Advanced High Availability Software For Linux - 5NINES Telecom Deployments

by *Robert Pettigrew and Noel Lesniak*

Motorola Computer Group

Abstract

This article provides an overview of the telecom industry's first advanced high availability software for Linux (HA-Linux) applications. Recently launched by Motorola Computer Group (MCG), the Linux solution provides "hot swap" and active/standby processor switchover capabilities for carrier-grade networking, wireless and Internet applications that run 24-hours a day, 365 days a year.

Sommaire

Cet article présente le premier logiciel à haute disponibilité de l'industrie des Télécommunications pour les applications Linux (HA-Linux). Motorola Computer Group (MCG) a récemment lancé sa solution Linux qui permet d'extraire à chaud et de changer le statut du processeur du mode actif au mode veille. Ce logiciel sera particulièrement utile pour les distributeurs de réseaux, spécialement pour les applications sans-fil et pour l'Internet qui sont disponibles 24 heures sur 24, 365 jours par année.

1.0 Introduction

The unprecedented adoption of the Internet as a business and consumer "utility" has resulted in explosive growth in the wireless and Internet sectors of the telecommunications industry. Convergence of data and voice - and soon multimedia communications - are driving the need for telecommunications embedded computer systems that must operate all day, every day, with almost no interruption. With Internet and e-commerce becoming as ubiquitous in our everyday lives as the telephone, loss of service has become intolerable since it represents lost revenue, a lost transaction or worse yet a lost customer.

Today the world's telecommunications infrastructure markets are clamouring for "off-the-shelf" computing solutions that enable them to cut costs, reduce time to market and meet international carrier grade telecommunications standards for 5NINES availability. 5NINES availability equates to 99.999% availability - five minutes or less of planned and unplanned downtime per year.

With its very low cost and freely available source code, Linux is being rapidly adopted as the platform of choice for new telecom applications. Until recently however, the software components required to implement a high availability Linux solution - supported by service, training and systems integration capabilities - were not available. MCG has built on its fault tolerant, embedded and telecom systems expertise to develop advanced high availability for Linux that meets the needs of mission critical applications requiring 99.999% service availability. Operations, Administration and Maintenance (OA&M) platforms, call servers, IP gateways, gatekeepers and home location registers are some of the applications where the benefits of high availability Linux is projected to have significant impact.

MCG's advanced HA-Linux framework (Figure 1), coupled with the CPX8000 family of CompactPCI-based systems enables the development of 5NINES-capable systems.

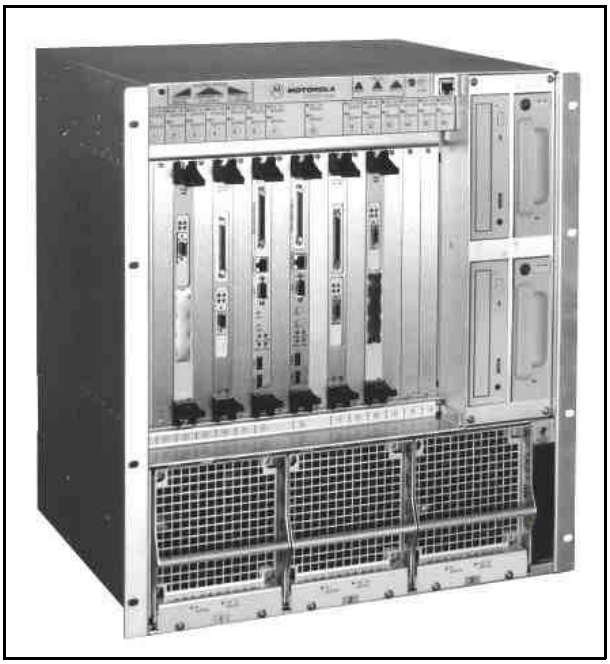

Figure 1: "Off-the-shelf" CPX8216 Carrier Grade HA Platform

2.0 Linux: Ready For Prime Time

Only seven years old, Linux is now the world's fastest growing operating system largely due to its stable, open architecture. Distributed under the GNU General Public License, a "free" license, anyone can modify the Linux source code and redistribute it. Although Linus Torvalds owns the trademark for the name and leads kernel development, no one owns or controls it and the system has therefore benefited from the contributions of many talented developers.

