

Content

1	Introduction.....	5
1.1.	IPv6.....	6
1.2.	6LoWPAN.....	6
1.3.	Organization of this Document.....	6
2	Functional Description.....	7
2.1.	Requirements Notation.....	7
2.2.	Terms.....	7
2.3.	Naming Conventions.....	8
2.4.	Library Architecture.....	8
2.5.	Operating Modes.....	9
2.5.1.	Device Mode.....	9
2.5.2.	Gateway Mode.....	10
2.5.3.	Sniffer Mode.....	10
2.6.	Operating System.....	11
2.6.1.	Initialization.....	11
2.6.2.	Normal Operation.....	11
2.6.3.	Power Modes.....	13
2.6.4.	Error Handling.....	14
2.7.	Firmware Version Information.....	14
2.7.1.	Vendor ID.....	15
2.7.2.	Product ID.....	15
2.7.3.	Major Firmware Version.....	15
2.7.4.	Minor Firmware Version.....	15
2.7.5.	Firmware Version Extension.....	15
2.7.6.	Library Version.....	15
2.8.	Addressing.....	15
2.8.1.	Address Types.....	16
2.8.2.	IPv6 Addresses.....	16
2.8.3.	IPv6 Address Auto-configuration.....	18
2.8.4.	Validation of Address Uniqueness.....	18
2.9.	Data Transmission and Reception.....	20
2.9.1.	User Datagram Protocol.....	20
2.9.2.	Data Transmission and Reception.....	20
2.9.3.	Address Resolution.....	22
2.9.4.	Recommendations.....	23
2.10.	Mesh Routing.....	23
2.10.1.	Multicast Traffic.....	24
2.10.2.	Unicast Traffic.....	24
2.10.3.	Mesh Routing Parameter Configuration Recommendations.....	24
2.11.	Network and Device Status.....	26

2.12.	Security.....	26
2.12.1.	Internet Protocol Security (IPSec).....	27
2.12.2.	Internet Key Exchange Version 2 (IKEv2)	28
2.12.3.	Recommendations	28
2.13.	Firmware Over-the-Air Updates.....	28
2.13.1.	Functional Description	29
2.13.2.	Firmware Constraints	30
2.14.	Memory Considerations.....	30
2.14.1.	Call Stack	31
2.14.2.	IDT Network Stack Dynamic RAM Requirements	31
2.14.3.	Using Dynamic Memory Allocation	32
2.15.	Supported Network Standards	33
3	Core-Library Reference	36
3.1.	Initialization	36
3.2.	Program Control	37
3.3.	Networking.....	41
3.3.1.	Address Management.....	41
3.3.2.	Socket and Datagram Handling	44
3.3.3.	Radio Parameters	49
3.3.4.	Gateway Mode Functions	51
3.3.5.	Miscellaneous	52
3.4.	Power Management	55
3.5.	Firmware Version Information	59
3.6.	Properties and Parameters.....	60
3.7.	Error Codes	61
4	UART Library Reference	62
4.1.	Symbol Reference	62
4.2.	Custom UART I/O Configuration	64
4.3.	Error Codes	65
5	GPIO Library Reference.....	66
5.1.	Symbol Reference	66
6	IPSec Library Reference	71
6.1.	Symbol Reference	71
6.2.	Error Codes	74
7	IKEv2 Library Reference	75
7.1.	Symbol Reference	75
7.2.	Library Parameters	76
8	Over-the-Air Update Library	77
8.1.	Library Reference	77
8.2.	Inclusion of the OTAU Library	79
8.3.	Error Codes	80

9	NetMA Libraries.....	81
9.1.	NetMA1 Library.....	81
9.1.1.	NetMA1 Library Symbol Reference	81
9.1.2.	Inclusion of the NetMA1 library	86
9.2.	NetMA2 Libraries.....	87
9.2.1.	NetMA2 Library Symbol Reference	87
9.2.2.	Inclusion of the NetMA2 Libraries	88
10	Accessing Microcontroller Resources.....	89
10.1.	Internal Microcontroller Configuration	89
10.2.	Backup Data Registers.....	89
10.3.	Interrupt Handlers.....	89
10.4.	Default I/O Configuration.....	92
11	Certification.....	96
11.1.	European R&TTE Directive Statements.....	96
11.2.	Federal Communication Commission Certification Statements	96
11.2.1.	Statements.....	96
11.2.2.	Requirements.....	96
11.2.3.	Accessing the FCC ID.....	97
11.3.	Supported Antennas.....	97
12	Alphabetical Lists of Symbols.....	98
12.1.	Functions and Function-Like Macros	98
12.2.	Data Types	99
12.3.	Variables and Constants	101
13	Related Documents.....	102
14	Glossary	103
15	Document Revision History.....	105

List of Figures

Figure 1.1	ZWIR451x Application Domain.....	5
Figure 2.1	Library Architecture.....	9
Figure 2.2	Application Interface into the Protocol Stack in Different Operating Modes.....	10
Figure 2.3	IPv6 Unicast Address Layout.....	17
Figure 2.4	IPv6 Multicast Address Layout	17
Figure 2.5	Resolving Address Conflicts in Local Networks	19
Figure 2.6	Working Principle of IPSec	27
Figure 2.7	Memory Layout of OTAU-Enabled Applications	29
Figure 2.8	Heap Memory Scattering	32
Figure 5.1	ZWIR_GPIO_ReadMultiple Result Alignment in ZWIR4512AC1 Devices	67
Figure 5.2	ZWIR_GPIO_ReadMultiple Result Alignment in ZWIR4512AC2 Devices	67
Figure 11.1	FCC Compliance Statement to be Printed on Equipment Incorporating ZWIR4512 Devices.....	97

List of Tables

Table 2.1	Naming Conventions Used in C-Code.....	8
Table 2.2	Event Processing Priority in the Main Event Loop.....	12
Table 2.3	Power Modes Overview.....	13
Table 2.4	Interrupts that Result in a System Reset.....	14
Table 2.5	Unicast Socket Examples.....	21
Table 2.6	Multicast Addressing Examples.....	22
Table 2.7	Stack Parameter Dynamic Memory Size Requiriements.....	31
Table 2.8	Supported RFCs and Limitations.....	33
Table 3.1	Configurable Stack Parameters and Their Default Values.....	60
Table 3.2	Error Codes Generated by the Core Library.....	61
Table 4.1	Error Codes Generated by the UART Libraries.....	65
Table 6.1	Error Codes Generated by the IPsec Libraries.....	74
Table 7.1	Overview of IKEv2 Library Parameters and Properties.....	76
Table 8.1	Error Codes Generated by the OTAU Library.....	80
Table 10.1	STM32 Interrupt Vector Table.....	90
Table 10.2	STM32 Default I/O Configuration.....	93

1 Introduction

This guide describes the usage of the 6LoWPAN application programming interface (API) for application development using ZWIR451x modules. These modules provide bidirectional IPv6 communication over an IEEE 802.15.4 wireless network. Using IPv6 as the network layer protocol allows easy integration of sensor or actor nodes into an existing Internet Protocol (IP) infrastructure without the need for additional hardware. Refer to the data sheet for the ZWIR451x module for detailed information about the module, including pin assignments.

Figure 1.1 ZWIR451x Application Domain

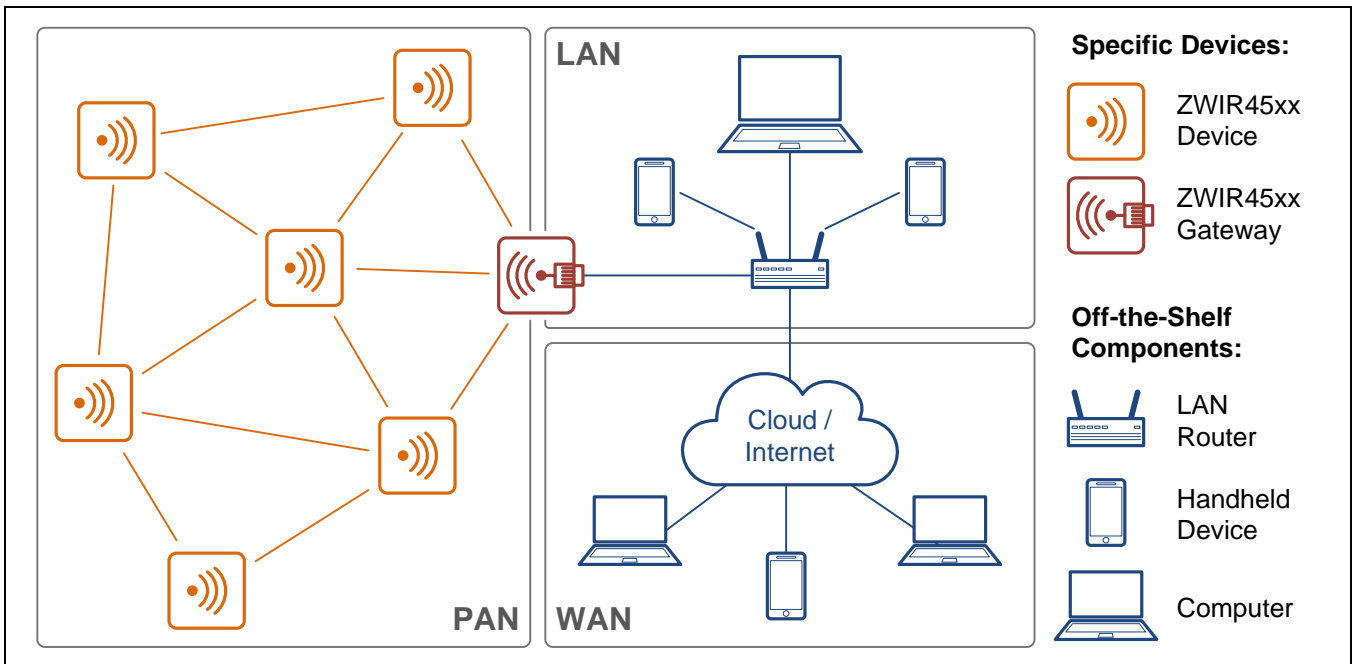


Figure 1.1 shows a typical network configuration. The Personal Area Network (PAN) is built from a set of various ZWIR451x modules. The network is connected via a border router to the local area network (LAN) and from there to a wide area network (WAN) such as the Internet. With this setup, each module can be accessed from anywhere in the world with just its unique IPv6 address.

The radio nodes are typically organized in a mesh topology. Routing of IP packets over this topology is handled by the software stack transparently for the user. The network allows dropping in new nodes or removing existing nodes without requiring manual reconfiguration. Routes to new nodes will be found automatically by the stack.

Application software runs on an STM32F103RC ARM® Cortex™ M3* microcontroller (MCU) on top of the ZWIR451x API. The MCU is clocked with up to 64 MHz and provides 256 kByte flash memory and 48 kByte RAM, which allows the implementation of memory and computationally intensive applications. The API provides functions to communicate with remote devices, access different I/O interfaces, and support power-saving modes.

* The ARM® and Cortex™ trademarks are owned by ARM, Ltd.

1.1. IPv6

IPv6 is the successor of the IPv4 protocol, which has been the major network protocol used for Internet communication over the past decades. One of the main advantages of IPv6 over IPv4 is its huge address area, which provides 2^{128} (about 3.4×10^{38}) unique addresses. This enormous address space allows assignment of a globally unique IP address to every imaginable device that could be connected to the Internet. Another advantage with respect to sensor networks is the stateless address auto-configuration mechanism, allowing nodes to obtain a unique local or global IP address without requiring a specific server or manual configuration.

The use of IPv6 makes it possible to connect sensor networks directly to the Internet. Basically this is possible with other network protocols, too, but those require a dedicated gateway that translates network addresses to IP addresses and vice versa. Usually this translation requires application knowledge and maintenance of the application state in the border router, and therefore changing the border router software might be required with each application update. The protocol gateway might also introduce an additional point of attack if secure communication between devices inside and outside of the PAN is required.

IDT's 6LoWPAN implementation supports IPSec, which is the mandated standard for secure communication over IPv6. The use of IPv6 throughout the whole network allows real end-to-end security.

1.2. 6LoWPAN

IPv6 has been designed for high bandwidth internet infrastructure, which does not put significant constraints on the underlying network protocols due to the vast amount of memory, computing power, and energy. In contrast, the IEEE 802.15.4 standard is intended for low data-rate communication of devices with very limited availability of all these resources. In order to make both standards work together, the 6LoWPAN standard (RFC 4944) has been developed by the Internet Engineering Task Force (IETF) to carry IP packets over IEEE 802.15.4 networks.

6LoWPAN adds an adaption layer between the link layer and network layer of the Open Systems Interconnection (OSI) reference model. This layer performs compression of IPv6 and higher layer headers as well as fragmentation to get large IPv6 packets transmitted over IEEE 802.15.4 networks. The 6LoWPAN layer is transparent for the user, and therefore on 6LoWPAN devices, the IPv6 protocol is used in exactly the same way as on native IPv6 devices. The presence of the 6LoWPAN adaption layer does not restrict IP functionality. The user of a 6LoWPAN system does not recognize the existence of the 6LoWPAN layer.

1.3. Organization of this Document

Sections 2 to 7 cover the API documentation, which is divided into two parts. The first part, covered by section 2 provides a functional description of the network stack. It explains the correlation of the different API functions and provides background information about stack internals. The second part is the function reference and is covered by sections 3 to 7. If familiar with the general stack functionality, the reader can simply use these sections to look up function signatures or basic usage information.

Section 8 explains how user applications can use the resources provided by the microcontroller and which resources are blocked by the operating system.

Terms set in **bold monospace** font can be clicked, activating a hyperlink to the section where a detailed definition of this term can be found.

2 Functional Description

The following subsections give a generic overview of the different functionalities of the firmware delivered with ZWIR451x modules. Background information is provided if required for proper use of the libraries. Usage recommendations are given for optimal performance in certain application configurations. A detailed description of the functions, types, and variables available for application programming is given in sections 3 through 7.

2.1. Requirements Notation

This document uses several words to indicate the requirements of IDT's 6LoWPAN stack implementation. The key words *MUST*, *MUST NOT*, *REQUIRED*, *SHALL*, *SHALL NOT*, *SHOULD*, *SHOULD NOT*, *RECOMMENDED*, *MAY* and *OPTIONAL*, set in italic small caps letters denote requirements as described below

MUST, *SHALL*, *REQUIRED* These words denote an absolute requirement of the implementation. Disregarding these requirements will cause erroneous function of the system.

MUST NOT, *SHALL NOT* These phrases mean that something is absolutely prohibited by the implementation. Disregarding these requirements will cause erroneous function of the system.

SHOULD, *RECOMMENDED* These words describe best practice, but there might be reasons to disregard it. Before ignoring this, the implications of ignoring the recommendation must be fully understood.

SHOULD NOT, *NOT RECOMMENDED* These words describe items that can impair proper behavior of the system when implemented. However, there might be reasons to choose to implement the item anyway. Implications of doing so must be fully understood.

MAY, *OPTIONAL* These words describe items which are optional. No misbehavior is to be expected when these items are ignored.

2.2. Terms

This document distinguishes between three types of functions: hooks, callbacks, and API functions. Basically, all three types are defined as normal functions in the programming language C, but they differ in the way that they are used.

API Functions are functions which are defined and implemented by IDT's 6LoWPAN stack. They provide a functionality that can be accessed by the user code. The declarations of API functions are provided in the header file belonging to the library in which the function is implemented.

Hooks are functions that provide the user the ability to extend the default behavior of the stack. They are called from the operating system (OS) to give the application the opportunity to implement custom features or reactions to events. The operating system provides a default implementation of the hook that is called if no custom hook is defined. The prototypes of all available hooks are defined in the header file belonging to the library in which the default implementation is located. `ZWIR_AppInitNetwork` and `ZWIR_Error` are examples of hooks.

Callbacks are also called from the operating system, but they must be registered explicitly at the OS. The function may have a custom name, but the signature must be matching. Callback functions are registered at the OS using API functions. One example for a function expecting a callback is `ZWIR_OpenSocket`. In contrast to hooks, callbacks do not have default implementations. For each callback function, there is a type declaration declaring how the signature of the user function should look.

2.3. Naming Conventions

For better readability of the code, all user accessible functions and types of the API comply with a set of naming conventions. Each identifier that is an element of the API is prefixed with “ZWIR_.” Function, variable, function argument and type identifiers are defined using “CamelCase” style. This means that each single word of a multiple word identifier starts with a capital letter. The remaining letters of the word are lower case. Preprocessor macros are defined using all capital letters.

Different style rules apply to functions, variables, and types definitions. Function-name and type-name identifiers start with a capital letter in the first word, while variable identifiers start with a lower case letter in the first word. Type names have an additional “_t” suffix. Variable names and function arguments are not differentiated in the naming conventions.

Table 2.1 Naming Conventions Used in C-Code

Identifier Type	Style
variableName, functionArgument	First word starts with lower case, all other words with capital letters.
FunctionName	All words start with capital letters (“CamelCase”).
TypeName_t	All words start with capital letters and name ends with “_t” suffix.
PREPROCESSOR_MACRO	All letters are capitalized.

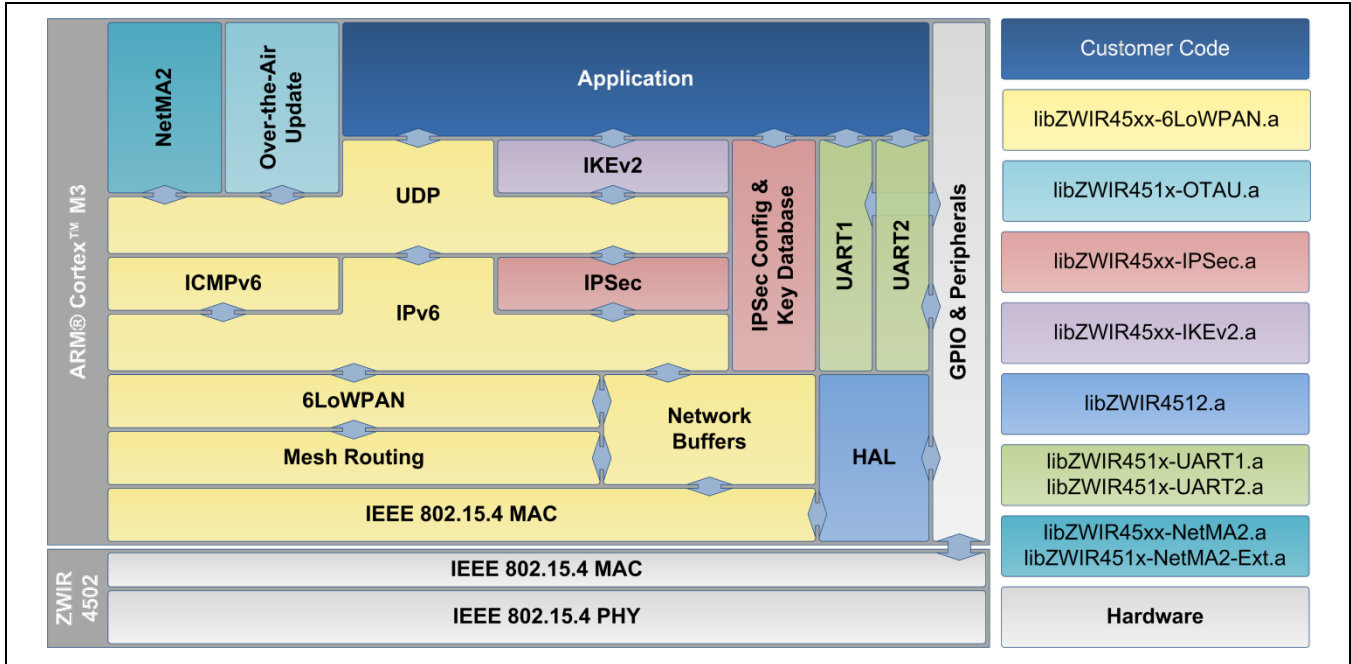
2.4. Library Architecture

ZWIR451x modules are freely programmable by means of an API that is implemented in a set of libraries. The libraries provide different functionality and can be linked into the user program. The use of the core library is mandatory, as it provides the operating system and all generic communication functionality. All other libraries are optional and can be linked depending on the requirements of the target application. Each library exposes a set of functions and types that are required to implement the desired functionality. The library architecture is depicted in Figure 2.1.

To make programming as easy as possible, the libraries make use of an event and command approach wherever possible. Using this approach, application code is not required to poll for data on the different interfaces. Instead, newly available data is passed to user-defined callback functions automatically. Timer hooks and callbacks are available, and they are automatically executed periodically or after expiration of a user-defined time interval.

Linking the library without any additional code will result in a valid binary that can be programmed on a radio module. Such binaries will not provide user-specific functionality. However, the nodes relay packets in mesh networks and respond to ping requests. In order to add functionality, several functions that have empty default implementations can be defined by the user.

Figure 2.1 Library Architecture



2.5. Operating Modes

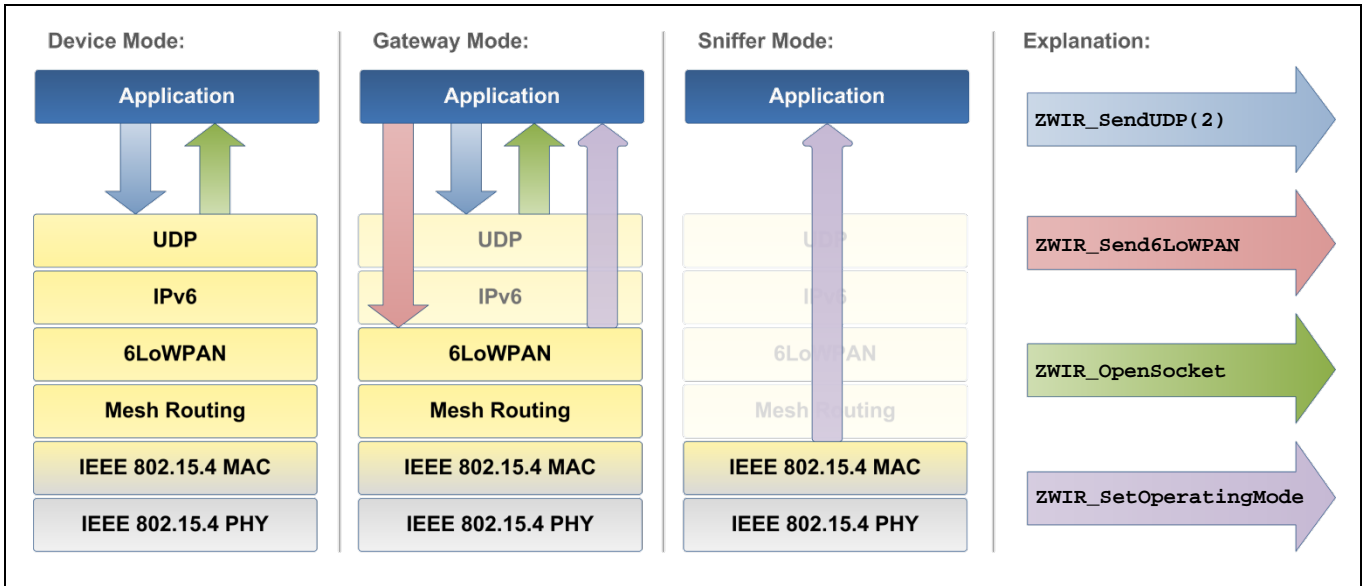
The API provides three operating modes: Device Mode, Gateway Mode and Sniffer Mode. The modes differ in how many of the protocol layers are processed by the network stack. All other API functionality remains the same. Setting the operating mode of a node must be done before any initialization of the API and the hardware. For this purpose, the `ZWIR_setOperatingMode` function is provided. Figure 2.2 shows how application code interfaces into the network stack in different operating modes. A description of the three different modes is provided in the next subsections.

2.5.1. Device Mode

The Device Mode is preconfigured since this is the most commonly used mode for ZWIR451x modules. Each node with sensing or acting functionality should use this operating mode. Full protocol processing is performed for incoming and outgoing data. This means that all header information is removed from incoming User Datagram Protocol (UDP) packets and only payload data is passed to the application. Accordingly, the application only has to provide payload data that should be sent over the network. The stack automatically adds all necessary header information.

Device-configured devices behave as normal IPv6 devices. Therefore address auto-configuration and neighbor-discovery is performed as defined by the IPv6 standard. Data are sent and received over UDP sockets. The functions `ZWIR_sendUDP` and `ZWIR_sendUDP2` serve as an interface to the network stack. Incoming data is passed to an application callback that must be registered when a socket is opened using `ZWIR_OpenSocket`.

Figure 2.2 Application Interface into the Protocol Stack in Different Operating Modes



2.5.2. Gateway Mode

The Gateway Mode is intended for use with modules that should work as protocol gateways. Protocol gateways change the physical media used for IPv6 packet transmission. This enables the integration of 6LoWPAN networks into Ethernet-based IPv6 networks for example.

In contrast to the Device Mode, not all network layers are processed in Gateway Mode. For any IPv6 packet that is received via the air interface, only the 6LoWPAN-specific modifications of the headers are removed, resulting in a packet containing all IPv6 and higher layer headers. This packet is passed to the receive callback function. Accordingly, all data that need to be sent over the network are assumed to have valid IPv6 and higher layer headers. Only 6LoWPAN-specific modifications will be applied to outgoing packets.

Gateway-configured devices do not perform address auto-configuration and neighbor-discovery as defined by the IPv6 standard. Moreover no router solicitation and router advertisement messages are generated automatically.

To enter the Gateway Mode, `ZWIR_SetOperatingMode` must be called from `ZWIR_AppInitHardware`. It is not possible to call `ZWIR_SetOperatingMode` from any other location in the code. `ZWIR_SetOperatingMode` accepts a callback function that is called upon reception of data in the gateway. Sending data is accomplished using the function `ZWIR_Send6LoWPAN`.

2.5.3. Sniffer Mode

The Sniffer Mode is provided to allow observation of raw network traffic. No protocol processing is performed. Thus the data passed to the application layer includes all header information. In contrast to the two other operating modes, all packets received over the air interface are passed to the application, regardless of the address to which the packet has been sent. This also includes MAC layer packets.

Sniffer Mode is useful for debugging purposes. It can be used to find out which devices in the network are transferring packets and which are not. Sniffer Mode devices do not generate any network traffic—not autonomously or user triggered. That is why there is only an interface from the stack to the application code, but not vice versa.

To enter Sniffer Mode, call `ZWIR_SetOperatingMode (ZWIR_omSniffer, YourCallbackFunction)` from `ZWIR_AppInitHardware`. The functions `ZWIR_Send6LoWPAN`, `ZWIR_SendUDP` and `ZWIR_SendUDP2` do not function in this mode. However, it is still possible to change the physical channel and the modulation scheme of the transceiver by calling `ZWIR_SetChannel` and `ZWIR_SetModulation`.

