digital™

AV-RGTAB-TE

DIGITAL VNswitch 900 Series
Version 3.4
Release Notes
October 1999

As warranted, Cabletron changes the firmware of this device to make functional enhancements or to correct reported problems. These Release Notes identify enhancements and changes to the firmware that impact end-user operations. They also contain firmware and software requirements and list updates in this release as well as known conditions and restrictions that apply to the operation of the modules.

The following example describes the firmware version number:

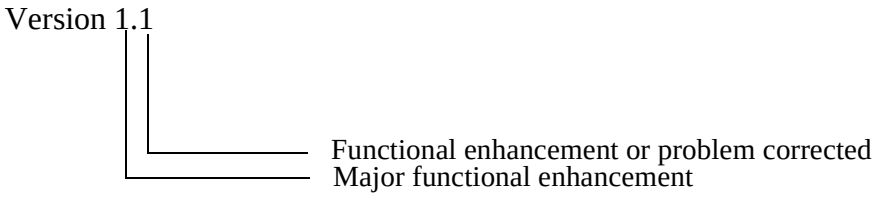

Note: The VNswitch product documentation is currently at Version 3.0. However, Version 3.0 product documentation, including these Release Notes, reflects the major functional release, Version 3.4.

Contents

Hardware and Firmware Support

The currently-supported released versions for the VNswitch are:

Version V3.0.1 - includes the following switch variants:

- DVNEA (VNswitch 900EA)
- DVNEE (VNswitch 900EE)
- DVNEF* (VNswitch 900EF)
- DVNEX (VNswitch 900EX)
- DVNFA (VNswitch 900FA)
- DVNLL (VNswitch 900LL)
- DVNXA (VNswitch 900XA)
- DVNXX (VNswitch 900XX)
- DVNFF (VNswitch 900FF)
- DVNFX (VNswitch 900FX)

* DVNEF-MM includes 12 10BaseT ports and 1 pair fixed ANSI MIC multimode fiber connectors
  DVNEF-MX includes 12 10BaseT ports and 1 MMI FDDI port

Version V3.1 - DVNGV (VNswitch 900GV)

Version V3.2 - DVNCC (VNswitch 900CC)

Version V3.3 - DVNCG (VNswitch 900CG)

DIGITAL MultiSwitch 900 Support

For VNswitch operation in the DIGITAL MultiSwitch 900, we recommend:

- For VNswitch 900EE, VNswitch 900EF, VNswitch 900EA, VNswitch 900EX, VNswitch 900XX, VNswitch 900LL, VNswitch 900FA, and VNswitch 900XA modules, you must use DIGITAL MultiSwitch 900 firmware V5.2 or higher.
- For VNswitch 900CC, VNswitch 900CG, VNswitch 900GV, VNswitch 900FF and VNswitch 900FX modules, you must use DIGITAL MultiSwitch 900 firmware V5.4 or higher.

*clear*VISN Support

To manage the VNswitch modules using the *clear*VISN application, we recommend:

- For VNswitch 900EE, VNswitch 900EF, VNswitch 900EA, VNswitch 900EX and VNswitch 900XX modules, you must use *clear*VISN V2.0 or higher.
- For VNswitch 900LL, VNswitch 900FA, and VNswitch 900XA modules, you must use *clear*VISN V2.1 or higher.
- For VNswitch 900FF and VNswitch 900FX modules, you must use *clear*VISN V2.2 or higher.
- For VNswitch 900CC, VNswitch 900CG, and VNswitch 900GV, you must use *clear*VISN V3.0 or higher.

New for Release V3.4

Release V3.4 is for all variants of VNswitch. This release supports IGMP-Snooping. IGMP-Snooping is a mechanism that limits IP Multicast traffic to only those LAN segments where there is a client that is interested in that particular IP Multicast group. Refer to *IGMP Snooping* on page 24 for more details.

New for Release V3.3

This section describes Release V3.3 specifics for the new VNswitch 900CG.

The 900CG combines features of both the 900CC and 900GV in a cost-effective, high-density Fast-Ethernet, Gigabit-Ethernet product. The 900CG supports eight Ethernet 10/100BaseTX ports and a Gigabit Ethernet uplink function with one Multiswitch 900 VNbus connection.

Configuration Management for the 900CG is similar to other VNswitches. The three VNswitch User Manuals (*Switch Management Guide, Technical Overview Manual*, and *Router Manual*) version 3.0, all apply to the 900CG. In addition, the management/configuration of the eight 10/100BaseTx ports is similar to the management/configuration of the 10/100BaseTx ports of the VNswitch 900XX documented in the *Switch Management Guide*.

The *Known Considerations and Restrictions* on page 7 for VNswitch 900CC and VNswitch 900GV also apply to VNswitch 900CG, as indicated.

Features in This Release

VNswitch V3.*x* firmware supports web-based management of VNswitch modules. The following features are included in this release. Refer to the *DIGITAL VNswitch 900 Series Switch Management* guide for further information.

- **Web-Based Management**
  The VNswitch module with V3.0 firmware includes a built-in web server and management application that allows you to configure and monitor the module over the Internet. For web access you must first assign an IP address to the module using the CLI. You can use either of the following web browsers: Netscape V4.0 or Internet Explorer V4.0. Following are the web management features supported in this release:

  — **System**
  General
  Reset
  Error Log (Crash Log and Diagnostic Log)
  Memory
  IP Host
  Upgrade Device
  Save/Restore Configuration
  SNMP (Communities, RMON- Statistics Configuration, RMON - History Configuration)

  — **Interfaces**
  Configuration
  Packet Counters
  Error Counters

  — **Bridging**
  General
  By Port
  Spanning Tree by VSD
  Address Filters
  Protocol Filters
  Forward Bridging Database
  VSD Configuration General

  — **IP Routing**
  General
  Enable/Disable Routing
  Addresses
  IP Counters
  Static Routers
  Access Control
  Filtered Routes
  Enhanced Proxy ARP

  — **Telnet to Module**

- **BGP4**
  This feature is an implementation of the latest version of the Border Gateway Protocol (BGP), BPG4, which is defined in RFC 1654. BGP is an exterior gateway routing protocol that allows the exchange of network reachability information among autonomous systems (AS). Routing information within each AS is shared using interior gateway protocols, such as RIP, OSPF, or Integrated IS-IS. For more information, refer to the *DIGITAL VNswitch 900 Series Router Management* guide.

- **T3/E3 and T1/E1 ATM ModPHYs**
  This release supports T3/E3 and T1/E1 ModPHYs for VNswitch ATM modules.

- **New Command Line Interface (CLI) Functionality**

  — **New ATM physical (ATM Phy Config>) commands for T3/E3 and T1/E1 ModPHY support**
  These commands include **ENABLE/DISABLE** commands, **LIST** commands, and **SET** commands.

  — **New counters for the ATM Phy>LIST STATUS command:**

    - Net up count
    - Net down count
    - Data direct and Bridge Tunnel VC count
    - Number of times the PHY went into loopback because the PHY was down

  — **Additional ATM enhancements**
  These enhancements include new **LIST SUMMARY** commands for ATM> and ATM Config> and improvements for the **Config>CLEAR ATM** command.

  — **Command Line Editing and Command Line Completion**
  This release allows you to edit commands on the command line, recall previously entered commands, and complete partially entered commands automatically.

- **Duplicate MAC**
  This release of the firmware supports existence of duplicate MAC addresses on separate VSDs.

- **Year 2000 Compliance**
  This release of the firmware is Year 2000 compliant.

Known Considerations and Restrictions

Read this section for special considerations and restrictions regarding the operation of your VNswitch 900 modules.

Ethernet

10/100 Mbit (DVNCC, DVNCG)

- When configured in Full-Duplex/100 Mbit mode, Ethernet ports on the 900CC/CG are configured to support 802.3 Pause Flow Control Frames. You ***cannot*** prevent the 900CC/CG from sending these frames. It is possible to remove this capability when auto-negotiating by using:

  FAST ETH/*XX*> set advertise-capability

  However, Ethernet ports on the 900CC/CG always send and honor 802.3 Pause frames when in 100 Mbit full duplex, regardless of how the capability is set.

- All Ethernet ports on 900CC/CG are RMON capable. There is support to display the RMON counters as part of the CLI command:

  MONITOR> interface statistics XX

- RMON Web pages are also supported via the WEB interface and corresponding standard RMON MIBs for any standard RMON Management application.

