


# User Documentation - PCIe-2SG1




# Table of contents

- [User Documentation - PCIe-25G1](#)
  - [Product description](#)
  - [Revision history](#)
  - [Warranty and support](#)
  - [Safety and regulatory information](#)
  - [Overview](#)
    - [Specifications](#)
    - [Product architecture](#)
    - [Recommended technical expertise](#)
  - [Planning](#)
    - [Power consumption and power budget](#)
    - [Cooling and thermal considerations](#)
    - [PCI mapping](#)
    - [Material, information and software required](#)
    - [Platform, modules and accessories](#)
    - [Validated server platforms and operating systems](#)
    - [Resources for monitoring via SMBus](#)
    - [PCIe-25G1 form factor details](#)
  - [Getting started](#)
    - [Getting started - Application installation and performance benchmarking](#)
  - [Configuring the server](#)
    - [Configuring basic UEFI options](#)
    - [Configuring server PCIe bifurcation](#)
  - [Installing](#)
    - [Mechanical installation and precautions](#)
      - [ESD protections](#)
      - [Unboxing](#)
      - [Installation and assembly](#)
    - [Software installation and deployment](#)
      - [Installing software required for the PCIe-25G1](#)
      - [Verifying installation](#)
      - [Common software installation](#)
  - [Configuring](#)
    - [Configuring server thermal parameters](#)
    - [Accessing the FRU](#)
  - [Operating](#)
    - [Monitoring](#)
      - [Monitoring sensors](#)
      - [Monitoring GPU resources using the OS](#)
    - [Maintenance](#)
      - [Upgrading SG1 GPU firmware](#)
  - [Troubleshooting](#)
    - [Troubleshooting GPU](#)
    - [Collecting diagnostics](#)
    - [Support information](#)
    - [Working with error messages](#)
  - [Document symbols and acronyms](#)

# Product description

Table of contents

- [Main applications](#)
- [Main features](#)



The insatiable demand for video puts a tremendous pressure on TV delivery infrastructures. The PCIe-2SG1 is an optimal solution to address density and power challenges, and further reduce total cost of ownership. This PCIe add-in card is optimized for video transcoding applications and provides the desired flexibility to further optimize software and create differentiated products. It is a next generation dual SG1 server GPU accelerator card.

## Main applications

- Video transcoding applications
- Decreasing the cost per stream, increasing the density and reducing the power footprint

## Main features

- Latest generation Intel® Server GPU providing an unparalleled feature set for video transcoding
- High quality CODECs with interlaced video support
- Extensive CODEC support (MPEG-2, AVC, HEVC, VP9, 4K/8K HEVC 10-bit)
- Open source software infrastructure including Media SDK and FFmpeg framework
- Transcodes up to 8 UHD/4K HEVC 60FPS 40 FULL HD AVC 30FPS stream per card\*
- PCIe Gen 4 x16 (bifurcation required to dual x8)
- High throughput PCIe Gen 4 interface to host processor
- Two SG1 Intel® GPUs per card with 8 GB DDR4 each
- PCIe form factor: half length, low profile (half height), single width
- Passive cooling
- < 65 W, no need for external power
- Operational with CentOS 7.4, CentOS 8.2, Debian 10.4 and Ubuntu 20.04

\*Cf. Intel benchmarks

## Revision history

Revision	Brief description of changes	Date of issue
1.0	First release version	May 2022
1.1	<ul style="list-style-type: none"><li>• Added SG1 firmware update information in Operating&gt;Maintenance&gt;Upgrading</li><li>• Added links to product web page for downloads</li></ul>	November 2022

# Warranty and support

Table of contents

- [Limited warranty](#)
- [Disclaimer](#)
- [Customer support](#)
- [Customer service](#)

## Limited warranty

Please refer to the full terms and conditions of the Standard Warranty on Kontron's website at:

[https://www.kontron.com/support-and-services/rma/canada/standard\\_warranty\\_policy\\_canada.pdf](https://www.kontron.com/support-and-services/rma/canada/standard_warranty_policy_canada.pdf).

## Disclaimer

Kontron would like to point out that the information contained in this manual may be subject to alteration, particularly as a result of the constant upgrading of Kontron products. This document does not entail any guarantee on the part of Kontron with respect to technical processes described in the manual or any product characteristics set out in the manual. Kontron assumes no responsibility or liability for the use of the described product(s), conveys no license or title under any patent, copyright or mask work rights to these products and makes no representations or warranties that these products are free from patent, copyright or mask work right infringement unless otherwise specified. Applications that are described in this manual are for illustration purposes only. Kontron makes no representation or warranty that such application will be suitable for the specified use without further testing or modification. Kontron expressly informs the user that this manual only contains a general description of processes and instructions which may not be applicable in every individual case. In cases of doubt, please contact Kontron.

This manual is protected by copyright. All rights are reserved by Kontron. No part of this document may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language or computer language, in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), without the express written permission of Kontron. Kontron points out that the information contained in this manual is constantly being updated in line with the technical alterations and improvements made by Kontron to the products and thus this manual only reflects the technical status of the products by Kontron at the time of publishing. Brand and product names are trademarks or registered trademarks of their respective owners.

©2022 by Kontron

## Customer support

Kontron's technical support team can be reached through the following means:

- By phone: 1-888-835-6676
- By email: [support-na@kontron.com](mailto:support-na@kontron.com)
- Via the website: [www.kontron.com](http://www.kontron.com)

## Customer service

Kontron, a trusted technology innovator and global solutions provider, uses its embedded market strengths to deliver a service portfolio that helps companies break the barriers of traditional product lifecycles.

Through proven product expertise and collaborative, expert support, Kontron provides unparalleled peace of mind when it comes to building and maintaining successful products. To learn more about Kontron's service offering—including enhanced repair services, an extended warranty, and the Kontron training academy—visit [www.kontron.com/support-and-services](http://www.kontron.com/support-and-services).

# Safety and regulatory information

Table of contents

- [General safety warnings and cautions](#)
  - [Elevated operating ambient temperature](#)
  - [Reduced air flow](#)
  - [CE mark](#)
- [General instructions on usage](#)
- [Quality and environmental management](#)
  - [Disposal and recycling](#)
  - [Waste electrical and electronic equipment directive](#)
- [Regulatory specifications](#)

## NOTICE

Before working with this product or performing instructions described in the getting started section or in other sections, read the Safety and regulatory information section pertaining to the product. Assembly instructions in this documentation must be followed to ensure and maintain compliance with existing product certifications and approvals. Use only the described, regulated components specified in this documentation.

## General safety warnings and cautions

Your new Kontron product was developed and tested carefully to provide all features necessary to ensure its compliance with electrical safety requirements. It was also designed for a long fault-free life. However, the life expectancy of your product can be drastically reduced by improper treatment during unpacking and installation. Therefore, in the interest of your own safety and of the correct operation of your new Kontron product, you are requested to respect the following guidelines.

### WARNING

To prevent a fire or shock hazard, do not expose this product to rain or moisture. The chassis should not be exposed to dripping or splashing liquids and no objects filled with liquids should be placed on the chassis cover.



### ESD sensitive device!

This equipment is sensitive to static electricity. Care must therefore be taken during all handling operations and inspections of this product in order to ensure product integrity at all times.

### CAUTION

All operations on this product must be carried out by sufficiently skilled personnel only.

### Elevated operating ambient temperature

If this product is installed in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be greater than the ambient temperature of the room. Therefore, be careful to install the product in an environment that is compatible with the maximum operating temperature specified by the manufacturer in the specifications.

### Reduced air flow

Do not compromise on the amount of air flow required for safe operation when installing this product in a rack. Side clearances must be respected.

### CE mark

The CE marking on this product indicates that it is in compliance with the applicable European Union Directives: Low Voltage, EMC, Radio Equipment and RoHS requirements.

## General instructions on usage

In order to maintain Kontron's product warranty, this product must not be altered or modified in any way. Changes or modifications to the product, that are not explicitly approved by Kontron and described in this user guide or received from Kontron Support as a special handling instruction, will void your warranty.

This product should only be installed in or connected to systems that fulfill all necessary technical and specific environmental requirements. This also applies to the operational temperature range of the specific board version that must not be exceeded. If batteries are present, their temperature restrictions must be taken into account.

In performing all necessary installation and application operations, only follow the instructions supplied by the present user guide. Keep all the original packaging material for future storage or warranty shipments. If it is necessary to store or ship the product, then re-pack it in the same manner as it was delivered. Special care is necessary when handling or unpacking the product.

## Quality and environmental management

Kontron aims to deliver reliable high-end products designed and built for quality, and aims to comply with environmental laws, regulations, and other environmentally oriented requirements. For more information regarding Kontron's quality and environmental responsibilities, visit <https://www.kontron.com/en/quality-management>.

### Disposal and recycling

Kontron's products are manufactured to satisfy environmental protection requirements where possible. Many of the components used are capable of being recycled. Final disposal of this product after its service life must be accomplished in accordance with applicable country, State and/or local laws or regulations.

### Waste electrical and electronic equipment directive

This product contains electrical or electronic materials. If not disposed of properly, these materials may have potential adverse effects on the environment and human health. The presence of this logo on the product means it should not be disposed of as unsorted waste and must be collected separately. Dispose of this product according to the appropriate local rules, regulations and laws.

WEEE directive logo



## Regulatory specifications

This product complies with the Low Voltage Directive, 2014/35/EU and EN 62368-1.  
This product meets FCC Part 15/ICES-003 Class A. It is designed to meet GR-1089 and GR-63.  
This product complies with the Electromagnetic Compatibility Directive 2014/30/EU.  
This product complies with CISPR 32 Class A and CISPR 35.

