Mark and Sweep Robocall Detection and Prevention

By Dan Weber, Alan Basinger, Dean Willis, and David Schwartz

Background of the Robocall Problem

Hi, This is Ann from Account Services. There's no problem with your current account, but …

Ann has recently been making calls from 214-504-1687. One of our team members receives over 10 misdials per day on a number differing by one row on the keypad -- mostly from angry people demanding not to be called back. No "do not call" list can solve this, because it's driven by human error in response to illegal robo-calls. Robo-calls are clearly a big problem.

The economics of spam drive robo-calls. What's more lucrative than getting something for nothing? There are number collectors, number validators, list-sellers, robo-call sourcers, and advertisers, all making a profit. Some are legitimate but misguided businesses, and others are fraudsters and thieves. None have a vested interest in complying with "do not call" lists.

The vast majority of robo-calls source from foreign VoIP devices, traverse the Internet, pass through middlemen, and deliver to our phone-numbered devices using PSTN termination services. These services are provided by large Internet telephony Service Providers, most of whom are unaware of the problem and have no vested interest to do anything about it.

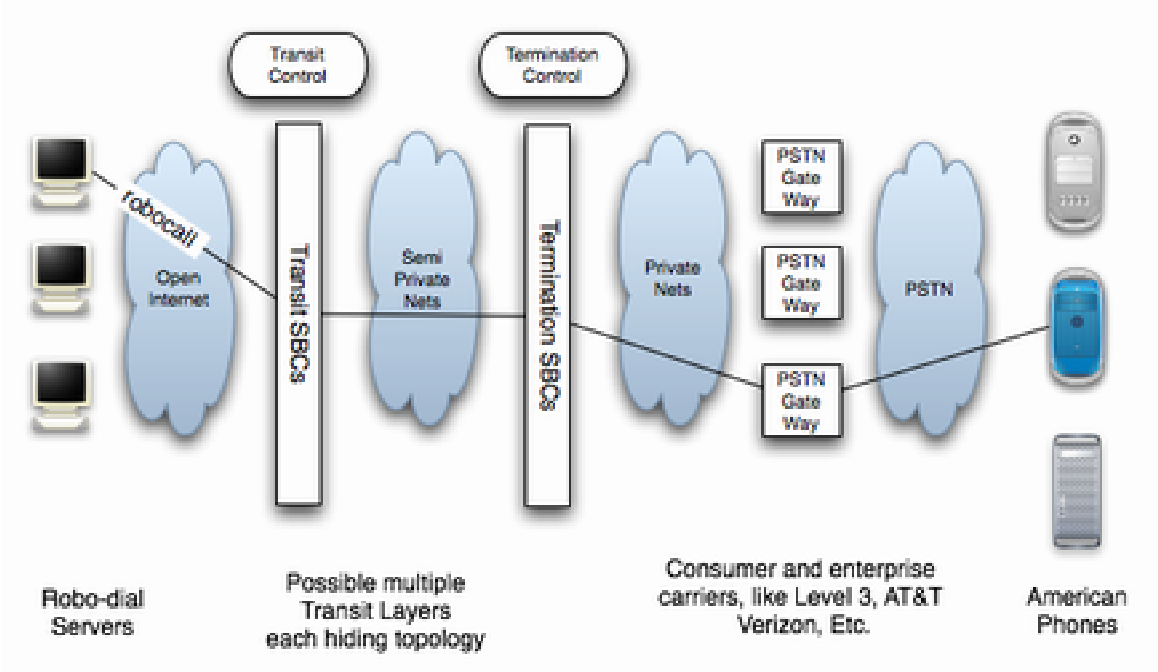

Figure 1: Topology of Robocall Delivery

Since most of the originations are foreign, there's little oversight available from US regulators at this level. However, transit (middlemen) operators are potentially regulatable, and the termination providers are typically large US companies like Level 3 that are clearly within the US regulatory scope. But these transit and termination providers lack the tools to really do anything, and have little incentive to invest in such tools given the economics -- in fact, they probably make money off of the robocalls, due to things like Feature Group D settlements. Consequently, costs charged back to the originators are very low (<1 cent/min), making large call volumes cost-effective and **very, very annoying**.

In short, advertisers seek to continue robodialing for one simple reason. It provides a very cost effective solution to sell products and services, many of which are fraudulent, to the world's largest consumer market, the American public.

