

Performance, Cost and Reliability Considerations in a MicroTCA Power System

October 2007

Technical Paper

A tutorial for designers considering their first use of products based on the MicroTCA standard, and for the experienced designer looking for detailed information regarding optimal solutions for the power module design.

Contents

1. Introduction	2
2. Historical Perspective	3
3. Architectural Overview	4
3.1 AdvancedTCA	4
3.2 MicroTCA	6
4. MicroTCA Power Module Overview	8
5. MicroTCA Power Module Design Considerations	10
5.1 Hold-up Capacitance	10
5.2 Input Voltage	13
5.3 Redundancy	16
5.4 Dual Input Feeds	18
6. Conclusions and Summary	21
7. Glossary	23
8. References	23

1. Introduction

Micro Telecommunications Computing Architecture (MicroTCA™) is a relatively new architectural specification for Information and Communications Technology (ICT) equipment, and is a complement to the Advanced Telecommunications Computing Architecture (AdvancedTCA™) specification, but with a different intended product market. After a brief summary of the history of these architectures and a comparison of the two, this paper will address the MicroTCA power system in more detail, with an emphasis on the MicroTCA power module. It will be shown that design considerations within the MicroTCA power module can have a significant effect on the performance, cost and reliability of the resulting MicroTCA system. These design details can be important to the OEM, because while the MicroTCA specification defines several mandatory requirements in terms of functionality, interfaces and thermal/mechanical design of the MicroTCA power module, it also allows for internal design flexibility for the MicroTCA power module supplier. The design choices made by the supplier or requested by the customer will affect the overall system performance.

This paper may serve as a general tutorial to MicroTCA power systems for those who are knowledgeable about contemporary power system design but are considering their first designs to the MicroTCA specification. It should also be useful for those who are already conversant with MicroTCA but are looking for more information on the details of the power system implementation and choices in the power module design. In either instance, the reader must consult the latest version of the actual MicroTCA specification for the latest requirements before finalizing a system design choice. While the information contained in this paper represents our considered opinion, there may be new developments over time that allow alternative approaches to become more viable.

2. Historical Perspective

MicroTCA, which was ratified by the PCI Industrial Computer Manufacturers Group (PICMG™) in July 2006, is the latest generation of open-architecture platforms developed by the PICMG for ICT equipment. It builds upon the heritage of previous architectures and technology, namely AdvancedTCA and AdvancedMC, maintaining much of the same functionality but with different system partitioning and with optimization for applications with lower power levels such as Customer Premises Equipment (CPE), and Edge and Access equipment. The AdvancedTCA specification has been in existence since 2002. AdvancedTCA carrier-boards are large planar structures operating from -48 V and containing both power control/conversion circuitry and some of the load electronics. Additional load electronics may be packaged in Advanced Mezzanine Card (AdvancedMC™) modules which are mounted to the carrier-boards. An AdvancedTCA system rack contains several of these carrier-boards. Up to 14 carrier-boards may be installed in a 13U high 19 inch shelf, while up to 16 carrier-boards may be used in an European Telecommunications Standards Institute (ETSI) 600 mm enclosure.

With MicroTCA, all load electronics are packaged on AdvancedMC modules. These mezzanine modules are identical to those used with AdvancedTCA systems, leveraging development costs between the two architectures, providing a migration path and providing economies of scale for the production of AdvancedMCs. Time-to-market and spares inventory costs are also reduced due to a smaller number of required unique module types. A key feature of the MicroTCA system is the power module, which contains the majority of the power conversion and control circuitry and eliminates the need for the large planar carrier-boards of the AdvancedTCA systems. MicroTCA systems may also be packaged in 19 inch racks, with a 6U height considered a large system. Smaller enclosures are also possible.

Figure 1 contains photographs of the two types of systems. The AdvancedMC module is a common element between them, plugging into the backplane of the MicroTCA rack and onto the carrier-board which is then plugged into the AdvancedTCA equipment rack. While our emphasis will be on the MicroTCA, a description of both architectures is given in the following section so that commonalities and differences can be better understood.




Figure 1 - AdvancedMC modules make use in both AdvancedTCA and MicroTCA systems

3. Architectural Overview

The information below is intended to provide a basic overview of the architecture and power partitioning of AdvancedTCA and MicroTCA systems. The actual system specifications should be consulted for the latest detailed information.

3.1 AdvancedTCA

Figure 2 shows a typical power architecture for an AdvancedTCA system. Some power conditioning takes place prior to the individual carrier-boards. AC/DC conversion and battery backup is usually accomplished in a redundant fashion at a centralized location. The resulting -48 V DC power is then distributed to individual AdvancedTCA shelves. At the shelf level, a Power Entry Module is used to provide filtering and transient suppression. Then the individual redundant -48 V feeds are connected to the shelf backplane, which acts as an interface between the shelf level power conditioning and the power circuitry contained on each carrier-board.

Each carrier board provides fusing, ORing diodes, inrush current limiting, filtering, hold-up capacitance

and input voltage monitoring for the -48 V input power. The result is a reliable and robust source of input power for the rest of the carrier-board power system. The reader may recognize the remainder of the on-board power system as an example of an intermediate bus architecture (IBA). The main isolated DC/DC converter is normally selected to have an output of 12 V, since there are readily available Intermediate Bus Converter (IBC) modules marketed for this application and the AdvancedMC modules require 12 V as their main input voltage. The AdvancedTCA specification limits the total power consumption to a maximum of 200 W per carrier-board.

The load circuitry power is referred to as “payload” power in the AdvancedTCA specification. A carrier-board may contain payload circuitry mounted directly to its printed circuit board (PCB), in which case one or more Point of Load (POL) regulators are used to convert the 12 V intermediate bus to the voltages required by the payload circuitry. Another option is to mount one or more AdvancedMC modules on the carrier-board. These modules, by definition, will require 12 V as their payload power voltage. Any required POL regulation will be accomplished within the AdvancedMC module.

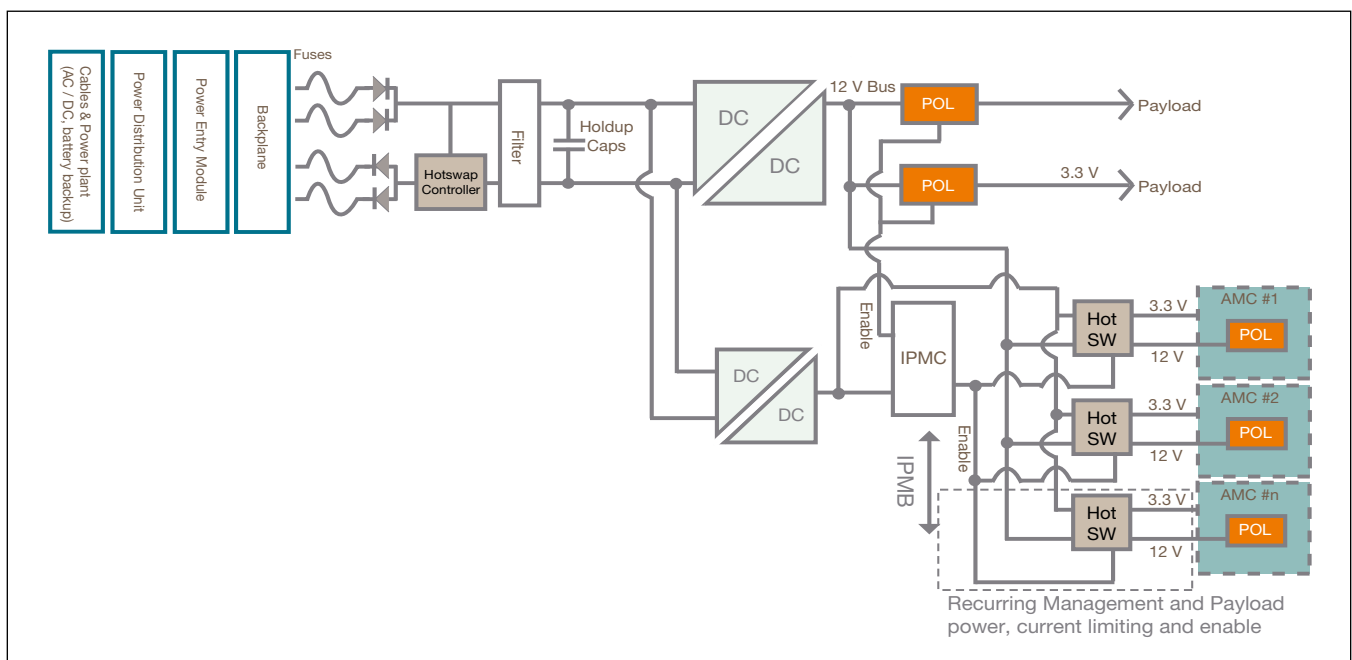


Figure 2 – Typical AdvancedTCA power system diagram




Figure 3 – AdvancedTCA carrier-board with AdvancedMC modules

Another function that must be provided on each carrier-board is power control, and each carrier-board will contain an Intelligent Platform Management Controller (IPMC) for this purpose. The specification requires that the control circuitry be active and in communication with shelf-level management to negotiate power-up rights before a maximum of 10 W of power is used by the carrier-board. This requirement could for example be accomplished by providing a separate 3.3 V isolated DC/DC converter on each carrier-board that is dedicated to powering the control circuitry, both the IPMC on the carrier-board and also the management power to each AdvancedMC that may be used. By this method, individual AdvancedMC modules can receive management power without the IPMC having powered up the full payload circuitry of the carrier-board. In addition, the carrier-board power control must provide the functions of voltage monitoring, current limiting, power sequencing and hotswap control

for each AdvancedMC location on the board.

