Software Anatomy of a Broadband Residential Gateway

Jack Manbeck, Senior Broadband Architect
Bill Witowsky, Chief Technical Officer
Broadband Communications Group, Texas Instruments

Introduction

Software is an essential ingredient for today's broadband residential gateway (RG) devices that provide consumers with high-speed access to the Internet and networking capabilities between computers and devices throughout the home (see Figure 1). Cable modems and DSL modems have evolved from simple bridge devices providing basic Internet access to full-blown routers integrating Voice over Internet Packet (VoIP) telephony services and wireless LAN access point functionality. This white paper:

- Provides an overview of the various software elements and software architecture that comprises today's broadband residential gateways
- Presents some of the key implementation challenges
- Looks at future trends and implications going forward

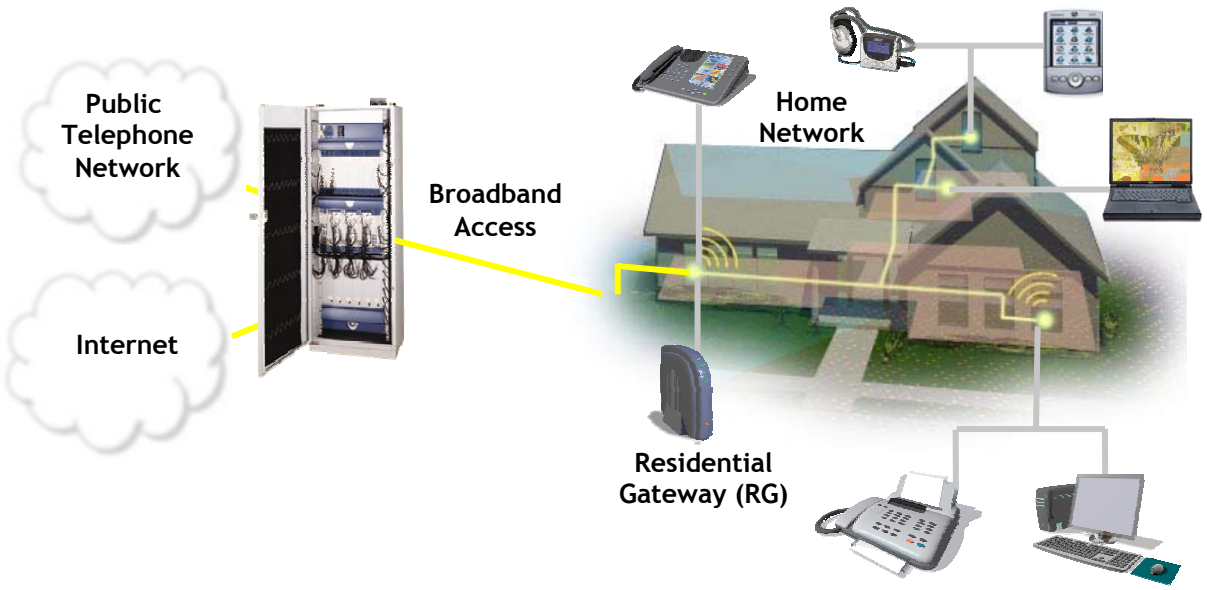

Figure 1: Residential Gateway Providing Broadband Access and Home Networking

Software is essential for providing differentiated features, flexibility and manageability in RG solutions and has been instrumental to enhancing broadband service deployments:

- Intelligent configuration along with remote network management capabilities has helped significantly reduce the number of truck rolls and customer churn.
- Software-based systems allow the equipment to be upgraded in the field to support new features and services.
- Software is being used to allow operators to offer tiered data services to attract more customers. It also allows equipment manufacturers to provide value-added features to end customers on a trial basis, in order to entice them to purchase these features as upgrades, e.g., advanced firewall, content filtering, etc.
- Finally, software is essential to readily achieve equipment interoperability and to address evolving industry standards.

A key enabler of these RG products is System-on-a-Chip (SOC) integration. The emergence of SOC solutions is a major factor in allowing manufacturers to meet consumer demands for price, performance, compactness and power efficiency. Increasing software content is being driven by SOC integration since these solutions combine highly programmable devices such as communication processors, digital signal processors and networking peripherals into a single chip. The complexity of these chips has in turn, shifted the burden for providing packet processing and networking software from the OEM/ODM to the semiconductor manufacturer. This is because robust software which includes interoperability and performance testing of the solution must be available to completely verify the silicon. Availability of product software from the silicon manufacturer also results in faster time to market and allows OEM/ODMs to focus on value-added features.

Figure 2 shows an example of a broadband access RG. The RG typically supports a broadband interface such as a DSL or cable modem, or it may be a standalone device that connects via Ethernet to a separate broadband access device. It has one or more local area network (LAN) interfaces such as Ethernet, 802.11 WLAN, USB, etc. and may support derived voice services via VoIP. A powerful communications processor is used to forward packets between the broadband interface and the LAN interfaces and to perform network management functions.

