

Session border controllers – a primer

Introduction

This paper provides the reader with a basic knowledge and understanding of the need for, and capabilities of, the network element known as a session border controller, or SBC. The paper is divided into two parts.

Part I describes what drove, and continues to drive, the need for SBCs, including:

- The respective abilities of the PSTN and IP networks, including the Internet, to effectively control and transport interactive communications
- The basic components and methods used in session-based IP communications
- Roles and limitations of other network elements used in IP communications

Part II provides SBC “basics,” including where in the network they’re typically deployed and the set of functions they deliver. The paper concludes with an overview of Acme Packet’s SBC products and platforms.

Part I

A tale of two networks

Voice and data services and applications have traditionally been delivered over disparate networks; voice over the Public Switched Telephone Network (PSTN) or over private networks based on Time Division Multiplexing (TDM) technology, and data over IP networks such as enterprise LANs, private WANS or the public Internet.

TDM networks such as the PSTN were created decades ago to provide seamless, reliable and secure global voice communications services—with an emphasis on the word *voice*. These networks deliver high reliability and security such that users have high confidence in utilizing them to share personal information and engage in activities such as banking and commerce. TDM networks, however, are limited in their ability to support high bandwidth video and other interactive multimedia services.

IP networks have historically provided global reach for a broad range of information services such as e-mail, web browsing, electronic commerce and research. IP is a data-oriented protocol which provides global addressing among computers. The service quality over IP networks, while adequate for these types of

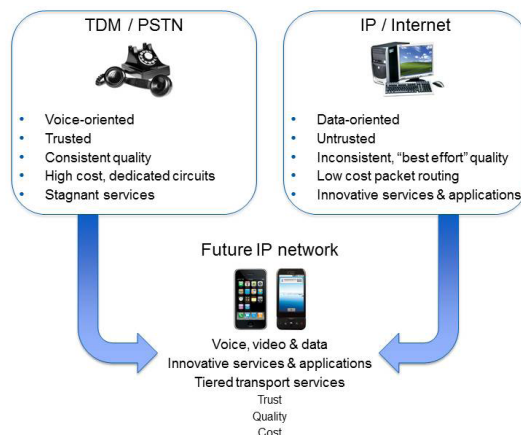


Figure 1: Evolution to next-generation end-to-end IP communications

information services, can vary significantly depending upon factors such as available bandwidth, web server activity, the number of active users at any given time and the activity being performed. Although IP networks are capable of cost-effectively transmitting any form of traffic that is IP-based, including interactive voice, video and data, many IP networks, especially the Internet, transmit only on a best-effort basis in which all forms of traffic have equal priority. This can result in significantly varying degrees of perceived quality for the same or similar types of traffic transmissions. In addition, IP communications, unlike those over “closed” networks such as the PSTN, are subject to disruptive and fraudulent behavior, including identity theft, viruses, unsolicited email (known as SPAM), unauthorized use, and attempts to circumvent or bypass security mechanisms associated with those services, known as hacking. Although Internet users have adopted many security measures to protect themselves, their networks and their websites, these measures currently are not adequate to provide highly secure, real-time interactive communications.

Evolution to a converged IP network

IP networks can be designed and operated more cost-effectively than TDM networks. In addition, IP networks are capable of delivering converged voice, video and data services and applications to businesses and consumers. Service providers are seeking to provide next-generation IP-based services to enhance their profitability by generating incremental revenue and by reducing subscriber turnover. Enterprises are searching for ways to unify their communications by seamlessly integrating voice, video, instant messaging and collaboration while reducing costs. Managing two distinct networks—TDM and IP—is not a viable economic proposition. As a result, service providers and enterprises have begun to migrate to a single IP network architecture to serve as the foundation for their next-generation services and applications. In order to successfully transition to a single IP network, however, they must maintain the same reliability, quality and security that have for decades exemplified their delivery of voice services.

Challenges of delivering interactive communications over IP

IP networks were designed initially to provide reliable delivery of data services such as file downloads and website traffic, which are not sensitive to latency, or time delay. If data packets are lost or misdirected, an IP network exhibits tremendous resiliency in re-transmitting and eventually executing the desired user request, which is a generally acceptable result for these types of services. However, IP networks historically have not been capable of guaranteeing secure delivery of high-quality interactive communications such as voice and video.

A session is a communications interaction that has a defined beginning and end, and is effective only when transmitted in real-time without latency or delays. In order to enable a session based communication, control of the session from its origination point to its defined end point is required. No single IP network extends far enough to enable that level of control, however, and the Internet lacks the fundamental quality of service and security mechanisms necessary to consistently deliver the security and quality users expect for real-time multimedia communications. In order to gain the acceptance of users, service providers and enterprises must be able to assure secure and high quality interactive communications end-to-end, as sessions traverse multiple IP networks.