Thus each carrier-board requires a high degree of power conditioning and control functionality, both for input power and for payload power located on either the carrier-board or on AdvancedMC modules. This high degree of functionality can result in a carrier-board assembly like the example shown in Figure 3. The carrier-board is a planar structure that is 280 mm deep and 322 mm high. In the example shown, the board contains DC/DC conversion from -48 V down to an intermediate bus voltage that provides payload power to two AdvancedMC slots. The bus voltage is also used to feed down-stream POL regulators for on-board payload circuitry.

3.2 MicroTCA

MicroTCA is intended to be a complement to AdvancedTCA rather than a replacement for it. The MicroTCA architecture offers several advantages for its intended markets. The smaller form factor and lower hardware cost makes it attractive for applications that require less processing power within an enclosure, for example, for edge, access and customer premises equipment. While smaller and more cost effective, the reliability and availability requirements for MicroTCA systems are typically just as stringent as those for equipment implemented with AdvancedTCA. The same basic functionality in terms of power conditioning and control is also required. One of the main differences between the two architectures is the degree of centralization and the physical partitioning of the power systems.

With AdvancedTCA, all power conversion functions are replicated on every carrier-board. There is also the flexibility to locate payload circuitry either on AdvancedMC modules, on the carrier-board PCB, or both. The MicroTCA approach simplifies this partitioning by requiring all payload circuitry to reside

in AdvancedMC modules and by centralizing all the main power conversion/control functions for a sub-rack into one or more MicroTCA power modules. The overall architecture of a MicroTCA system is shown in Figure 4. A complete MicroTCA system is defined in the specification as follows: “A minimum MicroTCA system is defined as a collection of interconnected elements consisting of at least one AdvancedMC, at least one MicroTCA Carrier Hub (MCH), and the interconnect, power, cooling and mechanical resources needed to support them.” The system as shown in the figure supports up to a maximum of 12 AdvancedMC modules which contain the payload circuitry, and each of these AdvancedMCs is specified to require from 20 to 80 watts of payload power. Again quoting from the MicroTCA specification, “They provide the functional elements needed to implement useful system functions. Examples of AdvancedMCs that could be installed into a MicroTCA shelf include CPUs, Digital Signal Processing devices, packet processors, storage, and various sorts of I/O AdvancedMCs (including metallic and optical line units, RF devices, and interfaces to other boxes).” The MCH function provides overall control for the interconnected AdvancedMCs. A second redundant

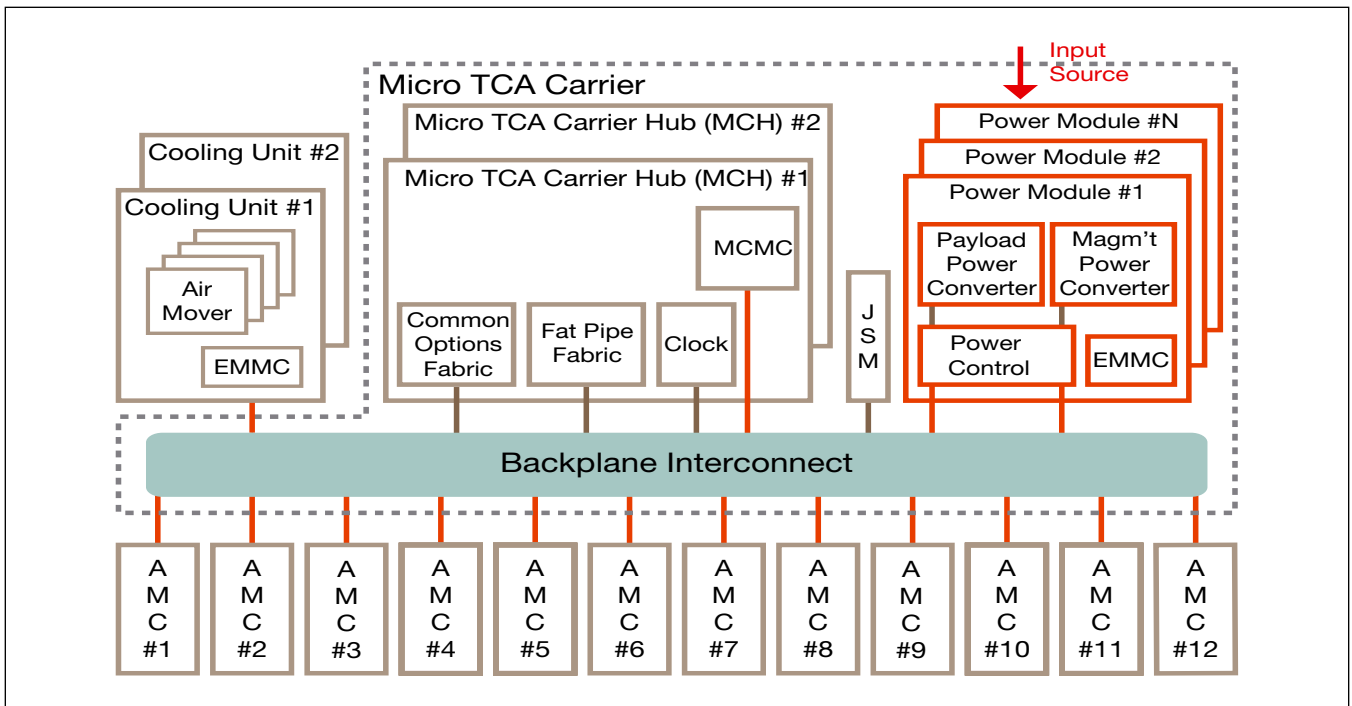


Figure 4 – MicroTCA overview with parts of the power system highlighted

MCH is often added for systems with high availability requirements. Similarly, a second redundant cooling system is sometimes used. The backplane is used as an interconnection mechanism for all of these elements. The power module is a very key element in the overall MicroTCA system. It serves as a centralized power conditioning, conversion and control block for the entire sub-rack. Anywhere from one to four power modules may be used in a single MicroTCA system with more than one being used either because of the power demand or because of a desire for redundancy.

The MicroTCA specification provides a very clear description of the power module purpose and functionality: “The MicroTCA power module(s) take the input supply and convert it to 12 V to provide payload power to each AdvancedMC. 3.3 V management power for AdvancedMCs is also supplied by the power subsystem. The power control logic on the power module performs sequencing, protection, and isolation functions. The Power Subsystem is controlled by the Carrier Manager which performs power budgeting to ensure adequate power is available prior to enabling Power Channels.” “Power modules also include the supervisory functions necessary to manage the Power Subsystem. They have circuitry to detect the presence of AdvancedMCs, MCHs, and Cooling Units (CU), and to energize individual power branches. Power modules also monitor the power quality of each branch and protect against overload. If a redundant power module is configured, it may automatically take over the Power Channel responsibilities of a failed primary power module.” In addition to supplying payload and management power to up to 12 AdvancedMC modules, the power modules must be capable of supplying payload and management power to up to two CUs and two MCHs. As a consequence, many power modules are designed to provide a total of 16 output power channels, or 32 channels if payload and management power are counted separately.

This is obviously a lot of functionality to include into one assembly, and as a consequence the power module is one of the most important elements of the overall MicroTCA system. It includes more power and control functionality than is required of the power elements on an AdvancedTCA carrier-board but is contained in a smaller package. Obviously the design of the power module will have a large impact on the overall system efficiency and reliability.



Figure 5 – MicroTCA sub-rack holding AdvancedMCs and power modules

The power module will be the focus of the remainder of this paper. First the overall functionality and partitioning will be described and then some of the important design details will be discussed. Figure 5 shows how power modules can be packaged in a MicroTCA sub-rack. The system shown uses two power modules and they are located at the extreme right and left ends of the upper row of plugged modules. DC input power is plugged to the connectors on the front panel of the power modules, while 12 V and 3.3 V payload and management power is connected to the MicroTCA backplane at the rear of the power modules.

4. MicroTCA Power Module Overview

MicroTCA places a significant amount of functional content into the power module, including:

- Input power ORing
- Hotswap control for input power inrush protection
- Input power filtering
- Power hold-up capacitance
- 48 V to 12 V DC/DC conversion (payload power)
- Input to output isolation
- 12 V to 3.3 V conversion (management power)
- Output power distribution
- Hotswap control for multiple Advanced MCs, CUs, MCHs
- Output power monitoring and control
- Output power protection circuitry

The consolidation of both power handling circuitry and system level control/management functionality from the large AdvancedTCA carrier-board into the relatively small MicroTCA power module means that the power module design, performance and reliability are all crucial to the success of the overall system. Figure 6 is a block diagram showing the content of a typical MicroTCA power module. Most of the functionality is similar to that of an AdvancedTCA carrier-board power system, but there are differences.

