IBM® PureData™ System for Transactions

Best practices

Integrating IBM PureData System for Transactions into a data center network

Charles Lai
Software Developer, Core Engineering Team

Philip Cho
Software Developer, Core Engineering Team

Mike Law
Test/Network Architect

Ian Hakes
Senior Technical Writer

Issued: May 2014

Table of Contents

Executive summary

IBM PureData System for Transactions is an expertly integrated system to help deploy optimized and continuously available mission-critical databases.

This paper covers the network requirements for integrating PureData System for Transactions into a data center environment and provides popular sample network topologies.

Using the graphical interfaces in the system, you can easily configure the network switches to integrate with a corporate network.

Network environment in IBM PureData System for Transactions

This section provides and introduction to basic networking concepts and describes the networking environment in IBM PureData System for Transactions.

Basic networking concepts

To integrate PureData System for Transactions with your data center, you need an understanding of basic networking concepts. The following concepts are used throughout this paper:

- Virtual Local Area Network (VLAN): A VLAN is an isolated broadcast domain that is created within a physical switch. Traffic passes from one VLAN to another only if you add a routing device.

- Port VLAN identifier (PVID): A PVID is a default VLAN identifier that is assigned to an access port. VLAN segments designate a PVID to ports that are connected to the VLAN.

- No-tagging (access) mode: In this mode, raw Ethernet frames are sent without any VLAN tagging information in the frame header. Only one VLAN is supported, and the PVID matches the VLAN for the port.

- Tagging (trunk) mode: In this mode, a VLAN identifier is placed in the frame header of a packet when it is sent. Tagged packets can pass traffic across multiple VLANs on a single port.

- Link aggregation: Link aggregation combines multiple network links into a single link to increase bandwidth and to improve redundancy. Examples of link aggregation methods are EtherChannel, which is used for static grouping, and Link Aggregation Control Protocol (LACP), which is used for dynamic grouping.

- Virtual Link Aggregation Group (VLAG): A VLAG allows for multichassis link aggregation and facilitates active-active uplinks of access layer switches. Using a VLAG with the Spanning Tree Protocol often helps avoid the wasted bandwidth that is associated with links that are blocked by the protocol.

- Spanning Tree Protocol: The Spanning Tree Protocol can prevent broadcast loops when multiple paths exist between two points on a network. One example of the protocol is Per VLAN Rapid Spanning Tree (PVRST).

- Layer 2: This data link layer is considered the switching layer and is processed by hardware. Layer 2 activities occur in the same broadcast domain or local network. Layer 2 finds adjacent partners by using MAC addresses.

- Layer 3: This network layer is considered the routing layer and is processed by software. This layer understands how to traverse multiple networks. Layer 3 finds adjacent partners by using IP addresses.

- Domain name server (DNS): A DNS translates domain names to IP addresses and vice versa. The DNS serves all devices that are connected to a data center's network.

- Network Time Protocol (NTP): The Network Time Protocol is used by system software to synchronize the date and time for all devices that are connected to a data center's network.

- Gateway server: A gateway server is a router or proxy server that routes data between a private network and the external network that is connected to the data center.

Networking hardware in PureData System for Transactions

PureData System for Transactions comes with a pair of IBM System Networking RackSwitch G8264 switches that act as the top-of-the-rack (TOR) switches. This pair is the main entry point to PureData System for Transactions. The G8264 switches are stacked by using four 40 Gb inter-switch links (ISLs). Forty-six ports are reserved for your data network, and two ports are reserved for your management network.

The system rack has either one IBM Flex chassis (for a small or medium system) or two IBM Flex chassis (for a large system). Each chassis has a pair of redundant IBM Flex System Fabric EN4093 10 Gb scalable switches. Between each pair of EN4093 switches and TOR switches is a aggregated link, which is a connection that is composed of eight aggregated 10 Gb/s links.

The IBM Flex System x240 compute nodes have an IBM Flex System EN4132 two-port 10 Gb Ethernet adapter. The IBM PureSystems Manager nodes have an IBM Flex System CN4054 four-port 10 Gb virtual fabric adapter.

Figure 1 depicts the networking hardware topology of a small system that uses PureData System for Transactions. For more information about the hardware specification of each PureData System for Transaction offering, see *Planning for the system: Components of the system* in the Information Center.

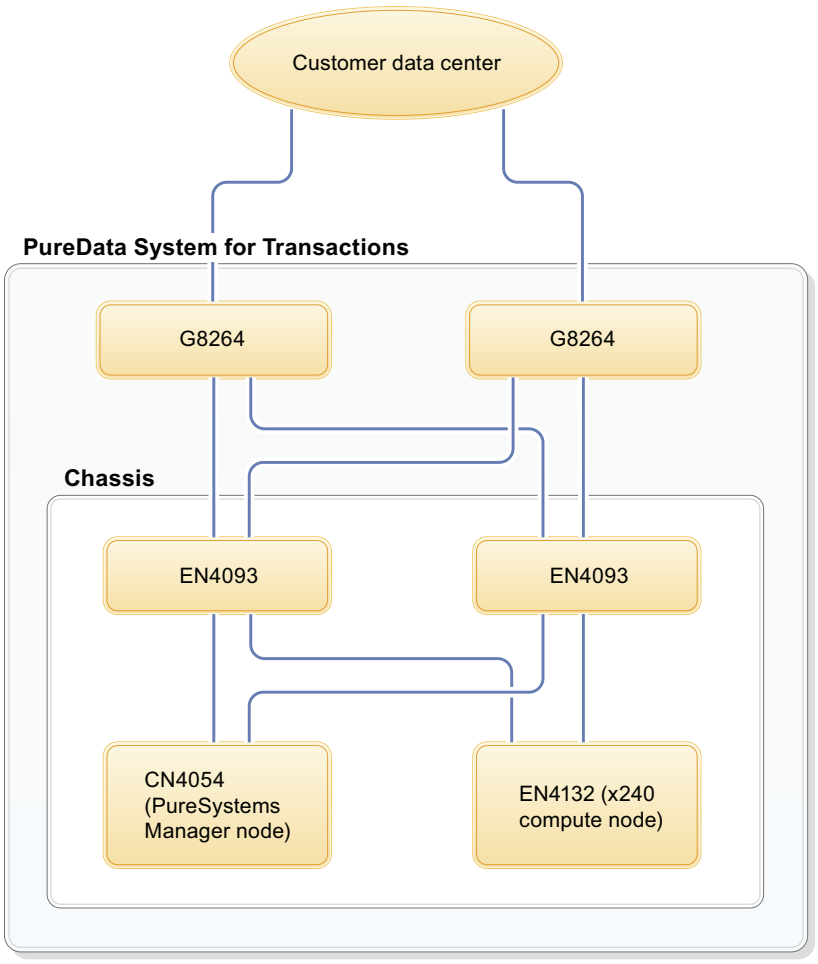

Figure 1. Networking hardware in PureData System for Transactions.

VLANs in PureData System for Transactions

The PureData System for Transactions network is composed of nine different VLANs. As shown in Table 1, there are seven internal VLANs and two external VLANs for customer data and management.

Table 1. PureData System for Transactions VLANs.

VLAN name	VLAN ID	Description
Console	3201	Used by the system console for communication between PureSystems Manager and compute nodes.
Hardware management	4091	Used by the system console and service laptop for communication between PureSystems Manager, the service laptop, and hardware components (IBM Flex System Chassis

VLAN name	VLAN ID	Description
		Management Module, IBM Integrated Management Module, IBM Storwize V7000 Disk System, and switches).
VLAG health check	4093	Used by TOR switches to ensure that the inter-switch link is healthy.
Top-of-the-rack inter-switch link (TOR ISL)	4094	Used to configure the VLAG on TOR switches.
DB2 RoCE	403	Used by the DB2 software for RDMA communication between a CF and members.
DB2 internal TCP	3105	Used by the DB2 software for TCP/IP traffic communication.
Cloud management	User-defined	Used by the workload console for communication between PureSystems Manager and a deployed instance. This VLAN is internal to the rack only.
Customer: Data	User-defined	Used by customers to connect client applications to deployed instances on compute nodes. This VLAN is unique from the management VLAN.
Customer: Management	User-defined	Used by customers to connect to the system console and workload console on PureSystems Manager. This VLAN is unique from the data VLAN.

Cloud groups and IP groups in PureData System for Transactions

Like IBM PureApplication System, PureData System for Transactions uses the concepts of a cloud group and an IP group. An IP group is a set of IP addresses that deployed instances use for their data network. A cloud group is a collection of compute nodes and its associated IP group.

