RC24802 (W0906-040) June 9, 2009
Computer Science

IBM Research Report

Securing Large and Distributed Enterprise VoIP Infrastructure Using Border Elements

Archana H. Rao, William J. Rippon
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Research Division
Almaden - Austin - Beijing - Cambridge - Haifa - India - T. J. Watson - Tokyo - Zurich

Securing Large and Distributed Enterprise VoIP Infrastructure using Border Elements

Archana H. Rao and William J. Rippon
IBM T. J. Watson Research Center, Yorktown Heights, NY, USA
{archarao, bjripp}@ us.ibm.com

***Abstract*—Despite an increasing enterprise adoption of IP-based real-time communication systems, large and distributed enterprises have not yet fully realized the benefits of such deployments. Our work explores the architectural challenges and security concerns of a fully IP-based real-time communication infrastructure in such large, multi-site, multi-vendor enterprises. In this paper, we propose a novel hierarchical architecture and analyze the importance of Border Elements in securing the VoIP infrastructure. To this end, we built a pilot environment and evaluated the suitability of Session Border Controllers as border elements in pure enterprise environments. Based on our deployment experiences and real-world evaluations, we posit that Session Border Controllers are an optimal solution for meeting the security and edge services requirements to enable transition to a fully IP-based voice infrastructure.**

***Index Terms*—Enterprise IP Telephony, VoIP Security, SBC, SIP Border Elements.**

I. Introduction

As the size and scope of real-time IP infrastructure deployments in enterprise environments increases, so too do the concerns related to the performance and security of these environments (Reference [1] provides an overview of threats to IP based voice systems). For geographically diverse enterprises, having islands of real-time IP infrastructure, each administered using local policies and connected by the PSTN backbone, offered many advantages including incremental deployability, ease of administration, acceptable reliability and limited security threats from external networks. Consequently, enterprise-wide core infrastructure and security mechanisms specific to the real-time infrastructure and services have been often minimal, or non-existent. The focus of architecture and implementation has been heavily geared towards individual site locations. However, the continued maturity in products and services supporting Session Initiation Protocol (SIP) [2] based real-time communication, particularly in the areas of security and reliability, are opening up avenues for the enterprises to move towards fully IP-based real-time communication infrastructure that can be large, distributed and well protected. However, this transition from islands of IP voice infrastructure interconnected via the PSTN, to a fully IP-based infrastructure, brings a new set of challenges, the foremost of which is security.

Securing the real-time IP infrastructure in large and distributed enterprise environments offers several interesting challenges and practical constraints. First, as the different enterprise sites may be located in different socio-economic or legislative jurisdictions, or controlled by different enterprise divisions, the site administrators should be able to retain the flexibility to implement local policies without compromising the global security and functionality of the infrastructure. Second, the solutions should co-exist with, and leverage when appropriate, the existing data and network security mechanisms like firewalls and intrusion prevention systems (IPS). Third, the solution should permit seamless connectivity between areas with differing security characteristics. Last, since it is unreasonable to expect large multi-site enterprises to instantly upgrade their entire infrastructure, the solution should have incremental deployment characteristics.

In this paper we describe a new architecture for SIP-based real-time communication infrastructure, in large multi-site enterprises, that may need to support multiple real-time services in multi-vendor environments. We discuss the importance of, and analyze the suitability of, border elements for realizing the security and enhanced functionality required in such large-scale distributed deployments. To this end, we develop a sample set of requirements for functionalities at the borders, and describe our evaluation and testing methodologies. We share our insights about the pilot deployment of this security architecture in our enterprise and reflect on the lessons learned. To our knowledge, this is the first such exploration of border elements in purely enterprise environments.

The remainder of the paper is organized as follows. In Section 2, we propose a hierarchical architecture for SIP-based real-time communication and provide a detailed analysis on using the Border Elements for security. Section 3 focuses on the structure and deployment scenarios of Session Border Controllers (SBC), while Section 4 describes a methodology for selecting the most suitable SBC for a given enterprise and progressing toward production implementation. We share our test and deployment experiences in Section 5 and provide our final conclusions in Section 6.