# Overview

- [Specifications](#)
- [Product architecture](#)
- [Recommended technical expertise](#)



# Specifications

Table of contents

- [Key hardware features](#)
- [Key software features](#)
- [Physical dimensions](#)
- [Packaging physical dimensions](#)
- [Shipping weights](#)
- [Environmental specifications](#)

When you select a server to install the PCIe-25G1 in, make sure it meets the requirements in terms of:

- Cooling airflow
- Power
- Physical dimensions


## Key hardware features

Feature	Description
Form factor	PCIe add-in card, half length, low profile (half height), single width Low profile and full height (factory installed) brackets included
Interface	PCIe Gen 4 x16 Bifurcation required to obtain two x8 logical ports, one for each GPU
GPU	Two Intel® Server SG1 GPU servers
Memory	8 GB LPDDR4x per GPU, 16 GB total Maximum memory speed of 4267 MT/s Memory bus width of 128 bits per GPU
Product cooling	Passive cooling, server-based forced air cooling required Refer to <a href="#">Cooling and thermal considerations</a> for details
Power inlet	12 V rail: 11.04 V to 12.96 V 3.3 V rail: 3.003 V to 3.597 V 3.3 Vaux: 3.003 V to 3.597 V
Power consumption	< 65 W No auxiliary power required Refer to <a href="#">Power consumption and power budget</a> for details
Server management	Temperature sensors and FRU ID EEPROM accessible via PCIe SMBus interface Refer to <a href="#">Resources for monitoring via SMBus</a> for details
Reliability	Designed for 24/7 operation
MTBF	931,489 hours at 30°C
Availability	Long-life availability (until January 2026)

## Key software features

Feature	Description
Compatible operating systems	CentOS 7.4, CentOS 8.2, Debian 10.4, Ubuntu 20.04
Software support	Intel® Media SDK, FFmpeg Linux kernel driver for SG1 support  Refer to <a href="#">Material, information and software required</a> for details

## Physical dimensions



Card	Measurements (mm [in])
Depth	167.65 [6.600]
Width	23.13 [0.911]
Height	69 [2.717]

## Packaging physical dimensions

Depth (mm [in])	Width (mm [in])	Height (mm [in])
347 [13.7]	178 [7.0]	56 [2.2]

## Shipping weights

Component	Weight (g)	Weight (lb)
Card weight with low profile bracket	183	0.41
Packaging (card + brackets + box + foam + bag)	402	0.89

## Environmental specifications

Environment	Specification
Operating temperature	0°C to 50°C, at card intake based on system air flow
Storage temperature	-40°C to 70°C
RoHS and WEEE	This product complies with EU directive 2012/19/EU (WEEE). This product complies with RoHS directive 2011/65/EU as modified by EU 2015/863. This product is designed to meet China RoHS Phase 1 (self-declaration and labeling).
Compliance/regulatory	Meets the following environmental, safety, and EMC requirements: <ul style="list-style-type: none"> <li>• FCC Part 15 / ICES-003 Class A</li> <li>• UL 62368-1 / CSA C22.2 No. 62368</li> <li>• CE mark <ul style="list-style-type: none"> <li>◦ Low Voltage Directive, 2014/35/EU and EN 62368-1</li> <li>◦ Electromagnetic Compatibility Directive 2014/30/EU</li> </ul> </li> <li>• CISPR 32 Class A and CISPR 35</li> </ul> Telcordia GR-63 (designed to meet) Telcordia GR-1089 (designed to meet)

# Product architecture

Table of contents

- [Description](#)
- [Block diagram](#)

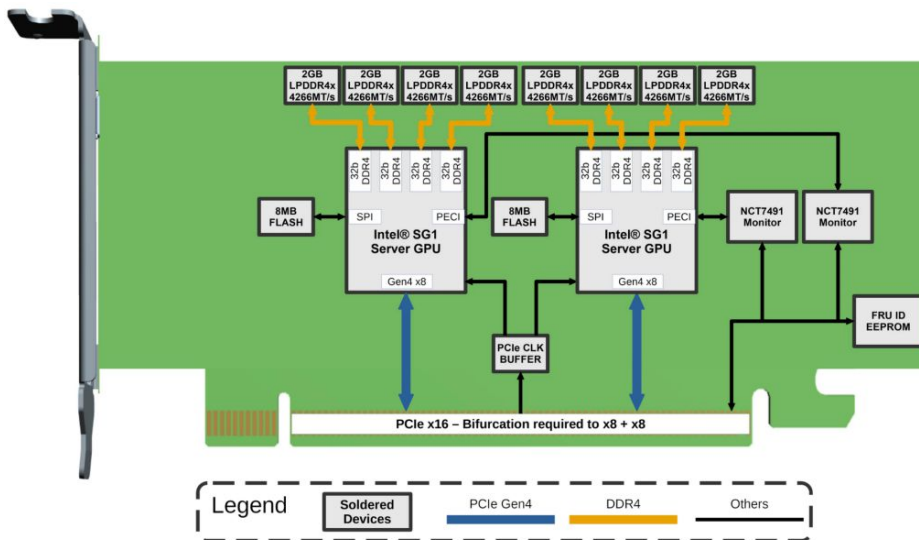
## Description

Kontron's PCIe-2SG1 card is a PCIe Gen 4 x16 with two Intel® SG1 GPUs on board. The standard x16 edge connector implementation is split into two x8 lane groups, each SG1 using one of the x8 group.

Each SG1 GPU is monitored by its dedicated hardware monitor (NCT7491). These monitors can be accessed through the SMBus port on the PCIe edge connector if the motherboard provides SMBus access.

The PCIe-2SG1 only uses the standard PCIe power rails provided by the edge connector. No auxiliary power connection is required.

## Block diagram



## Recommended technical expertise

Appropriate technical expertise is required to install, configure and operate this product.

More specifically, knowledge of the following domains is required:

- Linux kernel driver installation
- Intel® Media SDK or FFmpeg video CODEC use
- Server thermal configuration and monitoring

# Planning

- [Power consumption and power budget](#)
- [Cooling and thermal considerations](#)
- [PCI mapping](#)
- [Material, information and software required](#)
- [Platform, modules and accessories](#)
- [Validated server platforms and operating systems](#)
- [Resources for monitoring via SMBus](#)
- [PCIe-2SG1 form factor details](#)

# Power consumption and power budget

Table of contents

- [Power input voltage](#)
- [Power consumption for various use cases](#)

## Power input voltage

DC input voltage	12 V rail	3.3 V rail	3.3 Vaux
Nominal	12 V	3.3 V	3.3 V
Minimum	11.04 V	3.003 V	3.003 V
Maximum	12.96 V	3.597 V	3.597 V
<b>Power consumption</b>			
Maximum	60 W		

## Power consumption for various use cases

Use case	Power consumption (12 V rail)
Idle	9 W
Typical transcoding on both GPUs*	Average: 34.5 W Peak: 48 W
Maximum load**	60 W

\* Tested using Intel® Media SDK sample\_multi\_transcode utility to perform 8 parallel transcoding threads (4 on each GPU) from 4k60 UHD 75 Mbps to 4k60 UHD 75 Mbps.

NOTE:

- Tests were performed at an ambient temperature of 25°C .
- Power consumption was measured at the 12 V input rail. Power on 3.3 V rail is about 1 W under maximum load. \*\*

\*\* Tested using Intel® proprietary test tool for maximum SG1 power consumption.

# Cooling and thermal considerations


Table of contents

- [Cooling requirements](#)
- [Required cooling airflow](#)

## Cooling requirements

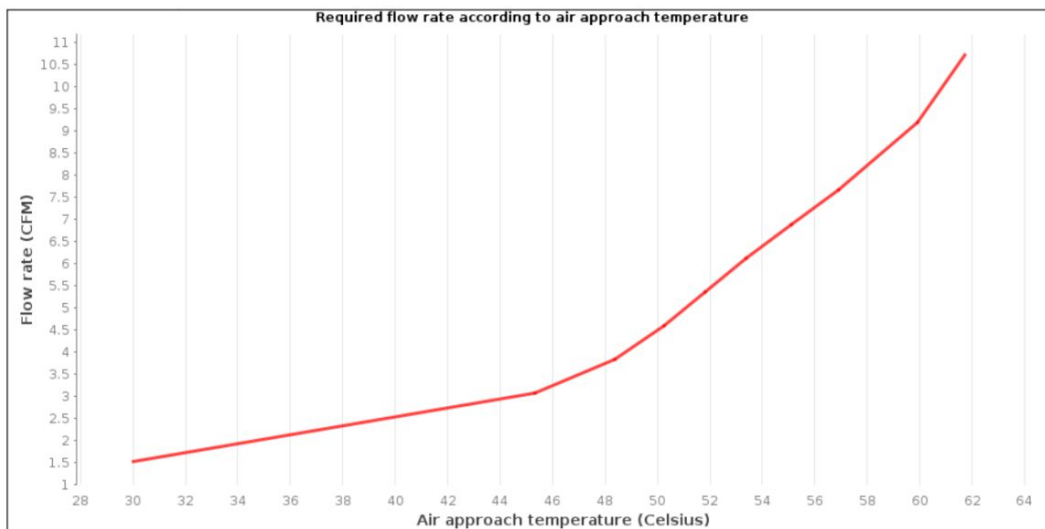
The PCIe-25G1 uses passive cooling and therefore requires server-based forced air cooling to operate below the maximum operating temperature of the GPUs. To avoid operating under the base frequency of the SG1 because of thermal throttling and to ensure long life of the product, the temperature of the GPU on-die digital thermal sensor (DTS) should stay under 95°C\* during card operation. See [Monitoring sensors](#) to learn more about monitoring the GPUs of the PCIe-25G1.

\*This is an absolute value. It is recommended to plan for an additional margin to account for system and product variations. The PCIe-25G1 supports bidirectional airflow. The CFM requirements are the same for both directions.



**NOTE :** If the PCIe-25G1 is not integrated into a server closed loop fan control system, fan speed must be set to an appropriate level so the card does not overheat when the server is running. As a preliminary safe configuration, Kontron recommends setting fan speed to the maximum. Please refer to [Configuring server thermal parameters](#) for suggestions tailored to different platforms.

## Required cooling airflow

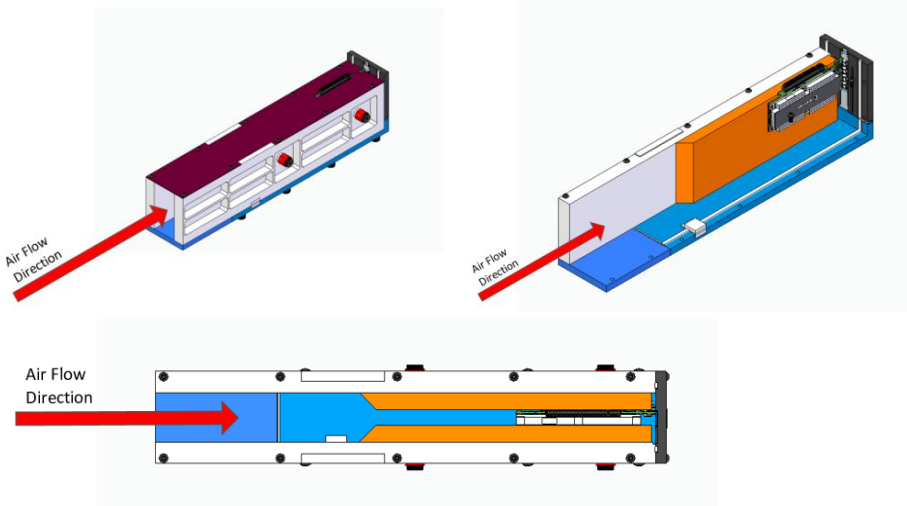


- **Flow rate** : Volumetric flow rate of the air measured in cubic foot per minute (CFM).
- **Air approach temperature** : Temperature of the forced air at the PCIe-25G1 card inlet (heat sink (forward airflow) or face plate (reverse airflow)).

