

Prosody S host media processing software

Prosody S is a host media processing (HMP) software package that gives developers and integrators a viable alternative to using DSP boards to integrate VoIP into their solutions. It brings granular scalability and cost-efficiencies to those familiar with traditional, enabling technology board-based designs and provides a software-only solution for IP-centric communications platforms.



In a purely IP context, Prosody S offers the same essential architecture as the board members of the product family and provides the same comprehensive mix of features as Prosody X boards for rich media processing in IP networks. It also enables the creation of flexible, resilient and scalable solutions, which can be distributed amongst different chassis and controlled by remote application servers.

Prosody S is licensed on a per channel basis, giving options from as low as a single channel up to very high densities, limited only by the performance of the host server. The product can be run on any host CPU (X86 or AMD) that is capable of running an operating system (Windows or Linux), which can, for example, be a PC, 1U or 2U industrial server, blade server, embedded system or laptop. This flexibility means solution providers can select the server that best suits their application, deployment environment, target market and price point.

Prosody S leverages Aculab's core expertise in combining complex technologies into a powerful and flexible proposition that assures developers a simple, clear migration path. Those looking to develop multimodal voice and data communications solutions will benefit from the feature rich architecture of the product line, which is already capable of handling multimedia call sessions.

Target applications

- Announcement server
- Conference server
- Wideband conference server
- Contact and call centre platforms
- IP-PBX
- Fax bureaux/fax server
- IP media server with transcoding capabilities
- IVR and voice portals
- Transcoding server
- Online gaming platforms
- Quality monitoring and test equipment
- Session border controller
- Unified communications platform
- Emergency services communications solutions
- Military embedded applications

Product features

- The widest variety of media processing resources for a complete range of communication solutions
- Thousands of media processing channels per host, depending upon host CPU performance
- Distributed architecture – provides system resilience; enables remote control via Ethernet
- The broadest range of codecs including:
 - G.711
 - G.723.1
 - G.726
 - G.728
 - G.729AB; G.729i
 - iLBC
 - GSM-FR, GSM-EFR, MS-GSM
 - AMR-NB
 - EVRC
 - SPEEX
- Wideband (HD Voice) codecs include:
 - G.722
 - G.722.1, licensed from Polycom®
 - AMR-WB (G.722.2)
 - Skype SILK
- Extensive fax support – including T.38 at up to V.34 speed and G.711 fax pass-through
- SIP, H.323, MRCP
- SIP authentication, SIP over TLS (SIPS), Secure RTP and symmetric signalling
- Unrestricted number of SIP Bridge channels for third party call control
- Windows and Linux versions are the same product and share the same API and development resources with Aculab's media processing hardware platform
- IPv6 support¹

Notes:

1. Planned release, contact your Account Manager for details

Product benefits

- Best fit from SME/SMB up to telco grade applications
- Provides 99.999% availability, system reliability and resilience – enabling service continuity
- Allows implementation of N+1 or 1+1 protection schemes with redundancy management
- Widest range of media processing functions – enables creation of highly functional solutions and reduces product variants needed
- Distributable amongst different chassis – cost-effective use of resources; resilience
- Ease of use through a single coherent and consistent API – faster time to market
- Feature rich architecture for IMS and mobility solutions
- A wide range of codecs enables connection to many endpoint device types
- Applications can be expanded or scaled to suit end-user channel count requirements
- Deploy software-only solutions on a choice of traditional server or virtual machine (VM) environments – use the platform best suited to the business need
- The low entry cost of a software-only option opens up new markets and increases margins in a competitive environment – lower cost of ownership and faster ROI
- Reliable, deployment proven technology means systems can be confidently specified for all application purposes
- Outstanding value for money is guaranteed through readily available software, high channel-count density capability, a range of support options and competitive pricing

Unrivalled density and feature set

Aculab's Prosody range offers a comprehensive selection of software modules or firmware algorithms that can run independently on the host CPU.

For Prosody S, the channel count is largely dependent upon the host CPU performance, some examples below:

- Benchmark tests for simple playback indicate a total significantly in excess of 2000 channels per server is possible
- A T.38 fax server built using a mid-range specification processor was able to achieve a channel density of 1600

Media processing functions

Assuming the appropriate hardware and software configuration, key media processing resources include record and playback with a range of compressions, DTMF tone handling, and fax transmission protocols. Each algorithm can be used separately or in combination to develop more sophisticated solutions, making Prosody S the clear choice for advanced speech processing development.

