Securing 21st Century VoIP and UC

A comparison of the UM Labs solution with legacy SBCs

Fact Sheet.

The UM Labs SIP Security Platform is designed to deliver effective security for 21st century voice over IP (VoIP) and Unified Communication (UC) applications. While most organisations are planning or actively implementing VoIP and UC deployments, many are not fully aware of the security risks associated with those applications.

Problem statement

While the organisations would not consider implementing a network data application without adequate security measures, many VoIP and UC deployments are poorly protected.

Story of the security platform revolution

As a result many companies are falling victim to call fraud and to other attacks targeting unprotected vulnerabilities. Call fraud alone is estimated to cost $40 billion per year [1].

The reality is that VoIP and UC are both IP applications. The security threats faced by these applications are IP security threats and need to be handled as such. The picture is complicated by the fact that many vendors are attempting to adapt a set of products known as Session Border Controller (SBC) to address VoIP and UC security requirements. SBCs were originally designed to meet a different requirement and are not the best solution for today's VoIP and UC security needs.

This white paper outlines the security requirements for VoIP and UC and shows how a product from UM Labs designed specifically to secure 21st century VoIP and UC applications offers a better solution than a Session Border Controller.

1. Background

UM Labs specialises in the research and development of security solutions for Voice over IP (VoIP) and other Unified Communication (UC) applications using the Session Initiation Protocol (SIP). The product which UM Labs developed entered a market which was previously dominated by a group of products known as Session Border Controllers. However, UM Labs deliberately chose not to adopt the SBC tag, naming their product the SIP Security Platform.

The reason for this choice is that UM Labs has taken a fresh view of the security challenges facing VoIP and UC users and service providers. VoIP and UC are developing rapidly both in terms of the functionality offered and in the diversity of deployment and usage models. While many early VoIP implementations were deployed within a trusted network, providing connectivity from desk phones to an IP-PBX, these same systems now link to SIP trunks and provide connectivity for remote and roaming users. Once remote connectivity is in place, there in a need to add additional functionality by implementing other real-time Unified Communications applications including video, Instant Messaging and presence.

The evolution of VoIP towards UC changes the connectivity requirements. Early VoIP deployments, contained within corporate networks, relied on legacy technologies such as ISDN for connectivity. The adoption of SIP trunks introduced the need for IP connections, with service delivery either over an MPLS connection or the public Internet. Once an IP connection is established, the natural next step is to extend VoIP services to mobile users and then to implement a full Unified Communication service. Both of these options mandate a connection to the public Internet.

The last quarter century's experience with running data applications over IP has shown that public IP networks are hostile places; connected systems continue to be attacked. VoIP and UC applications are no different. If anything, VoIP and UC system’s represent a more attractive attack target than a data application as there is the potential for direct financial gain through call fraud. All VoIP and UC systems with IP connections are at risk. As the range of services provided over these connections grows, so the risk increases.