II. Real-time IP Communication: Architecture and Security

A. Hierarchical SIP Architecture

The size and scope of enterprises, with the number of endpoints in the tens to hundreds of thousands, spread over diverse geographical regions, makes their environment somewhat similar to that of a carrier. A real-time IP architecture for such "carrier-prises" should satisfy certain characteristics. First, it should not require a clean-slate approach but instead build on the existing infrastructure without compromising on the intended global structure. Second, the correct functionality of the infrastructure at one site should have minimal, if any, dependency on the operational correctness of other sites. Third, it should allow the use of centralized services without restricting the flexibility at site level for the choice of components or services. Furthermore, it should be geared towards industry standards to provide a platform with the best opportunity for integration and interoperability. A secure, core infrastructure based on standards will also provide more flexibility and agility when it comes to adding new enterprise applications, services and components. We describe a sample hierarchical architecture that embodies these characteristics.

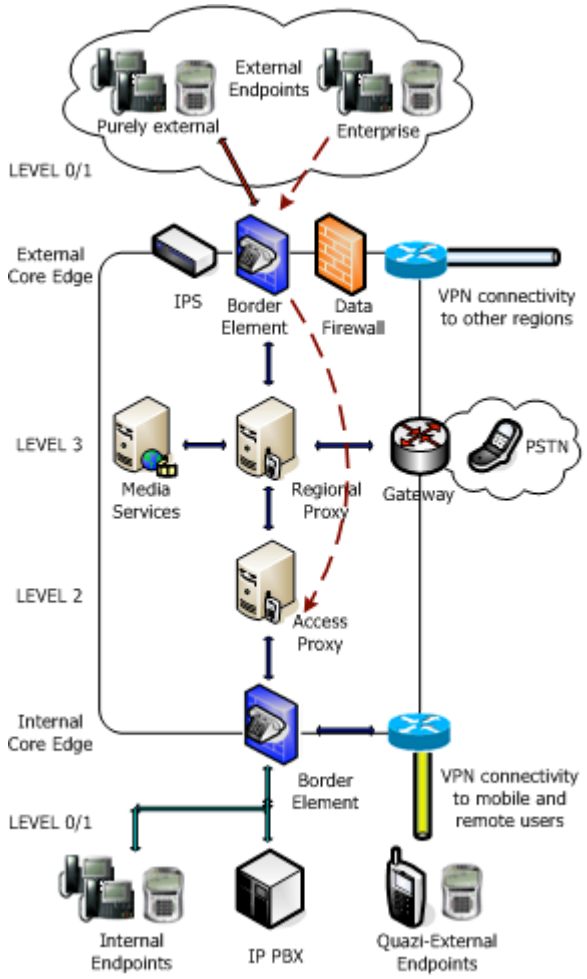

Figure 1. Hierarchical real-time IP architecture

Figure 1 shows a cross-section of the proposed structure of the real-time IP infrastructure for an enterprise, with the components and their functionalities being organized into multiple layers. The bottom and top most levels correspond to the communication endpoints or endpoint systems, located (Level 0/1) outside the boundary edge of the core enterprise infrastructure. Registration, authentication and call processing for these endpoints/systems pass through the border edge devices to communicate with proxies in the core, starting primarily the Access Proxy components at Level 2. The next layer, Level 3, contains the Regional Proxy, which is responsible for call admission and routing to/from; other enterprise access levels, other regional cores, the PSTN, the public Internet, or other external networks. As a result this level interfaces with the External Core Edge components and the level may also contain centralized core services, such as media infrastructure services (e.g. transcoding, monitoring, repair) and regional media gateways. This layered cross-section of the enterprise core will typically be replicated and interconnected, one or more times within or between geographic regions. Local or global applications, services, systems and endpoints are presented with a common, standard method for communication. Individual core cross-sections may even contain variations in configuration and products as long as the architecture and common interconnection methods are consistent. Reference [3] provides a more detailed description of this style of enterprise architecture.

