Fred Rehhausser, Lori S. Grob, Thomas Cheng,
Product Manager and Market Development Managers,
Sun Microsystems.

High-Availability Communications Requires Embedded Hardware and Software

As the telecomm market wrestles with the convergence of voice, data and video onto a single, reasonably priced line, the need for high availability cannot be overlooked - as Fred Rehhausser, Lori S. Grob, and Thomas Cheng of Sun Microsystems explain.

INTRODUCTION

In the not too distant past, AT&T dominated the telecomm market, providing end-to-end solutions that worked together remarkably well. Innovation proceeded at a stately pace, with new switch designs being deployed every 20 years or so. Today however, the market moves to 'Internet time' and consequently change has become the constant in telecomm, and successful players need to upgrade continuously to remain competitive and the industry is seeing an overwhelming fragmentation in software and hardware solutions.

This has led to major interoperability issues - especially as the telecomm market wrestles with the convergence of voice, data and video onto a single line. This is because, tired of the confusion caused by multiple services, customers are demanding that Internet access, broadcast television, video and other services be provided over the same medium.

Moreover, customers want this all bundled into a package that is as reasonably priced and reliable as their telephone service. In addition, today's networks were built for voice calls that last an average of three to four minutes. In contrast, the typical Internet connection is 30 minutes long, creating a tremendous strain on existing networks. This in turn is creating the opportunity for new competitors, many with much cheaper transmission costs, to enter the market with converged wireline, wireless, switching or access networks.

HIGH AVAILABILITY CRUCIAL

As embedded telecomm developers rush to fill the transmission gap created by the Internet and the convergence of customer wishes, they must address the exacting demands of high availability. With the increasing pervasiveness of the Internet, consumers are demanding the same concept of high availability inherent in the telecomm world. In other words, people want non-stop operation similar to the traditional telephone service, where each time you pick up a phone, there is always a dial tone - any time, any place, and under any condition.

High availability means adopting many technologies and mechanisms pioneered by the telecomm industry in their proprietary platforms, and then offering them in competitive commercial hardware and software platforms. Among these technologies and mechanisms are system features like hot swap.

Today's derivatives of hot swap specific to bus architectures include provisions for I/O fail over and CPU fail over, each of which allows the system to continue running while an I/O or CPU board is being replaced. I/O fail over enables a provider to replace I/O boards by switching the failed one out, replacing it and then switching it back in. With CPU fail over, when a system controller fails, the system controller for the entire set of cards has to be switched out and immediately replaced by another system controller - either in the same rack or another: safeguarding the integrity of the system.

Another important feature for high availability is the capacity to dynamically reconfigure services, allowing providers to update applications anywhere within their network with minimal to no downtime.

Hot restart is a third important facet of high availability. It addresses those times when an error occurs that forces a software application to fail and be restarted. Today, when this happens, nine times out of 10 the

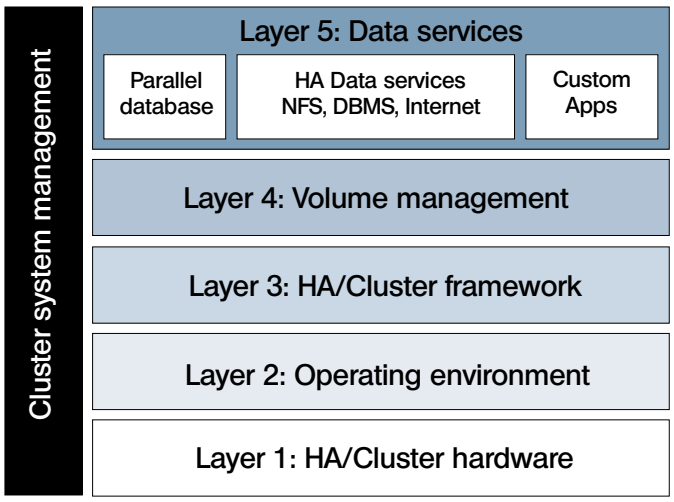

Figure 1. The majority of high availability architectures are built on a multi-layer cluster framework - one of the key elements of which is the volume manager: it, in conjunction with the rest of the framework, can be used by a host of tailored applications.

program must start from the beginning and any data accrued is lost. But in telecomm, forcing an application to completely start over can result in calls being dropped and the loss of services. With hot restart, a failed application will automatically reboot, starting near where the program left off.