However, unlike in PureApplication System, all compute nodes in PureData System for Transactions belong to one cloud group, and all deployed instances belong to one IP group. The relationship between a cloud group and an IP group is depicted in Figure 2.

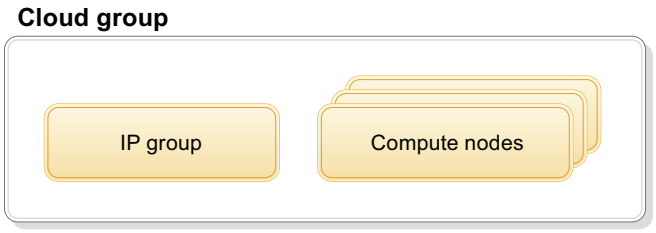

Figure 2. Relationship between a cloud group and an IP group.

For more information about cloud groups and IP groups, see *Managing application runtime environments in IBM PureApplication System* (http://www.ibm.com/developerworks/websphere/library/techarticles/1210_woolf/1210_woolf.html).

Sample network topologies

The following sample network topologies are used on customer systems.

Simple setup with stacked core switches

This setup is one of the most common. It is a highly available configuration that protects against a single point of failure in any individual port, cable, or switch. The TOR switches in PureData System for Transactions connect to a pair of core switches on the data and management networks that are linked by using stacking, a virtual switching system (VSS), or a virtual port channel (VPC). The aggregated link is formed from port 41 on each TOR switch to one port on each core switch in the customer data network. The management network has a similar aggregated link by using port 64 on the TOR switch.

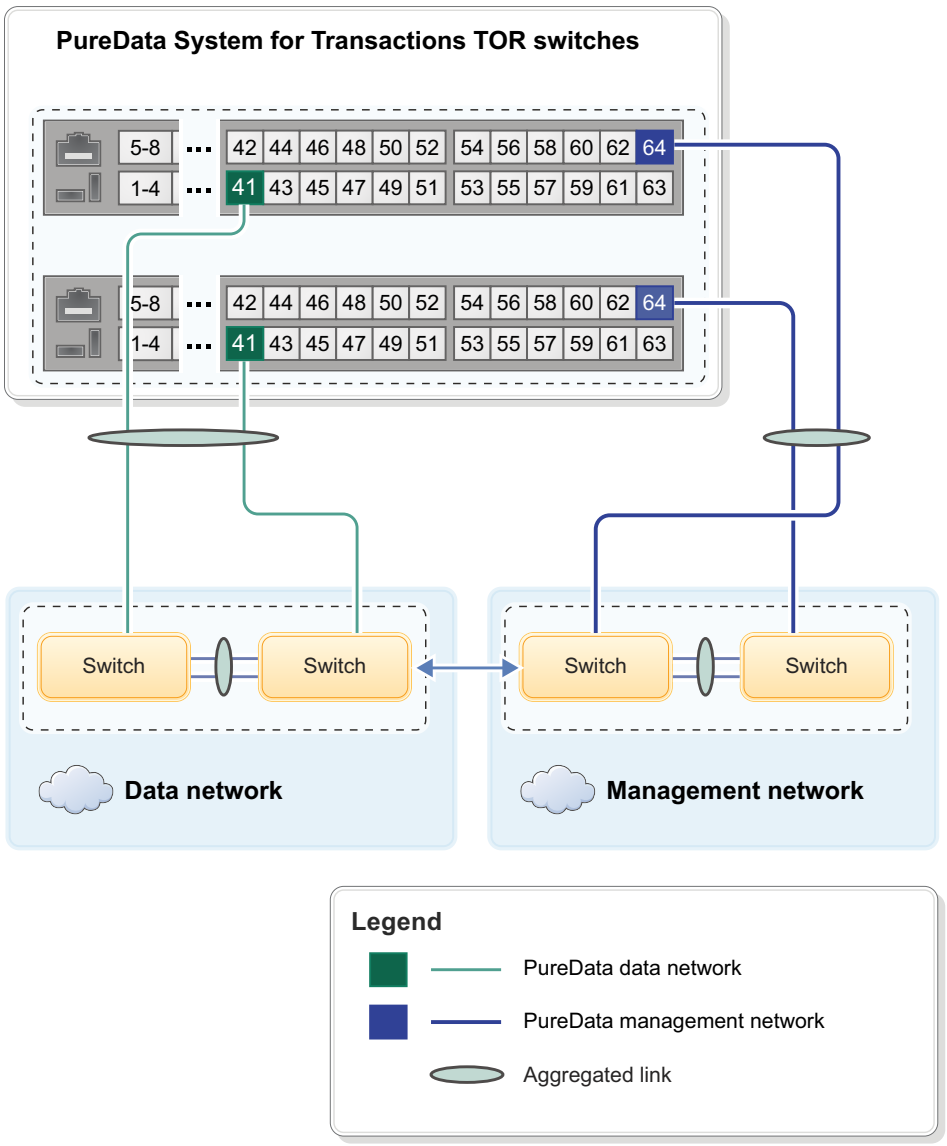

Figure 3. Simple setup with stacked core switches.

The aggregation method can be LACP or EtherChannel.

PureData System for Transactions supports Layer 2 switching only. If you need routing, you must connect the TOR switches to a router outside of PureData System for Transactions.

The customer management network needs access to the data network so that the administrator can access the database performance monitor UI from the workload console. Also, one or more NTP servers must be available on the customer management network, but they are not needed on the customer data network. In addition, one or more domain name servers must be available on both networks.

Redundant connections for high bandwidth requirement

If you have high network bandwidth requirements, create an aggregated link that has more than one port on each switch. Each port can be 1 Gb/s or 10 Gb/s. PureData System for Transactions includes transceivers of both types. In the example in Figure 4, four 10Gb/s ports are combined to form a single 40Gb/s connection.

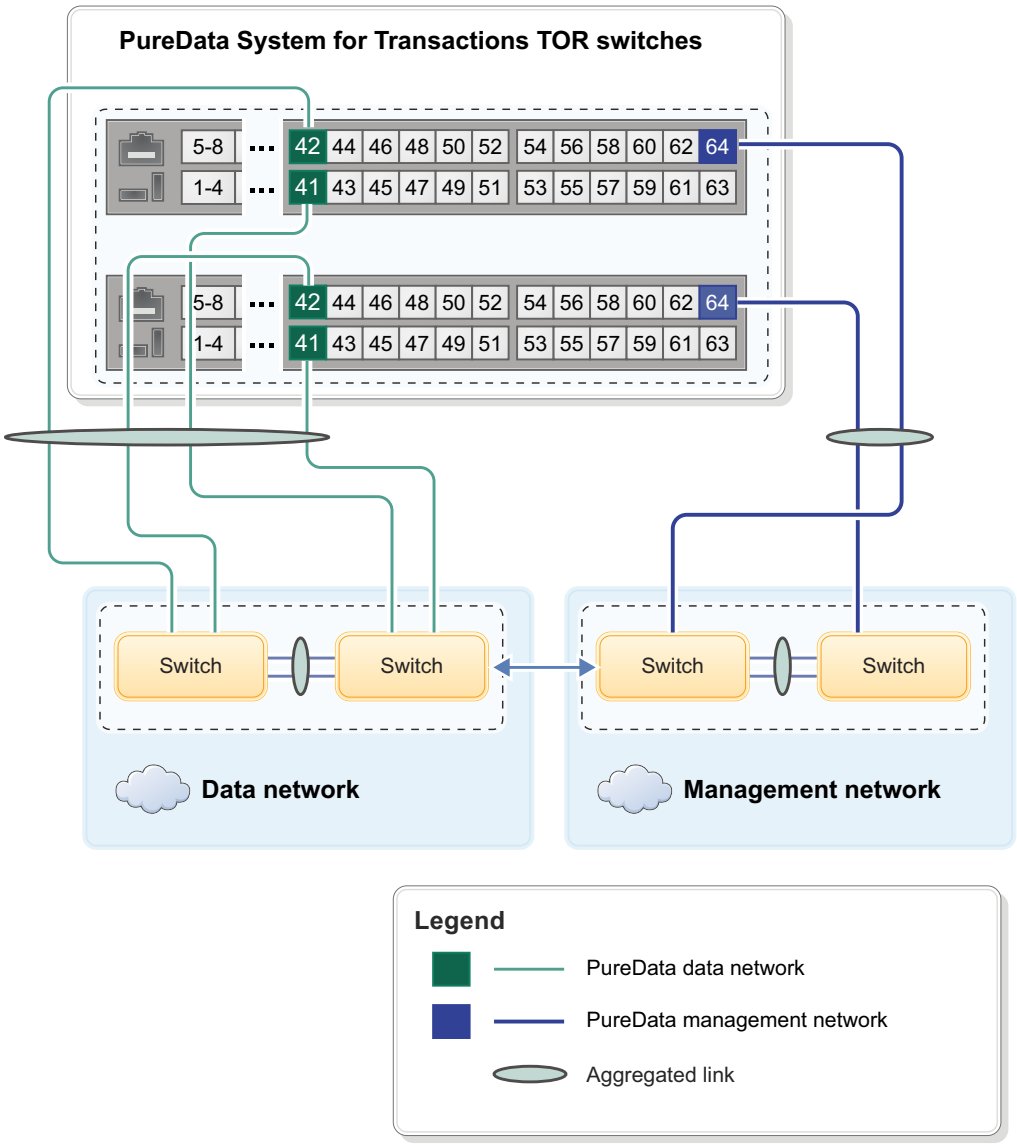

Figure 4. Redundant connections for high bandwidth requirement.

