

Distributed architecture for VoIP telephony solutions

White paper

Introduction

As the initial reservations about VoIP voice quality, security and reliability fade away, this technology becomes a cornerstone of next generation, enterprise level voice-based applications and telco services. Residential VoIP is experiencing a fast market penetration and has definitely ceased to be a novel phenomenon, and corporations are looking for better ways to leverage VoIP in next generation designs.

However, this new technology trend brings a new set of challenges, with many tough choices to be made. First, the complexity of voice solutions is on the increase and contemporary applications are hardly reminiscent of those of just a few years ago. VoIP and SIP are not optional any more, but 'must haves', while comprehensive support for the PSTN and SS7 is still needed. Second, marketing specifications are becoming even more demanding, with the need to reduce costs. Low price per channel, total cost of ownership, CAPEX, OPEX – this terminology suggests the key theme of requirements documents.

In this situation, solution architects constantly seek out the best practises for the integration of complex technologies, which allow seamless functionality increases without the risk of 'over engineering'. On the other hand, getting 'something for nothing' is another essential system design theme, mainly in regards to reducing the solution costs without compromising the performance.

Traditional voice-based solution architecture

Ahead of introducing the concept that holds the answers to the contemporary design challenges, it is imperative to realise the structure of the current generation of voice communication solutions and recognise the core of the existing problem in terms of functional limitations.

In the VoIP-less telephony world, with expensive call minutes and the mess of disparate CAS and CCS type PSTN protocols, most of the voice-based solutions had a common architecture, simply described as a host server with telephony boards. Whether voicemail, interactive voice response (IVR), call centre, fax service or a gateway, what most had in common was that the solution was integrated in one physical unit. Utilising this conventional architecture, voice and data processing hardware i.e., telephony boards with analog and digital network access interfaces, a database, and high level logic in the form of a control application were all carefully integrated into a computer server to create an enterprise level solution or into a Compact PCI chassis for higher density, telco level solutions. The interconnection between the cards was done via the proprietary telephony bus H.100/H.110, while the high level application controlled the boards using the PCI bus that is present in the chassis or the computer server. An example of a solution utilising the common architecture is shown in figure 1.

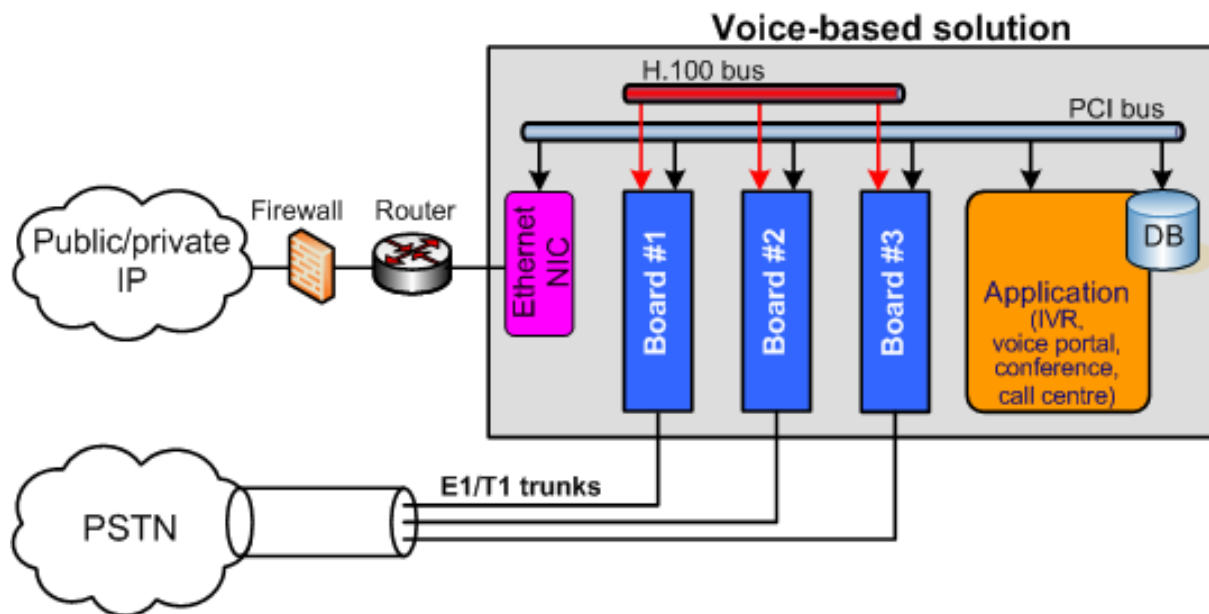


Figure 1 – Traditional voice-based solutions architecture

Although achieving the main objective of delivering a defined functionality set within a reasonable cost, employing the traditional architecture creates a significant headache to the developers when additional sophistication from such a solution is required. There are inherent problems with the solution scaling i.e., increasing the solution capacity with the minimal added cost. Implementation of redundancy, making the system fault tolerant, is technically cumbersome and almost doubles the CAPEX. Remote system management, hitless software upgrades, and hardware maintenance in the system without interrupting the service, add to the comprehensive list of system requirements that lead to increased architectural complexity and spirals the cost of solutions utilising the traditional architecture.

