ApplianX IP Gateway
as a
VoIP enabled Programmable switch

Author: Faye McClenahan
Head of Strategic Marketing

Issue 1 May 2007

Table of contents

Table of figures

Overview

The purpose of this document is to propose a way to use an ApplianX IP Gateway as an IP-enabled programmable switch.

The TDM/IP gateway

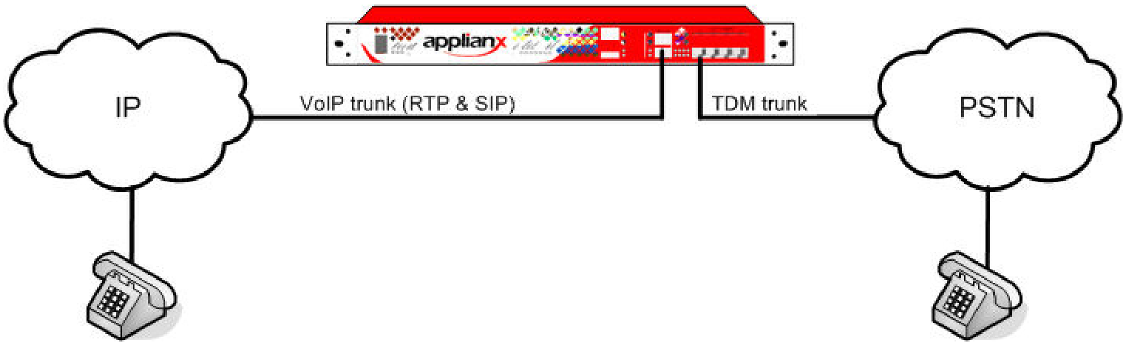

Figure 1: Applianx IP gateway

The diagram above shows the most direct architectural use of an IP gateway.

TDM calls connect on the TDM side of the gateway, and IP calls connect on the IP side of the gateway. Calls arriving on one side are forwarded to the other side after having been subject to various types of conversion:

- TDM G.711 media converted to RTP, using a codec negotiated and agreed with the called (or calling) IP terminal
- TDM call control is converted to IP call control (usually SIP), and vice-versa
- Often some manner of address translation is applied to maintain consistency within the called domain

Detailed gateway procedure

In a little more detail (without dissecting the SIP signalling), when a call arrives at the PSTN side of the gateway, a SIP Invite is generated and sent from the IP side. This Invite contains a list of voice codecs that the gateway is prepared to use, and has a called party URI that is often in the form: nnn@domain.com (having been derived from the incoming called party address of nnn). It also contains the address of the RTP port that the gateway is proposing to use for its media stream.

The response from the called SIP terminal will indicate its own RTP port, plus its choice of voice codec selected from those offered. The incoming TDM call is answered, and the bi-directional RTP media streams are established, each end sending its RTP to the port specified by the other end.

The architectural use of an IP gateway viewed in this manner is obvious – it is an autonomous gateway; a portal between an IP and a TDM domain, with a more or less complex, but fully self-contained mechanism for converting addresses.

However, with a little lateral thinking, it is possible to use the ApplianX IP gateway as a VoIP-enabled programmable switch, and this is described below. But first a little background on the nature and use of the programmable switch.

The programmable switch

The programmable TDM switch is essentially a ‘dumb’ switch – it is a TDM switch with little or no internal intelligence, and requires an external application to control it, and control is usually via a vendor-provided proprietary protocol or application programming interface (API).

A standard PBX is rather too smart to be called a ‘dumb’ switch in this context as it provides extensive automatic call handling. Many more capable PBXs provide computer telephony integration (CTI) links so that aspects of their operation may be controlled by external devices; but it is rarely the case that a PBX allows the level of total control afforded by a programmable switch.

Insofar as programmable switches are entirely defined by the way they are controlled, and that control is entirely in the hands of the application writer, they may be turned to a great number of uses. Consequently, programmable switch vendors have often played part in a ‘value network’ or ‘value chain’ where they sell the ‘iron’, but it is their resellers who are the application writers, who have the market specific knowledge and expertise, and produce applications to sell into specific market verticals.

