

---

# *CPCI-V240 User Manual and Installation Guide*

*Preliminary (May 2008)*



Solflower Computer, Inc  
3337 Kifer Road  
Santa Clara, CA 95051  
(408) 733-8100, Fax (408) 733-8106  
<http://www.solflower.com>

---

---

## **Credit and Trademarks**

SgPCI-2, SgCVME are trademarks of Solflower Computer, Inc.

SunOS, Solaris are registered trademarks of SPARC International, Inc. licensed exclusively to Sun Microsystems, Inc.

Universe is a registered trademark of Tundra Semiconductor Corporation.

StarGen SG2010 is registered trademark of StarGen Corporation.

This equipment generates, uses and can radiate radio frequency energy and if not installed and used in accordance with the Instruction Manual, may cause interference in radio communications.

Operation of this equipment in a residential area is likely to cause interference in which case the user at his own expense will be required to take whatever measures may be required to correct the interference.

This document presents information for users of Solflower Computer, Inc.'s Gateway Series.

Although the information contained within this document is considered accurate and characteristic of the subject product, Solflower Computer, Inc. reserves the right to make changes to this document and any product described herein to improve reliability, function or design. Solflower Computer, Inc. does not assume any liability arising out of the application or use of any product or circuit described herein.

No part of this document may be copied or reproduced in any form or by any means without the prior written consent of Solflower Computer, Inc.

Copyright 2008, Solflower Computer, Inc.

# Overview

---

## 1.1 INTRODUCTION

---

The CPCI-V240 Adaptor board set provides PCI and PCI-VME bus functions for a wide range of I/O applications. Its features including easy connection with inexpensive CAT5e cables from Host to VME Chassis, driver supports Solaris 8, 9, 10 SPARC and X86 for transparent bridging from PCI to VME.

---

## 1.2 Features and benefits:

---

- Compatible to 3.3V and 5V, PCI clock speed 33/66MHz.
- Connection is made by inexpensive CAT5e cables which can run up to 40 ft.
- Support Solaris 8, 9, 10, SPARC, x86 IA
- Available in multiple board formats: PCI, Compact PCI, PCI Express, VME and PMC
- Easy interface to VMEbus or custom bus with great flexibility
- Fully compliant with Local PCI Bus Specification Rev. 2.2
- Data rate up to 2.5 Gbps
- Integrates with SUN PCI software system, resulting a transparent environment for PCI, and Compact PCI
- Dynamic add/remove software instance of devices
- SUN 4 VME driver compatible

---

## 1.3 APPLICABLE DOCUMENTS

---

- VMEbus Specification Manual, Revision D IEEE STD 1014-9187
- Star-VME-2010PCI VME64x on CompactPCI PICMG 2.2 R1.0
- StarGen SG2010 Hardware Reference Manual
- Tundra Universe User Manual

---

## 1.4 GENERAL DESCRIPTION

---

CPCI-V240 board set contains three boards, SOLFpvmc driver CD media and 2 x CAT5e cables, and a user manual (this manual). These are the generic names:

- SgPCI-2 (2 channels PCI Host adapter board)
- SgCPCI (compact PCI adapter board)

- VME Interface board
- Optional SgPCI-LS (for PCI expansion subsystem)
- Solflower Solaris 2.x VME nexus driver
- 15 Ft. CAT5e cable
- CPCI-V240 User Manual and Installation Guide

The CPCI-V240 Adaptor boards utilize SG2010 PCI-StarFabric Bridges. SG2010 is capable of 66 MHz, 64-bit PCI bus operation, and provides up to two 2.5 Gbps full duplex links using two pairs of standard RJ45 connectors. Links are connected using standard CAT5e cables which can run up to 40 feet.

---

## **1.5 Selecting work area**

---

The board set is electronic equipment and must be handled with care. Extremes in temperature and humidity may affect the functionality and performance of the whole system.

---

## **1.6 Unpacking instructions**

---

Perform a visual inspection of the shipping carton for any handling damages. If any shipping carton is severely damaged, open the box immediately at the carrier agent's present. Carefully remove the contents and ensure that all the pieces are present. The printed circuit boards are wrapped in electrostatic-safe bags. We recommend that you save the shipping carton and the packing material for future use, in case the product must be shipped or returned to Solflower.

### **1.6.1 Identifying the system components**

Check the packing slip to make sure that the shipment is complete.

# CPCI-V240 Adaptor Board Set Components

Based on the StarGen SG2010 PCI-StarFabric Bridge, the CPCI-V240 Adaptor board set is available in four major product categories:

- SgPCI-2 for PCI base system: Solflower-gateway-PCI-n card install into any standard PCI workstation. SgPCI-2 supports: 2, 4, 5, or 8, StarFabric ports/channels.
- SgCPCI: Solflower-gateway Compact PCI interface board which provides transparent PCI to Compact PCI bridging function.
- VME interface board: Provides bridging between Compact PCI and VME
- SgPCI-LS (Optional): Solflower-gateway PCI Leaf interface with SYSCON enabled. This card install into a system slot of a PCI expansion box (e.g PCI expansion sub-system)

---

## 2.1 SgPCI-2 Series Description

---

### 2.1.1 SgPCI-2 Board

The SgPCI-2 board utilizes StarGen SG2010 PCI-StarFabric Bridge device. This device provides PCI to PCI bridge function (PCI-PCI) and Gateway function. SgPCI-2 is capable of 66/33 MHz, 64-bit/32-bit PCI bus operation, 3.3V operation and 5V tolerant, and provides two (2) 2.5 Gbps full duplex link ports. Ports are connected using standard CAT5e cables, which can run up to 40 feet. The Gateway functions can be programmed for up to two paths for expansion to two separated sub-systems.

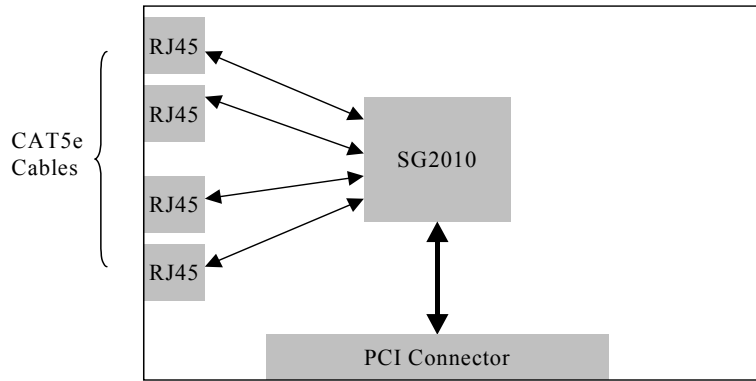
---

FIGURE 1

SgPCI-2 Board




FIGURE 2 SgPCI-2 block diagram



The 64/32bit 66/33 MHz PCI interface is compatible with PCI Local Bus Specification 2.2. Ports 0 and 1 operate at 2.5 Gbps full duplex each link. The following figure shows a SgPCI-2 board with two pairs of RJ45s at the front panel. The front panel label showing here is for two channels. One channel (link) goes to PCI expansion (P+/P-) and the other channel goes to VME (V+/V-). CAT5e cables shipped from Solflower are labeled with P+/P- or VME+/VME-. Simply connect the cables to the RJ45 connectors according to the labels. (i.e., VME+ of CAT5e cable goes to VME+ RJ45 connector, etc.)

FIGURE 3 SgPCI-2 Board



### 2.1.2 SgPCI-2 Electrical Characteristics

Power consumption:	10 Watts
Input power:	5VDC
Power consumption (max.)	5Volt@2000 mA. typical
PCI bus interface:	3.3 and 5V tolerant

---

## 2.2 SgPCI-LS Board (optional for PCI expansion)


---

The SgPCI-LS boards provides physical conversion from Fabric signals back to PCI bus signals, and allows the entire system to act like a single unit to the system management. SgPCI-LS is always installed in the system slot on a secondary PCI bus, and re-generates all required signals, including PCI clocks and Request/Grant signals. The Solflower PCI expansion system provides four additional 32/64-bit PCI slots and can be placed up to 40 feet away from the host system.