2.6. Operating System

The operating system is very light-weight and does not provide multi-threading. This means that any user-defined function that is called from the operating system is completely executed before control is passed back to the operating system. Therefore, the user is required to write cooperative code. Users must be aware that functions requiring long execution time will block the operating system kernel and might cause the kernel to miss incoming data, regardless of whether they are received over the air or any wired interface.

2.6.1. Initialization

During the operating system initialization phase, the different libraries and the MCU peripherals required for system operation are initialized. Startup initialization is done in two phases, each with its own hook for user application code. During the first phase, the internal clocks and the peripherals used by the stack are initialized and the random number generator is seeded. Also peripherals required by certain libraries are initialized if the corresponding library is linked into the project. After this, the `ZWIR_AppInitHardware` hook is called if present, enabling application code to initialize any additional hardware. The application can initialize its I/Os and peripherals in this function. `ZWIR_SetOperatingMode` must be called from here if the Gateway Mode is required. Sending data over the network or initializing network sockets is not possible from here, as the network stack is not initialized. However, functions controlling the physical parameters of the network (e.g., output power or physical channel) *SHOULD* be called from here.

During the second phase, the transceiver and the network stack are initialized. If the Normal Mode is selected, duplicate address detection (DAD) is also started and router information is solicited. DAD checks if the address given to the module is unique on the link. After finishing network initialization, `ZWIR_AppInitNetwork` is called. Application code can do its remaining initialization tasks such as setting up sockets at this point. Since DAD and router solicitation are started before the call to `ZWIR_AppInitNetwork`, it is recommended that physical parameters of the network are set up first in `ZWIR_AppInitHardware`. This ensures that DAD and RS are performed on the correct channel with correct settings.

2.6.2. Normal Operation

During normal operation, the operating system collects events from the different peripherals and the application and handles them according to their priority. Event processing priorities are fixed and cannot be changed. Events are processed highest priority first; the lowest number represents the highest priority. Table 2.2 lists all events with their priorities and triggered actions.

Table 2.2 Event Processing Priority in the Main Event Loop

Priority	Event	Triggered By	Effect
0	Application Event 0	Application Code	Call user-defined callback function
1	Transceiver Event	Transceiver Interrupt Request	Process transceiver request
2	Application Event 1	Application Code	Call user-defined callback function
3	Callback Timer Expired	SysTick Controlled Software Timer	Call user-defined callback function
4	Sleep Requested	Software	Sleep for the requested time
5	Received Data on UART1	UART1 Interrupt	Call user-defined callback function
6	Application Event 2	Application Code	Call user-defined callback function
7	10ms Timer Expired	SysTick Controlled Software Timer	Call <code>ZWIR_Main10ms</code>
8	100ms Timer Expired	SysTick Controlled Software Timer	Call <code>ZWIR_Main100ms</code>
9	1000ms Timer Expired	SysTick Controlled Software Timer	Call <code>ZWIR_Main1000ms</code>
10	Application Event 3	Application Code	Call user-defined callback function
11	Received Data on UART2	UART2 Interrupt	Call user-defined callback function
12	Sending Data Failed due to Resource Conflict	Network Stack	Retry sending
13	Application Event 4	Application Code	Call user-defined callback function

The operating system provides five application event handlers that can be used to process application events in the context of the operating system scheduler. Application event handlers *SHOULD* be used to react to asynchronous events requiring computationally intensive processing. Interrupts are a typical example for such events. If an interrupt occurs, the interrupt service routine (ISR) can trigger an event and delay the processing to an appropriate time. This ensures that multiple asynchronous events are handled in the order of their priority, without blocking interrupts.

Application events are triggered by calling `ZWIR_TriggerAppEvent` with the corresponding event number (0 through 4). When the OS scheduler reaches the user-triggered event, an application callback function is executed. Multiple calls to this function before the corresponding application callback is invoked will not cause multiple invocations of the application callback.

For each application event, an event handler callback function must be registered using `ZWIR_RegisterAppEventHandler`. If no event handler is registered for a certain event, triggering this event has no effect. In order to change an event handler, `ZWIR_RegisterAppEventHandler` must be called again with the new handler. Unregistering event handlers can be performed by calling the registration function with a NULL callback argument.

2.6.3. Power Modes

The stack supports different modes to reduce the power consumption of the device. In Active Mode, all module features are available. The Sleep, Stop, and Standby Modes reduce the power consumption by disabling different module functionalities. Each of the power-saving modes affects the behavior of the MCU and the transceiver and supports different wake-up conditions.

Table 2.3 Power Modes Overview

Mode	Wakeup		Clock		Context ¹⁾	I/O	Transceiver
	Source	Time	MCU Core	Peripherals			
Active			On	On ²⁾	Retained	As configured	On ³⁾
Sleep	Any IRQ	1.8 μ s	Off	On ²⁾	Retained	As configured	Off ⁴⁾
Stop	RTC IRQ External IRQ	5.4 μ s	Off	Off	Retained	As configured	Off ⁴⁾
Standby	RTC IRQ Wakeup Pin	50 μ s	Off	Off	Lost	Analog input	Off

¹⁾ Refers to the status of the RAM and peripheral register contents after wakeup – the backup registers of the MCU are always available.

²⁾ Clock is enabled for all peripherals that have been enabled by application code and all peripherals that are used by the library.

³⁾ Can be powered off by application code.

⁴⁾ Remains on if peripheral/transceiver is selected as wakeup source.

Active Mode is entered automatically after startup. In this mode, the MCU core and all peripherals used by the application are running and all functionality is available. The transceiver is typically on but can be switched off explicitly by a call to `ZWIR_TransceiverOff`. This mode has the highest power consumption.

In Sleep Mode, the MCU core is disabled but the MCU peripherals are still functioning if required. The transceiver can be switched on or off. Memory contents and I/O settings remain in the state that was active at the activation of the Sleep Mode. Waking up from Sleep Mode is possible on any MCU interrupt. After the wakeup event, the stack continues execution at the position it had been stopped. The power consumption in Sleep Mode is slightly reduced compared to Active Mode. If more significant reduction of the power consumption is required, the Stop or Standby Modes should be considered.

Stop Mode provides significant reduction of power consumption while still providing a short wakeup time and context saving. Depending on the application's requirements, the transceiver could remain enabled to wake up the module when a packet comes in (set the transceiver as the wakeup source). By default, the transceiver is disabled in Stop-Mode. The MCU core and all peripherals of the MCU are disabled in Stop Mode. Wakeup is only possible by the built-in real-time clock (RTC) or an external interrupt, triggered at any GPIO line. For that, the external interrupt must be configured appropriately.

Standby Mode is the lowest power mode. In this mode, the MCU is powered off and the transceiver is on standby. Only the MCU's internal RTC is running, serving as a wakeup source. Additionally, the external wakeup pin can be used to wake up the module. After wakeup, the memory contents of the MCU are lost and must be reinitialized the same as after normal power-on.

Any of the low power modes is entered by calling the function `ZWIR_PowerDown`. It can be chosen whether power-down is delayed until all pending events are processed or not. If delayed power down is chosen, the power-down procedure can be aborted by a call to `ZWIR_AbortPowerDown`. The wakeup sources for the different power modes are configured by `ZWIR_SetWakeupSource`.

2.6.4. Error Handling

The stack performs error handling in two different ways. The first is simply to reset the chip if an unrecoverable MCU exception occurs that caused an interrupt. For errors that are not caused by MCU exceptions, the stack provides a default handling, which may be overwritten by the application code.

The error handlers performing a system reset are triggered by one of the interrupts listed in Table 2.4. The reason for resetting the whole system is that in the case of normal operation, none of the listed interrupts should appear. However, if different behavior is desired, it is possible to overwrite the default implementation by providing the user's own interrupt service routines. See section 10.3 for details.

Table 2.4 *Interrupts that Result in a System Reset*

Resetting Interrupts
Non-maskable Interrupt
Hard Fault
Memory Management
Bus Fault
Usage Fault
Programmable Voltage Detector

In the case of a recoverable error, the `ZWIR_Error` hook is called by the operating system. The error number is passed as a function argument. In order to provide custom error handling, the application *MUST* provide an implementation of `ZWIR_Error`. The return value of the function determines whether the error has been handled by the application (return `true`) or if the default handler will be executed (return `false`).

2.7. Firmware Version Information

The ZWIR451x API provides the capability of including firmware version information in the stack. This information can be requested remotely afterwards and is required by the Firmware Over-the-Air Update Library. The complete firmware version consists of the Vendor ID, the Firmware ID, the Major Firmware Version, the Minor Firmware Version, and the Firmware Version Extension. These components are defined in the application code using global variables. The role of the different components is explained in the following subsections.

In addition to the firmware version information mentioned above, the stack provides additional version information for the library to which the application was linked. This version information consists of Major Stack Version, Minor Stack Version, and Stack Version Extension field.

2.7.1. Vendor ID

The Vendor ID is a 32-bit number that identifies the company that developed the device firmware. A Vendor ID must be requested from IDT. Each company must obtain its own Vendor ID before placing products on the market. The Vendor ID is set using the global variable `ZWIR_vendorID`. If this variable is not set, the firmware will use the Vendor ID `E966HEX`, which is reserved for experimental purposes and must not be used for production firmware.

2.7.2. Product ID

The Product ID is a 16-bit number identifying the product firmware. It is especially important for the Over-the-Air Update functionality but could also serve for remote identification of the device type. Refer to the *ZWIR451x Application Note – Enabling Firmware Over-the-Air Updates* for more information about the role of the Product ID in IDT's Over-the-Air Update Library.

The Product ID is set by defining the global variable `ZWIR_productID`. If this variable is not defined, the value will be read as zero.

2.7.3. Major Firmware Version

The Major Firmware Version is a version information field that is freely usable for application purposes. It is set by defining the global variable `ZWIR_firmwareMajorVersion`. If this variable is not defined, the value will be read as zero.

2.7.4. Minor Firmware Version

The Minor Firmware Version is a version information field that is freely usable for application purposes. It is set by defining the global variable `ZWIR_firmwareMinorVersion`. If this variable is not defined, the value will be read as zero.

2.7.5. Firmware Version Extension

The Firmware Version Extension is a version information field that is freely usable for application purposes. It is set by defining the global variable `ZWIR_firmwareVersionExtension`. If this variable is not defined, the value will be read as zero.

2.7.6. Library Version

IDT's firmware stack libraries have their own version information included. This information is compiled into the binary libraries and can be requested by the application code using the function `ZWIR_GetRevision`. Like the firmware version, the library version consists of the major and minor versions as well as extension information.

2.8. Addressing

Each module has three types of addresses: a PAN identifier, link layer address, and network layer address. This section describes the different address types and explains how they are used in the stack.

2.8.1. Address Types

The **PAN identifier** (PANId) is a 16-bit-wide number carrying an identifier of the network. Each device in the same network must have the same PANId. Nodes with different PANIds cannot communicate. A default PANId is preprogrammed in the network stack. The current PANId can be requested or changed using the functions `ZWIR_GetPANId` and `ZWIR_SetPANId`, respectively. The default value is `ACCAHEX`.

The **link layer address** is also referred to as the **MAC address** or **PAN address**. This address is used by the lower communication layers and does not need to be handled directly by the user. The link layer address must be unique in the network. Each ZWIR451x module has a predefined, hardware programmed address that is globally unique. The PAN address is 64-bits-wide. It can be requested and changed by the functions `ZWIR_GetPANAddress` and `ZWIR_SetPANAddress`. Changing the PAN address is not recommended as this could cause problems as described in section 2.8.4.

The third address type is the **network layer address**, which is equivalent to the **IPv6 Address**. These addresses are 128-bit-wide. They are used by the application to determine the destination that packets should be sent to or the source packets should be received from. Each device needs at least one IPv6 address to be reachable. However, multiple addresses can be assigned to each node. IPv6 addresses assigned to a node must be unique on the network. However, typically users do not need to handle IPv6 address assignment. IPv6 provides a mechanism that performs automatic address configuration. This mechanism is explained in section 2.8.3.

2.8.2. IPv6 Addresses

IPv6 addresses are 128-bit and therefore 16-bytes wide. As it would be impractical to use the byte-wise notation known from IPv4, IPv6 introduces a new notation. IPv6 addresses are represented by eight 16-bit hexadecimal segments that are separated by colons. An example for such address is

`2001:0db8:0000:0000:1b00:0000:0ae8:52f1`

The leading zeros of segments can be omitted as they do not carry information. Furthermore the IPv6 notation allows omitting a sequence of zero segments and representing it as double colon. With these rules, the above address can be written as

`2001:db8::1b00:0:ae8:52f1` or `2001:db8:0:0:1b00::ae8:52f1`

However, replacing multiple zero segments is not allowed, so the following address is invalid:


`2001:db8::1b00::ae8:52f1`

An IPv6 address consists of two components: a prefix and an interface identifier. The prefix specifies the network that a device is part of; the interface identifier specifies the interface of a device. A node with multiple network interfaces has multiple interface identifiers. The size of the prefix varies for different address types. In the IPv6 address notation, the prefix length can be appended to the address with a slash followed by the number of prefix bits. For example, the notation `2001:db8::/64` represents a network containing the addresses from `2001:db8::` to `2001:db8::ffff:ffff:ffff:ffff`.

IPv6 supports three kinds of addressing methodologies: unicast addressing, multicast addressing, and anycast addressing. Unicast and multicast addresses differ in how the prefix is formed. Anycast addresses can be used as target addresses and are handled like unicast addresses.

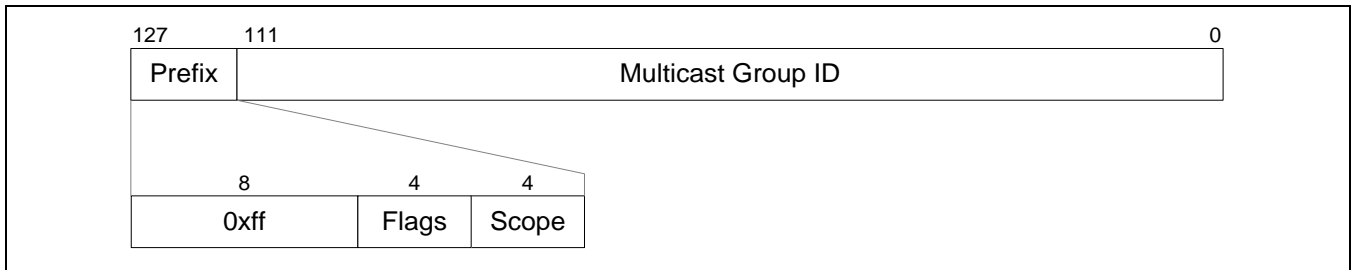
Unicast addresses are shown in Figure 2.3 and use 64 bits each for the prefix and the interface identifier. Unicast addresses exactly identify one single interface in a network. The prefix of the address determines the scope of the unicast address. If the prefix equals fe80::/64 this is a link-local unicast address. Link local addresses are valid only on the single link a node is connected to. The prefix of global unicast addresses is received via address auto-configuration from a router that is connected to the Internet. There are additional prefix configurations with limited scope that are not covered by this documentation.

Figure 2.3 IPv6 Unicast Address Layout



Multicast addressing allows sending out a single packet to multiple receivers. For this purpose, IPv6 provides multicast addresses. A multicast address can only be used as the destination address, never as the source address of a packet. The layout of a multicast address is shown in Figure 2.4. Multicast addresses have a 16-bit prefix with the most significant 8 bits set to FF_{HEX}, followed by two 4-bit fields for flags and the scope of the multicast packet. The remaining bits specify the multicast group ID.

Figure 2.4 IPv6 Multicast Address Layout



In this document, it is assumed that the flags field is always either 0000_{BIN} or 0001_{BIN}. 0000_{BIN} specifies that the multicast address is a well-known address. 0001_{BIN} marks the address as a temporarily assigned address that is not specified by Internet standards. These addresses must be used for custom multicast addressing. The scope field is always assumed to be 0010_{BIN}, representing the link-local scope. Other scopes are usable but must be supported by routers.

Two specific addresses should be noted as these are used very often. More information about their use can be found in section 2.9.2.

1. The unspecified address ::
All segments of this address are zero. It is used by receivers to listen to any sender. This address must never be used as destination address.
2. The link-local all nodes multicast address ff02::1
Packets sent to this address are received by all nodes in the network, so this multicast address is equivalent to broadcasting.

For more detailed information about IPv6 addressing refer to [RFC 4291](#) – “IP Version 6 Addressing Architecture.”

2.8.3. IPv6 Address Auto-configuration

IPv6 provides a stateless address auto-configuration mechanism. This mechanism allows the configuration of node addresses from information that is statically available on the node and information provided by routers. Router information is only required if global communication is required. Addresses for link-local communication can be derived from the link-layer address. This removes the need for manual configuration of addresses or dynamic host configuration protocol (DHCP) servers in the network.

The local information used for auto-configuration is the interface EUI-64 address. The EUI-64 address is a factory programmed link-layer address that IDT guarantees to be unique for each module. The EUI-64 address is often referred to as the MAC address. During network initialization, each node generates a unique link-local IPv6 address by putting the prefix fe80::/64 in front of the EUI-64 address with bit 1 of the most significant EUI-64 byte inverted. Assuming a link-layer address of 00:11:7d:00:12:34:56:78, the generated link-local IPv6 address would be fe80::211:7d00:1234:5678.

It is possible to switch off the link-local address generation by setting the stack parameter `ZWIR_spDoAddressAutoConfiguration` to zero before the stack initialization.

In addition to link-local address generation, nodes request router information during startup seeking a global prefix for building a globally valid address. Those requests are called router solicitations. Routers present on the link will respond to router solicitation messages of the node with router advertisements containing global prefix information. Taking this prefix and the EUI-64 address of the node, a global address is generated in the same way as for the link-local address. Router solicitation is done automatically during the startup phase.

If there is no router on the link, no global address will be assigned and only link-local communication is possible. In this case, the router solicitation messages can be suppressed by setting the stack parameter `ZWIR_spDoRouterSolicitation` to zero.


A host cannot rely on a generated address being unique, as there might be manually configured EUI-64 addresses on its link. Therefore, it must perform “duplicated address detection” (DAD) to be sure the generated address is unique. Duplicate address detection is mandatory for each address being attached to a node and is performed automatically by the network stack. It is described in more detail in the following section. Each address being assigned to an interface is subject to duplicate address detection. Addresses are not valid until duplicate address detection (DAD) is completed. Devices are not able to send or receive packets using a unicast address that has not been validated to be unique. After device startup (and therefore after assignment of the link-local unicast address), the network stack calls the hook `ZWIR_AppInitNetworkDone`, signaling that DAD on the link-local address has been completed and the address can be used. Applications should use this hook to send out initial packets.

2.8.4. Validation of Address Uniqueness

After a node has configured its own address, it performs Duplicate Address Detection (DAD) to check if the newly configured IPv6 address is unique on the link. For this purpose, the node starts to send neighbor solicitation (NS) messages to the address to be checked (to its own address). If another node replies to one of those messages or if another node also sends neighbor solicitation messages to this address, the assigned address is not unique and must not be used. In this case, the error handler hook `ZWIR_Error` is called with the error code `ZWIR_eDADFailed`. It is up to users to provide their own error handling mechanism for such cases.

The default implementation provided in the library only removes the failing address from the interface. If the failing address was the only address of the module, the module will not be reachable.

Figure 2.5 Resolving Address Conflicts in Local Networks



Application code can try to resolve the address conflict. One possible solution is to manually change the link-layer address of the node, using random numbers or a dedicated algorithm. **ZWIR_SetPANAddress** must be called with the new address and the network initialization must be restarted. This is done by calling **ZWIR_ResetNetwork**. The procedure can be repeated for an arbitrary number of times until a unique address is found.

Note that a duplicate address problem should not appear if each module in the network uses the factory programmed link-layer address. In this case, the link-layer address is guaranteed to be globally unique. Therefore, using the user's own addresses is not recommended. For some applications, the user knows that there are no duplicate addresses in the network. In such cases, the duplicate address detection mechanism can be disabled by setting the stack parameter **ZWIR_spDoDuplicateAddressDetection** to zero. This has the positive side effect of the immediate ability to send and receive packets using the user's own IPv6 address(es). Furthermore, less traffic is generated on the network.

2.9. Data Transmission and Reception

Data are transmitted using the User Datagram Protocol (UDP). If a destination node is not directly reachable from the source node, packets are routed over intermediate nodes automatically. Route setup is done transparently for the user. The following subsections describe the different aspects of data transmission and reception.

2.9.1. User Datagram Protocol

The User Datagram Protocol is used for data communication. UDP is a connectionless and lightweight protocol with the benefit of minimal communication and processing overhead. No connection has to be created, and no network traffic is required before data transmission between nodes can be started. Instead, communication is possible immediately. However, UDP does not guarantee that packets that have been sent are reaching the receiver. It is also possible that a single UDP packet is received multiple times. Furthermore, it is not guaranteed that the receiving order of packets at the destination is the same as the sending order at the source. This must be considered by the application programmer.

UDP uses the concept of ports to distinguish different data streams to a node. 65535 different ports can be distinguished in UDP. A port can be seen as the address of a service running on a node. Depending on the destination port of a packet, the network stack decides to which service the packet is routed on the receiver node. In IDT's 6LoWPAN stack, services that provide the callback functions to which network packets are passed are running in the application code. Each service has its own callback function.

2.9.2. Data Transmission and Reception

Data transmission is requested by calling `ZWIR_SendUDP` or `ZWIR_SendUDP2`. Both functions send a single UDP packet to a remote host. `ZWIR_SendUDP2` accepts the address and port of the remote device as a parameter, while `ZWIR_SendUDP` requires a socket handle specifying the destination parameters. For reception of data, a socket is required as well.

A socket is an object that stores the address of a remote device and the remote and local UDP ports used for communication. It can be seen as an endpoint of a unidirectional or bidirectional communication flow. Additionally, it is possible to specify a callback that is called when data is received over the socket. Sockets are opened and closed using the API functions `ZWIR_OpenSocket` and `ZWIR_CloseSocket`. The maximum number of sockets that can be open in parallel is defined by the stack parameter `ZWIR_spMaxSocketCount`.

Four parameters must be provided when a socket is opened:

- IPv6 address of the remote communication endpoint: This is the address that data should be sent to and/or received from. Data reception is only possible if the remote address is a unicast address or the unspecified address. If a multicast address is provided, only data transmission is possible.
- Remote UDP port: This is the UDP port that data are sent to. For reception of data, this port is ignored.
- Local UDP port: This is the UDP port that data are received on. Only packets that are sent to this port will be handled in the callback function. If this number is 0, no data is received.
- Receive callback function: This is a pointer to a function that is called when data from a remote device is received. If no data should be received, this pointer can be set to NULL.

The choice of whether `ZWIR_SendUDP` or `ZWIR_SendUDP2` should be used for communication depends on the characteristics of the network traffic between the communicating devices. `ZWIR_SendUDP2` is intended to send a

few packets to a remote device without expecting a response from the target device. The function accepts the remote address and UDP port together with the data to be sent. Internally the function will open a temporary socket that is immediately closed after sending out the packet, so a slight overhead is added. `ZWIR_SendUDP2` functions even if the maximum number of sockets is open. `ZWIR_SendUDP` must be used in cases where responses are expected from the remote device or data must be transmitted frequently.

The following subsections will explain unicast and multicast communication in more detail and give examples of how to use the IPv6 addresses and ports appropriately.

2.9.2.1. Unicast

Traffic that has only a single destination node is called unicast traffic. In order to send data unicast, the sender must open a socket with the remote address set to the IPv6 address of the intended receiver. The receiver must open a socket with the remote address field set to the sender's IPv6 address or to the unspecified address. The sender socket remote port field must match the receiver socket local port field in both cases. Table 2.5 shows some example socket configurations and notes whether communication is possible or not.

Table 2.5 Unicast Socket Examples

A	B	C	D	E
fe80::1:1:1:1	fe80::2:2:2:2	fe80::2:2:2:2	fe80::2:2:2:2	fe80::3:3:3:3
Rem. Addr.: fe80::2:2:2:2	Rem. Addr.: fe80::1:1:1:1	Rem. Addr.: fe80::1:1:1:1	Rem. Addr.: ::	Rem. Addr.: fe80::1:1:1:1
Rem. Port: 55555	Rem. Port: 44444	Rem. Port: 44444	Rem. Port: 44444	Rem. Port: 44444
Local Port: 44444	Local Port: 55555	Local Port: 33333	Local Port: 55555	Local Port: 55555
Sender	Recipients			
A	B receives packet. (Remote address and remote port of A match interface address and local port of B.) C does not receive packet. (Remote port of A does not match local port of C.) D receives packet. (Remote address of D matches all addresses; local port of D matches remote port of A.) E does not receive packet. (Remote address of A does not match interface address of E.)			
B	A receives packet. (Remote address and remote port of B match interface address and local port of A.) No other socket receives packet. (Interface addresses do not match remote address field of B; local ports do not match remote port of B.)			
C	A receives packet. (Remote address and remote port of C match interface address and local port of A.) No other socket receives packet. (Interface addresses do not match remote address field of C.)			
D	No socket receives packet. (Sending is not possible with an unspecified address as the destination.)			
E	No socket receives packet. (Remote addresses of sockets A, B, and C do not match interface address of E; local port of D does not match remote port of E.)			

2.9.2.2. Multicast

Multicast is used to send data to multiple nodes at the same time. For a multicast transmission, the sender must open a socket with the remote address set to a multicast IPv6 address. The semantics of ports is the same as for unicast communication. The receiver must open a socket with the remote address set to the IPv6 address of the sender or to the unspecified address. Note that a socket with a multicast remote address cannot be used for data reception.

IDT's implementation of the IPv6 multicast feature does not support explicit assignment of multicast groups to single nodes. Instead, if a packet is received that was sent to a temporary multicast address, the hook function `ZWIR_CheckMulticastGroup` is called by the network stack. This function must be implemented by the application code in order to use the multicast feature. The application can check the multicast group of the destination address and decide if it is part of it. This mechanism allows very flexible and application-tailored multicast addressing schemes. If the application does not provide the `ZWIR_CheckMulticastGroup`, temporary multicast addresses are rejected by the stack.

Table 2.6 Multicast Addressing Examples

A	B	C	D	E
fe80::1:1:1:1	fe80::1:1:1:1	fe80::1:1:1:1	fe80::x:x:x:x	fe80::x:x:x:x
Rem. Addr.: ff02::1	Rem. Addr.: ff02:x:x:x:x:x:x	Rem. Addr.: ff12:x:x:x:x:x:x	Rem. Addr.: fe80::1:1:1:1	Rem. Addr.: ::
Rem. Port: 55555	Rem. Port: 55555	Rem. Port: 55555	Rem. Port: x	Rem. Port: x
Local Port 44444	Local Port 44444	Local Port 44444	Local Port 55555	Local Port 55555
Sender	Recipients			
A	D and E receive packet. (Remote address matches interface address of A; local port of D and E matches remote port of A.)			
B	No socket receives packet. (If multicast group ID is not 1, packets are dropped by the receiver as well-known addresses must not be used by applications.)			
C	D and E receive packet if multicast group ID resolution is implemented in user code and returns "true." (Remote address of C is temporary link-local multicast group; local ports of D and E match remote port of C.)			