- The Activity LED behavior of Ethernet ports on the DVNCC/CG have a different behavior than other VNswitch 900 Ethernet Ports. While the Activity LED does indicate either Receive or Transmit traffic, the rate of blinking or the intensity of the Activity LED should not be used to estimate the level of activity.

- The DVNCC/CG uses RJ-45 connectors and a *cross-wiring* scheme. Please note that this is opposite to the wiring scheme used on DVNXX products. Use the appropriate cable or adapter to make a viable connection.

- The DVNCC/CG supports standard-based 802.3 Pause control frames. The only configuration is to both transmit Pause frames and honor the reception of Pause frames. In addition, the remote node connected to the DVNCC/CG port should be set to support the reception of 802.3 Pause frames.

- AT this time, the DEChub ONE-MX (DEF1H-MB) docking station is not supported on the DVNCC/CG.

Gigabit (DVNCG, DVNGV)

Configuration Management for the Gigabit Ethernet uplink is similar to other Ethernet interfaces on the VNswitch. All the CLI commands that are documented in the Switch Management guide that apply to Fast Ethernet interfaces, also apply to the Gigabit Uplink. Web management is also consistent between regular 10MB/100Mb Ethernet interfaces and the Gigabit Ethernet interface. RMON counters and MIB-support for RMON is also consistent. There are however some important differences for Gigabit:

- Gigabit Interface Converter (GBIC): The 900CG/GV uses a GBIC removable connector. The proper GBIC must be inserted for the Gigabit interface to work properly. Please see the *Installation and Configuration* manual (EK-DVNGV-IN.A0) for a detailed description for proper installation and the different GBICs supported.

- Support for 802.3 Pause Control Frames. The 900CG/GV supports the standard based 802.3 Pause control frames and the default configuration is to both transmit Pause-frames and honor the reception of Pause-frames. While these default values can be changed, it is recommended that the default values NOT be changed or serious performance degradation may occur. In Addition, the performance of the 900CG/GV might be significantly degraded if the remote node on the Gigabit Ethernet fiber does not fully support the reception of 802.3 pause frames.

- Full-duplex/speed: The 900CG/GV does not allow the default value of full-duplex to be changed to half-duplex. Also, the speed may not be changed from the default 1000 Mbit speed.

- RMON: The RMON-MIB, CLI, and the RMON Web-management for the 900CG/GV is consistent with other VNswitch RMON interfaces. However, it should be noted that the mib/ CLI/WEB are all based on counters that 'wrap' at 2**32. At 1000 Mbit/sec speeds, these counters (particularly for octet counters) can wrap quickly. The '1800 sec' utilization values that are displayed with the following command are likely invalid:

  Monitor> interface statistics 1

  The time utilization is calculated from the RMON history table. The interval time can be modified using the WEB-interface to something more reasonable, or this entry can just be deleted altogether.

- The Gigabit Ethernet port can NOT be redirected to any backplane channels.

- The DVNCG/GV supports Gigabit Ethernet Single-Mode Fiber (SMF) and Gigabit Ethernet Multi-Mode Fiber (MMF) via an industry-standard Gigabit Interface converter (GBIC). No other versions of GBIC are certified for use with the DVNCG/GV.

DVNGV

The Gigabit daughter card has some Field upgradable FPGA, which on some earlier DVNGV should be upgraded. AFTER upgrading to V3.4 (or later), to determine if the FPGA should be upgraded, from the Config> prompt issue the following command sequence:

```
Config> dvt
PE2K_DVT_XX> version
```

The output from this command has a section similar to:

Field Programmable Device Versions (hardware-version / firmware-update-version)
    DB-0 image Versions:
        G-DAUGHTER Control CPLD= (1 / 1)
        G-DAUGHTER Receive CPLD= (1 / 1)
        G-DAUGHTER Xmit CPLD= (2 / 2)

If the hardware version does not match the firmware version the FPGA should be upgraded when convenient. NOTE that the upgrade process takes about 8 minutes and the Vnswitch does not function for the entire time. When ready to upgrade, from the PE2K_DVT_XX> prompt issue the command:

```
PE2K_DVT_XX> r 81
```

Answer (yes) at the prompt, and after the Vnswitch restarts, repeat the above ‘version’ command to make sure that the upgrade was successful.

DVNLL

The power consumption of the VNswitch 900LL module necessitates these configuration considerations:

- A maximum of six VNswitch 900LL (DNVLL) modules can be placed in a DIGITAL MultiSwitch 900 chassis. This configuration requires three 140- or 163-watt power supplies, or four 140- or 163-watt power supplies for N+1 redundancy.
- The DEChub ONE and DEChub ONE-MX docking stations support the VNswitch 900LL module. However, a comparison of the power label on the DEChub ONE docking station and the power requirements of the VNswitch 900LL modules may indicate otherwise.

ATM

This section lists ATM known restrictions and conditions specific to this release.

Setting Line Attenuation and Transmission Power For DS3 Interfaces

The **set line-building-out** command, which sets the line attenuation, and the **set transmission-power** command are specific to the modPMD implementations. They are required for the E3/T3 physical media specifications.

It is strongly recommended that you use the factory default for the transmission power. If the transmission power is not properly set, the physical media does not conform to the T3 or E3 specifications. The factory default for T3 lines is low, and for E3 lines, the default is high.

ATMswitch 900F Interoperability Problems

To configure the VNswitch 900 ATM modules in a DIGITAL MultiSwitch 900 chassis for LAN interconnect operations with the DIGITAL ATMswitch 900F, follow these special instructions:

- Use ATMswitch 900, Version 1.2, firmware.
- Configure the ATMswitch 900 to operate with UNI V3.2 signaling, using the command:

```
UNI -v31 -p<port #> -s<slot #>
```

If you have two ATMswitch 900F modules in your MultiSwitch 900 chassis and you want to use two ports from each to connect to four VNswitch 900 ATM modules on the backplane, you must use two different sets of ports from each ATMswitch 900F module. For example, if you want to use ports 1 and 3 in the first ATMswitch 900F module, then you must use ports 2 and 4 in the second ATMswitch 900F module.

If you choose the same port numbers in both ATMswitch 900F modules, only the first instance of that port number is available for a connection. Once you have an ATMswitch 900F port 1 connected to another module (VNswitch 900EA), that signal set is in use. The second instance of an ATMswitch 900F port 1 connection to another VNswitch 900 ATM module cannot be completed because that signal set is already in use on the backplane.

Nonsupport of Nonzero VP Values in ATM

Nonzero Virtual Path (VP) values are not supported for ATM in this release.

No FLOWmaster Support for ATM

VNswitch 3.0 firmware does not support FLOWmaster for ATM networks.

Configuring Link Parameters

ATM link failures may be caused by an out-of-revision VNswitch ATM adapter card when connected to a backplane, or by an out-of-revision ModPHY card when connected to the front panel. For backplane connections, be sure that the VNswitch ATM adapter card is at revision 1 or higher, and for front panel connections, be sure that the ModPHY card is at the revision as shown in the following table:

ModPHY Type	Part Number	Revision Level
E1 (2 Mb/s) UTP/ScTP	DAGE1-AA	C01 or higher
E3 (34 Mb/s) coaxial	DAGGE-AA	E01or higher
DS1/T1 (1.54 Mb/s) UTP/ScTP	DAGT1-AA	A01 or higher
DS3/T3 (44 Mb/s) coaxial	DAGGT-AA	A01 or higher
STS-3c (155 Mb/s) UTP/ScTP	DAGGU-AA	E01 or higher
OC-3 (155 Mb/s) MMF	DAGGM-AA	C01 or higher
OC-3 (155 Mb/s) SMF	DAGGS-AA	C01 or higher

ATM Bridge Tunnels and LECs

If the default LEC or bridge tunnel does not come up, one of the ATM logical interfaces may already be configured. Reset to factory settings (or at the Config> prompt, enter **clear ATM**) to use the default configuration.

ATM E1/T1 LAN Emulation Performance

LAN Emulation does not function for an ATM E1 or T1 ModPHY due to low line rates.

Split Path Routing Learning Problems with LAN Emulation

In certain IP network configurations, outbound packets between a source and a destination may take one path while inbound packets may take a different path on their return. This is known as Split Path Routing. If ATM LAN Emulation (LANE) is used to construct these paths, throughput reductions are seen. The throughput is reduced to anywhere from one packet per second to ten packets per second depending on the BUS (Broadcast and Unknown Server) rate throughput set in the VNswitch. The default rate is one frame per second and can be modified for each LEC interface using the **set max unknown frame count** and **set max unknown frame time** commands.

