digital™

The VNswitch 900 Family

Multitechnology, Multilayer, GIGAbit switching solution
for the DIGITAL MultiSwitch 900 System

enVISN
DIGITAL Networking Architecture

Digital Equipment Corporation
Network Product Business

August 1997

CONTENTS

1. Executive Summary

The VNswitch 900 family is a central product in Digital Networks' ***en**terprise **V**irtual **I**ntelligent **S**witched **N**etworks (***en*VISN**)* architecture.

The VNswitch 900 product family significantly expands the performance and functionality of the ***MultiSwitch 900***. The VNswitch family of products leverages the power, flexibility, and ease of use inherent in this enterprise switching hub architecture, first introduced in August 1993.

Performance, configuration flexibility, and ease of management are the essence of Digital's MultiSwitch 900, the only enterprise-class hub chassis offering a software-defined and enabled, technology-independent backplane design. "Technology-independent" means the chassis can support many technologies, frame or cell-based, with a simple no-cost firmware upgrade, not an expensive hardware replacement. This unique architecture guarantees customers a built-in, non-disruptive growth path, with investment protection unmatched in the industry.

The VNswitch 900 products deliver high-density, switched Ethernet connections to a choice of high-speed uplinks including:

- Fast Ethernet — 100BaseTX and FX
- ATM at 155 Mb/s and other speeds
- FDDI

Other family members provide switching among Fast Ethernet, ATM and FDDI.

In addition, an associated firmware upgrade to the MultiSwitch 900 will support implementation of a new, high-performance interconnect, known as the ***VNbus***. This high-speed bus architecture provides a multitechnology switching fabric which delivers an aggregate throughput of 1.2 Gb/s across three VNbuses. In addition to the VNbus, this simple firmware upgrade extends the current MultiSwitch 900 backplane support to include up to six new native ATM channels and four additional Ethernet segments for a total of twelve Ethernet segments. The MultiSwitch 900 already supports up to four independent FDDI rings which can also be used by the VNswitch 900 family. See Figure 1.

Combined with the wide variety of products already available for the MultiSwitch 900 (e.g., Ethernet and Token-Ring repeaters, FDDI concentrators, port switches, WAN routers, remote access servers, and wireless LAN base stations) a total solution to most networking requirements is now provided all within a single chassis and one common management environment.

With ***clear*VISN** the entire chassis with all its modules is managed through a common suite of integrated graphical applications.

Figure 1

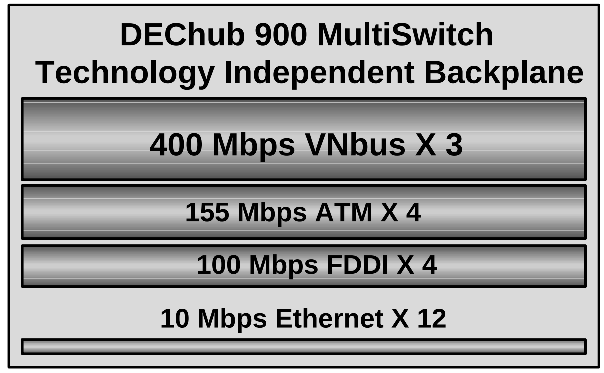

2. Functional Description

2.1 VNswitch 900 Product Overview

The VNswitch 900 is a powerful and flexible family of virtual network switches optimized for the MultiSwitch 900. The VNswitch 900 family supports switched Ethernet, Fast Ethernet, ATM, and FDDI connections. All members of the VNswitch 900 family are suitable for large site backbone operation as well as high-performance, "dedicated switched Ethernet" workgroups. Key features offered by all VNswitch 900 products include:

- Internal switching bandwidth of 1.2 Gb/s and up to 750,000 packets per second throughput
- Low-latency, line-rate switching on all ports between all technologies over a 400 Mb/s VNbus
- User-insertable I/O cards for a wide variety of Fast Ethernet, ATM, and FDDI media types
- Virtual LANs (VLANs) which interoperate with ATM Forum Emulated LANs
- Routing IP and IPX between VLANs providing up 100,000 pps per module
- 4 MB buffer per module that virtually eliminates packet loss and ensures optimum performance under heavy loads
- 8000 MAC addresses per module
- Full-featured bridge filtering and support for multiple IEEE 802.1d Spanning Trees
- SNMP and Telnet manageable, plus RMON and "Mirror Port" support

- ***clear*VISN** application support

2.2 VNswitch 900 Family Members

The initial wave of products in this family includes:

VNswitch 900EX	12 switched 10BaseT ports + two Fast Ethernet ports
VNswitch 900EA	12 switched 10BaseT ports + one ATM port
VNswitch 900EF	12 switched 10BaseT ports + one FDDI port pair
VNswitch 900EE	24 switched 10BaseT ports
VNswitch 900XX	4 switched Fast Ethernet ports
VNswitch 900LL	12 switched 10BaseFL ports + 12 switched Ethernet ports to the MultiSwitch 900 backplane
VNswitch 900XA	2 switched Fast Ethernet ports + 1 ATM port
VNswitch 900FA	1 FDDI port pair + 1 ATM port

All members of the VNswitch family also support the use of one of the 400 Mb/s VNbuses. See Figure 2.