IMPORTANT HARDWARE REQUIREMENTS

To keep up with users' expectations for new services and capabilities, embedded equipment and communications providers are scrambling to procure cost-effective, off-the-shelf hardware and software that supports their high-availability needs. Selecting the optimal combination of hardware and software components is a fast, effective way to create networking architectures that enable the convergence of voice and data technologies. By acquiring technology as opposed to creating it, providers can adapt quickly to emerging needs and market opportunities.

This allows them to focus their resources on value-added services. Two major issues that should be considered when using off-the-shelf hardware products are the supplier's understanding of the embedded telecomm market and their perceived longevity. That winnows the available field of suppliers down to a handful of vendors with the experience required to competently handle all the details of high availability.

The supplier of high availability technology is most likely to have a well defined road map for the future - to ensure a forward life cycle and upward compatibility as the developer moves into new designs.

Equally important to creating an embedded telecomm system using off-the-shelf solutions is firmly laying to rest any interoperability issues, and the best way to accomplish this is to build an infrastructure around the use of open system interfaces. Only by using solutions that adhere to industry-accepted interfaces can providers use third-party components to keep pace with user demands. An open architecture allows developers to deliver quickly new services to consumers at reasonable costs, in contrast to proprietary networks that are expensive to customise and therefore limited in the types of services and devices they can support.

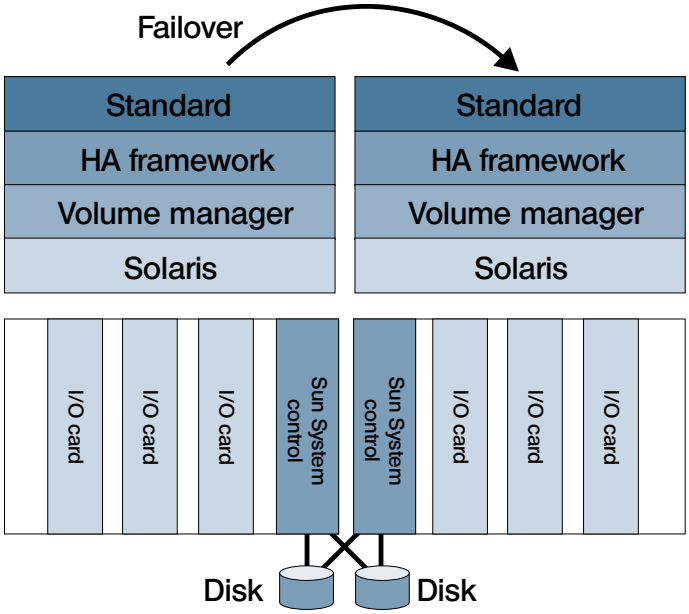

Figure 2. High availability is typically underpinned by dual redundancy. Here, call information, data and stored information is mirrored between the two disk drives.

ALWAYS THERE

The above high availability issues are some of the main reasons why compactPCI (cPCI) is rapidly overtaking the traditional VME form factor and fast becoming the dominant standard in telecomm hardware design - for everything from line cards to switches. First and foremost, cPCI is an open standard and has been since its inception. Just as importantly, cPCI offers an extremely compact, high-density form factor.

Like VME, cPCI boards can be rack mounted and provide extremely reliable connections to a common backplane with a very dense pin and socket connector. The PICMG, the open industry guiding group for the cPCI architecture, is actively defining software and hardware standards for hot swap, high availability and other telecomm requirements.

Compact PCI boards are available today that address the full range of telecomm needs, from providing basic system control and real-time support for building telecommunications infrastructures, to including VoIP gateways, communications servers and call centres. For example, to meet the most demanding telecomm requirements, single-slot, there are 64-bit-based cPCI boards with clock rates beyond 400 MHz and Mbyte memory caches. For less demanding applications, there are a range of cPCI boards designed to be used in more cost-sensitive embedded designs where performance and price must both be equally addressed.

CRITICAL SOFTWARE CONSIDERATIONS

As for high availability software, the challenge is to integrate the real-time applications typically found in telecomm systems with enterprise UNIX server applications and any legacy software applications. The goal is to expand service offerings on top of the traditional real-time operating systems (RTOS)-based applications to keep pace with customers' seemingly insatiable appetite for new capabilities.

The key criterion for both the RTOS and UNIX environments is their support of high-availability features like hot swap, an easy yet powerful development environment and sufficient real-time response. Another critical requirement is scalability: the operating system needs to be naturally extensible to allow it to be quickly configured and scaled to match the exact needs of the developer and service provider. To support high availability, the developers should be able to transparently configure (and reconfigure) their software with little or no downtime.