Current Solution Attempts

Robo-calls, also called “SPAM over Internet Telephony”, or SPIT, are analogous to the unsolicited commercial email plague on the Internet known as SPAM. To some extent, similar tools can be deployed to combat it. These SPAM blocking systems commonly perform Bayesian filter matching against common signatures, source IP addresses and host name verification, looking at time-frequency domain aspects of the sender, and so on. Indeed, a Google patent search for “spam telephony” reports 12,800 related patents and applications. Controlling SPIT has been widely discussed in the Internet Engineering Task Force (IETF). A Google search for “IETF and SPIT” returns over 500,000 hits, many of which are records of mailing-list discussions between some of the smartest people in the network world. This challenge is not the first attempt to solve the problem.

Despite all these patents and efforts, there are no widely available tools recognized as generally effective. Why? Because most of the proposed solutions fail to understand the economic realities, consumer behaviors, and underlying business practices that have made robo-calling so successful. Consequently, most widely-proposed solutions are impractical.

Proposed Solution	Why it Doesn't Work
Making SPIT and other unsolicited commercial calling illegal	When robo-calls are outlawed, only the outlaws have robo-calls. The senders are generally outside US enforcement scope, and US authorities can't touch them. They are essentially another form of "organized crime".
Do-Not-Call Lists	Since the SPITters don't care about the laws, they'll happily download the do-not-call lists, then call everybody on them. What's better than a list of pre-validated legitimate phone numbers?
Reliance on Caller-ID	Caller-ID is only a "best effort" to present the source telephone number. It is easily and frequently spoofed, wildly inaccurate, and unreliable as an authenticator.
Captcha/Challenge Response	Attempts to make callers "prove" they are human, rather than bots. This requires EVERY call to be screened, probably by a human (due to the use of improving speech-recognition technology), and vast changes in user behavior. This would take us back to the days when the operator had to connect all of our calls. It's not happening.
Smart Apps on the Called Phone	Smartphone apps could do some filtering, but all they have to work from is Caller-ID, or use challenge-response. Since Caller-ID is often spoofed, and challenge-response requires changes to the calling user's behavior, and not everybody has a smart-phone, this doesn't work on any scale.
Elimination of VoIP services	We're in the process of shutting down the circuit switched world in favor of VoIP. Changing this would be insanely expensive.

Moving everything to VoIP and authenticating senders	This could be a long-term solution, but the vast majority of the world’s users are operating from legacy devices today.
Individually-managed black-and-white lists	Difficult to use: either the lists are configured “tight” and block legitimate calls, or “loose” and let in SPIT. Given that they only work on Caller-ID, they’re not reliable.
End User Tagging Solutions (star codes)	As a standalone solution, it depends on information that is available to the carrier at the time the tag is made, and is often limited to time of day, duration, origination number, and destination number. It does not provide adequate information to easily identify a spammer consistently.

In summary, to be effective, a robo-call solution must:

1. Not require dramatic changes to consumer equipment
2. Not require dramatic changes to user behavior.
3. Not be dependent on any universal database of caller identification information
4. Not be dependent on international governmental participation
5. Not require the creation of new regulatory bodies or other political infeasibilities
6. Be incrementally implementable at the transit and termination ITSP level
7. Provide value to existing stakeholders including telecommunication companies, regulatory bodies, and corporate users
8. Lead to a general decline in the viability of the SPIT ecosystem by attacking the profit
9. Have a low false-positive rate; not block legal calls
10. Adapt to changing conditions and bypass attempts at the source

Solution

Solving the problem of robocalls takes a good understanding of how they work.

Take a moment to consider how a robocall works in today’s environment. First, an autodialer exists somewhere, likely in a foreign country. It takes a list of phone numbers and calls each one with a pre-recorded message. For the clear cost advantages, those calls go out to the internet as SIP or another VOIP technology towards their VOIP carrier. Then, more than likely their VOIP carrier forwards the call to its transit carriers trying to minimize its own cost for termination. At each termination point, the call pathway becomes masqueraded further by each transit carrier. By doing this, it can become difficult to track the true path of a call. Given this complex and convoluted path, it’s easy to understand how no solution can eliminate illicit robo-callers completely. Fighting a moving target means that they will continue to adapt and try to thwart any mechanism designed to stop them. But, what if it was possible to tag and identify the “bad players” such as the carriers, peers, and customers who knowingly or unknowingly let these annoying calls through? What possibilities could this entail?

