DIGITAL MultiSwitch 900

Multi-Technology, Gigabit Switching System

Digital Equipment Corporation
Network Product Business
August 1997

CONTENTS

1. THE VISION 3

2. PLATFORM OVERVIEW 3

2.1 Physical Design 3

2.2 Technology-independent, bandwidth-unlimited backplane 4
2.2.1 Backplane Layout 4
2.2.2 Backplane Bandwidth and Performance 6
2.2.3 Intra-Chassis Communications 7
2.2.4 DIGITAL MultiSwitch 900 'Citizenship' 7

2.3 Platform Manager — Intelligent management & control 8
2.3.1 Technological Neutrality 9
2.3.2 IP Access Services 9

2.4 LAN configuration strategy, flexibility, ease-of-use 10
2.4.1 Dynamic LAN Assignment 11

2.5 Availability/reliability/predictability 12
2.5.1 Fault Management 12
2.5.2 Power System 13

3. *clear*VISN NETWORK MANAGEMENT 14

4. FUTURE DIRECTIONS 15

1. The Vision

In the world of networks, many people would view the ideal network as one where the people responsible for keeping the business's lifeline up and running would *not* have to worry about the details, facts, figures, bandwidths, and performance numbers that they spend the vast majority of their time on today.

Certainly there is a level of reality that comes with managing a *real* network and all of its resources, including the necessary wires, fibers, active network electronics, PCs, servers, and users. However, should the people in charge really be burdened with LAN access methods, packet sizes, data link translations, bottlenecks, and catastrophic hardware faults?

People who are in the business of running manufacturing, software, education, government, service, or other environments should not need to worry about power supply failures and whether today's network platform can support tomorrow's bandwidth. Can't an approach be found to networking that provides the tools and capabilities to let managers become planners for future growth and change rather than continually responding to unforeseen and unpredictable events? It should be a goal to enhance their effectiveness and productivity by minimizing the reactive tasks that are required today, and enabling higher level, policy-based structures that help them approach a totally automated or "Lights Out" networking environment.

Digital Equipment Corporation has been focused on *exactly* that "Lights Out" approach to networking with its *en*terprise *V*irtual *I*ntelligent *S*witched *N*etworks (*en*VISN) strategy. *en*VISN is Digital's networking architecture and product strategy for the future. Software applications such as the *clear*VISN Management application suite provide the policy-based administration and aid in designing, configuring, and monitoring our customers' networks. Coupled with a flexible hardware platform such as the DIGITAL MultiSwitch 900, they provide the infrastructure needed to meet today's performance and technology needs as well as tomorrow's growth requirements including higher density, granularity, and segmentation capabilities; higher aggregate speed via wire-speed switching; increased efficiency and control through distributed routing; new technologies such as VLANs and ATM, etc.

This paper describes an Enterprise networking platform which protects our customers' investments *regardless of the direction that their business requirements take them!*

Let's begin by discussing the necessary prerequisites that such a *Virtual Network Platform* must have to provide a clear path towards this vision, and simultaneously describe in detail what makes the DIGITAL MultiSwitch 900 so unique.

2. Platform Overview

By necessity, any network platform structure consists of several basic functions, such as a backplane design, functional modules, some level of management and a power system.

Take these fundamentals of a chassis, add in a passive, bandwidth-unlimited, technology-independent backplane, allow software control over virtually every functional element, tie them together with unique, friendly packaging — and the DIGITAL MultiSwitch 900 is born.

2.1 Physical Design

The DIGITAL MultiSwitch 900, as shown in Figure 1, is a platform with 8 module slots, 4 power supply slots, a built-in platform manager and a liquid crystal display (LCD) for local message display. It may be mounted directly onto a wall to save space (since no cabling or slots are accessed from the rear) or in a standard 19-inch rack. Wall mounting is often a less expensive, easier means of physically installing the MultiSwitch 900 chassis since it can be accommodated in smaller data communications closets, and there is no need to invest in or use rack space.

Network modules are compact and are designed for ease-of-use. They are completely self-contained, have no exposed electronics other than the connectors for connecting to the backplane (described in section 2.2.1), and are packaged to be able to run either in the MultiSwitch 900 chassis or in a standalone 'single slot' power supply (DEChub ONE or DEChub ONE-MX)[1] with *no modifications*. This enables users to "start small and grow large," thus retaining their initial investment. The same simple module latching mechanism that is used to secure a module in the MultiSwitch 900 chassis is also used when installing it into one of these 'single slot' power supplies.

[1] DEChub ONE provides 90W of power, the DEChub ONE-MX provides the same 90W plus two open bays for MODular Physical layer Media Dependent (MOD-PMDs) in support of modules with FDDI uplinks.

