

# Reviewing the VoIP Threat Landscape



## **Introduction**

Voice over IP (VoIP) services are attracting a lot of interest and many organisations are implementing VoIP networks or considering their adoption. In all cases the motivation for implementing a new VoIP system or for migrating part of an existing phone network is to improve communication by streamlining existing applications, introducing new communication applications and by converging existing applications. In this context, convergence means linking existing communication applications such as voice telephony, email and Instant Messaging (IM) with newer ones such as video conferencing and providing closer integration with CRM and other back-office systems. The objective is to improve business efficiency and reduce operational costs.

The protocol that is delivering this goal is the Session Initiation Protocol (SIP). SIP along with some related protocols drives VoIP, Video and IM services and provides advanced services such as presence driven call routing.

This is a fast moving technology sector with continued development of new applications and services. In the excitement it is easy to forget that, as its name suggests, VoIP is an IP network service and that the full benefits of efficiency and convergence are gained when VoIP, IM, Video conferencing and presence based services are used on the global Internet to reach home workers and roaming users and well as business partners, customers and suppliers.

The down-side is that the Internet is a hostile place. Experience gained over the last 10 years or more running email and web servers has shown that connecting any application server to the Internet leaves that server open to attack.

This white paper examines the security risks that face deployments of VoIP and other real-time messaging applications with a focus on applications based on SIP.

## **Document Conventions**

The intention of this document is to provide information on the potential threats that face SIP based VoIP and other real-time communication applications and not to present a cook-book on how to attack VoIP systems. All of the threats described in this white paper are based on the results of a detailed threat analysis of VoIP and other SIP based applications. A large number of these threats have been implemented as applications and scripts for use in security testing or as part of a controlled vulnerability analysis. Unfortunately many of these threats have also been exploited by more unscrupulous individuals to attack VoIP systems.

All of the protocol examples shown in this white paper are based on real-life tests with many of them taken directly from network monitoring tools. In some cases uninteresting or irrelevant detail has been removed for readability and, where present, public IP addresses have been randomised.

When SIP commands (or to use the correct terminology, methods) are referred to in the text the convention of using uppercase has been followed.

## **VoIP Security Threats, Hype or Reality**

The standard response given to a discussion on VoIP security threats is "my network has not been attacked yet, and I have not heard of any other attacks, so is this all hype?". This is exactly the same reaction that greeted the first commercial firewalls when they were launched in the early 1990s yet who today would contemplate connecting a web server or email system to the Internet without a firewall? We are also seeing specialisation in perimeter security, with a strong market for email and web security products. These products are designed to supplement standard perimeter firewalls as it is recognised that the security requirements of those applications are too demanding for the average firewall to address fully.

VoIP is considerably more complex than either web or email and as we will see vulnerable to a large number of risks ranging from simple IP level threats through a complex set of application specific risks to content and usage related risks shared with other applications. No standard firewall currently available is able to address all of the risks and threats relates to SIP based

applications. If email and web servers warrant their own tailored security measures, then VoIP applications should certainly be given the same treatment.

The other standard response is “my VoIP network is completely contained within my LAN, I therefore don’t need security”. Unfortunately, unless the LAN used for VoIP services is completely isolated from all other systems and networks and especially from any desk-top workstation, this argument does not hold water. One of the factors driving the adoption of VoIP is the opportunity of linking voice telephony with other messaging applications such as email and IM and also providing closer integration with back-office and CRM systems. To achieve this goal of application convergence, voice and data services must share a common network. This means that the VoIP infrastructure is open to threats and attacks propagated via other vectors. For example an email or web download may contain a worm that specifically targets the VoIP system and exploits one of the VoIP application threats. This means that VoIP specific application level security measures are still needed *even if local policy blocks VoIP calls to and from the outside world.*

To restrict VoIP use to the local network is to overlook one of the real benefits. In the mid 1990s, email quickly developed from a basic text based communication tool restricted to the technical community to the feature rich global service we all use today. VoIP systems are following the same development path. This is because VoIP and specifically SIP based VoIP can deliver fast, cheap and feature rich communication services. As the number of deployed VoIP networks and the level of connectivity of those networks continues to grow, VoIP specific perimeter security controls are becoming an essential component of any network.

### **The Risk is Real**

The security threats outlined in the white paper are real. VoIP networks are under attack today. Don’t expect news of these attacks to hit the headlines. With a few notable exceptions IP network security breaches are not widely publicised as this is the last thing that a company targeted by an attack wants. The fact is, that all IP network applications are open to attack, either direct attacks aimed at the clients and server systems comprising that application or indirect attacks via some other application.

It is also a fallacy to assume that all attacks originate from the outside world or that all risks stem from malicious actions. Internal users are equally or more likely to be motivated to attack network systems and attempt to exploit weaknesses than external attackers. This is particularly true with threats such as call eavesdropping when local users may have most to gain from listening to confidential discussions. Many of the threats discussed in this paper, and particularly the registration and other flooding threats can arise from configuration errors or abnormal events such as a power outage or temporary network failure. When the fault is fixed, the VoIP application server can be swamped with registration requests from previously isolated clients.