Our solution called “Robocall Mark and Sweep” (RMS -- US Patent Pending) is a high performance adaptive detection and classification system for SPIT. Inspired by SpamAssassin and Spamhaus, it can provide real-time scoring and classification of incoming calls based off predefined criteria including:

1. Call duration
2. Number validity
3. Caller Name/ID Lookups
4. IP Black-and-white lists
5. Geolocation
6. Calling patterns such as Random and Sequential dialing
7. Destination Number
8. Acoustic Fingerprinting
9. Do Not Call list participation
10. Text derived from speech

This scoring/classification event is then delivered to carrier systems which in turn follow rules and criteria defined by the carrier on how to process the call. In this manner, RMS is a multi component system that operates on both egress and ingress telephony traffic, and is designed to be used within the constraints of existing telephony infrastructure.

Now, think about “Ann from Account Services” for a moment. When calls from that autodialer come in, it ultimately terminates on the PSTN through at least one VOIP carrier if not a series of VOIP carriers. Each carrier in the pipeline has rules and criteria they track for. For instance, if a series of calls with a call duration of 10 seconds or less come in rapid succession, geolocation of the originating IPs is China or Nigeria, and the number dialed participates in a Do-Not-Call list, it’s a reasonably safe assumption that this is an illicit call. RMS then monitoring the call flows provides an event notification in real time to the carrier. At its sole discretion, the carrier then can take action. It could immediately notify their customer or peer for potential fraud. It could throttle calls. It could redirect the call to a real person to validate that the caller is human. Any number of options are available, and it allows the carrier to make the best decision given its needs and circumstances.

To understand how RMS works, think about a simple hub and spoke architecture. At the center of the hub, a central processing system (CPS) or automated validation center correlates data to generate actionable events and reports. This information is received from passive (ingress) and active (egress) listener devices. The passive listeners include devices such as Sonus equipment that generate call detail records, packet capture devices designed for call probing and so forth. With an active listener, such as a honeypot, it listens for inbound calls like a voicemail box with a predefined prompt that sounds human. Acoustic fingerprinting technologies applied to the messages then provide consistent identification of the robo-dialers. The combination of acoustic fingerprinting with other predefined criteria allows for an adaptive learning system similar to those used in existing SPAM defense systems. With human defined rules, bayesian filtering, and neural networks as methods of correlation and adaptation, the CPS makes an informed decision to classify and score a telephone call.

Upon detection of fraudulent activity, RMS generates events customized to the carrier needs. Like described previously, an event can be used as a trigger for any number of actions including but not limited to, routing table adjustments, throttling, call redirection or blocking and so forth. The decoupled design creates flexibility to adapt and adjust to continuously evolving customer requirements in the field. This allows the carrier to take responsibility for deciding what actions to take, if any, based off of the events that occur. Likewise, it reduces the liability of using any one robo-call blocking method.

This solution is also capable of supporting end user tagging as a source of input events assuming appropriate support exists from within the telephony infrastructure. End user tagging using a star code is another simple example of generating events to RMS for processing. Provided it delivers the previous origination number, destination number, timestamp and some other unique identifier, RMS can lookup the previous call metrics to match the call and mark it as SPIT, thus further enhancing the “training” of the system.

Our design recommends that the honeypot receive calls from unused or recycled phone numbers, as well as, receive calls from syndicated known bad numbers. The latter could effectively be implemented by adding known bad numbers to the Do-Not-Call list and cycling them on a regular basis. The objective of the honeypot is to capture and identify as many known spammers as possible in a way that is not easily detectable to the robo-caller. The information retrieved in both the acoustic fingerprint as well as the recording allows a trained human operator or a support desk to evaluate whether or not a call is illegal as well as provide classification to which type of call it is including: political, healthcare, non-profit, etc. This output is then used to “train” the system in identifying known spammers to include in blacklists. Likewise, it will populate whitelists with known good callers. For instance, a call from your neighborhood pharmacy or doctor’s office would be easily identified and whitelisted.

Via network taps and monitoring ports from switches at the edge of telephony networks, packet capture appliances provide ingress events to the CPS which captures criteria such as source IP address, origination and destination numbers, and other metrics on all calls. This allows for a drop-in installation within carrier networks without large changes in infrastructure and functionality. In conjunction with other sources of data such as device logging provided by carrier endpoints themselves, information can be derived from all legitimate and illegitimate calls that make it through the system.