Using the default number of transceivers that are included with a rack, an extra-small or small system can deliver up to 60 Gb/s of bandwidth by using

a single aggregated link, a medium system can deliver up to 120 Gb/s of bandwidth, and a large system can deliver up to 240 Gb/s of bandwidth.

This cross-connection that is shown in the diagram allows the two TOR switches to remain highly available even if one of the two core switches fails. The cross-connection also allows the two core switches to remain highly available even if one of the two TOR switches fails.

Multiple VLANs using outside router

If you have a dedicated backup network that is separate from the data network, a router must exist outside PureData System for Transactions to direct the traffic between the two networks. Because PureData System for Transactions supports only one IP group, the deployed instances have only one external network interface. Both the data network and the backup network must be routable to the instances.

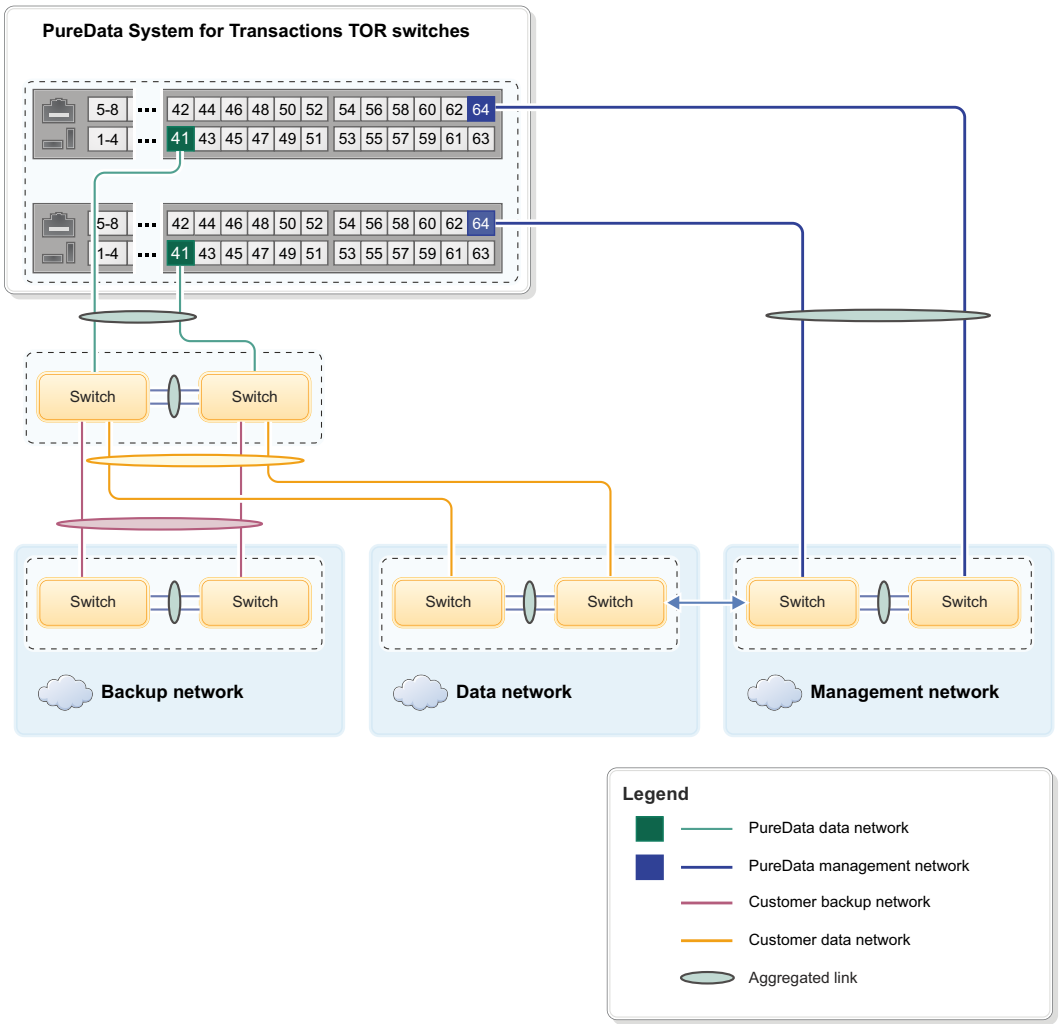

Figure 5. Multiple VLANs using outside router.

You can apply a similar strategy if there are more than two networks, for example, due to the addition of a dedicated network for disaster recovery. Use of Quality of Service (QoS) control is recommended if you have a strict bandwidth requirement for a certain network. You must implement QoS control outside PureData System for Transactions.

Non-stacked core switches with spanning tree

If you have a pair of core switches that you did not link together by using stacking, a VSS, or a VPC, you must create two aggregated links: one that is connected to the primary core switch and one that is connected to the secondary core switch. To prevent network loops, you must enable the Spanning Tree Protocol to ensure only one switch is active at a time.

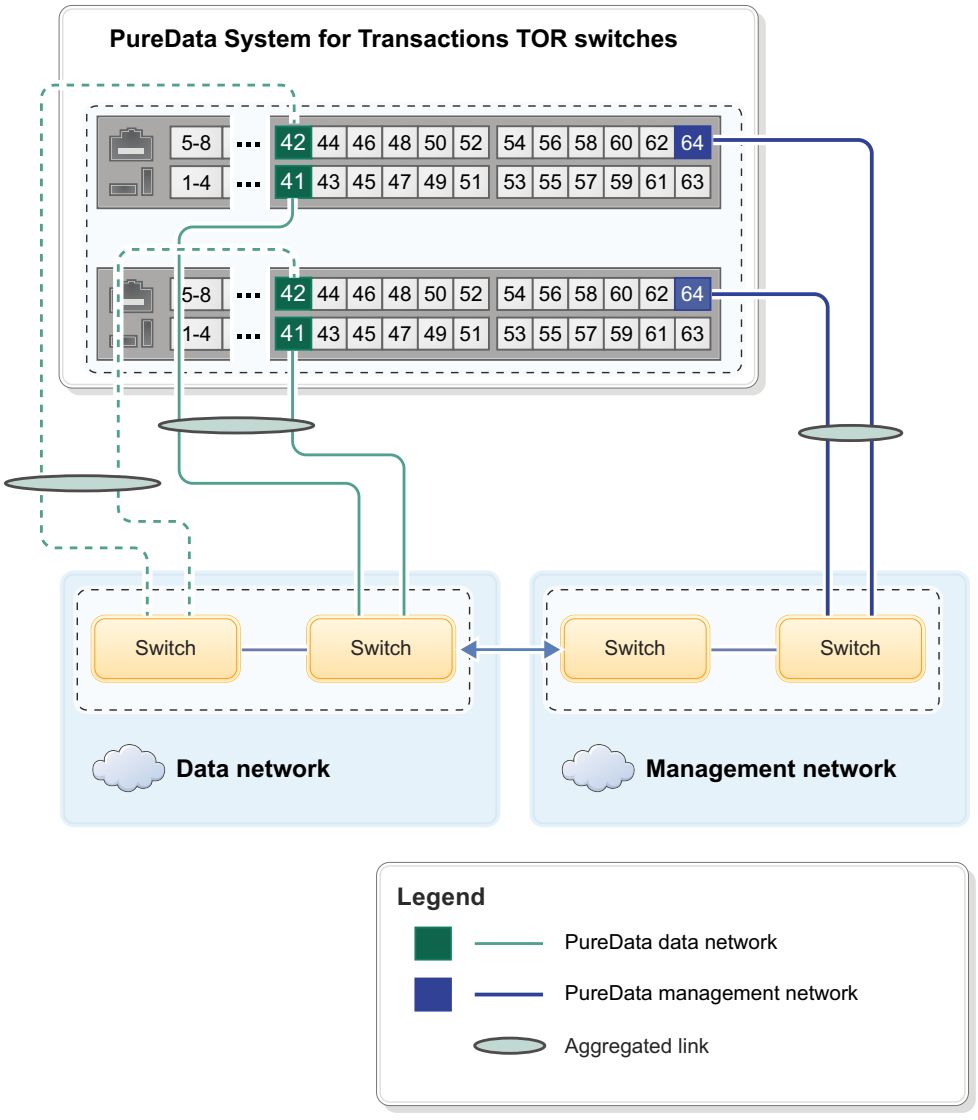

Figure 6. Non-stacked core switches with spanning tree protocol enabled.