### **The Economic Impact**

The most compelling reason to take these security issues seriously is the potential economic impact when just one of these threats is exploited in a successful attack or when a network incident triggers a threat. The impact of the threats discussed in this document ranges from total loss of service on the targeted system, through disruption of calls, to the leakage of confidential information. There are also the consequential risks of indirect financial loss when for example a caller ID spoofing attack is used as part of a phishing campaign and sensitive information is acquired by the attacker.

### **Threat Taxonomy**

The security threats and challenges that face VoIP networks can be classified into three main groups. These are:

- IP Network level threats. These are generic threats that apply to any application connected to the Internet
- VoIP Application Specific threats. These are threats that are specific to VoIP and related applications and in many cases specific to the underlying protocol

- Converged threats. These are threats that are a consequence of application convergence or which are shared with related applications, for example threats common to VoIP, email and IM applications

Each of these groups includes a range of specific threats. Examples of these specific threats and their consequences are discussed below.

### **IP Network Level Threats**

IP network level threats face any application that is connected to an IP network such as the Internet or includes campus and corporate networks. Servers and clients for Email and web applications have faced these threats for ten years or more, and the range of threats and the technologies used to exploit them is well understood. However new threats are constantly emerging as is evident from vulnerability tracking services such as Security Focus [1].

IP network level threats fall into 3 main categories

- Malformed packet attacks
- Flooding attacks
- Buffer overflow attacks

#### **Malformed Packet Attacks**

These attacks attempt to swamp an application server with malformed or illegal packets. There have been a number of examples of this kind of attack. One example, Jolt2 was particularly effective against Windows systems as it tied up system resources eventually causing all other services to fail. Other examples include ping-of-death which sends malformed (over length) ping requests. Vulnerable systems crash or slow down to the extent that applications cease to function correctly.

#### **Flooding Attacks**

These attacks rely on sending legitimate packets, sending such high volumes that the targeted system is so busy processing the requests that it is unable to process anything else. Even if the targeted system is able to continue to process requests, it can become so slow that applications cease to function correctly. Examples of flooding attacks include SYN Floods and UDP flood attacks.

#### **Buffer Overflow Attacks**

A buffer overflow attack exploits a software bug or coding error. A buffer overflow condition occurs when an application attempts to store data in a memory buffer that is too small for that data. If this operation is not checked and stopped, then the stored data will over-write some other memory location. If that location contained executable code then a carefully crafted message can replace that executable code with code written by the attacker. Network applications are particularly at risk from these threats as an attacker can easily send arbitrary length data to an application server. The consequences can be serious, allowing an attacker to gain control over a targeted server. Buffer overflow vulnerabilities are common, virtually every network application server ever written has been found to be vulnerable to one or more of these vulnerabilities. Sendmail, a widely used email server, has been the subject of numerous alerts, but perhaps the most well-known example is the Code Red worm that targeted Microsoft's IIS.

#### **Impact on VoIP Systems**

It's easy to lose sight of the fact that all VoIP systems are IP applications and therefore potentially at risk from these IP network level threats. If anything, VoIP systems are more vulnerable to these threats than applications such as web and email because VoIP services are much more sensitive to processing delays or network latency. The quality of a voice call relies on fast and consistent packet delivery. Any degradation to the quality of a voice call will be noticed much more quickly than a short delay in delivering an email message or loading a web page.

#### **VoIP Application Specific Threats**

VoIP application specific threats arise from the misuse and abuse of the VoIP protocols. When SIP is used to drive VoIP or related services, the end-devices or User Agents (UA) exchange a series of messages with application servers or with each other. A device like a SIP hardware

phone or a soft-phone capable of making VoIP calls uses 3 primary messages for making simple calls. These are:

- A SIP REGISTER message. This message informs an application server that a device is available to place and receive calls and associates a device's network address with the end-user's identity. This 2<sup>nd</sup> step is important for mobile or roaming users as their network address will change as they move from location to location, but their personal identity (defined in a SIP URI) will not.
- A SIP INVITE message. This message is sent to initiate a phone call.
- A SIP BYE message. This message is sent to terminate a phone call.

All SIP messages are sent as readable text in a format that looks a little like an email header. The following example shows a REGISTER message which links a SIP Uniform Resource Identifier or URI, (sip:404@borderware.co.uk) to an IP address (192.168.4.102) and also defines a personal name ("Peter Cox").

```
REGISTER sip:borderware.co.uk SIP/2.0
Via: SIP/2.0/UDP 192.168.4.102:49670;branch=z9hG4bK-d87543-f06b6c686b-1--d87543
Max-Forwards: 70
Contact: <sip:404@192.168.4.102:49670;rinstance=e3cda44dd1775e65>
To: "Peter Cox" <sip:404@borderware.co.uk>
From: "Peter Cox" <sip:404@borderware.co.uk>;tag=92469122
Call-ID: OTEzY2IyNjZhNTBiZjkwM2IwZTk1ZmRlNWNjMjJhYzY.
CSeq: 1 REGISTER
Expires: 45
Allow: INVITE, ACK, CANCEL, OPTIONS, BYE, REFER, NOTIFY, MESSAGE, SUBSCRIBE, INFO
User-Agent: eyeBeam release 1006e stamp 33793
Content-Length: 0
```

As SIP messages are text based and in most cases are transmitted in clear text, they can be modified, spoofed or intercepted. A number of application specific attacks are possible by manipulating SIP packets.

