Concurrent Technologies Inc
Suite 130, 3840 Packard Road,
Ann Arbor, MI 48108, USA

Telephone: (734) 971 6309
Fax (734) 971 6350
E-mail info@gocct.com
Web www.gocct.com

Hybricon
12 Willow Road
Ayer, MA 01432 USA

Telephone: (978) 772 5422
Fax (978) 772 2963
E-mail info@hybricon.com
Web www.hybricon.com

HTF122A

VITA 31.1: Leveraging PICMG 2.16 Packet Switching technology in VME64x based systems

Technology Focus

WP Revision History: -	
#0	Original document

1. Executive Summary

This paper will provide an introduction to the ANSI approved VITA 31.1-2003 Gigabit Ethernet on VME64x Backplanes standard. VITA 31.1 is based on the PICMG 2.16 CompactPCI Packet Switching Backplane specification, and in fact it adopts the same connectors and pin-out for the Fabric Slots that are used in PICMG 2.16. This allows VITA 31.1 adopters to utilize PICMG 2.16 Fabric Switch boards that are already commercially available from multiple vendors.

VITA 31.1 is one of the newest additions to the VME family of specifications based on the IEEE 1101.10/1101.11 mechanical standards that are used for rugged, reliable systems in applications for Military/Aerospace, Industrial, and Telecom markets, among others.

This specification provides, for the first time, a standardized way to implement an Ethernet Switch Fabric in a VME64x backplane without the use of cables, while using standardized backplane connector pin-outs. This provides a clean, efficient way to implement multiprocessor architectures that leverage available Ethernet hardware, IP software, and clustering software.

There is no doubt that industry is moving toward switch fabric architectures to supplement, and possibly replace, traditional busses. Switch Fabrics are attractive because they abstract the data transfer and get away from the software complexities of address register mapping. They also minimize the number of interconnect signals required, and eliminate bus loading skew problems that limit the performance of parallel busses.

In today's embedded world, customers are fast moving away from expensive custom system designs and are looking for ease of system integration (software and hardware), scalable performance and system reliability.

2. VME background

The ANSI/VITA 1.1 VME64x standard defines 6U card sizes. 3U card sizes are supported by the legacy IEEE VME standard, and there are additional ANSI/VITA standards for 9U VME form factors. However, the 6U form factor is the mainstream VME technology. Each VME64x card slot has (3) backplane connectors; (2) DIN style connectors plus (1) 2MM hard metric style connector.

Existing VME64x systems suffer from a number of shortcomings:

1. The existing VME bus supports limited bandwidth.
2. The VME bus is a single point of failure, affecting system availability.

These limitations prompted VITA (VME International Trade Association) to form the VITA 31.1 working group to investigate a better method of transferring large amounts of data between cards, with an emphasis on packet traffic.

3. Introduction to the VME64x Packet Switching Backplane

Switched serial interconnects are a better match for the functional densities of today's board-level functions. Today, the 6U x 160mm VME board can be a self-contained computing node, where even some I/O based boards have the ability to process data on-board. Peer-to-peer communications between these intelligent devices is better supported in a switched serial environment than over a bus. Plus, it is not just the technical focus that is changing – customers are looking for ease of system integration (software and hardware), scalable performance and system reliability.

A developing market segment in VME64x applications is the addition of fabric cards with a connection for a packet switching backplane to P0 connectors. The VMEbus becomes the control plane while large data packets are moved on the P0 data plane. The P0 connector is a 2MM HM connector that has a maximum bandwidth of about 1.2 GHz. Therefore, customer applications can use up to Gigabit Ethernet, in a star or dual star connection, in the original P0 connector. To ratify this, as a standard, ANSI approved VITA 31.1-2003 Gigabit Ethernet on VME64x in June 2003. This standard defines the pin-outs of two Ethernet connections on the 2MM HM connector. VITA 31.1 pin-outs, connector type and layout are compatible to PICMG 2.16 switched Ethernet fabric cards. This means that existing PICMG 2.16 switch fabric cards can be used in a VME64x system to interconnect VITA 31.1 computing nodes within a VITA 31.1 based chassis.

Recognizing that VITA 31.1 provides a reusable serial interconnect standard, which is related to the mature PICMG 2.16 standard, Concurrent Technologies (as a single board computer manufacturer) and Hybricon (as a backplane and chassis manufacturer) both have products available today. See section 8.