Figure 1. DIGITAL MultiSwitch 900

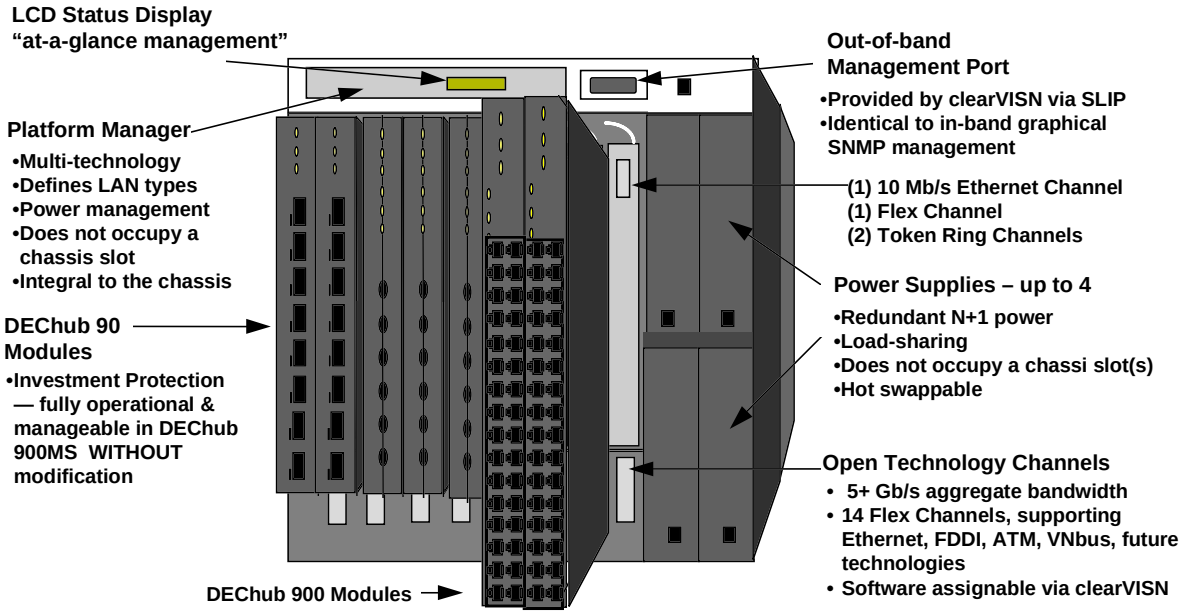

Digital believes that the design of the chassis should be as simple as possible and the network modules should be doing the work. Therefore, the financial investments for this system are mainly in the network modules. This should not be interpreted, however, to mean that the chassis is simple in the sense of lacking functionality. As will be demonstrated throughout this paper, the MultiSwitch 900 chassis is *elegantly simple* while providing the user with a *dynamic* and *powerful* foundation that forms the basis — without any hardware upgrades — for *increased performance* and *functionality* in the future.

2.2 Technology-independent, bandwidth-unlimited backplane

2.2.1 Backplane Layout

The MultiSwitch 900 backplane is completely passive and consists of over 200 lines of etch. The male conductors that drive the backplane all reside on the individual modules, thus ensuring that a bent pin condition affects only that module and not the backplane. Each module consists of two connectors, one of 48 pins and the other of 160 pins.

The services provided by each of the connectors are shown pictorially in Figures 2 and 3.

Figure 2. 48-pin Connector Backplane Service

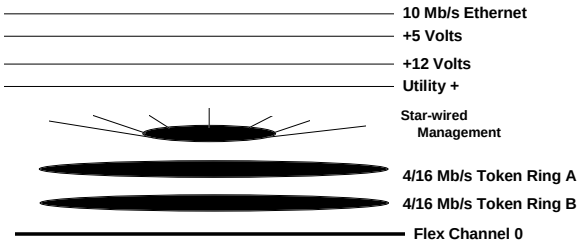

This 48-pin connector supports all of the functions of the DEChub 90[2] plus additional capabilities. The "10 Mb/s Ethernet" channel maintains backward compatibility with all existing DEChub 90 modules[3] and enables our customers to retain their investment in these modules by migrating them forward into the MultiSwitch 900 chassis. A key added feature of this portion of the backplane is the star-wiring of the management bus, which supports variable speed, dedicated management connections from the Platform Manager (to be described in section 2.3) to the up to 8 modules installed in the chassis.

[2] DEChub 90 is a single-technology (Ethernet), 8-slot workgroup hub that supports "half height" modules with a single 48-pin connection on the backplane.

[3] DEChub 90 modules range in function from simple repeaters and bridges to WAN routers, dialup access servers, wireless access points, and RMON probes.

Notice also the addition of the two Token Ring channels. These channels are dedicated for Token Ring LANs and support either 4 Mb/s or 16 Mb/s implementations. The family of Token Ring DECmau 900 modules all have the capability of autosensing the speed at which the ring is set.

A truly innovative addition to the 48-pin connector is the Flex Channel, denoted by "Flex Channel 0" in Figure 2. The Flex Channel implementation is a technology-independent, high-bandwidth channel that can support FDDI, Ethernet, ATM, Fast Ethernet, VNbus[4], or any other technology that becomes implemented across the backplane. The technologies that are driven on the Flex Channels[5] are determined by the modules that the user inserts into the MultiSwitch 900 chassis and configures via the clearVISN Network Management application suite[6].

When modules are installed in a MultiSwitch 900 chassis, they inform the Platform Manager of the type of connectivity they support. Through the *clear*VISN MultiChassis Manager application (point-and-click), the user instructs the Platform Manager which modules should be placed on which LAN. This implementation allows a flexible, software-controlled connection and interconnection of modules within the chassis, regardless of the technology. The Platform Manager keeps track of all of this information. The user simply decides how to solve his/her interconnect needs without having to be concerned with the technology that enables the Flex Channel.

[4] VNbus is a low latency, high performance 400 Mb/s channel. Up to three VNbuses are currently supported in every MultiSwitch 900chassis (with no hardware backplane upgrade). VNbus technology is a major enabling feature which effectively uses the power of the MultiSwitch 900 chassis; it allows high-density, low-latency "Any-technology to Any-technology" switching with extensive VLAN support and is discussed in detail in the "The VNswitch 900 Family — High-performance switching solutions for the DIGITAL MultiSwitch 900 Enterprise hub" white paper (EC-N6761-42).

[5] There are an additional 14 Flex Channels on the 160-pin connector, to be discussed later.

[6] clearVISN is Digital's Policy-based network management application suite which enables the Network Manager to become more effective by managing "the network" rather than individual devices. Currently the applications consist of Stack Manager, MultiChassis Manager, VLAN Manager, RMON Manager, Recovery Manager, Router Manager, Flash Loader and Router Configurator.

On the face of it, the previous description of technology-independent Flex Channel(s) may seem to be somewhat mundane since the concept is so simple; however, this technology produces some of the most important benefits for Digital's customers today.

These benefits include:

- *Investment protection* since no hardware backplane upgrades are needed to support new technologies in the future
- *Increased performance* since, as new modules are developed, they will be able to make more efficient use of these Flex Channels, thus gaining bandwidth by simply using more sophisticated signaling techniques
- *Ease of migration* as new technologies may be run on the same Flex Channels that are shipping today
- *Simplicity and flexibility* since it provides the foundation to reconfigure a MultiSwitch 900 chassis through simple point-and-click operations without having to worry about *which* Flex Channel or *how many* Flex Channels are being used. This significantly lowers operational and training costs

This Flex Channel technology is also used extensively on the 160-pin backplane connections. This connector provides for the placement of 12 additional Flex Channels and power management associated with each channel. While the placement allows for 12 positions for the Flex Channels, Channels 11 and 12 have been segmented between slots 4 and 5 to allow 14 independent connections on the matrix backplane, shown below.

