Dell Ethernet Pass-Through Module Manual

File Name: Dell Ethernet Pass-Through Module Manual.pdf
Size: 3691 KB
Type: PDF, ePub, eBook
Category: Book
Uploaded: 25 May 2019, 14:14 PM
Rating: 4.6/5 from 658 votes.

Status: AVAILABLE

Last checked: 3 Minutes ago!

In order to read or download Dell Ethernet Pass-Through Module Manual ebook, you need to create a FREE account.

Download Now!

eBook includes PDF, ePub and Kindle version

Book Descriptions:

We have made it easy for you to find a PDF Ebooks without any digging. And by having access to our ebooks online or by storing it on your computer, you have convenient answers with Dell Ethernet Pass-Through Module Manual . To get started finding Dell Ethernet Pass-Through Module Manual , you are right to find our website which has a comprehensive collection of manuals listed.
Our library is the biggest of these that have literally hundreds of thousands of different products represented.

Book Descriptions:

Dell Ethernet Pass-Through Module Manual

This product has been tested and validated on Dell systems. It is supported by Dell Technical Support when used with a Dell system. Site Terms and Privacy Policy. Sep 23 29Our payment security system encrypts your information during transmission. We don’t share your credit card details with thirdparty sellers, and we don’t sell your information to others. Used Very GoodPlease try again.Please try again.In order to navigate out of this carousel please use your heading shortcut key to navigate to the next or previous heading. Please try your search again later.To calculate the overall star rating and percentage breakdown by star, we don’t use a simple average. Instead, our system considers things like how recent a review is and if the reviewer bought the item on Amazon. It also analyzes reviews to verify trustworthiness. Unfortunately, I have never seen one of these before, let alone configure one, so it is testing my patience. I have it on the network, I can access the web interface, and I have Windows Server 2016 installed on one of the blades. Within the blade I have a variety of mezzanine cards that match up to the configuration of the enclosure, so all good there. Basically I have a quad gigabit Ethernet card as well as a quad fiber card. Windows sees 8 NICs. However, when I plug the respective passthrough port on the enclosure into our switch, only ONE of the NICs lights up with activity. The other three show media unplugged. Why is this I have DHCP enabled everywhere in the web interface where DHCP can be enabled except for the enclosure address that is static. Since this ONE mezzanine card controls all four ports that Windows sees, and since theres only ONE Ethernet passthrough port on the enclosure for this blade, how do I get the other three Ethernet to activate. Another oddity. we have two vlans. One that is a 10.20.1.X range, and one that is a 10.20.2.X range. They are configured on our Cisco switch to communicate through IVR.http://sterenstein.ru/userfiles/indiglo-44250a-manual.xml

- **dell ethernet pass-through module manual, 1.0, dell ethernet pass-through module manual.**

Every computer on this switch can talk to every other computer, regardless of which vlan theyre on IVR works fine. However, I plug this M1000e into the switch on the 10.20.1.X vlan and it CANNOT talk to the 10.20.2.X vlan. This is an annoying problem if anyone has a clue whats going on I would appreciate it!!! I suspect the issue lies in some obscure configuration or change needed through the web interface, though some part of me thinks that the answer lies in a command line change that is needed on the CMC interface. I just dont know enough about this to figure out a solution, and the information online is sparse. ANY thoughts would be appreciated. Thanks in advance. Dave Each Fabric has 2 slots associated with it example B1 and B2. When you install a multiport network adapter in a blade and plug it into the chassis nic port 1 will map to B1 side and nic port 2 will map to B2. In your case you have a quad port nic so 2 ports will be presented to each fabric slot. But you have a 16 port direct passthrough module in B1 and it is not capable of providing 2 1gbe connections through a single external port its a direct onetoone passthrough and only has 16 internal facing connections one for each blade. And your B2 slot is empty so you are unable to connect anything to the other side 1gbe ports on your blade. You would need to move them to slots B1 and B2 then uplink them to your external rack switch using the 40gbe ports which have several options like a breakout cable from 40gbe to 4x10gbe assuming your rack switch has 10gbe capable ports.It would appear dell does as wellThats why I came here in hopes someone has actually done a configuration and knows how to sidestep some of these issues. Thanks though. Dave You mention when you plug in the respective passthrough port only 1 nic shows a connection. There is no single passthrough port that can map to 4 nic ports.I then installed the Gigabit card and 4 NICs showed up