This situation occurs because traffic flow in these paths is unidirectional. In ATM LAN Emulation, the creation and subsequent use of a direct virtual circuit (VC) between LANE clients depends on traffic in both directions to prevent the VC from being aged out. With unidirectional traffic, the direct VC gets aged out (aging time = bridge aging time) and any subsequent traffic is sent through the BUS. The V1.0 LANE standard states that unknown unicast traffic on the BUS is limited to a maximum rate of ten frames per second. This problem can be avoided by reconfiguring the network to eliminate split paths.

ATM Link Down; Ethernet and VNbus UP

When the ATM PHY LEDs blink amber and ATM module LED blinks green (fatal ATM physical layer error), the ATM link is down, but the Ethernet and VNbus parts of the device are still functional. Make sure that the ATM daughter card is up to the ECO revision level 2.

ATM FDDI Bridge Tunnel to ATM Ethernet Bridge Tunnel Connection *(Not Recommended)*

You can manually configure an ATM FDDI bridge tunnel on one VNswitch module to connect to an ATM Ethernet bridge tunnel on another VNswitch module. While this configuration does allow the bridge tunnel to come up, the mismatch in tunnel types causes unwanted and unpredictable results. For this reason, we recommend that you do not attempt this configuration.

VNswitch 900FA Defaults to ATM Ethernet Bridge Tunnel

The plug-and-play values for the VNswitch 900FA module default to an ATM Ethernet bridge tunnel, even if two VNswitch 900FA modules are connected together. If an ATM FDDI bridge tunnel is desired, this must be manually configured.

Copying ATM Configurations to Another Module is Not Supported

VNswitch v3.0 does not support copying ATM configurations to another module. If you back up a module's configuration database to a server with the intention of later restoring it to a VNswitch module other than the one it was originally saved from, you must clear the ATM configuration (**Config> clear atm**) before saving the configuration.

ATM LEC Maximum Frame Size is Not Settable from the Web Management Interface

Do not use the VNswitch Web Management interface to set the ATM LEC maximum frame size. This causes the module to crash. You must use the CLI (**ATM/n LEC Config> set max_frame_size**) command to set ATM LEC maximum frame size.

Default DEC ATM MIB Values Changed

In the V3.0 release, the default values for the following DEC ATM MIB values are changed:

- adpReceiveBuffers = 0 (used to be 20)
- adpMaxReceiveBufferCounter = 20 (used to be 0)

RMON

Support for RMON Capable Ethernet Daughter Cards (10BaseT and 10BaseFL)

This release of the firmware supports new versions of the 12-port 10BaseT and 12-port 10BaseFL modules. These newer versions provide support for RMON Ethernet counters. To determine if your module has this new functionality, check the product revision against the following table:

Module	Revision
DVNEA-MX	H01 or higher
DVNEE-MA	F01 or higher
DVNEF-MM	F01 or higher
DVNEF-MX	F01 or higher
DVNEX-MX	J01 or higher
DVNLL-MA	B01 or higher

Note: All DVNCC and DVNGV devices include RMON capabilities.

If your unit is already installed, you can issue the Monitor> **list all** command. The hardware version of the VNswitch is displayed in the HW= field. The current Ethernet adapter card has the designator E and a newer RMON capable Ethernet card has an e designator. For example, the hardware field for an older version VNswitch 900EF might look like HW=70F0E1. This same VNswitch with a RMON capable card would look like this: HW=70F0e0. Note that the REV-id for the newer RMON capable card starts at zero (0).

Note: The power requirements for the modules are slightly higher to accommodate this new functionality. The power ratings are listed on the serial label for the product.

RMON Statistics Counter Differences (10BaseT or 10BaseFL)

RMON statistics conform to the RMON RFC, except for the following counter differences:

- The octet (etherStatsOctets) and packet (etherStatsPkts) counters do not include all error packets.
- The number of collisions in the collision counter (etherStatsCollisions) is understated.
- The undersize packet (etherStatsUndersizePkts) value is always listed as 0 and included as part of the fragment counter (etherStatsFragments).
- The oversize packet (etherStatsOversizePkts) value is always listed as 0 and included as part of the jabbers counter (etherStatsJabbers).
- The broadcast packet (etherStatsBroadcastPkts) value is always listed as 0 and included as part of the multicast counter (etherStatsMulticastPkts).

To display RMON statistics for interface 7 using the CLI; for example, perform the following command:
`Monitor>`**interface statistics 7**

RMON Restrictions with NetScout Manager

When using the NetScout Manager application, the following restrictions apply:

- Since the VNswitch supports only the four basic RMON groups, NetScout applications that rely on the other groups, such as Top N Talkers, return an error.
- The interface number supplied should be equal to the port number of the desired port. For example, to monitor Ethernet port 2 use interface number 2.
- The interface name returned using the **Test Agent** command is incorrect and should be ignored.
- When monitoring Ethernet Port 1, ignore NetScout Manager applications displaying information pertaining to an FDDI type device. Note that the percentage utilization is wrong.
- Using the Switch option results in the wrong set of ports. For example, the VNswitch 900EE ports will be defined as ports 2 through 25 instead of ports 1 through 24.
- The reporting of multicast packets in the short term history display has a significant variance between intervals. The smaller the interval time, the greater the variance. For accurate interval multicast packet counts, it is recommended to have an interval time of at least 10 minutes.

RMON Interface Utilization

The "line utilization" value may be incorrect for the next time period (default 30 seconds) if either:

- The Ethernet link changes state to Down, *or*,
- The **monitor> clear counters** command is issued.

Routing

Momentary Routing Performance Impact for Large IP Networks

It is possible that certain customers with large IP networks that have also enabled **IP Access Controls** on the VNswitch 900 may observe a momentary performance impact for routed traffic. If this performance impact does occur, it is most likely to happen during a Spanning Tree topology change in the network. This momentary slow down affects routing only and does **not** affect the bridging performance of the VNswitch. The most noticeable symptom is that a CLI session to the VNswitch is slower than normal.

Clearing the ARP Cache when Routing Enable/Disable is Changed

If you change your routing configuration from enabled to disabled, or from disabled to enabled, you must clear your ARP cache (`Config>` **clear arp**) and reconfigure the ARP configuration data.

BGP and RIP Route Limitations

You cannot import subnetted routes (routes with a mask that is not the natural mask for a Class A, B or C network) learned from the BGP protocol into RIP. Note that non-subnetted BGP routes can be imported into RIP without any problem.

*clear*VISN

IP Services Module Address

If you take the IP address of any module that is serving as the IP Services module for the DIGITAL MultiSwitch 900 and reassign it to the VNswitch module, you must remove (power-cycle) the other module (from which the IP address was taken) before this VNswitch module can assume the new address and operate as an IP Services module.

VNswitch Configuration on the VNbus

The primary mechanism for configuring the VNbus is *clear*VISN LAN-interconnect. If this management tool is unavailable, you may use the VNbus-AutoConnect feature of the DIGITAL MultiSwitch 900. The VNbus-AutoConnect feature should not be enabled when the *clear*VISN management tool is used to configure any LAN-interconnect.

Configuring Two Ethernet Ports on the Same VNswitch to the Same Backplane Channel

The VNswitch module does not support the mapping of two or more Ethernet ports on the same VNswitch module to the same MultiSwitch 900 backplane channel. In the VNswitch V3.0 release, if this operation is attempted, the second port remains in the DOWN state until it is configured to either a different backplane channel and manually enabled *or* to a front panel port and manually enabled.

Configuring VNswitch IP Addresses in a MultiSwitch 900

A problem occurs when IP addresses are configured for a VNswitch module installed in a Multiswitch 900 and the module is managed by the *clear*VISN MultiChassis Manager (MCM). The VNswitch module requires a restart when a new or changed IP address is entered.

Problems Configuring VNswitch 900 ATM Modules Using *clear*VISN MCM

Problems and solutions configuring the VNswitch 900 ATM modules with the *clear*VISN MultiChassis Manager (MCM) are as follows:

- **Problem:** Enabling a configured LEC does not result in the LEC being moved to the enable window. No error or warning indication is given.

  Solution: Verify that no LAN name conflict exists. The VNswitch 900EA does not allow more than one LEC with the same LAN name to be enabled. This includes multiple LECs with blank (i.e. default) LAN names.

- **Problem:** Enabling a configured FDDI bridge tunnel does not result in the tunnel being moved to the enable window. No error or warning indication is given.

  Solution: The FDDI bridge tunnel may recently have been configured without performing a required restart of the VNswitch 900 ATM module. To restart the module use MCM's RESET button in the VNswitch Switch Summary view, or use the **VNswitch 900EA Command Line Interface Restart** command from the configuration menu.