2.9.3. Address Resolution

Before unicast data can be transmitted from one device to another, the sending node must determine the link-layer address of the receiver. This is called address resolution and is done automatically by the sender's network stack, using the neighbor discovery protocol (NDP). NDP replaces the address resolution protocol (ARP), which was used in IPv4 networks. Address resolution starts on demand and is transparent to the user.

If a data packet is to be sent to a receiver for which the link-layer address is not known, the sender performs address resolution to find the link-layer address of the receiver. The result is added to the neighbor cache and the data packet is sent out. The maximum size of the neighbor cache is configurable using the stack parameter `ZWIR_spNeighborCacheSize`. Note that changing this parameter at runtime will result in the loss of all cache entries, regardless of whether the neighbor cache size is increased or decreased.

Neighbor cache entries are valid only for a limited time. After this time, the accessibility of the neighbor must be verified. This is also automatically done by the NDP and is beyond the scope of this document. The NDP's `NeighborRetransTime` parameter can be adjusted with the stack parameter `ZWIR_spNeighborRetransTime`. This time defines the timeout in ms between retransmitted NDP packets. The lifetime of neighbor cache entries is defined by routers attached to the network. If the network does not have any routers, a default reachable time is used. IDT's network stack provides the stack parameter `ZWIR_spNeighborReachableTime` for configuring this time. In cases where routers advertise lifetime information, this information is given precedence over the stack parameter.

Note that the default value for this constant (see Table 3.1) significantly differs from the Ethernet default setting of 30 seconds. This is to reduce communication overhead generated by neighbor reachability detection messages. It is possible to disable the timeout completely. This is done by setting `ZWIR_spNeighborReachableTime` to zero.

The exact specification of the neighbor discovery protocol can be found in [RFC 4861](#) – “Neighbor Discovery for IP version 6 (IPv6)”.

2.9.4. Recommendations

The 6LoWPAN protocol performs IPv6 header compression to make the transmission of IPv6 packets more efficient over IEEE 802.15.4 based networks. The header compression mechanism assumes that the interface identifier (i.e., the lower 64 bits of the IPv6 address) is generated from the link-layer address of the device. In such situations, the header compression mechanism is capable of eliding link-local IPv6 addresses completely from the compressed header. Therefore, using manually assigned IPv6 addresses is *NOT RECOMMENDED*. Instead the IPv6 addresses generated by address auto-configuration after device startup *SHOULD* be used.

In order to achieve maximum compression of global IPv6 addresses, it is possible to define compression contexts using the parameters `ZWIR_spHeaderCompressionContext1`, `ZWIR_spHeaderCompressionContext2` and `ZWIR_spHeaderCompressionContext3`. These parameters define frequently occurring prefixes that should be compressed by the 6LoWPAN header compression mechanism. If such prefixes are defined, it must be ensured that each device in the network uses the same configuration of these parameters!

In addition to IPv6 header compression, the 6LoWPAN layer can also compress the UDP header. This is done if the source and/or destination port is in the range of 61616 to 61631. Thus, if the application does not explicitly require another port range, these ports *SHOULD* be used to maximize the data transmission efficiency.

2.10. Mesh Routing

IDT's 6LoWPAN stack enables devices to work in a mesh network topology. If the distance between two communicating devices is too wide for direct radio transmission, packets are routed over intermediate devices – known as “hops” – automatically. Routes through the mesh are detected transparently for the application. Nodes can take two roles in a mesh network scenario: they can act as endpoints only, or they can provide relaying service. In this documentation, devices are named endpoints or relays depending on their configuration. If the stack configuration is not changed by the application, each node is configured as a relay with a maximum hop count of four.

IDT's mesh routing protocol functions on top of the MAC layer, immediately below the network layer.

2.10.1. Multicast Traffic

Network layer multicast traffic is handled by broadcast messages on the mesh and lower layers. Receivers of a mesh broadcast message can forward it depending on their configuration. The decision of whether the message is forwarded is determined based on the configuration parameter **ZWIR_spMaxHopCount**. This value determines the upper limit of hops that a broadcast packet can take through the network. Each broadcast packet carries its hop count in its mesh routing headers. This field is incremented each time the packet is forwarded by a relay. When a node receives a packet with a hop count that is equal to or higher than **ZWIR_spMaxHopCount**, the node does not forward the packet. Otherwise, the hop count is incremented and the packet is forwarded. Thus nodes can be forced to function as endpoints by setting **ZWIR_spMaxHopCount** to zero. Any other configuration makes a node function as a mesh network relay.

2.10.2. Unicast Traffic

For unicast traffic all nodes, hence endpoints and relays, maintain a routing table. The routing table stores the MAC address of the next hop to be taken to a specific destination MAC address. After power on, reset, or network reset, the routing table does not contain any entries. A node requiring unicast communication must set up a route to each of its unicast destination nodes. This is done on demand and transparent for the application.

When the transmission of a unicast packet is requested and there is no matching routing table entry for the destination, the packet to be sent is queued and the route discovery process is started. A Route Request (RReq) is broadcasted into the network, requesting a route to the destination address of the unicast packet. Nodes receiving an RReq check whether the requested address matches their own address or not. If not, the packet is retransmitted by relays and ignored by endpoints, respectively. If the node's own address is matched, a Route Reply (RRep) message is sent to the source hop of the RReq packet. Nodes receiving an RRep packet create/update a record in their routing tables, storing the source address of the RRep packet as the next hop to the requested destination.

Unicast packets always take the same route through the network as long as the route is not removed from the routing tables. Routing table entries are removed if one of the following conditions occurs:

- The route has not been used for **ZWIR_spRouteTimeout** seconds.
- The route has been failing for **ZWIR_spRouteMaxFailCount** times.
- The route was oldest when creation of a new route was required but the routing table was full.

If a hop fails for **ZWIR_spRouteMaxFailCount** times (no acknowledge is sent by the hop), the sender considers the route as broken and sends an informative packet to the originator of the packet. The originator then reinitiates the route discovery process, searching for an alternative route.

2.10.3. Mesh Routing Parameter Configuration Recommendations

In order to maximize the network performance for different application scenarios while maintaining a high level of stability and without wasting resources, the different routing parameters should be configured according to the application's characteristics. All parameters are listed below with explanations of the basic function of the parameter and recommendations for their setting in different application scenarios.

ZWIR_spMaxHopCount

This parameter determines whether a node acts as endpoint or as a relay and constrains the forwarding of multicast packets. With **ZWIR_spMaxHopCount** set to zero, the node acts as a communication endpoint. Note that the node is still able to communicate with remote nodes over multiple hops. Only the ability to forward packets is constrained by this parameter!

In order to make a node function as a mesh network relay, **ZWIR_spMaxHopCount** *MUST* be set to a value greater than zero. However, the value *SHOULD NOT* be chosen arbitrarily, but it *SHOULD* reflect the actual size of the network. The optimal value is the number of hops required to reach the farthest remote communication partner. If no mesh routing is required, setting **ZWIR_spMaxHopCount** to zero will improve the performance.

Limiting the number of nodes working as a relay in the network is strongly recommended. As a rule of thumb, a relay should not have more than ten other relays within direct reachability. Otherwise, the network latency and the packet loss are very likely to increase.

If a large number of relays is desired, using the **ZWIR_spRouteRequestMinRSSI** parameter should be considered for limiting the amount of traffic generated during the route discovery process.

ZWIR_spRoutingTableSize

This parameter configures how many routes can be kept “alive” concurrently. Thus, this parameter defines the number of nodes that the device can communicate with before needing to drop and reestablish routes. The routing table is required in endpoints and relays! On endpoints, the table size should be equal to or larger than the number of remote nodes that the device is intended to communicate with. On relays, this number should be increased by the number of nodes for which the relay service is provided.

The routing table is stored in the RAM and therefore limited by the RAM size. The RAM for the routing table is quasi-statically allocated before the **ZWIR_AppInitNetwork** hook is called. Therefore it is recommended to define the size of the routing-table in **ZWIR_AppInitHardware**. Otherwise a network reset has to be performed in order to get the change into effect.

ZWIR_spRouteTimeout

This parameter defines how many seconds an idle route is kept in the routing table. The default value is 3600 seconds. The idle time counter is restarted each time the route is used. Typically the route timeout parameter does not explicitly affect memory consumption or application performance. However, in frequently changing network configurations, reduction of the timeout value may be advantageous, as old routes do not have to be tested and found to be defective before a new route is established.

ZWIR_spRouteMaxFailCount

This parameter controls how often a route can fail before it is considered to be dead. Depending on the network characteristics, this value should be set to a rather low value between zero and five. The higher the probability of unreachability of a relay or endpoint, the lower this value should be configured. In networks with frequent changes of positions of nodes or a rapidly changing environment, the probability of unreachability is high and therefore this variable should be low. In contrast, fixed installations of nodes and relays can select a higher value, as the unreachability of a node/relay is very likely to be temporary.

ZWIR_spRouteRequestAttempts

This parameter configures how many attempts are made to set up a route to a remote device. By reducing this number, the application can reduce the network load caused by failing route discovery attempts. However, reducing this number will increase the chance of a failing route discovery when it would be physically possible.

ZWIR_spRouteRequestMinLinkRSSI and ZWIR_spRouteRequestMinLinkRSSIReduction

Propagation of electromagnetic waves is influenced by a multitude of external parameters. As a result, radio transmission sometimes appears to behave randomly. Typically this is caused by subtle changes in the external environment. The occurrence of random behavior might increase noticeably in mesh network topologies. For one logical connection of two nodes, there are typically multiple physical links included, all of which have an independent failure probability.

In order to make links more robust against loss of connection due to environmental variations, the parameters **ZWIR_spRouteRequestMinLinkRSSI** and **ZWIR_spRouteRequestMinLinkRSSIReduction** are provided. These parameters allow specifying link quality constraints on each physical link of the whole route. With such constraints in place, smaller environmental changes will impair the routes less, as the signal quality on any link is less likely to drop below the sensitivity level of the module.

2.11. Network and Device Status

The API provides functions for discovering the network and requesting the device status. Network discovery is performed using the **ZWIR_DiscoverNetwork** function. This function broadcasts a message to all devices in the PAN and makes the answers available to the user. For each device, the hop-distance, the link-quality, and all assigned IPv6 addresses are returned.

The node status is returned by **ZWIR_GetTRXStatistic**. The returned data structure contains information such as sent and received packet and byte count and failing transmission attempts. However, the most important value is the sender duty cycle. This value is important, as frequency regulations require nodes to keep their transmission duty cycle lower than 1%. It is the responsibility of the application code to ensure that this number is not exceeded.

2.12. Security

Most applications require secure communication in order to protect sensitive data and to protect actors from unauthorized accesses through attackers. For that reason, IDT provides an implementation of the Internet Protocol Security Suite (IPSec) and the Internet Key Exchange protocol version 2 (IKEv2). IPSec is used to encrypt and authenticate data, while IKEv2 is used to manage the keys used for encryption and authentication. IPSec and IKEv2 are standardized by the Internet Engineering Task Force (IETF). Both protocols are mandated to be used for encryption and key management in IPv6. The implementation of the protocols is provided in two separate libraries.

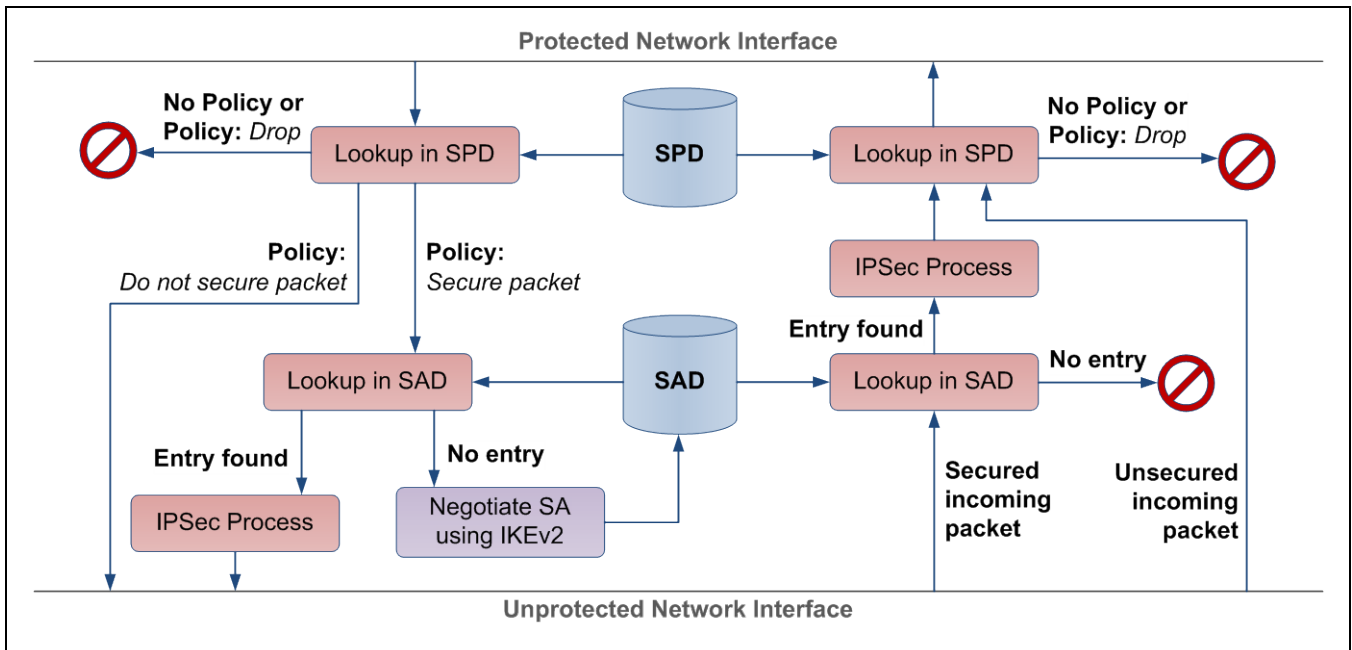
2.12.1. Internet Protocol Security (IPSec)

IDT provides an IPSec implementation in conjunction with its communication libraries. IPSec is a protocol suite for encryption and authentication of data sent over an IP network. IPSec is supported by virtually all modern operating systems. The encapsulating security payload (ESP) and authentication header (AH) protocols are supported for data encryption and authentication, respectively. Data encryption ensures confidentiality of information transmitted over the network. Authentication is applied to ensure that data are not modified along the way and that the sender is the entity that it claims to be.

In order to use the security features of the stack, the *libZWIR45xx-IPSec.a* library must be included in the project and must be configured appropriately. IPSec maintains a Security Policy Database (SPD) that contains rules for how outgoing and incoming traffic must be handled. For each incoming and outgoing packet, the stack checks the SPD for a matching rule that contains information on how the packet should be handled. The rules can direct the network stack to either drop, bypass or process the packet in the security module. Bypassing a packet means that no security processing is applied. Rules can be applied to single addresses or complete subnets.

Each item in the SPD requiring security processing contains a pointer into the Security Association Database (SAD). Each item of this database contains the required information for encryption and decryption of packets. This information includes keying material and the algorithm to be used for encryption and decryption.

Figure 2.6 Working Principle of IPSec



The SAD items can be configured manually or automatically. For automatic configuration, the Internet Key Exchange protocol is used. This protocol is implemented in a separate library and is described in the following section. In either case, SPD entries for incoming and outgoing traffic must be configured by the application. This is done using the function `ZWIRSEC_AddSecurityPolicy`. If manual configuration should be used, the function `ZWIRSEC_AddSecurityAssociation` must be called on both communicating devices, setting the security parameters for the connection.

2.12.2. Internet Key Exchange Version 2 (IKEv2)

The Internet Key Exchange version 2 protocol can be used for automatic creation of keying material for secured connections. This protocol is implemented in the `libZWIR45xx-IKEv2.a` library. If IKEv2 is used, no manual configuration of the SAD is required. Instead, keying material is negotiated automatically on demand. If application code attempts to send data to a remote node and the corresponding SPD entry requires security processing of this data, it is checked to determine if a security association (SA) is assigned to the SPD entry. If no entry exists, IPSec requests the establishment of a security association from the IKEv2 daemon.

IKEv2 first attempts to set up a secure communication channel over which keying material is exchanged. This channel is set up using the Diffie-Hellmann-Key-Exchange algorithm.

Both communicating parties use a Pre-Shared Key (PSK) for mutual authentication. The PSK is registered using the `ZWIRSEC_AddIKEAuthenticationEntry` function. After setup of the secure channel, keying material for the security association to be created is exchanged.

2.12.3. Recommendations

IDT strongly recommends using the security features provided by the network stack. Security is not only required to prevent data from being visible for third parties—more critical is active attacks on a network. Most applications will suffer from such attacks. In the best case, applications might behave erratically; however, in the worst case, perilous behavior of actors can be caused by an attack. Attacking possibilities are manifold: packets can be changed on their way to the destination; they can be sent again; invalid packets infiltrate into the network; or packets can be simply blocked by an attacker. IPSec can protect against all of these attacks. IPSec in conjunction with IKEv2 further increases the security, as the keying material can be renewed on a regular basis.

2.13. Firmware Over-the-Air Updates

IDT provides an over-the-air update (OTAU) library. This library extends the application with functionality for the reception and processing of update packets, as well as functionality for replacing the existing code with a new version. The update mechanism incorporates recovery mechanisms, ensuring the proper recovery after occurrence of an error during the update process.

The OTAU firmware library is designed to require minimal interaction of the firmware programmer. However, it places some constraints on the firmware in order to ensure reliability of the OTAU function.

2.13.1. Functional Description

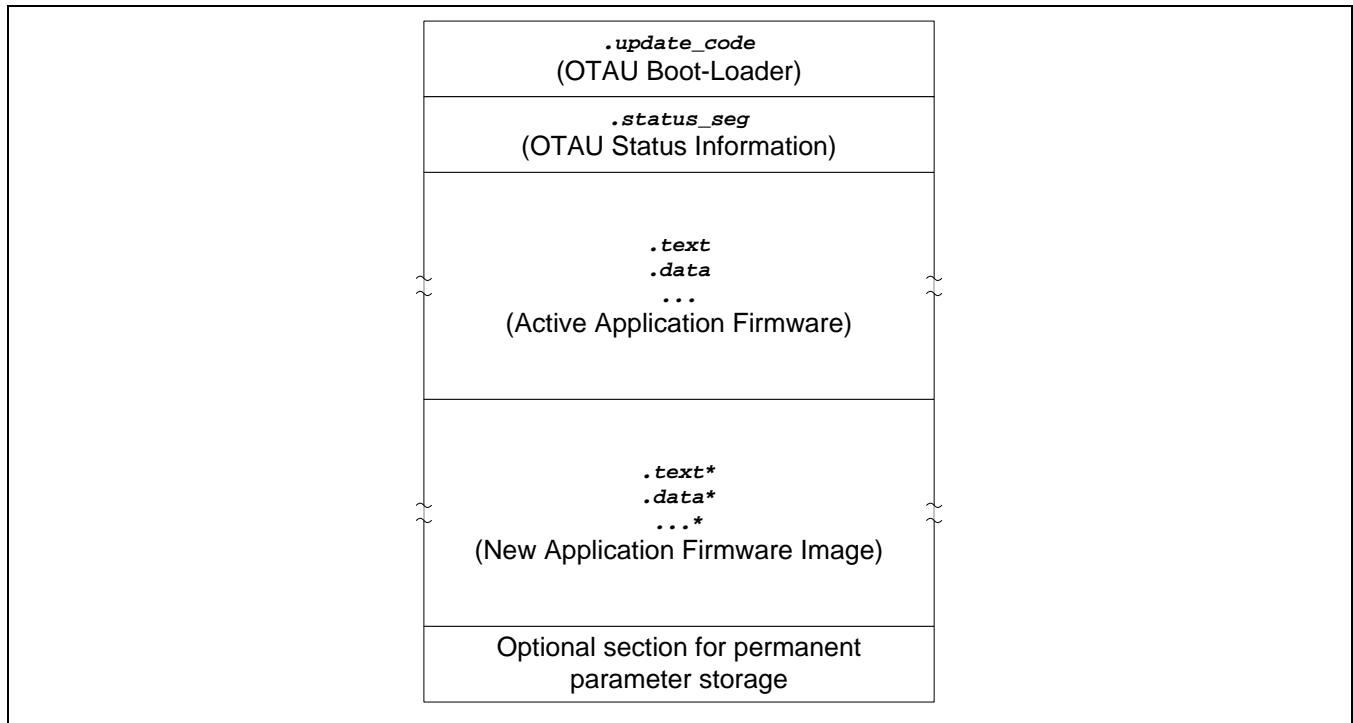
Integrating the firmware over-the-air update (OTAU) adds two components into the user application. The first one is a service for the reception and processing of OTAU-related network traffic. The second one is a boot-loader that replaces the old firmware image with the new one after complete reception and verification of all update traffic. The boot-loader component is located in a program section called `.update_code`. The location of this segment *MUST* be the first flash memory page(s). During the update process, the boot-loader is not replaced!

A second segment that is dedicated to OTAU-enabled code is the `.status_seg`. This segment resides directly behind the boot-loader component and stores status information about the firmware update.

Including the sections referenced above, the application's memory layout would be as shown in Figure 2.7. The application code is located after the `.update_code` and `.status_seg` sections. Optionally the OTAU memory layout can incorporate a section for the storage of permanent parameters. This section *MUST* be located at the end of the flash. In contrast to non-OTAU-enabled applications, the amount of memory available for the application is limited to less than one half of the microcontroller's total flash memory size. This is because the space for the buffering of the full new firmware image must be provided.

The OTAU network service is started through a call to the `ZWIR_OTAU_Register` function. The function takes the UDP port to be used by the OTAU network service as the argument. Calling this function is all that is required to enable the reception and processing of firmware over-the-air updates. All other update parameters are controlled by the update server, which is typically a computer in the network.

Figure 2.7 Memory Layout of OTAU-Enabled Applications



A new firmware image to be loaded into the device is typically received in small chunks of data. Whenever a chunk of data is received, the corresponding flash location in the new application firmware image portion of the flash is updated. If this is the first chunk written to a flash page, the flash page is completely erased before the chunk is written to it. All packets corresponding to the same firmware update *MUST* include the same version information and the size of data chunks *MUST* be the same for all packets. Packets containing fragments of a different size than the first fragment are ignored once the update is started. Packets containing different version information trigger a complete re-initialization of the update.

2.13.2. Firmware Constraints

The OTAU network service uses a dedicated UDP port for the reception and transmission of OTAU-related packets. The application *MUST NOT* use the same port for any other purpose. If this limitation is ignored, the application behavior is determined by the behavior of sockets being reopened with the same parameters.

The contents of the `.update_code` and `.status_seg` segments must not be changed in any way by the application. This means the application *MUST NOT* explicitly place functions or data in either of these sections!

In order to allow the OTAU from a firmware-version **A** to a firmware version **B**, the firmware versions *MUST* share the following properties:

- The call stack of A and B must be located at the same RAM position.
- The call stack of A and B must be of the same size.
- The flash memory layout of A and B must be the same (e.g., no optional section as shown in Figure 2.7 can be added or removed).

2.14. Memory Considerations

Applications must manage their memory consumption – especially with respect to RAM (random access memory). There are basically three components that contribute to the overall RAM size requirements of the application:

1. Statically allocated memory
2. Dynamically allocated memory
3. Call stack

The call stack is required by the application to store the return addresses from function calls, function arguments, and variables stored locally in functions. The RAM size reserved for the call stack can be configured in the linker script. Applications utilizing IDT's network stack need a minimum of 2kB of call stack. In order to leave some flexibility for the user application, the default stack size configured in the linker script is 5kB. The call stack resides at the lower end of the RAM area.

Static memory is consumed by globally declared and local statically declared variables. IDT's network stack requires less than 8kB of static memory in a minimal configuration. The static memory consumption can easily be determined by examining the map file generated by the linker. Static memory is allocated immediately behind the call stack.

The third component, the dynamically allocated memory, is allocated in the unused area between static memory and the end of RAM. Memory in this area is typically allocated at runtime using C's `malloc` function. Memory can be freed using the `free` function. Some lists and buffers used by IDT's network stack are allocated dynamically. There is no tool support for automatic determination of the dynamic memory size requirements of applications. The size of dynamic memory used by IDT's network stack depends on the configured parameters. This is explained in more detail in section 2.14.2.

2.14.1. Call Stack

IDT's network stack places the call stack of the application at the lower end of the RAM. This enables detection of stack overflows. The stack grows downwards from its topmost address towards the beginning of the RAM. If a stack overflow occurs, the MCU tries to access an address that is not in the RAM area and the MCU will generate a Bus-Fault interrupt. Note that during interrupt handling, the interrupt handler function does not have a working stack. Thus, it is not secure to use local variables or calling subroutines. The Bus-Fault interrupt default handler performs a system reset.

2.14.2. IDT Network Stack Dynamic RAM Requirements

IDT's network stack has a number of configurable parameters that require allocation of memory at runtime. During system startup and after reset, memory for these variables is allocated dynamically on the heap. Table 2.7 shows on how these parameters influence the dynamic RAM requirements of the application.

Table 2.7 Stack Parameter Dynamic Memory Size Requirements

Note: The "Size" column specifies the size of a single element. It must be multiplied with the configured parameter value. For each row, an additional 32-byte element is required for storing the allocation record if not otherwise noted.

Parameter	Size [bytes]	Element Count		Notes
		Min	Default	
ZWIR_spRoutingTableSize	28	1	8	
ZWIR_spNeighborCacheSize	60	1	8	
ZWIR_spMaxSocketCount	28	4	8	
-	238	-	6	This memory is allocated for internal buffers for which size cannot be configured.
-	48	-	2	

In addition to the quasi-statically allocated memories above, IDT's network stack dynamically allocates memory at runtime if packets must be sent to destinations for which address resolution and route discovery have not been performed. One packet can be buffered for each destination node. Allocation is performed if enough memory is available. If no memory is left, the packet to be sent is dropped, but the address resolution and route detection procedure is initiated in any case. Thus, even if the packet is not being buffered, the next packet being sent is likely to arrive at the destination node.


Application developers must always ensure that the parameter settings allow proper allocation of all quasi-static memory. If parameters are chosen too large, stack initialization will fail and report the failure as error **ZWIR_eMemoryExhaustion**. The default handling of this error is a system reset. Thus, if the parameters causing the memory exhaustion are set during system startup, this will result in an infinite loop. The error is detected easily with a custom implementation of the **ZWIR_Error** function.