4. Features and benefits of VITA 31.1

	Features	Benefits
Ease of Integration	VITA 31.1 VME 64x Backplane. Providing conventional connections for VME board's P1, P2 and P0 connections.	Legacy and non-VITA 31.1 VME boards can still be used alongside VITA 31.1 boards, whether computing or I/O nodes.
	Concise P0/J0 pin-outs, for the dual Ethernet connections, require only 16 connections per Node Slot.	VITA 31.1 does not require P1/J1 or P2/J2. Remainder of P0/J0 can be used for customer based rear I/O.
	Standards based architecture (I.e. Ethernet, commodity silicon, IP etc.)	Uses widely accepted industry standard, Ethernet 10/100/1000 Base-T. Faster time to market by utilizing existing IP software and Ethernet hardware. 20 years invested in the development of the network and the TCP/IP stack can now be leveraged in a switched serial environment. Faster time to market by re-using PICMG 2.16 CompactPCI Fabric Switch boards that are already commercially available
Scalable Performance	Point-to-Point interconnections across the backplane.	Scalable system performance: Up to 2 Gigabits/sec per slot. Up to 40 Gigabits/sec per chassis.
	By utilizing the switch fabric architecture a standard system can support many user configurations and options.	Scalable cost: from single fabric card/10Mbit/sec node cards configurations up to dual fabric cards/Gigabit/sec node cards Scalable bit rate/slot (10-1000Mbit/sec) using Ethernet "Auto Negotiation". Upgradeable architecture by simply adding/changing fabrics/nodes. "Virtual backplane" capable allowing for multi-chassis architectures.
Reliability	1 or 2 fully independent fabrics.	In a solely bussed system the bus is a single point of failure. Switched serial interconnects are inherently easier to make redundant.
	Point-to-Point interconnections across the backplane.	"Fine" granularity in point-to-point architecture, single point of failure equals 1 slot.

5. Ease of System Integration

Industry Standards

The use of IEEE 802.3 10/100/1000 Base-T Ethernet as the data link layer is part of the PICMG philosophy of using industry standard, well-understood technology, and adapting it for reliable, rugged systems. In the case of VITA 31.1, Ethernet is used to transfer traffic between Node Cards and Switch Fabric Cards. Ethernet was chosen as the data-link layer because it is cheap, simple, reliable, and software is well proven. These are the same reasons why Ethernet is the networking technology of choice for networks worldwide.

Gigabit Ethernet on VME64x

The backplane for VITA 31.1 is designed for backwards compatibility with legacy VME cards as well as providing a standard for specifying a set of differential pin-out connections on P0/J0 for one or two Gigabit Ethernet links. The impact on available I/O pins is minimal. Node Slots/Cards utilize pins on the 2MM HM P0/J0 connectors, while Fabric Slots/Cards utilize pins on the 2MM HM P3/J3 – P5/J5 connectors (PICMG 2.16 CompactPCI Fabric Slot pin-out using 2MM HM connectors).

It is interesting to note that the VME bus on P1/J1 and P2/J2 is optional. In fact, these backplanes could have no VME bus whatsoever, relying on the point-to-point Ethernet Switch Fabric for data transfer.

Switch Fabric Architecture

VITA 31.1 defines a Switch Fabric Architecture, as contrasted with a Bus Architecture (e.g. VME). A Switch Fabric is an interconnected network of switching devices and the topology is a star (not a bus) as shown in Figure 5.1. The fabric (board) contains many input and output ports, and the switch fabric transports the data from the inputs to the outputs. This means that in order to travel from one VME slot to another, data passes through an active switching device. In a VITA 31.1 (or PICMG 2.16) system, the data travels via a packet switching backplane and the active switching devices are located on the Fabric Card(s). The other cards that connect to the Fabric Card(s) are called Node Cards.

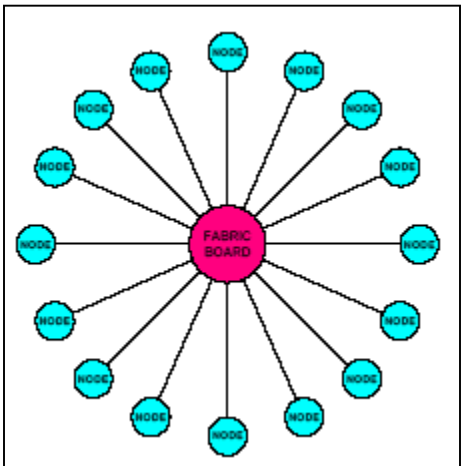

Figure 1: Switch Fabric Topology

Therefore two new types of Slots/Cards are supported: Fabric Slots/Cards and Node Slots/Cards. The Switch Fabric architecture supports point-to-point connections (or links) from every Node Card to centralized Fabric Card(s). System performance is not diminished with extra Nodes. The overall performance is scalable, as each Node Card can utilize the full bandwidth of its own point-to-point link. Plus, the availability of 2 Fabric Slots/Cards supports redundant links, increasing the system reliability. See Section 6 on Scalable Performance and Section 7 on System Reliability.

6. Scalable Performance

Point-to-Point Architecture

To maximize bandwidth, a point-to-point architecture is used, utilizing active switching elements to route data between many processing cards, called Node Cards. See Figure 1. These active switching elements are located on special cards, called Fabric Cards, which are located in special backplane slots, called Fabric Slots.