B. Analysis of Border Elements

Our architectural model proposes a trust model for the real-time traffic, in which connection requests into the core infrastructure (as represented by the box around components in Figure 1) are from either trusted or authenticated endpoints. Major security zones are defined to more easily facilitate security policies that define inter-zone communication requirements and lead to standard interfaces, implementation and operation. In this example, the internal core (Dark Blue) and other internal environments (Light Blue) are a Blue Zone, DMZ or enterprise controlled interconnects (e.g. VPN) are a Yellow Zone and any purely external environment are a Red Zone.

With this background, we now analyze the security characteristics of the proposed architecture. Since all the ingress and egress connection traffic must flow through the interfaces of either Level 2 or Level 3 proxies, it is sufficient to focus our attention on these interfaces. Namely it is the borders toward the external or internal endpoints/systems that need to be secured (Note: Due to the closed nature of the PSTN, the security threats from the media gateway are virtually non-existent from an IP perspective). Owing to the complexity of the IP voice protocols together with the real-time constraints of the media traffic, an appropriate, security-compatible architecture cannot rely solely on traditional data security mechanisms (e.g. data firewall). This is where we identify the importance of border elements that are purpose-built to support security for real-time IP communications.

Choosing an appropriate border element is crucial to the

TABLE I
COMPARISON OF BORDER ELEMENTS

Border Functionalities	ALG	PMG	VAF	SBC
Traffic shaping/policing	No	Limited	Partial	Yes
Topology hiding/privacy	Yes	Yes	Yes	Yes
Authentication, Authorization, Accounting	No	Partial	Partial	Yes
Call Admission Control	Limited	Limited	Yes	Yes
Protocol inspection and manipulation	Limited	Partial	Yes	Yes
Media traffic management	No	No	No	Yes
Media encryption	No	No	No	Yes
PSTN circuit intefraces	No	Yes	No	No
Capability mismatch fixup	Yes	Yes	Yes	Yes
Intrusion detection prevention	No	No	Yes	Yes
Traffic mirroring	No	No	No	Yes

security and operation of the overall system. On one hand it needs to support all the required functionality, while on the other hand it cannot become a bottleneck to performance of the real-time IP system. In this section, we analyze the applicability of four candidate border elements namely, Application Level Gateways (ALG), PSTN Media Gateways (PMG), Voice-Aware Firewalls (VAF), and Session Border Controllers (SBC), against our requirements. Table I summarizes at a high level, how these solutions compare against one another[1].

A VAF typically consists of a SIP proxy built into the firewall that is able to handle pure proxy functionality and address translations, in addition to normal firewall functions. The PMG has a SIP endpoint on one side and a PSTN circuit connection on the other. While it has some SIP and media capabilities it does not provide end-to-end IP functionality. The ALG does provide end-to-end IP connectivity for SIP and media sessions, with address translation, but it does not provide true call processing capabilities or other advanced features. The SBC however, not only provides end-to-end IP connectivity and advanced security, its operation as a Back-to-Back User Agent (B2BUA) [2], also provides a rich set of signaling and media features. For these reasons, SBCs have already been heavily utilized in the carrier space, either as part of service provider's deployment of a formal IMS (IP Multimedia Subsystem) environment [4,5,6] or as border elements to specific applications and services.

Since the requirements of a border element in large and distributed enterprises are not fundamentally different from that of a carrier/provider, it is no surprise to see that a solution that arose specifically for the provider borders, also satisfies the enterprise requirements. There are some potential reasons why the adoption of SBCs for enterprise environments has been delayed. First, enterprises have typically employed PSTN gateways at site boundaries and avoided fully IP-based enterprise-wide deployments, which meant that the security and functionalities at the border were not as critical as they will be in fully converged enterprises. Second, the SBCs have been comparatively more expensive than the other border elements deployed in enterprises (e.g. firewalls). Third, the learning curve associated with integrating a new element into the infrastructure (especially where data and voice teams are separate). However, with the continued advancements in real-time IP technologies enabling enterprise-wide real-time IP infrastructure deployments, we see SBCs providing a crucial platform for implementing the new requirements for border security, integration and interoperability. Thus we foresee a trend toward widespread adoption of SBCs in enterprise environments.