In 1994, Linux visionary Mark Bolzern, president and founder of WorkGroup Solutions, Inc. and a member of the non-profit Linux International board, predicted Linux would be as revolutionary to operating systems, as the PC was to hardware, by providing a common standard everyone could access [1]. Data released by industry analyst IDC suggests he may be right. Last year IDC reported that in 1999 Linux server license distributions exceeded 1.4 million copies, doubling the 1998 shipments. At the same time all shipments of Unix server licenses totaled only 839,000, growing just 1.3% over 1998 (Figure 2). IDC further predicted that Linux shipments would grow at a CAGR of 25% through 2003. Among the many major technology companies supporting the Linux O/S are Dialogic, IBM, Informix, Intel, Motorola, Natural Microsystems, and Oracle.

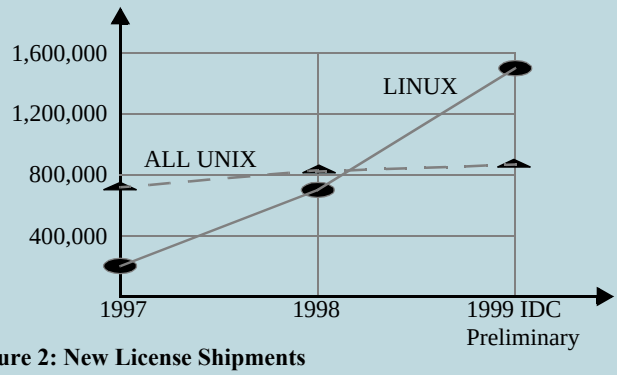

Figure 2: New License Shipments

While there has been considerable discussion about the challenge Linux represents in the server and desktop markets, MCG's introduction of HA-Linux now makes it a compelling choice for Original Equipment Manufactures (OEMs) building for mission critical industry sectors like data storage, medical imaging and telecom. Not only does Linux provide a full-featured O/S with proven reliability, the open source model enables lower cost, greater control and simplified licensing.

3.0 5NINES In The Telecom Environment

5NINES availability is critical for a growing number of telecom applications such as 911 and Emergency Services; Call Control and Setup; IP Telephony; Packet Switching; billing; OA&M; and e-commerce where the web store never closes.

Even so, application software and operating system maintenance and upgrades are always necessary. For example, software updates are done, on average, twice a year in the telecommunications industry. This could result in up to 60 minutes of downtime annually, potentially putting 5 NINES beyond reach. Achieving 5NINES (five minutes or less of downtime a year), for both planned and unplanned events, requires:

- Increasingly more stringent qualifying component reliability and a service-minded system design - a 3NINES platform,
- A combination of redundancy throughout the hardware architecture (including the CPU), combined with fault management and warm or hot restart models - 4NINES applications, and
- An architecture that supports hot upgrade of software - 5NINES.

4.0 MCG's Advanced HA-Linux Solution

MCG's HA-Linux - with “hot swap” capability and support for both Intel- and PowerPC-based system platforms - is the first Linux offering for carrier-grade networking, wireless and Internet applications that require 5NINES availability. HA-Linux refers to an integrated set of Linux kernel and application level components that provide high-availability (HA) functionality on MCG's CPX8000 Series of carrier grade, NEBS/ETSI compliant platforms.

The CPX8000 Series is CompactPCI-based. CompactPCI has been widely adopted by the telecom industry, bringing the power and performance of low-cost computing technology to this mission critical sector. A primary driver has been the ability to hot swap boards. When coupled with component redundancy schemes, it offers inherent fault recovery and online service leading to a HA system.

CPX8000 is a family of carrier grade, NEBS compliant system platforms that feature an advanced HA architecture. The core of this advanced architecture is the hot swap controller and bridge module that allows the system processors to access both I/O domains under control of HA-Linux. This patent pending facility enables mission critical applications to rapidly switch over to a standby system processor and to access both I/O domains after a system processor failure making the CPX8000 ideal for 5NINES applications.

The advanced features of HA-Linux allow OEMs to develop and network operators to deploy and operate applications that meet or exceed 5NINES. Planned upgrades of hardware and software can be implemented without disrupting services.