Distributed architecture

Understanding the deficiencies of the traditional architecture of voice-based solutions makes a convincing reason to improve.

The new approach is to implement a distributed architecture. Distributed architecture is considered to be the best practice for a number of reasons. It offers scalability, redundancy, remote management and high service continuity at a minimal added cost. It utilises the inherent strength of IP to allow remote devices to communicate distantly and operate as if they were all present locally. It is imperative to acknowledge that although the distributed architecture is a significant breakthrough in the way voice solutions are designed, it is merely an adoption of the practices used for at least three decades, in what is called 'distributed computing' and well known to everyone who ever used a UNIX-type operating system.

The fundamental approach in distributed architecture is the creation of discrete functions within a defined functionality scope. These functional elements exist in separate physical entities and communicate via IP. A solution can have each function in either single or multiple presentations (to eliminate a single point of failure). The minimal level of distribution is achieved by the separation of media processing elements from the call control (SIP service) and applications. Having the call control operating on the media processing cards takes the developers to square one and prohibits achieving the design goals.

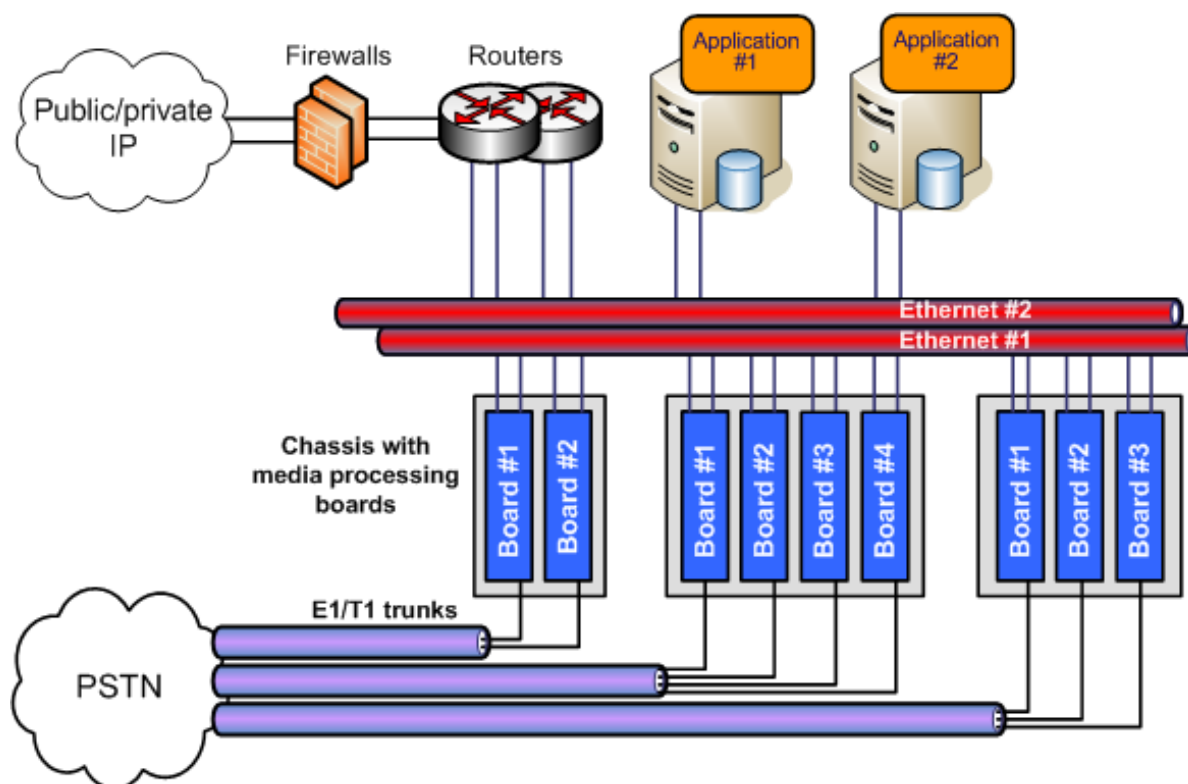


Figure 2 – Distributed voice-based solutions architecture

An example solution based on the distributed architecture is shown in figure 2. In this particular scenario, every function exists in multiple physical entities, including the application, media processing boards, firewalls, routers and the local area network (switches).

