VoIP Reliability in Managed Service Deployments

Technical White Paper

Introduction

This White Paper introduces the Aspen 365 family of network appliances and explains how service providers offering a managed VoIP solution targeted at small businesses can utilize them to provide superior quality and reliability.

Different flavors of VoIP service providers, along with their differing deployment models, are considered. Some are Internet Service Provider (ISP) neutral, some offer bundled Internet access with a VoIP package, and others offer a bundled private network link of some sort. The Aspen 365 series addresses all these scenarios, using a comprehensive and flexible software feature set, which is remotely manageable and secure.

Aspects of VoIP Reliability

There are 3 major aspects to reliability in managing a VoIP service for business customers:

1. The overall availability and quality of voice communications.
2. The security of the customer's VoIP equipment, and equally important, the security of the VoIP provider's infrastructure. If either is compromised by attackers, the VoIP service ceases to be reliable.
3. Meeting expectations of business customers in ease of provisioning and rapid time to recover when problems are discovered i.e. serviceability.

In the remainder of this paper, these criteria and the various features of the Aspen 365 that address them are explored and discussed. The last section explains how the Aspen 365 can be deployed in a variety of scenarios.

1. Availability and Quality of VOIP

Complete control of both voice and data traffic is provided. A key goal is to give providers a complete spectrum of choices as to (a) if voice and data will be allowed to mix on the same Internet link (b) in the case of allowed mixing, how that mixing is to be controlled (c) in the case of fail over and backup, how backup is controlled.

A range of user policies from ultra-conservative (no mixing allowed) to ultra-liberal (aggressive mixing of voice and data with QoS controls) is supported.

A managed VoIP provider may opt for conservative methods of voice-data mixing, while an enterprise building its own solution may opt to be more aggressive with voice-data

mixing in order to optimize on the number and cost of Internet links. All these issues are articulated in greater detail in the next section.

1.1 Controls and Flexibility

Can a network link be tagged as “preferred” for voice? How should backup of the voice link be arranged? How many links should be dedicated to voice only, and how many to data? The Aspen 365 series allows a total of 8 network links to be configured for voice and data with various load-sharing policies decided by the user or service provider.

When a voice link fails, is voice traffic allowed to fail over to a link carrying data traffic? If yes, can a quality of service control be put in effect giving higher priority to Voice traffic?

Conversely, when a link reserved for only data traffic fails, can the data traffic be allowed to fail over to the voice link?

Fail over policies can be symmetrical for data and voice, or they can be one-way voice fail over only. The Aspen 365 supports either choice.

1.2 Graceful “Fail Back” After a Previously Failed Link Recovers

In cases where the service provider has configured via policy controls a “preferred” link for voice, the Aspen 365 will strictly enforce this policy. If the voice link fails, new voice calls will be diverted to one or more backup links that have been configured. The question that arises is – what happens when the “preferred” voice link becomes operational again? Since the existing media streams at that instant (if any) are successfully running on the backup link, attempting to move them immediately to the “preferred” voice link will result in the voice calls being broken.

The Aspen 365 in this situation implements the correct actions - which is to not move existing calls being currently transported on the backup link, wait for them to complete, and in the meanwhile if any new voice calls do get set up, place these new calls on the “preferred’ voice link. This behavior is sometimes referred to as “graceful failback”. Without “graceful failback” voice calls are disrupted twice, once when the preferred voice link fails, and a second time when the preferred voice link recovers.

1.3 Reliable Detection of ISP States

It is vital to avoid false positives. A false positive is when an ISP or network link carrying VoIP is falsely declared as “DOWN”. The Aspen 365 implements various patent pending ISP state detection algorithms that avoid the known pitfalls of less sophisticated equipment, which tend to rely on ICMP tests, or other simple mechanisms. Because the Aspen software monitors and uses live application traffic as a background test on each link, the risk of “false positive” syndrome is minimized if not altogether eliminated. The VoIP provider can also specify custom application level probes, based on UDP or TCP, to accurately determine that the VoIP infrastructure (Registration Servers, etc) is accessible via any specific path. Often, for security reasons, such servers wisely disable

ICMP. It also happens at times that some ISP along the Internet path may temporarily disable ICMP.