Managing session-based communications

To provide secure and high quality interactive communications, IP networks must be able to manage and integrate the communication flows that comprise a session. Each session includes three sets of bidirectional communication flows:

- **Session signaling messages**, which are used to initiate, modify or terminate a session;
- **Media flows**, which are data packets containing the actual media being exchanged; and
- **Media control messages**, which are used to compile information to report on quality of service levels.

A session is initiated using signaling messages. These messages establish a virtual connection between the participants' personal computers, IP phones or other end-points. In addition, they negotiate the IP addresses used for the session's media streams and control messages as well as the algorithms, referred to as codecs, used to digitize analog voice and video. Various codecs are available for voice and video transmission, each offering trade-offs between quality and bandwidth efficiency. Once the call is initiated, media streams and control messages flow in both directions between participants. Signaling messages also are used to transfer a call, place a call on hold and terminate a session.

The management of session based communications is complicated by the following characteristics of today's IP networks:

- The identities of the participants are difficult to ascertain and security needs are complex.
- The number of session signaling protocols, codecs and related standards continues to grow.
- Addressing schemes are not consistent or compatible across networks.
- Bandwidth and signaling element resources are finite.
- Interactive communications service provider business models and regulatory compliance requirements continue to evolve and require network flexibility.

Additionally, unlike typical data communications, not all session based communications can be treated with the same priority. For example, an emergency services (911) call or a high quality enterprise video conference should take priority over a person viewing TV entertainment.

Limitations of existing network elements

Successful session based communications require tight integration between signaling and media control. However, existing network elements such as softswitches, IP Private Branch Exchanges, or IP PBXs, unified communication servers, routers and data firewalls do not provide the control functions required for session based communications.

Softswitches, IP PBXs and UC servers

Softswitches, IP PBXs and unified communications (UC) servers set-up interactive communication sessions using signaling protocols such as, Session Initiation Protocol (SIP), H.323, Media Gateway Control Protocol (MGCP) and H.248. Session agents associated with these protocols process only signaling messages while performing a variety of signaling based functions, such as user registration, authentication, authorization and session routing based upon telephone numbers or SIP addresses. These session agents currently do not provide functions related to, for example, media control for interactive communication sessions or protection against signaling based denial of service and distributed denial of service, (DoS/DDoS), attacks. DoS/DDoS attacks prevent network equipment from receiving legitimate network traffic by flooding the network with unrequested information or inundating the equipment with non-compliant protocol messages.

Routers

Routers make simple routing decisions for discrete IP packets based upon IP addresses. Routers do not participate in call signaling, and therefore, are unable to recognize the multiple individual data packets that comprise a session. Without signaling intelligence, routers currently are unable to perform key border control functions such as signaling overload prevention or session routing based upon quality and cost requirements. Routers may use a number of QoS (quality of service) technologies, such as Multi Protocol Label Switching, (MPLS), Differentiated Services, (DiffServ), and Resource Reservation Protocol, (RSVP), to grant preferential treatment to certain IP packets. However, routers using these technologies are currently incapable of classifying all the communications flows associated with a single session and handling those communications flows correctly as a single entity. Without the ability to identify the multiple individual packets that comprise a session, control call signaling, or understand the access link capacity and utilization, the router is unable to make any call admission or rejection decisions. As a result, the router will continue to send packets along a path even though the session should have been rejected because the quality was insufficient for the requested session. When this overloading occurs, not only is the quality of the session associated with that packet unacceptable, but other sessions using that same path also will suffer degradation.

Data firewalls

Data firewalls are the most common security element in IP networks. Firewalls work by allowing into the network only traffic that has been requested from inside the network and by presenting a single IP address for all of the personal computers, phones and other devices behind it. The firewall effectively blocks session-based communications because it does not allow incoming calls from unknown end points. Like the router, the firewall is unable to group together the multiple media flows associated with a single session and apply consistent policy. Furthermore, firewalls are not capable of identifying and protecting against service overloads or DoS/DDoS attacks on other signaling elements such as the softswitch.

An introduction to session border controllers

Session border controllers, or SBCs, enable the delivery of secure and high quality interactive communications across multiple IP networks.

For service providers, these include the separate IP networks that comprise fixed line, mobile and cable networks. SBCs are deployed at the borders between IP networks, such as between two service providers or between a service provider and its enterprise, residential or mobile customers.

For enterprises, SBCs are used to interconnect communications “islands” that exist within the enterprise, connect the enterprise to a wide-area service designed for interactive communications (e.g. SIP trunk), or enable “federations” between multiple enterprises for B2B communications. Enterprise SBCs (or E-SBCs) also enable selected remote locations or mobile workers to securely access enterprise interactive communications services via the Internet.