---

FIGURE 4

SgPCI-LS Board



### 2.2.1 SgPCI-LS Electrical Characteristics

Power consumption:	10 Watts
Input power:	5VDC
Power consumptionmax.	5Volt@2000 mA. typical
PCI bus interface:	3.3 and 5V tolerant

---

## 2.3 SgCPCI Comact PCI Bus Bridge

---

### 2.3.1 SgCPCI board

The Solflower SgCPCI board provides PCI-Compact PCI Bridging function. It is a 6U Compact PCI form factor board installed in the system slot of a standard Compact PCI subsystem. The SgCPCI can drive up to 4 Compact PCI slots. The on-board SG2010 PCI-StarFabric Bridge provides bridging between PCI and Compact PCI bus.

---

FIGURE 5

SgCPCI board





### 2.3.2 SgVME Electrical Characteristics

Power consumption:	15 Watts
Input power:	5VDC
Power consumptionmax.	5Volt@3000 mA. typical
PCI bus interface:	3.3 and 5V tolerant

---

## 2.4 Solflower VME interface board

---

The Solflower VME interface board provides bridging function between Compact PCI and VME. The on-board Tundra Universe II chip provides high performance 64 and 32 bit PCI to VME interface.

Solflower Computer Inc. also provides PCI-VME bridge driver: Solaris 8, 9, 10 SOLFpvme. SOLFpvme allows Sun-4 device drivers to be used with PCI-based Sun workstation and X86 IA system without changing the code in VME applications.

Some features of the Tundra Universe II VME-to-PCI bridge include:

- Fully support VME64 Rev D, board requires only 5VDC.
- Block mode data transfer up to 70MB/sec
- Support Posted Writes through 128B FIFO for quick bus release
- Built-in DMA engine for bidirectional PCI-VME and VME-to-PCI transfers
- All VME options such as Request Modes, Bus time-out is software programmable
- Up to seven window-mappings 32MB each for cover larger VME address spaces
- 32 bit or 64 bit, 33 MHz PCI bus interface; Programmable DMA controller;
- Sustained transfer rates up to 70 Mbytes per second;
- Automatic initialization for slave only application;
- Full VMEbus system controller functionality;
- Supports the latest generation of VME application.

### 2.4.1 SgVME Electrical Characteristics

Power consumption:	15 Watts
Input power:	5VDC
Power consumptionmax.	5Volt@3000 mA. typical
PCI bus interface:	3.3 and 5V tolerant

FIGURE 6

PCI-VME Bridge Interface Board



---

## 2.5 CPCI-VME Chassis

---

### 2.5.1 INTRODUCTION

The Solflower cPCI-VME 516 is a complete system, built on 19" rack-mount, in one single enclosure. All of the components are modular designed for easy integration and services.

The cPCI-VME 516 integrates two busses:

The CompactPCI bus for the host system (SgCPCI, CompactPCI peripheral boards)

- The VME bus for industrial I/O cards. These cards are connected to the CPU through a cPCI-VME Bus Converter.

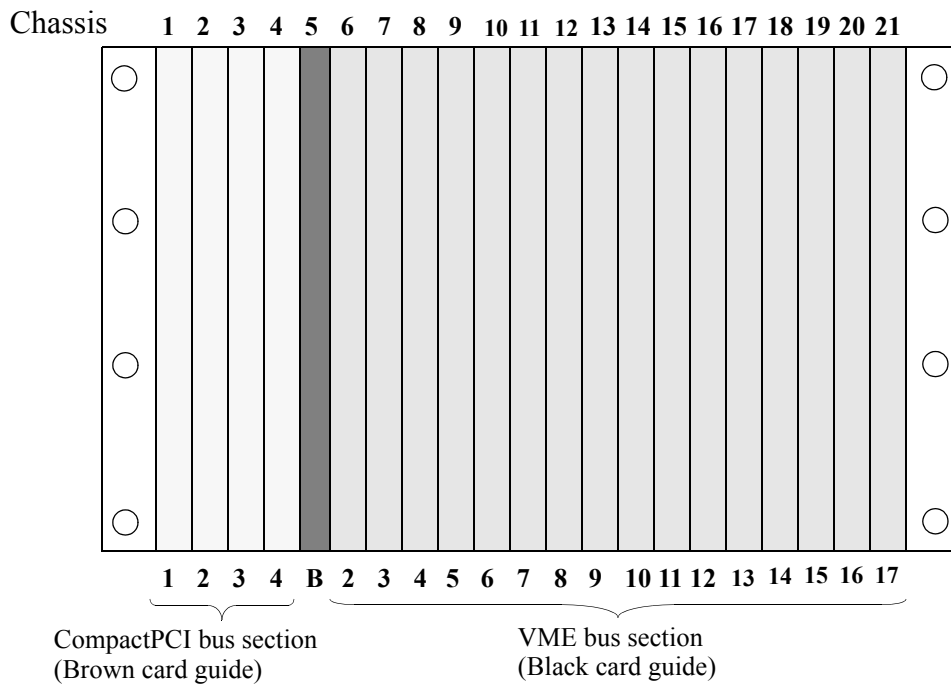
FIGURE 7

cPCI-VME 516 Chassis



FIGURE 8

cPCI-VME 516 Chassis



**TABLE 1** Slot assignment from left to right

SLOT NUMBER	BUS FUNCTION
1	CompactPCI system slot (for SgCPCI board)
2 to 4	Three CompactPCI slots (#2 to #4)
5	Solflower cPCI-VME Bridge (system controller for VME)
6 to 21	16 VME slots (#2 to #17)

**FIGURE 9** cPCI-VME 516 Chassis

### 2.5.2 GENERAL DESCRIPTION

The cPCI-VME 516 Chassis provides up to five CompactPCI slots and sixteen VME 6U slots that feature front and rear I/O connectivity. It has also built-in fan for cooling. The VME interface bus bridge, based on Tundra's Universe IIB chip, is a 6U card. It includes all translation logic between PCI bus and VME bus, provides on-board DMA engine for fast data moving across the busses.

The Solflower cPCI-VME 516 is primarily built out of aluminum. All the surfaces are painted or alodined for corrosion resistance. The system is 8U in high, 19 inches in overall width and 20 inches in depth. Flanges are provided for mounting the cPCI-VME 516 on a rack.

### 2.5.3 Slot 1- System slot for CompactPCI CPU

The system slot supports either 3U or 6U CompactPCI CPU form factor. The bus is designed for 32/64-bit and runs at 33MHz clock.

### 2.5.4 Slot 2-4 CompactPCI 64-bit /33MHz bus

The J1 and J2 connectors are standard CompactPCI bus 64-bit/33MHz. Refer to CompactPCI specification R2.1 for pin out assignment and more information. The J3, J4 and J5 are free for user.

### 2.5.5 Slot 5 - Solflower PCI-VME bridge Interface Board

The heart of the PCI-VME bridge is a Tundra Universe IIB chip. All information and programming documentation can be found at <http://www.tundra.com>.

The Universe bus bridge is a single chip in a 313 pin BGA package. All the PCI I/O signals are 5V level. The VME signals are buffered with F245 and/or F244 TTL logic in compliance with VME bus loading specification.

### 2.5.6 Slot 6-21 VME 64 bus

The total of 16 VME64 slots are available to the user. Each slot has its Bus Request, Bus Grant and IACK jumpers, except the last VME64 slot. Slot 5 is dedicated to the VME system controller. That means the first available VME slot is slot 6 for the peripheral. There is no other VME system controller allowed in the system.

### 2.5.7 System back plane

The cPCI-VME-516 system backplane is an 8-layer stripline board with outer ground plane for low noise signaling.

---

## 2.6 SgPCI-2, SgCPCI boards Miscellaneous

---

### 2.6.1 Operating Voltages (regulated from 5 Volt source)

1.5 Volt

3.3 Volt

5 Volt


### 2.6.2 RJ45 connectors and LEDs

There are two pairs of RJ45 connectors on both SgPCI-2 PCI board. Each pair of RJ45s make up one port which connects to a peer system. At the corners of a RJ45 there are two LEDs as shown in the following Figure. Each pair of RJ45s has a total of four LEDs indicating four LVDS receiving pairs of that link. These LEDs indicate the link status of that link. Table 1 describes LED's status.