With AdvancedTCA, there is a Power Entry Module (PEM) which provides some input power conditioning and protection in the form of transient protection and filtering before the power is distributed to the carrier-board. With MicroTCA, there is no PEM, so “raw” DC input power is supplied directly to the input connectors on the front of the power module. This means that the functionality of the PEM must be included in each power module. AdvancedTCA carrier-boards contain fuses on the -48 V inputs. In the case of MicroTCA, fusing is typically done in a power distribution unit by providing fuses for each of the cables that

distribute input power to the front of the power modules. Consequently, no internal fuses are normally contained within a power module. Otherwise, the front-end functions are very similar to AdvancedTCA – ORing diodes, EMI filtering, inrush current limiting and hold-up capacitance. While AdvancedTCA is specified to always contain dual redundant -48 V input power feeds, MicroTCA may be configured with either a single feed or with two redundant feeds.

The power module contains a single -48 V to 12 V isolated DC/DC converter, similar to the approach taken with AdvancedTCA, but with power levels up to 600 W. A POL regulator from the 12 V output is used for generation of management power. Additional POLs, within the AdvancedMC modules are used to derive their needed low voltage power from the 12 V payload power. The control mechanism for MicroTCA power modules is the Enhanced Module Management Controller (EMMC), which monitors and controls both management and payload power for all of the AdvancedMCs, CUs and MCHs configured in the system.

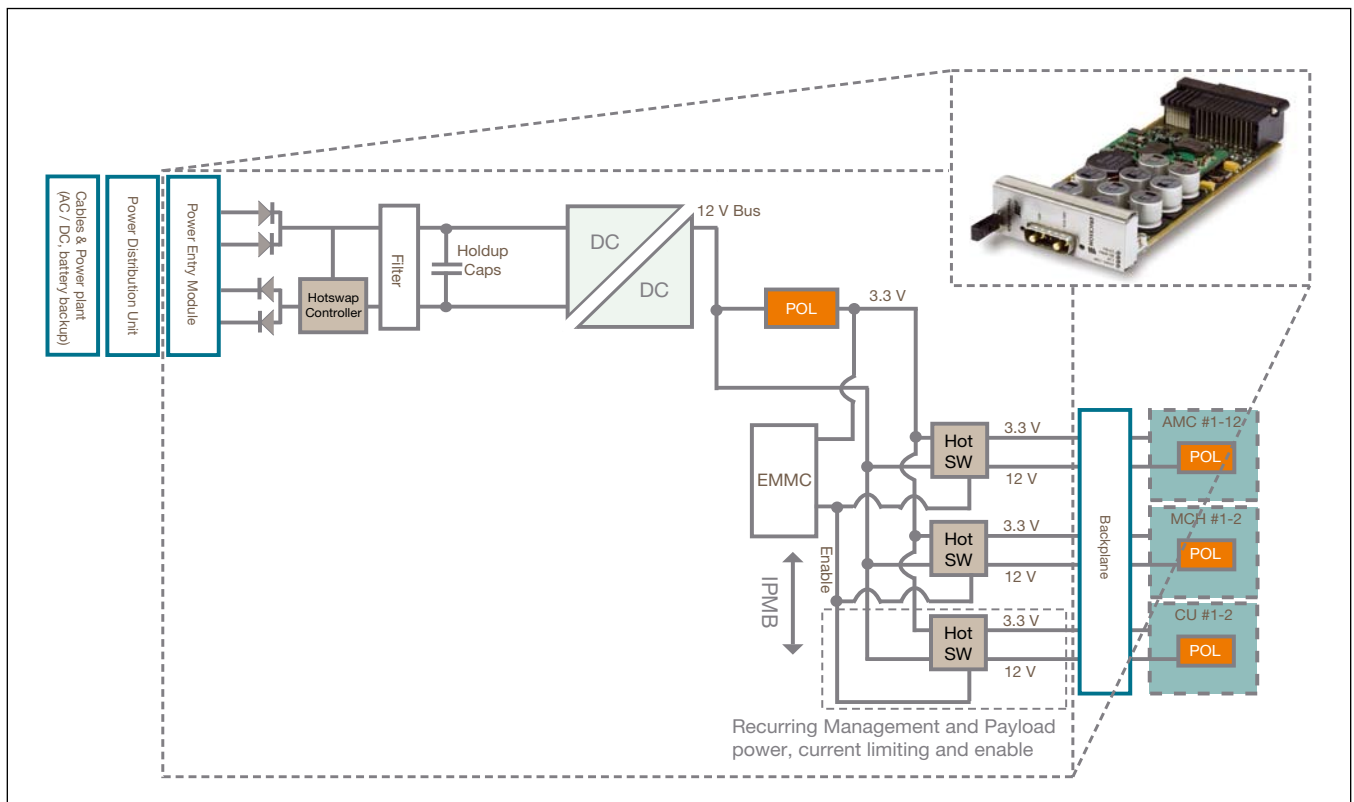


Figure 6 – Typical MicroTCA power module block-diagram

A closer view of an actual MicroTCA power module is shown in Figure 7. This particular power module is configured in a single-width full-height form factor with approximate external dimensions of 73.5 by 186.6 by 28.9 mm. It is capable of supplying and controlling power to 12 AdvancedMCs, 2 CUs and 2 MCHs, for a total of 32 output voltage channels.

The repartitioning of the distributed power conversion functions from several AdvancedTCA carrier-boards into a relatively few (one to four) centralized MicroTCA power modules results in a higher power conversion density, with some power modules delivering up to 600 W. As a result, high efficiency design is extremely important both for packaging purposes and also to achieve the reliability requirements for these systems. Input power inrush, EMI control, and hold-up design are also more difficult in the MicroTCA environment since it must be done to the same

standards but at a maximum power level of 600 W rather than the 200 W seen on an AdvancedTCA carrier-board. The control and management demands are also challenging, with the need to interface with up to 32 output voltage channels as well as with the shelf-level MCHs.

The above demands make the design and selection of the MicroTCA power module a critical element in the success of the overall system. In the following section we will explore in more detail some aspects of the power module design that can help insure successful system designs. While MicroTCA systems may be configured with other input voltage sources such as 24 VDC or universal AC, the following discussion will be limited to the most commonly used telecom -48 VDC input voltage.




Figure 7 – Example of MicroTCA power module



Figure 8 – Ericsson MicroTCA power module ROA 117 5078/1

5. MicroTCA Power Module Design Considerations

The MicroTCA specification actually contains three levels of designation, which in everyday language can be called “shall”, “should” and “may”. That is, it defines a level of absolute requirements and then two additional levels that are intended as recommendations and guidelines but which allow for some flexibility as to the implementation of MicroTCA component assemblies and systems so that they are best suited for the actual application. This flexibility is an advantage to the system designer who must balance performance, reliability and cost and in most cases must make some sort of trade-off between these attributes. The specification and design of MicroTCA power modules represents an example of this flexibility in that their parameters may vary and still meet the provisions of the MicroTCA specification. Ericsson has undertaken a detailed study of some of these parameters in order to determine how the selected performance level affects other attributes of the power module, including cost. This information is essential for the OEM system designers when speci-

fying and selecting MicroTCA power modules. The design considerations to be discussed in this paper are hold-up capacitance, input voltage, redundancy and dual input feeds.

The baseline hardware used for this study is an Ericsson MicroTCA power module with an output power rating of 355 W as shown in Figure 8. The basic specifications for this unit, with its standard implementation, are as follows:

Output Power	355 W
Input Power	385 W
Output Voltage Channels	16 x 12 V & 16 x 3.3 V
Efficiency at 50 % load	95 %
Normal Input Voltage (full performance)	-40.5 V to -57 V
Hold-up (54 V in)	10 ms
Conducted Emissions	Class B
Package form factor	Single-width Full-height (6HP)

The term “cost unit” will be used to define cost. This unit is normalized based on the Q2-2007 cost estimate for the prototype power module, which is assigned a value of 400 cost units for the entire bill of material. If a design change results in a savings of 20 cost units, this would represent a 5 % reduction in the total bill of material cost. Thus the cost conclusions stated in this paper, although represented in relative terms, should provide some guidance for the system designer when making power module specification trade-off decisions.

5.1 Hold-up Capacitance

Inside a typical MicroTCA power module there is a number of fairly large electrolytic capacitors. The presence of most of this capacitance is due to one of the recommendations in the MicroTCA specification which addresses operating through an input power interruption that is caused by, for example, a short circuit fault on the input voltage bus.

The specification defines this outage under a worst-case scenario as a reduction in input voltage to a level of 5 V for a duration of 10 ms, and the power module is expected to operate without interruption through this fault condition. Power module designers normally accomplish this by including several “hold-up” capacitors on the -48 V input after the ORing

diodes. The power module operates from the stored energy in these capacitors for up to 10 ms until the normal input voltage is restored.