Avoiding false positives in VoIP networks is important because in the worst case, the link flips up and down repeatedly, resulting in broken voice, or in the best case, the voice traffic is needlessly placed on a backup link mixed with unpredictable data traffic.

In addition, the Aspen 365 has the ability to detect a third state, in addition to just an “ISP UP” or “ISP DOWN” state, and this is the “ISP UNSTABLE” state. These are links which repeatedly flip up and down. A link detected as unstable, after a threshold number of UP-DOWN transitions has been reached, will be kept in a down state for a longer length of time and subject to more stringent performance tests before being brought to the UP state.

Lastly, alternate ISP links can be chosen for reasons other than simple hard failures or instability. The Aspen 365 is also optionally capable of re-routing RTP media streams to superior alternate links, based on real-time measurement and statistics collection of actual RTP stream performance over each link. For instance, a link susceptible to a steady packet loss may pass tests for Link Up/Down but fail tests for adequate voice quality.

1.4 Quality of Service and Priority Queues

The Aspen 365 includes selective enabling, on each link, of a priority queue for RTP streams tagged as belonging to a specific VoIP provider. These streams are identified by tracking the state-transitions to RTP stemming from SIP exchanges with a specific provider’s infrastructure. It is important to note that this tagging and classification is independent of the DiffServ marking on the packet. In other words, even if the voice and data streams are identically marked as DiffServ Best Effort, a common default behavior, the Aspen 365 will correctly identify the RTP streams that require high priority service.

The RTP packets in the priority queue will always be serviced first; when it is empty the packets in the non-priority queue will then be serviced. This ensures that, for the duration of voice fail over when voice and data traffic mix, the voice quality will always be maximized. If anything, the data traffic may degrade slightly during this period, but that is a tradeoff that usually makes sense for most customers.

1.5 Scaling and Quality of Service

A single DSL Internet link can usually handle a small number of concurrent VoIP phone calls. In provisioning for larger customers, the ability to split the IP phone traffic over multiple DSL links can often be an advantage in cost and availability terms, as DSL is very inexpensive in comparison to T1. The Aspen N356 includes a variety of policy controls for load balancing the IP phone traffic. Static placement on individual links – such as by the IP address or phone number of the phone, as well as more dynamic methods (round robin or available bit rate load balancing) can be applied on up to 8 DSL or any other broadband (such as cable modem, fixed wireless) links. This feature may be of interest to providers that bundle Internet DSL (or other) links together with the managed VoIP service.

2. VoIP Security for Businesses and Providers

Providers need to ensure (1) that their business customers are protected adequately from Internet attack and fraud, and (2) that their own infrastructure is protected from attack and fraud using the customer premises as a backdoor for Internet or physical-premises intrusion attacks.

2.1 Stateful Firewall for VoIP

An optional VoIP firewall is included in the Aspen 365, to be used when a separate VoIP firewall does not exist (In cases where the customer runs their own VoIP firewall behind the Aspen 365, this feature can be disabled leaving enabled only the VoIP traffic management features).

In many deployments by VoIP providers, the customer may often have an existing firewall for the Data VLAN. In deploying new IP phones and a second private or Internet link bundled with the VoIP service, the built in firewall in the Aspen 365 enables the following security benefits:

- Automatic creation of a separate VLAN for Voice, which ensures that any security breaches that might occur inside the Data VLAN are kept isolated from the Voice VLAN
- A stateful firewall that tracks SIP exchanges and only opens up UDP/RTP ports based on a validated SIP state transition involving authorized SIP servers, and which closes these ports after the IP phone call has concluded. This stateful firewall ensures also that intrusion attempts and ports scans from outside are blocked and will fail.
- Security management with optional use of a built in DHCP server supporting multiple IP address pools, along with Network Address Translation (NAT) and Port Address Translation (PAT).
- Continuous validation of the hardware MAC address and associated IP address and PSTN number associated with an IP phone, useful in fraud detection and prevention.