---

FIGURE 10

**RJ45 and LEDs**



**TABLE 2.****LED display**

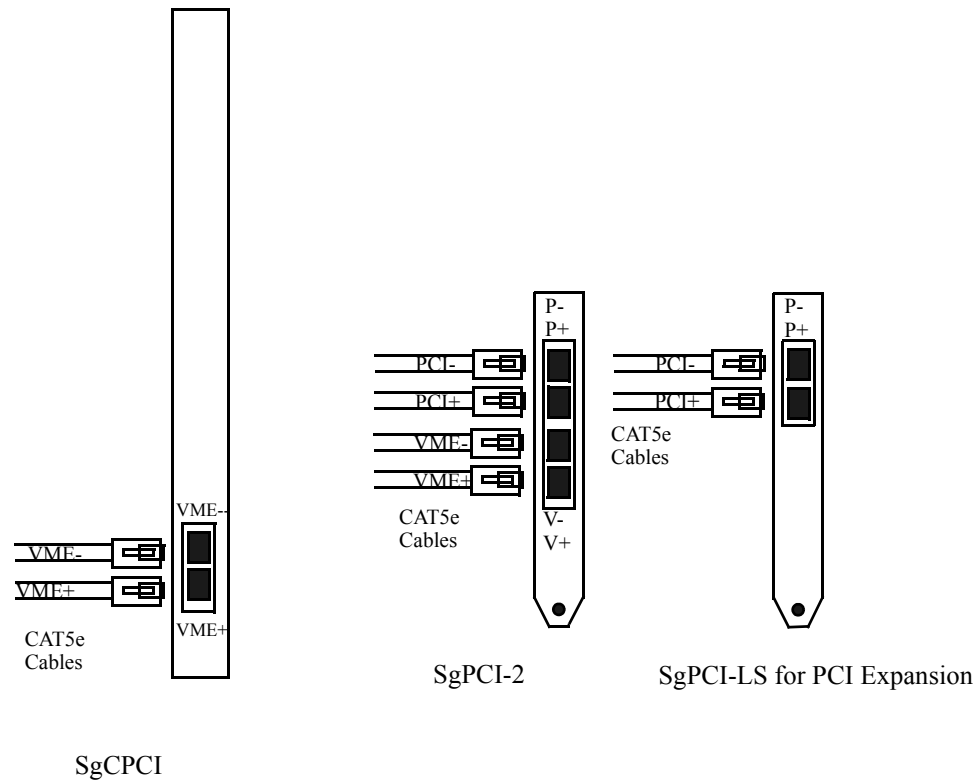
<b>LED States</b>	<b>Description</b>
LEDs are off	Link is unsynchronized and traffic is disabled
LEDs are on	Link is synchronized and traffic is enabled
LEDs are flashing	Link is synchronized but traffic is disabled

During normal operation, LEDs start blinking if CAT5e cable was disconnected or VME power sub-system power was down. This indicates the whole system must be re-boot to reinitialize. Access to hardware such as VME devices's address spaces while LEDs blinking may lead to hanging or crashing the system.

## 2.7 Cable Connection

There are four (4) RJ45 Jacks mounted on the front panel of the SgPCI-2 boards, one pair for each port. Be sure to connect the CAT5e cables according to labels. Please refer to Fig 7. After reset or recycle power to the host, LEDs should be ON indicating the system is ready. Always power on first the leaf system where SgCPCI or SgPCI-LS was installed BEFORE your host system. In the case when the whole system is connected to the same power source, it is ok to turn on the main switch which will power up the leaf system and the host at once.

FIGURE 11 RJ45 connectors at the front panels

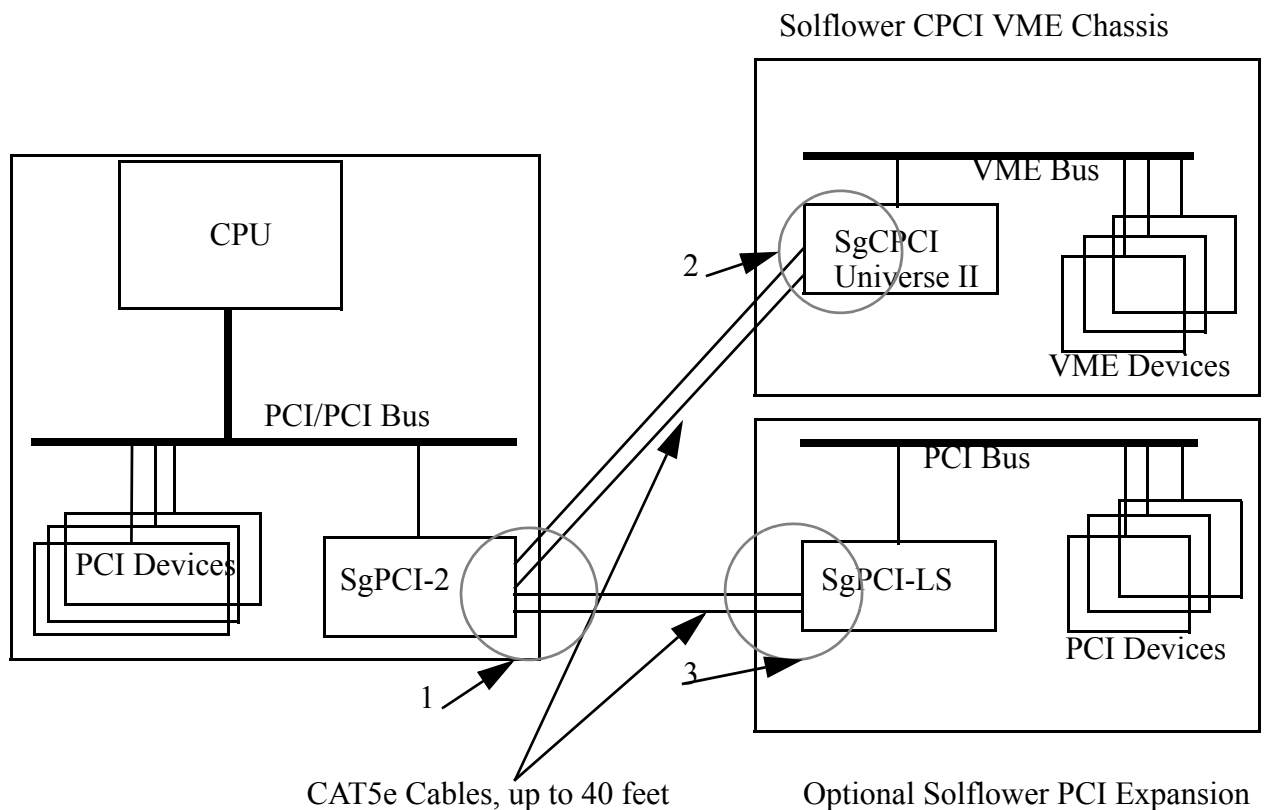


## 2.8 Sample System Connections

### 2.8.1 SgPCI-2, SgCPCI, and VME interface Application

The following figures shows a typical application of SgPCI-2, SgCPCI, and VME interface board. Optional PCI expansion is also shown here.

FIGURE 12 SgPCI-2, SgCPCIe-2, and SgVME+SgPCI-LS application



This is a typical application of the SgPCI, SgCPCI, and VME interface boards with Solflower's Compact PCI VME Chassis and optional PCI expansion chassis.

The Optional Solflower PCI Expansion Chassis is a stand-alone PCI Expansion Chassis. It has four 64-bits PCI slots and one 32 bit slots. One of the slots is dedicated for the system controller, such as a SgPCI-LS board. The other four slots are standard PCI slots.

The Solflower Compact PCI VME chassis provides expansion from Compact PCI to VME. The Solflower SgCPCI board and VME interface board translate traffics between PCI and VME buses. The SgCPCI board is seated at the system slot (first slot on your




left) of the chassis, and acts as a system controller on behalf of the Compact PCI bus. VME interface board bridges between Compact PCI bus and VME bus. Cable connections are also shown in the following three figures for the parts circled in the above figure.

---

FIGURE 13

Circle 1: CAT5e cable connections with labels at SgPCI-2 front panel



---

FIGURE 14

Circle 2: CAT5e Cable connections at SgCPCI with labels





FIGURE 15

Circle 3: CAT5e Cable connections at SgPCI-LS with labels



---

## 2.9 SunFire V240

---

SunFire V240 server is a 2U rack-optimized server designed for high availability. Its features include:

- One 1.34 GHz or two 1.5 GHz UltraSPARC processors
- Four Gigabit Ethernet ports
- Three PCI slots
- Two redundant power supplies
- Advanced Lights Out Manager (ALOM)
- For more information please refer to [www.sun.com](http://www.sun.com)

CHAPTER 3

# Installation Procedures

It is important that certain steps should be followed in order to have a trouble-free installation. Certain steps might need to be carried out in a specific order, otherwise the system might not be able to provide proper services. This chapter describes detail procedures for hardware installation as well as software installation with SunFire V240 machine.

