

Enter Once, Share Everywhere:

User Profile Management in Converged Networks

Arnaud Sahuguet* Rick Hull Daniel Lieuwen Ming Xiong
Bell Laboratories, 600 Mountain Ave., Murray Hill, NJ 07974

Abstract

Interoperation between network types (telephony, wireless, internet, etc.) is becoming increasingly feasible, leading to the so-called “converged network”, and to a broad new family of end-user focused “converged services”, that combine different kinds of network connectivity (voice, text-based instant and email messaging, real-time presence and location information, and access to numerous web-based services). A crucial but still largely unresolved issue in providing convenient end-user access to and control of these services is to provide easy access and sharing of end-user profile data, including data about the user’s devices, services enabled, billing arrangements, address, calendar and preferences.

This paper surveys the kinds of profile data that are currently held in the various networks and where/how that data is stored, identifies key standards group activities working to enable profile data sharing, and proposes a data management framework (GUP^{ster}) that can be used across and within networks and organizations, to facilitate sharing of profile data to support converged services. GUP^{ster} combines ideas from the federated architecture for data integration and the Napster approach for peer-to-peer sharing, in order to satisfy the high level requirements coming from standards bodies such as 3GPP GUP and the Liberty Alliance. The paper describes how the GUP^{ster} framework can be supported, and identifies key topics for future research.

1 Introduction

The mobile telephony revolution has changed the behavior of millions of people. People can now reach and be reached anywhere, anytime. First limited to traditional voice networks (wireless and PSTN¹), this revolution is now spreading across network boundaries to include the internet and other networks, giving rise to so-called *converged networks*. WAP was a first (viewed by many as unsuccessful) way to bridge mobile telephony with the Web. The 2.5G and 3G wireless networks which followed already deliver email connectivity, instant messaging, and web access.

The convergence of networks has quite naturally given birth to a new breed of services – the so called *converged services* – which try to make the most of the network infrastructure in order to deliver value to end users and revenue to operators. These converged services revolve around providing information to end-users, providing connections between end-users with each other and/or with automated systems working on behalf of enterprises, and enabling end-users to invoke commercial, financial or other services outside of the network proper.

Providing these services entails the storage, sharing, and use by the converged network of a broad array of information about end-users, including, *e.g.*, the correspondence between users and their devices; awareness of presence, location, and availability of devices and their users; and awareness of billing models and plans (*e.g.*, pre-paid vs. post-paid). As network-hosted services become more complex, end-users will insist on highly personalized renditions of the services, which will entail the storage and access of large amounts of personal data (*e.g.*, device specific information, usage preferences, address book, buddy lists, calendar, ...) This paper con-

*Contact author: sahuguet@lucent.com

¹Public Switched Telephone Network

siders the data management technologies that will be needed to manage end-user profile information in the context of converged services. In particular, the paper describes how and where profile data is currently stored in the various networks, surveys relevant industry standards groups, identifies key requirements for profile data management, and proposes a framework that fulfills these requirements.

The framework presented here, called **GUP^{ster}**, is based on three central ideas. First is that a standardized schema for (most) user profile information will emerge from the activities of the 3GPP Generic User Profile (GUP) standards body [6]. Second is the adaptation of the peer-to-peer paradigm pioneered by Napster [3], whereby profile data will continue to reside in widely distributed networks and locations, but information about *where* to find the profile data will be easily accessible, e.g., from a centralized meta-data manager analogous to the Napster server. And third is the adaptation of techniques from federated data management [16, 24, 29] to the specialized context of profile data management. A variety of other technologies are also used in the **GUP^{ster}** framework, including technologies around XML, access control, and data replication and synchronization.

The key contribution of this paper is to identify profile management for converged services as a crucial data management problem for the coming years, and to provide a promising framework and structure for attacking this problem. While pointing the way towards a solution, the paper also raises a number of open research questions on how to flesh out the framework.

This paper is focused primarily on the data management aspects arising in profile data for converged services. It is beyond the scope of this paper to provide detailed consideration of other technologies needed to support converged services, such as preference and policy management; workflow and collaborative systems; distributed systems; secure data transmission; standards for sharing information and control across networks and devices; new forms of call models and session management; and continued work on natural language interfaces.

The remainder of this paper is organized as follows. We first consider requirements for profile data management in support of converged services, by describing motivating examples and then listing the high-level requirements (Section 2). We then

overview the current state of the art in terms of data management for user profile in the various domains/network and give an overview of the 3GPP GUP (Generic User Profile) initiative (Section 3). Section 4 presents a high-level description of the **GUP^{ster}** framework, and Section 5 gives variations of the basic framework, more detail, and a discussion of how **GUP^{ster}** fulfills the requirements given in Section 2. We present some related work and open issues (Section 6) before we offer some future work and our conclusion in Section 7.

2 Converged Services: Examples & Requirements

This section overviews the requirements that a profile data management framework for converged services should satisfy. We first present two illustrative examples, and then provide a listing of the essential requirements.

2.1 Example 1: Roaming Profile

Let us consider a corporate employee Alice working for Lucent Technologies. She owns a SprintPCS cell phone that she uses for both business and personal matters. She stores her phone book and phone preferences (*e.g.*, ring tones, speed keys) on the phone. Alice's SprintPCS contract offers WAP capabilities, and she can store WAP bookmarks on her phone. When she travels abroad she also carries a GSM mobile phone operated by Vodafone. The SIM card she puts inside the phone contains her "European" phone book and phone preferences. Alice also owns a personal data assistant. Her personal address book and calendar information are synchronized with her Yahoo! account. Her corporate address book and calendar are securely hosted by Lucent.

As we can see from this scenario, data is spread all over the place, in different networks. With existing technology, Alice's access to data is quite restricted, making it difficult or impossible for her to

- access her corporate calendar when she is traveling in Europe
- share her address and phone book among SprintPCS, Vodafone and Yahoo!
- keep her personal data and preferences if she

decides to switch from SprintPCS to AT&T, (note that this is usually not an issue in Europe because the data is stored on the SIM card that can be transparently exchanged between devices)

A framework for profile data management should enable easy and efficient access, sharing, update, and synchronization of this data.

2.2 Example 2: Selective reach-me

Mobile telephony makes it possible to contact people anytime, anywhere. This reach-me service can be customized in many ways. First, a callee can make (or program) the decision of accepting a call or not. Operators already offer basic call screening capabilities based on caller ID. Second, a callee can define the best way(s) to contact her.

