



HTF125 A

VITA 41.x - VXS: Adding
High Performance Switch
Fabrics to VME Based
Systems

Technology Focus

Executive Summary

This paper will provide an introduction and overview of the upcoming VITA 41.x - VXS Family of VMEbus Switched Serial Draft Standards. VXS is part of the VME Renaissance initiative, and maintains a level of compatibility with VME, VME64, and VME64x legacy products, as long as the VME64x P0/J0 connectors are not populated. This allows many VXS adopters to utilize a mix of new VXS boards along with legacy VME boards that are already commercially available from multiple vendors.

VXS 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.

The VXS draft standard is scheduled to go to ANSI ballot in the first quarter of 2005. This new specification will provide, for the first time, a standardized way to implement a high performance 10 Gbps switch fabric in a VME-style backplane without the use of cables, and using standardized backplane connector pinouts. This provides a clean, efficient way to implement multiprocessor architectures that utilize leading-edge high-speed serial interconnects such as InfiniBand and Serial RapidIO.

This is also the dawn of a new era for VME, as system architectures move towards serial point-to-point interconnections and switch fabric topologies to supplement, and eventually replace, parallel buses such as VME in the long term.

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 in a high availability system.

These limitations have prompted VITA (VME International Trade Association) to form the VXS working group to investigate a better method of transferring large amounts of data between cards, with an emphasis on leading edge high-speed serial switch fabric traffic.

Introduction to the VXS Backplane

The VXS draft standard has the following high-level features:

- Standards based architecture (i.e. InfiniBand, Serial RapidIO, Ethernet, PCI Express)
- Scalable system performance: Up to 720 Gbps/chassis (1 chassis)
- Scalable reliability: 1 or 2 fully independent fabrics
- Scalable cost: single fabric or dual fabric
- “Fine” granularity P2P architecture, single point of failure (SPF) = 1 slot
- “Virtual backplane” capable allowing for multi-chassis architectures
- Single backplane design can support all configurations and options

Key aspects of VXS include:

1. Point-to-point interconnections across the backplane
2. Uses standard switch fabric protocols, defined in various layered VXS standards:
 - o VITA 41.0 is the base VXS standard
 - o VITA 41.1 for InfiniBand Fabric
 - o VITA 41.2 for Serial RapidIO Fabric
 - o VITA 41.3 for Gigabit Ethernet Fabric
 - o VITA 41.4 for PCI Express Fabric
 - o VITA 41.10 for Live Insertion
 - o VITA 41.11 for Rear Transition I/O
 - o VITA 41.x for additional switch fabrics; more to come in the future...
3. Utilizes a switch fabric architecture

Switch Fabric Architecture

VXS defines a switch fabric architecture, as contrasted with a bus architecture (e.g. VME). A switch fabric is an interconnected network of switching devices. The fabric 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 slot to another, data passes through an active switching device; in a VXS system, these active switching devices are located on the switch card(s). The other cards that connect to the switch card(s) are called payload cards.

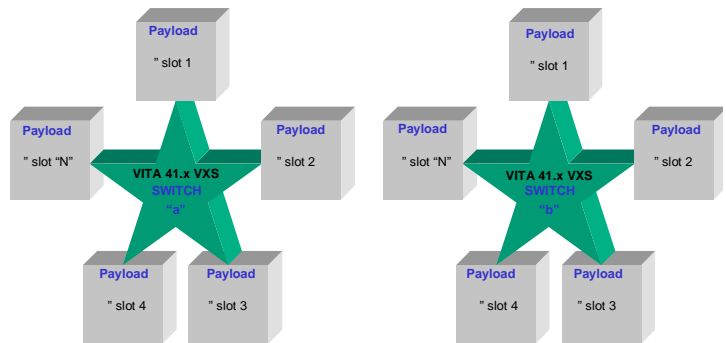
Two new types of slots/cards are supported: switch slots/cards and payload slots/cards. The switch fabric architecture supports connections from every payload card to centralized switch card(s). The availability of 2 switch slots/cards supports redundant connections for high availability applications.

In order to support the high-speed switch fabric signaling, a high-speed MultiGig RT connector is used in place of the VME64x 2MM HM type connector in the J0/P0 position on the payload slots/boards.