2.2 Network Admission Control for IP Phones

By constant checking of the hardware MAC address on any IP phone, an optional knob is available to block both SIP and RTP traffic from any device whose hardware address is unauthorized. Such traffic will not be forwarded into the provider's IP infrastructure.

This feature applies to providers that are able to maintain lists of the MAC addresses associated with the IP phones they ship out to business customers, and maintain some kind of registration procedure that requires enforcement of a MAC address-Phone number (PSTN) association.

Admission control can be enabled using either (1) the phone's MAC address, or (2) the phone's manufacturer name.

Because blocking such traffic can have drastic effects, this feature will need to be explicitly enabled at configuration time.

2.3 Intruder Detection via SIP Auto Detect

The Aspen 365 auto-detects all SIP traffic, regardless of the SIP/UDP port being used. An option is provided to transmit alert messages to the provider's centralized monitoring system when such devices are detected.

Alerting is one step short of blocking, which is described in 2.2 above. It is also an option that needs to be explicitly enabled.

3. Serviceability and Provisioning

The Aspen 365 has been engineered for ease of serviceability by VoIP service providers, which includes ease of IP phone provisioning as well. The dual CPU architecture of the Aspen 365 includes a built in VoIP protocol analyzer tool, as well as real-time network measurement tools for jitter, packet loss and other metrics affecting call-quality. These can run concurrently without impacting the network performance of the Aspen 365 while under potentially heavy voice and data load.

By cutting down on time spent resolving problems, and minimizing the need for escalated intervention, providers can reduce overall costs while improving customer retention. For example: a problem as simple as a loose Ethernet cable attached to an IP phone can be discovered and rectified over the phone.

3.1 IP Phone Provisioning

An optional feature for bulk provisioning is provided, whereby a reserved DHCP pool of IP addresses and an automatically created VLAN are made available to support first-time install procedures. An option to create an automatic VLAN based on the name of the phone manufacturer simplifies the process (since the need for the user or the provider's installer to type in the hardware MAC address is eliminated)

3.2 Secure Remote Management

Secure remote access (via SSH, SCP) into the Aspen 365 management processor is provided; the packet forwarding processor is non-IP addressable and acts as a built-in firewall protecting the management processor from unauthorized access. Multiple levels of operations staff user can operate in different shells concurrently. An IP access control list can be used to limit access if desired to specific source computers or servers.

Up to 8 different Internet links can be used to access the appliance, to either troubleshoot an IP phone remotely or to "push" out configuration files in an automated manner. This ensures survivable remote access during various problem conditions that may afflict individual links.

Simple problems – such as a device not properly connected or powered off – can be easily detected by first level technical support remotely inspecting the Aspen 365 management console.

3.3 Real-Time VoIP Protocol Tracing

At times, a VoIP service request over the phone will require escalation to a protocol-knowledgeable technician. The non-technical customer reporting the problem with phone service will communicate the phone (PSTN) number having the problem. The Aspen 365 allows a technician to remotely start a real time trace of all IP packet activity on any such identified IP phone – including protocol traces showing SIP and RTP state transitions. Filters by IP address and PSTN numbers are included, to ease the effort of isolating a specific phone.

3.4 Pro-Active Real Time Measurements

The Aspen 365 packet forwarding processor is capable of continuously measuring the RTP payload, in order to proactively measure packet loss, round-trip time, payload receiver bit rate, jitter and other call-quality parameters. It does this as a background task on all VoIP calls in progress. This can be enabled as a proactive means of measuring call quality and detecting problems in advance of customers calling to complain. Alerts on various conditions can be generated to a centralized management console.