In a converged network situation (see previous example), a user usually has more than one way to be reached. At work for instance, Alice can be reached (1) on her office phone (and voice mail if she is not in), (2) via email, (3) via instant messaging, (4) via VoIP (if she runs a soft phone like MSN messenger), etc. At home, she can be reached on her home phone. Potentially she can be reached anywhere on her cell phone, depending on the coverage. When she is near a WiFi hot-spot she can be reached on her laptop via email, IM, and VoIP. Depending on Alice's location, the choice of the *best* communication medium varies: quality, price, nature of the communication, etc.

The selective reach-me service permits the network to optimally route a call (taken in its broadest sense) to reach Alice. To do so, the service needs to aggregate information for all the networks Alice is in contact with:

- location information, on/off air from the Wireless network
- call status from the PSTN network
- presence information (*e.g.*, IM status, entries in a log web site) from the Internet
- call status from the VoIP network
- calendar information from Yahoo! or from corporate intranet
- devices owned by Alice with their respective capabilities.

Based on all these information (some static, some dynamic; some provided by the network, some provided by Alice), the service should be able to make decisions such as:

- during working hours (9am-6pm), if Alice presence is "available" (verified with IM): call office phone first, and then try soft phone
- from 8-9am and 6-7pm, Alice is commuting: call cell-phone
- on Fridays, Alice is working from home: call home phone

Importantly, the access and processing of the disparate and distributed data must have fast response time, so that a selective reach-me decision can be rendered in just a few seconds.

2.3 A Listing of the Essential Requirements

This subsection lists requirements that should be satisfied by a profile management framework for converged networks.

To set the stage, we note that in the converged network, both profile data and services will be widely distributed. This is illustrated in Figure 1, which shows the world of profile data management from the perspective of a Wireless Service Provider (WSP). In particular, some services will be hosted within the WSP: *e.g.*, Presence and Availability Management (PAM), pre-paid billing services (Pre-Pay). Others will reside outside. Also, different subsets of the profile data will be accessed by different external applications.

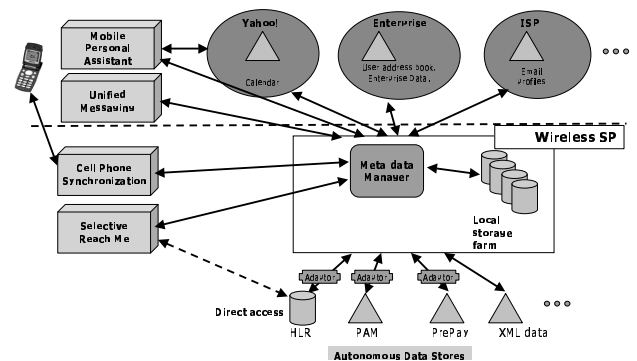


Figure 1: Converged Network: profile and services

We now list the key requirements on a profile management framework. These are based on re-

quirements identified by key standards bodies, and on experience at Bell Labs in providing support for existing and emerging converged services.

1. **Common Data Model:** Profile data should be accessible through a common data model, regardless of where or how it is physically stored or generated. The importance of support for *easy* access to profile data cannot be over-emphasized, given the network operators' intense interest in encouraging 3rd party developers to create the bulk of applications that will use and build upon converged services.
2. **Data Richness:** End-user profile information will become increasingly rich and varied² as the services being supported become increasingly rich and varied. In particular, data may be deeply nested, involve complex integrity constraints, and for each end-user have different statistical characteristics.
3. **Data Transformation:** Data sources on the various networks have radically different structures and data models. Sharing data across these repositories, and also between network and terminal devices, will require efficient data transformation capabilities.
4. **Data Placement:** The placement of data will be dictated by end-user desires (the end-user may not trust some entities to hold profile data) and by optimization needs. In many cases profile data will be stored redundantly (*e.g.*, telephone book may be stored in the end-user's phone, with a "primary" copy held by an internet portal, and with a cached copy held by a wireless service provider, to provide fast synchronization with the end-user's phone). Transformations may be used between redundant copies of data.
5. **Data Integration:** Many converged services will use profile data about a single end-user that is spread across multiple sources. Because the profile data may be distributed in different ways for each end-user, common tools should be provided to perform needed data integration. We expect data integration of profile data to be simpler than in the traditional setting, because profile-related queries do not typically require exotic joins³: most of them are lookup queries like "retrieve presence information for Alice", "retrieve Alice's appointments for today", "retrieve Alice's buddies who are available", etc.
6. **Data Reconciliation:** Profile management must include mechanisms for reconciliation of slightly inconsistent data (*e.g.*, for synchronizing address book on phone with address book in network, or merging of address books from disparate places). End-users should be able to provision the policies used to reconcile profile data.
7. **Data Synchronization:** As indicated above, profile data will be cached, which implies the need for convenient synchronization mechanisms, and triggers to indicate when data has become stale. The need for synchronization is most apparent with handhelds, but it is also needed between different networks and/or enterprises.
8. **Access to Meta-data:** Because profile data will continue to be widely distributed, there must be mechanisms available for discovering where relevant profile data can be found. In particular, the schema of profile data along with an indication, for a given user, of where/how the actual profile can be accessed must be made available to qualified applications. Looking forward, as this kind of meta-data becomes re-ified and widely available, it will be natural to expand on the traditional meta-data representations (*e.g.*, SQL's data definition capabilities, or XML Schema), to include information about data placement, rules for data reconciliation, etc.
9. **Access Control:** The end-user should be in control of what, when, how, and to whom profile data is shared. Mechanisms must be provided so that end-users can specify (possibly intricate) policies about access control (*e.g.*, that presence data is revealed to co-workers only at times when the end-user is "at work").
10. **Security:** Data transmission should be secure; in the traditional wireline network this security was essentially guaranteed because the network was wholly controlled by the service provider; in the converged network encryption techniques will be needed. Authorization

²Just look at the information stored in Palm Pilot.

³For example, here is a query that we would not need to support: "Among all phone users who were on the NJ transit train to Penn Station arriving at 7:13pm, I want the

one who received a call while he was already on the phone: he put the first call on hold for less than 30 seconds before he resumed the call."

mechanisms need to be supported.