In the diagram, the dotted green line shows the inactive path. Only port 64 on each TOR switch is reserved for the customer management network. The pair of ports, formed from port 64 on each TOR switch, can connect to only one switch at a time and thus does not support automatic failover.

Minimal setup using single core switch

The minimal network requirement for PureData System for Transactions is a single Layer 2–capable network switch for a combined data and management network. This topology is not recommended, because of the lack of redundancy. However, you could configure a temporary test system this way.

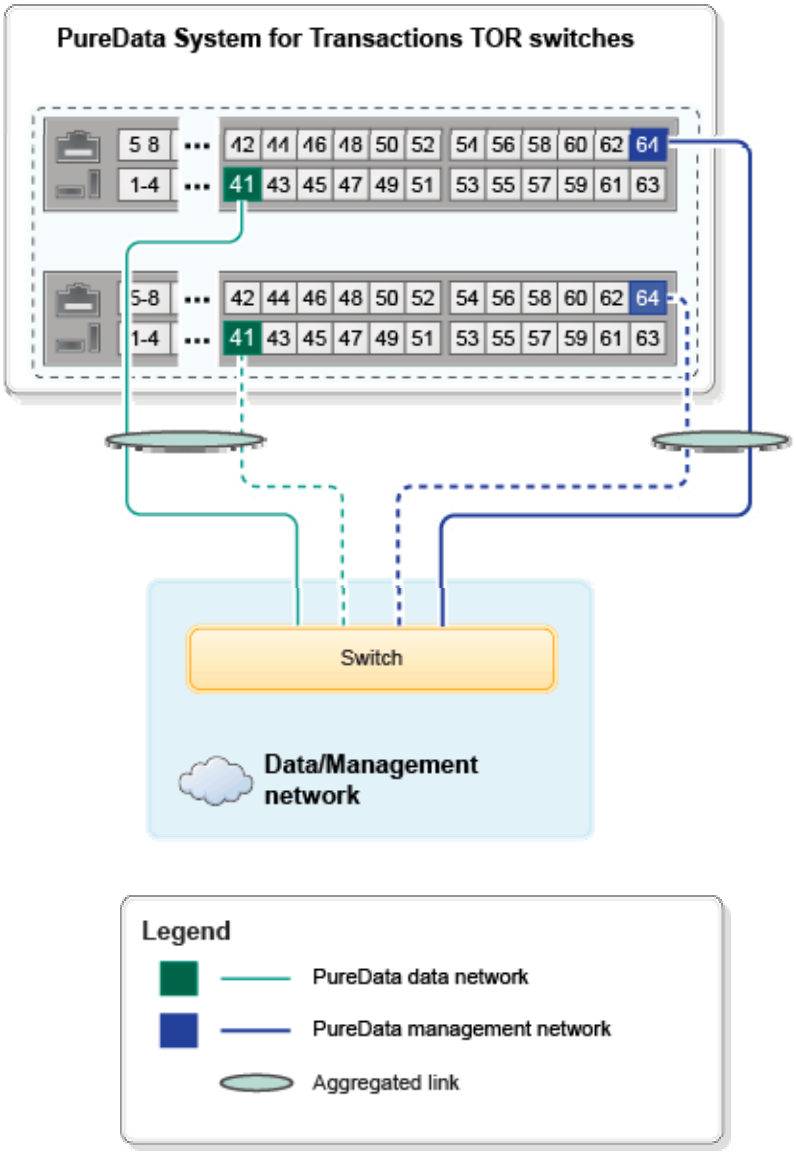

Figure 7. Minimal setup using single core switch.

The switch must be able to support 1 Gb/s or 10 Gb/s link speed and link aggregation using either LACP or EtherChannel.

Although you must configure link aggregation, you can choose to connect only one of the two cables in the link aggregation. Both aggregation methods allow the network to continue to function with only a single link.

In the diagram, the dotted line indicates the cables that you do not have to connect.

PureApplication System and PureData System for Transactions integration

If the application in PureApplication System must drive a workload on the database in PureData System for Transactions, you must use one of the following approaches:

- Configure the PureApplication System data network so that it can route data through the customer network to the PureData System for Transactions data network.

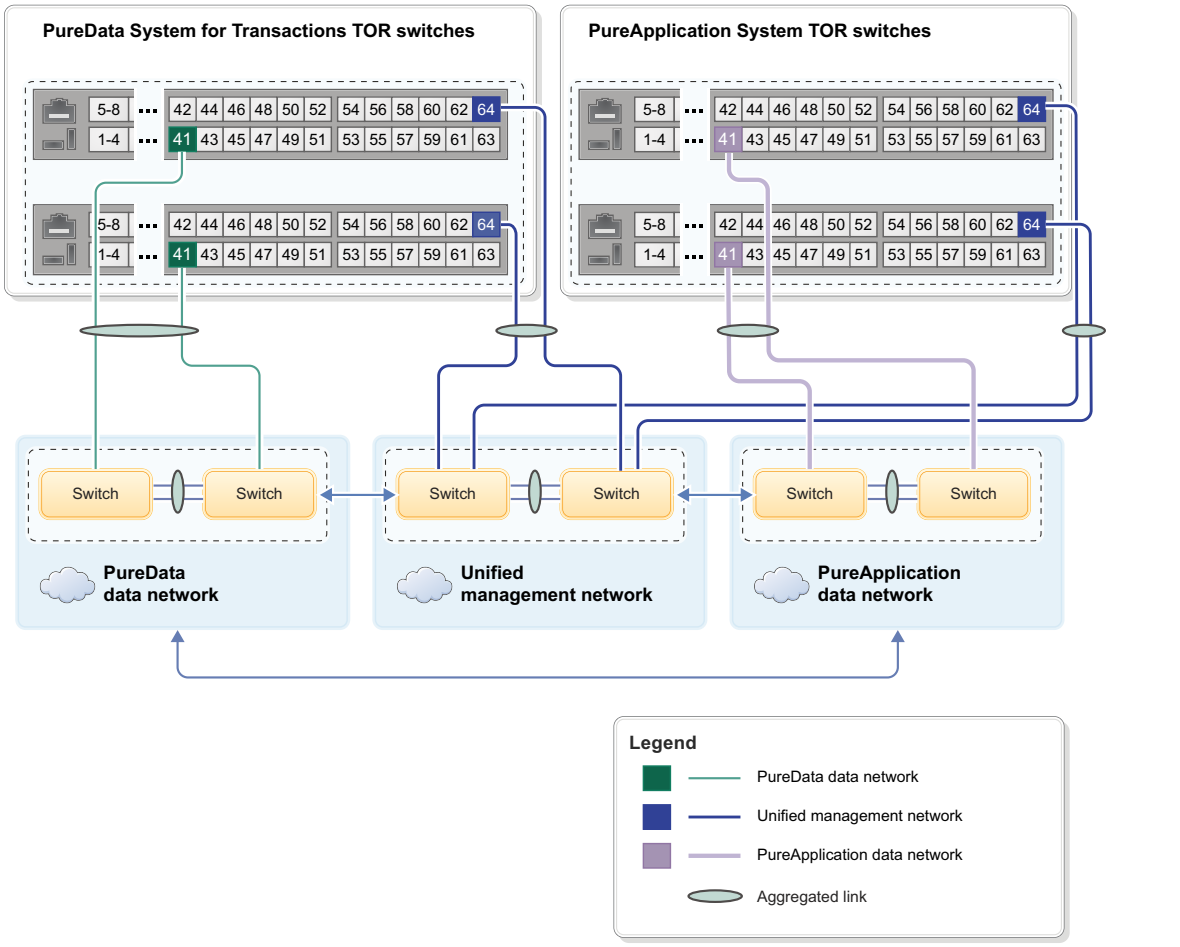

Figure 8. PureApplication System and PureData System for Transactions integration. This diagram shows a connection via the customer network.

