Teamwork Creates "Made In Canada", Next Generation Telecom Solution

by *W. M. (Mike) Campbell, Nortel Networks,*
Paul Holt, Noel Lesniak and Jerry Sidman,
Motorola Computer Group,

This article describes the technical challenges faced by the engineering teams behind a Canadian technology success story, the solutions each of the teams arrived at, and the new global telecom solution that resulted from their collaboration.

Cet article traite des défis techniques qu'ont du relever des équipes d'ingénieurs responsables d'une percée technologique canadienne, des solutions trouvées par chaque équipe et de la nouvelle approche globale en télécommunication qui a résulté de leur collaboration.

1. Introduction

Global deregulation of telecommunications, privatization, and the growing number of multiple service providers continues to stimulate competition and growth within the telecommunications sector. Established and emerging service providers alike are demanding world-class technology and reliable network solutions to meet the needs of all network applications.

Extensive market research conducted by Nortel Networks showed that service providers wanted enhanced operation, administration, maintenance and provisioning (OAM&P) capabilities. Telephone companies needed to significantly enhance their ability to manage expanded billing and service tasks in near real time - obligations largely being driven by the global market changes, such as local number portability and deregulation.

Nortel's 1995 analysis of where network evolution was heading also considered advances in the internet, intranet and TCP/IP capabilities and trends in computing technology where computing costs were dropping and computing capacity was increasing. Central office switches continue to get larger, with features and capabilities expanding at an increasing rate. As a result, overall data generation is growing, more billing information is produced and data warehousing initiatives are underway.

2. State Of The Art

Worldwide deregulation of the telecommunications industry has attracted many new service providers to the business who have limited expertise in operating and administering central office switches. Established carriers and network operators are also beginning to augment high capacity, centralized "Operational Support Systems" (OSS), with a client/server architecture, thus there is more distributed computing available on the desktop and in the network.

Based on this customer intelligence, Nortel's challenges included:

- Connectivity via TCP/IP, not point to point or X. 25
- Significant increase of data through put in and out of switches
- Deregulation/competition driving simplification
- Traditional user moving from ASCII terminal to workstation/PC on the desktop

3. Specification Objectives -- Nortel

The Nortel SuperNode Data Management (SDM) team concluded that adding a dedicated Operations, Administration, Maintenance and Provisioning (OAM&P) platform to the DMS architecture was the logical next step in the product's evolution. The team's challenge was that call processing is a "real time" application, while these new requirements were more akin to a "data processing" environment.

The solution was a dedicated server to provide high data throughput, storage, processing and communications that would not jeopardize the DMS' market-leading reliability and robustness. The computing platform also needed to meet the stringent requirements of "NEBS" imposed on central office equipment.

Nortel customers take the reliability of the telephone network very seriously, and Nortel had designed the DMS switch for complete fault tolerant operation from its inception. The fact that most people can't recall when their phone service didn't work, because of a central office switch problem, is a testimony to the company's commitment to quality and reliability. Because few computer vendors could meet Nortel's stringent requirements, the team debated whether to build a fault tolerant system in-house or settle for a less reliable simplex product from a traditional computer vendor.

While Nortel was defining its specific SDM product, Motorola Computer Group was developing a fully fault tolerant platform designed for the telecom central office market to be called the FX Series. When Nortel approached MCG both teams soon realized they held the key for the other's success. Nortel knew it could bring the product to market far faster using a soon-to-be available fault tolerant, central office compliant system that could support commercial software products (e.g. protocol stacks, databases, expert systems and security packages).

Motorola saw the opportunity to collaborate with a partner who had extensive fault tolerant experience and who also understood the requirements for computing platforms in central office telecom applications, a collaborator that also had an immediate need for such a product. It was a perfect fit.

Confident in each other's commitment and culture of engineering excellence, Nortel and Motorola commenced collaborative development of the first truly fault tolerant, central office compliant AIX/Power PC-based product. Despite the significant engineering challenges they knew lay ahead, the teams set an ambitious target for product introduction of the SuperNode Data Manager in 1997.

As the teams set out to meet their goal, Nortel immediately brought MCG together with customers who would eventually use the equipment to ensure the fault tolerant system would not only serve an application platform purpose but would also meet the operational needs of customers in central office working environments.