11. **Data Provisioning:** As the network-hosted profile data becomes richer and more voluminous for each end-user, it is increasingly important that end-users can provision (*i.e.*, insert, change, or delete) their profile data at will (self-provisioning). The provisioning should be accessible through a variety of interfaces, including large-screen web browser, hand-held wireless device, and even voice.

Provisioning interfaces should be automatically generated and should provide some guarantees (*e.g.*, constraint checking). Users should have the impression of “enter once, use everywhere”, *i.e.*, the distribution and heterogeneity of the actual profile data should be transparent, except in connection with privacy of the data and trust in the organizations holding the data and/or meta-data.

The area of provisioning is often neglected in the data management community, where people assume wrongly that the data is already there.

12. **Reliability:** Telecommunication services must be responsive and simply cannot go down; wire-line telephony is near real-time and 99.999% uptime is the norm. By merging telephony networks (PSTN and wireless) with other networks, users’ expectation is to find the same quality of service with more features.
13. **Scalability and Performance:** These are key issues that need to be addressed. Converged services have to be scalable to support millions of subscribers. Therefore, data has to be partitioned and stored at different sites. Converged services have to be engineered based on the fact that the weakest link(s) will be part of the non-managed networks (*e.g.*, Internet). Many telecommunication services require real-time performance, *e.g.*, in case of call delivery, response time of a transaction has to be within hundreds of milliseconds. Various optimization techniques will have to be used to overcome these limitations (*e.g.*, query optimization, caching, pre-fetching)

3 Profile Management in Converged Networks

In this section, we give an overview of the environment within which a profile management framework must reside. In 3.1 we discuss where profile data is found in the various networks today. In 3.2, we briefly discuss the activities of 3GPP GUP, the primary standards body currently working in the area of profile data.

Perhaps the most important idea in this section is that we expect the 3GPP GUP standards group to be successful in creating and promoting a standardized schema that captures a substantial amount of end-user profile data; this standardized schema will become a key component of any future framework for profile management.

3.1 Overview of Profile Data in the Different Networks

3.1.1 PSTN networks

The Public Switched Telephone Network is the collection of interconnected public telephone networks, based on circuit-switching. In order to avoid a fully-meshed architecture, switches are used to connect together end-points (see Figure 2).




Figure 2: A PSTN network

The role of the switch is to take care of the call control signaling, the media termination and the logic. In current PSTN deployed architectures, the

switch is a multi-purpose box (*e.g.*, Lucent 5ESS or Nortel Meridian, or the emerging softswitches) which encapsulate all these functionalities. User profile information is stored inside the switch itself, which makes it hard to access and extend.

User profile information stored in switches depends on the nature of the services supported by the switch itself: call forwarding number, call barring numbers, caller id flag, 800-number resolution (when the user is a company), etc.

Historically, provisioning of profile data in the PSTN has been quite cumbersome. Most provisioning must be performed manually by network operators rather than the end-user. In some cases (*e.g.*, to set call forwarding numbers) the end-user can self-provision through a phone's keypad, a severely limited interface. Technology is now emerging for providing a web-based interface for self-provisioning of this data, but is not yet widely available.

3.1.2 Wireless networks

An abstract view of a wireless network is presented in Figure 3. Devices are connected to the network by a *Radio Access Network*. Subscriber information is stored by the *Home Location Register* (HLR), maintained by the subscriber home carrier. The HLR contains permanent subscriber information and the relevant temporary information of all subscribers permanently registered in the HLR. For each subscriber, its subscriber profile, location and authentication information are stored in HLR. Subscriber profile contains identity information, telephone numbers, and mobile user preferences related to services, such as call forwarding, barring, roaming, etc.

A Mobile Switching Center (MSC) is a switch (analogous to 5E or softswitch for PSTN) used for call control and processing. The MSC also serves as a gateway to the PSTN network. The MSC interacts with HLR to retrieve the information necessary for switching.

A Visitor Location Register (VLR) maintains temporary subscriber information (snapshot of the master copy stored in the HLR) in order to handle requests from subscribers who are covered by the VLR. When a user moves from one cell (managed by a RAN) to another, a different VLR may be used. The new VLR will send this new location information to the HLR by sending a location up-

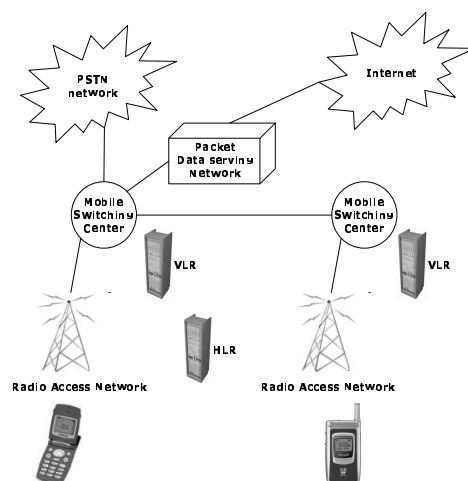


Figure 3: A Wireless Network

date request to the HLR. The HLR will cancel the location information in the old VLR after it receives new location information.

A simplified interaction between the various components can be described as follows. When a user initiates a call, the current location and locally responsible MSC of the callee must be determined. In order to route the call to this MSC, the routing information to the callee must be obtained. The HLR is the only entity in the network that can potentially supply such routing information. Therefore, a subscriber's HLR must be interrogated for each connection setup to the mobile subscriber.

HLR is indispensable in wireless networks. It has to perform database operations for different types of events/messages initiated from wireless networks, for authentication (using AAA⁴ servers), registration, call delivery, short message services, billing, etc. A typical HLR stores information for millions of users in main memory relational databases. Most read-only queries performed by HLR are simple lookup queries (*e.g.*, location of the callee). Updates can be initiated by mobile subscribers (*e.g.*, update transaction when a mobile subscriber's location is changed) or provisioning centers (*e.g.*, when a user changed her call forwarding number).

The next generation of wireless networks (3G) includes both voice and data. User profile information related to data services (*e.g.*, VPN configuration, data access quality of service, service level

⁴Authentication, Authorization, Accounting

agreements, etc.) is becoming more and more important.

We deliberately do not include here the IP-based wireless network based on 802.11 (*a.k.a.* WiFi) hot spots. For most of them no user profile information is stored on the access point. Consolidation efforts lead by companies like Boingo and NetNearU will probably make this change.

