Ubiquitous Home-Based Services

Jean-Charles Grégoire
INRS-EMT
Montreal, CANADA
gregoire@emt.inrs.ca

***Abstract*—Remotely accessing services or content based at home is increasingly required as high speed wireless networks become more widespread and mobile terminals more capable. Still, providing such access in reliable and secure fashion presents challenges, especially since media is involved. We explore here how this can be done in a SIP-based framework, taking into account more recent developments in media architectures such as IMS, from extension to home-based (DSL, cable) access to new means of exchanging information between end users through messaging. We demonstrate how MSRP is used to that effect.**

Keywords - SIP; SDP; MSRP; Home monitoring; Home Services.

I. Introduction

Cheap, ubiquitous Internet access, from hotspots to greater affordability of wireless (3G and beyond) services, means that users can be connected to the Internet in almost continuous fashion. This however does not translate into universal access to services as specific, remote access terminals (e.g. RIM's blackberry) remain the norm. We thus tend to see the creation of mobile-device specific variants of common services, or services created specifically for mobility (again, RIM's service)[2], [6]. Even for newer devices (e.g. the iPhone), there tends to be a distinction between a hotspot-based use and a cellular-based use.

We can argue that there are really two different markets at play here, one based on the mobile terminal, "always" connected, the other that of the mobile computer, served by hotspots in a context such that, for the user, connectivity is indistinguishable from the home network, at least as long as massive data transfers are not involved. In the case of the mobile terminals, the restrictions in the nature of the service which can be accessed are manifest: while some are infrastructure-based, most applications are essentially terminal-based, with simple client-server behaviour, and activated on demand or periodically, typically the "app" market for new devices. In either case, user to user (IP–based) communications are elusive.

The emergence of middleware for mobile services, such as the IP Multimedia Subsystem (IMS) puts another twist on this issue, as they allow the creation of new services with proper mechanisms to overcome restrictions that mobility and/or restricted bandwidth access can impose. These operator-based services can come in competition with Internet-based services and this is actually a topic of some controversy, although this is not our focus here.

Both models, Internet-based service specific or operator-based middleware multi–service, present restrictions in the delivery of services. In the first case, we depend on a silo model, where only services deployed on Internet servers are available, with little—or proprietary—means for extension. While this model serves some applications such as social networks or personal communications rather well, it has clear limits in terms of integration (e.g., [3], [9]. The middleware-based model is more flexible in that respect, but users themselves usually have no possibility to provide personal extensions. In both cases, access to personal information is quite limited, restricted to repositories, or confined within applications (e.g. pictures).

Our focus here is on providing access to home-based services or information from remote terminals, as well as allowing home-based applications to communicate remotely with owners, in a secure way, where both parties can mutually authenticate and protect their communications.

In this paper, we show how current SIP–related features actually provide most of the required support for such services, with minimal extra effort. Such an approach has advantages over network-based services as it can more easily enable direct (user to user) communications. It also avoids holding information in the network for the user, which may have security and legal ramifications. Finally, it also allows us to take advantage of established mechanisms to bypass devices which restrict communications.

In section II, we start with an overview of the different elements upon which our argument is built. Section III presents our view of home-based services. Section IV discusses all the issues which need to be resolved. Section V illustrates how messaging mechanisms can be used to transport various forms of date. Our solution is discussed in Section VI and we draw our conclusions in Section VII.

II. Background

In this section we present the key elements required to understand the foundations of our work. We assume that the reader will be familiar with most of the technological underpinnings and we keep this presentation succinct.

A. Home Monitoring Services

Remote access to home services from a wireless terminal is hardly a new concept. For example, we find in [6] a description of the use of off–the–shelf protocols and programming tools to implement alarm monitoring. More recently, wireless operators

 ISBN: 978-1-61208-133-5

have started to offer such services, again centred around monitoring and alarms. In AT&T's case[2], for example, the application was proprietary and required users to deploy specific hardware, which included a remotely operable video camera and various sensors. Motion, door and window activity, water leakage, and temperature changes are cited as common examples.

Building a home sensor network is certainly no longer a challenge, and it is also straightforward to program alarms based on monitored values. The issue is rather the interconnection of this network — or a home-based driving application — with the remote user. In AT&T's case, the application had a web interface and the user had to connect to the server remotely via IP to access the services, essentially enabling access to a web server from any terminal, including cellular phones with such a capacity.

