

Telecom/Datacom Solutions

OVERVIEW

Telecommunications and Data Communications were once separate markets, with their own supply chains, ecosystems, and customers. Today, communications is one category, consuming both telephony and data connectivity. The convergence of these two independent technologies has formed a category more powerful than the two separate ones, providing new opportunities for growth, and new ways for users to communicate.

Although once dominated by Power PC architecture, telecommunications equipment, and data communications equipment continues to migrate to X86, especially because of the steeper growth curve of X86 performance over Power PC, and the sharper decline of power usage of X86 compared with Power PC. Within a few years, it is estimated that the majority of this category will employ X86 architecture, not the original Power PC.

Along with X86 (and X64, a compatible, close relative) comes a huge ecosystem of suppliers that is much larger than the Power PC ecosystem. Phoenix Technologies has numerous BIOS and UEFI solutions available for ODMs and OEMs.

Phoenix Technologies has actually been enabling this market for several decades, partnering with Intel in its

High Availability and Communications initiatives. Phoenix customers in this category include Cisco, Alcatel, Nortel Networks, Lucent, Nokia, Apple and many others.

At the high end, Cisco uses Phoenix Technologies' solutions to manage its MDS-9000, its high-end disk switch fabric. The MDS-9000 incorporates Phoenix Embedded BIOS[®] firmware in each blade, in the chassis, POSTing in under one second, increasing availability. Each blade also includes Phoenix Technologies' High Availability firmware stack, allowing the system to support blade fail-over within 0.5 seconds. The availability offered by the MDS-9000 series raised the industry's bar considerably, enabling the product to be a big success in its market.

At the opposite end of the spectrum, Apple licensed Phoenix firmware in its AirPort wireless hub, because of its quick boot support and configurable footprint, increasing usability and saving cost in a high-volume product.

Phoenix Technologies has been driving the internet's infrastructure, from the core to the edge, since the mid-1990's. For your next telecommunications or data communications project, contact Phoenix to make it a more successful one.

APPLICATIONS

Metro Area Networks (MAN)

The internet involving packet switching over POTS phone lines ten years ago has been replaced by connected metropolitan area networks. At the Metropolitan Area Network (MAN) level, availability and performance are



primary concerns. Because the system may involve so many system components at the board-level, health monitoring and early warnings become critical for maintaining the overall health of the city network.

Phoenix Technologies provides firmware in MAN-class switches for DWDM, CWDM, and ROADM networks—today's contemporary standards, including Mahi Network's TRANSITION[™] optical switches.

Contact Phoenix to learn more about how we can help you with carrier-grade switching at the metro area level.

Carrier Grade Server

Carrier Grade Servers (CGSs) uniquely focus on supporting computing applications in the communications market. Originally only the domain of common carriers (which is where the term originated), central offices are now distributed and voice is merging with data, making it essential that CGSs be employed at locations emanating outward from the core of the communications infrastructure.

Today CGSs can be found in central and local communications switch offices, as well as corporate call centers and even medium to large-scale business communications networks.

Convergence of Voice and Data

The ongoing convergence of telecommunications and data communications is the single most important factor driving requirements in CGS devices. The different historical approaches to switching make convergence an evolutionary process, not a revolutionary, or disruptive process, or convergence won't happen.

Originally, telecommunications involved switching of calls, wherein the (actually large) connection time was dwarfed by the time spent during the voice call. Data communications, on the other hand, switches at the packet level, even permitting each packet in a stream of packets in a communication to go along different paths to reach their destination. Further, data switching supports the notion of broadcasting and multicasting, whereas telecommunications has not evolved the same capabilities on a large scale, beyond simple call conferencing.

Strategy migration to permit convergence requires backwards compatibility with older equipment, and open systems to support multiple suppliers. ATCA, MicroATCA, and AdvancedMC are all examples of open system definitions that permit multiple suppliers to drive convergence forward in a common direction.

The most widely-available open platform available today is the PC and X86 architecture. Although Power PC previously held a performance/watt lead, X86 is becoming more commonplace, especially with the introduction of silicon specifically designed to address the CGS market.

Along with PC and X86 architecture, BIOS is the firmware standard that makes all X86 platforms capable of running



all industry standard operating systems, from Windows to Linux to real-time operating systems like VxWorks. Embedded BIOS® with StrongFrame® Technology allows the flexibility to build "perfect fit" BIOSes that eliminate traditional IT-type features to keep the points of failure at a minimum, and keep a focus on performance.