3.1.3 Voice-over-IP (VoIP) Networks

VoIP networks – as the name implies – deliver voice using the Internet protocol (IP). The delivery of voice is initiated using a signaling protocol like SIP or H.323. Endpoints are IP phones or softphones (all software) which can talk the signaling protocol and process digitized voice using codecs. VoIP networks can be connected to other networks (mainly PSTN) using gateways that (1) perform the encoding/decoding of voice and (2) translate SIP addresses into phone numbers.




Figure 4: A SIP-based VoIP Network

Unlike PSTN and Wireless networks, VoIP networks are not network-centric from either a call control point of view or a user profile management point of view. This permits a wide range of implementation strategies, which are still being explored by the industry. In general, much of the “intelligence” (and therefore profile information) in VoIP networks is stored at the end-point. For instance, IP phones (soft or hard) usually store address book, call logs, forwarding preferences, etc. Some profile data is found in the network, however. In the case of SIP, network components consist almost exclusively of SIP registrars or SIP proxies. SIP registrars simply store a mapping between a SIP address (a VoIP phone number) and the corresponding IP address of the endpoint. SIP proxies are used for message

routing and may store some user information. In future SIP-based services, proxies will probably make use of user profile information coming from other databases. And moving further in this direction, some emerging HLRs will support SIP, and act as a centralized SIP profile and location management center, in much the same way as it serves for wireless calls [12].

3.1.4 Web

The Internet consists of machines connected via the IP protocols, some of them are servers, some of them are clients. We can distinguish between the public Internet (the Web) and the private Internet (the numerous intranets managed by corporations), usually isolated from the Internet by firewalls.

A broad variety of profile data is stored in the machines that form the internet, including

- bookmarks
- address book
- computer preferences
- calendar information (iCal, vCal, Exchange)
- e-commerce profile data (shipping addresses, wish lists, buying patterns, etc.)
- employee information stored in directory servers
- presence information (e.g. instant messaging client, connection to DHCP servers, etc.)
- web site log files
- edge-router caching and routing policies
- cross network info: ISP info about a user being connected or not and its IP address and calling phone number (in the case of a modem).

Note that solutions to unify web-accessible profile data have been proposed. Netscape *roaming profiles* are based on (1) a schema describing a limited set of user profile information and (2) a client/server protocol (LDAP or HTTP). Roaming profiles make it possible for a user to store her preferences (address book, bookmarks, cookies, browsing history) in a central server and have access to it from any client.

In this section we have seen that user profile information is spread across many machines and organizations on the various networks (see Figure 5).

Network	Locations of Profile Data
PSTN	Class 5 switches, billing systems
Wireless	HLR, VLR, MSC, billing systems
VoIP	end-user device, SIP registrar/proxy, AAA
Web	end-user device, ISP, portal, e-merchant, enterprise, edge-router, ...

Figure 5: Where profile data is stored

3.2 Overview of 3GPP GUP Goals

Several standards bodies and consortiums address the issue of profile data. Some view preference information as a kind of data that should be maintained alongside profile information. We focus here primarily on GUP. Other efforts will be presented as part of related work (Section 6).

3.2.1 3GPP Generic User Profile (GUP)

The 3rd Generation Partnership Project (3GPP) is a collaboration agreement that was established in December 1998. The scope of 3GPP is to produce Technical Specifications and Technical Reports for a 3rd Generation Mobile System based on evolved GSM core networks and the radio access technologies that they support.

As part of its numerous efforts, 3GPP has started the Generic User Profile effort as a way to standardize the management of user profile information. Because of the enormous number of wireless users and the strong interest that the wireless sector has in continuing to host, or at least to provide access to, end-users subscriber data, we expect that the GUP effort will flourish, and will become one of the dominant forces in helping to standardize the area of profile management.

Although 3GPP GUP is focused primarily on profile management for the 3G network, we expect that GUP will address profile data that resides elsewhere but is relevant to 3G (*e.g.*, calendars, which can play an important role in preference around privacy). The GUP working group is also establishing links with other standards groups that are concerned with profile data, including 3GPP2 and Liberty Alliance [2]. For this reason, we expect the GUP activity to be a dominant force in profile management for converged services in general, not just for 3G wireless.

3.2.2 GUP vision

The GUP vision starts by looking at the current problems:

- **great quantities of data, spread all over:** As mentioned previously on this section, data is stored in many different entities (*e.g.*, network components, IT information systems, user devices). Numerous and incompatible schemas and models are used to represent and store the data. Moreover, data is not often reused – requiring redundant storage, which leads to inconsistencies and wasteful re-entry.
- **terminal management:** A key focus of GUP is to make the management of end user terminals (provisioning and synchronization) easier. GUP has already identified SyncML [4] as the protocol for synchronization.
- **customer care:** End user problems must be able to be diagnosed and fixed via the network.
- **multiple protocols:** There is a need to support protocols for communication between network components, between network components and terminals, and between terminals (*e.g.*, phone ↔ laptop).
- **incompatibilities across domains:** Various domains usually use incompatible data models, schemas and protocols. Requirements are also radically different (*e.g.*, real-time, reliability, security, etc.).

3.2.3 GUP requirements

Even though GUP is still in its early stage, the following requirements for GUP have been identified:

- harmonized data description
 - common data model
 - extensible
 - compatible with GUP information model

The information model consider a user profile as a collection of profile components. A component is used as a unit of storage and access control. Components are linked together by the identity they refer to. See Figure 6 for the corresponding UML diagram.
- harmonized query/update interface
- common transport mechanism




Figure 6: GUP information model

- common security protocols
- privacy

Although not a formal requirement of GUP, it is very likely that the GUP group will adopt XML as the underlying format for the common data model and for data exchange.

It should be clear that the requirements we have listed in Section 2 are a natural generalization of the GUP requirements to support services across the converged network, not just across the 3G wireless network.

At present, the GUP group has not specified a reference architecture that can be used to provide simple and focused access to profile information.

4 GUP^{ster} in a nutshell

This section introduces the high-level GUP^{ster} vision, in its most basic form. The following section presents a variety of extensions and alternative formulations.

4.1 Napster + GUP = GUP^{ster}

To summarize in one sentence what GUP^{ster} is all about, we can say that GUP^{ster} is to user profile components what Napster was to music files.

In Napster, users willing to share music files join the Napster community by registering their files on the Napster central server. The server stores meta-data about files and users. When a user wants to

retrieve a given file, it sends a request to the Napster server and gets back list of peers (other users) who have registered this very file. The user can then ask for the file directly from the peers.