Figure 2

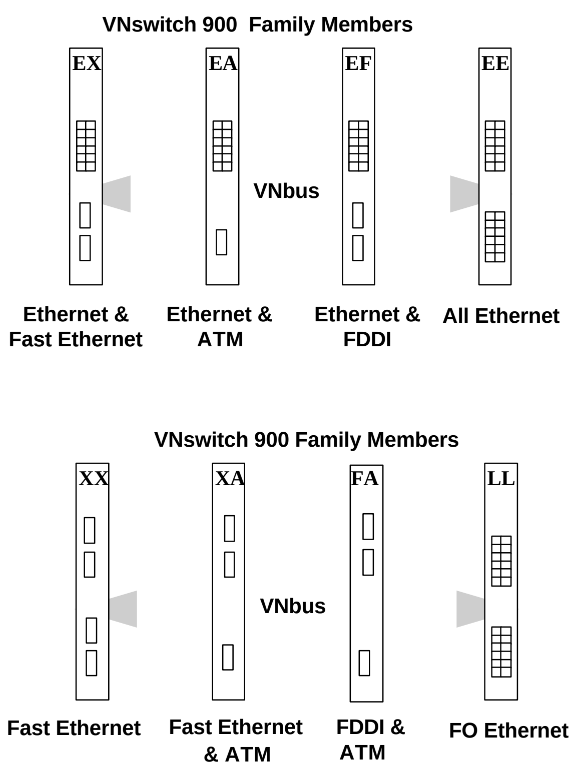

In addition, all VNswitch productsthat support Ethernet may connect to up to 12 Ethernet segments in the backplane of the MultiSwitch 900. FDDI and ATM backplane connections are also supported by this family.

The VNswitch 900LL, with 12 ports of 10BaseFL, (plus 12 backplane Ethernets) supporting a total of 24 switched Ethernets. Subsequent product releases will also include: "FF" (two FDDI) and "FX" (FDDI + Two Fast Ethernet). These future products are natural extensions given the modular internal architecture of the product. Refer to Product Architecture.

All LAN interfaces (Ethernet, Fast Ethernet, and FDDI) on the VNswitch family support **full-duplex** communication out the front ports. This can be used on point-to-point links for increased performance. ATM links are full-duplex at all times. The FDDI port pair supports both DAS and SAS operation, and dual homing.

Any mixture of VNswitch modules can be configured within a MultiSwitch 900 to support a wide variety of network requirements. Combined with the high-performance software routing option the result is an advanced *multilayer switching* system unprecedented in the market today.

2.3 Media Interface Flexibility

Media flexibility is provided on the VNswitch 900 family with front-inserted, modular I/O cards for high-speed links supporting a wide variety of Fast Ethernet, ATM, and FDDI cable media types. The I/O cards allow for hot-replacement with one of the same LAN type. The variety of media supported includes: Unshielded Twisted-Pair Category 5 (UTP-5), Multi-Mode Fiber (MMF) and Single-Mode Fiber (SMF)

Fast Ethernet Modular Media Interfaces (MMIs):

100BaseFX on MMF with SC connectors — 2 Km
100BaseTX (UTP-5) with RJ-45 connectors — 100 Meters

FDDI Modular Media Interfaces (MMIs):

FDDI on SMF, SC connectors — 60 Km
FDDI on MMF, SC connectors — 2 Km
CDDI (UTP-5), RJ-45 connectors — 100 Meters

ATM Modular Physical Interfaces (ModPHYs):

OC3c (155 Mb/s Sonet) on SMF, SC connectors — 25 Km
OC3c (155 Mb/s Sonet) on MMF, SC connectors — 2 Km
STS3c (155 Mb/s Sonet) on UTP-5, RJ-45 — 100 Meters
DS3 (45 Mb/s) on coax — requires CSU for WAN
E3 (34 Mb/s) on coax — requires CSU for WAN

In addition, a variant of the VNswitch 900EF with a fixed configuration of two ANSI MIC connectors for FDDI over MMF supporting distances of 2 Km is provided. This variant also has FDDI Optical Bypass Relay (OBR) support.

If FDDI or ATM connections are made over the backplane (e.g., to access an FDDI concentrator or ATM switch in the MultiSwitch 900) then *no* FDDI MMI or ATM ModPHY is required, lowering overall product costs.

The ATM ModPHYs used on the VNswitch 900EA are the same as those used on the GIGAswitch/ATM and other ATM products from DIGITAL. The ATM Forum standard UTOPIA interface is used internally to the unit to achieve this flexibility with a single physical "plug hole."

2.4 The VNbus for High-Speed Interconnect

Up to three VNbuses may exist in a one MultiSwitch 900. Each operates at 400 Mb/s for a total VNbus capacity of 1.2 Gb/s.

Each VNswitch can attach to only one VNbus at a time (or to none), and can be configured through software to attach to either of two: the upper or the lower. The lower is divided into two separate buses, one on each one-half of the hub covering 4 slots each. See Figure 3.

Figure 3

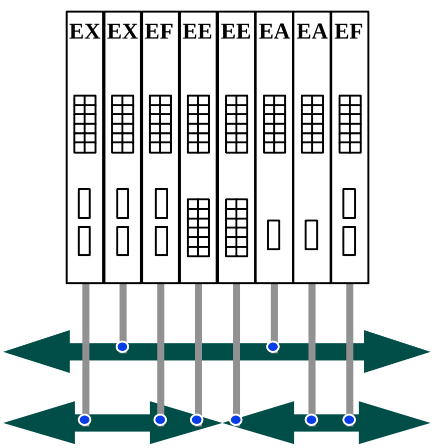

Three VNbuses

The use of the VNbus ties the members of the family together and allows for advanced capabilities:

- Automatic bridging or routing among Ethernet, Fast Ethernet, FDDI and ATM ports regardless of the modules on which the ports are located
- VLANs that span across multiple VNswitch modules, containing broadcast traffic within the VLAN (e.g., a designated port group)
- Inter-VLAN routing among multiple VNswitches from a single-switch module with routing code

The VNbus works with the *existing* MultiSwitch 900 chassis; hardware upgrade to the chassis is *not* required. The VNswitch modules themselves implement the VNbus on the MultiSwitch 900 backplane.