In addition to the ability to “hot swap” system components like processors, I/O controllers, power modules & fans, HA-Linux provides:

- Active stand-by system and CPU support, enabling applications to continue service while switching to a standby processor,
- Management of telecom alarms and component status LED's, allowing easy and error free operations by craft personnel,
- Simple network management protocol (SNMP) support with an SNMP agent, system MIB and Event Manager MIB, enabling system state changes and events to be monitored and managed by the network and operations management system, and
- An inter-system communications interface that allows applications in the primary and backup system processors to communicate with each other to facilitate rapid switch over to the backup.

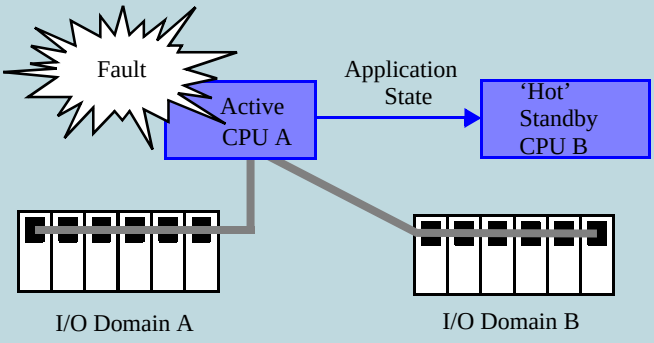

Figure 3: Advanced HA Architecture

4.1 Intel x86 and PowerPC Support

HA-Linux capabilities are offered on MCG's Intel x86- and PowerPC-based processor modules, giving OEM's the choice of processor architecture that is best for their application. HA-Linux runs with the Red Hat distribution in the Intel x86 environment and with the Linux PPC kernel in the PowerPC environment.

4.2 “Hot Swap” of All System Components

The CPX8000 system architecture features redundant “hot swap” components for all active system elements (field replaceable units). HA-Linux controls and manages these redundant “hot swap” components to enable applications to continue providing service even when a system component has failed. This “hot swap” support allows any system component to be switched over to a backup component in order to continue operation. Then the faulty component is replaced with a spare and subsequently returned to active operation.

4.3 System Processor Switch Over

HA-Linux enables high availability applications to switch over from the active system or CPU module to a standby system or CPU module in order to continue service. HA-Linux uses the CPX8000's Hot Swap Controller and Bridge to switch the CPX8000's I/O domain to the standby CPU module in order to rapidly begin application processing on the backup CPU module with access to all the system's I/O components. HA-Linux provides heartbeat, messaging, and checkpoint services to aid in developing high availability applications for this environment.

4.4 Telco Alarms

HA-Linux fully supports the critical, major, and minor visual and dry contact telco alarms on MCG's carrier grade, NEBS compliant system platforms. These telco alarms give craft and operations personnel easy to understand indications of system status conditions that may require their attention. The telco alarms are turned on and off based on rules in the configuration database and component status changes as they occur in the system. These alarm changes can also be communicated through an interface to SNMP.

4.5 In Service and Out of Service LED's

HA-Linux supports In Service (Green) and Out of Service (Red) LED's associated with each component in the CPX8000 system, allowing craft or operations personnel to easily see the status of components in the system and to quickly and correctly identify out of service conditions. The service status changes can also be communicated through an interface to SNMP.

4.6 SNMP V3 Agent, System MIB and Event Manager MIB

HA-Linux includes an SNMP agent, System MIB and Event Manager MIB so the system can be easily integrated into an operations/ network management network and managed remotely. The UCD SNMP v3 agent is provided along with a system MIB and an event manager MIB. This allows operators and network managers to access system and component statuses and to remotely configure and operate the system. Additional UCD functionality provides information on processes, disks, memory, and load average plus shell commands and error handling capabilities.

4.7 Diskless Operation

To support cost effective embedded system configurations that are diskless, HA-Linux provides a network boot facility that allows system slot and non-system slot processors to network boot using standard network boot mechanisms like “bootp” and “tftp”.