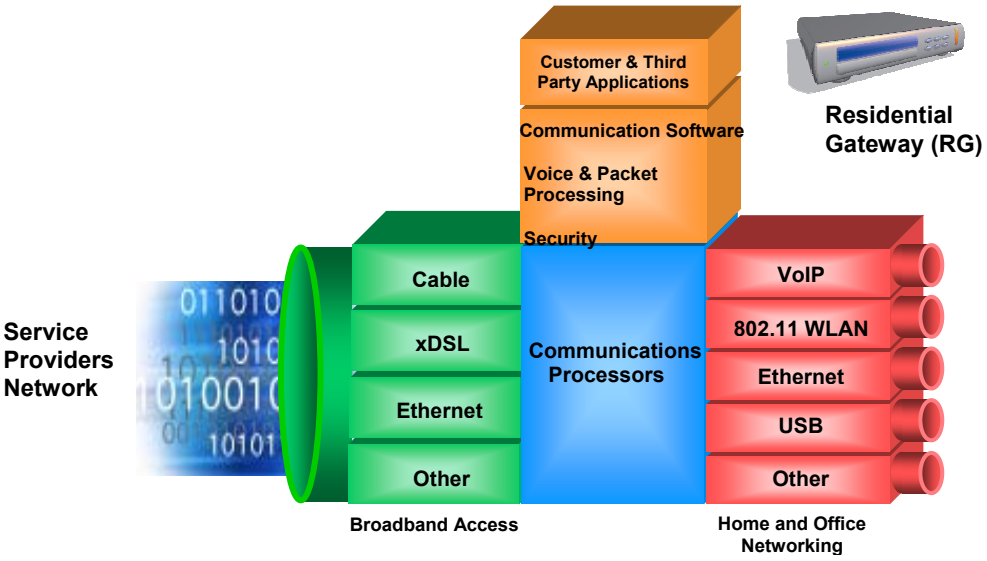

Figure 2: Broadband Access Residential Gateway (RG)

Residential Gateway Software Architecture

Figure 3 shows the software that comprises a DSL residential gateway and is an example of the amount of software that resides in today's products. In this example, the RG consists of a DSL broadband interface, Ethernet Interfaces, a USB interface, 802.11 WLAN access point interface and POTS interfaces that support VoIP telephony services.

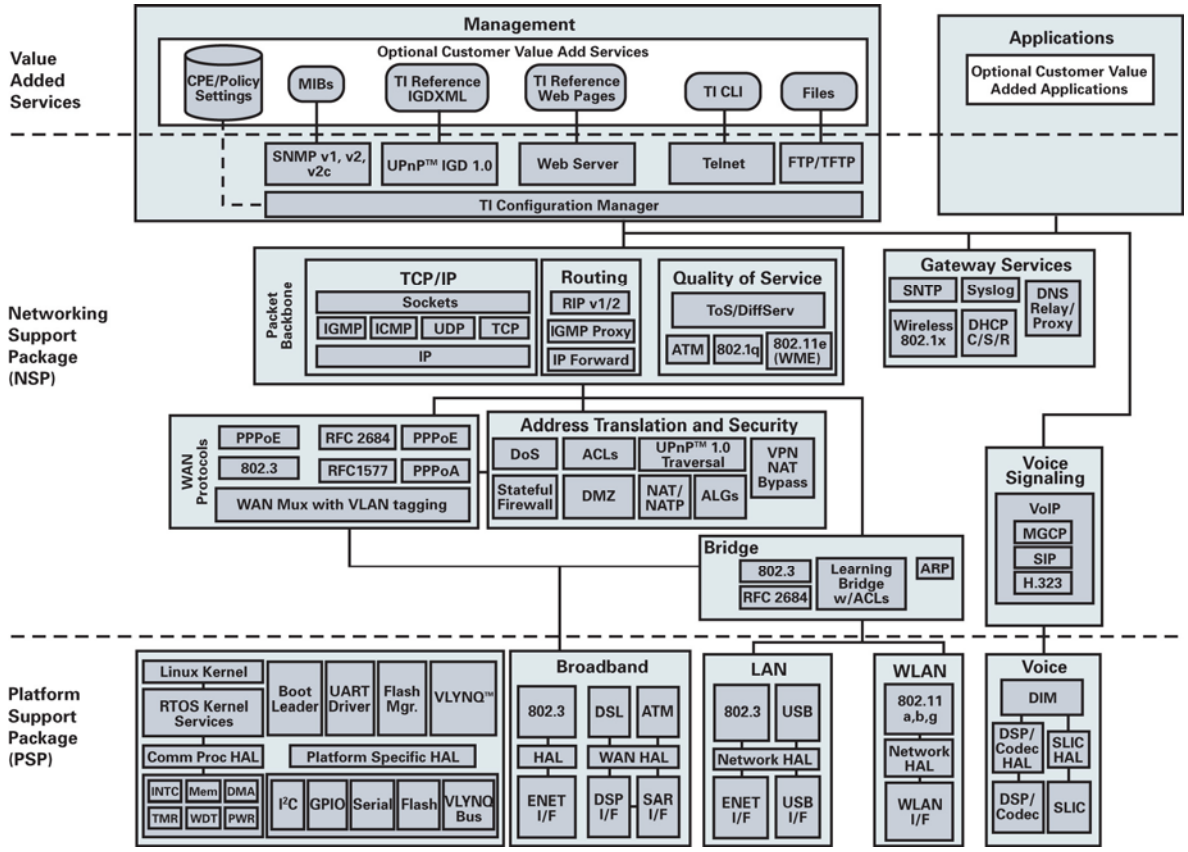

Figure 3: DSL Residential Gateway Software

The **platform support package (PSP)** provides the necessary components for implementing higher layer software features and functions on top of the hardware. Specifically, it abstracts both the SoC and the underlying platform hardware from the higher level software. The PSP typically contains the following:

- A boot loader to perform initial board level diagnostics and then load the executable image.
- Device drivers for all silicon and hardware-related peripherals. These device drivers use a hardware abstraction layer (HAL) that decouples the device drivers and higher level software from the specifics of the silicon/hardware. Drivers are provided for:
    - Networking peripherals. This includes the broadband interface (DSL or cable), LAN interfaces (Ethernet, USB, Wireless LAN, etc.) and telephony interfaces.

  - Communications processor peripherals such as
    - Timers
    - Interrupts
    - DMA
  - Reference platform peripherals such as
    - Flash memory for program and configuration storage
    - General purpose input/output (GPIO) pins controlling LEDs, etc.
  - Other peripherals and buses such as
    - UART
    - PCI

- Pre-port support for leading real-time operating systems (RTOS) such as VxWorks, Linux or Windows® CE. An abstraction layer is used to decouple the platform software from RTOS dependencies to facilitate portability.

The **networking support package (NSP)** performs all of the network packet processing and management and includes the following functionality:

- Bridging

  The bridge contains 802.3d source transparent and 802.1q VLAN bridging, allowing it to be used in those environments where Layer 3 routing of data packets is not required. When bridging, decisions are made about where to send data packets based upon the destination MAC address in the Layer 2 MAC header.

- Routing

  Routing determines over which interface a data packet should be sent. These decisions are typically made using the destination IP address and a set of routing tables within the RG. Traditionally, the RG has not had a need to support sophisticated routing protocols since in general there is only one WAN connection, and all traffic is merely passed to/from the WAN interface to the various LAN interfaces.

  The need for more sophisticated and flexible routing is becoming necessary, as services and applications such as triple-play devices begin to grow. The ability to be able to tag specific forms of data and route it over specific physical or logical interfaces is necessary. This allows for the separation of the data by the service provider's network. To accomplish this, policy routing and other mechanisms become necessary.

One mechanism for routing packets over logical/physical interfaces is VLAN port tagging on the WAN interfaces. This allows the RG to add or remove specific VLAN tags for individual data flows. For example, the RG could tag voice, video and data and then route the traffic separately over a combination of physical and/or logical interfaces, allowing the network to more efficiently route the data to its proper destination.

- WAN Protocols

  The WAN protocols provide the necessary encapsulation, and possibly authentication, when communicating with the service provider's broadband remote access server (B-RAS). Primarily, the RG should support two sets of WAN protocols. One set is Ethernet-based and includes 802.3/802.1q and PPP over Ethernet protocols. The second set is DSL-based. Using these two sets, the RG can be adapted to several different configurations, allowing it to be used for cable, DSL, fiber, satellite and other deployments.

- Address Translation and Security

  Network Address Translation

  The most common use of network address translation (NAT) for a residential gateway is to enable the LAN devices to use one set of IP addresses for internal (LAN) traffic and another set for external or (WAN) traffic. The internal addresses are never seen by people on the Internet, and the external addresses are never seen by the LAN device. In that respect, the role of the RG is to replace (or translate) network packet IP addresses between the LAN and the WAN. Figure 4 shows the most common scenario, in which the LAN devices sit behind the residential gateway[1].

[1] In the world of acronyms, this type of configuration is often referred to as M:1, Basic NAT and/or Masquerading.