The derivation of this recommendation can be better understood by referring to Figure 9, which represents a fairly typical MicroTCA system. The system shown consists of one cabinet with two shelves. Each shelf contains one power module, which receives its -48 V power feed from a Power Distribution Unit (PDU) via a separate power cable. Assume that a short circuit occurs on the -48 V feed to shelf 1 as shown in the diagram. Each cable is individually protected in the PDU, so the fuse or circuit breaker will open and isolate the fault condition from the remainder of the system. But the fault clearing is not instantaneous – there will be some time period before the fuse opens during which the fault current will pull the -48 V for the entire cabinet to a voltage below the operating voltage specification. Consequently, the power module in shelf 2 must operate through an interruption in its input voltage. The hold-up time of 10 ms at a voltage of 5 V is recommended in the MicroTCA specification to represent the worst-case condition in terms of the fault clearing time and the robustness (voltage drop) of the -48 V power distribution system.

The above specification is certainly a safe one and will result in a reliable system. But there are circumstances in which the same system reliability and per-

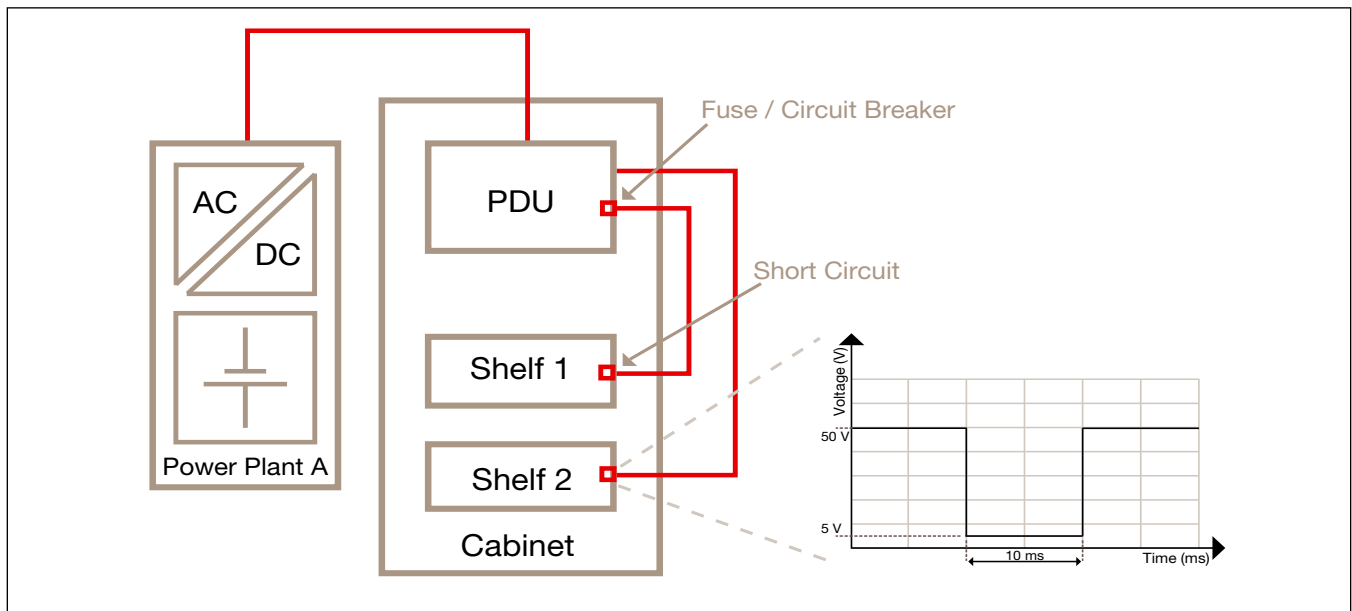


Figure 9 – Short circuit event resulting in voltage drop in power feed

formance can be obtained with less hold-up time and consequently less capacitance. Some examples are:

- If the cabinet in the application only contains one shelf and one power module, the fault condition presented above is not a valid scenario, since the only power module would be inoperative after the fuse opens anyhow. In this situation, no hold-up capacitance at all could be an acceptable design.
- The system designer should understand the fault clearing behavior of the devices in the PDU. There are fuses and circuit breakers available that will open in less than 10 ms. If, for example, the design guarantees fault clearing in 5 ms, the required amount of stored energy and the number of hold-up capacitors could be cut in half.
- The specification assumes that the power module is operating at full rated output power during the fault condition. In most properly designed systems this is not the case, as some design margin is allowed for by running the power modules at less than full load. If the maximum actual load is less than the rated maximum load, then the amount of hold-up capacitance for a given hold-up time will be reduced accordingly.
- Some system designers use what is referred to as Two-Step High Ohmic Distribution (TS-HOD). With this technique, the -48 V cabling contains a pre-defined resistance. This will limit the fault current and the resulting voltage drop will not go below -40.5 V, which ensures that the power module stays within its normal operating range.

These examples show situations in which the hold-up capacitance can be reduced without any degradation in system performance or reliability. Note, however, that they all depend upon the knowledge of the system designer about the actual application. Once this understanding is obtained, the system designer can work with the power module supplier to specify the appropriate hold-up time.

What are the cost impacts of the hold-up time specification? The power module used in the study is rated at 10 ms hold-up time at full load at an input voltage of -54 V, which is the nominal input voltage for a -48 V system with battery back-up. A short circuit fault during a static input voltage of less than that would represent a double fault condition which is not normally designed for. The hold-up capacitors in the subject power module are Nichicon 63 V LS series electrolytics. They consume 1100 mm² of PCB area, which represents 10 % of the total PCB. The cost of the hold-up capacitors is 2 cost units. Using fewer hold-up capacitors would only have a small effect on component cost but a very positive impact on PCB component area. The latter improvement could come to good use in other parts of the design.

Another design approach which is not pursued or quantified here would be the addition of a voltage boost circuit in the front end of the power module. Such a booster would charge the capacitors at a higher voltage such as -72 V. Then higher voltage capacitors would be used and a lower amount of capacitance would be required for a given hold-up time since the stored energy is proportional to the square of the capacitor charge voltage. Some of the savings would however be off-set since a given capacitor holds less capacitance when rated for a higher voltage. Plus, the design would also have to carry the extra circuitry for the booster.

5.2 Input Voltage

Another parameter that needs to be specified by the system designer is the input operating voltage range. As a general rule, the narrower the voltage range, the more the power module design can be optimized in terms of performance, efficiency and cost. Most often, the input voltage is specified at the normal -48 V telecom range of -40.5 to -57 VDC, with a nominal value of -54 V. Some systems are required to also operate over the range of the less commonly used -60 V telecom power system, which can vary

from -50 to -72 VDC. In this part of our study, we attempted to quantify the performance and cost impacts of extending the input voltage range to address both -48 V and -60 V systems versus designing for -48 V only.

The input and hold-up capacitors must of course be resized to operate at the higher input voltage by changing the capacitor voltage rating from 63 V to 80 V. Higher voltage capacitors have less volumetric efficiency in terms of capacitance per volume, so the higher voltage capacitors will require additional volume and PCB area within the power module. It is true that when operating from the -60 V system less total capacitance would be required for hold-up due to the larger amount of energy stored in a given capacitance. But since the analysis considers both -60 V and -48 V power sources, the amount of capacitance must be predicated on the worst case scenario (i.e. the -48 V application). It is worth mentioning that this calculation was based on 80 V capacitors, while ones rated for 100 V would most likely be needed for the sake of design margin. That would make the higher input range even less favorable.

The results of our study are shown in Figure 10. The data in the 40.5 – 57 VDC column represents the baseline design as presented in the previous section. When the capacitors are changed to 80 V rated devices to accommodate the -60 V system requirements, both the PCB area and the cost increase as shown in the rightmost column. An additional 550 mm² of PCB area are required and the cost of the power module increases by approximately 0.5 cost units. If a voltage booster topology as described in the previous section were used, this analysis would not apply.

Hold-up Capacitance		
Operating Range	40.5-57 VDC	40.5-72 VDC
Nominal Input	-54 VDC	
Duration	10 ms	
Output Power	355 W	

Nichicon LS-series	63 V	80 V
PCB area	1100 mm ²	1650 mm ²
PCB area of PM	10 %	14 %
Cost	2 units	2.5 units

Figure 10 – Study results

The effect of the maximum input voltage on the power module efficiency was also studied. Specifically, we examined the efficiency differences that could be measured in the main 48 V to 12 V DC/DC converter. As a general rule, a lower input voltage means that the primary switch transistors can have a lower voltage rating, which should imply a lower on resistance and less power losses. A PKM 4304B PI isolated DC/DC converter was used for this study. This design uses 100 V primary power transistors, suitable for both -48 V and -60 V power systems. As an experiment the 100 V devices were replaced with 60 V devices from the same supplier and product line. These 60 V transistors will only support operation with -48 V systems. There is a 2.5 milliohm reduction in on-resistance with the 60 V devices which should theoretically result in a 0.3 W reduction in power losses at maximum load. The results are shown in Figure 11.