In December 1997, Motorola Computer Group (MCG) recognized the team from Nortel Networks with the:

"Motorola Design of the Year Award"

This award is given annually for the most innovative product designed in Canada by an original equipment manufacturer, based on MCG Technology.

4. Nortel Challenges

Foremost among Nortel's challenges was the successful - physical and operational - integration of the existing system and the new OAM&P applications. The application software had to interface gracefully with the call processing system and all of the underlying messaging, maintenance, alarm and user interface subsystems had to be integrated to provide a seamless subsystem extension to the overall DMS system. Hardware packaging, power, grounding and physical alarms likewise needed to be integrated.

The requirement to address marketplace expectations and also consider the needs of a large installed base presented challenges for the SDM team. Changing system interfaces from traditional RS232 asynchronous or X. 25 interfaces to a TCP/IP based LAN interface - while maintaining full support for the legacy interfaces - would also be a challenge. Because the new platform was targeting an installed customer base and not all of the switches within a given network would have this new capability, it had to be optional.

Customers implementing this new capability needed a transition approach that supported simultaneous, dual operation of various interfaces. For example, given the critical nature of billing data and the multitude of record types and formats, a flash cut-over was not considered a viable transition for customers to move from an X. 25 based interface to an FTP based Ethernet interface.

The user interface posed similar challenges. Customers did not want new commands for existing capabilities on which employees were already trained. Yet, they wanted to move from command line, ASCII terminals to graphical, point-and-click workstations and/or PC-based terminals. Customers were also clear that while legacy screens had to retain the same look and feel, new capabilities should only be provided as a graphical user interface (GUI).

Also associated with the user interface transition was the subtle challenge of security. Often overlooked, user ID's and passwords not be passed as clear text across a LAN. This is akin to sending credit card numbers over the internet - some people do not worry about information being intercepted, while others would never risk it without encryption security. Despite these challenges, the first Beta system was delivered within 18 months. Dividing the implementation into focused, manageable activities was part of Nortel's development strategy.

The key hardware development focused on creating a new I/O interface module to be housed within the SDM which terminated the transmit and receive fiber links on the new subsystem and converted the DMS proprietary interface to link with the FX Series fault tolerant system.

The SDM team built application programming interfaces (APIs) - comprised of software in the DMS Compute Module and structured interfaces in the AIX-based SDM - to interface with the OAM&P subsystems such as billing, performance data, event/log messages, alarms, provisioning data, file systems, etc. These APIs were critical to achieving wide acceptance across all of Nortel's switching products and ensuring customers around the world would benefit from a fully integrated and supported platform.

This design now enables development teams around the world to add market-specific applications required for a particular service provider. Application software can be added, removed or augmented as required without incurring development in the underlying hardware, operating system or messaging subsystems.

The architecture team relied heavily on the client server architecture and used commercially available Distributed Computing Environment (DCE) software to address requirements such as name server simplification, centralized user ID/password administrator, time of day server, and the major challenge of security and password encryption. The concept of Remote Procedure Calls (RPC) made it possible to support the needs of network management and OSS systems very efficiently while also enabling an efficient approach for graphical user interface development to utilize the same RPC-based interfaces.

Utilizing web servers and browsers also provides the ability to address security, while coupling user access via intranets or internet to any DMS central office equipped with the SuperNode Data Manager (Fig.1).

5. Specification Objectives - MCG

A central challenge in developing the FX Series was the addition of fault tolerant features to the operating system while still maintaining