in Windows.http://www.dean-cpa.com/files/adminpic/indico-user-manual.xml

I then installed the Fiber Channel card and 4 more NICs showed up in Windows. Currently we only have one ethernet connection from the passthrough port of the blade straight to to a gigabit ethernet switch in a rack. As mentioned, one of the NICs is now active and talking. None of the other enclosure modules are being used at this point. Basically, the only ports I can actually plug into are on the ethernet passthrough module, as well as the fiber channel module which I believe is also a passthrough module. Understanding that one passthrough port maps to one blade, how do I take advantage of the other 3 ethernet ports that are shown in Windows. Likewise, Im assuming that when I get the fiber plugged in for that blade, there will only be one fiber NIC active. If these cards supply 4 ports each giving me 8 total NICs shown in Windows, how do I break out these NICs for use. How do I configure this to work if each blade only maps to one passthrough port Ive gone through the manuals but without context they dont mean much. There is a lot of very productspecific terminology that I am unfamiliar with i.e. I had never heard of a mezzanine card before yesterday never in my life have I come across such a thing, likewise, I had no knowledge that they had to be installed a certain way in correlation to the enclosure. Our head networking guy didnt have a clue about the orientation of the mezzanine cards I found that gem on a forum. This is why Im coming to this community I know people here have had experience with these units and Im trying to learn based on the experience of other members. Thank you all in advance for any information you can pass along. Dave There is a section on page 37 that explains how the fabrics work. If you have further questions then you need to provide the model of the blades you are using. Thanks Each Fabric has 2 slots associated with it example B1 and B2.

You would need to move them to slots B1 and B2 then uplink them to your external rack switch using the 40gbe ports which have several options like a breakout cable from 40gbe to 4x10gbe assuming your rack switch has 10gbe capable ports.Firstly, if you look at the back of your chassis, those you named are switches or Passthrough ports. The m1000e would have 6 NIC fabrics A1 A2, B1, B2, C1, C2.to know the diff is that switches usually have only 2x or 4x 10Gbps OR 10x 1Gbps uplink per module while passthrough would have 16 NIC ports one per NIC. Now I have a path forward, thanks to both of you!! Regarding my other question, the other little annoyance.Just to recap, we have two VLANS that are connected through IVR, and every computer on this switch can talk between the two VLANS. However, this Dell unit is only talking on the VLAN of the particular IP address I give it i.e. if I give it a 10.20.1.X IP, it only talks on that VLAN, or if I give it a 10.20.2.X IP, it only talks on that VLAN it will not ping anything on the alternate VLAN, nor can any computer that happens to have an alternate IP as this Dell ping back into it. The switch is configured properly we have been using this setup for a number of years without issue. I know the command line interface has a lot of features and complexity, so Im just wondering if something additional needs to be setup on the server side to regain functionality. Logic dictates that it wouldnt need to, as the switch is managing the IVR, and with any other system we connect it just works. So Im a bit stumped on this as well. Once again, thank you for the insight and CONTEXT. Dave. It includes the following sections Blades servers are designed to reduce the space, power, and cooling requirements within the data center by providing these services within a single chassis.