General

No Frame Interval Functionality

The “No Frame Interval” functionality is not supported in this firmware release.

Out-of Band Management (OBM)

The OBM baud rate cannot be set to 4800 when the VNswitch 900 is in the MultiSwitch 900. However, 4800 is a valid speed when the module is in a DEChub ONE docking station.

Clearing SNMP Configuration

When clearing SNMP configuration (using **Clear all** or **Clear SNMP** commands), you must restart the module for the clear to take effect.

Nonsupport of FDDI Port Redirection to Back of DEChub ONE-MX

This release does not support redirection of the FDDI port to the back of the DEChub ONE-MX docking station.

Duplicate MAC Address

For networks in which the same MAC address appears within more than one VSD (Duplicate MAC), it is necessary to configure the duplicated addresses as static duplicate MAC addresses using the command **BRIDGE config>set duplicated-address** <duplicate MAC address>.

Diagnostic Failures

If a hardware problem is detected during diagnostics, an entry is made in the diagnostic log. You can view this log from the web-based application (System Error Log window), from the Command Line Interface (CLI) using Monitor> **err-log list diag**, or using **Option 7** from the VNswitch 900 Installation Menu. If any diagnostic entries are present, it is a hardware failure and the module should be replaced.

Displaying Event Log Messages

When you display events using the Main> **events** command or indirectly using the Config>**set output console** command, be aware that you can retrieve event log messages only once. That is, once an event log message is displayed, it cannot be viewed again. Therefore, if you want to save event log messages for later analysis, save the display output using an appropriate method, such as logging/saving a terminal session.

Ping Packets Greater than 1500 Bytes

The VNswitch 900 module does not reply to ping packets that are greater than 1500 bytes.

Incorrect False Carrier Sense Counter Values

The false carrier sense counter (FCSR) is incorrectly incrementing for the VNswitch 900EX, VNswitch 900XX, VNswitch 900XA, and the VNswitch 900FX modules and, therefore, the counter value may not be valid.

Firmware Upgrades

Refer to Chapter 10 in the *DIGITAL VNswitch 900 Series Switch Management* guide for instructions on how to perform firmware upgrades, including a boot block software installation. You can use the Command Line Interface (CLI) or the web-based management application to upgrade the module.

If you are using the CLI to perform the upgrade, do not log out from the Main prompt during the upgrade.

Note: You must install the V3.0 boot block software first. Then install the Version 3.0 firmware.

The VNswitch modules do not support firmware upgrades using the DIGITAL MultiSwitch 900 Downline Upgrade menu option. You can perform firmware upgrades for the VNswitch modules using the CLI **reload** or **load remote** commands or using the *clear*VISN Flash Loader application or Web management.

Note: The **reload** and **load remote** commands rely on IP Host Services being configured.

If you are using an OpenVMS system and VMS UCX (V4.0 and earlier) as the TFTP load server for the firmware upgrade, the TFTP load may fail. As a workaround, convert the firmware image file format from Fixed-512 to Stream_LF record format.

Upgrading to VNswitch Version 3.0

Note: Before upgrading to the VNswitch Version 3.0 firmware, you need to install the boot block software, dvnbb301.bin. For more information, refer to the section titled *Firmware Upgrades*.

You can use the Command Line Interface (CLI), VNswitch Web management or *clear*VISN Flash Loader to upgrade the module.

The VNswitch 3.0 images can be upgraded from a V2.0.6 image only. If an attempt is made to upgrade a module that is running a version earlier than V2.0.6, the update fails with either a **timeout** or **out of memory** error message.

Upgrading from V1.1/V1.5 to V3.0

To upgrade from V1.1 or V1.5 to V3.0, you must upgrade to V1.6.2 first, then to V2.0.6. When upgrading to V1.6.2, please review the V1.6.2 Release Notes as they contain important information with respect to the upgrade process.

Note: Before upgrading to the VNswitch Version 3.0 firmware, you need to install the boot block software, dvnbb301.bin. For more information refer to the section titled *Firmware Upgrades*.

The VNswitch 3.0 images can be upgraded from a V2.0.6 image only. If an attempt is made to upgrade a VNswitch module that is running a version earlier than V2.0.6, either via the *clear*VISN Flash Loader program or via the CLI **load remote** or **reload** commands, the update fails with either a timeout or out of memory error message.

Clearing the ARP Cache

You must clear your ARP cache (`Config>` **clear arp**) and reconfigure the ARP configuration data at the completion of your upgrade, unless routing was enabled prior to the upgrade and it remains enabled at the completion of the upgrade.

Incorrect Boot Block Upgrade Procedure

Chapter 10 in the *DIGITAL VNswitch 900 Series Switch Management* guide incorrectly states that upon the completion of a boot block upgrade, the yellow and green LEDs remain lit until you power-cycle the module. With boot block version V3.0, you do not have to power-cycle the module, the yellow and green LEDs turn off automatically.

Documentation

The following documentation supports the VNswitch Version 3.0 firmware release:

- *DIGITAL VNswitch 900 Series Technical Overview*
- *DIGITAL VNswitch 900 Series Switch Management*
- *DIGITAL VNswitch 900 Series Router Management*

These documents exist in Adobe Acrobat online readable and printable (PDF) format on the documentation CD-ROM that ships with the module.

VNswitch MIB Support

The VNswitch supports the following MIBs. If a MIB is defined in more than one RFC, the supported RFC is listed first and the other RFCs are listed on a separate line. The MIB handlers do not support SNMP set requests unless otherwise noted.

MIB	RFC/GROUP
mib-2	iso(1).org(3).dod(6).internet(1).mgmt(2).mib-2(1)
	rfc-1213 rfc-1158 -> rfc-1213
	system(1) (set) interfaces(2) ifAdminStatus(7) (set)
	at(3) ip(4)
	ipDefaultTTL(2) (set) icmp(5)
	tcp(6)
	udp(7)
	egp(8)
	transmission(10) (interface mibs) snmp(11)
ethernet	.mib-2(1).transmission(10).dot3(7)
	rfc-1643 rfc-1284 -> rfc-1398 -> rfc-1623 -> rfc-1643
	dot3StatsTable(2) dot3CollEntry(5)
	dot3Tests(6) <oid pointers> dot3Errors(7) <oid pointers> dot3ChipSets(8) <oid pointers>
fddi	.mib-2(1).transmission(10).fddi(15).fddimib(73)
	rfc-1512 rfc-1285 -> rfc-1512
	fddimibSMT(1) fddimibMAC(2) fddimibPATH(3) ddimibPORT(4)
ds1	.mib-2(1).transmission(10).ds1(30)
	rfc-1406
	dsx1ConfigTable(6) (set) dsx1CurrentTable(7) dsx1IntervalTable(8) dsx1TotalTable(9) dsx1FarEndCurrent(10) Not supported dsx1Interval(11) Not supported dsx1total(12) Not supported dsx1FracTable(13) Not supported

MIB	RFC/GROUP
ds3	.mib-2(1).transmission(10).ds3(30) rfc-1407 dsx3ConfigTable(5) (set) dsx3CurrentTable(6) dsx3IntervalTable(7) dsx3TotalTable(8) dsx3FarEndConfigTable(9) Not supported dsx3FracTable(13) Not supported
rmon	.mib-2(1).rmon(16) rfc-1757 statistics(1) (set) history(2) (set) alarm(3) (set) event(9) (set)
mau	.mib-2(1).snmpDot3MauMgt(26) draft-ietf-hubmib-mau-mib-03.txt dot3RpMauBasicGroup(1) Not applicable dot3IfMauBasicGroup(2) dot3BroadMauBasicGroup(3) Not applicable dot3IfMauAutoNegGroup(5)
sonet	.mib-2(1).transmission(10).sonetMIB(39) rfc-1595 sonetObjects(1) sonetMedium(1) sonetSection(2) sonetSectionIntervalTable(2) sonetLine(3) sonetLineIntervalTable(2) sonetFarEndLine(4) Not supported sonetObjectsPath(2) sonetPath(1) sonetPathCurrentTable sonetPathIntervalTable sonetFarEndPath(2) Not supported sonetObjectsVT(3) Not supported sonetVT(1) Not supported sonetFarEndVT(2) Not supported