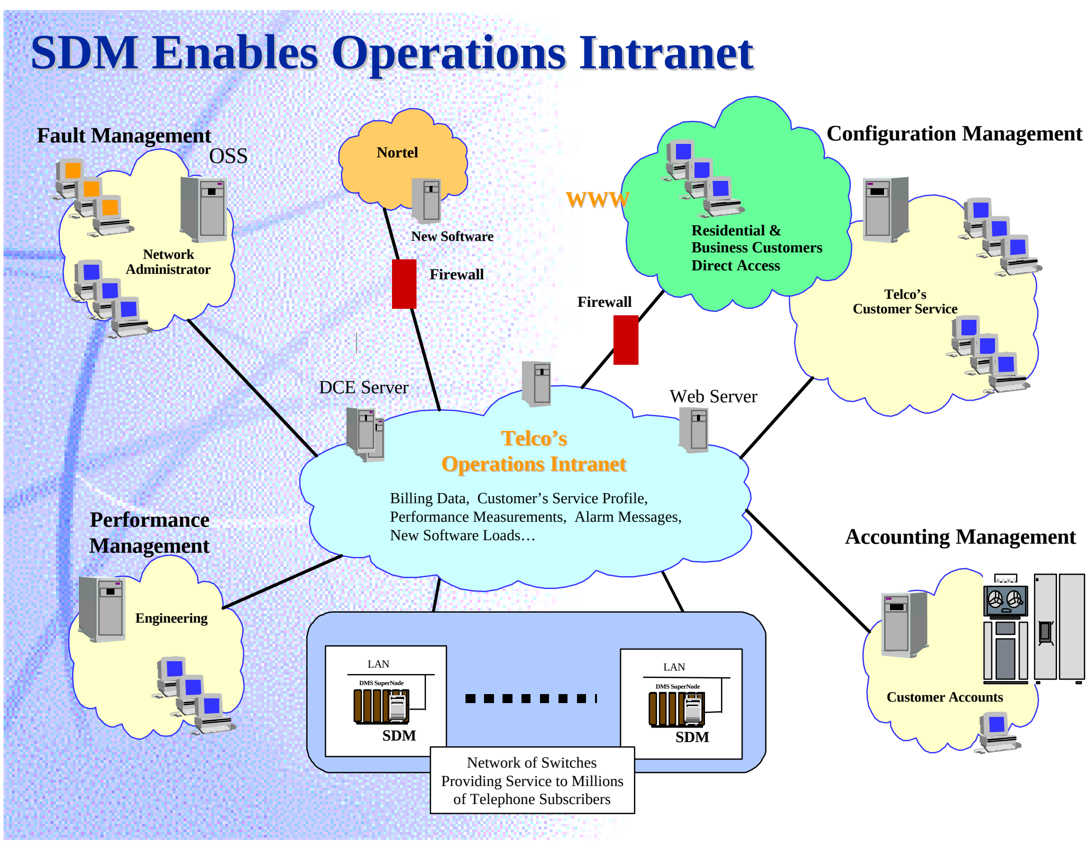

Figure 1: SDM Enables Operations Intranet

full binary compatibility with IBM's UNIX offering called AIX. This achievement would provide Nortel with a fault tolerant system capable of running standard, off-the-shelf software from IBM and other companies.

The areas of AIX requiring enhancement for fault tolerant operation were:

- Support for a trusted processing core
- Hardening and switchover capability for device drivers
- Intelligent rules based response for the AIX object data manager
- Enhancements to the AIX logic volume manager to manage the FX Series mirror disk subsystem
- Addition of Split Mode™ functionality

A. Support for Trusted Processing Core

The trusted processing core contains the PowerPC and main memory in an FX Series CPU Module (Figure 2). These are configured for reliability with redundant CPU Modules running in cycle-by-cycle lockstep and checking on each other.

When CPU Modules execute I/O transactions to or from the I/O domains, they check their results. Any detected faults in the processor core cause an exception, and the corresponding exception handler performs the initial recovery from the fault. This mechanism enables all application programs, device drivers, and the operating system kernel to be free of features relating to processor core fault tolerance. To support this architecture Motorola added an AIX kernel extension and specified firmware to manage the CPU Module through synchronization faults, fault diagnosis and isolation, and subsequent re-synchronization.

B. Hardening and Switchover Capability for Device Drivers

Standard UNIX device drivers are not hardened and do not support switchover between redundant devices. Adding this functionality to FX Series device drivers represented a major enhancement. Hardening is implemented in the physical layer of the device since it deals directly with the hardware. Switchover is involved with maintaining the service through the logical driver and is usually implemented in the logical layer.

Main features which shape the architecture of an FX Series device driver are:

- **Hardening:** The most critical step in the process of achieving continuous service availability is implemented so faults in the underlying hardware will not crash or hang the system. The FX Series assumes hardware will fail in strange and complex ways and the driver manages these errors.
- **Fault Detection:** The driver is the most logical place for the first level of fault detection and, in most cases, errors can be detected virtually at the point of hardware failure. Placing switchover functionality in the driver ensures a short time line from fault detection to recovery.