For simplicity, all VNswitch modules in a MultiSwitch 900 may reside on the *upper* VNbus with full connectivity among them. For performance or security sensitive environments, some modules can be directed to one of the lower VNbuses (as in Figure 3). All VNbuses can be operated concurrently. These form physically separate networks within a single MultiSwitch 900. These separate networks may be joined, if desired, with bridged or routed connections via FDDI, ATM, or Fast Ethernet.

The VNbus is a low-latency bus with fairness guarantees for all modules on the bus. Large LAN frames sent over the VNbus are broken into "fragments" so that small frames are not forced to wait for larger ones to complete transmission. The latency of using the VNbus is less than three microseconds. Under conditions of heavy load, there are flow control mechanisms implemented between modules to ensure that frames will not be lost if several VNswitch 900 modules are all transmitting to the same destination switch over the VNbus.

Frames are transmitted over the VNbus in either Ethernet or FDDI frame format. No translation of the frame type is performed unless the source and destination ports are of different types. (Bridged connections over the ATM port are also based upon Ethernet and FDDI frame formats, per industry standards such as LANE and RFC 1483).

2.5 Configuration Flexibility and Use with Other DEChub Modules

For high-density 10BaseT switching requirements, one MultiSwitch 900 may be configured entirely with VNswitch 900EEs for a total 192 10BaseT switched ports. Typically though, some number of 100 Mb/s or higher-speed links are desired for servers or backbone connections. Any mix of EX, EA or EF, EE, XX, LL, XA and FA modules can be located in the MultiSwitch 900.

Full backplane interoperability with the complete family of DEChub 90 and MultiSwitch 900 switches, repeaters, routers, etc., and support for industry standards, ensures smooth migration from existing configurations and protects investments.

Any mixture of dedicated or shared 10BaseT ports can be created conveniently in the MultiSwitch 900 by mapping some number of the 10BaseT switch ports (up to 12) to the Ethernet segments in the backplane of the MultiSwitch 900. Sharing of the FDDI switch port on a VNswitch 900EF is also conveniently done using FDDI concentrators in the DEChub supporting both FDDI dual rings and trees over the backplane.

Alternatively, a single VNswitch 900 module may be used standalone in a low-cost single-slot chassis (the DEChub ONE) to support a small high-performance workgroup or a dozen or two shared media stackable hubs distributed about a site. It can be used with or without connection to a higher speed backbone.

2.6 Building Gigabit Backbones

For very large or very high-performance networks the VNswitch 900 family provides an ideal edge-switch for multi-Gigabit backbones using DIGITAL award-wining ***GIGAswitch/ATM*** (with 10.4 Gb/s of full-duplex throughput and up to 52 ports at 155 Mb/s) or ***GIGAswitch/FDDI*** (with 3.4 Gb/s of full-duplex throughput and up to 34 ports of full-duplex FDDI). Alternatively, any standards-compliant multi-Gigabit switch can be used. Full-duplex Fast Ethernet may also be used for backbone switch connections. See Figure 4.

Figure 4

High-Performance Configuration

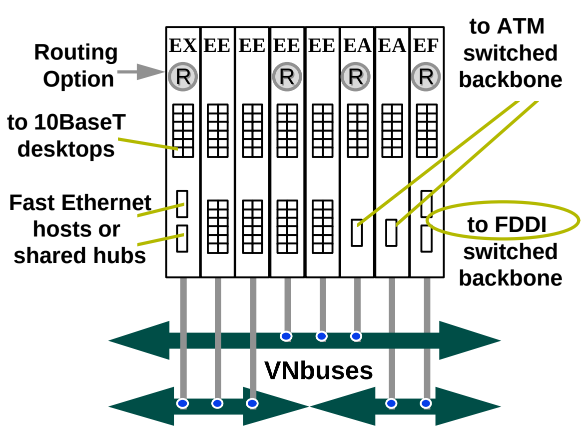

If only modest additional backbone capacity is needed the new ***ATMswitch 900*** family can be used. It provides 1.2 Gb/s of bandwidth and 8 ports of 155 Mb/s in a single module for the MultiSwitch 900. Up to four VNswitch 900EA modules may connect to it over the backplane, or all eight OC-3c ports can be accessed via the front bezel, which includes two ATM ModPHY ports. ATMswitch 900 supports multiple ATM Forum LAN Emulation Servers, flow control, dynamic routing and other advanced ATM features—same as supported on the GIGAswitch/ATM.

VNswitch 900s may be connected to each other, (and to the GIGAswitch and ATMswitch family) using multiple high-speed links (*with all links actively forwarding*) to support even the most demanding switched environments requiring many hundreds of Mb/s between chassis. This is accomplished using the multidomain VLAN and multiple Emulated LAN capabilities of these products.

Since VLANs and ATM Emulated LANs are under software control, they also facilitate intra-building network reconfigurations as well as the many small changes required as individual users move about.

The VNswitch 900 family will provide support for several classes of VLANs in accordance with the ***en*VISN** architecture and emerging industry standards. Support for multiple VLAN classes allows networks to be configured easily in a variety of ways to meet specific management objectives (e.g., security concerns, network stability, and the automatic isolation of “noisy” protocols).

3. Advanced Product Features

3.1 VLANs and VLAN Secure Domains (VSDs)

Virtual LANs (VLANs) are defined as separate *broadcast domains*. With the VNswitch it is possible for each VLAN to also have its own separate *spanning tree domain*. This is done by having just one VLAN per “VLAN Secure Domain” (VSD).

