

# What Went Wrong? Negative Results from VoIP Service Providers

Extended Abstract for IPTComm 2011 Industrial Talk Track

Mark R Lindsey  
ECG, Inc.  
505 N Toombs St, Valdosta, Georgia, 31601 USA  
lindsey@e-c-group.com

## Keywords

VoIP, SIP, Security, Interoperability, IP Telecommunications

Will the extensibility of SIP, the security environment of the Internet, and software development practices of a few vendors turn VoIP into the next great technology to be replaced? Has the spirit of open, interoperable, democratized telecommunication been lost in the haste to upgrade old telephone networks?

In a survey of 36 North American Service Provider networks that use SIP for voice services, we discovered a striking set of common problems that may signal a path away from the original spirit of VoIP's birth.

## 1. METHODOLOGY

The survey was conducted as part of Network Analysis, Security Audit, Network Design, and related projects within the SPs between 2005 and 2011. In each project, the SP described to us the services they offer, the VoIP servers they use, the IP routing equipment in use, and the number of customers or concurrent calls they support. We then interviewed personnel at those SPs about (a) types of devices to be used, (b) expected call paths and features, (c) DNS domains and servers, (d) traffic exchanged with other organizations, (e) codecs (e.g., G.711 and G.722) and media types (audio, video, and faxing), (f) the location of NAT boundaries, (g) the role of SBCs and firewalls.

Each project required three to ten days of work with the SP to collect sufficient detail. The 36 SPs represented a mixture of service providers in various regions of North America and the Caribbean. All but one provided some form of traditional telephone service or tandem call routing. See Table 1 for an overview of the survey set.

## 2. COMMON PROBLEMS

We identified three major classes of negative results that occur in numerous cases. These relate to the naming of devices and users, security, and interoperation.

### 2.1 Naming

First, the naming (i.e., addressing) of devices and users within networks becomes a problem because of distrust of DNS and abandonment of the SIP design philosophy. Because DNS already provides a robust naming mechanism, academics and standards designers expected DNS to be used to identify SIP users and services [5, 6]. Yet not one of the SPs uses DNS as anticipated in the standards, i.e., for flexible and reliable routing of calls between users, both inside and outside the SP's control.

31 of the SPs reported that their own DNS systems are not considered reliable enough for critical telecommunications purposes, such as business telephone service or 911 service. Approximately half of the SPs do use DNS solely for SIP server location for SIP registration of SIP endpoints, and fewer still used DNS only for call PSTN call routing using ENUM. In those cases static IP address-based bindings are also used as backup.

Part of this addressing problem is the rampant use of private IP addresses. Among neophyte network designers, RFC1918 [3] Private IP addresses are the only option considered for Critical equipment. SPs in this situation often identify a user by a SIP Address of Record having only local significance; e.g.,  
`sip:2293160013@10.0.0.5`.

Two of the SPs reviewed reported that integration projects between VoIP deployments owned or acquired by that same SP had failed because both systems used overlapping private IP address space.

Using private IP addresses on SIP servers creates artificial boundaries where Network-Address-Translation (NAT) [8] has to occur. Since the NAT implementations available on routers and firewalls in the market do not properly handle SIP and attendant media streams, and NAT was used in some form in every SP reviewed, a Session Border Controller (SBC) [2] was often present to unravel the knot.

<i>SP #</i>	<i>Com-mercial</i>	<i>Resi-dential</i>	<i>Geogra-phy</i>	<i>Ser-ver NAT</i>	<i>Fire-wall</i>	<i>End-point NAT</i>	<i>Good DNS</i>
1	C		SEUS		F	E	
2			Global	N	F		
3			Global	N		E	
4	C	R	SEUS	N	F	E	D
5	C		SWUS	N	F	E	D
6	C	R	MWUS		F	E	
7	C		MAUS	N	F	E	
8	C		MWUS	N	F	E	
9	C		USA	N	F	E	
10	C		NEUS	N	F	E	D
11		R	SEUS		F		
12	C		Canada	N	F	E	
13	C	R	SEUS		F		
14	C	R	Carribbean	N	F	E	D
15	C	R	MWUS	N	F		
16	C		NEUS	N	F	E	D
17	C		SEUS		F	E	
18	C		SEUS	N	F	E	
19	C		USA	N	F	E	
20	C		Canada	N	F	E	
21	C		Canada	N	F	E	
22	C		SEUS	N	F	E	
23	C	R	Canada	N	F	E	
24	C	R	MWUS	N	F	E	
25	C	R	NWUS		F	E	
26	C		USA	N	F	E	
27	C		MWUS		F	E	
28	C	R	MWUS		F	E	
29		R	Canada		F	E	
30	C		SEUS	N	F	E	
31	C	R	MWUS		F	E	
32	C	R	SWUS	N	F		
33	C		Global	N	F	E	
34	C		NEUS		F	E	
35	C		NEUS		F	E	
36	C		SEUS		F	E	