Applications for programmable switches

Typical applications may be:

Prepaid and calling card
- Alternate routing
- Call centre automation
- Tandem (class 4) switching
- Adjunct and edge switches
- Enhanced services platforms
- CO (exchange) switches for competitive carriers and mobile operators
- Protocol conversion

Programmable switch architecture

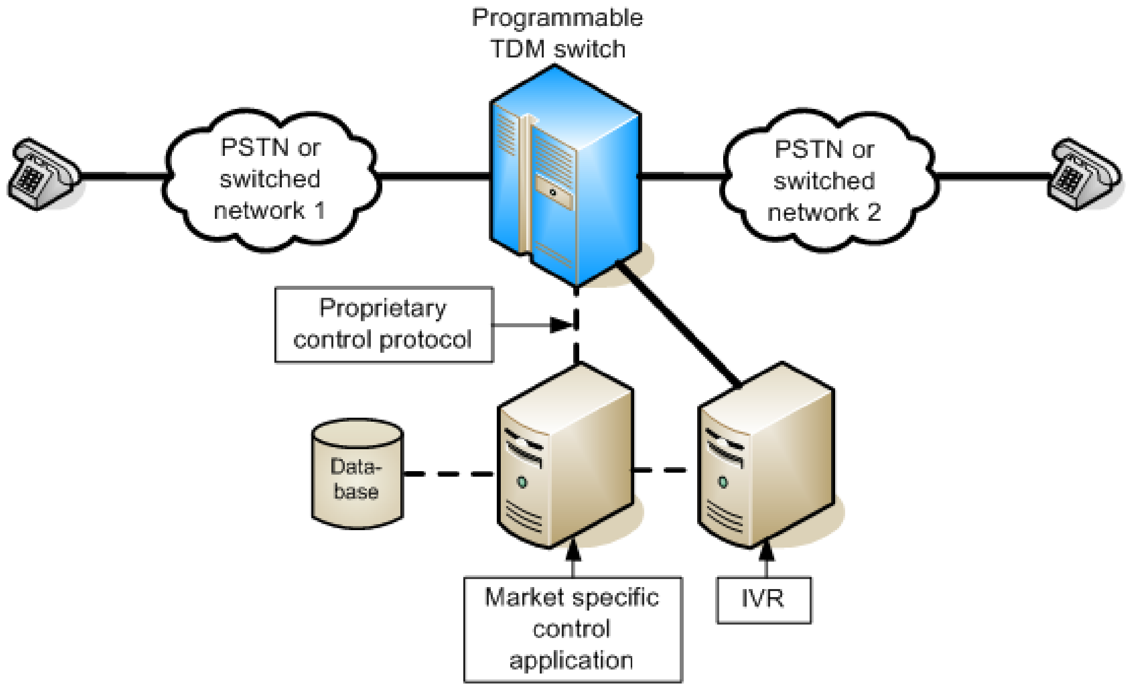

Figure 2: A system using a TDM programmable switch

The architecture of a system incorporating a programmable switch is fairly straightforward, as is illustrated as the diagram above.

A call arrives, the application is notified, then makes whatever call routing decisions it might, perhaps as a result of database lookups (and including the possibility of initially routing the call to a media server for enrolment or authentication), then the application forwards the call to its eventual destination, which may be over an outbound trunk group, or may be back into the network from whence it arrived.

Switch capacity (in terms of channel count) can vary from a few hundreds, to tens of thousands of ports, but the overall architecture remains largely the same. On a more detailed level though, there are always architectural concerns about the scalability of such a system – typically there needs to be a substantial non-blocking switching fabric, to ensure calls can be routed wherever they need to go. In a TDM switch, this switch fabric will typically be a TDM backbone, and due to the nature of such fabrics, the backbone will often need to be replicated to ensure resilience in event of a failure.

The IP gateway as a programmable switch

The diagram below shows how an IP gateway can be used as a programmable switch:

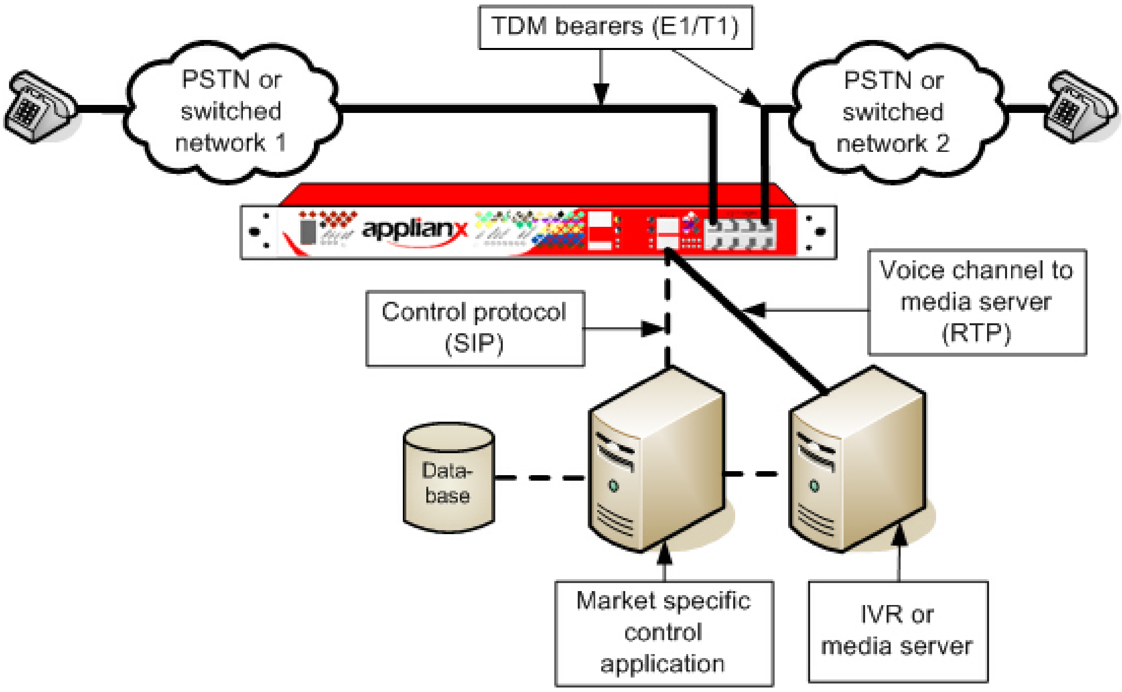

Figure 3: An Applianx IP Gateway as an IP-enabled programmable TDM switch

Essentially the overall arrangement is identical to figure 2, except this time, the control protocol is no longer proprietary – it is SIP!

Perhaps a few words of explanation are needed. So using the approach of the previous descriptions, a call arrives at the PSTN side of the gateway, and a SIP Invite is generated and sent from the IP side, but in this case, the Invite arrives at the control application (which would now more appropriately be called the SIP application server). The Invite contains a list of voice codecs that the gateway is prepared to use, and has a called party URI that is often in the form: nnn@domain.com (having been derived from the incoming called party address of nnn), along with the calling party address (if supplied). It also contains the address of the RTP port that the gateway is proposing to use for its media stream.

Once again, the control application makes whatever call routing decisions it might, but in forwarding the call, the control application sends another SIP Invite back to the IP gateway and uses an appropriately formed SIP URI to send the new called party address, offering the RTP port of the incoming leg of the call as an originating port, and G.711 as a voice codec in order to avoid quality impairments that might be caused by unnecessary transcoding.

The response by the gateway to the second Invite will indicate an RTP port address for that leg of the call, and when subsequently answered back and the RTP streams are bi-directionally connected, the RTP never actually leaves the gateway – it is simply cross-connected within the Ethernet switch in the gateway.

In undertaking this call forwarding process, the SIP Control Application is acting as a SIP 'back-to-back user agent', the procedures for which are described in RFC 3725, 'Best current practices for 3rd Party call control'. Note that the SIP application is entirely and separately in control of both legs of the call at all times, and can effect call transfers and alternate service stages at any point during the call, just by using standard SIP procedures.

As before, an IVR stage may be incorporated, either by initially routing the call (in the manner just described) to a TDM-based IVR, or to an IP media server (as illustrated in figure 3). In addition, either the incoming or outgoing leg of a call may be routed across IP in its entirety – ApplianX IP Gateway is thus fulfilling its destiny as an IP gateway!