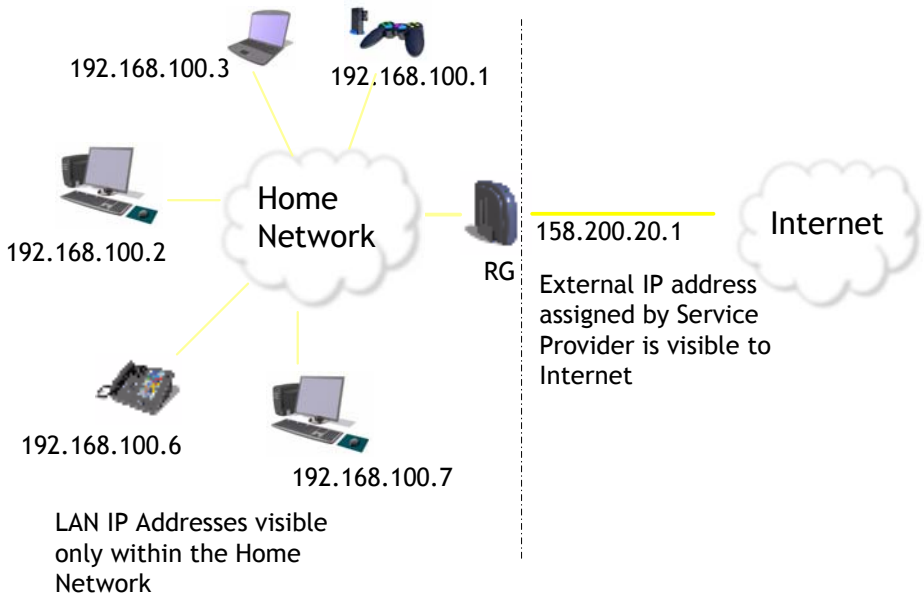

Figure 4: Network Address Translation with a Residential Gateway

In this scenario, the LAN devices are assigned what are called private or non-routable IP addresses from the RG's DHCP server[2]. The RG must perform two basic functions to allow the LAN devices to use a set of addresses that are different from the ones assigned by the service provider. They are address replacement and routing of incoming packets to a device on the LAN. Address replacement involves changing the source IP address for IP packets from the LAN to the WAN and vice versa for packets from the WAN to the LAN. Routing of incoming packets involves deciding which of the LAN devices should receive the incoming data packet. This would either be a response to a request from a LAN device or an unsolicited packet.

Translating IP Addresses

When an RG replaces IP addresses in a packet, it typically does so in one or more of the well-defined IP (IP/TCP/UDP) headers. This is an easy function to perform as the fields are generally well-known, and the RG knows exactly where to look. There are times though, when an application on the LAN will decide that it needs to communicate its IP address to the outside world. Common examples of this are instant messaging and voice telephony. When this occurs, the application is "unknowingly" including its internal address as

[2] Although any address may be used, there are specific address ranges that have been set aside from RFC 1918 for this particular use.

data that it sends in the packet. When this occurs, the feature or function of the application that uses the address typically fails as the address is not "known" outside of the LAN environment. To solve this problem, an RG will typically implement an application level gateway (ALG). This ALG is software on the RG that understands a specific application or protocol and "knows" where to look in the data to replace these internal addresses with the external ones. This requires the RG to have an ALG for each such application. While ALGs solve the problem, they increase processing complexity which can affect packet throughput and require additional memory (Flash and RAM). For these reasons, an RG will typically only support ALGs for very common applications. A potential solution to this problem is being dealt with by a new set of standards developed by the Universal Plug-n-Play (UPnP™) Forum[3], which is discussed further in this article.

Routing Translated Packets

When a connection is made from the LAN to the WAN, a signature is created that the RG can use for mapping back incoming responses. This is a typical scenario for applications such as web browsing and works well on most RGs. What happens, however, if the connection is not initiated by a LAN device? For example:

- An application may register its IP address so that it can receive calls back at a later time. Some examples are video cameras, voice telephony and peer to peer file sharing.
- Some applications can host a service such as a Web server or game server that expects to receive unsolicited requests from the WAN.

In these examples, there is no signature for the RG to use when it receives a packet from the WAN. When this occurs, the RG may not know to which application it should route the incoming packet. Thus, the application, or that specific application feature, fails to work.

This is a fundamental problem with NAT and has generally been solved using several different mechanisms which include port forwarding and/or a demilitarized zone (DMZ). These mechanisms, though, introduce other concerns. Port forwarding requires a user to configure the RG with the protocols and ports an application is using and on which LAN device. Not only is this not user friendly, leaving ports open when the application may not be running presents an increased exposure to probing by hackers. A DMZ is a catch-all case. It generally says, if I do not know who this packet should go to, send it to

[3] Texas Instruments is a member and active participant in the Universal Plug-n-Play Forum.

this IP address which is associated with this device. Clearly this does not work in all cases either. These are problems that are not easily solvable without some sort of standard. One such standard that is addressing these needs is the Internet gateway device specification by the UPnP Forum.

Universal Plug-n-Play and NAT

The UPnP Forum has developed the Internet Gateway Device Specification to meet the limitations and problems associated with network address translation. Briefly, the specification allows applications to automatically perform the functions of port forwarding and to query the RG for its external address if it must be included as data. Among the many benefits provided by the spec are:

- Removes the need for users to know what ports/protocols their applications use and on which devices, making both the RG and their applications seamless and easy to use.
- Removes the burden of developing and testing ALGs.
- Solves the issue of unsolicited packets.

This in turn leads to increased user satisfaction (it works out of the box) and lower costs for both ownership and support. To find out more about the UPnP Forum and its associated standards for seamless networking, please visit their Web site at: (http://www.upnp.org).

Security and the Residential Gateway

Security can be a complicated subject as it involves many factors and several layers, each with a different focus. To understand what security features the RG should support, one must first have a model to help understand its role. Such a model is shown in Figure 5.