MIB	RFC/GROUP
bridge	.mib-2(1).dot1dBridge(17) (multiple spanning tree support) rfc-1493 rfc-1286 -> rfc-1493 & rfc-1525 dot1dBase(1) dot1dStp(2) (set) dot1dSr(3) Not applicable dot1dTp(4) (set) not implemented: dot1dStatic destination address filtering dot1dStaticTable(1) traps
interfaces	.mib-2(1).ifMIB(31).ifMIBObjects(1) rfc-1573 ifStackTable(2)
digital	.private(4).enterprises(1).dec(36).ema(2) mib-extensions-1(18)
elan	elanext(1).efddi(1) (set) elanext(1).ebridge(4) (set) ebrIfSpanTable Not supported ebrTwoPortStatic Not supported ebrTwoProtoFilt Not supported ebrNTP Not supported
hub	dec_elan_vendor_mib_v30.mib decHub900(11).pubCommon(2) pcomHub(2) pcomLed(3) (set) pcomLoad(4) (set) pcomSnmpAuth(5) pcomSnmpAuthTrap(1) (set)
atm	.mib-2(1).atmMIB(37) rfc-1695 atmInterfaceConfTable(2) (set) atmInterfaceDs3PlcpTable(3) atmInterfaceTCTable(4) atmTrafficDescrParamTable(5) atmVplTable(6) atmVclTable(7) atmVpCrossConnectIndexNext(8) Not supported atmVpCrossConnectTable(9) Not supported atmVcCrossConnectTable Not supported aal5VccTable
comet	comet-mib(2) cinterface(1)

MIB	RFC/GROUP
vlan	vlan_v1.mib pe2000(33).bridgeGroup(1) (set) bridgeGroupPortTable(4) bridgeGroupNameTable(5) bridgeGroupPeBusTagTable(7) bridgeGroup atomics
proteon	.private(4).enterprises(1).dec(36).ema(2). mib-extensions-1(18).cometBROUTERS(20).proteon-mib(1) no rfc - proteon mib text fully supported including sets admin(1).oid(1) admin(1).status(2) admin(1).els(3) admin(1).xface(4) admin(1).private(5) (no documentation) nvram(1) reset(2) xface(2) proto(3)
atm	dec_atm.mib atmExpand(17) ad(1) dxatm(2) Not supported
atm bridge tunnel	decAtmBridgeTunnel.mib decAtmBridgeTunnel(28)
atm lec	.private(4).enterprises(1).atmForum(353). atmForumNetworkManagement(5).atmfLanEmulation(3) leClientMIB(1).leClientMIBObjects atmLecClient.mib lecConfigTable(1) (set) lecStatusTable(2) Not supported in v1.0 lecMappingTable(3) Not supported in v1.0 lecStatisticsTable(4) Not supported in v1.0 lecServerVccTable(5) Not supported in v1.0 lecAtmAddressTable(6) Not supported in v1.0 lecMacAddressTable(7) Not supported in v1.0 lecRouteDescrTable(8) Not supported in v1.0 leArpTable(9) Not supported in v1.0 leRDArpTable(10) Not supported in v1.0

MIB	RFC/GROUP
icom	internal(0).intCommon(1)
	int-common.mib
	The icom MIB objects are normally only visible to the MAM on the chassis backplane serial line
	icomHlap(1) icomRoot(2) Table not populated icomHub(3) icomStatus(4) icomTrap(5) Not supported icomIps(6) icomEnviron(7) icomPower(8)
	icomIntProtInstrumentation(9) Not supported icomBp(10)
	icomBpTotalConfigChanges(1) icomBpIfNumEntries(2) icomBpIfTable(3) icomBpPortDescrTable(4) Missing in walk icomBpIfSubtypeNumEntries(5) icomBpIfSubtypeTable(6) icomBpSignalSetNumEntries(7) icomBpSignalSetTable(8) icomBpConnNumEntries(9) icomBpConnTable(10) icomSlot(11) icomEntity(12) icomRemotePoll(14) Not supported icomLigo(15)
	icomLast Not supported

IGMP Snooping

Overview

IGMP Snooping is a bridging mechanism which limits traffic to only LAN segments which are interested in seeing that traffic. IGMP Snooping is not IGMP, nor is it an extension of IGMP. Its function is to limit traffic based on knowledge gathered by watching (for example "Snooping") IGMP messages.

This document is a supplement to the VNSwitch Management Guide and assumes an understanding of the following:

- Bridging
- VNSwitch VSDs
- IP Multicast
- IGMP

IGMP Snooping Operation

Purpose of IGMP Snooping

The goal of IGMP Snooping is to limit the multicast traffic to only those LAN segments that are actually interested in hearing the messages. Normally, IP Multicast traffic is flooded throughout the whole LAN. This is wasteful since the traffic needs only to reach the lan segments that have joined the particular Multicast Group. The interested parties consist of two groups: segments which contain clients that have joined the group and multicast routers which may need to forward the packet to other domains.

Basic Operation

IGMP Snooping watches IGMP traffic on all interfaces. It also watches protocol-specific MC traffic. IGMP Snooping keeps track of the following traffic:

- * IGMP General Queries
- * IGMP Group-specific Queries
- * IGMP Membership Reports ("Joins" and "Leaves")
- * DVMRP, PIM, MOSPF, and CBT "Hello" messages

IGMP Snooping supports IGMP V1 and IGMP V2.

IGMP Snooping does not send multicast traffic to interfaces which 'have no interest' in multicast traffic for that MC group.

To implement IGMP Snooping, a bridge would need to know the location of all multicast routers and of clients that have joined a specific group. Normal Bridging/VLAN filters are applied before transmission on any of the indicated links.

Listening Mode

When IGMP Snooping is first enabled on a VSD, multicast traffic must be flooded. This is to ensure that any existing multicast members on the LAN will have their traffic continued. This is Snooping's 'Listening Mode'.

Once the Listening Mode is complete, all multicast members have been learned, and IGMP Snooping begins.

Automatic Detection of Multicast Routers

IGMP Snooping automatically detects multicast routers running the following multicast routing protocols:

- * DVMRP
- * PIM-SM
- * PIM-DM
- * MOSPF
- * CBT

Each of these multicast routing protocols has a mechanism by which the multicast router identifies itself to the network. Typically, these are called Hello Messages. Hello messages are sent out periodically, and the VNSwitch watches for these messages to detect the presence of the multicast router. Once the router(s) are detected, the ports on which routers were detected are included in the set of ports to which multicast traffic will be sent.

Snooping in Networks with No IP Multicast Routers

IGMP Snooping relies on an IGMP Querier to send queries to the LAN so that individual LAN segments will send joins and/or leaves. In a network with no IP Multicast routers, there is no IGMP Querier present. Therefore, if a VNswitch detects that there is no Querier present, it assumes IGMP Querier functionality. Only General Queries (not Group Specific Queries) are sent out. This ensures that Snooping will see the Join messages and Leave messages. When/if IGMP Queries are received from another device, the VNSwitch stops acting as IGMP Querier.

In order to take on the role of the IGMP Querier, the VSD must have an IP address assigned to it. If there is no IP address assigned to the VSD and there is no IGMP Querier, then the VNSwitch will not send Queries, and IP multicast will be broken. Therefore, in networks with no IP Multicast router, IGMP Snooping should not be enabled on VSDs with no IP address.

Transmission of Packets

Transmission of IGMP Control Traffic

The following 256 IP addresses are used for control traffic and are reserved:

224.0.0.0 -> 224.0.0.255

IGMP Snooping is not performed on frames destined to these addresses i.e. traffic to these addresses is flooded normally. Each of these 256 addresses has 31 other IP addresses with the exact same Ethernet address. This represents a total of 8192 IP multicast addresses upon which IGMP Snooping is not performed.

These 8192 IP multicast groups correspond to the following 256 Ethernet group addresses:

01-00-5E-00-00-00 --> 01-00-5E-00-00-FF

These 8192 IP multicast addresses are:

224.0.0.0 -> 224.0.0.255	224.128.0.0 -> 224.128.0.255
225.0.0.0 -> 225.0.0.255	225.128.0.0 -> 225.128.0.255
226.0.0.0 -> 226.0.0.255	226.128.0.0 -> 226.128.0.255
227.0.0.0 -> 227.0.0.255	227.128.0.0 -> 227.128.0.255
228.0.0.0 -> 228.0.0.255	228.128.0.0 -> 228.128.0.255
229.0.0.0 -> 229.0.0.255	229.128.0.0 -> 229.128.0.255
230.0.0.0 -> 230.0.0.255	230.128.0.0 -> 230.128.0.255
231.0.0.0 -> 231.0.0.255	231.128.0.0 -> 231.128.0.255
232.0.0.0 -> 232.0.0.255	232.128.0.0 -> 232.128.0.255
233.0.0.0 -> 233.0.0.255	233.128.0.0 -> 233.128.0.255
234.0.0.0 -> 234.0.0.255	234.128.0.0 -> 234.128.0.255
235.0.0.0 -> 235.0.0.255	235.128.0.0 -> 235.128.0.255
236.0.0.0 -> 236.0.0.255	236.128.0.0 -> 236.128.0.255
237.0.0.0 -> 237.0.0.255	237.128.0.0 -> 237.128.0.255
238.0.0.0 -> 238.0.0.255	238.128.0.0 -> 238.128.0.255
239.0.0.0 -> 239.0.0.255	239.128.0.0 -> 239.128.0.255

Transmission of Non-Control Multicast Traffic

For all other multicast addresses, the multicast packets are bridge-forwarded only to ports which contain either a multicast router and/or a group member.