Figure 3. 160-Pin Connector Backplane Services

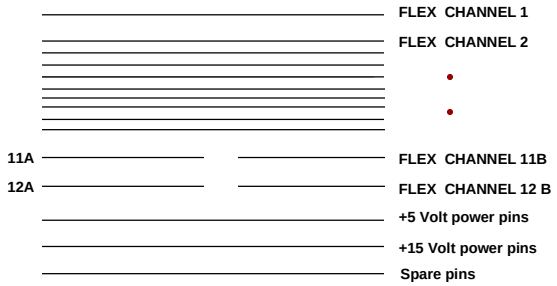

These Flex Channels are identical in function to the single Flex Channel on the 48-pin connector.

They are allocated by the Platform Manager to meet the needs of the modules installed in the chassis — essentially they can be thought of as a general "bandwidth pool" to draw from. When allocated to

specific network modules, the channels may take on whatever characteristics the module requires and are controlled using *clear*VISN MultiChassis Manager. As mentioned earlier, Flex Channels allow communication between modules for standard LAN connections such as Ethernet, FDDI, ATM, as well as the VNbus, etc.; however, other implementations of buses, LANs, point-to-point links and bandwidth pools are possible, i.e. it is *truly* technology-independent.

This type of flexible implementation allows the modules to use as much or as little of the chassis channels as needed. It thereby allows much flexibility for a user in his/her tradeoffs of functionality and performance. This freedom gives a very powerful level of control and flexibility to our customers who may require single collision domain repeaters today, personal switched Ethernet connected to an FDDI or Fast Ethernet server tomorrow, and ATM to the desk in two years (or next week). All environments are supported by the same original MultiSwitch 900 chassis that first shipped in 1993 — no backplane upgrade required, only firmware!

Figure 4. A Sample Configuration Using Technology-Independent Flex Channels

While the above Flex Channel technology actually allows the Platform Manager to assign the backplane connections on a pin-by-pin basis, this detailed level of pin assignment is of course transparent to the network manager and simply enables him/her to more

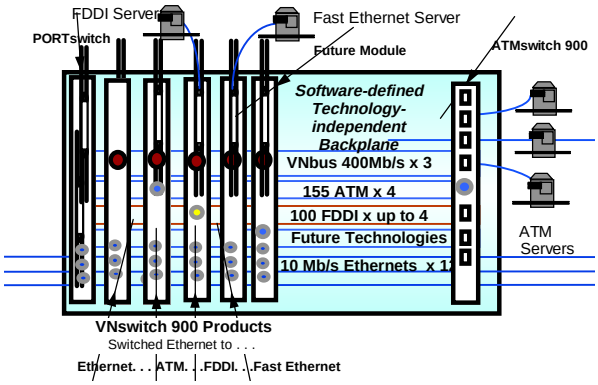

effectively use the dynamic bandwidth of the backplane.

This design approach allows the basic chassis structure to have a considerably longer life (several generations of network modules) than fixed configuration or fixed bus backplanes which may reach their maximum capacity and require hardware upgrades or possibly a whole new platform to support new technologies, faster speeds or both. Thus, our customers' investment is protected indefinitely. The next section delves more deeply into this concept.

2.2.2 Backplane Bandwidth and Performance

The backplane bandwidth of the MultiSwitch 900 chassis has no set limit. This is a design feature of the chassis and key to its forward migration to *any* network technology. Essentially it has a bandwidth-unlimited capacity that is a function of the number of pins used by the modules and the signaling driven over those pins.

Consider the backplane as a set of wires or etch. The speed at which they signal is a function of how the installed *modules* use these wires. Current signaling rates include 10 Mb/s for Ethernet[7], 25 Mb/s for FDDI[8] and early signal integrity work has shown that with simple signaling, it is easy to push this interconnect style to 50 Mb/s for each signal interconnect on the backplane. At a data rate of 50 Mb/s, the 100 backplane interconnect lines[9] provide approximately 5 Gigabits of bandwidth. While this number is often used as the *current* bandwidth capacity of the chassis, it is clear that the maximum capacity of the MultiSwitch 900 chassis has yet to be reached.

As new modules are developed, new technology will drive the *same* backplane harder and more intelligently and thus obtain more interconnect bandwidth. Personal Ethernet applications[10] certainly increase the interconnect bandwidth used. ATM and Fast Ethernet solutions increase it again. It has been proven that ATM can run on the existing backplane using signaling at over 75 Mb/s per line. This pushes the effective chassis data rate to about 7.5 Gigabits. In the future, it is expected that modules will signal at the OC-3 155 Mb/s ATM rate on each pin of the interconnect. If so, this would push the backplane capacity even higher to over 10 Gigabits[11]. Future modules will continue to prove for some time to come that the backplane capacity of the MultiSwitch 900 chassis is only limited by the signaling limit of each line of etch on its technology-independent chassis.

---

[7] First generation Ethernets used six pins per Ethernet LAN. Second generation Ethernet LANs are supported on only two pins!

[8] Five pins are used for each FDDI channel. 4B/5B encoding yields 100 Mb/s.

[9] 100 of the 160 available pins are used for actual data traffic today. As shown on Figure 3, others are used for power, clocking and spares (or future use).

[10] Dedicated 10 Mb/s channel per user.

[11] There is a non-linear relationship between speed per pin and overall capacity because not all pins are used in the same way, e.g. some are dedicated, others are not.