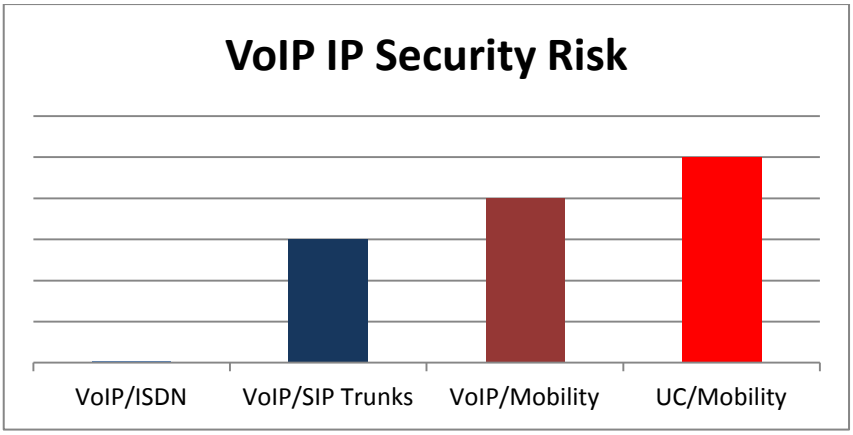

2. Security Requirements

VoIP and UC applications are complex network applications driven by complex protocols. The Session Initiation Protocol which drives most VoIP and UC real-time applications is significantly more complex than data applications such as Web and email. As always, complexity increases the likelihood of security vulnerabilities. For this reason and because of the necessity of running VoIP and UC applications over IP networks, it is essential to provide robust security for all VoIP and UC deployments. As a generation of security R&D has shown, the only effective method of securing applications on IP networks is to use a layered approach.

Effective security must address threats at the following layers:

- IP Network Level
- Application Level
- Content Level

2.1. IP Network Security

The need for IP Network security for data application is well established. There is a whole industry devoted to developing data firewalls to protect against threats at this level. VoIP and UC applications run on the same IP networks as data applications and therefore need the same protection.

The security threats at the IP Network level include:

- Denial of Service attacks (Dos) and distributed DoS attacks.
- Flooding attacks
- Malformed packet attacks
- Port scanning and service enumeration attacks

DoS attacks, attacks designed to disrupt a network services, are a growing problem. Businesses with an Internet presence are a common target where the attack can be motivated by a political protest or for financial gain. Earlier this year ING bank in the Netherlands suffered multiple attacks.[2]

The obvious question is: if firewalls are designed to protect data applications from DoS attacks, can they do the same for VoIP and UC applications? The answer is no, because the protocols used for VoIP and UC are not *firewall friendly.* Configuring a firewall to handle SIP and the related protocol used to handle audio and video streams in calls, the Real-time Transport Protocol (RTP), means opening up a large port range. This reduces the firewall's security to a level where a competent firewall administrator would not want to apply the necessary configuration.

The only effective way to implement the necessary IP security controls is as part of a comprehensive VoIP and UC security product.

2.2. Application Level Security

Application level security controls threats aimed directly at the VoIP and UC protocols and applications. The complexity of these protocols means that there is a long list of potential threats. These threats can be combated only by implementing a range of security controls directed at the application level. In VoIP and UC terms this means targeting security controls at the protocol messages responsible for functions such as tracking the status of

connecting devices and managing calls. The security threats at the application level include:

- Denial of Service attacks (Dos) and distributed DoS attacks.
- Flooding attacks
- Malformed message attacks
- Directory harvesting attacks
- Authentication attacks
- Call fraud attacks
- Protocol violation attacks

DoS attacks at the application level include flooding attacks, where the targeted system is flooded with requests, and also more subtle attacks where smaller numbers of invalid messages are sent with the aim of disrupting a service.

2.3. Content Security

The content delivered by VoIP and UC applications includes voice and video calls, text and other content delivered via Instant Messaging applications and *Meta content* such as presence information indicating the availability status of colleagues. All of these content types are potential attack targets.

The most obvious attack is eavesdropping, listening in on voice and video calls or monitoring presence data to gather information on the identity of users. There are also a number of more subtle attacks including:

- Media injection (replacing or disrupting voice or video streams)
- Media level denial of service attacks
- Call Hijacking attacks (taking over one leg of an established call)

2.4. Integrating the Security Layers

While it is important that any VoIP and UC security product addresses each of the identified layers, there must also be coordinated feedback between the layers. For example if the application security layer validates a request to initiate a voice call, that layer should also notify both the IP network security layer and the content security layer so that the voice streams for that call are permitted but all other voice streams including potential content level attacks are blocked.

3. The UM Labs Design

The UM Labs SIP Security Platform was designed to provide the security needed to protect today's VoIP and UC applications. This design uses a layered security approach which has proven effective in securing data applications. To address the complexity of the VoIP and UC protocols and applications, there is a high degree of linkage between the various layers.