Most of the attacks listed in this section are possible for two reasons; firstly that SIP allows any request to be processed without authentication and secondly that there are no mandatory checks on the source of a message.

While SIP includes an authentication service, its use is optional in that the protocol definition allows all SIP requests to be processed without authentication. Most VoIP systems will insist on authentication for at least some operations, a SIP REGISTER command is normally not accepted without authenticating the registering device. However, many VoIP systems will allow a number of other types of request to be processed without authentication. This coupled with the fact that most systems do not check the source of SIP requests means that an attacker can send a SIP request to a VoIP phone or SIP application server, the request will then be processed with potentially devastating results.

### **SIP REGISTER Flooding**

All SIP devices send REGISTER requests when they start-up and at intervals thereafter. In some cases devices need to register as often as every 30 or 60 seconds to maintain firewall NAT mappings [2]. In networks with a large number of deployed phones, the processing load imposed on the application servers can easily reach a point where the application server is too busy processing REGISTER requests to handle new calls. This problem is exacerbated when some abnormal event, for example a power interruption or temporary network failure results in a higher than normal number of registrations.

Malicious REGISTER floods are an even worse problem, it is trivial for an attacker to construct a fake REGISTER request and flood an application server with these requests. A simple minded attack would flood an application server with multiple copies of the same spoofed REGISTER request, such as the following:

```
REGISTER sip:borderware.co.uk SIP/2.0
Via: SIP/2.0/UDP 192.168.4.199:5050;branch=z9hG4bK-FAKE
Max-Forwards: 70
Contact: <sip:911@192.168.4.199:5060;rinstance=e3cda44dd1775e65>
To: "Ambulance Chaser"<sip:911@borderware.co.uk>
From: "Ambulance Chaser"<sip:911@borderware.co.uk>;tag=92469122
Call-ID: OTEzY2IyNjZhNTBiZjkwM2IwZTk1ZmRlNWNjMjJhYzY.
CSeq: 1 REGISTER
Expires: 3600
Allow: INVITE, ACK, CANCEL, OPTIONS, BYE, REFER, NOTIFY, MESSAGE, SUBSCRIBE,
INFO
User-Agent: SIP Script V1.0
Content-Length: 0
```

A more sophisticated attack might use multiple user names to create unique spoofed requests. The application server will spend time looking up each user in its database and then send back a "Not Found" error message which of course the attacker will ignore. Even if the targeted system requires that registrations are authenticated, the attack is still likely to succeed without the attacker supplying any authentication credentials as the targeted system will be forced to look up each user and construct an authentication challenge for that user.


A successful registration attack can be disastrous. In a controlled test carried out by Borderware's professional services team, a flood of 15,000 spoofed registrations per second generated by a simple script swamped a SIP application server to the point where phones lost dialling tone and the voice quality of calls in progress degraded to a point where the system was unusable.

### De Registration Attack

A more subtle version of a registration attack relies on a variant of the standard REGISTER command where the *Expires* field is set to zero. This is used legitimately for a User Agent such as a soft-phone to indicate that it is shutting down and that no more calls should be sent.

All an attacker needs to do is to construct a REGISTER request with a zero valued Expires field:

```
REGISTER sip:borderware.co.uk SIP/2.0
Via: SIP/2.0/UDP 192.168.4.102:49670;branch=z9hG4bK-d87543-f06b6c686b-1--
d87543
Max-Forwards: 70
Contact: <sip:404@192.168.4.102:49670;rinstance=e3cda44dd1775e65>
To: "Peter Cox"<sip:404@borderware.co.uk>
From: "Peter Cox"<sip:404@borderware.co.uk>;tag=92469122
Call-ID: OTEzY2IyNjZhNTBiZjkwM2IwZTk1ZmRlNWNjMjJhYzY.
CSeq: 1 REGISTER
Expires: 0
Allow: INVITE, ACK, CANCEL, OPTIONS, BYE, REFER, NOTIFY, MESSAGE, SUBSCRIBE,
INFO
User-Agent: eyeBeam release 1006e stamp 33793
Content-Length: 0
```



This attack targets a single user agent at a time, in this case extension 404. The target of a successful attack will receive no calls, will have no dial-tone and will therefore be unable to make outbound calls. While some SIP application servers will require authentication before

processing a REGISTER request, many do not and are therefore vulnerable to an attack that can effectively disable all phones. Even if the application server enforces authentication, unless all devices use strong passwords, there is a risk that a dictionary attack will discover the passwords for at least some of the phones.