4. Aspen 365 – How it Fits in VoIP Deployment Models

Having a Layer 2 switch personality while performing advanced VoIP security and traffic management functions enables the Aspen 365 to be highly flexible in the way it fits into new business customer LANs. The next few sections serve to illustrate this flexibility in deployment and application.

4.1 Provisioning with Existing Data LAN Bypass

In this model, the VoIP provider bundles a private DSL or frame relay or T1 link with the managed voice service. The existing Internet link is used purely as a backup to the private link; the 365 ensures that voice stays on the private link, and will gracefully handle all failure and recovery events (see Section 1.2).

Model 1: Managed VoIP – Private Links Bundled, with Data Bypass

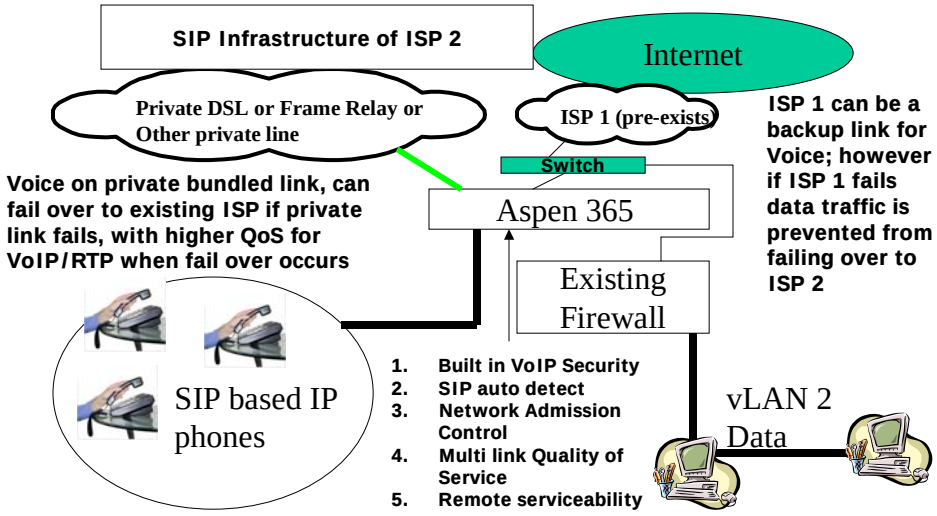

Fig 1

In addition, data traffic will never fail over to the private voice link by design, since (see Fig 1 above) the data firewall is physically by-passed.

4.2 Provisioning without Data LAN By Pass

The deployment model in Fig 1 above does have the limitation that when the voice link fails, voice traffic moves to the backup data link, where it has no quality-of-service guarantees. A deployment model that removes this limitation is shown below.