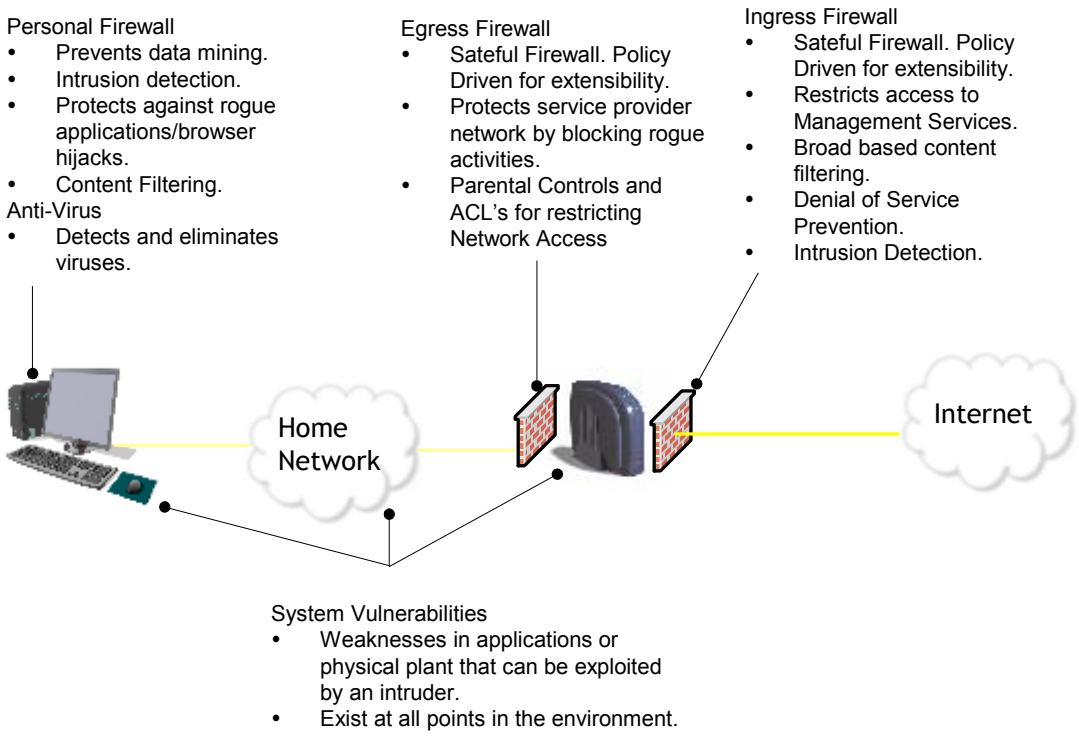

Figure 5: The Role of the Residential Gateway in Providing a Secure Environment

As can be seen from the model, the RG should:

- Provide an ingress stateful packet inspection (SPI) firewall that:
  - Prevents denial of service (DoS) attacks on hosted services.
  - Monitors traffic for intrusion detection with a positive feedback mechanism to alert the user to the condition and to provide steps they should take should such an occurrence happen.
  - Restrict access to management mechanisms of the RG to authorized users.
  - Can filter content on a broad scale for all LAN devices.
  - Is extensible through a policy language to quickly add new features and functions.
- Provide an egress SPI firewall that:
  - Protects the service provider network from rogue applications users may have inadvertently downloaded (example: distributed DoS attacks).
  - Has access control lists for applications such as parental controls.

SPI Firewall

SPI firewalls generally provide two important benefits:

1. They are typically more secure then just ACLs or NAT because they monitor and understand the protocols being used and can reject packets that do not appear to belong to a valid session.

2. They have a policy language that, along with the increased understanding of protocols and sessions, allows more sophisticated ACLs to be written in a timely fashion to check for a wider variety of attacks and so forth.

As stated above, a basic SPI understands the IP protocols being used and can track the state of connections between a WAN and LAN device. As such, it checks to be sure that packets being sent or received belong to one of those devices. In addition to this, an SPI can be extended with application level gateways (ALGs) to monitor other protocols and open or close ports that may be used. A very common protocol that is supported is FTP. Going a step further, the SPI has the same sort of general concerns that NAT does, in that an application may receive an unsolicited packet on a port which the SPI does not believe to be part of the session. When this occurs, that feature of the application fails to work unless the firewall ports are opened with some other mechanism (e.g. port forwarding or UPnP). For performance reasons, an SPI typically shares functions with Address Translation to reduce the number of touches per packet.

- Gateway Services

  Gateway services provide necessary functions required by all RGs.

  ### Dynamic Host Configuration Protocol (DHCP)

  There are three DHCP services. They are client, server and relay. The DHCP server is used for assigning IP addresses to the LAN devices. A critical feature, for usability, in the DHCP server is its ability to re-assign the same IP address to a LAN device. This is necessary when port mappings for NAT are used as it allows the mappings to remain valid. The DHCP relay agent is used for forwarding DHCP client requests over the WAN. While not often used, it can be a critical component in video deployments where an IP set top box must request an IP address from the service provider (SP) instead of the RG. The DHCP client is used for assigning IP addresses to some WAN interfaces when the protocol does not have a mechanism for negotiating one.

Domain Name Server (DNS)

The DNS proxy/relay service handles resolution of the LAN device's DNS requests and will typically allow the user to enter a user friendly URL name (e.g. myrouter) to get to the RG's web pages (as opposed to its IP address). Caching of the DNS requests makes web browsing by the LAN device more responsive.

(S)NTP

The network time protocol service provides the ability to be able to set the date and time for the RG without the need for costly hardware solutions.

SYSLOG

The SYSLOG service provides the ability for components within the RG to generate standard SYSLOG messages. These messages can then be captured by a SYSLOG application running either locally or remotely in the network.

802.1x

The 802.1x service provides advanced authentication, especially for wireless applications.

- Telephony Voice Services

  Voice services provide the necessary core components for Voice over IP applications and are broken down into three well-defined areas: Call processing, CO emulation and voice packet processing.