The curve presented was produced with data obtained in a laboratory setting using an enclosure with dimensions respecting the PCI Express Card Electromechanical Specification (Version 1 of the 5th revision). Data shown on the graph corresponds to the minimum airflow through the inlet of the test enclosure (59.35 mm wide x 120 mm high) (see the images below) to keep the card below its maximum operating temperature. During this test, the PCIe-25G1 card was stressed at maximum power consumption using the Intel® proprietary test tool. The data obtained does not cover all use cases. It should be used as a reference for estimation purposes during



thermal design. Since the thermal performance of each host server will be different, validations must be performed with the PCIe-2SG1 installed in the server in which it will be used. Another option is to contact a Kontron representative for thermal performance validation.



## PCI mapping

Two SG1 GPUs are installed on the PCIe-2SG1. The following table lists the devices for one SG1 instance. Once bifurcation has been configured correctly, the list below will be present for both SG1 GPUs with different starting bus numbers. Refer to [Configuring server PCIe bifurcation](#) for further information.

**NOTE:** In the first column, the [BUS] variable represents the value assigned by the firmware of the server hosting the PCIe-2SG1.

Bus: Device. Function	Vendor ID	Device ID	Component	Description
[BUS]:00.00	0x8086	0x4910	PCI bridge	PCI Express Upstream Switch device
[BUS+1]:01.00	0x8086	0x490f	PCI bridge	PCI Express Downstream Switch device
[BUS+1]:04.00	0x8086	0x490f	PCI bridge	PCI Express Downstream Switch device
[BUS+1]:05.00	0x8086	0x490f	PCI bridge	PCI Express Downstream Switch device
[BUS+2]:00.00	0x8086	0x4907	VGA compatible controller	Actual GPU device within the SG1
[BUS+3]:00.00	0x8086	0x490d	Audio device	High Definition Audio within the SG1
[BUS+4]:00.00	0x8086	0x490e	Memory controller	

# Material, information and software required

Table of contents

- [Material and information required](#)
  - [Component installation and assembly](#)
- [Software Required](#)

## Material and information required

### Component installation and assembly

Item_1	PCIe I/O brackets with two M3x4mm screws (one low profile and one full height), included in the box
Item_2	T10 Torx screwdriver

## Software Required

Item_1	SG1 software support package including Linux kernel driver and Media SDK NOTE: Contact <a href="#">Customer support</a> to determine the appropriate software package for your application.
--------	--

## Platform, modules and accessories

This section provides the complete list of compatible parts and components that can be ordered from Kontron.



The PCIe-2SG1 card can be installed in MSP nodes used in MS2900 series platform. Refer to [kontron.com](http://kontron.com) for information in this specific product.

Description	Kontron P/N	Illustration
MP_AS_L-BRACKET_PCIE_HH_SG1_ST	1068-0388	
MP_AS_L-BRACKET_PCIE_FH_SG1_ST	1069-6016	

# Validated server platforms and operating systems

Table of contents

- [Validated server platforms](#)
  - [Intel® M50CYP1UR204 \(formerly Coyote Pass\)](#)
    - [Configuration](#)
  - [Intel® M50CYP2UR208 \(formerly Coyote Pass\)](#)
    - [Configuration](#)
  - [HPE® DL360 Gen10](#)
    - [Configuration](#)
- [Validated operating systems](#)
  - [Status description](#)
  - [OS certification status](#)

## Validated server platforms

### Intel® M50CYP1UR204 (formerly Coyote Pass)

VALIDATED

The validation of this server platform has been completed internally.

#### Configuration

CPU	Dual Intel® Xeon® Gold 6330
Memory	16x DDR4 16GB DIMM
Number of PCIe-25G1 cards	2

### Intel® M50CYP2UR208 (formerly Coyote Pass)

VALIDATED

The validation of this server platform has been completed internally.

#### Configuration

CPU	Dual Intel® Xeon® Gold 6330
Memory	16x DDR4 16GB DIMM
Number of PCIe-25G1 cards	4

### HPE® DL360 Gen10

PLANNED

The validation of this server platform is planned.

#### Configuration

CPU	Dual Intel® Xeon® Silver 4210R
Memory	8x DDR4 16GB DIMM
Number of PCIe-25G1 cards	3

## Validated operating systems

### Status description

Status legend	Description
<b>CERTIFIED</b>	The product is certified by the OS vendor as compliant hardware
<b>VALIDATED</b>	The product was internally tested
<b>TESTED CERT</b>	The unit passed the certification tests, but the official OS vendor certificate was not published
<b>PLANNED</b>	Certification is planned
<b>IN PROCESS</b>	Certification has started

## OS certification status

Operating system	Status
CentOS 7.4	<b>VALIDATED</b>
CentOS 8.2	<b>VALIDATED</b>
Debian 10.4	<b>VALIDATED</b>
Ubuntu 20.04	<b>VALIDATED</b>

# Resources for monitoring via SMBus

Table of contents

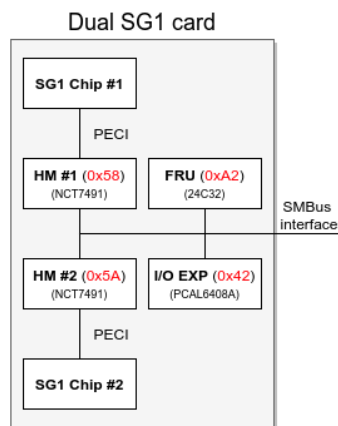
- [Resources for monitoring temperature sensors via SMBus](#)
- [Monitoring GPU temperature](#)
  - [Programming the registers](#)
  - [Getting GPU temperature](#)
- [PCIe-2SG1 card FRU information](#)
- [Reserved SMBus resources](#)

## Resources for monitoring temperature sensors via SMBus



The contents of this page describe elements for a custom BMC integration. Contact the OEM server manufacturer for more information about this type of integration.

Two temperature sensors are used to get the temperature of the two SG1 GPUs. These sensors can be monitored through the SMBus. Each sensor reading is obtained from a hardware monitor device on the PCIe-2SG1. Typical servers implement control of the PCIe add-in card SMBus in the BMC. In these cases, the BMC firmware must integrate functions to read the PCIe-2SG1 sensors described in this section to manage cooling resources and potentially display SG1 GPU temperatures. **If the PCIe-2SG1 is not integrated into a server closed loop fan control system, the speed of the fan must be set to 100% so the card does not overheat when the server is running.** Communicate with the server manufacturer to add this support. Within the PCIe-2SG1, NCT7491 hardware monitors from onsemi are used to retrieve GPU temperature data. The hardware monitors retrieve data via the PECL interface of each SG1 GPU. An EEPROM, also accessible via the SMBus, contains information about the card installed in the system. The following diagram illustrates the entire SMBus chain on the PCIe-2SG1. Refer to [Product architecture](#) for more information.



## Monitoring GPU temperature

Each GPU of the PCIe-2SG1 is connected to a dedicated NCT7491 hardware monitor through its PECL interface. Refer to the NCT7491 datasheet for complete information on the registers.

Device	Description	SMBus address
NCT7491	Thermal monitoring for GPU1	0x58
NCT7491	Thermal monitoring for GPU2	0x5A

### Programming the registers

**NOTE:** Kontron provides this information for reference purposes as the actual method will depend on the server in which the PCIe-2SG1 card is installed.

To retrieve the temperature using the SMBus, the registers of each hardware monitor must be programmed adequately.

The SMBus of PCIe slots are typically controlled by the motherboard BMC. In this case, the BMC will need to issue the described register programming using its I2C tools. It is also possible that the PCIe slot SMBus is controlled by an SMBus host in the server's hardware. In

this case, Linux I2C tools could be used to implement the register programming. In each case, the SMBus host controller could then read the GPU temperatures.

Typical programming steps could be as follows:

Step_1	Write <b>0x10</b> to register <b>PECI Config 3</b> (offset 0x37). <ul style="list-style-type: none"> <li>• Bit 7: CPU writes are disabled</li> <li>• Bit 5-4: Update rate is 2/sec</li> <li>• Bit 1: Retry is enabled</li> <li>• Bit 0: Processor scanning is enabled</li> </ul>
Step_2	Write <b>0x32</b> to register <b>PECI0 Address</b> (offset 0x00). <ul style="list-style-type: none"> <li>• Bit 7-0: PECI address of the GPU</li> </ul>
Step_3	Write <b>0x10</b> to register <b>Configuration 1</b> (offset 0x40). <ul style="list-style-type: none"> <li>• Bit 4: Enable and start PECI monitoring</li> </ul>
Step_4	Repeat steps 1 to 3 for every hardware monitor.

## Getting GPU temperature

Once the register programming is done, the value stored in register **PECIO** (offset 0x33) will be the GPU temperature. This value is stored in 2's complement format. It is always negative and represents the thermal margin before reaching the maximum junction temperature of the GPU ( $T_{jmax}$ ). To obtain the value of  $T_{jmax}$  required for programming, contact [Customer support](#).

To get the real GPU temperature in degrees C, the register value must be added to the  $T_{jmax}$  value using the following formula:

$$\text{GPU temperature} = (\text{value from PECIO}) + T_{jmax}$$

## PCIe-2SG1 card FRU information

PCIe-2SG1 manufacturing information can be retrieved using the embedded 24C32 I2C EEPROM accessible via the PCI Express SMBus interface described above.

Device	Description	SMBus address
24C32	Serial EEPROM	0xA2

This EEPROM contains a FRU structure as per the IPMI FRU Information Storage Definition document. Here is a content example:

```

hexdump -C sg1_mod.dat
00000000 01 01 00 00 21 00 00 dd 01 00 00 00 00 00 00 00 |.....!.....|
00000010 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
*
00000100 00 00 00 00 00 00 00 00 01 08 00 c7 4b 6f 6e 74 |.....Kont|
00000110 72 6f 6e c9 50 43 49 45 2d 32 53 47 31 c9 31 30 |ron.PCIe-2SG1.10|
00000120 36 38 2d 30 37 30 30 c1 32 ca 30 39 38 37 36 35 |68-0700.2.098765|
00000130 34 33 32 31 c0 cf 46 52 55 50 43 49 45 2d 32 53 |4321..FRUPCIe-2S|
00000140 47 31 2d 30 31 c1 00 c5 00 00 00 00 00 00 00 00 |G1-01.....|
00000150 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
*
000003f0 11 04 00 00 00 00 00 00 01 01 00 00 21 00 00 dd |.....!.....|
00000400

```

Based on the IPMI FRU Information Storage Definition, the information above translates to:

Product Manufacturer: Kontron  
Product Name: PCIe-2SG1  
Product Part Number: 1068-0700  
Product Version: 2  
Product Serial: 0987654321  
Product Asset Tag:  
Product FRU ID: FRUPCIe-2SG1-0

## Reserved SMBus resources

The SMBus also provides an I/O Expander (PCAL6408A). This device is only for Kontron's internal use.