**Table 1: Selected characteristics of the SPs. Key:** *C*. Commercial SPs offer telephone service to enterprises. *R*. Residential SPs offer traditional home-phone landline telephone service. *SEUS*, *MAUS*, *NEUS*, *MWUS*, *NWUS*, *SWUS*, *USA* Within the USA, SPs may be in the Southeast, Mid-Atlantic, Northeast, Midwest, Northwest, or Southwest, or cover the entire country. *Global* service providers offer service in many countries, or via wireless networks in the oceans. *N*. Service providers with Server NAT have RFC 1918 [3] IP addresses on SIP Registrars, Media servers, and other devices. *F*. A key vendor of VoIP equipment at the SP requires a firewall to block most traffic from reaching the VoIP equipment. *E*. The SP supports SIP Phones or other endpoints that have private IP addresses. *D*. “Good DNS” SPs consider their DNS resolvers reliable enough to run Emergency 911 or other critical services.

## 2.2 Security Against Attack

We find that VoIP server security is lamentable, but vendors and SPs are mostly satisfied. (Here we focused on an SP’s equipment’s ability to perform properly under an onslaught of traffic, or when sent malformed SIP or RTP packets.) SPs typically would not use web server or mail server software that you cannot connect to the Internet, but they generally expect much less from VoIP equipment and software vendors. Purveyors of Carrier-VoIP equipment typically recommend their products to be protected from the Internet. These products are designed to run only within a limited, secured environment. They require the use of another security device, such as firewall or SBC, to limit the inbound traffic so that only friendly, trusted endpoints can send traffic.

Thus vendors require a “Walled Garden,” and expect that VoIP components can only communicate with other devices within the garden, or through a border device like the SBC. But a device like the SBC is not the best way to secure software running elsewhere. The SBC may function simply to restrict traffic to certain IP addresses. Or, the SBC may detect potential attackers based on the SIP response codes sent from the “trusted” side: too many SIP failure response codes, and the source making those failed attempts may be considered insecure.

But the SBCs themselves are becoming ever more complex; at least one model now provides Turing-complete programmability. As this complexity grows, why should the SBC be considered more trustworthy than other VoIP servers?

Further, the SBC can become a serious bottleneck: in addition to providing server capacity to provide some service, SPs have been forced to provide SBC capacity to validate every SIP signaling message and every media packet that passes through the Wall of the Garden. This can dramatically increase the cost and complexity of running an SP.

The Walled Garden also reduces the overall value of the deployments, because of the difficulty of adding new services and features is significant. For example, many SPs wish to add video calling to their platform, but cannot readily do so because of the barriers between their networks. These barriers were erected to protect and defend their critical core services.

SPs are rewarded for reliability and robustness, and vendors are rewarded for enhanced features and functionality. Thus SPs and their Vendors build networks that minimize the interactions with unknown devices.

## 2.3 Interoperability

Nearly all of the SPs reported some difficulty verifying interoperability among large numbers of VoIP devices. The SIP standards permit combinations of possible features, and sometimes multiple ways to provide the same feature. There is flexibility within the basic SIP standard, and more flexibility within its many subsequent enhancements and drafts. Because no two devices must support the same mechanisms in this growing library of options, certain SIP devices may not be useable within an SP simply because of arbitrary choices made for other devices earlier in a SP’s lifetime. For example, when providing a Caller Privacy feature, an SP

may have chosen equipment A that only supports the SIP *From* header, while a newer device B may only work properly with *P-Preferred-Identity*.