The most important consequence of the architecture is *not* any specific bandwidth number however. The major advantage of this architecture, and thus the major benefit to our customers, is that it allows the provision of inexpensive interconnects for some technologies and yet still provides methods for extracting more bandwidth as needed for higher performance interconnects. *No other chassis can provide the price/performance scalability from a simple Ethernet repeater interconnect through multiple FDDI, ATM, or VNbus interconnects on the same backplane.*

The user invests in the MultiSwitch 900 chassis and pays only for the reliable, inexpensive, passive etch interconnect with power and full manageability. The modules can then be chosen for the bandwidth/performance required for the network at hand. Our customers know that as the network grows, the chassis can accommodate that growth by simply choosing and configuring modules to migrate forward to meet the new requirements.

2.2.3 Intra-Chassis Communications

Technology independence also guided the hardware design and the firmware architecture of the MultiSwitch 900 chassis so that *intrachassis* management communication was implemented in the most flexible and efficient manner.

The Platform Manager communicates with each module in the chassis via an efficient, direct mechanism, that is bi-directional serial lines. This mechanism was used since the speed of the direct connections could be negotiated between the Platform Manager and each module, and the address overhead could be optimized by the point-to-point nature of the star-topology wiring. Subsequently, the Platform Manager *always* knows which module is at the end of each communications line.

Star wiring also allows for simple fault tolerance in that no module could prevent communications with another, either by a hardware fault or communication error. Each channel is very reliable, yet very simple, and independent of the module's technology. This star-wired management structure is scalable to larger or smaller applications both in terms of line speeds as well as platform applications with fewer or more slots.

2.2.4 DIGITAL MultiSwitch 900 'Citizenship'

All MultiSwitch 900modules have many common features and characteristics ranging from: having the same physical look and feel; using the same setup ports/cables whether in chassis or standalone; using the same power mechanisms whether in the chassis or standalone; detecting environmental changes in the same manner; having similar setup menus; and sharing *common code* wherever possible, thus allowing for greater standardization and compatibility across the family.

In particular, each DIGITAL MultiSwitch 900 "citizen" must provide a series of services, including:

- IP Access Services (see section 2.3.2)
- Dynamic LAN Assignment services (see section 2.4.1) whereby it shares each module's backplane interface requirements and capabilities with the Platform Manager and executes the Platform Manager's commands
- Serial driver for communication with the Platform Manager
- Timer services for implementing time-out functions
- Buffer services to receive bursts of packets from the Platform Manager
- Non-volatile storage to store the management status and error logs
- Device services to allow a user to disable/reset/enable each line card
- Status services to record overall device status
- Controlled shutdown services to deactivate a device when required
- Identification services to indicate the identity of the module at initialization
- Authentication services to provide the SNMP information needed to access the device
- Power services to provide information on power usage to the Platform Manager
- Out-of-band management (OBM) services whereby the Platform Manager acts as an IP Access Server for the module
- Hub Console Carrier Protocol (HCCP) services to enable the redirection of the chassis setup port to the setup port of an installed network module

These common services are independent of the module's network function, whether it is a repeater, concentrator, mau, bridge, switch, router, wireless access point or remote access server. They are also independent of the LAN/WAN technology that the module is providing to the users whether it is Ethernet, Fast Ethernet, Token Ring, FDDI, ATM, ISDN, Frame Relay, SLIP, PPP, dialup or T1 leased line services.

The Layer 2 LAN switches that need to support "Distributed Routing" or "Layer 3" switching can do so with a simple firmware upgrade. In particular, the

DECswitch 900 family is a set of multilayer LAN Switches which are used to interconnect Ethernet, FDDI, and Token Ring segments. The VNswitch 900 family of multilayer switches provides the platform for "Any-to-Any" (Ethernet, Fast Ethernet, FDDI, ATM) switching with common high-speed VNbus backplane support. The RouteAbout family of products provide routed WAN connectivity in the form of RouteAbout Access (Remote Ethernet to WAN; Remote Token Ring to WAN) and RouteAbout Central (High-Density Ethernet to WAN). In fact, the same core common routing code is used on all of the above devices. All are interoperable and compatible with one another, thus reinforcing another of the design goals of the MultiSwitch 900 chassis and its modules, that is *investment protection*.

Another key benefit that users gain from the architected DIGITAL MultiSwitch 900 solution is the knowledge that once they have installed one module, they can setup the next module in essentially the same manner. Each DIGITAL MultiSwitch 900 "citizen" implements a setup menu that minimally has the following six options:

- Reset to factory defaults
- Reset with current settings
- Show current settings
- Set in-band IP address
- Set SNMP R/W Community String
- Set out-of-band IP address

Once each module is set up to run within the chassis or standalone environment, there are typically additional module-specific or technology-specific configuration options that the network manager may fine tune for that environment using *clear*VISN's GUI interface. These include setting up ports for HDX/FDX, adjusting flow control parameters, selectively filtering by protocol/address, setting the broadcast limiting threshold, creating VLAN broadcast domains for security reasons, etc.

The point is that the user goes through the same initial setup procedure, using the same user interface, the same setup cable, connecting to the same setup port, and can expect the same common characteristics from each of the DIGITAL MultiSwitch 900 "citizens," regardless of module. The *consistency* and *predictability of operation* across the DIGITAL MultiSwitch 900 product family allow users to be more effective in day-to-day operations.

2.3 Platform Manager — Intelligent management & control

The Platform Manager is a technology-independent SNMP management agent that is built into every MultiSwitch 900 chassis. This independence allows the management of Ethernet, Fast Ethernet, Token Ring, FDDI, and ATM modules today as well as technologies in the future. It is essentially a communications processor which monitors, configures, and controls the chassis. It also provides a central point of control for allocation of chassis resources to those network modules installed in the MultiSwitch 900 chassis.

Think of the Platform Manager as a behind-the-scenes, minimum requirement for any MultiSwitch 900 chassis tasked with managing both itself and its network modules. That is why all MultiSwitch 900 chassis contain a Platform Manager — to provide intelligent networking! When viewed in this light, the Platform Manager should be seen as an enabler, or agent, in assisting the optimal connection of users to resources and the subsequent monitoring and control of that connection, regardless of technology.

