



Managing Power for Advanced Managed CompactPCI® Platforms

Introduction

CompactPCI® has grown in popularity over the last several years, gaining tremendous acceptance as the architecture of choice for equipment manufacturers developing high availability embedded applications. Due to its standards-based hot-swap and high availability capabilities, it is an ideal architecture for critical military and communications applications. Packet-based advanced managed platforms take this architecture further by employing comprehensive management and a suite of highly compatible and feature-complete components to deliver high availability platforms that are application-ready. However, when developing with these platforms, there are several power architecture issues to consider:

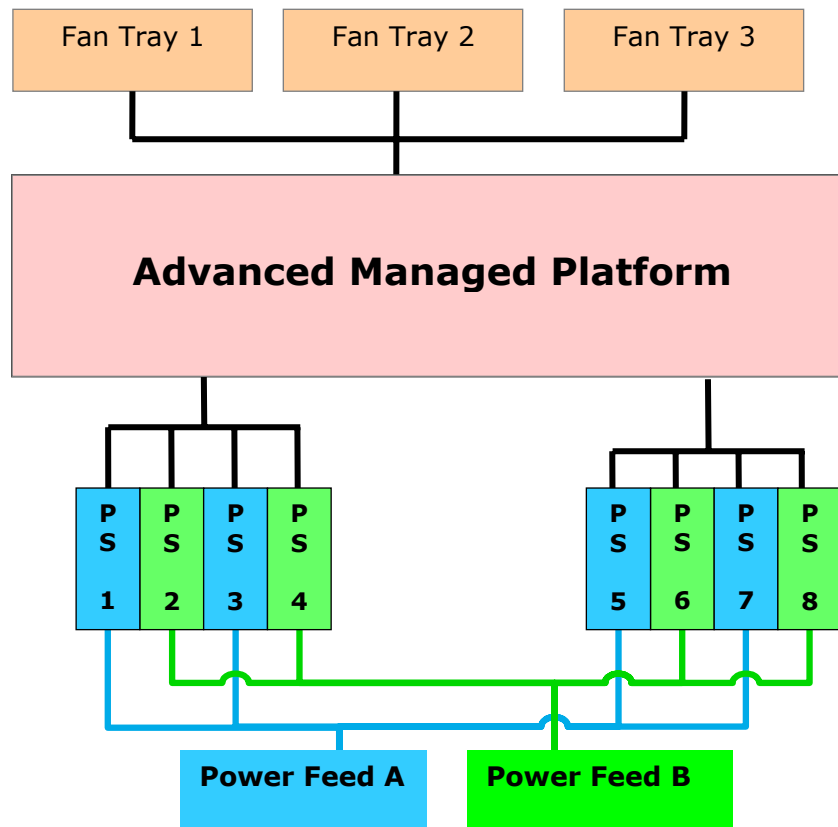
1. Calculating the power budget can be a complex issue. High performance boards tend to draw large amounts of power, so one must take into consideration all the elements that draw power, as well as the limited current for each of the four voltage rails: 5V, 3.3V, 12V and -12V.
2. Advanced managed platforms support multiple power configurations to balance between costs, power and redundancy.
3. A new breed of CompactPCI power supplies offers PICMG 2.9 Intelligent Platform Management Interface (IPMI). By adding intelligent management, power supplies can be monitored more closely, even to the point of predicting and preventing failures before they occur.
4. Isolated power input architecture provides another level of high availability and reliability to the platforms.

Advanced Managed Platform Power Architecture

12U Platform Power Architecture

Before we continue into the details, it is important to gain a good understanding of an advanced managed platform power architecture.

Figure 1: Power architecture for 12U advanced managed platforms



12U advanced managed platforms with DC input use two isolated -48V power feeds to deliver power into the platform. For high reliability, the two feeds are not ORED together with diodes. The power architecture is N+N, where N=4. Power feed A delivers power to the odd numbered power supplies and Power feed B to the even numbered supplies. This provides two levels of redundancy: power feed failure and power supply failure.

