:  
hvi.ReleaseHw();  
Console.WriteLine("Releasing Hardware...");
```

HVI Logging and Troubleshooting

PathWave Test Sync Executive comes with an integrated logger that you can use for troubleshooting.

The HVI Logger is aimed at producing information that is useful for support engineers. It provides information that is additional to the [Sequence Representation](#) output.

The logger has the following features:

- The level of logging is configurable.
- You can force flush messages.
- The output can be configured to go to the console or to an output file.
- You can configure the logger from environment Variables or in a `.env` configuration file.
- You can instruct some instruments to produce logs.

The logger can produce the following levels of logging information, where each level also includes all the information in the levels below it:

Logger level	Description
Trace	Produces trace information that is useful to support engineers.
Debug	Produces debug information that is useful to support engineers. This level also provides the Sequence Representation output.
Info	Produces generally useful information.
Warning	Logs anything that can potentially cause application oddities, but are automatically recovered.
Error	Logs any errors that are fatal to an operation, but not the service or application.
Fatal	Logs any errors that forces a shutdown of the application.
Off	Does not log anything.

Logger Configuration

The logger is configured with environment Variables. The following table describes the Variables:

Environment Variable	Values	Description
<code>HVI_LOGGER_LEVEL</code>	<ul style="list-style-type: none"> • Trace • Debug • Info • Warning • Error • Fatal • Off 	<p>This value indicates the level of information printed to the log.</p> <p>The information printed out contains the information for the level specified and all of the levels below it. For example, if the level is set to Debug, all messages except Trace are printed to the log.</p> <p>By default, the level is set to Error, so only Error and Fatal are printed.</p>
<code>HVI_LOGGER_OUTPUT_PATH</code>	<p>Any existing valid path in your system, For example: <code>C:\tmp</code></p>	<p>This Variable disables console output and tells the logger to save the log to a file at the specified location.</p> <p>The file with the log messages is called: <code>HVILog_hviLogger_[num1]_[num2].log</code>, where <code>num1</code> is the date and time, and <code>num2</code> is the thread ID.</p>
<code>HVI_LOGGER_FORCE_FLUSH</code>	1 or 0	<p>This Variable forces the log messages to be flushed to the output every time a message is logged. Enable this if you want to troubleshoot a program that is crashing, so that all messages before the crash shall be written. Do not enable this option in any other cases, because it impacts the performance of the execution.</p>
<code>HVI_LOGGER_EXTENDED</code>	<p>"*", "ALL", or a comma separated list.</p> <p>For example: <code>M9032,M9546</code></p>	<p>This Variable enables the logging output of instruments managed by HVI.</p> <p>An output file for each instrument is generated in the path specified with <code>HVI_LOGGER_OUTPUT_PATH</code>.</p> <p>The file is saved as: <code>{MODEL}_{Chassis Slot for M903x}_{date}.log</code></p> <p>See the section <i>Logger Extended mode Supported Instruments</i> for a list of supported instruments.</p>

NOTE

By default the configuration for the logger is:

- Logging level: Error.
- Output: console.
- Force flush: disabled.
- Logger Extended: disabled.

.env Configuration File

The logger configuration can be also configured from a `.env` file. The configuration values are stored in the file as KEY=VALUE pairs and you can use `#` for comments.

The `.env` file must be located in the same folder as the HVI script to be executed. HVI parses the `.env` file and sets all the environment Variables found for that script.

The following shows an example `.env` file:

`.env`

```
# The hvi logger level: Trace, Debug, Info, Warning, Error, Fatal, Off.
HVI_LOGGER_LEVEL=Fatal
#
# Set this parameter to write the logs to a file instead of being printed to the console
HVI_LOGGER_OUTPUT_PATH=C:\tmp\hviLogs
#
# Set this parameter to force flush the log every new line instead of doing it at the end.
# This helps you to identify the line of code before a crash.
HVI_LOGGER_FORCE_FLUSH=0
#
# Activates the Logger for all HVI controlled instruments. The supported models are the
# System Synchronization Modules (M9032,M9033), the High Performance Reference Clock Source,
# (M9546)
# or "ALL"
HVI_LOGGER_EXTENDED=ALL
```

Logger Extended mode Supported Instruments

PathWave Test Sync Executive can control a number of different instruments. The environment Variable `HVI_LOGGER_EXTENDED` activates logging output from the instruments that support it. The way the logs are produced depends on the instruments, some instruments produce individual log files whereas other instruments combine log files together into a single file.

The supported models for release 2023 are:

Model	Description
M9546x	High Performance Reference Clock Source
M9032, M9033	System Synchronization Modules

NOTE

Some Instruments, like the Keysight M5000 PXIe family for example, might include native logging facilities that cannot be controlled by PathWave Test Sync Executive, for more information, see your instrument documentation.

Recommended Logger settings for contacting support

If you require support for PathWave Test Sync Executive, a log file will help the support team to rapidly diagnose any problems.

If you want to contact support, first generate a log with the following settings:

- `HVI_LOGGER_LEVEL=Trace`
- `HVI_LOGGER_FORCE_FLUSH=1`
- `HVI_LOGGER_OUTPUT_PATH=C:\Logs` or another path¹

¹ The path must be an existing valid path.

Chapter 9: TSE Service and Multi-Host support

This chapter describes TSE Service, the features it provides and how it enables Multi-Host support.

For detailed description of how to install and configure a TSE Service based system, see Chapter 9 in the *PathWave Test Sync Executive System Setup Guide*.

This chapter contains the following sections:

- [About TSE Service and Multi-Host Systems](#)
- [TSE Service Operation Modes](#)
- [TSE Service Configuration](#)
- [Accessing Remote Resources](#)
- [Using TSE Service in an Application](#)

About TSE Service and Multi-Host Systems

PathWave Test Sync Executive release 2023 includes TSE Service and supports Multi-host systems.

TSE Service Overview

TSE Service offer extended capabilities for system initialization and configuration:

1. Extend the use of TSE to Multi-Host architecture (see *Free-Running mode*). If you want a system with more than 6 PXIe chassis or remote connectivity, then you require a Multi-Host system.
2. Enables the user to define system configuration using one or more **.yaml** files to automate the complete system initialization (a time-consuming task), and execute it at host boot-time (see *Leader-Follower mode*). This can be used with both single and multi-host systems.

TSE Service Running Modes

TSE Service can be configured to operate in one of 2 different modes:

Free running Mode:

This mode is mainly intended for Multi-Host systems because it enables a client application to access resources distributed across multiple hosts. The client application is responsible for defining the topology of the system using the SystemDefinition class to add chassis, SSMs, HVI engines, etc, and to run the system initialization.

Leader-Follower Mode:

In this mode TSE Service in the leader host automates the complete system initialization, and this initialization is executed at boot time, speeding up the applications execution later. The system is defined in the leader in a **system_definition.yaml** configuration file. The client application does not need to define or initialize the system, it just creates the application SystemDefinition that connects to the leader TSE Service and gets all the system information automatically.

About Multi-Host Systems

There is a current limitation with PCs in that their bus systems can only connect to up to 6 PXI chassis per PC host. To enable highly scalable systems that extend to more than 6 chassis, Multi-Host capabilities are required to coordinate the operations in a single application across multiple hosts. TSE Service implements a remote access infrastructure that enables Multi-Host operation. In Multi-Host operation, a single HVI Sequence implemented in one application in one host, is deployed to all instruments across several hosts, to achieve a large-scale system that runs fully synchronized.

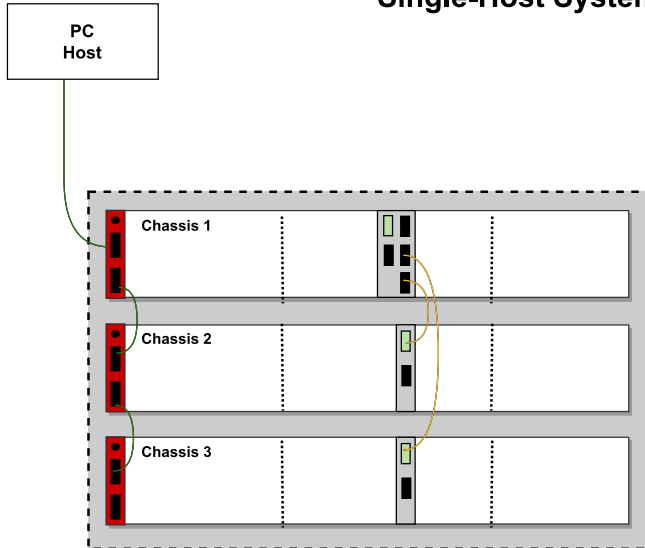
Multi-Host systems are similar to the Multi-chassis systems possible with TSE 2022, but they use a number of hosts to increase the number of chassis you can use. Each host is connected to a number

of chassis. As with the existing Multi-chassis systems, The chassis are connected together with *System Synchronization Modules (SSMs)*.

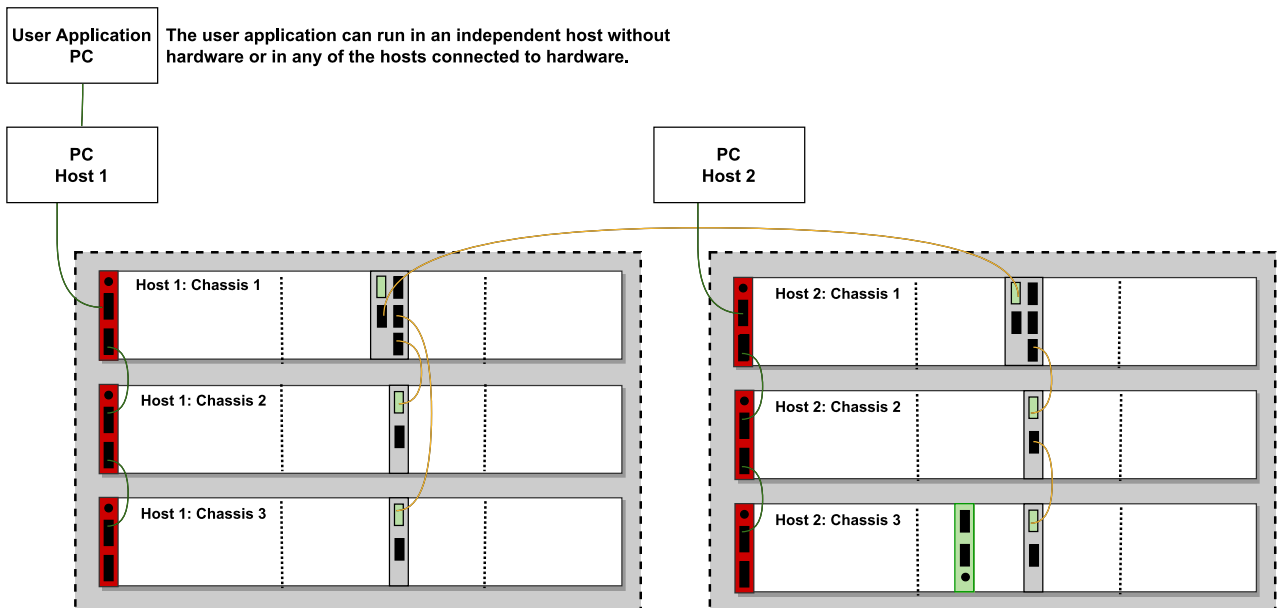
TSE Service provides the capabilities required to enable Multi-host systems. *Keysight Distributed Infrastructure (KDI)* is the underlying service that supports this feature. On top of KDI, PathWave Test Sync Executive adds support for very large systems using the same existing programming paradigm, enabling an easy transition for existing users.

The following diagram shows a Single-Host and a Multi-Host system:

Single-Host System



Multi-Host System



NOTE

For a Multi-Host system to work, the following must be installed:

- TSE Service
- [OPTIONALLY] Keysight Distributed infrastructure (KDI)

Typically, these are installed when you install PathWave Test Sync Executive. For more information, see *Chapter 2: Install PathWave Test Sync Executive* in this document and Chapter 9 in the *PathWave Test Sync Executive System Setup Guide*.

TSE Service Operation Modes

You can configure TSE Service to operate in one of 2 different modes for different use cases:

- Free-Running mode.
- Leader-Follower mode.

Free-Running Mode

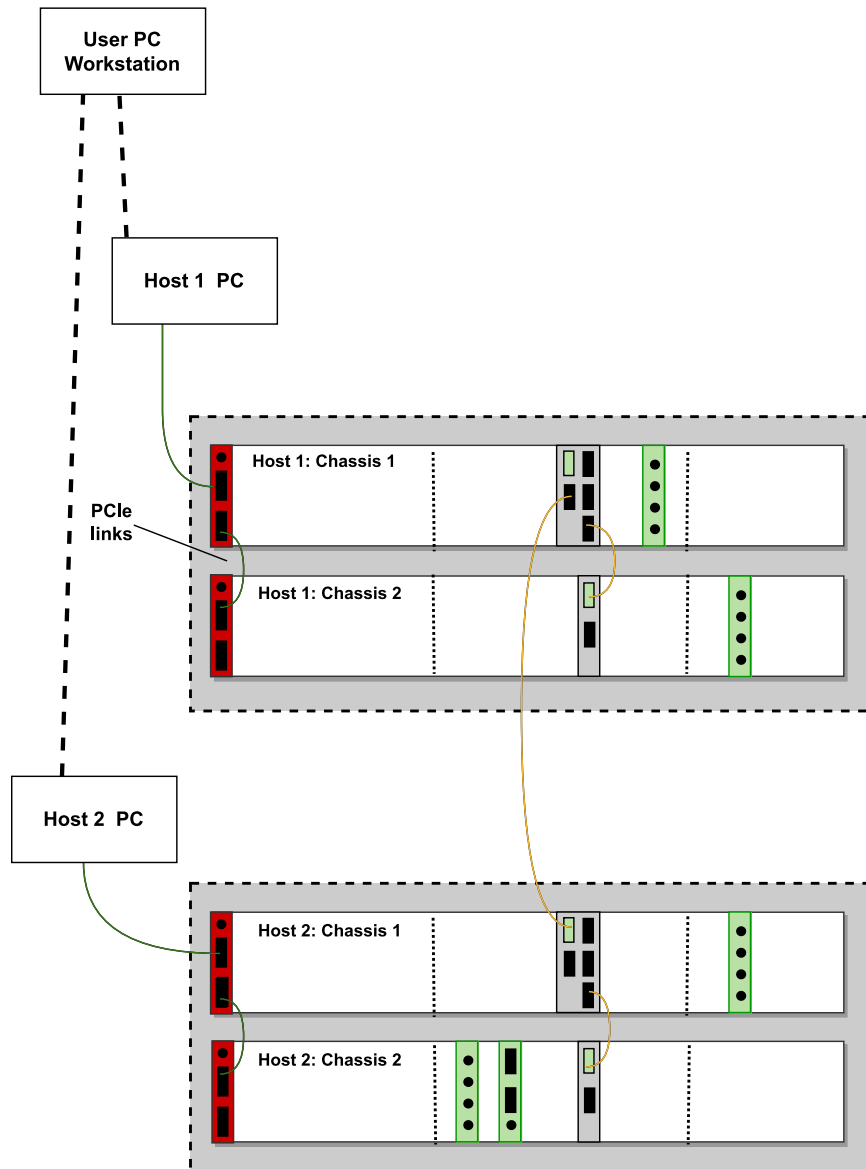
This mode is mainly intended for Multi-Host systems because it enables a client application to access resources distributed across multiple hosts. The client application is responsible for defining the topology of the system using the SystemDefinition class to add chassis, SSMs, HVI engines, etc. and to run the system initialization.

User applications can:

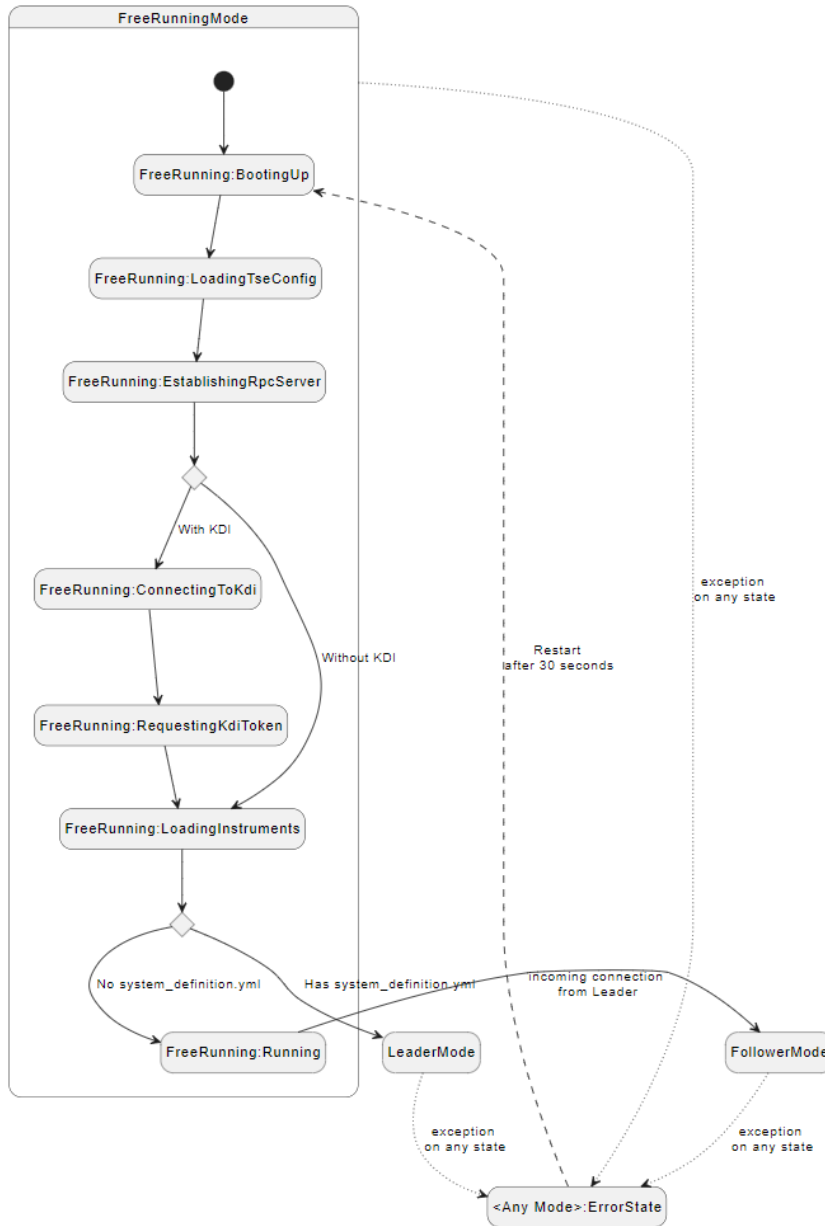
- Access remote resources like chassis and *System Synchronization Modules* (SSMs) using the TSE Resource IDs.
- Access remote instruments using the KDI Resource ID.
- Configure the system as required using the TSE and instruments API.

TSE Service Free-Running mode is configured by means of the `tse_config.yml` file located in each host. See [TSE Service Configuration](#) for more details.

The following diagram shows a Multi-Host system in Free-Running mode. In Free-Running mode the client (or user) application, that can run in a separate host other than the hosts connected to hardware, connects to the TSE Service in the different hosts to use the chassis and Sync Module resources in each host.



The following state machine diagram shows the TSE Service boot up sequence in Free-running mode:



Leader-Follower Mode

The Leader-Follower mode builds on the Free-Running mode to provide simplified usage for client applications. A system configured in Leader-Follower mode has a predefined system configuration that is fully initialized at start-up. This configuration is specified in the **system_definition.yml** configuration file located in one of the hosts, called the Leader TSE Service.


The Leader TSE Service system definition can include resources distributed in other hosts running TSE Service, called Follower TSE Services. In the `system_definition.yml` configuration file the user can specify:

- System Topology.
 - This includes information from the chassis and interconnects (SystemSync connectivity) across them.
- Clocking configuration.
- Optionally, the instruments to include in the automatic start-up initialization.

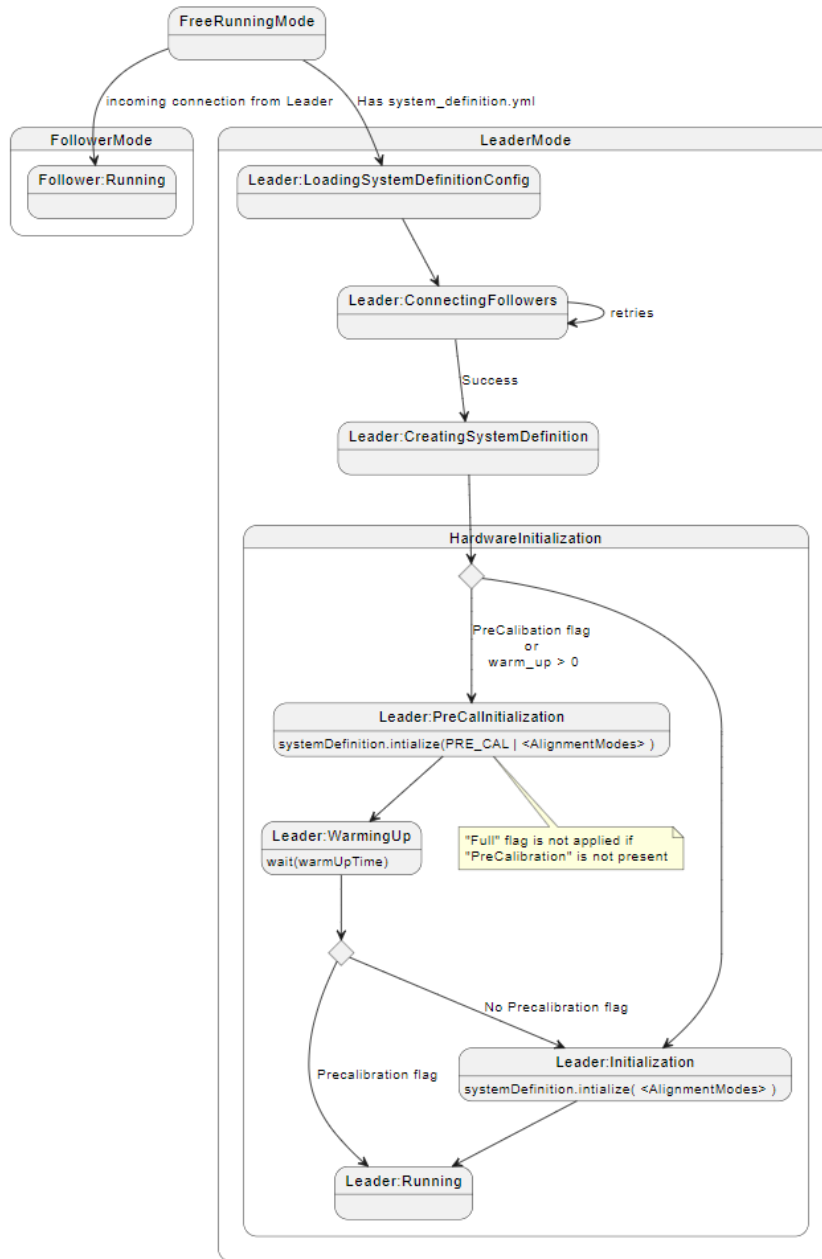
See TSE Configuration section for details on the `system_definition.yml` configuration options.

In this mode, client applications specify the Leader TSE Service when creating the application SystemDefinition instance and only need to specify the HVI Engines and related resources, but there is no need to specify topology and clocking because that's inherited from the Leader TSE Service. This is a simpler and quicker way to use a system because it enables client applications to exploit TSE real-time capabilities without having to first specify chassis, SSMs, clocks, etc. or having to wait for a full system initialization.

The following diagram shows a Multi-Host system in Leader-Follower mode. In Leader-Follower mode the client (or user) application, that can run in a separate host other than the hosts connected to hardware, connects to only to the Leader TSE Service.



The state diagram below shows the initialization steps once TSE Service has completed the Free-running initialization:



It is important to note that TSE Service always start in **Free-Running mode** and performs the free-running initialization per the local `tse_config.yml` file. It is in this step where all TSE Service hosts open chassis, SSMs, Instruments, etc. depending on the local `tse_config.yml` file configuration. Then, only if the `system_definition.yml` file is found locally, this TSE Service switch to the **Leader mode**, in this mode it creates a SystemDefinition instance including all elements specified in the `system_definition.yml`, topology, clocking, and instruments if specified. Before creating the SystemDefinition, the

Leader TSE Service connects to other TSE Services, the Follower TSE Services, that own the elements specified in the topology section. When a TSE Service received a connection request from the Leader, it switches from the Free-Running into the **Follower mode**, TSE Services accept only one Leader connection when in free-running mode.

Configuring and using TSE Service in an application

For a detailed description of how to configure and use TSE Service in an application, see Chapter 9 in the *PathWave Test Sync Executive System Setup Guide*.

TSE Service Configuration

This section describes the contents of the TSE Service configuration files. For examples on how to use these files see the *PathWave Test Sync Executive System Setup Guide*.

TSE Service requires specific configuration files. These files are in YAML (**.yaml**) format. See the YAML documentation for specifications of the YAML format.

There two configuration file:

- TSE Service Config File: **tse_config.yaml**
- System Definition File: **systemDefinition.yaml**

These files must be located in:

C:\ProgramData\Keysight\PathWave Test Sync Executive 2023\TseService\config

TSE Service Config File

Every host running TSE Service requires its own **tse_config.yaml** file. This applies to any of the configurations, Free-Running or Leader-Follower mode. It must be located in:

C:\ProgramData\Keysight\PathWave Test Sync Executive 2023\TseService\config

The Free-Running mode is the startup TSE Service mode and is configured through the **tse_config.yaml** file in each host. The TSE Service config file enables you to configure:

- TSE Server TCP port (this is strongly recommended if KDI is not used to launch TSE Service).
- Hardware resources: chassis, *System Synchronization Modules* (SSMs) and instruments.
- Enables you to specify simulated hardware instances.

TSE Service opens the specified resources at boot-up enabling remote access for PXI Chassis and SSMs, and also saving significant time when running the application.

See *Chapter 9 in the PathWave Test Sync Executive System Setup Guide* for a complete example of how to configure the **tse_config.yaml** for a system to run in Free-Running mode.

TSE Server TCP Address

You can set the TCP address and port for TSE Service manually, or let the system set it automatically by setting it to **0.0.0.0:0**.

It is important to select a port that is available and permitted by your organization's IT policies. The auto detection of the port is recommended only when using KDI to start the TSE Service, because the KDI infrastructure enables TSE Service to resolve the actual port at runtime. When KDI is not used to start TSE Service, Keysight strongly recommends that you specify the TCP port manually, and in the

case of multiple network interfaces, also the IP address, since you will need to specify the TCP port when using the TSE-TCP resource IDs.

For example:

```
# Configure automatic port for the TSE Service server
listen_address: 0.0.0.0:0

# Configure autodetect IP with port 8674 for the TSE Service server
listen_address: 0.0.0.0:8674
```

Infrastructure (Chassis, System Sync Modules and High Performance Reference Clocks)

The "chassis" section enables you to list the chassis to open at start-up, individually by VISA Resource IDs, or to autodetect them. This section is optional, you only need to list the chassis. The SSM and High Performance Clock are opened automatically with the chassis.

```
chassis:
- PXI0::1::BACKPLANE # Chassis 1
- PXI0::2::BACKPLANE: "OptionX=???" # Chassis 2 opened with options
- PXI0::3::BACKPLANE: "OptionX=???" # Chassis 2 opened with options
  ssm_options: "OptionY=???" # OPTIONALLY can also specify options for the SSM in the
chassis
```

Can also specify chassis in simulation mode:

```
chassis:
- PXI0::1::BACKPLANE: "Simulate=True,
DriverSetup=Model=GenericPxieChassis,EnhancedTrigger=True" # Chassis 1 opened in simulation
  ssm_options: "Simulate=true, DriverSetup=Model=M9033A" # OPTIONALLY can also specify
options for the SSM in the chassis
- PXI0::2::BACKPLANE: "Simulate=True,
DriverSetup=Model=GenericPxieChassis,EnhancedTrigger=True" # Chassis 2 opened in simulation
  ssm_options: "Simulate=true, DriverSetup=Model=M9032A" # OPTIONALLY can also specify
options for the SSM in the chassis
```

NOTE For PXI chassis simulation you must specify the "GenericPxieChassis" model.

Alternatively, you can configure TSE Service to autodetect and open all PXI chassis connected to that host, note that this is not supported for simulation:

```
- autodetect_all_pxi
```

or include options as well:

```
- autodetect_all_pxi: "OptionX=true"
```

NOTE You are not required to explicitly specify SSMs and High Performance Clock Modules, because these are detected and opened automatically with the chassis.

Instruments Section

The "Instruments" section enables you to list instruments to be opened automatically at start-up, these instruments can be later be inherited by the Leader TSE Service to include them in the automatic initialization. Similarly, as for Chassis, you can list the instruments individually using the VISA Resource ID:

```
instruments:
- PXI0::CHASSIS1::SLOT7::INSTR:
- PXI0::CHASSIS2::SLOT6::INSTR:
- PXI0::CHASSIS1::SLOT3::INSTR: "Option1=true"          # OPTIONALLY can specify initialization options
```

You can also specify instrument in simulation mode:

```
instruments:
- PXI0::CHASSIS1::SLOT7::INSTR: "Simulate=true,DriverSetup=Model=M50XXA"
- PXI0::CHASSIS2::SLOT6::INSTR: "Simulate=true,DriverSetup=Model=M53XXA,Option1=true"
- PXI0::CHASSIS1::SLOT3::INSTR: "Simulate=true,DriverSetup=Model=M52XXA,Option2=false"
```

Alternatively, you can configure TSE Service to autodetect and open all PXI instruments in that host and connected to the chassis defined in the chassis section (this option does not work for simulation):

```
- autodetect_all_pxi
```

or include options as well:

```
- autodetect_all_pxi: "Option1=true" # OPTIONALLY can specify initialization options
```

NOTE

For most instruments, except the M3xxxA family, KDI must be installed to enable the instrument to be opened by TSE Service and shared by the user application. Do not use the Instruments section if you want the application to open the instruments locally with exclusive access, instead of using the remote capability for shared access.

Instrument and Chassis Simulation

To work in simulation, you must specify all chassis and instruments that must be opened in simulation using the '`simulate=true`' and the appropriate model in addition to any other options required for the chassis or instrument to work properly in simulation. Autodetection is not supported for opening Chassis or Instruments in simulation. Resources must be specified using VISA_RESOURCE_IDs, do not use any other form of Resource IDs.

NOTE

Keysight recommends you not to include hardware and simulated instruments together in the same TSE Service Config file. It is not recommended to specify simulated instruments and also specific the '`autodetect_all_pxi`', because it also opens any hardware resources detected.

KDI User and Password

It is recommended to launch TSE Service automatically with KDI, but in those cases where that is not possible and TSE Service must be launched manually, the `tse_config.yml` file must include the KdiUser and KdiPassword if the instrument section opens instruments other than the M3xxxA family.

```
# Must specify KDI User and Password for the instrument section to open instruments for shared
access (default operation)
# Not needed for M3xxxA instrument family
kdi_user: user123456
kdi_password: pass123456
```

When KDI User and Password are included in the `tse_config.yml` for a given TSE Service, all remote operations performed by this TSE service instance will be authenticated using the specified KDI User and Password.

The System Definition YAML File

The Leader host in a Leader-Follower system requires a `system_definition.yml` file. This file is only used in the Leader in systems configured in Leader-Follower mode and must be located in:

C:\ProgramData\Keysight\PathWave Test Sync Executive 2023\TseService\config

Once TSE Service completes the free-running initialization if the `system_definition.yml` file is found, it switches into the Leader mode. This file defines the entire system configuration, including:

- System topology.
 - This includes information about the chassis and how they are connected with SSMs and SystemSync Connectivity.
- Clocking.
- Sync resources.
- Warmup time.
- Instruments (optional).

System Topology

To add the topology of the system, you list the chassis in the same order they are physically connected. Chassis are physically connected together with *System Synchronization Modules* (SSMs). The chassis are listed with TSE Resource IDs (can also use VISA resource IDs for chassis connected to the same host as the Leader TSE Service).

1. First list the leader with all follower chassis directly connected downstream from it.
2. Add an entry for each other chassis with all follower chassis directly connected downstream from it.

The following table shows the format for listing the chassis:

Format	Example
- chassis: <chassis or host address>PXIO::<chassis number>::BACKPLANE	
downstream:	- chassis:
- <System Sync Downstream Connector Number ¹ >: <chassis or host address>PXIO::<chassis number>::BACKPLANE	tse://host1/PXIO::1::BACKPLANE
- <System Sync Downstream Connector Number ¹ >: <chassis or host address>PXIO::<chassis number>::BACKPLANE	downstream:
	- 1: tse://host1/PXIO::2::BACKPLANE
	- 2: tse://host2/PXIO::1::BACKPLANE

(1) The Connector Number is the number of the physical connector on the upstream SSM.

The following example code shows the leader first with two downstream chassis and one of these downstream chassis also having two additional downstream chassis, note that the indentation must follow the YAML rules:

```
topology:
- chassis: tse://host1/PXIO::1::BACKPLANE
  downstream:
    - 1: tse://host1/PXIO::2::BACKPLANE
    - 2: tse://host2/PXIO::1::BACKPLANE
- chassis: tse://host2/PXIO::1::BACKPLANE
  downstream:
    - 1: tse://host2/PXIO::2::BACKPLANE
    - 2: tse://host2/PXIO::3::BACKPLANE
```

For a complete example, see the *PathWave Test Sync Executive System Setup Guide*.

Client applications can define and use a subset of the system's topology. This is supported via the optional `support_client_subtopologies` flag when it is enabled. This flag is optional, the default is `true`.