[1] The functionalities supported by products under the same category but by different vendors could vary considerably; so this comparison is not absolute but only indicative of the widely available commercial products.

III. SESSION BORDER CONTROLLERS

SBCs are purpose-built, voice-aware, security devices that provide a variety of functions to enable or enhance session-based multimedia services. These devices play a vital role in establishing secure communication between endpoints that are located in networks with differing security and network characteristics. SBCs can manipulate the session information on each side of the connection to ensure security and interoperability, as well as providing other features that wouldn't otherwise be realizable.

Even though most SBCs (as a B2BUA) break the end-to-end philosophy [7] of the Internet and impact the session parameter negotiations in SIP, these devices are indispensable in complex large-scale networks. This has prompted the IETF to start an effort to enumerate the functionalities of SBCs [8]. There are several major ways that an SBC could be delivered. An SBC could be packaged as software-only solutions, as appliances running general purpose or proprietary OS, as a function of another component like an IPS, firewall or router, or finally, as embedded part of an IP voice application (such as call process systems for IP Telephony).

In the following subsections, we discuss some potential deployment scenarios for SBCs, in relation to the proposed hierarchical architecture.

A. Deployment Scenarios

1) Securing enterprise connections to Internet

When placed at the enterprise border to the Internet, SBC securely negotiates and manages the sessions between the enterprise realm and the external resources. This not only protects against security threats originating on the outside, but also mitigates the risk of insider attacks being propagated to the Internet. Apart from enabling communication with any external SIP endpoint, it also facilitates registration and routing for enterprise endpoints that are connected to external networks. The enterprise endpoints can be securely authenticated through the SBC and seamlessly place calls as if they were directly connected to the enterprise network.

2) Securing enterprise connections to service providers

While services like IP access to the PSTN will certainly be

done through a service provider, it is often the case that external service providers may also provide services like conferencing, contact centers and unified messaging. Access to such services could be provided over direct IP circuits to the provider, through VPNs (e.g. IPsec, MPLS), or a service provider may implement all or part of a service on the enterprise premise. SBC in this scenario would be placed on the border between the internal enterprise environment and the network access to the services of the provider.

3) As part of the enterprise core infrastructure

An interesting and perhaps to date least considered deployment scenarios would be to place an SBC internally within an enterprise, as opposed to a traditional border between two separate entities. A security boundary could be created encompassing the core SIP infrastructure of the enterprise and the SBC would manage connectivity between the core and edge devices or services within the enterprise. We see several drivers for such placement. First, to protect the core SIP infrastructure from any type of malicious or non-malicious attacks experienced from inside the enterprise [9,10]. Second, to address the device capability mismatches by providing protocol inter-working such as SIP-to-H323 and IPv4-to-IPv6 SIP infrastructure. Third, the SBC could be utilized to protect organizational or regional subsections of the enterprise infrastructure. Last, the SBC could be used to create a secure platform specifically for interfacing with mobile and wireless users of the enterprise.

4) Securing enterprise users on their home network

For enterprises that permit their employees to work from home, we see a potential for having a smaller footprint SBC solution, perhaps embedded in a SOHO[2] router, to enable the employees to use their home Internet connectivity to accomplish everything that they were able to do inside of the enterprise network. At a minimum, such deployments would solve NAT problems for employees working from home.

B. Co-existence with Other Elements

The ideal behavior in the converged network of an enterprise would be to have the SBC co-exist with other components, leveraging the strengths of each component to provide a better overall solution. SBCs would only take on the real-time communication traffic, leaving other data traffic to firewalls, IPS, routers, etc. However, there is no reason why these other security and control mechanisms could not be utilized to further enhance the real-time communication services. For example, routers, firewalls and IPS components could all be utilized as pre-verification mechanisms vetting specific types of issues that they are well suited to, alleviating potential performance issues on the SBC.

