z/VM Version 7.1

New Function APAR

Virtual Switch Priority Queuing

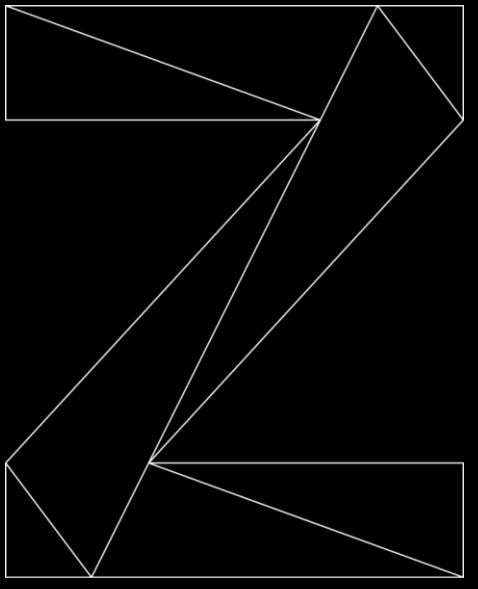

Virtual Switch Priority Queuing

- Introduces multiple priority levels for uplink transmissions from a Virtual Switch
- Allows VSwitch management communication (including IVL) to operate at highest priority to ensure better management
- Three optional user priority levels allow:
  - Different SLAs for different groups of guests
  - Combining different priority workloads onto fewer, or a single VSwitch
    - Multiple priorities allow z/VM to use additional logical processors/cores for outbound data transmissions
- The benefit of Priority Queuing is realized when multiple users are simultaneously transmitting data to the VSwitch's uplink port

Priority queuing levels for a guest are *not* a hard limit setting.
This function only impacts the priority of the data that gets sent to the physical network. It has no effect on how fast it goes over the fabric and has no impact on incoming data.

Virtual Switch Priority Assignments

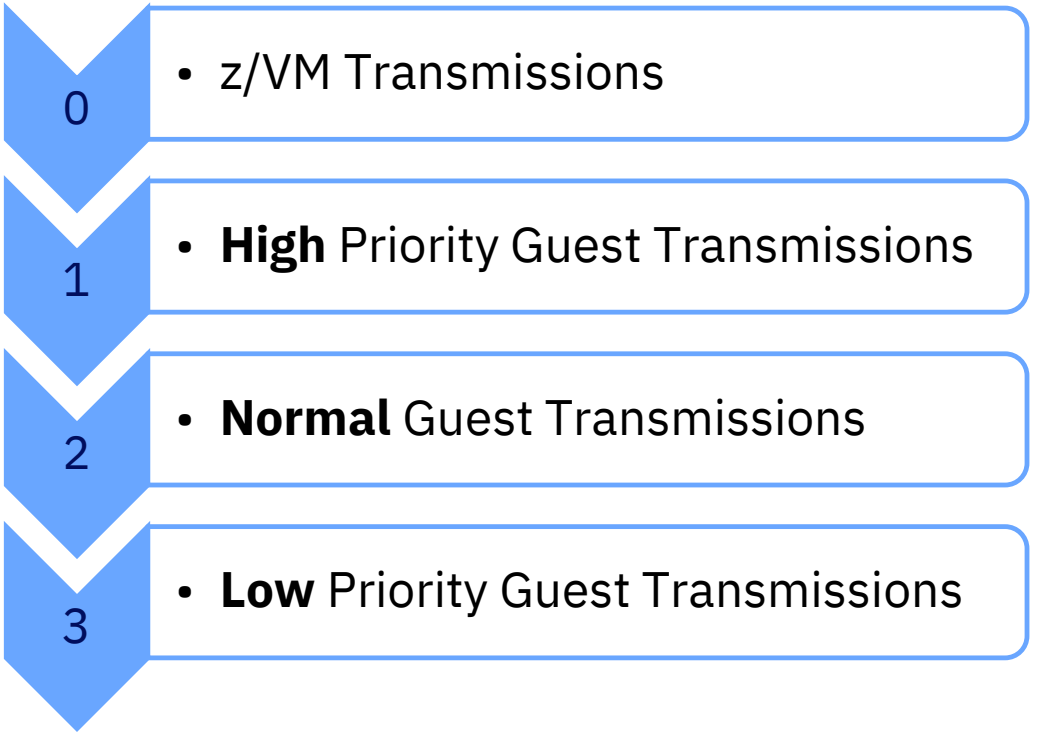

- 4 QDIO output queues established on the uplink port.
z/VM transmissions include ARP, LACP transmissions, IVL management data.
Best practice is to assign users with different transmission types to different priorities. For example, guests that are doing backup or streaming workloads could be assigned Low priority. Guests with interactive workloads could be assigned higher priority.