In addition to standard media processing functions, Prosody S also supports a broader set of higher-level technologies, including transcoding between various speech codecs, N-way (wideband and narrowband) conferencing, Group 3 and T.38 fax processing, G.711 fax pass-through, call progress analysis, live speaker detection and packet forking.

Being the enabling technology for a wide range of applications, from IMS, through convergence of IP with mobile or cable networks, to mobile messaging, Prosody S is capable of accommodating the comprehensive feature set required for creation of many sophisticated solutions.

Signalling protocols

Aculab's Prosody S customers also have access to a broad choice of signalling protocols used for call control in VoIP, PSTN, IMS and mobile networks. Aculab offers a complete implementation of SIP and H.323 under a cost free licence, and integration with 3rd party speech engines is enabled through the MRCP protocol. In addition, developers are given an option of using 3rd party signalling stacks to facilitate multi-vendor solutions integration.

Application programming interfaces

All functionality is accessed via a consistent set of APIs that are common to both Prosody S and Prosody X media processing boards. Consequently, an application developed using Aculab's API can be easily ported to any of the products in the Prosody portfolio.

Call Control is handled in two ways. The standard approach is to use the extended SIP API to control the Aculab SIP server directly (Figure 1). This exposes a high level of SIP detail to the application and provides support for specialised SIP functionality.

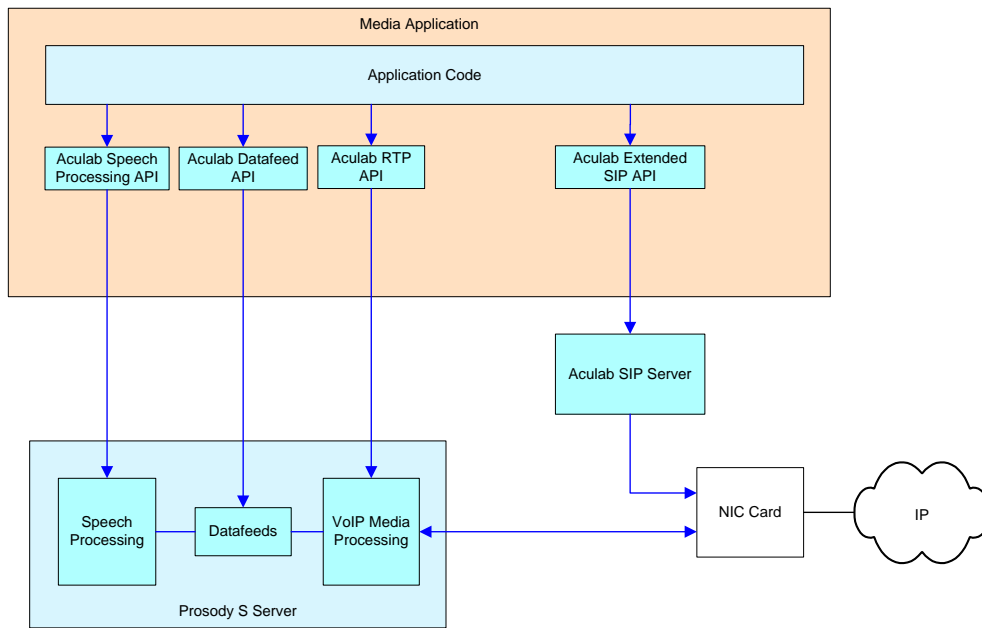


Figure 1 - Media application using the extended SIP API

For simple applications, the Prosody S media application can use the Aculab generic call control API to control the Aculab IP signalling servers and Prosody S server, see Figure 2. The media handler plugin (MHP) module controls both the IP signalling and the VoIP media processors (VMP) for each call.