However, while conceptually straightforward, remote access to home–based servers is blocked by many operators, and IP addresses may change through time. Furthermore such a form of remote–access is open to various forms of attacks, as typically befalls web servers.

B. Other services

While sensors/actuators and video surveillance are the most often cited examples of home applications, there are many other possibilities we can imagine, such as access to various forms of content, including audio and video, or pictures. Such access can take different forms, as we shall see later. Accessing content directly from the home is important to alleviate such issues as protecting copyrighted, personal or sensitive information.

C. SIP & SDP

The Session Initiation Protocol (SIP)[10] is the foundation of media services. SIP is a signalling protocol which supports negotiation of parameters for the establishment of an end-to-end session for multimedia communications. The Session Description Protocol (SDP)[7] is used to present parameters.

While SIP was originally proposed for multimedia services, we must take notice that it resolves many issues that arise in home connectivity and enriched, interactive end-to-end communications. The challenge is to identify whether it offers all the flexibility we need for home services and, in the next section, we clarify our expectations in that respect.

III. Home Services

There is no single definition of what home services can be, so we must define what we mean in this context. We have seen earlier examples of monitoring, alarms and surveillance. We broaden this definition with entertainment. We must insist here that we focus on remote services, namely services which must be accessible (but not exclusively) remotely.

Figure 1 presents a schematics view of home services and their connection to the outside world. We consider a network for home devices, with possibly separate dedicated networks for sensors based on proximity technology such as variants of 802.15 (Bluetooth, Zygbee). Communications with the Internet go through a gateway device, which acts as an SIP User Equipment (UE); this device would integrate other functions described below.

Note also that we can have internal home communications as well as communications between the home and an external user. Home communications can be device to device, device to person or person to device. These communications need not be SIP-based, and can be supported through proprietary means. We shall come back to this issue later.

A. Remote services

Home Monitoring: Monitoring is a classical example of remote home automation. This includes remotely receiving alarms notification, reading sensor, setting actuators but also possibly reading documents, such as a shopping list of a family memo and receiving a video stream.

From an Internet-based service perspective, such services do not present many challenges. Access and security are the key issues, but the functionality required to manipulate sensors and actuators and the network resources required are readily available.

Home Entertainment: We mean here access to media sources, such as music and video, from a home server, not unlike what is achievable through Apple's iTunes software in a LAN.

Such an offering is more challenging. We need to be able to browse directories and activate transmission of a specific content. It may be necessary to choose a suitable codec—or suitable parameters/profile—for the medium. Depending on the quality desired, as well as the degree of interactivity required, bounded bandwidth and delay constraints may exist.

B. Some support

We require to make some assumptions about support functions for these services.

Connectivity: We assume that all services are supported by a home IP network, wired and/or wireless. Monitoring devices on a wireless sensor network could be accessible indirectly, i.e. through a control centre which itself would be part of a home network.

Access: For uniformity, we suppose that internal/external access to services is organized through a home-based portal. It receives requests and redirects them to the appropriate device and answers back to the query device. It must also keep track of whether requests are internal—within the home, or external. In the latter case, it would also have to act as a relay for media communications.

Presence: Because alarms are to be sent unrequested, it is important to know whether the user is inside or outside the home to notify her with the suitable means. The portal must therefore also register presence information for the user and forward requests accordingly. Our assumption is that, unless the user is registered internally, the portal will attempt to reach her externally. In any case, all events will always be logged and the logs available for consulting.