The users benefit since the Platform Manager controls all of the modules in the chassis, and there are *no* required Network Management modules in the chassis stealing valuable slots. A sampling of typical wiring closets often reveals at least three technology types and multiple generations of interconnect technology. Do users really want a management architecture which is dependent or biased towards a single LAN technology? Digital believes not. The Platform Manager is a service provider to the installed modules and *not* a limitation to the future growth of the chassis.

The Platform Manager has the following basic functions but does much more:

- Provides a Setup Port (RS232/RJ45) for local configuration to set IP addresses, request downline upgrades, configure individual modules, and reset its nonvolatile storage
- Responds to SNMP queries for platform statistics
- Identifies (via IP/MAC address) and provides power management for all installed modules
- Performs network interconnections between modules
- Supports local status information via LCD
- Supports SLIP-based OBM via RS232-DB9 port for the Platform Manager and all installed modules

The Platform Manager's capabilities allow for full management of the chassis and all modules in the

chassis via a *single* IP address, a very critical benefit for users who need to conserve address space. The network manager can access it either via its in-band or out-of-band IP address to manage the chassis and also provide access to all of the SNMP agents in the chassis.

As an alternative, management applications can communicate directly to network modules via their own IP address as long as the network manager has assigned that address and there is an in-band communication path to that module. [12]

This gives the network manager the *flexibility* to configure as much redundancy into the management infrastructure as needed. While able to act as the single management interface for the chassis and all of its resident modules, the DIGITAL MultiSwitch 900 Platform Manager does not, however, connect directly to any given LAN technology which will shown in the following sections as a distinct and unique advantage.

2.3.1 Technological Neutrality

A key feature of the MultiSwitch 900 chassis is that the backplane is "technology-neutral"; that is, it is not biased toward or away from any particular LAN type, topology, or physical network media. It is *not* simply an Ethernet hub, a router platform, or a UTP wiring center, but rather a platform that can provide a home for *any* network technology module, *without* compromising one module for another's required functions.

If the design goal had been "be all things to all technologies," the resultant platform could have (some would say "would have") actually been mediocre at *all* of these tasks. The technology independence concept guided the design of the fundamentals of the platform. The DIGITAL MultiSwitch 900 backplane, built from an architecture which is totally technology independent, inherently has the capability of supporting any of today's technologies as well as tomorrow's with *no backplane hardware changes.*

2.3.2 IP Access Services

One aspect of the technological neutrality of the Platform Manager is that it does not have a specific Media Access Control (MAC) for any LAN type. To have included one would have created both a bias toward that LAN type and a burden for applications where that LAN type was not needed for the application or configuration.

IP Access Services is a communications mechanism that allows in-band access to *any* technology within the platform. Only one network module, configured as the IP Access Services module, is needed to provide this service for the DIGITAL MultiSwitch 900 and all of the other modules in that chassis.

The IP Access Services module forwards IP packets (on behalf of the Platform Manager) between the management channel and the local network interface(s) of that module. This module binds its local MAC address to the Platform Manager's IP address and responds to ARP requests on behalf of the Platform Manager.

The IP Access Services module, of course, receives all packets destined for its own MAC address. It then forwards those IP packets which are destined for the Platform Manager's IP address to the Platform Manager via the backplane management channel. Packets which are destined for the IP Access Services module's own IP address are kept and processed locally. Packets which are destined for other modules within the same chassis are sent to the Platform Manager for forwarding to the proper module.

For example, if your IP Access Services module were an Ethernet switch in slot 8, and you were managing an FDDI concentrator in slot 4, the IP Access Services module would receive the SNMP message from the *clear*VISN MultiChassis Manager station in-band via the Ethernet switch's 10 Mb/s Ethernet connection and immediately forward it to the Platform Manager. The Platform Manager would forward that message to the FDDI concentrator in slot 4. The IP Access Services function is similar to a network access server acting as a SLIP router for a remote PC. Another way of thinking of it is that the module is acting as a network interface adapter card for the Platform Manager — in this case the adapter just has other jobs to do as well!

Figure 5. IP Access Services Information Flow

[12] It also allows these modules to be fully manageable even when configured as a standalone device (with DEChub ONE or DEChub ONE-MX).

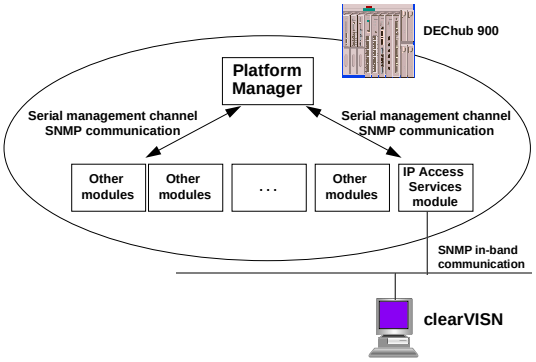

There is very little overhead incurred by the module when providing IP Access Services for the Platform Manager. Additionally, the architecture does not *force* a network module to provide this service — it is set up under the *control* and *choice* of the network manager. A key point of the architecture that bears emphasis is that the network modules are meant to provide the connectivity and other functions required for any network topology while the Platform Manager provides the value-added features to those network modules to complete the system.

This partitioning of functions also has other benefits. It is possible to have multiple IP Access Services modules active at the same time, thus providing either a *redundancy* function or alternatively offering *simultaneous communications* from network management applications in different subnets.

In the case where in-band communication between your Network Management Station and the IP Access Services module has failed for some reason, the Platform Manager provides the user with an OBM port whereby the chassis and all of the installed modules may be managed in exactly the same way (point-and-click GUI, SNMP), by the same *clear*VISN network management applications.

Lastly, from the perspective of MultiSwitch 900 chassis *availability*, the network modules in the chassis are not dependent upon a working Platform Manager to perform properly. The Platform Manager is independent from actual user connectivity. Testing this independence is as easy as upgrading the Platform Manager's flash memory anytime[13], even during normal working hours — during this process, not a single user packet is lost! Once upgraded, full management capability of the chassis and its modules is resumed. As mentioned in section 2.3, direct in-band management of each module *can* be performed via its SNMP agent, even during the Platform Manager's upgrade operation as long as there is an in-band connection from Network Management Station to each module —the *flexibility* is there to meet any customer need.