The switch slots/boards have a totally new connector complement, also based on high-speed MultiGig RT connectors, plus a dedicated

power connector. All slots also have new keying and alignment blocks.

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 switch fabric for data transfer. However, given the importance of legacy VME compatibility in many applications, it is unlikely that VXS backplanes without a VMEbus will actually be a popular option.



High availability systems strive to eliminate single points of failure. The VXS specification was designed to carry significant amounts of traffic, and so reliability and availability is a key issue. Each payload card can connect to two switch cards, with one switch card used as a redundant system element. If one switch card fails, traffic can be routed through the other one without the system going down.

VME Migration Path and Compatibility

VME hardware and software is available off-the-shelf. In many applications, this allows system vendors to improve Time-To-Market by re-using available VME hardware and software.

Because the VXS payload board/slot P0/J0 connector is different than the VME64x board/slot P0/J0 connector, some VME64x boards will not plug into a pure VXS backplane, but this can be easily handled with a hybrid VXS/VME64x backplane. Refer to the table below for a summary of compatibility between VXS and various styles of VME.

Compatibility Matrix for VXS with VME, VME64, and VME64x

Backplane Module	VME64 ANSI/VITA 1 (also applies to VME)	VME64x ANSI/VITA 1.1 without J0	VME64x ANSI/VITA 1.1 with J0	VXS VITA 41.x	Mixed Technology Backplane
Connectors:	J1 & J2: 3 Row DIN J0: None. Some backplanes have stiffening bars between J1 and J2.	J1 & J2: 5 Row DIN J0: None.	J1 & J2: 5 Row DIN J0: 2 mm HM	J1 & J2: 5 Row DIN J0: 7 row RT2 Switch slots: Use RT1 and RT2 data connectors, Positronic power connector	Application-specific mix of different slot types from VME64, VME64x, and VXS
VME64 ANSI/VITA 1 (also applies to VME)	COMPATIBLE	COMPATIBLE	COMPATIBLE	COMPATIBLE	<p>COMPATIBLE – DEDICATED SLOTS</p> <p>Requires a backplane with specific slots dedicated to VME64, VME64x, VXS board types. This mix is application specific.</p> <p>VME software, VMEbus electrical interface, and VMEbus protocol are compatible between board types.</p>
VME64x ANSI/VITA 1.1 without P0	PARTIALLY COMPATIBLE VME64x functions (3.3 volts, etc.) will not operate in a VME64 backplane.	COMPATIBLE	COMPATIBLE	COMPATIBLE	
VME64x ANSI/VITA 1.1 with P0	PARTIALLY COMPATIBLE P0 connector may interfere with stiffener bars on VME64 backplane between J1 and J2. VME64x added functions (3.3 volts, etc.) will not operate in a VME64 backplane.	COMPATIBLE	COMPATIBLE	A VME64x module with a P0 connector will not plug into a VXS backplane.	
VXS (VITA 41.x) payload board	PARTIALLY COMPATIBLE VXS payload board's P0 connector may interfere with stiffener bars on VME64 backplane between J1 and J2. VXS/VME64x added functions (3.3 volts, switch fabric, etc.) will not operate in a VME64 backplane.	PARTIALLY COMPATIBLE A VXS payload board will plug into a VME64x backplane without J0. VXS switch fabric will not operate in a VME64x backplane.	A VXS module (payload board or switch board) will not plug into a VME64x backplane with J0.	COMPATIBLE With Payload Slots	
VXS (VITA 41.x) switch board	A VXS Switch board will not plug into a VME64 backplane.	A VXS Switch board will not plug into a VME64x backplane.	A VXS Switch board will not plug into a VME64x backplane.	COMPATIBLE With Switch Slots	

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 payload cards. These active switching elements are located on special cards, called switch cards, which are located in special backplane slots, called switch slots. The VXS card types are:

- ✓ A central interface card to perform switching of serial fabric signals to and from payload slots on the backplane as well as incoming and outgoing traffic (switch card(s)). Note that the system may have one or two switch cards.
- ✓ A distributed set of processing cards (payload cards).

Since the VXS specification was designed for packet traffic, traffic can be routed from one card to any other card via a switch card.

The VXS specification was designed for a 20-slot backplane. A VXS backplane can support a maximum of two switch cards and 18 payload cards, but it is possible to add in another slot that is a standard VME64 or VME64x slot.