  The driver categorizes hardware and software errors as:

  - Data corruption errors: Designers must isolate areas of the driver sensitive to data corruption from bus errors or other causes, and develop protection.
  - Passive and assertive protocol violations: Passive protocol violations usually occur when an I/O module has failed or halted and is simply not responding to external events. In fault tolerant environments hardware can fail at any time, hence any part of the driver that waits in a loop-type structure for an external event on the I/O module should be protected with a time-out.

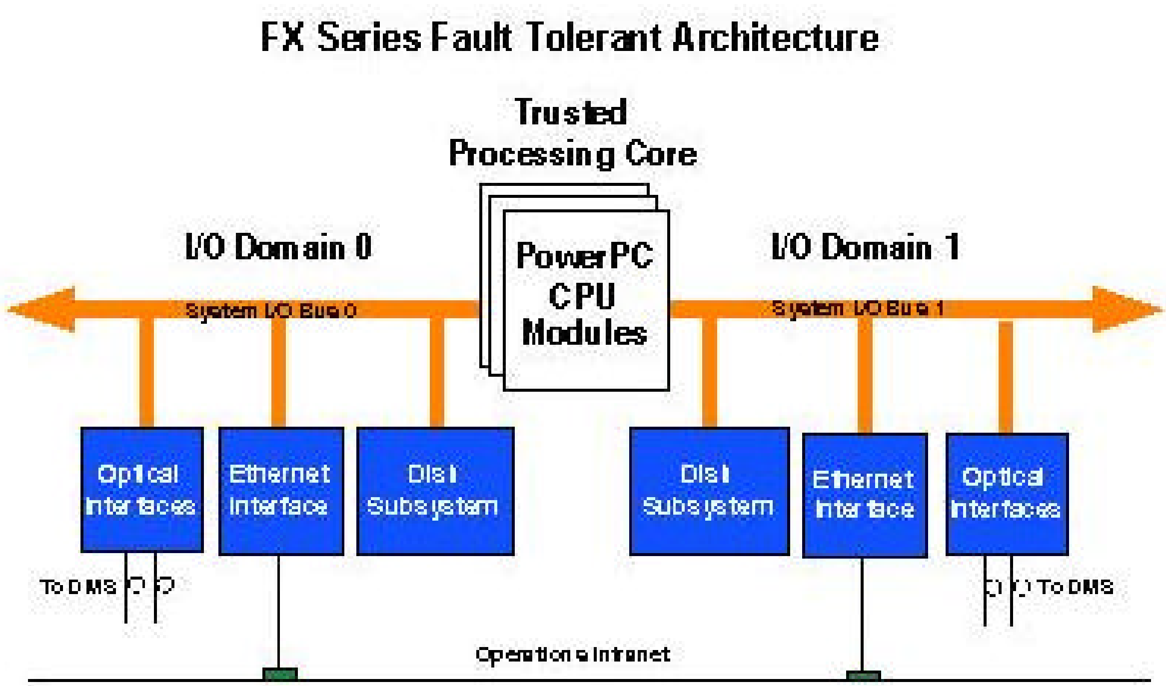

Figure 2: FX Series Fault Tolerant Architecture

The DS512 I/O module had to be implemented in a redundant, fault tolerant fashion to ensure no single point of failure associated with connectivity, and it had to be "hot swappable" like all components of the system. New software drivers were developed to support this DS512 control hardware and the proprietary messaging protocols used within the DMS architecture.

An assertive protocol violation may be the first symptom of a data corruption failure mode, in most cases, a simple check against expected behavior is all that is required

- Bus related errors: After accessing a device capable of causing a bus error, the driver checks to determine if such an event has occurred. If so, the system protects itself from taking an actual processor exception and sets a flag the driver can examine. The driver responds by rewinding to the last known good state and, if appropriate, repeats the transaction using a redundant module.

- **Time-out errors**: I/O device failures that would otherwise be missed until the I/O subsystem was accessed are detected by Watchdogs which can be implemented in either hardware or software. Software watchdogs are either static or dynamic. Static watchdogs periodically access a data area on the I/O device containing static data to provide an indication of an overall failure of the bus interface or of the module being physically removed

  Dynamic watchdogs use counters or timers in a critical area of the protocol to indicate it is running, usually linked to a time-out in the driver. An absence of activity for a certain time signals an error. Hardware based watchdogs include a free running counter on the I/O device that is periodically reset. If the firmware fails, the counter expires and triggers an interrupt line to alert the driver or leave a message for the driver to pick up.