Device	Description	SMBus address
PCAL6408A	I/O Expander	0x42

# PCIe-2SG1 form factor details


Table of contents

- [Card reference](#)
- [Card dimensions](#)
- [Card keep-outs](#)
- [Card retainer](#)
- [PCI defeat tab](#)

This page provides more information about the PCIe-2SG1 form factor. It describes implementation in the context of the PCI-SIG Low Profile Half Length. This page should be used to ensure compatibility with the system hosting a PCIe-2SG1 card.

## Card reference

The following figure identifies terminology used in this page to describe the card edges of the PCIe-2SG1.



## Card dimensions

Refer physical dimensions in the [Specifications](#) section. A standard x16 PCIe connector has been implemented on the south edge of the card.

## Card keep-outs

The PCIe-2SG1 does not observe a 3-mm keep-out on the north edge of the card, as required by the PCIe CEM specification (rev 4.0) to accommodate chassis-level card retention.






Image source: PCI express Card Electromechanical Specification, Revision 4.0, Version 1.0, p.140

The PCIe-2SG1 does observe the 1-mm keep-outs on the north, south, east and west edges.



## Card retainer

A card retainer, commonly referred to as a "hockey stick," has been implemented on the PCIe-2SG1 card for compatibility with host system retention mechanisms.



Associated keep-outs have also been respected.





Image source: PCI express Card Electromechanical Specification, Revision 4.0, Version 1.0, p.140

## PCI defeat tab

A PCI defeat tab has been implemented to prevent insertion into legacy PCI slots.



## Getting started

- [Getting started - Application installation and performance benchmarking](#)

# Getting started - Application installation and performance benchmarking

## Table of contents

- [Introduction](#)
  - [Assumptions](#)
- [Unboxing the platform](#)
  - [What's in the box](#)
- [Planning](#)
  - [Material and information required](#)
  - [Software required](#)
- [Configuring the platform](#)
  - [Disabling the Compatibility Support Module \(CSM\)](#)
  - [Enabling Memory Mapped IO \(MMIO\) mapping above 4 GB](#)
  - [Setting Memory Mapped IO \(MMIO\) size to 256 GB or higher](#)
  - [Setting PCIe link speed](#)
  - [Configuring bifurcation to x8x8 on the slot where the PCIe-25G1 is connected](#)
  - [Disabling the host server integrated GPU \(iGPU\) \(server dependent\)](#)
- [Installing components](#)
  - [Installing or changing an I/O bracket](#)
  - [Installing the PCIe-25G1 card in a server](#)
- [Installing related software](#)
  - [Installing the Linux kernel driver](#)
  - [Installing Intel® Media SDK for SG1](#)
- [Verifying installation](#)
  - [Verifying PCI Express link status](#)
  - [Running a sanity check test](#)
- [Benchmarking an application](#)
- [Monitoring a PCIe-25G1 card](#)

### NOTICE

Before working with this product or performing instructions described in the getting started section or in other sections, read the Safety and regulatory information section pertaining to the product. Assembly instructions in this documentation must be followed to ensure and maintain compliance with existing product certifications and approvals. Use only the described, regulated components specified in this documentation. Use of other products/components will void the CSA certification and other regulatory approvals of the product and will most likely result in non-compliance with product regulations in the region(s) in which the product is sold.

## Introduction

This getting started section describes the installation and integration steps required to start operating a PCIe-25G1 card inside a server.

### Assumptions

The scenario described in this getting started section is based on the following assumptions:

- The host server has a x16 PCIe Low Profile Half Length compatible free slot that supports bifurcation settings (x8x8)
- The host server has enough airflow to cool the passive heat sink of the PCIe-25G1
- The host server has a Linux OS already installed and compatible with the SG1 software support package

## Unboxing the platform



### ESD sensitive device!

This equipment is sensitive to static electricity. Care must therefore be taken during all handling operations and inspections of this product in order to ensure product integrity at all times.

## What's in the box

The box includes:

- One PCIe-25G1 with a full height I/O bracket installed using two M3x4mm screws
- One low profile I/O bracket (no screws)

Step_1	Carefully remove the PCIe card and the low profile I/O bracket from the box.
Step_2	Remove the PCIe card from the ESD bag.

## Planning

### Material and information required

Item_1	PCIe I/O brackets with two M3x4mm screws (one low profile and one full height), included in the box
Item_2	T10 Torx screwdriver

### Software required

Item_1	SG1 software support package including Linux kernel driver and Media SDK NOTE: Contact <a href="#">Customer support</a> to determine the appropriate software package for your application.
--------	--

> Proceed with the server configurations required prior to installing the PCIe card in the server.

## Configuring the platform

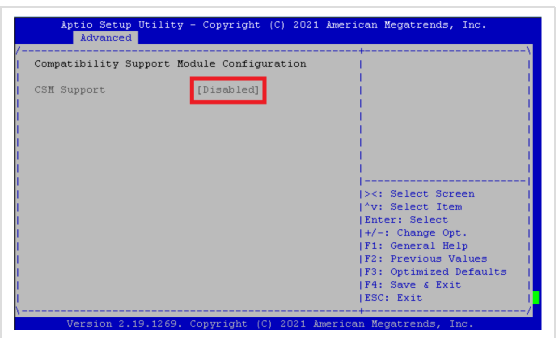
The following configurations are typically required for a PCIe-25G1 to work correctly within a server. Other configurations may be required depending on the host server. Note that all the configurations are server dependent.

**NOTE:** The following procedures were documented using a PCIe-25G1 card installed in an MSP8070 series node included in an MS2900 series platform. An AMI UEFI is used as an example.

To perform the configurations described in this section, access the UEFI menu using the preferred method based on the host server.

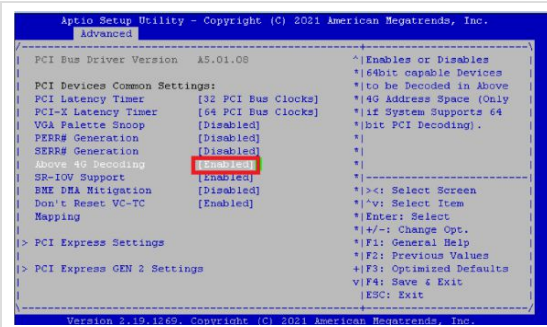
### Disabling the Compatibility Support Module (CSM)

Step_1	From the UEFI setup menu, navigate to the <b>Advanced</b> menu and select <b>CSM Support</b> . Set the option to <b>Disabled</b> .
--------	---



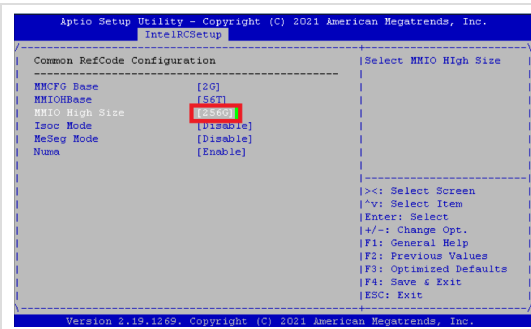
### Enabling Memory Mapped IO (MMIO) mapping above 4 GB

Step\_1 From the UEFI setup menu, navigate to the **Advanced** menu, then **PCI Subsystem Settings** and select **Above 4G Decoding** .  
Set the option to **Enabled** .



## Setting Memory Mapped IO (MMIO) size to 256 GB or higher

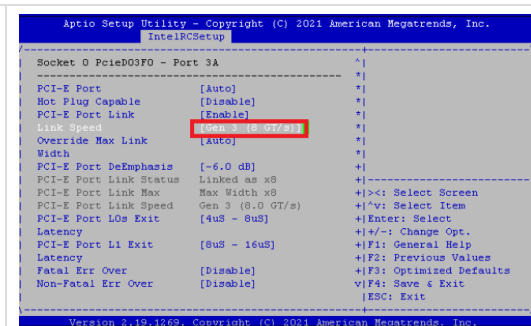
Step\_1 From the UEFI setup menu, navigate to the **IntelRCSetup** menu and select **MMIO High Size** .  
Set the option to **256G** .



## Setting PCIe link speed

The PCIe-25G1 link speed can be set to: 2.5 GT/s (Gen 1), 4 GT/s (Gen 2), 8 GT/s (Gen 3) or 16 GT/s (Gen 4). For further details, refer to [Specifications](#) .

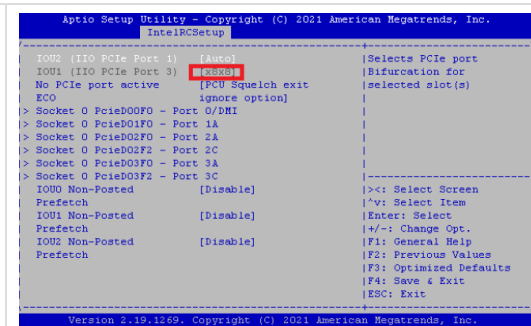
Step\_1 From the UEFI setup menu, navigate to the **IntelRCSetup** menu and select the port where the PCIe-25G1 is connected. Select **Link Speed** .  
Set the option based on the functionality required.



## Configuring bifurcation to x8x8 on the slot where the PCIe-25G1 is connected

Step\_1 From the UEFI setup menu, navigate to the **IntelRCSetup** menu and select the port where the PCIe-25G1 is connected. Select **IIO PCIe Port 3** .  
Set the option to **x8x8** .

**NOTE:** Refer to [Configuring server PCIe bifurcation](#) for more details on how to configure this settings on other platforms.



## Disabling the host server integrated GPU (iGPU) (server dependent)

Some servers have a processor with an integrated GPU. If this is your case, it must be disabled. Refer to the server manufacturer's documentation.

> You now have the material, software and configurations required. Proceed with the installation of the PCIe card in the server.

## Installing components



### ESD sensitive device!

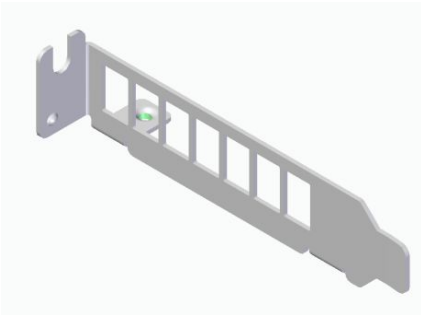
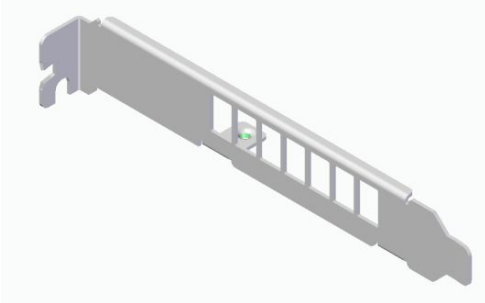
This equipment is sensitive to static electricity. Care must therefore be taken during all handling operations and inspections of this product in order to ensure product integrity at all times.