```
# [OPTIONAL] When enabled, this enables the client applications to use a subset of the system
topology. By default, it is true.
support_client_subtopologies: true
```

Sync Resources

You must list Synchronization Resources. For information about Sync Resources, see [Synchronization Resources](#).

For example:

```
sync_resources:
- PXI_TRIGGER_1
- PXI_TRIGGER_3
```

Clocking

You can optionally specify the clock reference, mode and frequency. For more information about clocking see [Clocking](#), [Non-HVI Clocks](#), and [System Clocking Configuration](#).

System Reference Clock

The permitted values can be one of:

Parameter	Option	Description	Default
reference:	ssm	The SSM internal clock source	x
	chassis	The internal chassis clock	-
	hpc	The <i>High Precision Refence Clock Source</i> (HPRCS), available as an option in some Keysight chassis.	-
mode:	internal	Use the internal clock.	x
	external	Use an external clock as a reference.	-
frequency:	-	Use if <code>mode=external</code> , This is the frequency of the external reference.	-

For example:

```
clocking:
# If this section is not included, TSE uses the defaults ssm and internal
reference: ssm # Permitted values are: ssm|chassis|hpc. Default is: ssm.
# the SSM used is the one of the leader chassis defined in the topology
section
mode: external # Permitted values are: internal|external. Default is: internal.
frequency: 100e6 # Permitted only with mode == external
```

Inputs, Outputs and Synchronization Clocks

You can optionally specify the clocks used in the system:

- External Analog Clocks.
- Clock Outputs.
- Non-HVI Core Clocks.
- Non-HVI System Clocks.

These are specified by giving the name and the frequency.

External Analog Clocks

External Clocks are used to enable the external analog/RF clock reference in instruments. When running the initialization TSE Service will configure instruments to use external analog/RF reference for those instruments that support the frequency specified in the "*external_analog_clock*" entry:

```
external_analog_clocks: #OPTIONAL
- 2.4e9
- 4.8e9
```

Clock Outputs

Use the "clock_outputs" section to enable specific Chassis and SSMs clock outputs. The clocks are listed with the following options:

Option	Description	Required
<code>output:</code>	The output being specified	Yes
<code>enabled:</code>	If the output is enabled, true or false	Yes
<code>frequency:</code>	Frequency of the clock output	Not for outputs with a single specific frequency.

For example:

```
clock_outputs: #OPTIONAL
- chassis: tse://host1/PXI0::1::BACKPLANE
  output: FPref1Out
  enabled: true
  frequency: 100e6
- chassis: tse://host2/PXI0::2::BACKPLANE
  output: FP2.4GHzOut
  enabled: true
```

NOTE

The options available depend on the chassis or SSM that is generating the clock, for more information see the relevant documentation.

Non-HVI Core Clocks and Non-HVI System Clocks

If the system has specific synchronization constraints, use the entries `non_hvi_core_clocks` and `non_hvi_system_clocks` to add non-HVI clocks as required by the system, see [Non-HVI Clocks](#) for more details:

```
non_hvi_core_clocks:  #OPTIONAL - In Hz
- 15e6
- 10e6
non_hvi_system_clocks:  #OPTIONAL - In Hz
- 1e6
```

Initialization Modes

The TSE Service Leader performs a complete system initialization at start-up. Depending on the Warm-Up-Time defined, TSE Service performs two initializations:

1. [Only when Warm-Up-Time > 0] Pre-calibration initialization.
2. [Only when Warm-Up-Time > 0] Wait warm-up time.
3. Perform final initialization (this step requires calibration data available).

By default, step #3 uses the default initialization options as described in [System Initialization in the SystemDefinition Object](#). You can add specific initialization options with the "initialize_alignment_modes" section:

```
# The Alignment Mode flags to apply during system initialization
initialize_alignment_modes:
- ForceClockMonitoring
- Full
# - ResetCalibration
# - PreCalibration
# - DisableClockMonitoring
```

NOTE

Note that before running a normal initialization using calibration data, you must execute once an initialization to calculate the calibration data using the "*ResetCalibration*" initialization mode. To run an initialization without calibration data, specify the "*PreCalibration*" flag.

Warmup Time

You can specify the warmup time, in seconds.

```
initialization_warmup_time_seconds: 30
```

When this time is specified and different than zero, TSE Service performs two initializations:

1. Pre-calibration initialization.
2. Wait warm-up time.
3. Perform initialization using calibration data.

The warm up is required for instruments to achieve the steady operation temperature. Check the instrument documentation for recommended warm-up time.

Instruments

You can explicitly list instruments in the `system_definition.yml`, so these instruments are included in the automatic start-up system initialization/alignment. These are listed with KDI Resource IDs.

```
instruments: #These instruments are added to the Leader TSE Service SystemDefinition to be
initialized.
- kdi://localhost/PXI0::CHASSIS1::SLOT3: Option1=1,      # include explicitly by instrument
address and options.
- kdi://localhost/PXI0::CHASSIS2::SLOT12: Option1=1,    # include explicitly by instrument
address and options.
```

The `inherit_from_tse_config` flag instructs TSE Service to include all the instruments listed or auto-detected per the `tse_config.yml` file in each host (leader and followers) that belong to the topology specified. This can simplify the `system_definition.yml` because you don't need to explicitly list all instruments across multiple hosts and instead rely on each host configuration.

For example:

```
instruments: #These instruments are added to the Leader TSE Service SystemDefinition to be
initialized.
- inherit_from_tse_config # includes all instruments found for the Topology specified and
defined in the tse_config.yml files of each host
```

Accessing Remote Resources

The standard way to open local resources (instruments, chassis, Sync Modules, etc) is to use a *VISA Resource ID*. With TSE Service and KDI there are other special resource identifier formats that enable access to resources distributed across the network:

- TSE and TSE-TCP Resource IDs.
- KDI Resource ID.

TSE Service enables you to remotely access resources: chassis and *System Sync Modules* (SSMs) by means of the TSE and TSE-TCP resource IDs. While KDI Resource ID is supported by some instruments product families to enable remote and multiple access to instruments.

TSE Resource IDs to access Chassis, SSMs and TSE Service instances

In a system or user application with TSE Service, you can access chassis, *System Synchronization Modules* (SSMs), and TSE Service instances distributed across multiple hosts using the TSE Resource IDs. TSE Service and TSE Resource IDs can also be used in a Single chassis system.

The following table shows the 2 types of TSE Resource IDs available:

TSE Resource ID Format	Use Case	Example
<code>tse://<host>/<VISA_RESOURCE_ID></code>	Requires TSE Service launched by KDI. Requires <i>KdiUser</i> and <i>KdiPassword</i> when used in the user application (see KDI Resource ID usage).	For <code>VISA_RESOURCE_ID = PXI0::1::BACKPLANE</code> • <code>tse://myhost/PXI0::1::BACKPLANE</code>
<code>tse-tcp://<host>:<port>/<VISA_RESOURCE_ID></code>	Does not require KDI.	For <code>VISA_RESOURCE_ID = PXI0::CHASSIS1::SLOT10::INSTR</code> <code>tse-tcp://myhost:7587/PXI0::CHASSIS1::SLOT10::INSTR</code>

Where <host> can be specified using one of these forms:

1. Host Name (or Device Name), this is the recommended for most users. For example: **MyLabPc1**.
2. Full Host Name, For example: **MyLabPc1.NetworkDomain**.
3. IP address, this is not recommended, For example: **10.127.1.89**.

With TSE-TCP Resource IDs, the application must know the hostname and the TCP port used by TSE service to connect to or access resources controlled by that TSE Service, for example, to add a remote chassis or Sync Module.

With TSE Resource ID, the service discovery is resolved by KDI and there is no need to specify upfront the TSE Service TCP port. As a result of using KDI, when using TSE Resource IDs in the client application it is required to specify the KDI User and Password options.

NOTE KdiUser, KdiPassword options are mandatory when using TSE Resource ID.

Accessing Chassis and System Sync Modules

To open/access the chassis, you must use the TSE or TSE-TCP Resource IDs explained above, depending on whether TSE Service is launched with KDI or not. For instance:

Chassis Resource ID specified in tse_config.yml in testNode2	TSE and TSE-TCP Resource ID
PXI0::1::BACKPLANE	tse://testNode2/PXI0::1::BACKPLANE tse-tcp://testNode2:8674/PXI0::1::BACKPLANE

For the System Synchronization Module, the resource ID is derived from the corresponding chassis resource ID, for instance:

Chassis Resource ID specified in tse_config.yml in testNode2	Corresponding SSM TSE and TSE-TCP Resource ID
PXI0::1::BACKPLANE	tse://testNode2/PXI0::CHASSIS1::SLOT10::INSTR tse- tcp://testNode2:8674/PXI0::CHASSIS1::SLOT10::INSTR

Where *SLOT10* corresponds to the timing slot in the PXI0::1::BACKPLANE chassis (an 18-slot chassis).

The following snippet illustrates in code the examples above:

Opening Chassis & SSMs with TSE

```
# TSE Resource IDs rely on KDI infrastructure to resolve hosts and TCP ports
# Must specify a valid user registered in KDIS
kdi_user_option = 'KdiUser=hviuser1234,KdiPassword=hviuser1234password'
```

```
#
# Add Chassis using TSE Resource IDs
mySystemDefinition.chassis.add('tse://host1/PXI0::1::BACKPLANE', kdi_user_option) #Chassis 1 in
Host 1
#
# Add SSMS for each chassis with TSE Resource IDs (SSMs are opened already by TSE Service)
# the slot must be the timing slot of the PXI Chassis (for 18-slot chassis, it is slot 10)
primarySSM = mySystemDefinition.interconnects.add_sync_module
('tse://host1/PXI0::CHASSIS1::SLOT10::INSTR', kdi_user_option)
```

Opening Chassis & SSMS with TSE-TCP

```
# Or add Chassis using TSE-TCP Resource IDs => must know the port the TSE Service is using
mySystemDefinition.chassis.add('tse-tcp://host1:8674/PXI0::1::BACKPLANE') # Chassis 1 in Host 1
#
# Add SSMS for each chassis with TSE-TCP Resource IDs (SSMs are opened already by TSE Service)
# must know the TCP port the TSE Service is using
# the slot must be the timing slot of the PXI Chassis (for 18-slot chassis, it is slot 10)
primarySSM = mySystemDefinition.interconnects.add_sync_module('tse-
tcp://host1:8674/PXI0::CHASSIS1::SLOT10::INSTR') #SSM for Chassis 1 in Host 1
```

Accessing TSE Services

In Leader-Follower mode, the client application creates a client system definition that connects to a Leader TSE Service, to establish this connection the TSE or TSE-TCP Resources IDs must be used. The following snippet demonstrates this:

Create a Client System Definition with TSE

```
# TSE Resource IDs rely on KDI infrastructure to resolve hosts and TCP ports
# Must specify a valid user registered in KDIS
kdi_user_option = 'KdiUser=hviuser1234,KdiPassword=hviuser1234password'
#
# Create system definition client connected to the TSE Service Leader using TSE Resource ID
sys_def = pyhvi.SystemDefinition("MyAppSystemDef","tse://host1", kdi_user_option)
```

Create a Client System Definition with TSE-TCP

```
# Create system definition client connected to the TSE Service Leader using TSE-TCP Resource ID
sys_def = pyhvi.SystemDefinition("MyAppSystemDef","tse-tcp://host1:8674")
```

KDI Resource ID to open instruments

KDI enables you to open or access instruments across the network using the KDI Resource ID. With the appropriate options, KDI also enables you to share the same instrument across applications and processes. For this, the first call to open a given resource launches an independent process where the hardware or service session is opened. Following calls to open the same resource with KDI will connect to the existing session.

TSE Service makes use of KDI to open and initialize instruments at boot-up and get them ready to be used by the user application later. It is important to note that TSE Service performs the first open of the hardware resource, so applications do not need to apply any of the specific options required in the first call to open a resource with KDI.

NOTE

When configuring TSE Service to open instruments at boot-up, you do not need to specify any of the options listed in the *First KDI open of a Resource*.

To open an instrument with KDI, the process is the same as without KDI, except:

1. Use the KDI Resource ID instead of the VISA one. The KDI Resource ID is built by adding a prefix to the VISA Resource ID:
 - `kdi://<host>/<VISA_RESOURCE_ID>`.
2. For Authentication the initialization options must include:
 - `KdiUser` and `KdiPassword`, these are the user and password you set when you added a user (client) in KDIS, see explanation above on how to configure KDIS.
 - `KdiUrl` in general not needed, but it depends on the system configuration, see details below.
3. To open multiple KDI instances to the same Hardware you must include:
 - `AllowMultipleClientAttach=1`

If you intend to use the HVI Engine of the instrument with TSE API, at the *First opening* of a KDI Resource, you must also include:

- `HviServer =>` for example, for automatic IP and Port: `HviServer=HVITCP:[::]:0`

This first-time option is ignored in following open calls with KDI, as long as the specific resource is kept open.

For example, if the instrument is not opened by TSE Service:

Opening instrument with KDI

```
myRemoteInstrument = keysight_ktmodule.KtModule('kdi://testNode1/PXI0::CHASSIS1::SLOT7::INSTR',  
1, 1, 'Simulate=0,DriverSetup=,KdiUser=user1234, KdiPassword=pass1234,  
AllowMultipleClientAttach=1, HviServer=HVITCP:[::]:0, other_instrument_specific_options...')
```

You must always include the KDI user and password as defined in KDIS configuration when you open instruments in your application.

When the instrument is already opened by TSE Service, the user application must specify the KDI username and password and `AllowMultipleClientAttach=1`, other options are in general not needed, since they are passed by TSE Service:

Opening instrument with KDI

```
myRemoteInstrument= keysight_ktmodule.KtModule('kdi://testNode1/PXI0::CHASSIS1::SLOT7::INSTR',
1, 1, 'Simulate=0,DriverSetup=,KdiUser=user1234, KdiPassword=pass1234,
AllowMultipleClientAttach=1')
```

NOTE KdiUser, KdiPassword and AllowMultipleClientAttach options are mandatory.

NOTE All instances of the same instrument must use the same AllowMultipleClientAttach value, otherwise, you will get this error:

Could not start or attach to station service KtCornerstone 0.3.6914. Error message: Cannot connect to existing instance (PXI0::CHASSIS1::SLOT2::INSTR) because AllowMultipleClientAttach has not been set by both requestor (false) and instance (true).

NOTE Simulate=false is the default option if not specified. For IVI compliant drivers the option "**DriverSetup=**" is mandatory after the IVI standard options and before instrument specific ones. "DriverSetup=" is not required for Python drivers, it is ignored.

KdiUrl initialization option

KdiUrl is an optional initialization option specific for instrument drivers, it must be specified as an initialization option for instruments when the KDI client configuration file does not include the *upstream-manager* entry, see *Configure KDI Clients to find KDIS* for more details:

Opening instrument with KDI

```
myRemoteInstrument = keysight_ktmodule.KtModule('kdi://testNode1/PXI0::CHASSIS1::SLOT7::INSTR',
1, 1, 'Simulate=0,DriverSetup=,KdiUser=user1234, KdiPassword=pass1234,
KdiUrl=wss://localhost:9090/ws, AllowMultipleClientAttach=1, other_instrument_specific_
options...')
```

NOTE If KdiUrl is needed and not specified, you will get an error like:

Couldn't get test station manager URL through options parameter:Simulate=0,DriverSetup=,LogLevel=Info,KdiUser=hviuser1234,KdiPassword=hviuser1234,HviServer=HVITCP[:]:0

Building the correct Remote Resource ID for multiple access

When specifying a Remote resource ID for multiple access, TSE, TSE-TCP or KDI, it is mandatory that all instances uses the same Resource ID format, in particular all Remote Resource ID used in different instances that refer to the same Instrument must use the same VISA Resource ID when building the Remote Resource ID.

The table below show the different VISA Resource IDs supported for PXI Chassis and Instruments, and the one that is used by TSE Service which must be used by any application working with TSE Service:

Type	VISA Resource ID format	TSE Support
PXI Chassis Example: Chassis #1	PXI0::1::BACKPLANE	Supported
	PXI0::35-0:0::INSTR	Not Supported
	PXI35::0:0::INSTR	Not Supported
PXI Instrument (also System Sync Modules) Example: Instrument in Chassis #1 and Slot 12	PXI0::CHASSIS1::SLOT12::INSTR	Supported
	PXI0::CHASSIS1::SLOT12::INDEX0::INSTR	Not Supported
	PXI0::30-0.0::INSTR	Not Supported
	PXI30::0:0::INSTR	Not Supported

NOTE

When the instrument or chassis is opened by TSE Service at boot up, the Remote Resource ID (KDI, TSE or TSE-TCP) used in the client application must include the same VISA_RESOURCE_ID as listed in the `tse_config.yml` file (or TSE Service Log). See *TSE Service Free-Running Mode* section for more details on instrument management. For instance, if the client application KDI Resource ID does not match the one used by TSE Service, you will get this error:

```
Failed to instantiate PXI0::CHASSIS1::SLOT2::INDEX0::INSTR. Error:
<StdException>#std::runtime_error#Visa error 0xbfff000f. VI_ERROR_RSRC_
LOCKED: Operation failed due to locked resource
```

Using remote instruments HVI engines in a user application

In order to use TSE capabilities in a user application, you must add the HVI Engines of the desired instruments in the application systemDefinition instance using the *HVI Engine unique ID* provided by the instrument drivers. HVI-capable instruments expose in the instrument drivers the *hvi* interface which includes properties with all the available HVI resources, in particular the *HVI Engine Unique IDs*, the following snippet illustrates how an HVI engine is added to the systemDefinition instance:

```
my_instrument = keysight_ktmodule.KtModule('PXI0::CHASSIS1::SLOT7::INSTR', False, False,
InstrOptions)
```

```
#  
my_sys_def = keysight_tse.SystemDefinition('Hvi') # optionally add Leader TSE Service Resource  
ID  
#  
my_sys_def.engines.add(my_instrument.hvi.engines.main_engine, "MyEngineAlias")
```

TSE has built-in multi-process and multi-host capabilities which enable a user application to add HVI Engines and exploit TSE capabilities of instruments opened in other processes and hosts, the information required for this is encoded in the *HVI engine Unique IDs*. An application in a different process or host just needs to know the *HVI Engine Unique ID* of a specific instrument, to use it regardless of the process or host where the instrument is opened.

By default when an instrument is opened, the HVI Engine only supports in-host multi-process access implemented using shared memory for performance. To enable access from other hosts the **HviServer** option must be specified when opening the instrument, for example, for automatic IP and Port: "**HviServer=HVITCP[:]0**". The **HviServer** option opens a TCP server that allows TSE applications out of host to access and control the HVI Engine and HVI capabilities in that instrument. The following snippet illustrates how to open an instrument locally with multi-host access for the TSE capabilities:

Opening instrument with KDI

```
my_instrument = keysight_ktmodule.KtModule('PXI0::CHASSIS1::SLOT7::INSTR', 1,  
1, 'Simulate=0,DriverSetup=,HviServer=HVITCP[:]0, other_instrument_specific_options...')
```

NOTE

"**HviServer=**" initialization option must be specified when opening instruments to allow multi-host access/control of TSE capabilities from a different host. This is required for instruments located in hosts different from the host where the application implementing the systemDefinition instance is executed.

Using TSE Service in an Application

With the introduction of TSE Service, there are a number of different options to configure the system in an application to exploit TSE capabilities:

- No TSE Service.
- TSE Service in Free-Running mode.
- TSE Service Leader-Follower mode.

No TSE Service

You create a SystemDefinition object and define the complete system in it, this is the same as the method used in TSE 2022 and previous releases. You use the call:

```
SystemDefinition(<name>)
```

For example:

```
# Define SystemDefinition  
my_system = kthvi.SystemDefinition("MySystem")
```

Once the SystemDefinition object is created, you must define the complete system, with chassis, interconnects, clocks, sync resources, HVI Engines, etc.

TSE Service Free-Running mode

The setup of Free-Running mode is similar to the *No TSE Service* setup in that you must define the chassis, clocks, topology etc. however, in this case you only need to define what you plan to use.

The application code to create and configure the SystemDefinition object, when using TSE Service in Free-Running mode, is the same as you use for no TSE Service. The advantages of using TSE Service are:

1. Access remote in other hosts like chassis and System Sync modules (SSMs) using the TSE Resource IDs.
2. Optionally, you can configure instruments to be open at TSE Service boot-up using KDI Resource IDs to reduce the instrument driver load time in the application. For local instruments this may have the disadvantage that some calls to the instrument driver may be slightly slower because they go through the TCP remote connection.

The user application must specify and initialize all the components they want to use in the application in the System Definition object using the TSE API.

After boot-up, and once the TSE Service is running and completed initialization, the user (or client) application must:

1. Open all instruments.
2. Create a System Definition and add the Chassis and SSMs.
 - a. Use TSE Resource IDs to access Chassis and SSMs in other hosts. See the [Accessing Remote Resources](#) Section for more details.
3. Complete the System Definition configuration as usual:
 - a. Specify the topology of the system by indicating the physical System Sync connections between the SSMs.
 - b. Configured the clocking, Sync Resources, etc. as required.
 - c. Add the Instrument HVI Engines, etc.
 - d. Trigger the system initialization with `systemDefinition.initialize(...)`.
4. Create the Sequencer, HVI instance, and rest of the application as usual.

For a full description of how to do this with examples, see Chapter 9 of the *PathWave Test Sync Executive System Setup Guide*.

NOTE

Your application using the TSE API can run in any host, you do not need to run the application in the hosts connected to the hardware. If your application is used in a host with no hardware, then the TSE Service is not required to be running on this host. Only the Hosts with hardware connected to them require the TSE Service to be running.

TSE Service Leader-Follower mode

Starting an application in a system in Leader-Follower mode is simpler than Free-Running mode because you are not required to add chassis, SSMs, clocks or define the topology, all this is defined once in the `system_definition.yml` file. Additionally, applications run faster because the complete system initialization can be executed automatically at windows boot-up before the application is executed. The application only needs to open the instruments and add the HVI engines into the SystemDefinition instance to get the complete hardware configuration.

When using TSE Service in leader-follower mode the user application creates a client SystemDefinition instance which connects to the Leader TSE Service. To create a client SystemDefinition instance, you must specify the TSE or TSE-TCP Resource ID of the leader TSE Service:

- `SystemDefinition(<name>, <leader_address>)`
- `SystemDefinition(<name>, <leader_address>, <options>)`

The following snippets demonstrates how to create a client SystemDefinition. Note that when using TSE Resource IDs, the options must include the KDI User and Password.

```
# To access remote resources with the KDI infrastructure must use a valid user registered in KDIS
```

```
KdiUserOption = 'KdiUser=hviuser1234,KdiPassword=hviuser1234password'  
#  
# Create system definition client connected to the TSE Service Leader  
sys_def = pyhvi.SystemDefinition("Hvi","tse://host1", KdiUserOption)  
  
# Create system definition client connected to the TSE Service Leader  
sys_def = pyhvi.SystemDefinition("Hvi","tse-tcp://host1:8674")
```

If the connection fails, or the TSE Service is not ready in the Running state in the Leader mode, the operation will throw an exception with information on the TSE Service state if the connection to the TSE Service was established.

Once the SystemDefinition object is created, the application must add the HVI Engines and configure the engine resources, there is no need for the application to configure Chassis, SSMs, interconnects or clocking, since all this is done by the Leader TSE Service per the **systemDefinition.yml** file. For a full description of how to do this with examples, see Chapter 9 of the *PathWave Test Sync Executive System Setup Guide*.

Chapter 10: HVI Time Management and Latency

This chapter describes HVI time management and latency. It introduces the concepts involved and describes the timing and latencies of Statement execution, how they impact the overall execution timing of Sequences, and the constraints on the Start Delay and duration of Statements. It also provides latency information for the different Statements and instructions.

This chapter contains the following sections:

- [Timing Concepts](#)
- [Sync Statement Timing](#)
- [Local Flow-Control Statement Timing](#)
- [HVI Instruction Timing](#)
- [Minimum Start Delay Calculation for Local Flow-Control and Sync Statements](#)
- [Errors in Start Delay or Duration specification](#)

Timing Concepts

This section describes timing concepts. It contains the following sections:

- [Synchronization Clocks, Signals, and Modes](#)
- [General Timing Concepts](#)
- [Sync and Local Flow-Control Statement Timing Concepts](#)
- [HVI Instruction Timing Concepts](#)

Synchronization Clocks, Signals, and Modes

This section introduces the main concepts of clocks, signals and propagation delays involved in HVI time management. It includes the following sections:

- **Clocks Overview**
 - HVI Engine clock
 - HVI Engine cycle
 - HVI Common clock
- **Synchronization Signals**
 - Core clocks
 - System clocks
- **Physical Propagation Delay**
 - Sync Period Calculation

Clocks Overview

The following list describes the main concepts that apply to all Statement types:

HVI Engine clock

This is the HVI Core clock of the Engine in an instrument. There is one HVI Engine clock per instrument.

HVI Engine cycle

An Engine cycle is the timeframe in which the HVI Engine can fetch, dispatch or execute instructions. One Engine cycle is equal to the period of the Engine clock. For example, for an Engine that runs at 100 MHz, the duration of an Engine cycle will be equal to 10 ns.

HVI Common clock

This is not a real clock, it is a definition to calculate timing for Sync Statements within HVI Sequences. It can be seen as a clock that has its rising edge aligned with all HVI Engines clocks rising edges.

Therefore, its frequency is equal to the GCD of the frequencies of all the HVI Engine clocks:

$HVI_Common_Clock_{Frequency} = GCD\{HVI_Engine_Clock_1_{Frequency}, HVI_Engine_Clock_2_{Frequency}, \dots, HVI_Engine_Clock_N_{Frequency}\}$, where N is the number of Engines added to HVI.

The period can be calculated in two ways:

- as the LCM of HVI Engine Cycles (the periods of all the HVI Engine clocks):

$HVI_Common_Clock_{Period} = LCM\{HVI_Engine_Clock_1_{Period}, HVI_Engine_Clock_2_{Period}, \dots, HVI_Engine_Clock_N_{Period}\}$, where N is the number of Engines added to HVI

- or, just the inverse of the HVI Common Clock frequency:

$HVI_Common_Clock_{Period} = 1 / HVI_Common_Clock_{Frequency}$

For example, if the Engines added to HVI have the following HVI Engine Clock frequencies {100MHz, 187.5MHz, 300MHz}, the HVI Common Clock frequency/period will be:

$HVI_Common_Clock_{Frequency} = GCD\{100MHz, 187.5MHz, 300MHz\} = 12.5MHz,$

$HVI_Common_Clock_{Period} = 1 / HVI_Common_Clock_{Frequency} = 80ns$

When calculating Sync Statement timing, the HVI Common Clock period is used to round some timing magnitudes to the next HVI Common Clock period (below **ncc** stands for **next common clock**):

$round_{ncc}(TimeValue) = ceil(TimeValue / HVI_Common_Clock_{Period}) * HVI_Common_Clock_{Period}$

NOTE

When working with fractional or periodic time values or periods to avoid problems with the numerical precision it may be better to use the frequency value instead following this simple equation:

$HVI_Common_Clock_{Frequency} = 1 / HVI_Common_Clock_{Period}$

Synchronization Signals

HVI uses different periodic digital signals for synchronization purposes. The definition of those digital signals depends on platform and instruments signals. Platform signals are the CLK100 and CLK10 signals in a PXI platform such as a PXI chassis. Instruments have different clock signals inside that are classified as core clocks or system clocks.

In cases where an instrument does not include an HVI Engine, these are known as non-HVI core clocks or non-HVI system clocks. For more information, see [Non-HVI Clocks](#).

Core clocks

These are clocks used by the HVI Engine embedded in an HVI-enabled instrument, including:

- HVI Core clocks. These are clocks automatically reported to HVI by an HVI-enabled instrument. These include the HVI Engine clocks of all engines in the SystemDefinition instance
- Non-HVI Core clocks. These are clocks users report explicitly to HVI using the PathWave Test Sync Executive API, for instruments without HVI, or Devices Under Test. These are included so they can be taken into account during the synchronized execution of an HVI sequence. Non-HVI Core clocks must also include any other clock period you require the execution of an HVI sequence to be synchronized with.

These clocks are used in the calculation of both the SYNC and SYNC_BASE periods.

System clocks

These are the relevant clocks that drive the internal logic of individual instruments, including:

- HVI System clocks. These are clocks automatically reported to HVI by an HVI-enabled instrument.
- Non-HVI System clocks. These are clocks users reports explicitly using the PathWave Test Sync Executive API to account for specific synchronization requirements of instruments without HVI, or Devices Under Test.

These clocks are taken into account in the alignment process, because they affect the calculation of the SYNC_BASE period.

System clocks are not Core clocks.

Platform and instrument clock signals contribute to define the HVI Sync signals according to the following definitions.

The period (and inversely, also the frequency) of the SYNC signal is defined as:

$$\text{Sync_Period} = N \times \text{LCM}(\text{all instrument core clocks}), N \text{ such that } N \times \text{LCM}(\cdot) \geq \text{PhysicalPropagationDelay}$$

The period (and inversely also the frequency) of the SYNC_BASE signal is defined as:

$$\text{Sync_Base_Period} = \text{LCM}(\text{CLK10}, \text{Sync_Period}, \text{all instrument system clocks})$$

where, in the above formulas, LCM(.) stands for the *Least Common Multiple* operation.

Both SYNC and SYNC_BASE periods must be equal to or greater than the *Physical Propagation Delay* value for the relevant multi-chassis topology, these are given in the following table. If the LCM(all instrument core clocks) is smaller than that, then you must take the next multiple after the LCM(.) that is equal to or greater than the Physical Propagation Delay. The next multiple is the actual Sync Period value for your system. Once that Sync Period value is obtained, you can use it in the SYNC_BASE LCM formula to estimate the SYNC_BASE Period, which is automatically also greater than the Physical Propagation Delay.

Physical Propagation Delay

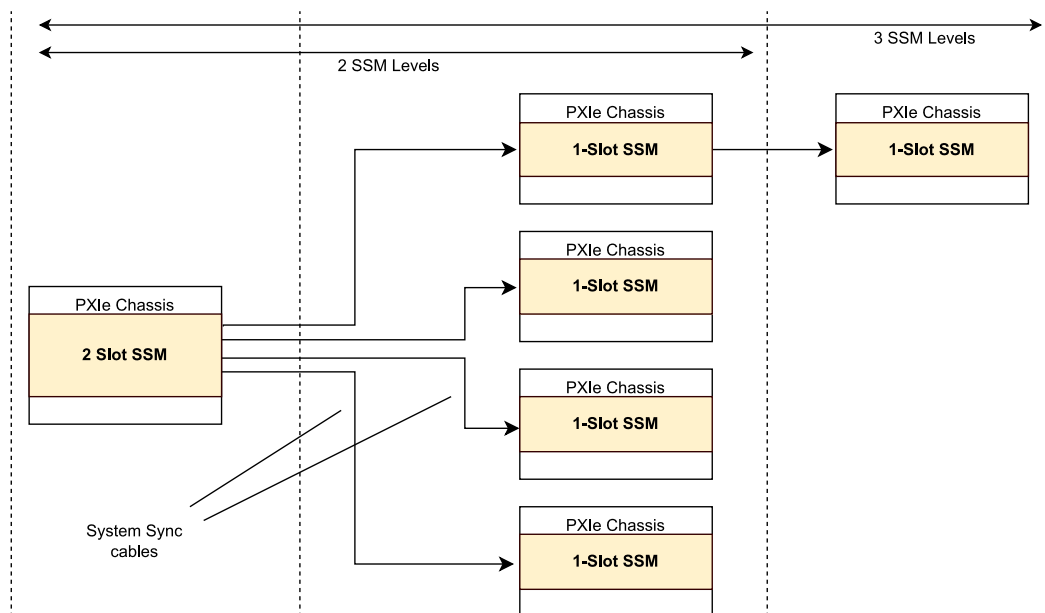
The Physical Propagation Delay corresponds to the amount of time (expressed in nanoseconds) that a PXIe Trigger needs to cover the path between any given pair of segments in a topology. This value is used when running Sync Statements because it provides information about how long the execution signaling between modules takes.

The topology is defined by the number of chassis in your system and how they are connected to each other with the System Sync cabling.

The System Sync cabling distributes clocks, Triggers, and data from the Leader SSM to the followers, possibly going through intermediate followers. The number of System Sync hops between the Leader SSM and each Follower determines what is known as the SSM level. The Leader is SSM level 1, all SSMs connected with 1 hop to the Leader SSM are Level 2 SSMs, those with 3 hops are Level 3 SSMs, and so on.

In the case of the M9033A SSM, there can be up to 4 followers connected to a single SSM, so there can be up to 5 chassis in system with 2 SSM levels. If you connect additional SSMs to the level-2 SSMs, this creates a 3rd level. In this arrangement you can add one additional chassis, this is because the PathWave Test Sync Executive supports up to 6 chassis in a single host system. To use more chassis, you must use a Multi-Host system, for more information see [Chapter 9: TSE Service and Multi-Host support](#).

The following diagram shows a 6 chassis system with 3 SSM levels:



The following table shows the Physical Propagation Delay values for different numbers of chassis and SSM levels:

Number of Chassis	Number of SSM levels	Physical Propagation delay ^{1,2} (ns)	Notes
1 chassis	-	100	
2 chassis	-	200	
3 chassis	-	300	
>3 chassis	2 SSM levels	300	Maximum 5 chassis
>3 chassis	3 SSM levels	400	Maximum 6 chassis with PathWave Test Sync Executive in a single host system.

- 1** Upper bound on the time it takes for a PXIe Trigger to travel from the furthestmost segments
- 2** Ensure your M904x chassis has version 5 or higher firmware revision for the Left and Right Trigger Bridges. See the hardware revision in your chassis Software Front Panel (SFP).

Sync Period Calculation

The Sync Period must always be greater than or equal to the Physical Propagation Delay. To obtain the actual Sync Period value, you first calculate the *Least Common Multiple* (LCM) of all HVI and non-HVI core clock periods added to the SystemDefinition. Secondly, you compare the LCM with the Physical Propagation Delay and take the next multiple of the LCM that is greater than or equal to the Physical Propagation Delay. This is what was also conveyed by the previous Sync Period formula.

The base unit of time measurement on an HVI Engine is the period of its own HVI Engine Clock, but the Physical Propagation Delay is expressed in nanoseconds. To be able to use it, each Engine must express it in Clock cycles, so a conversion is required:

```
Propagation_delay_cycles = Round(Physical_Propagation_Delay/Hvi_Engine_Clock_period)
```

For example, to calculate the Sync frequency for instruments A and B use the formula:

```
Sync = LCM(all instrument core clocks)
```

Instrument A Core clock = 100 MHz, period = 10 ns

Instrument B Core clock = 300 MHz, period = 3.333 ns.

Since 10ns is a multiple of 3.333 ns, the LCM is 10ns. If your instruments are all in 1 chassis, the Physical Propagation Delay constrained by the propagation delay is 100 ns (per the values in the previous table). Therefore, you need to take the next multiple of the LCM = 10 ns which is also equal or greater than 100 ns. This gives the final value of the Sync Period as 100 ns and the Sync signal frequency is 10MHz.

NOTE

You can find the instrument System and Core clocks in the documentation of each instrument.

General Timing Concepts

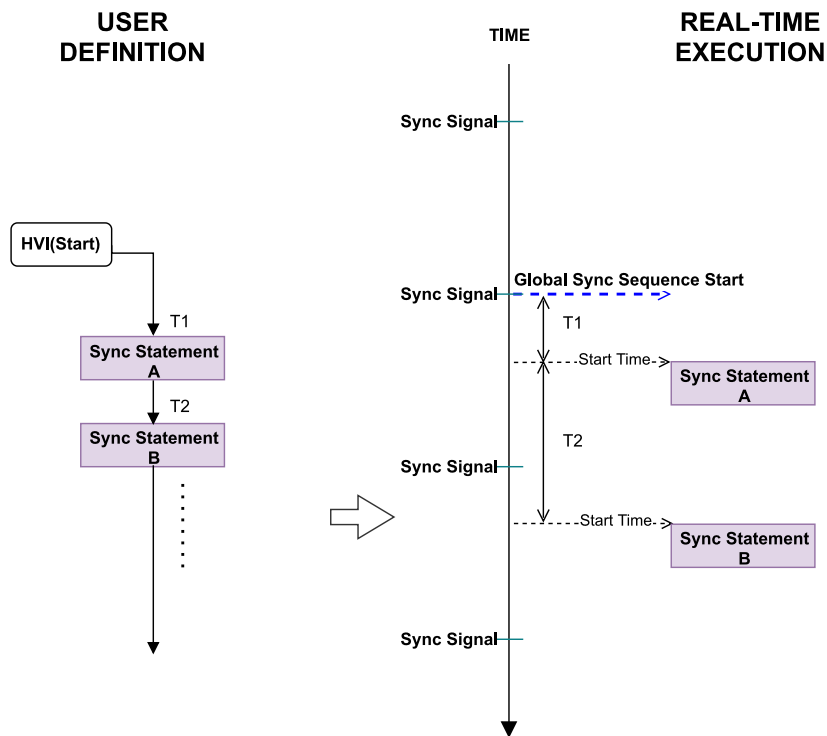
This section introduces general timing concepts that apply to HVI Sequences and Statements. It includes the following sections:

- Global Sync Sequence Start
- Start Time of Statement execution
- Start Delay
- Execution Time
- Internal Sequences

Global Sync Sequence Start

The Global Sync Sequence start is time 0 for the HVI execution, that is the timing point when the HVI Sequences will start executing at the same time in all the Engines added in the SystemDefinition. This timing point is aligned with the arrival of the Sync signal, that is, it always matches the rising edge of the Sync signal (in PXIe systems aligned with the PXIe-SYNC100 signal).

The following diagram shows the Global Sync Sequence Start.



Start Time of Statement execution

The relative time in nanoseconds from the HVI Execution Start Time to the start of the execution of a Statement.

Start Delay

This is the user-defined delay value from the Start Time of the previous Statement to the Start Time of the current Statement. This value can be expressed in seconds or one of its fractions, down to picoseconds. Generally, the valid range is from 0 to +infinity, however the exact range and granularity of this value is defined by the following:

- The acceptable values are multiples either of the **HVI Engine Clock** period (in Local Statements) or, of the **HVI Common Clock** period (in sync Statements).
 - o For example, for a Local Statement for an HVI Engine with Clock frequency of 100MHz, the *clock period* is 10 ns, so the acceptable values are the multiples: 0 ns, 10 ns, 20 ns, etc.
 - o As another example, for a sync Statement, if there are three Engines added to HVI with the frequencies {100MHz, 187.5MHz, 300MHz}, the HVI Common Clock frequency will be 12.5MHz and the *period* is 80 ns, so the acceptable values are the multiples: 0 ns, 80 ns, 160 ns, etc.
 - o The acceptable margin of the value is defined in the Error and Warning Margins section below.

- The minimum possible value is affected by the Start-Latency of the current Statement and the End-Latency of the previous Statement. Formulas to calculate the minimum values are provided in the Timing Tables.
- The maximum possible value is only limited by the actual representation of the value in hardware and software. While this limit in hardware is instrument-dependent, in software it is defined as: The maximum value that can be represented in a signed 64-bit integer value.

The following sections explain how to calculate the Start Delay. When compiling the Sequence, the compiler will report any timing violation and suggest a closer correct value.


NOTE

If you do not specify a valid Start delay, the compiler generates an error and indicates the minimum valid minimum value.

Execution Time

This is the time interval from the Start time until the End time of the Statement. This interval is determined by constraints and inherent limits of the instrument, such as propagation delays and resource availability. Sync and Flow-control Statement execution cannot overlap with other Statements, so in these cases the execution time must be added to the timing calculation. The Start delay of the next Statement from a flow-control or Sync Statement is measured from the end-time of the Statement.

The following diagram shows these concepts in an HVI diagram:



Internal Sequences

Some Sync and all Local Flow-Control Statements are broken into internal Sequences for execution in HVI Engines.

Sync and Local Flow-Control Statement Timing Concepts

This section describes Sync and Local Flow-Control Statement Timing. It includes the following sections:

- [Latency](#)
- [Duration Property](#)

There are several additional concepts and parameters you must be aware of to calculate timing, especially for specifying Start Delays and the Duration property of Statements.

NOTE

The knowledge of these concepts can assist you to understand HVI timing and accurately specify proper values for these timing properties, however it is not mandatory to use them at development time. This is because all limitations are checked by HVI at the time of compilation and any violation is reported with information provided about how it can be resolved. This enables you to focus on its Sequence creation without worrying about complex timing calculations.

Latency

Latency Parameters

The latency parameters are defined for all Sync and Local Flow-Control Statements. They impose a minimum value to the Start Delays of the Statements used in a Sequence:

Start-Latency

This is the minimum number of clock cycles a Sync or Local Flow-Control Statement requires to start execution.

Entry-Latency

This is the minimum number of clock cycles a Local Flow-Control Statement requires to start the execution of the internal Sequence. This imposes a minimum value on the Start Delay of the first Statement of the internal Sequence.

End-Latency

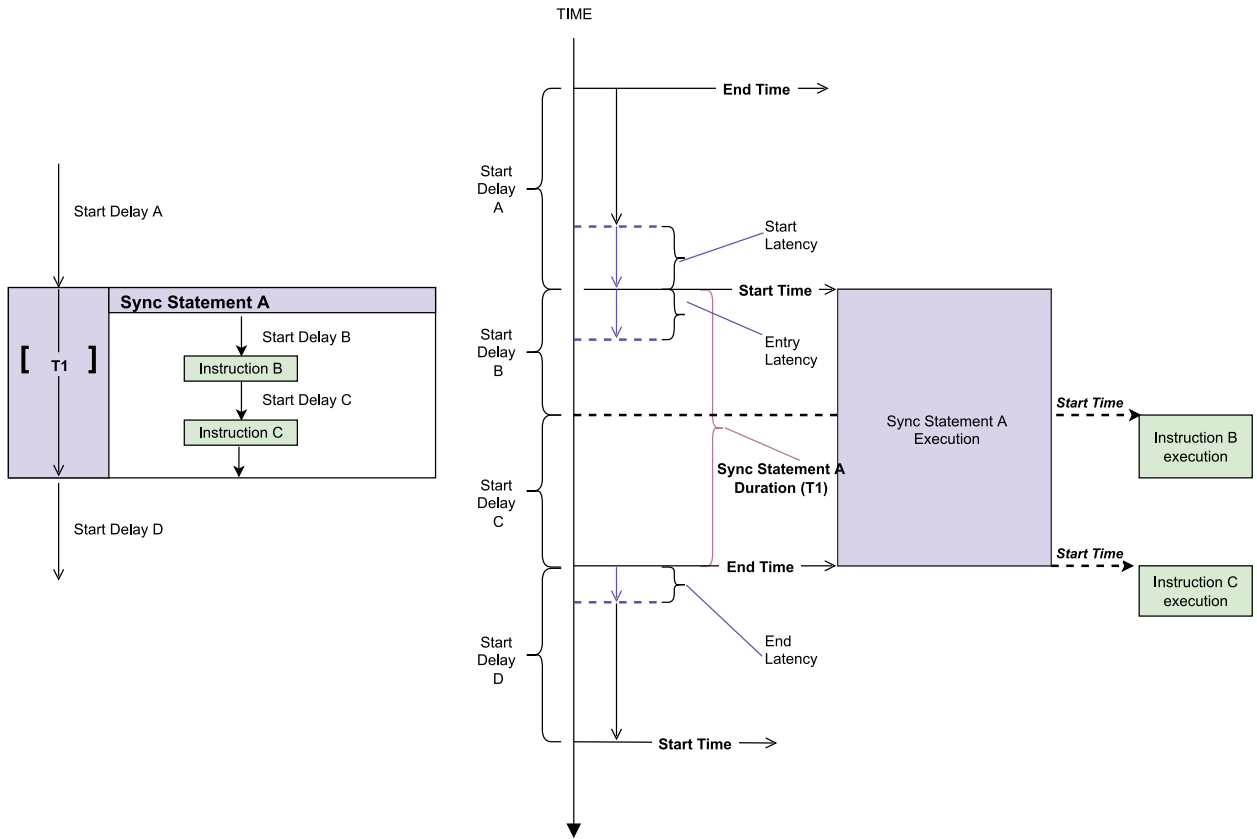
This is the minimum number of clock cycles a Statement requires to exit its execution, before another Statement can be executed.

Iteration Latency (loop Statements)

For loop Statements only, this is the minimum number of cycles a loop Statement requires to start another execution of the internal Sequence after one iteration is completed. This imposes a minimum value on the Start Delay of the first Statement of the internal Sequence.

The exact definitions of Start latency, Entry latency and End latency depend on the type of Statement. Latency values are used in [Sync Statement Timing](#) and [Local Flow-Control Statement Timing](#) and [HVI Instruction Timing](#). The Latency values are listed in [Appendix C: Timing Tables](#).

The following diagram shows the Start, Entry and End Latencies and how they relate to Start Delays:



Duration Property

The Sync Statements and Local Flow-Control Statements If and While include a `duration` property that you can set. The `duration` property enables you to specify the time interval that a Statement takes to execute.

This value can be expressed in seconds or one of its fractions, down to picoseconds. Generally, the valid range is from 0 to +infinity, however the exact range and granularity of this value is defined by the following:

- The acceptable values are multiples either of the **HVI Engine Clock** period (in Local Statements) or, of the **HVI Common Clock** period (in sync Statements).
 - o For example, for a Local Statement for an HVI Engine with Clock frequency of 100MHz, the *clock period* is 10 ns, so the acceptable values are the multiples: 0 ns, 10 ns, 20 ns, etc.
 - o As another example, for a sync Statement, if there are three Engines added to HVI with the frequencies {100MHz, 187.5MHz, 300MHz}, the HVI Common Clock frequency will be

12.5MHz and the *period* is 80 ns, so the acceptable values are the multiples: 0 ns, 80 ns, 160 ns, etc.

- The acceptable margin of the value is defined in the Error and Warning Margins section below.

- The minimum possible value is affected by internal operations of the Statement. For Statements that contain internal Sequences, the minimum is affected also by the Start-Delay and the Duration of the internal Statements. Formulas to calculate the minimum values are provided in the Timing Tables.
- The maximum possible value is only limited by the actual representation of the value in hardware and software. While this limit in hardware is instrument-dependent, in software it is defined as: The maximum value that can be represented in a signed 64-bit integer value.

NOTE

For the loop Statements Local while and Sync while, the **duration** property specifies the execution time of 1 iteration. This means that the overall execution time of a while Statement depends on the number of iterations that are executed. The total execution time is **duration** multiplied by the number of iterations.

If the **duration** is set to a **fixed-time** interval, then the execution time of the Statement shall match the value specified in the **duration** property. If this time cannot be matched an error is generated. For example, this can happen with an if-Statement when more time is required to complete the Statements inside a branch than the duration specified.

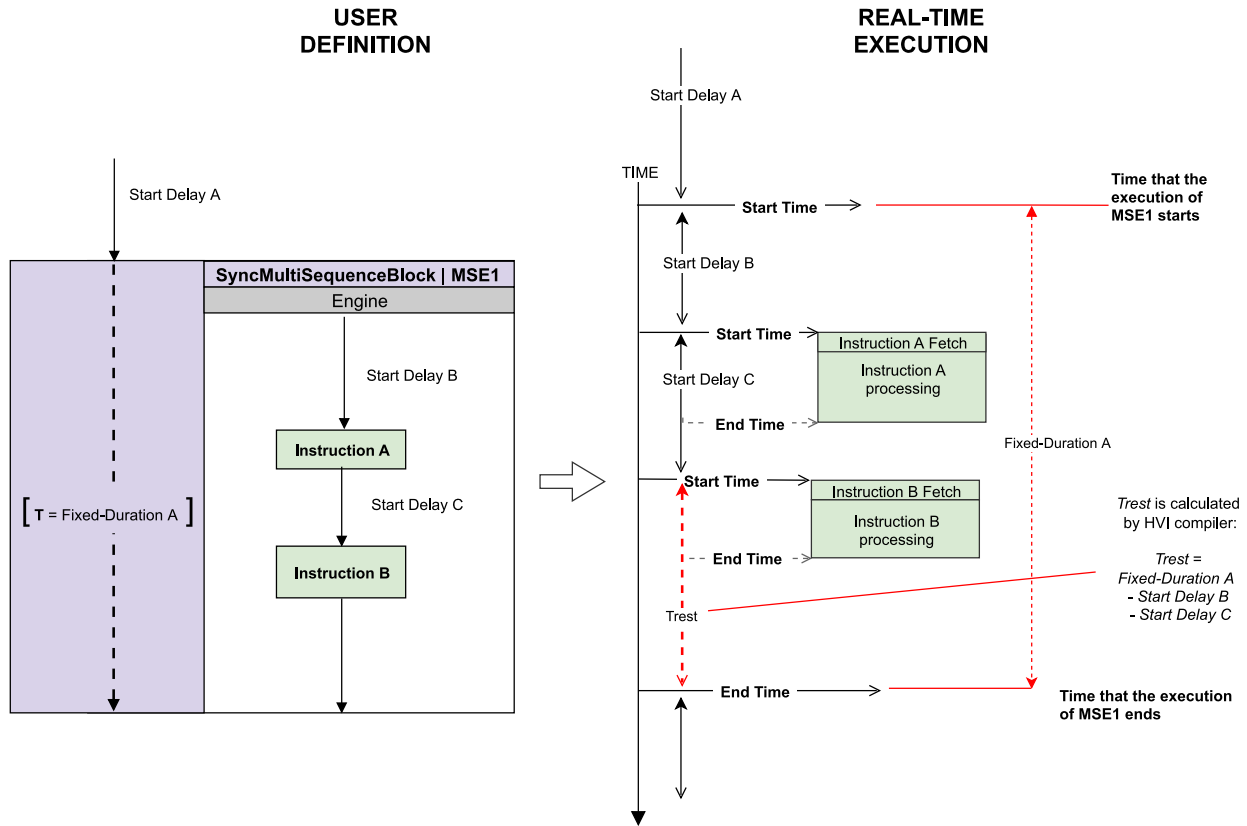
The **duration** property cannot be set to fixed value if there is a flow control Statement inside that has an unknown duration.

If the **duration** is set to a **minimum-time** interval, then the execution time of the Statement is the minimum possible given by the Statements inside.

NOTE

By default, if not specified, **duration** property is set to **minimum-time**.

The following diagram shows how the **duration** property is applied to a Sync Multi-Sequence Block:



Python code for the preceding diagram:

```


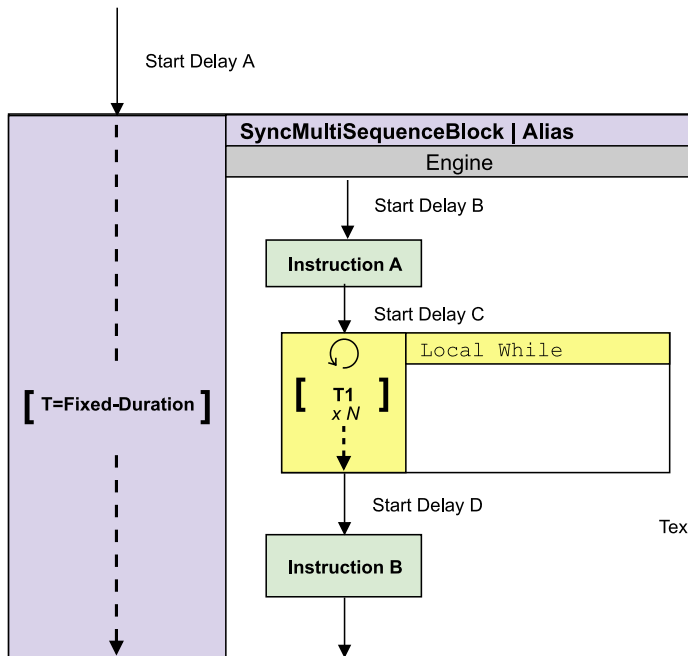
fixed_duration_A = time.Duration(100)
mse1 = sequencer.sync_sequence.add_sync_multi_sequence_block('mse1', start_delay_A)
mse1.duration = fixed_duration_A
sequence = mse1.sequences['Engine1']
instructionA = sequence.add_instruction("instructionA", start_delay_B, sequence.instruction_set.action_execute.id)
instructionB = sequence.add_instruction("instructionB", start_delay_C, sequence.instruction_set.action_execute.id)
    