IV. Adopting SBC in Enterprises

This section outlines a five-phase process for adopting an SBC as a new component in an enterprise production environment. The process begins with the compilation of SBC requirements, based on the best estimates of current and future needs, into a ***Request for Proposals*** (RFP) document, which is circulated to all potential vendors. Apart from serving as a communiqué to the vendors, the RFP would also serve as an extensive test case document to track testing and evaluation in later phases of the project. Responses from interested parties go through a screening process, which leads to a short-list of candidate SBC solutions that would be involved in the second phase of the project, the ***Lab Testing*** phase. The *Lab Testing* phase focuses on hands-on testing and evaluation. The third phase introduces a candidate solution into the ***Living Lab*** environment, where real users would ascertain the SBC's impact on the voice environment as well as other production systems. Feedback and experiences from the *Lab Testing* and *Living Lab* phases would be leveraged in the next phase, the institution of a formal ***Pilot Deployment***. The final phase would be the integration of the SBC into a ***Production System.*** The remainder of this section focuses on the *RFP*, *Lab Testing* and *Living Lab* phases.

A. Request For Proposal (RFP)

There are a wide variety of ways that the requirements for an enterprise SBC could be grouped and organized for the RFP. The IETF has published an Internet Draft [8], which defines the high level services that an SBC should provide for SIP based, real-time communications. There are also standards related to IMS, which discuss the SBC functionality in the IMS environments. However, the RFC does not provide the level of detail or adjunct services needed for an enterprise deployment, and the IMS specifications are more complex than what is likely required for current enterprise implementations. Furthermore, the RFC and IMS documents do not provide any performance benchmarks, which would be crucial for planned production deployments.

The approach we took was to organize the requirements into four main categories; (1) *Security, (2) Services Functionality* (which is further subdivided into primary and secondary, based on whether they directly relate to real-time communication services vs. providing ancillary services), (3) *Interoperability* and (4) *Performance.* Individual requirements within these categories were characterized as *Mandatory*, *Current*, *Future*, *Investigative* or *Informational*.

B. Lab Testing

The goal of the second phase is to evaluate candidate SBCs as per the requirements outlined in the RFP. Central to this task is the creation of an evaluation testbed that emulates the intended communications architecture along with its components and services, and the ability to generate the desired use-case scenarios with a minimum of physical resources. We accomplished this by creating a lab environment through a combination of purpose-built test lab resources, inter-connectivity with existing production environments, and extensive virtualization of networks and servers. Building a multi-site, multi-layered infrastructure, we emulated real-time traffic loads (including attacks) from four regional areas, and also connected the test lab with our voice service provider. Integrating the loaner equipment from the candidate vendors into the test environment, we exercised all

[2] Small Office Home Office

of the test cases listed in the RFP (Refer to Section 5 for further details).

Testing and automation tools are critical to this phase, and those employed typically include call generators, traffic generators, and call analyzers. Apart from the ones provided by specific vendors, we used commercial tools like Emprix Hammer [11], Spirent Smartbits [12] and open source tools like SIPp [13], PROTOS [14], Sip Bomber [15] and SiVus [16] for call generation, intrusion attacks (e.g. DoS), spoofing and fuzzing.

C. Living Lab

Living Lab refers to a concept in our enterprise, where new and innovative solutions are mixed with production environments containing real users and real traffic. This approach provides earlier and more widespread availability of new services, allowing a variety of stakeholders to leverage the technology and achieving a more comprehensive assessment of the overall impact. The *Living Lab* can be utilized for; test, prototype, pilot, pre-production and even production level solutions. This differs from pure test-bed environments, where the users are significantly limited and traffic may be simulated. It also differs from formal pilot deployments where the scope of deployment and the number of users are significantly limited.

V. Test, Evaluation & Deployment Experiences

We share the experiences[3] of our SBC efforts from the Lab Testing and the Living Lab phases. The evaluations are organized along the following four dimensions – security, performance, functionality and interoperability.

A. Security

Our security requirements encompassed two aspects – first, the security of the real-time infrastructure and services, and second, the security of the SBC component itself.