 ISBN: 978-1-61208-133-5

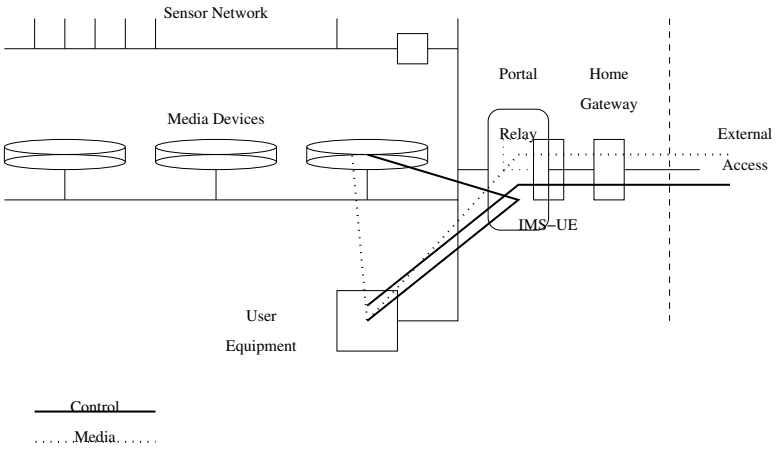

Figure 1. Home Services

Inter-Networking: To allow external access, the portal must be reachable from the outside network. As we have seen in examples above, this can sound like a simple statement, but there are practical limits: network connectivity is one, as is security and possibly quality of service.

Networking restrictions can be imposed by the operator or also the networking equipment used to connect the home network to the Internet: use of private addresses or firewall blocking impose typical limits. Networking devices will typically allow traffic to originate from the home network towards the Internet, but block incoming requests for connection.

When private addresses are used translation devices (i.e. NAT boxes) may impose restrictions on application traffic. It is necessary to translate private to public addresses, including communication ports. If we have SIP traffic, for example, this requires that its SDP content be modified.

IV. Issues

The simple challenge we are confronted with is to provide access to home services remotely. Some would argue that it could easily be done directly, in a typical Internet end–to–end (e2e) model, but there are many practical restrictions to such a model and we propose that there are benefits to take advantage of the access to the IMS infrastructure and its features. Most important, network-based support is required to circumvent restrictions imposed by the presence of middleboxes, which we have mentioned before but revisit below in closer relationship with SIP/SDP. Beyond transport-level connectivity, we must also consider user to user connectivity, i.e. that either home and user can initiate communications at any time.

A. Middleboxes

Middleboxes are network devices which impair communications in some way, either for security reasons, such as firewalls, or for address reuse, such as NATs. Each create specific problems. In the first case, TCP connection establishment can be blocked in one direction and authorized in the other. Still, once established, traffic can freely flow on both directions although this may require modification of signalling content—in our case SDP bodies.

Extensions to SDP provide support to help alleviate the problem; they are specific annotations in SDP bodies which are read and possibly manipulated by middleboxes. For TCP transport, it is possible to set an attribute (*a=setup:*) with the values of *active*, *passive*, *actpass* or *holdconn*. These values announce whether or not the end point can set up the connection or not, does not care one way or the other (actpass), or whether the establishment should be suspended for the time being.

Another attribute, *a=connection:*, allows to specify whether a new connection must be established or an existing one can be reused. It supports modifying the parameters of an established connection without having to tear down and re-establish a new TCP session.

We must note that the protocol does not support the establishment of several TCP connections for the same medium. On the other hand, the secured form of TCP, TLS, is also available for transport.

B. Presence & Reachability

Alarms are sent from the home to the user and the user can contact her home to access sensor status and media. This requires that:

- User and home must have names well-known to each other,
- The home knows whether the user is “present” in the network and,
- Both user and home can initiate connections, which implies that,
- Both user and home can access each other’s address.

Names are important because home and user need to be able to reach each other, i.e. initiate data transfer. This is done trivially if both are customers of the same service network, but generalized with URIs. Presence should also indicate whether the user is reachable at all.

 ISBN: 978-1-61208-133-5

C. Relays

Middlebox traversal can be sufficient to achieve user to network communications, but may not be sufficient to achieve end to end connectivity, e.g. if both end users are behind firewalls. In this case, application relays in the network have to be used. Such architecture is commonly used by services such as Skype[12] and are also fundamental to the architecture of the Asterisk[11] soft PBX.

The use of such relays raise several issues of security. They require proper authentication, protection against hijacking or DoS. Note that there is also a chicken and egg issue at work here: To enable a relay, there must be a way to discover it to force its presence on exchanges. This can be done through a separate discovery process, or through registration mechanisms *à la* SIP: either communications are permanently enabled between both ends, or an enabling signalling channel is established which allows to negotiate and setup proper connections.

The relay may provide added value to the communication. Minimally, it can be buffering and flow control, in case of mismatched performance in the links. Media conversion (transcoding) can also be performed.

D. Information transmission

The remaining issue is the transmission of information from end to end. This includes:

- Commands and values for sensors and alarms;
- Menu, menu selection;
- A/V streaming and streaming control, e.g. play/pause.