The one exception to this is Join messages. For each Query interval, only the first Join for a group is forwarded to ports with multicast routers. Subsequent Joins for that group in the same Query interval are recorded by the VNSwitch, but they are not forwarded. This is necessary because flooding Joins suppress other hosts on other LAN segments from sending their own Joins, which Snooping must see in order to work.

Example of IGMP Snooping

The following is an example scenario to illustrate the operation of IGMP Snooping[SWR1].

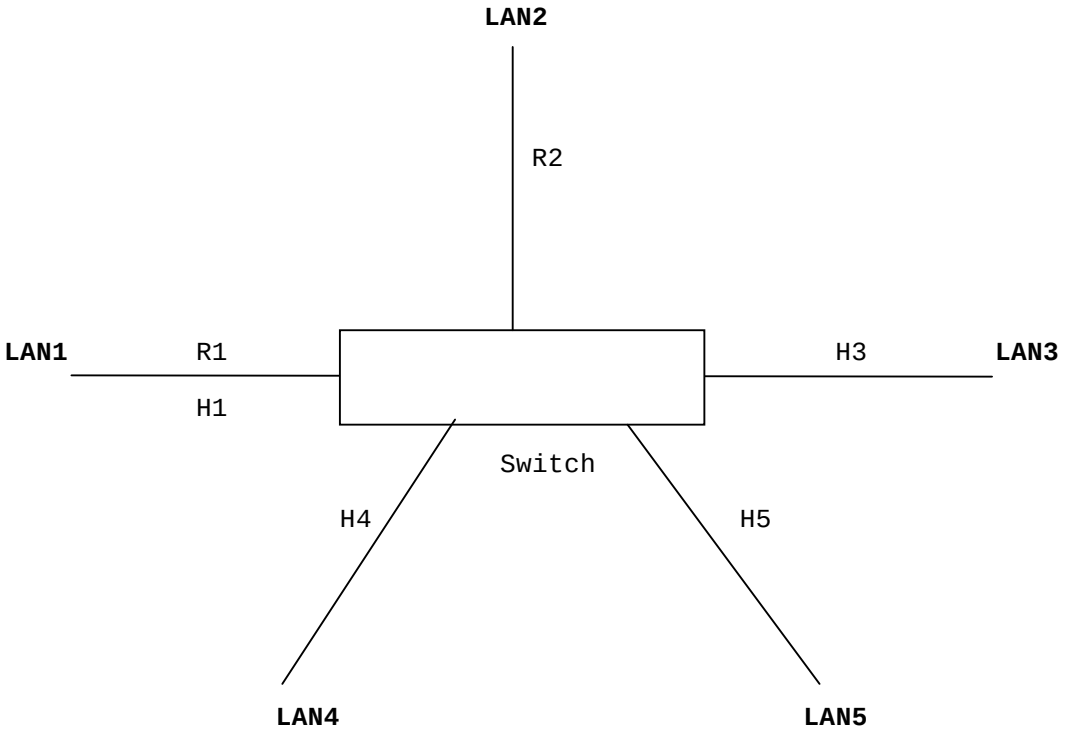

- * LAN1 through LAN5 are all on the same VSD
- * R1 and R2 have been detected as multicast routers
- * R1 is the IGMP Querier
- * No hosts have joined groups

1) R1 issues a General Query
Switch floods this out all ports. This is flooded because it was issued to a 224.0.0.X address.

2) H1 joins group 224.1.1.1
Switch bridges the Join to LAN2 only. This is because LAN2 contains a multicast router.

3) H3 sends a data packet to 224.1.1.1
Switch bridges packet to LAN1 and LAN2. This is because LAN1 contains a client and a multicast router, and LAN2 contains a multicast router.

4) H4 joins group 224.1.1.1
Switch records reception of this packet but does not forward it to any segment. This is because there is no lan segment that needs to know about this Join. R1 and R2 have already heard about this group, and there is no other segment on this group.

5) H4 joins group 224.2.2.2
Packet is bridged to LAN1 and LAN2. This is because both LAN1 and LAN2 have multicast routers.

6) H4 sends a data packet to 224.2.2.2
Packet is bridged to LAN1 and LAN2.

7) H5 joins group 224.2.2.2
Switch does not forward packet to any segment.

8) H4 sends a data packet to group 224.2.2.2
Packet is bridged to LAN1, LAN2 and LAN5.

IGMP-SNOOPING on a Global VSD

When IGMP-SNOOPING is enabled on a Global-VSD (for example, when running in a MultiSwitch 900 over the VNBus), it is important to enable IGMP-SNOOPING on all other Vnswitches that are part of the Global VSD. It is possible when the same IP-Multicast address is in use on multiple Global-VSDs that some IP-Multicast traffic may flood unnecessarily over the VNbus. If IGMP-SNOOPING is enabled on the other Vnswitches, this unneeded flooding is just harmlessly dropped.

IP-Multicast Address Filters

Normally if IGMP-SNOOPING is enabled, there is no need to configure Address Filters that relate to IP-Multicast, for example 01-00-5e-XX-XX-XX. Static (or Temporary) filters should NOT be used as they interfere with the proper operation of IGMP-SNOOPING. However, if there is need for an address filter, for example to limit flooding for an IP-Multicast address in a VSD that does not have IGMP-SNOOPING enabled, a PERMANENT filter should be used. The effect of the address filter is to restrict the ports that IGMP-SNOOPING can flood to.

When an address filter is created, the original Allowed Ports can be seen via the command, 'list bridge-filter' (in the IGMPS subsection of Monitor) or via the WEB IGMP Snooping page. What is seen in the Forwarding Database view or the Address Filter view under BRIDGING is the current list of Allowed Ports, which is a logical ANDing of the Configured Allowed Ports and what IGMP-SNOOPING has determined are the current flooding ports.

When creating an IP-Multicast Address Filter, it is important to NOT enable 'Rate Limiting'. Rate limiting works on a switch-wide basis, so that every IP-Multicast packet that is rate-limited counts against the total Rate Limit of the switch (default 400 pkts/sec). A single IP-Multicast data stream that is rate-limited might cause normal broadcast traffic to be rate limited.

Normally, the Allowed Port mask is updated by IGMP-SNOOPING only for IP-Multicast addresses that it has seen an IGMP-Report(Join) for. If an address filter is created that has no IGMP-Report activity, it is possible for the Allowed Port Mask to be stale and not reflect all IGMP-SNOOPING changes; for example, the detection of a new IP-Multicast Router.

Changes to the CLI

Config IGMPS Sub-section

The IGMP Snooping section of the CLI Config is entered with the command IGMPS, for example:

Config> IGMPS

IGMPS Config>

The IGMPS Config subsystem has the following IGMPS-specific commands:

- * Set
- * Enable
- * Disable
- * Add
- * Delete
- * List

Note: Most of these commands require a VSD number in the argument list. These commands accept VSD numbers for VSDs that do not exist. Configuration having VSD numbers for VSDs that do not exist has no effect until the VSD is created.

Set Command

The set command is used to create/modify IGMP Snooping VSD configuration entries.