Disconnect the power supply cord before servicing the product to avoid electric shock. If the product has more than one power supply cord, disconnect them all.

## Installing or changing an I/O bracket

There are two I/O brackets available, and the procedure to install or remove them is identical.

Low profile	
Full height (factory installed)	

Step_1	(Optional) If an I/O bracket is already installed on the PCIe add-in card, remove it. Remove the two M3x4mm screws using a T10 Torx screwdriver.
Step_2	Install the I/O bracket selected by fastening the two M3x4mm screws using a T10 Torx screwdriver (6 lbs-in torque).

## Installing the PCIe-2SG1 card in a server

Install the PCIe-2SG1 card according to the instructions provided in the server manufacturer's documentation.

> You now have the correct material configuration. Proceed with software installation in your server.

## Installing related software

### Installing the Linux kernel driver

Follow the instructions provided in the SG1 software support package downloaded.

### Installing Intel® Media SDK for SG1

Follow the instructions provided in the SG1 software support package downloaded.



# Verifying installation

## Verifying PCI Express link status

This command will verify the link status of the two GPUs installed on the PCIe-2SG1 card. There will be two results, one for each bifurcation.

The command requires administrative privileges (sudo).

Step_1	From the OS, open a command line interface.
Step_2	<p>Confirm the link with the PCIe-2SG1 card is active and at the correct speed.</p> <pre>LocalServer_OSPrompt:~# lspci -vvv -d 8086:4910   grep -E "4910 LnkSta"</pre> <p><b>NOTE:</b> The SG1 GPU supports a 16GT/s (if the system is rated for PCIe Gen 4). For most applications, an 8 GT/s (with a server rated for PCIe Gen 3 for example) can be sufficient. If the speed or width is lower than expected, check the UEFI settings. For more information please refer to <a href="#">Troubleshooting GPU</a>.</p> <pre># lspci -vvv -d 8086:4910   grep -E "4910 LnkSta" [PCIe_BUS_ADDRESS] PCI bridge: Intel Corporation Device 4910 (prog-if 00 [Normal decode]) LnkSta: Speed 8GT/s, Width x8, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt- LnkSta2: Current De-emphasis Level: -6dB, EqualizationComplete+, EqualizationPhase1+ [PCIe_BUS_ADDRESS] PCI bridge: Intel Corporation Device 4910 (prog-if 00 [Normal decode]) LnkSta: Speed 8GT/s, Width x8, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt- LnkSta2: Current De-emphasis Level: -6dB, EqualizationComplete+, EqualizationPhase1+</pre>

## Running a sanity check test

The Intel SG1 software support package contains a Media SDK package that includes a sanity check script. Refer to [Software required](#) for details.

Running the following command executes a quick transcode test on both GPUs to verify their installation.

Step_1	From the OS, open a command-line interface.
Step_2	Extract the Media SDK package.
Step_3	Change to the root directory of the extracted package.
Step_4	<p>Confirm that two GPU "cards" are detected for every PCIe-2SG1 installed in the system. The script output uses the term card to refer to a chip.</p> <pre>LocalServer_OSPrompt:~# ./media_sanity_test.sh</pre> <p><b>NOTE:</b> Both GPUs should be detected with a PASS status. If multiple PCIe-2SG1 cards are installed in the system, more render instances could be available.</p> <pre># ./media_sanity_test.sh .....Sanity test begin..... Intel GPU card 0: PCI:0000:07:00.0, ID:0x4907 card:/dev/dri/card1 render:/dev/dri/renderD128 transcode sanity PASS  Intel GPU card 1: PCI:0000:0c:00.0, ID:0x4907 card:/dev/dri/card2 render:/dev/dri/renderD129 transcode sanity PASS  .....end.....</pre>

> The installation is now completed. Proceed with application benchmarking.

## Benchmarking an application

Install your application and proceed with benchmarking.

## Monitoring a PCIe-2SG1 card

Refer to [Monitoring sensors](#) for monitoring strategies for the PCIe-2SG1 card.

# Configuring the server

# Configuring basic UEFI options

Table of contents

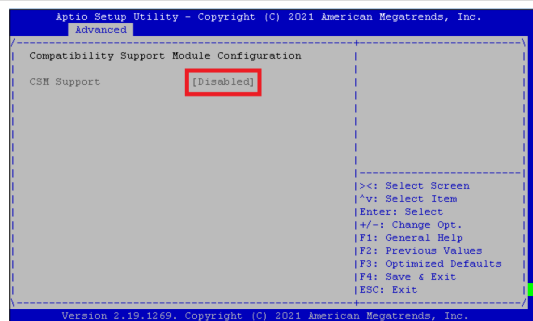
- [Disabling the Compatibility Support Module \(CSM\)](#)
- [Enabling Memory Mapped IO \(MMIO\) mapping above 4 GB](#)
- [Setting Memory Mapped IO \(MMIO\) size to 256 GB or higher](#)
- [Setting PCIe link speed](#)
- [Configuring bifurcation to x8x8 on the slot where the PCIe-25G1 is connected](#)
- [Disabling the host server integrated GPU \(iGPU\) \(server dependent\)](#)

The following configurations are typically required for a PCIe-25G1 to work correctly within a server. Other configurations may be required depending on the host server. Note that all the configurations are server dependent.

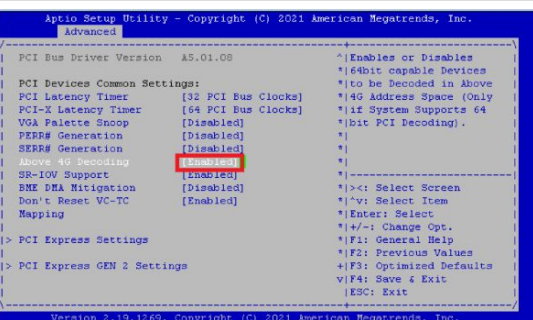
**NOTE:** The following procedures were documented using a PCIe-25G1 card installed in an MSP8070 series node included in an MS2900 series platform. An AMI UEFI is used as an example.

To perform the configurations described in this section, access the UEFI menu using the preferred method based on the host server.

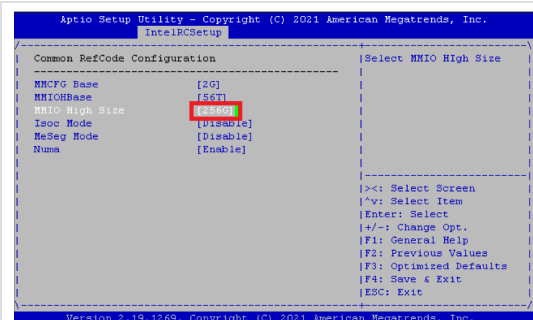
## Disabling the Compatibility Support Module (CSM)

<p>Step_1 From the UEFI setup menu, navigate to the Advanced menu and select CSM Support . Set the option to Disabled .</p>	
---	---

## Enabling Memory Mapped IO (MMIO) mapping above 4 GB

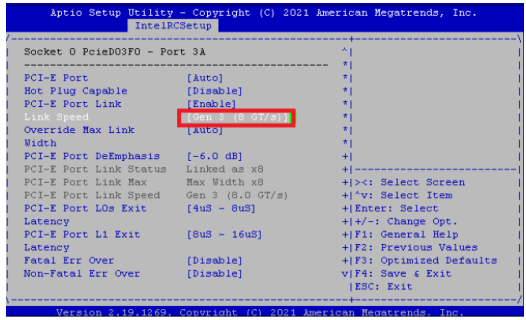
<p>Step_1 From the UEFI setup menu, navigate to the Advanced menu, then PCI Subsystem Settings and select Above 4G Decoding . Set the option to Enabled .</p>	
---	--

## Setting Memory Mapped IO (MMIO) size to 256 GB or higher

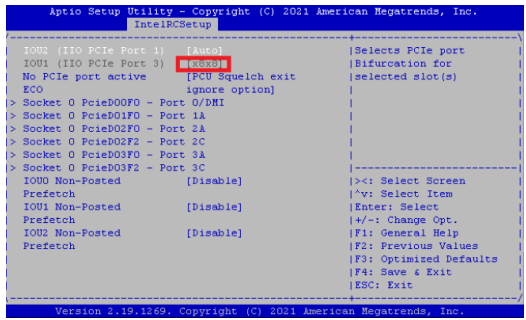
<p>Step_1 From the UEFI setup menu, navigate to the IntelRCSetup menu and select MMIO High Size . Set the option to 256G .</p>	
--	--

## Setting PCIe link speed

The PCIe-25G1 link speed can be set to: 2.5 GT/s (Gen 1), 4 GT/s (Gen 2), 8 GT/s (Gen 3) or 16 GT/s (Gen 4). For further details, refer to [Specifications](#).

Step_1	From the UEFI setup menu, navigate to the IntelRCSetup menu and select the port where the PCIe-25G1 is connected. Select <b>Link Speed</b> . Set the option based on the functionality required.	
--------	--	--

## Configuring bifurcation to x8x8 on the slot where the PCIe-25G1 is connected

Step_1	<p>From the UEFI setup menu, navigate to the IntelRCSetup menu and select the port where the PCIe-25G1 is connected. Select <b>IIO PCIe Port 3</b> . Set the option to x8x8 .</p> <p><b>NOTE:</b> Refer to <a href="#">Configuring server PCIe bifurcation</a> for more details on how to configure this settings on other platforms.</p>	
--------	---	---

## Disabling the host server integrated GPU (iGPU) (server dependent)

Some servers have a processor with an integrated GPU. If this is your case, it must be disabled. Refer to the server manufacturer's documentation.

# Configuring server PCIe bifurcation

Table of contents

- [Setting a PCIe bifurcation on a DELL® server](#)
- [Setting PCIe bifurcation on an Intel® server](#)
- [Setting PCIe bifurcation on an HPE® server](#)

For the PCIe-25G1 to work properly inside a system, the x16 PCIe slot where the card is installed needs to be split into two x8 lane groups. This is done using the PCIe bifurcation setting on the host server. Note that all the configurations are server dependent.

## Setting a PCIe bifurcation on a DELL® server

These are the steps for a DELL® R6515 server.

Step_1	Access the BIOS Setup Utility during boot.
Step_2	Select <b>System Bios Settings&gt; Integrated Devices&gt; Slot Bifurcation</b> .
Step_3	Select the proper slot and change the bifurcation to <b>x8x8 Bifurcation</b> . <b>NOTE:</b> Server documentation will contain a description of the various parameters available.