Availability

The second most important factor in CGS devices is availability. While the internet can still survive when nodes go down, failure of a communications device such as a CGS means denial of service in a big way; therefore, availability (in terms of percentage of system operation versus total real time) is critical. Common measures of availability are specified in terms of nines:

NINES OF AVAILABILITY	PERCENT UPTIME	MAX ANNUAL DOWNTIME
Four	99.99%	52 minutes
Five	99.999%	5 minutes
Six	99.9999%	30 seconds
Seven	99.99999%	3 seconds

CGS devices incorporate redundancies at many levels, including power supplies, buses, fail-over blades, and other components. Yet, if a CGS crashes, simply the time it takes for the operating system to load and initialize can eclipse the maximum annual downtime. Consider that if the firmware takes more than 30 seconds to boot, the system will not be able to offer six nines of availability, but instead be limited to five.

Of course, simply noticing that a CGS has crashed can take time, and compromise the system's effective availability. Simple hardware monitoring can offer some relief, but smart monitoring that observes trends in bus errors, growing thermal problems, and software failures at the operating system can offer preventative health care so that the CGS never enters a critical state affecting the system's overall uptime.

Embedded BIOS® with StrongFrame® Technology offers the fastest POST in the industry (85 milliseconds, which can vary depending on hardware). But that's only the first step; getting the OS loaded needs optimized native-mode disk services available only from a carrier-grade BIOS like Embedded BIOS®; operating system files are loaded using the same fast, UDMA transfer speeds normally only supported by modern operating systems themselves. Getting through POST fast, and booting the operating system at near-theoretical limit speeds assures the fastest possible recovery from a failure, once it does occur.

With Firmware® Technology and its High Availability Monitor running in parallel with the OS as an overall system monitor, all of the operating system's processes, the hardware's buses, and measured properties (such as fan speeds, thermal monitors, and voltages) can be monitored independently of the operating system. The OEM can define low and high watermarks for normal operation, so that when those tolerances are exceeded by a signal source, the firmware can take action, such as alerting an application on the OS side, alerting another blade in the system, notifying a user or other system via email, or even become manageable via SNMP, Telnet, or HTTP clients on the internet.

Dense Blade

As the digital economy grows, more and more users need access to increasing amounts of computing. The density of computing equipment, starting with a simple PC on every desk in the 1980's, has increased by orders of magnitude, in order to scale the resources so that end users don't need to administer the equipment.



User administration can involve upgrades of failing hardware, failing operating systems, or replacement of whole systems. When this happens, a once unified set of hardware across an enterprise becomes more entropic, increasing the total cost of ownership.

Consolidating computing into blades increases density through form factor (usually, vertically-positioned blades within a 19" rack or other enclosure). It also standardizes on form factors and system design, allowing multiple sources who follow the same industry standards to make competitive solutions that help move the market forward.

Centralizing computing into a common array of blades also allows the blades to be centrally-administered (including hardware maintenance, software maintenance, and offering non-stop power and network feeds that would not easily be available in a distributed approach.)

Finally, with virtualization technologies, a single array of hardware blades can run virtual machines that allows migration of jobs from one blade to the next seamlessly, and without the end user actually being aware of the migration. Virtualization permits an administrator to manage the load evenly and move work away from failing blades, so they can be replaced without interrupting user activities.

Density of computing in blades gives rise to a performance/power dissipation tradeoff. The challenge is to minimize power consumption when no computing is performed, and then offer as much computing capacity as possible when required, while still fitting into the thermal operating envelope of the blade and enclosure. Performance and power management are highly-coupled and dynamic in blade systems.

The density of blades today will be viewed as low as the density of simple office PCs in the next few years, as demands for higher and higher density grow exponentially. New applications not possible without higher densities will appear, driving the trend further.

ODMs developing dense blade solutions need firmware that optimizes the dynamic range of computing within a fixed thermal operating envelope. Embedded BIOS® with StrongFrame® Technology provides the requisite firmware required to maintain safe operating levels while optimizing performance as needed.

To the extent that availability is also important in dense blades as it is in Carrier Grade Servers, Phoenix Technologies offers High Availability solutions including Firmware® Technology High Availability Monitor, and the fastest POST in the industry (as little as 85ms—less than a 10th of a second).

