GPRS Facts for the Internet Application Developer – Part I

by John Pagonis, July 2003

Modern GSM networks, now offer 'always-on' packet-based Internet traffic on mobile devices. Such GSM networks are called 2.5G GSM networks and the service infrastructure and protocols that enable such packet-based Internet access on modern GSM mobile phones is called GPRS; which stands for General Packet Radio Service.

GPRS builds on top of the existing GSM standards and brings evolutionary but widespread changes and features to the GSM network. At this moment GPRS is offered in most GSM networks across the globe.

In this series of articles we will familiarize ourselves with the GPRS network from the perspective of the Internet user and developer, treating the GPRS network as a 'black box' that brings certain IP traffic and services to mobile phones.

Although we do not need to be concerned with the complex internals of the implementation of a GPRS network, we will discuss certain aspects that will help understand what GPRS has to offer and how we can best utilise and interface to it.

APNs, PDP context and IP addressing

Although GPRS gives to the user the impression of 'always-on' Internet connectivity, in reality mobile terminals need to explicitly create data sessions, on demand, to the GPRS part of the GSM network, in order to enable data (as opposed to voice) traffic to flow from and to the mobile phones.

For each such session to a Packet Data Network (PDN), such as the Internet, some sort of session description information needs to exist to enable agreement on the characteristics of any data flow between the GSM/GPRS network and the mobile terminal.

That session description is called the Packet Data Protocol (PDP) context. A PDP context is effectively a record that includes descriptions about the PDN type (e.g. IPv4), the PDN address assigned (e.g. IPv4 address 192.168.116.5), QoS (quality of service) requirements and an Access Point Name (APN) network identifier, through which traffic is to be routed to the PDN (Internet in this case). In most cases today, terminals will provide an empty IP address during such 'PDP attach' requests, which will prompt the network to allocate them an IP address in combination with the requested APN.

A GPRS network can assign to the mobile device a static or dynamic (and in most cases private) IP address.

A mobile terminal can have more than one PDP context active and by that token more than one IP address. Therefore mobile terminals can be multi-homed (i.e. appearing to have more than one network interface).

Moreover a mobile terminal can decide which PDP contexts to activate at any time. This of course can be influenced, by the user or the needs of a particular application. For example if a mobile device is not multi-homed (e.g. phones like the Nokia 3650 based on Symbian OS v6.1, or the Sony Ericsson P800 based on Symbian OS v7.0) and an application needs to access the Internet through an operator's Internet APN, then the application under the control of the user may decide to disconnect from the WAP APN, for example, in order to route traffic to the Internet. Whereas on a multi-homed device (such as forthcoming phones based on Symbian OS v7.0s) the WAP browser could use a different APN to the one used by the Web browser, while both applications receive traffic.

APN resolution and routing

As we said, every time a mobile terminal negotiates a GPRS data session with the network, it needs to specify the APN to use for that session. That APN effectively denotes a route to a PDN and in most cases is specified in terms of a Network Identifier (usually entered in the device by the user or preset by the operator).

That APN Network Identifier which is a string resembling a Domain Name Service URL (e.g. serv.myoperator.com) in fact gets resolved internally in the GPRS network, using operator DNS servers, to what is the IP address of an external network interface of a gateway. That gateway interfaces from the operator's private IP network, to the public Internet (or even some other intranet if set-up that way from the operator). Hence the IP address of its external interface is used by the GPRS network to denote the gateway to terminals, in order to route traffic in and out of the operators' intranet.

The significance of the choice of APN from a terminal's application is that the activated PDP will receive a (possibly) different IP address and routing from the operator's network. For example if for general internet access a terminal is mistakenly set up to use an APN which was specified from the operator to be used for WAP or MMS services only, there will be no routing of traffic (intended to go to the Internet) further than these gateways. That is why a user may see various APNs configured in their device settings.

To the GPRS user the choice of APN effectively makes the choice of route to a network like the Internet, or a service like MMS or WAP.