### Call Flooding Attack

SIP based VoIP applications make calls by sending an INVITE request. This request identifies the caller and the call target. Just as email systems must be able to accept email messages from any user, so an open VoIP system must be able to accept calls from any source. The consequence of this is that inbound calls from users in non-local domains cannot be authenticated and that the VoIP application has to accept any inbound call. This is analogous to the email systems where the default behaviour of email servers is to accept any inbound message.

```
INVITE sip:202@borderware.co.uk SIP/2.0
Via: SIP/2.0/UDP 192.168.4.102:34516;branch=z9hG4bK-4607e324f48-1-d87543-
Max-Forwards: 70
Contact: <sip:404@192.168.4.102:34516>
To: "202"<sip:202@borderware.co.uk>
From: "External Caller"<sip:404@example.com>;tag=9c1e4e74
Call-ID: YjQyOTU5ZGExODlmNjVmOTNmNTA3ZmI4Y2Y5NTg5NWE.
CSeq: 1 INVITE
Allow: INVITE, ACK, CANCEL, OPTIONS, BYE, REFER, NOTIFY, MESSAGE, SUBSCRIBE,
INFO
Content-Type: application/sdp
User-Agent: soft-phone
Content-Length: 417
```

A malicious attacker can flood system with calls by using simple script which drives a software application that sends SIP INVITE requests and makes calls. There are two variants of the call flooding attack; the attacker can hang-up as soon as the target phone answers or the attacker can send a recorded audio stream.

A flood of silent calls is a very effective denial of service attack. If a phone rings but the caller hangs up as soon as the phone is answered only to ring again a second or two later, the end-user is unable to make any calls or to receive any legitimate calls. The standard reaction will be to leave the phone off-hook or just to un-plug it. A flood of calls playing a recorded audio stream is a spam attack. The problem of SIP spam is discussed in more detail in the section on converged threats.

### SIP BYE Attack

VoIP calls are terminated by one of the call participants sending a SIP BYE request. Many VoIP application servers will process a BYE request without requiring authentication. This means that it is easy to construct a BYE request and send it to the application server which will then terminate the call. Tests by Borderware have shown that for many SIP application servers, the details in the BYE request do not need to be particularly accurate. In many cases the requests is processed even if the only accurate part of the request is the Call-ID. The Call-ID is relatively easy to track down as the SIP INVITE and subsequent SIP replies transmit the ID in clear text.

```
BYE sip:202@218.67.54.101:5060 SIP/2.0
Via: SIP/2.0/UDP 10.0.0.116:40822;branch=z9hG4bK-d8754a5b8
Max-Forwards: 70
Route: <sip:202@192.168.4.8:5060;transport=UDP;lr>
Contact: <sip:404@10.0.0.116:40822>
To: "202"<sip:202@borderware.co.uk>;tag=as793d0acc
From: "John Q Random"<sip:fred@borderware.co.uk>;tag=12683a24
Call-ID: 72fc645@192.168.19.12
CSeq: 3 BYE
User-Agent: soft-phone
Reason: SIP;description="User Hung Up"
Content-Length: 0
```

A BYE attack is highly disruptive because it will terminate any successfully targeted call. A significant volume of BYE attacks can render the entire VoIP system unusable. A well planned

BYE attack focused on an organisation such as a call centre could plant a worm or Trojan on a key system which would listen for new calls, capture the Call-ID wait a few seconds and then terminate each call with a BYE request. A successful attack of this kind would quickly close down the call centre.

### SIP CANCEL Attack

The SIP CANCEL attack is similar to a SIP BYE attack, except that it relies on a CANCEL request rather than a BYE. A SIP CANCEL cancels a pending INVITE before the call set up is complete. A CANCEL attack terminates a call before it completed rather than terminating a call in progress. The impact of a CANCEL attack is similar to that of a BYE attack.

### Identity Spoofing

The discussion on call flooding established that calls are made by sending INVITE requests and that application servers generally have to accept requests from non-local domains without authentication. The details of an INVITE request can be spoofed just as easily as an email sender can spoof his identity. In the SIP world, spoofing a caller-ID is as easy as faking the From: header field in the INVITE request.

```
INVITE sip:202@borderware.co.uk SIP/2.0
Via: SIP/2.0/UDP 192.168.19.152:62522;branch=z9hG4bK-4c2e3a-1726
Max-Forwards: 70
Contact: <sip:404@192.168.19.152:62522>
To: "202"<sip:202@borderware.co.uk>
From: "Bank Manager"<sip:CustomerCare@BankofEngland.co.uk>;tag=f51e5774
Call-ID: YWQ1N2M2YTk3NDRlNjZmNjBiZTNhZmNjYjY3NTJhYWU.
CSeq: 1 INVITE
Allow: INVITE, ACK, CANCEL, OPTIONS, BYE, REFER, NOTIFY, MESSAGE, SUBSCRIBE,
INFO
Content-Type: application/sdp
User-Agent: soft-phone
Content-Length: 419
```

Several years of experience with email has taught us that the From: address displayed in our inbox is not to be relied on, but users tend to be much more trusting of a phone's Caller-ID display. The risk of this attack is that a user will be tricked into revealing confidential information in the belief that the caller's claimed identity is confirmed by the Caller-ID.