There was indeed a reduction in power losses at full load with the 60 V devices, but it was relatively small. However when operated at less than 50 % load, the efficiency was actually less when using the 60 V transistors. The change in the shape of the efficiency

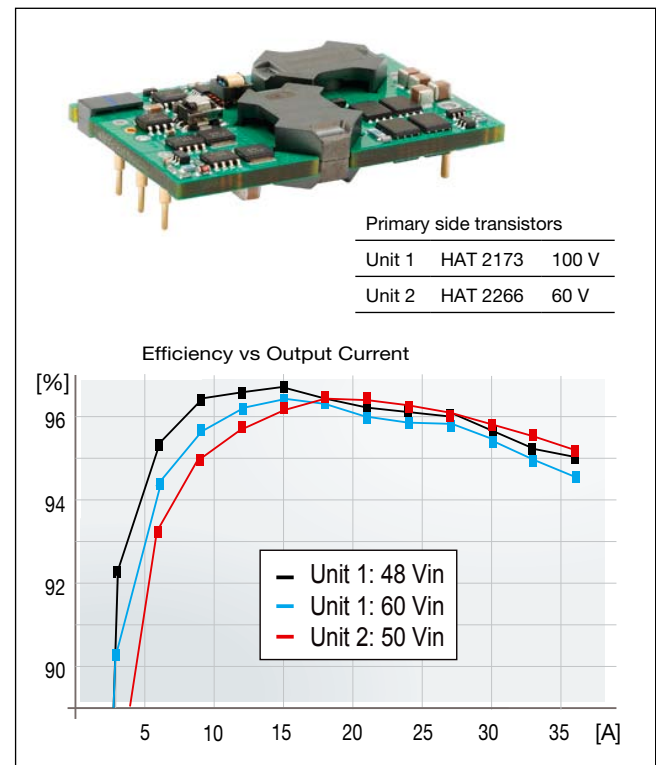


Figure 11 – PKM4304B PI and Efficiency dependent of input voltage

curve is probably due to the voltage levels at which the lower rated transistors turn on and off rather than the actual DC resistance. The baseline design has been optimized for 100 V transistors. Expending significant additional engineering time with the 60 V devices might change the outcome shown. Nonetheless, for now we cannot conclude that the broader input voltage range required to address both -48 V and -60 V systems has any significant penalty on operating efficiency for this particular device.

5.3 Redundancy

The MicroTCA specification includes provision for redundant power modules to increase system availability in critical applications. When needed, this capability can function quite well and achieve the system availability goals. It is important to understand, however, that power modules designed for redundant operation are inherently more complex and costly than power modules intended for stand-alone operation. The basic power module redundancy approach used with MicroTCA payload and management power channels will first be described for the benefit of those who may not be familiar with it. Then two aspects of the power module design that are affected by the redundancy decision, payload power channel control and DC/DC converter performance will be discussed. The intent is to inform the OEM

designer about the size, efficiency and cost impacts of redundancy to the power module design so that the redundancy approach is used when it is actually needed and the benefits of a non-redundant power module may be enjoyed in systems with lower availability requirements.

An example of a 2+1 redundant MicroTCA power module implementation is shown in Figure 12. In this system, two power modules are used to supply both payload and management power to a total of 16 output channels. A third power module is normally in a stand-by state and is available to provide power to any of the 16 channels (32 voltage outputs) in the event of a fault in either of the two main power modules. The MicroTCA specification contains very specific requirements for the implementation of power module redundancy. Techniques such as power paralleling and current sharing are not intended to be used, and only one power module may deliver current to any load channel at any given time. This restriction can be seen in the system shown in the figure. Power module 1 supplies the normal power to only channels 1 through 8, while power module 2 does the same for channels 9 through 16. The redundant power module 3 can supply power to any of the 16 output channels, but only in case of failure in one of the primary power modules or if one of the two has been disabled. This architectural restriction was established so that the maximum overcurrent possible to any channel is limited. If two power modules were paralleled, the

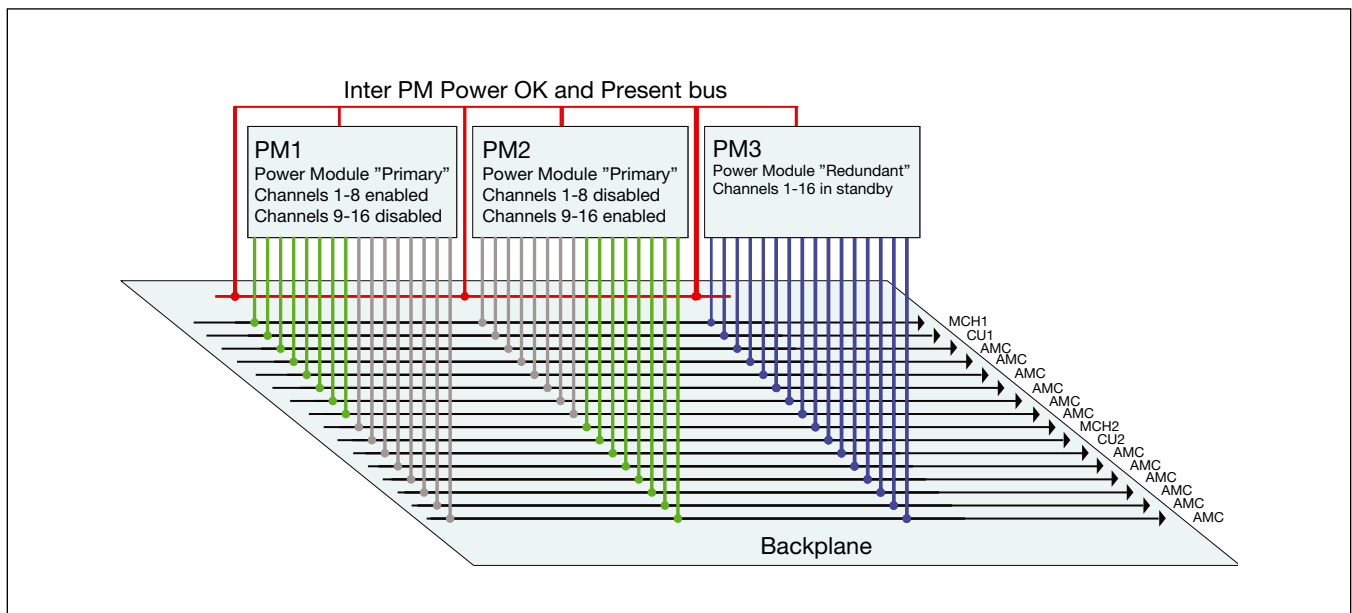


Figure 12 – Example of 2+1 redundant MicroTCA power module implementation

maximum fault current could be doubled, which would expose the system backplane and connectors to excessive current and possibility of damage.

The MicroTCA specification requires that any given power module be identified to the system as either a primary power module or a redundant power module. A given power module within the system may transition between these two roles as decided by the MCH, but one power module could not maintain both roles at the same time. In the event of a failure in any output channel of a primary power module, the redundant power module will take over responsibility for all output channels of that primary power module, not just the failed channel. Automatic transition between a failed primary power module and the redundant power module is accomplished by the settings of their output voltages. Primary power modules are set to a higher output voltage than redundant power modules, the two nominal settings being perhaps 12.5 V and 11.5 V. This output ORing allows instantaneous and automatic transition in the event of a failure due to the power module with the higher output voltage delivering power to the loads. This technique also imposes much more stringent voltage budgets and output regulation requirements on power modules (including the primary power modules) used in redundant systems. This impact on the power mod-

ule design will be discussed later in this section.

To understand the impacts of redundancy on the power module output channel control, it is easiest to first examine a typical implementation for a non-redundant power module, which is shown in Figure 13. The drawing is for only one payload channel, but all of the elements shown other than the DC/DC converter and the EMMC would need to be replicated for each of the payload power channels as well as for each of the management power channels – i.e. up to 32 times. Due to the significantly lower current levels, the management power channels do not pose as much of a design challenge, so concentrating on a single payload channel is sufficient to gain an understanding of the situation. There is a single DC/DC and a single EMMC in each power module, and these functions are shared with all 32 output channels.