To prevent damage to the board and system electronic components: Do not handle the board without ESD precaution. Solflower Computer recommends user to wear an anti-static strap or similar. Do not install or remove the board when your system is powered on.

## 3.1 Boards and Cables Installations

---

For a new installation of CPCI-V240 boards, (meaning there is no previous installation of the boards and SOLFpvme driver in the same system) the steps list below should be followed.

1. If not done so yet, shut down your host system with “sync; halt”, turn power off
2. Disconnect system’s AC power.
3. Open the V240 cover

---

FIGURE 16

Unlocking the latches--Left and Right



**FIGURE 17**

**Unscrew the Lever**



Undo the philip screw with a screw driver until you can lift the lever.

**FIGURE 18**

**Lift Lever**



Lift the lever with one hand and pull the cover with the other hand to open the lid.

FIGURE 19

Remove card lock



Untighten the nob and remove the card lock. Remove the PCI slot covers (not shown)

FIGURE 20


Remove the plastic air duct



Push and pull the air duct then lift it up to remove it for access to the PCI slot.

FIGURE 21


Install SgPCI-2 in PCI0



Locate PCI slot “PCI0” (the bottom PCI slot) and install SgPCI-2 board into the PCI slot. Be sure the board is firmly seated.

FIGURE 22

Install Serial board (optional)




Locate PCI slot “PCI1” (the middle PCI slot) and install Serial board into the PCI slot. Be sure the board is firmly seated.



FIGURE 23


Install Visicom Video card



Locate PCI slot “PCI2” (the top PCI slot) and install Visicom video card into the PCI slot. Be sure the board is firmly seated.

FIGURE 24

Locking the boards



Replace the card lock. Press it tight and turn the nob to lock it.

FIGURE 25

Replace Air Duct



Replace the air duct and make sure it fits.

FIGURE 26

Finishing the installation



Rotate the lid back to close position. Lock it with the latches on both side. Tighten the screw in the catch on the lid.




FIGURE 27 Checking installation



The boards installed should look like the above figure.

---


FIGURE 28 Connecting cables



Connect cables according to the picture above. (Serial board not present)

FIGURE 29


Install SgCPCI and VME interface boards



Install SgCPCI board to the system slot in VME Chassis. This is the leftmost slot. Press the board in firmly and carefully while making sure no bent pins. Install VME interface board in the 5th slot.  
If there is a Solflower PCI expansion system, install SgPCI-LS to the system slot at Solflower expansion box.

FIGURE 30

Connect VME Chassis to V240



Connect Solflower Compact PCI VME Chassis and Optional Solflower PCI expansion to your host system with CAT5e cables. For cable connections, please refer to section 2.7 and 2.8.

Power on Solflower Compact PCI VME Chassis (and PCI expansion chassis) first, then turn on your host Computer power. If host Computer and Chassis are connected to the same power source, the main switch will power on the host and the chassis at the same time. This is also permitted. At this point you should see the link LEDs are on and solid at the front panels of both SgCPCI board and SgPCI-2 board.

Boot up your host system to Solaris environment, and insert PVME Driver CD. Please make sure the CD is for the right machine and the right Solaris version.

At Solaris prompt, cd to your cdrom. e.g., /cdrom/cdrom0 directory. You should see the packages such as SOLFpvme and SOLFlib listed in the directory. Execute “pkgadd -d .” Choose the packages you want to install. Hit enter to install all packages.

After SOLFpvme and SOLFlib packages are installed successfully, type q to quit, reboot the system with “reboot” right away. Please do not install any leaf VME device driver at this point.

After the host system boots up, you can install your leaf VME device drivers. Edit your leaf VME device driver.conf file with “parent = pvme”.

---

### 3.2 SunFire V240 Rack Installation (Example)

---

FIGURE 31

Install Rail (Left)



- Position a mounting bracket against the AC3K chassis so that the holes are aligned with the screw holes on the side of the AC3K chassis.
- Attache the mounting bracket with the supplied 10-32 x 1/2in screws
- Repeat step 1 and 2 to install the mounting bracket back sides.


FIGURE 32

Install Rail (Right)



FIGURE 33

Slide in V240



**Caution: To avoid personal injury and damage to the server, this procedure requires minimum of two people because of the weight of the server.**

- Raise the server so that the inner glides are aligned with the slide rails.
- Insert the inner glides into the slide rails, then push the server into the rack until the server stops (approx. 12 in)
- Pull and hold the inner glide release green buttons on both side to slide the server in all the way until the green handles lock on the front of the mounting brackets (Fig 34-35)

FIGURE 34

Hold the green Clips and slide in V240



FIGURE 35

Completing installation



**Verify the operation of the slide rails:**

**WARNING: Removing the server completely out of the rack requires two person. Attempting this step alone could result in equipment damage or personal injury.**

- Press and hold the green handlers and slowly pull the server out of the rack until the server stops (approx. 18 in)

Note: If you want to remove the server completely out of the rack, then pull and hold the inner glide green release buttons (Fig. 34-35) while you continue pull the server.

- Pull and hold the inner glide green release buttons (Fig. 34-35) while you push the server back into the rack.

FIGURE 36

Open Lid for Future Service



Lift server lid for future service when needed.

# Software installation for SgVME

---

## 4.1 Introduction

---

The pvme device driver is a 32- and 64-bit Solaris 8, 9, 10 nexus driver for the Solflower PCI-to-VME bridge.

The pvme device driver permits the installation, configuration, and use of VME devices on Sun Microsystems UltraSPARC platforms that support the PCI bus, as well as on selected Intel or AMD platforms that run the Solaris 8, 9, and 10 operating system.

Most VME device drivers need only be written to conform to the Device Driver Interface (DDI) definition from Sun, and they should work transparently, with no need to be aware that they are communicating to Solaris through a bus bridge.

---

## 4.2 Features of the pvme Device Driver

---

- Complete Solaris Nexus driver, providing leaf driver management and translation for VME device mappings, interrupts, and DMA.
- Direct VMEbus access, though the export of VME address space device files.
- A special programmable DMA engine that transparently provides high-performance DMA transfers to and from VME address spaces when the VME device files are accessed by read(2) and write(2) operations.
- VMEbus power-off detection and management for user processes that access the pvme device files directly.
- Selectable operation modes for VME block transfer, arbitration, and posted writes.
- Loadable and unloadable Nexus driver.



### 4.3 Components

---

The pvme release from Solflower is comprised of binary executable files, libraries, and documentation. The table below provides a list of the major components.

Component Name	Component Description
/usr/kernel/drv/pvme	The loadable pvme device driver binary file.
/usr/kernel/drv/sparcv9/ pvme	(UltraSPARC specific) 64-bit version of pvme device driver.
/usr/kernel/drv/amd64/pvme	(AMD Specific) 64-bit version of pvme device driver.
/usr/kernel/drv/pvme.conf	The pvme driver configuration file. May be used to set a variety of control options governing the pvme nexus driver. Among them are options to assign at which SPARC processor interrupt level the pvme device interrupts, enables for use of posted writes, etc.
/opt/SOLFpvme/etc/ pvme.tab	Description for how to create the symbolic links in the /dev directory for the pvme device files. When you run pkgadd(1M) to install the pvme device driver, the installation scripts automatically appends these lines to the system /etc/devlink.tab file.
/opt/SOLFpvme/vmemem	A directory that contains a simple VME pseudo driver. This driver provides access to a VME memory card. This is not necessarily the way you would want to write such a driver, but it illustrates how to access and configure a VME device under the pvme nexus driver.
/opt/SOLFpvme/drv_test/ pvme/pvme_test	A directory that contains a test program and source code that shows you an example of how to open the pvme device files and perform various operations, such as mmap a chunk of VME memory or dump the contents of the pvme control registers. Included is a simple VME memory verification test.
/opt/SOLFpvme/drv_test/ vmemem	A directory that contains a simple test program and source code that shows you an example of how to interact with the ioctl(2) interface exported by the vmemem device driver.
/opt/SOLFpvme/doc	A directory that contains documentation on the pvme and vmemem drivers, as well as information about test programs.
/opt/SOLFlib/drv_lib	A directory that contains source code and a binary library for building the test programs for the pvme and vmemem drivers.
/opt/SOLFlib/drv_util	A directory that contains include files and a binary library for building the vmemem device driver. Certain of the test programs may use some of the include files, also.
/opt/SOLFpvme/bin	The home of the a_pvme and r_pvme scripts, which are useful for installing and removing the pvme driver on a running system. Also contains the test programs pvme_test.