Most GPRS users will be behind Network Address Translation (NAT) Gateways
Today all GPRS providers operate their networks, for consumers, behind NAT gateways and dynamically allocate private IP addresses to the mobile terminals on each PDP activation. Some operators state that for corporate usage they do allocate static private or public IP addresses without NAT. Therefore unless working in an exceptional situation, application and middle-ware developers should always assume that their software will run on devices that are behind NATed gateways.

The significance of running applications behind NAT is multifaceted and obtrusive many times. Network Address Translation on a gateway typically prevents inbound connections (which is true for dynamic NAT), since there are no externally visible IP addresses. To complicate matters further, there are many different NAT implementations with different application-level knowledge and different binding time-outs for ports and sessions.

This latter property implies that when a GPRS terminal sends a UDP packet to a remote host on the Internet, the NAT gateway may (or may not) allow that host to send back data during a limited period of time; after which the gateway will re-allocate that port (and effectively the session resources) to another GPRS terminal. Some NAT policies allow for the session information (and thus allocated port) to remain associated with the GPRS terminal as is, as long as traffic flows through it. Such inactivity timeouts though, vary between operators quite significantly.

Another problem is that some (traditionally PC-based) applications unfortunately encapsulate in their payload (transmitted over UDP/TCP) their sender address, in their communication to a remote host. When such communication goes through a NAT gateway, that does not have knowledge specific to that application, the UDP/TCP payload stays as is. As a result, the remote host receives a "send from" address (from that UDP/TCP payload as opposed to IP source) that is private to the GPRS operator's address space. This has as an effect that the remote host will not be able to respond back to the sender.

In the case where the NAT gateway has application-level specific knowledge, the payload part which specifies the sender address is transparently remapped, so the remote host may send a reply which again will be quietly remapped from the gateway; without any of the two communicating parties being affected or knowing that such an operation took place. Again such application-specific knowledge from NAT gateways is not consistent across operator's or NAT gateway vendors.

In general developers need to keep in mind that on GPRS their software will run behind NAT. For that matter it is advisable to familiarize oneself with NAT and what constitutes NAT-friendly application protocols and designs. As a rule-of-thumb using TCP (or TCP-based session protocols like HTTP, SyncML and WebDAV) and avoiding the need for any peer connections or source address encapsulation, should be sufficient to avoid any NAT-related problems.

Roaming
Today it is very common for users to utilize roaming between GSM operators across the globe for voice communication. As GPRS usage now picks up, operators can have roaming agreements. Today when a terminal is visiting another operator's GPRS network, that network and terminal can discover and route the terminal's traffic to its home operator's APN; thus the user does not have to change their GPRS phone settings. To achieve that, operators must utilize some kind of inter-operator IP backbone (whether that is physical or virtual).

When roaming and while attached to the visited GPRS network, a terminal can use an APN provided by its home network, or the user may choose the APN provided by the visited network. Effectively the user's choices are limited by the agreement with their operator and how their subscriber records have been set up in the operator's network.

E.g. assuming GPRS when roaming is enabled, the subscriber information includes info about whether:
- the user is allowed to use visited network access points
- the user has a choice to select home or visited APNs
- the user has to use the access points in their home network

For roaming users billing is also an issue which should always be clarified with the operator. Billing information may be exchanged directly between operators or through a clearinghouse and most likely will be a combination of charging tariffs. Usually the cost is high, as most operator agreements make sure that the user is charged by both operators.

Therefore developers have to keep in mind that, under roaming conditions, although users may have GPRS connectivity, it will most likely be expensive for users to make use of GPRS, if it is at all available. In the case of some consumer applications that may not be crucial, but in many corporate scenarios such GPRS connectivity may be vital. Thus developers have to consider bandwidth utilization and efficient protocol and application design (for example data delivery may be prioritized and only high priority data delivered under roaming conditions). At the same time, the user may decide to chose visiting operator APNs instead, which should be considered by applications when connecting to the Internet autonomously.

Keeping in mind…
In general developers need to remember that on GPRS their applications will run behind NAT and that utilizing TCP- and HTTP-based protocols will make things easier. The significance of APNs is effectively in the gateway selection process and in that peer-to-peer application scenarios need to be mediated through a server since incoming connections are not possible in almost any cases.