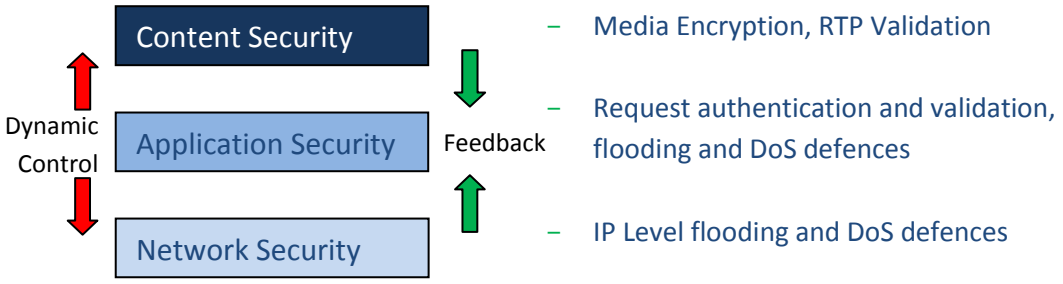

3.1. IP Network Security Layer

The IP network security layer in the UM Labs product is implemented as a robust firewall module designed to conform to the security requirements detailed in the US Government *Protection Profile for Firewalls*. This profile details the security controls needed to protect networks and applications in environments with elevated security requirements.

The firewall module included in the UM Labs product provides defences against the attacks which Internet connected VoIP and UC systems will inevitably suffer, including low-level flooding attacks and Denial of Service (DoS) attacks.

The complexity of VoIP and UC protocols means that simply including a robust firewall module is not enough. This design of these protocols makes them *firewall unfriendly,* as anyone who has tried to configure a standard firewall to handle VoIP will agree. The problem is that media streams (audio and video) do not use fixed network ports but rather dynamically negotiate these ports. This means that a standard firewall must be pre-configured with a large number of open ports. This reduces the firewall's security and means that there is no mechanism to block content level attacks. The UM Labs design addresses this problem by placing the firewall module under the control of the application level security layer. This has several benefits:

- Maintains a secure firewall configuration at all times
- Ensures that media streams for valid calls are permitted and blocks all other streams
- Hides the complexity of the firewall configuration needed to support VoIP and UC from the system administrator and simplifying system management.

3.2. Application Security Layer

The application security layer is responsible for managing and controlling all VoIP and UC application requests. This includes registering end-user devices, setting up voice and video calls, sending Instant Messages and managing presence. The security controls at this level include

- Validation of all application requests
- Comprehensive authentication services
- Encryption, ensuring the confidentiality of all calls requests (protecting the identity of caller and call recipient), protecting presence information and protecting the content of Instant Messaging.

The application security layer in the UM Labs product is implemented as an application proxy. Proxies are widely used in the data security industry and offer an efficient mechanism for securing network applications. The proxy used by UM Labs is both *transaction* and *dialog* stateful. This means that the proxy maintains state information for each transaction (for example

sending an Instant Message and receiving an acknowledgement that the message was received) and for each dialog (for example a call which includes a set-up transaction and a termination transaction). This state information enables the UM Labs product to ensure that every processed request is linked to a fully authenticated call or other operation. This makes many of the common denial of service attacks for example call termination and call hijacking impossible.

When the UM Labs product process a call with the application proxy, the proxy relays the call request monitoring and validating the request as outlines above. The use of a proxy contrasts with most Session Border Controllers which use a technique known as back-to-back user agent (B2BUA). A B2BUA processes a call by accepting the incoming call and then making a separate call to the destination. This process imposes a significant overhead but offers no security advantage. As an example, two B2BUA based SBCs from different vendors claim call set-up rates of between 8 and 84 calls per second. An equivalent UM Labs system can handle more that 500 calls per second.