## Setting PCIe bifurcation on an Intel® server

These are the steps for Intel® M50CYP servers.

Step_1	Access the BIOS Setup Utility using the <b>F2</b> key during boot.
Step_2	Select <b>Advanced &gt; Integrated IO Configuration</b> .
Step_3	Select <b>PCIe Slot Bifurcation Setting</b> .
Step_4	Select the proper Riser/slot and change the setting to <b>x8x8</b> . <b>NOTE:</b> Server documentation will contain a description of the various parameters available.

## Setting PCIe bifurcation on an HPE® server

These are the steps for HPE® ProLiant Gen10 servers.

Step_1	Access the <b>UEFI System Utilities</b> using the <b>F9</b> key during boot.
Step_2	Select <b>System Configuration &gt; BIOS/Platform Configuration (RBSU) &gt; PCIe Device Configuration &gt; Advanced PCIe Configuration &gt; PCIe Bifurcation Options</b> .
Step_3	Select the proper Riser/slot and change the setting to <b>Bifurcate</b> . <b>NOTE:</b> Server documentation will contain a description of the various parameters available.

# Installing

- [Mechanical installation and precautions](#)
  - [ESD protections](#)
  - [Unboxing](#)
  - [Installation and assembly](#)
- [Software installation and deployment](#)
  - [Installing software required for the PCIe-25G1](#)
  - [Verifying installation](#)
  - [Common software installation](#)


## Mechanical installation and precautions

- [ESD protections](#)
- [Unboxing](#)
- [Installation and assembly](#)

# ESD protections

Electrostatic discharge (ESD) can damage electronic components (e.g. disk drives and boards).

Look for this warning in the documentation as it indicates that the device is ESD sensitive and that precautions must be taken.


	<p><b>ESD sensitive device!</b> This equipment is sensitive to static electricity. Care must therefore be taken during all handling operations and inspections of this product in order to ensure product integrity at all times.</p>
---	---

We recommend that you perform all the installation procedures described in the documentation at an ESD workstation. If this is not possible, apply ESD protections such as the following:

- Wear an antistatic wrist strap attached to a chassis ground (any unpainted metal surface) on the equipment when handling parts.
- Touch the metal chassis before touching an electronic component (e.g. a DIMM or board).
- Keep a part of your body (e.g. a hand) in contact with the metal chassis to dissipate the static charge while handling the electronic component.
- Avoid moving around unnecessarily.
- Use a ground strap attached to the front panel (with the bezel removed).
- Read and follow the safety precautions provided for a specific component by the manufacturer.



# Unboxing

	<p><b>ESD sensitive device!</b> This equipment is sensitive to static electricity. Care must therefore be taken during all handling operations and inspections of this product in order to ensure product integrity at all times.</p>
---	---

## What's in the box

The box includes:

- One PCIe-2SG1 with a full height I/O bracket installed using two M3x4mm screws
- One low profile I/O bracket (no screws)

Step_1	Carefully remove the PCIe card and the low profile I/O bracket from the box.
Step_2	Remove the PCIe card from the ESD bag.

# Installation and assembly

Table of contents

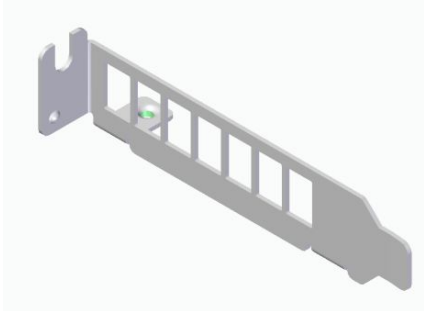
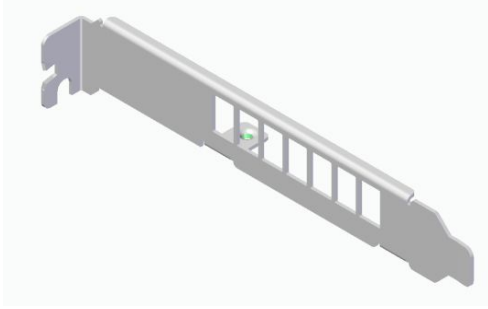
- [Installing the PCIe-25G1 card in a server](#)
- [Installing or changing an I/O bracket](#)

## Installing the PCIe-25G1 card in a server

Install the PCIe-25G1 card according to the instructions provided in the server manufacturer's documentation.

## Installing or changing an I/O bracket

There are two I/O brackets available, and the procedure to install or remove them is identical.

Low profile	
Full height (factory installed)	

Step_1	(Optional) If an I/O bracket is already installed on the PCIe add-in card, remove it. Remove the two M3x4mm screws using a T10 Torx screwdriver.
Step_2	Install the I/O bracket selected by fastening the two M3x4mm screws using a T10 Torx screwdriver (6 lbs-in torque).

# Software installation and deployment

- [Installing software required for the PCIe-25G1](#)
- [Verifying installation](#)
- [Common software installation](#)

# Installing software required for the PCIe-2SG1

Table of contents

- [Prerequisites](#)
- [Installing the Linux kernel driver](#)
- [Installing Intel® Media SDK for SG1](#)
- [Installing the sg\\_monitor tool as an option for GPU monitoring](#)

## Prerequisites

1	The host server has a Linux OS already installed and compatible with the SG1 software support package.
2	The appropriate SG1 software support packages (as determined by Kontron) are available.

Visit the [PCIe-2SG1 product page](#) for SG1 software support packages. Contact [Customer support](#) for more software support package options.

Relevant section:

[Material, information and software required](#)

## Installing the Linux kernel driver

Follow the instructions provided in the SG1 software support package downloaded.

## Installing Intel® Media SDK for SG1

Follow the instructions provided in the SG1 software support package downloaded.

## Installing the sg\_monitor tool as an option for GPU monitoring

There are two options to monitor GPU temperature:

- Via the SMBus of the PCIe slot in which the card is installed. Refer to [Resources for monitoring via SMBus](#) and contact the OEM server manufacturer for support.
- Via the OS using the sg\_monitor tool. For further details on monitoring, refer to [Monitoring GPU resources using the OS](#).

If you choose to monitor using the OS, proceed with the installation of the sg\_monitor tool.

Step_1	Extract the files of the sg_monitor .zip package on the host server.
Step_2	Access the directory where the files were extracted. The package name is provided as an example.  LocalServer_OSPrompt:~# cd sg_monitor_pv1.5/sg_monitor_sample/
Step_3	Compile the sg_monitor tool. LocalServer_OSPrompt:~/sg_monitor_pv1.5/sg_monitor_sample# make

# Verifying installation

Table of contents

- [Verifying PCI Express link status](#)
- [Running a sanity check test](#)

To perform the following verifications, access the OS using the preferred method based on the host server.

Relevant sections:

- [Material, information and software required](#)
- [Configuring basic UEFI options](#)

## Verifying PCI Express link status

This command will verify the link status of the two GPUs installed on the PCIe-25G1 card. There will be two results, one for each bifurcation.

The command requires administrative privileges (sudo).

Step_1	From the OS, open a command line interface.
Step_2	<p>Confirm the link with the PCIe-25G1 card is active and at the correct speed.            LocalServer_OSPrompt:~# lspci -vvv -d 8086:4910   grep -E "4910 LnkSta"</p> <p><b>NOTE:</b> The SG1 GPU supports a 16GT/s (if the system is rated for PCIe Gen 4). For most applications, an 8 GT/s (with a server rated for PCIe Gen 3 for example) can be sufficient. If the speed or width is lower than expected, check the UEFI settings. For more information please refer to <a href="#">Troubleshooting GPU</a>.</p> <pre># lspci -vvv -d 8086:4910   grep -E "4910 LnkSta" [PCIe_BUS_ADDRESS] PCI bridge: Intel Corporation Device 4910 (prog-if 00 [Normal decode]) LnkSta: Speed 8GT/s, Width x8, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt- LnkSta2: Current De-emphasis Level: -6dB, EqualizationComplete+, EqualizationPhase1+ [PCIe_BUS_ADDRESS] PCI bridge: Intel Corporation Device 4910 (prog-if 00 [Normal decode]) LnkSta: Speed 8GT/s, Width x8, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt- LnkSta2: Current De-emphasis Level: -6dB, EqualizationComplete+, EqualizationPhase1+</pre>

## Running a sanity check test

The Intel SG1 software support package contains a Media SDK package that includes a sanity check script. Refer to [Software required](#) for details.

Running the following command executes a quick transcode test on both GPUs to verify their installation.

Step_1	From the OS, open a command-line interface.
Step_2	Extract the Media SDK package.
Step_3	Change to the root directory of the extracted package.
Step_4	<p>Confirm that two GPU "cards" are detected for every PCIe-25G1 installed in the system. The script output uses the term card to refer to a chip.</p> <pre>LocalServer_OSPrompt:~# ./media_sanity_test.sh</pre> <p><b>NOTE :</b> Both GPUs should be detected with a PASS status. If multiple PCIe-25G1 cards are installed in the system, more render instances could be available.</p> <pre># ./media_sanity_test.sh .....Sanity test begin..... Intel GPU card 0: PCI:0000:07:00.0, ID:0x4907 card:/dev/dri/card1 render:/dev/dri/renderD128 transcode sanity PASS  Intel GPU card 1: PCI:0000:0c:00.0, ID:0x4907 card:/dev/dri/card2 render:/dev/dri/renderD129 transcode sanity PASS  .....end.....</pre>

## Common software installation

### Required software tools

Tool	Description
pciutils	Tool used to manage PCIe cards connected to the platform.

# Configuring

- [Configuring server thermal parameters](#)
- [Accessing the FRU](#)

# Configuring server thermal parameters

Table of contents

- [Configuring parameters when the PCIe-2SG1 has custom user-controlled GPU thermal sensors](#)
  - [Setting sensor thresholds in the BMC](#)
- [Configuring parameters when the PCIe-2SG1 has no custom user-controlled GPU thermal sensors](#)
  - [Statically setting minimum fan speed](#)
    - [Setting minimum fan speed on a DELL® R6515 server](#)
    - [Setting minimum fan speed on Intel® M50CYP servers](#)
      - [Setting the minimum fan speed using the BIOS Setup Utility](#)
      - [Setting the minimum fan speed using IPMI](#)
    - [Setting minimum fan speed on HPE® ProLiant Gen10 servers](#)
  - [Dynamically setting fan speed based on actual SG1 GPU temperature](#)

Relevant section:

[Cooling and thermal considerations](#)

There may be thermal parameters to adjust in the server to account for the installation of the PCIe-2SG1 card. The method will vary depending on whether there is a custom integration of the card with the BMC of the server and user-controlled GPU thermal sensors are available.