The circuit block shown between the EMMC and the current sense resistor and output control transistors is most often some kind of hotswap control IC. Often these ICs can handle multiple channels, so the number of chips required will vary but the described functionality is independently required for each of the up to 32 output voltage channels. Each channel has two semiconductor switches in series. The switch to the left is the pass device and the one on the right is

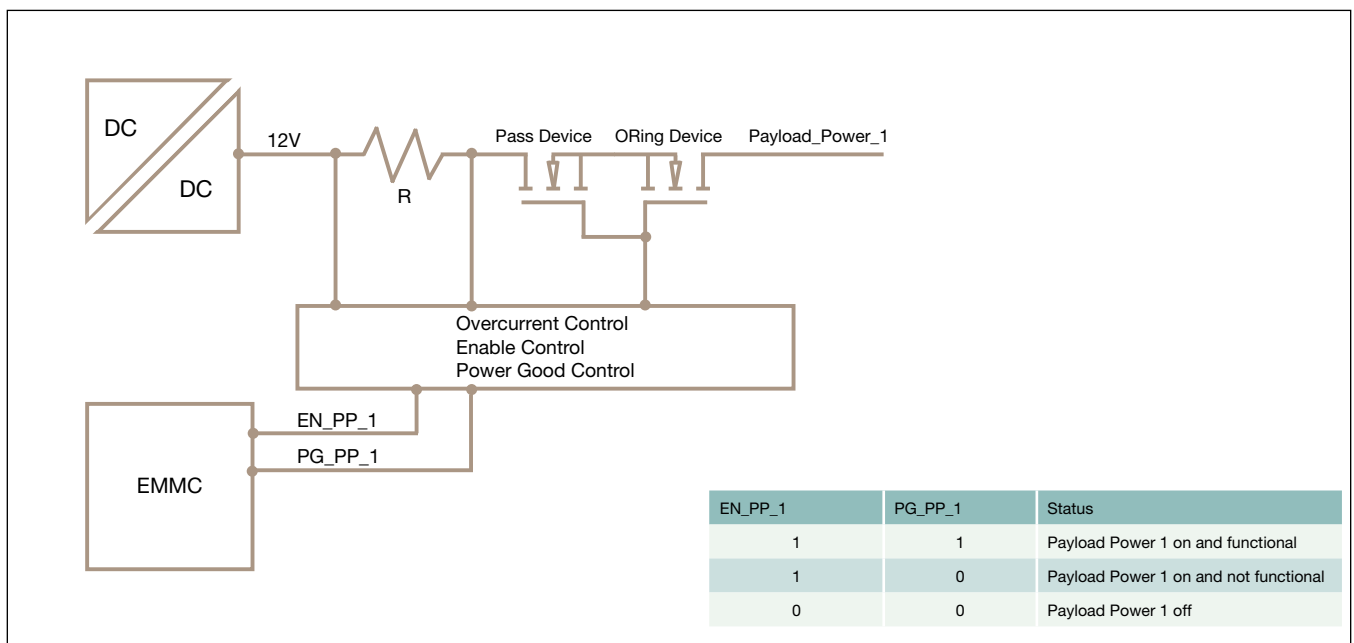


Figure 13 – Payload channel that only allows for non-redundant operation

the ORing device. The ORing device prevents current from flowing in the reverse direction from the load into the power module. The pass device is used to enable or inhibit the output current and also to limit the value of the current to provide for functions such as soft start for hotswap and fault current limiting.

Since this is a non-redundant power module, its outputs will be either on or off. There is no need for it to be in a stand-by state ready to take over for another power module. Therefore both of the transistors can be driven by the same control line as shown in the figure. This results in a fairly simple implementation, with only two control lines (enable and power good), and only a total of three defined conditions for this payload power channel. The three conditions are:

- Channel OFF
- Channel ON and functional
- Channel ON with fault

Note that non-redundancy does not limit the number of power modules in a given sub-rack. There can be more than one, but each power module needs to be assigned to specific AdvancedMCs, CUs or MCHs and there would be no hand-over between the power modules in the event of a failure.

The typical implementation of the output channels for redundant operation is much more complex, as shown in Figure 14. The same circuit elements are used, but there are more interconnections and control states. The EMMC might need to be connected to the 12 V DC/DC converter so that the converter output voltage can be set to the appropriate value as a function of the power module's definition as either primary or redundant. In figure 14 this connection is exemplified by a Power Management Bus (PMBus™). The PMBus may also connect to the hotswap control function in order to obtain data collection capability from the output channels. There is also a control line between the EMMC and the hotswap control function that sets the primary/redundant definition to the hotswap controller. Note that for the redundant implementation, the ORing switch is driven separately from the pass device. If the specific power module in question is defined as a redundant device, this ORing switch is normally turned off, and the intrinsic diode of the switch is reverse biased due to the lower voltage setting of the redundant DC/DC converter. This intrinsic diode will become forward biased to automatically deliver current from the redundant power module in the event of a failure in the primary power module. The ORing transistor is then turned on by the control logic to reduce the on-resistance of the

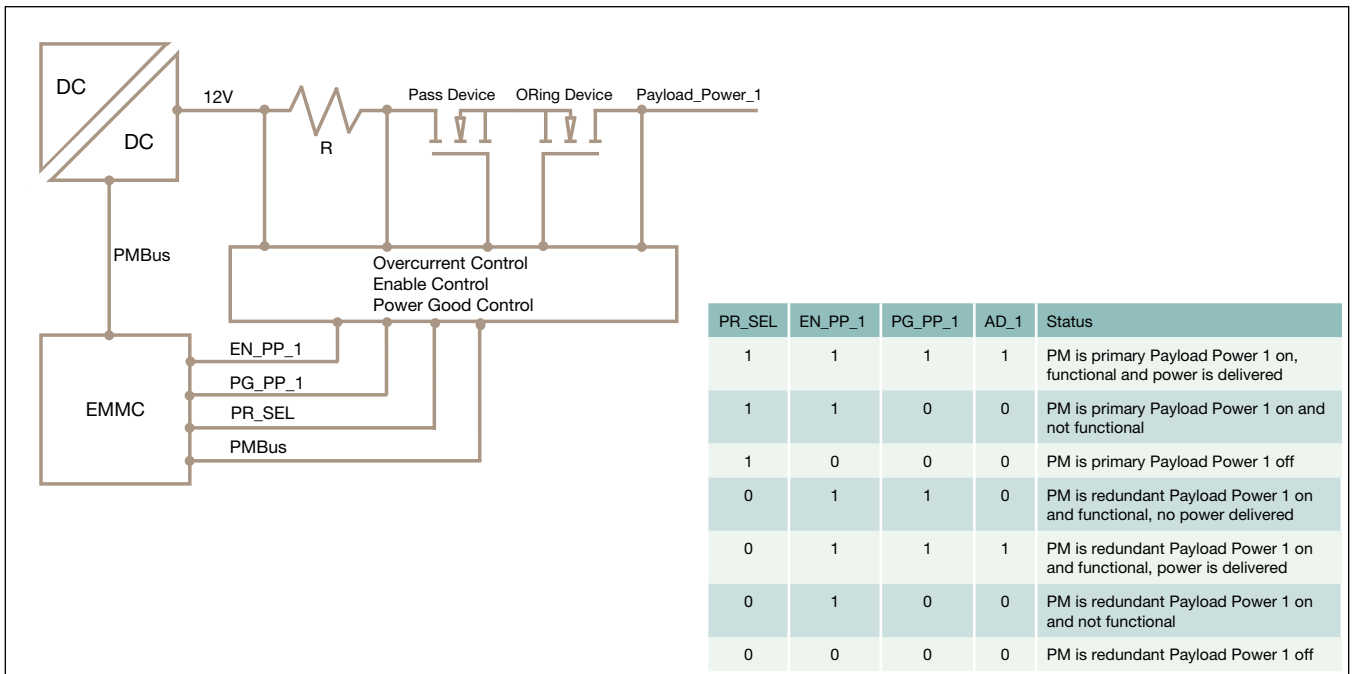


Figure 14 – Payload channel that allows for redundant operation

connection and to reduce power dissipation in the concerned component.

The net of all this is that the redundant implementation results in considerable extra complexity. There are now four interface lines between the EMMC and hotswap function instead of two, and the number of defined states that need to be controlled is now seven rather than three. The PMBus connection to the DC/DC converter may also be needed. In addition, the accuracy of the current limiting needs to be much higher when using redundancy. Compared with the non-redundant setup shown in Figure 13, the redundant solution requires an additional 300 mm² of PCB area to accommodate the hardware for the increased control complexity. This represents about 2.5 % of the total PCB area in the power module. The control cost impact for the redundant solution is about 10 cost units over and above that required for the non-redundant power module. These estimates are for the 16 payload channels. The impact of the lower current management channels is relatively minor. It should also be noted that the above analysis is predicated on the hotswap hardware available in 2006. As semiconductor vendors develop more highly integrated and flexible channel control devices for the MicroTCA market this situation may change. In addition to the cost differentials for each power module as described here, a redundant solution will of course require at least one additional power module compared to a non-redundant system.

We will now examine the impacts of redundancy on the 12 V DC/DC converter. The basic MicroTCA specification defines the tolerance range for the AdvancedMC module input voltage as 10 V to 14 V. Since the load module will operate at any voltage in this range, the 12 V DC/DC converter could have a +/- 10 % tolerance in a non-redundant system. In a redundant system, the situation becomes more challenging. In order to keep the voltage budgets of both the primary and the redundant power modules within the same overall range at the AdvancedMC inputs without possibility of overlap, the tolerance ranges for the primary power module would be approximately 12.25 V to 12.95 V and the range for the redundant power module from 11.6 V to 12.0 V. These ranges include the effects of line and load regulation as well as temperature. This means that the DC/DC converter in a power module intended for operation in a redundant system must have a +/- 2 % output

voltage tolerance. Going from a +/- 10 % to a +/- 2 % regulation tolerance has a significant impact on the DC/DC converter design.

A redesign of the baseline converter used in this study was not done to quantify this impact, but a meaningful analysis can be obtained by looking at two other DC/DC converters in the Ericsson product line.

Figures 15 and 16 summarize the parameters of two Ericsson DC/DC converters with 12 V outputs and




Figure 15 – Performance comparison with and without feedback loop.

approximately the same input voltage range. They are both very contemporary designs and are highly regarded as representing industry-leading performance in terms of efficiency and power density given their respective design assumptions. They both have exactly the same form factors and total volume. The PKM 4304 is more loosely regulated with only feed-forward regulation from the input line voltage and no load regulation feedback loop. This greatly simplifies the module's control system, but does create a droop in its output load characteristic as shown in the figure. The additional space freed up by the less complex control system was used to enhance the power train resulting in high efficiency (95.3 %) and output power (380 W). This converter would be well suited for usage in a power module not intended for redundant applications.