Dense blades are also starting to require additional support at the firmware level—provisioning. As more and more blades enter the data center, the need to install and configure operating systems and applications on each newly-installed blade becomes more and more critical. The Phoenix provisioning solution is Platform Update Facility, a Firmware® Technology component that can initialize and restore all the configurable bits in a blade, including CMOS settings, Flash, disk partitions and the partition table itself, and even individual files and ranges of files.

Platform Update Facility can also automatically be invoked by High Availability Monitor, when it is detected that the operating system is no longer operating or has become corrupted, offering automatic recovery from operating system failures.

Switch/Router

Switches and routers provide the high-speed packet switching required by data networks. Whereas the term router is commonly used to refer to a more generic platform that switches packets, the term switch is often used to distinguish the packet switching architecture as incorporating a switching fabric supported by ASIC hardware.

The switching problem exists at virtually every layer in the ISO model; however, it is most common to see switches at layers 2 and 3 on Ethernet networks. Layer 4-7 switches provide content switching to perform load balancing among disks, web servers, or other application-specific components. These switches are sometimes referred to as "disk fabric" OR "web fabric"—an extension of the nomenclature used when referring to suitably programmed ASICs as "switching fabric".

At the most fundamental level, switches are specialized computing engines with a fixed control policy; they do not run user applications. Because their internals can therefore be proprietary and not open, OEMs can customize the design for optimal switching performance and functionality. OEMs may elect to focus on security, or manageability, or customization for enterprise customers.

In order to support a wide range of switching solutions, OEMs may need to draw on multiple silicon technologies spanning multiple silicon vendors. StrongFrame® Technology provides the build framework that makes it possible to maintain multiple versions of each product SKU, all sharing a common set of versioned silicon and board support modules. This support for multiple ongoing product lifecycles saves costs



associated with managing multiple trees of firmware and other applications, while reducing risk by reusing components already written.

For proprietary switching solutions, OEMs can use Phoenix Embedded BIOS® with StrongFrame® Technology to reduce the complexity of the firmware to the minimal control paths necessary to initialize the hardware and load the operating system. Embedded BIOS® scales from a full-featured server, desktop, or UMPC BIOS down to the smallest bootloader, depending on over 1,000 OEM-tunable build options.

Gateway

Building on the data switch concept is the Gateway, which offers front-line access to a data network. Gateways are fundamentally switches that provide security (i.e., NAT) and dissimilar network connectivity (i.e., fiber to Ethernet, or Ethernet to Wireless). Commonly, gateways also implement services as added value, such as DHCP for SOHO networks. Any property not shared by two networks connected by a gateway is a potential issue for the gateway to manage, and a potential user requirement.

Gateways are not necessarily as performance-focused as might be expected from a switching perspective. Whereas a corporate backbone implemented by a gigabit layer 3 Ethernet switch is measured in terms of aggregate switching capacity, gateways are measured by their effectiveness at joining networks.



The attributes that make gateways desirable to users, then, are ease of use, configurability, and reliability. Wireless LAN gateways are an example of systems that need to work within a huge range of user-configured networks, and be configured by novice and expert users alike. In this device category, Phoenix helped Apple develop the Apple AirPort wireless hub, by providing its Embedded BIOS® firmware to boot the device in less than one second, and allow use of standard x86 building blocks while keeping BOM costs low.

Phoenix Technologies can bring its Embedded BIOS® with StrongFrame® Technology to bear on gateway design problems, minimizing your BOM by reducing the size of the firmware, along with providing value-add features and configurability as necessary, through its industry-leading Quick Boot capabilities.



Phoenix Technologies Ltd.

Headquarters:
915 Murphy Ranch Road
Milpitas, CA 95035 USA
408.570.1000 main
408.570.1001 fax

www.phoenix.com

Embedded Products:
915 118th Avenue SE, Suite 320
Bellevue, WA 98005 USA
425.576.8300 main
425.576.8334 fax
800.850.5755 toll free
embedded_sales@phoenix.com

Copyright© 2008 Phoenix Technologies Ltd. All rights reserved. Phoenix Technologies®, the Phoenix logo, Embedded BIOS®, Firmbase®, StrongFrame® and Firmware for Every Application™ are Trademarks or Registered Trademarks of Phoenix Technologies Ltd. The names of actual companies and products mentioned herein may be the trademarks of their respective owners. Information subject to change without notice.
DS252.00 11.08