### Call Transfer Attacks

SIP provides a number of methods for controlling call transfer. A SIP REFER request sent to a phone directs that phone to place a call to a supplied number or SIP URI. A REFER request can be spoofed just as easily as an INVITE, a BYE or a REGISTER request. A REFER can even be sent to a phone that does not have an active call.

```
REFER sip:202@sip.borderware.co.uk SIP/2.0
Via: SIP/2.0/UDP 192.168.19.152:62522;branch=z9hG4bK-4c2e3a-1726
Max-Forwards: 70
Contact: <sip:404@192.168.19.152:62522>
To: "202"<sip:202@borderware.co.uk>
From: "Bank Manager"<sip:CustomerCare@BankofEngland.co.uk>;tag=f51e5774
Call-ID: YWQ1N2M2YTk3NDRlNjZmNjBiZTNhZmNjYjY3NTJhYWU.
CSeq: 17 REFER
Allow: INVITE, ACK, CANCEL, OPTIONS, BYE, REFER, NOTIFY, MESSAGE, SUBSCRIBE,
INFO
Refer-To: <sip:Phisher@hacker.org>
Content-Length: 0
```

An attacker can send a spoofed REFER request to a phone, if the phone or the associated application server does not correctly validate the REFER request then an existing call could be transferred to an attacker's phone.

The consequences of this attack are obvious. If a call can be transferred to an attacker's phone or a recording device then any confidential information discussed during the call will be revealed to a 3<sup>rd</sup> party. The only down-side of this attack (from the attackers point of view) is that unless the attack takes place right at the start of the call, participants are likely to notice the

transfer. One participant will be cut-off while the other suddenly finds himself speaking to someone else.

A variant of this attack overcomes this problem. The same REFER message is used, but the original call is maintained, the attacker is effectively conferencing himself into the call. The success of this variant of the call transfer attack depends on the behaviour of the targeted phone(s) and on the VoIP application server in use.

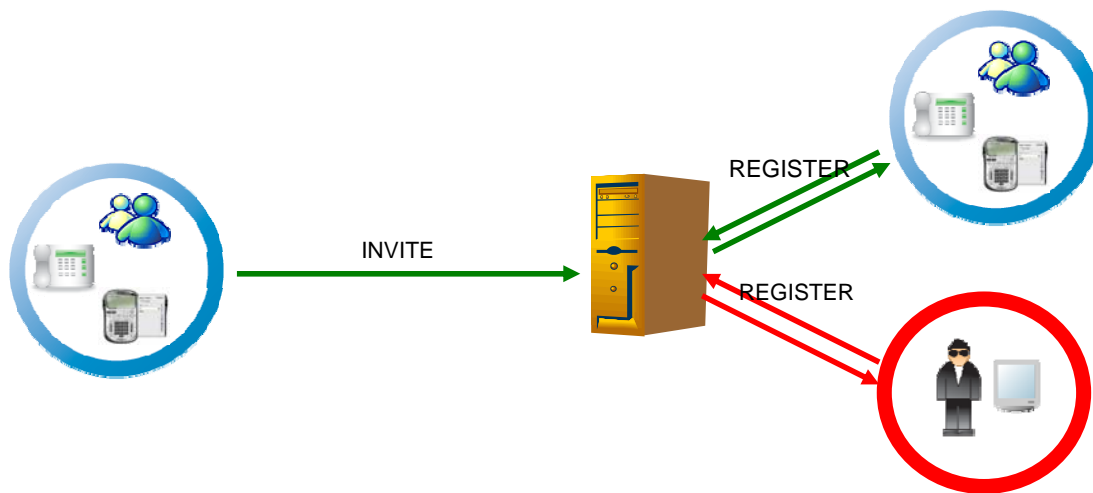
Even if this attack fails to go unnoticed by its victim, having your calls arbitrarily transferred to a 3<sup>rd</sup> party is a very effective DoS attack.

### Call Hijacking

Call hijacking attacks are similar in their effect to call transfer attacks, in that they misdirect calls. Call hijacking uses a different approach which if successful is more likely to go unnoticed. One method of call hijacking relies on a combination of a de-registration attack followed by a registration attack, both of which are discussed above. If an attacker can de-register a phone and then register his own device calls intended for the victim of the attack will be directed to the attacker's phone.

It's interesting to note that the first automated telephone exchange was invented in 1889 by Almond B. Strowger, a Kansas undertaker who apparently discovered that his local telephone operator was diverting his calls to another business which just happened to be owned by the operator's husband [3]. This was an early example of call hijacking which relied on an entirely manual method. VoIP call hijacking is faster, easier and will have an even greater commercial impact.

An alternative and stealthier method of call hijacking is to skip the de-registration phase of the attack, and simply to register a second device using the same identity. The result will depend on the details of the configuration and operation of the SIP application server, but some systems will direct calls to both the original device and to the attacker. If the attacker builds an application that passively listens to and records calls, then this becomes a very effective technique for call eavesdropping.