- Create an aggregated link to connect the PureApplication System TOR switch directly to the PureData System for Transactions TOR switch. This reduces the latency of network traffic from one system to another.

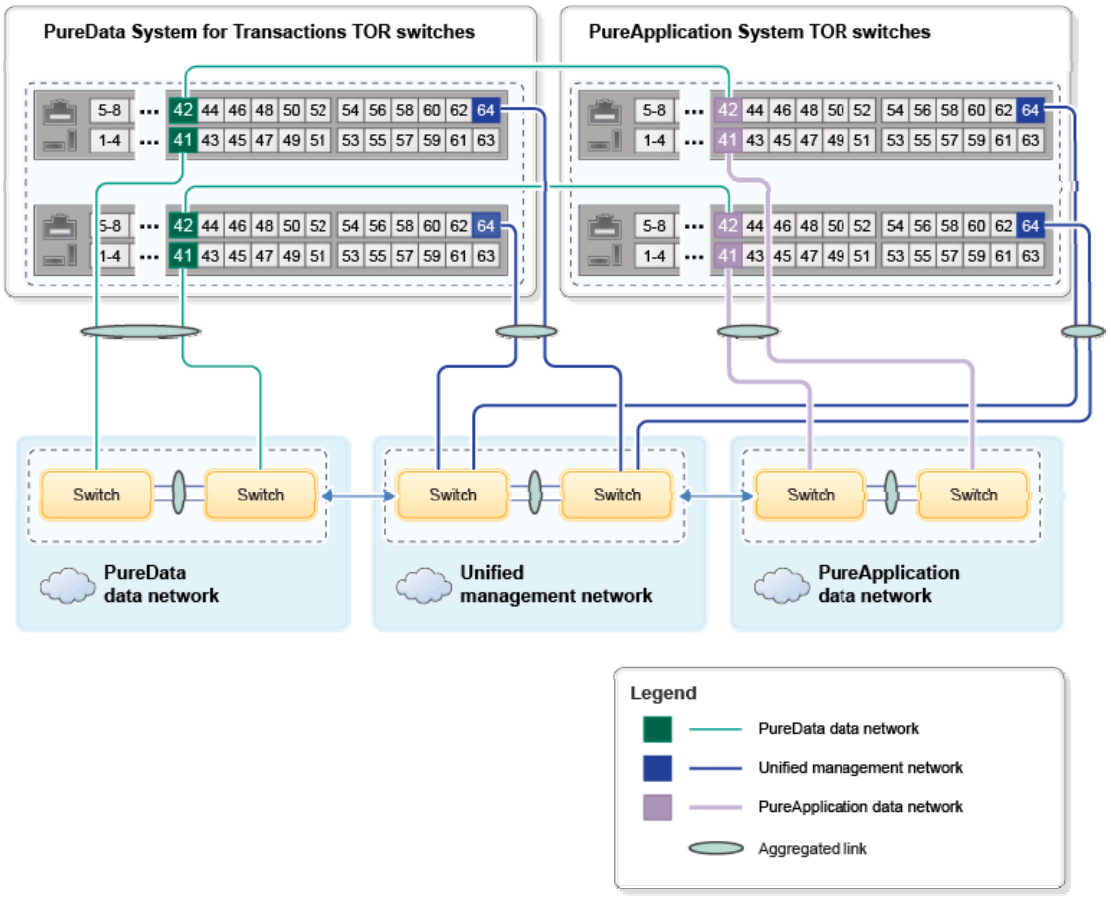

Figure 9. PureApplication System and PureData System for Transactions integration. This diagram shows a direct connection from one system to the other.

Both systems can share the same management network so that PureApplication System can manage PureData System for Transactions.

A unique VLAG Tier ID is required for the PureApplication System and the PureData System for Transactions.

Requirements for Technical Delivery Assessment

The Technical Delivery Assessment (TDA) is a checklist of items that the IBM client technical professional (CTP) completes with you before the system is set up. This document contains four sections that are related to the network integration of PureData System for Transactions with the data center; these require signoff from your network administrator. The four sections are Data Network, Management Network, IP Group Setup, and Firewall Ports. The values in the TDA directly translate to the selections that you make during the initial network setup process.

Data Network

In the Data Network section, you must provide a VLAN ID for the network that will connect client applications to the deployed instances on the compute nodes. Only one data VLAN is supported.

The VLAN ID for the data network must not conflict with the IDs that are set for the predefined internal VLANs (403, 3105, 3201, 4091, 4093, and 4094) or the VLAN ID that you set for the management network. Reusing the internal VLAN IDs in your corporate network does not impact PureData System for Transactions.

Depending on the design of the overall network, you can enable or disable PVRST for the Spanning Tree Protocol. The Spanning Tree Protocol can prevent broadcast loops when multiple paths exist between two points on a network.

Figure 10 depicts the data network settings pane where you define the data network VLAN.

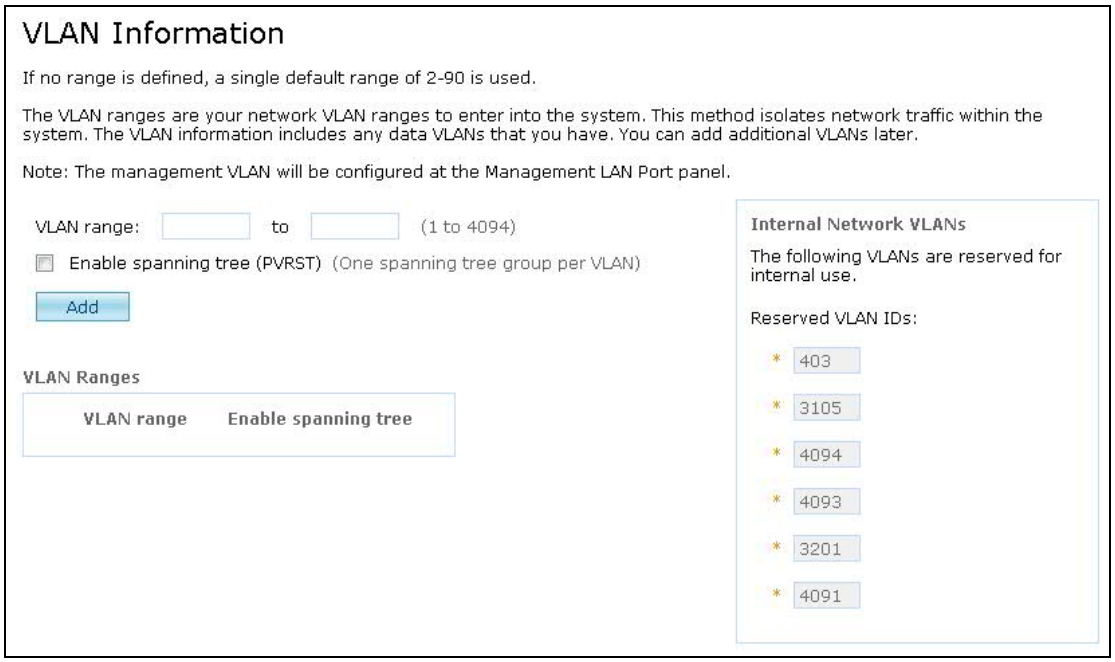

Figure 10. PureData System for Transactions system console: VLAN information pane.

Forty-six ports on the TOR switch are reserved for your data network connections. Table 2 displays the number of transceivers that are included with the system. IBM sales representative must be consulted if you want to change the number of transceivers or their positions.

Table 2. Transceivers and their positions in PureData System for Transactions

System size	Number of 1 Gb SFP RJ45 copper transceivers	Number of 10 Gb SFP+ SR optical transceivers
Extra small	4, on ports 62 – 63	6, on ports 41 - 43
Small	4, on ports 62 – 63	6, on ports 41 - 43
Medium	10, on ports 59 – 63	12, on ports 41 - 46
Large	22, on ports 53 – 63	24, on ports 41 - 52

Figure 11 depicts the data network settings pane where you assign the data ports.