3.3. Content Security Layer

The content security layer is responsible for the security of the content of VoIP and UC sessions. Content includes media streams (voice and video) and the content of Instant Messages. This layer is implemented as an application level proxy designed specifically to handle media streams (both voice and video). The proxy is designed to be extremely efficient in forwarding media streams with minimal delay while maintaining the necessary security controls and providing any conversion needed to ensure that the media streams are sent to their destination in an appropriate form.

The content security layer is under the control of the application security layer. The application security layer ensures that the content security layer permits only media sessions from validated calls and applies encryption where required. These measures prevent unauthorised eavesdropping, guard against call hijacking, media injection and other DoS attacks.

3.4. Comparing Proxy and B2BUA Architectures

Vendors of products based on back-to-back user agent architecture will undoubtedly argue that a B2BUA has advantages but in reality a well-designed proxy can deliver all of the functions of a B2BUA and do so much more efficiently. The UM Labs SIP Security Platform provides all of the features commonly attributed to a B2BUA including:

- Interworking including protocol normalisation
- Media control
- Topology hiding
- Call management, including deal call detection and generation of call data records
- Transport conversion

The proxies, for application level and content level security, used by the UM Labs SIP Security Platform deliver these features while maintaining the higher level of security provided by the layered security architecture. The proxies also are significantly more efficient than B2BUA architecture. For example a UM Labs SIP Security Platform running on hardware suitable for supporting up to 500 concurrent active calls can accept new calls at rate exceeding 500 calls per second. Rates for a comparable B2BUA based SBC range from 8 to 80 calls per second. Comparing media latency, the UM Labs is twice as efficient as a standard SBC adding 0.03 milliseconds of latency compared to 0.06 for a typical SBC.

Where the UM Labs SIP Security Platform really scores is its ability to handle attacks. One SBC claims that the fact that the CPU usage of their system is elevated to only 14% when the system is faced with a flood of spoofed responses at a rate of 500 per second as an advantage. The UM Labs SIP Security Platform can handle a flood of 2,000 spoofed responses per second with the CPU remaining at 99.9% idle.

3.5. Interoperability

The formal specifications of the protocols used by VoIP and UC applications are very broad and include a large number of options. This means that interoperability between different vendor's equipment and between enterprise systems and external services is far from guaranteed. Interoperability problems remain one of the main challenges facing the deployment of VoIP and UC systems. The UM Labs SIP Security Platform provides a solution to this challenge.

The interoperability features provided by the SIP Security Platform are implemented in both the application level and content layer modules. The interoperability features are provided at two levels:

- Commonly used operations such as converting protocol headers (for example changing a call destination from an internal to external form) which are implemented in the management GUI
- Less commonly used operations which can be applied using a simple yet flexible scripting language

4. Session Border Controllers

To quote Wikipedia [4], a Session Border Controller:

> *is a device regularly deployed in Voice over Internet Protocol (VoIP) networks to exert control over the signalling and usually also the media streams involved in setting up, conducting, and tearing down telephone calls or other interactive media communications.*

This definition reflects the origins of SBCs; they were originally designed to manage traffic flow on the connection between two service providers. While this interconnection has some security requirements these requirements are not comparable to those of today's VoIP and UC deployments. The security focus of an SBC is much more in partitioning traffic between two controlled service provider networks where there is some level of trust between the networks. This security model is not adequate for protecting VoIP and UC applications as they are deployed and used today. This limitation is further reflected in the formal requirements definition for SBCs issued by the Internet Engineering Task Force (IETF) [5].

This document focuses on controlling call set-up and media sessions and does not address security requirements at other levels.

While some SBCs have attempted to address a wider set of security needs, none of these products were designed specifically to handle VoIP and UC security when these applications are deployed on un-trusted IP networks. As a result, SBCs do not offer integrated security architecture which is a core feature of the UM Labs SIP Security Platform.