No single RTOS meets every development need - rather there are different levels of requirements depending on the telecomm application. At the line-

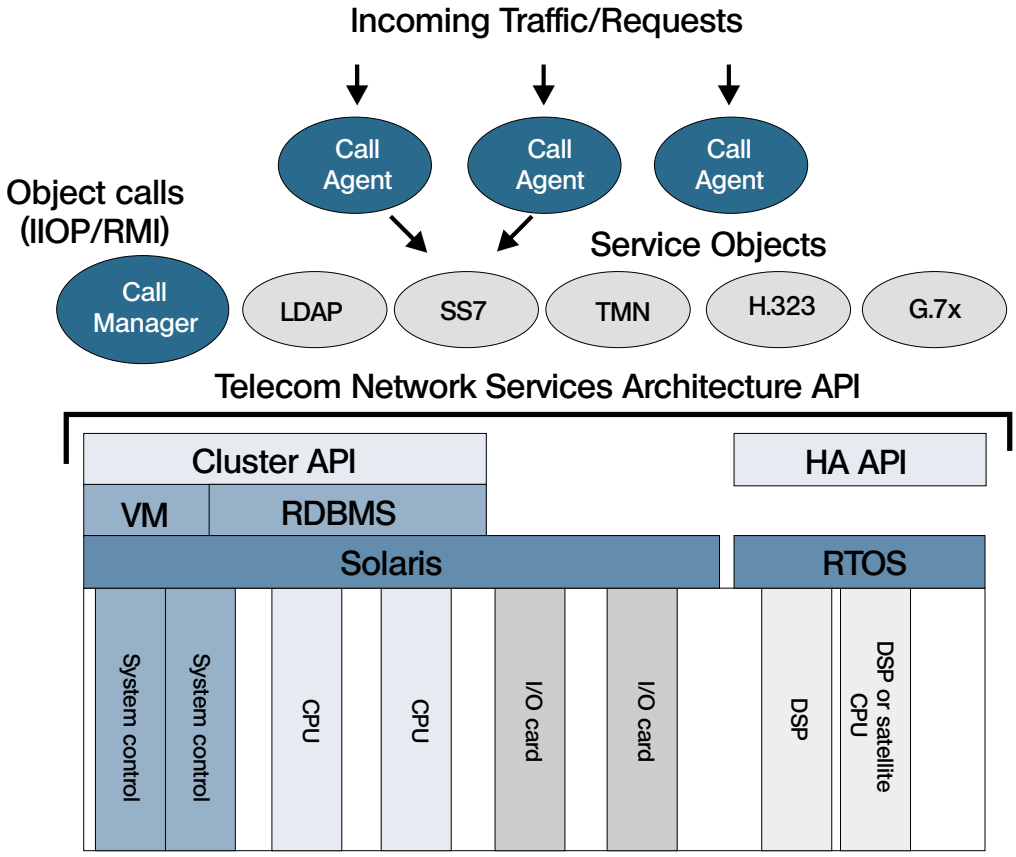

Figure 3. True clustering and flexible hardware. The Telecom Network Service Architecture API deals with all objects (for example the H.323 voice over IP protocol) and is an umbrella API to lower (hierarchically) API's and RTOS's down to core hardware.

card controller level the most important criterion for an RTOS is that it has a robust, flexible development environment. For more complex, mission-critical applications (like switches), a sophisticated, highly configurable RTOS is required.

INTEGRATION OF HARDWARE-SOFTWARE SOLUTIONS IDEAL

If you are buying in high availability, find a vendor that offers a comprehensive set of hardware and software solutions designed to meet all the current and future needs of the telecomm market. A single source leads to superior reliability, integration and serviceability because the hardware and software are integrated and thoroughly tested by the supplier, instead of being constructed piece by piece by the developer.

This enterprise approach to the telecomm challenge ensures that all the pieces will fit together smoothly, giving equipment manufacturers and service providers a range of products within a highly integrated, ready-to-use platform, which they can use to rapidly and cost-effectively build and deploy new products and value-added services.

It is important that this system integrator still offer all the benefits associated with open standards technology, without sacrificing innovation.

Only by assembling a full complement of development solutions can equipment and service providers hope to provide the high-availability priorities the telecomm market demands. The most viable way the converged network arena can match the level of traditional availability is to use off-the-shelf embedded hardware solutions based on open standards - combined with today's sophisticated RTOS's ■

Fred Rehhausser is CompactPCI Product Line Manager, Sun Microsystems; Lori S. Grob is RTOS Market Development Manager, Sun Microsystems; and Thomas Cheng is Market Development Manager, Telco OEM Solutions, Sun Microsystems Inc.