2.4 LAN configuration strategy, flexibility, ease-of-use

LAN Configuration/Reconfiguration supports the simple creation/deletion of multiple Ethernet, FDDI, Token Ring, ATM, VNbus backplane LANs today.

The backplane interconnect strategy for the MultiSwitch 900 chassis is very simple. All connectivity between modules occurs across the backplane and is under software control.

This capability allows front panel ports to be used for user or LAN segment connectivity with some high-speed links dedicated to backbone/server connections, or alternatively some ports may be redirected to the backplane[14] for inter-module communication; it is all under the control of the network manager. Certain modules allow additional configuration control by grouping capabilities at the physical layer with port and/or group switching while other modules allow for datalink layer grouping. Also, most LAN switches can be upgraded with routing software to allow software control of Layer 3 functions such as IP subnet connectivity, firewalls, and/or inter-VLAN communication.

Sample in-hub[15] configurations support many user environments, some of which are displayed in Figures 6 and 7. Multiple shared Ethernets interconnected by a homogeneous 10/10 switch; Personal switched Ethernets accessing one or more Fast Ethernet or FDDI servers through high-speed VNbus based switched connections; Token Ring clients accessing an Ethernet server through a switched Ethernet-Token Ring connection; Ethernet LANE clients accessing an ATM server; native ATM clients accessing a server farm on a connected GIGAswitch/ATM, etc. The key point is that all of these configurations, plus more, are supported on the MultiSwitch 900 chassis using the user-friendly *clear*VISN application suite.

[13] This is typically a 6 or 7 minute process.

[14] LAN Switches such as the DECswitch 900 and the VNswitch 900 allow the assignment of ports to either the front panel or the backplane.

[15] All switching is performed by the appropriate switch/router module so that users are not forced to provide an external switch/router to interconnect technologies

Figures 6 and 7. Sample *clear*VISN LAN Interconnect Windows

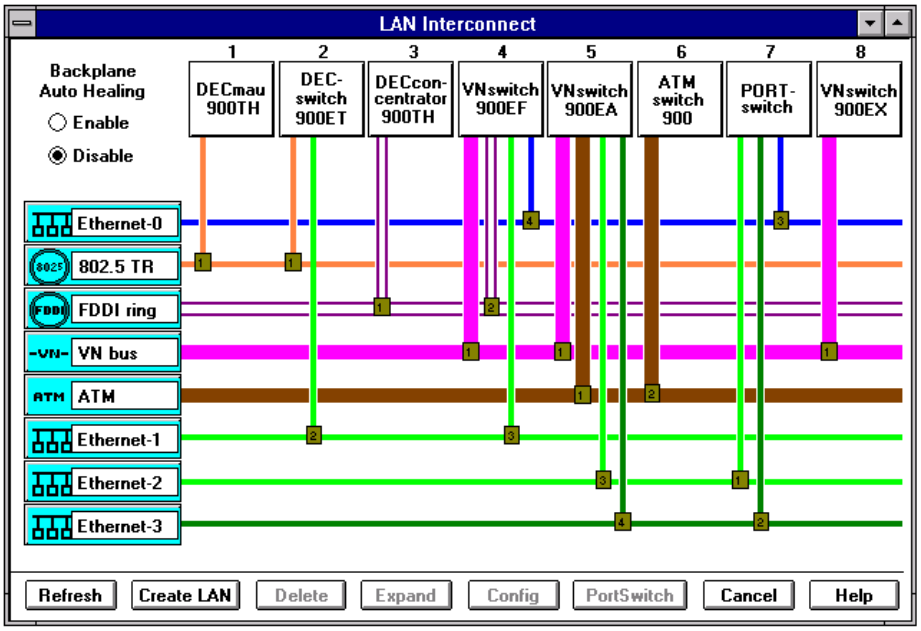

Supports any technology

2.4.1 Dynamic LAN Assignment

After creating the appropriate LANs on the DIGITAL MultiSwitch 900 backplane, users then have the need for simple configuration management and scalable performance. Dynamic LAN Assignment allows the network manager to configure and assign individual modules (or ports within modules) to one or more backplane channels through the point-and-click, drag-

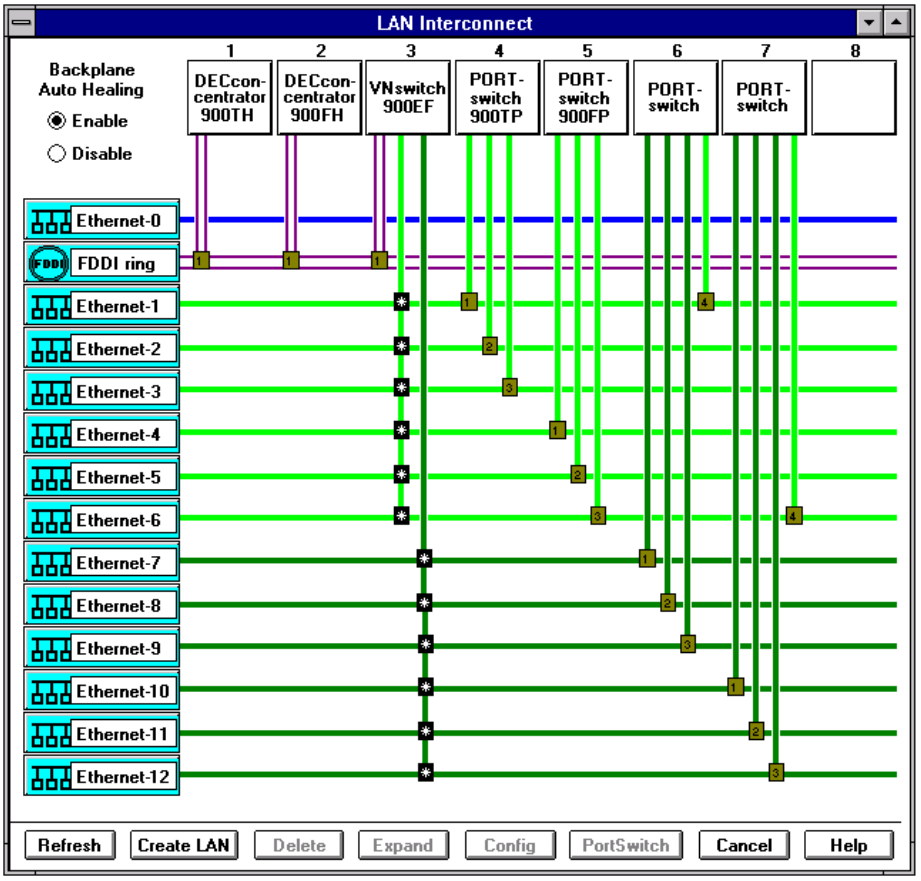

Supports a large number of Ethernet LANs

and-drop mechanism illustrated above.