Model 2: Managed VoIP – Private Links Bundled – Enhanced QoS

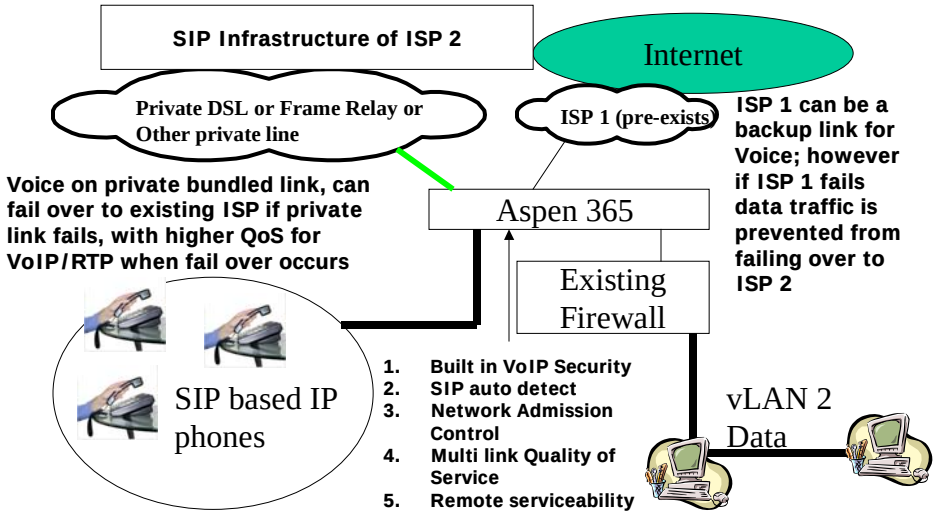

Fig 2

Here, both the Data and Voice VLANs are behind the Aspen 365 appliance. Whenever the voice link fails, data and voice will mingle on the pre-existing Internet link reserved for data originally. The Aspen 365 implements a priority queue for the RTP based voice traffic, thus shaping traffic so that voice gets the best link service, at the cost of some possibly slight degradation to the data traffic performance.

Additionally, the Aspen 365 has been engineered with 2 features that make the data firewall placement a risk-free process:

(1) By having a layer 2-switch personality, the data traffic from the existing firewall is simply bridged through to its natural router hardware address by the Aspen 365. The software will never place the data traffic on the voice link unless it is explicitly configured to do so. (An option exists where if the data link fails, data traffic recovers on the voice link, but this option is disabled by default).

(2) A second feature, which is useful at install time, is **"wire mode"** – this allows the box to be installed physically (while powered off or on) without disrupting the data network or existing firewall. During the install procedure, application data traffic continues to flow through the Aspen 365 as though it was a physical cable, bypassing the software completely.

4.3 Provisioning with Multiple ISP Links

This type of deployment model (see Fig 3 below) is suited to an ISP that is able to provide multiple DSL or cable-modem links, together with path diverse fixed wireless services in a bundle, together with a managed VoIP offering. This kind of a provider adds value by providing resilience and bandwidth for the data network traffic as well, at little risk of disruption to data or voice traffic.

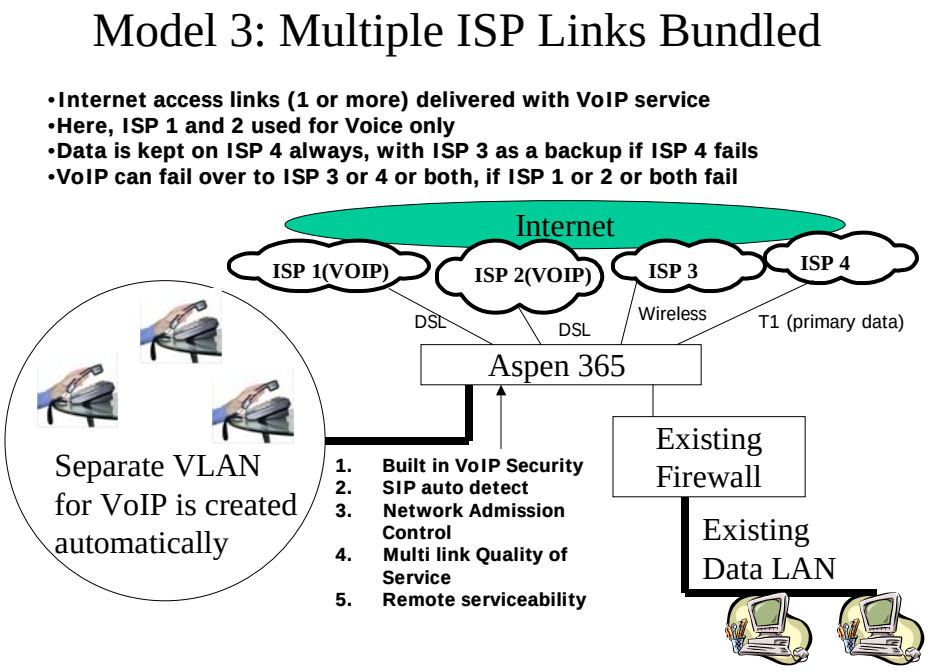

Fig 3

Here, because of the larger number of IP phones and expected call volume, 2 DSL links are bundled in for the VoIP traffic as “preferred” links. The Aspen 365 intelligently load balances the VoIP media streams on the 2 DSL links.