1) Security features for real-time communication services

We performed detailed evaluations during the lab testing and living lab phases, to verify the ability of the candidate SBCs to meet our enterprise specific benchmarks, and also to exercise the solutions robustness. While the performance requirements may vary across the enterprises, we identified the following as some of the key security features for the hierarchical enterprise architecture we proposed earlier –

- Access Control Lists (ACL) for comm. services
- Pre-processing of SIP headers (validity checking)
- Intrusion mitigation (e.g. DoS, protocol attacks, etc.)
- Defining Trusted Peers with configurable thresholding
- Ability to restrict application service interfaces to only the necessary protocols (typically SIP, RTP, RTCP)
- Logging and accounting, to discover potential attacks or other security issues (e.g., the number of incomplete session initiation attempts)

It is important that these security features and general functionality for the application services be completely bi-directional in nature. In other words, we are not looking for products that exhibit an "inside-out" type of security, where the inside is believed to be relatively free of security concerns.

2) Self-protection

As a border element, the ability for an SBC to protect itself against attacks and disruptions, from any unwanted source (internal or external), is crucial to its acceptance in enterprise environments. Here are some of the key mechanisms that will be relied on in the SBC to achieve the required levels of compliance and resiliency –

- Dedicated management interfaces (no real-time traffic)
- ACLs to avoid processing unnecessary traffic
- Control plane protection mechanisms
- Support for key network management interfaces (e.g. SNMP), logging and accounting (including auditing)
- Individual, Multi-Level Admin accounts
- Secure management protocols (e.g. SFTP and SSH)

As an example to reinforce the above recommendations our testing revealed that certain management interactions, even valid ones, could significantly degrade the real-time communication services provides by the SBC (call drops, processing delays, etc.)

B. Performance

Since the SBC forces all the real-time traffic of the enterprise to pass through a common ingress/egress point, it is easier to control and manage the traffic. However, many functions activated in the SBC will typically have a negative impact on performance. While a suitably provisioned SBC can handle the average expected load extremely well, we have observed performance degradations in excess of 50%, when operating under the maximum expected load, with one or more of the following features turned on, or when dynamic changes in these areas were performed –

- Access Control Lists
- SIP header translation entries
- Anchoring or transcoding media
- Logging and Traffic mirroring
- Stateful clustering (mirroring session state)
- Quality of Service
- IPsec and TLS

Therefore, we recommend careful consideration of features required and extensive testing to assess the full impact of enabling (or modifying) the desired features. Furthermore, media offload should be utilized in call scenarios where only signaling traffic is required, as long as the security requirements are being satisfied.

Of equal importance is the overall scalability of the SBC solution. Performance levels and feature impact will vary between vendors. Clusters of SBCs or load balancing mechanisms may be required to achieve necessary throughput. Call admission controls should be in place to limit call rates and volumes from defined sources and destinations (undefined sources/destinations can be placed in a general pool with appropriate call levels). These controls will dampen the impact of overload situations. Alerting and monitoring would provide information for proactive responses to rising call levels.

[3] Although our observations may not be universally true across all the SBCs in the market; we are fairly confident of having covered many of the significant SBC vendors, to draw generic inferences

C. Functionality

Our functionality evaluation focused on two aspects – correctness and completeness of the border services, and the ability to meet our enterprise-specific benchmarks. From our experiences, we infer that the following SBC services and functionalities are some of the key items that would be relevant for large and distributed enterprise environments –

- Operation as a Back-to-Back UA
- Topology hiding and privacy
- Registration forwarding
- Protocol (SIP, SDP) validation and manipulation
- Advanced routing & translation (including DNS and ENUM routing)
- Call Admission Control and Quality of Service
- Multiple transport / network layer connectivity options
- Multiple physical and virtual interfaces to increase flexibility in supporting multiple usage profiles

We found that most of the current generation SBCs supported these functionalities quite well. However, we identify a set of essential features that may be absent or only partially supported. While none of these may be "show-stoppers" for initial implementations, we think that including these functions would increase the acceptance factor.