Each IGMP Snooping VSD entry has five fields:

Status	Whether IGMPS is enabled on that VSD. If disabled, multicast traffic is flooded normally. [Default value: Disabled]
Fast Query	Whether Fast Query is enabled on that VSD. If enabled, the VNSwitch sends out IGMP Queries every 11 seconds while in Listening Mode. [Default value: Enabled]
Querying Interval	Amount of time (seconds) to wait (after seeing or sending a Query) before the VNSwitch sends the IGMP Query. Setting this to 0 disables the sending of Queries. Range of values: 11 --> 65520 [Default value: 385]

Listening Interval Amount of time (seconds) after IGMP Snooping is first enabled on the VSD (either at boot time or through CON or Web enabling) to listen to traffic before IGMP Snooping begins. Multicast traffic is flooded normally during this window. This value should be large enough to allow all hosts to join their multicast groups and allow the multicast routers to send their Hello messages. The Listening Interval begins when IGMPS is first enabled (either through enabling it through configuration or from being enabled at power-up).
Range of values: 0 --> 65535
[Default value: 395]

Router Timeout Interval Amount of time (seconds) before timing out a multicast router learned via a Hello message. Upon reception of a Hello message from a multicast router, this value is compared against the Holdtime value inside the Hello message. The larger of the two is taken as the amount of time before timing out this port as having a multicast router. The Router Timeout Interval applies to all ports in the VSD.
Range of values: 0 --> 65535
[Default value: 0]

The Status field is set through the command, **enable** (explained below).

The other four fields are set through the **set** command:

```
IGMPS Config>set ?
ALL parameters
LISTENING interval
QUERYING interval
ROUTER interval
FAST Query Status
```

When the 'ALL' option is used, the user is prompted to enter the Vsd, Router, Listening, Querying, and Fast Query values, for example:

```
IGMPS Config>set all
VSD number []? 1
Router Timeout Value [0]?
Listening Interval [395]?
Querying Interval [385]?
Enable Fast Query? (Yes, No): [Yes]
```

When any of Listening, Querying, Router Timeout or Fast Query is entered, the user is prompted to enter that field, for example:

```
IGMPS Config>set listening
VSD number []? 2
Listening Interval [395]? 240
```

Enable Command

This enables IGMPS on a VSD:

Syntax: **enable** <vsd_num>

Like the **set** command, snooping VSD IGMPS entries can be created with the enable command.

Snooping VSD entries created with the **enable** command will have the Status field set to Enabled. The other fields will take on their default values.

The enable command takes effect immediately; a restart is not necessary.

Disable Command

The **disable** command disables IGMPS on a VSD:

Syntax: **disable** <vsd_num>

The disable command never creates or deletes a VSD entry; it simply sets the status field of an existing Snooping VSD entry to disabled.

The disable command takes effect immediately; a restart is not necessary.

Add Command

IGMP Snooping automatically detects multicast routers running DVMRP, PIM, MOSPF, and CBT. Routers that are running other multicast protocols can have their presence known through configuration.

The add command is used to configure the presence of these multicast routers.

The add command may use either the Mac Address of the multicast router or a Portlist for the multicast router.

a. Using the Mac Address of the Router

Syntax: add macaddr <Vsd_num> <Mac_addr> <Num_addrs>

Vsd_num:	The VSD
Mac_addr:	The unicast MAC address of the router [Entered as XX-XX-XX-XX-XX-XX]
Num_addrs:	The number of addresses in the range Range of values: 1 --> 255.

During a window in which that Mac addresses has not been learned (or has timed out), this configuration entry will have no effect.

When that mac address is learned/known, any multicast traffic to be forwarded will be sent to the port on which that mac address was learned (along with any other port(s) to which the traffic must go).

The following example shows how to configure the switch to assume that there is a multicast router on VSD #3 with a mac address of 08-00-2B-A2-6F-DE:

```
IGMPS Config>add macaddr
VSD Number [1]? 3
Please enter a range of MAC addresses
Base MAC Address (in 12-digit hex) []? 08-00-2B-A2-6F-DE
Number of MAC Addresses in the range (255 Max) [1]? 1
```

Since the mac address of the router is known, the Num_addrs field can simply be left at the default value, 1.

When the exact mac address of the router is unknown, the Num_addrs field is used to define the size of a range of addresses. The following example shows how to configure the switch to assume that there is a multicast router on VSD #6 with a Mac address somewhere in the range of 08-00-2B-A2-6F-00 and 08-00-2B-A2-6F-0F (inclusive):

```
IGMPS Config>add macaddr
VSD Number [1]? 6
Please enter a range of MAC addresses
Base MAC Address (in 12-digit hex) []? 08-00-2B-A2-6F-00
Number of MAC Addresses in the range (255 Max) [1]? 16
```

b. Using a Portlist

Syntax: add portlist <vsd_num> <port_list>

Vsd_num:	The VSD
Port_list:	The list of ports. This can be a range of ports, or "none" or "all". The range of acceptable ports is 0 --> 254. Port 0 is for the VNBus. Choosing "all" selects ports 0 through 254.

The following example shows how to configure the switch to realize that there are multicast router(s) on VSD #4, but only on ports 7 and 8, for example:

```
IGMPS Config> add portlist

    VSD Number [1]? 4
    Allowed port numbers, "None", or "all" []? 7,8
```

The Port_list can contain ports that are not in the VSD. However, this has no effect until those ports are added into the VSD.

Note: Once ports are added to a portlist entry, they cannot be removed (for instance 'add portlist' commands can only augment the existing list of ports). Removal of a port from a portlist entails deleting the portlist entry and re-entering the desired set of ports into a new portlist entry. The 'delete portlist' command is described below.

c. Which Method to Choose

Using the Mac Address has the advantage of flexibility: a router can be moved around from port to port and Bridging/Snooping will keep up with the movement. It also has the advantage that the Mac address, in actuality, need not be a multicast router. It could be any unicast address that wants to send and/or receive multicast traffic.

Using the Portlist has the advantage of ease of use: the Mac Address does not need to be looked up and/or gleaned from anywhere. All that needs to be known is the set of port(s) on which the multicast router(s) reside.

The following is true for both methods:

For VSDs that span VNSwitches in a MultiSwitch 900, defining a router (using either the portlist or the mac address) on a VNSwitch does not ensure that multicast traffic from other VNSwitch(es) will come to it unless port 0 is added into the router portlist on those other VNSwitch(es).

In other words, adding port 0 into a portlist informs that VNSwitch there is at least one multicast router somewhere else on the VNBus.

Delete Command

Syntax: Delete [vsd | macaddr | portlist]

a. Delete VSD

Syntax: delete vsd <vsd_num>

Using delete vsd deletes a Snooping VSD entry (Status, Fast Query Status, Query Interval, Listening Interval, Router Interval). However, any multicast router configuration (for example, Mac Address or Portlist entries) for that VSD is left untouched. For example:

IGMPS Config> delete vsd 1

b. Delete MacAddr

Syntax: delete macaddr <vsd_num> <mac_address>

This deletes the macaddr entry for the <vsd_num>, <mac_address> combination, for example:

IGMPS Config> delete macaddr 2 01-2f-45-6a-7d-5c

c. Delete Portlist

Syntax: delete portlist <vsd_num>

This deletes the portlist entry for <vsd_num>, for example:

IGMPS Config> delete portlist 1

List Command

The list command has the following syntax:

list [routers | vsd | all]

a. List Routers

Syntax: list routers [<vsd_num> | all]

This lists the multicast routers that are configured in IGMP Snooping. There are two tables: mac addresses and router portlists.

The 'list routers <vsd_num>' command displays entries for vsd_num. For example:

IGMPS Config> list routers 1

CONFIGURED MAC ADDRESSES

MAC Address	Addresses in range	VSD
0f-93-4f-13-6f-3c	1	1

CONFIGURED ROUTER PORTLISTS

VSD	Portlist
1	6

The 'list routers all' command lists all entries for all VSDs, for example:

IGMPS Config> list routers all

CONFIGURED MAC ADDRESSES

MAC Address	Addresses in range	VSD
0f-93-4f-13-6f-3c	6	4
0c-02-e7-22-d4-1c	1	6

CONFIGURED ROUTER PORTLISTS

VSD	Portlist
1	6
2	7

b. List VSD

Syntax: list vsd [<vsd_num> | all]

'list vsd <vsd_num>' lists all of the VSD information (Status, Query Interval, Listening Interval, Router Interval) for vsd_num, for example:

IGMPS config>list vsd 1

VSD	IGMPS Status	Fast Query Status	Query Int. (seconds)	Listening Int. (seconds)	Router Int. (seconds)
1	Enabled	Enabled	11	55	40

The 'list vsd all' command lists all of the VSD information for all VSDs, for example:

IGMPS config>list vsd all

VSD	IGMPS Status	Fast Query Status	Query Int. (seconds)	Listening Int. (seconds)	Router Int. (seconds)
1	Enabled	Enabled	11	55	40
2	Enabled	Enabled	385	55	0
3	Enabled	Enabled	385	55	0
4	Enabled	Disabled	385	55	0

c. List All

This shows a combination of all router and all vsd information, for example:

IGMPS config>list all

CONFIGURED MULTICAST ROUTERS

MAC Address	Addresses in Range	VSD
00-00-F8-8F-FC-4D	250	1
00-00-F8-6C-C3-56	251	3

CONFIGURED ROUTER PORTLISTS

VSD	Portlist
1	9

CONFIGURED PARAMETERS

VSD	IGMPS Status	Fast Query Status	Query Int. (seconds)	Listening Int. (seconds)	Router Int. (seconds)
1	Enabled	Enabled	11	55	0
3	Disabled	Enabled	385	55	0
4	Enabled	Disabled	385	55	0

IGMPS config>

Clear Command

The clear command is entered from the Config sub-section (not the Config>IGMPS subsection). This command is used to clear all configuration information for IGMPS.

```
Config>clear igmps
You are about to clear all configuration information
*** WARNING *** This will invoke an automatic RESTART
Are you sure you want to do this (Yes or [No]): no
aborted
```

Maximum Bridging Entries

The maximum number of bridging entries is set to 100. This means that IGMP Snooping can work with a list of up to 100 MAC multicast group addresses. When Snooping's list of Mac addresses exceeds 100, multicast traffic is flooded normally for addresses beyond that upper limit, and their flood masks appear as 'all'.