Enabling VSwitch Priority Queuing

- Priority Queuing is enabled in OSA-Express hardware by default
  - IOCP or dynamic I/O change is required to disable
- IVL VSwitches always exploit priority queuing if not disabled
- Exploitation must be enabled for non-IVL VSwitches
  - **`DEFINE VSWITCH`** command/config statement
- Set guest priority (default is NORMAL)
  - **`NICDEF`** directory statement
- If you want to relocate a guest that is using priority other than NORMAL, then the VSwitch on the target system must also be enabled for priority queuing
  - Or set guest priority to NORMAL before relocating guest

How to get VSwitch Priority Queuing

- Availability – May 2019
  - z/VM 7.1

Component	APAR	PTF	RSU
CP	VM66219	UM35465	TBD
TCP/IP	PH04703	UI62768	TBD
DirMaint	VM66223	UV99352	TBD

Publications

New editions of the following publications are available:

- SC24-6267-01: Connectivity
- SC24-6268-02: CP Commands and Utilities Reference
- GC24-6270-01: CP Messages and Codes
- SC24-6271-01: CP Planning and Administration
- SC24-6272-01: CP Programming Services
- GC24-6286-02: General Information
- GC24-6294-02: Migration Guide
- SC24-6301-02: Performance
- SC24-6281-01: Directory Maintenance Facility Commands Reference

Additional Information

OSA-Express Priority Queuing Support

- **Current Approach**

  All outbound traffic to the external network is transmitted at the same priority.

  - Data transmission is in a FIFO order
  - z/VM and guest transmissions have the same priority

- **New Approach**

  Multiple OSA QDIO Output Queues will be established and activated on the VSwitch UpLink Port

  - OSA Express Feature must be configured to support priority queuing via IOCP (Default is PQ_ON). This can be set with DEFINE CHPID/PATH.
  - Will be enabled on the IVL VSwitch
  - Can be enabled on Global and Conventional VSwitch UpLink Ports
  - System Administrator determines transmission priority of each guest's transmission to the external network

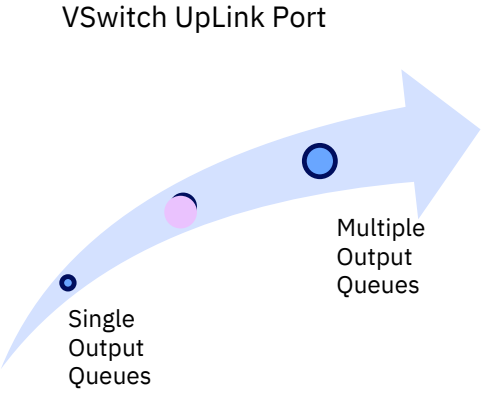

PQ is capability of QDIO architecture to define multiple output queues for a single network connection. Each queue is weighed by a priority on how often it gets serviced by OSA Express.

Rationale

VSwitch management transmissions are sent to the physical network ahead of other transmissions to ensure VSwitch management is always maintained.

The IVL VSwitch

- IVL Network Management data sent on highest priority queue
- Error recovery traffic sent on lower priority queue

Other VSwitches

- Management data on highest priority queue
- Guest transmissions sent on the three lower priority queues

Externals

DEFINE VSWITCH new UPLINK operand

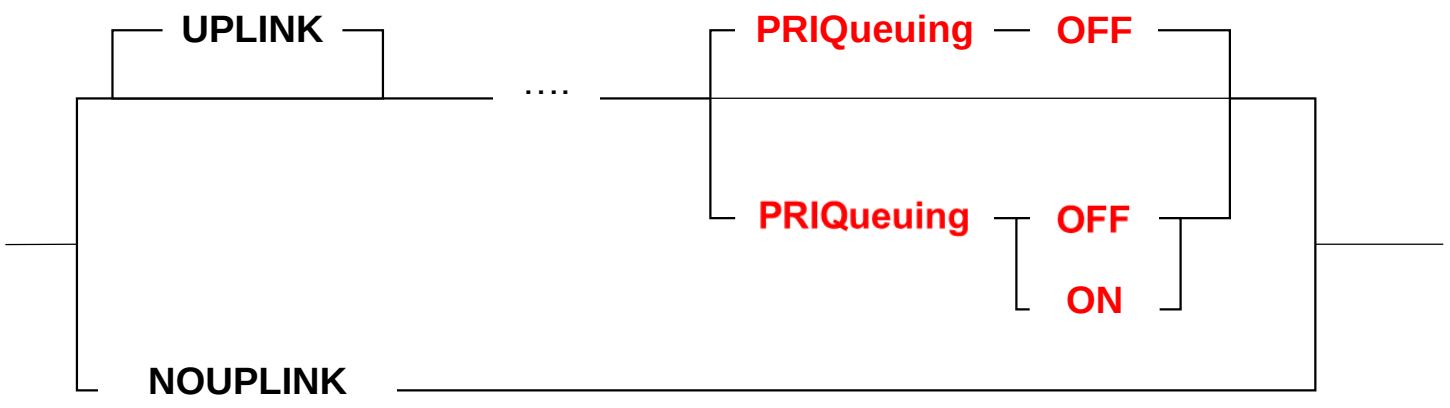

For an IVL VSwitch, PRIQueuing ON is the default and only allowed option.