SBCs are the only network element capable of integrating the control of signaling messages and media flows used by interactive communications. This capability complements the roles and functionality of routers, softswitches and data firewalls that operate within the same network. SBCs support a broad range of communications applications, providing key control functions for enterprises and service providers alike to uniquely ensure: security; interoperability and service reach maximization; quality of experience (QoE), availability and service level agreement (SLA) assurance; service revenue optimization and cost management; and regulatory compliance, while also supporting next-generation services and applications.

The evolution to SBC-enabled IP communications

Prior to the advent of the SBC, IP network infrastructure equipment, such as those discussed above were able to initiate and route undifferentiated data, but lacked the ability to target specifically the management of interactive communications sessions. The development of the SBC, unlike many emerging networking products, was not initially catalyzed by standards bodies, but rather by the pragmatic needs of service providers and enterprises.

To date, SBCs, now more frequently recognized by standards bodies, have also been deployed around the world to support next generation interactive communications services and applications such as VoIP, videoconferencing, instant messaging (IM) and presence, as well as the routing of voice conversations over both private and public IP networks, including the Internet.

SBC deployment at access, interconnect and trunking borders

SBCs are deployed at the borders of IP networks. The border between two service providers is referred to as an interconnect border; the border between a service provider and its enterprise, residential or mobile customers is referred to as an access border. Enterprises also deploy E-SBCs between their IP network and their service provider’s network, referred to as the trunking border.

The border between enterprise data centers and their employees is referred to as the enterprise access border. SBCs act as the source and destination for all signaling messages and media streams entering and exiting the network. To that end, SBCs complement rather than replace existing network and service infrastructure.

At all borders, SBCs sit in front of session agents, such as softswitches, IMS Call Session Control Function (CSCF) elements, IP-enabled Mobile Switching Centers (MSC), IP PBXs, unified communications servers and application servers, and make call acceptance or rejection decisions. This function protects the session agent from both malicious signaling attacks initiated by hackers and non-malicious overloads, as well as, ensuring calls are only accepted when adequate network quality and softswitch resources are available.

At many borders, SBCs are deployed and function in parallel with data firewalls. The data firewalls protect web and application servers and PCs from attacks while the SBC protects session agents.

SBCs augment the simple discrete packet-by-packet routing decisions routers make. Unlike routers, SBCs classify flows as interactive communication sessions and make more intelligent routing decisions to ensure secure, high quality communications.

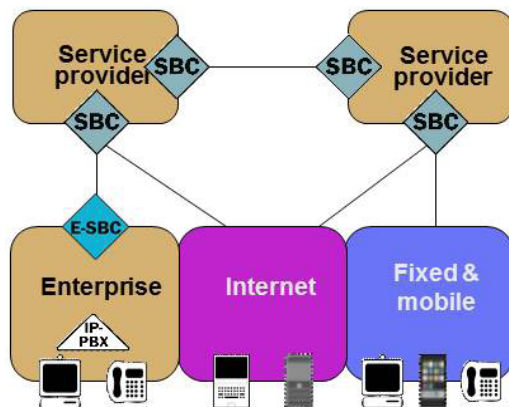


Figure 2: SBCs deployed at enterprise and service provider network borders

- **Interoperability and service reach maximization.** Our products extend the reach of IP communications by resolving differences between the many types of networks and devices supported. Critical features include: NAT traversal, which is the ability to enable communication sessions to be carried over existing data firewall and NAT devices; bridging of private and public IPv4 and IPv6 address spaces, interworking between VPNs, signaling, encryption and transport protocols; transcoding; and number, address and response code translations.
- **Quality, availability and SLA assurance.** Our products support a number of features designed to guarantee session capacity and quality. These features include: re-routing around failed links or upstream elements; admission control based upon signaling element load, bandwidth availability (including policy server interfaces) and observed quality of service; quality of service marking and mapping; and quality of service reporting.
- **Revenue optimization and cost management.** Net-Net OS includes features that help service providers maximize revenues and both service providers and enterprises minimize network CAPEX and OPEX. These include protection against revenue leakage from service theft, including bandwidth policing, quality of service theft protection, accounting, session timers, least-cost routing and load balancing. Support for virtual SBCs is one of several Net-Net OS capabilities designed to help enterprises and service providers minimize costs.
- **Regulatory compliance.** Our products support compliance with government mandated regulations worldwide, including emergency services such as E-911, the Government Emergency Telecommunication Service (GETS) and lawful intercept such as CALEA in the United States.