There is a full-duplex connection between each payload card and each switch card (four pairs for transmit, and four pairs for receive, as well as grounds.)

Each payload card can connect to both switch cards. The dual switch cards can be used as redundancy or as additional resources.

VXS Layered Switch Fabric Standards

- ✓ **VITA 41.1 InfiniBand**—Designed for ultra-high bandwidth for servers, with data rates of 16 Gbps per 4x link, full duplex. It can also be used for carrying a mix of traffic types including Voice and Video with Quality of Service (QOS) provisions. Infiniband operates

at 2.5 Gbps per pair. The VXS VITA 41.1 draft standard supports data rates of up to 16 Gbps per 4x link, full duplex. Infiniband silicon is commercially available now.

- ✓ **VITA 41.2 Serial RapidIO**—RapidIO has been around for a few years, and has been deployed in a Parallel RapidIO physical layer version, but Serial RapidIO is the new high-speed serial physical layer using the same RapidIO protocol. RapidIO is designed for communications inside a system rather than between systems. Serial RapidIO operates at 1.25 Gbps, 2.5 Gbps, or 3.125 Gbps per pair. The VXS VITA 41.2 draft standard supports data rates of up to 20 Gbps per 4x link, full duplex. Serial RapidIO silicon is available now.
- ✓ **VITA 41.3 Gigabit Ethernet**—Gigabit Ethernet has been around for a number of years, and low cost silicon is widely available. The VXS VITA 41.3 draft standard supports data rates of up to 8 Gbps per 4x link, full duplex. Ethernet is a popular choice for loosely coupled systems because of the maturity of the technology and the availability of standard software.
- ✓ **VITA 41.4 PCI Express**—PCI Express has been selected by the PCI-SIG as the next generation successor to the PCI bus. The PCI Express Base specification supports transparent PCI operation, and the Advanced Switching specification, based on StarFabric technology, supports a mix of traffic types including Voice and Video with Quality of Service (QOS) provisions. PCI Express operates at 3.125 Gbps per pair. The VXS VITA 41.4 draft standard supports data rates of up to 20 Gbps per 4x link, full duplex. PCI Express Base silicon is available now, with PCI Express Advanced Switching silicon to be available in early 2005.

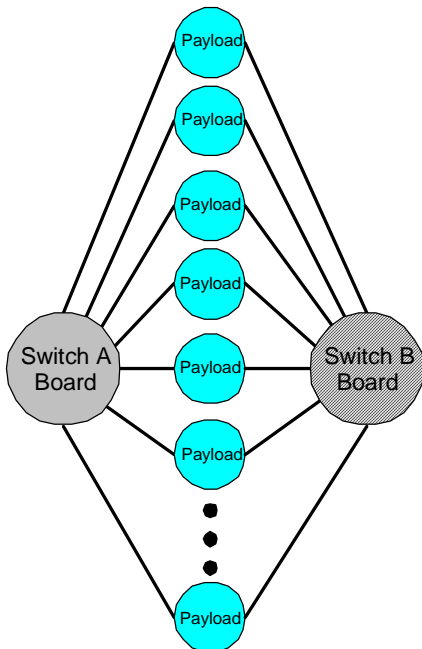
The high performance MultiGig RT connectors include additional shielding in order to provide electrical isolation between

the differential pairs of connections between the payload cards and the switch cards, and trace routing must be done carefully on the boards and backplanes in order to maintain signal integrity at these high speeds.

Each payload card can support up to a 20 Gigabit connection, full duplex, to each switch card. Thus the total bandwidth for the system is up to 720 Gbps, full duplex, for a system with 18 payload cards and 2 switch cards.

Switch Fabrics in General

There is no doubt that industry is moving toward switch fabric architectures to supplement, and eventually replace, traditional buses. 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 buses. The following table contrasts switch fabrics with VME and CompactPCI, and the figure contrasts bus structures with fabrics.