### Authentication Attacks

Many of the attacks outlined in this section can be at least mitigated by authentication. SIP includes an authentication mechanism based on the HTTP Digest mechanism widely used to authenticate web site access [4]. This authentication mechanism uses a challenge/response model. In SIP applications most requests submitted to the application server can be challenged by a "407 Proxy Authentication Required" message. This message will include a randomised string, which is the challenge. For example a REGISTER request may be challenged by:

```
SIP/2.0 407 Proxy Authentication Required
Via: SIP/2.0/UDP 192.168.19.12:5060;branch=z9hG4bK52a0c13a8c0342d3db5
Proxy-Authenticate: Digest
nonce="1164468517:9a8afdf6672cddcde55c2b6683316583",
    algorithm=MD5, realm="borderware.co.uk", qop="auth"
To: <sip:fred@borderware.co.uk>;tag=7c176d42
From: <sip:fred@borderware.co.uk>;tag=5431705
Call-ID: SL-rjzvmerv-494d71bf@192.168.19.12
CSeq: 15484 REGISTER
Content-Length: 0
```

The challenge or nonce is included in the Proxy-Authenticate header field. On receiving this challenge, the client will re-issue the original request and include the Proxy-Authenticate header adding the response to the header. The response is calculated as the MD5 hash of a string generated by combining the nonce, the name of the user requesting the service, the user's password and possibly other information.

```
REGISTER sip:borderware.co.uk SIP/2.0
Via: SIP/2.0/UDP 192.168.19.12:5060;branch=z9hG4bK52a0c13a8c0342e41cb
Max-Forwards: 70
To: <sip:fred@borderware.co.uk>
From: <sip:fred@borderware.co.uk>;tag=5431705
Call-ID: SL-rjzvmerv-494d71bf@192.168.19.12
CSeq: 15485 REGISTER
Expires: 3600
Allow: INVITE, ACK, CANCEL, OPTIONS, BYE, REFER
Contact: <sip:fred@192.168.19.12>
User-Agent: SIP Library
Proxy-Authenticate: Digest username="fred", realm="borderware.co.uk",
    algorithm=MD5, uri="sip:borderware.co.uk", qop=auth, nc=00000001,
    cnonce="f1318ad95cc54bd6",
    nonce="1164468517:9a8afdf6672cddcde55c2b6683316583",
    response="e33a5bedbb4abcdab5083515138645a3"
Content-Length: 0
```

When the 2<sup>nd</sup>, authenticated request, is received by the application server the server checks the response by repeating the MD5 calculation using its stored value for the user's password. If the calculated response matches the value submitted in the 2<sup>nd</sup> request, then the client making the request clearly has the correct password and the request can be processed.

Calculating the response is a computationally expensive task for the server, it has to look up the user name, extract the password from a database, combine this password with the original challenge and other information and then calculate an MD5 checksum. An attacker can exploit this to run a DoS attack. The attacker can generate large numbers of REGISTER requests, and respond to each challenge with a randomised or fixed response. Note that the attacker does not need to go to the expense of calculating MD5 checksums, any random response will suffice as the attacker does not have any valid passwords, so all responses will fail. However the application server still has to check each response before rejecting it.

An attack of this type can easily have the same result as a simple registration flood; the application server will be kept busy checking bogus authentication requests and will have less time to process new calls and to handle existing calls.

### Dictionary Attacks

Although SIP includes an authentication mechanism, that mechanism is, as with any security system, only as strong as its weakest link. In this case the weakest link is the password. Unless long randomised passwords are used the system will be vulnerable to a dictionary attack. Dictionary attack tools have long been a part of a malicious attacker's tool kit. These tools work by submitting repeated authentication attempts using a pre-defined word list as a password source or possibly just sequencing through 4, 5 or 6 digit PINs.

There is at least one SIP specific dictionary attack tool available on the Internet, this tool works on a captured challenge/response session (such as the sequence shown in the discussion on authentication attacks). The advantage (to the attacker) of this off-line approach is that the targeted system is not alerted by multiple authentication failures.

The consequences of a dictionary attack are obvious, unless strong passwords are used on all systems, then at least some passwords may be recoverable leaving the system open to attack and misuse.

### Call Relay, Toll Fraud

The idea of getting something for nothing has universal appeal and the various techniques for making free of charge calls, or calls at someone else's expense, have had a long history. In the late 1950's it was discovered that playing toy whistles down phone lines could allow free trunk calls to be made [5].

While the phone companies have long since closed these loopholes, the advent of VoIP opens up some new opportunities for toll fraud. The simplest technique, borrowed from email, is call relay. Call relay means persuading someone else's VoIP system to make a call on your behalf. Technically it's as simple as sending the appropriately formatted INVITE request. If this results in a free call to another VoIP user, then there is little financial gain from this attack. However an excessive number of relayed calls can consume system resources making it a denial-of-service attack. If a chargeable call via a PSTN gateway results then the attacker can make free calls to any location at the expense of the owner or operator of the targeted system.