This example solution has three levels of redundancy for each element, also called protection. First, is the connectivity protection, which is achieved by having dual redundant Ethernet interfaces, each connected to a different Ethernet switch. The connectivity protection is not always achievable on the PSTN side, therefore only VoIP calls will benefit from automatic failover to the alternative connectivity link. Second, is board protection, which can be implemented utilising the venerable N+1 or N+n schemes. A solution can have alternative media processing boards that are designated to handle the media originally allocated to the main boards. Implementing an 'active-active' model is always desirable and simple 'random allocation' or 'round robin' load balancing techniques will help to minimise the impact of the board level failover. Finally, we have application level protection, which is achieved by having at least two application hosts sharing the database(s) and call control state information using handshake mechanisms. As software elements are usually the most prone to failure, application level protection is seen as vital.

Another example (see figure 3) shows a true, boundary-less distribution of a voice solution between remote geographical locations. An example of a solution benefiting from such an 'extreme' form of distribution may be a call centre for a bank, or any other large institution, with a major customer service centre offshore and branches in numerous geographical locations. Implementing a VoIP-based solution, which utilises the distributed architecture, allows not only significant cost reduction for the voice calls carried over the public Internet or a VPN instead of the PSTN, but also high service continuity, because of the ability to employ remote control applications or media processing resources in case of the failure of an element in one of the geographical locations.

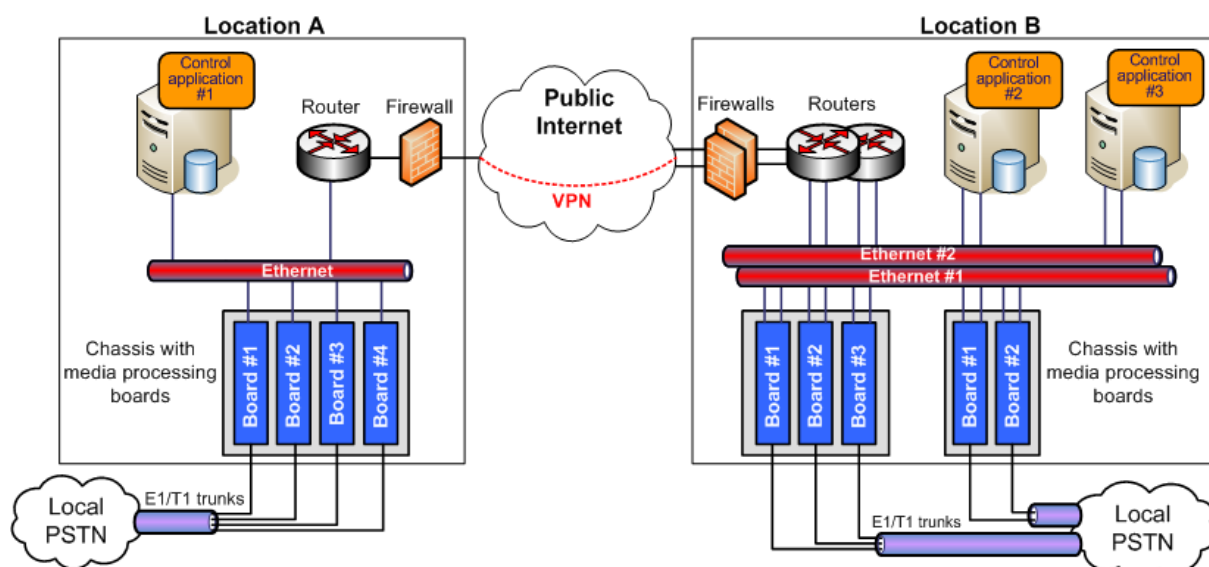


Figure 3 – Architecture of a voice-based solution distributed between remote geographical locations

Practicing distribution

The decision to build a solution with distributed architecture makes a fast ROI. And here is a list of the major reasons, which allow this technological advancement to make a real impact on a service provider's business.

Firstly, and probably contrary to the common belief, distributed systems are inherently cheaper to build, because of the ability to use a set of commodity, optimal performance and quality hosts. The need to deploy industrial-grade or special computer blades is eliminated and solution control applications can be hosted on a standard server machine connected to the enterprise LAN.

Secondly, these solutions are much easier, faster and therefore cheaper to develop and test, since the communications over the IP protocol are simple to trace and analyse with a wide range of freely available software-only tools.