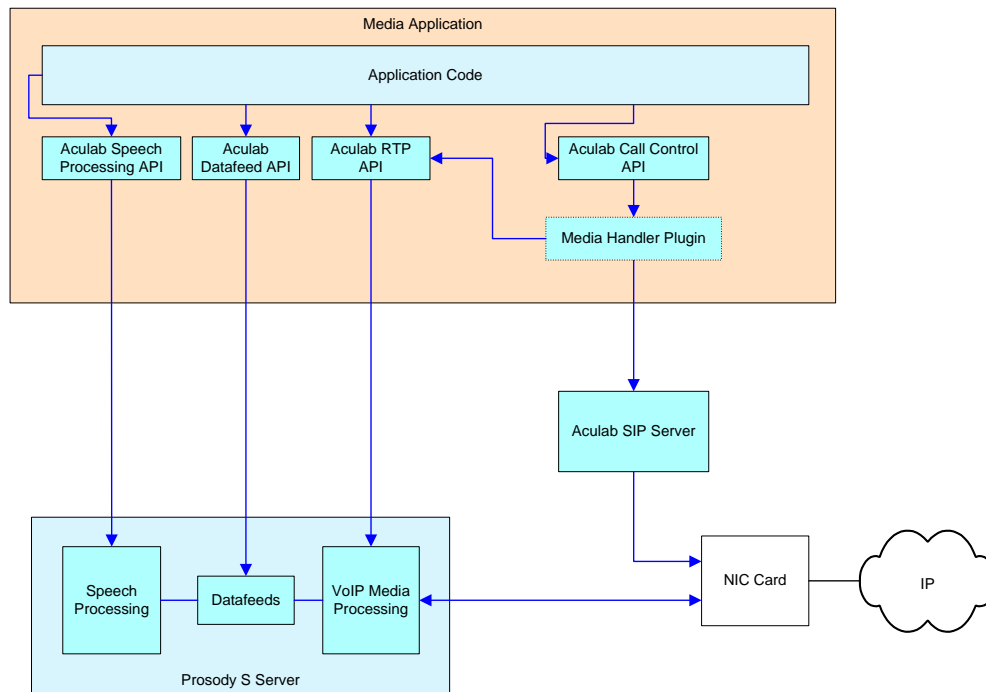


Figure 2 - Media application using the generic call control API

Scalability

Scalability is essential for every telecom service application, to support future growth with minimal added costs. Prosody S features linear scalability, therefore, allowing seamless growth from one to thousands of channels in a solution. In addition, a unique benefit of Aculab's media processing products is the capability to create solutions with distributed physical architecture, making the solutions scale cost-effectively beyond a single node.

Furthermore, Prosody S software has flexible licensing mechanisms, allowing users to purchase a license for the exact channel count they need for a system and cost-effectively scale to higher density when needed. A great advantage of Prosody S is that it doesn't preclude solution providers who only require a few channels, however it has the flexibility to scale from single channels up to many hundreds of full duplex, full function channels per system.

Telco grade functionality

Prosody S meets the particular functionality demands of ISPs and voice carriers by providing the reliability, resilience and scalability necessary for these customers. The basic requirement for the creation of resilient solutions is eliminating single points of failure by introducing redundant functional components. Prosody S supports sophisticated redundancy management to allow the creation of applications with 'five 9s' availability (99.999%), independent of the final solution scale or channel count.

Distributed architecture

All variants of the Prosody family allow the flexible allocation of media processing resources, which can be shared between several high level control applications, where one of them could be on the local host and others on remote machines, connected to the enterprise LAN.

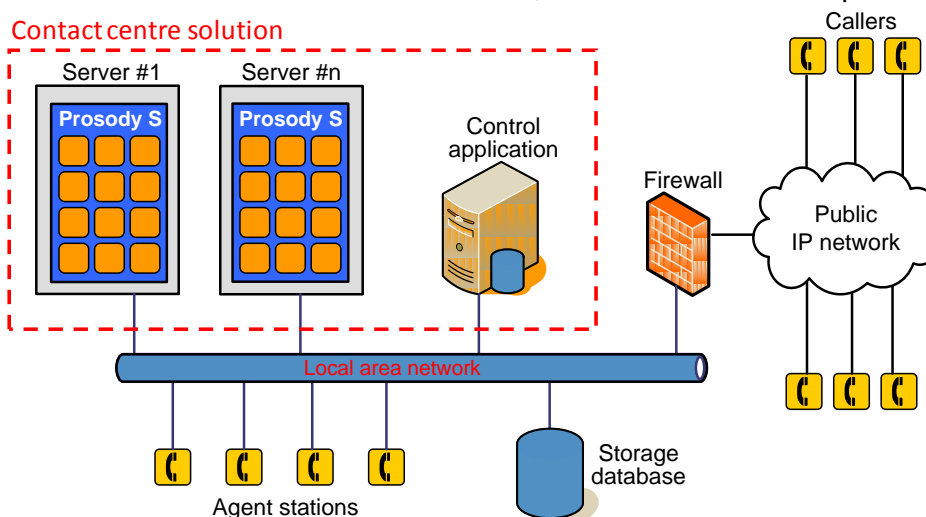


Figure 3 - Software, distributed amongst several servers, can act as one solution

In case of failure of the main high level application, application control can be switched to a remote standby alternative, providing service continuity by implementation of either N+1 or 1+1 protection schemes. Additionally, Prosody S can be distributed between several physical hosts and load balancing software may be implemented to minimise the impact of a failover event. Automatic failover mechanisms are also supported by the Prosody family, making the solution failure tolerant, from the low-level telephony server up to the application part.