An advanced managed platform can handle any single failure at a given time. If a single power feed fails, the system continues to receive power from the four supplies powered from the healthy feed. If a single power supply fails, the system continues to receive adequate power from the healthy supplies. If the ambient temperature rises above 40°C, the power supplies and the platform (connectors, cables harnesses, midplane, etc.) can continue to operate for up to 96 hours up to 55°C ambient, as defined by NEBS. Four 325W power supplies deliver 1300W output power. It is important to note that the 325W rating actually means 325W of output power; the drop of power due to efficiency is already accounted for. The following table lists the effects of various types of power faults with eight 325W power supplies.

Table 1: Power architecture for 12U advanced managed platforms with eight 325W DC power supplies

Failure	Healthy Feeds	Healthy Supplies	Watts
No faults	2	8 (4+4)	1300W output
Power supply fault	2	7 (4+3 or 3+4)	1300W output
Feed A fault	1 (Feed B)	4	1300W output
Feed B fault	1 (Feed A)	4	1300W output
>40°C ambient	2	8 (4+4)	1300W output (up to 55°C for 96 hrs)

The AC power architecture is similar, except that there are four power feeds. AC input power is provided via four IEC connectors labeled A1, B1, A2 and B2. The A1 and B1 connections provide redundant power feeds to power supplies 1-4. For full redundancy, these connections must have independent circuit breakers and come from different sources (or at least different phases). The A2 and B2 connections provide redundant power feeds to power supplies 5-8.

4U Platform Power Architecture

4U advanced managed platforms use N+1 power architecture. In this case, there are a total of three power supplies, where one supply is used as a back up should another one fail. Using the same 325W DC power supplies, the platform can deliver 650W of redundant output power. However, the cooling architecture in this chassis limits the amount of power that can be utilized in this chassis to 500W of redundant power.

Power Budgeting

When dealing with high performance applications in a CompactPCI platform, it is important to develop a power budget analysis to ensure all components operate correctly. High performance applications will tax both the power and cooling architectures supported in the platform. In this section, we will focus on the power. This analysis can also provide a rough estimate of how much average power can be delivered to each platform slot.

Poor power budgeting skills look something like this: "Well the platform delivers a total of 1300W of redundant power, and I've got boards that draw a maximum of 50W of power each, so I can install 26 boards in the chassis." Not only did we not account for all the infrastructure elements that draw power, such as fan trays, switches and shelf managers, but we didn't account for the fact that there are four voltage rails in the platform, and each have limited power.

To properly budget power for 12U platforms, it is best to develop a spreadsheet to account for all the components that draw power. Performance Technologies has developed Platform Power Budget worksheets. You can access the spreadsheets at: http://www.pt.com/Wpaper_Power_exceltool.html. The following sections provide detail on how to properly budget power in the platform using these spreadsheets.

Step 1: Maximum Current Available for Each Platform

CompactPCI defines four independent and limited voltage rails that deliver power to the platform: 5V, 3.3V, 12V and -12V. It is important to analyze power budgeting at each voltage rail in addition to overall power of the chassis (e.g. 1300 watts), because one specific configuration of boards in a chassis may not tax the 5V rail but could significantly tax the 3.3V rail. Each voltage rail is independent of each other and does not share current. So as soon as one rail is maxed out, you have reached the limit of boards you can integrate into the chassis that use that specific voltage. Step one is to calculate the maximum current for each of the four voltage rails from the total set of power supplies in the platform. Remember to discount the redundant supply(s).

See Step 1 in one of the following example worksheets at: http://www.pt.com/Wpaper_Power_exceltool.html.

Step 2: Peak Power Rating per Component

Step 2 is to list the maximum current rating for each component and board that will be integrated into the platform. For each unique board, write down its maximum current draw in amps (A) for each voltage rail. It is important when calculating power budget to use the maximum power draw for each component to calculate worst case scenario. If its datasheet does not provide current ratings for each rail, contact the manufacturer. Be sure to include the maximum draw on the V (I/O) pins in with the 3.3V or 5V rail depending on the V (I/O) configuration of the system. The maximum power rating should take high environment temperature conditions into consideration. Typically, 40°C ambient temperature is used for each board and component. NEBS calls out that the shelf shall operate up to 40°C continuously or 55°C for 96 hours (where the one failure is the ambient temperature). The NEBS specification also assumes only one failure at a given time. Thus, if ambient is at 55°C, all power supplies are still operational and can power the boards successfully, even if they draw more current. It is important to also account for the resistance found in cables, such as wire harnesses that are found inside the platform. Their resistance increases as the temperature increases.