A protocol is therefore required to carry this information.

V. MSRP

The Message Session Relay Protocol (MSRP[4]) is a protocol to support session-oriented instant messaging. It is text-based, connection-oriented and supports exchange of arbitrary (binary) MIME-encoded content. Unlike SIP's page–mode messages, MSRP allows messages of any length and structure.

Unlike other messaging protocols, MSRP is integrated with SIP and its offer-answer mechanism, and thus blends naturally into IMS. Note here that we have three protocols present in MSRP exchanges:

- SIP carries the information required to negotiate the exchange between endpoints, possibly through relays;
- SDP is used to capture this information, including data format, ports, transport used, etc.;
- MSRP formats the IM messages, supports chunks, fragmentation, success reports, etc.

The specific use of SDP and MSRP is illustrated below.

A. Basic MSRP Operations

The following example, borrowed from [4], illustrates key elements of the use of MSRP; it is a typical first step in a SIP transaction between Alice and Bob.

```
INVITE sip:bob@biloxi.example.com SIP/2.0
To: <sip:bob@biloxi.example.com>
From: <sip:alice@atlanta.example.com>;tag=786
Call-ID: 3413an89KU
Content-Type: application/sdp

c=IN IP4 atlanta.example.com
m=message 7654 TCP/MSRP *
a=accept-types:text/plain
a=path:
  msrp://atlanta.example.com:7654/jshA7weztas;tcp
```

The c field sets the address (Internet, IPv4) of the source point. The m field specifies an IM protocol, based on MSRP, and the port used for communications. The a fields contain MSRP-specific information, including encoding supported. The presence of "path" information in mandatory.

The field values "TCP/MSRP" and "TCP/TLS/MSRP" have been added to the SDP protocol for explicit support of MSRP. They support two forms of transport for MSRP content, one plain TCP the other one encrypted.

Note that, with MRSP and unlike other use of SDP, the attributes—and more specifically the a=path attributes—rather than the information contained on the c and m lines are to be used to determine where to connect. Also note that a TCP connection can be used for several different transfers.

Bob's answer could be the following:

```
SIP/2.0 200 OK
To: <sip:bob@biloxi.example.com>;tag=087js
 From: <sip:alice@atlanta.example.com>;tag=786
Call-ID: 3413an89KU
Content-Type: application/sdp

c=IN IP4 biloxi.example.com
m=message 12763 TCP/MSRP *
a=accept-types:text/plain
a=path:msrp://biloxi.example.com:12763/
   kjhd37s2s20w2a;tcp
```

The answer contains Bob's contact information which matches Alice's, i.e. IP address (or name) and port, together with protocol.

And Alice's final answer:

```
ACK sip:bob@biloxi SIP/2.0
To: <sip:bob@biloxi.example.com>;tag=087js
From: <sip:alice@atlanta.example.com>;tag=786
Call-ID: 3413an89KU
```

We see that this exchange follows the normal SIP 3–way handshake of INVITE, OK and ACK. After this, both Alice and Bob can open a TCP connection and exchange MSRP messages over it. MSRP has SEND methods and acknowledgement. The SEND method supports sending fragments of large messages. It is also possible to specify the nature of the content of the message.

```
MSRP a786hjs2 SEND
To-Path: msrp://biloxi.example.com:12763/
   kjhd37s2s20w2a;tcp
From-Path: msrp://atlanta.example.com:7654/
   jshA7weztas;tcp
Message-ID: 87652491
Byte-Range: 1-25/25
Content-Type: text/plain

Hey Bob, are you there?
```

 ISBN: 978-1-61208-133-5

```
-------a786hjs2$
```

All messages sent are acknowledged with a copy of the transaction identifiers present in the message header, as well as a copy of information present in the SDP body: to-path and from-path.

MSRP has several provisions for reporting on message sent. It is possible to request in the header whether or not a report should be sent in situations of success or failure. Reports use the message ID to differentiate between different transmissions.

B. URIs, Paths and Relays

MSRP endpoints are identified by URIs, with am msrp (or msrps, when carried by TLS) prefix, as seen in the example above. From-Path, To-Path fields in MSRP contain sequences of URIs, which are relays to the final destination. Beyond the protocol used, URIs have features we are accustomed to from other uses of SIP. Rather than a fully qualified name, it is also possible to use IP addresses.