Other Net-Net OS features include the following:

- **Multi protocol support.** A broad range of signaling protocols to enable interworking, load balancing and routing, and decomposed SBC control.
- **Programmable signaling manipulation.** Powerful and flexible capabilities to inspect and manipulate any field in SIP headers as well as protocols transported by SIP such as SDP and ISUP. In this way, the Net-Net SBC addresses a practically unlimited range of signaling interworking and interoperability challenges.
- **High availability.** Protection against loss of service in the event of hardware or software failures. The checkpointing of media, signaling and configuration state is designed to ensure no loss of active calls, or support for new call requests.
- **Management.** A comprehensive collection of element management tools and operational support system interfaces.
- **Architectural flexibility.** Support for different architectural models including the decoupling of signaling and media control functions (known as a decomposed SBC) or the more popular integrated SBC model, which tightly coordinates signaling and media controls within a single system. Net-Net SBCs can also be configured to deliver dedicated functionality such as transcoding, or integrate with other Net-Net products such as our Net-Net SG multiservice security gateway (MSG) or Net-Net PD policy exchange controller (PEC) for comprehensive multiservice capabilities.

Acme Packet SBC hardware platforms

Acme Packet's Net-Net platforms address a broad range of performance, capacity and bandwidth requirements. Each of these platforms is more fully discussed below:

- The **Net-Net 3820** platform is our solution for small to mid-sized service providers, government defense and security-focused agencies, small to medium enterprises and smaller sites within larger organizations. Hardware-based transcoding, encryption and QoS measurement options deliver high-end SBC functionality on this 1RU platform to meet the critical requirements for next generation session delivery networks.
- The **Net-Net 4500** platform delivers unmatched performance and capacity for service providers, enterprises, government organizations and contact centers in a 1RU form factor. Optional hardware-based transcoding, high-capacity encryption and QoS measurement position the Net-Net 4500 for deployment in a wide variety of services and applications
- The **Net-Net 9200** platform offers our highest levels of performance, availability and capacity to service provider and large enterprise VoIP/UC deployments in a single 7 RU chassis-based system. The multiprocessor Net-Net 9200 platform features high performance transcoding and transrating for a wide selection of wireline and wireless codecs.
- The **Net-Net 14000** platform, introduced in 2011, is a highly-scalable solution for large next-generation service providers. Based on industry-standard Advanced Telecommunications Computing Architecture (ATCA), the Net-Net 14000 integrates Acme Packet's custom hardware in ATCA-compliant modules to minimize cost of ownership in central office environments while offering the flexibility and scale to support millions of subscribers per shelf.




Figure 4: Acme Packet SBC platforms

Acme Packet Net-Net Central

Net-Net Central, our next-generation management platform for service providers and enterprises, delivers highly-scalable configuration, fault, performance and security management for Acme Packet products. Its flexible high-availability architecture accommodates small to very large networks and provides extensibility for hosting advanced management applications and services. Through multiple dashboard and configuration views, Net-Net Central facilitates flow-through provisioning, capacity planning and comprehensive performance and fault-monitoring with “at-a-glance” status indicators that simplify real-time network-wide management. Through standard interfaces including SNMP, SFTP, XML and SOAP, Net-Net Central also integrates with OSS/BSS ecosystems to deliver advanced service fulfillment, service assurance, billing and mediation.

About Acme Packet

Acme Packet enables the delivery of trusted, first class interactive communications—voice, video and other real-time multimedia sessions—and data services across IP network borders. Our Net-Net family of session border controllers, session-aware load balancers, multiservice security gateways and session routing proxies supports multiple applications in service provider, enterprise and contact center networks—from VoIP trunking to unified communications to hosted enterprise and residential services to fixed-mobile convergence. They satisfy critical security, service assurance and regulatory requirements in wireline, cable and wireless networks; and support multiple protocols—SIP, H.323, MGCP/NCS, H.248 and RTSP—and multiple border points—service provider access and interconnect, and enterprise access and trunking.



100 Crosby Drive
Bedford, MA 01730 USA

t +1.781.328.4400
f +1.781.425.5077
www.acmepacket.com

© 2011 Acme Packet, Inc. All rights reserved. Acme Packet, Session-Aware Networking, Net-Net and related marks are trademarks of Acme Packet. All other brand names are trademarks or registered trademarks of their respective companies.

The content in this document is for informational purposes only and is subject to change by Acme Packet without notice. While reasonable efforts have been made in the preparation of this publication to assure its accuracy, Acme Packet assumes no liability resulting from technical or editorial errors or omissions, or for any damages resulting from the use of this information. Unless specifically included in a written agreement with Acme Packet, Acme Packet has no obligation to develop or deliver any future release or upgrade or any feature, enhancement or function.

09/01/11