Functionality, technology, and speed of transmission, is always a function of the actual network module. For example, DECrepeater 900 modules provide a single collision domain such that all users on that module may be assigned to only one LAN segment. On the other hand, PORTswitch 900 modules (sometimes referred to as configuration switches) support multiple collision domains and thus allow more granular port-level assignment of users to LAN segments. LAN switches such as the VNswitch 900 and DECswitch 900 families may be configured to either switch traffic between LAN segments at wire speed or route that traffic with the increased control that Layer 3 functions bring with it.

The main benefit that Dynamic LAN Assignment brings to the network manager is simplified moves, adds, and changes and optimized bandwidth as it supports the reassignment of individual users to new LAN segments as well as the configuration and/or reconfiguration of the switched internetwork with a simple "point and click" process from a *clear*VISN MultiChassis Manager console. The result is that there is no need to visit wiring closets or to physically move cables when a change is required.

Finally, Dynamic LAN Assignment provides an important configuration checking function in that it will either disallow or warn the user of technically incorrect configurations such as connecting an ATMswitch 900 to an Ethernet channel, creating a repeater loop, etc.

The power of the Dynamic LAN Assignment capability can be seen in a simple example where a user or set of users need to be isolated as an isolated Ethernet workgroup in the morning, and require connectivity to a new Fast Ethernet, FDDI or ATM backbone or server in the afternoon. Flexibility of LAN design and immediate scalable bandwidth availability become easily attainable goals by simply adding and configuring the appropriate module — hardware backplane upgrades are *not* part of the picture.

2.5 Availability/reliability/predictability

2.5.1 Fault Management

The MultiSwitch 900 chassis protects the user's environment from faults by providing redundant configuration capabilities, isolation of faults and the enabling of detection of faults through the *clear*VISN application suite.

One particularly unique fault management capability is Backplane Auto Healing, which consists of patching around a failed module and reconnection upon module reinsertion.

Reconnection upon module reinsertion is applicable for all technologies.[16] If a module fails and is replaced by an identical module, all previous connections are re-established.

The other aspect of Auto Healing is patching around a failed module. This is only relevant for FDDI networks. The MultiSwitch 900 chassis of course supports the innate fault tolerant capabilities of standard FDDI networks across the backplane in both ring and tree configurations; however, it *adds* considerable value by supporting this unique feature.

Figure 8. Port Types Used in FDDI

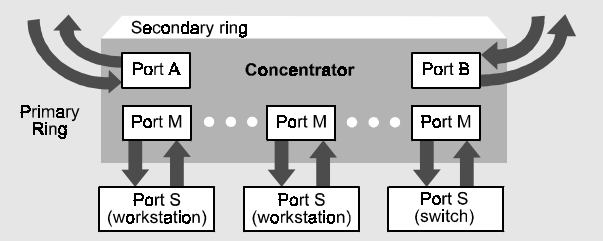

Ring Configuration

When a concentrator or switch in a dual ring is removed (hot swapped), the FDDI ring wraps according to FDDI standards. If the removed module has both A and B ports[17] that are connected to the DIGITAL MultiSwitch 900 backplane, the Platform Manager detects the removal and reconnects the two neighboring (upstream/downstream) FDDI stations as long as the Auto Healing feature has been enabled. Similarly, if a network module is detected as being unavailable (e.g. faulty), the Platform Manager will also patch-out the module and reconnect the appropriate stations.

If, at a later time, a replacement concentrator/switch is installed into the same slot, it will be automatically reinserted into the ring through the Auto Healing

[16] Reconnection of 802.5 Token Ring networks occurs regardless of the state of Backplane Auto Healing.

[17] There are four FDDI port types: the "A" port connects to incoming primary ring and outgoing secondary ring and is part of a Dual Attachment Station (DAS) or a Dual Attachment Concentrator (DAC); the "B" port connects to outgoing primary ring and incoming secondary ring and is part of a DAS or DAC; the "M" port is only implemented on concentrators and connects to a Single Attachment Station (SAS), DAS, Single Attachment Concentrator (SAC) or DAC; the "S" port connects a SAS or a SAC to a concentrator (DAC, SAC).

feature. This provides a distinct advantage to FDDI modules that are configured into a dual ring in a MultiSwitch 900 chassis versus those configured in other chassis. In the case of the MultiSwitch 900 chassis FDDI network with Auto Healing enabled, the single wrapped ring will automatically unwrap and reconnect around the failed module thus maintaining the original dual ring, whereas other chassis implementations which do not implement Auto Healing will simply wrap the network into a single, less resilient, ring.

It is important to remember that ring wrap occurs faster than backplane healing. Therefore, ring wrap occurs first after an FDDI module is removed or is detected as unavailable. The Platform Manager firmware thereby automatically reconnects the disconnected DAS stations through the DIGITAL MultiSwitch 900 backplane, thus allowing the dual ring to unwrap.

The above well-defined process provides the user with another example of MultiSwitch 900 chassis *availability* and *predictability*, this time in Fault Management situations. Fault recovery (unwrapping of the dual ring in the backplane) is supported on all DIGITAL MultiSwitch 900 FDDI modules.