https://labroclub.ru/blog/i-coo-targo-manual

Blade server systems are a key component of data center consolidation that help reduce costs and provide a platform for improving virtualization, automation, and provisioning capabilities. Both copper and optical passthrough modules are available that provide access to the blade server controllers. It is therefore important to understand the internal connectivity provided by the blade server chassis before discussing the external ramifications of passthrough deployments. Various blade system architectures are available from various vendors. The following section describes two generic blade server systems, which illustrate many of the design features found in these various

architectures Figure 34 shows this design, which supports a total of 32 independent channels for the eight blade server slots. This octopus cable allows multiple servers to be supported by a single output cable that connects to the external network with transmit and receive pairs dedicated to each blade server controller. Instead, it simply passes the blade server signaling to the external network porttoport. It includes the following topics The network infrastructure provides the level of availability required by these applications through device and link redundancy and a deterministic topology. Servers are typically configured with multiple NIC cards and dualhomed to the access layer switches to provide backup connectivity to the business applications. Implementing blade servers with passthrough technology allows nondisruptive deployment. Passthrough deployments do not alter the fast convergence and deterministic traffic patterns provided by Layer 2 and Layer 3 technologies. The connection established between the external network device and the blade server by the passthrough module is neither switched nor blocked. The modules simply expose the blade server NICs to the network without affecting the Layer 2 or Layer 3 network topologies created through spanning tree or routing protocols.

When using passthrough modules, the blade servers function as servers that happen to be located inside a blade server chassis. Each blade server achieves an increased level of accessibility by using NIC teaming software. NIC teaming lets you create a virtual adapter consisting of one to eight physical NIC interfaces, which can typically support up to 64 VLANs. NIC teaming is a high availability mechanism that can provide both failover and local load balancing services. There are the following three primary modes of NIC teaming The standby interface becomes active if the primary NIC fails because of probe or physical link problems. Fault tolerance can be achieved in the following two ways NFT provides a greater level of network availability. However, neither configuration optimizes the bandwidth available to the blade server. This feature lets the server utilize more of the available bandwidth. Passthrough technology permits this configuration, as shown in Figure 37. The dotted green lines are standby interfaces that only transmit traffic. A hashing algorithm, usually based on the source and destination IP addresses, determines which NIC is responsible for transmitting the traffic for any given transaction. The standby controller becomes responsible for both ingress and egress traffic only if the primary NIC fails. This requires the switch to load balance traffic across the ports connected to the server NIC team. This mode of NIC teaming provides link redundancy and the greatest bandwidth utilization. However, the access switch in this design represents a single point of failure. Figure 38 shows a channel established between the blade servers and the access switch. Scalability allows increases in services or servers without requiring fundamental modifications to the data center infrastructure.The access layer should provide the port density to support the data center servers and the flexibility to adapt to increased demands for bandwidth or server capacity.

For more information on data center scalability, see Design Goals, or the Cisco Data Center Infrastructure SRND at the following URL. From this point of view, blade server passthrough modules reduce the administrative complexity of the data center. Passthrough modules do not require configuration, which eliminates configuration errors on the devices and reduces the need for configuration backups. The blade server chassis may also provide limited diagnostic information and the ability to enable or disable external ports on the module. The availability of these features and the level of diagnostic information depends on the manufacturer of the blade system. However, SOL requires the use of an integrated switch and is not currently available with the IBM passthrough modules. SOL leverages the trunk that exists between the management module and the Ethernet blade switch to allow console access to the individual blade servers. The configuration examples use Cisco Catalyst 6500s as the aggregation layer platform.The modular access provides port density

and 10 Gigabit Ethernet uplinks to the aggregation layer where intelligent services such as security, load balancing, and network analysis reside. There is no single point of failure in this network topology. The access layer switches are dualhomed to the aggregation layer switches, which provide redundant network paths. Spanning tree manages the physical loops created by the uplinks between the aggregation and access switches, assuring a predictable and fast converging topology. This allows the server farm to scale as it addresses future data center needs without an exponential increase in administrative overhead see Figure 310 . Each pair of access switches supports 12 fully populated blade server systems housed in a set of three racks. In this example, each blade system requires 32 ports on the access switch for the passthrough links from the blade servers.