Here the customer is shown as having a primary Internet link (pre-existing) called ISP 4 for data traffic only. The VoIP provider – in addition to providing VoIP service, bundles in 2 DSL links (ISP 1 and ISP 2 in Fig 2) and a high speed fixed wireless link (ISP 3). The wireless link can serve as a backup for either the primary data link or the primary voice links.

The Aspen 365 has a flexible set of link policy management knobs. For example, voice traffic can be programmed to fail over to the wireless first, then the T1 if the wireless were to be down or even degraded due to higher packet loss and signal noise. Additionally, data traffic can be load balanced on both the wireless and T1 link based on application type and various thresholds (such as overflow and utilization).

4.4 VoIP Provider that is ISP Neutral

In this model, the managed VoIP provider provides no bundled links; the customer independently buys multiple ISP links to ensure maximum voice uptime.

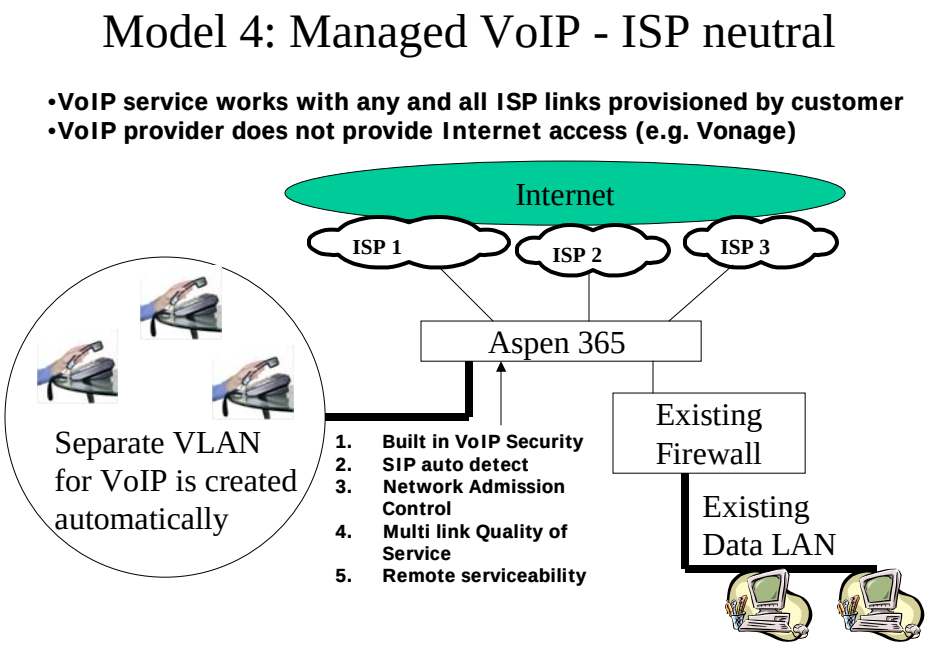

Fig 4

In this case, the Aspen 365 is either managed by the provider or by the business customer. In Fig 4, the customer purchases 3 Internet links.

Traffic is managed in such a way that the goals of multi-link quality of service, maximum uptime and optimum bandwidth cost are attained.

ISP 1 could be primary for voice, ISP 2 primary for data, and ISP 3 a backup for either one of the first 2, Alternate policies are also possible; the policy management framework is very flexible for both load balancing and fail over options.

Conclusion

We have introduced the Aspen 365, focused on VoIP traffic management and managed VoIP providers, or enterprises using a managed VoIP service. The Aspen 365 delivers reliability of VoIP traffic, including its quality and security. The security of the managed provider's infrastructure has built in protections that can be enabled if needed. Serviceability features and real time measurement of VoIP payload streams have been engineered for rapid problem troubleshooting, as well as pro-active service engineering prior to problems being reported.