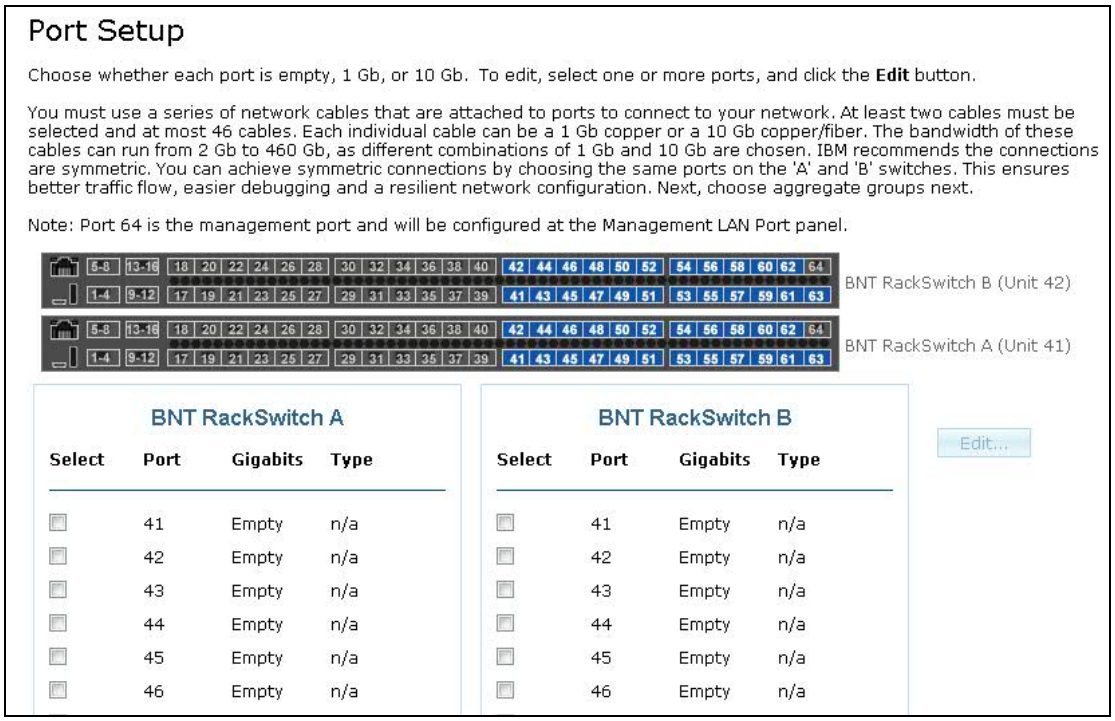

Figure 11. PureData System for Transactions system console: port setup pane.

You can create one or more aggregated links from any of the ports with transceivers that are listed in Table 2. Each aggregated link is composed of one or more pairs of ports. Ports cannot operate in stand-alone mode and can be used only when the link is aggregated. A single aggregated link cannot be composed of a combination of 1 Gb/s and 10 Gb/s ports.

The ports can auto-negotiate only at 1 Gb/s and 10 Gb/s. Speeds of 10 Mb/s or 100 Mb/s are not supported.

If you do not use ports during your initial setup, you can reassign them later.

Figure 12 depicts the data network settings pane where you assign aggregated links for the data network.

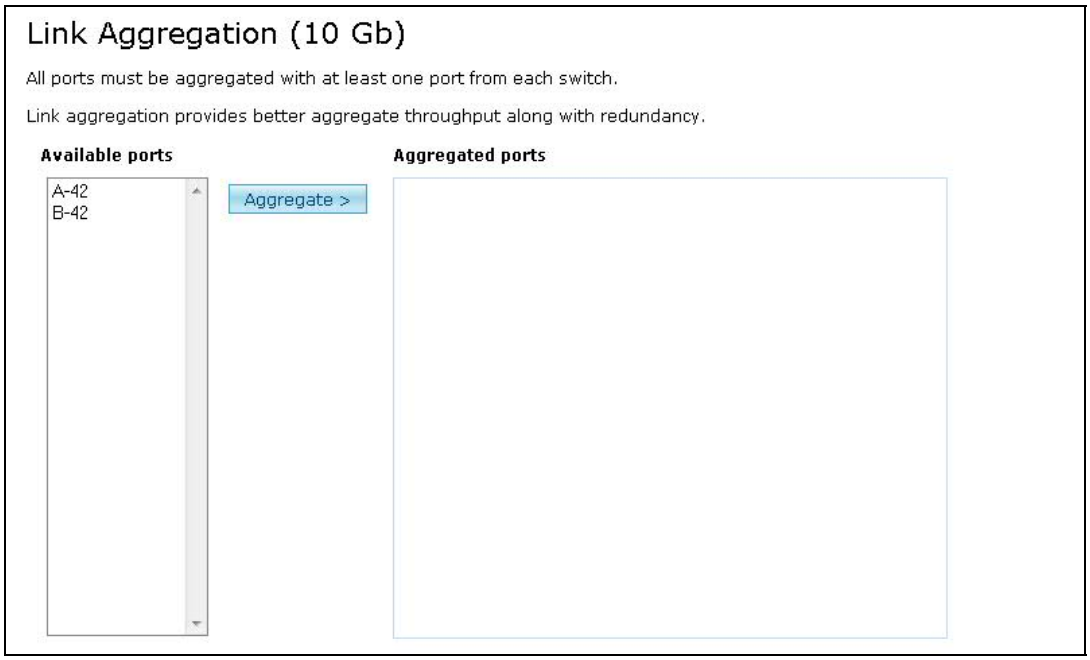

Figure 12. PureData System for Transactions system console: link aggregation pane

An aggregated link must be composed of ports from both G8264 switches. The G8264 switches are stacked by using four 40 Gb ISL with a default VLAG Tier ID of 10. You must change this value if there are other switches with the same VLAG Tier ID that you want to connect to PureData System for Transactions.

Figure 13 depicts the data network settings pane where you assign the VLAG Tier ID for the data network.

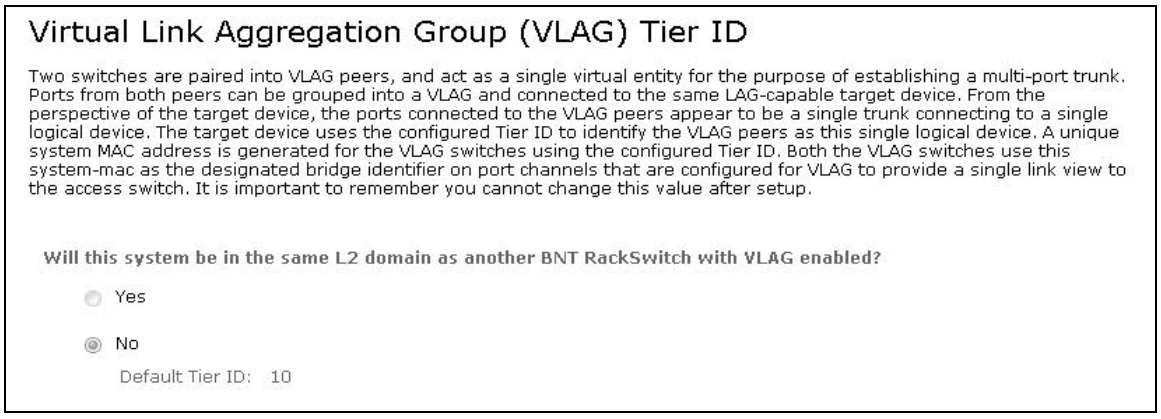

Figure 13. PureData System for Transactions system console: VLAG Tier ID pane.

When you configure your TOR switches in PureData System for Transactions, you must also configure your corporate switches that are

connected to the system. The following tables list the possible configurations for PureData System for Transactions and the configurations for the corresponding corporate switches[1]. You must configure the settings on your corporate switches in a way that is similar to the examples.

Table 3. Flow control mode options.

PureData System for Transactions configuration	IBM 8264 definitions	Cisco definitions
Flow control is enabled	`interface port 11` `flowcontrol both`	`set port flowcontrol receive 1/11 on` `set port flowcontrol send 1/11 on`
Flow control is disabled	`interface port 11` `no flowcontrol both`	`set port flowcontrol receive 1/11 off` `set port flowcontrol send 1/11 off`

Table 4. VLAN mode, tag default port VLAN ID, and PVID options.