In GUP^{ster}, data stores willing to share user profile components join the GUP^{ster} community by registering their components on the GUP^{ster} server. The server stores meta-data about data stores and components (*e.g.*, location, access control, etc.). When an application (*e.g.*, client application, data store, etc.) wants to retrieve a given component, it sends a request to GUP^{ster} and gets back a list of data-stores which have registered this very component. The application can then ask for the component directly.

4.2 Overview of the GUP^{ster} architecture

The GUP^{ster} architecture is presented in Figure 7 and consists of client applications, GUP-enabled data-stores and the GUP^{ster} server.

Client applications are applications which need to access user profile information for both query and update. Selective reach-me presented in Section 2.2 is a good example.

Data-stores are components which store and manage user profile information, such as HLRs, presence servers, portal sites, etc. Data stores need to be GUP-enabled in order to participate in the GUP community. Concretely, this means that an adapter is put on top of the data store to offer a GUP-compliant interface (protocol and data model)⁵.

The GUP^{ster} server is the central repository of meta-data regarding user profile components. Among other things, it stores coverage (how the GUP schema is mapped onto existing data stores) and access control information.

Note that “central repository” has to be understood from a logical point of view and may be implemented as a constellation of connected servers. In the simplified context of this section, we envision GUP^{ster} being supported in a manner similar to the UDDI registry [5], *i.e.* a family of mirrored servers hosted by a consortium of enterprises and

⁵As of this writing, the details of the interface have not been finalized yet. The data model will be XML-based and the protocol will probably be SOAP or HTTP.




Figure 7: GUP^{ster} top-level architecture

freely available to all users.

4.3 GUP^{ster} in action

We assume that we have some data stores willing to share user profile components and that these components support the GUP interface. (Figure 8 shows an example of GUP^{ster} integrating data from multiple data stores.) We now present a simplified scenario for GUP^{ster} in action. Later, we will discuss the access control aspect.

Like for Napster, the first step is for a data store to register to GUP^{ster} the components it is willing to share. For instance, Yahoo! will tell GUP^{ster} that it stores the address book of Arnaud and the address book and game scores of Rick. Sprint PCS will inform GUP^{ster} that it stores Arnaud's address book and game scores. Data stores can also unregister components.

For each user, GUP^{ster} maintains the *coverage* of profile components by data stores. In the case of Arnaud, the coverage would look like ⁶:

Coverage	
□ /user[@id='arnaud']/address-book	⇒ { gup.yahoo.com, gup.spcs.com }
□ /user[@id='arnaud']/presence	⇒ { gup.spcs.com }

⁶We assume that the GUP schema is defined with elements such as `user`, `address-book`, `presence`, `buddy-list`, etc.

□ /user[@id='arnaud']/buddy-list	⇒ { gup.spcs.com }
□ /user[@id='arnaud']/payment	⇒ { gup.citibank.com }

A coverage is a mapping between sub-trees of the GUP schema (expressed as XPath expressions) and data-stores. Note that a given profile component can be mapped to multiple data-stores.

When a client application wants to retrieve or update a profile component, it needs to send the request to GUP^{ster}. For instance, Arnaud wants to synchronize the address book on his cell-phone. The cell phone (through an application running on the cell phone) will send a request to GUP^{ster}. GUP^{ster} will rewrite the request and send it back to the client.

In our example, GUP^{ster} will return to the client application something like:

Referral from GUP ^{ster}	
gup.yahoo.com/user[@id='arnaud']/address-book gup.spcs.com/user[@id='arnaud']/address-book	

where `||` has to be understood as a choice [22]. GUP^{ster} does not return any data, just a referral to be used by the client application.

The client application will then use the referral (one of them, or both) to get the data directly from the GUP data stores.

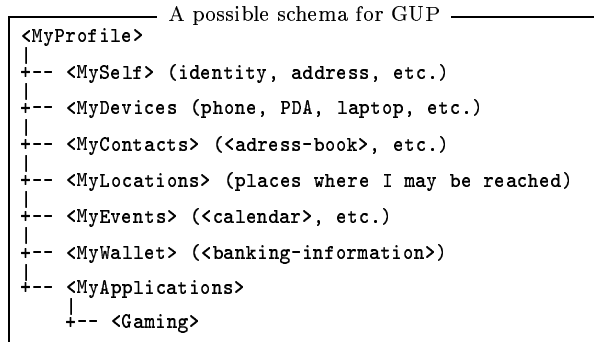
4.4 Schema management

A key aspect of GUP^{ster} is to integrate profile components from a large number of data stores living in different networks.

As we already noted, since the kind of queries we plan to support are not join-based, the “local as view” and “global as view” approaches essentially converge in this context. We assume that a global schema will be defined and maintained by a standard body (*e.g.*, 3GPP, W3C).

This is the responsibility of the GUP^{ster} server and the various data stores to make sure that the latest version of the schema is the one being used. Note that by design, the schema can be made more tolerant (or not) to evolutions (*e.g.*, using optional elements or attributes).

Here is a possible top-level schema for user profiles. The exact definition (and extensions) of the schema will be the responsibility of a standard body.



4.5 Schema coverage

The schema coverage is a mapping between sub-trees of the GUP schema and user profile components available at the data-stores (see Figure 8). We assume that a profile component can always be defined as a sub-tree of the GUP schema. We exclude cases where, for instance, two profile components could not be defined as sub-trees but their join could be. The language we use to define coverage is a subset of XPath [28] with child- and attribute-axis only and limited predicates, in order to have a canonical way to navigate the tree.

For a given user, the GUP^{ster} server will store a list of mappings between a sub-tree of the GUP schema and one or more data stores. We have already presented some examples of coverage in the

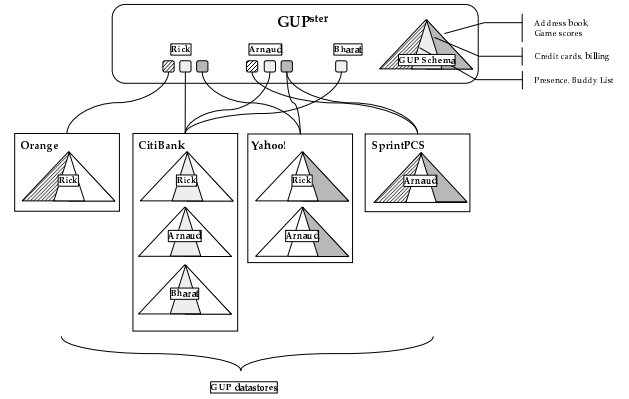


Figure 8: Schema coverage

previous section.