## Configuring parameters when the PCIe-2SG1 has custom user-controlled GPU thermal sensors

This configuration procedure can be used if a custom integration to the BMC was programmed by the OEM server manufacturer. Refer to [Resources for monitoring via SMBus](#) for information.

### Setting sensor thresholds in the BMC

If user-controlled sensors are available for the SG1 GPUs in the BMC, here are suggested threshold settings that could be implemented.

LowerCritical (LCR)	UpperNonCritical (UNC)	UpperCritical (UCR)	UpperNonRecoverable (UNR)
0°C	75°C	95°C	110°C

The values recommended above are provided as a reference only. Validations must be performed with the PCIe-2SG1 installed in the server in which it will be used.

To avoid thermal throttling situations, the fan manager should be triggered by at least the UpperNonCritical (UNC) threshold. Refer to the server manufacturer user guide for BMC threshold configuration and to [Cooling and thermal considerations](#) for more thermal information on the PCIe-2SG1.

## Configuring parameters when the PCIe-2SG1 has no custom user-controlled GPU thermal sensors

This configuration procedure is to be used if no custom integration to the BMC was programmed by the OEM server manufacturer. The additional cooling required to account for the PCIe-2SG1 can be [statically set in servers](#) or [dynamically determined using data obtained through the sg\\_monitor tool](#).

### Statically setting minimum fan speed

These procedures can be used to increase the minimum fan speed to account for the integration of the PCIe-SG1 card.

The procedures are server specific and are detailed for a few specific servers. The list is not exhaustive and multiple methods (Web UI, CLI, UEFI) could be available.

### Setting minimum fan speed on a DELL® R6515 server

These are the steps in the Web UI for a server using iDRAC. This server has not yet been validated with the PCIe-2SG1 card, so no specific parameter can be recommended by Kontron for the minimum fan speed. Refer to the graph provided in section [Cooling and thermal considerations](#).



Step_1	Access the Web UI of the server's iDRAC.
Step_2	Select Configuration > System Settings > Hardware Settings > Cooling Configuration .
Step_3	Select the proper value for the Fan Speed Offset parameter. <b>NOTE:</b> Server documentation will contain a description of the various offset parameters available.

### Setting minimum fan speed on Intel® M50CYP servers

Server model	Number of PCIe-25G1 cards installed	Minimum fan PWM offset value recommended
M50CYP2UR204	2	18*
M50CYP2UR208	4	40*

\*The fan PWM offsets suggested should prevent GPU throttling on these servers operating at a maximum room temperature of 35°C. The values were obtained while running the Intel® proprietary test tool for maximum SG1 power consumption and Prime95 torture test for the CPU. **These suggested values are for reference purposes only**, many factors could influence the minimum fan speed requirements on a server (CPU load, CPU SKU, GPU load, etc.). Clients need to validate the minimum fan speed required on their system with their application.

### Setting the minimum fan speed using the BIOS Setup Utility

Step_1	Access the BIOS Setup Utility using the F2 key during boot.
Step_2	Select Advanced > System Acoustic and Performance Configuration > Fan PWM Offset .
Step_3	Enter the desired Fan PWM Offset .

### Setting the minimum fan speed using IPMI

Step_1	Open a command line interface on the host OS (local) or on a remote server with the ipmitool utility installed (IOL).
Step_2	Issue a set fan speed offset command. LocalServer_OS Prompt:~# ipmitool raw 0x30 0x8c [OFFSET]  <b>NOTE:</b> The [OFFSET] can have a value between 0 and 100. The example above uses ipmitool locally on the host server.
Step_3	(Optional) Use the following command to confirm the offset programmed. LocalServer_OS Prompt:~# ipmitool raw 0x30 0x8d

### Setting minimum fan speed on HPE® ProLiant Gen10 servers

These are the steps for a server using iLO. This server has not yet been validated with the PCIe-25G1 card, so no specific parameter can be recommended by Kontron for the minimum fan speed. Refer to the graph provided in section [Cooling and thermal considerations](#).

Step_1	Access the Web UI of the server's iLO.
Step_2	Select Power & Thermal> Fans or Fans & Cooling Modules
Step_3	Click the Edit button (pencil icon).
Step_4	Enter the Minimum Fan Speed (in %) and click OK .

### Dynamically setting fan speed based on actual SG1 GPU temperature

It might be possible to dynamically override the minimum fan speed value used by the BMC fan manager based on the GPU temperature retrieved using the [sg\\_monitor tool](#) to mimic threshold based sensors. This could be done by sending a custom IPMI OEM raw command or via REST API, depending on the server's implementation. Refer to the server manufacturer's documentation.

## Accessing the FRU

The PCIe-Z5G1 is equipped with an EEPROM containing FRU data. A BMC with access to the PCIe slot's SMBus may be able to discover FRU data for the PCIe cards via SMBus scan processes.

Refer to the server manufacturer's documentation to see if this is supported on your device.

Refer to [Resources for monitoring via SMBus](#) for more details about the implementation of the PCIe-Z5G1 FRU.

# Operating

- [Monitoring](#)
  - [Monitoring sensors](#)
  - [Monitoring GPU resources using the OS](#)
- [Maintenance](#)
  - [Upgrading SG1 GPU firmware](#)

# Monitoring

- [Monitoring sensors](#)
- [Monitoring GPU resources using the OS](#)

# Monitoring sensors

Table of contents

- [Monitoring GPU temperature](#)
  - [Via the PCI Express SMBus interface](#)
  - [Via the PCI Express bus](#)

## Monitoring GPU temperature

### Via the PCI Express SMBus interface

Two temperature sensors are used to get the temperature of the two SG1 GPUs. These sensors can be monitored through the SMBus. Each sensor reading is obtained from a hardware monitor device on the PCIe-2SG1.

Refer to [Resources for monitoring via SMBus](#).

### Via the PCI Express bus

A tool can be used to monitor the PCIe-2SG1 temperature, power consumption, memory usage, etc.

Refer to [Monitoring GPU resources using the OS](#).

# Monitoring GPU resources using the OS

Table of contents

- [Important metrics](#)
- [Reading monitored values using sg\\_monitor](#)

Relevant section:

[Installing software required for the PCIe-2SG1](#)

PCIe-2SG1 cards can be monitored directly using the `sg_monitor` tool. The tool displays information about the GPU temperature, power consumption, memory usage and more.

Contact [Customer support](#) to get the software package for this tool.

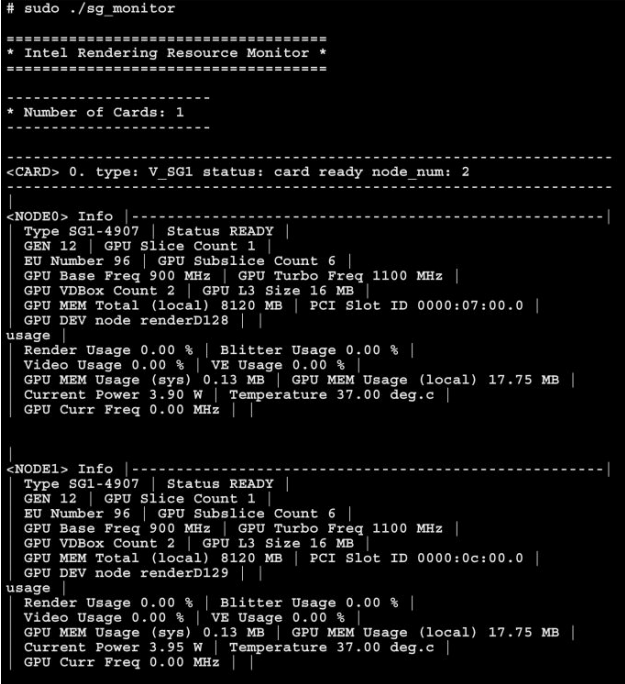
## Important metrics

Here is a list of key metrics to monitor on the PCIe-2SG1 card:

- **Current Power** : Instantaneous power consumption of a GPU. Summing both GPU power values together gives an estimate of the total power consumption of the card. See [Power consumption and power budget](#) for more information about this topic.
- **GPU Curr Freq** : Instantaneous clock frequency of the GPU.
- **Temperature** : Instantaneous maximum on-die digital thermal sensor (DTS) value of a GPU. See [Cooling and thermal considerations](#) for more information about this topic.

## Reading monitored values using sg\_monitor

NOTE: API calls are available for this tool. For information, consult the README file inside the tool's package.

Step_1	Access the directory where the <code>sg_monitor</code> tool was installed. The package name is provided as an example. LocalServer_OSPrompt:~# cd sg_monitor_pv1.5/sg_monitor_sample/
Step_2	Run the tool with the following command. LocalServer_OSPrompt:~# sudo ./sg_monitor 

# Maintenance

- [Upgrading SG1 GPU firmware](#)

# Upgrading SG1 GPU firmware

Table of contents

- [Upgrading SG1 GPU firmware from the server operating system](#)
  - [Prerequisites](#)
  - [Procedure](#)

## Upgrading SG1 GPU firmware from the server operating system

### Prerequisites

1	The PCIe-2SG1_GPU-FWIGSC_v[x].[y].tar.gz archive provided by Kontron is on a filesystem accessible to the Linux OS installed on the platform hosting the PCIe-2SG1 card(s).
2	The igsc tool requires that the driver for the SG1 GPU be installed in the Linux OS installed on the platform hosting the PCIe-2SG1 card(s). This driver is available via Intel SG1 software support package.

Visit the [PCIe-2SG1 product page](#) for the latest firmware archive as well as SG1 software support packages.  
Contact [Customer support](#) for more software support package options .

Relevant section:

[Installing software required for the PCIe-2SG1](#)

### Procedure

The source repository can be consulted for information on build requirements and procedure. The repository can be found at: <https://github.com/intel/igsc.git>

The update package archive readme file will give the igsc repository tag version used in testing.

Note that some commands in this procedure need to be run with root privileges. The examples below use **sudo** to attain root privileges.

Step_1	Access the operating system and open a command line interface.
Step_2	Uncompress the .tar.gz archive on the Linux OS installed on the platform. LocalServer_OSPrompt:~# tar -zxvf [FILE_NAME].tar.gz
Step_3	Clone the igsc repository to a known directory. LocalServer_OSPrompt:~# git clone --branch [VERSION_TAG] https://github.com/intel/igsc.git [CLONE_DIRECTORY_NAME]
Step_4	Follow the repository instructions to build the igsc tool.
Step_5	Access the folder created by the firmware update archive in Step_2. LocalServer_OSPrompt:~# cd [FILE_NAME]
Step_6	Rename the included pre-compiled igsc tool. LocalServer_OSPrompt:~/[FILE_NAME]# mv igsc igsc.archive
Step_7	Create a symlink to the igsc binary compiled in Step_4. LocalServer_OSPrompt:~/[FILE_NAME]# ln -s [CLONE_DIRECTORY_NAME]/builddir/src/igsc .
Step_8	Check currently running versions and also ensure all the installed SG1 GPU are detected LocalServer_OSPrompt:~/[FILE_NAME]# sudo ./igsc list-devices --info
Step_9	Run the update script LocalServer_OSPrompt:~/[FILE_NAME]# sudo ./update.sh The script will list SG1 GPU found and display update progress for each in turn
Step_10	Reboot the server.