  The call processing subsystem detects the presence of new calls, collects addressing information and maintains the state of a call. In this respect, support for various telephony signaling standards such as SIP, MGCP and H.323 are required.

  The CO emulation subsystem provides for control of the SLIC and functions such as DTMF tone detection, dial tone generation, caller ID, pulse detection (rotary dial) and ringing.

  The voice packet processing subsystem contains the necessary software to convert between analog voice and packets. This includes voice compression, e.g. (G.711, G.723, G.729, AMR, etc.), echo cancellation, tone detection and generation, packetization, jitter processing and packet loss concealment.

The ability of the software to address other critical concerns when the VoIP services are integrated with the RG is equally as important as the quality of the voice software. The concerns that must be addressed are:

- The call processing subsystem needs to interact with the stateful firewall to open and close ports as needed based on the calls in progress.
- Management and provisioning for popular SP VoIP deployments is necessary as each SP has unique provisioning for the phones and service.
- Quality of service (QoS) must be applied to the voice packets along with a scheme for fragmenting other data packets to reduce head of line blocking. Head of line blocking is a common cause of poor voice quality as it may add significant latency or delay to voice samples.

- Quality of Service

  Applications such as VoIP and audio/video streaming are beginning to place increasing demands on residential gateways and the allocation of network bandwidth. These applications have, among other requirements, specific bandwidth needs or latencies that must be met in order to provide a reliable and pleasurable user experience. QoS in the larger sense is being used to help meet these needs.

  The RG sits between the LAN devices and the WAN and between different LAN segments (such as wireless and wired Ethernet). In fact, some may even view the RG as the focal point for networking in the home. As such, it needs to be able to support the following for QoS:

  1. Have the ability to be able to predictably classify particular data flows and assign QoS treatment for them.
  2. Have support for MAC layer QoS schemes such as 802.11e/802.1q.
  3. Have support for various Layer 3 QoS schemes and packet tagging such as ToS and DiffServ.
  4. The ability to mark classified data flows based on policy settings.
  5. Have the ability to transition Layer 3 QoS to MAC layer settings or convert one MAC layer setting to another (e.g. 802.11e to 802.1q and vice versa).
  6. Have the ability to police an interface to keep it within its contractual limits.

- Network Management

  Management for RGs is becoming more diverse and more complex. Much of this complexity exists, not so much from the protocols themselves, but from the diversity of management schemes and protocols that exist around the world and the levels that they interact with at the RG. (For example: CableLabs®, DSL Home™, UPnP, SNMP, Web-Based, CLI and so forth.) This diversity demands a flexible architecture that can be adapted to a wide range of needs.

  Such architecture is shown in Figure 6. As can be seen, management of an RG can involve three distinct sources, any of which may be present at the same time. There can be provisioning of the RG by the customer, using the RG's Web interface. This is the most common method of management today. The second form of management comes from the devices within the home network itself. These devices can use the newer standards, such as UPnP and the Digital Living Network Alliance (DLNA), to gain access to the wide area network through the RG and to set up QoS for incoming and outgoing content. The third source of management today comes from the service providers themselves. This may involve the use of standards such as those by the DSL Forum, or they may involve the use of service provider specific requirements. By providing a scaleable architecture, the RG can adapt to the requirements of today's market and needs by SPs without sacrificing modularity, footprint and time to market.