Parameter	VME, CompactPCI	Switch Fabric
Topology	Bused bi-directional multi-drop	Point-to-point
Types	VME – asynchronous CompactPCI – synchronous	Centralized – connect from each node slot to one or more dedicated fabric slots Mesh – connect from each slot to every other slot
Signal Type	Single ended signals with TTL-like levels	Differential signals typically LVDS
Data Width	32 – 64 bit address/data + control	1 – 16 bits
Clocking	VME – bused CompactPCI – point-to-point	Point-to-point connection, where clock is typically embedded in serial bit stream with 8b/10b encoding Clock is occasionally a separate point-to-point connection, one per N data connections
Data rate	VME – 10 ~ 320 Mbytes/s CompactPCI – 132 ~ 528 Mbytes/s	Ethernet – 0.02 ~ 20 Gbps per link Infiniband – 16 Gbps per 4X link Serial RapidIO – 4 ~ 20 Gbps per 4X link PCI Express – 20 Gbps per 4X link StarFabric – 4 Gbps per 4X link

All of these backplanes have their own niche. However, the VXS backplane has been designed to provide a VME family migration path into high-performance switch fabrics.

As a key part of the widely touted VME Renaissance initiative, VXS also has strong support from leading companies in the VME industry. This makes VXS very attractive for OEMs that build systems using VME64x today, but need higher interconnect bandwidth.

A word about VITA 41/VXS and VITA 46

The VITA Standards Organization is developing a new standard, VITA 46, scheduled for release in 2005. VITA 46 abandons the VME DIN connector in favor of a higher density connector system that supports the VMEbus protocol as well as switch fabrics and high speed rear I/O, but

does not provide legacy support in a pure VITA 46 backplane. Support for legacy VME and/or VXS boards will require custom mixed backplanes with selected slots supporting the VME and/or VXS boards, but the VMEbus signals can be interconnected between these various slot types.

When available, VITA 46 will provide advantages in some applications such as high performance multiprocessing and DSP. However, there will be a cost penalty for this extra performance.

VXS boards, backplane, and chassis are available now for users who need to move up to a switch fabric; for those who need the features provided by VITA 46, they will need to wait for a little while for the standards to finalize, and for companies to develop products.

Chassis Implications

Although the VXS standards do not require higher power consumption, VXS also cites the ANSI/VITA 1.7 standard which allows higher power supply currents, and higher power dissipation, on the VXS boards.

Vendors will likely develop VXS products that have higher power dissipation than typical VME64x boards, in order to fully utilize the high performance switch fabric capabilities of VXS. Also, the new family of VITA 42.x XMC Switched Mezzanine Card draft standards, designed to support switch fabrics on mezzanine modules, will allow higher power dissipation on mezzanines vs. existing PMC mezzanines.

The net result is that VXS boards, along with XMC mezzanines, will likely consume more power than existing VME64x boards with PMC mezzanines, thus requiring better-designed chassis that can adequately cool the higher power boards and mezzanines.

Why you should talk to Hybricon

Hybricon is a VITA/VSO member and a PICMG Executive member that is very active in many VITA/VSO and PICMG subcommittees and is very well known and respected within the VITA/VSO and PICMG communities. Our involvement in writing the specifications means that Hybricon is well versed in all approved VITA/VSO and PICMG specifications as well as draft specifications.

Hybricon has been providing high performance signal integrity simulation services to our customers for over 15 years, including many high-performance switch fabric backplane designs. This includes signal integrity work for both PICMG and VITA. This gives Hybricon a solid background in high performance backplane design. This is new technology; why trust your critical design to anyone else?

Hybricon's VME-based products cover a wide range of backplanes, power planes, card cages, rack mount chassis and desktop enclosures. While we do offer a broad range of standard products, Hybricon specializes in custom solutions. We offer a broad range of engineering services including sophisticated signal integrity analysis for custom buses or multi-gigabit serial links using 2D/3D field solvers and circuit simulators as well as sophisticated thermal/air flow analysis using CFD analysis tools. With higher power boards, robust thermal analysis and design is a necessity; this feeds into one of Hybricon's strengths.

For more information on our plans for VXS backplanes, or any other Hybricon product, please e-mail us at info@hybricon.com, visit us on the Web at www.hybricon.com, or call Sales at (978) 772-5422.

Product Links:

- o VXS Backplanes
http://www.hybricon.com/vme64x/vita_41.html
- o RME21 10U Enclosure with 85W/slot cooling capacity
<http://www.hybricon.com/coolslot/rme21.html>