# Troubleshooting

- [Troubleshooting GPU](#)
- [Collecting diagnostics](#)
- [Support information](#)
- [Working with error messages](#)

# Troubleshooting GPU

Table of contents

- [No GPU detected or partial GPU detection](#)
- [Poor GPU performance](#)

If troubleshooting steps fail to resolve the problem, contact [Kontron](#).

## No GPU detected or partial GPU detection

If during software installation verification no GPU or only one GPU is detected, perform the following.

Step_1	Verify that the card is properly seated in the server. To do so, follow the server manufacturer's instructions.
Step_2	Power on the server and check all the UEFI configurations. Refer to <a href="#">Configuring basic UEFI options</a> .
Step_3	Make sure that drivers are correctly installed by <a href="#">Running a sanity check test</a> .

## Poor GPU performance

If GPU performance is not as expected, the device could be throttling due to thermal issues. The PCIe bus bandwidth or device mapping may also be incorrect.

Step_1	Verify installation and check all the UEFI configurations. Refer to <a href="#">Verifying installation</a> and to <a href="#">Configuring basic UEFI options</a> .
Step_2	<a href="#">Verify PCIe link speed and width</a> .
Step_3	<a href="#">Check GPU resources</a> .
Step_4	Verify that sufficient airflow is supplied to the PCIe-25G1 card. Refer to <a href="#">Cooling and thermal considerations</a> .
Step_5	<a href="#">Monitor PCIe-25G1 thermal sensors</a> to confirm that GPUs are operating below their throttling temperature.

## Collecting diagnostics

Refer to [Monitoring sensors](#) for information on sensors.

Refer to [Monitoring GPU resources using the OS](#) for information on GPU resources.

## Support information

Kontron's technical support team can be reached through the following means:

- By phone: 1-888-835-6676
- By email: [support-na@kontron.com](mailto:support-na@kontron.com)
- Via the website: [www.kontron.com](http://www.kontron.com)

# Working with error messages

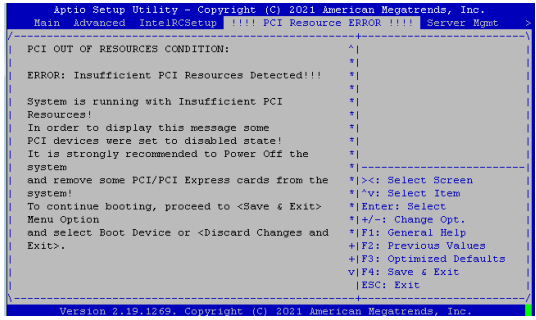
Table of contents

- [Unexpected errors](#)
  - [PCI resource error](#)

## Unexpected errors

### PCI resource error





If MMIO was not configured properly, unexpected PCI errors could be raised.


Step_1	Example of a PCI resource error recorded in the UEFI/BIOS.	 <pre>APTIO Setup Utility - Copyright (C) 2021 American Megatrends, Inc. Main Advanced IntelRSTSetup *    PCI Resource ERROR *    Server Mgmt ----- PCI OUT OF RESOURCES CONDITION: ERROR: Insufficient PCI Resources Detected!!! System is running with Insufficient PCI Resources! In order to display this message some PCI devices were set to disabled state! It is strongly recommended to Power Off the system and remove some PCI/PCI Express cards from the system! To continue booting, proceed to &lt;Save &amp; Exit&gt; Menu Option and select Boot Device or &lt;Discard Changes and Exit&gt;. -----  &gt;&lt;: Select Screen  ^v: Select Item  Enter: Select  +/-: Change Opt.  F1: General Help  F2: Previous Values  F3: Optimized Defaults  F4: Save &amp; Exit  ESC: Exit ----- Version 2.18.1268, Copyright (C) 2021 American Megatrends, Inc.</pre>
Step_2	Refer to <a href="#">Configuring basic UEFI options</a> and validate MMIO settings.	



# Document symbols and acronyms


## Symbols


The following symbols are used in Kontron documentation.

	DANGER indicates a hazardous situation which, if not avoided, will result in death or serious injury.
	WARNING indicates a hazardous situation which, if not avoided, could result in death or serious injury.
	CAUTION indicates a hazardous situation which, if not avoided, may result in minor or moderate injury.
	NOTICE indicates a property damage message.

	<p><b>Electric Shock!</b></p> <p>This symbol and title warn of hazards due to electrical shocks (&gt; 60 V) when touching products or parts of them. Failure to observe the precautions indicated and/or prescribed by the law may endanger your life/health and/or result in damage to your material.</p> <p>Please also refer to the "High-Voltage Safety Instructions" portion below in this section.</p>
---	--

	<p><b>ESD Sensitive Device!</b></p> <p>This symbol and title inform that the electronic boards and their components are sensitive to static electricity. Care must therefore be taken during all handling operations and inspections of this product in order to ensure product integrity at all times.</p>
	<p><b>HOT Surface!</b></p> <p>Do NOT touch! Allow to cool before servicing.</p>

	<p>This symbol indicates general information about the product and the documentation.</p> <p>This symbol also indicates detailed information about the specific product configuration.</p>
---	--

	<p>This symbol precedes helpful hints and tips for daily use.</p>
---	---

## Acronyms

ACPI	Advanced Configuration and Power Interface
AI	Artificial Intelligence
API	Application Programming Interface
ASIC	Application Specific Integrated Circuit
BIOS	Basic Input/Output System
BMC	Baseboard Management Controller
BSP	Board Support Package
CBIT	Continuous Built-In Test
CE	Community European (EU mark)
CLI	Command-Line Interface
CPU	Central Processing Unit
CRMS	Communications Rack Mount Servers
CSA	Canadian Standards Association
DC	Direct Current
DDR4	Double Data Rate Fourth Generation
DHCP	Dynamic Host Configuration Protocol

DHCP	Dynamic Host Configuration Protocol
DIMM	Dual Inline Memory Module
DRAM	Dynamic Random Access Memory
DTS	Digital Thermal Sensor
DU	Distributed Unit
ECC	Error Checking and Correcting
EEPROM	Electrically Erasable Programmable Read-Only Memory
EMC	Electromagnetic Compatibility
EMI	Electromagnetic Interference
ESD	Electrostatic Discharge
ETSI	European Telecommunications Standards Institute
ETSI	European Telecommunications Standards Institute
eUSB	Embedded Universal Serial Bus
FCC	Federal Communications Commission
FH/FL	Full Height/Full Length
FPGA	Field Programmable Gate Array
FRAU	Field Replaceable Unit
FRU	Field Replaceable Unit
Gb, Gbit	Gigabit
GB, Gbyte	Gigabyte – 1024 MB
GbE	Gigabit Ethernet
GND	Ground
GPI	General Purpose Input
GPIO	General Purpose Input/Output
GPO	General Purpose Output
GPS	Global Positioning System
GPU	Graphics Processing Unit
GUI	Graphical User Interface
HDD	Hard Disk Drive
Hz	Hertz – 1 cycle/second
I/O	Input/Output
I <sup>2</sup> C	Inter-Integrated Circuit Bus
iBMC	Integrated Baseboard Management Controller
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IMU	Inertial Measurement Unit
IOL	IPMI over LAN
IPMB	Intelligent Platform Management Bus
IPMI	Intelligent Platform Management Interface
IRQ	Interrupt Request Line
KB, Kbyte	Kilobyte – 1024 bytes
KCS	Keyboard Controller Style

KEAPI	Kontron Embedded Application Programming Interface
KVM	Keyboard, Video, Mouse
LAN	Local Area Network
LED	Light-Emitting Diode
LP	Low Profile
LPC	Low Pin Count
LVDS	Low Voltage Differential SCSI
MAT	Maximum Ambient Temperature
MB, Mbyte	Megabyte – 1024 KB
MCU	Microcontroller
MEC	Multi-Access Edge Computing
MXM	Mobile PCI Express Module
NCSI	Network Communications Services Interface
NEBS	Network Equipment-Building System
NIC	Network Interface Card, or Network Interface Controller, or Network Interface Controller port
NMI	Non-Maskable interrupt
NOS	Network Operating System
NVMe	Non-Volatile Memory Express
OCXO	Oven-Controlled Crystal Oscillator
OS	Operating System
OTP	Over-Temperature Protection
OVP	Over-Voltage Protection
PBIT	Power On Built-In Test
PCH	Platform Controller Hub
PCI	Peripheral Component Interconnect
PCIe	Peripheral Component Interconnect Express
PECI	Platform Environment Control Interface
PIRQ	PCI Interrupt Request Line
PMbus	Power Management Bus
PMM	POST Memory Manager
PnP	Plug and Play
POST	Power-On Self Test
PSU	Power Supply Unit
PTP	Precision Time Protocol
PXE	Preboot eXecution Environment
RAID	Redundant Array of Independent Disks
RAN	Radio Access Network
RAS	Reliability, Availability, and Serviceability
RDIMM	Registered Dual In-Line Memory Module
RDP	Remote Desktop
RMM	Remote Management Module



RoHS	Restriction of Hazardous Substances
SAS	Serial Attached SCSI (Small Computer System Interface)
SATA	Serial Advanced Technology Attachment
SCSI	Small Computer Systems Interface
SDRAM	Synchronous Dynamic RAM
SEL	System Event Log
SFP+	Small Form-factor Pluggable that supports data rates up to 10.0 Gbps
SMBus	System Management Bus
SMS	Server Management Software
SNMP	Simple Network Management Protocol
SOC	System on a Chip
SOL	Serial over LAN
SSD	Solid State Drive
SSH	Secure Shell
THOL	Tested Hardware and Operating System List
TPM	Trusted Platform Module
TUV	Technischer Überwachungs-Verein (A safety testing laboratory with headquarters in Germany)
UART	Universal Asynchronous Receiver Transmitter
UEFI	Unified Extensible Firmware Interface
UL	Underwriter's Laboratory
USB	Universal Serial Bus
UV	Under-Voltage
V	Volt
VA	Volt-Ampere (volts multiplied by amps)
Vac	Volts Alternating Current
Vdc	Volts Direct Current
VDE	Verband Deutscher Elektrotechniker (German Institute of Electrical Engineers)
VGA	Video Graphics Array
vRAN	Virtualized Radio Access Network
VSB	Voltage Standby
W	Watt
WEEE	Waste Electrical and Electronic Equipment
Ω	Ohm