A VSD is a set of ports (on one or more switches) operating with its own separate IEEE 802.1d spanning tree process and logically isolated from other ports on the same switch by blocking all unicast and multicast traffic between VSDs. Spanning tree events (e.g., topology changes) in one VSD do not affect other VSDs, even if they are on the same switch or in the same hub. A VSD, as the name implies, allows for the highest degree of stability and security between VLANs. Only physically separate LANs unconnected in any way, would be more secure and isolated from each other.

Each VNbus can be divided into 63 VSDs. Each VNswitch on the VNbus can participate in a maximum of 32 VSDs on that VNbus; thus each port on a VNswitch could be placed in a different VSD, but joined in a VSD with ports on other switches. This allows multiple parallel active connections among VNswitches and various Gigabit backbone switches.

Figure 5

63 VSDs with one VLAN each

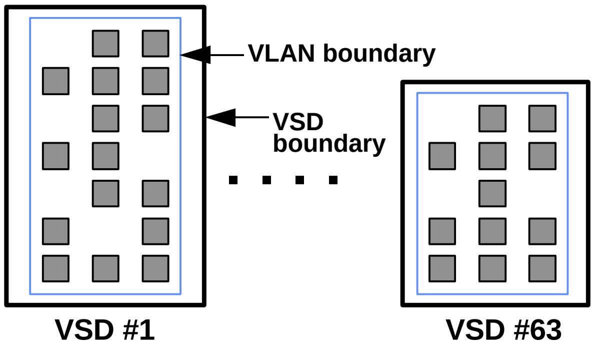

Non-overlapping, protocol independent, static with strong security. Based on port only.

Figure 6

3 special VSDs with up to 64 VLANs in each

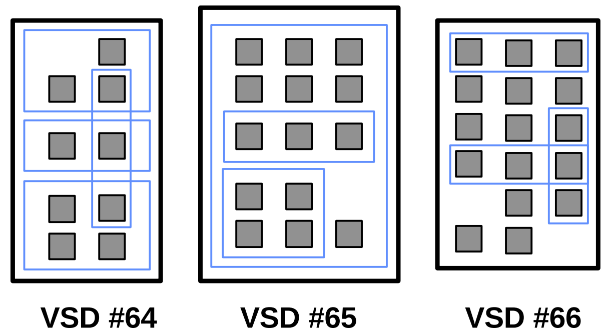

Overlapping, protocol sensitive, dynamic. Based on port, MAC address and protocol.

Within these VSDs the VNswitch 900 will implement just one VLAN per VSD for a total of 63 port-based (Class 1) VLANs, as shown in Figure 5. *The VSD and VLAN boundaries include exactly the same set of ports*. Hence, each VLAN will be secure from others. These can be mapped into ATM Emulated LANs and aligned easily to distinct user groups or IP or IPX networks, so that broadcasts do not impact stations in other networks. The use of DHCP (an IETF standard) can make this alignment automatic for IP. These VLANs can also be used to isolate stations that use unique or non-routable protocols or to isolate security-sensitive systems from the rest of the network.

While there is simplicity and clear advantages to having each VLAN be defined as a set of distinct (non-overlapping) ports in its own secure spanning tree domain, in other situations multiple VLANs per VSD are desirable. In the future three additional VSDs will be supported over the VNbus (for a total of 66 VSDs). Within each of these three VSDs it will be possible to create up to 64 ***advanced functionality VLANs***, providing a total of 255 VLANs over the VNbus.

Within each of these special VSDs the VLANs can be ***overlapping*** and ***protocol sensitive***. They can also support ***dynamic*** VLAN capabilities: that is, stations can be included in a VLAN based on their MAC address or protocol type. The port to which they connect need not specifically be assigned to any VLAN, as long as that port is in the VSD.
See Figure 6.

With overlapping port-based VLANs a server can be placed into multiple different VLANs, even though it has only one MAC address. In this way, groups of users can be isolated from each other, but all users can access that server without using routers. Alternatively, servers can be placed in distinct VLANs and users' workstations can be located in one or more VLANs — one for each server to which they are allowed access.

Overlapping VLANs also allow port-based VLAN boundaries to be "protocol sensitive" so that only traffic of a specified protocol type is blocked at the VLAN boundary. This is useful, for instance, to create a large VLAN for a non-routable protocol (e.g., LAT) allowing it to be bridged throughout most of a building or campus. Yet there may still be separate VLANs for each IP subnet within the campus since IP subnets communicate via routing. This provides the functions of a typical "brouter," but with much higher performance for bridged traffic.

The dynamic features of these advanced VLANs are useful, for instance, when a set of laptop computers (with known MAC addresses) need to be able to roam a campus and have access to the VLAN where their server(s) reside, without worrying about being on the "right side of the router."

Similarly, a VLAN for a particular protocol type (e.g., IPX) can be established dynamically *without* having to specify all the ports that may have IPX client stations on them. This effectively isolates the relatively "noisy" IPX stations (which use broadcasts heavily) from the rest of the network, even though they may move frequently.

Within the MultiSwitch 900, VSDs and VLANs are implemented over the VNbus. Between MultiSwitch 900s (or between standalone VNswitches) several options for carrying multiple VSDs and VLANs are available. One simple approach is to use a separate physical high-speed link for each VLAN. This approach, while attractive for its ability to bring added inter-chassis capacity, can become cumbersome as the number of VSDs and VLANs grows large. Therefore, an option to **trunk** multiple VSDs and VLANs per physical FDDI, Fast Ethernet, or ATM link is desirable.

“VLAN trunk” support will be provided in accordance with the emerging new IEEE 802.1 standards for VLAN trunks when that standard is complete. This standard defines a LAN frame format which includes the size and location of a VLAN “tag.” This will be implemented in future releases of the VNswitch 900 firmware for all high-speed links.