```

You must not set the **duration** property of a Statement A to a **fixed-time** if the Statement A contains a flow control Statement with an unknown duration (e.g. Local Wait-For-Event, Local Wait-For-Time, Local While, etc.). Doing so will result into an error at compilation.

The following diagram shows a while loop that generates an error if the user would try to set to a fixed-value the duration of the Sync Multi-Sequence Block that contains a Local While Statement:

USER DEFINITION

REAL-TIME EXECUTION



Real-time execution not possible

The `SyncMultiSequenceBlock` is set to use fixed-duration, however, it contains statements that have **unknown execution time** at compile stage

HVI Instruction Timing Concepts

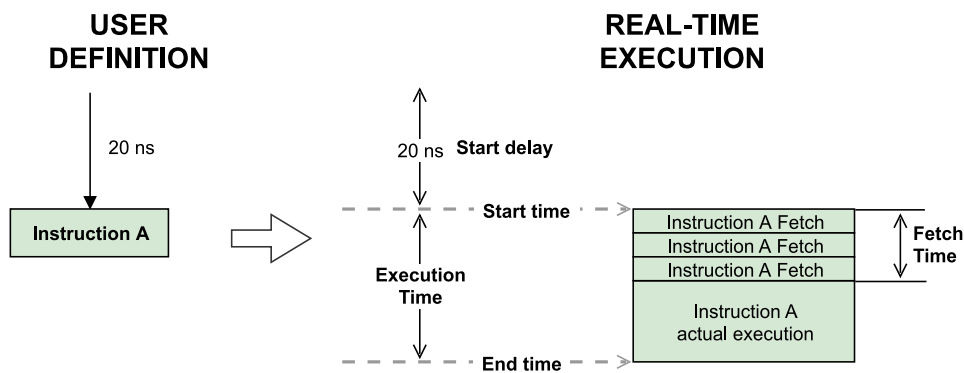
The following section explains HVI Instruction timing. It includes the following sections:

- **Fetch Time**

Fetch Time

This is the time interval required by the HVI Engine to fetch and dispatch an HVI Instruction for processing. The Fetch time consumes HVI Engine execution cycles. A Statement may take several HVI Engine cycles to complete the fetch before processing can start. The number of cycles a fetch takes depends on the Statement or instruction characteristics, for instance, the number of parameters.

The following diagram shows the fetch time with other timing definitions:



Sync Statement Timing

This section describes Sync Statement timing. It contains the following sections:

- [About Sync Statement Timing](#)
- [How to use the Timing Tables for Sync Statements](#)
- [HVI Start](#)
- [Sync Multi-Sequence Block Timing](#)
- [Sync While](#)
- [Sync Register-Sharing](#)
- [Sync FPGA Data-Sharing](#)

About Sync Statement Timing

Sync Statements consume HVI Engine execution time and cannot overlap their execution with other Statements. Their start and end is synchronized and happens at the same HVI Common Clock cycle across all the Engines participating in the system. The Start delay of a Sync Statement is measured from the end of previous Sync Statement to the start of the current one.

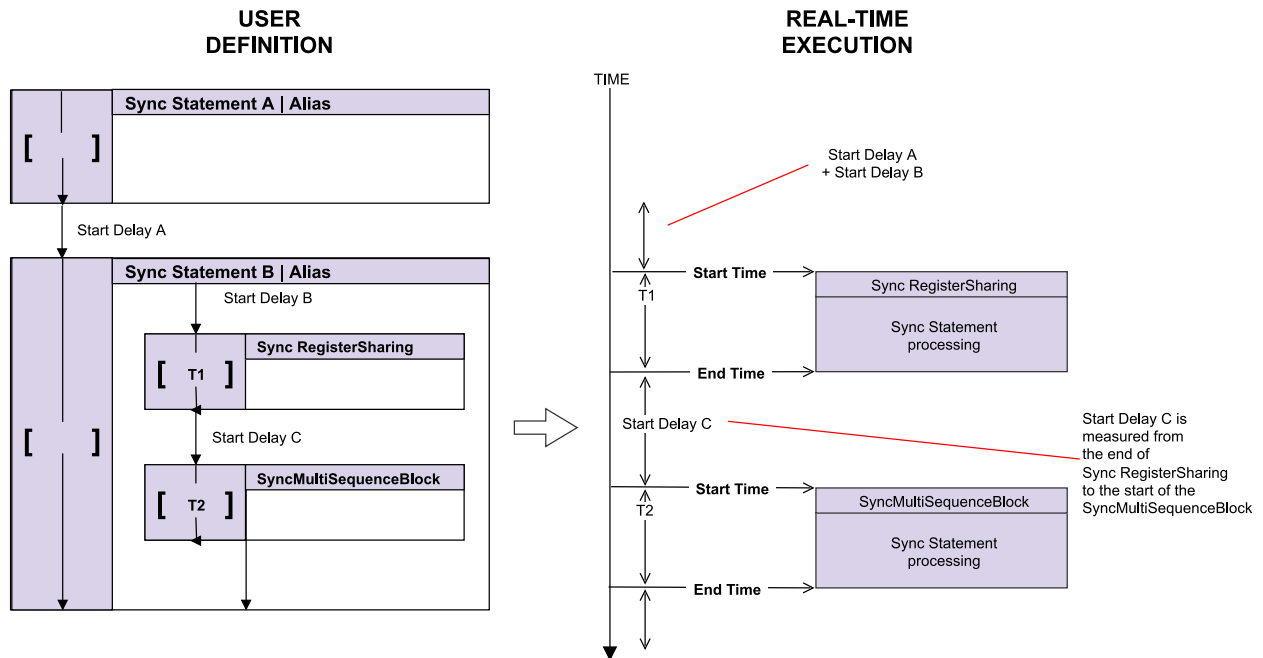
The following diagram shows the timing between a number of Sync Statements including a Sync Register-Sharing Statement and Sync Multi-Sequence Block Statement.

The diagram shows two Sync Statements A and B. Sync Statement B is a container for two further Sync Statements: Sync Register-Sharing and Sync Multi-Sequence Block. The times indicated are **Start Delay A**, **Start Delay B**, **Start Delay C**, T1, and T2.

The time between the end of Sync Statement A and the start of Sync Register-Sharing is **Start Delay A + Start Delay B**.

The time between the end of Sync Register-Sharing and the start of Sync Multi-Sequence Block is **Start Delay C**.

Sync Register-Sharing and Sync Multi-Sequence Block timing:



How to use the Timing Tables for Sync Statements

All the timings provided in the tables below are expressed in HVI Engine Clock cycles. To use them to calculate minimum start delay values, see the explanations in [Minimum Start Delay Calculation for Flow-Control and Sync Statement](#). Start delays of Local Statements must be multiple of the HVI Engine Clock, whereas Sync Statements must be multiples of the HVI Common Clock, see [General Timing Concepts](#).

Leader Engine

In some of the Sync Statements, one of the Engines that leads the Statement operation. For example, in a Sync While Statement, the Engine that leads is the one where the condition is evaluated. For the context of Timing Latency calculation, we are going to call this Engine the *Leader Engine*.

Rounding Delays

When a latency value needs to be applied to multiple Engines, we must round the Engine cycles to the next HVI Common Clock cycle. We do this using the following formula (below **ncc** stands for **n** ext **c** ommon **c** lock):

$$\text{round}_{\text{ncc_cycles}}(\text{TimeValue}_{\text{EngineCycles}}) = \text{ceil}(\text{TimeValue}_{\text{EngineCycles}} * \text{HVI_Engine_ClockPeriod} / \text{HVI_Common_ClockPeriod}) * \text{HVI_Common_ClockPeriod} / \text{HVI_Engine_ClockPeriod}$$

NOTE

In the case that all the Engines are running at the same frequency, the HVI Engine Clock cycles and the HVI Common Clock cycles will have the same value for all the Engines. Therefore, you can skip the rounding calculation because it has no effect:

$$\text{round}_{\text{ncc_cycles}}(\text{TimeValue}_{\text{EngineCycles}}) == \text{TimeValue}_{\text{EngineCycles}}$$

Matching Delays

Some parts of the latency may need to be aligned between Engines. In order to achieve this, we use the following formula:

$$\text{match}\{\text{TimeValue}_{\text{Engine1Cycles}}, \text{TimeValue}_{\text{Engine2Cycles}}, \dots, \text{TimeValue}_{\text{EngineNCycles}}\} = \text{round}_{\text{ncc_cycles}}(\max\{\text{TimeValue}_{\text{Engine1Cycles}}, \text{TimeValue}_{\text{Engine2Cycles}}, \dots, \text{TimeValue}_{\text{EngineNCycles}}\})$$

In the previous formula, all the **TimeValues** have to first be converted to the **EngineCycles** of the target Engine so that the **max** can be applied among similar quantities. This can be done using this formula:

$$\text{TimeValue}_{\text{TargetEngineCycles}} = \text{ceil}(\text{TimeValue}_{\text{OtherEngineCycles}} * \text{HVI_OtherEngine_ClockPeriod} / \text{HVI_TargetEngine_ClockPeriod})$$

NOTE

In the case that all the Engines are running in the same frequency, the calculation of the match formula is just the time value of the Leader Engine:

$$\text{match}\{\text{TimeValue}_{\text{Engine1Cycles}}, \text{TimeValue}_{\text{Engine2Cycles}}, \dots, \text{TimeValue}_{\text{EngineNCycles}}\} == \text{TimeValue}_{\text{LeaderEngineCycles}}$$

Physical Propagation Delay

The Physical Propagation Delay corresponds to the amount of time (expressed in nanoseconds) that a PXIe Trigger needs, to cover the path between any given pair of segments in a topology. This value is used when running Sync Statements because it provides information about how long the execution signaling between modules takes.

The Physical Propagation Delay is expressed in nanoseconds and its value depends on the topology. A table with the values is defined on the page [Synchronization Clocks, Signals, and Modes](#).

In the context of the Timing Tables, the value is expressed in cycles of the HVI Engine, as it is shown in the timing tables. To be able to use it, a conversion is required:

$$\text{Propagation_delay}_{\text{EngineCycles}} = \text{Round}(\text{Physics_Propagation_Delay}_{\text{Seconds}} / \text{HVI_Engine_Clock}_{\text{Period}})$$

HVI Start

This is the time 0 for the HVI execution. It always matches the rising edge of the Sync signal (in PXIe systems aligned with the PXIe-SYNC100 signal).

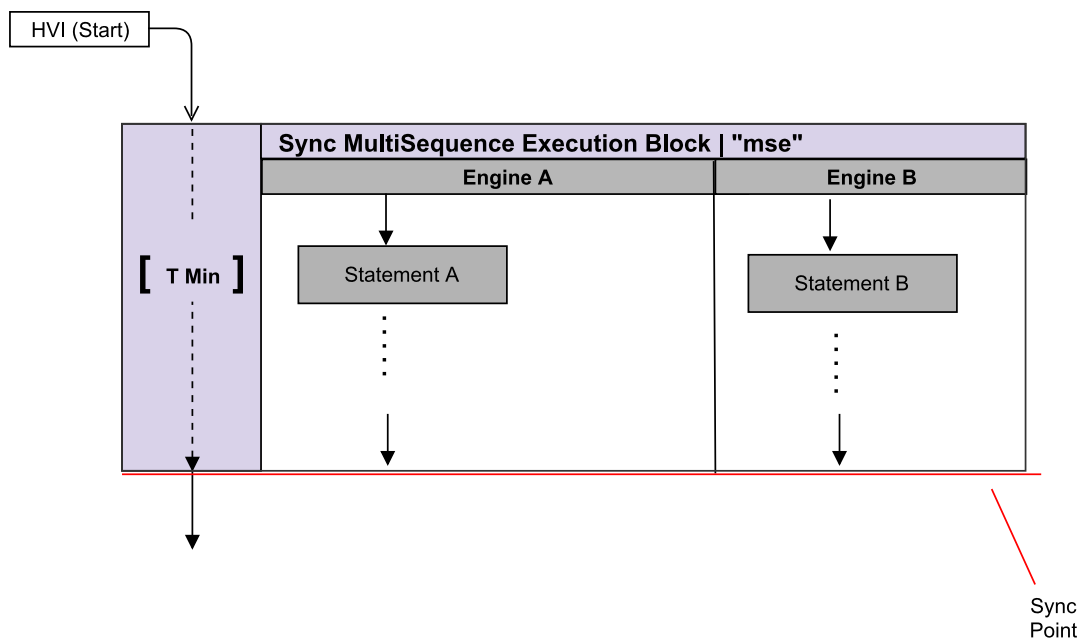
HVI start basic timing value:

Parameter	Time (cycles)
End-Latency	2

Sync Multi-Sequence Block Timing

In a synchronized multi-Sequence block, you can define the Statements that the HVI Engines execute in parallel with other Engines.

Local Sequences start and end their execution within the Sync Multi-Sequence Block synchronously.



HVI automatically calculates the execution time of each local Sequence and adjusts the execution of all local Sequences within the Sync Multi-Sequence Block so that they all can end together deterministically.

The individual Sequences can have different execution times, so HVI automatically adjusts the timing of each individual Sequence. The final time is calculated automatically.

There are two cases for the **Sync-Point** that are treated in different ways by HVI:

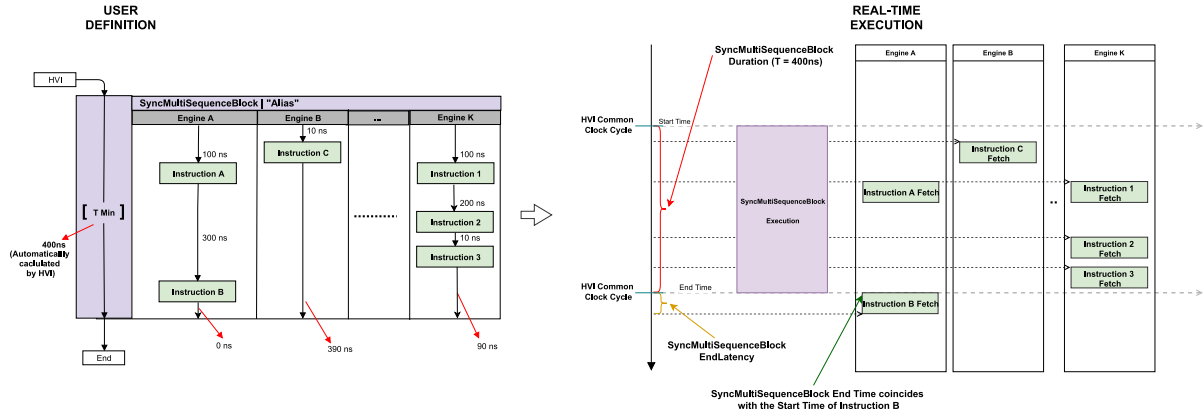
- **Timed-Sync:** When the execution time is known at HVI compilation time for all Local Sequences within the Sync Multi-Sequence Block.
- **Triggered-Sync:** When the execution time is unknown at HVI compilation time for one or more Local Sequences within the Sync Multi-Sequence Block.

Timed-Sync (Sync Multi-Sequence Block containing Local Sequences with known total execution time)

For Sync Multi-Sequence Blocks that contain HVI Instructions or Local Flow-Control Statements with execution times that are known at HVI compilation time, the HVI compiler accounts for the different Sequence execution times during compilation and then adjusts the final times. This ensures all of the Local Sequences reach the end of the Sync Multi-Sequence Block at the same time.

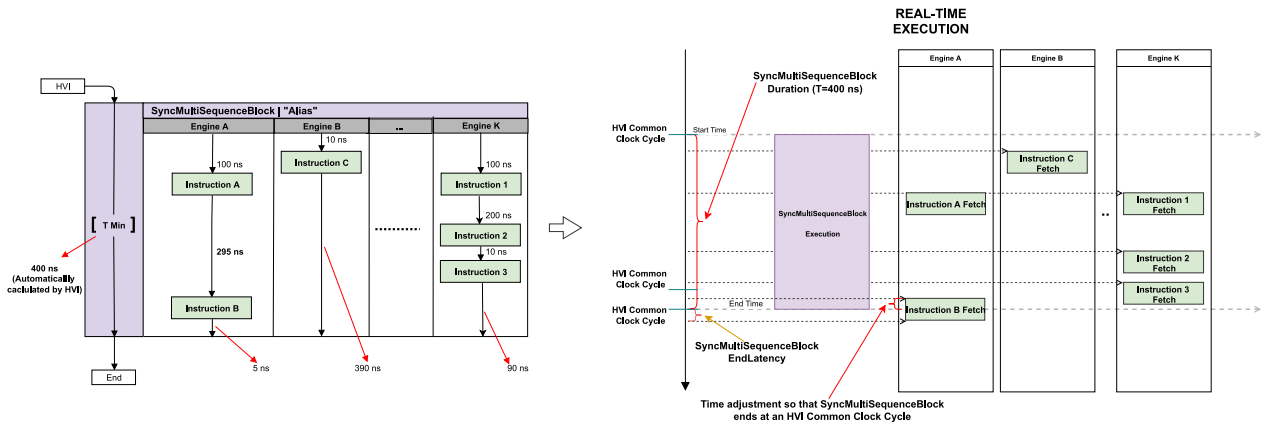
When the execution time (duration property) of the Sync Multi-Sequence Block is not specified, the compiler adjusts the total execution time to be the minimum possible to allow the execution of the longest Local Sequence. Note that in the case that the Engines participating in the system do not share the same frequency, HVI will automatically adjust the duration (or execution time) of the Sync Multi-Sequence Block Statement to a multiple of the HVI Common Clock.

For example, in the diagram below, the total time for Engine A is 400 ns. HVI calculates the times required for the other Engines to finish at the same time. For Engine B this is 390 ns, for Engine K this is 90 ns.



Using the previous example but assuming that the Engine A runs at 200 MHz, while the rest Engines run at 100 MHz, the common clock cycle will happen at multiples of 10 ns. In the following diagram we can notice two things:

- The duration of the longest Sequence (Engine A) is 395 ns, which is not a multiple of a common clock cycle. Therefore, HVI will adjust the end of the Sync Multi-Sequence Block to the next Common Clock Cycle at 400 ns and then make sure the Sequences of all the Engines match this time.
- The start of the end-latency of the Statement will not start from the start time (395ns from the beginning of the Sync Multi-Sequence Block Statement) of the last Statement of the longest Sequence (Engine A) because it is not at a common clock cycle. Rather, it will start from the next Common Clock Cycle, at 400 ns.

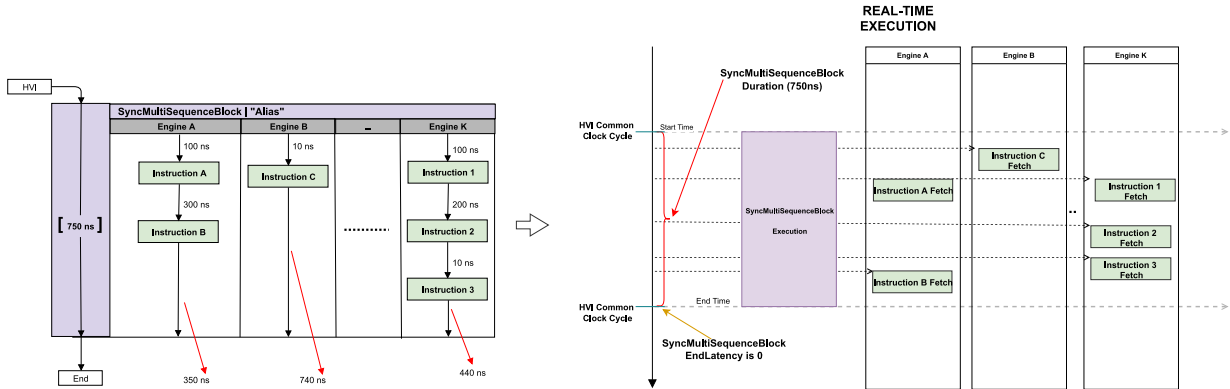


Sync Multi-Sequence Block with a specific execution time (duration property)

When the execution time (duration property) of the Sync Multi-Sequence Block is specified, the compiler verifies that the specified execution time is enough to allow the execution of the longest Local Sequence, if not an error is generated. Note that in the case that the Engines participating in

the system do not share the same frequency, the specified execution time (duration property) must be a multiple of the HVI Common Clock.

In the following diagram, the times of the HVI Instructions and the delays between them are known, so the timing between them and for the entire block can be calculated. In this case the total time is specified at 750 ns. The HVI calculates the times required for all the other Engines to finish at the same time. For Engine A this is 350 ns, for Engine B this is 740 ns, for Engine K this is 440 ns.



Triggered-Sync (Sync Multi-Sequence Block containing Local Sequences with unknown execution times)

In some cases, one or more of the local Sequences within the Sync Multi-Sequence Block include a Local Flow-Control Statement that has an execution time that is unknown at HVI compilation time. At the point in the Local Sequence where the unknown execution time is encountered, the Local Sequence becomes de-synchronized and **an active triggering process is required** at the end of the Sync Multi-Sequence Block to re-synchronize the execution of all HVI Engines. This guarantees that all the HVI Engines then continue execution at exactly the same point after the **Triggered-Sync point**. The execution resumes in all HVI Engines at the same time, aligned with a sub-sequent Sync pulse, which forces the execution to be aligned to a multiple of the Sync period of the main Sync signal. Triggered-sync points require the use of Trigger resources assigned in the `SyncResources` property in the `SystemDefinition` instance and the main Sync signal.

Possible cases of the unknown execution time is when one of the Local Sequences contain:

- A Local Wait-for-time Statement with an HVI Register defining the wait time at runtime.
- A Local Wait-for-Event Statement.
- A Local While Statement.
- A Local If Statement with unmatched branches, that take different execution times.

NOTE Note that specifying the execution time (duration property) in this scenario is not allowed and will lead to a compilation error.

Triggered-Sync delay

A Triggered-sync point adds a delay to the Sequence timing that has four parts. Two of them are constant and the other two vary depending on the last Statement and its position compared to the Sync pulse time. The formula to calculate the delay is:

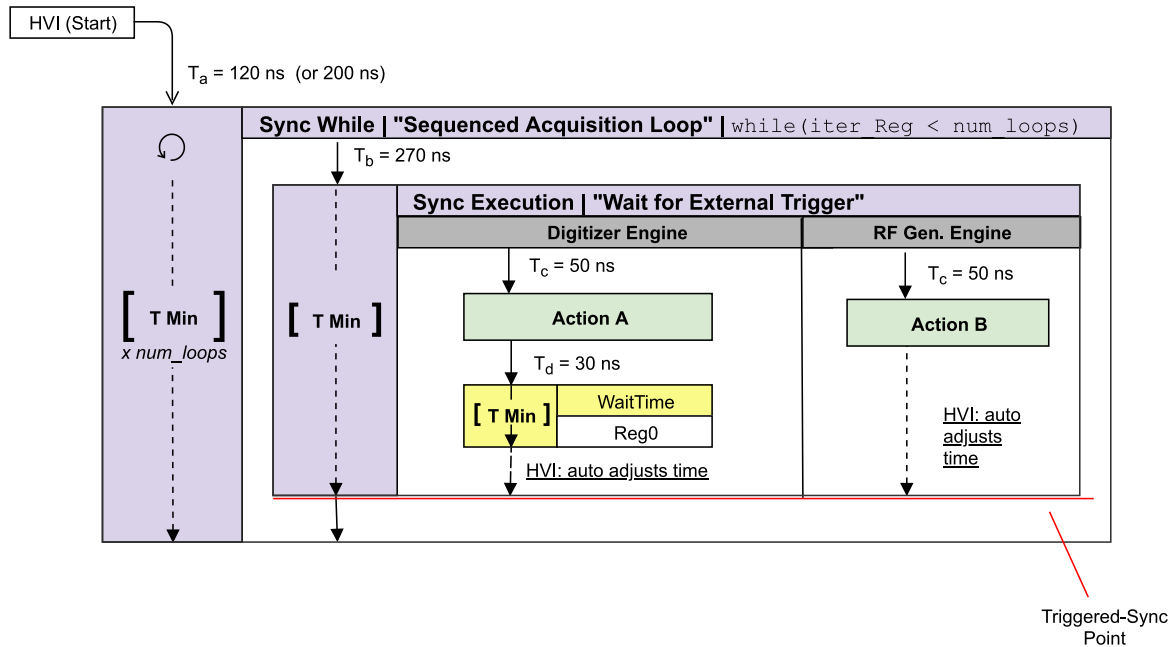
$$\text{triggered_sync_delay} = \text{end_latency} + \text{sync_overhead} + \text{edge_offset} + \text{sync_period}$$

where:

- `end_latency` is the End-latency of the last Statement before the resync. If the last Statement is an HVI Instruction, this is equal to its Fetch time.
- `sync_overhead` is constant per instrument. Its value is 3 cycles.
- `edge_offset` is the time interval from the end of the `sync_overhead` to the sync-pulse edge. This time can vary depending on the position of the last Statement compared to the Sync pulse time.
- `sync_period` is constant per configuration and is calculated by the equation defined in [Synchronization Clocks, Signals, and Modes](#).

Example of timing management with Triggered-Sync point

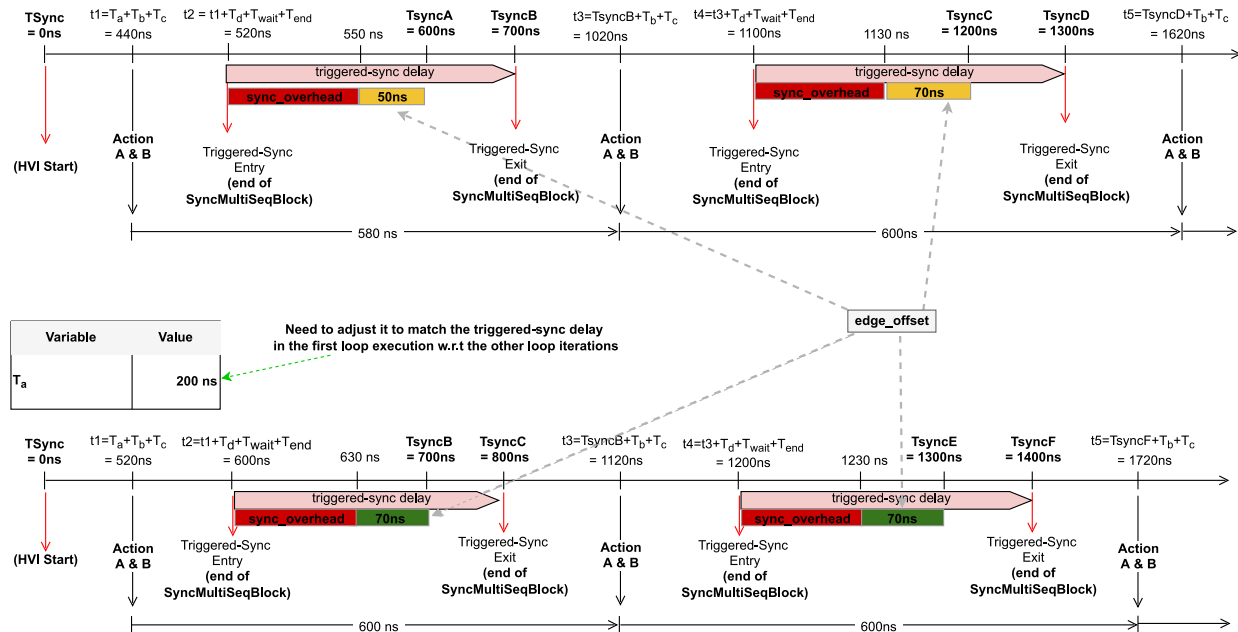
The following diagram shows an example with a simple Sequence where the Triggered-sync point is marked in red. The Triggered-sync point is at the end of the Sync Multi-Sequence Block and it is required because there is a WaitTime Statement and the time for this cannot be determined at compile time.



The following table shows the Variables and their execution times:

Variable	Value	Description
T_a	120 ns	Start delay of Sync-while Statement
T_b	270 ns	Start delay of Sync Multi-Sequence Block
T_c	50 ns	Start delay of Action A HVI Instruction
T_d	30 ns	Start delay of Wait-for-time Statement
Reg0	4	The HVI Register used for the Wait-for-time
T_{WAIT}	40 ns	The total wait time based on the value on the value of Reg0
T_{END}	10 ns	End-latency of Wait-for-time Statement
T_{sync_period}	100 ns	Sync period for 1 chassis
$T_{sync_overhead}$	30 ns	Sync overhead

The following diagrams shows the execution timeline for the first 3 iterations of the Sequence shown in the previous diagram, it is important to note that the first Triggered-sync aligns the execution with the Sync pulse and consequently the duration (or execution time) for the following cycles will be different, this effect is in some cases seen as a skew (or jitter) in the 1st cycle. A way to eliminate the first cycle variation is to adjust the Sync While start time.



Jitter when waiting for external Events or Triggers

The triggered-sync is controlled by the SYNC signal. This means that repeated executions, for example, inside a Sync While loop of an Sync Multi-Sequence Block that contains a WaitEvent, may show jitter of the Sync Multi-Sequence Block actions with respect to the Event that is in the WaitEvent condition. In these cases, the maximum skew variation (or jitter) is the maximum time difference between Trigger Events and the Sync Period. The Variable skew (or jitter) value can be:

- 0 => when the Trigger Events have the same time delay with respect to the SYNC signal.
- Sync period => When the Trigger Events are asynchronous and at a rate that is not multiple of the Sync period.
- If more than one synchronization signal is used (SYNC, SYNC_BASE, etc.), the largest will dominate:
 - A Sync Multi-Sequence Block always aligns its start to the SYNC signal, so at least the jitter for an asynchronous Event will be equal to the Sync period.
 - For example, if you also re-sync the Wait-For-Event with the SYNC_BASE (by using the SYNC_BASE SyncMode), and the Trigger is asynchronous to the SYNC_BASE, then the jitter will be equal to SYNC_BASE.
- The SYNC and SYNC_BASE periods depend on any Non-Hvi clocks (core/system) added to the systemDefinition using the TSE API.

Sync Multi-Sequence Block Timing Tables

Timing value for Sync Multi-Sequence Blocks:

Execution time (cycles) (1)
$\text{round}_{\text{ncc_cycles}} \left(\text{sum}_{\text{for_all_internal_statements}}(\text{StartDelayCycles}) + \text{sum}_{\text{for_all_internal_flow_control_statements}}(\text{DurationCycles}) \right) \quad (2)$

The following tables shows latency values for Sync Multi-Sequence Blocks:

Parameter	Description	Time(cycles)
Start-Latency	Minimum start-delay for Statement	1
Entry-latency	Minimum start-delay for first Statement inside any of the contained Sequences	1
End-Latency	Minimum start-delay for the next Statement	$\text{round}_{\text{ncc_cycles}}(\text{End-Latency}_{\text{Last-statement-of-longest-branch}}(3) - 1)$ * if the last Statement of the longest branch is not starting from a common clock cycle (see section Sync Multi-Sequence Block Timing and Time Matching in Sync Statement Timing), the formula is updated to: $\text{round}_{\text{ncc_cycles}}(\text{End-Latency}_{\text{Last-statement-of-Longest-branch}}(3) - 1 - \text{DistanceToNextCommonClock})$ where: - <i>DistanceToNextCommonClock</i> is the number of Engine Cycles from the start of the last Statement to the following common clock cycle.
		timed-sync (5) Minimum Duration 0
		triggered-sync (5) Fixed Duration 0

Fixed-Duration	Minimum fixed-duration for Statement	$\text{round}_{\text{ncc_cycles}}([\text{max}_{\text{for_all_Sequences}}[\text{Sequence-Duration}]])(4),$ <p>where <code>Sequence-Duration</code> is calculated as follows:</p> $\text{sum}_{\text{for_all_internal_statements}}(\text{Start-Delay}) + \text{sum}_{\text{for_all_internal_flow_control_statements}}(\text{Duration}) + \text{End-Latency}_{\text{Last-statement}} - 1$
----------------	--------------------------------------	---

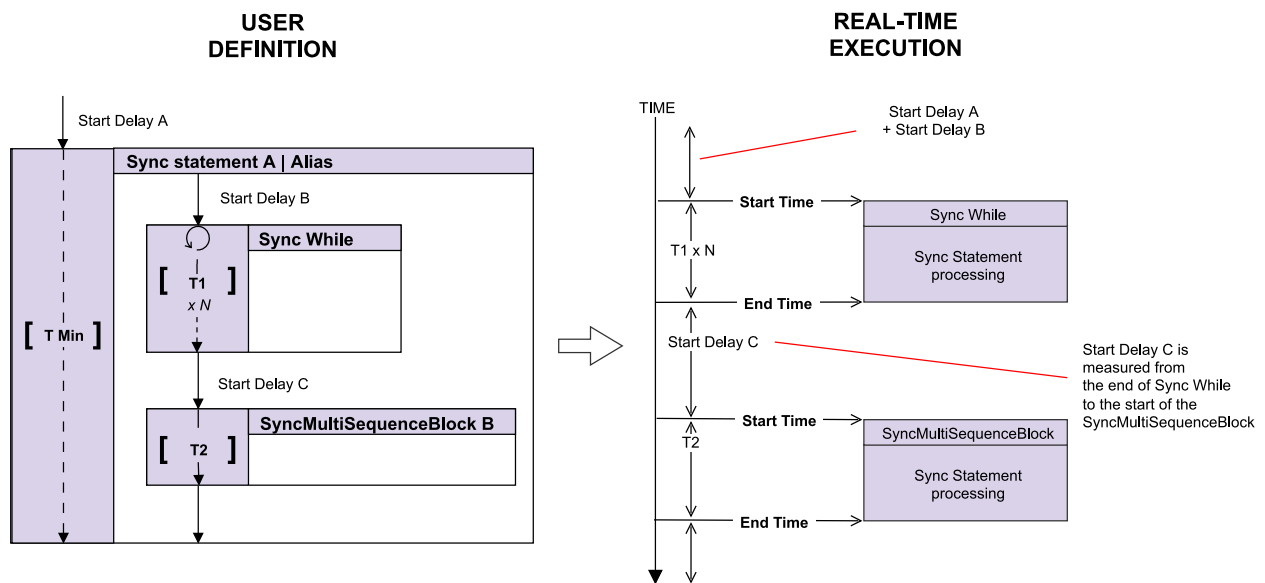
- (1) The values provided here apply if the duration property of the Statement is set to Minimum (default). If a fixed-duration has been set, then the Execution time is equal to that value.
- (2) The values are only calculated for the branch that is being executed, if there are multiple branches available.
- (3) If the Sequence is empty, the value is 0.
- (4) If the Sequence is empty, then the duration is 0.
- (5) Triggered-sync is required if any of the Sequences in a Sync Multi-Sequence Block contains a Statement that has unknown execution time at compile time.

Sync While

For the Sync While flow-control Statement, the timing is different compared to other Sync Statements. The Sync While Statement continues operation while a condition is met. It stops executing when the condition is no longer met.

The following diagram shows a Sync While Statement with other Sync Statements. The time for an iteration of Sync While is $T_2 \times N$, where T_2 is the time per iteration and N is the number of iterations. The time cannot be indicated exactly on a diagram or in code because the number of iterations is not known until runtime.

The time for the containing Statement Sync Statement A cannot be indicated because it contains a flow-control Statement. This is indicated by the dotted line and the time indicated as T_{min} .



Sync While Timing Tables

Timing value for Sync While Statement:

Execution time (cycles) (1)
$\text{round}_{\text{ncc_cycles}} \left(\#Iterations * [\text{sum}_{\text{for_all_internal_statements}}(\text{StartDelay}_{\text{Cycles}}) + \text{sum}_{\text{for_all_internal_flow_control_statements}}(\text{Duration}_{\text{Cycles}})] \right)$

The following tables shows latency values for the Sync While Statement:

Parameter	Description		Time (cycles)
Start-Latency	Minimum start-delay for Statement		<ul style="list-style-type: none"> Leader Engine: $6 + \#Register_Conditions$ Follower Engine(s): 2
Entry/Iteration latency	Minimum start-delay for first Statement inside the while loop	Minimum Duration	$\text{match}\{LatencyA_{LeaderEngine}, LatencyA_{FollowerEngine1}, \dots\} + 2 + \text{End-Latency}_{Last-statement}(2)$ <p>where <i>LatencyA</i> is :</p> <ul style="list-style-type: none"> Leader Engine(3): $12 + \#Register_Conditions + Instrument_SyncResources_Latency(4) + Propagation_delay_{Cycles}$ Follower Engine(s): 2
		Fixed Duration	$\text{match}\{LatencyB_{LeaderEngine}, LatencyB_{FollowerEngine1}, \dots\} + 2$ <p>where :</p> <ul style="list-style-type: none"> $LatencyB = LatencyA - 1$ <i>LatencyA</i> as defined above

Parameter	Description		Time (cycles)
End-Latency	Minimum start-delay for next Statement outside the while loop	Minimum Duration	$\text{match}\{\text{LatencyA}_{\text{LeaderEngine}}, \text{LatencyA}_{\text{FollowerEngine1}}, \dots\} + \text{match}\{\text{LatencyC}_{\text{LeaderEngine}}, \text{LatencyC}_{\text{FollowerEngine1}}, \dots\} + \text{End-Latency}_{\text{Last-statement}}(2),$ <p>where:</p> <ul style="list-style-type: none"> • <i>LatencyA</i> as defined above • <i>LatencyC</i> is 2 for each Engine.
		Fixed Duration	$\text{match}\{\text{LatencyB}_{\text{LeaderEngine}}, \text{LatencyB}_{\text{FollowerEngine1}}, \dots\} + \text{match}\{\text{LatencyC}_{\text{LeaderEngine}}, \text{LatencyC}_{\text{FollowerEngine1}}, \dots\},$ <p>where:</p> <ul style="list-style-type: none"> • <i>LatencyB</i> as defined above • <i>LatencyC</i> as defined above
Fixed-Duration	Minimum fixed-duration for Statement	Sync-While Branch with at least one Statement inside	$\text{round}_{\text{ncc_cycles}}(\text{sum}_{\text{for_all_internal_statements}}(\text{StartDelay}_{\text{Cycles}}) + \text{sum}_{\text{for_all_internal_flow_control_statements}}(\text{Duration}_{\text{Cycles}}) + 1 + \text{End-Latency}_{\text{Last-statement}})$
		Empty Sync-While Branch	$\text{match}\{\text{LatencyB}_{\text{LeaderEngine}}, \text{LatencyB}_{\text{FollowerEngine1}}, \dots\} + \text{match}\{\text{LatencyC}_{\text{LeaderEngine}}, \text{LatencyC}_{\text{FollowerEngine1}}, \dots\} + 1,$ <p>where:</p> <ul style="list-style-type: none"> • <i>LatencyB</i> as defined above • <i>LatencyC</i> as defined above

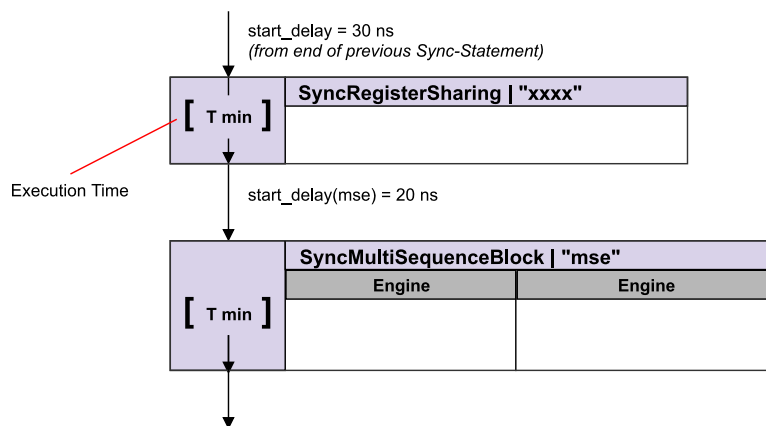
Parameter	Description	Time (cycles)
Register Evaluation Latency	Time to evaluate the Register condition	Leader Engine (Only): <ul style="list-style-type: none"> • From start: 2 • For each iteration: $-(3 + \#Register_Conditions)$

- (1) This value applies if the duration property of the statement is set to Minimum (default). If a fixed-duration has been set, then the Execution time is equal to that value.
- (2) If the sequence is empty, the value of `End-LatencyLast-statement` is 0.
- (3) In the context of this statement, Leader is the engine that contains the Register or Registers used in the while condition.
- (4) `Instrument_SyncResources_Latency` is an instrument specific value. For more information see the instrument documentation.

Sync Register-Sharing

The Sync Register-Sharing Statement execution time must be accounted for when calculating the Sync Sequence timing.

The following diagram shows Sync Register-Sharing Statement followed by a Sync Multi-Sequence Block:



Sync Register-Sharing Timing Tables

Sync Register-Sharing latency does not depend on the number of bits shared. For more information on this functionality, see [Sync Statements](#).

Timing value for Sync Register-Sharing Statement:

Execution time (cycles) (1)
$\text{round}_{\text{ncc_cycles}}(5 + \text{Propagation_delay}_{\text{cycles}})$ (2)

Latency values for Sync Register-Sharing Statement:

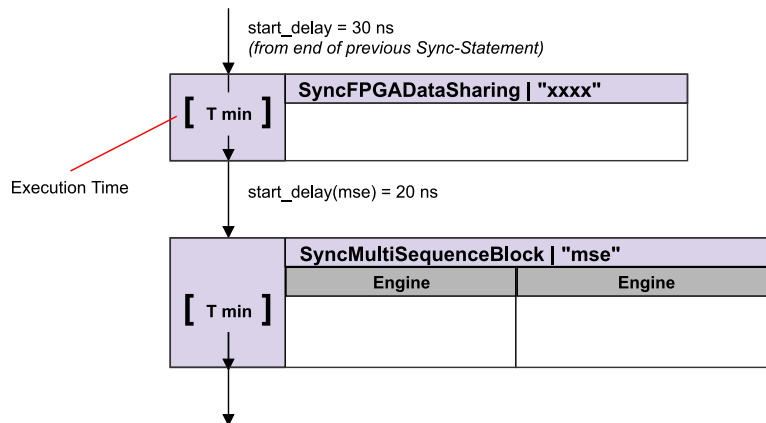
Parameter	Description	Time (cycles)
Start-Latency	Minimum start-delay for Statement	1
End-Latency	Minimum start-delay for the next Statement	0
Fixed-Duration	Minimum fixed-duration for Statement	$\text{round}_{\text{ncc_cycles}}(5 + \text{Propagation_delay}_{\text{cycles}})^{(2)}$
Register Evaluation Latency	Time to evaluate the Register condition	-1

(1) The value provided here applies if the duration property of the Statement is set to Minimum (default). If a fixed-duration has been set, then the Execution Time is equal to that value.

(2) This latency needs to be calculated only on the Leader Engine. In the context of this Statement, Leader is the Engine that contains the Register(s) used as source.

Sync FPGA Data-Sharing

The Sync FPGA data-sharing Statement enables you to share data between the FPGA-Sandboxes on different instruments in the same or different chassis.



Sync FPGA Data-Sharing Timing Table

Calculating execution times for Sync FPGA Data-Sharing can be a complex process. This is because Sync FPGA Data-Sharing execution time depends on a number of factors:

- Instrument specific delay characteristics.
- The topology of your system.
- The amount of data be transferred.
- If the transfer of data is in a single chassis or if it is between different chassis.
- The scheduling of multiple transactions.

Latency values for Sync FPGA Data-Sharing Statement:

Parameter	Description	Time (cycles)
Start-Latency	Minimum start-delay for Statement	1
End-Latency	Minimum start-delay for the next Statement	0

Estimating the execution time of a Sync FPGA Data-Sharing Statement

Sync FPGA Data-Sharing execution time depends on several factors:


- Instrument specific delay characteristics.
- The topology of the system, for instance, the number of chassis and the System Sync connectivity topology across chassis.
- Location of the source and destination instruments, such as if the data transfer happens in a single chassis or through multiple chassis.
- The amount of data to be transferred.
- The scheduling of multiple transactions.

In order to assist you estimate of the execution time, the following examples are provided that abstract some of the complexities, these should enable you to make good estimates of execution time.

Latency Equations

To estimate the Sync FPGA Data-Sharing Statement execution time, a number of equations are provided below.

The following image shows the example topology used in the examples:



The formula to calculate the time it takes to send a single nibble (4-bits) of data from Instrument 1 to Instrument 2 (in the same chassis) is:

$$T_{\text{single_chassis}} = T_{\text{tx_latency}} + T_{\text{link_latency}} + T_{\text{ssm_latency}} + T_{\text{link_latency}} + T_{\text{rx_latency}} = T_{\text{tx_latency}} + 2 * T_{\text{link_latency}} + T_{\text{ssm_latency}} + T_{\text{rx_latency}}$$

The formula to calculate the time it takes to send a single nibble of data from Instrument 1 to Instrument 3 (in a different chassis) is:

$$T_{\text{two_chassis}} = T_{\text{tx_latency}} + T_{\text{link_latency}} + T_{\text{ssm_latency}} + T_{\text{link_latency}} + T_{\text{ssm_latency}} + T_{\text{link_latency}} + T_{\text{rx_latency}} = T_{\text{tx_latency}} + 3 * T_{\text{link_latency}} + 2 * T_{\text{ssm_latency}} + T_{\text{rx_latency}}$$

To generalize this to N number of SSMs hops, the formula is:

$$T_{\text{N_chassis}} = T_{\text{tx_latency}} + (N_{\text{SSM}} + 1) * T_{\text{link_latency}} + N_{\text{SSM}} * T_{\text{ssm_latency}} + T_{\text{rx_latency}}$$

Since the previous equations are for sending only one nibble of data, if you want a transaction with more bits, this will be split into multiple of 4-bit transactions happening one after the other on consecutive clock cycles. For N number of data bits (which must always be multiple of 4-bits), the equation is:

$$T_{\text{transaction_duration}} = T_{\text{tx_latency}} + (N_{\text{SSM}} + 1) * T_{\text{link_latency}} + N_{\text{SSM}} * T_{\text{ssm_latency}} + T_{\text{rx_latency}} + N_{\text{num_bits}} / 4$$

Constants Estimation

In the formulas above, the constants of the different latency equations are *instrument-specific*. See your instrument or *System Synchronization Module* (SSM) documentation for the exact latency values. The following values are reference values for the product-specific latencies:

Variable	Reference Value (Clock Cycles)
$T_{tx_latency}$	4
$T_{rx_latency}$	3
$T_{SSM_latency}$	4

In addition to the instrument/SSM specific latency the calculations must account for the link latency that depends on the link characteristic (PXIe backplane or System Sync/Link cable length) and receiving instrument implementation. See your instrument or SSM documentation for the exact latency values. The following is a reference value for the link latencies:

Variable	Reference Value (Clock Cycles)
$T_{link_latency}$	12

Example Scenarios

As described above, a Sync FPGA Data-Sharing Statement can contain multiple transactions. A transaction can go to one or multiple destinations, that is, deliver the same data to multiple Rx endpoints, but independently of the number of destinations, a transaction has only one transmission point and operation. PathWave Test Sync Executive optimizes the Sync FPGA Data-Sharing Statement timing by parallelizing as much as possible the execution of different transactions. Depending on the number of transactions, the system topology and the overlap in terms of sources and destinations of the different transactions, the level of parallelization can vary. In the following examples we show some typical use cases.

Example 1: Single transaction with multiple destinations

The simplest use case for Sync FPGA Data-Sharing Statement consists of a single data sharing transaction. A transaction can go to one or more destinations, that is, deliver data to multiple Rx endpoints, but independently of the number of destinations, a transaction has only 1 Tx endpoint and operation. The execution time for a single-transaction Sync FPGA Data-Sharing Statement is given by the time required to complete the final Rx operation, where *tr1* is the transaction, *rx1*, *rx2* ...*rxN* are the Rx endpoints:

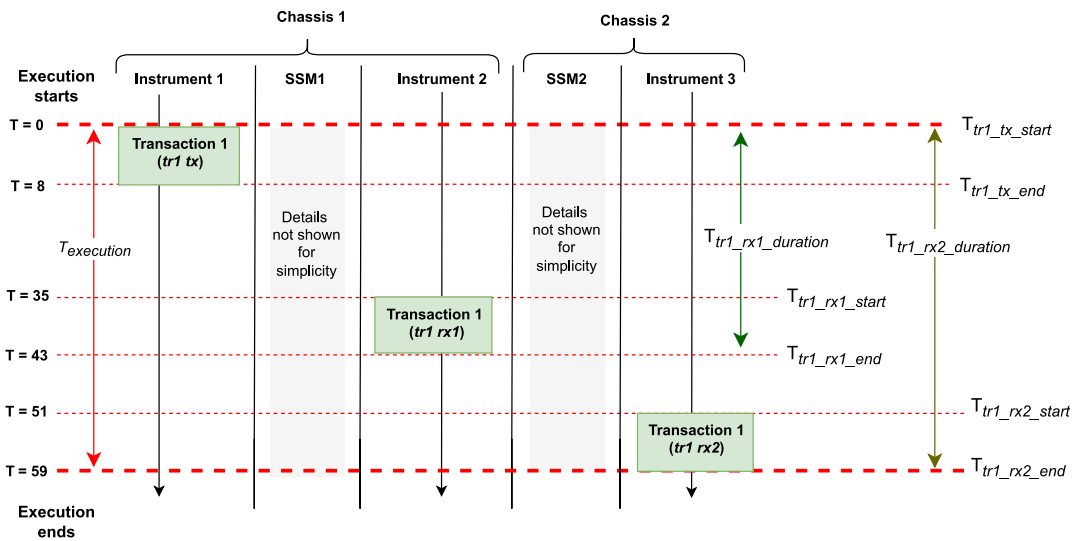
$$T_{\text{execution}} = \max(T_{tr1_rx1_end}, T_{tr1_rx2_end}, \dots, T_{tr1_rxn_end})$$

The following example shows how to calculate the Sync FPGA Data-Sharing Statement execution time for a single-transaction (*tr1*) with 2 destinations (*rx1* and *rx2*):

The following code snippet shows a Sync FPGA Data-Sharing Statement sharing to two destinations:

```
# SyncFpgaDataSharing definition with a single transaction to 2 destinations (Rx)
#
## Transaction 1 (tr1)
instrument1_tx = keysight_tse.FdsPortAddress(source_port, source_address)
instrument2_rx1 = keysight_tse.FdsPortAddress(dst1_port, dst1_address)
instrument3_rx2 = keysight_tse.FdsPortAddress(dst2_port, dst2_address)
fpga_data_sharing_st.transactions.add(instrument1_tx, [instrument2_rx1 , instrument3_rx2], num_bits_to_share)
```

The following diagram shows the timing (execution starts at end of cycle 0).



tr1_tx refers to the transmission of the data from the transmission port.

tr1_rx1 and *tr1_rx2* refer to reception of the data at the two receive ports.

tr1_rx1_duration is the total time from the beginning of transmission tr1_tx to the end of reception of the data at receive point 1 (tr1_rx1_end).

tr1_rx2_duration is the total time from the beginning of transmission tr1_tx to the end of reception of the data at receive point 2 (tr1_rx2_end).

Single Tx operation start:

$$T_{tr1_tx_start} = T_{execution_start} = 0 \text{ cycles}$$

$$T_{tr1_tx_end} = T_{tr1_tx_start} + N_{num_bits}/4 = 32/4 = 8 \text{ cycles}$$

Timing for tr1_rx1:

$$T_{tr1_rx1_duration} = T_{tx_latency} + (N_{SSM}+1)*T_{link_latency} + N_{SSM}*T_{ssm_latency} + T_{rx_latency} + N_{num_bits}/4 = 4 + (1+1)*12+1*4+3+32/4 = 43 \text{ cycles}$$

$$T_{tr1_rx1_end} = T_{tr1_tx_start} + T_{tr1_rx1_duration} = 0 + 43 = 43 \text{ cycles}$$

Timing for tr1_rx2:

$$T_{tr1_rx2_duration} = T_{tx_latency} + (N_{SSM}+1)*T_{link_latency} + N_{SSM}*T_{ssm_latency} + T_{rx_latency} + N_{num_bits}/4 = 4 + (2+1)*12+2*4+3+32/4 = 59 \text{ cycles}$$

$$T_{tr1_rx2_end} = T_{tr1_tx_start} + T_{tr1_rx2_duration} = 0 + 59 = 59 \text{ cycles}$$

The execution time of the Sync FPGA Data-Sharing Statement is:

$$T_{execution} = \max(T_{tr1_rx1_end}, T_{tr1_rx2_end}) = \max(43, 59) = 59 \text{ cycles}$$

NOTE

Note that if only one destination is used, for instance only Instrument 2, then

$$T_{execution} = T_{tr1_rx1_end} = 43 \text{ cycles}$$

Example 2: Multiple simultaneous transactions

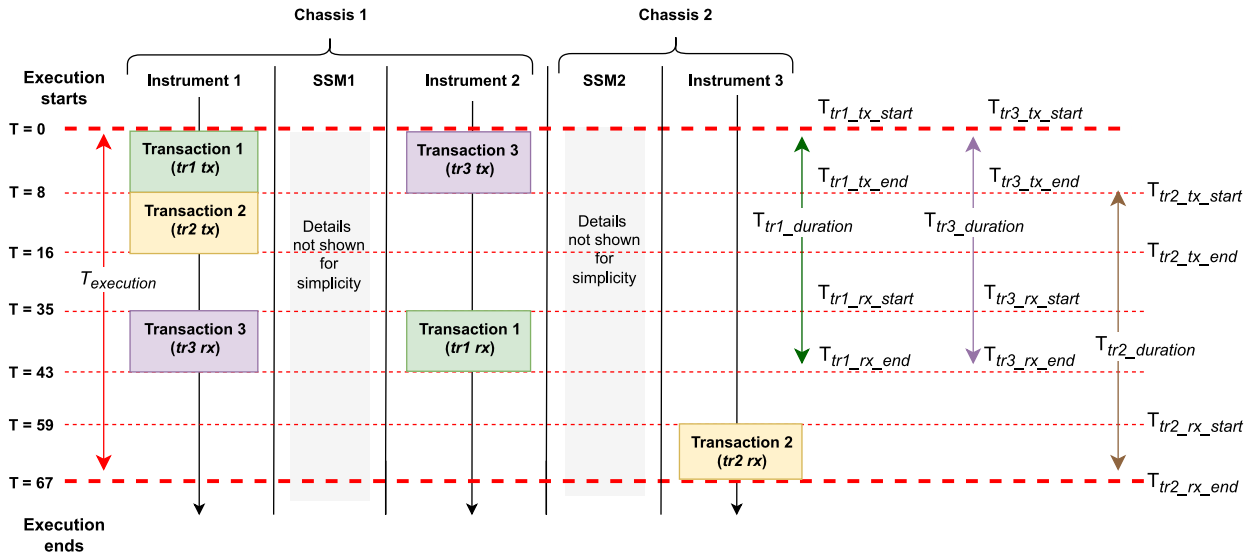
When sending data from the same transmitter, the transactions are executed in series. This means that the next transaction starts transmission as soon as the last nibble of the previous transaction is transmitted. Therefore, to estimate the execution time of the Sync FPGA Data-Sharing Statement, you must delay the individual transactions accordingly, then compare their end times and you pick the highest one.

In this example there are two different transactions originating from the same Tx port in Instrument 1. In the first transaction (tr1), 32 bits is sent from Instrument1 to instrument2 and in the second transaction (tr2), a different data packet of 32 bits is sent from Instrument1 to Instrument3. At the same time as the first transaction, Instrument2 does a separate transaction3 (tr3), sending 32 bits of data to Instrument1. Unlike example 1, these are all different transactions that send different data packets.

The following code snippet shows the transactions:

```
# SyncFpgaDataSharing definition with 3 transactions to 3 destinations
#
# Sources
instrument1_tx = keysight_tse.FdsPortAddress(source1_port, source1_address)
instrument2_tx = keysight_tse.FdsPortAddress(source2_port, source2_address)
# Destinations
instrument1_rx = keysight_tse.FdsPortAddress(dst1_port, dst1_address)
instrument2_rx = keysight_tse.FdsPortAddress(dst2_port, dst2_address)
instrument3_rx = keysight_tse.FdsPortAddress(dst3_port, dst3_address)
#
# Transaction 1
fpga_data_sharing_st.transactions.add(instrument1_tx, [instrument2_rx], num_bits_to_share)
#
# Transaction 2
fpga_data_sharing_st.transactions.add(instrument1_tx, [instrument3_rx], num_bits_to_share)
#
# Transaction 3
fpga_data_sharing_st.transactions.add(instrument2_tx, [instrument1_rx], num_bits_to_share)
```

The following diagram shows the execution time of the Sync FPGA Data-Sharing Statement as well as the timings of each individual transaction:



tr1_tx, tr2_tx, and tr3_tx, refers to the transmission of the data from the transmit ports.

tr1_rx, tr2_rx, and tr3_rx, refer to reception of the data at the receive ports.

tr1_duration is the total time from the beginning of transmission (tr1_tx) to the end of reception of the data at receive point 1 (tr1_rx_end).

tr2_duration and tr3_duration are the total times for tr2 and tr3 respectively.

The timing for Transaction1 (tr1) is the same as tr1_rx1 in example 1:

$$T_{tr1_end} = T_{tr1_start} + T_{tr1_duration} = 0 + 43 = 43 \text{ cycles}$$

For Transaction2 (tr2), Instrument1 sends data to Instrument 3.

tr2 can start the data transmission as soon as the transmission of Transaction1 has ended. The duration for t2 is 59 cycles, the same as tr1_rx2 in example 1.

Therefore, the timing for Transaction2 is:

$$T_{tr2_start} = T_{tr2_tx_start} = T_{tr1_tx_end} = 8 \text{ cycles}$$

$$T_{tr2_end} = T_{tr2_start} + T_{tr2_duration} = 8 + 59 = 67 \text{ cycles}$$

Transaction3 (tr3) sends data from Instrument 2 to instrument 1. This is sent and received at the same time as tr1, the duration is the same as tr1_rx1 in example 1.

$$T_{tr3_end} = T_{tr3_start} + T_{tr3_duration} = 0 + 43 = 43 \text{ cycles}$$


The execution time of the Sync FPGA Data-Sharing Statement is:

$$T_{execution_end} = \max(T_{tr1_end}, T_{tr2_end}, T_{tr3_end}) = \max(43, 67, 43) = 67 \text{ cycles}$$

Example 3: Optimizing the timing of multiple simultaneous transactions by reordering transactions

Changing the order of the transactions can affect the execution time of the Statement. If example 2 is modified so tr2 is sent before tr1, some time is saved.

```
# SyncFpgaDataSharing definition with a 3 transactions to 3 destinations
# transaction 1 and 2 are reversed
#
# Sources
instrument1_tx = keysight_tse.FdsPortAddress(source1_port, source1_address)
instrument2_tx = keysight_tse.FdsPortAddress(source2_port, source2_address)
# Destinations
instrument1_rx = keysight_tse.FdsPortAddress(dst1_port, dst1_address)
instrument2_rx = keysight_tse.FdsPortAddress(dst2_port, dst2_address)
instrument3_rx = keysight_tse.FdsPortAddress(dst3_port, dst3_address)
#
# Transaction 2
fpga_data_sharing_st.transactions.add(instrument1_tx, [instrument3_rx], num_bits_to_share)
#
# Transaction 1
fpga_data_sharing_st.transactions.add(instrument1_tx, [instrument2_rx], num_bits_to_share)
#
# Transaction 3
fpga_data_sharing_st.transactions.add(instrument2_tx, [instrument1_rx], num_bits_to_share)
```



tr1_tx, tr2_tx, and tr3_tx, refers to the transmission of the data from the transmit ports.

tr1_rx, tr2_rx, and tr3_rx, refer to reception of the data at the receive ports.

tr1_duration is the total time from the beginning of transmission (tr1_tx) to the end of reception of the data at receive point 1 (tr1_rx_end).

tr2_duration and tr3_duration are the total times for tr2 and tr3 respectively.

In this case Transaction2 starts first, The duration for t2 is 59 cycles, same as tr1_rx2 in example 1:

$$T_{\text{transaction2_start}} = T_{\text{transaction2_transmission_start}} = T_{\text{execution_start}} = 0 \text{ cycles}$$

This time transaction2 starts at 0:

$$T_{\text{tr2_tx_end}} = N_{\text{num_bits}}/4 = 8 \text{ cycles}$$

$$T_{\text{tr2_end}} = T_{\text{tr2_start}} + T_{\text{tr2_duration}} = 0 + 59 = 59 \text{ cycles}$$

Transaction1 starts as soon as Transaction2 finishes transmission. The duration for tr1 remains at 43 cycles (same duration as tr1_rx1 in example 1), however it starts 8 cycles later and the timing becomes:

$$T_{\text{tr1_start}} = T_{\text{tr1_tx_start}} = T_{\text{tr2_tx_end}} = 8 \text{ cycles}$$

$$T_{\text{tr1_end}} = T_{\text{tr1_start}} + T_{\text{tr1_duration}} = 8 + 43 = 51 \text{ cycles}$$

Transaction 3 (tr3) is the same as tr3 in example 2. The duration is 43 cycles, same as tr1_rx1 in example 1. The timing is:

$$T_{\text{tr3_end}} = T_{\text{tr3_start}} + T_{\text{tr3_duration}} = 0 + 43 = 43 \text{ cycles}$$

The execution time of the Sync FPGA Data-Sharing Statement is:

$$T_{\text{execution_end}} = \max(T_{\text{tr2_end}}, T_{\text{tr1_end}}, T_{\text{tr3_end}}) = \max(59, 51, 43) = 59 \text{ cycles}$$

In this case, by starting Transaction2 first, you can save 8 cycles compared to example 2.

$$T_{\text{difference}} = T_{\text{example3_end}} - T_{\text{example2_end}} = 67 - 59 = 8 \text{ cycles}$$

Example 4: Multiple simultaneous transactions with a resource conflict

Transactions are sent across different types of link between instruments. These can be DSTARB/C links between SSMs and instruments that are in the same chassis or point-to-point SystemSync links between two SSMs located on different chassis. The difference is that DSTARB/C links allow only a single data path, whereas the SystemSync connections have multiple data paths. The multiple data paths in SystemSync connections allow the SSMs to route multiple transactions simultaneously. On the contrary, DSTARB/C links can only send one transaction at a time. Different data packets cannot use the same link simultaneously and must be queued instead. Therefore, DSTARB/C links are much more prone to **transaction conflicts**.

In principle, when all the transactions being sent are from different Tx ports, HVI tries to execute them in parallel. If this is possible, the execution time of the Sync FPGA Data-Sharing Statement is just the maximum of the duration of the individual transactions. However, this is not always possible because in some case **there is a conflict** if two or more transactions try to use the same path at the same time, for example, if they both arrive at the same Rx port at the same time. In these cases, the HVI delays some of the transactions to avoid the conflict, where any delay added is kept as small as possible.

The following example is similar to example 3, except for transaction 3 which in this case, goes from instrument 2 to instrument 3.

Transaction 2 (tr2) and transaction 3(tr3) both start at the same time and send data to SSM1. This is possible because both transactions are sent to different ports on SSM1.

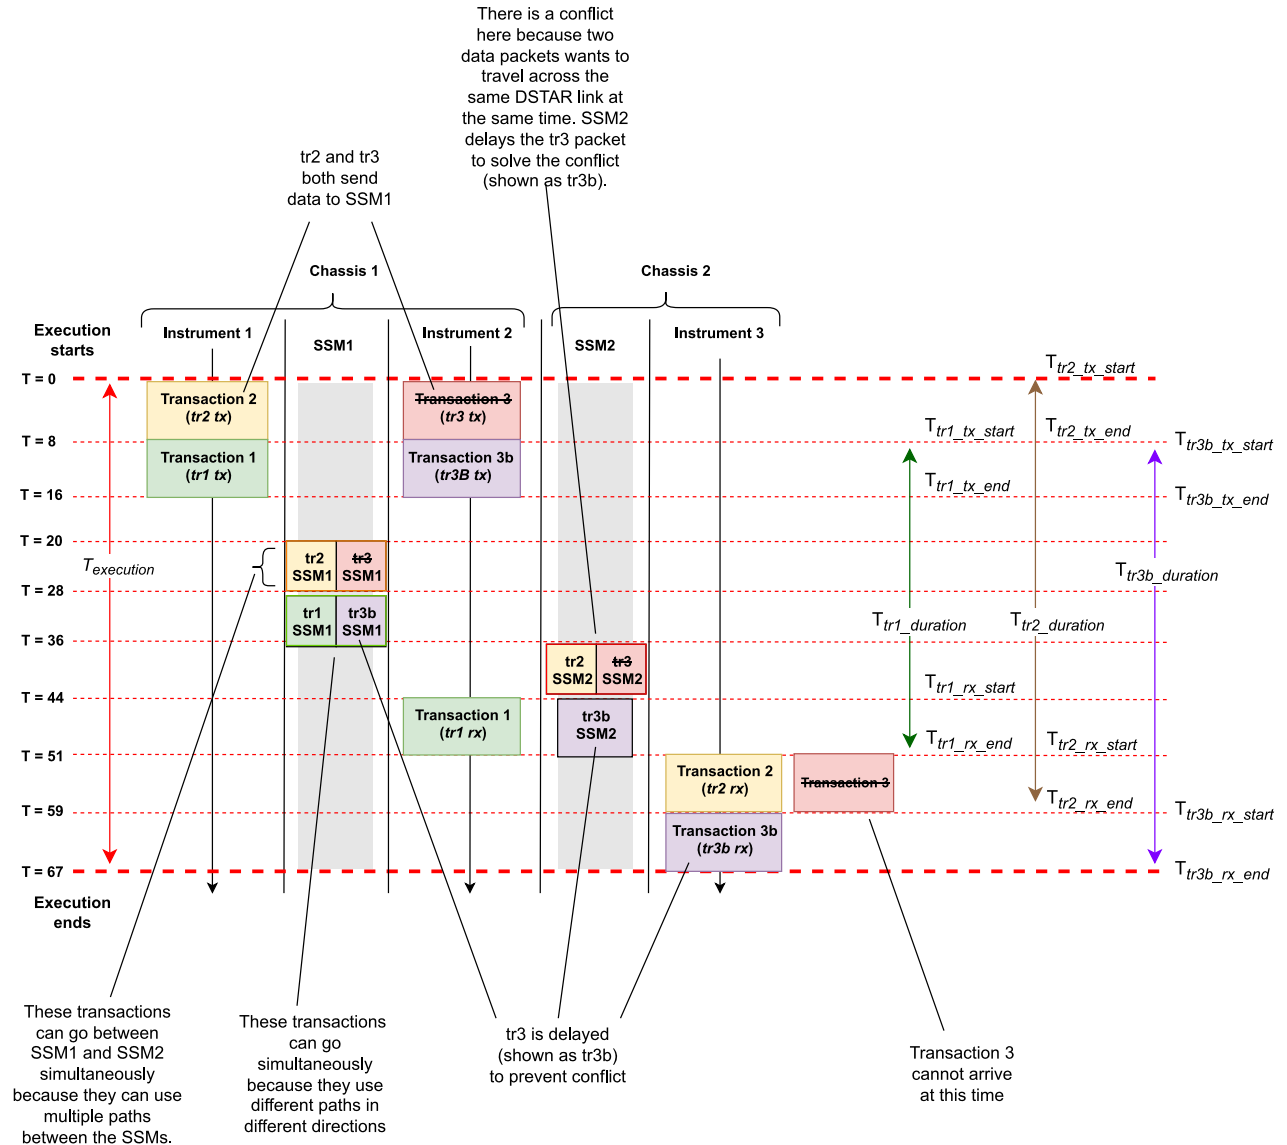
SSM1 then sends this data to SSM2, the SSMs have multiple paths for data between them so both transactions can be sent at the same time.

SSM2 then sends the data for tr2 and tr3 to instrument 3, but there is a **conflict** at this point because the data from both transactions cannot be sent at the same time to the same DSTARB port. DSTARB/C links have only one data path and the packets cannot cross it simultaneously, as explained before. To resolve this conflict, the HVI delays transaction3 at compilation time. This is shown as Transaction 3b (tr3b) in the diagram. The delay enables the transaction to be sent to instrument 3 with no conflicts.

```
# SyncFpgaDataSharing definition with a 3 transactions to 2 destinations
# transaction 1 and 2 are reversed
# transaction 3 goes to instrument 3
#
# Sources
instrument1_tx = keysight_tse.FdsPortAddress(source1_port, source1_address)
instrument2_tx = keysight_tse.FdsPortAddress(source2_port, source2_address)
# Destinations
instrument1_rx = keysight_tse.FdsPortAddress(dst1_port, dst1_address)
instrument2_rx = keysight_tse.FdsPortAddress(dst2_port, dst2_address)
```

```

instrument3_rx = keysight_tse.FdsPortAddress(dst3_port, dst3_address)
#
# Transaction 2
fpga_data_sharing_st.transactions.add(instrument1_tx, [instrument3_rx], num_bits_to_share)
#
# Transaction 1
fpga_data_sharing_st.transactions.add(instrument1_tx, [instrument2_rx], num_bits_to_share)
#
# Transaction 3 - This will be delayed
fpga_data_sharing_st.transactions.add(instrument2_tx, [instrument3_rx], num_bits_to_share)
    