PureData System for Transactions configuration	IBM 8264 definitions	Cisco definitions

No tagging with PVID 101[2]	Old configuration (firmware levels 7.2.x - 7.4.x):
	
	`interface port 11`
	`no tagging`
	`pvid 101`
	`.`
	`vlan 1`
	`no member 11`
	`vlan 101`
	`enable`
	`member 11`
	
	New configuration (7.6.x or later[3]):
	
	`interface port 11`
	`switchport access vlan 101`
		`interface Ethernet1/11`
		`switchport mode access`
		`switchport access vlan 101`

Tagging enabled and tag default port VLAN ID disabled with PVID 1	Old configuration (firmware levels 7.2.x - 7.4.x):
	
	`interface port 11`
	`tagging`
		`interface Ethernet1/11`
		`switchport mode trunk`
		`switchport trunk encapsulation dot1q`
		`switchport trunk allowed vlan 1,101`

[1] In the examples, assume that the port number on the customer uplink switch is port 11 and the customer data VLAN ID is 101, unless otherwise specified.
[2] Because frames are not tagged, the PVID of your core switch ports does not need to match the PVID in PureData System for Transactions.
[3] The IBM System Networking RackSwitch G8264 added support for Industry Standard CLI (ISCLI) syntax as of firmware level 7.6.

	no tag-pvid pvid 1 . vlan 1 member 11 vlan 101 enable member 11 New configuration (firmware level 7.6.x or later): interface port 11 switchport mode trunk switchport trunk allowed vlan 1,101	
Tagging enabled and tag default port VLAN ID disabled with PVID 101	Old configuration (firmware levels 7.2.x - 7.4.x): interface port 11 tagging no tag-pvid pvid 101 . vlan 1 member 11 vlan 101 enable member 11 New configuration (firmware level 7.6.x or later): interface port 11 switchport mode trunk switchport trunk allowed vlan 1,101 switchport trunk native vlan 101	interface Ethernet1/11 switchport mode trunk switchport trunk encapsulation dot1q switchport trunk native vlan 101 switchport trunk allowed vlan 1,101
Tagging enabled and tag default port VLAN ID enabled with PVID 101	Old configuration (firmware levels 7.2.x - 7.4.x): interface port 11 tagging tag-pvid pvid 101 . vlan 1 no member 11 vlan 101 enable member 11	vlan dot1q tag native . interface Ethernet1/11 switchport mode trunk switchport trunk encapsulation dot1q switchport trunk allowed vlan 101

	New configuration (firmware level 7.6.x or later):
	
	`interface port 11`
	`switchport mode trunk`
	`switchport trunk allowed vlan 101`
	`vlan dot1q tag native`
	`switchport trunk native vlan 101`	

Table 5. Link aggregation method[4].

PureData System for Transactions configuration	IBM 8264 definitions	Cisco definitions

LACP	`interface port 11`
	`lacp mode active`
	`lacp key 1500`
	`.`
	`vlag enable`
	`vlag tier-id 20`
	`vlag isl vlan 4094`
	`vlag hlthchk peer-ip 192.168.70.82`
	`vlag isl adminkey 1000`
	`vlag adminkey 1500 enable`
		`interface Ethernet1/11`
		`switchport mode trunk`
		`switchport trunk allowed vlan 101`
		`channel-group 30 mode active`

EtherChannel	`portchannel 30 port 11`
	`portchannel 30 enable`
	`.`
	`vlag enable`
	`vlag tier-id 20`
	`vlag isl vlan 4094`
	`vlag hlthchk peer-ip 192.168.70.82`
	`vlag isl adminkey 1000`
	`vlag portchannel 30 enable`
		`interface Ethernet1/41`
		`switchport mode trunk`
		`switchport trunk allowed vlan 101`
		`channel-group 30 mode on`

Figure 14 and 15 depict the data network settings pane where you assign link aggregation options.

[4] The link aggregation methods for PureData System for Transactionsand the corporate switches must match. Otherwise, there will be no network connectivity.

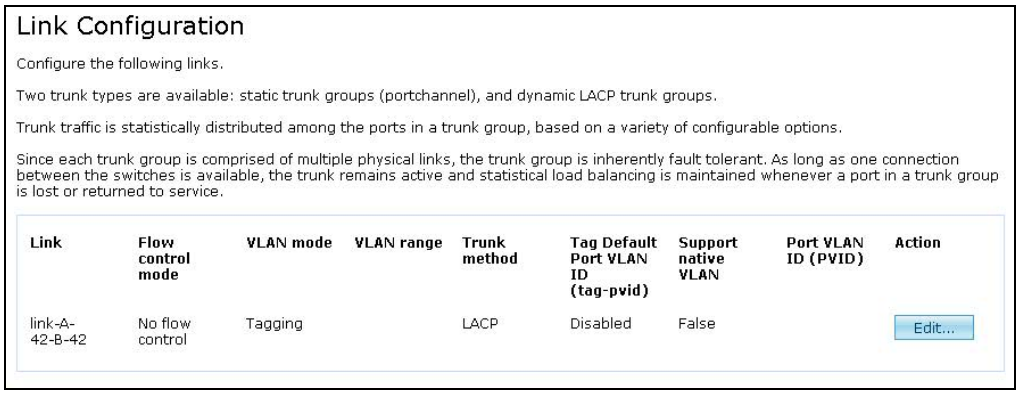

Figure 14. PureData System for Transactions system console: link aggregation summary pane.

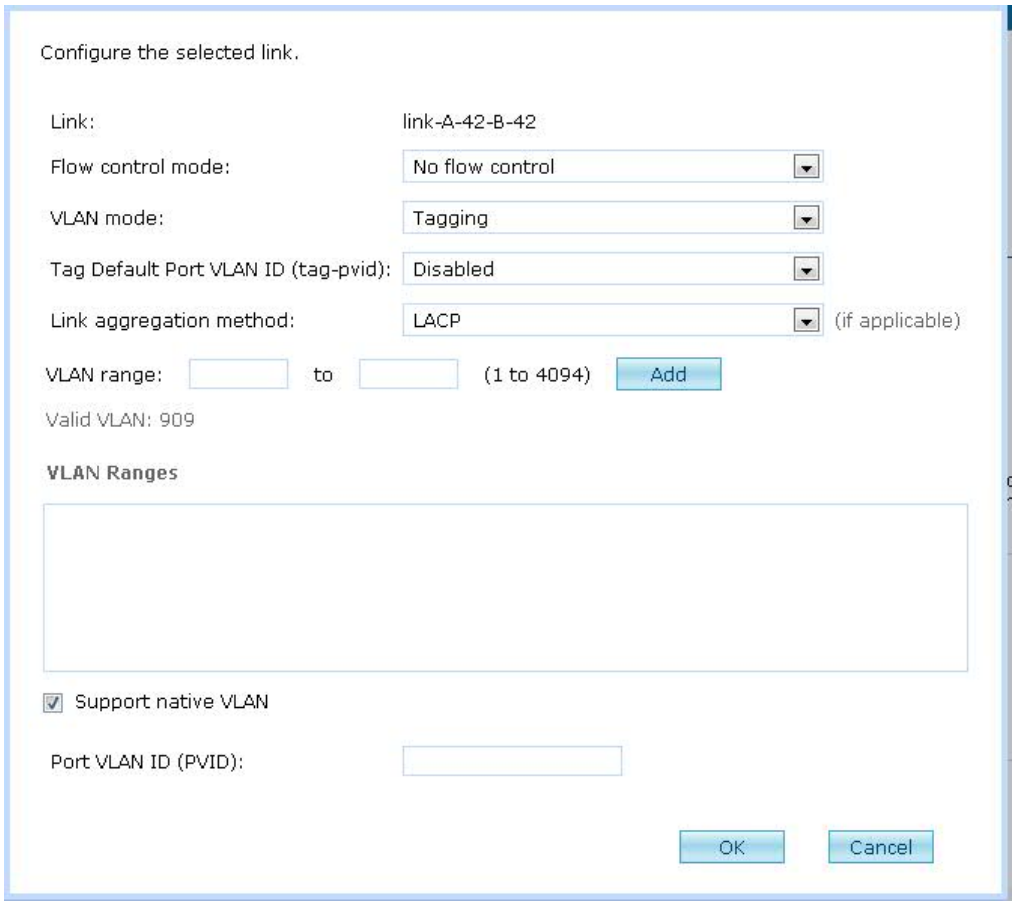

Figure 15. PureData System for Transactions system console: link aggregation options pane.

Management Network

Two 1 Gb ports are reserved for the management network. In the TDA, the options in the Management Network section for link aggregation (flow control mode, VLAN mode, link aggregation method, and spanning tree) are similar to those that are outlined in the Data Network section. The management network VLAN ID must not conflict with the IDs for the predefined internal VLANs (403, 3105, 3201, 4091, 4093, and 4094) or the VLAN ID for the data network.

Figure 16 depicts the management network settings panel where link aggregation options are assigned.