See Step 2 in one of the following example worksheets at: http://www.pt.com/Wpaper_Power_exceltool.html.

Step 3: Power Budget Analysis

Now let's configure these boards into the platform and calculate total power budget. Create a table that lists each board, the quantity of unique boards in the platform and the total current consumed per voltage rail. A simple example is shown below.

Board or Component	Qty in Platform	5V	3.3V	12V	-12V	Total Power (W)
Board 1	2	6A*	2A*	.20A*	0A	39W
Board 2	5	10A	30A	0A	0A	149W
Fan Tray	3	0A	0A	9A	0A	108W
Total Current Draw		16A	32A	9.2A	0A	296W
Total Current Available **		18A	30A	12A	2A	N/A
Delta		2A	(2A)	2.8A	2A	N/A

*Note: Board 1 uses 3A of 5V, but since there are two boards configured in the platform, the total is 6A. Do the same with the 3.3V and 12V rail.

**The numbers here would come from Step 1. These numbers are fictional and purely for this example.

- **Total Current Draw:** each value in the Total Current Draw row represents the sum of each voltage column.
- **Total Current Available:** the maximum current available from the platform we calculated in step 1.
- **Delta:** the difference between Total Current Available and the Total Current Draw from components and boards for each voltage rail.

See Step 3 in one of the following example worksheets at:

http://www.pt.com/Wpaper_Power_exceltool.html.

Temperature Coefficient

The resistance of pure metals, such as silver, copper and aluminum, increases as the temperature increases. The amount of increase in the resistance of a conductor per degree rise in temperature above 20°C is called the temperature coefficient of resistance. For copper, the value is approximately 0.004041. This applies to the entire length of wire and for each degree of temperature rise above 20°C.

$$R = R_{ref}[1 + \alpha(T - T_{ref})]$$

Where:

R = Conductor resistance at temperature "T"

R_{ref} = Conductor resistance at reference temperature T_{ref}, usually 20°C, but sometimes 0°C

α = Temperature coefficient of resistance for the conductor material

T = Conductor temperature in degrees C

T_{ref} = Reference temperature that α is specified at for the conductor material

For a 12U platform running at maximum power of 1300W at 40°C, the two wire harnesses will consume 82.2 W of power.

Redundancy in 12U Platforms

Eight Power Supplies

As discussed in the **Advanced Managed Platform Power Architecture** section of this paper, implementing eight 325W DC power supplies in a 12U platform delivers 1300W of redundant output power at two levels:

- **Power supply redundancy:** if one power supply fails, the others can still deliver full load to the platform.
- **Power feed redundancy:** if one power input feed fails, the other four supplies can deliver a full load to the platform.

Six Power Supplies

Depending on your specific configuration, you may opt to go with six power supplies. After going through the power budgeting exercise above, if you have determined that the total power budget required in the chassis is less than 975W (using DC power supplies) and the current per rail analysis is valid, then six supplies may be a valid option.

When loading only six supplies into a 12U platform, the valid configuration is as follows:

- Power bay 1: slot 1, slot 2, slot 4 [no power supply in slot 3]
- Power bay 2: slot 5, slot 7, slot 8 [no power supply in slot 6]

This is important, because the platform can't handle a failure of feed A if you load supplies 1, 2, 3, 4, 5 and 7 – you'll lose all four odd supplies. Slots 3 and 6 must have a 3U x 8HP air management blade.

Table 2: Power redundancy for platforms with six 325W DC power supplies

Failure	Healthy Feeds	Healthy Supplies	Watts
No faults	2	6 (3+3)	975W output
Power supply fault	2	5 (3+2 or 2+3)	975W output
Feed A fault	1 (Feed B)	3	975W output
Feed B fault	1 (Feed A)	3	975W output
>40°C Ambient	2	6 (3+3)	975W output (up to 55°C for 96 hrs)