Figure 9 has a more interesting example where Arnaud’s address book is split between his corporate and personal address books.

A request for `/user[@id='arnaud']/address-book` should return a referral to both data stores as well as a way to merge to two XML fragments. We will discuss this issue in Section 6.

4.6 Access control

A critical aspect of the sharing of user profile components is privacy: users are willing to grant access to their profile information (“enter once, use everywhere”), provided they remain in control of who can access this information and when. Access control raises two issues: (i) how can the user specify her privacy shield, *i.e.* the set of access control rules/policies for her profile information; (ii) how are these policies being enforced.

Privacy shield

Conceptually the idea of a privacy shield is quite simple: should a given request be granted access to the data it asks for? The problem is to define what we mean by a request. The privacy is defined in terms of policies to be applied whenever some user profile is requested through GUP^{ster}.

As part of his privacy shield, for example, a corporate user may want to establish the following policies: *any co-worker can access my presence information during working-hours; my boss and my fam-*

□ /user[@id='arnaud']/address-book/item[@type='personal'] \mapsto { gup.yahoo.com } □ /user[@id='arnaud']/address-book/item[@type='corporate'] \mapsto { gup.lucent.com }
--

Figure 9: GUP^{ster} representation of address book split across two sites

ily can access my presence information at any time; my family can access my personal address book and calendar.

In GUP^{ster}, a request consists of two facets: a context and a path. The path defines what components of the user profile are asked for. The context provides some information about the context of the request, i.e. identity of the requester (*e.g.*, third party application, end user, etc.), purpose of the request (*e.g.*, plain request, caching request, subscription-based request, etc.), etc. We envision the context to be an XML document as well, defined using a *request context schema*.

The path can be defined as an XPath expression (returning nodes) over an instance of the GUP schema. The context can also be defined as an XPath expression (returning a boolean) over an instance of the request context schema.

We plan to reuse as much as we can from the emerging XACML [20] (XML Access Control Meta Language) standard. Unfortunately, the notion of request context in XACML is too limited (restricted to principals) to define a sufficiently rich privacy shield. As mentioned above, we will need to define an XML schema for the context information as well.

Policy infrastructure

Defining policies is one point; evaluating them and enforcing them is another. Issues to be solved are:

- who stores the policies
- who evaluates them
- who enforces them

Various standard bodies ([18, 15]) have defined an abstract architecture for policy management, identifying some key roles:

- policy repository: in charge of storing policies
- policy administration point: in charge of provisioning the rules (i.e. letting the user access and modify the rules that define her privacy

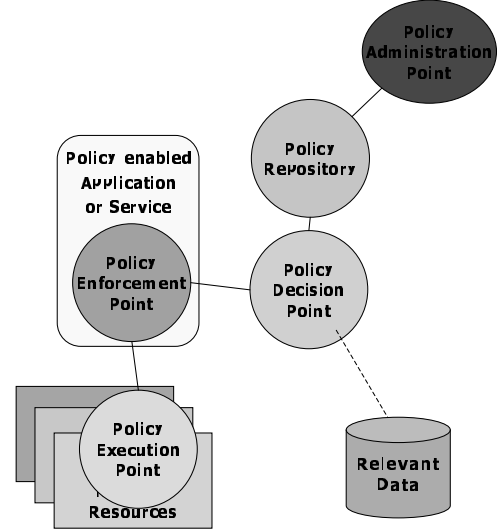


Figure 10: Policy Infrastructure

shield) and other administrative tasks (*e.g.*, checking that the rules are valid)

- policy decision point: in charge of rendering a decision based on a rule set and a context. The decision point only returns a decision and has absolutely no side-effect on the environment
- policy enforcement point: in charge of asking for a decision and enforcing it (i.e. taking the required actions). Note that in some cases, the enforcement point uses a policy execution point

The details of the architecture are presented in Figure 10.

Applied to our architecture, we envision the following distribution of roles (in the next section we will present variations where the roles can be assigned slightly differently). GUP^{ster} will be at the same time the administration point (allowing end-users to provision their policies), the policy repository (storing policies), the decision point (computing the decision) and the enforcement point (applying the decision, *i.e.*, sending or not a referral to the client). The data stores will be execution points. For 3rd party application, the policy infras-

structure will be transparent: the application will send a request to GUP^{ster} and get back (or not) a referral.

5 GUP^{ster} Unleashed

In the previous section, we have presented the basic GUP^{ster} where all meta-data is stored in a centralized way and where the server only returns referrals. We now motivate and explore some variations of this architecture.

5.1 Architectural Variations

Maintaining and managing the profile meta-data is a central component of the GUP^{ster} paradigm. In the previous section, we described one architectural approach for how and where the meta-data might be managed, namely, a centralized meta-data manager (MDM). That approach was inspired directly from Napster, and assumes a UDDI-like universally available, mirrored meta-data store.

We now explore two dimensions along which alternatives to centralized MDM may be found. These alternatives share many of the same technological challenges as centralized MDM (*e.g.*, how to specify data partitions, how to efficiently convert real-time data into the GUP^{ster} -compliant format), and also involve some new challenges (*e.g.*, security for the meta-data access, hierarchical versions of GUP^{ster}).

5.1.1 Motivation for a revised architecture

The primary reason to examine architectural alternatives is that two things are unpredictable at present: (i) the degree of privacy about meta-data that end-users and enterprises will insist on, and (ii) the business model around meta-data management that will emerge.

With regards to privacy, consider for a moment Microsoft's Hailstorm initiative, which proposed a framework whereby end-user profile data (including calendar, address book, credit card, shipping addresses, etc.) could be stored at a central location by a single organization (MSN in this case) and could be shared on a selective basis with e-merchants at the "click of a button". As announced

(see [30]), this initiative was dropped (and the underlying technology re-directed to a different business model), because consumers were not willing to have a substantial amount of their profile controlled by MSN. It seems that consumers are unwilling to have all of their meta-data stored in a universally available store managed by single corporation. Further, large enterprises will probably have quite stringent requirements about security on what individuals and organizations can hold or have access to what end-user data. They will also have business reasons not to share data with certain other repositories (*e.g.*, to make it harder for customers to leave a wireless service provider).