A Call Relay or Toll Fraud attack has an immediate economic impact. The owner or operator of the attacked gateway is billed for someone else's calls. While many VoIP systems will require authentication for INVITE requests to non-local destinations, we have already seen techniques that can recover weak passwords. It only takes one recovered password to leave a VoIP server vulnerable to this fraud.

### Malformed Message Attacks

SIP is a complex protocol with many different message types. Each message should conform to rules that govern its format and the allowable fields and options in that message. These rules are defined in a number of Internet Request For Comment (RFC) documents. The principal document is RFC 3261 [6]. Sending messages which do not conform to the published standard is a well-known testing technique. This technique is known as protocol fuzzing. When protocol fuzzing is misused by an attacker it becomes a malformed message attack. The attack relies on sending large numbers of malformed message to a SIP application server. At best the server's resources are tied up in processing these bogus messages, at worst the message triggers a failure in the server or leaves it in an unstable state where an attacker can use other techniques to gain control of the system.

There are a number of libraries of malformed messages produced for testing purposes, some commercial, others freely available including one set published as an RFC [7].

### RTP Injection Attack

So far we have focused on SIP application level attacks and threats. However SIP is responsible for call management, it does not carry call data (voice or video). Call data is carried by the Real Time Protocol (RTP) [8]. In a VoIP call, the RTP end-points are negotiated as part of the call set-up. The initial call request, a SIP INVITE includes a payload which specifies the local RTP end-point.

```
INVITE sip:401@borderware.co.uk SIP/2.0
Via: SIP/2.0/UDP 192.168.19.12:5060;branch=z9hG4bK5323c13a8c072535991
Max-Forwards: 70
To: <sip:401@borderware.co.uk>
From: <sip:fred@borderware.co.uk>;tag=5449563
Call-ID: SL-iucdrxyd-1edd7f6f@192.168.19.12
CSeq: 15525 INVITE
Expires: 240
Allow: INVITE, ACK, CANCEL, OPTIONS, BYE, REFER
Contact: <sip:fred@192.168.19.12>
User-Agent: SIP Library
Content-Type: application/sdp
Content-Length: 270
```

```
v=0
o=fred 21283 21430 IN IP4 192.168.19.12
s=SIP-LIB-UNIX
c=IN IP4 192.168.19.12
```

```
t=0 0
m=audio 8000 RTP/AVP 0 2 3 4 5 8
```

In this example the local endpoint is port 8000 on IP address 192.168.19.12. The local system will expect to receive a RTP data stream at this address and port number. A subsequent acknowledgement (a SIP "200 OK" message) will define the remote RTP end point, in this case 192.168.4.8:12038.

```
SIP/2.0 200 OK
Via: SIP/2.0/UDP 192.168.19.12:5060;branch=z9hG4bK5323c13a8c072535991
Record-Route: <sip:401@192.168.4.8:5060;transport=UDP;lr>
Contact: <sip:401@192.168.4.28>
To: <sip:401@borderware.co.uk>;tag=as608203e8
From: <sip:fred@borderware.co.uk>;tag=5449563
Call-ID: SL-iucdrxyd-ledd7f6f@192.168.19.12
CSeq: 15525 INVITE
Allow: INVITE, ACK, CANCEL, OPTIONS, BYE, REFER, SUBSCRIBE, NOTIFY
Content-Type: application/sdp
User-Agent: Asterisk PBX
Content-Length: 180
```

```
v=0
o=root 6606 6606 IN IP4 192.168.4.8
s=session
c=IN IP4 192.168.4.8
t=0 0
m=audio 12038 RTP/AVP 8 0
```

In a normal call, RTP would flow between these two end points.

This sequence of an INVITE followed by an OK is easily monitored. Once an attacker knows one or both of the end points is it trivial to inject an alternative RTP stream by sending a sequence of RTP packets to the appropriate IP address and port. One method of doing this is to use a software application which converts a WAV file or other recording to a sequence of RTP packets. The result will depend on details of the configuration and operation of each of the end-points, but there is a real risk of at least one of the callers receiving the injected RTP stream rather than hearing the actual conversation.

### Call Monitoring and Eavesdropping

We have already seen one mechanism to monitor calls and eavesdrop on conversations; this was the call transfer attack. A more direct approach is simply to run a standard network packet sniffer and capture a conversation. Most packet sniffers include tools for extracting an RTP stream from other traffic and are even able to convert this stream to a WAV file for easy listening.

### Converged Threats

Converged threats are threats that are shared with other network applications, particularly other messaging applications such as Email, IM and browser based applications.

### VoIP Spam

We are all familiar with email spam. The volume of email spam has now grown to a point where it threatens the viability of email as a service. The impact of VoIP spam will be much greater than email spam. Wading through a hundred or so voice mail messages with the latest stock tips or promising cut-price prescription drugs will be much more time consuming than deleting spam from an email inbox.

Fortunately, VoIP spam is not yet a big problem. The reason is simple, the number of deployed and publicly addressable VoIP phones has not yet reached the critical mass needed for VoIP spam to be cost effective. Spamming is a commercial operation where spammers are paid by the number of responses. The response rate to email spam is very small, but spam is very cheap to send. Spamming becomes a worth while exercise if the target population is big enough. Sending VoIP spam is just as easy and just as cheap as sending email spam, all it needs for a the target population to grow to a critical mass before email spammers turn their attention to VoIP.