2.14.3. Using Dynamic Memory Allocation

Applications requiring dynamic memory allocation can freely use the functions **malloc** and **free** for allocation and de-allocation of RAM at runtime, respectively. However, the application developers must be aware of the limited availability of RAM on the device. Each allocated block consumes an additional 32-byte block on the heap for the allocation record.

Due to memory fragmentation effects, it cannot be guaranteed that memory allocation is successful, even if the total amount of free heap space is sufficient for an allocation request. If memory blocks are allocated and freed frequently, the free space in memory might become scattered over the whole heap, not providing any free block large enough for holding a requested block. Figure 2.8 demonstrates this with a simple example that has 3kB of free memory but does not allow the allocation of a 2kB memory block with **malloc**.

Figure 2.8 Heap Memory Scattering



2.15. Supported Network Standards

Table 2.8 lists RFCs that are supported by IDT's network stack, and it specifies the limitations that apply with respect to these RFCs.

Table 2.8 Supported RFCs and Limitations

RFC	Limitations
Internet Protocol Version 6 (IPv6) Specification	
2460	<ul style="list-style-type: none"> • Hop-by-hop options header <ul style="list-style-type: none"> ○ Only Pad1 and PadN options are supported (as specified in RFC). ○ Other options will cause unrecognized option processing as proposed in RFC. • Routing extension header <ul style="list-style-type: none"> ○ If "Segments Left" > 0, packets are ignored and an Internet Control Message Protocol (ICMP) error message parameter problem is sent. However, the use of the type 0 routing header has been deprecated by RFC 5095! • Fragmentation extension header <ul style="list-style-type: none"> ○ Not supported – packets received with this extension header are silently dropped. ○ Specification requirement of receiving 1500-byte packets is not supported. However, use of fragmentation is discouraged by the RFC! • Destination options header <ul style="list-style-type: none"> ○ Only Pad1 and PadN options are supported (as specified in RFC). ○ Other options will cause unrecognized option processing as proposed in RFC. ○ Packets with next header 59 are dropped. ○ Traffic class and flow label are always set to zero in packets sent from 6LoWPAN nodes.
Security Architecture for the Internet Protocol	
4301	<ul style="list-style-type: none"> • Tunnel mode is not supported. • ESP SA with both null encryption and no integrity algorithm is allowed. • Events are not logged. • Local IPv6 address cannot be used as a selector. • No sequence counter overflow handling. • SA lifetime is handled by IKE and only time controlled. • Certificates are not supported for IKE authentication. • No ICMP error messages processing and generation. • Fragmentation and reassembly is not supported.
IP Authentication Header	
4302	<ul style="list-style-type: none"> • AH is not supported.

RFC	Limitations
IP Encapsulating Security Payload (ESP)	
4303	<ul style="list-style-type: none"> • SPI 0 to 255 are not reserved. • Anti-replay service is not active. • The sequence number will cycle. • ESN is not supported. • Dummy packets are not supported. • Traffic flow confidentiality (TFC) padding is not supported. • Auditing is not supported.
Cryptographic Algorithm Implementation Requirements for ESP and Authentication Header (AH)	
4835	<ul style="list-style-type: none"> • Supports ONLY NULL encryption and AES-CTR. • Supports ONLY NULL authentication and AES-XCBC-MAC-96.
Neighbor Discovery for IP Version 6 (IPv6)	
4861	<ul style="list-style-type: none"> • Only host functionality is implemented. • Destination cache as proposed by the standard is not available. However, there is no need for this as the purpose of the cache is to reduce the time required for the next hop determination, but this is already done in a fraction of the time that message transport requires. • No checking of the linkMTU option is performed – however this is not required as 6LoWPAN only supports the minimum linkMTU of 1280 and never sends larger packets. • If no reachable router is in the router list, default router selection is not performed in a round-robin manner as proposed by the standard; instead the first entry found is taken. However, reachable routers still have precedence over routers whose reachability is unknown (see section 6.3.6 of RFC4861). • Variables are mainly implemented as constants. • During address resolution, packets must be queued until address resolution is complete. The memory for this is allocated dynamically at runtime. If memory allocation fails, the packet is not queued! Only a single packet will be queued. A queued packet is not replaced with the latest one if more than one packet needs to be buffered during the address resolution process. • Changes of the link-layer address are not advertised as proposed in section 7.2.6 of RFC4861. However, this is not required, as the node will also change its IPv6 address. • Anycast neighbor solicitation is not supported. • Redirect messages are not supported (see section 8 of RFC4861). • Nodes that use multicast do not use the Multicast Listener Discovery (MLD) protocol to announce the groups in which they are members.
IPv6 Stateless Address Auto-configuration	
4862	<ul style="list-style-type: none"> • Maximum number of NS for DAD is limited to 1. • Address deprecation is not implemented. Behavior is the same as preferredLifetime==validLifetime. However, router advertisements with a preferredLifetime>validLifetime are ignored by the device.

RFC	Limitations
Transmission of IPv6 Packets over IEEE 802.15.4 Networks	
4944	<ul style="list-style-type: none"> IDT's implementation does not support one of the specified header compression algorithms proposed by the RFC. Instead, it implements the RFC draft-hui-6lowpan-hc version 01 (see "6LoWPAN Compression of IPv6 Datagrams" below). 03_{HEX} is used as the dispatch value for this compression format. Mesh functionality specified in the RFC is not implemented. Instead, a proprietary link-layer mesh implementation is provided. Only 64-bit link-layer addresses are supported at this time.
IPv6 Configuration in Internet Key Exchange Protocol Version 2 (IKEv2)	
5996	<ul style="list-style-type: none"> Supports only the negotiation of an ESP in transport mode between two protected endpoints. The Diffie-Hellman group cannot be changed. Only one initial exchange can occur at the same time. Only one pair of child SAs can be negotiated with one IKE SA. Windowing is not supported. Timeouts are defined by user. The critical flag is ignored. Cookies are not supported. Implementation provides only one proposal. Only packets for port 500 are accepted. Extensible Authentication Protocol (EAP) is not supported. IP Compression (IPComp) is not supported. Network address translation (NAT) traversal is not supported. Only ID_IPV6_ADDR is supported. Only shared key message integrity code is supported. Vendor ID payload is not supported. Configuration payload is not supported.
Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)	
4307	<ul style="list-style-type: none"> Supports only 768 MODP (Modular Exponential) Group. Supports only ENCR_AES_CBC. Supports only PRF_AES128_CBC. Supports only AUTH_AES_XCBC_96.
6LoWPAN Compression of IPv6 Datagrams	
draft-hui-6lowpan-hc-01	<ul style="list-style-type: none"> ISA100_UDP Header Compression is not implemented (see section 3.3 in draft specification).

3 Core-Library Reference

3.1. Initialization

The core library provides two different hooks that can be used to initialize the application during system startup. The first hook, named `ZWIR_AppInitHardware`, is called before network initialization. The second one, named `ZWIR_AppInitNetwork`, is called afterwards.

```
void
ZWIR_AppInitHardware ( ZWIR_ResetReason_t  resetReason )
```

This hook is called after power on and after reset to configure the module peripherals. The `resetReason` argument specifies the reset source that triggered the execution of this function. If required, the operating mode of the module *SHOULD* be set from this function.

Note: Most API functions must not be called from this function. The documentation specifies which API functions can be called from `ZWIR_AppInitHardware`.

```
void
ZWIR_AppInitNetwork ( ZWIR_ResetReason_t  resetReason )
```

This hook is called when the default networking parameters have been initialized after power-on or reset. It should be used to open sockets and initialize further network parameters if required. However, it must not be used for sending out data, as the node has not completed Duplicate Address Detection (DAD); refer to section 2.8.4 for further details) at this point. The `resetReason` argument specifies the reset source that triggered the execution of this function.

```
void
ZWIR_AppInitNetworkDone ( ZWIR_ResetReason_t  resetReason )
```

This hook is called after successful completion of the DAD procedure. It is also called when DAD is disabled. In this case, the function is called immediately after the call to `ZWIR_AppInitNetwork`. The `resetReason` argument specifies the reset source that triggered the execution of this function. All API functions can be called from this function.

```
void
ZWIR_SetOperatingMode ( ZWIR_OperatingMode_t      opMode,
                       ZWIR_RadioReceiveCallback_t  callback )
```

This function sets the operating mode of the device. Device Mode, Gateway Mode, or Sniffer Mode can be selected (refer to section 0) with the `opMode` argument. The `callback` argument is used in Gateway Mode and Sniffer Mode to specify the function to be called on the reception of data. If the callback is `NULL`, no function will be called. If Device Mode is selected `callback` is ignored. It is *RECOMMENDED* that `ZWIR_SetOperatingMode` only be called from `ZWIR_AppInitHardware`.

```
typedef enum { ... } ZWIR_OperatingMode_t
```

Type enumerating the different operating modes of the device. Possible values include:

<code>ZWIR_omNormal</code>	Device Mode
<code>ZWIR_omGateway</code>	Gateway Mode
<code>ZWIR_omSniffer</code>	Sniffer Mode

```
typedef enum { ... } ZWIR_ResetReason_t
```

Type enumerating the different reasons for system reset. Possible values include:

<code>ZWIR_rPowerOnReset</code>	Reason: The device has been powered on after being switched off.
<code>ZWIR_rStandbyReset</code>	Reason: The device is waking up from standby mode
<code>ZWIR_rIndependentWatchdogReset</code>	Reason: The MCU's independent watchdog (IWDG) was triggered.
<code>ZWIR_rSoftwareReset</code>	Reason: A software reset has been performed
<code>ZWIR_rPinReset</code>	Reason: System reset was triggered by pulling low the reset pin.
<code>ZWIR_rWindowWatchdogReset</code>	Reason: The window watchdog has triggered
<code>ZWIR_rLowPowerReset</code>	Reason: The supply voltage dropped below the specified threshold.

3.2. Program Control

The API does not provide the concept of a central main function as is typical in traditional C programs. Only three timing-driven hooks, named `ZWIR_Main10ms`, `ZWIR_Main100ms`, and `ZWIR_Main1000ms`, are provided; they are called periodically after 10, 100 and 1000 milliseconds, respectively. Sensing and acting by the user application should be implemented in these hooks. For more fine-tuned time control, a user-programmable callback timer is available. This timer can be programmed at 1ms increments. Initialization and de-initialization of this freely programmable timer function is via `ZWIR_StartCallbackTimer` and `ZWIR_StopCallbackTimer`.

The `ZWIR_Main10ms`, `ZWIR_Main100ms`, and `ZWIR_Main1000ms` hooks can be defined to implement application behavior that has to be executed periodically. The default implementations of these hooks do nothing, so these hooks can be left undefined if they are not required.

In addition to the fixed period main functions, the API also provides a freely configurable callback timer. The timer is started using the `ZWIR_StartCallbackTimer` function. It is possible to provide a data pointer to this function, which is passed to the callback when the timer expires. This allows for delayed data processing. Whether the timer is triggered just once or periodically can be selected.

The timer is stopped with the `ZWIR_StopCallbackTimer`.

```
void  
  ZWIR_Reset ( void )
```

This function causes a software reset of the system. Both, the microcontroller and the transceiver are reset. The complete startup sequence is executed. If this function is called while the transceiver receives or transmits data, the packet will be lost.

```
void  
  ZWIR_ResetNetwork ( void )
```

This function resets the radio transceiver and reinitializes the network stack. After a call to this function, IPv6 address auto-configuration is restarted and manually assigned addresses are lost. Also routing and address resolution information are lost.

```
void  
  ZWIR_Main10ms ( void )  
void  
  ZWIR_Main100ms ( void )  
void  
  ZWIR_Main1000ms ( void )
```

These hooks are called with a period of 10, 100 and 1000 milliseconds. On timeslots that are multiples of 10ms or 100ms, the shorter period function has priority over the longer period. This means that `ZWIR_Main10ms` is called before `ZWIR_Main100ms`, which is called before `ZWIR_Main1000ms`. All three functions are called immediately at system startup.

Note: These functions are not suitable if exact timer behavior is required. A constant execution interval cannot be guaranteed nor is it guaranteed that the function is executed at each intended time instant.

```
void
ZWIR_StartCallbackTimer ( uint32_t          timeout,
                          ZWIR_TimeoutCallback_t callback,
                          void*           data,
                          bool            periodic )
```

If this function is called, the freely programmable timer is initialized and started. The function provided with the *callback* argument will be called about *timeout* ms after the call to `ZWIR_StartCallbackTimer`. The value provided with *data* will be passed to *callback* when it is called. If the *periodic* flag is set to one, *callback* is called periodically; otherwise *callback* is called just once. If this function is called while the timer is running, the timer will be reprogrammed and the previous programming will be lost.

Note: The callback timer is not suitable if exact timer behavior is required. Consider using a MCU timer peripheral for exact timing.

```
void
ZWIR_StopCallbackTimer ( )
```

This function stops a running callback timer. If no timer is running, nothing will happen.

```
void
ZWIR_TriggerAppEvent ( uint8_t eventId )
```

This function allows the processing of application events with a certain operating system priority. This function notifies the operating system of the presence of an application event and schedules the appropriate callback function for execution. Typically this is used to execute computationally intensive code in response to an interrupt.

```
void
ZWIR_RegisterAppEventHandler ( uint8_t          eventId,
                               ZWIR_AppEventHandler_t handler )
```

This function registers an application event handler callback for a certain application event in the operating system. If this function is called more than once with the same *eventId*, the callback function provided with the last call will be in effect.

```
typedef void ( * ZWIR_AppEventHandler_t ) ( void )
```

This function pointer type defines the signature of a callback function that is executed in response to an application event.

```
typedef void ( * ZWIR_TimeoutCallback_t ) ( void* data )
```

Function pointer type for the callback function that should be called if the callback timer expires.

```
ZWIR_RevisionInfo_t
ZWIR_GetRevision ( void )
```

This function returns a structure containing detailed version information. This information must be provided if support requests are sent to IDT.

```
typedef struct { int8_t    majorRevision
                int8_t    minorRevision
                int16_t   versionExtension } ZWIR_RevisionInfo_t
```

Type for objects carrying version information. If problems are encountered while using the stack, request this structure using `ZWIR_GetRevision` and provide the information obtained to IDT with an error report.

```
void
ZWIR_NetEventCallback ( ZWIR_NetEvent_t event )
```

This function is called when a network event occurs. The event type is passed in the `ZWIR_NetEvent_t event` argument.

```
typedef enum { ... } ZWIR_NetEvent_t
```

Net events enumeration type accepted by `ZWIR_NetEventCallback`. Possible values:

- `ZWIR_neAppReceive`
Event: UDP receive callback function is called.
- `ZWIR_neAppTransmit`
Event: UDP transmit function is called.
- `ZWIR_neIPv6Receive`
Event: IPv6 stack receive function is called.
- `ZWIR_neIPv6Transmit`
Event: IPv6 stack transmit function is called.
- `ZWIR_neMACReceive`
Event: Receiver receive function is called.
- `ZWIR_neMACTransmi`
Event: Receiver transmit function is called.


```
bool  
ZWIR_Error ( int32_t errorCode )
```

This function is called when a recoverable library error is encountered. The error-code is passed in the *errorCode* argument. If *true* is returned, the error is assumed to be processed and no action will be taken by the stack. If *false* is returned, the default error handler will be executed.

```
int32_t  
ZWIR_Rand ( void )
```

This function returns a random number. The sequence of numbers generated by this function is actually only pseudo-random, as a linear feedback shift register with a generator polynomial is used. However, by calling **ZWIR_SRand** with zero as argument, the random number generator is seeded with a true random number.

```
int32_t  
ZWIR_SRand ( int32_t seed )
```

This function seeds the random number generator. If zero is provided as *seed*, a true random number is generated from thermal noise. Any value other than zero is used to initialize the generator directly. This allows creating reproducible application behavior for debug purposes. Using non-zero seed values in production code is not recommended.

3.3. Networking

A ZWIR451x node can join any IPv6 network. Each device automatically gets an IPv6 address that is computed from the link-layer address of the module. Additionally, the user's own IPv6 addresses can be assigned to the module. The modules can communicate bidirectionally using the UDP protocol.

3.3.1. Address Management

The API provides a set of functions for managing the different addresses of a 6LoWPAN module. A module has three different types of addresses: the PAN identifier, the link-layer address, which is also called the PAN address, and a set of IPv6 addresses.

3.3.1.1. PAN Identifier

The PAN identifier is a 16-bit value that determines the personal area network that the module belongs to. All devices in a PAN must have the same PAN identifier. Devices with different PAN identifiers cannot communicate with each other. They cannot even use each other as a network relay. Each device has exactly one PAN identifier. It can be read and altered using the **ZWIR_SetPANId** and **ZWIR_GetPANId** functions, respectively. The default value of the PAN identifier is `ACCAHEX`.

```
uint16_t
ZWIR_GetPANId ( void );
```

Reads and returns the PAN identifier of the module.

```
void
ZWIR_SetPANId ( uint16_t panId );
```

Sets the PAN identifier of the node to the value provided in *panId*. The value is retained until the next reset or until Standby Mode is entered. After waking up, the factory-programmed value is active again until the next call to this function.

3.3.1.2. Link-Layer Address

The link-layer address, also called the PAN address, is a 64-bit-wide value that identifies the device in the PAN. All communication between devices in a PAN is based on the link-layer address. Higher layer protocols, such as IPv6, resolve their addresses into link-layer addresses. A globally unique link-layer address is programmed into each device during manufacturing. Changing this address is not recommended, as this could cause the address to be no longer unique. Nevertheless, the functions `ZWIR_SetPANAddress` and `ZWIR_GetPANAddress` allow reading and writing the PAN address.

```
ZWIR_PANAddress_t const*
ZWIR_GetPANAddress ( void )
```

Reads and returns the link-layer address of the module.

```
void
ZWIR_SetPANAddress ( ZWIR_PANAddress_t const* panAddress )
```

Sets the link-layer address to the value provided in *panAddress*. Changing the link-layer address of the module is not recommended for the reasons given in section 2.8.4. However, if manual assignment of addresses is required, calling `ZWIR_SetPANAddress` from `ZWIR_AppInitHardware` is recommended. The assigned *panAddress* value is retained until the next system reset or deep sleep. After waking up, the factory-programmed value is active again until the next call to this function.

Note: Changing the link-layer address of a device during normal operation will typically cause a loss of all incoming packets for a period of time. The standard allows sending unsolicited neighbor advertisements as an option if the link-layer address changes. This feature is not included in IDT's 6LoWPAN stack.

```
typedef uint8_t ZWIR_PANAddress_t [8]
```

Data type for representation of link-layer addresses. Bytes are stored in network byte order, which is big endian. This means that the highest-order byte is stored first.

3.3.1.3. IPv6 Addresses

Although manual assignment of IPv6 addresses is not required by most applications, the API provides the function `ZWIR_SetIPv6Address`. This function can be used to assign additional IPv6 addresses to an interface. Up to three addresses can be assigned in total, but the first address is always allocated by the automatically configured link-local address. The function `ZWIR_GetIPv6Addresses` can be used to request all addresses assigned to an interface.

The `ZWIR_CheckMulticastGroup` hook is called if a multicast packet is received that does not belong to the all nodes multicast group or the node solicited address multicast group. This function must be implemented by the application if multicast addressing is used.

bool

```
ZWIR_SetIPv6Address ( ZWIR_IPv6Address_t const* ipv6 )
```

Add the address provided in *ipv6* to the network interface.

The function returns *true* if the operation was successful or *false* otherwise. The function fails if the maximum number of IPv6 addresses is assigned to the interface already.

uint8_t

```
ZWIR_GetIPv6Addresses ( ZWIR_IPv6Address_t const* ipv6Buffer,
                       uint8_t maxCount )
```

Request a set of addresses assigned to the interface. The *ipv6Buffer* argument must carry a pointer that points to a buffer that is able to store at most *maxCount* IPv6 addresses. The function will store *maxCount* addresses in this buffer if the interface has at least *maxCount* addresses assigned. If the number of assigned address is lower than *maxCount*, the function will store all available addresses. The return value determines the number of addresses that have actually been stored. Thus, if 0 is returned, the interface has no IPv6 address assigned. This could be due to failing duplicate address detection (DAD).

void

```
ZWIR_SetDestinationPANId ( uint16_t pandID )
```

This function is used for changing the destination PAN identifier temporarily. The configured value remains in effect until `ZWIR_ResetDestinationPANId` is called or the device is reset.

```
void
  ZWIR_ResetDestinationPANId ( void )
```

This function resets the PAN identifier of the device to the last value that had been configured before it was changed using `ZWIR_SetDestinationPANId`.

```
bool
  ZWIR_CheckMulticastGroup ( ZWIR_IPv6Address_t const* ipv6 )
```

This hook is called whenever a multicast packet is received that contains a multicast group that is not known by the network stack. The user implementation must decide if the node is part of the multicast group provided with the multicast group ID (the lower 112 bytes) in the IPv6 address. *True* must be returned if the node belongs to the multicast group, *false* otherwise.

```
typedef union { uint8_t   u8   [ 16 ],
                uint16_t  u16  [  8 ],
                uint32_t  u32  [  4 ] } ZWIR_IPv6Address_t
```

Data type for representation of IPv6 addresses. Bytes are stored in network byte order, which is big endian; i.e., the highest-order byte is stored first. Therefore using the `u16` or `u32` elements might cause unexpected results especially when printing. Consider changing the byte order if host byte order is required.

3.3.2. Socket and Datagram Handling

The functions `ZWIR_OpenSocket` and `ZWIR_CloseSocket` are provided for opening and closing sockets, respectively. Datagrams are sent over sockets using the `ZWIR_SendUDP`, `ZWIR_SendUDP2` and `ZWIR_Send6LoWPAN` functions, depending on the operating mode of the device.

Incoming datagrams are handled by user-defined callback functions, which must be assigned to sockets using the `ZWIR_OpenSocket` function. The API functions `ZWIR_GetPacketSenderAddress`, `ZWIR_GetPacketDestinationAddress`, `ZWIR_GetPacketSenderPort` and `ZWIR_GetPacketHopCount` are provided for requesting the address and port of a sender.

ZWIR_SocketHandle_t

```
ZWIR_OpenSocket ( ZWIR_IPv6Address_t const*  remoteAddr
                 uint16_t                    remotePort,
                 uint16_t                    localPort,
                 ZWIR_RadioReceiveCallback_t rxHandler )
```

This function opens a new socket to a remote host. The *remoteAddr* and *remotePort* arguments specify the IPv6 address and the UDP port of the remote host, respectively. The *localPort* argument specifies the port on which incoming data is accepted. If *localPort* is set to zero, an unused port from the range from 4096 through 32000 is chosen. If incoming data should be received, a pointer to a callback function must be passed in the *rxHandler* argument. If no data is expected, the value of *localPort* does not matter and *rxHandler* must be set to NULL. In order to receive packets from arbitrary remote hosts, the unspecified address can be passed to the function. In this case, the socket is not suitable for sending packets.

On success, the function returns a socket handle that can be used with datagram handling functions. The function fails if the maximum number of sockets is already opened or if a socket with the same parameters for *remoteAddress* and *localPort* already exists. In this case NULL is returned.

void

```
ZWIR_CloseSocket ( ZWIR_SocketHandle_t  socket )
```

Open sockets are closed using this function. If a *socket* is invalid or has already been closed, the function has no effect. Closing a socket has no effect on any previously sent packets, even if transmission is not yet completed.

bool

```
ZWIR_SendUDP ( ZWIR_SocketHandle_t  socket,
               uint8_t const*        data,
               uint16_t               length )
```

UDP datagrams are sent over a specific socket using this function. The socket denoted by the *socket* argument determines the destination address and port of the datagram. The *data* and *length* arguments specify the payload. The maximum packet size is 1232 bytes. The function returns a non-zero value if the packet to be sent was successfully queued in the output queue; otherwise zero is returned. If zero is returned, there is not enough room in the output queue to buffer this packet. In this case, control must be passed back to the operating system.

Note: A non-zero return value does not automatically denote the successful delivery of the packet. Successful delivery can only be verified by response packets sent on the application level.

Note: Calling this function in a while loop waiting for a non-zero result will deadlock the system if a packet cannot be queued. After a zero result, control must always be passed to the operating system. Otherwise the output buffers will never be freed and this function continues to fail, resulting in a deadlock.

Note: This function cannot be used in the Gateway Mode; use `ZWIR_send6LOWPAN` when in Gateway Mode!

```
bool
ZWIR_SendUDP2 ( uint8_t*      data,
                uint16_t     size,
                ZWIR_IPv6Address_t* remoteAddress,
                uint16_t     remotePort )
```

This function sends an UDP packet without the need for opening a socket. The destination address and destination port are provided in the *remoteAddress* and *remotePort* arguments. The local UDP port is selected arbitrarily by the network stack. The maximum packet size is 1232 bytes. The function returns a non-zero value if the packet to be sent was successfully queued in the output queue; otherwise zero is returned. If zero is returned, there is not enough room in the output queue to buffer this packet. In this case, control must be passed back to the operating system.

Note: A non-zero return value does not automatically denote the successful delivery of the packet. Successful delivery can only be verified by response packets sent on the application layer.

Note: Calling this function in a while loop waiting for a non-zero result will deadlock the system if a packet cannot be queued. After a zero result, control must always be passed to the operating system. Otherwise the output buffers will never be freed and this function continues to fail, resulting in a deadlock.

Note: This function cannot be used in the Gateway Mode; use `ZWIR_send6LoWPAN` when in Gateway Mode!

```
ZWIR_IPv6Address_t const*
ZWIR_GetPacketSenderAddress ( void )
```

This function returns a pointer to the IPv6 source address of the last received packet. It can be used reliably in the RX callback function. Using this function outside of the RX callback function might cause unpredictable results.

Note: This function cannot be used in Gateway Mode.

```
ZWIR_IPv6Address_t const*
ZWIR_GetPacketDestinationAddress ( void )
```

This function returns a pointer to the IPv6 destination address of the last received packet. It can be used reliably in the RX callback function. Using this function outside of the RX callback function might cause unpredictable results.

Note: This function cannot be used in Gateway Mode.

```
uint16_t
ZWIR_GetPacketSenderPort ( void )
```

This function returns the sender port of the last received packet. It can be used reliably in the RX callback function. Using this function outside of the RX callback function might cause unreliable results.

Note: This function cannot be used in Gateway Mode.

`uint8_t`

`ZWIR_GetPacketHopCount (void)`

Returns the number of hops the last received packet has taken.

Note: Using this function outside of the RX callback function might cause unreliable results.

`int32_t`

`ZWIR_GetLastRSSI (void)`

Returns the receive signal strength indicator (RSSI). The value approximately corresponds to the receive power level in dBm. Using this function outside of the RX callback function might cause unreliable results.