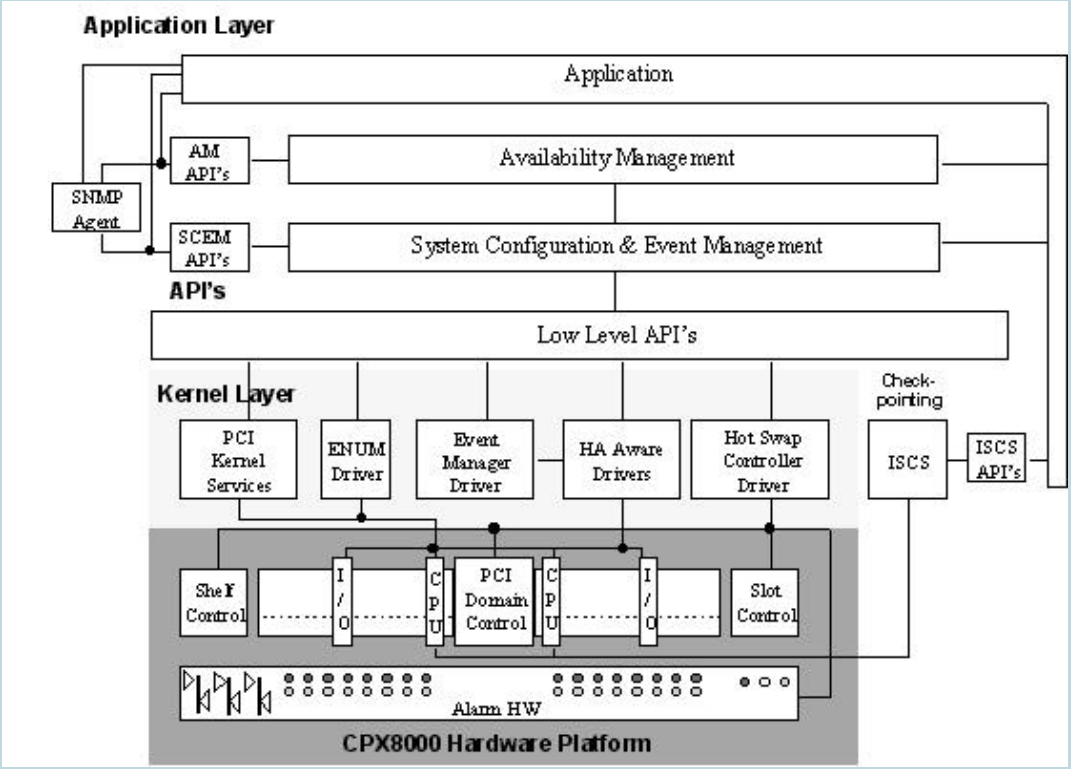

Figure 4: HA-Linux Software Architecture

4.8 System Configuration and Event Manager

The System Configuration and Event Manager (SCEM) is fundamental to the HA operation of the system. It provided integrated configuration, event, alarming, and availability management capabilities. The SCEM is composed of Linux kernel and application components

4.9 Linux Kernel Components

SCEM Linux kernel components include the following:

- PCI Kernel Services supports full dynamic insertion and deletion of PCI devices into the PCI configuration tree,
- The ENUM Hander enables the system to respond to the removal and insertion of components using predetermined rules and policies by signaling the insertion or removal of non-system slot components,
- The Hot Swap Controller Driver provides an interface to the Hot Swap controller and Bridge module in a multi-domain system. Supported functions include the power on and off of components, connect and disconnect of PCI devices, access to LED's and alarms, access to domain control signals, and event notification of state changes for system components. A low-level API interface is provided,
- HA-Aware Device Drivers are Linux drivers that provide comprehensive error checking, possible switchover to backup devices upon failure, latent fault checking to determine that devices are operational, and watchdog timers to check devices that may become inactive,
- The Event Manager Driver provides a convenient mechanism for a kernel module to send event messages to the SCEM, and
- The V-Term Driver allows the system processor to interrogate local components, run tests, and set configuration options on non-system slot processor cards. For example, this allows the system to determine the Ethernet addresses assigned to the network interfaces of each non-system slot PowerPC-based processor board.

4.10 Linux Application Components

SCEM Application Components include the following:

- The Availability Manager manages asynchronous events that can effect the availability of the system. In a distributed environment this will encompass a cross shelf approach that can be integrated with cluster management,
- The Shelf Event Manager is primarily concerned with asynchronous events like hardware and software component failures within a shelf, it responds to these events based on policies and distributes event messages to processes,
- The Event Manager API Library provides application program interface (API) services for applications and utilities that need to communicate with the event manager. API services allow easy access to event manager status and the attributes of devices and objects, sending and receiving event messages, and to log messages to the event manager's log file,
- The Configuration Manager allows the addition or removal of components and objects to and from the system and maintains a model of component dependencies and hierarchies in order for the system to make decisions in response to events,
- The Alarm Manager controls the telco visual and dry contact alarms and component “In Service” and “Out of Service” visual alarms using a configurable rule set,
- HA-Aware Device Driver Methods are used to define, undefine, configure, unconfigure, and change a device, and
- User Interface Utilities provide interfaces for system operators and network managers to display, change, & manage devices and objects in the system. Both local & remote interfaces are available.