- **Switchover:** This process of moving operations from one module to another provides a mechanism to recover from a fault using redundant hardware. In some subsystems data can be lost during fault detection and switchover. To maintain data integrity, a transport protocol is layered on top of the unreliable data transfer mechanism provided by the I/O subsystem, either within the driver or externally. It detects and corrects data transfer errors and can also notify the subsystem's device driver that an error has occurred and a switchover should be performed.

- **Event Reporting:** The driver reports state changes to a configuration management subsystem.

C. Intelligent Rules Based Response for the AIX Object Data Manager

The AIX Object Data Manager (ODM) is a device database used for device configuration and management. It is divided into two main sections -- pre-defined and custom data. Pre-defined data contains information on all devices that could exist on the system and custom data contains information about all of the devices actually in the system. At boot time the custom database is populated from the pre-defined database depending on what devices are seen by "probe" functions.

While a standard ODM works well in a relatively static system, it is not effective in fault tolerant systems where devices may "come and go" without any reboot and where the system must respond intelligently to changes in device state. Motorola needed to enhance the ODM to handle dynamic changes proficiently and to respond to those changes through a rule set. This was accomplished through a combination of enhanced configuration methods and creating the Configuration Management System (CMS) -- a finite state machine that models the configuration of the system.

As well as invoking device specific actions, CMS also manages all visual status indications and telecom alarms. This ensures rapid and consistent response across all devices, a significant feature for telecom fault tolerant environments. The combination of the ODM and the CMS provides a very powerful system management and fault response system while still maintaining compatibility with AIX's ODM.

D. Enhancements to the AIX Logic Volume Manager

The FX Series system can support multiple mirrored disk subsystems. Motorola decided early in the development process to use the AIX Logical Volume Manager (LVM) to support the underlying mirroring functionality. This saved time and also reused a proven, integrated mirroring subsystem.

While the LVM was sufficient to manage a mirrored configuration in a standard system, Motorola found the fault tolerant FX Series needed particular characteristics to manage mirrored disks. Because FX Series