An endpoint that uses one or more relays will indicate that by putting a URI for each device in the relay chain into the SDP path attribute. The final entry will point to the endpoint itself. The other entries will indicate each proposed relay, in order.

Since both ends of communications can be isolated behind security devices, it may be necessary to communicate through relays, not unlike what is done for SIP. In our specific case, we would consider the use of a single relay. In the following section, we see how it can be inserted in the communication, and its practical benefits.

VI. Discussion

We propose that both a home user agent and the remote user are both customers of the same IMS infrastructure. End to end communication establishment is done by the basic mechanisms of IMS. Both parties know each other's name and correct authentication is guaranteed by IMS. e2e signalling is thus quite trivially established between parties. The issues remaining are the transparency of the home services (for the IMS infrastructure) and the support for information exchange.

Services: Home services and their nature are essentially transparent for the IMS operator: media exchanges can be no different from typical usage, while notifications, menus and operations are embedded in MSRP messages and encoded in, say, XML, in a simple command–parameter format. The following example shows a sequence of sensor information.

```
<?xml version="1.0" encoding="ISO-8859-1"?>
  <status date=31/01/2009>
    <sensor>
      <name code="1">Kitchen</name>
      <value>empty</value>
    </sensor>
    <sensor>
      <name code="2">Living Room</name>
      <value>empty</value>
    </sensor>
      ...
  </status>
```

A generic application can associate operations and (GUI) presentation based on XML documents exchanged through a SIP UA. This application has a home and a remote flavour. At home, it interacts with devices and with the user either through the UA or through a local menu. On the remote terminal, it is only interacting with the user.

We are not investigating the application any further here as it presents no specific challenge.

Communications: The main hurdle we face is the possible presence of devices restricting the establishment of communications in one direction. While IMS-related standards[5] are designed to circumvent such restrictions, we may still require that a relay be present in the network; this relay acts as a back to back user agent (B2BUA), typical in SIP architectures. Note that this relay has two dimensions: signalling, and media. IMS is structured in such a way that a signalling relay is not necessary, beyond what is supplied by the CSCF. Yet in some circumstances, the use of a B2BUA has been mandated (e.g. [1]).

Media is a more critical issue, especially when TCP is used for transport, which is also why MSRP relays [8] were created. Typical UDP-based SIP communications are initiated from the user to the network, with the first REGISTER message, which would be allowed to traverse NATs and firewalls and set the path for future SIP exchanges. TCP connections must be initiated from one side only. Our alternative is either to use a B2BUA, or simply an MSRP relay.

The relay issue is important for another reason: the provider must not hold any personal information for the user, unlike typical Internet "service in the cloud" models, beyond subscription information. We must therefore exclude architectures where a storage server would act as a temporary repository. Note however that communications between home and relay, and user and relay can be encrypted, but other solutions must be found if strict end-to-end confidentiality is required.

We propose that a B2BUA would be required for all communications, i.e., media and data. While some forms of communication could be authorized by middleboxes and not others, it is simpler to use a single connectivity model for all.

Relay Discovery: An issue with the B2BUA is to 1) decide whether or not it is necessary and 2) discover its location.

For the first problem, it is simpler to impose its systematic use, as we have just discussed. For the second one, S-CSCF filters must be used to route the call through the B2BUA. This can simply be done by assigning homes a special class of URIs, and recognizing a communication between the user and the home.

Configuration: Home User Agent and Remote User must share some information for proper inter-operation, such as list of known devices/sensors and other supported media services, e.g. audio, video or pictures. We would typically create a remote configuration based on that of the home application and transfer it to the remote terminal.

 ISBN: 978-1-61208-133-5

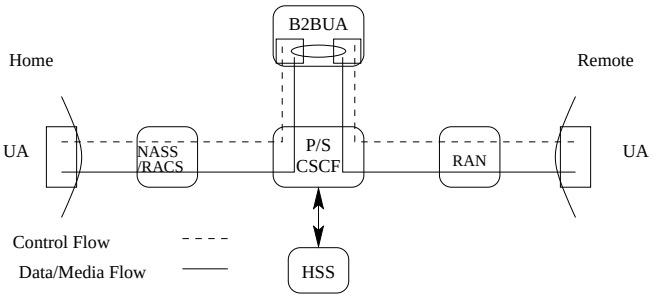

Figure 2. e2e View

They would also obviously know each other's URI and could also share keys for mutual authentication and message encryption.