Initially, only ATM LAN Emulation (LANE) will be used for trunking VLANs between chassis. This provides the most straightforward and only standards-based “VLAN trunking” available today.

The VNswitch 900 family can, of course, be used without taking advantage of any of its VSD or VLAN capabilities. In this case, all the VNswitch ports in the hub are in one Spanning Tree domain and one broadcast domain — the “default” VSD and VLAN. In fact, this is how all VNswitches initially boot-up from the factory automatically.

3.2 ATM and Emulated LANs (ELANs)

The ATM Forum standard for LAN Emulation (LANE V1.0) provides a highly effective and easy-to-use backbone capability for transporting multiple logical LANs over a single link and set of switches, each with its own spanning tree domain or VSD. The VNswitch 900 family exploits this standard approach by allowing each VNswitch 900EA to be a member of up to *16* different Emulated LANs (ELANs). It does this by supporting up to 16 concurrent LAN Emulation Clients.

If there are more than 16 VLANs in a hub to be mapped into the ATM backbone (and joined with VLANs in other hubs or switches) multiple VNswitch 900EAs can be installed per MultiSwitch 900.

In an ATM backbone, of course, there may be many more than 16 Emulated LANs. Dozens of ELANs are possible using the LAN Emulation Server capabilities of the GIGAswitch/ATM.

Because ELANs are a standard way to implement VLANs, hosts attached directly to the ATM backbone can be assigned to participate in any one ELAN (or possibly more if they have multiple MAC addresses and necessary host O/S support). See Figure 7.

Figure 7

Standards-based VLANs via ATM Backbone

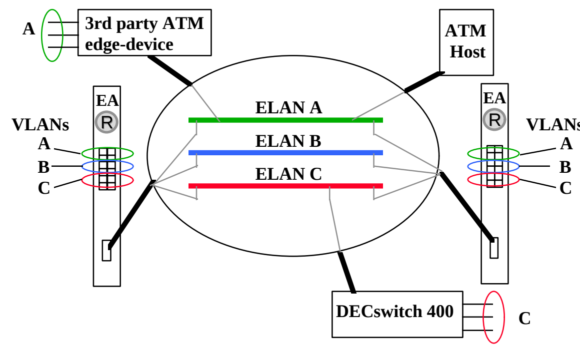

This contrasts with the proprietary VLAN tagging schemes used sometimes that make it difficult or impossible for hosts on the backbone to be selectively a member of the appropriate VLAN. The other major benefit is that any third-party ATM edge-switch supporting LANE can also participate in one or more ELANs (and hence the VLAN).

Using the VNswitch family in the MultiSwitch 900 all types of LANs (including FDDI and Fast Ethernet) can participate in an ATM ELAN, via the 400 Mb/s VNbus into a VNswitch 900EA module.

In addition to LAN Emulation, the VNswitch 900EA will support RFC 1483 “Bridge Tunnels” (Ethernet and FDDI format) and will support RFC 1577 “Classical IP.”

This is possible because the physical ATM port on the In fact, these can be supported concurrently with LANE over the same physical ATM port.

VNswitch 900EA actually provides **16 virtual ATM ports** that can be configured in any combination of LANE Client, Bridge Tunnel or Classical IP.

Shared across these 16 virtual ATM ports, there are up to 1024 Virtual Circuits (either SVCs or PVCs). This allows the creation of very large Emulated LANs because a “full-mesh” of SVCs is needed among all edge-devices and ATM attached hosts within one ELAN. For instance, an ATM ELAN with 60 edge-devices or ATM hosts requires 60 SVCs on each edge device and host, plus a few control VCs. Support for 1024 VCs allows 16 ELANs this size, with each VNswitch 900EA supporting hundreds or thousands of 10BaseT attached stations that can be in any one of these different ELANs. Fewer but larger ELANs can also be formed (e.g., one ELAN with 1000 ATM-attached participants).

SVCs are supported with both UNI 3.0 and UNI 3.1 options. In the future advanced flow control (*FLOWmaster*) can be used to prevent cell loss while providing instant access to the full bandwidth of the ATM link.

The VNswitch 900EA will also support four ATM Virtual Paths (VPs) for connection to public ATM networks: VPI=0, plus three non-zero VPs.

3.3 Multilayer Switching and the Routing Software Option

To communicate between VLANs the VNswitch 900 family puts routing where it's needed: *in the switch*. The result is called "*multilayer switching.*" It supports unicast LAN frame forwarding (i.e., Layer 2 bridging) at line rate within VLANs at up to 750,000 pps for all protocols, while only the inter-VLAN traffic that requires packet header handling (i.e., Layer 3 routing) for IP or IPX goes to the routing process which performs at up to 100,000 pps.

Each software routing option deployed supports up to 32 "virtual interfaces", and therefore may connect up to 32 different VLANs on the VNbus. *No* physical ports on the switch are consumed by the routing software.

The major benefits of multilayer switching are:

- Unicast bridging performance is not compromised by the use of Layer 3 routing between VLANs
- No external routers are required (which are themselves often quite expensive) and which consume physical ports on the switches, typically one per bridged domain or VLAN
- Higher overall performance at a much lower total network cost

This "designed-in" ability to run routing software adds minimally to the cost of the switch. "MIPS are cheap," and the large memory configurations often associated with high-performance routers are optional (required only when routing between many thousands of Layer 3 networks).

This unprecedented power and economy is achieved by employing an advanced internal architecture and several specialized integrated circuits as dedicated processing "engines" each carefully targeted to optimize the performance of selected networking layers. See heading 4.