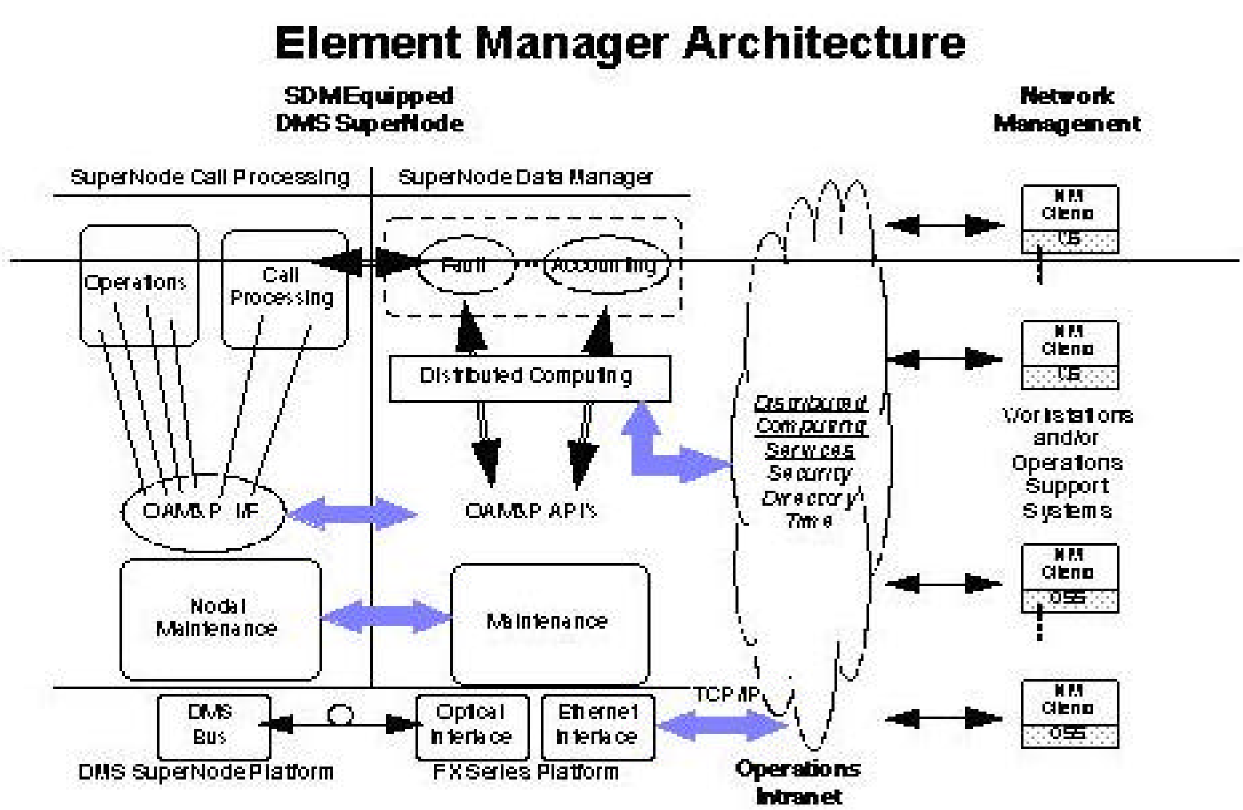

Figure 3: Element Manager Architecture

disks can be replaced without rebooting the system and matched pairs were somewhat derived from the geographical position of the modules, the LVM sub system needs to respond differently. This was addressed by creating a FT Volume Group with an associated set of rules for configuration and management. The FT Volume Group sits above the normal AIX LVM volume group and presents an interface that allows users to manage and configure the disk subsystem at the FT Volume Group Level.

E. Addition of Split Mode Functionality to AIX

Because upgrading operating systems, application software, or CPU Modules in standard UNIX-based systems can cause significant out-of-service time - often for many hours - the FX Series required Split Mode operation to enable half of the system to be taken out of service and upgraded while the other half continues to provide service. When confidence in the new half is established, service can be transferred with minimum downtime - as little as 16 seconds. To support Split Mode on the FX Series, Motorola had to add split mode functionality to AIX.

F. Device Configuration

Device configuration analysis utilities were added to ensure the system could be safely split and still carry out the split operation. Later these utilities would bring the system back to full fault tolerant mode and reconstruct the device tree.

G. Split Mode Communication

To manage the Split Mode, Motorola added a low level driver interface and application level library to enable the applications and Split Mode management utilities to communicate with each other.

H. Split Mode Management Utilities

There are a number of steps in the Split Mode process during which Split Mode utilities, system functions, and user applications need to be coordinated. This process is managed by the Split Mode Management Utilities that provide features such as registered notification of Split Mode states and fallback to previous states. It also enables easy integration of OEM applications into the Split Mode environment.

I. Boot Modifications

Both the firmware and AIX boot functionality were modified to support booting of the FX Series while in Split Mode. This prevents the upgraded half of the system from grabbing the devices being used by the half of the system still providing service.

6. Joint Testing Benefits Collaboration

The MCG fault tolerant program benefited greatly from the joint test and qualification activity shared between the Nortel and Motorola test organizations. The test teams were able to focus on different aspects of the test plan. Nortel's team focused on the operational aspects and how the system would handle a broad range of error events which can occur in the central office. The Motorola team was then able to focus on the operation of the system architecture itself under heavy stress and fault insertion scenarios. The joint effort allowed complete test coverage in the shortest possible time.

By February 1997, a fully functional beta system was installed and operational in the lab of a major Canadian customer before either company had formally announced their respective products and less than 18 months from the start of the project. After initial testing with this lead customer, the products were officially released for in-service application in May 1997.

7. Conclusion

The combination of Nortel's switching experience and application focus and MCG's NEBS compliant, fault tolerant computing platform expertise has enabled the switch world to be connected via high speed TCP/IP links to telco intranets.

The SDM represents the next generation in telecommunications operational interface capabilities. For the first time, service providers around the world can access a highly reliable and continuously available server over their "operations intranet" that meets their accelerating customer-service requirements both efficiently and cost-effectively.

Customer acceptance of the new technology has clearly demonstrated the collaboration was a success, with product now being deployed in Canada, the U. S., Europe, Asia, Australia and the Caribbean. The basic system designed and introduced in Canada has become a building block for Nortel designs around the world to add new functionality for specific customer requirements.