---

## 4.4 Installation

---

The pvme device driver is delivered on cdrom as a Solaris 2.x package.

The pkgadd(1m) operation should copy and install the pvme driver for you. If a situation occurs where you need to circumvent the installation performed by pkgadd(1m) and perform the installation by hand, the following summarizes the operations performed by the package installation wrapper.

1. Because the pvme driver is not essential for boot-up of the Solaris operating system, by convention, the driver executable, pvme and the driver configuration file, pvme.conf, are copied to the /usr/kernel/drv directory. These components could also have been placed under the /kernel/drv or /platform/sun4u/kernel/drv paths.

2. Entries specifying the device files to create for pvme are appended to the system /etc/devlink.tab file. Before inserting the requisite lines for pvme, we make a backup of /etc/devlink.tab, called /etc/devlink.tab.BAK, so the original file may be recovered if problems occur.

The SOLFpvme package installation scripts append the lines for pvme automatically. We provide a file with the SOLFpvme package called pvme.tab that contains the lines to be inserted. This file will be resident in the SOLFpvme package directory, once installed. Typically, the pvme.tab file may be found at /opt/SOLFpvme/etc/pvme.tab.

The following operations mimic the installation operations with respect to the specification of device links:

```
# cp /etc/devlink.tab /etc/devlink.tab.BAK
```

```
# cat /etc/devlink.tab.BAK /opt/SOLFpvme/etc/pvme.tab > /etc/devlink.tab
```

3. Once the pvme driver files are placed under /usr/kernel/drv, and the device link specifications have been added to /etc/devlink.tab, the pvme driver is installed and loaded via the Solaris add\_drv(1M) command. This is the syntax used by the package installations scripts to install and load the driver:

```
# add_drv -c vme -i "pci10e3,0" pvme
```

The system /etc/driver\_aliases file is updated by this operation as well.

Assuming that the operations above complete successfully, the pvme driver is now loaded. Henceforth, if the Solaris system is shut down and rebooted, the pvme device driver will be automatically loaded.

To undo this operation requires that the system administrator deinstall the driver, using the rem\_drv(1M) command.

4. When the pvme driver is loaded by the Solaris operating system, it prints a banner on the system console. This banner contains information about the pvme product, including version, build date, copyright, and Solflower-specific driver library components employed for this release of the product.

Below is an example of such a banner printed as display on the console. Note that this same information is also sent to the system log file (i.e., syslog).

```
PVME - PCI-to-VME bridge nexus driver, version 8.9(1)
Copyright (c) 1997-2002 by Solflower Computer, Inc.(2)
Built: Wed Aug 28 10:43:23 PDT 2002 (3)
Solflower Driver Kit, Module: DRV-util, Version 1.0(4)
Copyright (c) 2002 by Solflower Computer, Inc.(5)
Built: Tue Aug 27 13:53:10 PDT 2002 (6)
```

Notes:

- (1)PVME driver version information
- (2)Copyright notice
- (3)Build date
- (4-6)Solflower Driver Kit module:  
version, copyright, and build date

**Note: If SOLFpvme driver was installed in a different PCI-VME subsystem than Solflower SgVME-PCI, the installation will stop with an error message " not a Solflower subsystem".**

---

## 4.5 Theory of Operation

---

### 4.5.1 T-of-O: Introduction

The primary tasks for which the pvme device driver is responsible are two-fold.

First, pvme provides memory address translation between the VME bus and PCI bus memory spaces. Second, it manages VME interrupts and vectors. Since the pvme device is a PCI device, it maps associated VME actions and operations into analogous ones for the PCI bus, where possible.

#### **4.5.2 VME-to-PCI address translation**

When the pvme driver is installed, it examines the PCI bus device tree to determine which addresses are already occupied and thus how much memory is available for accessing VMEbus devices. If other PCI devices are installed in the same PCI space, the amount of PCI space available for VMEbus access is reduced.

However, even when the free PCI memory space is heavily fragmented, the pvme driver can, in many cases, work around holes in the PCI memory space caused by other devices. The pvme driver accomplishes this by using slave image windows.

Slave image windows are a PCI-to-VME address translation mechanism provided by the pvme hardware; they are in many respects analogous to memory management mapping hardware. A slave image window operates much like an MMU page table entry: it translates memory accesses on a PCI memory address range into corresponding memory accesses on a VME memory address range.

The pvme hardware provides up to 8 slave image windows for mapping address ranges on the VMEbus. Once initialized, a slave image window may provide PCI-to-VME address translation for any VME A32, A24, or A16 space.

By using the pvme slave image windows, the driver can hide the effects of fragmentation of the PCI address space.

For example, PCI frame buffers may occupy multiple megabytes of space for their bit-planes. The pvme device and driver CANNOT use the PCI space occupied by these bit-planes for mapping or accessing the VMEbus. However, using the Solaris device tree, the pvme driver is able to detect the presence of such a device, and can use the PCI memory ranges that are lower or higher than the occupied regions of memory.

The default sizes assigned to each slave image window are predefined as a function of the type of VME address space being mapped. (Note that the default size for each VME address space slave image window may be overridden by adding directives to the pvme.conf configuration file.)

The availability of multiple slave image windows provides flexibility. If a mapping operation requires more than 32 MBytes of contiguous A32 memory space, more than one pvme slave image window will be used. Mapping requests made to the pvme driver will silently detect whether a new slave image window is required, and if so, will allocate and initialize the window so that the desired VMEbus addresses can be mapped. Unless numerous "sparse" mappings are made, the user should not be aware that mapping windows are being used.

The pvme device driver assigns the PCI memory physical addresses for accessing VMEbus address space by using hardware base and bound registers of the slave image windows. These base and bound registers permit us to define space on the PCI bus for our use that did not show up when our device was initialized by the Sun Open Boot PROM.

The pvme device driver must allocate the PCI memory space to use for these slave image windows from that memory space not used by other PCI devices on the same bus as the pvme hardware. So, in a sense, the pvme is competing for space against its sibling

PCI devices. Consequently, if your system needs for VME address space are large, it is best to put the pvme hardware on the least occupied PCI bus available. If there is only one PCI bus on the system, of course, this is not a consideration.

### 4.5.3 VME interrupt management

The VME bus supports 7 interrupt priority levels. The PCI nexus can support only 1 interrupt. Thus, all VME interrupts are in a sense "funneled" into a single PCI interrupt. When an VME interrupt occurs, the pvme driver scans for pending VME interrupts from highest to lowest priorities. Thus, in theory, higher priority VME interrupts are therefore given preference over lower priority VME interrupts.

However, short term priority inversion of interrupt handlers is possible with this scheme. There is a timing window where a lower priority VME interrupt may arrive and activate its interrupt handler, and thereby block the servicing of a higher priority VME interrupt until the low priority interrupt handler completes. When the interrupt handler completes, the higher priority interrupt will occur and be serviced. The duration of potential interrupt priority inversion is therefore limited by the duration of the other VME interrupt handlers.

VME devices use an interrupt vector scheme to determine the source of any VME interrupt. The VME vector is a unique number between 1 and 255. Each device, when an interrupt acknowledgement cycle is run by the pvme hardware, emits its unique identifying vector on the VME bus. The pvme hardware receives this vector, and the pvme driver uses it to find the appropriate handler for that VME vector.

This implies that there can be only one handler for each unique VME vector.

### 4.5.4 Bus Error Monitoring

If there is no "Target Abort" or "Bus Error" signal terminate a non-existing-memory cycle, there are two options to generate Bus Error from pvme driver level:

The first option is to call one of pvme driver's ioctl after each memory access cycle. Here is an example of how to implement it:

```
/*
 * vaddr is the virtual address of VME memory.
 * out is the virtual address of a user buffer.
 */
static int
dev_peek32(int fd, uint32_t * vaddr, uint32_t * out)
{
    int tabort = 0;

    *out = *vaddr;
    if (dev_ioctl(fd, PVME_TARGET_ABORT, &tabort) == 0) {
        return (tabort ? -1 : 0);
    }
    return (0);
}
```

For more detail, please take a look at the files pvme\_dev.c and pvme\_mem.c in the directory SOLFpvme/drv\_test/pvme/pvme\_test. The pvme\_test program also

implements this as the "-b" option. For example, "pvme\_test -b -i -s 100 -o 30000000" reads 100 bytes at address 0x30000000 and checks if bus error occurs. "pvme\_test -b -W -s 100 -o 55000000" writes 100 bytes at address 0x55000000 and checks if bus error occurs.