Because 32 IP addresses map to one multicast MAC address, Snooping potentially has an upper limit of 3200 IP addresses.

Monitor IGMPS Sub-section

The IGMP Snooping section of the CLI Monitor is entered with the command **IGMPS**, for example:

Monitor> IGMPS
IGMPS>

The IGMPS Monitor subsystem has the following IGMPS-specific commands:

- * List
- * Delete

List Command

Syntax: list [vsd | groups | bridge-filter | all]

a. List Vsd

Syntax: list vsd [<vsd_num> | all]

The list vsd <vsd_num> lists

This command lists VSD Snooping information for that vsd_num. It does not list information for the VSD entry if the VSD has not been defined, for example:

```
IGMPS>list vsd 1
```

---------- IGMPS Information for VSD #1 ----------

IGMPS Status:	Enabled	MC-Router Portlist:	2,3,9
FAST QUERY Status:	Enabled		
Query Interval:	120	Query Timer:	9
Listening Interval:	55	Listening Timer:	0
Last IGMP Querier:	192.168.103.101	Max Response Timer:	8
Total MC-MAC Addr:	4	Total MC-IP Addr:	7

Router Port	Detection	Router Timers Initial, Remaining
2	Configured in Portlist	
3	Query from 192.168.103.101	135, 79
3	Hello(DVMRP) from 192.168.103.101	200, 31
9	00-00-79-9D-A8-04 in Range	

MC-Router Portlist shows an aggregation of the ports which:

a) have been configured as a multicast router (either with a macaddr or with a portlist) and

b) the ports which have had multicast routers learned on them (either through Hello messages or Queries).

Query Timer shows the amount of real time before the next Query is sent.

Listening Timer shows the amount of real time before the Listening Interval is complete.

Last IGMP Querier shows the IP address of the last sender of an IGMP Query. This can be "None", "Self", or <IP Address>.

Max Response Timer shows the amount of real time for IGMP hosts to respond in this Query Interval.

Total MC-MAC Addr shows the total number of multicast mac addresses for this VSD.

Total MC-IP Addr shows the total number of multicast IP addresses for this VSD.

The display for Router Ports shows how each router was detected. Detection can be via Queries, Hellos, Configured Portlists, and Configured Mac Addresses. For Queries and Hellos, the source IP address is displayed, as are the initial and remaining holdtime values.

The 'list vsd all' command lists the same information, but it lists for all VSDs.

b. List Groups

Syntax: list groups [vsd | mac | ip]

This command lists IP multicast group information.

List Groups Vsd

Syntax: list groups vsd <vsd_num>

This command lists the multicast groups (displayed with the MAC address) that have joined.

The 'list groups vsd <vsd_num>' command shows this information for vsd_num, for example:

IGMPS>list groups vsd 1

MAC Address	VSD	#IP	Join Ports	Flood Ports
01-00-5e-10-10-10	1	3	4-6	4-6, 10-13
01-00-5e-11-11-11	1	1	18	10-13, 18

Total MAC Addresses: 2
Total IP Addresses: 4

The **# of IP Add**r column shows the number of IP addresses (i.e. Multicast Groups) that have joined. This ranges from 1 to 32.

The **Join Ports** column shows the port on which the group was joined.

The **Flood Ports** column shows the allowed ports for this group.

The 'list groups vsd all' command shows the same information but for all VSDs, for example:

IGMPS>list groups vsd al1

MAC Address	VSD	#IP	Join Ports	Flood Ports
01-00-5e-10-10-10	1	3	4-6	4-6, 10-13
01-00-5e-10-10-10	4	1	20-22	20-22, 28
01-00-5e-10-10-22	5	1	30	30,32

Total MAC Addresses: 3
Total IP Addresses: 5

List Groups Mac

Syntax: list groups mac <mac_addr>

This lists all multicast groups that have joined mac_addr.

Mac_addr is the Mac multicast group

IGMPS>list groups mac 01-00-5e-10-10-10, for example:

MAC Address	VSD	#IP	Join Ports	Flood Ports
01-00-5e-10-10-10	1	3	4-6	4-6, 10-13

Joined IP Address(es)

224.16.16.16 225.16.16.16 226.16.16.16

MAC Address	VSD	#IP	Join Ports	Flood Ports
01-00-5e-10-10-10	4	1	20-22	20-22, 28

Joined IP Address(es)

224.16.16.16

Joined IP Address(es) lists all IP multicast addresses that have joined on the MACAddr/VSD combination.

List Groups Ip

Syntax: list groups ip <ip_addr>

This does the same as 'list groups mac mac_addr', but takes an IP multicast group address instead.

IGMPS>list groups ip 225.16.16.16, for example:

MAC Address	VSD	#IP	Join Ports	Flood Ports
01-00-5e-10-10-10	1	3	4-6	4-6, 10-13

Joined IP Address(es)

224.16.16.16 225.16.16.16 226.16.16.16

MAC Address	VSD	#IP	Join Ports	Flood Ports
01-00-5e-10-10-10	4	1	20-22	20-22, 28

Joined IP Address(es)

225.16.16.16

c. List Bridge-Filter

The LIST BRIDGE-FILTER command is used to list the Permanent Bridge Filters configured that are IP-Multicast Addresses. IGMP-SNOOPING uses these address filters to limit which ports a particular IP-Multicast address will be flooded out. IGMP-SNOOPING displays the flood mask for any IP-Multicast address based upon IGMP-SNOOPING updates, for example, IGMP-Joins, and so forth. The actual flood mask for a given address is the IGMP flood mask logically ANDed with a Permanent Bridge Filter, if it exists, and is displayed via the 'BRIDGE> list database static' command.

IGMPS>list bridge-filter

MULTICAST CONFIGURED ADDRESS FILTERS

Address	Allowed Port(s)
01-00-5e-07-07-07	6,10-12
01-00-5e-08-08-08	7-12

IGMPS>

Delete Command

Syntax: delete group [vsd | ip_addr | mac_addr]

This command removes multicast groups. All information about the group is removed, and IP multicast on those groups is broken until new information about the groups is learned (typically, when the next Query Interval begins and information is re-learned).

a. Delete Group Vsd

Syntax: delete group vsd <vsd_num>

This command deletes all IP multicast groups on vsd_num, for example:

IGMPS> delete group vsd 1

b. Delete Group Ip_addr

Syntax: delete group ip <ip_address> <vsd_num>

This command deletes multicast group ip_address from vsd_num. It deletes it only from vsd_num and only that IP address. The other 31 IP addresses that may or may not be involved are untouched, for example:

IGMPS> delete group ip 224.10.10.10 1

c. Delete Group Mac_addr

Syntax: delete group mac <mac_address> <vsd_num>

This command deletes all multicast groups for mac_address. All 32 IP addresses that map to this mac_address will have their information removed, for example:

IGMPS> delete group mac 01-00-5e-0f-22-04 2

HTTP Support

IGMP Snooping can be configured and monitored with a web browser. IGMP Snooping is listed under the Bridging section.

Accessing Online Information

Online Services

To locate product-specific information, refer to one of the following World Wide Web sites:

Americas:	http://www.networks.digital.com *or* http://www.cabletron.com
Europe	http://www.networks.europe.digital.com
Asia Pacific	http://www.networks.digital.com.au

Documentation Comments

If you have comments or suggestions about this document, email them to:

TechWriting@cabletron.com