`ZWIR_PANAddress_t*`

`ZWIR_GetSourcePANAddress (void)`

Returns the source PAN address of the latest received packet.

Note: Using this function outside the receive callback might cause unreliable results.

`ZWIR_PANAddress_t*`

`ZWIR_GetDestinationPANAddress (void)`

Returns the destination PAN address of the last received packet.

Note: Using this function outside the receive callback might cause unreliable results.

`ZWIR_SocketHandle_t`

`ZWIR_GetPacketRXSocket (void)`

Returns the socket handle of the socket on which the last packet was received.

`ZWIR_IPv6Address_t*`

`ZWIR_GetFailingAddress (void)`

Returns the last address for which address resolution failed. This function is typically called from `ZWIR_Error` when the `ZWIR_eHostUnreachable` error was reported. If no address resolution error occurred since the last reset, the result is undefined.

```
bool  
ZWIR_Send6LoWPAN ( ZWIR_PANAddress_t const* remoteAddr,  
                  uint16_t const* data,  
                  uint8_t const length )
```

This function is used to send complete IPv6/UDP packets to the remote host with the link-local address *remoteAddr*. No UDP or IPv6 header processing is performed on the packet. Instead it is passed directly to the 6LoWPAN processing layer. The *data* argument must point to the first header byte of the IPv6/UDP packet; *length* specifies the size including all headers. This is useful in conjunction with the Gateway Mode.

Note: Calling this function in a while loop waiting for a non-zero result will deadlock the system if a packet cannot be queued. After a zero result, control must always be passed to the operating system. Otherwise the output buffers will never be freed and this function continues to fail, resulting in a deadlock.


```
typedef
void ( * ZWIR_RadioReceiveCallback_t ) ( uint8_t* data,
                                         uint16_t length )
```

Function pointer type for the callback function that should be called on reception of data over an UDP socket.

```
typedef
void* ZWIR_SocketHandle_t
```

Type representing a socket.

3.3.3. Radio Parameters

The radio module provides the capability to alter the physical radio channel and the transmit output power. This is accomplished using the `ZWIR_SetChannel`, `ZWIR_SetModulation`, and `ZWIR_SetTransmitPower` functions. Changes to the radio parameters take effect immediately. The changes are reset by `ZWIR_Reset`, but not by `ZWIR_ResetNetwork`. Changing radio parameters during the transmission/reception of a packet will very likely cause the loss of the packet.

To verify the transceiver parameter, the corresponding get function `ZWIR_GetChannel`, `ZWIR_GetModulation` or `ZWIR_GetTransmitPower` can be called.

```
void
ZWIR_SetChannel ( ZWIR_RadioChannel_t channel )
```

Sets the module to the radio channel specified by *channel*.

Note: If this is done while a transmission or reception is ongoing, the transmitted or received packet will be lost. It is recommended that this function be called from `ZWIR_AppInitHardware`.

Note: Local regulations limit the use of the spectrum. The user *MUST* only select channels that are allowed to be used in the area where the application is going to be installed! Check with local authorities to determine which part of the spectrum the application is allowed to use.

```
ZWIR_RadioChannel_t
ZWIR_GetChannel ( )
```

Returns the current `ZWIR_RadioChannel_t` *channel*

```
void
ZWIR_SetModulation ( ZWIR_Modulation_t modulation )
```

This function is used to change the modulation scheme to the value specified with the *modulation* argument.

Note: If this is done while a transmission or reception is ongoing, the transmitted or received packet will be lost. It is recommended that this function be called from `ZWIR_AppInitHardware`.

```
ZWIR_Modulation_t
ZWIR_GetModulation ( )
```

Returns the current `ZWIR_Modulation_t modulation`

```
void
ZWIR_SetTransmitPower ( int power )
```

Sets the transceiver output power to the value specified by `power`. The valid range of values depends on the channel being selected. In the European frequency band, a transmission power of -10dBm to 5dBm can be selected; in the US band, -10dBm to 10dBm are permitted. Values that are too low or too high are automatically adjusted to the closest valid value.

Note: If this is done while a transmission is ongoing, the transmitted packet is very likely to be lost. For that reason, it is recommended that this function be called only from `ZWIR_AppInitHardware`.

```
int
ZWIR_GetTransmitPower ( )
```

Returns the current transmit power.

```
typedef enum { ... } ZWIR_RadioChannel_t
```

Radio channel enumeration type accepted by `ZWIR_setChannel`. Possible values include

<code>ZWIR_channel0,</code>	<code>ZWIR_eu868</code>	EU Band, 868.3 MHz
<code>ZWIR_channel1,</code>	<code>ZWIR_us906</code>	US Band, 906 MHz
<code>ZWIR_channel2,</code>	<code>ZWIR_us908</code>	US Band, 908 MHz
<code>ZWIR_channel3,</code>	<code>ZWIR_us910</code>	US Band, 910 MHz
<code>ZWIR_channel4,</code>	<code>ZWIR_us912</code>	US Band, 912 MHz
<code>ZWIR_channel5,</code>	<code>ZWIR_us914</code>	US Band, 914 MHz
<code>ZWIR_channel6,</code>	<code>ZWIR_us916</code>	US Band, 916 MHz
<code>ZWIR_channel7,</code>	<code>ZWIR_us918</code>	US Band, 918 MHz
<code>ZWIR_channel8,</code>	<code>ZWIR_us920</code>	US Band, 920 MHz
<code>ZWIR_channel9,</code>	<code>ZWIR_us922</code>	US Band, 922 MHz

ZWIR_channel110, ZWIR_us924	US Band, 924 MHz
ZWIR_channel1100, ZWIR_eu865	EU Band, 865.3 MHz
ZWIR_channel1101, ZWIR_eu866	EU Band, 866.3 MHz
ZWIR_channel1102, ZWIR_eu867	EU Band, 867.3 MHz

```
typedef enum { ... } ZWIR_Modulation_t
```

Enumeration of modulation schemes accepted by `ZWIR_SetModulation`. Possible values include

ZWIR_mBPSK	Binary Phase Shift Keying
ZWIR_mQPSK	Offset Quadrature Phase Shift Keying

3.3.4. Gateway Mode Functions

The ZWIR45xx network stack provides the Gateway Mode to allow easy implementation of network bridges. Therefore, the stack only performs network processing up to and including the 6LoWPAN layer. Thus, when a packet comes in, a full IPv6 packet is passed to the receive callback. This allows the implementation of a perfectly transparent network bridge. However, if the bridge should be updateable or configuration parameters should be set remotely, it would be desirable to have the opportunity of performing higher layer processing of incoming packets that are addressed to the bridge. For this purpose, the types and functions defined in this section are provided.

```
bool
ZWIR_GatewayProcessPacket ( uint8_t* data,
                             uint16_t size );
```

This function must be called to cause the network stack to process the network and higher layer protocols of an incoming packet. The function is intended to be called from the receive callback of Gateway Mode devices. The *data* and *size* arguments of the receive callback should be passed unmodified to this function.

```
void
ZWIR_GatewaySetOutputFunction ( ZWIR_GatewayOutputFunction_t fn )
```

Devices operating in Gateway Mode typically have more than one network interface. When higher layer protocol processing is in place, it must be decided which network interface is used for sending out packets. This decision must be determined by an application callback that must be registered at the network stack using this function. NULL resets to default output function.

```
typedef
uint8_t ( *ZWIR_GatewayOutputFunction_t ) ( uint8_t*      data,
                                             uint16_t      size,
                                             ZWIR_PANAddress_t* address );
```

This type defines the signature of functions to be used as the output function in Gateway Mode. The task of this function is determining to which interface an outgoing packet must be sent and calling the corresponding output function for this interface. The decision is typically based on the PAN address of the destination node that is provided in the *address* argument. The *data* argument carries a pointer to the IPv6 packet to be sent; the *size* argument determines the size of this packet.

3.3.5. Miscellaneous

```
void
ZWIR_LocalBroadcast ( uint8_t value )
```

Disable or enable rebroadcasting of broadcast packets to send a local broadcast. If enabled, all broadcast packets will be sent with the maximum hop count so they will not be rebroadcast.

Note: The usual way to use this function is to enable the local broadcast, to send a broadcast packet, and to disable the local broadcast.

```
void
ZWIR_MulticastPreferExistingRepeater ( uint8_t value )
```

Disable or enable sending broadcast packets that will be rebroadcasted by nodes with existing multi-hop routes (repeater) only. If enabled all broadcast packets will only be rebroadcast from devices with existing multi-hop routes. This feature reduces the rebroadcast traffic in larger networks and ensures that existing repeaters are most likely to be used for multi-hop routes.

Note: If all retries of the route discovery and the address resolution fail and the “multicast prefer existing repeater” option is enabled, the route discovery and the address resolution will restart and use all possible hop routes. This ensures communication during the network start up.

```
ZWIR_TRXStatistic_t
ZWIR_GetTRXStatistic ( void )
```

This function returns statistic information about transmission and reception. The returned data structure contains the number of packets, the number of bytes received and transmitted, the number of re-transmissions, and the number of CRC failures on reception. Furthermore, the transmit duty-cycle is included. The counters are reset either on reset or if `ZWIR_ResetTRXStatistic` is called.

Checking the duty cycle should be performed on a regular basis in order to meet the duty cycle requirements at the operation site of the device. Contact the local authorities to find out if duty-cycle limitations apply in the target market(s).

Note: All values in the structure might be higher than expected. This is due to the overhead communication that is required for address resolution and route discovery. Refer to section 2.10 to determine if it is possible to optimize constant settings in order to reduce overhead traffic to a minimum.

```
void
ZWIR_ResetTRXStatistic ( void )
```

This function resets all values of the transceiver statistics to 0. This function has no effect on ongoing transfers.

```
typedef struct {
    uint32_t  txBytes
    uint32_t  txPackets
    uint32_t  rxBytes
    uint32_t  rxPackets
    uint32_t  txFail
    uint32_t  dutyCycle
} ZWIR_TRXStatistic_t
```

This structure is returned by `ZWIR_GetTRXStatistic`. The values contained are counted starting from reset, network reset (initiated by `ZWIR_ResetNetwork`) or a call to `ZWIR_ResetTRXStatistic`. In addition to data sent from the application code, the fields contained in the structure also consider packets that are sent in the background (e.g., route and neighbor discovery). The `dutyCycle` field contains the ratio of time spent sending and the time elapsed since the occurrence of one of the above events. In order to obtain the actual duty cycle percentage, divide `dutyCycle` by 1000.

```
void
ZWIR_SetDutyCycleWarning ( uint32_t          percentage,
                           ZWIR_DutyCycleCallback_t  callback)
```

This function defines a duty cycle threshold for the transmit duty cycle that, if exceeded, triggers a call to the callback function specified with `callback`.

If `percentage` is set to 0 or `callback` is set to NULL, no warning will be triggered. Otherwise, callback will be called once per second as long as the 1h duty cycle exceeds the value defined with the `percentage` argument. The `percentage` parameter defines the duty cycle threshold for triggering duty cycle warnings (the value must be specified as 1/1000 percent, thus `percentage = 1000` sets a threshold of 1%).

The `callback` parameter defines the function that should be called if the duty cycle threshold is exceeded.

```
typedef
void ( * ZWIR_DutyCycleCallback_t ) ( uint32_t  percentage )
```

Function pointer type for the callback function that should be called if the duty cycle warning threshold is exceeded. The `percentage` parameter contains the duty cycle in 1/1000 percent.

```
void
ZWIR_SetPromiscuousMode ( bool enable )
```

This command puts the device into Promiscuous Reception Mode. This means that on the MAC layer, all packet filtering is disabled. In Promiscuous Mode, the device receives packets sent to all PAN identifiers and all PAN addresses, regardless of its own PAN identifier and PAN address configuration. Filters on higher protocol layers are still active.

The Promiscuous Mode should not be used in normal operation. It might be appropriate for gateways, and it is required for sniffers.

```
bool
ZWIR_CreateAlternativeAddressList ( uint16_t size )
```

This function is used in Promiscuous Mode to allocate memory for an alternative PAN address list. The device will treat each packet sent to one of the addresses in the alternative PAN address list in the same way as if the packet had been sent to the device's own PAN address. The *size* argument determines the maximum number of entries in the list. The function returns **true** on success or **false** otherwise.

```
bool
ZWIR_AddAlternativeAddress ( ZWIR_PANAddress_t* address,
                           ZWIR_AlternativeAddressType_t type )
```

This function adds a PAN address to the alternative address list. The *address* argument specifies the address to be added and the *type* argument specifies the type of the address. The function returns **true** if the address was added successfully; **false** is returned if no alternative address list has been allocated (refer to `ZWIR_CreateAlternativeAddressList`). If *address* is already in the address list, **true** is returned. If there is no free item in the alternative address list, the item that has not been used for the longest time is overwritten.

```
ZWIR_AlternativeAddressType_t
ZWIR_IsAlternativeAddress ( ZWIR_PANAddress_t* address,
                          ZWIR_AlternativeAddressType_t type )
```

This function checks whether the *address* of *type* is in the alternative PAN address list or not. *type* is used as filter. The type of address stored in the address list is logically AND'ed with *type*. The function returns the type of the stored address if available or `ZWIR_aatNone` otherwise.

```
typedef enum {...} ZWIR_AlternativeAddressType_t
```

This type is used by `ZWIR_IsAlternativeAddress` and `ZWIR_AddAlterantiveAddress` as address type, address filter, or return value. Possible values are

```
ZWIR_aatNone    0x00    Address not found (only with ZWIR_IsAlternativeAddress)
```

<code>ZWIR_aatEUI64</code>	<code>0x01</code>	Address is a EUI64 address
<code>ZWIR_aatEUI48</code>	<code>0x02</code>	Address is a EUI48 address
<code>ZWIR_aatAny</code>	<code>0x03</code>	Only to be used as a filter in <code>ZWIR_IsAlternativeAddress</code>

bool

`ZWIR_ExternalClockEnable (bool enable)`

This function is used to select whether the external or the internal clock is used as the system clock. The external clock is much more precise, but it is not possible to use Sleep Mode or turn off the transceiver when the external clock is used. The return value indicates whether the transceiver clock is used or not. It is not possible to use the external clock if the transceiver is switched off, so switching to the external clock from `ZWIR_AppInitHardware` will fail.

void

`ZWIR_SetPAControl (ZWIR_PAControl_t pacontrol)`

The ZWIR module provides two pins to control an external power amplifier. These pins are toggled depending on RX or TX. This function will enable or disable the PA pins and adjust the switching lead time.

typedef enum {...} ZWIR_PAControl_t

Enumeration of selectable PA control configurations. Possible values are

<code>ZWIR_paOff</code>	<code>0x00</code>	PA pins are switched off
<code>ZWIR_pa2us</code>	<code>0x10</code>	PA control is on with 2 μ s lead time
<code>ZWIR_pa4us</code>	<code>0x11</code>	PA control is on with 4 μ s lead time
<code>ZWIR_pa6us</code>	<code>0x12</code>	PA control is on with 6 μ s lead time
<code>ZWIR_pa8us</code>	<code>0x13</code>	PA control is on with 8 μ s lead time

char const*

`ZWIR_GetFCCID (void)`

This function returns the FCC ID of the module. The FCC ID is returned as a NULL-terminated string.

3.4. Power Management

The MCU on the module can operate at different clock rates. The lowest possible clock frequency matching the needs of the application should be selected in order to work as power efficiently as possible. The operating system frequency is set using `ZWIR_SetFrequency`.

In addition to the clock speed modification, the API provides the function `ZWIR_SetRTC` to set and `ZWIR_GetRTC` to get the RTC counter. Furthermore the function `ZWIR_SelectRTCSource` selects the RTC clock source. The RTC clock interval is 1s and is suitable for UNIX time. Due to the inaccuracy of the internal RC oscillator, an external crystal should be used for RTC applications. The RTC clock source selection remains unless the backup domain is powered via the VBAT pin. The internal clock is used as the default clock source.

The radio module supports the Sleep, Stop and Standby Modes (see section 2.6.3). The wakeup conditions are adjustable for each power mode individually. By default, the system continues its execution after an RTC alarm. The state of the transceiver depends of the selection of the transceiver interrupt or event as the wakeup source. When the transceiver interrupt or event is masked, the transceiver will be switched off automatically.

All three power modes can be executed immediately or delayed to send out all buffered packets. After entering the Standby Mode, all RAM content is lost and the microcontroller will be reset.

The functions `ZWIR_SetWakeupSource`, `ZWIR_PowerDown` and `ZWIR_AbortPowerDown` are used to configure, enter, or leave the low power modes.

```
void  
  ZWIR_SetFrequency ( ZWIR_MCUFrequency_t  freq )
```

This function sets the clock speed of the MCU core. Peripheral clocks are not changed.

```
void  
  ZWIR_SetRTC ( uint32_t value )
```

This function sets the internal RTC counter to the given value.

```
uint32_t  
  ZWIR_GetRTC ( void )
```

This function reads the current RTC counter and return the value.

```
uint8_t  
  ZWIR_SelectRTCSource ( ZWIR_RTCSource_t source );
```

This function selects the RTC clock source. Possible return values are

- 0 Selection fails. This occurs if the external clock sources is selected but no crystal connected.
- 1 Clock source was successfully changed.
- 2 Clock source was not changed because source was already selected.

```
void  
  ZWIR_AbortPowerDown ( void )
```

This function stops all delayed power down actions.


```
void
ZWIR_PowerDown ( ZWIR_PowerDownState_t powerDownMode,
                 uint32_t time )
```

This function changes the power mode of the system immediately or after sending all buffered fragments. The *powerDownMode* argument defines the next power down mode. The *time* parameter specifies the power down time. For all modes, *time* is given in seconds. If the RTC alarm is not selected as the source, the time parameter will be ignored.

```
void
ZWIR_SetWakeupSource ( ZWIR_PowerDownState_t powerDownMode,
                      uint64_t wakeupSource )
```

This function sets the wakeup condition for a power mode. The *powerDownMode* argument defines the power down mode to be configured and the *wakeupSource* parameter specifies the event(s) that will cause the module to enter active mode again. Depending on the value of *powerDownMode*, the *wakeupSource* parameter is interpreted differently. The settings being applied to different registers will be revoked when exiting power down mode.

In Sleep Mode, each interrupt can be used to wake up the system. Accordingly, the *wakeupSource* parameter is interpreted as an interrupt mask. The interrupt mask allows selecting one or more of the lower 64 interrupts to be selected as a wakeup source. The bits correspond to the interrupt position according to the Nested Vectored Interrupt Controller (NVIC) documentation in the *STM32 Reference Manual*.

For Stop Mode, only events are supported as a wakeup source. Therefore, the *wakeupSource* parameter is used to configure the external interrupt/event controller's event mask register (EXTI_EMR). The external interrupt controller limits the wakeup sources for Stop Mode to the external pins, the programmable voltage detector, the real-time clock, and the USB wakeup function. Refer to the *STM32 Reference Manual* for further information about the external interrupt controller.

Exit from Standby Mode is only possible using the real-time clock (RTC) or the external wakeup pin. Wakeup by RTC is selected if '1' is passed as *wakeupSource*, the WKUP pin is selected by '2', and an argument of '3' selects both.

If an invalid wakeup source is selected, the default wakeup source, which is the RTC, is set.

```
void
ZWIR_Sleep ( uint16_t sleepTime )
```

This function puts the system into Sleep Mode. The *sleepTime* argument controls the duration of Sleep Mode and is given in 10ms multiples. This means that a *sleepTime* value of 100 puts the system into Sleep Mode for 1 second. During Sleep Mode, all memory contents are retained. Waking up the system from sleep is not possible.

Note: This function is deprecated – use *ZWIR_PowerDown* instead.

```
void
ZWIR_Standby ( uint32_t  standbyTime )
```

This function puts the system into Standby Mode. The *standbyTime* argument controls the duration of Standby Mode and is given in seconds. In order to consume minimal power, almost all power domains of the MCU are disconnected. Memory contents are not retained during Standby Mode. The system can be awakened before expiration of the standby timer if the external wakeup pin of the module is triggered.

Note: This function is deprecated – use `ZWIR_PowerDown(ZWIR_pStandby, standbyTime)` instead

```
void
ZWIR_TransceiverOff ( void )
```

This function switches the transceiver off manually. Attempts to send data using one of the functions `ZWIR_SendUDP`, `ZWIR_SendUDP2` and `ZWIR_Send6LoWPAN` while the transceiver is switched off will fail and the sent data will be lost. To turn the transceiver on, the functions `ZWIR_TransceiverOn` or `ZWIR_ResetNetwork` can be used. The transceiver is re-enabled after a system reset.

```
void
ZWIR_TransceiverOn ( void )
```

This function switches the radio transceiver on manually after having it switched off using `ZWIR_TransceiverOff`. If the transceiver is already active, this function does nothing.

```
typedef
enum { ... } ZWIR_PowerDownState_t
```

MCU frequency enumeration type accepted by `ZWIR_PowerDown` and `ZWIR_SetWakeupSource`. Possible values include

Value	Power Mode	Wait Until Sent?	Possible Wakeup Sources
<code>ZWIR_pSleep</code>	Sleep	No	IRQ 0 to 63
<code>ZWIR_pSleepAfterActivities</code>	Sleep	Yes	IRQ 0 to 63
<code>ZWIR_pStop</code>	Stop	No	EXTI 0 to 18
<code>ZWIR_pStopAfterActivities</code>	Stop	Yes	EXTI 0 to 18
<code>ZWIR_pStandby</code>	Standby	No	RTC, wakeup pin (WKUP pin)
<code>ZWIR_pStandbyAfterActivities</code>	Standby	Yes	RTC, wakeup pin (WKUP pin)

```
typedef
enum {...} ZWIR_RTCSource_t
```

RTC source enumeration type accepted by `ZWIR_selectRTCSource`. Possible values are

<code>ZWIR_rIntern</code>	0x00	Use the internal RC oscillator as RTC clock source
<code>ZWIR_rExtern</code>	0x01	Use an external crystal as RTC clock source

```
typedef
enum { ... } ZWIR_MCUFrequency_t
```

MCU frequency enumeration type accepted by `ZWIR_setFrequency`. Possible values include

<code>ZWIR_mcu8MHz</code>	8 MHz
<code>ZWIR_mcu16MHz</code>	16 MHz
<code>ZWIR_mcu32MHz</code>	32 MHz
<code>ZWIR_mcu64MHz</code>	64 MHz
<code>ZWIR_mcu8MHzLowPower</code>	8 MHz with disabled PLL
<code>ZWIR_mcuUserFrequency</code>	Customer MCU clock settings

3.5. Firmware Version Information

Each productive firmware version *MUST* include a valid set of version information. The complete set consists of major and minor version number, version extension, vendor ID, and product ID. For more detailed information about these elements refer to section 2.7.

Version information is included in the firmware by global definition of the variables listed below. If these variables are not defined by the application code, they will contain default values.

```
uint32_t ZWIR_vendorID = 0xe966
```

This variable defines the Vendor ID. A vendor ID is assigned by IDT. It *MUST* be defined appropriately in production code!

```
uint16_t ZWIR_productID = 0
```

This variable identifies a product or firmware type, respectively. It *MUST* be defined appropriately for each firmware type. It is used by the firmware over-the-air update mechanism to distinguish different firmware types.

```
uint8_t ZWIR_firmwareMajorVersion = 0
```

This variable defines the major version number of the firmware.

```
uint8_t ZWIR_firmwareMinorVersion = 0
```

This variable defines the minor version number of the firmware.

```
uint16_t ZWIR_firmwareVersionExtension = 0
```

This defines version extension information of the firmware.

3.6. Properties and Parameters

The behavior of the built-in network stack functionality is configurable to some extent by a set of parameters that are changed using the function `ZWIR_SetParameter`.

```
int32_t
ZWIR_SetParameter ( ZWIR_SystemParameter_t parameter,
                   int64_t value )
```

This function changes the setting of a single network stack parameter. Configuration changes are effective immediately when this function is called from `ZWIR_AppInitHardware`. Otherwise, the new value is buffered until `ZWIR_ResetNetwork` is called.

```
int64_t
ZWIR_GetParameter ( ZWIR_SystemParameter_t parameter )
```

This function queries the current value of a stack parameter. The function returns the parameter value that is currently in effect. Thus, if `ZWIR_GetParameter` is called after `ZWIR_SetParameter` but before `ZWIR_ResetNetwork` has been called, the previous parameter value is returned and not the value set with `ZWIR_SetParameter`.

```
typedef enum { ... } ZWIR_SystemParameter_t
```

This enumeration names the different parameters that can be configured using `ZWIR_SetParameter`. Possible names are listed in Table 3.1 below:

Table 3.1 Configurable Stack Parameters and Their Default Values

Enumerator	Size	Default	Description
<code>ZWIR_spRoutingTableSize</code>	1	8	Refer to section 2.10.3.
<code>ZWIR_spNeighborCacheSize</code>	1	8	Refer to section 2.9.3.

Enumerator	Size	Default	Description
ZWIR_spMaxSocketCount	1	8	Refer to section 2.9.2.
ZWIR_spRouteTimeout	2	3600 (s)	Refer to section 2.10.3.
ZWIR_spNeighborReachableTime	2	3600 (s)	Refer to section 2.9.3.
ZWIR_spMaxHopCount	1	4	Refer to section 2.10.3.
ZWIR_spRouteMaxFailCount	1	3	Refer to section 2.10.3.
ZWIR_spRouteRequestMinLinkRSSI	1	-128 (dBm)	Refer to section 2.10.3.
ZWIR_spRouteRequestMinLinkRSSIReduction	1	0 (dB)	Refer to section 2.10.3.
ZWIR_spDoDuplicateAddressDetection	1	1	Refer to section 2.8.4.
ZWIR_spDoRouterSolicitation	1	1	Refer to section 2.8.3.
ZWIR_spRouteRequestAttempts	1	4	Refer to section 2.10.3.
ZWIR_spHeaderCompressionContext1	8	0	Refer to section 2.9.4.
ZWIR_spHeaderCompressionContext2	8	0	Refer to section 2.9.4.
ZWIR_spHeaderCompressionContext3	8	0	Refer to section 2.9.4.
ZWIR_spNeighborRetransTime	1	3000 (ms)	Refer to section 2.9.3.
ZWIR_spDoAddressAutoConfiguration	1	1	Refer to section 2.8.3.

3.7. Error Codes

The error codes listed in Table 3.2 are generated by the core library and passed to the `ZWIR_Error` hook if it is implemented in the application code.