A warning message will be displayed to inform the customer to configure priority queuing. For migration reasons, if the OSA-Express Features used by an IVL VSwitch Uplink Port are configured without priority queuing the OSA WILL be allowed to establish an active network connection.

SET VSWITCH new UPLINK operand

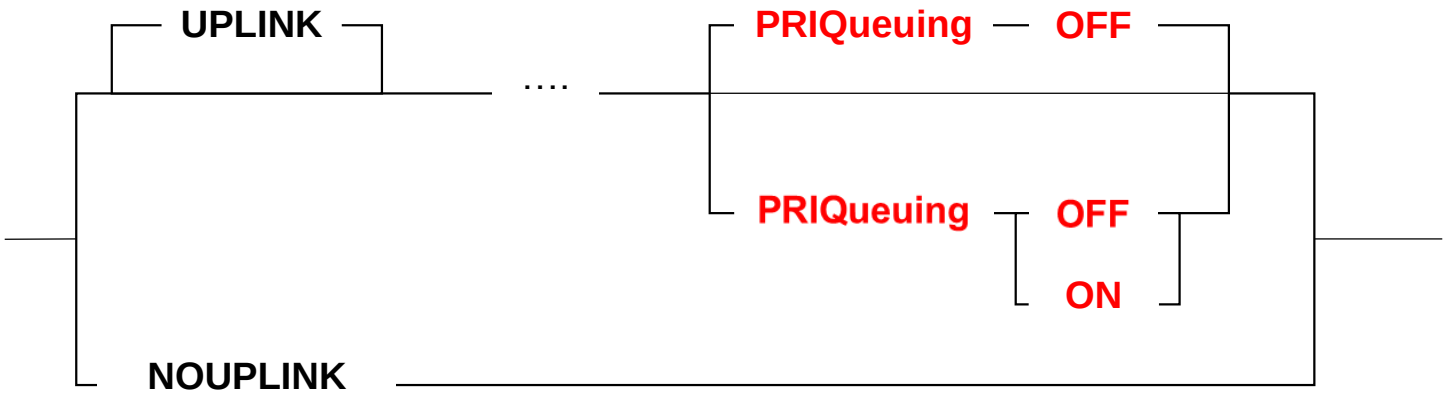

Allow customers to dynamically turn Priority Queuing ON or OFF.
VSwitch must be disconnected before making this change.