The second option is to call pvme\_test program before any memory access: "pvme\_test -t <ms>". This needs to be called only once to start a daemon monitoring bus error. "ms" is how many milliseconds the daemon should wait before checking status of the bus again, and it must be between 100ms to 1000ms. The default wait time is 500ms. Use "pvme\_test -t 0" to stop the daemon.

Note: Reading from a memory location to determine if it is valid is safer than writing to it since writing to an invalid address can panic the system. Writing also has posted-write option to be considered. To detect bus error correctly while writing, posted-write must be disable. Posted-write can be disable or enable by an entry in the file pvme.conf: post-writes="yes|no";

For more information about the pvme.conf file, go to the next section "pvme.conf sample file".

---

## 4.6 pvme.conf File Examples

---

/usr/kernel/drv/pvme.conf is the pvme driver configuration file. It may be used to set a variety of control options governing the pvme nexus driver. Among them are options to assign at which processor interrupt level the pvme device interrupts, enables for use of posted writes, setting the mapping memory space for VME devices, etc. Here is a sample pvme.conf file for SUN Fire V240, Ultra 40, Ultra 20, and Ultra 45.

```
# "@(#)pvme.conf 1.28 05/04/28 Solflower Computer, Inc."

#

# Copyright © 1997-2003 Solflower Computer, Inc.

#

# The property "interrupt-priorities" controls which

# SPARC processor interrupt level is employed by the

# PVME device for interrupts.

# Permissible values are:

# interrupt-priorities=3;

# interrupt-priorities=4;

# interrupt-priorities=6;

# interrupt-priorities=8;

# interrupt-priorities=9;
```

```
#
# At this time, only one interrupt-priorities value may be specified.
#
interrupt-priorities=4;
#
# Parameter Ranges
# =====
#
# max-retry: 1 to 0xf
# post-writes="yes" | "no";# access to VME use posted writes
# posted-writes: 0 to 0xf# size of VME posted writes, if enabled
# request-level: 0 to 3
# request-mode: "demand" | "fair"
# release-mode: "release-on-request" | "release-when-done"
# vme-reset="on" | "off";
# supervisor="on" | "off";# enable | disable supervisor AM generation
### the default is "off"
# bus-timeout= 1 to 7;
# arb-timeout= 1 to 2;
# dtack-enable="rescind";
# arb-mode="priority" | "round-robin";
#
# ctl-mode="syscon";
#
# pci-abs= 0 or 1;# PCI aligned burst size: 0=32 bytes 1=64 bytes
# dvma-size=0x100000 to 0x800000;
#####
```

```
# The following control the use of VME block mode for VME accesses
# from the UltraSPARC.

# The dma-blt option, when turned "on" causes read(2) or write(2)
# system calls on /dev/pvme/vmeXXdXX to use block-mode transfers.
# Similarly, the mem-blt option, when turned "on" causes
# user accesses to mmap(2) VME memory to perform block-mode
# transfers.

# Synopsis:
# dma-blt="on" | "off";
# mem-blt="on" | "off";
#
# If there is more than 1 PVME board in the same VME backplane,
# the ctl-mode="syscon" must be enabled for ONLY 1 PVME board.
# You can't have more than 1 VME system controller.

debug=0;
ctl-mode="syscon";# comment out if another board is the system controller
max-retry=8;
request-level=3;
request-mode="demand";
# release-mode="release-on-request";
release-mode="release-when-done";# this is for ADI
vme-reset="off";# don't reset devices on VME bus when installed.

bus-timeout=7;
arb-timeout=2;
dtack-enable="rescind";
arb-mode="priority";
pci-abs=1;
```

```
reg-delay=10;
latency-timer=0x10;
power-fail="on";
# dvma-size=0x100000;
#
# Memory access controls
#
post-writes="yes";
posted-writec=4;
dma-blt="off";
mem-blt="off";
supervisor="off";# don't generate supervisor AM
vme-posted-write="yes";
#
# PCI CSR parity control
#
pvme-parity="off";
#
# PCI LMISC CWT
#
pci-cwt=5;
#
# VME RAI controls
# Ordinarily, the VME RAI controls, which permit the PVME registers
# to be visible to others in VME space are not used.
#
# Setups that need to have the PVME registers accesible on
```



```
# the VME bus from another processor board should set these
# as needed.
#
# Note that when the RAI is in use, ALL the rai- options below
# must be set to one of the legal values listed. If one is omitted,
# the results will vary from unpredictable to non-functional.
#
# rai-enable="on" | "off";# Access controls of PVME registers
# rai-pgm="both" | "prog" | "data";# AM controls for PVME register access
# rai-super="both" | "priv" | "user";# AM controls for PVME register access
# rai-vas="a32" | "a24" | "a16";# VME Address space for PVME register
## access
#
# If you want to enable the PVME RAI, here are some reasonable
# defaults. This example enables the registers as being available
# as both program and data and to both supervisor and user access,
# and places the registers in A32 space.
#
# rai-enable="on";# PVME registers accessible on VME
# rai-pgm="both";# both program and data
# rai-super="both";# both supervisor and user
# rai-vas="a32";# A32 space
#
# VME RAI Base Address Register initialization
#
# Assign address where PVME registers live in VME space,
# when PVME RAI is enabled.
```

```
#  
#Note this depends on the addresses of other devices on VME.  
#In addition, the A32 and A24 DVMA windows live at  
#location 0x0..0x100000 in each address space.  
#  
#Address collisions should be avoided, so select  
#your rai-bar address to avoid address ranges in use.  
#Note also that only the highest 20 bits of rai-bar  
#address are significant. Thus, your rai-bar addresses  
#should be of the form: 0xXXXXX000 where X is in [ 0..f ]  
#  
# rai-bar=0x1000000;# PVME registers live at VME 0x1000000  
# virtual-dma=0xd0000000,0x10000000;# PCI DMA <base>, <size>;  
#  
#virtual-dma=<base>,<size>;  
#virtual-dma=0xc0000000,0x20000000;# Newer Sun machines  
#virtual-dma=0xfe000000,0x007ffff;# Ultra-60  
#  
#  
# The following lines is needed for the V240  
#  
virtual-dma=0xc0000000,0x20000000;# Newer Sun machines  
#  
# VME slave map controls: enable/disable response to VME DMA addresses  
#  
# Synopsis:  
# vme-slave32="on|"off" ;
```

```
# vme-slave24="on|"off" ;
# vme-slave16="on|"off" ;
#
vme-slave32="off";# LTX specific
vme-slave24="off";# LTX specific
vme-slave16="off";# LTX specific
#
# The following is an override for the ddi_intr_hilevel(9F)
# function of the PVME nexus driver. If you set hi-level=1,
# it will cause ddi_intr_hilevel to return TRUE for children
# of PVME.
#
#
# hi-level=0 | 1;
#
hi-level=0;# do not return hi-level for VME children
#
# A32/A24/A16 mapping window length
#
# If not specified, the length of each mapping window
# defaults to the below values.
#
a32-length=0x02000000;#
a24-length=0x01000000;#
a16-length=0x00010000;#
#
# The SunBlade 100 has extremely restrictive rules about available
```

```
# space for subsidiary PCI mappings, so it is necessary to use these
# directives just to be able to map a VME device.
#
# Suggested SunBlade 100 Settings:
#
#a32-length=0x00020000;# minimum: 0x00010000 (64 KBytes)
#a24-length=0x00010000;# minimum: 0x00010000 (64 KBytes)
#a16-length=0x00010000;# minimum: 0x00010000 (64 KBytes)

##### PLEASE READ THE NOTE FIRST #####

##### NOTE: EEPROM must be programmed by Solflower Computer in order to
##### use the following setups. In SUN Fire V240 case, the following setups may
##### not be needed.

# the following 5 lines are needed for Ultra 40 slots PCIe3 and PCIe2, uncomment them
#when use in Ultra 40

#bus-base=0x82000000;

#bus-size=0xe000000;

#base=0x8200;

#limit=0x9000;

#pci-exclude=0x20000000,0x62000000,0xc0009000,0x17000,0xc0030000,0x3fad0000;

# the following 5 lines are needed for Ultra 20, uncomment them when use in Ultra 20

#bus-base=0xb3000000;

#bus-size=0xd000000;

#base=0xb300;

#limit=0xc000;

#pci-exclude=0x20000000,0x93000000,0xf0009000,0xfaf7000;
```

#The following 4 lines are Ultra-45 for PCI-E 0, PCI-E 1, & PCI-E 2, uncomment them  
#when use in Ultra45

#bus-base=0x02000000;

#bus-size=0x0e000000;

#base=0x0200;

#limit=0x1000;

---

## 4.7 PVME Device Files

---

The pvme nexus driver creates 13 device files in /dev to permit access to the VMEbus and the pvme control registers.