5. Deployment Options

Virtually all SBC products are delivered as hardware appliances. This is partly because of their origin as service provider interconnection devices and partly because of the architecture. Most SBCs are built on a custom hardware platform. The only deployment option for these SBCs is to install hardware and the only upgrade option is to install additional hardware or to replace that hardware with a larger system. This precludes true cloud deployments.

The UM Labs SIP Security Platform is a software product; the software runs on standard processors and offers a choice of deployment options. These include:

- An appliance deployment using standard low cost servers rather than custom hardware
- A virtualised deployment running in a private cloud
- A true cloud deployment running in a service such as Amazon's AWS or Microsoft's Azure

The UM Labs software architecture ensures a high level of throughput and efficiency without the need for custom hardware. In many cases the UM Labs product is significantly more efficient than an SBC or comparable capacity. This is particularly so in the UM Labs' capability to handle flooding and denial or service attacks.

6. UM Labs/SBC Comparison

There are six fundamental differences between the UM Labs SIP Security Platform and between a legacy SBC. These are:

1. The product design. The UM Labs product was designed specifically to handle VoIP and UC security for IP network connections and to address today's security problems. This means providing security for a range of connection types including SIP trunk connections and remote and roaming user connections. The UM Labs product achieves this with a layered security architecture. In contrast, most SBCs have evolved from a design where the primary requirement was to enable interconnections between service providers. The security needs for this class of connection are completely different from those of today's integrated VoIP and UC networks which connect service providers, enterprises and individuals. As a consequence SBCs do not provide the same level of security controls offered by UM Labs.

2. The product architecture. SBCs rely on an architecture known as back-to-back User Agent (B2BUA). Put simply, when an intermediate device receives a call request, it answers that call and establishes a second call to the destination. This architecture reflects the telecommunications background of SBCs. A B2BUA design introduces a high level of overhead. The UM Labs product is built on a proxy architecture. Proxies are widely used in IP data security and provide a very efficient method of relaying an application session while maintaining full control over that session. A proxy is far more efficient than a B2BUA, for example while B2BUA based SBCs claim to process between 8 (Cisco CUBE) and 84 (Acme Packet) new calls per second, an equivalent UM Labs product can process in excess of 500. The SBC vendors will doubtless claim that B2BUA technology offers more security or a greater level of control than a proxy; this is simply not the case. Firstly the use of proxies is well documented in the SIP standard [3], secondly the UM Labs proxy implementation can provide all of the security measures that are possible with a B2BUA.

3. Implementation. The UM Labs product is implemented as software. The software is designed to run on a range or processor types and to take advantage of multi-processor, multi core systems. This ensures a very high degree of scalability. The system can be delivered on a low cost platform for small office and can scale to meet service provider loads. Most SBCs rely on custom designed hardware. This approach limits scalability and increases cost. In addition it is much more difficult for hardware based systems to adapt to changing needs.

4. Adapting to evolving needs. One of the consequences of the SBC's reliance on custom hardware is that they are slow to adapt to evolving needs. Most SBCs were designed when the requirement was to handle voice calls. To meet this need the hardware was built with standard processors to handle signalling (call setup) and Digital Signal Processors (DSP) to handle media (the voice streams). Today's requirements are for a mix of voice, video, presence and Instant Messaging. This changes the proportion of signalling to media traffic. Products built assuming the typical signalling to media ratio of voice may need an expensive upgrade to handle new traffic patterns. The UM Labs SIP Security Platform is software based and uses standard hardware. The software is optimised to run on multi-core processors and will automatically adjust to changing traffic patterns allocating more cores to signalling or media processing as requirements change.

5. Encryption. The UM Labs SIP Security Platform was designed from the outset to provide encryption for both signalling and media channels. This means that when performance figures are quoted for the UM Labs product, these figures assume that VoIP and UC traffic is encrypted by the product. In contrast, most SBCs quote performance figures without encryption. When encryption is enabled on a SBC there is often a significant performance penalty.