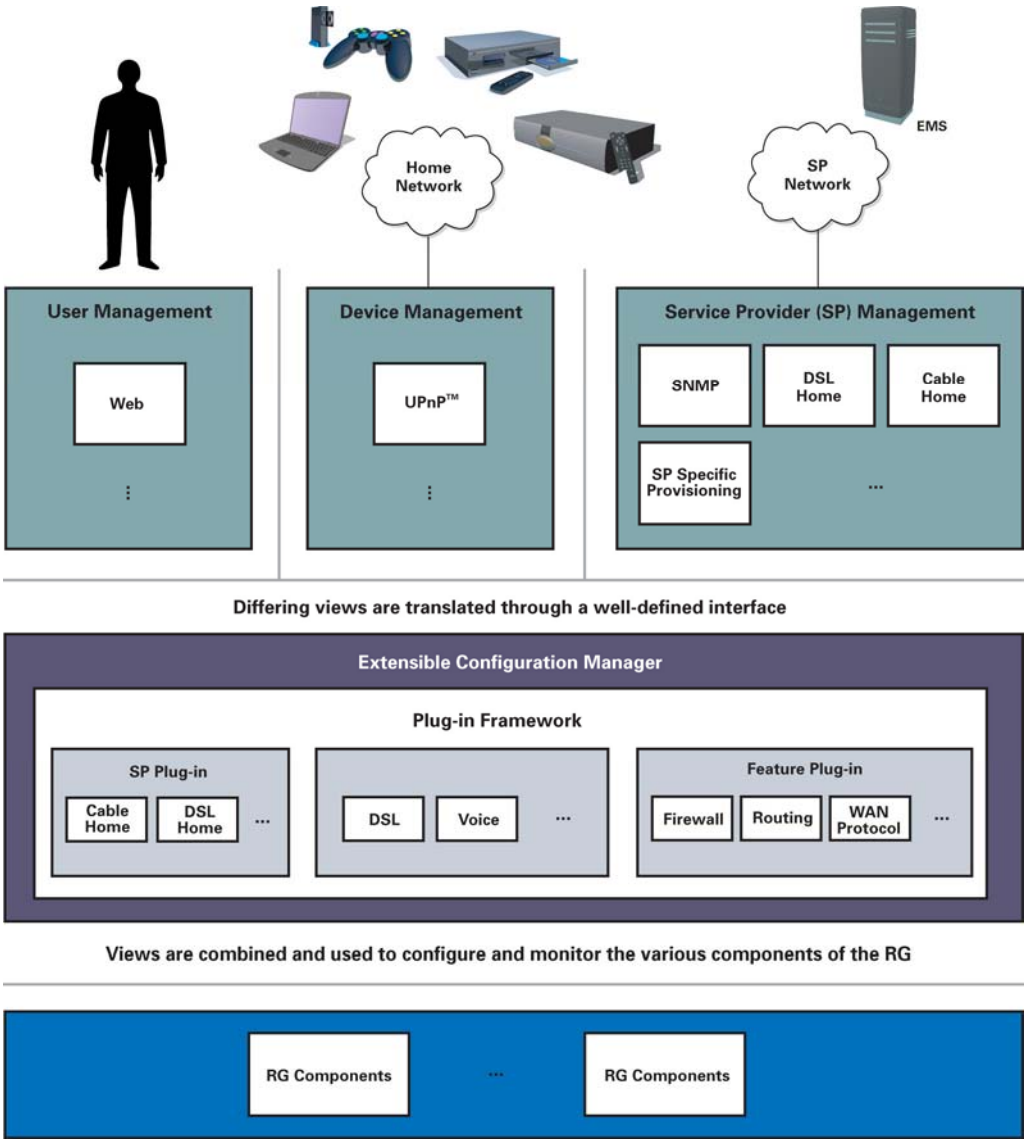

Figure 6: Flexible Management Architecture

Policy Management

Much of what is done on an RG comes down to protocols, ports and IP addresses. To make the configuration of the RG simpler, policy management is required. This allows SPs and users to select operations based on a more understandable identifier (my XYZ Instant Messenger) and then to set the proper operations for NAT, the firewall and even QoS.

RG Software Implementation Challenges

While many RG software solutions are based on open source code such as the Linux kernel and open source networking software, there are still many software implementation challenges for the RG solution provider, including:

- System performance
- Interoperability
- Product hardening and reliability
- Network management
- Extensibility

System performance is a key implementation concern. The primary purpose of an RG is to reliably move data packets between the broadband access interface and the LAN interfaces connected to the RG. It is extremely important to provide acceptable throughput in terms of data packets per second. Today's RGs perform routing, NAT and firewall functions, thus increasing the per packet processing required. Furthermore, broadband access speeds are continuing to increase from ADSL1 (8 Mbps) to ADSL2+ (24+ Mbps), and the LAN interfaces are also getting faster: 10 Mbps Ethernet going to 100 Mbps and even 1 Gbps; WLAN going from 11 Mbps to 54 Mbps and even 100+ Mbps. It is important to do so in a cost-effective manner, while meeting these ever-increasing packet processing requirements.

Not only is the RG expected to meet increasing throughput demands, but the system architecture will also need to address application concerns that demand the RG be able to process simultaneously different types of flows. These different types of flows could be a data stream over wireless through the RG's AP with a video stream through the broadband interface to a PVR on the LAN and a VoIP call all occurring at the same time. The ability of the system to be able to schedule the various software components to process these flows, mark them and prioritize them for QoS and to do so without affecting their quality is critical to deploying a system that has a good customer experience and functions as intended.

This must all be done with respect to the cost of FLASH memory and RAM. Increasing functionality and throughput demands require larger program and buffer memory, but these add to the hardware cost. Memory reduction techniques and system level optimizations are required to keep the software size to an acceptable footprint while still providing headroom for new features.

Interoperability is also extremely important. As broadband becomes mainstream, there are less homogenous environments. While standards bodies work to specify interoperability, there is always some degree of ambiguity in specifications and implementations that must be flushed out through extensive testing. Carriers and service providers demand interoperability. Consumers expect their electronic purchases to work. Interoperability is required for both networking interfaces, e.g. DSL operability

with the central office DSLAM, as well as service provider network management systems for provisioning and remote monitoring.

Due to SOC integration, the burden of interoperability testing has shifted to the semiconductor provider. This requires large investments in system test labs that perform quality assurance testing, interoperability testing, performance testing and certification of the complete solution (software, silicon, hardware reference platform and documentation) prior to delivery to customers. Interoperability requirements and issues are greatly a function of the geographical location where the equipment is installed as well as the service provider's specific requirements.

Product reliability and hardening is essential to successful service deployments. Software must be tested extensively for robustness: no crashes, memory leaks, reliable network management including software download capability for future remote upgrades, consistent APIs, etc. Cutover of new features must be carefully tested for internal software compatibility as well as memory footprint.