All these issues are beyond the IMS infrastructure's influence, however.

Integration: Figure 2 illustrates an example of end-to-end connectivity. The home is connected to the IMS infrastructure via a Network Attachment Subsystem (NASS), associated with a Resource and Administration Control Subsystem (RACS) which can perform network security operations. At the other extremity, the user would have a mobile terminal exploiting a radio access network (RAN).

Home-User SIP sessions are switched by the CSCF function towards the B2BUA which bridges requests and connects data/media flows. As we have explained above, the use of the B2BUA can be transparent to the users and inserted in the signalling path through the S-CSCF filters.

The functionality required of the B2BUA for the data/media path is minimal, and content dependent. Audio/Video codecs are negotiated end-to-end and media frames, carried over UDP, need only be relayed towards their destination.

MSRP data is carried over TCP and presents a different problem. While it would be possible to collect TCP segments and relay them directly, it is more appropriate to collect well-formed messages and forward them, as would an MSRP relay. Again, it is possible to use encryption to keep message content private if necessary.

Overall, we see that the infrastructure we need for our communications is well within the IMS model. Since filtering is involved, the participation of the IMS operator is required, although we should put a caveat there: all IMS services (A/V communications, Messaging) are straightforward, except that operator support is required to overcome networking restrictions imposed in some domains. While we can imagine that, in some circumstances, offered IMS services could be integrated into a suitable application, it may also well be the case that a B2BUA would have to be deployed in the operator's network, with a matching service offering. Considerations for a suitable business model are beyond the scope of this paper, however.

VII. Conclusions

We have shown a SIP-based model to support home-based services and how it is possible to use an IMS infrastructure to deploy such basic tools. Beyond established A/V services, the use of MSRP, for data exchange, combined with a B2BUA in the operator's network are sufficient to allow the user to safely exchange information between home and remote locations. The application itself can be designed independently, for example on an XML basis, while benefiting from IMS' services. The scheme proposed is overall rather straightforward and would support applications of various degree of complexity.

Further work is required to study how to support streaming more efficiently, or closer to an Internet model, since we have here IMS' interactive model. We believe this would require special support in a network B2BUA.

Finally, we should be able to bridge the gap between home-internal and home-external communications, if only to be able to transparently reuse the same devices. This is also the focus of further investigations.

References

[1] Open Mobile Alliance. Instant message using simple. [OMA-TS-SIMPLE_IM-V1_0-20080312-D, Sep. 2008.
[2] AT&T. AT&T launches remote home monitoring video service nationwide. last visited March 15th, 2011, http://www.att.com/gen/press-room?pid=4800&cdvn=news&newsarticleid=23003.
[3] Paula Bernier. Verizon tests home monitoring and control service. Last visited March 15th, 2011, http://www.techzone360.com/topics/techzone/articles/134094-verizon-tests-home-monitoring-control-service.htm, January 2011.
[4] B. Campbell, R. Mahy, and C. Jennings. The Message Session Relay Protocol (MSRP). Request for Comments (RFC) 4975, September 2007.
[5] ETSI. Telecommunications and internet converged services and protocols for advanced networking (tispan); resource and admission control sub-system (racs): Functional architecture. ETSI Standard 282 003 B3.2.0, Nov. 2008.
[6] David Fox. Home monitor on a cell phone, o'reilly, last visited march 15th, 2011.
[7] M. Handley, V. Jacobson, and C. Perkins. SDP: Session Description Protocol. Request For Comment (RFC) 4566, July 2006.
[8] C. Jennings, R. Mahy, and A.B. Roach. Relay extensions for the message session relay protocol (msrp). Request for Comments (RFC) 4976, September 2007.
[9] Meye. A ground-breaking home-monitoring service. Last visited March 15th, 2011, http://www.meye.com.my/.
[10] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks M. Handley, and E. Schooler. SIP: Session Initiation Protocol. Request For Comment (RFC) 3261, June 2002.
[11] www.digium.com. Asterisk is a registered trademark of digium inc.
[12] www.skype.com. Skype is a registered trademark of skype ltd,.

 ISBN: 978-1-61208-133-5