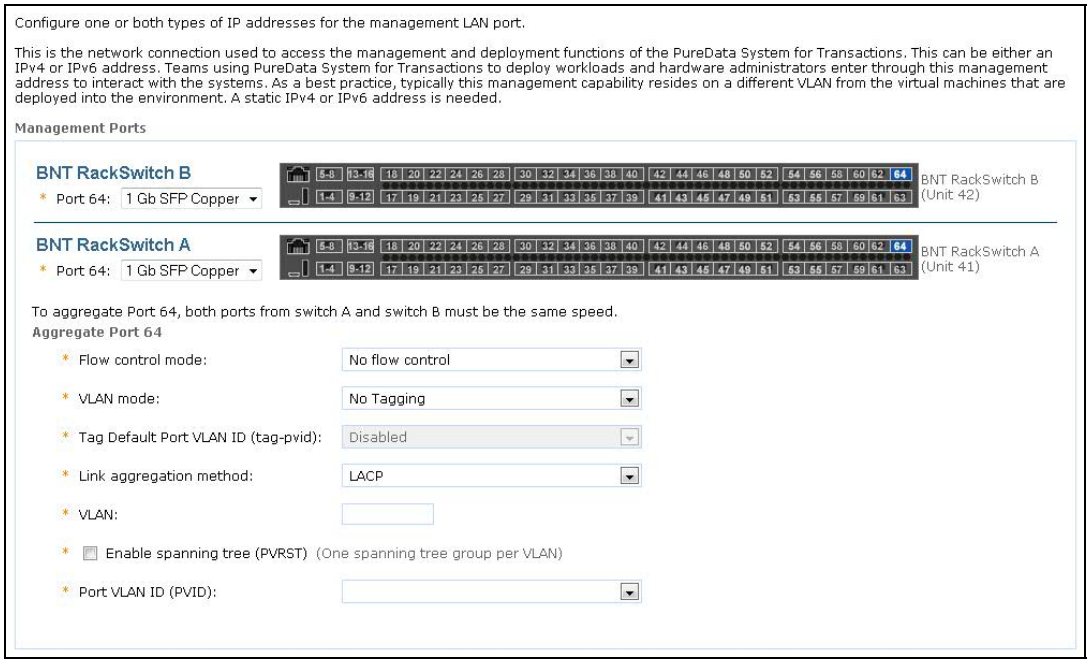

Figure 16. PureData System for Transactions system console: management network link aggregation pane.

In the Management Network section, you can choose the networking details for PureSystems Manager. The network configuration of PureSystems Manager requires three IP addresses:

- An IP address for PureSystems Manager that plays the primary role in bay 2
- An IP address for PureSystems Manager that plays the secondary role in bay 1
- A floating IP address for two PureSystems Managers

Both IPv4 and IPv6 addresses are supported. You can specify other details, including the host name for PureSystems Manager, a domain name, the prefix length or subnet mask, and a default gateway.

Figure 17 depicts the management network settings pane where you assign IP addresses for PureSystems Manager.

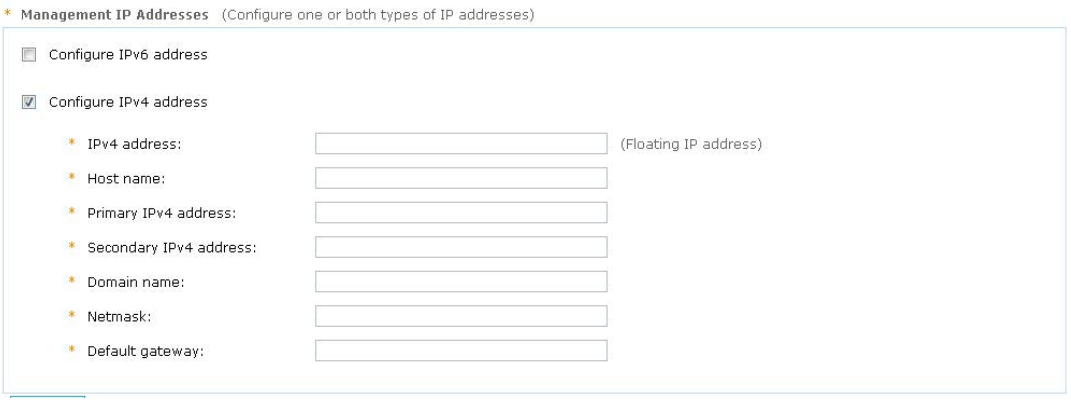

Figure 17. PureData System for Transactions system console: management node networking pane.

At least one DNS and one NTP server must be accessible to PureSystems Manager on the management network. You must set up the DNS with both forward and reverse lookup for the host names and IP addresses that you select in the Management Network section and the next section (IP Group Setup). The NTP server synchronizes the date and time for all devices that are connected to the data center's network.

IP Group Setup

This section continues to describe the data network setup and includes the IP networking details for the deployed instances on the compute nodes.

The data VLAN field refers to the VLAN ID in the Data Network section. The cloud management VLAN is an additional VLAN that is reserved for rack-internal communication between cloud group compute nodes. . You can assign the cloud management network a VLAN ID that does not conflict with the data VLAN ID, management VLAN ID, or reserved internal VLAN IDs.

For the IP group, you can reserve a number of IP addresses that is equal to the number of compute nodes in your PureData System for Transactions system. You can choose either IPv4 or IPv6 addresses. These addresses cannot conflict with addresses for PureSystems Manager or addresses in the

following reserved ranges: 10.153.0.0 - 10.153.7.255, 192.168.70.0 - 192.168.70.255 and 169.254.197.0 - 169.254.197.255.

You can select other networking details, including the netmask, subnet address, gateway address, and DNS addresses, in the IP Group Setup section.

Table 6 lists the number of compute nodes and IP addresses that you must reserve in the IP group for each system size.

Table 6. Number of compute nodes and IP addresses to reserve for each rack size

System size	Number of compute nodes	Number of IP addresses to reserve in IP group
Extra small	2	2
Small	6	6
Medium	12	12
Large	24	24

Firewall Ports

PureData System for Transactions uses several network ports. Your network administrator must confirm that they are open. For the management network, you must enable the ports to allow access between PureSystems Manager and the management console clients.

Table 7. Ports to enable to allow access between PureSystems Manager and management console clients

Port number	Port usage
80	HTTP PureData management network
443	HTTPS PureData management network

For the data network, you must enable the following ports between the compute nodes and the management console clients, application servers, or any other clients that need access to the database instance. In addition, the database management feature requires PureSystems Manager to be accessible to the compute nodes.

Table 8. Ports to enable between compute nodes and management console clients, application servers, or any other clients that need access to the database instance

Port number	Port usage
22	SSH
80	HTTP
443	HTTPS
8888	Management interface for applications and shared services
5801	VNC launch shortcut
5901	VNC
50000	Optim Performance Manager database monitoring
55000	Optim Performance Manager database monitoring

You can also define the ports for SMTP (25) and LDAP (389). You must enable them for PureSystems Manager and the compute nodes.

You must enable additional ports for the DB2 instances on compute nodes when you create new instances through pattern deployment.

Best practices

- To integrate IBM PureData System for Transactions into a data center network, select one of the sample network topologies in this paper.
- Confirm that the reserved IP addresses and VLANs do not conflict with the data center network.
- Verify that the NTP server is available on the customer management network.
- Verify that the DNS is available on both the customer data and management networks. Configure it with both forward and reverse lookup for the host names and IP addresses.
- Ensure that the configuration of the core switches corresponds to the network options that you selected in the TDA.
- Verify that the ports that you require for IBM PureData System for Transactions are opened on the firewall.

Conclusion

Solid and reliable network communications are critical to getting the most from PureData System for Transactions. The system console interfaces make virtual LAN setup, port selection, link aggregation, and VLAG Tier ID setup straightforward. Configuration of the data network, management network, IP groups, and firewall ports for your system is also quick and efficient.

Integrating PureData System for Transactions into your network by using the best practices that are presented in this paper will help ensure the best possible network response and throughput.

Contributors

Leslie McDonald
Technical editor, IBM Canada

Sandra Seenauth
Visual designer, IBM Canada

Challen Pride-Thorne
Visual designer, IBM Canada

Jessy Chung
Visual designer, IBM Canada

Serge Boivin
Senior technical writer, IBM Canada

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

Without limiting the above disclaimers, IBM provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any recommendations or techniques herein is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Anyone attempting to adapt these techniques to their own environment does so at their own risk.

This document and the information contained herein may be used solely in connection with the IBM products discussed in this document.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments

may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

Windows is a trademark of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Contacting IBM

To provide feedback about this paper, write to psdbsys@ca.ibm.com.

To contact IBM in your country or region, check the Directory of Worldwide Contacts (http://www.ibm.com/planetwide).

To learn more about IBM Information Management products, go to http://www.ibm.com/software/data/.