VoIP spam will become a problem unless pre-emptive action is taken. VoIP spam is so disruptive that it will quickly lead to end-user dissatisfaction with the entire VoIP system. To avoid this potential problem any VoIP security system should include spam controls.

A recently published Internet draft on SIP spam puts it well:

*“Spam, defined as the transmission of bulk unsolicited email, has been a plague on the Internet email system, rendering it nearly useless. Many solutions have been documented and deployed to counter the problem. None of these solutions is ideal. However, one thing is clear: the spam problem would be much less significant had solutions been deployed ubiquitously before the problem became widespread.” [9].*

The same source includes a good discussion on the economics of VoIP spam.

#### **Malicious and Unwanted Content**

Viruses, Trojans and other malicious content are a familiar email problem. While not a direct threat to VoIP system, as it is difficult to usefully transfer a virus as part of an RTP stream, this will certainly become a problem as VoIP is linked to Instant Messaging and other applications to form part of a larger set of converged applications. Unwanted or offensive content is as much or more of a problem with VoIP and particularly video applications.

#### **Policy Control**

Most organisations have local security policies and acceptable usage policies which apply to all traditional messaging applications. Policy defines the types of information and content that can be sent or received and should apply equally to VoIP applications as to email and IM. There is little point in investing in email security systems that check email content and prevent malicious employees sending out confidential information to a competitor, if the same employee can make a VoIP call and read the information out. The need to extend policy controls to VoIP and other real-time messaging systems will grow as those applications become more tightly integrated with traditional communication systems.

#### **Multiple Attack Vectors**

As VoIP systems join email and web servers as part of a broader set of converged applications, and as IM and Video conferencing is linked in, the risk of an attack via one application server affecting another system grows. Each system can be attacked from multiple vectors. As an example, a worm attached to an email could target a VoIP system, either attacking the VoIP application server or monitoring the traffic flowing to and from a user's workstation and recording calls made on a soft-phone.

Application convergence means that we can no longer treat each application server as an island; security controls must take a broader view.

#### **Addressing the Threats**

There can be little doubt that the potential threats that face VoIP applications are sufficiently serious to require specialised security controls. These controls must address the complete range of threats identified in this document covering IP network level threats, VoIP application threats and Converged Threats. A detailed discussion on security technologies is outside the scope of this paper. For more details, see *Reviewing SIP Security Technologies* a white paper from Borderware Technologies.

## References

1. <http://www.securityfocus.com>
2. NAT Traversal for SIP Applications, In publication.
3. <http://www.btplc.com/Thegroup/BTRegions/Scotland/History/1880s.htm>
4. HTTP Authentication: Basic and Digest Access Authentication.  
<http://www.ietf.org/rfc/rfc2617.txt>
5. <http://en.wikipedia.org/wiki/Phreaking>
6. SIP: Session Initiation Protocol.  
<http://www.ietf.org/rfc/rfc3261.txt>
7. Session Initiation Protocol (SIP) Torture Test Messages.  
<http://www.ietf.org/rfc/rfc4475.txt>
8. RTP: A Transport Protocol for Real-Time Applications.  
<http://www.ietf.org/rfc/rfc3550.txt>
9. The Session Initiation Protocol (SIP) and Spam  
<http://www.ietf.org/internet-drafts/draft-ietf-sipping-spam-03.txt>
10. Reviewing SIP Security Technologies  
[http://www.borderware.com/pdfs/WP\\_SIP\\_0906.pdf](http://www.borderware.com/pdfs/WP_SIP_0906.pdf)

### About BorderWare

BorderWare Technologies makes Internet communications safe. The company's messaging security, privacy and compliance solutions enable customers to mitigate the risks and threats associated with Internet communications. Founded in 1994, BorderWare has developed partnerships and affiliations with some of the industry's most prominent companies including Mitel, Avaya, Marconi/Ericsson, Ubiquity, Sun Microsystems, Research in Motion (RIM), Kaspersky Labs, McAfee, RSA Security, and Symantec. More than 8,000 customers in 65 countries have selected BorderWare's solutions for their superior security, scalability, business continuity and lower total cost of ownership.



Headquarters. +1.905.804.1855 | Toll Free. +1.877.814.7900 | US Federal Office.. +1.866.211.6789 | Europe. +44.20.8759.1999  
[www.borderware.com](http://www.borderware.com)

#### About this document

Borderware Technologies Inc. disclaims any and all liability for damages, costs, lost profits, fines, fees or financial penalties of any kind suffered by any party acting or relying on the general information contained herein.

©2006 BorderWare Technologies Inc. Any product photos shown are for reference only and are subject to change without notice. Internet Communications Made Safe, BorderWare Intercept, MXtreme, SIPassure, S-Core and related marks are trademarks of BorderWare Technologies Inc. Other product and/or company names mentioned are trademarks and/or registered trademarks of their respective holders. November 2006