Many factors will be involved in the emergence of a business model (or models) that will arise to support a GUP^{ster} -style framework. A key question, of course, is how will the organizations that provide meta-data management be funded. Various alternatives exist, ranging from the end-user or the data requester paying for each use of the services explicitly (through a billing mechanism similar to current phone usage), to having the service bundled with other services (with the costs thereby "hidden" from the end-user), *e.g.*, as part of an internet portal or part of wireless phone service.

5.1.2 Alternatives to MDM

- *User-level distributed MDM*: A simple variation of centralized MDM is to assume that different end-users will want their meta-data managed by different organizations, perhaps their wireless or internet service provider, their bank, an internet portal, their employer, or even their home computer.

How will applications discover where to find the meta-data for a given end-user? Several alternatives exist. One approach is a UDDI-like universally available "white pages" that includes links from personal identifiers to meta-data repositories. The other extreme is that applications will be able to find the meta-data for a given user only if that user informs the application of where to find it. Perhaps a compromise will emerge — a universal white pages but with the option for people to have "un-listed" pointers to their meta-data management.

- *Hierarchical MDM*: A further refinement of user-level distributed MDM is to allow the meta-data management for *each user* to be distributed.

For example, a user might specify his primary

meta-data manager to be their wireless service provider, but have meta-data about his credit card and other banking information be stored by a bank, and his meta-data about internet games be stored by an internet portal. In this example, the wireless service provider will “know” that the end-user has banking meta-data and internet gaming meta-data, but will “know” essentially nothing about it. As indicated above, this may be especially relevant for enterprises, who might insist that all enterprise-relevant profile meta-data of employees be managed directly by the enterprise.

This suggests that there will not be a “one size fits all” implementation of GUP^{ster} nodes, but rather a handful of general implementations, and then several more narrowly targeted implementations.

5.2 Query Processing Variations

In the proposed architecture, GUP^{ster} is not doing any query processing only query rewriting. Its role is to combine coverage information and access control information to rewrite a client’s request accordingly. GUP^{ster} does not return any data, only referrals. This means that a GUP^{ster} server is easier to implement and can serve more client concurrently.

But there might be some situations where having GUP^{ster} do some query processing would be useful. In the case of a client application with very limited capabilities (*e.g.*, a cell phone), the client may not be able to perform some required data transformation (like our address book example where the data is split among two stores and requires a merge). Offering a larger variety of *distributed query patterns* [22] like chaining, referral, recruiting (where the request is actually migrated to a different node) will be needed. GUP^{ster} should probably also offer some caching to make the access to user profile component faster.

Another important dimension concerns subscriptions. In the current architecture, GUP^{ster} is a reactive (pull-based) not pro-active (push-based) system. It is always possible to push-enable a pull-based system using polling, but this may not be very efficient. In our case, every polling request needs to be checked to enforce the end-user’s privacy shield. Having the subscription handled by GUP^{ster} internally would save this extra work.

These variations can always be hard-coded into

the various nodes of the system (the GUP^{ster} server and the various data stores). But it seems that a notion of *mobile query process* (as proposed in [22]) needs to be defined in order to capture the full range of query processing that may be needed to deploy such a distributed architecture.

5.3 How GUP^{ster} addresses the GUP requirements

We now revisit the requirements presented in Section 2.3 and show how they are being addressed by GUP^{ster}.

Data richness: Just like Napster has a built-in schema to “talk” about music files, GUP^{ster} maintains the schema which defines the structure of the various user profile components and their relationships.

The structure of the GUP schema is defined using XML as the underlying data-model. We will motivate the choice of XML (compared to LDAP for instance) in Section 6.

Data stores willing to join the GUP^{ster} community need to export (virtually or physically) user profile components as XML, according to the GUP schema.

Data transformation: When the structure of the data hosted on the data store is not directly compatible with the GUP schema, we assume the existence of some wrappers/mediators in charge of transforming the data into the right structure. The transformation can be virtual or physical.

In some extreme cases (*e.g.*, end user terminals with very limited query capabilities), we could imagine GUP^{ster} itself (or a component part of it) doing some further transformation if needed.

Data placement: We need to distinguish here between (i) placement decisions made by users (based on trust concerns) and (ii) placement decisions made by the infrastructure (based on performance concerns).

For (i), the placement is dictated by the registration of components by data stores. For instance a user will instruct Yahoo! that it is responsible for storing its address book. Yahoo! will in turn register this component to GUP^{ster}:
 /user[@id='Arnaud']/AddressBook =>
 www.yahoo.com

For (ii), we can imagine that data stores operated by a given operator will be replicated, transparently for the rest of the network. Yahoo! for instance will store address book information in multiple servers (*e.g.*, `us-east.yahoo.com`, `us-west.yahoo.com`, `www.yahoo.co.uk`, etc.) and requests sent to `www.yahoo.com` will be routed to the closest Yahoo! store available. GUP^{ster} can also offer some caching services.

Data integration and Reconciliation: Data integration is provided by using adaptors to convert profile data from multiple repositories into a common format. Reconciliation can be handled by prioritizing sites or by some more sophisticated method when GUP^{ster} data is requested.

Data synchronization: The GUP working group has already agreed to use SyncML as the synchronization protocol. But this is only one aspect of the problem because SyncML [4] is only a transport protocol. Issues like synchronization semantics need to be addressed. GUP^{ster} does not provide specific solution for this yet.

Security and access control: GUP^{ster} does not plan to innovate in the field of security but will try to reuse the best existing solutions. An interesting issue is where the access control should be performed. We think that GUP^{ster} should be in charge of access control because it offers a single point of access. Having access control at the level of the data-stores would require keeping access control policies in sync.

Here is a possible way to enforce access control. When an application sends a request to GUP^{ster} for a given component, GUP^{ster} checks whether or not access is granted. It rewrites the query accordingly (for instance only a subset of the information asked for can be returned) and signs it, including a timestamp. The application can send the rewritten and signed query to the corresponding data store(s). The store will check the time-stamp and the signature and eventually return the data. We assume that data store will only accept queries which have been signed by GUP^{ster}.

Reliability: Reliability will be achieved by having the logical single entry point be implemented by a constellation of GUP^{ster} servers. The reliability of data-stores themselves is independent of GUP^{ster}.