A single hub of VNswitch 900s may have from zero to eight routing options deployed. Figure 8 shows two modules with routing software options configured.

Figure 8

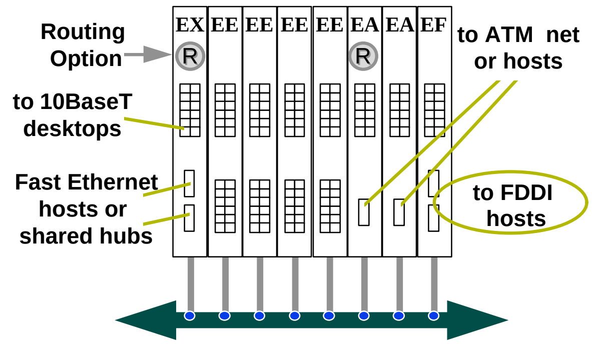

A routing option employed in just one or two VNswitch 900 modules will typically be sufficient for all LAN routing of IP or IPX required within a single chassis, or even across an entire large site supporting many thousands of users.

Routing for IP will be available in October 1997 with IPX available in the future.

The IP routing software option for the VNswitch 900 supports both RIP and OSPF routing protocols. Also supported are filtering options based on network, host address, or application (e.g., FTP); BOOTP Relay for DHCP support; automatic start-up; and quick configuration management tools. Base memory capacity of 8 MB of DRAM (16 MB in ATM versions) plus 4 MB of FLASH is adequate to support routing table entries for a couple thousand Layer 3 networks.

The routing function does not reside in any particular fixed relationship to physical ports on the modules. For instance, it is *not* necessary to route all the traffic going out over FDDI, Fast Ethernet, or ATM links (though that is certainly an option). It depends on how the VLANs are configured (i.e., which types of ports are included in the same VLAN). See Figure 9.

Figure 9

Routing with Virtual Interfaces to each VLAN

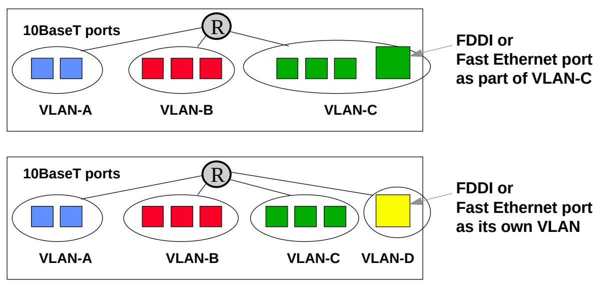

If a dedicated port into the router is desired, this is done simply by creating a single port VLAN and linking it to one of the router's virtual interfaces. If desired, any number of distinct physical ports for the router can be configured, up to 32 (spread across several VNswitch 900 modules in one MultiSwitch 900).

Any type of physical port (Ethernet, Fast Ethernet, or FDDI), can be configured in this way—as either a dedicated router port or as a member of a VLAN attached to one of the router's virtual interfaces. For the ATM port this logic extends to *each of the 16 virtual ATM ports* within one ATM physical port.

Therefore, an ATM virtual port, (e.g., acting as a LAN Emulation Client) may be joined into a VLAN with physical ports on the front of the modules (Ethernet, Fast Ethernet, FDDI, etc.). This is how ATM Emulated LANs are easily mapped to VLANs that include any number of different types of ports.

Alternatively, an ATM virtual port can be configured to have direct private access into a virtual interface on the router. This is how Classical IP is configured on the router. Likewise, a bridge tunnel between two routers may be constructed. Or, an Emulated LAN may be constructed whose only direct participants are virtual interfaces on routers in various VNswitch 900 modules. In any of these ways the VNswitch 900EA can provide both LAN and WAN routing capabilities over ATM.

Using the multiple virtual ATM ports some traffic may be bridged across the ATM cloud with very high performance (with tunnels or ELANs) while other traffic is being routed. This allows for complete flexibility in how to set up communications between workgroups in different locations over ATM.

If ATM links are not used for the WAN, connection to traditional WAN routers (e.g., to the RouteAbout family), can be made via backplane Ethernets or via front panel connections (e.g., using FDDI, Ethernet) to the DECNIS router family or any vendor's WAN router. Fast Ethernet connections to any vendor's LAN or WAN router are also supported.

Use of the routing capabilities of the VNswitch 900 family is purely optional. All VLAN and bridging features function as described *with or without* the routing option deployed.

3.4 Layer 2 Switching: Operation, Performance, and Features

The VNswitch 900 family operates as an IEEE 802.1d compliant store-and-forward bridge, filtering runts and frames with bad CRCs. With forwarding capacity up to 750,000 pps full line rate performance is obtained for unicast traffic across all ports for both transparent bridging (e.g., Ethernet-to-Ethernet, FDDI-to-FDDI). Line rate performance is also delivered for Ethernet-to-FDDI translation bridging. IP packet fragmentation for FDDI packets larger than 1500 bytes, and support for effective bridging AppleTalk and "raw IPX" between Ethernet and FDDI is also provided.

Multicast/broadcast forwarding can be accomplished at up to 100,000 pps, if desired (e.g., for multicast applications). However, in most networks limiting the rate of multicast forwarding is desirable. The VNswitch 900 provides a user setable **rate-limiting** feature for containing multicast traffic.

A large centralized dynamic buffer of 4 MB provides optimal congestion control for communication from high- to low-speed ports. This buffer is dynamically allocated as needed, virtually eliminates potential packet loss and retransmissions, thus increasing overall effective throughput. Switches with small dedicated per port buffers (e.g., 192 KB and less) may suffer from buffer overruns and lost packets when 100 Mb/s ports concentrate traffic into a 10 Mb/s port. This problem is solved in the VNswitch.