Table 3.2 Error Codes Generated by the Core Library

C – Identifier	Code	Default Handling
ZWIR_eDADFailed	100 _{HEX}	Node shutdown (permanent deep-sleep, node can only be restarted with external reset).
ZWIR_eProgExit	101 _{HEX}	Node shutdown (permanent deep-sleep, node can only be restarted with external reset).
ZWIR_eReadMACFailed	102 _{HEX}	System reset triggered.
ZWIR_eMemoryExhaustion	103 _{HEX}	System reset triggered. Important: This error is only triggered when allocation fails for the memories required by the network stack. Failing allocation attempts from the application code must be detected by checking the allocation result for NULL!
ZWIR_eHostUnreachable	104 _{HEX}	Ignore – the packet causing this failure is dropped.
ZWIR_eExtClockPowerDown	105 _{HEX}	Ignore – the node will not enter power-down mode.
ZWIR_eRouteFailed	106 _{HEX}	Ignore – the node received information about a broken route.

4 UART Library Reference

The *libZWIR451x-UART1.a* and *libZWIR451x-UART2.a* libraries provide functions for easy access to the UART interfaces provided by the microcontroller. Each ZWIR451x module has two UART interfaces. IDT provides two separate libraries, one for each UART interface. Both libraries expose exactly the same interface. Symbol names only differ in the number after “UART,” which is part of each symbol name. Simply replace the question mark in the following function and type names with 1 or 2 depending on which UART is being used.

It is possible to use only one of the UARTs or both in parallel. Consider the different priorities of the interfaces when selecting the UART (refer to section 2.6.2). The initialization of a UART interface is performed automatically if the UART library is linked into the project. The UARTs can be used after hardware initialization but not inside of `ZWIR_AppInitHardware`.

Note: If UART I/Os are configured by the user application inside `ZWIR_AppInitHardware`, the UART library retains the configured GPIOs without changes.

The UART’s receive buffer is 256 bytes. If data are received on the UART interface, data are kept in the buffer until read by `ZWIR_UART?_ReadByte`. If the buffer is full and more data are received, `ZWIR_Error` is called with `ZWIR_UART1_eOvf1` as the argument. The UART libraries use an event-based programming approach. Instead of relying on polling the UART interfaces, a callback function must be specified, which is called automatically when data is available in the receive buffer. This is done using the function `ZWIR_UART?_SetRXCallback`. If this function is not called, the UART receiver remains disabled, saving some power.

4.1. Symbol Reference

`bool`

```
ZWIR_UART?_SendByte ( uint8_t data )
```

Single bytes are sent via the interface using this function. The byte to be sent is provided in *data*. The function returns *true* if the byte was successfully placed in the transmit buffer or *false* otherwise. If bytes are in the buffer, they are written immediately until the buffer is empty. However, note that significant time can elapse between a call to `ZWIR_UART?_SendByte` and the actual sending of the byte if there is already data in the buffer.

`uint8_t`

```
ZWIR_UART?_Send ( uint8_t* data,
                  uint16_t dataSize )
```

A block of bytes is written to the buffer and transmission is started. The *data* argument must point to the data to be written. *dataSize* determines the number of bytes to be transferred. The return value contains the number of bytes that have actually been written. It can be lower than *dataSize*.

```
bool
ZWIR_UART?_ReadByte ( uint8_t* data )
```

This function reads a single byte from the receive buffer. The read byte is stored to the location to which *data* points. The function returns *true* if a byte is successfully read and *false* otherwise.

```
void
ZWIR_UART?_SetRXCallback ( ZWIR_UART_RXCallback_t callback )
```

This function registers a callback function that is called if data is received on the UART. The *callback* argument is a pointer to the function to be called. The UART receiver is not started until *ZWIR_UART?_SetRXCallback* is called. If NULL is passed as *callback* argument, the receiver is disabled.

```
void
ZWIR_UART?_Setup ( uint32_t baudRate,
                  uint32_t parameters )
```

This function configures the UART peripheral of the microcontroller. The *baudRate* argument configures the speed of the transmission line. Its value must be given in bits per second. The *parameters* argument is a set of flags that configure the parity bit generation, the stop bit generation, and controls whether flow control is used or not. The parameters argument is generated from a binary OR combination of one constant from each block described below. Default values can be omitted. In order to set all values to their default settings, a zero can be passed in the *parameters* argument.

The following configuration options are available for parity configuration:

ZWIR_UART_NoParity	No parity bit is transmitted (default).
ZWIR_UART_OddParity	Odd parity bit is transmitted/checked.
ZWIR_UART_EvenParity	Even parity bit is transmitted/checked.

The following configuration options are available for stop-bit configuration:

ZWIR_UART_Stop_1	One stop-bit is transmitted at the end of a frame (default).
ZWIR_UART_Stop_2	Two stop bits are transmitted at the end of a frame.

The following configuration options are available for flow-control configuration:

ZWIR_UART_NoFlowControl	Flow control is disabled (default).
ZWIR_UART_HWFlowControl	Use hardware flow control with the CTS and RTS pins.

Note: A call to this function drops all bytes that are still in the transmission buffer. If this function is called during an active transmission, the active transmission is very likely to fail.

Note: When flow control is enabled, the configuration of the CTS and RTS pins of the corresponding UART interface are configured as input and alternative push/pull output, respectively. This configuration overwrites the existing configuration of these pins. When flow control changes from enabled to disabled, the pin configuration of the CTS and RTS pins is not changed.

```
bool
ZWIR_UART?_IsTXEmpty ( void )
```

This function returns *false* if there are still bytes in the UART transmit buffer and returns *true* otherwise.

```
uint16_t
ZWIR_UART?_GetAvailableTXBuffer ( void )
```

This function returns the number of free bytes in the UART TX buffer.

```
typedef
void ( * ZWIR_UART_RXCallback_t ) ( void )
```

This is the type definition for a UART callback function. This type is used with `ZWIR_UART?_setRXCallback`. The callback does not accept any elements and does not return any.

```
ZWIR_UART?_PRINTF
```

This macro is provided for convenience. If high-level functions for text output like `printf` are used, an appropriate low-level function must be provided that can output characters to a device. This macro defines a low-level output function writing to the UART interface. This macro must only be used once in the whole application source code. It must not be put inside of function definitions and should not be put in header files. It is not possible to use both, `ZWIR_UART1_PRINTF` and `ZWIR_UART2_PRINTF` in the same project.

4.2. Custom UART I/O Configuration

During system startup, the UART library configures the GPIOs appropriately to provide the configured functionality. UART output pins (i.e., TX and RTS if flow control is enabled) are configured as 2MHz push/pull alternative outputs; input pins are configured as floating inputs. If this configuration is not appropriate for the application, an alternative configuration can be supplied in the `ZWIR_AppInitHardware` hook. If an alternative configuration is supplied the UART does not overwrite this configuration.

Note: It is required that all pins are configured appropriately. The UART library does not configure any pin if one of the UART pins is not in its default state.

Note: Inappropriate manual pin configurations could lead to a non-functional UART.

4.3. Error Codes

The error codes listed in Table 4.1 are generated by the UART libraries and passed to the `ZWIR_Error` hook if it is implemented in the application code.

Table 4.1 Error Codes Generated by the UART Libraries

C – Identifier	Code	Default Handling
libZWIR451x-UART1.a		
<code>ZWIR_UART1_eOvfl</code>	210 _{HEX}	Ignore
<code>ZWIR_UART1_eParity</code>	211 _{HEX}	Ignore
<code>ZWIR_UART1_eFrame</code>	212 _{HEX}	Ignore
<code>ZWIR_UART1_eNoise</code>	213 _{HEX}	Ignore
libZWIR451x-UART2.a		
<code>ZWIR_UART2_eOvfl</code>	220 _{HEX}	Ignore
<code>ZWIR_UART2_eParity</code>	221 _{HEX}	Ignore
<code>ZWIR_UART2_eFrame</code>	222 _{HEX}	Ignore
<code>ZWIR_UART2_eNoise</code>	223 _{HEX}	Ignore

5 GPIO Library Reference

The GPIO library provides a convenient interface for accessing and controlling the GPIO ports of the module. It allows configuring GPIOs to be used as application programmable inputs or outputs, and it is possible to enable or disable special functions of certain ports such as the JTAG ports. All functions from the GPIO library are also accessible using the MCU's peripheral control registers.

The intention of the GPIO library is to provide a high-level, lightweight, convenient interface to the GPIOs. For that reason, the GPIO library functions do not implement parameter checking. It is the responsibility of the application to ensure that appropriate parameters are used. All microcontroller GPIO pins that are not connected to one of the modules I/O pins are locked so that it is impossible to change their configuration accidentally.

5.1. Symbol Reference

```
void
ZWIR_GPIO_ConfigureAsOutput ( ZWIR_GPIO_Pin_t           pins,
                              ZWIR_GPIO_DriverStrength_t driver,
                              ZWIR_GPIO_OutputMode_t    mode )
```

This function registers one or multiple pins as output. The *pins* argument specifies the pin to be configured. If multiple pins are to be configured, provide a binary OR'ed combination of the enumeration values corresponding to the pins. The *driver* argument determines the driving strength of the pin; the *mode* argument determines the output mode. If multiple pins are specified, all pins will be configured in the same way.

Be sure to use a combination of `ZWIR_GPIO_Pin_t` enumeration values to specify which pins are to be configured. Otherwise the wrong pins might be configured, resulting in a configuration that could cause damage to the system.

```
void
ZWIR_GPIO_ConfigureAsInput ( ZWIR_GPIO_Pin_t           pins,
                              ZWIR_GPIO_InputMode_t    mode )
```

This function registers one or multiple pins as input. The *pins* argument specifies the pin to be configured. If multiple pins are to be configured, provide a binary OR'ed combination of the enumeration values corresponding to the pins. The *mode* argument selects the configuration of the inputs. If multiple pins are specified, all pins will be configured in the same way.

Important: Use a combination of `ZWIR_GPIO_Pin_t` enumeration values to specify which pins are to be configured. Otherwise the system might not behave as expected.

```
bool
ZWIR_GPIO_Read ( ZWIR_GPIO_Pin_t pin )
```

This function reads the input value of a single pin. The function does not care if the pin is configured as input or output pin. If the pin is configured as analog input, the return value is undefined.

```
uint32_t
ZWIR_GPIO_ReadMultiple ( ZWIR_GPIO_Pin_t pins )
```

This function reads the input value of multiple pins. For this function, it does not matter if the *pins* are configured as inputs or outputs. If a pin is configured as an analog input, the return value at the corresponding bit is undefined. The result is aligned as shown in Figure 5.1 and Figure 5.2. In order to extract single bit results, the return value of the function can be binary OR'ed with the *ZWIR_GPIO_Pin_t* enumeration values.

Figure 5.1 ZWIR_GPIO_ReadMultiple Result Alignment in ZWIR4512AC1 Devices





Figure 5.2 ZWIR_GPIO_ReadMultiple Result Alignment in ZWIR4512AC2 Devices



Note: Only pins from the same GPIO bank are read at the same time. If pins do not share the same GPIO bank, there will be a time difference between the accesses to their input registers. All pins belonging to the same GPIO bank are highlighted with the same color in Figure 5.1 and Figure 5.2.

```
void
ZWIR_GPIO_Write ( ZWIR_GPIO_Pin_t pins,
                  bool value )
```

This function sets the output value of one or multiple pins. All pins specified in the *pins* argument are set to *value*.

Important: The output is written regardless of the pin configuration! If one of the pins was configured as a pull-up or pull-down input, writing to the output register of this pin can change the pull-up/pull-down configuration accidentally!

```
void
ZWIR_GPIO_Remap ( ZWIR_GPIO_RemapFunction_t function,
                  int32_t value )
```

This function is used to control the remapping of GPIO pins to system functions. It allows configuring whether the JTAG pins are used for debugging purposes or as normal GPIO pins. The *value* argument must be one of the options defined by the enumeration type `ZWIR_GPIO_SWJRemapValue_t`.

```
typedef enum { ... } ZWIR_GPIO_Pin_t
```

This enumeration type assigns a name to each pin. Some GPIO operations allow specifying multiple pins. This is done by OR-combining the pins.

Enumeration	MCU I/O Port	
	ZWIR4512AC1	ZWIR4512AC2
ZWIR_Pin1	A7	A7
ZWIR_Pin2	A6	A6
ZWIR_Pin3	A5	A5
ZWIR_Pin4	A4	A4
ZWIR_Pin5	A3	A3
ZWIR_Pin6	A2	A2
ZWIR_Pin7	A1	A1
ZWIR_Pin8	A0	A0
ZWIR_Pin9	C13	C13

ZWIR_Pin12	A10	-
ZWIR_Pin13	A9	A10
ZWIR_Pin14	-	A9
ZWIR_Pin16	A11	-
ZWIR_Pin17	A12	C14
ZWIR_Pin18	-	C15
ZWIR_Pin19	B3	A11
ZWIR_Pin20	A13	A12
ZWIR_Pin21	A15	B3
ZWIR_Pin22	A14	A13
ZWIR_Pin23	B7	A15
ZWIR_Pin24	B6	A14
ZWIR_Pin25	-	B7
ZWIR_Pin26	-	B6

```
typedef enum { ... } ZWIR_GPIO_DriverStrength_t
```

This enumeration value specifies the driving strength of GPIO output pins.

ZWIR_GPIO_dsLow	Low driving strength
ZWIR_GPIO_dsMedium	Medium driving strength
ZWIR_GPIO_dsHigh	High driving strength

```
typedef enum { ... } ZWIR_GPIO_OutputMode_t
```

This enumeration value specifies the mode of GPIO output pins.

ZWIR_GPIO_omPushPull	Application controlled push/pull output
ZWIR_GPIO_omOpenDrain	Application controlled open drain output
ZWIR_GPIO_omAlternativePushPull	Hardware controlled push/pull output
ZWIR_GPIO_omAlternativeOpenDrain	Hardware controlled open drain output

```
typedef enum { ... } ZWIR_GPIO_InputMode_t
```

This enumeration value specifies the mode of GPIO input pins.

<code>ZWIR_GPIO_imAnalog</code>	Analog input (default configuration)
<code>ZWIR_GPIO_imFloating</code>	Floating input
<code>ZWIR_GPIO_imPullUp</code>	Pull-up input
<code>ZWIR_GPIO_imPullDown</code>	Pull-down input

```
typedef enum { ... } ZWIR_GPIO_RemapFunction_t
```

This enumeration type is used to specify which remapping is to be changed with `ZWIR_GPIO_Remap`.

<code>ZWIR_GPIO_rfSWJ</code>	Configure remapping of the JTAG/SWD pins (Pin19 – Pin 22)
------------------------------	---

```
typedef enum { ... } ZWIR_GPIO_SWJRemapValue_t
```

In calls to `ZWIR_GPIO_Remap`, this enumeration value specifies which configuration is used for JTAG/SWD remapping.

<code>ZWIR_GPIO_swjrEnableSWJ</code>	Enable full JTAG/SWD support
<code>ZWIR_GPIO_swjrSWOnly</code>	Enable SWD support only
<code>ZWIR_GPIO_swjrDisableSWJ</code>	Disable JTAG/SWD support

6 IPsec Library Reference

The IPsec library provides functions to manage security policies and security associations. A security policy is added using `ZWIRSEC_AddSecurityPolicy`. Security Associations can be added using the function `ZWIRSEC_AddSecurityAssociation`. For more detailed information about security policies and security associations, refer to the *ZWIR451x Application Note – Using IPsec and IKEv2 in 6LoWPANs*.

6.1. Symbol Reference

`uint8_t`

```
ZWIRSEC_AddSecurityPolicy ( ZWIRSEC_PolicyType_t      type,
                           ZWIR_IPv6Address_t*      remoteAddress,
                           uint8_t                 prefix,
                           ZWIR_Protocol_t         proto,
                           uint16_t                 lowerPort,
                           uint16_t                 upperPort,
                           ZWIRSEC_SecurityAssociation_t* securityAssociation )
```

A call to this function adds a security policy to the IPsec security policy database. The *type* argument determines the traffic direction and how packets must be handled. The combination of the *remoteAddress*, *prefix*, *protocol*, *lowerPort*, and *upperPort* arguments specify the traffic that is affected by this policy. See section 2.12 and the *ZWIR451x Application Note – Using IPsec and IKEv2 in 6LoWPANs* for more details. The function returns the security policy index of the newly created security policy. If there is an error, `FFHEX` is returned.

The last argument specifies the security association to be used by this policy. A security association must be created using `ZWIRSEC_AddSecurityAssociation` before `ZWIRSEC_AddSecurityPolicy` is called. If IKEv2 should be used for generating the security association automatically, pass `NULL` as the *securityAssociation* argument.

`void`

```
ZWIRSEC_RemoveSecurityPolicy ( uint8_t spi )
```

This function removes the security policy with the index *spi* from the security policy database. If no index is stored at *spi*, the function does nothing.

`ZWIRSEC_SecurityAssociation_t*`

```
ZWIRSEC_AddSecurityAssociation ( uint32_t      securityParamIdx,
                                uint8_t       replayCheck,
                                ZWIRSEC_EncryptionSuite_t* encSuite,
                                ZWIRSEC_AuthenticationSuite_t* authSuite )
```

This function adds a security association to the security association database manually. Use this function before calling `ZWIRSEC_AddSecurityPolicy` if not using IKEv2 for automatic key exchange.

The *securityParamIdx* argument is a unique number identifying the security association. The *encSuite* and *authSuite* parameters specify the encryption and authentication algorithms and keys. The *replayCheck* parameter determines if replay checks are performed for this security association. The internal replay check counter can only be reset by re-creating the specific security association. Refer to section 2.12.1 and the links for `ZWIRSEC_EncryptionSuite_t` and `ZWIRSEC_AuthenticationSuite_t` for more details.

The function returns a pointer to the security association descriptor if it was created successfully. If there is an error, the function returns NULL.

```
void
ZWIRSEC_RemoveSecurityAssociation ( ZWIRSEC_SecurityAssociation_t* sa )
```

This function removes the security association pointed to by *sa*.

```
typedef enum { ... } ZWIRSEC_PolicyType_t
```

IPSec policy type enumeration. Possible values include

<code>ZWIRSEC_ptOutputApply</code>	Secure outbound traffic with IPSec
<code>ZWIRSEC_ptOutputBypass</code>	Bypass outbound traffic unsecured
<code>ZWIRSEC_ptOutputDrop</code>	Drop outbound traffic
<code>ZWIRSEC_ptInputApply</code>	Unsecure inbound traffic with IPSec
<code>ZWIRSEC_ptInputBypass</code>	Bypass inbound traffic unsecured
<code>ZWIRSEC_ptInputDrop</code>	Drop inbound traffic

```
typedef enum { ... } ZWIR_Protocol_t
```

IPSec protocol enumeration. Possible values include

<code>ZWIR_protoAny</code>	Apply IPSec rule to any protocol
<code>ZWIR_protoTCP</code>	Apply IPSec rule to TCP packets only
<code>ZWIR_protoUDP</code>	Apply IPSec rule to UDP packets only
<code>ZWIR_protoICMPv6</code>	Apply IPSec rule to ICMPv6 packet only


```
typedef struct {
    ZWIRSEC_EncryptionAlgorithm_t  algorithm
    uint8_t                        key   [ 16 ]
    uint8_t                        nonce [  4 ]
} ZWIRSEC_EncryptionSuite_t
```

This structure carries all encryption related information. It is used to pass encryption information to `ZWIRSEC_AddSecurityAssociation`.

```
typedef struct {
    ZWIRSEC_AuthenticationAlgorithm_t  algorithm
    uint8_t                            key [ 16 ]
} ZWIRSEC_AuthenticationSuite_t
```

This structure carries all authentication related information. It is used to pass authentication information to `ZWIRSEC_AddSecurityAssociation`.

```
typedef void*  ZWIRSEC_SecurityAssociation_t
```

Objects of this type are returned by `ZWIRSEC_AddSecurityAssociation`. They are passed to `ZWIRSEC_AddSecurityPolicy`.

```
typedef enum { ... } ZWIRSEC_EncryptionAlgorithm_t
```

Enumeration of algorithms available for encryption; possible values include

<code>ZWIRSEC_encNull</code>	no encryption
<code>ZWIRSEC_encAESCTR</code>	AES [†] Counter Mode based encryption

```
typedef enum { ... } ZWIRSEC_AuthenticationAlgorithm_t
```

Enumeration of algorithms available for authentication; possible values include

<code>ZWIRSEC_authNull</code>	no authentication
<code>ZWIRSEC_authAESXCBC96</code>	Extended AES128 CBC [‡] Mode based authentication.

[†] AES – Advanced Encryption Standard
[‡] CBC – Cyclic Block Cipher

6.2. Error Codes

The error codes listed in Table 6.1 are generated by the IPsec libraries and passed to the `ZWIR_Error` hook if it is implemented in the application code. `ZWIRSEC_eDroppedICMP` indicates that an ICMP packet was dropped due to an IPsec rule and `ZWIRSEC_eDroppedPacket` indicates that any other (non-ICMP) packet was discarded due to an IPsec rule. `ZWIRSEC_eUnknownSPI` indicates that an IPsec packet was received but no associated security association was found. With active replay check, `ZWIRSEC_eReplayedPacket` indicates a replayed packet.

IPsec indicates authentication vector mismatches (corrupted packet) with `ZWIRSEC_eCorruptedPacket`.

Table 6.1 Error Codes Generated by the IPsec Libraries

C – Identifier	Code	Default Handling
libZWIR45xx-IPsec.a		
<code>ZWIRSEC_eDroppedICMP</code>	500 _{HEX}	Ignore
<code>ZWIRSEC_eDroppedPacket</code>	501 _{HEX}	Ignore
<code>ZWIRSEC_eUnknownSPI</code>	502 _{HEX}	Ignore
<code>ZWIRSEC_eReplayedPacket</code>	503 _{HEX}	Ignore
<code>ZWIRSEC_eCorruptedPacket</code>	504 _{HEX}	Ignore

7 IKEv2 Library Reference

IKEv2 is used for IPsec key management. Using IKEv2, it is possible to limit the lifetime of a security association and automatically regenerate it with new keys. In order to add IKEv2 functionality to an application, the IKEv2 library must be linked into the project. If this is done, the IKEv2 daemon is automatically registered as the key management engine for IPsec. All IPsec SAs, negotiated with IKEv2, have activated replay checking.

The only task to be performed by the application is adding the suitable authentication entries to the IKEv2 authentication database. This is done using the `ZWIRSEC_AddIKEAuthenticationEntry` function.

7.1. Symbol Reference

`uint8_t`

```
ZWIRSEC_AddIKEAuthenticationEntry (  ZWIR_IPv6Address_t*  remoteAddress,
                                     uint8_t              prefixLength,
                                     uint8_t*             id,
                                     uint8_t              idLength,
                                     uint8_t*             presharedKey )
```

Calling this function adds an authentication entry to the IKE authentication database. The *remoteAddress* argument contains the IPv6 address of the remote device; *prefixLength* contains the prefix length of *remoteAddress*. The device identifier is given in *id*. Its length is specified in *idLength*.

The *presharedKey* argument carries a pointer to the pre-shared key that is used for authentication.

The function returns *true* on success or *false* otherwise. A *false* return indicates there is no room in the authentication database.

`uint8_t ZWIRSEC_ikeRetransmitTime = 10`

This is a weak constant defining how many seconds IKE waits for a reply before retransmission is initiated. The time should be long enough to enable IKE processing at the receiver. This value largely depends on the clock frequency. Set the value accordingly. The predefined value of 10 seconds is only suitable for a receiver clock frequency of 32MHz or 64MHz. The value can be redefined by definition of the variable `ZWIRSEC_ikeRetransmitTime` with an appropriate value in the application code.

`uint32_t ZWIRSEC_ikeRekeyTime = 86400`

This is a weakly defined variable that controls the interval at which the IKE connection must be rekeyed. The default setting corresponds to one week. In order to change this value, the variable `ZWIRSEC_ikeRekeyTime` must be defined with an appropriate value in the application code.

```
uint32_t ZWIRSEC_ikeSARekeyTime = 604800
```

This weakly defined variable controls the interval during which security associations remain valid before rekeying is required. The default setting corresponds to one day. In order to change this value, the variable `ZWIRSEC_ikeSARekeyTime` must be defined with an appropriate value in the application code.

7.2. Library Parameters

Table 7.1 shows a summary of IDT's IKEv2 library parameters and properties.

Table 7.1 Overview of IKEv2 Library Parameters and Properties

Property	Value
Authentication database size	5
Number of sockets used	2
Parameter	Value
<code>ZWIRSEC_ikeSARekeyTime</code>	604800
<code>ZWIRSEC_ikeRekeyTime</code>	86400
<code>ZWIRSEC_ikeRetransmitTime</code>	10s

8 Over-the-Air Update Library

This library implements the Firmware Over-the-Air Update functionality. The only function exported from this library is used to register the Over-the-Air Update daemon in the system.

8.1. Library Reference

```
void
ZWIR_OTAU_Register ( uint16_t localPort )
```

Register the Over-the-Air Update daemon and configure the UDP port on which the daemon is listening.

```
void
ZWIR_OTAU_Status ( ZWIR_OTAU_StatusCode_t status )
```

This hook is provided for debugging purposes. If it is implemented in the application code, each OTAU event will be reported through this function, providing progress and error indications. Refer to the list of possible `ZWIR_OTAU_StatusCode_t` values for more detailed information about event indicators.

```
typedef enum { ... } ZWIR_OTAU_StatusCode_t
```

This enumeration defines event indicators that are supplied to the `ZWIR_OTAU_Status` hook. Possible values include

ZWIR_sInvalidPacketHeader

An OTAU packet with invalid header contents was received. Possible reasons are OTAU packet is too short; Product or Vendor ID of the OTAU packet does not match the corresponding device IDs; or Firmware Version of the OTAU packet is not higher than the installed version.

ZWIR_sUpdateInProgress

An OTAU packet was received after the update execution countdown has been started.

ZWIR_sInvalidPacketCRC

An OTAU packet with an invalid packet CRC was received.

ZWIR_sUnknownPacketType

A packet with an invalid packet type has been received on the OTAU socket.

ZWIR_sInvalidFragmentSize

Fragment size is not a multiple of the page size.

ZWIR_sInvalidDataPacket

OTAU packet does not belong to the same firmware version as the ongoing OTAU, or fragment is not inside the update segment.

ZWIR_sFragmentWriteError

Fragment could not be written to flash.

ZWIR_sInvalidCRCPacket

OTAU packet does not belong to the same firmware version as the ongoing OTAU, or fragment is not inside the CRC segment.

ZWIR_sCRCWriteError

CRC could not be written to flash.

ZWIR_sInvalidExecutePacket

OTAU packet does not belong to the same firmware version as the ongoing OTAU, or the packet contains an invalid page count.

ZWIR_sFirmwareImageVerifyFailed

Integrity of the new firmware could not be verified.

ZWIR_sInitNewFirmwareDone

Preparation for the new firmware is done, and the second flash segment was erased successfully.

ZWIR_sWriteFragmentDone

Fragment was successfully written to flash.

ZWIR_sWriteCRCDone

CRCs were successfully written to flash.

ZWIR_sFirmwareImageVerifyDone

Integrity of the new firmware was verified successfully.

ZWIR_sScheduleUpdate

Update execution was scheduled.