PKM 4313C PI	
Input Range	38 - 75 VDC
Size	36.8 x 57.9 mm (WxL)
Vout tolerance	+/- 2.5 %
Efficiency	93.3 % (48 V in, max load)
Output Power	204 W

PKM 4304B PI	
Input Range	36 - 75 VDC
Size	36.8 x 57.9 mm (WxL)
Vout tolerance	- 10.8 / +4.2 %
Efficiency	95.3 % (48 V in, max load)
Output Power	380 W

Figure 16 – Summary of performance parameters

The PKM 4313, in the same sized physical package, contains output voltage feedback and features output voltage regulation of +/- 2.5 %, making it suitable for usage in a power module for redundant applications. But there are penalties for this enhanced performance. The efficiency is 93.3 %, significantly lower than that of the PKM 4304. Also, the maximum power output is 204 W. The power density is only 54 % that of the PKM 4304. We can conclude from this that power modules used in redundant systems will have higher power losses than those intended for non-redundant systems, and that their internal packaging will be more challenging.

These numbers are of course continuously improving as new technologies are being developed. The fact, however, remains that tighter control of the output voltage will require additional control circuitry in the DC/DC converter, thus affecting power density and efficiency.

5.4 Dual Input Feeds

The unique physical partitioning of MicroTCA actually provides some opportunity for enhanced system availability by means of redundant DC/DC conversion. Other types of systems based on a -48 V bus voltage in the backplane normally utilize redundant -48 V power feeds, but contain only one 48 V to low voltage DC/DC converter on each carrier-board. This DC/DC converter represents a single-point-failure source with no redundancy. Providing redundancy in the form of a second DC/DC converter on each carrier-board would be prohibitive in terms of cost and board space, since it would need to be replicated for every carrier-board in the system.

MicroTCA provides a neat and effective solution for this dilemma. MicroTCA equipment racks are normally designed with provision for two power modules. If each power module position is fed from a separate -48 V power feed, the power modules can be easily configured into a redundant arrangement so that complete redundancy is provided for both loss of one power feed and for all power conditioning, conversion and control functions servicing all AdvancedMCs in the system.

This can all be accomplished with only one additional power module assembly, rather than the multiple additional high power DC/DC converters that would be required in a system based on a -48 V backplane. This approach gives OEMs using MicroTCA an opportunity to offer complete power system redundancy and extremely high rack level system availability all within a small enclosure and at a reasonable cost. It is important to emphasize that the -48 V backplane analogy is only included to explain the architectural differences. Actual effects on system level reliability cannot be compared as a DC/DC converter feeds a single board in one case, while in the other a complete shelf. Further examination of the redundancy options available with MicroTCA should help clarify

the possible system design decisions. One very key point to keep in mind is that providing power feed redundancy does not automatically imply that power modules will require dual power feed input capability.

A generalized drawing of a dual power feed situation is shown in Figure 17. Here the cabinet and shelves are supported with dual feeds. The main question to be addressed next is how to utilize these dual feeds. Three possibilities will be considered, as follows:

- One dual input power module
- Two redundant single input power modules
- Two redundant dual input power modules

One dual input power module - This is a non-redundant power module implementation, where a single power module supplies the entire shelf. The power module has dual input power feeds, which does provide for redundancy in the event of a loss of one of the power feeds, but no redundancy exists for the loss of the main DC/DC converter. That is, the DC/DC converter represents a single point-of-failure. The system designer should have the data to determine if this is a viable design direction, but it could be argued that the failure rate of a power feed may be less than that of the DC/DC converter, making one of

the other options below more attractive.

Two redundant single input power modules - This is a 1+1 power module redundancy implementation where power feed A is connected to one power module and power feed B to the other. Both power modules only have a single input power feed connection, and only one power module is required to supply all of the system loads. Redundancy is provided for both DC/DC converters and input power feeds. This solution is an improvement over the previous scenario because of the DC/DC converter redundancy.

Two redundant dual input power modules - This solution differs from the previous one only in that both power feeds go to both power modules and both power modules require dual input power feed capability. As with the previous method, this approach provides redundancy for both single point DC/DC converter and power feed failures. It does offer additional protection from multiple point failures, essentially providing 1+3 redundancy for input power feed faults along with the 1+1 redundancy for power module faults. This really represents protection against multiple failures at the same time, such as up to three blown fuses or cable faults in the PDU or a simultaneous failure of a power feed and a DC/DC converter. There may be a few systems where this

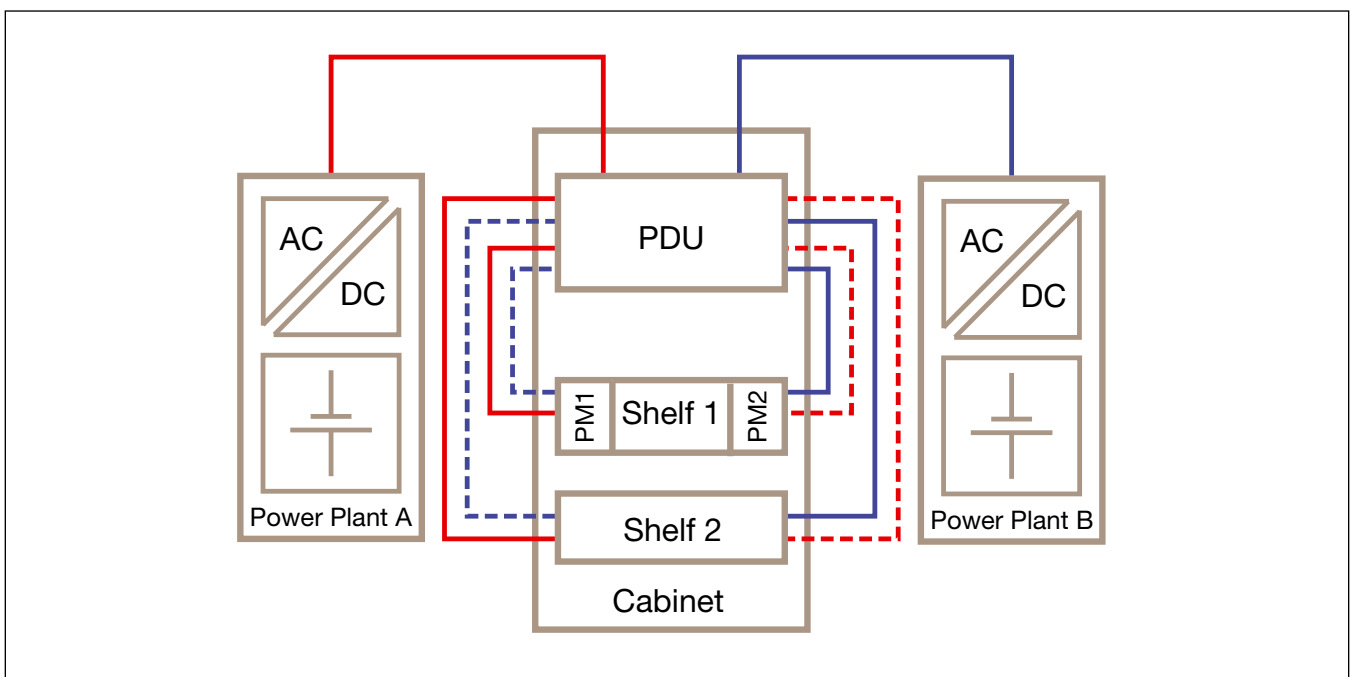


Figure 17 – Dual power feed general set-up

level of protection is desirable, but many MicroTCA applications will achieve their availability targets with only protection from a single simultaneous failure.

While only the system designer will be in a position to make the above trade-offs for a particular application, it is our contention that many MicroTCA systems requiring dual power feed redundancy may be best served by the second option. It provides single failure protection for both input power feeds and for DC/DC converters but does not require dual feeds to each power module. This analysis only applies to the assumed 1+1 power module redundancy. In other cases the conclusions will vary. For example in a 3+1 redundant power module system with single feed power modules, the loss of one DC power source will disable two power modules, resulting in an output current overload to the remaining two modules. The cost, efficiency and size impacts of providing dual input feed capability to a power module will be examined below. The system designer must essentially trade off these impacts vs. the capability of protection from multiple simultaneous failures.

A comparison of the front end implementations for single feed and dual feed power modules is shown in Figure 18. The single feed system uses an “active diode” with a forward resistance of only 12 milliohms for reverse polarity protection. This device is actually a transistor connected to simulate a diode with only connection to the input voltage (no external control needed). Providing input polarity protection for the dual feed input is considerably more complex, requiring a total of four conventional diodes to reliably address all the provisions of the MicroTCA specification. The difference in power loss and efficiency is significant, with the dual feed input dissipating 10 W in diode losses vs. only 1 W for the single feed power module. The net efficiency reduction for the dual feed module is 2.7 %. Using the dual feed design requires an additional 750 mm² of PCB area as well as increasing the cost by 12 cost units. Much of this additional cost is driven by the need for a second input power connector. These penalties may outweigh the advantage of protection from some multiple simultaneous failures in many system designs.