Dualhoming the blade servers to each modular access switch means that each access layer switch must provide 64 ports per rack or 192 ports for three racks.In addition, the acceptable oversubscription ratio for the applications must be taken into account. For more information on scaling the Layer 2 and Layer 3 topologies in the data center, see the Data Center Infrastructure SRND at the following URL. In the previous example, 120 integrated blade switches would be required, assuming two blade switches per chassis, to support an equivalent number of blade servers. The ramifications of introducing this number of devices into the data center network are obvious. Specifically, the Layer 2 and Layer 3 topologies expand and become more difficult to manage. In addition, there is an increase in the network administration required to support the integrated switch. Using a modular access layer switch reduces these operational and logical complexities. Cabling represents an obstruction that restricts the airflow within the data center and may adversely affect the temperature of the room. When using passthrough technology and blade servers, the design and use of an effective cable management system within the facility is necessary to mitigate these issues. The 1RU access layer switches provide the port density and uplink connectivity to the aggregation layer required by the blade servers. This design allows the cable density created with the passthrough modules to remain within the rack, which helps contain the potential problems. This is a distinct advantage compared to the modular access layer model discussed previously. The uplinks from the 1RU access layer switches provide a common, consolidated connection to the aggregation layer and its services. This essentially reduces the number of cable runs required between the rack and the aggregation layer switches. In Figure 311, Aggregation1 is the primary root switch with all links forwarding.

Aggregation2 is the secondary root and provides an alternative traffic path for application traffic in the event of a failure. The solid black lines represent links that are forwarding and the dotted red lines represent links in a spanning tree blocking state. The sold yellow lines represent each link to the active NIC while the dotted green lines show the links to each interface in a standby state. In this configuration, the NIC teaming software offers subsecond convergence at the server. Redundant power and hotswappable fans are recommended to improve Layer 1 availability. A 1RU switch cannot provide the same level of port density for server connectivity as a modular switch. As a result, the number of switching devices in the data center is increased, compared to the solution using modular switches. This in turn increases the spanning tree domain as well as the administrative overhead required to implement and maintain the solution. In this instance, each data center rack houses three blade systems providing 32 individual passthrough connections to the internal blade servers. The blade servers are dualhomed over these passthrough connections to a pair of 1RU access switches located in the rack. Three blade systems with 16 dualhomed servers per chassis require 96 ports. To provide for network fault tolerance, each 1RU access layer rack switch should supply 48 ports for server connectivity. The modular aggregation layer switch must furnish the uplink port density for the 1RU access layer. A Catalyst 6509 would suffice in this scenario. A pair of aggregation layer switches can support 12 1RU access switches that are dualhomed over ten Gigabit Ethernet uplinks. In this example, 288 servers are supported in the 1RU access switch model with modular aggregation layer support. In addition, the acceptable oversubscription ratio for the

applications must be taken into account.

For more information on scaling the Layer 2 and Layer 3 topologies in the data center, see the Data Center Infrastructure SRND 2.0 at the following URL. This affects the scalability of this design by requiring greater uplink port density to connect to the aggregation layer switches. The example shown in Figure 312 would require 36 integrated blade switches, assuming two blade switches per chassis, to support an equivalent number of blade servers. In addition, the design reduces the number of cables required to provide external connectivity to the rack. As previously discussed, the modular access layer switch design requires fewer network devices and topology changes but uses more cabling. The following configurations are described To allow for more granular STP calculations, enable the use of a 32bit value instead of the default 16bit value. The longer path cost better reflects changes in the speed of channels and allows STP to optimize the network in the presence of loops. The following two types of interswitch connections can be used to provide this connectivity Each of these pointtopoint links between the switches is a trunk because they typically carry more than one VLAN. Each aggregate switch uses this feature to create a port channel across the line cards. The use of multiple line cards within a single switch reduces the possibility of the pointtopoint port channel becoming a single point of failure in the network. The following is an example of the interface configuration between the aggregate and access layer switch with Root Guard enabled In other words, a trunking NIC should be attached to a trunking switch port. Enable PortFast for the edge devices in the spanning tree domain to expedite convergence times. This feature protects the STP topology by preventing the blade server from receiving BPDUs. A port disabled using BPDU Guard must be recovered by an administrator manually. Enable BPDU Guard on all server ports that should not receive BPDUs.