An expert system manages alarms and events associated with the entire switch and all peripherals, simplifying maintenance. A database with high speed OSS interface and GUI-based human interface simplifies the provisioning of features and services assigned to end customer lines- often in excess of 100,000 lines per switch. High capacity storage and real time delivery makes it possible to handle well over one million billing records per hour as well as all performance monitoring information associated with the service levels provided.

This initiative will quickly lead to telecom providers offering customers direct access, via the internet, to retrieve their latest billing data or review the service profile associated with their telephone service. It's also expected that the SDM initiative will spawn new opportunities for service providers and third-parties building on the fast, efficient availability of data generated by telephone calls. These opportunities will range from real time billing systems, fraud detection systems, customer usage profiling, along with new and improved end-user access to their own feature profiles and line usage information via simple internet - point and click - from their home computers.

The collaboration between Nortel and Motorola is an excellent example of how such initiatives can produce win/win solutions for carriers, network operators, and their customers.

8. Legend - List of Acronyms Used

API	Application Programming Interface
CMS	Configuration Management System
CPU	Central Processing Unit
DMS	Digital Multiplex System
FTP	File Transfer Protocol
GUI	Graphical User Interface
I/O	Input / Output
LAN	Local Area Network
LVM	Logical Volume Manager
MCG	Motorola Computer Group
NEBS	Network Equipment Building Standards
OAM&P	Operation, Administration, Maintenance &Provisioning
ODM	Object Data Manager
OEM	Original Equipment Manufacturers
OSS	Operational Support Systems
SDM	SuperNode Data Manager
TCP/IP	Transmission Control Protocol, Internet Protocol

9. References for Further Reading

[1]. FX Series Fault Tolerant Systems Architecture Overview

[2]. FX Series Split Mode Overview and User's Guide

Both above documents are available in PDF format at the following URL: http:\\www.mcg.mot.com/literature/PDFLibrary.

[3]. SuperNode Data Manager Planning Guide is available by calling 1-800-4-NORTEL or at www.nortel.com

DMS SuperNode Data Manager

Nortel's fully-featured DMS SuperNode switch (see also cover picture) is deployed in thousands of central offices around the world and meets the needs of the most demanding network applications. Capable of handling in excess of one million calls per-hour, the DMS family of switches support:

- Local (residential and business)
- Intra LATA Toll (long distance)
- Inter-exchange & Global Carriers
- International Gateway
- Wireless
- Signaling Transfer Point (STP)
- Operator Services

Photo of DMS SuperNode Data Manager

The Fault Tolerant SDM SuperNode is fully integrated into the DMS system line-up, power, grounding and alarm subsystem. The SDM is maintainable to the circuit pack level without shutting down the dual power feeds and is FCC EMI, NEBS and Zone 4 Earthquake compliant. Software is upgraded while the system remains in service.

About the Author(s)

W. M. (Mike) Campbell, has worked for Nortel Networks for 20 years in various roles. He is currently Director, Business Line Management, responsible for DMS OAM&P Evolution and Core Hardware and Software.

Mike is an Engineering Science graduate from the University of Toronto and has an MBA from York University.

Paul Holt, VP/Regional Director, Central US & Canada Motorola Computer Group, is responsible for sales, distribution, support, marketing and finance for Motorola Computer Group, Technical Products Division for the Canadian and Central U. S. A. region. Major centres under his management include Chicago, Cleveland, Minneapolis, Milwaukee, and St. Louis in the U. S. A. and Vancouver, Ottawa, Calgary, Montreal and Toronto in Canada.

Paul received a MBA from York University and a Bachelor of Arts in Economics from the University of Western Ontario.

Jerry Sidman, Engineering Section Manager Motorola Computer Group, is responsible for all AIX and UNIX development at Motorola Computer Group for VME and FX Series Fault Tolerant computer systems. In addition to managing a number of development departments he leads the software process improvement effort for MCG's Technical Systems business unit. Jerry joined Motorola in 1992 and has held various engineering positions at Motorola.

He holds a Bachelor's Degree in Electrical Engineering from Newcastle University, and a MBA from Durham University.

Noel Lesniak, Product Line Manager, Telecommunications and Fault Tolerant Systems, has been with Motorola for 18 years in various roles. He is currently product line manager in the Motorola Computer Group with responsibility for the FX Series family of fault tolerant system platforms.

He attended the University of Southern California prior to serving at the U.S. Military Academy at West Point.