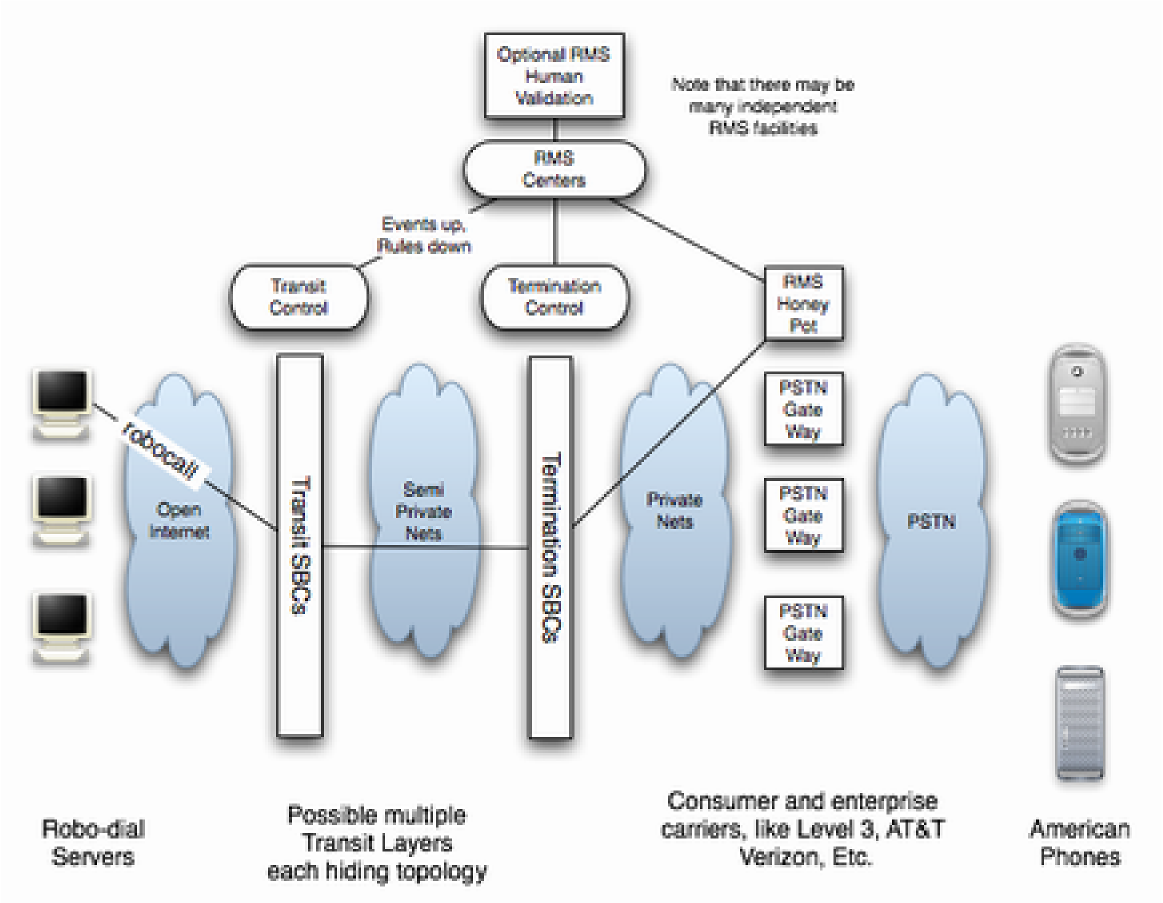

Figure 2

Combined with other egress and ingress solutions, Robocall Mark and Sweep provides a comprehensive solution to handle SPIT at all points in the call process including post call workflows such as end user tagging.

Deployability

The Robocall Mark and Sweep system plugs into the transit and terminating ITSP infrastructure without significant changes. Only three modifications are required:

1. Implementation of RMS listeners within carrier networks by means of rackspace and appropriate network connectivity.
2. Data provided to our system by the carrier such as:
    1. Feeding the call logs into the analysis system
    2. Diverting "honeypot" numbers (which can be selected from the unassigned number pools, plus advertised decoy numbers) to the automated validation system.
    3. SIP Messaging off the wire.

3. Carrier system acceptance of feedback from RMS to places that can include among others:
    1. Least cost routing engines
    2. Call center support queues
    3. Session border controllers

Additionally, terminating ISPs can provide user tagging of unwanted calls as optional input to the automated validation centers. False-positives can be further reduced by human operator review of high-probability samples as a follow-on stage after automated validation.