Filtering of LAN frames is also conducted at line rates and may take place in several ways. Specified protocol types or specified MAC addresses (source and destination) can be filtered at the input port. Protocol filtering can be inclusive or exclusive (i.e., filter just those listed or filter everything *except* those listed).

If the frame passes the input filters, the forwarding logic depends on whether the destination address is a "known" unicast address, an "unknown" unicast address or a multicast/broadcast address. It also depends on the presence of VLANs.

The basic filtering rule is as follows: a unicast with a known destination address (DA) will be forwarded to the port on which that DA was learned, unless that port is in a different VLAN. If it is an unknown DA or multicast/broadcast address the frame will be flooded to all ports within the VLAN.

Latency in the VNswitch family (measured as "last bit in, to first bit out") is under 16 microseconds for 64 bytes packets. Use of the VNbus between modules adds less than three microseconds. Therefore, the total latency is less than 19 microseconds from Ethernet port-to-Ethernet port between two modules. Ethernet-to-FDDI translation can be done with an additional latency as low as 10 microseconds.

Learning of a new MAC address by the bridging function is dynamic as in most modern switches. Static entries (manually entered) are also supported and dynamic learning can be disabled on a per-port basis.

However, unlike some low-end "Personal Ethernet" switches the number of MAC addresses that can be learned on any given port is *not* limited to an uncomfortably small number resulting in frequent flooding of "unknown" MAC addresses to all ports.

Each VNswitch 900 supports dynamic learning of over 8,000 MAC addresses; these may be learned on any port in any quantity. Inactive MAC addresses are aged-out according to timers adjustable by the network manager. If the limit of 8,000 is approached, then accelerated aging-out of older MAC addresses occurs automatically, rapidly freeing up the bridge learning capacity for new MAC addresses. This means up to 8,000 concurrently active users may reside within one VLAN (or the 8,000 may be spread across several VLANs on one or more switches). Total end-stations supported per VLAN can be much more because all will not be concurrently active.

Of course, the use of the router option between VLANs on different switches separates bridge "learning domains" and allows the construction of still much larger networks. If separated by routing, a hub with eight VNswitches can support 64,000 MAC addresses.

3.5 Management and *clear*VISN

Remote management for the VNswitch 900 is provided via both Telnet and SNMP. These are supported for *in-band* and *out-of-band* operation (using a dedicated management port) for high-security environments. Password and community string security is provided for Telnet and SNMP respectively, for in-band and out-of-band operation. A local console port is also provided.

Telnet allows for command-line control of all routing, bridging and VLAN functions, plus access to a wide range of performance statistics and status information. The Telnet access supports three different secure levels of management users.

SNMP provides access to all statistics of the VNswitch and control over almost all functions except advanced routing configuration options and user access lists. SNMP also provides control over reconfigurations of the backplane of the MultiSwitch 900. SNMP is the preferred protocol used within Digital's suite of ***clear*VISN** GUI-based management applications.

When two or more VNswitches are placed in a MultiSwitch 900, the "upper" VNbus (across the full width of the chassis) is automatically created. A module can be attached or detached from the VNbus via the Telnet command line or via SNMP with ***clear*VISN.** Use of the lower VNbuses requires ***clear*VISN** control.

The ***clear*VISN** Multi-Chassis Manager provides this control, plus the ability to easily establish ATM, FDDI, and Ethernet backplane connections to other switches, routers or repeaters. The Multi-Chassis Manager provides front-bezel views of the all modules in the MultiSwitch 900 (including LEDs) and allows "double-click" access to module specific configuration, status, and performance details.

With the ***clear*VISN** VLAN Manager advanced VLAN configurations (including those involving ATM Emulated LANs) are easily created, visualized, and reconfigured.

Both of these ***clear*VISN** applications are separately priced options available for Windows, Windows NT, and UNIX platforms. Included (free) with the routing software option is the Router Configurator, a GUI-based tool for conducting and maintaining simple routing configurations "off-line" and then downloading as needed.

Other ***clear*VISN** applications supporting the VNswitch 900 family include Recovery Manager, Alarms Manager, Flash Loader, Net Designer, and Router Manager. See ***clear*VISN** documentation for descriptions of all these and other applications. All ***clear*VISN** applications are integrated under a common database and can easily share information with each other.

Industry-standard MIBs supported in the VNswitch 900 family include MIB-II (RFC 1213), ifExtensions (RFC 1573), dot1d (RFC 1493), dot3 (RFC 1398), FDDI (RFC 1512), SONET (RFC 1595), AToM (RFC 1695), DS3/E3 (RFC 1407), and Etherlike (RFC 1643), 100BaseX, LAN Emulation Client, and ILMI**.**

Additional MIBs supporting a number of DIGITAL value-added features for Extended LANs, VLANs, LANE, ATM, and the MultiSwitch 900 are also provided.

Integral RMON support for RMON groups, Events and Alarms will be provided as a firmware upgrade to the product. Also, "mirror port" capability for non-disruptive monitoring of traffic between port pairs (by attaching an RMON Probe to the mirror port) will also be provided for support of all RMON groups. RMON History and Statistics will be provided in future hardware releases.

The switch normally boots up from its internal non-volatile FLASH memory, which stores VLAN, routing and other configuration data. The FLASH can also be upgraded remotely via TFTP. The switch can also boot from a BOOTP server, if necessary.

4. Product Architecture

4.1 Modular Internal Design

Each VNswitch 900 module for the MultiSwitch 900 is built with a modular internal design. There is a common "mother board" and a common set of "daughter boards" for all VNswitch 900 products. The mother board can support *two* daughter boards. The mother board contains the major frame and packet processing engines, the major DRAM, SRAM and FLASH memory subsystems, and the interface to the VNbus. Each daughter board supports a set of homogeneous data link types.