The commands required to enable this feature are as follows To configure Port Security, configure the maximum number of MAC addresses expected on the port. The NIC teaming driver configuration the use of a virtual MAC address must be considered when configuring Port Security. This device may be a firewall, a load balancer, or a router. Using a redundancy protocol, such as HSRP, protects the gateway from becoming a single point of failure and improves data center network availability. HSRP allows the two aggregate switches to act as a single virtual router by sharing a common MAC and IP address between them. To enable HSRP, define a Switched VLAN Interfaces SVI on each aggregate switch and use the HSRP address as the default gateway of the server farm. The priority command helps to select this router as the active router by assigning it a higher value. If you continue browsing the site, you agree to the use of cookies on this website. See our User Agreement and Privacy Policy.If you continue browsing the site, you agree to the use of cookies on this website. See our Privacy Policy and User Agreement for details.If you wish to opt out, please close your SlideShare account. Learn more. You can change your ad preferences anytime. Why not share! Bachelor of Commerce Honours in Information Systems Parte Practica.Now customize the name of a clipboard to store your clips. It is also possible to connect a virtual KVM switch to have access to the mainconsole of each installed server.The blade servers, although following the traditional naming strategy e.g. M520, M620 only blades supported are not interchangeable between the VRTX and the M1000e.Current versions of the enclosure come with midplane 1.1 and it is possible to upgrade the midplane.Next to this is a powerbutton with powerindication. Basic status and configuration information is available via this display.

To operate the display one can pull it towards one and tilt it for optimal view and access to the navigation button. For quick status checks, an indicator light sits alongside the LCD display and is always visible, with a blue LED indicating normal operation and an orange LED indicating a problem of some kind.The rearside is divided in 3 sections top here one insert the 3 managementmodules one or two CMC modules and an optional i KVM module. At the bottom of the enclosure there are 6 bays for powersupply units. A standard M1000e operates with three PSUsThere are also older blades like

the M605, M805 and M905 series.An M420 server only supports a single Mezzanine card Mezzanine B OR Mezzanine C depending on their location whereas all halfheight and fullheight systems support two Mezzanine cards.RAM memory options via 12 DIMM slots for up to 192 Gb RAM DDR3. A maximum of two onblade hotpluggable 2.5inch harddisks or SSDs and a choice of builtin NICs for Ethernet or converged network adapter CNA, Fibre Channel or InfiniBand.Two external and one internal USB ports and two SD card slots. The blades can come preinstalled with Windows 2008 R2 SP1, Windows 2012 R2, SuSE Linux Enterprise or RHEL.Supported on both the M1000e and PowerEdge VRTX chassis. The server uses iDRAC 9.CPU can be two quadcore or 6core Xeon 5500 or 5600 with the Intel 5520 chipset.Memory via 32 DDR3 DIMM slots offering up to 512Gb RAM. Onboard up to two 2,5 inch HDD or SSDs. The blade comes with a choice of onboard NICs and up to two mezzanine cards for dualport 10Gb Ethernet, dualport FCoE, dualport 8Gb fibrechannel or dual port Mellanox Infiniband.These on board NICs connect to a switch or passthrough module inserted in the A1 or the A2 bay at the back of the switch. For redundancy one would normally install switches in pairs the switch in bay A2 is normally the same as the A1 switch and connects the blades onmotherboard NICs to connect to the data or storage network.