There are techniques that would allow for transistors to be used even in dual input configurations, thus eliminating the power dissipation generated by the diodes. Such an implementation would however require a rather complex control system to ensure that

the transistors are turned on when they are supposed to be turned on. And even more importantly, turned off when they are supposed to be turned off. Making sure that this is the case at the same time as providing reverse polarity protection and eliminating cross conduction between the two power feeds is quite a challenge using transistors. Diodes on the other hand, would naturally provide the robustness and reliability needed.

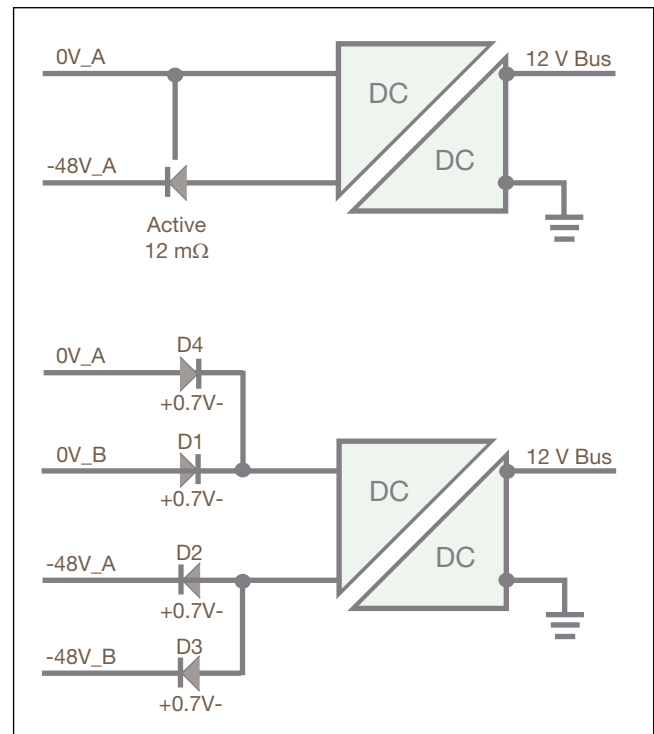


Figure 18 – Single feed and Dual feed power modules

Single Input	
PCB area	200 mm ²
PCB area of Power Module	1.7 %
Loss @ 380 W; 54 V in	1 W
Efficiency loss	0.3 %
Cost	9 units

Dual Input	
PCB area	950 mm ²
PCB area of Power Module	8 %
Loss @ 380 W; 54 V in	10 W
Efficiency loss	3 %
Cost	21 units

6. Conclusions and Summary

It is difficult to form valid generalizations about many of the topics discussed in this paper because the MicroTCA power module cannot be viewed as a stand-alone entity, but rather must be considered as an element in the overall system. Consequently, the system requirements for the particular application being considered will be the primary driver for determining the correct answer in any given situation. As such, the system designer will be the person who determines the requirements for system elements such as power modules. The intent of this paper was to communicate to the system designer how some of these decisions can affect the cost, performance, efficiency and power density of the power module so that more informed decisions are possible. Keeping in mind that they will not be universally applicable, the following conclusions are offered in the spirit of some general guidelines that may be useful. Figure 19 contains a summary of the power module impacts.

<p>Hold-up capacitance (10 ms short-circuit)</p> <ul style="list-style-type: none"> - 10 % of power module area and 2 units in cost - Need for hold-up related to multiple power modules on the same -48 VDC supply - Amount of hold-up dependent on power consumption, input voltage and short-circuit duration
<p>Input voltage (-48 VDC and -60 VDC)</p> <ul style="list-style-type: none"> - Minor effects on DC/DC converter efficiency - Slight increase in cost and real-estate of hold-up capacitance
<p>Redundancy (redundant operation of power module)</p> <ul style="list-style-type: none"> - 2.5 % increase in real-estate and 10 units in cost added - More advanced control logic - Significant reduction in DC/DC converter power density and efficiency
<p>Dual input feeds (-48 VDC into power module)</p> <ul style="list-style-type: none"> - 10 W of power loss - 21 units in cost and 8 % of power module area - Reliability improvement in non-redundant system - Reliability improvement negligible in redundant system

Figure 19 – Power module impact summary

- MicroTCA is an exciting and successful complement to the more established AdvancedTCA, offering significant benefits for smaller, lower powered, lower cost systems.
- MicroTCA is capable of high-availability redundant power system solutions, and can even provide enhanced availability by means of redundant DC/DC conversion.
- Commercially available MicroTCA components and system elements, including power modules, will continue to evolve based on user demand and technology enhancement. Probable trends include increased silicon integration, higher packaging density and lower production costs and pricing.
- Because it contains an extensive amount of power conditioning, conversion and control functions, the power module is a very key element in achieving a successful and reliable MicroTCA system design.
- System design decisions can affect the performance, efficiency, size and cost of a MicroTCA power module.
- Hold-up Capacitance - The 2 cost units required to implement the full 10 ms hold-up requirement is not in itself a major deterrent. The capacitors used are high quality low failure rate units and conservative design derating is used, so the reliability impact of the number of capacitors is also not a concern. The main impact to the overall power module design is the amount of PCB area needed for the implementation of the hold-up function, which is already approximately 10 % of the total PCB in a 355 W output power module. With the power output of a single-width full-height power module projected to reach 600 W in the future, a corresponding increase in hold-up PCB area will create significant design challenges. Using one or more of the hold-up specification approaches discussed in this paper can mitigate this concern.

- **Input Voltage** - Designing the power module for both -48 V and -60 V power systems rather than only for -48 V does not result in significant impacts on either cost or performance. There may be a slight change to the efficiency curve and an incremental impact to the amount of hold-up capacitance required. Perhaps a more significant impact of this decision, which is outside the scope of this paper's analysis, is the complication that including -60 V may have in terms of safety agency approval. It will entail designing for operating voltages going above Safety Extra Low Voltage (SELV) standards which will require additional testing and greater creepage and clearance distances inside the power module.
- **Redundancy** - The cost differential for providing redundant power module operation is presently somewhat significant at 10 cost units. This cost, along with the PCB real estate requirement, is expected to diminish over time as more highly integrated hotswap semiconductor solutions become available. A larger cost to the system designer, which may not diminish, is the need for more extensive software development, performance verification and interoperability testing for the significantly more complex redundant power module implementation. From a power module perspective, the most significant impact of providing redundant operation is the substantially tighter regulation requirements imposed on the DC/DC converter. The tighter requirements are technically very feasible even with today's technology, but result in significantly lower power density and conversion efficiency. And, of course, using a redundant power module configuration will require the cost of at least one additional power module.
- **Dual Input Feeds** - Of the four power module design areas explored, this one has perhaps the highest overall impact. Dual feed inputs to a power module add approximately 12 cost units as well as 9 watts of additional power dissipation and a correspondingly lower efficiency. For many MicroTCA systems these penalties need not be incurred since a solution is offered that provides dual feed redundancy and redundant DC/DC operation at the shelf level without the need for dual feeds to any power module.

This paper will hopefully be a useful guide to some of the design decisions that need to be made when configuring a MicroTCA power system. The content, however, should not be considered as the final authority on the issues presented. The system designer should always consult with the latest version of the appropriate MicroTCA specification when determining design requirements. Ericsson intends to continue its investment in developing industry-leading MicroTCA power solutions as well as to continue our commitment to providing open dialog with our customers about the design trade-offs inherent with this exciting new architecture.

Trademarks

Ericsson and the Ericsson logotype is the trademark or registered trademark of Telefonaktiebolaget LM Ericsson. All other product or service names mentioned in this document are trademarks of their respective companies.

7. Glossary

AdvancedMC™, AMC	Advanced Mezzanine Card
AdvancedTCA™, ATCA	Advanced Telecommunications Computing Architecture
PCB	Printed Circuit Board
CPE	Customer Premises Equipment
CU	Cooling Unit
EMMC	Enhanced Module Management Controller
ETSI	European Telecommunications Standards Institute
IBA	Intermediate Bus Architecture
IBC	Intermediate Bus Converter
IC	Integrated Circuit
ICT	Information and Communications Technology
IPMB	Intelligent Platform Management Bus
IPMC	Intelligent Platform Management Controller
IPMI	Intelligent Platform Management Interface
MCH	MicroTCA Carrier Hub
MicroTCA™	Micro Telecommunications Computing Architecture
PDU	Power Distribution Unit
PEM	Power Entry Module
PICMG™	PCI Industrial Computer Manufacturers Group
PMBus™	Power Management Bus
POL	Point of Load
SELV	Safety Extra Low Voltage
TS-HOD	Two-Step High Ohmic Distribution

8. References

1. AdvancedMC base specification R2.0, PICMG, 15 November 2006
2. AdvancedTCA base specification R2.0 ECN001 & ECN002, PICMG, 26 May 2006
3. MicroTCA base specification R1.0, PICMG, 6 July 2006