Because every implementation chooses only a subset of the standards, and only certain parts of each standard, the task of proving interoperability and functionality of a specific feature can be monumental.

SIP users are suffering from the problems described by Marshall Rose [4, Section 4.5] and by Christopher Alexander [1]: A change in any of the endpoints in a VoIP network can affect the performance of all of the others. To survive despite these drawbacks, Industry bodies such as SIP Forum are developing subsets of SIP [7] to simplify it by – in effect – making mandatory many optional behaviors and prohibiting others. If a closed-set approach is required for interoperation, then SIP's goal of dynamic endpoint negotiation has failed.

Because they are paid for features and not openness, SPs make the rational choice to prove interoperability only among a small set of endpoints, and exclude other interactions among SIP devices. Service Providers need to be convinced that better interoperability through endpoint negotiation will improve their networks and reduce their operational expenses by obviating the need for so much active protocol manipulation in the network.

### 3. CONCLUSION

Use of VoIP technology has grown rapidly largely because of its rapid adoption by Service Providers. Vendors have sprung up to satisfy that demand. But the goals and objectives of Service Providers and their Vendors do not appear aligned with the goals of open IP Telecommunications.

Service Providers are paid to offer useful features reliably. They continue a model of network design over a century in the making. In that model, devices are connected with point-to-point links, and the pairwise interoperation between connected devices is all that matters. Therefore, interoperability is not judged based on compliance to a protocol, but rather performance under tests for a pair of connected devices. Service providers are little motivated to ensure endpoints can negotiate features or functionality directly with without the active assistance of the Service Provider core; after all, the network core has long been a place where traffic is actively “translated” from one system to work in another.

Changes in traditional telephony networks are made slowly and carefully; so the indirect naming through DNS is not an advantage. SS7 point codes are numeric and they work well so IP addresses will do just fine as well. And Voice traffic does not enter the network from unpredictable sources, but rather through a predetermined set of trunk groups that provide passage into the Walled Garden.

VoIP Equipment Vendors are paid primarily to enable the Service Providers to succeed. So they make products to meet the traditional network practices of those service providers. Because Service Providers only expect to build Walled Gardens, Vendors accept that lower standard, and ship equipment only safe to use in a Walled Garden. VoIP Equipment Vendors are not strongly motivated to sell equipment to end users when that cuts out their primary customer, i.e., Voice Service Providers; therefore, they their products do not easily enable person-to-person communication across open IP networks.

End users of Voice services pay for *reliable* phone calls with useful features that behave the same from one year to the next. Therefore they gladly buy services that exactly and reliably replicate the networks and services of Telephony's first century.

Standards Developers have made protocols that grant flexibility to implementors who only implement tiny portions, but declare compliance. Developers pick and choose parts of the standards they wish to implement to meet their most immediate customer requirements. To the extent the VoIP standards allow for such laxity, standards developers must improve them. Stricter standards with provable compliance criteria will lead to greater freedom and predictability for users of the components that implement those standards.

These negative results highlight the way the VoIP could be sidelined to meet the goals of a specific segment of users. If this pattern continues, VoIP as used will eventually be narrowed to a limited set of features, and it will smell a lot like Q.931 or ISUP. These results identify rich areas of research for improving current practice, so that Service Providers and End Users can benefit from the promise of IP Telecommunications.

### 4. REFERENCES

- [1] Alexander C., *Notes on the Synthesis of Design*. Harvard University Press, 1964.
- [2] Hautakorpi, J. et al. *Requirements from Session Initiation Protocol (SIP) Session Border Control (SBC) Deployments*. IETF RFC 5853, 2010.
- [3] Rekhter, Y., et al. *Address Allocation for Private Internets* IETF RFC 1918, 1996.
- [4] Rose, M., *On the Design of Application Protocols*. IETF RFC 3117, 2001.
- [5] Rosenberg, J. et al. *SIP: Session Initiation Protocol*. IETF RFC 3261, 2002.
- [6] Rosenberg, J. et al. *Session Initiation Protocol (SIP): Locating SIP Servers*. IETF RFC 3263, 2002.
- [7] SIPForum. <http://www.sipforum.org/sipconnect>.
- [8] Srisuresh et al., *IP Network Address Translator (NAT) Terminology and Considerations*. IETF RFC 2663, 1999.