Redundancy in 4U Platforms

Three Power Supplies

Table 3: Power redundancy for platforms with three 250W DC power supplies

Failure	Healthy Feeds	Healthy Supplies	Watts
No faults	2	3 (2+1)	500W output*
Power supply fault	2	2	500W output*
Feed A fault	1 (Feed B)	3 (2+1)	500W output*
Feed B fault	1 (Feed A)	3 (2+1)	500W output*
>40°C Ambient	2	3 (2+1)	500W output (up to 55°C for 96 hrs)

*4U platforms have a maximum power and cooling specification of 500W. When loading only two supplies into the 4U platform, the valid configuration is as follows:

- Slot 1: 3U x 8HP air management board
- Slot 2: power supply
- Slot 3: power supply

Two Power Supplies

Table 4: power redundancy for platforms with two 250W DC power supplies

Failure	Healthy Feeds	Healthy Supplies	Watts
No faults	2	2 (1+1)	250W output
Power supply fault	2	1	250W output
Feed A fault	1 (Feed B)	2	250W output
Feed B fault	1 (Feed A)	2	250W output
>40°C Ambient	2	2 (1+1)	250W output (up to 55°C for 96 hrs)

In a two supply configuration, the platform supports 500W with feed redundancy, but no supply redundancy.

Intelligent Management

To achieve higher levels of high availability, equipment manufacturers are requiring all active components to have intelligent management information. Intelligent management allows the system (OA&M) managers to understand the current condition of a component and predict failures before they occur. Power supplies that do not support the IPMI standard typically only provided the DEG# thermal warning and the FAL# fail signal. Intelligent power supplies monitor temperature, output voltages, output currents, thermal warning status and failed status. Additionally, intelligent management provides IPMI-based controller reset, geographic address and FRU information, such as part number, version number, serial number and manufacturing date. Since power supplies are typically one of the first components to experience failure, some system managers like to perform routine replacement of power supplies after a specific number of years, months or services. By reading the FRU information of each power supply through an ISM, the system manager can remotely detect which power supplies need replacing and set the power supplies LED red via the ISM.

Isolated Power Feeds

Newer central office (CO) environments are based on redundant -48V DC power sources, where power is supplied to equipment by two on-site power plants. AC power, delivered by the public utility, is rectified to -48V DC power within these on-site power plants (although it can range from -40 to -60V DC). These, in turn, supply the CO equipment. In the event that local AC power is interrupted; a bank of batteries within these power plants provides DC power until the motor generators take over. The benefit of supplying CO power in this manner is that when the equipment is designed to take advantage of it, it supports extremely high availability. In order to retain this level of redundancy through to the platform level, several platforms use an N+N power architecture. One set of power supplies (N) receives power from one feed, and the second set of power supplies receives power from the other feed. This keeps the feeds completely isolated from each other. In contrast, diodes tying two feeds together can compromise the integrity of the redundant feeds.

Conclusion

Fully understanding the power architecture of a CompactPCI platform is critical to the success of your application. Many companies have already undergone the power budgeting exercise, and hopefully you were successful with it as well. Most of the time with advanced managed platforms, there is ample power to drive the most demanding applications. But it is important to also consider future applications with even greater power requirements. With this in mind, your platform can be implemented for more products yet to come.

Configuring the chassis with fewer supplies not only saves cost, (for example, using six supplies in a 12U platform) but it also allows for scaling up to eight supplies for future applications that demand more power. It is important to realize that you don't have to sacrifice redundancy when configuring with fewer power supplies.



The importance of intelligent management in power supplies can't be stressed enough. For any high availability application, it is just as critical to have intelligent management in your power supplies as it is in your payload boards. Power supplies are active components, and they need to be monitored to predict failure.

Finally, the concept of isolated power inputs has quickly become an important issue to central office operators. As cost of ownership has increasingly become a main driver for their operations, any steps taken to prevent a breach of redundancy helps to decrease the cost of adding more infrastructure components to prevent failures or worse yet, downtime.

The power architecture of an embedded platform must be taken seriously when selecting and designing your application. Hopefully this paper has provided you with tools to understand it and to appreciate the architecture of advanced managed platforms.

Tony Romero

###

Performance Technologies

For more information, visit Performance Technologies' Web site at www.pt.com or contact sales@pt.com.

To learn about embedded hardware, enabling software and system-level products, visit: <http://www.pt.com/products/>.