- Support for IPv6
- Support for SDP header manipulation
- Advanced loop-back routing detection
- Support for real-time statistics for media traffic
- Advanced signaling and media timer configurations
- Ability to rate-limit on all SIP methods
- More advanced alerting, logging & auditing capabilities
- Distinct admin accounts for any content replication

D. Interoperability

SBCs can play a crucial role in solving the interoperability issues between endpoints/systems in multi-vendor environments, by acting as both *'protocol police'* and *'protocol doctor'*. . The SBC should maintain very high levels of compliance with standards and can be used to identify and correct limitations of other components. SBCs can achieve these goals through it's placement as a B2BUA in the signaling and, or media paths. They force the SIP stacks and media engines to adhere to the standards, but can also manipulate messages if necessary to provide interim connectivity for those that do not comply with the standards or lack essential functionality. For example, performing protocol manipulation or translation (either signaling or media), when the two endpoints are incompatible.

This protocol policing function when first introduced to the existing components in the enterprise environment may expose previously unknown problems with a devices protocol stacks. Pre-production testing is key to identify and correct these issues, which can significantly impact operation. For example, a SIPUA in the test environment was not properly adjusting SIP CSEQ field [2] values, but current devices in the path were still processing calls. The introduction of the SBC exposed the protocol error and did not forward these packets. In our experience to date, we find the pros for interoperability far outweigh the cons.

VI. CONCLUSION

New architectures and implementations will appear as large and distributed enterprises continue their march towards a fully IP based, converged, real-time communications environment. This still relatively new phenomenon exposes a lack of well-understood and well-deployed real-time security mechanisms in enterprise environments. We propose a novel hierarchical real-time IP architecture and infrastructure that captures all the enterprise requirements and enables a smooth transition. We identify the importance of Border Elements in such setups and for the first time, and explore the applicability of Session Border Controllers in the enterprise space. Discussing the migration of our enterprise to a fully IP-based real-time network, we provide a real-life case study for a large, multi-site, multi-vendor environment, covering architectural aspects, design challenges, product features and deployment experiences. We find that commercially available SBCs in today's market are mostly adequate for the needs of large enterprises. They provide not only the requisite security and edge services, but also acting as the glue to solve many interoperability issues.

ACKNOWLEDGMENT

We are grateful to numerous people in various IBM organizations, as well as to the vendors, providers and partners who contributed directly or indirectly to this effort.

REFERENCES

[1] W. Rippon, "Threat Assessment of IP Based Voice Systems", IEEE Workshop on VoIP Management and Security, 2006
[2] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, N. Handley, and E. Schooler, "SIP: Session initiation protocol", RFC 3261, Internet Engineering Task Force, May 2002
[3] W. Rippon, "Enterprise Architecture for IP Based Voice Systems", IBM Research Technical Report, August 2008
[4] 3GPP, "TS 23.228: IP Multimedia Subsystem (IMS); Stage 2", 9/2002.
[5] 3GPP, "TS 24.228: Signaling flows for the IP Multimedia call control based on SIP and SDP", September 2002.
[6] 3GPP, "TS 24.229: IP Multimedia Subsystem (IMS); Stage 3", 9/2002.
[7] J.H. Saltzer, D.P. Reed, and D.D. Clark, "End-To-End Arguments in System Design", ACM TOCS, Vol 2, pp. 277-288, November 1984.
[8] J. Hautakorpi, G. Camarillo, R. Penfield, A. Hawrylyshen, and M. Bhatia, "Requirements from SIP Session Border Control Deployments", Internet Engineering Task Force, Internet-Draft
[9] Nimda Worm/Virus Report, http://www.incidents.org/react/nimda.pdf.
[10] Security starts from the inside out, http://www.computerworld.com/securitytopics/security/story/0,10801,98431,00.html
[11] Empirix Hammer, http://www.empirix.com/
[12] Spirent Smartbits, http://www.spirent.com/
[13] SIPp, http://sipp.sourceforge.net/
[14] PROTOS Test Suite for SIP, http://www.ee.oulu.fi/research/ouspg/protos/testing/c07/sip/
[15] SIP Bomber, http://metalinkltd.com/downloads.php
[16] SIP Vulnerability Scanner, http://www.vopsecurity.org/