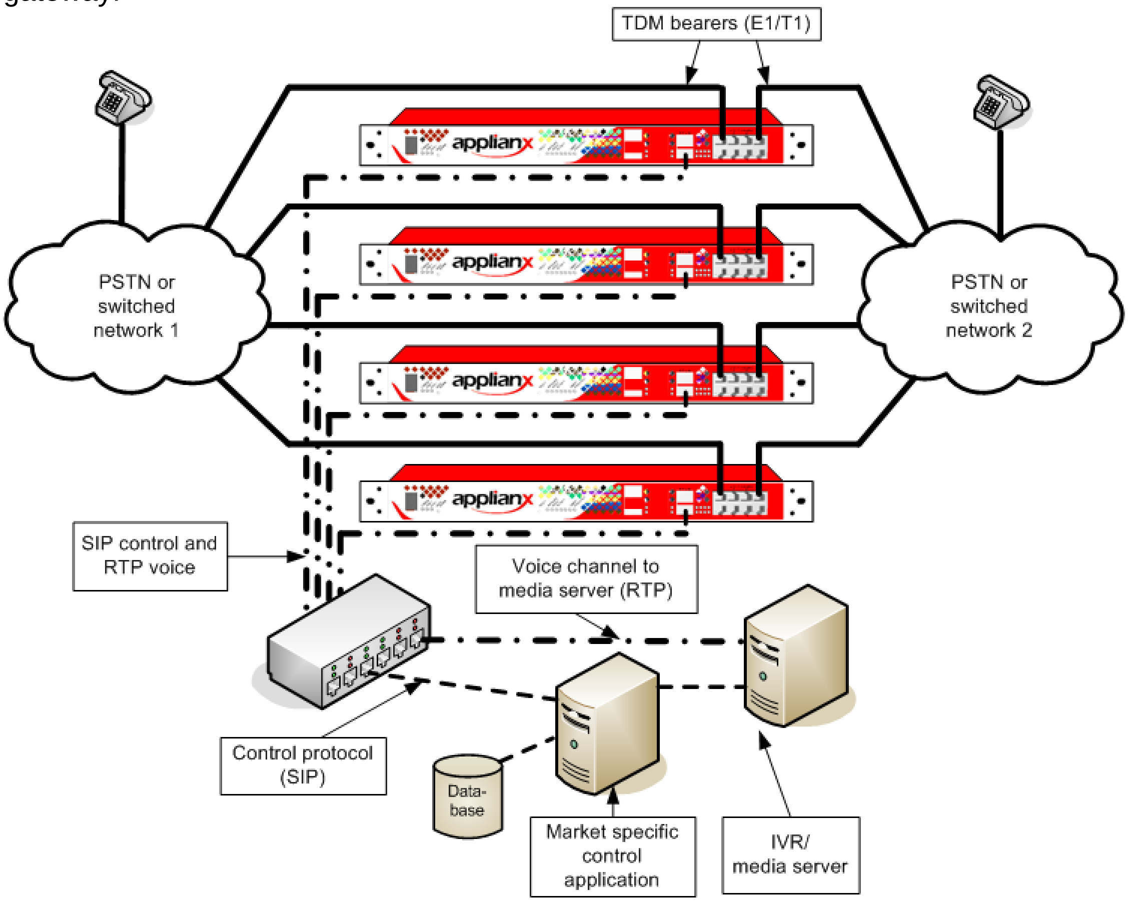

Figure 4: An Applianx IP Gateway as a scalable, IP-enabled programmable TDM switch

Scalability

The illustrations have used a single 1U Applianx IP Gateway as a controlling switch – that device has a maximum of 8 E1 or T1 trunks, so may establish a maximum of 240 separate TDM calls. Insofar as two separate ports are required for each call, this means it can actually switch a maximum of 120 simultaneous calls, however, such a system is readily scalable, as shown in figure 4. (As it is an important architectural element, the Ethernet switch is now shown).

When a particular call enters and leaves the system via trunks on the same IP gateway, the RTP is switched internally to the gateway as before. But now, if the point of ingress and egress are on different gateways, the RTP traffic is switched by the external IP switch – the programmable switch's switch fabric is IP, and it is self-routing!

From the standpoint of the SIP control application, it is aware of the address of the SIP port for each IP gateway, aware of the calls entering and present on each device, and can address each device as a separate entity. With a little ingenuity within the address translation table of each gateway, the control applications can even cause particular calls to leave the system via specific trunks.

The only real limit to system scalability is the capacity of the SIP control application to handle SIP messaging, and it is entirely possible (and indeed wise, from a resilience standpoint) to cluster that function.

Resilience

As there is always a concern about reliability when such systems start to grow in size, it is worth noting that ApplianX IP Gateway has dual IP ports for SIP signalling and RTP traffic, and these may be connected as a dual-redundant switch fabric, so if one Ethernet switch fails (or some interconnect wiring fails), the system will fall back to the alternate.

Resilience (apart from the robustness of each individual gateway) is asserted at a shelf level, as failure of a single 1U gateway will not affect the others, although the traffic capacity represented by its trunks will be lost.

SS7

A configuration similar to figure 4 may be constructed, but using a pair of Applianx SS7 Signalling Nodes in addition to the devices as shown. This offers advantages to embedding SS7 within a pair of gateways, as it prevents a (theoretical) failure of a gateway from simultaneously impacting SS7 signalling capacity (although the dual-redundant arrangement should ensure that the system as a whole keeps running). In addition, SS7 signalling often takes a separate path to the voice bearers, and separate SS7 signalling nodes readily provide that.

Conclusion

This note has described how ApplianX IP Gateway may be used in an extremely flexible manner, as an enabling element of a scalable, IP-enabled programmable TDM switch.