The two VITA 31.1 card types are:

- ✓ A central interface card to perform switching of Ethernet signals to and from Node Slots on the backplane as well as incoming and outgoing traffic (Fabric Card(s)). Note that the system may have one or two Fabric Cards. The Fabric Card can be a conventional unmodified PICMG 2.16 Fabric Switch Card.
- ✓ A distributed set of processing cards (Node Cards).

Since the VITA 31.1 specification is designed for packet based traffic (like PICMG 2.16), traffic can be routed from one Node Card to any other Node Card via a Fabric Card across a packet switching backplane.

There is a full-duplex Gigabit connection between each Node Card and each Fabric Card (two pairs for transmit, and two pairs for receive, as well as grounds.)

The VITA 31.1 specification is designed for a 21-slot chassis. Therefore, a VITA 31.1 chassis can support a maximum of two Fabric Cards and 19 Node Cards. Thus the total bandwidth for the system is up to 40 Gigabits/sec.

The Node Cards and Fabric Card(s) use auto-negotiation (as defined in IEEE 802.3) to select compatible rates between 10/100 and 1000 Mb/s Ethernet.

Each Node Card can connect to both Fabric Cards. The dual Fabric Cards can be used as redundancy or as additional resources. See Section 7.

Performance Analogy…

A switched Gigabit Ethernet infrastructure deployed on a VITA 31.1 (or PICMG 2.16) platform delivers roughly the equivalent of a 32-bit 33 MHz PCI bus segment to each slot in the backplane.

Across 19 Node slots the aggregate is in the order of 2 Gbytes/sec, an approximate 20-to-1 ratio when compared to a basic 32-bit 33 MHz PCI Bus.

In VME systems the gain is even more dramatic, compared with the 40 Megabyte aggregate bandwidth of a basic 32-bit VME bus.

7. System Reliability

Redundancy and Availability

To improve a system's reliability, and hence availability, the objective is to eliminate single points of failure. The VITA 31.1 specification is designed to carry significant amounts of traffic, and so reliability is a key issue. The VITA 31.1 backplane is designed to accommodate up to two Fabric Cards. Each Node Card can connect to the two Fabric Cards, in a dual star topology, with one Fabric Card used as a redundant system element. See Figure 2. If one Fabric Card fails, traffic can be routed through the secondary card without the system going down.

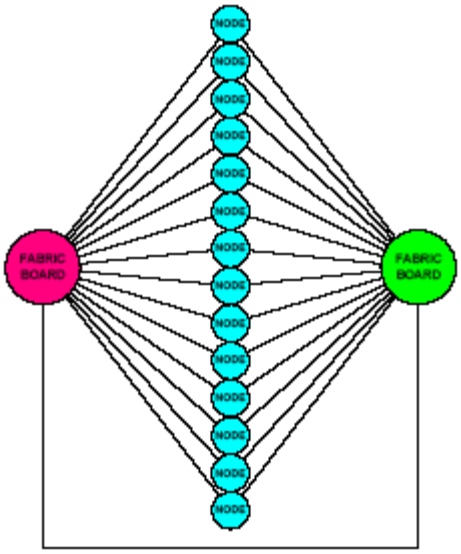

Figure 2: Dual Switch Fabric Topology

Some system designers will refer to the term "high availability". The full requirements of a high availability system are hot swap capabilities, fault tolerance and failover, redundant internal communication between system components and software to monitor and control redundancy. VITA 31.1, and therefore this document, does not cover all the requirements of high availability. VITA 31.1-2003 defines the pin assignments and interconnection methodology for implementing a 10/100/1000 Base-Tx Ethernet switched network on a VME64x backplane.

8. Product Offerings from Concurrent Technologies and Hybricon

Pentium M based VMEbus Single Board Computer supporting extensive I/O and VITA 31.1

- The VP 305/01x, a new addition to Concurrent Technologies VME SBC family
- This single slot board supports either the 1.6GHz or 1.1GHz Intel® Pentium M processor
- The VP 305/01x rich I/O feature set includes extensive front and rear I/O access that can easily be expanded with various I/O boards
- By supporting VITA 31.1, Concurrent Technologies has enabled a system designer, to reap the benefits of a packet-based switching architecture
- For harsher environments – an extended, -25°C to +70°C, operating temperature version is available

VITA 31.1 backplanes and chassis available now

- VITA 31.1 Backplanes: Standard version 21 slots; custom configurations 2-21 slots
- High performance, low noise backplane leverages Hybricon's PICMG 2.16 signal integrity analysis and implementation experience to achieve the highest performance
- VITA 31.1 backplanes are available in various types of Hybricon enclosures
- Hybricon's RME21 enclosure (show at left) supports cooling for 85 Watts/slot, with up to 1800W power supplies.
- By supporting VITA 31.1, Hybricon has enabled a system designer, to reap the benefits of a packet-based switching architecture