4.11 Inter System Communication Services

Inter-System Communications Services (ISCS) provides a mechanism for data communication between two system domains in order to facilitate rapid application switch over to a standby system or system processor. The domains can exist within a shelf in the case of the CPX8000 or across multiple shelves.

The ISCS is critical in managing high availability applications in that it provides mechanisms for applications to checkpoint, heartbeat, and manage fail-over events, as well as providing:

- Application to application message and data communication,
- Remote program execution,
- File transfer, and
- Logging services.

API's exist for applications to access these services.

5.0 Conclusion

The need for high availability embedded solutions is becoming increasingly important to OEMs designing systems for the data storage, medical imaging, telecom and other industries that use a growing number of mission critical applications in their business.

HA-Linux development will focus on advancements in scalability and availability of CPX8000 systems through the addition of clustering, availability management, backplane messaging & network management. Although HA-Linux systems support ethernet & ATM, additional communication protocols will be supported with HA-Aware drivers.

As the CPX8000 family of systems grow and the underlying architecture and hardware is enhanced, HA-Linux will take advantage of these enhancements to provide increasingly higher availability while providing new features and functions that further enable the "application readiness" of the platform for the telecom industry.

6.0 References

[1]. An article about Linux and its significance, (1994) Mark Bolzern, President, WorkGroup Solutions, http://www.li.org/li/resources/papers/1994-linuxsig/lxsig.txt

2.0 Glossary

5NINES	- 99.999% availability
HA	- High availability
I/O	- Input/output
IDC	- International Data Corporation
IP	- Internet Protocol
ISCS	- Inter-system Communications Services
MIB	- Management Information Base
O/S	- Operating System
OA&M	- Operations, Administration and Maintenance
SCEM	- The System Configuration and Event Manager
SNMP	- Simple Network Management Protocol

About the Authors

As Canadian Engineering Manager with Motorola Computer Group (MCG), **Rob Pettigrew** is responsible for ensuring the requirements of Canadian customers are reflected through innovative product design. To a large degree, his role involves managing partnerships with OEMs such as Nortel, Newbridge, and Glenayre that reduce development time and result in breakthrough products for the telecommunications sector. He manages teams of engineers across Canada who ensure OEMs in the telecommunications sector succeed in the intensely competitive global telecom market

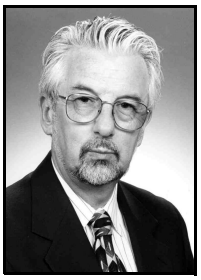

Noel Lesniak has been with Motorola for 20 years in various roles. He is currently business manager in the Motorola Computer Group's Telecom Business Unit with responsibility for High Availability Software Platforms.

Noel led the team to determine the Computer Group's Linux strategy and to drive its high availability Linux initiatives. Prior positions included Business Manger for Unix and Linux Telecom Platforms and Product Line Manager for the FX Series of fault tolerant, NEBS compliant systems.

IEEE Applauds Great Engineering Achievements

Nominated by 29 professional engineering societies, the top 20 list of Great Engineering Achievements was selected and ranked by a distinguished panel of the nation's top engineers, working in anonymity to ensure objectivity. The top 20 list was announced as part of National Engineers Week 2000 at the National Press Club by Neil Armstrong, an engineer and astronaut whose moon landing in 1969 was listed in #12, "Space Exploration".

List of the Top 20 Great Engineering Achievements of the 20th Century

1. Electrification
2. Automobile
3. Airplane
4. Safe and abundant water
5. Electronics
6. Radio and Television
7. Agricultural Mechanization
8. Computers
9. Telephone
10. Air Conditioning and Refrigeration
11. Interstate Highways
12. Space Exploration
13. Internet
14. Imaging Technologies
15. Household Appliances
16. Health Technologies
17. Petroleum and Gas Technologies
18. Laser and Fiber Optics
19. Nuclear Technologies
20. High Performance Materials

Pender M. McCarter, APR, Fellow PRSA
Director, Communications and Public Relations, p.mccarter@ieee.org