ZWIR_sStartUpdate

Update execution was started.

8.2. Inclusion of the OTAU Library

Including the OTAU library requires special care to ensure that the OTAU will remain functional with different firmware versions, which might be compiled with different compiler versions. The compiled binary program is divided into two sections: the first section that replaces the old firmware with the new one and a second section that handles the OTAU packet processing and storing the received data in the device's flash memory. The first section will remain in the device unchanged over the device's lifetime. Therefore this section is referred to as the invariant section. The second section is replaced with each update, allowing the addition of new functionalities to the OTAU.

The binary data of the first section *MUST* reside at the beginning of the flash memory. For that purpose, the linker script defines the following sections: *.update_code*, *.interface_seg* and *.status_seg*. The location, order, and contents of these sections in the linker script *MUST NOT* be changed! The *.update_code* section contains the code for replacing the old with the new firmware image. The *.status_seg* section is a reserved flash area for storing OTAU status information. The *.interface_seg* section contains a list of pointers into the invariant code section. This is required to ensure that the OTAU functions for firmware versions compiled with different compiler versions.

For the parameter configuration of the application, take into account that the OTAU will consume one of the available network sockets for reception of OTAU data. Consequently, when defining the system parameter *ZWIR_spMaxSocketCount*, this must be considered.

For production code, it is highly recommended that security policies be defined, protecting the UDP port used by the OTAU to prevent devices from being hacked into and made dysfunctional by external attackers. If no protection is installed, intruders who are aware of the network could replace the device's firmware with their own firmware.

8.3. Error Codes

During module startup, the OTAU library performs different checks to verify whether the firmware is configured and linked appropriately. If one of these checks fails, the corresponding error is reported through **ZWIR_Error** and the OTAU functionality is disabled. The error code **ZWIR_eInvalidVID** is reported if the Vendor ID assigned in the firmware is invalid. Refer to the sections 2.13 and 3.5 for further information. The error code **ZWIR_eInvalidOTAUInterface** typically indicates a linking problem, causing an incorrect position for the function pointer table into the invariant OTAU code.

Table 8.1 Error Codes Generated by the OTAU Library

C – Identifier	Code	Default Handling
libZWIR45xx-IPsec.a		
ZWIR_eInvalidVID	400 _{HEX}	Ignore
ZWIR_eInvalidOTAUInterface	401 _{HEX}	Ignore

9 NetMA Libraries

IDT provides a protocol called Network Monitoring and Administration (NetMA) protocol, providing different functionalities to analyze the network. As of network stack version 1.9, there are two versions of this protocol available: NetMA1 and NetMA2. NetMA1 is deprecated and replaced with NetMA2. However, NetMA1 is still available to keep compatibility with older stack releases. Using NetMA1 for new projects is not recommended. Network stack releases prior to 1.9 had NetMA1 included in the core library. As of network stack release 1.9, NetMA functionality has been removed from the core library, providing dedicated libraries instead.

9.1. NetMA1 Library

The NetMA1 library provides functionality for discovery of network devices, determination of node configuration parameters, and determination of routes through the network. The full functionality can be used from a computer that is connected to the network or directly from PAN nodes. For the use with PAN nodes, the NetMA1 library provides API functions for convenient access to its functionality.

9.1.1. NetMA1 Library Symbol Reference

```
void
ZWIR_DiscoverNetwork ( ZWIR_DiscoveryCallback_t  callback,
                      uint8_t                    responseInterval )
```

This function initiates network discovery. A network discovery request is broadcasted to all nodes in the network. The *callback* argument is a pointer to a function that should be called to pass replying node information to the application. The information provided includes the hop distance, the RSSI, IPv6 address count, and all IPv6 addresses of the node. The *responseInterval* argument specifies a maximum time interval within which a responding device must answer the request. The actual response time is chosen randomly. *responseInterval* should be increased with increases in the number of devices in the network. If zero is specified as *responseInterval*, the default response time of three seconds is used.

```

void
ZWIR_NetMA_RemoteParameterRequest ( ZWIR_IPv6Address_t*    address,
                                     ZWIR_NetMA_RPRCallback_t callback,
                                     ZWIR_NetMA_RPRFields_t  fields = 0,
                                     ZWIR_NetMA_Flags_t      flags = 0,
                                     uint8_t                 respInterval = 3,
                                     uint8_t                 queryId = 0,
                                     uint8_t                 hopLimit = 0 )

```

Use this function to obtain configuration data remotely. Data can be requested from a single device or from multiple devices. Different parameters control the answering of the requested devices. The function is actually a macro providing the capability to use the function like a C++ style function with default arguments.

The *address* parameter specifies the device(s) from which data is requested. The function provided with *callback* will be called when a response is received. The remaining arguments are optional. As in C++ all arguments before the last one to be specified must be provided. Assuming the application requires specifying the *respInterval* argument, *fields* and *flags* must be specified as well.

The *fields* argument controls which sets of information are requested from the remote device. Refer to the documentation for `ZWIR_NetMA_RPRFields_t` to determine which options are available. *flags* limits the scope of the request, which is especially useful in conjunction with multicast addressing. For example, it is possible to send requests only to devices configured in Gateway Mode using the *flags* argument. *respInterval* specifies the maximum number of seconds that a device waits before it sends its response. The actual response interval is chosen randomly between zero and *respInterval* seconds in order to avoid collisions of responses sent from multiple devices at the same time.

The *queryId* is used to distinguish between different queries. A device that has successfully responded to a remote parameter request will not respond to another remote parameter request with the same *queryId*. However, the *queryId* is only considered when the corresponding flag in the *flags* argument is set.

The *hopLimit* argument specifies the maximum number of hops allowed when devices respond to the request. If *hopLimit* is left unspecified or explicitly set to zero, all devices will respond regardless of their hop distance to the requesting device. However, the *hopLimit* is only considered when the corresponding flag in the *flags* argument is set.

```

typedef
void ( *ZWIR_NetMA_RPRCallback_t ) ( ZWIR_NetMA_RPRFields_t  fields,
                                     ZWIR_NetMA_RemoteData_t* data )

```

This type defines the signature of functions to be used as remote parameter request response callbacks.

```
void
ZWIR_NetMA_SetPort ( uint16_t port )
```

Using this function, the application can change the *port* used by the NetMA protocol. The default UDP port is 1356.

```
void
ZWIR_NetMA_Trace ( ZWIR_PANAddress_t* routeDestination,
                  ZWIR_IPv6Address_t* routeSource,
                  ZWIR_NetMA_TraceCallback_t callback )
```

This function allows examining routes through the network. Routes to *routeDestination* can be examined from the node calling this function or from a starting point that is passed as in the *routeSource* argument. Selecting a different starting point is required for requests coming from nodes not implementing IDT's network stack, e.g. computers in a network.

Responses to route requests are received by the application through the function defined by *callback*. If *callback* is NULL, the trace request will not be executed. Returned route information contains the list of all hops to *routeDestination* together with the forward RSSI of each hop. Take into account that response times might be significant long if long routes are examined.

Note: The *ZWIR_NetMA_Trace* function will not create routes between *routeSource* and *routeDestination*, but can generate routes between the requesting device and *routeSource* if required.

```
typedef
void ( * ZWIR_DiscoveryCallback_t ) ( uint8_t hopCount,
                                     int8_t rssi,
                                     uint8_t addrCount,
                                     ZWIR_IPv6Address_t* addresses )
```

Function pointer type for the callback function that should be called if network discovery reply packets are received.

```
typedef
void ( * ZWIR_NetMA_TraceCallback_t ) ( uint8_t hopCount,
                                       ZWIR_NetMA_HopInfo_t* hopInfo )
```

This type defines the function signature that is required by functions to be used as callback for the *ZWIR_NetMA_Trace* function.

```
typedef
struct {
    ZWIR_PANAddress_t  address;
    int16_t            linkRSSI;
} ZWIR_NetMA_HopInfo_t
```

Objects of this type are passed to the trace route callback function. Each object contains the address of a hop and the RSSI value of the forward path from the previous hop/source node to this node. Note that the RSSI of the return path might be slightly different.

```
typedef
enum { ... } ZWIR_NetMA_RPRFields_t
```

This enumerator is used with `ZWIR_NetMA_RemoteParameterRequest` to specify which sets of information must be included in the response. The values can be binary OR'ed to request multiple sets of information at the same time. The following values are available:

<code>ZWIR_NetMA_rprfMACAddress</code>	0x0100	include	<code>ZWIR_NetMA_RemoteMACAddr_t</code>
<code>ZWIR_NetMA_rprfFirmwareVersion</code>	0x0200	include	<code>ZWIR_NetMA_RemoteVersion_t</code>
<code>ZWIR_NetMA_rprfConfig</code>	0x0400	include	<code>ZWIR_NetMA_RemoteConfig_t</code>
<code>ZWIR_NetMA_rprfIPv6Addresses</code>	0x0800	include	<code>ZWIR_NetMA_RemoteIPv6Addr_t</code>
<code>ZWIR_NetMA_rprfTRXStatistics</code>	0x1000	include	<code>ZWIR_NetMA_RemoteStatus_t</code>

```
typedef
enum { ... } ZWIR_NetMA_Flags_t
```

This enumerator defines flags to be used in conjunction with remote parameter requests. The flags limit the scope of the request. Possible values include

<code>ZWIR_NetMA_fDevice</code>	0x04
<code>ZWIR_NetMA_fBridge</code>	0x08
<code>ZWIR_NetMA_fQueryID</code>	0x10

ZWIR_NetMA_fHopCountLimitation 0x20

```
typedef
struct {
    ZWIR_NetMA_RemoteMACAddr_t*   macAddr;
    ZWIR_NetMA_RemoteIPv6Addr_t*  ipv6Addr;
    ZWIR_NetMA_RemoteConfig_t*    config;
    ZWIR_NetMA_RemoteVersion_t*   version;
    ZWIR_NetMA_RemoteStatus_t*    status;
} ZWIR_NetMA_RemoteData_t
```

This structure contains pointers to the remote parameters received in response to a remote parameter request. The fields correlate with the `ZWIR_NetMA_RPRFields_t` enumerators. For each requested field, the corresponding pointer should be set in the response. Fields that have not been requested result in a NULL pointer of the corresponding structure element.

```
typedef
struct {
    uint16_t          panID;
    ZWIR_PANAddress_t panAddr;
} ZWIR_NetMA_RemoteMACAddr_t
```

This structure type carries a remote device's configured PAN ID and its PAN address.

```
typedef
struct {
    uint8_t          count;
    ZWIR_IPv6Address_t addresses [ ];
} ZWIR_NetMA_RemoteIPv6Addr_t
```

This type defines a structure carrying all IPv6 addresses of a device. The `count` argument defines how many addresses are contained in the structure; the `addresses` element contains the actual addresses. The size of memory required by this structure varies – it depends on the number of contained addresses. The size of this structure cannot be determined using C's `sizeof` operator.

```
typedef
  struct {
    uint16_t  routeTimeout;
    uint16_t  routingTableSize;
    uint16_t  neighborReachableTime;
    uint8_t   neighborCacheSize;
    uint8_t   maxNetfloodHopCount;
    uint8_t   maxSocketCount;
    uint8_t   routeMaxFailCount;
    int8_t    routeRequestMinLinkRSSI;
    uint8_t   routeRequestMinLinkRSSIReduction;
    uint8_t   routeRequestAttempts;
    uint8_t   channel;
    uint8_t   power;
    uint8_t   modulation;
    uint8_t   doDuplicateAddressDetection;
    uint8_t   doRouterSolicitation;
  } ZWIR_NetMA_RemoteConfig_t
```

Objects of this type are used to report the configuration of the remote device.

```
typedef
  ZWIR_TRXStatistic_t  ZWIR_NetMA_RemoteStatus_t
```

This type defines the remote status as being equal to the transceiver statistics type.

```
typedef
  struct {
    uint32_t  vendorID;
    uint16_t  productID;
    uint8_t   firmwareMajorVersion;
    uint8_t   firmwareMinorVersion;
    uint16_t  firmwareVersionExtension;
    uint8_t   libraryMajorVersion;
    uint8_t   libraryMinorVersion;
    uint16_t  libraryVersionExtension;
  } ZWIR_NetMA_RemoteVersion_t
```

This type bundles all version information defined in a device.

9.1.2. Inclusion of the NetMA1 library

Including the NetMA1 library in an application does not require any specific measures. However, note that NetMA1 will use one of the available network sockets for reception of data. When defining the system parameter `ZWIR_spMaxSocketCount`, this additional socket must be taken into account.

9.2. NetMA2 Libraries

NetMA2 is an advancement of the NetMA1 protocol that is not downward compatible. It provides extended functionality including network discovery, topology discovery, remote parameter readout and modification, parameter storage, route tracing and routing table readout. In addition to the extended functionality, NetMA2 provides advanced filtering, allowing restriction of the scope of multicast requests sent into the network. The NetMA2 functionality is implemented in two separate libraries: one that provides only readout remote readout functionality and one that provides all functionality for changing and storing network parameters.

The protocol specification is open and can be found in the *ZWIR45xx Application Note – The NetMA Protocol*. In order to keep the footprint of the libraries as small as possible, NetMA2 does not provide API functions for the different NetMA2 functionalities. However, due to the open protocol specification, applications can implement the required functionality. For that purpose, NetMA2 requests must be sent explicitly using the function `ZWIR_sendUDP2` and a custom implementation of the `ZWIR_NetMA_ResponseHandler` must be provided for processing the corresponding responses.

9.2.1. NetMA2 Library Symbol Reference

```
void
ZWIR_DiscoverNetwork ( ZWIR_DiscoveryCallback_t  callback,
                      uint8_t                    responseInterval )
```

This function initiates network discovery. A network discovery request is broadcasted to all nodes in the network. The *callback* argument is a pointer to a function to be called to pass replying node information to the application. The information provided includes the hop distance, the RSSI, IPv6 address count, and all IPv6 addresses of the node. The *responseInterval* argument specifies a maximum time interval within which a responding device must answer the request. The actual response time is chosen randomly. *responseInterval* should be increased with increases in the number of devices in the network. If zero is specified as *responseInterval*, the default response time of three seconds is used.

```
void
ZWIR_NetMA_SetPort ( uint16_t  port )
```

This function is used to change the port used by the NetMA protocol. By default port 61356 is used.

```
bool
ZWIR_NetMA_ResponseHandler ( uint8_t const*  data,
                              uint16_t       size )
```

This hook is provided to allow the reception of NetMA2 packets in the application. It is required to receive NetMA2 responses triggered by NetMA2 requests sent by the application. The raw NetMA2 packet is contained in the data pointed to by *data*. The packet size is provided with *size*. If the implementation returns false, the NetMA2 stack continues packet processing. If true is returned, the NetMA2 packet does not do further processing. If no implementation is provided, NetMA2 requests will be processed normally, but it is impossible to receive responses to NetMA2 requests.

Note: This function is called for any packet received on the configured NetMA2 port. Thus, NetMA2 requests will trigger the execution of this function as well. In order to keep NetMA2 fully functional, the application must return false for each NetMA2 request that has not been handled.

```
bool
ZWIR_NetMA_Filter ( uint8_t* data,
                    uint16_t size,
                    uint8_t* dataOffset )
```

This function applies the filtering rules defined by the NetMA2 protocols and returns true if the packet has to be filtered (i.e., the packet must be ignored) or false otherwise. The *data* and *size* determine the raw packet data and size. The value pointed to by *dataOffset* is written by the function to indicate the offset of the first byte after the NetMA2 header.

Note: This function is required for NetMA2 requests only, as NetMA2 responses typically do not use filters. If no custom handling of NetMA2 requests is provided by the application, this function does not need to be called as the NetMA2 request default handlers call this function automatically.

9.2.2. Inclusion of the NetMA2 Libraries

The NetMA2 functionality that does not require write access to stack parameters or flash memory is included in the library *libZWIR45xx-NetMA2.a*. Including this library does not require special measures. However, note that the library requires a network socket for reception of data. When defining the system parameter *ZWIR_spMaxSocketCount*, this additional socket must be taken into account.

If the application requires functionality to configure devices remotely, the library *libZWIR45xx-NetMA2-Ext.a* must be included in the project. This library requires having *libZWIR45xx-NetMA2.a* included as well. Furthermore, it must be ensured that the linker script contains the section *.nvDataMemory*. This section is provided with the default linker script of stack release 1.9. When upgrading from a lower stack version, be sure to replace the linker script with the newest version.

10 Accessing Microcontroller Resources

Many applications can be optimized by using the rich internal resources provided by the ZWIR451x module's STM32 microcontroller. Typically, this does not cause problems, but some caution must be taken when this is considered. No resources must be used that are already occupied by the operating system (OS). Furthermore some of the MCU configuration parameters must not be altered. Refer to the next section for a complete list of resources that are used by the OS.

The library does not provide dedicated functions for configuration of the microcontroller peripherals. This must be done by third-party libraries or by programming the appropriate configuration registers directly. Names for interrupt handlers are predefined by the library. If interrupt handlers are required, a function with the library-determined name must be implemented by the library user.

10.1. Internal Microcontroller Configuration

By default the internal STM32 is clocked from its internal 8MHz oscillator (HSI). Depending on the clock speed setting (`ZWIR_SetFrequency`), the system clock (SYSCLK) is taken from the HSI directly or from the phase-locked loop (PLL) output (PLLCLK). If the clock is taken from the PLL, the PLL source is HSI/2 and the PLL multiplier is 16, so SYSCLK has a frequency of 64MHz. The Advanced High-Performance Bus (AHB) clock (HCLK) is configured according to the selected CPU frequency (see `ZWIR_SetFrequency`). The APB1 clock (PCLK1) frequency is always 4MHz. The APB2 clock (PCLK2) frequency is always 8MHz.

Important recommendation: The frequencies of APB1 and APB2 *SHOULD NOT* be changed! Doing so would result in improper timing behavior of the operating system and could result in system breakdown.

The Cortex® System Timer (SysTick) is used as the operating system base timer. It is configured to issue an interrupt each millisecond. The real-time clock (RTC) is used for the different sleep modes.

All microcontroller GPIO pins that are not connected to one of the module's I/O pins are locked so that it is impossible to change their configuration accidentally.

Important: The configuration of the external interrupt line 0 (EXTI0) and the configuration of the SPI2 peripheral of the microcontroller *MUST NOT* be changed. Otherwise the interfacing between MCU and transceiver might be impaired.

10.2. Backup Data Registers

The STM32 provides 42 backup data registers with a size of 2 bytes each. The stack uses the first register (BKP→DR1) for internal configurations. The user application *SHOULD NOT* use or modify this register.

10.3. Interrupt Handlers

The API library comes with a set of predefined interrupt handlers that is sufficient for the built-in functionality but does not go beyond it. For all other interrupts that are not required, default handlers are provided that typically do nothing or perform a reset in the event of an error. Most of the interrupts are defined as weak symbols in the library. This means that the default implementations of the handlers can be overwritten by simply defining the interrupt handler symbol in the user's application code. Only those interrupts that are required for proper operation of the stack are not defined as weak and therefore cannot be overwritten. An attempt to overwrite these handlers will result in a linker error.

Table 10.1 lists the interrupts for the STM32. For each interrupt, the handler name, the default priority, and the default behavior is shown and whether or not the interrupt can be overwritten.

Table 10.1 STM32 Interrupt Vector Table

Interrupt		Implementation			
Id	Name	Handler	Priority	Fix	Default Behavior
0	Reset	Reset_Handler	-3	Yes	Perform system reset
1	NMI	ZWIR_ISR_NMI	-2	No	Perform system reset
2	HardFault	ZWIR_ISR_HardFault	-1	No	Perform system reset
3	MemManage	ZWIR_ISR_MemManage	0	No	Perform system reset
4	BusFault	ZWIR_ISR_BusFault	1	No	Perform system reset
5	UsageFault	ZWIR_ISR_UsageFault	2	No	Perform system reset
6	SVCall	ZWIR_ISR_SVCall	3	No	NULL
7	DebugMonitor	ZWIR_ISR_DebugMonitor	4	No	NULL
8	PendSV	ZWIR_ISR_PendSV	5	No	NULL
9	SysTick	ZWIR_ISR_SysTick	6	Yes	Used as operating system timer
10	WWDG	ZWIR_ISR_WWDG	7	No	NULL
11	PVD	ZWIR_ISR_PVD	8	No	Perform system reset
12	TAMPER	ZWIR_ISR_TAMPER	9	No	NULL
13	RTC	ZWIR_ISR_RTC	10	Yes	Reserved for OS use
14	FLASH	ZWIR_ISR_FLASH	11	No	NULL
15	RCC	ZWIR_ISR_RCC	12	No	NULL
16	EXTI0	ZWIR_ISR_EXTI0	13	Yes	Handle transceiver service request
17	EXTI1	ZWIR_ISR_EXTI1	14	No	NULL
18	EXTI2	ZWIR_ISR_EXTI2	15	No	NULL
19	EXTI3	ZWIR_ISR_EXTI3	16	No	NULL
20	EXTI4	ZWIR_ISR_EXTI4	17	No	NULL
21	DMA1_Channel1	ZWIR_ISR_DMA1_Channel1	18	No	NULL
22	DMA1_Channel2	ZWIR_ISR_DMA1_Channel2	19	No	NULL
23	DMA1_Channel3	ZWIR_ISR_DMA1_Channel3	20	No	NULL
24	DMA1_Channel4	ZWIR_ISR_DMA1_Channel4	21	No	NULL
25	DMA1_Channel5	ZWIR_ISR_DMA1_Channel5	22	No	NULL
26	DMA1_Channel6	ZWIR_ISR_DMA1_Channel6	23	No	NULL
27	DMA1_Channel7	ZWIR_ISR_DMA1_Channel7	24	No	NULL
28	ADC1_2	ZWIR_ISR_ADC1_2	25	No	NULL

Interrupt		Implementation			
Id	Name	Handler	Priority	Fix	Default Behavior
29	USB_HP_CAN_TX	ZWIR_ISR_USB_HP_CAN1_TX	26	No	NULL
30	USB_LP_CAN_RX0	ZWIR_ISR_USB_LP_CAN1_RX0	27	No	NULL
31	CAN_RX1	ZWIR_ISR_CAN1_RX1	28	No	NULL
32	CAN_SCE	ZWIR_ISR_CAN1_SCE	29	No	NULL
33	EXTI9_5	ZWIR_ISR_EXTI9_5	30	No	NULL
34	TIM1_BRK	ZWIR_ISR_TIM1_BRK	31	No	NULL
35	TIM1_UP	ZWIR_ISR_TIM1_UP	32	No	NULL
36	TIM1_TRG_COM	ZWIR_ISR_TIM1_TRG_COM	33	No	NULL
37	TIM1_CC	ZWIR_ISR_TIM1_CC	34	No	NULL
38	TIM2	ZWIR_ISR_TIM2	35	No	NULL
39	TIM3	ZWIR_ISR_TIM3	36	No	NULL
40	TIM4	ZWIR_ISR_TIM4	37	Yes	NULL
41	I2C1_EV	ZWIR_ISR_I2C1_EV	38	No	NULL
42	I2C1_ER	ZWIR_ISR_I2C1_ER	39	No	NULL
43	I2C2_EV	ZWIR_ISR_I2C2_EV	40	No	NULL
44	I2C2_ER	ZWIR_ISR_I2C2_ER	41	No	NULL
45	SPI1	ZWIR_ISR_SPI1	42	No	NULL
46	SPI2	ZWIR_ISR_SPI2	43	Yes	Used by network stack
47	USART1	ZWIR_ISR_USART1	44	No	NULL [§]
48	USART2	ZWIR_ISR_USART2	45	No	NULL ^{**}
49	USART3	ZWIR_ISR_USART3	46	No	NULL
50	EXTI15_10	ZWIR_ISR_EXTI15_10	47	No	NULL
51	RTCAlarm	ZWIR_ISR_RTCAlarm	48	Yes	Reserved for OS use
52	USBWakeup	ZWIR_ISR_USBWakeup	49	No	NULL
53	TIM8_BRK	ZWIR_ISR_TIM8_BRK	50	No	NULL
54	TIM8_UP	ZWIR_ISR_TIM8_UP	51	No	NULL
55	TIM8_TRG_COM	ZWIR_ISR_TIM8_TRG_COM	52	No	NULL
56	TIM8_CC	ZWIR_ISR_TIM8_CC	53	No	NULL
57	ADC3	ZWIR_ISR_ADC3	54	No	NULL

[§] Implementation is provided if libZWIR451x-UART1.a is linked into the project

^{**} Implementation is provided if libZWIR451x-UART2.a is linked into the project

Interrupt		Implementation			
Id	Name	Handler	Priority	Fix	Default Behavior
58	FSMC	ZWIR_ISR_FSMC	55	No	NULL
59	SDIO	ZWIR_ISR_SDIO	56	No	NULL
60	TIM5	ZWIR_ISR_TIM5	57	No	NULL
61	SPI3	ZWIR_ISR_SPI3	58	No	NULL
62	UART4	ZWIR_ISR_UART4	59	No	NULL
63	UART5	ZWIR_ISR_UART5	60	No	NULL
64	TIM6	ZWIR_ISR_TIM6	61	No	NULL
65	TIM7	ZWIR_ISR_TIM7	62	No	NULL
66	DMA2_Channel1	ZWIR_ISR_DMA2_Channel1	63	No	NULL
67	DMA2_Channel2	ZWIR_ISR_DMA2_Channel2	64	No	NULL
68	DMA2_Channel3	ZWIR_ISR_DMA2_Channel3	65	No	NULL
69	DMA2_Channel4_5	ZWIR_ISR_DMA2_Channel4_5	66	No	NULL

10.4. Default I/O Configuration

Table 10.2 shows the default I/O configuration that is set when the device is powered on and no manual changes are made to the I/Os. In some cases, the I/O configuration depends on the libraries included and the library settings. For these I/Os there are multiple lines, showing the configuration depending on the library and their configuration.