In general, the user should not need to access the pvme control registers.

The VMEbus device files are potentially useful when you wish to mmap(2) a chunk of VME address space and access it directly with a user program.

### 4.7.1 VME Device Files

There are 12 device files exported by the pvme driver that represent the various VME address spaces, plus associated data widths:

Device File -----Description

/dev/pvme/vme16d64 ----Used to access A16 space with D64 width.

/dev/pvme/vme24d64 ----Used to access A24 space with D64 width.

/dev/pvme/vme32d64 ----Used to access A32 space with D64 width.

/dev/pvme/vme16d32 ----Used to access A16 space with D32 width.

/dev/pvme/vme24d32 ----Used to access A24 space with D32 width.

/dev/pvme/vme32d32 ----Used to access A32 space with D32 width.

/dev/pvme/vme16d16 ----Used to access A16 space with D16 width.

/dev/pvme/vme24d16 ----Used to access A24 space with D16 width.

/dev/pvme/vme32d16 ----Used to access A32 space with D16 width.

/dev/pvme/vme16d8 ----Used to access A16 space with D8 width.

/dev/pvme/vme24d8 ----Used to access A24 space with D8 width.

/dev/pvme/vme32d8 ----Used to access A32 space with D8 width.

#### **4.7.2 PVME Control Registers**

The pvme nexus driver exports another device file for direct access to the nexus's control registers:

Device -----Description

/dev/pvme/nexctl ----Used to access the pvme control registers.

This file is present for diagnostic purposes, such as to examine the current control register settings of the pvme. It is not ordinarily something that the user has a need to use. Since it permits direct access to the pvme device registers, it is possible to change pvme control register settings with this mechanism. The user is advised to exercise great caution if he desires to try this, since an incorrect value inserted into one of the pvme control registers has the potential to put the pvme into a mode that could result in a system panic.

#### **4.7.3 Device File Example**

A typical program would access one of these VME spaces from a user program in the following manner:

A typical program would access one of these VME spaces from a user program in the following manner:

```
char * vme_map(char * vme_device, int vme_off, int bytes)
{
    intfd;

    fd = open(vme_device, O_RDWR);

    printf("[Mapping 0x%x bytes at VME offset 0x%x]\n", bytes, vme_off);

    return (mmap((caddr_t)0, bytes, PROT_RW, MAP_SHARED, fd, vme_off));
}
```

```
#define VME_OFFSET 0x00100000

#define VME_BYTES 0x00010000

int main()

{

char *vme_mem;

char vme_byte;

/*

* map VME_BYTES of VME memory at location

* VME_OFFSET in VME A32 space.

*/

vme_mem = vme_map("/dev/pvme/vme32d32", VME_OFFSET, VME_BYTES);

/*

* Now read the first byte at the

* specified VME offset.

*/

vme_byte = vme_mem[0];

}
```

## 4.8 PVME Test Program

---

The tests included in the `drv_test/pvme` directory include

`pvme_test`

The `pvme_test` program allows users verify the state of VME mapping and interrupt activities by examining the register state of the Universe chip. A few examples are given below.

`pvme_pages`

This is a simple VME memory mapping and access test.

### 4.8.1 A `pvme_test` Example

To print out all available options of the `pvme_test` program, use the `-h` option to `pvme_test`:

```
# pvme_test -h
```

```
[PVME_TEST - Version: 2.0 Date: Sun Sep 9 13:05:55 PDT 2001 ]
```

```
-A - tests don't ask the user for input
```

```
-a - set VME address space for memory and DMA tests
```

```
-B - print pvme_test build date and version
```

```
-b - VME memory accesses check for bus errors
```

```
-C - run a memory comparison test
```

```
-c - reset the VME bus
```

```
-D - disconnect busops
```

```
-d - run the DMA test
```

```
-g - generate a VME software interrupt
```

```
-h - get this message
```

```
-H - get ownership of the VME bus
```



- I - clear pending interrupts
- i - read VME memory
- L - run the memory test in store/load mode
- l - Parent memory limit check
- m - run the memory test
- M - run VME mapping tests
- n - toggled interrupt recording switch
- N - print interrupt record
- O - Check the power on state of the VME bus
- o - set the starting address offset for memory and DMA tests
- P - VME switched on: toggle state
- p - print register state
- R - reconnect busops
- r - PCI resource map limits
- s - set the size to use for DMA and memory tests
- S - set the stride length for memory tests
- T - memory test mode: increment or walking ones
- t - monitor target abort with the specified time interval
- U - release ownership of the VME bus
- V - device registers
- v - set the initial test pattern for memory tests
- W - write VME memory
- w - CPU uses 64, 32, 16, or 8 bit memory operations (DEFAULT 32)
- X - select the VME HW-enforced Data Width to use for memory & DMA tests

#### 4.8.2 A pvme\_test example: Examining Universe Registers

```
# /opt/SOLFpvme/bin/pvme_test -p

[opening /dev/pvme/nexctl]

VID 10e3 DID 0 COMM 107 <IO><MAE><ME><SERR> STATUS 200 <DSTMED>

CLASS revid 2 prog 0 sub 80 base 6

Cache line size 0 Latency Timer 0 Header Type 0 BIST 0

BAR[0] ffafd000 BAR[1] 0000d001 BAR[2] 00000000

BAR[3] 00000000 BAR[4] 00000000 BAR[5] 00000000

CIS Pointer 0 SUBSYSTEM: VID 0 ID 0 CAP 0 ROM 0

MAX LAT 0 MIN GNT 3 ILINE b IPIN 1

PCI SI[0] CTL c0400000 BAR ffa00000 BD ffa10000 TO 00600000

PCI SI[1] CTL c0820000 BAR 20000000 BD 22000000 TO 78000000

PCI SI[2] CTL 00000000 BAR 00000000 BD 00000000 TO 00000000

PCI SI[3] CTL 00000000 BAR 00000000 BD 00000000 TO 00000000

-----

PCI SI[4] CTL 00000000 BAR 00000000 BD 00000000 TO 00000000

PCI SI[5] CTL 00000000 BAR 00000000 BD 00000000 TO 00000000

PCI SI[6] CTL 00000000 BAR 00000000 BD 00000000 TO 00000000

PCI SI[7] CTL 00000000 BAR 00000000 BD 00000000 TO 00000000

-----

VME SI[0] CTL 80f20000 BAR 00000000 BD 00fff000 TO 00000000

VME SI[1] CTL 80f10000 BAR 00000000 BD 00fff000 TO 00000000

VME SI[2] CTL 00000000 BAR 00000000 BD 00000000 TO 00000000

VME SI[3] CTL 00f00000 BAR 00000000 BD 00000000 TO 00000000

-----

SPEC CYCL CTL 00000000 ADDR 00000000 EN 00000000 CMP 00000000 SWP 00000000

PCI LMISC 05 SSI 00000000 ERR_LOG CMD 70000000 ADDR 00000000
```

PCI INTR ENABLE 000081fe STATUS 00004000 MAP 0..1 00000000..00000000

-----

VME INTR ENABLE 00000000 STATUS 00000000 MAP 0..1 76543210..00000001

VME INTR STATUS-ID fe000000

IRQ1 00000000 IRQ2 00000000 IRQ3 00000000 IRQ4 00000000

IRQ5 00000000 IRQ6 00000000 IRQ7 00000000

-----

DMA DCTL 00000000 DTBC 00000000 DLA 00000000 DVA 00000000

DMA CPP 00000000 DGCS 00000000 DLLUE 00000000

-----

VME CTL MAST 84e01000 MISC 76060000 MSTAT 00260000 USER\_AM 40400000

-----

VME RAI CTL 00000000 BAR 00000000

VME CSR CTL 00000000 CSR\_TO 00000000 AM\_ER 24000000 ADDR\_ER 992fff84

VME CSR CLR 40000000 SET 40000000 BAR 00000000

-----

---

## 4.9 Known Limitations

---

If a VME master transfers data to PCI memory at the same time as the CPU attempts to access VME space, a PCI timeout may occur that causes the system to panic with the message "too many retries".