Tree Configuration

Auto Healing is also supported in a similar manner when treed FDDI networks are configured across the backplane. Per FDDI standards, when a concentrator or a switch in a tree is removed (hot swapped), the FDDI tree splits into two independent networks. If the removed module has both M and S ports that are connected to the DIGITAL MultiSwitch 900 backplane, the Platform Manager detects the removal and reconnects the two neighboring (upstream/downstream) FDDI stations as long as Auto Healing has been enabled. If a network module is detected as being unavailable, the Platform Manager will also patch-out the module and reconnect the appropriate stations. If, at a later time, a replacement concentrator/switch is installed into the same slot, it will be automatically reinserted into the tree. In a similar manner to the dual ring case, Auto Healing of FDDI trees provides the distinct advantage that DIGITAL MultiSwitch 900 modules will be re-connected around failed modules across the backplane, whereas FDDI networks configured in chassis which do not implement Auto Healing would simply split into two independent networks.

2.5.2 Power System

Often relegated to the ranks of very reliable**,** important, but unexciting, the power system is a very integral system within the MultiSwitch 900 chassis. Typical power systems allow universal international operation, hot swappable units, redundant configurations and possibly even front access. The DIGITAL MultiSwitch 900 system provides all of these standard features along with a modular, highly available load-sharing (N+1) system, a fully managed power operation, and other significant attractions which will be described herein.

The MultiSwitch 900 chassis supports a maximum of four power supplies. Power may be added to the chassis incrementally, as needed, to provide power for the modules required by the network configuration. *All* power supply modules operate in a power sharing mode — this provides a tremendous reliability advantage over a system where the entire load is handled by the primary supply with the backup supply in “warm” standby.

The Platform Manager plays a key role in the allocation and control of power within the MultiSwitch 900 chassis. It obtains status from the power supplies for basic operation such as: Installed, Input OK, 48 V OK, and Outputs OK. In addition, the Platform Manager knows the power supply type and determines its resultant power output capability. It subsequently allocates the appropriate amount of correctly operating power to each installed network module. This power allocation process prevents system overloads and potential brownout conditions while permitting a flexible, incremental, “buy what you need” power system.

The Platform Manager controls the uninterrupted power services to all modules when/if a power supply fails in a (N+1) redundant power configuration. It also controls the well-defined “power shedding” feature which is discussed below. Power shedding supports the graceful degradation of chassis modules in the case where a power supply fails in a non-redundant power configuration — the result is total *predictability*.

Redundancy in the power system is obtained by adding one more power supply beyond the N supplies required to support the DEChub and all of its configured modules. (N+1) redundancy in the full capacity case is enforced by having the Platform Manager implement the policy that “in a system of four power supplies, the total power allocated shall not exceed the capabilities of any three units.” This allows for a failure of any of the supplies with no loss of communications. Thus, if one had a configuration which required three power supplies and a fourth were inserted, power output would be shared across all four units — if one of the supplies failed, the MultiSwitch 900 chassis and all associated modules would continue with *nonstop* operation since the three

remaining supplies would continue to provide 100% of the required power.

In MultiSwitch 900 chassis configurations where the customer has made the choice to implement no redundancy, if a power failure occurs, power will be recovered from the network modules by powering down modules from left to right, a term often called “power shedding.” This design gives the network manager complete control over the prioritization of the shedding of modules in the configuration — it is as simple as picking their slot number!

Additionally, upon replacement of the lost supply, the Platform Manager automatically brings the entire configuration back to full operation. This simple power shedding mechanism keeps the power system within its capabilities, allows simple effective control of the system in a very *predictable* manner, and provides another example of the *architected* nature of the MultiSwitch 900 chassis.

The Platform Manager has full visibility into the status/operation of the power system at all times and that information is provided to the user through a local LCD display (general chassis status, power system status, available system power) on the front panel and of course remotely through the *clear*VISN MultiChassis Manager application.

Additional capabilities of the system include redundant power cord operation with the use of Digital’s POWERswitch (dual AC input to single AC output) and/or the standard 48 Volt DC power connector for those customers who have 48 volt inputs, often used for telecommunications equipment.

3. clearVISN Network Management

clearVISN is Digital's strategy toward "Lights Out Network Management," delivering on the enVISN promise of centralized, policy-based network management with distributed intelligence. There are four major components of the clearVISN program.

- Policy-based network management — Policies are high-level rules and procedures that simplify the management of complex network devices and technologies. The clearVISN policy-based methodology treats the network as a whole, thus freeing the network manager from dealing with the complexity of individual devices and technologies.
- Distributed intelligent agents — Intelligent agents in Digital's network devices implement policies and move control closer to that device. This capability allows for very scalable management domains.
- Common core — All clearVISN applications come with a "common core". This core provides common installation and user interfaces plus data integration and event infrastructure giving clearVISN customers true control over that information.
- Multivendor legacy management — clearVISN facilitates the migration toward switch-based virtual networking while giving network managers better control of their routed inter-networks by providing consistent GUI management of Digital’s as well as a wide array of other vendors’ routers and switches.

4. Future Directions

Imagine a network which readjusts bandwidth needs based on the needs of the current applications and users. The architecture of the MultiSwitch 900 chassis provides the fundamental underpinning of the structure needed to create just such a network infrastructure. From power cords to backup network links, ring wrapping and load sharing, this environment lets you build the reliability and availability you need for your application.

Today you can configure and re-configure your MultiSwitch 900 chassis under software control from across the room or across the world. The future plan is to enable an environment that is completely "Lights Out." This goal will be achieved by working in concert with Digital's *clear*VISN Network Management development team to take full advantage of the MultiSwitch 900 chassis' unique software controlled backplane.

In closing, it is important to emphasize a powerful subtlety in that last statement. That is, that the DIGITAL MultiSwitch 900 is a chassis built *not just for* management of today's diverse environments, but *additionally* to enable *increasingly* sophisticated algorithms to monitor the environment for performance, configuration, fault, security and accounting information, to determine the best allocation (and reallocation) of resources, and then to actively *adjust* or *configur*e your network for both maximum efficiency and optimal performance.