```



tr1_tx, tr2_tx, tr3_tx, and tr3b_tx, refers to the transmission of the data from the transmit ports.

tr1_rx, tr2_rx, and tr3b_rx, refer to reception of the data at the receive ports.

tr1_duration is the total time from the beginning of transmission (tr1_tx) to the end of reception of the data at receive point 1 (tr1_rx_end).

tr2_duration and tr3b_duration are the total times for tr2 and tr3b respectively.

SSM1 and SSM2 indicate the System Synchronization Modules. In the boxes, these indicate when data is transmitted from the SSM.

In this example, first 32 bits of data from Instrument1 to Instrument3 (transaction2, tr2), and then 32 bits of data from Instrument2 to Instrument3 (transaction3, tr3). By checking the **topology in the system diagram**, you can see that both transactions pass through SSM1 and then SSM2 before they reach Instrument3.

To see if there is going to be a collision, calculate the time when each transaction is at the exit of SSM2, if it started at time 0.

For Transaction2, this is:

$$T_{tr2_at_SSM2_start} = T_{tx_instrument_latency} + T_{link_latency} + T_{ssm_latency} + T_{link_latency} + T_{ssm_latency} = 4 + 12 + 4 + 12 + 4 = 36 \text{ cycles}$$

$$T_{tr2_at_SSM2_end} = T_{tr2_at_SSM1_start} + N_{num_bits}/4 = 36 + 32/4 = 44 \text{ cycles}$$

For Transaction3 the numbers are the same, so there is an overlap from clock cycle 36 to 44. To resolve this, you must delay the second transaction, tr3, to start at SSM2 when tr2 ends. That is, you must delay tr3 by:

$$T_{tr3_delay} = T_{tr2_at_SSM2_end} - T_{tr3_at_SSM2_start} = 44 - 36 = 8 \text{ cycles}$$

In the diagram the delayed transaction3 is shown as transaction3b (tr3b). Given this delay, you can now calculate the timing for both transactions:

For Transaction2, you don't need to change anything, so the timing is the same as in example 3:

$$T_{tr2_end} = T_{tr2_start} + T_{tr2_duration} = 0 + 59 = 59 \text{ cycles}$$

As soon as Transaction2 finishes transmission, Transaction3b starts.

Transaction3b starts 8 cycles after tr2 and sends its data to SSM1 and then onto SSM2. SSM2 can then pass the data Instrument3. The duration remains at 59 cycles, so the timing is:

$$T_{tr3b_start} = T_{tr2_start} + 8 = 0 + 8 = 8 \text{ cycles}$$

$$T_{tr3b_end} = T_{tr3b_start} + T_{tr3b_duration} = 8 + 59 = 67 \text{ cycles}$$

Transaction1 (tr1) goes through SS1, at time=8 which is also when tr3b goes through SSM1.

However, these transactions go in different direction and use different ports so there is no conflict, therefore tr1 and tr3b can go ahead.

For Transaction1, the timing is the same as in example 3:

$$T_{tr1_end} = T_{tr1_start} + T_{tr1_duration} = 8 + 43 = 51 \text{ cycles}$$

The execution time of the Sync FPGA Data-Sharing Statement is:

$$T_{execution_end} = \max(T_{tr1_end}, T_{tr2_end}, T_{tr3b_end}) = \max(51, 59, 67) = 67 \text{ cycles}$$

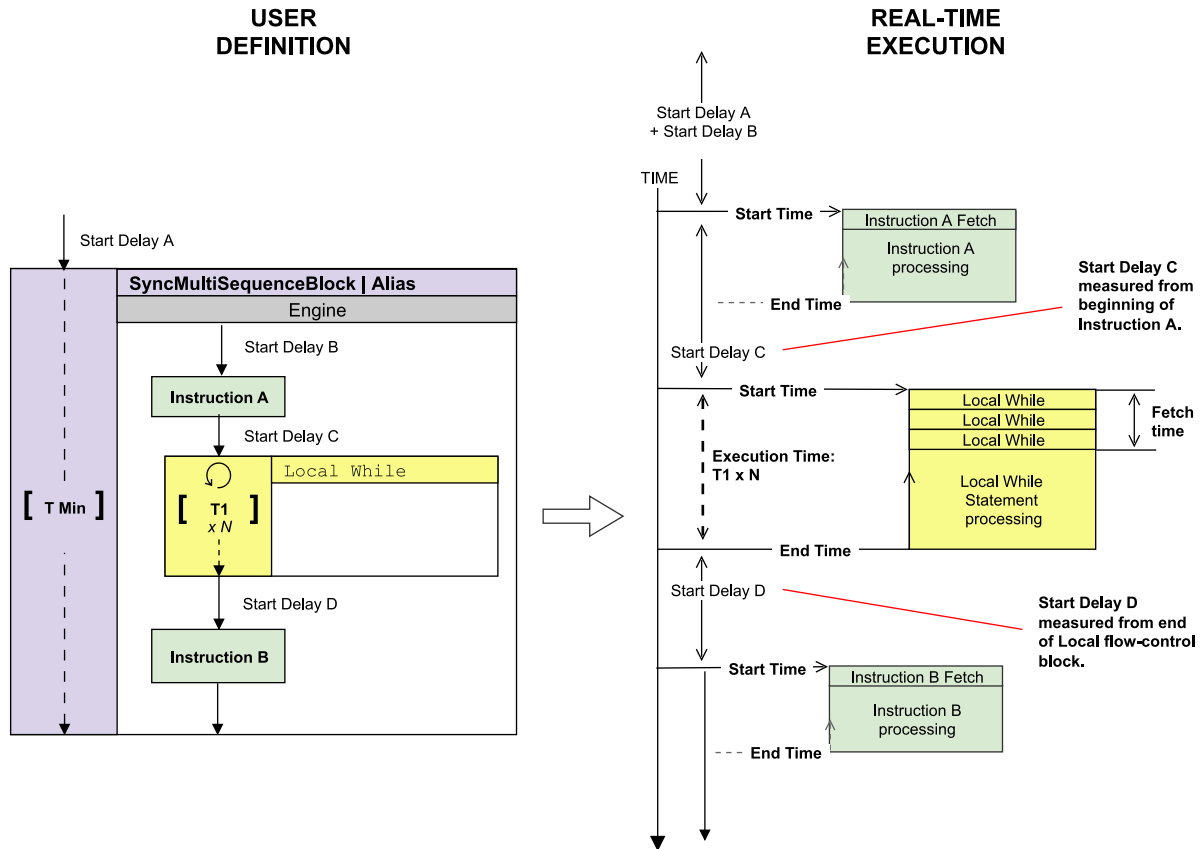
Local Flow-Control Statement Timing

This section describes Local Flow-Control Statement timing, it contains the following sections:

- **Local While**
 - **Local While Statement Timing Tables**
- **Local If**
 - **Local If with matched branches**
 - **Local If Statement Timing Tables**
- **Local Wait (for Event or time in register)**
 - **Local Wait-For-Time Statement Timing Tables**
 - **Local Wait-For-Event Statement Timing Tables**
- **Local Delay Statement**
 - **Local Delay Statement Timing Tables**

Local Flow-Control Statements and Sync Statements consume HVI Engine execution time and do not overlap their execution. When you are calculating the timing of a Sequence, you must consider the execution time of these Statements.

The following diagram shows the timing for a Sync Multi-Sequence Block that contains a pair of HVI instructions and a Local While:



Local While

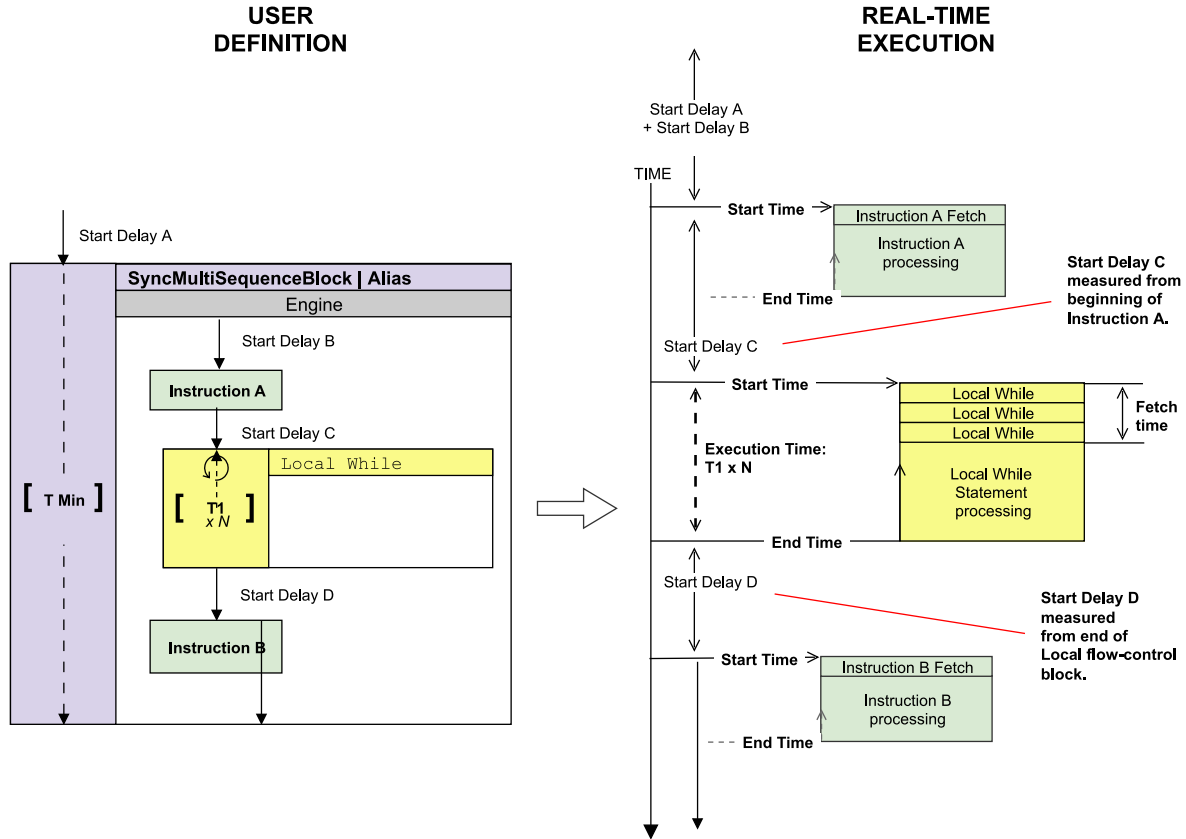
The Local While Statement continues execution while a condition is met and finishes the execution when the condition is no longer met. This has the same timing as Sync while Statements.

The following diagram shows a Local While Statement with other instructions.

The total execution time for a Local While is $T1 \times N$, where $T1$ is the iteration time and N is the number of times it iterates. The time cannot be indicated exactly on a diagram or in code because the number of iterations is not known until runtime.

For Statements coming after a Local While Statement, the Start Delay is measured from the end of the Local While Statement. In the following diagram, Start Delay D is measured from the end of the Local While Statement.

The dotted line indicates that the execution time of the Local While block $T1$ is not known at compile time.



Local While Statement Timing Tables

Local While timing value:

Execution time (cycles) ⁽¹⁾
$\#Iterations * [\sum_{for_all_internal_statements}(Start-Delay) + \sum_{for_all_internal_flow_control_statements}(Duration)]$

⁽¹⁾ This value applies if duration property of the Statement is set to Minimum (default). If a fixed-duration has been set, then this is the Execution time is equal to that value.

Local While latency values:

Parameter	Description		Time (cycles)
Start-Latency	Minimum start-delay for the Statement		$5 + \#Register_Conditions$
Entry/Iteration latency	Minimum start-delay for first Statement inside the while loop	Minimum Duration	$8 + \#Register_Conditions + End-Latency_{Last-statement}$
		Fixed Duration	$8 + \#Register_Conditions$
End-Latency	Minimum start-delay for the next instruction outside the while loop	Minimum Duration	$8 + \#Register_Conditions + End-Latency_{Last-statement}$
		Fixed Duration	$8 + \#Register_Conditions$
Fixed-Duration	Minimum fixed-duration for Statement		$[\sum_{for_all_internal_statements}(Start-Delay) + \sum_{for_all_internal_flow_control_statements}(Duration) + End-Latency_{Last-statement}]^{(1)}$
Register Evaluation Latency	Time to evaluate the register condition		<ul style="list-style-type: none"> From start: 3 For each iteration: $-(2 + \#Register_Conditions)$

(1) If the branch is empty, then the duration is equal to the Entry-Latency of the branch.

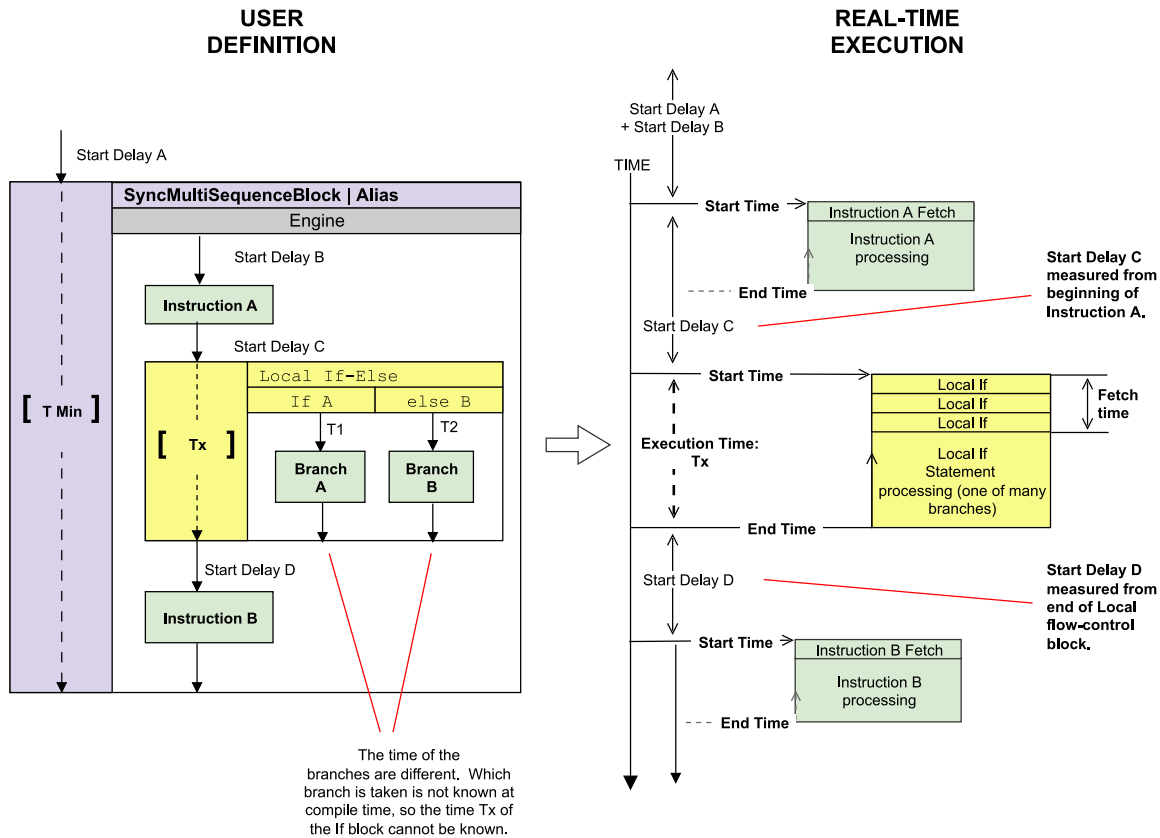
Local If

For Local If Statements (if-elseif-else), the following Start Delay is measured from the end of the Local If Statement. The time taken is only known at runtime, so it is not possible to indicate them on a diagram or in code. This is the same as while Statements.

This following diagram shows the timing of Local If Statements. The Start Delay D is measured from the end of the Local If Statement.

The Local If has two branching options with times T1 and T2. These times can be different. Since the choice of branch is not known at compile time, the time for the Local If block cannot be known.

The line for the Local If block is dotted. This indicates that the execution time of the Local If block Tx is unknown. The time of the containing block is also therefore unknown, and it is also dotted. The time of the Sync Multi-Sequence Block is indicated as T min.

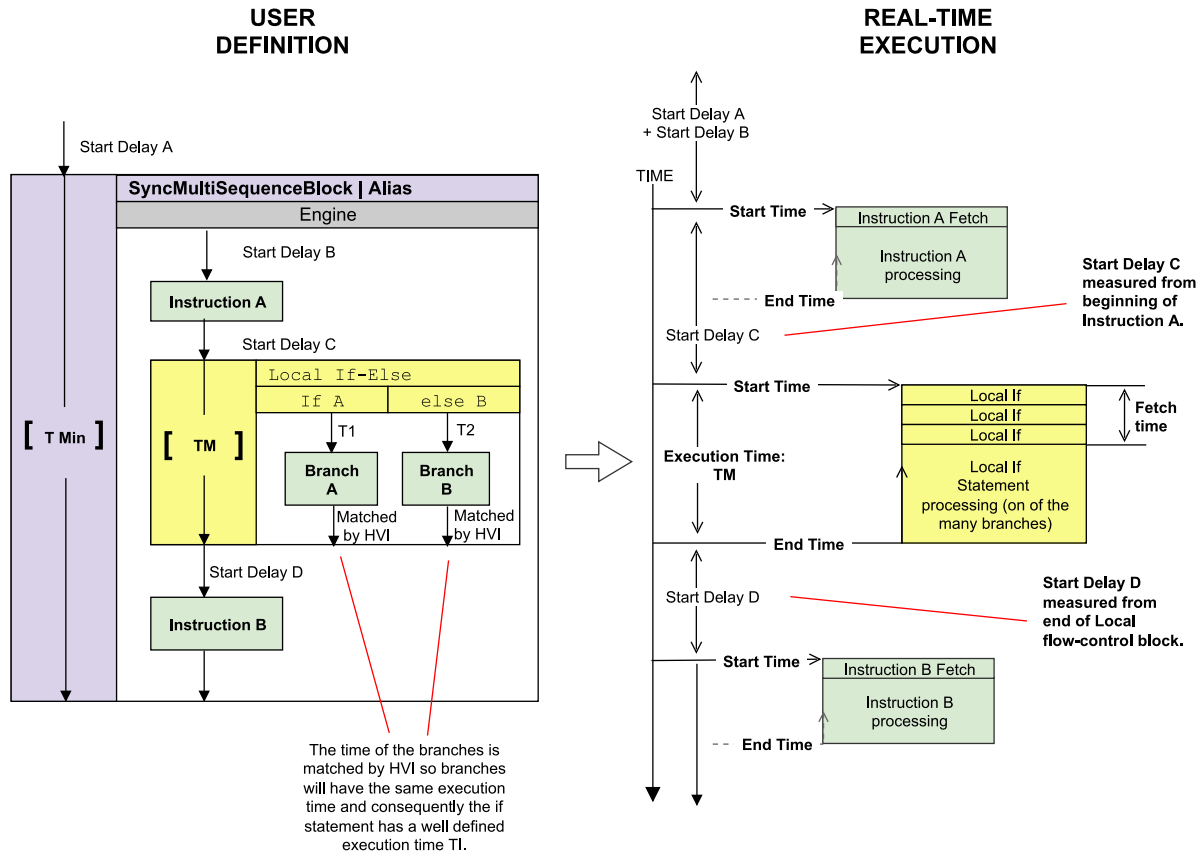


Local If with matched branches

In the following diagram the branches in the If and else branches are matched. This ensures the Local If ends at the same time irrespective of the branch taken.

The total branch time is marked with the time TM, this represents the matched time. The choice of branch is not known at compile time, but since the times are matched the time TM is known.

The times are known at compile time so the timelines of the Local If block and the Sync Multi-Sequence Block that contains it are both solid.



Local If Statement Timing Tables

For if Statements with multiple If / Else-If / Else branches, the Entry delays are the same for all branches.

If the match-branches attribute is enabled, the HVI ensures that the execution of all of the branches have the same overall delay. If match-branches is not enabled, some branches might take less time than others.

The If Statement latency depends on the number or register-conditions used: `#Register_Conditions`.

Local If timing value:

Execution time (cycles) ⁽¹⁾ ⁽²⁾
<code>sum_for_all_internal_statements(Start-Delay) + sum_for_all_internal_flow_control_statements(Duration) ⁽³⁾</code>

⁽¹⁾ The value provided here applies if the duration property of the Statement is set to Minimum (default). If a fixed-duration has been set, then the Execution time is equal to that value.

(2) This value is only calculated for the branch that is executed, if there are multiple branches available.

(3) If the branch is empty, the execution time becomes $\text{Entry-Latency}_{\text{branch}} - 1$.

The following table shows Local If latency values:

Parameter	Description		Minimum time (cycles)
Start-Latency	Contributes to the minimum-possible start-delay for the Statement		$5 + \#Register_Conditions_IfBranch$
Entry-latency	Contributes to the minimum-possible start-delay for first Statement in branch #		<ul style="list-style-type: none"> • If-Branch: 3 • Else-If-Branch: <ul style="list-style-type: none"> - For each else-if branch, we need to add: <ul style="list-style-type: none"> ◦ $6 + \#Register_Conditions_Else_If_Branch$ ◦ For the 1st else-if branch we will have: <ul style="list-style-type: none"> ◦ $2 + \#Register_Conditions_IfBranch + 7 + \#Register_Conditions_Else_If_Branch_1$ ◦ For the 2nd else-if branch we will have: <ul style="list-style-type: none"> ◦ $2 + \#Register_Conditions_IfBranch + 7 + \#Register_Conditions_Else_If_Branch_1 + 7 + \#Register_Conditions_Else_If_Branch_2$ ◦ and so on, so forth... • Else-Branch: Equal to last Else-If-Branch value
End-Latency	Contributes to the minimum-possible start-delay of the next Statement outside the if Statement	Matching Branches <i>disabled</i>	$3 + \max_{for_all_Branches}[End-Latency_{Last-statement}]$ (1)
		Matching Branches <i>enabled</i>	$3 + End-Latency_{Last-statement-of-longest-branch}$ (2) Where longest branch means the branch with longer execution time.
	Fixed-Duration	1	

<p>Fixed-Duration</p>	<p>Minimum fixed-duration for Statement</p>	$2 + \max_{\text{for_all_branches}}[\text{Branch-Duration}] \text{ (3)}$ <p>Where Branch-Duration is calculated as follows:</p> $[\text{sum}_{\text{for_all_internal_statements}}(\text{Start-Delay}) + \text{sum}_{\text{for_all_internal_flow_control_statements}}(\text{Duration}) + \text{End-Latency}_{\text{Last-statement}}] \text{ (4)}$
<p>Register Evaluation Latency</p>	<p>Time to evaluate the register condition</p>	<p>3</p> <p><i>Then, for registers used in the condition of any else-if branch, we need to subtract:</i></p> <ul style="list-style-type: none"> - $6 + \#Register_Conditions_Else-If-Branch$ <p><i>Therefore, for the 1st else-if branch we will have:</i></p> <ul style="list-style-type: none"> - $3 - (6 + \#Register_Conditions_Else-If-Branch_1)$ <p><i>For the 2nd else-if branch we will have:</i></p> <ul style="list-style-type: none"> - $3 - (6 + \#Register_Conditions_Else-If-Branch_1) - (6 + \#Register_Conditions_Else-If-Branch_2)$ <p><i>and so on, so forth...</i></p>

(1) If the maximum end latency used in this equation corresponds to the if-branch, and the calculated latency is greater than 4, then the **End-latency** is the calculated value minus 1.

(2) If the longest branch is the if-branch, then the **End-latency** is the calculated value minus 1.

(3) If the maximum branch duration used in the equation corresponds to the if-branch, then the duration is the calculated value minus 1.

(4) If a branch is empty, then the branch duration is equal to the Entry-latency of the branch.


Local Wait (for Event or time in register)

For Local Wait Statements, the following Start Delay is measured from the end of the Local Wait Statement. As with Sync while Statements, the time taken is only known at runtime, so it is not possible to indicate them on a diagram or in code.

The following diagram shows the timing of a Local Wait Statement. The following Start Delay D is measured from the end of the Local Wait Statement.

The execution time of the Local Wait Statement T1 is not known at compile time, this is indicated by the dotted line.

The time of the Sync Multi-Sequence Block is indicated as T min. The dotted line indicates an unknown time.



Local Wait-For-Time Statement Timing Tables

A Wait-for-time Statement blocks HVI execution in a Local Sequence until a specific amount of time passes. This amount of time is defined in a register that is specified as an argument in the Wait-for-time Statement. The value of the register specifies the number of cycles to wait.

Local Wait-for-time Statement timing value:

Execution time (cycles)
<i>RegisterValue</i>

Local Wait-for-time Statement latency values:

Parameter	Time (cycles)
Start-Latency	1
End-Latency	1
Register Evaluation Latency	1

Local Wait-For-Event Statement Timing Tables

A Local Wait-for-Event Statement blocks HVI execution in a Local Sequence until an Event occurs. Events sources can be the Trigger IOs, or internal to the instrument (including User FPGA-Sandbox Events).

Local Wait-for-Event Statement timing values:

Event type	Execution time (cycles)	Fetch time (cycles)
Internal Event	$\text{MAX}(\text{Event_Arrival_Time}(1) + \text{Instrument_Event_Latency}(2) + 1, \text{Fetch_Time}) + 1$	3
Trigger IO	$\text{MAX}(\text{Event_Arrival_Time}(1) + \text{Instrument_Event_Latency}(2) + \text{Instrument_Event_Condition_Latency}(3), \text{Fetch_Time}) + 1$	$1 + \text{Instrument_Event_Condition_Latency}$

(1) *Event_arrival_time* is:

- Internal Events
 - $\text{Event_Arrival_Time} = \text{Internal_Event_Generation_Time} - \text{WaitForEvent_Start_Time}$
- External Events
 - $\text{Event_Arrival_Time} = \text{Event_At_Module_Connector_Time} - \text{WaitForEvent_Start_Time}$
 - The Event time can be measured at the front panel or PXIe backplane connector depending on the Event.

(2) **Instrument_Event_Latency** is the delay from the Event source until the Event state is available inside the HVI Engine. Events sources can be the Trigger IOs, or internal to the product (including User FPGA-Sandbox Events). It is an instrument and Event specific value. Refer to the instrument documentation for more information.

(3) **Instrument_Event_Condition_Latency** is the time needed for the condition evaluation to be executed once the Event has settled inside the HVI Engine. It is an instrument specific value. Refer to the instrument documentation for more information.

NOTE


The *Event_Arrival_Time* can be a negative value if the Event enters the module before the Wait-For-Event instruction Start Time. A number of scenarios are shown in the diagrams below.

Local Wait-for-Event latency values:


Parameter	Time (cycles)
Start-Latency	0
End-Latency	1

The following diagrams shows scenarios where the execution time of a Wait-For-Event Statement can vary:

Case 1: Event arrival + propagation delay after Fetch-Time completion



Case 2: Event arrival + propagation delay completes before Fetch-Time completion




Local Delay Statement

The Local Delay Statement delays the execution of a Local Sequence for a time you specify. The default unit is nanoseconds but the delay is specified in any unit of seconds. The delay is fixed and cannot be changed during HVI execution, so the delay value must be known at the time of creating the HVI Sequence.

The delay Statement works in a similar way as the Start Delay Statement parameter, however the difference is that the Start Delay can only be specified before the other Statements in a Sequence. The delay Statement enables you to place a fixed delay at the end of Sync Multi-Sequence Block or a flow control Statement.

Unlike a wait-for-time Statement, the delay Statement does not introduce a de-synchronization and therefore does not trigger a resynchronization. This therefore avoids the timing overhead introduced by the Triggered re-synchronization point.

The following diagram shows the timing of a delay Statement **Delay Z**:



Local Delay Statement Timing Tables

A Delay Statement delays HVI execution in a Local Sequence until a specific amount of time passes. This amount of time is specified in a parameter in the Statement.

Local Delay Statement timing value:

Execution time (cycles)
Delay Specified

Local Delay latency values:

Parameter	Time (cycles)
Start-Latency	0
End-Latency	1

HVI Instruction Timing

This section describes HVI Instruction timing, it contains the following sections:

- [HVI Instruction Parameters](#)
- [Instruction position](#)
- [Overlapping instruction execution](#)
- [Trigger Write](#)
- [Action Execute](#)
- [Arithmetic Logic Unit Instructions](#)
- [FPGA-User Sandbox Instructions](#)
- [FPGA-Instruction Statement](#)
- [Instrument-Specific HVI Instruction Timing Values](#)
- [HVI Instruction Position Mapping](#)
- [Examples of HVI Instruction Timing Calculation across Sync and Local Flow-Control Statements](#)

The following sections list the fetch and execution latency for HVI Instructions. Unless stated otherwise, all times are in HVI Engine clock cycles. The HVI Engine clock frequency is instrument-specific. For information about the HVI Engine clock frequency and Instrument-Specific HVI Instruction latencies, See your instrument documentation.

HVI Instruction Parameters

HVI Instructions have a number of parameters and properties you must be aware of for calculating timing:

TriggerIO groups and Action groups

The following additional parameters are used for calculating timing for some HVI Instructions.

Triggers and Actions are organized into groups and the timing can change depending on these:

TriggerIO groups

Trigger Inputs / Outputs are organized together in groups of 16 called TriggerIO groups. Any number of TriggerIO groups can be written at the same time.

Action groups

HVI Actions are organized together in groups of up to 16 called Action groups. Any number of Action groups can be executed synchronously.

Instruction position

This section describes the concept of *instruction position*.

NOTE This is provided for your information, you are not typically required to program an HVI at this level of performance.

Instructions are broken into internal-instructions that the compiler maps onto the HVI Engine hardware. During one HVI Engine cycle, the HVI Engine can fetch, dispatch and execute multiple instructions in parallel.

Instructions can be scheduled for execution together, however, depending on the Instructions involved, this is not always possible because of the inner structure of the HVI Engine. To understand why, you must understand the concept of instruction position.

An HVI Engine is a processor with a set of execution pipelines, each of which has a numbered *position*. The individual internal-instructions are mapped across the different pipeline positions for execution.

For parallel instruction fetching to be possible, the internal-instructions must use different positions inside the instruction register of the HVI Engine. If two internal-instructions are using overlapping positions, then they cannot be fetched in parallel. The positions where each internal-instruction is to be mapped depends on the instruction. This means the hardware can only execute certain internal-instruction in specific positions. The internal-instructions are mapped by the compiler. This process is not user programmable.

A table with the per-instruction mapping is provided in the documentation for each instrument.

The following diagram provides an example table.

	Position										
Instruction	1	2	3	4	5	6	7	8	9	10	n (n > 10)
A					5 - 7						
B	1 - 4										
C								8 - 9			
D					5 - 7			8 - 10			
E	1 - 7										

From the table, you can see that:

- Instruction A can be mapped, one at a time, to positions 5 to 7.
- Instruction B can be mapped, one at a time, to positions 1 to 4.
- Instruction C can be mapped, one at a time, to positions 8 to 10.
- Instruction D can be mapped, two at a time, to positions 5 to 7, positions 8 to 10, or both.
- Instruction E can be mapped, one at a time, to positions 1 to 7.

At compile time, HVI maps the instructions to be executed to their respective supported positions. If an instruction cannot be mapped to its supported position because another instruction is already mapped there, HVI generates an error and informs the user. For example:


- If an instruction is A is followed by a second instruction A at the same time, HVI assigns the first instruction A to positions 5-7, but generates an error with the second instruction A because positions 5-7 are already used.
- If an instruction D is followed by another instruction D at the same time, HVI will assign the first instruction D to positions 5-7 and will then assign the second instruction D to positions 8-10. However, if there is a third instruction D to be fetched at the same time, HVI generates an error because neither possible position for D are available for the third instruction.
- If there are instructions A, B and C at the same time, HVI assigns them to positions 5-7, 1-4 and 8-10, respectively, without any issue.
- If there are instructions A, B and D at the same time, HVI assigns them to positions 5-7, 1-4 and 8-10, respectively, without any issue. If, however, the order was B, D and A, then HVI assigns B to positions 1-4, D to positions 5-7 and, then, HVI generates an error because positions 5-7 are not be available for instruction A.
- If there is an instruction E, then if it is fetched at the same time with any of the instructions A, B or E, then HVI generates an error. However, if it is fetched in parallel with C or D, then there will be no issue.

NOTE

When there is no fetching in parallel, An HVI Engine is capable of executing instructions in parallel irrespective of their instruction position.

Overlapping instruction execution

The following diagram shows Instruction B and Instruction E are executed in parallel, even though they are using the conflicting positions in the instruction register (positions 1-4 are overlapping as seen in the table earlier). This is possible, as long as the Start delay T3 for instruction E is such that its fetch cycle does not coincide with the fetch cycles of instruction B. The green dotted line indicates the minimum extent that T3 should have.



Overlapping instruction fetching

An HVI Engine is capable of fetching and executing multiple instructions in parallel, providing their instruction positions are not overlapping. Most instructions have only 1 fetch cycle, but it is possible for instructions to require multiple fetch cycles. Refer to the instructions timing tables for details on the fetch cycles of the different instructions.


The following figures show examples of instruction fetching in parallel. For the instruction positions that are being used by each instruction, the values are from the previously defined example table.

Example A. In this example it is assumed that:

- Start delay $T1 > 0$ cycles.
- $T2 = 0$ cycles.
- $T3 = 1$ cycle.
- $T4 > 3$ cycles.

At real-time execution, after the $T1$ delay has passed, Instruction A and Instruction B are fetched at the same time, since the Start delay $T2$ for instruction B is equal to 0. Then, after one cycle, that is, the Start delay $T3$, instruction C is fetched before the fetching of Instructions A and B is completed. Finally, after delay $T4$ from the beginning of instruction C, instruction D is fetched.

As shown in the diagram, instructions A and B are fetched in parallel for 2 Engine cycles and instructions A, B and C are fetched in parallel for 1 Engine cycle. Looking at the table, instructions A, B and C can be fetched in parallel as they are not using the same positions. Instruction D is fetched later, so there is no conflict in the available positions.

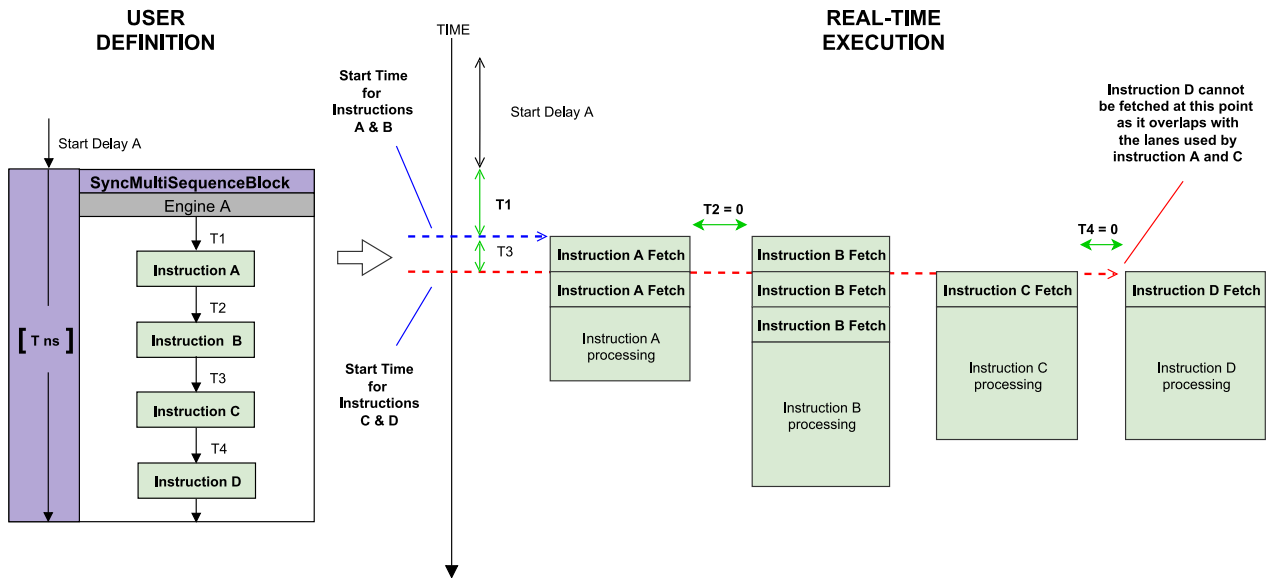


Example B. In this example start delay $T1 > 0$ cycles, $T2 = 0$ cycles, $T3 = 1$ cycle, and $T4 = 0$ cycles.

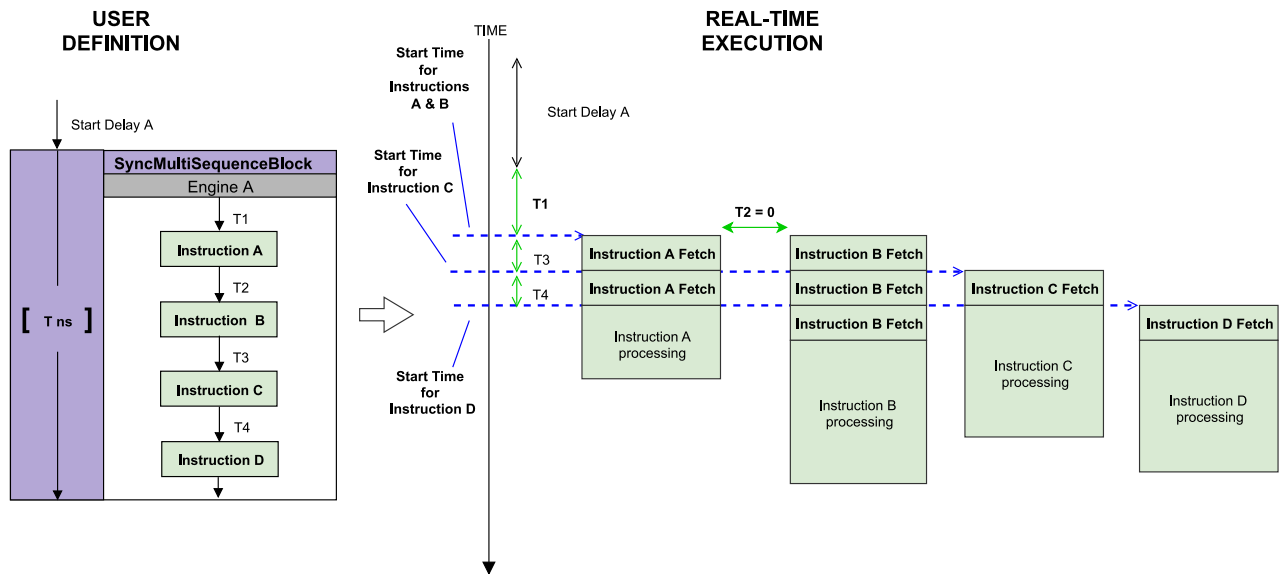
At real-time execution, after the $T1$ delay has passed, instructions A and Instruction B are fetched at the same time because the Start delay $T2$ for instruction B is equal to 0.

Then, after one cycle, that is the Start delay $T3$, instruction C is being fetched and at the same time ($T4 = 0$) instruction D is fetched.

Compared to the previous example, in this case, instruction D cannot be placed to either positions 5-7 (assigned to instruction A) or positions 8-10 (assigned to instruction C), so it is not possible to fetch instruction D at the same time. as A and C. **This example generates an error during the HVI compilation.**



One way to fix the issue is to increase the Start delay T4 of instruction D so that it is not fetched at the same time as instruction A and C. This can be done by increasing T4 by at least 1 cycle. This is shown in the following figure:

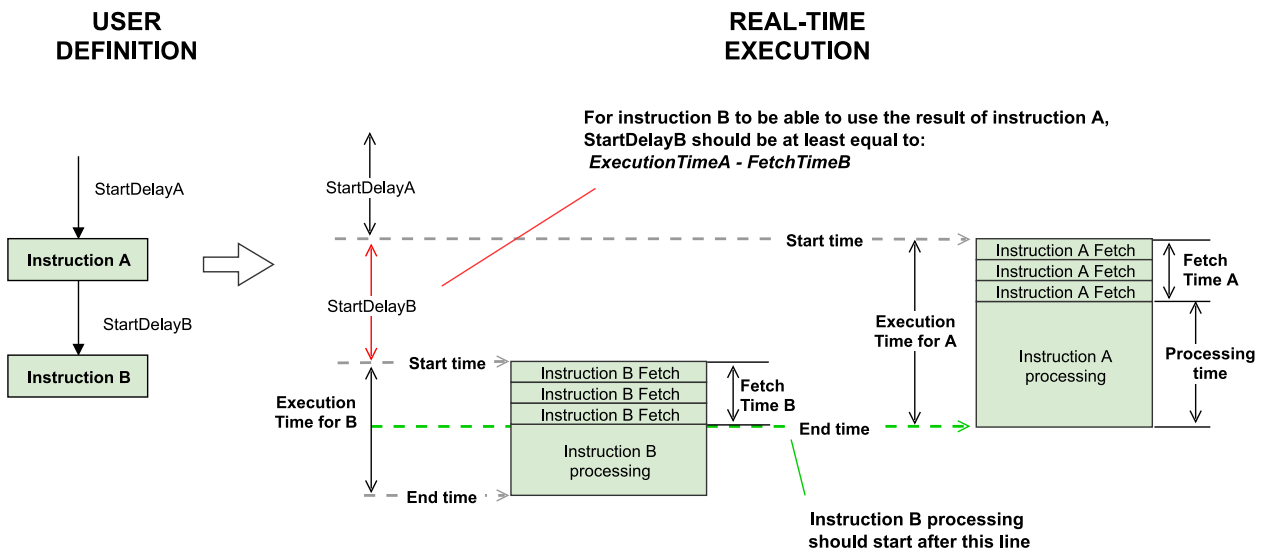


Overlapping instruction execution with result dependencies

HVI is capable of processing instructions in parallel. This is a powerful capability, but it can lead to unexpected results when there are dependencies between the instructions, that is, when one instruction depends on the result of the other. For example, an instruction might update the value of an HVI register and the following instruction might need to use that updated register value. To avoid unexpected results, the user needs to ensure that the delay between the independent and the dependent instructions is big enough so that the processing of the independent (Instr1) is completed before or when the processing of the dependent instruction (Instr2) start. The minimum delay to achieve this can be expressed with the following formula:

$$\text{MinDelay_Instr1_to_Instr2} = \text{Instr1_ExecutionTime} - \text{Instr2_FetchTime}$$

The following diagram shows an example with two HVI Instructions and the timing when executed by the HVI Engine. Assuming that instruction B is using the result of instruction A, you must ensure that the value of StartDelayB is greater or equal to the Processing Time of instruction A, minus the Fetch Time of instruction B. This way, the processing of instruction B will start after the end of processing of instruction A.



NOTE

It is important to consider the effects of overlapping instructions with dependencies, because HVI does not track dependencies. This is because in some cases it is desirable to implement pipelines of operations and exploit the fact that the next instruction uses the previous value of a register, before the previous operation is completed. It is your responsibility to ensure you have specified sufficient Start delay between instructions with dependencies.

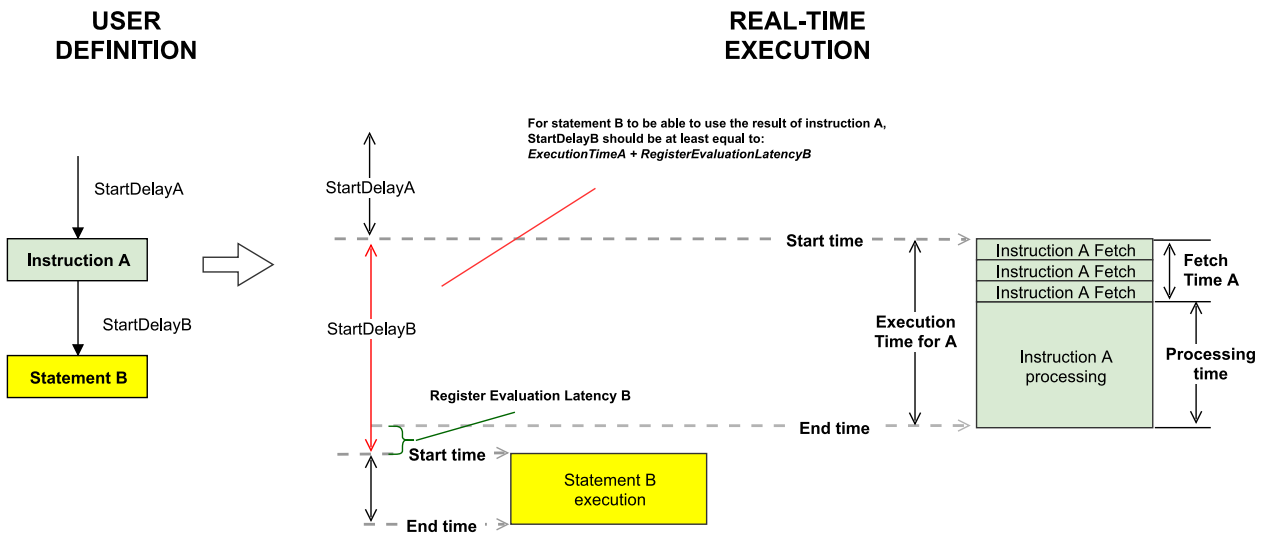
Overlapping instruction execution with Sync or Flow-Control Statement with result dependencies

For the case that the result of an instruction is used from a sync or a flow-control Statement (e.g. register used in the condition of a While or a Sync While), the RegisterEvaluationLatency of that Statement need to be taken into account. Therefore, the formula is updated to:

$$\text{MinDelay_Instr_to_Statement} = \text{Instr_ExecutionTime} + \text{Statement_RegisterEvaluationLatency}$$

NOTE If the flow-control (or sync) Statement comes right after the instruction from which it needs the result, this imposes a minimum value for the StartDelay of that Statement. The final StartDelay to be used should be the maximum between the MinDelay calculated with the previous formula and the MinStartDelay applicable (see [Minimum Start Delay Calculation for Local Flow-Control and Sync Statements](#))

If there are more Statements/instructions between the flow-control (or sync) Statement and the instruction from which it needs the result, then the MinDelay imposes a minimum to the sum of the StartDelays of all the intermediate Statements/instructions and the flow-control (or sync) Statement.



Example cases with instruction result dependencies

The following examples show how to calculate the minimum delay required when the result of HVI Instructions is used by Flow-Control Statements. The latency information is provided in the **Timing Tables** and in the instrument documentation.

Example 1: Instruction "ADD" followed by a Local If Statement

In this example an **Add** instruction writes to a register and the new value of the register is used for the **if** condition.

1. **Reg1 = RegN + 10** (Add).
2. **If(Reg1 > 10)** (the **if** uses the result of the previous **Add** instruction).

In this case, the minimal delay between the **If** and the previous **Add** using the fetch and execution timing is calculated with this equation:

$$\text{MinDelay_If} = \text{Add_ExecutionTime} + \text{If_RegisterEvaluationLatency} = 8 + 3 = 11 \text{ cycles}$$

Example 2: Instruction "ADD" inside a While Statement

In this example there is an **Add** instruction that writes to a register and the new value of the register is used by the while condition.

1. **Reg1 = 0**
2. **While(Reg1 < 1)** (the **While** uses the result of the internal **Add** instruction).
3. **Reg1 = Reg1 + 1** (Add).

In this case, the minimal delay between the **Add** inside the **while** and the condition check for executing one more iteration is calculated with the following equation.

$$\text{MinDelay_While} = \text{Add_ExecutionTime} + \text{While_RegisterEvaluationLatency} = 8 - 3 = 5 \text{ cycles}$$

To add this extra time at the end of the internal while Sequence a Delay Statement can be added. The Delay Statement will need at least 4 cycles of delay which with the EndLatency of the Delay Statement, will give the total the added delay of 5 cycles.

Trigger Write

Trigger Inputs / Outputs are organized together in groups of 16 called TriggerIOs. Each value can be ON or OFF.

Any number of TriggerIOs can be written at the same time.

- #TriggerIOGroupsON is the number of TriggerIOGroups that contain values set to ON.
- #TriggerIOGroupsOFF is the number of TriggerIOGroups that contain values set to OFF.

The Fetch time of the instruction depends on the number of different TriggerIO groups included in the instruction for the two possible values (#TriggerIOGroupsON or #TriggerIOGroupsOFF).

The following table provides some examples.

Triggers ON	Triggers OFF	#TriggerIOGroupsON	#TriggerIOGroupsOFF	Execution time (cycles)	Fetch time (cycles)
1, 2		1	0	2	1
1, 2, 17, 18		2	0	2	1
1, 2	3, 4	1	1	2	1
1, 2, 17, 18	3, 4	2	1	3	2
1, 2, 17, 18	3, 4, 19, 20	2	2	3	2

See your instrument documentation for information about instrument-specific TriggerIO definitions.

Example Trigger write basic timing values:

Instruction	Execution time (cycles)	Fetch time (cycles)
TriggerWrite	Instrument_Trigger_Execution + (#TriggerWriteGroups - 1)	#TriggerWriteGroups

#TriggerWriteGroups = $\text{ceil}[(\text{TriggerIOGroupsON} + \text{TriggerIOGroupsOFF})/2]$, where

- #TriggerIOGroupsON is the number of TriggerIOGroups that contain values set to ON.
- #TriggerIOGroupsOFF is the number of TriggerIOGroups that contain values set to OFF.

NOTE

Trigger execution time is instrument-specific. For Trigger execution timing information, see your instrument documentation.

Action Execute

The Action-execute HVI instruction synchronously executes a list of HVI Actions defined by the user. HVI Actions are organized in groups called ActionGroups that can contain up to 16 Actions. Each instrument defines its own groups of Actions. See the instrument documentation for information about instrument Action definitions and the way they are grouped. Any number of HVI Actions can be executed synchronously, regardless of the group that each Action user belongs to.

However, the number of Action groups included in the Action-execute instruction (#ActionGroups) affects both the Fetch time and the Execution time of the instruction, as shown by the equations in the following table.

Example Action execute basic timing values:

Instruction	Execution time (cycles)	Fetch time (cycles)
ActionExecute	$\text{Instrument_Action_Execution} + \text{INT}[(\#ActionGroups-1) / 2]$	$1 + \text{INT}[(\#ActionGroups - 1) / 2]$

Where INT is the integer part of a decimal number, for instance $\text{INT}(1.0)=\text{INT}(1.5)=1$.

NOTE

Action execution timing is instrument-specific. For Action execution timing information, see your instrument documentation.

Arithmetic Logic Unit Instructions

Arithmetic Logic Unit (ALU) instructions are the register add, subtract or assign operations that are available in the HVI-native instruction set.

ALU instructions basic timing values:

Instruction	Execution time (cycles)	Fetch time (cycles)
Add	8	1
Subtract	8	1
Assign	5	1

FPGA-User Sandbox Instructions

The access latency of the FPGA-Registers and memory map from HVI depends on the implementation of the specific instrument. The following table summarizes the latency for all FPGA-Read FPGA-Write instructions. For the specific value of `Instrument_HVI_FPGA_Latency`, see your instrument documentation.

Example FPGA-User Sandbox operations basic timing values:

Instruction	Execution time (cycles)	Fetch time (cycles)
FpgaArrayRead	$2 * \text{Instrument_HVI_FPGA_Latency} + 4$	1
FpgaArrayRead (Address from HviRegister)	$2 * \text{Instrument_HVI_FPGA_Latency} + 6$	1
FpgaArrayWrite	$\text{Instrument_HVI_FPGA_Latency} + 2$	1
FpgaArrayWrite (Address or data from HviRegister)	$\text{Instrument_HVI_FPGA_Latency} + 4$	1
FpgaRegisterRead	$2 * \text{Instrument_HVI_FPGA_Latency} + 4$	1
FpgaRegisterWrite	$\text{Instrument_HVI_FPGA_Latency} + 2$	1
FpgaRegisterWrite (Address or data from HviRegister)	$\text{Instrument_HVI_FPGA_Latency} + 4$	1

NOTE FPGA-User Sandbox timing is instrument-specific. For FPGA-User Sandbox timing information, see your instrument documentation.

- NOTE**
- Consecutive FPGA-Read instructions must be issued with at least 1 cycle of delay between them.
 - If an FPGA-Instruction that uses an HVI register is issued before an FPGA-Instruction that does not use an HVI register, the delay between both instructions must be at least 3 cycles.

FPGA-Instruction Statement

FPGA-Instruction Statement latency depends on a number of factors:

- Instruction fetch time.
- Time to fetch data from any HVI registers it uses.
- Instrument-specific delays.

Apart from fetch time and the first two execution cycles spent inside the HVI Engine, the rest of the latency is defined by the instrument, this is condensed into the single parameter `Instrument_FpgaInstruction_Latency`. See your instrument documentation for information about the HVI Engine clock frequency and FPGA-Instruction timing information.

FPGA-Instruction Statement timing values:

Instruction	Execution Time (cycles)	Fetch Time (cycles)
FPGA-Instruction	<code>Instrument_FpgaInstruction_Latency + 2</code>	1

Instrument-Specific HVI Instruction Timing Values

Instrument-Specific HVI Instruction latency depends on a number of factors:

- Instruction fetch time.
- Time to fetch data from any HVI registers it uses.
- Instrument-specific delays.

Apart from fetch time and the first two execution cycles spent inside the HVI Engine, the rest of the latency is defined by the instrument and condensed into the single parameter `Instrument_LocalInstruction_Latency`. See your instrument documentation for information about the HVI Engine clock frequency and instrument-specific instruction timing information.

Instrument-Specific HVI Instruction timing values:

Instruction	Execution Time (cycles)	Fetch Time (cycles)
Instrument-Specific HVI Instruction	<code>Instrument_LocalInstruction_Latency + 2</code>	1

HVI Instruction Position Mapping

The following table show the instruction positions that each HVI-native HVI Instruction uses during fetch time, for more information see [HVI Instruction Timing Concepts](#). For instrument custom instructions, see your instrument documentation:

HVI Native Instruction	Positions					
	1	2	3	4	5	...
ActionExecute		Y			-	-
Add		Y			-	-
Assign		Y			-	-
Fpga Array-Read		Y			-	-
Fpga Array-Write		Y			-	-
FPGA-Register-Read		Y			-	-
FPGA-Register-Write		Y			-	-
Subtract		Y			-	-
TriggerWrite		Y			-	-

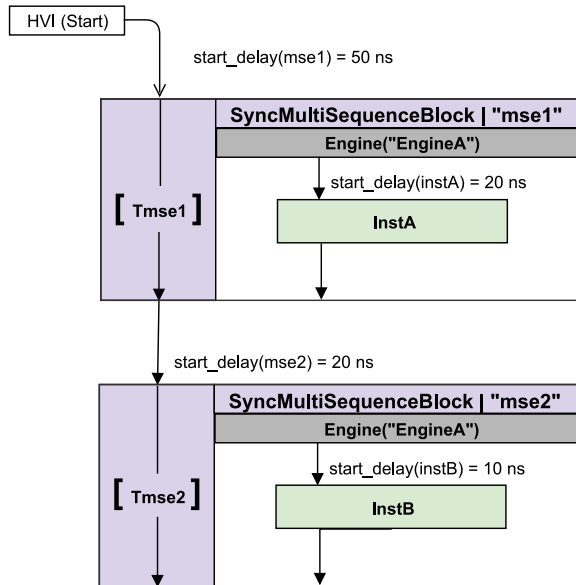
Examples of HVI Instruction Timing Calculation across Sync and Local Flow-Control Statements

This section shows basic examples of HVI Instruction across within Sync and Local Flow-Control Statements and how the timing is calculated.

HVI Instruction timing across Sync Multi-Sequence Blocks example

This example shows a pair of Sync Multi-Sequence Blocks each with a HVI Instruction each. A diagram and the code and timing calculations are shown.

The following is a diagram of the example:



The code for the example:

```
mse1 = sequencer.sync_sequence.add_sync_multi_sequence_block("mse1", 50)
seq = mse1.sequences['EngineA']
instA = seq.add_instruction("instA", 20, seq.instruction_set.trigger_write.id)
#
mse2 = sequencer.sync_sequence.add_sync_multi_sequence_block("mse2", 20)
seq = mse2.sequences['EngineA']
instB = seq.add_instruction("instB", 10, seq.instruction_set.trigger_write.id)
```

The timing calculations for the example:

InstA Execution Start time from HVI-Start ($InstA_start$):

$$InstA_start = start_delay(mse1) + start_delay(instA) = 50ns + 20ns = 70ns$$

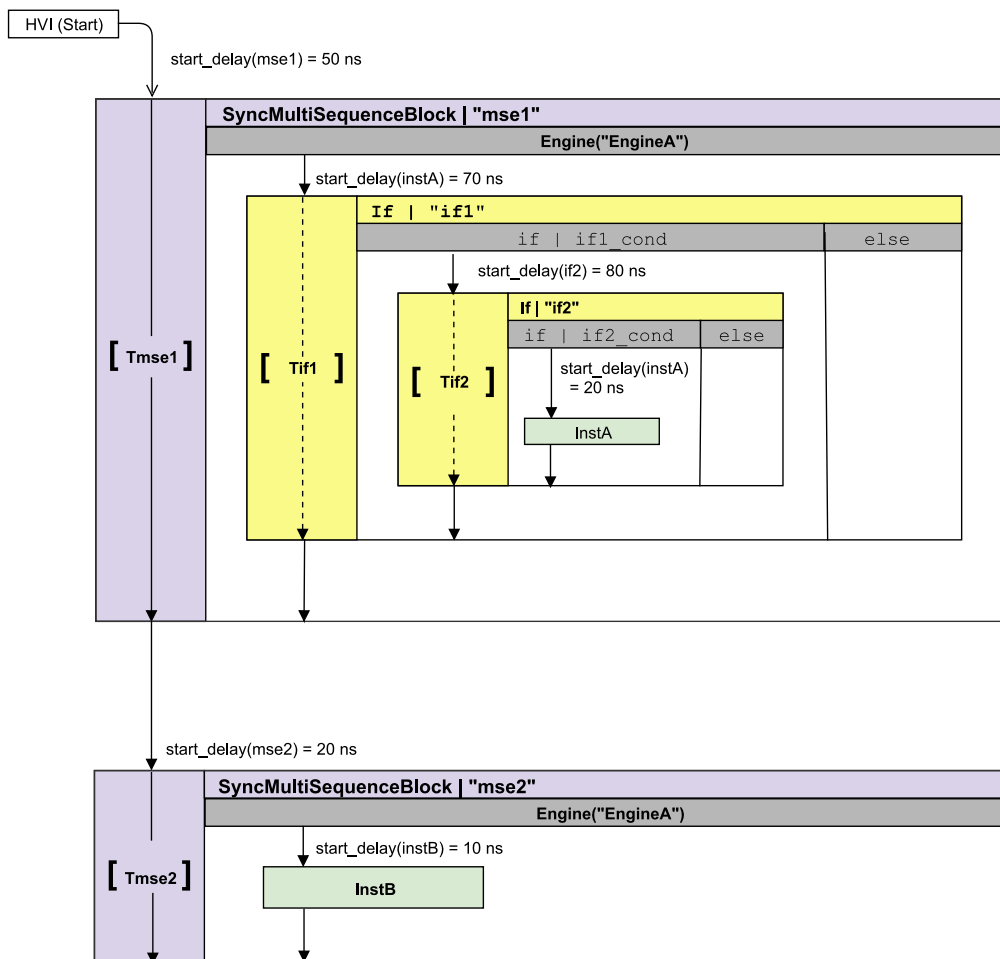
Time from InstA to InstB (T_InstA_InstB):

$$T_InstA_InstB = start_delay(mse2) + start_delay(instB) = 20ns + 10ns = 30ns$$

HVI Instruction timing across Sync Multi-Sequence Blocks and Local If example

This example shows cascaded Local If Statements within a Sync Multi-Sequence Block followed by a HVI Instruction in a Sync Multi-Sequence Block. The code and timing calculations are also shown:

The following is a diagram of the example:



The code for the example:

```
mse1 = sequencer.sync_sequence.add_sync_multi_sequence_block("mse1", 50)
seq = mse1.sequences['EngineA']
if1 = seq.add_if('if1', 70, if1_cond, True)
if1_branch_seq = if1.if_branch.sequence
if2 = if1_branch_seq.add_if('if2', 80, if2_cond, True)
if2_branch_seq = if2.if_branch.sequence
instA = if2_branch_seq.add_instruction("instA", 20, seq.instruction_set.trigger_write.id)
#
mse2 = sequencer.sync_sequence.add_sync_multi_sequence_block("mse2", 20)
seq = mse2.sequences['EngineA']
instB = seq.add_instruction("instB", 10, seq.instruction_set.trigger_write.id)
```

The timing calculations for the example:

The formula to calculate the `InstA` execution start time from HVI-Start, `InstA_start` is:

$$\text{InstA_start} = \text{start_delay}(\text{mse1}) + \text{start_delay}(\text{if1}) + \text{start_delay}(\text{if2}) + \text{start_delay}(\text{instA}) = 50\text{ns} + 70\text{ns} + 80\text{ns} + 20\text{ns} = 220\text{ns}$$

The formula to calculate time from InstA to InstB, T_{InstA_InstB} is:

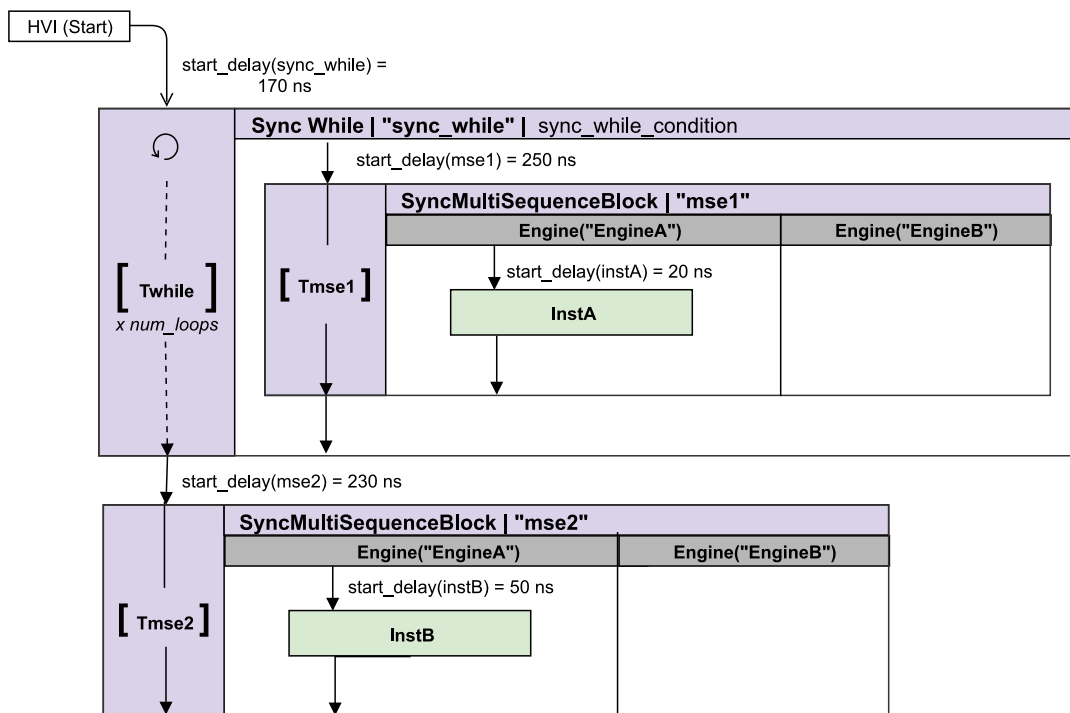
$$T_{InstA_InstB} = start_delay(mse2) + start_delay(instB) = 20ns + 10ns = 30ns$$

NOTE The `end_latency(mse1)` is accounted for in the `start_delay(mse2)`, this imposes a minimum value.

HVI Instruction timing across Sync While and Sync Multi-Sequence Blocks example

This example shows how time is calculated for a Sync while Statement that contains a Sync Multi-Sequence Block and a single instruction:

The following diagram shows the example:



The following block shows the example code:

```
sync_while = sequencer.sync_sequence.add_sync_while('sync_while', 170, sync_while_condition)
mse1_sequence = sync_while.sync_sequence.add_sync_multi_sequence_block("mse1", 250).sequences
['EngineA']
instA = mse1_sequence.add_instruction("InstA", 20, seq.instruction_set.assign.id)
#
mse2_sequence = sequencer.sync_sequence.add_sync_multi_sequence_block("mse2", 230).sequences
['EngineA']
instB = mse2_sequence.add_instruction("InstB", 50, seq.instruction_set.assign.id)
```

The following are the equations used to calculate the timing in the example:

InstA Execution Start time from HVI-Start, $InstA_start$:

$$InstA_start = start_delay(sync_while) + start_delay(mse1) + start_delay(instA) = 170ns + 250ns + 20ns = 440ns$$

Sync Multi-Sequence Block Execution time, $Tmse1$:

$$Tmse1 = SequenceTime = 20ns$$

Sync while Execution time for 1 loop when looping, $Twhile_loop$:

$$Twhile_loop = Twhile = \{start_delay(mse1) + Tmse1\} = \{250ns + 20ns\} = 270ns$$

Time from InstA to InstA in consecutive repetitions, $Tloop_InstA$:

$$Tloop_InstA = Twhile_loop$$

Time from InstA to InstB (the last Sync while execution), T_InstA_InstB :

$$T_InstA_InstB = start_delay(mse2) + start_delay(instB) = 230ns + 50ns = 280ns$$

NOTE

The `end_latency(sync_while)` is accounted for in the `start_delay(mse2)`. This imposes a minimum value.

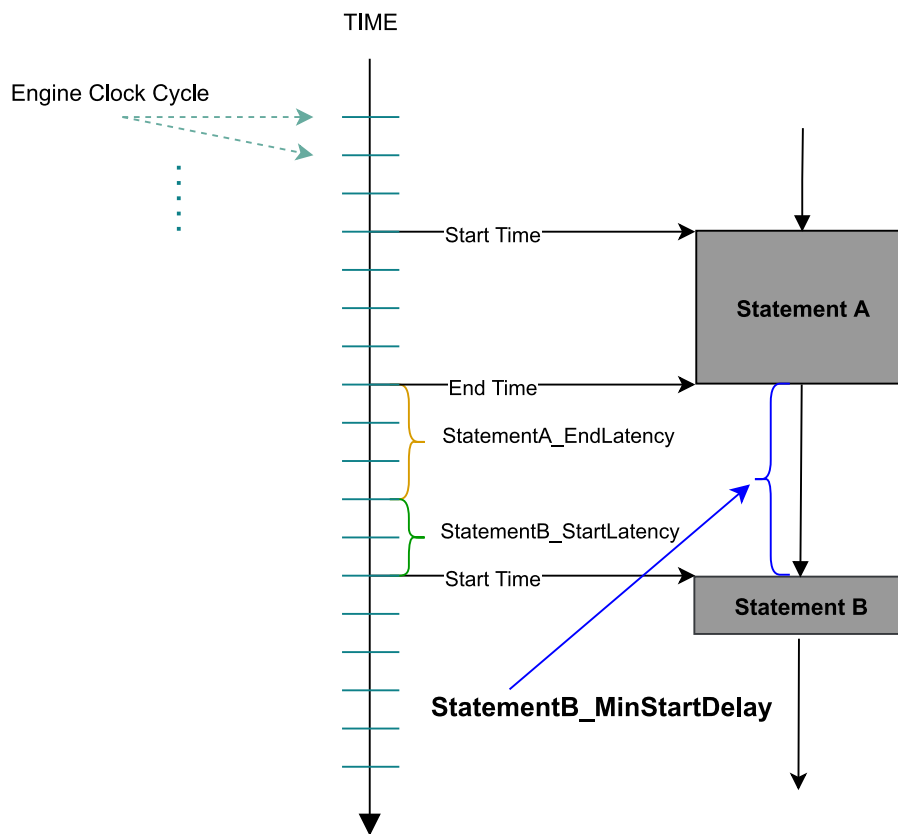
Minimum Start Delay Calculation for Local Flow-Control and Sync Statements

This section describes minimum Start Delay calculations for Local Flow-Control and Sync Statements, it contains the following sections:

- Minimum Start Delay for Local Flow-Control Statements
- Minimum Start Delay for Sync Statements

To calculate the minimum valid Start Delay for a given Local Flow-Control or Sync Statement, the general rule is to add the End-Latency of the previous Statement with the Start-Latency of the current Statement:

$$\text{Statement_MinStartDelay}_{\text{EngineX}} = \text{PreviousStatement_EndLatency}_{\text{EngineX}} + \text{Statement_StartLatency}_{\text{EngineX}}$$



From this general rule, we can distinguish 2 subcases:

- First Statement of the global HVI Sequence: Use the End-Latency of the HVI Start.
- First Statement of a sub Sequence: Instead of the End-Latency use the Entry-Latency of the parent Statement.

The values for each latency can be found in the Timing tables for Local Flow-Control Statements and Sync Statements.

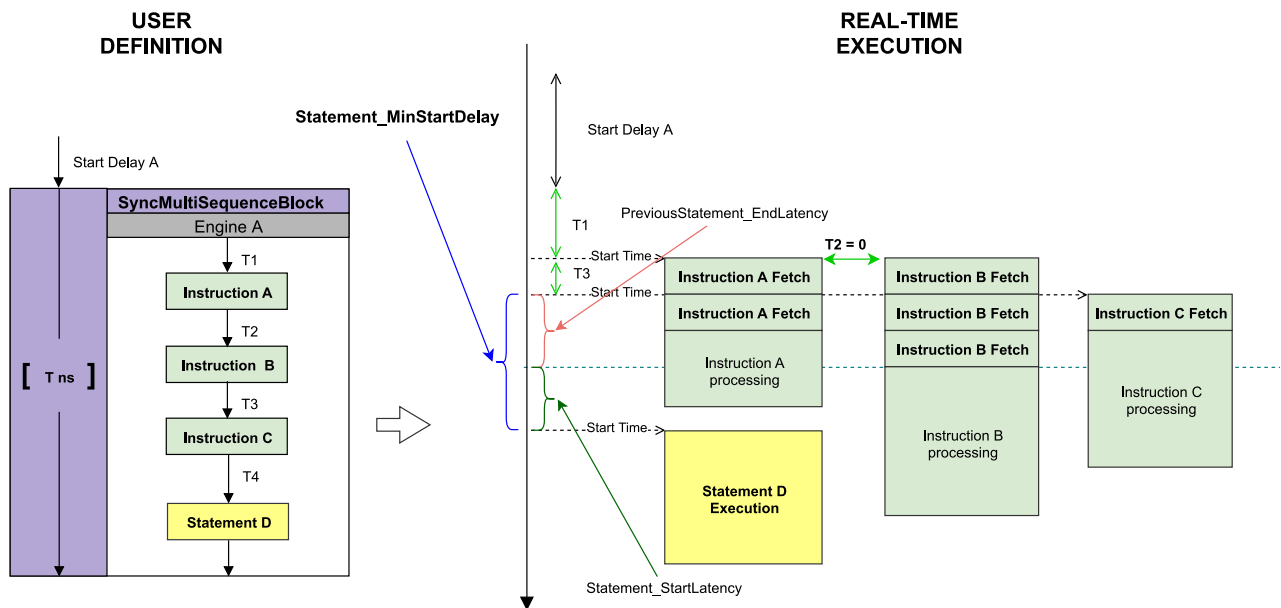
Minimum Start Delay for Local Flow-Control Statements

The Local Flow-Control Statements are following the general rule described above to calculate the minimum Start Delay.

Minimum Start Delay after HVI Instructions

As explained earlier, to find out the minimum Start-Delay of a Statement, it is important to know the end-latency of the previous Statement. If the previous Statement of a Local Flow-Control Statement is one or more HVI Instructions, the end-latency is calculated as the remaining fetch-cycles of all the HVI Instructions starting from the beginning of the last instruction.

This can be seen in the following figure. Starting from the beginning of Instruction C (last HVI Instruction), we calculate the remaining fetch cycles of all the instructions executed before Statement D. From the picture we see that there is one fetch cycle where instructions A, B, C are executed together and then, one more fetch cycle for instruction B. So, in total, there are 2 remaining fetch cycles, therefore, the end-latency is 2 cycles.



End-Latency of Local Flow-Control Statements with internal Sequence (If and While)

The End-Latency of Local Flow-Control Statements with internal Sequence, like If and While Statements, depend on the End-Latency of the last Statement of their internal Sequences. When the last Statement of the internal Sequence is one or more HVI Instructions, for the calculation of the End-Latency, the same principle applies as described in the previous section, i.e. the end-latency is calculated as the remaining fetch-cycles of all the HVI Instructions starting from the beginning of the last instruction.


Minimum Start Delay for Sync Statements

For the **Sync Statements**, the minimum Start Delay is the maximum of all the minimum Start Delays calculated for each HVI Engine rounded to the next multiple of the *HVI Common Clock* period:

$$\text{Statement_MinStartDelay} = \text{round}_{\text{ncc}}(\max\{\text{Statement_MinStartDelay}_{\text{Engine1}}, \text{Statement_MinStartDelay}_{\text{Engine2}}, \dots, \text{Statement_MinStartDelay}_{\text{EngineN}}\})$$

where Roundncc $\text{round}_{\text{ncc}}(.)$ is an operation that rounds to the next multiple of the HVI Common Clock period.

In the following diagram, we show graphically the minimum Start Delay calculation process between two Sync Statements in a system that has two Engines with different frequencies, EngineA and EngineB:



Example: Minimum Start Delay from HVI Start to Sync While

In this example we show how to calculate the minimum Start Delay value acceptable for a Sync While Statement that is placed as the first Statement of the HVI root SyncSequence.



```
# Create SystemDefinition object
system_definition = keysight_tse.SystemDefinition("MySystemDefinition")
system_definition.engines.add(instrument_1.hvi.engines.leader, "HVI_Engine_1")
```

```

system_definition.engines.add(instrument_1.hvi.engines.leader, "HVI_Engine_2")
...
# Create Sequencer object
sequencer = keysight_tse.Sequencer("MySequencer", system_definition)
# Iteration counter register for "HVI_Engine_2"
iteration_counter = sequencer.sync_sequence.scopes["HVI_Engine_2"].registers.add("MyRegister",
keysight_tse.RegisterSize.SHORT)
iteration_counter.initial_value = 0
# Define sync while condition
num_loops = 5
sync_while_condition = keysight_tse.Condition.register_comparison(iteration_counter, keysight_tse.ComparisonOperator.LESS_THAN, num_loops)
SyncWhile_MinStartDelay = ... # The calculation for the minimum is explained below
sequencer.sync_sequence.add_sync_while("MySyncWhileStatement", SyncWhile_MinStartDelay, sync_while_condition)
    
```

We assume the following values to be used in the calculations:

Variable	Value	Description
#Register_Conditions	1	<i>In the example above, we used only one register condition for the Sync While</i>
Engine1Period	5 ns	<i>We assume HVI Engine 1 to run at frequency of 200 MHz which results in a 5 ns period</i>
Engine2Period	3.333 ns	<i>We assume HVI Engine 2 to run at frequency of 300 MHz which results in a 3.333 ns period</i>
HVI_Leader_Engine_ClockPeriod	3.333 ns	<i>The leader Engine for the Sync While in the example is HVI Engine 2, since the register used in the condition of the Sync While belongs to that Engine. So, the frequency of the leader Engine is that of Engine 2</i>
HVI_Common_ClockPeriod	10 ns	<i>This is the result of the GCD of the Engines included in HVI: $LCM\{Engine1Period, Engine2Period\}$</i>

Engine 1:

Using the timing table of HVI Start, we calculate the End Latency for this Engine:

HVI Start		
Parameter	Time (cycles)	Result (n s)
HviStart_EndLatency _{Engine1}	2	10

Using the timing table of the Sync While Statement, we calculate the Start Latency for this Engine:

Sync While		
Parameter	Time (cycles)	Result (n s)
<code>SyncWhile_StartLatency_{Engine1}</code>	Follower Engine: 2	10

Therefore, the minimum for this Engine is:

$$\text{SyncWhile_MinStartDelay}_{\text{Engine1}} = \text{HviStart_EndLatency}_{\text{Engine1}} + \text{SyncWhile_StartLatency}_{\text{Engine1}} = \mathbf{20 \text{ ns}}$$

Engine 2:

Using the timing table of HVI Start, we calculate the End Latency for this Engine:

HVI Start		
Parameter	Time (cycles)	Result (n s)
<code>HviStart_EndLatency_{Engine2}</code>	2	$6.66\bar{6}$

Using the timing table of the Sync While Statement, we calculate the Start Latency for this Engine:

Sync While		
Parameter	Time (cycles)	Result (n s)
<code>SyncWhile_StartLatency_{Engine1}</code>	Leader Engine: 6 + <code>#Register_Conditions</code>	$23.33\bar{3}$

Therefore, the minimum Start Delay for this Engine is:

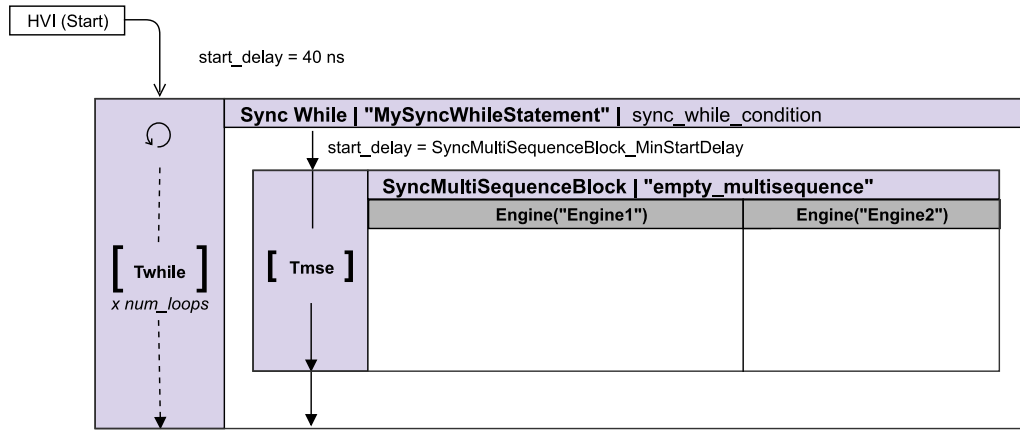
$$\text{SyncWhile_MinStartDelay}_{\text{Engine2}} = \text{HviStart_EndLatency}_{\text{Engine2}} + \text{SyncWhile_StartLatency}_{\text{Engine2}} = \mathbf{30 \text{ ns}}$$

Minimum Start Delay:

$$\text{SyncWhile_MinStartDelay} = \text{round}_{\text{ncc}}(\max\{\text{SyncWhile_MinStartDelay}_{\text{Engine1}}, \text{SyncWhile_MinStartDelay}_{\text{Engine2}}\}) = \mathbf{30 \text{ ns}}$$

Example: Minimum Start Delay for the first Statement inside a Sync While

Continuing from the previous example, in this example we show how to calculate the minimum Start Delay value acceptable for a Sync Multi-Sequence Statement that is placed as the first Statement of a Sync While internal SyncSequence.



#...