Manageability is extremely important for successful service deployment. In particular, it is important to efficiently support the following capabilities:

- Configuration including auto provisioning
- Customer support through remote diagnostic
- Ability to field upgrade software to offer new features and to fix bugs
- Ability to provide the home user with an intuitive user interface that hides the complexities of configuring a router

The software must be extensible to make it easy to add new interfaces and value-added services. This requires a modular software architecture with well-defined APIs and complete documentation. Abstraction of the hardware, RTOS and management interface (data maintained independent of management protocol, i.e. SNMP, Web server, command line, etc.) is also important.

Future Trends and Implications

SOC integration will continue to enable higher speeds and more interfaces that software will be required to support. Newer standards such as ADSL2+, VDSL2, channel bonding, etc. will greatly increase the number of packets per second that must be processed.

Data services are migrating from best-effort data to multimedia services including video distribution and video conferencing which will require end-to-end QoS. Peer-to-peer traffic will require more symmetrical bandwidth processing. Security concerns will require greater dependence on stateful firewall processing, packet inspection and filtering as well as data encryption.

Broadband access has moved from multiple PC sharing to connecting many types of devices in the home for audio and video streaming, home automation control, network gaming, etc. These devices will need proxy network management and plug-and-play capabilities enabled by the RG. Reliability of the network and RG will become increasingly important as people become highly dependent on these devices in their everyday life.

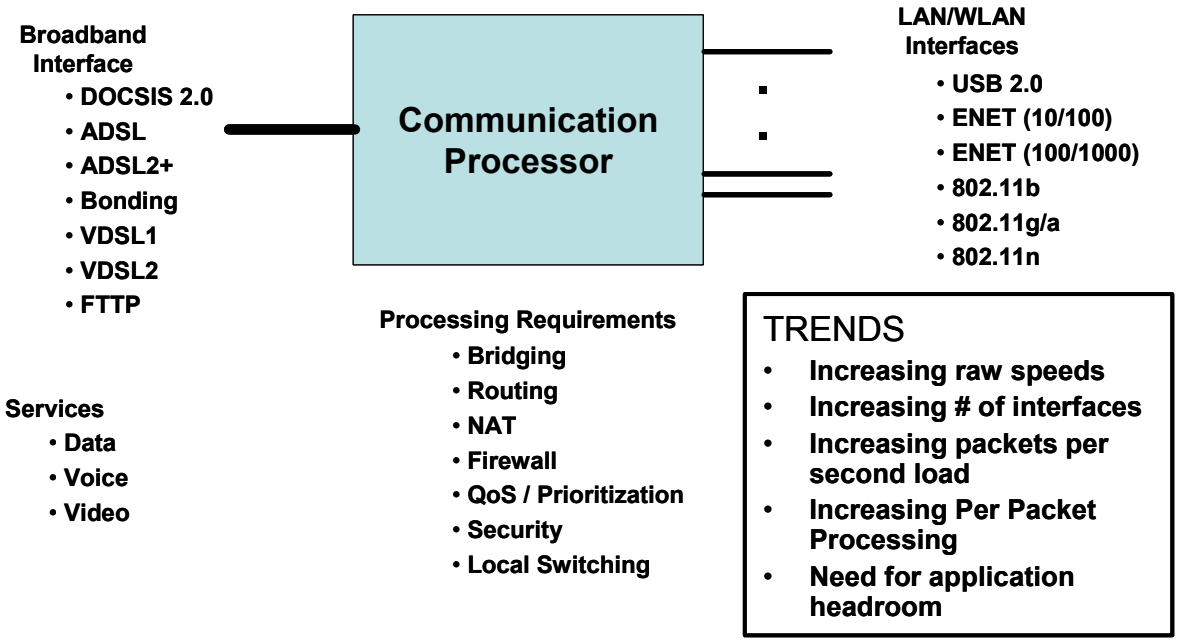

Figure 7: RG Future Trends

Summary

In summary, today's broadband solutions have evolved from simple modems to highly complex RGs capable of supporting a multitude of functions. At the heart of the RG is the software which implements these functions. It is extremely important to have a solid software architecture that is modular and extendable while providing performance, reliability and manageability.

About the Authors

Jack Manbeck is the senior broadband architect and works in the office of the chief technical officer for the Broadband Communications Group of Texas Instruments, where he is involved with the development of broadband access and home networking solutions including DSL, cable modem, WLAN and VoIP. Prior to working for the CTO, Mr. Manbeck served as the director for the Software Platform Technology Center (SPTC), which he helped to create. Prior to TI, Mr. Manbeck worked for several consumer product companies including US Robotics and 3Com, where he was the manager in charge of broadband modem and router development. Mr. Manbeck received his BS in computer science from Old Dominion University and holds a number of patents.

Bill Witowsky is the chief technical officer for the Broadband Communications Group of Texas Instruments, where he is involved with the development of broadband access and home networking solutions including DSL, cable modem, WLAN and VoIP. He was named a TI Senior Fellow in July 2003. Prior to TI, Mr. Witowsky was a cofounder of Telogy Networks, a company pioneering Voice over Packet technologies, where he served as the senior vice president of engineering and chief technical officer. Mr. Witowsky was also a director of engineering at Hughes Network Systems, where he was involved in the development of packet switching equipment and satellite communications systems. Mr. Witowsky received his BS in electrical engineering from Stevens Institute of Technology and his MS in computer science from Johns Hopkins. Mr. Witowsky holds a number of patents and is a member of IEEE and ACM.