Table 10.2 STM32 Default I/O Configuration

MCU	Pin		Configuration	Drive	Comment
Port	ZWIR4512AC1	ZWIR4512AC2			
A0	8	8	Analog Input	-	Freely configurable by application, but not to be used as external interrupt source
			Floating Input	-	With libZWIR451x-UART2.a with flow control enabled
A1	7	7	Analog Input	-	Freely configurable by application
			2 MHz Push/Pull Alternative Output	x	With libZWIR451x-UART2.a with flow control enabled
A2	6	6	Analog Input	-	Freely configurable by application
			2 MHz Push/Pull Alternative Output	x	With libZWIR451x-UART2.a
A3	5	5	Analog Input	-	Freely configurable by application
			Floating Input	-	With libZWIR451x-UART2.a
A4	4	4	Analog Input	-	Freely configurable by application
A5	3	3	Analog Input	-	Freely configurable by application
A6	2	2	Analog Input	-	Freely configurable by application
A7	1	1	Analog Input	-	Freely configurable by application
A8	-	-	2 MHz Push/Pull Output	0	Configuration Locked
A9	13	14	Analog Input	-	Freely configurable by application
			2 MHz Push/Pull Alternative Output	x	libZWIR451x-UART1.a
A10	12	13	Analog Input	-	Freely configurable by application
			Floating Input	-	libZWIR451x-UART1.a
A11	16	19	Analog Input	-	Freely configurable by application
			Floating Input	-	libZWIR451x-UART1.a with flow control enabled
A12	17	20	Analog Input	-	Freely configurable by application
			2 MHz Push/Pull Alternative Output	x	libZWIR451x-UART1.a with flow control enabled
A13	20	22	Pull-up Input	-	Initially configured as JTAG pin Freely configurable by application when remapping is enabled

MCU	Pin		Configuration	Drive	Comment
	ZWIR4512AC1	ZWIR4512AC2			
Port					
A14	22	24	Pull-down Input	-	Initially configured as JTAG pin Freely configurable by application when remapping is enabled
A15	21	23	Pull-up Input	-	Initially configured as JTAG pin Freely configurable by application when remapping is enabled
B0	-	-	Pull-up Input	-	Configuration locked, application <i>MUST NOT</i> modify the output data register
B1	-	-	Floating Input	-	Configuration locked
B2	-	-	Floating Input	-	Configuration locked
B3	19	21	50 MHz Push/Pull Output	-	Initially configured as JTAG pin Freely configurable by application when remapping is enabled
B4	-	-	Floating Input	-	Configuration locked
B5	-	-	Floating Input	-	Configuration locked
B6	24	26	Floating Input	-	Freely configurable by application
B7	23	25	Floating Input	-	Freely configurable by application
B8	-	-	Floating Input	-	Configuration locked
B9	-	-	2 MHz Push/Pull Output	1	Configuration locked
B10	-	-	2 MHz Push/Pull Output	0	Configuration locked
B11	-	-	2 MHz Push/Pull Output	0	Configuration locked
B12	-	-	10 MHz Push/Pull Output	1	Configuration locked
B13	-	-	10 MHz Push/Pull Alternative Output	x	Configuration locked
B14	-	-	10 MHz Push/Pull Alternative Output	x	Configuration locked
B15	-	-	10 MHz Push/Pull Alternative Output	x	Configuration locked
C0	-	-	2 MHz Push/Pull Output	1	Configuration locked
C1	-	-	2 MHz Push/Pull Output	1	Configuration locked
C2	-	-	2 MHz Push/Pull Output	1	Configuration locked
C4	-	-	Analog Input	-	Configuration locked
C5	-	-	Analog Input	-	Configuration locked

MCU	Pin				
Port	ZWIR4512AC1	ZWIR4512AC2	Configuration	Drive	Comment
C6	-	-	Analog Input	-	Configuration locked
C7	-	-	Analog Input	-	Configuration locked
C8	-	-	Analog Input	-	Configuration locked
C9	-	-	Analog Input	-	Configuration locked
C10	-	-	Analog Input	-	Configuration locked
C11	-	-	Analog Input	-	Configuration locked
C12	-	-	Analog Input	-	Configuration locked
C13	9	9	Analog Input	-	Freely configurable by application
C14	-	17	Analog Input	-	Freely configurable by application
C15	-	18	Analog Input	-	Freely configurable by application

11 Certification

11.1. European R&TTE Directive Statements

The ZWIR4512 module has been tested and found to comply with Annex IV of the R&TTE Directive 1999/5/EC and is subject of a notified body opinion. The module has been approved for Antennas with gains of 4 dBi or less.

11.2. Federal Communication Commission Certification Statements

11.2.1. Statements

This equipment has been tested and found to comply with the limits for a **Class B Digital Device**, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from where the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

This device complies with part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) This device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Modifications not expressly approved by ZMD AG could void the user's authority to operate the equipment.

The internal/external antennas used for this mobile transmitter must provide a separation distance of at least 20cm from all persons and must not be co-located or operating in conjunction with any other antenna or transmitter.

11.2.2. Requirements

The ZWIR4512 complies with Part 15 of the FCC rules and regulations. In order to retain compliance with the FCC certification requirements, the following conditions must be met:

1. Modules must be installed by original equipment manufacturers (OEM) only.
2. The module must only be operated with antennas adhering to the requirements defined in section 11.2.3.
3. The OEM must place a clearly visible text label on the outside of the end-product containing the text shown in Figure 11.1 below.

IMPORTANT: The compliance statement as shown in Figure 11.1 must be used without modifications for both ZWIR4512 product versions as the FCC ID covers the ZWIR4512AC1 and the ZWIR4512AC2!

Figure 11.1 FCC Compliance Statement to be Printed on Equipment Incorporating ZWIR4512 Devices

Contains FCC ID: COR-ZWIR4512AC1
This device complies with part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) This device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

11.2.3. Accessing the FCC ID

ZWIR451x modules are capable of showing their FCC-ID electronically. C-API applications can read the module's FCC-ID through the function `ZWIR_GetFCCID`. Due to space constraints, the FCC ID is not printed on the module. Host devices incorporating this module must be marked according to the above guidelines.

11.3. Supported Antennas

The FCC compliance testing of the ZWIR4512 has been carried out using the MEXE902RPSM antenna from PCTEL Inc. This antenna has an omnidirectional radiation pattern at an antenna gain of 2 dBi. In order to be allowed to use the module without re-certification, the product incorporating the ZWIR4512 module must either use the antenna mentioned above or must use an antenna with an omnidirectional radiation pattern and a gain being less than or equal to 2 dBi.

12 Alphabetical Lists of Symbols

12.1. Functions and Function-Like Macros

ZWIR_AbortPowerDown	57	ZWIR_MulticastPreferExistingRepeater	52
ZWIR_AddAlternativeAddress	54	ZWIR_NetEventCallback	41
ZWIR_AppInitHardware	37	ZWIR_NetMA_Filter	87
ZWIR_AppInitNetwork	37	ZWIR_NetMA_RemoteParameterRequest	81
ZWIR_AppInitNetworkDone	37	ZWIR_NetMA_ResponseHandler	87
ZWIR_CheckMulticastGroup	45	ZWIR_NetMA_RPRCallback_t	82
ZWIR_CloseSocket	46	ZWIR_NetMA_SetPort (NetMA1)	82
ZWIR_CreateAlternativeAddressList	54	ZWIR_NetMA_SetPort (NetMA2)	87
ZWIR_DiscoverNetwork (NetMA1)	81	ZWIR_NetMA_Trace	82
ZWIR_DiscoverNetwork (NetMA2)	87	ZWIR_OpenSocket	45
ZWIR_Error	42	ZWIR_OTAU_Register	78
ZWIR_ExternalClockEnable	55	ZWIR_OTAU_Status	78
ZWIR_GatewayProcessPacket	51	ZWIR_PowerDown	57
ZWIR_GatewaySetOutputFunction	52	ZWIR_Rand	42
ZWIR_GetChannel	49	ZWIR_RegisterAppEventHandler	40
ZWIR_GetDestinationPANAddress	48	ZWIR_Reset	39
ZWIR_GetFailingAddress	48	ZWIR_ResetDestinationPANId	44
ZWIR_GetFCCID	56	ZWIR_ResetNetwork	39
ZWIR_GetIPv6Addresses	44	ZWIR_ResetTRXStatistic	53
ZWIR_GetLastRSSI	48	ZWIR_SelectRTCSource	57
ZWIR_GetModulation	50	ZWIR_Send6LoWPAN	48
ZWIR_GetPacketDestinationAddress	47	ZWIR_SendUDP	46
ZWIR_GetPacketHopCount	47	ZWIR_SendUDP2	46
ZWIR_GetPacketRXSocket	48	ZWIR_SetChannel	49
ZWIR_GetPacketSenderAddress	47	ZWIR_SetDestinationPANId	44
ZWIR_GetPacketSenderPort	47	ZWIR_SetDutyCycleWarning	53
ZWIR_GetPANAddress	43	ZWIR_SetFrequency	56
ZWIR_GetPANId	43	ZWIR_SetIPv6Address	44
ZWIR_GetParameter	61	ZWIR_SetModulation	50
ZWIR_GetRevision	41	ZWIR_SetOperatingMode	38
ZWIR_GetRTC	57	ZWIR_SetPAControl	55
ZWIR_GetSourcePANAddress	48	ZWIR_SetPANAddress	43
ZWIR_GetTransmitPower	50	ZWIR_SetPANId	43
ZWIR_GetTRXStatistic	53	ZWIR_SetParameter	61
ZWIR_GPIO_ConfigureAsInput	67	ZWIR_SetPromiscuousMode	54
ZWIR_GPIO_ConfigureAsOutput	67	ZWIR_SetRTC	57
ZWIR_GPIO_Read	68	ZWIR_SetTransmitPower	50
ZWIR_GPIO_ReadMultiple	68	ZWIR_SetWakeupSource	57
ZWIR_GPIO_Remap	69	ZWIR_Sleep	58
ZWIR_GPIO_Write	69	ZWIR_SRand	42
ZWIR_IsAlternativeAddress	55	ZWIR_Standby	58
ZWIR_LocalBroadcast	52	ZWIR_StartCallbackTimer	40
ZWIR_Main1000ms	39	ZWIR_StopCallbackTimer	40
ZWIR_Main100ms	39	ZWIR_TransceiverOff	58
ZWIR_Main10ms	39	ZWIR_TransceiverOn	59

ZWIR_TriggerAppEvent	40	ZWIR_UART2_PRINTF	65
ZWIR_UART1_GetAvailableTXBuffer	65	ZWIR_UART2_ReadByte	64
ZWIR_UART1_IsTXEmpty	65	ZWIR_UART2_Send	63
ZWIR_UART1_PRINTF	65	ZWIR_UART2_SendByte	63
ZWIR_UART1_ReadByte	64	ZWIR_UART2_SetRXCallback	64
ZWIR_UART1_Send	63	ZWIR_UART2_Setup	64
ZWIR_UART1_SendByte	63	ZWIRSEC_AddIKEAuthenticationEntry	76
ZWIR_UART1_SetRXCallback	64	ZWIRSEC_AddSecurityAssociation	72
ZWIR_UART1_Setup	64	ZWIRSEC_AddSecurityPolicy	72
ZWIR_UART2_GetAvailableTXBuffer	65	ZWIRSEC_RemoveSecurityAssociation	73
ZWIR_UART2_IsTXEmpty	65	ZWIRSEC_RemoveSecurityPolicy	72

12.2. Data Types

ZWIR_AlternativeAddressType_t	55	ZWIR_GPIO_rfSWJ	71
ZWIR_aatAny	55	ZWIR_GPIO_SWJRemapValue_t	71
ZWIR_aatEUI48	55	ZWIR_GPIO_swjrDisableSWJ	71
ZWIR_aatEUI64	55	ZWIR_GPIO_swjrEnableSWJ	71
ZWIR_aatNone	55	ZWIR_GPIO_swjrSWOnly	71
ZWIR_AppEventHandler_t	40	ZWIR_IPv6Address_t	45
ZWIR_DiscoveryCallback_t	83	ZWIR_MCUFrequency_t	59
ZWIR_DutyCycleCallback_t	53, 54	ZWIR_mcu16MHz	59
ZWIR_ErrorCode_t		ZWIR_mcu32MHz	59
ZWIR_eDADFailed	62	ZWIR_mcu64MHz	59
ZWIR_eExtClockPowerDown	62	ZWIR_mcu8MHz	59
ZWIR_eHostUnreachable	62	ZWIR_mcu8MHzLowPower	59
ZWIR_eMemoryExhaustion	62	ZWIR_mcuUserFrequency	59
ZWIR_eProgExit	62	ZWIR_Modulation_t	51
ZWIR_eReadMACFailed	62	ZWIR_mBPSK	51
ZWIR_eRouteFailed	62	ZWIR_mQPSK	51
ZWIR_GatewayOutputFunction_t	52	ZWIR_NetEvent_t	41
ZWIR_GPIO_DriverStrength_t	71	ZWIR_neAppReceive	41
ZWIR_GPIO_dsHigh	71	ZWIR_neAppTransmit	41
ZWIR_GPIO_dsLow	71	ZWIR_neIPv6Receive	41
ZWIR_GPIO_dsMedium	71	ZWIR_neIPv6Transmit	41
ZWIR_GPIO_InputMode_t	71	ZWIR_neMACReceive	41
ZWIR_GPIO_imAnalog	71	ZWIR_neMACTransmi	41
ZWIR_GPIO_imFloating	71	ZWIR_NetMA_Flags_t	84
ZWIR_GPIO_imPullDown	71	ZWIR_NetMA_fBridge	84
ZWIR_GPIO_imPullUp	71	ZWIR_NetMA_fDevice	84
ZWIR_GPIO_OutputMode_t	71	ZWIR_NetMA_fHopCountLimitation	84
ZWIR_GPIO_omAlternativeOpenDrain	71	ZWIR_NetMA_fQueryID	84
ZWIR_GPIO_omAlternativePushPull	71	ZWIR_NetMA_HopInfo_t	83
ZWIR_GPIO_omOpenDrain	71	ZWIR_NetMA_RemoteConfig_t	85
ZWIR_GPIO_omPushPull	71	ZWIR_NetMA_RemoteData_t	84
ZWIR_GPIO_Pin_t	70	ZWIR_NetMA_RemoteIPv6Addr_t	85
ZWIR_GPIO_RemapFunction_t	71	ZWIR_NetMA_RemoteMACAddr_t	84

ZWIR_NetMA_RemoteStatus_t.....	85	ZWIR_protoTCP.....	73
ZWIR_NetMA_RemoteVersion_t.....	86	ZWIR_protoUDP.....	73
ZWIR_NetMA_RPRFields_t.....	83	ZWIR_RadioChannel_t.....	50
ZWIR_NetMA_rprConfig.....	84	ZWIR_channel0.....	50
ZWIR_NetMA_rprFirmwareVersion.....	83	ZWIR_channel1.....	50
ZWIR_NetMA_rprIPv6Addresses.....	84	ZWIR_channel10.....	51
ZWIR_NetMA_rprMACAddress.....	83	ZWIR_channel100.....	51
ZWIR_NetMA_rprTRXStatistics.....	84	ZWIR_channel101.....	51
ZWIR_NetMA_TraceCallback_t.....	83	ZWIR_channel102.....	51
ZWIR_OperatingMode_t.....	38	ZWIR_channel2.....	50
ZWIR_omGateway.....	38	ZWIR_channel3.....	50
ZWIR_omNormal.....	38	ZWIR_channel4.....	50
ZWIR_omSniffer.....	38	ZWIR_channel5.....	50
ZWIR_OTAU_ErrorCode_t.....		ZWIR_channel6.....	51
ZWIR_eInvalidOTAUIInterface.....	80	ZWIR_channel7.....	51
ZWIR_eInvalidVID.....	80	ZWIR_channel8.....	51
ZWIR_OTAU_StatusCode_t.....	78	ZWIR_channel9.....	51
ZWIR_sCRCWriteError.....	79	ZWIR_eu865.....	51
ZWIR_sFirmwareImageVerifyDone.....	79	ZWIR_eu866.....	51
ZWIR_sFirmwareImageVerifyFailed.....	79	ZWIR_eu867.....	51
ZWIR_sFragmentWriteError.....	79	ZWIR_eu868.....	50
ZWIR_slInitNewFirmwareDone.....	79	ZWIR_us906.....	50
ZWIR_slInvalidCRCPacket.....	79	ZWIR_us908.....	50
ZWIR_slInvalidDataPacket.....	79	ZWIR_us910.....	50
ZWIR_slInvalidExecutePacket.....	79	ZWIR_us912.....	50
ZWIR_slInvalidFragmentSize.....	78	ZWIR_us914.....	50
ZWIR_slInvalidPacketCRC.....	78	ZWIR_us916.....	51
ZWIR_sScheduleUpdate.....	79	ZWIR_us918.....	51
ZWIR_sStartUpdate.....	79	ZWIR_us920.....	51
ZWIR_sUnknownPacketType.....	78	ZWIR_us922.....	51
ZWIR_sUpdateInProgress.....	78	ZWIR_us924.....	51
ZWIR_sWriteCRCDone.....	79	ZWIR_RadioReceiveCallback_t.....	49
ZWIR_sWriteFragmentDone.....	79	ZWIR_ResetReason_t.....	38
ZWIR_PAControl_t.....	56	ZWIR_rlIndependentWatchdogReset.....	38
ZWIR_pa2us.....	56	ZWIR_rlLowPowerReset.....	38
ZWIR_pa4us.....	56	ZWIR_rPinReset.....	38
ZWIR_pa6us.....	56	ZWIR_rPowerOnReset.....	38
ZWIR_pa8us.....	56	ZWIR_rSoftwareReset.....	38
ZWIR_paOff.....	56	ZWIR_rStandbyReset.....	38
ZWIR_PANAddress_t.....	43	ZWIR_rWindowWatchdogReset.....	38
ZWIR_PowerDownState_t.....	59	ZWIR_RevisionInfo_t.....	41
ZWIR_pSleep.....	59	ZWIR_RTCSource_t.....	59
ZWIR_pSleepAfterActivities.....	59	ZWIR_rExtern.....	59
ZWIR_pStandby.....	59	ZWIR_rlIntern.....	59
ZWIR_pStandbyAfterActivities.....	59	ZWIR_SocketHandle_t.....	49
ZWIR_pStop.....	59	ZWIR_SystemParameter_t.....	61
ZWIR_pStopAfterActivities.....	59	ZWIR_spDoAddressAutoConfiguration.....	62
ZWIR_Protocol_t.....	73	ZWIR_spDoDuplicateAddressDetection.....	61
ZWIR_protoAny.....	73	ZWIR_spDoRouterSolicitation.....	61
ZWIR_protolCMPv6.....	73	ZWIR_spHeaderCompressionContext1.....	62

ZWIR_spHeaderCompressionContext2.....	62	ZWIR_UART1_eOvfl.....	66
ZWIR_spHeaderCompressionContext3.....	62	ZWIR_UART1_eParity.....	66
ZWIR_spMaxHopCount.....	61	ZWIR_UART2_ErrorCode_t	
ZWIR_spMaxSocketCount.....	61	ZWIR_UART2_eFrame.....	66
ZWIR_spNeighborCacheSize.....	61	ZWIR_UART2_eNoise.....	66
ZWIR_spNeighborReachableTime.....	61	ZWIR_UART2_eOvfl.....	66
ZWIR_spNeighborRetransTime.....	62	ZWIR_UART2_eParity.....	66
ZWIR_spRouteMaxFailCount.....	61	ZWIRSEC_AuthenticationAlgorithm_t.....	74
ZWIR_spRouteRequestAttempts.....	62	ZWIRSEC_authAESXCBC96.....	74
ZWIR_spRouteRequestMinLinkRSSI.....	61	ZWIRSEC_authNull.....	74
ZWIR_spRouteRequestMinLinkRSSIReduction.....	61	ZWIRSEC_AuthenticationSuite_t.....	74
ZWIR_spRouteTimeout.....	61	ZWIRSEC_EncryptionAlgorithm_t.....	74
ZWIR_spRoutingTableSize.....	61	ZWIRSEC_encAESCTR.....	74
ZWIR_TimeoutCallback_t.....	41	ZWIRSEC_encNull.....	74
ZWIR_TRXStatistic_t.....	53	ZWIRSEC_EncryptionSuite_t.....	74
ZWIR_UART_FlowControl_t		ZWIRSEC_ErrorCode_t	
ZWIR_UART_HWFlowControl.....	64	ZWIRSEC_eCorruptedPacket.....	75
ZWIR_UART_NoFlowControl.....	64	ZWIRSEC_eDroppedICMP.....	75
ZWIR_UART_Parity_t		ZWIRSEC_eDroppedPacket.....	75
ZWIR_UART_EvenParity.....	64	ZWIRSEC_eReplayedPacket.....	75
ZWIR_UART_NoParity.....	64	ZWIRSEC_eUnknownSPI.....	75
ZWIR_UART_OddParity.....	64	ZWIRSEC_PolicyType_t.....	73
ZWIR_UART_RXCallback_t.....	65	ZWIRSEC_ptInputApply.....	73
ZWIR_UART_StopBits_t		ZWIRSEC_ptInputBypass.....	73
ZWIR_UART_Stop_2.....	64	ZWIRSEC_ptInputDrop.....	73
ZWIR_UART_StopBits_t		ZWIRSEC_ptOutputApply.....	73
ZWIR_UART_Stop_1.....	64	ZWIRSEC_ptOutputBypass.....	73
ZWIR_UART1_ErrorCode_t		ZWIRSEC_ptOutputDrop.....	73
ZWIR_UART1_eFrame.....	66	ZWIRSEC_SecurityAssociation_t.....	74
ZWIR_UART1_eNoise.....	66		

12.3. Variables and Constants

ZWIR_firmwareMajorVersion.....	60	ZWIR_vendorID.....	60
ZWIR_firmwareMinorVersion.....	60	ZWIRSEC_ikeRekeyTime.....	76
ZWIR_firmwareVersionExtension.....	60	ZWIRSEC_ikeRetransmitTime.....	76
ZWIR_productID.....	60	ZWIRSEC_ikeSARekeyTime.....	77

13 Related Documents

IETF Documents
Internet Protocol, Version 6 (IPv6) Specification
IP Version 6 Addressing Architecture
Security Architecture for the Internet Protocol
Internet Key Exchange (IKEv2) Protocol
Neighbor Discovery for IP Version 6 (IPv6)
IPv6 Stateless Address Autoconfiguration
Transmission of IPv6 Packets over IEEE 802.15.4 Networks

IDT Documents
<i>ZWIR4512 Data Sheet</i>
<i>ZWIR451x Application Note – Using IPSec and IKEv2 in 6LoWPANs*</i>
<i>ZWIR451x Application Note – Enabling Firmware Over-the-Air Updates*</i>
<i>ZWIR45xx Application Note – The NetMA Protocol*</i>

Visit www.IDT.com/ZWIR4512 or contact your nearest sales office for the latest version of these documents.

Note: Documents marked with an asterisk () require a free customer login account for access.

14 Glossary

Term	Description
6LoWPAN	IPv6 over Low Power Wireless Personal Area Networks
AES	Advanced Encryption Standard
AH	Authentication Header
API	Application Programming Interface
ARP	Address Resolution Protocol
BPSK	Binary Phase Shift Keying
CBC	Cyclic Block Cipher
DAD	Duplicate Address Detection
DHCP	Dynamic Host Configuration Protocol
EAP	Extensible Authentication Protocol
ESP	Encapsulating Security Payload
GPIO	General Purpose Input/Output
ICMP	Internet Control Message Protocol
IETF	Internet Engineering Task Force
IKEv2	Internet Key Exchange version 2
IPSec	Internet Protocol Security
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
MAC	Media Access Control
MLD	Multicast Listener Discovery
MTU	Maximum Transmission Unit
LAN	Local Area Network
MCU	Micro Controller Unit
MODP	Modular Exponential
NA	Neighbor Advertisement
NAT	Network Address Translation
NetMA	Network Monitoring and Administration
NDP	Neighbor Discovery Protocol
NS	Neighbor Solicitation
NVIC	Nested Vectored Interrupt Controller
OSI	Open Systems Interconnection

Term	Description
PAN	Personal Area Network
PLL	Phase-Locked Loop
PSK	Pre Shared Key
QPSK	Quadrature Phase Shift Keying
RA	Router Advertisement
RC	Resistive Capacitive
RFC	Request for Comments (a type of technical document maintained by the Internet Engineering Task Force)
RS	Router Solicitation
RSSI	Receive Signal Strength Indicator
RTC	Real-Time Clock
SA	Security Association
SAD	Security Association Database
SP	Security Policy
SPD	Security Policy Database
SWD	Serial Wire Debug
TFC	Traffic Flow Confidentiality
TRX	Transceiver
UART	Universal Asynchronous Receiver Transmitter
UDP	User Datagram Protocol
WAN	Wide Area Network
WPAN	Wireless Personal Area Network

15 Document Revision History

Revision	Date	Description
1.00	October 1, 2010	Initial Version
1.10	December 13, 2010	Added I/O pin descriptions for SAM3S based modules (e.g. ZWIR4522-I). Added interrupt vector table for ZWIR452x-I modules (e.g. ZWIR4511-I). Corrected declaration of ZWIR_DiscoveryCallback_t. Replaced invalid declaration ZWIR_GetPacketRSSI with ZWIR_GetLastRSSI. Minor edits.
1.20	March 27, 2011	Renamed document to <i>ZWIR451x Programming Guide</i> . Removed parts of documentation dedicated to SAM3S-based modules. Adapted documentation to library release 1.2. Cross-linked all symbols.
1.30	July 5, 2011	Minor revisions for clarity.
1.40	November 17, 2011	Added libGPIO documentation. Minor revisions for clarity.
1.60	June 18, 2012	Added documentation of new functionality provided with API version 1.6. Many text improvements for clarity. Fixed error in Table 2.1. Added documentation for NetMA functions and types.
1.61	July 27, 2012	Minor edits.
1.62	August 31, 2012	Added documentation of FCC-ID readout command. Added R&TTE & FCC conformity statements.
1.90	August 24, 2014	Added documentation of new functionality provided with API version 1.9. Added UART error codes Multiple text improvements for clarity.
	April 12, 2016	Changed to IDT branding.

IMPORTANT NOTICE AND DISCLAIMER

RENESAS ELECTRONICS CORPORATION AND ITS SUBSIDIARIES (“RENESAS”) PROVIDES TECHNICAL SPECIFICATIONS AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES “AS IS” AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT OF THIRD-PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for developers who are designing with Renesas products. You are solely responsible for (1) selecting the appropriate products for your application, (2) designing, validating, and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, or other requirements. These resources are subject to change without notice. Renesas grants you permission to use these resources only to develop an application that uses Renesas products. Other reproduction or use of these resources is strictly prohibited. No license is granted to any other Renesas intellectual property or to any third-party intellectual property. Renesas disclaims responsibility for, and you will fully indemnify Renesas and its representatives against, any claims, damages, costs, losses, or liabilities arising from your use of these resources. Renesas' products are provided only subject to Renesas' Terms and Conditions of Sale or other applicable terms agreed to in writing. No use of any Renesas resources expands or otherwise alters any applicable warranties or warranty disclaimers for these products.

(Disclaimer Rev.1.01 Jan 2024)

Corporate Headquarters

TOYOSU FORESIA, 3-2-24 Toyosu,
Koto-ku, Tokyo 135-0061, Japan
www.renesas.com

Trademarks

Renesas and the Renesas logo are trademarks of Renesas Electronics Corporation. All trademarks and registered trademarks are the property of their respective owners.

Contact Information

For further information on a product, technology, the most up-to-date version of a document, or your nearest sales office, please visit www.renesas.com/contact-us/.