```
SyncWhile_MinStartDelay = 30 # As calculated earlier
sync_while_statement = sequencer.sync_sequence.add_sync_while("MySyncWhileStatement",
SyncWhile_MinStartDelay, sync_while_condition)
SyncMultiSequence_MinStartDelay = ... # The calculation for the minimum is explained below
sync_while_statement.sync_sequence.add_sync_multi_sequence_block("empty_multisequence",
SyncMultiSequence_MinStartDelay)
```

In addition to the Variables provided in the previous example, we assume the following values to be used in the following calculations:

Variable	Value	Description
Instrument_SyncResources_Latency	0 cy	This is an instrument dependent value. We assume it to be 0 for this example.
Propagation_delayCycles	30 cy	Assuming that we use only 1 chassis in this example, the Propagation delay would be 100 ns. Translating it to cycles of the leader Engine of the Sync While, that gives us 30 cycles.
End-Latency _{Last-statement}	0 cy	The last Statement is an empty Sync Multi-Sequence Block which, according to the timing tables, will have 0 cycles of end-latency.

Engine 1:

Using the timing table of the Sync While Statement, we calculate the End Latency for this Engine:

HVI Start		
Parameter	Time (cycles)	Result (n s)
SyncWhile_EntryLatency _{Engine1}	$\text{match}\{\text{LatencyA}_{\text{Engine1}}, \text{LatencyA}_{\text{Engine2_inEngine1Cycles}}\} + 2 + \text{End-Latency}_{\text{Last-statement}}$	160*

*Calculation:

$$\text{LatencyA}_{\text{Engine1}} = 2 \text{ cy (Engine1 cycles)}$$

$$\text{LatencyA}_{\text{Engine2}} = 12 + \#\text{Register_Conditions} + \text{Instrument_SyncResources_Latency} + \text{Propagation_delay_cycles} = 43 \text{ cy (Engine2 cycles)}$$

$$\text{LatencyA}_{\text{Engine2_inEngine1Cycles}} = \text{ceil}(43 * \text{Engine2Period} / \text{Engine1Period}) = 29 \text{ cy (Engine1 cycles)}$$

$$\text{SyncWhile_EntryLatency}_{\text{Engine1}} = (\text{match}\{2,29\} + 2 + 0) * \text{Engine1Period} = (\text{round}_{\text{ncc}}(29) + 2) * 5 = 160 \text{ ns}$$

Using the timing table of the Sync Multi-Sequence Statement, we calculate the Start Latency for this Engine:

Sync While		
Parameter	Time (cycles)	Result (n s)
SyncMultiSequence_StartLatency _{Engine1}	1	5

Therefore, the minimum for this Engine is:

$$\text{SyncMultiSequence_MinStartDelay}_{\text{Engine1}} = \text{SyncWhile_EntryLatency}_{\text{Engine1}} + \text{SyncMultiSequence_StartLatency}_{\text{Engine1}} = 165 \text{ ns}$$

Engine 2:

Using the timing table of the Sync While Statement, we calculate the End Latency for this Engine:

HVI Start		
Parameter	Time (cycles)	Result (n s)
SyncWhile_EntryLatency _{Engine2}	$\text{match}\{\text{LatencyA}_{\text{Engine1_inEngine2Cycles}}, \text{LatencyA}_{\text{Engine2}}\} + 2 + \text{End-Latency}_{\text{Last-statement}}$	$156.66\bar{6}^*$

*Calculation:

$\text{LatencyA}_{\text{Engine1_inEngine2Cycles}} = \text{ceil}(\text{LatencyA}_{\text{Engine1}} * \text{Engine1Period} / \text{Engine2Period}) = 3 \text{ cy (Engine2 cycles)}$

$\text{LatencyA}_{\text{Engine2}} = 43 \text{ cy (Engine2 cycles)}$

$\text{SyncWhile_EntryLatency}_{\text{Engine2}} = (\text{match}\{3,43\} + 2 + 0) * \text{Engine2Period} = (\text{round}_{\text{ncc}}(43) + 2) * 3.33\bar{3}$
 $= 156.66\bar{6} \text{ ns}$

For an explanation of the `match()` function, see [Sync Statement Timing](#).

Using the timing table of the Sync Multi-Sequence Statement, we calculate the Start Latency for this Engine:

Sync While		
Parameter	Time (cycles)	Result (n s)
SyncMultiSequence_StartLatency _{Engine2}	1	$3.33\bar{3}$

Therefore, the minimum Start Delay for this Engine is:

$\text{SyncMultiSequence_MinStartDelay}_{\text{Engine2}} = \text{SyncWhile_EntryLatency}_{\text{Engine2}} + \text{SyncMultiSequence_StartLatency}_{\text{Engine2}} = 160 \text{ ns}$

Minimum Start Delay:

$\text{SyncMultiSequence_MinStartDelay} = \text{round}_{\text{ncc}}(\text{max}\{\text{SyncMultiSequence_MinStartDelay}_{\text{Engine1}}, \text{SyncMultiSequence_MinStartDelay}_{\text{Engine2}}\}) = 170 \text{ ns}$

Errors in Start Delay or Duration specification

This section describes Error and Warnings.

Errors when setting Start Delays and how to deal with them

The previous section explains how to calculate the Start Delay for each type of Statement. Depending on the type of Statements as well as the number of Engines with different frequencies, this can be a complex and error-prone procedure.

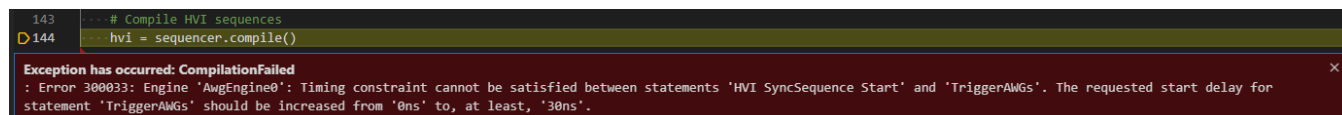
In the case of an error, the HVI compiler will validate the provided values and generate a message with the minimum Start Delay applicable.

Example: Fixing a Start Delay with a compiler message

In the following example code, a Sync Multi-Sequence Block is added as the first Statement to a Sync Sequence. To get the compiler to generate an error that indicates the minimum Start Delay value that can be used to add the Sync Multi-Sequence Block, you can set the Start Delay to 0 as shown in the following example code snippet.

```
# Create SystemDefinition object
my_system = kthvi.SystemDefinition("MySystem")
...
# Create Sequencer object
sequencer = kthvi.Sequencer("MySequencer", my_system)
# Add a Sync Multi-Sequence Block (SMSB) with a 0 ns Start Delay
sync_block = sequencer.sync_sequence.add_sync_multi_sequence_block("TriggerAWGs", 0)
```

When the Sequencer object is compiled, the compiler detects that **0 ns** does not comply with the minimum latency for the Sync Multi-Sequence Block in the HVI Sequence. It returns an error message stating that the specified Start Delay shall be at least **30 ns**. See the following image for an example of the returned error message.



```
143  ... # Compile HVI sequences
D 144  hvi = sequencer.compile()

Exception has occurred: CompilationFailed
: Error 300033: Engine 'AwgEngine0': Timing constraint cannot be satisfied between statements 'HVI SyncSequence Start' and 'TriggerAWGs'. The requested start delay for statement 'TriggerAWGs' should be increased from '0ns' to, at least, '30ns'.
```

The reason why a minimum latency of 30 ns is required is explained in this chapter in the section [Sync Statement Timing](#). In a similar way, you can set any Start Delay to 0 ns and let the compiler error messages provide the minimum latency required for each of those Start Delays. The reasons behind the specific minimum values advised by the compiler are explained in the rest of this chapter.

NOTE**Limitations:**

The compiler provides an indication of the minimum Start Delay for each Engine. If you are using different instruments these can be different values. That is because the clock cycle duration is different in each instrument. For example, a minimum latency of 3 cycles lasts 30 ns on M3xxxA instruments and 10 ns on M5xxxA instruments. In case the compiler error messages suggest different values, ***pick the highest value of those indicated and round it to the next HVI Common Clock cycle*** to set a Start Delay value that can accommodate the requirements for all the different instruments that are included in the HVI.


Error and Warning Margins Related to Timing Resolution

PathWave Test Sync Executive implements a policy for error and warning margins when you specify the timing for a Start Delay or a duration.

The following table shows example values for an instrument with a 300MHz clock ($3.\overline{3}$ ns clock period):

Range Type	Range	Example	Description
No Error or Warning	± 10 ps	3.323ns to 3.343ns	If you set a value with ± 10 ps error from the exact clock period multiplier value, no error or warning is generated.
Warning	± 100 ps	3.233ns to 3.323ns, or 3.343ns to 3.433ns	If you set a value between ± 10 ps and ± 100 ps of the exact clock period multiplier value, a warning is generated.
Error	> 100 ps	0.000ns to 3.233ns, or 3.433ns to 6.566ns	If you set a value with more than ± 100 ps error from the exact clock period multiplier value, an error is generated.

The following diagram shows an example where the exact clock period multiplier value is $3.\overline{3}$ ns, this is the same as the example in the table.



To calculate the margins for other period multiplier values, warnings are $\pm 10\text{ps}$ from the exact value and errors are $\pm 100\text{ps}$ away from the exact value.

Appendix A: Supported Instruments

PathWave Test Sync Executive supports a number of instruments and PXIe chassis, these require specific minimum software and firmware versions to work with PathWave Test Sync Executive.

The software and firmware version requirements for the supported instruments and chassis are listed on-line here: [Instrument and Chassis Software and Firmware Requirements for KS2201A](#) .

Product specific documentation

For product-specific information and documentation please refer to the product pages.

Firmware is available at [Keysight PXI Products](#), on the **Technical Support** page for your specific instruments, see the **Drivers, Firmware & Software** tab.

M3000 Series

The M3000 series (SD1) software provides drivers, programming libraries and software front panels for the M3000 series.

Instruments are shipped with the latest versions of firmware and SD1 software. To use an older instrument with PathWave Test Sync Executive, the firmware and SD1 software must be upgraded to the versions recommended in the product page following the guidelines at the link above. SD1 software is available at [Keysight SD1 Software](#).

Other Instruments

Instruments are provided with their own drivers, programming libraries, and software front panels, and are shipped with the latest versions of firmware and software. To use an older instrument with PathWave Test Sync Executive, the firmware and software must be upgraded to the versions recommended in the product page following the guidelines at the link above.

PXIe Chassis

The Chassis software provides drivers, programming libraries and software front panels for the Keysight chassis.

Chassis are shipped with the latest versions of firmware and software. To use an older chassis with PathWave Test Sync Executive, the firmware and software must be upgraded to the versions recommended in the product page following the guidelines at the link above.

Compatibility with M3601A

M3601A is an older generation of HVI technology that is only programmable by the M3601A Hard Virtual Instrument Design Environment. PathWave Test Sync Executive is a new generation and is not backward compatible with the M3601A generation.

Both PathWave Test Sync Executive and M3601A work with the M3000 series of PXIe products. However, PathWave Test Sync Executive requires newer firmware while M3601A requires older firmware.

Appendix B: Additional Documentation and Examples

This appendix lists the PathWave Test Sync Executive Programming Examples and additional documentation that you can download from the [KS2201A Programming Examples](#) page.

NOTE

The Programming Examples are often updated so ensure you check for the latest versions.

Programming Example 1: Multi-Channel Sync Playback using M32xxA Arbitrary Waveform Generators

In Programming Example 1, PathWave Test Sync Executive is used to program multiple M3xxxA Arbitrary Waveform Generators (AWGs). The AWDs synchronously output a front panel trigger pulse followed by a previously queued waveform. All instruments run fully synchronized and Actions across the instruments can be controlled at the timing resolution of the M3xxxA AWDs, which is 10ns.

Programming Example 2: Synchronous Signal Generation and Acquisition using M3xxxA PXI Instruments

In Programming Example 2, a M3102A digitizer performs sequenced acquisition of heterogeneous signals generated by multiple M320xA AWDs. The first AWD generates a train of RF pulses and the other AWDs output a queued arbitrary waveform. By using PathWave Test Sync Executive, each cycle of the digitizer measurements is precisely synchronized with the AWD output signals.

Programming Example 3: PathWave Test Sync Executive Integration with PathWave FPGA

This Programming Example shows how to use Keysight PathWave Test Sync Executive together with Keysight PathWave FPGA. A custom FPGA block is designed using Keysight PathWave FPGA and loaded into the FPGA-Sandbox of two modular instruments. The two instruments execute HVI Sequences that can communicate with the custom FPGA blocks programmed into the FPGA-Sandbox of the module FPGA. Using an HVI Port, the HVI Sequence can read/write values in any HVI Port Register inserted among the custom FPGA blocks. This example also shows how the HVI Sequence and FPGA-Sandbox of an instrument can communicate by using Actions and Events. The exchanged information can also be written to PXI lines.

Programming Example 4: Real-Time Pulsed Characterization of a Device-Under-Test

In this Programming Example, an M3202A AWG and an M3102A digitizer are used to perform a real-time pulsed characterization experiment on a Device Under Test.

A pool of different waveforms is loaded to the AWG RAM. The digitizer uses the register-sharing functionality to select a real-time the waveform to be played by the AWG at each iteration of the experiment. The selected waveform is used by AWG CH1 and CH2 to play I-Q modulated pulses and re-play them after a Variable delay. In the same iteration, AWG CH3 and CH4 play a second burst of I-Q pulses after another Variable delay. The second burst pulse length can be increased after each iteration. The experiment can be repeated for a user-defined number of loops, allowing you to choose the delay between each loop and the delay necessary for example to let the DUT return to its equilibrium state. Example use cases for this programming example include power amplifier characterization for 5G mobile communications and quantum bit characterization experiments for physics applications. In the physics case, the AWG generates the control and readout pulses necessary for characterization of quantum bits.

Programming Example 5 - Synchronized Multi-Channel Mixed-Signal Generation using M3xxA PXI Instruments

In this Programming Example, KS2201A PathWave Test Sync Executive is used to program multiple M3xxx Arbitrary Waveform Generators to synchronously generate mixed signals. Each instrument can be programmed to output either a front panel marker pulse or a previously queued waveform. All signal channels run fully synchronized and Actions across instruments can be controlled with the timing resolution of the M3xxx AWGs, which is 10ns.

Programming Example 6 - Synchronized MIMO Measurements using M5302A Digital I-O and M3xxxA PXI Instruments

In this programming example, PathWave Test Sync Executive is used to program multiple M5302A Digital I/O (DIO) and M3xxxA PXI instruments. By using HVI (Hard Virtual Instrument) capabilities, DIO instruments can output a pulsed signal from any of their Front Panel (FP) SMB trigger ports and M320xA AWGs can synchronously play a previously queued waveform. Multiple M3102A Digitizers can also be included in the same HVI to synchronously capture all the generated analog and digital signals. This way the example can showcase a Multiple-Input Multiple-Output (MIMO) measurement setup having all his input and output channels fully synchronized.

Programming Example 7 - RF Sweeps using M320x AWGs M5300 RF AWGs and M9046 Chassis

In this programming example, PathWave Test Sync Executive is used to define a real-time algorithm to be executed by the FPGA (Field Programmable Gate Array) of Arbitrary Waveform Generators (AWGs). This enables the AWG channels to be used to output pulsed signals that are swept in amplitude and frequency, to perform a pulsed characterization of a Device-Under-Test.

System Setup Guide

This document describes the different ways you can set up a single or multi-chassis system, with clocking and communications.

Transitioning from M3601A HVI Programming Environment to KS2201A PathWave Test Sync Executive

This Transition Guide is intended for M3601A users and explains how to translate an M3601A project into TSE API Python code programmed using Keysight KS2201A PathWave Test Sync Executive.

Appendix C: Timing Tables

This appendix lists statement timing tables for quick access. It contains the following sections:

- [Sync Statement Timing Tables](#)
- [Local Flow-Control Statement Timing Tables](#)
- [HVI Instruction Timing Tables](#)

Sync Statement Timing Tables

Sync Multi-Sequence Block

Timing value for Sync Multi-Sequence Blocks:

Execution time (cycles) (1)
$\text{round}_{\text{ncc_cycles}} \left(\text{sum}_{\text{for_all_internal_statements}}(\text{StartDelayCycles}) + \text{sum}_{\text{for_all_internal_flow_control_statements}}(\text{DurationCycles}) \right) \quad (2)$

The following tables shows latency values for Sync Multi-Sequence Blocks:

Parameter	Description		Time(cycles)
Start-Latency	Minimum start-delay for Statement		1
Entry-latency	Minimum start-delay for first Statement inside any of the contained Sequences		1
End-Latency	Minimum start-delay for the next Statement	timed-sync (5) Minimum Duration	$\text{round}_{\text{ncc_cycles}}(\text{End-Latency}_{\text{Last-statement-of-longest-branch}}(3) - 1)$ * if the last Statement of the longest branch is not starting from a common clock cycle (see section Sync Multi-Sequence Block Timing and Time Matching in Sync Statement Timing), the formula is updated to: $\text{round}_{\text{ncc_cycles}}(\text{End-Latency}_{\text{Last-statement-of-Longest-branch}}(3) - 1 - \text{DistanceToNextCommonClock})$ where: - <i>DistanceToNextCommonClock</i> is the number of Engine Cycles from the start of the last Statement to the following common clock cycle.
		timed-sync (5) Fixed Duration	0
		triggered-sync (5)	0

Fixed-Duration	Minimum fixed-duration for Statement	$\text{round}_{\text{ncc_cycles}}([\text{max}_{\text{for_all_Sequences}}[\text{Sequence-Duration}]])(4),$ <p>where <i>Sequence-Duration</i> is calculated as follows:</p> $\text{sum}_{\text{for_all_internal_statements}}(\text{Start-Delay}) + \text{sum}_{\text{for_all_internal_flow_control_statements}}(\text{Duration}) + \text{End-Latency}_{\text{Last-statement}} - 1$
----------------	--------------------------------------	---

- (1) The values provided here apply if the duration property of the Statement is set to Minimum (default). If a fixed-duration has been set, then the Execution time is equal to that value.
- (2) The values are only calculated for the branch that is being executed, if there are multiple branches available.
- (3) If the Sequence is empty, the value is 0.
- (4) If the Sequence is empty, then the duration is 0.
- (5) Triggered-sync is required if any of the Sequences in a Sync Multi-Sequence Block contains a Statement that has unknown execution time at compile time.

Sync While

Timing value for Sync While Statement:

Execution time (cycles) (1)
$\text{round}_{\text{ncc_cycles}} (\#Iterations * [\text{sum}_{\text{for_all_internal_statements}}(\text{StartDelay}_{\text{Cycles}}) + \text{sum}_{\text{for_all_internal_flow_control_statements}}(\text{Duration}_{\text{Cycles}})])$

The following tables shows latency values for the Sync While Statement:

Parameter	Description		Time (cycles)
Start-Latency	Minimum start-delay for Statement		<ul style="list-style-type: none"> Leader Engine: $6 + \#Register_Conditions$ Follower Engine(s): 2
Entry/Iteration latency	Minimum start-delay for first Statement inside the while loop	Minimum Duration	$\text{match}\{LatencyA_{LeaderEngine}, LatencyA_{FollowerEngine1}, \dots\} + 2 + \text{End-Latency}_{Last-statement}(2)$ <p>where <i>LatencyA</i> is :</p> <ul style="list-style-type: none"> Leader Engine(3): $12 + \#Register_Conditions + Instrument_SyncResources_Latency(4) + Propagation_delayCycles$ Follower Engine(s): 2
		Fixed Duration	$\text{match}\{LatencyB_{LeaderEngine}, LatencyB_{FollowerEngine1}, \dots\} + 2$ <p>where :</p> <ul style="list-style-type: none"> $LatencyB = LatencyA - 1$ <i>LatencyA</i> as defined above

Parameter	Description		Time (cycles)
End-Latency	Minimum start-delay for next Statement outside the while loop	Minimum Duration	$\text{match}\{\text{LatencyA}_{\text{LeaderEngine}}, \text{LatencyA}_{\text{FollowerEngine1}}, \dots\} + \text{match}\{\text{LatencyC}_{\text{LeaderEngine}}, \text{LatencyC}_{\text{FollowerEngine1}}, \dots\} + \text{End-Latency}_{\text{Last-statement}}(2),$ <p>where:</p> <ul style="list-style-type: none"> • LatencyA as defined above • LatencyC is 2 for each Engine.
		Fixed Duration	$\text{match}\{\text{LatencyB}_{\text{LeaderEngine}}, \text{LatencyB}_{\text{FollowerEngine1}}, \dots\} + \text{match}\{\text{LatencyC}_{\text{LeaderEngine}}, \text{LatencyC}_{\text{FollowerEngine1}}, \dots\},$ <p>where:</p> <ul style="list-style-type: none"> • LatencyB as defined above • LatencyC as defined above
Fixed-Duration	Minimum fixed-duration for Statement	Sync-While Branch with at least one Statement inside	$\text{round}_{\text{ncc_cycles}}(\text{sum}_{\text{for_all_internal_statements}}(\text{StartDelay}_{\text{Cycles}}) + \text{sum}_{\text{for_all_internal_flow_control_statements}}(\text{Duration}_{\text{Cycles}}) + 1 + \text{End-Latency}_{\text{Last-statement}})$
		Empty Sync-While Branch	$\text{match}\{\text{LatencyB}_{\text{LeaderEngine}}, \text{LatencyB}_{\text{FollowerEngine1}}, \dots\} + \text{match}\{\text{LatencyC}_{\text{LeaderEngine}}, \text{LatencyC}_{\text{FollowerEngine1}}, \dots\} + 1,$ <p>where:</p> <ul style="list-style-type: none"> • LatencyB as defined above • LatencyC as defined above

Parameter	Description	Time (cycles)
Register Evaluation Latency	Time to evaluate the Register condition	Leader Engine (Only): <ul style="list-style-type: none"> From start: 2 For each iteration: $-(3 + \#Register_Conditions)$

- (1) This value applies if the duration property of the statement is set to Minimum (default). If a fixed-duration has been set, then the Execution time is equal to that value.
- (2) If the sequence is empty, the value of `End-LatencyLast-statement` is 0.
- (3) In the context of this statement, Leader is the engine that contains the Register or Registers used in the while condition.
- (4) `Instrument_SyncResources_Latency` is an instrument specific value. For more information see the instrument documentation.

Sync Register Sharing

Timing value for Sync Register-Sharing Statement:

Execution time (cycles) (1)
<code>round_{ncc_cycles}(5 + Propagation_delay_{Cycles}) (2)</code>

Latency values for Sync Register-Sharing Statement:

Parameter	Description	Time (cycles)
Start-Latency	Minimum start-delay for Statement	1
End-Latency	Minimum start-delay for the next Statement	0
Fixed-Duration	Minimum fixed-duration for Statement	$\text{round}_{\text{ncc_cycles}}(5 + \text{Propagation_delay}_{\text{cycles}})^{(2)}$
Register Evaluation Latency	Time to evaluate the Register condition	-1

(1) The value provided here applies if the duration property of the Statement is set to Minimum (default). If a fixed-duration has been set, then the Execution Time is equal to that value.

(2) This latency needs to be calculated only on the Leader Engine. In the context of this Statement, Leader is the Engine that contains the Register(s) used as source.

Sync FPGA Data Sharing

Latency values for Sync FPGA Data-Sharing Statement:

Parameter	Description	Time (cycles)
Start-Latency	Minimum start-delay for Statement	1
End-Latency	Minimum start-delay for the next Statement	0

Local Flow-Control Statement Timing Tables

Local While Statement Timing Tables

Local While timing value:

Execution time (cycles) ⁽¹⁾
$\#Iterations * [\text{sum}_{\text{for_all_internal_statements}}(\text{Start-Delay}) + \text{sum}_{\text{for_all_internal_flow_control_statements}}(\text{Duration})]$

⁽¹⁾ This value applies if duration property of the Statement is set to Minimum (default). If a fixed-duration has been set, then this is the Execution time is equal to that value.

Local While latency values:

Parameter	Description		Time (cycles)
Start-Latency	Minimum start-delay for the Statement		$5 + \#Register_Conditions$
Entry/Iteration latency	Minimum start-delay for first Statement inside the while loop	Minimum Duration	$8 + \#Register_Conditions + End-Latency_{Last-statement}$
		Fixed Duration	$8 + \#Register_Conditions$
End-Latency	Minimum start-delay for the next instruction outside the while loop	Minimum Duration	$8 + \#Register_Conditions + End-Latency_{Last-statement}$
		Fixed Duration	$8 + \#Register_Conditions$
Fixed-Duration	Minimum fixed-duration for Statement		$[\sum_{for_all_internal_statements}(Start-Delay) + \sum_{for_all_internal_flow_control_statements}(Duration) + End-Latency_{Last-statement}]^{(1)}$
Register Evaluation Latency	Time to evaluate the register condition		<ul style="list-style-type: none"> From start: 3 For each iteration: $-(2 + \#Register_Conditions)$

(1) If the branch is empty, then the duration is equal to the Entry-Latency of the branch.

Local If Statement Timing Tables

Local If timing value:

Execution time (cycles) (1)(2)
$\sum_{for_all_internal_statements}(Start-Delay) + \sum_{for_all_internal_flow_control_statements}(Duration) (3)$

- (1) The value provided here applies if the duration property of the Statement is set to Minimum (default). If a fixed-duration has been set, then the Execution time is equal to that value.
- (2) This value is only calculated for the branch that is executed, if there are multiple branches available.
- (3) If the branch is empty, the execution time becomes $\text{Entry-Latency}_{\text{branch}} - 1$.

The following table shows Local If latency values:

Parameter	Description		Minimum time (cycles)
Start-Latency	Contributes to the minimum-possible start-delay for the Statement		$5 + \#Register_Conditions_IfBranch$
Entry-latency	Contributes to the minimum-possible start-delay for first Statement in branch #		<ul style="list-style-type: none"> • If-Branch: 3 • Else-If-Branch: <ul style="list-style-type: none"> - For each else-if branch, we need to add: <ul style="list-style-type: none"> ◦ $6 + \#Register_Conditions_Else_If_Branch$ ◦ For the 1st else-if branch we will have: <ul style="list-style-type: none"> ◦ $2 + \#Register_Conditions_IfBranch + 7 + \#Register_Conditions_Else_If_Branch_1$ ◦ For the 2nd else-if branch we will have: <ul style="list-style-type: none"> ◦ $2 + \#Register_Conditions_IfBranch + 7 + \#Register_Conditions_Else_If_Branch_1 + 7 + \#Register_Conditions_Else_If_Branch_2$ ◦ and so on, so forth... • Else-Branch: Equal to last Else-If-Branch value
End-Latency	Contributes to the minimum-possible start-delay of the next Statement outside the if Statement	Matching Branches <i>disabled</i>	$3 + \max_{for_all_Branches}[End-Latency_{Last-statement}]$ (1)
		Matching Branches <i>enabled</i>	$3 + End-Latency_{Last-statement-of-longest-branch}$ (2) Where longest branch means the branch with longer execution time.
	Fixed-Duration	1	

<p>Fixed-Duration</p>	<p>Minimum fixed-duration for Statement</p>	<p>$2 + \max_{\text{for_all_branches}}[\text{Branch-Duration}]$ (3)</p> <p>Where Branch-Duration is calculated as follows:</p> <p>$[\text{sum}_{\text{for_all_internal_statements}}(\text{Start-Delay}) + \text{sum}_{\text{for_all_internal_flow_control_statements}}(\text{Duration}) + \text{End-Latency}_{\text{Last-statement}}]$ (4)</p>
<p>Register Evaluation Latency</p>	<p>Time to evaluate the register condition</p>	<p>3</p> <p><i>Then, for registers used in the condition of any else-if branch, we need to subtract:</i></p> <ul style="list-style-type: none"> - $6 + \#Register_Conditions_Else-If-Branch$ <p><i>Therefore, for the 1st else-if branch we will have:</i></p> <ul style="list-style-type: none"> - $3 - (6 + \#Register_Conditions_Else-If-Branch_1)$ <p><i>For the 2nd else-if branch we will have:</i></p> <ul style="list-style-type: none"> - $3 - (6 + \#Register_Conditions_Else-If-Branch_1) - (6 + \#Register_Conditions_Else-If-Branch_2)$ <p><i>and so on, so forth...</i></p>

(1) If the maximum end latency used in this equation corresponds to the if-branch, and the calculated latency is greater than 4, then the **End-latency** is the calculated value minus 1.

(2) If the longest branch is the if-branch, then the **End-latency** is the calculated value minus 1.

(3) If the maximum branch duration used in the equation corresponds to the if-branch, then the duration is the calculated value minus 1.

(4) If a branch is empty, then the branch duration is equal to the Entry-latency of the branch.

Local Wait-For-Time Statement Timing Tables

A Wait-for-time Statement blocks HVI execution in a Local Sequence until a specific amount of time passes. This amount of time is defined in a register that is specified as an argument in the Wait-for-

time Statement. The value of the register specifies the number of cycles to wait.

Local Wait-for-time Statement timing value:

Execution time (cycles)
<i>RegisterValue</i>

Local Wait-for-time Statement latency values:

Parameter	Time (cycles)
Start-Latency	1
End-Latency	1
Register Evaluation Latency	1

Local Wait-For-Event Statement Timing Tables

A Local Wait-for-Event Statement blocks HVI execution in a Local Sequence until an Event occurs. Events sources can be the Trigger IOs, or internal to the instrument (including User FPGA-Sandbox Events).

Local Wait-for-Event Statement timing values:

Event type	Execution time (cycles)	Fetch time (cycles)
Internal Event	$\text{MAX}(\text{Event_Arrival_Time}(1) + \text{Instrument_Event_Latency}(2) + 1, \text{Fetch_Time}) + 1$	3
Trigger IO	$\text{MAX}(\text{Event_Arrival_Time}(1) + \text{Instrument_Event_Latency}(2) + \text{Instrument_Event_Condition_Latency}(3), \text{Fetch_Time}) + 1$	$1 + \text{Instrument_Event_Condition_Latency}$

(1) *Event_arrival_time* is:

- Internal Events
 - $\text{Event_Arrival_Time} = \text{Internal_Event_Generation_Time} - \text{WaitForEvent_Start_Time}$

- External Events
 - $Event_Arrival_Time = Event_At_Module_Connector_Time - WaitForEvent_Start_Time$
 - The Event time can be measured at the front panel or PXIe backplane connector depending on the Event.

(2) **Instrument_Event_Latency** is the delay from the Event source until the Event state is available inside the HVI Engine. Events sources can be the Trigger IOs, or internal to the product (including User FPGA-Sandbox Events). It is an instrument and Event specific value. Refer to the instrument documentation for more information.

(3) **Instrument_Event_Condition_Latency** is the time needed for the condition evaluation to be executed once the Event has settled inside the HVI Engine. It is an instrument specific value. Refer to the instrument documentation for more information.

NOTE The *Event_Arrival_Time* can be a negative value if the Event enters the module before the Wait-For-Event instruction Start Time. A number of scenarios are shown in the diagrams below.

Local Wait-for-Event latency values:

Parameter	Time (cycles)
Start-Latency	0
End-Latency	1

Local Delay Statement Timing Tables

Local Delay Statement timing value:

Execution time (cycles)
Delay Specified

Local Delay latency values:

Parameter	Time (cycles)
Start-Latency	0
End-Latency	1

HVI Instruction Timing Tables

Trigger Write

Example Trigger write basic timing values:

Instruction	Execution time (cycles)	Fetch time (cycles)
TriggerWrite	$\text{Instrument_Trigger_Execution} + (\#\text{TriggerWriteGroups} - 1)$	$\#\text{TriggerWriteGroups}$

$\#\text{TriggerWriteGroups} = \text{ceil}[(\text{TriggerIOGroupsON} + \text{TriggerIOGroupsOFF})/2]$, where

- $\#\text{TriggerIOGroupsON}$ is the number of TriggerIOGroups that contain values set to ON.
- $\#\text{TriggerIOGroupsOFF}$ is the number of TriggerIOGroups that contain values set to OFF.

NOTE Trigger execution time is instrument-specific. For Trigger execution timing information, see your instrument documentation.

Action Execute

Example Action execute basic timing values:

Instruction	Execution time (cycles)	Fetch time (cycles)
ActionExecute	$\text{Instrument_Action_Execution} + \text{INT}[(\#\text{ActionGroups}-1) / 2]$	$1 + \text{INT}[(\#\text{ActionGroups} - 1) / 2]$

Where INT is the integer part of a decimal number, for instance $\text{INT}(1.0)=\text{INT}(1.5)=1$.

NOTE Action execution timing is instrument-specific. For Action execution timing information, see your instrument documentation.

Arithmetic Logic Unit Instructions

ALU instructions basic timing values:

Instruction	Execution time (cycles)	Fetch time (cycles)
Add	8	1
Subtract	8	1
Assign	5	1

FPGA User Sandbox Instructions

Example FPGA-User Sandbox operations basic timing values:

Instruction	Execution time (cycles)	Fetch time (cycles)
FpgaArrayRead	$2 * \text{Instrument_HVI_FPGA_Latency} + 4$	1
FpgaArrayRead (Address from HviRegister)	$2 * \text{Instrument_HVI_FPGA_Latency} + 6$	1
FpgaArrayWrite	$\text{Instrument_HVI_FPGA_Latency} + 2$	1
FpgaArrayWrite (Address or data from HviRegister)	$\text{Instrument_HVI_FPGA_Latency} + 4$	1
FpgaRegisterRead	$2 * \text{Instrument_HVI_FPGA_Latency} + 4$	1
FpgaRegisterWrite	$\text{Instrument_HVI_FPGA_Latency} + 2$	1
FpgaRegisterWrite (Address or data from HviRegister)	$\text{Instrument_HVI_FPGA_Latency} + 4$	1

NOTE

FPGA-User Sandbox timing is instrument-specific. For FPGA-User Sandbox timing information, see your instrument documentation.

FPGA-Instruction Statement

FPGA-Instruction Statement timing values:

Instruction	Execution Time (cycles)	Fetch Time (cycles)
FPGA-Instruction	$\text{Instrument_FpgaInstruction_Latency} + 2$	1

Instrument-Specific Local Instruction Statement Timing Values

Instrument-Specific HVI Instruction timing values:

Instruction	Execution Time (cycles)	Fetch Time (cycles)
Instrument-Specific HVI Instruction	$\text{Instrument_LocalInstruction_Latency} + 2$	1



This information is subject to change without notice.

© Keysight Technologies 2020-2023

Edition 2023_U0_00, June, 2023

Keysight Technologies, USA

KS2201-90000

www.keysight.com