There are six different daughter boards, each listed below:

- Ethernet with 12 fixed 10BaseT ports
- Fast Ethernet with two flexible I/O ports
- FDDI with one pair of flexible I/O ports
- FDDI with one pair of fixed ANSI MIC Multi-Mode Fiber ports
- ATM with one flexible I/O port
- Ethernet with six fixed 10BaseFL ports

These can be combined with the mother board to produce a wide variety of self-contained switching solutions.

4.2 Three Processing Engines

The mother board contains three processing engines, each optimized to deliver high performance for a specific set of networking or management tasks. These are:

Application Processor (AP) — The AP is a high-performance 32-bit CISC processor. It is used for management interfaces (SNMP and Telnet), 802.1d Spanning Tree operation, and other high-level functions in the switch. When the routing software option is deployed it also performs RIP and OSPF route calculations.

Fastpath Processor (FP) — The FP is a high-performance 32-bit RISC processor. It is used for multicast bridging, FDDI-to-Ethernet translation bridging, assisting with advanced VLAN features, and IP fragmentation for FDDI. When the routing software option is deployed it provides the high-performance packet forwarding engine for Layer 3 routing, by using a high-speed cache of active routes.

Data Moving Engine (DME) — The DME is a sophisticated custom ASIC designed by DIGITAL to perform unicast transparent bridge forwarding (i.e., switching) at up to 750,000 pps, advanced bridge filtering, and core VLAN functions at very high speeds. The DME can support up to 48 bridged ports.

The system architecture is shown in Figure 10.

Figure 10

VNswitch 900 Internal System Architecture

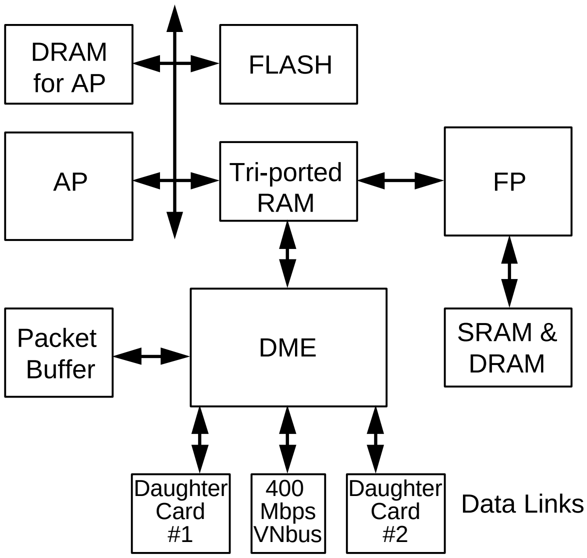

All three processing engines work together in a highly complementary fashion. Because advanced Layer 2 functions use the same processor as Layer 3 packet forwarding, there is no major component of the switch that is unnecessary for "bridging-only" environments. If used heavily for Layer 3 routing, the load from advanced Layer 2 functions is typically much less. This frees the power of the FP to be used mainly for Layer 3 routing. Therefore, regardless of the mix of loads, the system is balanced, optimized, and very cost-effective.

In addition, there are a number of highly sophisticated custom ASICs designed by DIGITAL to multiplex each different data link type efficiently into the DME.

4.3 Memory Subsystems

Each major processing engine has its own dedicated memory subsystem designed for the speed of that processor and the unique demands upon it.

The AP comes with 8 MB of DRAM (16 MB for ATM version) plus 4 MB of non-volatile FLASH. The DRAM is used for data structures required by the bridging and routing application code. Also, it supports the collection of alarms and events data and performance statistics. The FLASH memory is used to store the configuration and operational code, and to run the code.

The FP comes with 1 MB of DRAM plus 0.5 MB of very high-speed SRAM. These are used for running the small kernel of code needed to conduct Layer 3 packet forwarding and filtering operations, and operations related to advanced Layer 2 functions. Over 16,000 Layer 3 destination addresses can be held in the cache for the FP, ensuring a high-probability of a cache hit for LAN backbone routing environments.

The DME has a dedicated 4 MB packet buffer shared dynamically across all ports on the switch. This is used for speed matching (between different links) and congestion control. This central buffering strategy provides much better performance under load than dedicated per-port buffers used on some switch designs.

Finally, there is a section of very high-speed tri-port RAM used for communications between all processor/memory subsystems.

4.4 Internal Addresses

A total of 256 unique internal MAC addresses are assigned and shipped with each VNswitch available for management and other special purposes.

(**Note:** These should not to be confused with the 8,000 MAC addresses the bridging function learns dynamically.)

To support Telnet and SNMP management, the VNswitch allows an IP address to be assigned to it with a unique MAC address for this purpose. One MAC address is also used for each physical port and each ATM virtual port.

A unique MAC address for each virtual router interface is also assigned. This fully distinguishes the router's virtual interfaces from the physical ports of the switch (and the ATM virtual ports), and allows separate tracing of traffic into and out of the router by MAC address.

4.5 Priority Queues

Management traffic and Spanning Tree packets into and out of the switch are given a higher priority so that manageability and Spanning Tree performance are not compromised under heavy load conditions.

In addition, the DME chip supports several classes of priority queues for user traffic, including one that allows regular servicing of high-bandwidth "constant bit rate" data flows to support realtime multimedia networking. Also provided is a priority class for urgent user data. However, these features require additional firmware and management application support before they can be used. These will be implemented consistent with the emerging IEEE 802.1 standards for Layer 2 priority queuing and various methods of distinguishing the traffic which is to be given priority.