This problem arises because Solaris initializes the PCI bus arbiter to use an insufficient number of retries when contention occurs. The result is that if a VME master holds the bus for a long time when the CPU is trying to access the PCI space, the PCI bus arbiter may give up too early.

Before you attempt to access the VMEbus with the CPU, you should make sure that no VMEbus device is performing a transfer of data to PCI memory. Typically, a way to trigger this problem is to use `mmap(2)` to map a chunk of VME memory into a program's address space and then to access the mapped VME memory while a VME device is performing a DMA transfer. Such simultaneous attempts to access the VMEbus should be avoided by the user as much as possible.

# Trouble shooting

In trouble shooting section, we assume SgPCI-2, SgCPCI and SgPCI-LS were properly installed and seated securely.

*Link LEDs are off*

Make sure the cables are connected correctly according to the label on the CAT5e cables and board front panels. Check if the +/- CAT5e cables are swapped. Turn on the Solflower 21 Slots Chassis (and PCI expansion system) first, then power up the Host. Make sure power supplies of the Solflower 21 Slot VME Chassis is good. If the LEDs are still OFF, replace the parts one at a time: CAT5e cable, SgPCI-LS (if there is any), SgVME, then SgPCI-2.

*Link LEDs are blinking, Links were not established.*

Be sure to turn on the Solflower 21 Slots VME Chassis (and PCI expansion system) before the Host. The Host is responsible for enumerating the SgPVME and SgPCI-LS (if any) and all PCI buses behind the SgPCI-2 card. If the Host is ON before the expansion box, the enumeration will complete without seeing that expansion PCI path. Issue reset-all command at OBP to force new Host enumeration cycle.

*At SUN machine's OBP level, one or more leaves are not seen.*

Try “reset-all” at the OK prompt. Due to power up delay or some other reasons, the leaves were in some unknown state when the OBP is probing the device tree. Thus OBP will not recognize them. “reset-all” will reset the devices tree, and restart a new probing.

*Cable connections are correct, power sequence was correct, tried “reset-all”, but the link LEDs are still blinking.*

Check the board model # used in the system. SgPCI-2 (P/N 621-001-2CN) can be connected with SgPCI-LS (P/N 519-001-2CN) and SgVME (P/N621-002-SVME), but not SgPCI-2(P/N 621-001-2CN) . These boards physically are almost identical but configuration is set differently from factory as default, and should not be altered on the field.

*My application program cannot communicate with my leaf VME devices*

a. Solflower SgVME driver SOLFpvme is written strictly for used with Solflower SgVME-PCI Adaptor board set. SOLFpvme will not attach to if the hardware configuration is different than Solflower suggested setting.

b. Make sure your device driver.conf file is set to “parent = pvme”

- c. Check and see if the PCI slot the installed SgPCI-2 is a suitable slot. Please refer to Chapter 4.
- d. if a, b,c wasn't the case, try to reinstall SOLFpvme driver. Please refer to section 3.2
- e. Make sure the configuration in /usr/kernel/drv/pvme.conf file is correct.

### *System get stuck at boot up*

In some rare instances at boot up, after SOLFpvme driver is loaded by the system at OBP, but it's not ready to attach the leaf VME device drivers on the VME bus, the system will get stuck at that point trying to allocate resource for leaf VME device(s). In this case we need to do the following:

Disconnect CAT5e cables from SgPCI-2 board; reboot the system; after the system boot up, remove all your leaf VME device drivers; remove SOLFpvme driver by “/opt/SOLFpvme/bin/r\_pvme”; reconnect the CAT5e cables to SgPCI-2 board; reboot again; add SOLFpvme driver by “/opt/SOLFpvme/bin/a\_pvme”; add your leaf device drivers.

Make sure the /usr/kernel/drv/pvme.conf file has the right configurations.

CHAPTER 6

# Physical and Environmental Specifications

---

## 6.1 Mechanical Dimensions

---

SgPCI-2:  
Board Dimension: 6.5 in x 4 in (L x W)  
Front Panel Dimension: 4.75 in x 0.72 in (L x H)

SgPCI-LS:  
Board Dimension: 6.5 in x 4 in (L x W)  
Front Panel Dimension: 4.75 in x 0.72 in (L x H)

SgCPCI:  
Board Dimension: 9.2 in x 6.8 in (L x W)  
Front Panel Dimension: 10.3 in x 0.8 in (L x H)

VME Interface Board:  
Board Dimension: 9.2 in x 6.8 in (L x W)  
Front Panel Dimension: 10.3 in x 0.8 in (L x H)

---

## 6.2 Operating Environment

---

### 6.2.1 SgCPCI Board

Operational Humidity: 5-95% RH non condensing at 40 C

Non-Operational Humidity: 5-95% at 40 C

Operating Altitude: 10,000ft

Non-Operating Altitude: 40,000 ft.

Cooling Requirement: 80 LFM minimum

Power Consumption: 5V only @4.5A max

MTBF 100,000 poh

### **6.2.2 SgPCle-2, SgPCI-2 and SgPCI-LS**

Operational Humidity: 5-95% RH non condensing at 40 C

Non-Operating Humidity: 5-95% at 40 C

Operating Altitude: 10,000ft

Non-Operating Altitude: 40,000 ft.

Cooling Requirement: 80 LFM minimum

Power Consumption: 5V only @4.5A max

MTBF 100,000 poh

### **6.2.3 VME Interface Board**

Operational Humidity: 5-95% RH non condensing at 40 C

Non-Operating Humidity: 5-95% at 40 C

Operating Altitude: 10,000ft

Non-Operating Altitude: 40,000 ft.

Cooling Requirement: 80 LFM minimum

Power Consumption: 5V only @4.5A max

MTBF 100,000 poh

## Ordering Informations

### AC3K Upgrade Kit:

- 19" Rack-mount Sunfire V240 -- P/N: SV240-2CPU-2GHZ
- SgCPCI card with 7' CAT-5e cables -- P/N: 623-CPCI-2CN
- 2 x16' USB Extension cables -- P/N: CBL-USBX-A5M
- 6' SCSI HD cable -- P/N: CBL-VDH68MM-6FT
- USB Key board and Mouse -- P/N: SUN-X3738A
- Visicom Front panel -- P/N: PNL-PCi004

### General:

- 2 Channels PCI Express/VME host card (SgPCIE-2)--P/N:621e-001-2CN
- 2 Channels PCI/VME host card (SgPCI-2)--P/N:621-001-2CN
- SgCPCI Compact PCI to VME board--P/N:623-CPCI-2CN
- VME Interface board--P/N:CP1006
- StarPCIE-VME combo bridge (SgeVME)--P/N: 621e-002-SVME
- StarPCI-VME combo bridge (SgVME)--P/N: 621-002-SVME
- PCI Expansion System Controller board (SgPCI-LS)--P/N:519-001-2CN
- CBL-CAT5e,7 7Ft CATe 5 cable--P/N: CBL-CAT67
- CBL-CAT5e,15 15Ft CATe 5 cable--P/N: CBL-CAT615
- CBL-CAT5e,40 40Ft CATe 5 cable--P/N: CBL-CAT640
- Solaris nexus driver for SPARC--P/N: 621-PVME-006 VME-4
- VME-4 Solaris nexus driver for X86--P/N:621-PVMEX86-007
- User Manual--P/N: 621-PVME-UM01



CHAPTER 8

# Warranty, Maintenance, and Technical Support

Solflower products are warranted against defective materials and workmanship within the warranty period of one year from date of invoice. Within the warranty period, Solflower will, free of charge, repair or replace any defective unit covered by this warranty, shipping prepaid. A Return Material Authorization number (RMA # must be obtained from Solflower prior to the return of any defective product.

Solflower's warranty is limited to the repair or replacement policy described above and neither Solflower or it agent shall be responsible for consequential or special damages related to the use of their products.

Any questions or requests for repair/technical support should be directed to:

Solflower Computer, Inc.  
3337 Kifer Road  
Santa Clara, CA 95051  
Phone: (408) 733-8100  
Fax: (408) 733-8106  
E-Mail: [info@solflower.com](mailto:info@solflower.com)  
URL: <http://www.solflower.com>