6. Deployment. The UM Labs SIP Security Platform's software architecture offers a choice of deployment options. The software may be installed on a standard Intel Architecture server, may be deployed in a virtualised environment or may be installed in a commercial cloud services such as Amazon's AWS. Standard SBCs with their dependency on custom hardware do not offer this flexibility.

6.1. UM Labs/SBC Comparison Summary

Feature	Legacy SBC	UM Labs
Architecture and Deployment		
Core architecture	Back-to-back User Agent	Proxy. Significant efficiency gains with no security or functionality penalty
Implementation	Custom hardware	Software, optimised for multi-core systems
Scalability		
Session Capacity	Varies, typically 100 – 5,000	30 to 100,000+
Call setup rate	Varies, typically 8 – 100	500+
Encryption penalty[1]	Up to 90%	Zero
Security Features		
IP Level Security	Limited	Extensive protection against flooding and DoS attacks.
Application level security	Typically limited to simple rate controls	Complex rate controls (by SIP method) plus defences against DoS attacks
Session Authentication	Typically limited to INVITE, REGISTER and BYE (where supported)	Authenticates any SIP request or can delegate authentication while monitoring authentication state

[1] Performance penalty when encryption enabled

Feature	Legacy SBC	UM Labs
Co-ordination between the various security controls	Limited	Linkage between each security control enhancing the overall security
Performance under attack		
CPU Load triggered by Spoof SIP Response attack	14% at 500 responses/second (Dialogic)[6]	< 0.1% at 2,000 responses/second
CPU Load triggered by rogue RTP attack	20% at 175,000 frames/second (Dialogic)	< 0.1% at 150,000 frames/second
Deployment		
Hardware appliance	Custom Hardware	Standard Intel Architecture Hardware
Virtualisation and Cloud deployment	No	Fully supported for private and public clouds (including AWS and Azure)
Management		
Management Options	Complex GUI or CLI	Simple GUI or customisable interface to allow control from an Network Control Centre

7. References

1. Telecoms fraud costing operators $40 billion annually. http://tinyurl.com/qesmjml
2. ING again targeted in DoS attack. Dutch News. http://tinyurl.com/pu8z7ww
3. SIP: Session Initiation Protocol, RFC 3261. http://tools.ietf.org/html/rfc3261
4. Session Border Controller. http://en.wikipedia.org/wiki/Session_border_Platform
5. Requirements from Session Initiation Protocol (SIP) Session Border Control (SBC) Deployments, RFC 5853. http://tools.ietf.org/html/rfc5853
6. Lab testing summary report, March 2012. Report SR1120308 http://www.miercom.com/pdf/reports/20120308.pdf

About UM-Labs

UM Labs is a pioneer and leader in Voice over IP/SIP Video/BYOD and Unified Communications security. The company markets a family of cost effective SIP Security Platform which make connecting VoIP/Video systems to the public internet easy and secure. As a software generated solution this can be implemented on any Intel or Arm technology either as soft implementation or SPaaS service.

Confidentiality, integrity, and authenticity of voice communications over the internet (VoIP/Video) are critical considerations for most businesses. Driven by lower bandwidth costs and the promise of increased flexibility, VoIP is quickly becoming a critical tool in the business-to-business landscape. Significant growth in SIP Trunking and consolidation of voice and data traffic over the public internet are raising new security and interoperability concerns that were previously overlooked. To solve these problems, UM Labs has developed a family of cost effective SIP Security Platform which can be easily delivered as a service via the PaaS Toolkit or plugged into existing networks to enable SIP connectivity, security and voice encryption.

UM-Labs today works through our partners to deliver these innovative solutions and the ‘Innovation in Security Showcase’ becomes their blue print to demonstrate the advancement of these solutions and allows them to fit directly into the Unified Communications value chain.

Contact sales@UM-Labs.com the sales team will arrange a presentation and you can receive case studies at the same time.