Keeping high service continuity as the paramount objective, Prosody S utilises the inherent

strengths of the distributed IP architecture to provide, as applicable, effective link and application protection mechanisms.


Virtual Machine support

Prosody S has also been tested and successfully deployed in Virtual Machine (VM) environments, enabling the use, if required, of a single server for multiple instances of Prosody S, and the creation of hosted/cloud applications. Tests have been carried out with the following, and the test result details are available from your Aculab Account Manager, should you be interested in a VM deployment:

- Citrix XenServer 5.5.0
- VMWare ESXi 4.0.0
- Microsoft HyperV 2008 R2
- Xen 4.0.0

Licensing

Prosody S uses a very simple licensing scheme. It is licensed on a per channel basis including all features, making deployments simple and quick. A web-based licensing system is used to allow customer self-service for licence key creation once the initial licence credits have been purchased.



The screenshot shows the 'Manage Licences' page in the Aculab system. It includes a search bar, filters for start and end dates, and a table of licence records. The table columns are Licence ID, Company, Custom User Tag, Product, Feature, Quantity, Created date, and Duration. There are also quick links and filter options on the right side of the page.

Licence ID	Company	Custom User Tag	Product	Feature	Quantity	Created date	Duration
training-004-5	Aculab Plc	Upgraded old licence key: training-004-4/9999_-_Aculab_Plc/9b6b-9d28-548f/PROSODY3/01/3fc8-3921-8362-9625-cc9f-16d7-(9999) b79b-3285	PROSODY3	n/a	3 VMP points	09-May-2011 4:15 pm	n/a
training-004-3	Aculab Plc	9999 -	PROSODY3	n/a	2 VMP points	09-May-2011 4:14 pm	n/a
training-004-2	Aculab Plc	Upgraded old licence key: training-004-1/9999_-_Aculab_Plc/9b6b-9d28-507f/PROSODY3/02/ae76-65b1-6694-9678-f68b-ac48-(9999) b0d1-cc36	PROSODY3	n/a	6 VMP points	09-May-2011 4:11 pm	n/a

Figure 4 - screenshot of customer licences from licensing system

The licensing system provides a single interface to manage Prosody S and other Aculab products requiring licences, such as SIGTRAN and Dual Redundant SIP.

A full SIP stack is available to complement Prosody S. The SIP stack can be downloaded at any time, via our [software downloads](#) page. Prosody S customers using the SIP stack should adhere to a 1:1 relationship of media channels to signalling channels. Sufficient Prosody S channels should be licensed to support the maximum quantity of media or signalling channels required for use at any one time.

Technical summary

		Prosody S
VoIP and telephony functionality		
Audio channel capacity		CPU dependent, see table below
VoIP protocols	Signalling and control	SIP, SIPS ⁴ , SDP, H.323, MRCP, VoiceXML, CCXML – see website for further details
	Media	RTP, Secure RTP ⁴ , RTCP, RTCP XR ¹ ; with variable frame size
Voice compression		G.711 Annex I & II, G.723.1A, G.726, G.728, G.729A, G.729AB, G.729D, G.729E, G.729i, GSM-FR, GSM-EFR, MS-GSM, AMR-NB, EVRC, iLBC, Speex, TETRA, Skype SILK, G.722, G.722.1, licensed from Polycom [®] , G.722.2/AMR-WB Additional codecs supported – please contact Aculab for further information
Jitter buffer		Adaptive, with configurable upper limit
Lawful intercept support (LI/CALEA)		Yes, via RTP forking and packet replication
Additional functionality		User-configurable DSCP (ToS byte); DHCP; Transparent data over RTP (IETF RFC 4040)
Media processing functionality		
Conferencing		N-way matrix conferencing; narrowband and wideband modes; mixing loudest/active speakers; independent volume and gain control for each participant; personalised mix for each participant (e.g., for call centre coaching, network gaming, voice chat, etc.); active speaker detection; active speaker notification via CSRC; DTMF events suppression; HD Voice support
Predictive dialling / call progress analysis		Robust and accurate live speaker detection (e.g., differentiating between a human response and that from an answering machine); simultaneous signal categorisations on a per channel basis; DTMF, tone and call progress (ringing, busy/engaged, fax, SIT etc.) detection; speech energy detection; complete cause code functionality
Audio recording and playback		Recording and playback to local and remote hosts; multiple file formats; fast/slow pitch invariant replay
Audio gain control		Automatic (AGC) or programmable for each channel
Transcoding		Any-to-any voice codec ¹ ; full-duplex channels; rate matching; narrowband/wideband conversion (up/down sampling)
Fax handling		T.30 and T.38 fax termination at up to V.34 speeds, pass-through, relay and gateway; fax over G.711; automatic fax detection and notification
DTMF handling		DTMF detection and generation; inband; pass-through; DTMF relay and user indications (RFC 2833, RFC 4733); DTMF out-of-band (SIP INFO, RFC 2976)
Speech - ASR, TTS SVI		Integration with 3rd party voice engines via MRCP v1 and v2; Interoperability tested with Loquendo, Lumenvox, Nuance, Telisma and Verbio
Stream connection		CALEA / lawful intercept support for RTP streams; packet forking, switching and media replication (fan out)
Additional functionality		Tone generation; universal tone detection; call progress tone detection; pulse/rotary dial detection; grunt detection; voice activity detection (VAD); comfort noise generation (CNG); packet loss concealment (PLC); silence suppression; live speaker detection; voice morphing/pitch change
Operating systems supported		Operating system support for Windows and Linux:
		<ul style="list-style-type: none"> • Windows XP, Windows Server 2003 (32 Bit) • Windows Server 2008 R2 (64 Bit) • Windows 7 (32/64 Bit) • Linux(2.6 kernel and gcc 4.X compiler)
Other functionality		
Software licensing		Granular channels per host; no upper limit ² ; single licence covers all functionality Users of Prosody S are also able to take advantage of Aculab's SIP stack provided that a 1:1 relationship of media channels to signalling channels is adhered to