Thirdly, the intrinsic scalability, yet again enabled by the use of IP communications rather than local PCI and H.100/H.110 buses, abolishes the need for solution re-engineering when the service capacity grows beyond that originally planned. Distributed solutions are almost linearly scalable and require extra capital expenditure mainly for additional media processing resources, while other elements of the system are completely re-used.

Additionally, the ability to distribute is a colossal benefit in terms of deployment flexibility. With plural elements communicating via managed LAN or even public IP, the same solution design can be instantly installed in disparate deployment scenarios, entailing only minimal software configuration.

Finally, and probably the major rationale, is the simplicity and cost efficiency in achieving telco grade reliability and service continuity. Eliminating the need to employ specialised proprietary or open standards hardware, developers of solutions based on the distributed architecture are saving tens of thousands of \$/£/€/¥ per system, whilst achieving the same level of element redundancy and automatic failover. The greatest benefit of the distributed architecture is that it allows the achievement of the highest level of service availability using commodity, lower cost computer servers and media processing hardware in PCI or PCI Express form factor. Nevertheless, for core network solutions, like IMS MRFP, the use of the 'best of breed' carrier-grade hardware platforms is still recommended. As such, CompactPCI and AdvancedTCA form factors allow higher integration density, which leads to reduced operating costs.

Creating distributed solutions with Aculab's products

Developers of VoIP and PSTN telephony solutions can benefit from Aculab's application development platforms, which offer a wide assortment of media processing functions, TDM connectivity cards and signalling protocols. Aculab offers hardware and software variants of media processing platforms providing the developers with the ability to create high density, feature-rich and cost-effective solutions.

In particular, the Prosody X product line of DSP-based media processing platforms is based around an IP-core and was specifically designed to support the creation of multimedia solutions with a distributed architecture. Because of the internal use of IP for media and signalling transport, Prosody X products are intrinsically scalable, support straightforward implementation of redundancy, load balancing and automatic failover, and also provide remote management capability utilising standard SNMP or Aculab's proprietary API.

In addition, the introduction of Prosody S v3.0, a new version of Aculab's software only media processing product, brings developers the benefits of support for distributed architecture without the need for proprietary telephony hardware. The product is a leading choice for the creation of multimedia solutions on a variety of 3rd party commodity computing hardware and can be used for very low density enterprise level applications up to very high density, carrier-class telecom services on AdvancedTCA or CompactPCI hardware platforms.

Conclusion

The ability to physically distribute VoIP telephony and IMS solutions amongst separate hosts is seen as a real step forward in voice solutions technology. The use of the distributed architecture allows the leveraging of the core strengths brought by IP and realising plural business benefits, including scalability, redundancy, remote management and high service continuity.

Equipment vendors and systems integrators, who design and build VoIP solutions for the enterprise market segment, are constantly looking to differentiate themselves from the competition. Finding it difficult to sustain barriers to entry in the maturing VoIP industry, these companies have a straightforward option to progress their business by adopting an innovative solution architecture.

Aculab's Prosody family of media processing platforms offers innovative organisations the ability to create high density, scalable and redundant solutions that deliver broad functionality, provide high ROI and in this way fulfil the contemporary market needs.



About Aculab

Aculab enables developers and systems integrators to produce a variety of high performance communications solutions. Aculab's portfolio offers an exceptional mix of capabilities that are easy to integrate and bring real value – reduced costs, increased customer satisfaction and competitive advantage. A complete range of open standards building block technologies for use within telco or enterprise environments, as well as essential support services, are offered. Products include media resources, digital network access, VoIP, fax, speech processing and conferencing. Support is available to help developers through each stage of their product's life cycle including pre-sales consultancy, technical support, training and marketing.

Copyright and other notices

© 2007 Aculab plc. All rights reserved.

Aculab, The Groomer, Prosody, "Connected With Aculab", Aculab Worldwide Digital Connectivity and Aculab Your Connection To The Future are registered trademarks of Aculab plc. Prosody S, TiNG and GroomerII are trademarks of Aculab plc. All other product or company references or registered and unregistered trademarks are the sole property of their respective owners.

For more information, visit Aculab's website at www.aculab.com

Aculab plc, Lakeside, Bramley Road, Mount Farm, Milton Keynes. Buckinghamshire. MK1 1PT. United Kingdom. Tel: +44 (0) 1908 273 800

The information in this white paper is provided by Aculab plc, and is for informational purposes only. Nothing in this publication forms any part of any contract.

The information contained herein is based on material, which Aculab, based on its best efforts, believes to be reliable, but no representation is made as to its completeness or accuracy. Aculab make no warranties, express or implied, in this document. E&OE