Incentives

No SPIT control system is likely to succeed without engagement by the telecommunications companies. While legislative or regulatory mandates may be effective in starting action, these have not been shown to produce a high level of enthusiasm. In other words, without clear incentives, carriers are likely to drag their feet on implementation.

The first-order effect of Robocall Mark and Sweep is that it provides a first-ever tool for identifying the “bad players” who are sourcing high SPIT volumes. This gives transit and termination providers the tools needed to enforce their usage and service-level agreements. It also puts peer pressure on the bad players to reform. This in turn leads to public pressure: since end-users hate SPIT, they’re going to pressure their providers to be actively engaged in blocking it, and may “vote with their feet” in favor of operators who are more successful. This provides an incentive for operators to enforce SPIT controls.

The second-order effect of Robocall Mark And Sweep is the potential it offers for ITSPs to provide additional services (for a fee) to their users, including consumer-grade anti-SPIT services, commercial-level reporting and tracking functions, as well as other potential value-adds.

Response to Challenge Questions

Does it work? (weighted at 50%)	

How successful is the proposed solution likely to be in blocking illegal robocalls? Will it block wanted calls? An ideal solution blocks all illegal robocalls and no calls that are legally permitted. (For example, automated calls by political parties, charities, and health care providers, as well as reverse 911 calls, are not illegal robocalls.)	While no solution will block 100% of illegal robocalls, our solution will provide 100% blocking of identifiable robocalls and is flexible enough to insure no wanted calls are blocked. The solution should be highly effective in blocking repeated illegal robocalls, or robocalls from repeat offenders. Legal automated calls are protected by whitelists and optionally by human validation.
How many consumer phones can be protected? What types of phones? Mobile phones? Traditional wired lines? VoIP land lines? Proposals that will work for all phones will be more heavily weighted.	All consumer phones serviced by a Telephony Internet Service Provider can be protected. It can provide protection to all PSTN phone types, which include VoIP, Wired landlines, and Wireless/Cellular phones.
What evidence do you already have to support your idea? Running code? Experiments? Peer-reviewed publications?	We are using a proven and widely deployed method of identification and marking solution based upon proven e-mail spam fighting techniques. Like SPAM, robocalls or SPIT (Spam over IP Telephony) is very similar in action. The processes used to combat one are easily adapted to fight the other. And, like email spam we can do it at the carrier level.
How easy might it be for robocallers to adapt and counter your scheme? How flexible is your scheme to adapt to new calling techniques? How have you validated these points? Remember that the real test of a security system is not whether or not you can break it; it's whether or not other people can.	Any solution once deployed will cause robocallers to adapt. Our solution accounts for this and adapts accordingly. Combining a honeypot to capture and identify illegal robocallers with acoustic fingerprinting we can quickly keep up with any changes in calling techniques. These processes are validated by years of effective SPAM prevention and are currently in use.
Is it easy to use? (weighted at 25%)	
How difficult would it be for a consumer to learn to use your solution?	Consumers use the basic functions transparently, with no training or notice.
How efficient would it be to use your solution, from a consumer's perspective?	Extremely efficient; most consumers would never notice the system.
Are there mistakes consumers might make in using your solution, and how severe would they be?	No. Consumers can't make mistakes in a system they have no direct interaction with.

How satisfying would it be to use your solution?	As the system is transparent, most users will neither be satisfied or dissatisfied -- they just won't get as many robocalls. The optional "end-user tagging" or a star code add-on may provide visceral satisfaction for its users.
Would your solution be accessible to people with disabilities?	Yes. It would be completely accessible with no additional changes to existing means of accessing the current telecommunications network.
Can it be rolled out? (weighted at 25%)	
What has to be changed for your idea to work? Can it function in today's marketplace? (E.g., Does it require changes to all phone switches world-wide, and require active cooperation by all of the world's phone companies and VoIP gateways, or can it work with limited adoption?) Solutions that are deployable at once will be more heavily weighted, as will solutions that give immediate benefits with even small-scale deployment.	Three changes are required: Implementation of our system within the carrier's network, Data provided to our system by the carrier, and Carrier processing of our system feedback (spam calls routed to appropriate systems) The system can be deployed incrementally on a per-operator basis in today's marketplace, and begins to provide its value immediately.
Is deployment economically realistic?	Using off the shelf inexpensive hardware the systems can be deployed cost effectively. Also it can provide carriers with an opportunity for financial gain.
How rapidly can your idea be put into production?	The idea can be put into production in a phased manner, starting within a few months of acceptance. Automated systems can scale easily using low cost hardware and virtualization platforms.