Channel count summary – Prosody S version 3

Typical figures tested on Intel Core 2 Extreme (X6800) running at 2.93GHz²

Windows XP, service pack 3

Features	Feature detail	Voice codecs ¹									Fax
		G.711	G.723.1A	G.726	G.728	G.729AB	AMR-NB	EVRC	iLBC	TETRA	T.38
Music on hold playback	10 different music replays	2200	130	650	100	400	180	260	170	100	-
Session recording	2-party conversation	350	55	155	25	115	70	90	70	40	-
Recording and playback	Full-duplex channels	630	115	310	55	230	135	175	135	85	-
	Full-duplex channels, with DTMF detection	580	115	300	55	220	130	170	130	80	-
Matrix conferencing ³	Full-duplex channels	730	120	310	55	250	140	180	150	90	-
Transcoding	To/from G.711 codec, full-duplex channels	N/A	115	300	55	210	145	170	130	85	-
Fax		-	-	-	-	-	-	-	-	-	750

SUSE Linux version 10.3

Features	Feature detail	Voice codecs ¹									Fax
		G.711	G.723.1A	G.726	G.728	G.729AB	AMR-NB	EVRC	iLBC	TETRA	T.38
Music on hold playback	10 different music replays	4000	130	750	115	650	205	270	185	155	-
Session recording	2-party conversation										-
Recording and playback	Full-duplex channels	850	120	310	65	310	155	195	145	130	-
	Full-duplex, with DTMF detection	750	120	300	65	310	150	190	145	125	-
Matrix conferencing ³	Full-duplex channels	800	120	320	65	360	150	195	155	130	-
Transcoding	To/from G.711 codec, full-duplex	-	120	340	65	340	165	200	155	130	-
Fax		-	-	-	-	-	-	-	-	-	1600

Notes:

- Aculab does not grant the right to practice the following standards: G.722.1, licensed from Polycom[®], G.722.2 (AMR-WB), G.726, AMR-NB, EVRC, iLBC, GSM-FR, GSM-EFR and MS-GSM. To seek the right to practice the standards please contact the appropriate intellectual property rights (IPR) holders. If you intend to use the iLBC codec, it is available under an open source ([3-clause BSD license](#)) license as a part of the open source [WebRTC](#) project. For IPR related to the G.722.2, AMR-NB and EVRC codecs, please contact the VoiceAge Corporation (licensing@voiceage.com). For IPR related to the G.723.1A and G.729AB codecs, please contact Sipro Lab Telecom (www.sipro.com) or the DSP Group (www.dspg.com). For IPR related to the ITU-T G.722.1 codec, licensed from Polycom[®], please contact Polycom (www.polycom.com); if you or your customer is a conference service provider, you must display Polycom's Licensed Trademark in your product.
- Maximum achievable channel count on a host depends on CPU performance and memory resources.
- Conferencing tests were carried out using 32-party conferences.
- Software functionality encompassing 'strong encryption' and subject to export restrictions, see [customer note](#).

For more information, please contact your Account Manager or view our website

APB0316 issue 2.3, April 2012