NICDEF Directory Statement

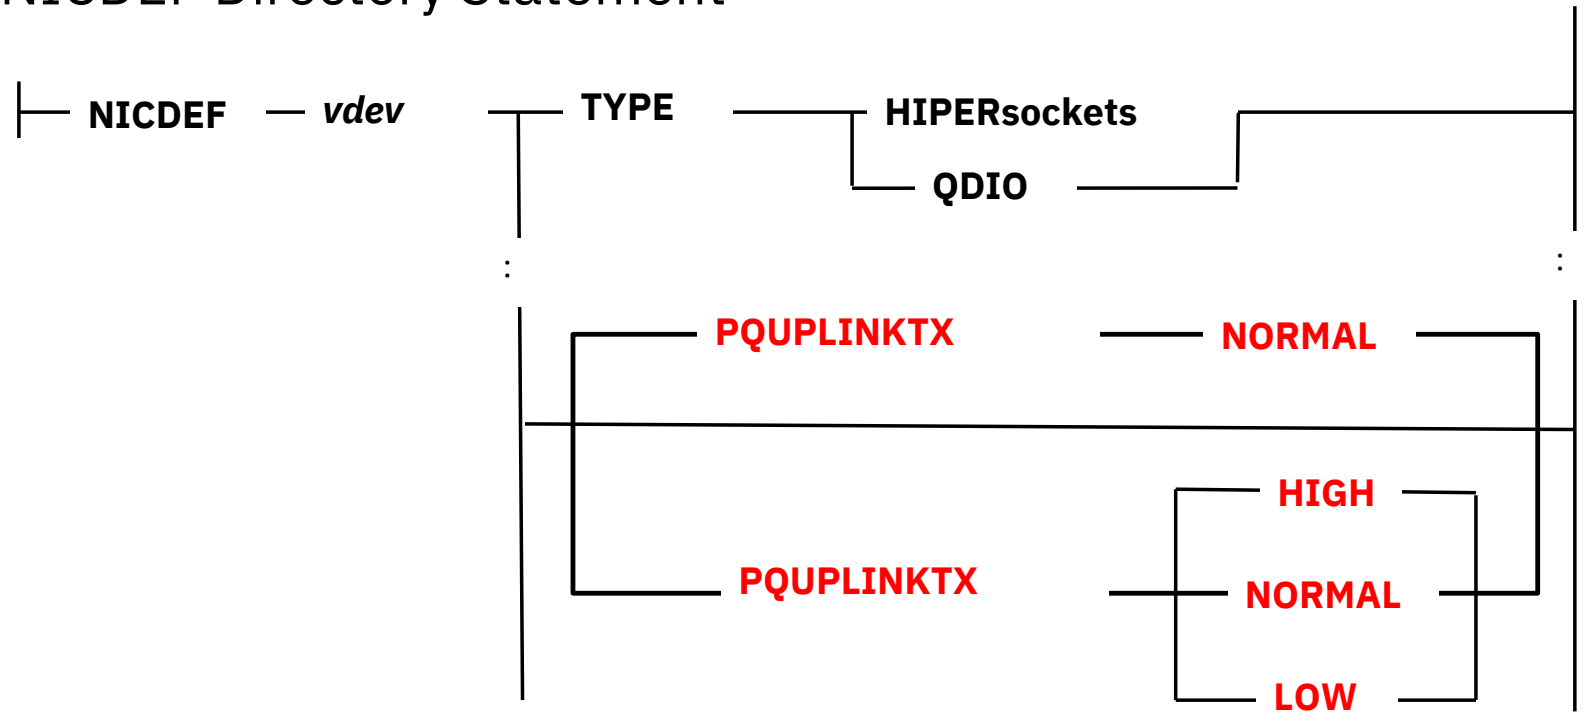

If PQUPLINKTX is specified for VSwitch without priority queuing, the NIC is allowed to couple with normal priority, and priority setting will be saved in case PQ is turned on later.

SET VSWITCH Priority Assignment

Grant and Portnumber Options:

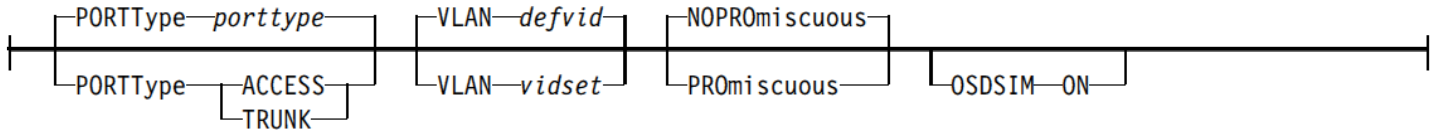

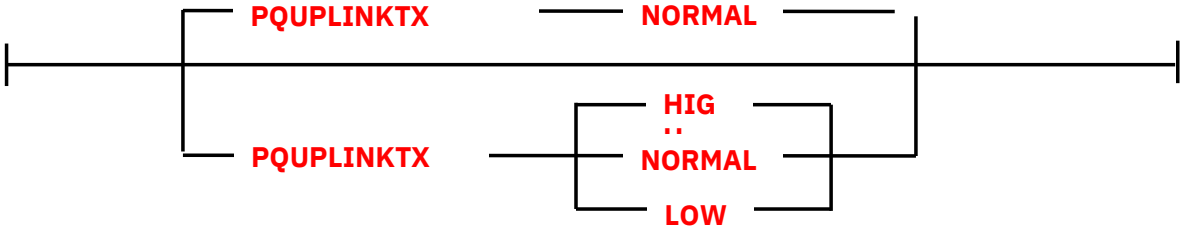

IVL VSwitch does not allow grant or portnumber so priority pquplinktx does not apply. For IVL, priority of outbound transmissions is handled by z/VM.

For other VSwitch types, existing connections will be updated.

Config Statements

- **DEFINE VSWITCH**
  - New UPLINK operand PRIQUEUING

  DEFINE VSWITCH VSW1 ETHERNET UPLINK RDEV 63D0.P01 PRIQUEUING ON

- **MODIFY VSWITCH**
  - New PQUPLINKTX option for GRANT/PORTNUMBER

  MODIFY VSWITCH VSW1 GRANT USER1 PQUPLINKTX HIGH
  MODIFY VSWITCH VSW1 PORTNUMBER 1 USERID USER2 PQUPLINKTX LOW

Relocation

- Priority Queuing does NOT need to be on both systems.
- If the source guest is exploiting PQ (using high or low priority queue) , then the destination VSwitch must have PQ enabled to relocate.
- If the source guest is not exploiting priority queuing, then relocation will succeed to destination VSwitch with or without PQ.