The same applies to adding a Fibre Channel host bus adapter or a Fibre Channel over Ethernet FCoE converged network adapter interface.It is also possible to use completely diskless blades that boot via PXE or external storage. But regardless of the local and bootstorage the majority of the data used by blades will be stored on SAN or NAS external from the bladeenclosure.The management of the SAN goes via the chassismanagement interface CMC.Also the Mseries cant be running outside the enclosure it will only work when inserted in the enclosure.The M6348 can be stacked with other M6348 but also with the PCT7000 series rackswitches.Although this new stackingoption is also introduced in the same firmware release for the PCT8024 and PCT8024f one cant stack blade PCM and rack PCTversions in a single stack.On the linklevel PCM switches support link aggregation both static LAGs as well as LACP. As all PowerConnect switches the switches are running RSTP as Spanning Tree Protocol, but it is also possible to run MSTP or Multiple Spanning Tree. Another feature is to use linkdependency. One can, for example, configure the switch that all internal ports to the blades are shut down when the switch gets isolated because it loses its uplink to the rest of the network. All PCM switches can be configured as pure layer2 switches or they can be configured to do all routing both routing between the configured VLANs as external routing. Besides static routes the switches also support OSPF and RIP routing.All ethernet extension modules for the MXL can also be used for the rack based N4000 series fka Power connector 8100.The MXL can either forward the FCoE traffic to an upstream switch or, using a 4 port 8Gb FC module, perform the FCF function, connecting the MXL to a full FC switch or directly to a FC SAN.Cisco offers a range of switches for bladesystems from the main vendors.One can stack up to 8 Catalyst 3130 switches to behave like one single switch.

This can simplify the management of the switches and simplify the spanning tree topology as the combined switches are just one switch for spanning tree considerations. The 3130 switches come standard with IP Base IOS offering all layer 2 and the basic layer 3 or routingcapabilities. Users can upgrade this basic license to IP Services or IP Advanced services adding additional routing capabilities such as EIGRP, OSPF or BGP4 routing protocols, IPv6 routing and hardware based unicast and multicast routing.Such FEXs were already available for HP and Fujitsu blade systems, and now there is also a FEX for the M1000e blade system.The PCM8428 is the only full Fibre Channel over Ethernet capable switch for the M1000e enclosure that offers 16 x enhanced Ethernet 10Gb internal interfaces, 8 x 10Gb enhanced Ethernet external ports and also up to four 8Gb Fibre Channel interfaces to connect directly to a FC SAN controller or central Fibre Channel switch.It uses either the B or C fabrics to connect the Fibre Channel mezzanine card in the blades to the FC based storage infrastructure. The M5424 offers 16 internal ports connecting to the FC Mezzanine cards in the bladeservers and 8 external ports. From factory only the first two external ports 17 and 18 are

licensed additional connections require extra Dynamic Ports On Demand DPOD licenses. The switch runs on a PowerPC 440EPX processor at 667 MHz and 512 MB DDR2 RAM system memory.Standard license offers 12 connections which can be increased by 12 to support all 24 ports.There is the SFS M7000e InfiniBand switch from Cisco.Unlike the M4001 switches where the external ports are using QSFP ports for fiber transceivers, the 2401 has CX4 copper cable interfaces.For example if only a few of the bladeservers do use fibrechannel storage one dont need a fully manageble FC switch one just want to be able to connect the internal FC interface of the blade directly to ones existing FC infrastructure.

A passthrough module has only very limited management capabilities. Other reasons to choose for passthrough instead of enclosure switches could be the wish to have all switching done on a one vendor infrastructure; and if that isnt available as an M1000e module thus not one of the switches from Dell Powerconnect, Dell Force10 or Cisco one could go for passthrough modulesThe M1000e offers out of band management a dedicated VLAN or even physical LAN for management. One would normally connect the Ethernet links on the CMC avoiding a switch in the enclosure. Often a physically isolated LAN is created for management allowing management access to all enclosures even when the entire infrastructure is down. Each M1000e chassis can hold two CMC modules.It is also possible to access the enclosure management via a serial port for CLI access or using a local keyboard, mouse and monitor via the iKVM switch. It is possible to daisychain several M1000e enclosures.To access the management system one must open the CMC Webgui via https using the out of band management IP address of the CMC. When the enclosure is in stand alone mode one will get a general overview of the entire system the webgui gives one an overview how the system looks in reality, including the statusleds etc. By default the Ethernet interface of a CMC card will get an address from a DHCP server but it is also possible to configure an IPv4 or IPv6 address via the LED display at the front of the chassis. Once the IP address is set or known the operator can access the webgui using the default rootaccount that is built in from factory.Via the CMC interface one can configure blades in the system and configuring iDRAC access to those servers. Once enabled one can access the iDRAC and with that the console of the server via this webgui or directly opening the webgui of the iDRAC.