Scalability and performance: As the logical single point of entry, before forwarding the queries, GUP^{ster} is able to filter out spurious ones (*e.g.*,

queries which do not fit with the GUP schema, queries which do not satisfy the access control requirements).

Since GUP^{ster} does not store any data, GUP^{ster} does not require a heavy duty database. GUP^{ster} needs to rewrite queries using coverage and access control information. By putting these function at the level of GUP^{ster}, we can keep the data stores as is (except for the GUP adapter on top of any data store) and expect very little overhead because of GUP^{ster}⁷.

The use of multiple distributed query patterns (*e.g.*, chaining, referral, recruiting) will permit minimizing the transport cost of result information.

The Napster analogy can give us some useful insights about scalability and performance. At its peak, Napster had more than 50m users with more than a few millions connected at the same time.

6 Related Work

Profile Management efforts

There is a lot of on-going efforts in the area of user profile management. Both Sun (pushing for Liberty Alliance [2]) and Microsoft (pushing for Passport aka TrustBridge aka MyServices aka Hailstorm) are fighting over the next standard for network identity. Both initiatives are currently focusing more on the authentication aspect (single sign-on) than on the issue of data management. Liberty Alliance – as of this writing – has not proposed any data model and schema for user profile information. Microsoft's user profile only consists of the traditional user information stored by an ISP.

A few years back, the DEN initiative [10] proposed a suite of schemas to describe network components and devices. Netscape roaming profiles are a direct extension of those.

W3C Composite Capability/Preference Profiles (CC/PP) [27] initiative started in 1999 aims at creating *“a general, yet extensible framework for describing user preferences and device capabilities”*.

⁷The GUP^{ster} architecture does not preclude network components from communicating with each other the way they used to. For instance, the presence component can retrieve location information from the HLR through GUP^{ster} or directly from the HLR.

The framework is based on RDF but does not seem to have received much support so far.

LDAP-based approaches

When deciding on a representation/transport representation for GUPster, XML and LDAP were natural contenders. LDAP has long been used for storing profile data and standard objectclasses exist for detailed user profiles. Furthermore, DEN [10] proposed even more detailed extensions.

Advantage of LDAP:

- LDAP is built for extensibility. Objects are modeled with "aspects" and can always implement a new objectclass (and hence a set of fields). XML would require power-sets of basic types to get similar flexibility — and way to migrate objects/interfaces from one type to another to boot.
- LDAP already defines a number of standardized objectclasses which are well-understood and used by many.
- LDAP's hierarchical structure makes LDAP directories very easy to scale — it is straightforward to move arbitrary sub-trees to different servers.

Advantages of XML:

- XML data can have arbitrary structure. LDAP objects are very simple (and flat): each entry in the LDAP tree is a set of name/value pairs. Each of the values can be set valued, but only for atomic types.
- XML comes with more sophisticated query languages. The LDAP query language is very limited and can only return sets of single objects.
- LDAP does not offer any typing mechanisms. Typing is not used so much for sanity checking input as for deciding which comparison function to use (*e.g.*, if a field is a phone number type, then 908-582-4393 and (908) 582-4393 should compare as equal despite their different representation).
- XML is being enriched by new Emerging standards. The fact that XML Schema supports annotations, and the proposed Schema Adjunct Framework [26], can provide a foundation for building up a profile management framework and associated tools.

We picked XML not just because it was a requirement from GUP). We also plan to leverage the LDAP/DEN schemas (translating them into XML schemas) in developing GUP^{ster} profiles and to provide tools to wrap LDAP sites.

Netscape Roaming profiles illustrate nicely the limitations of the LDAP approach. Address book, bookmarks and calendar are best described as nested structures, which are hard to capture in the LDAP data model. The workaround used by Netscape is to create new LDAP objectclasses that store the information as binary objects. From the client point of view, the objectclass is just a container for the information.

The advantage of this approach is simplicity and extensibility. For instance, I can store my MP3 play list in my roaming profile: this is a binary object with its own format and LDAP does not need to know anything about it.

This opaque storage has two major drawbacks. First these opaque objects can only be accessed (retrieved or updated) as a whole. Second, it is not possible to combine information from two separate objects (*e.g.*, combining calendar information with address book information to find the phone number of the people I am having a meeting with) because they don't share the same data model. XML puts all the objects on the same ground (they all share the same data model).

Personalization through policies

Personalization of converged services requires more than providing simple ways to input and access profile data; it requires the ability to specify and enforce a broad range of preferences.

Technology for storing and sharing preferences and their associated policies can build on the GUPster framework described in this paper. But considerable research is still needed in terms of how to provision and process preferences. Some user preferences are rather simple, *e.g.*, that a user prefers to use one credit card for book purchases and another for travel-related expenses, and can easily be captured (and shared) using GUPster. But other preferences may be quite elaborate, *e.g.*, about when and to whom a user's presence information should be revealed, or about whether and how voice communication should be established as in the selective reach me example.

Standards bodies (e.g., Parlay policy management API [15], OPES [18]) and research activities are developing approaches to support the specification and enforcement of such elaborate preferences, typically through the use of rules engines. A variety of rules semantics might be used for different application areas, including systems with no chaining and systems that support production system style semantics [17].

XML integration

A lot of work has been done in the domain of XML data integration.

Silkroute [13, 14] enables the applications to specify virtual XML views over a relational database, and allows these views to be queried using XQuery. The IBM Xperanto [9, 23] project and the Enosys XML Integration Platform [21, 1] further allow the application to define and query XML views on data spanning multiple data sources.

As mentioned earlier (Section 2.3), query-oriented data integration of profile data is simpler because profile queries typically do not involve joins. As such, with respect to querying, these systems seem to offer more than GUP^{ster}. However, querying is just one aspect of data integration as supported in GUP^{ster}; apart from being queried, the data across the different data sources in GUP^{ster} needs to be updated (provisioned) in an integrated manner. None of the systems mentioned above handles integrated updates. Furthermore, GUP^{ster} also handles data placement (including caching) and reconciliation aspects (see Section 2.3) that are not addressed in any prior framework to the best of our knowledge.

The notion of coverage in GUP^{ster} requires some algorithms to decide query containment of XPath expressions (or subset of XPath), as studied in [11]. Some recent work on keys for XML [7] can also be applied to define coverage. New operators for merging XML components are also relevant (like Deep Union [8] or Merge [25]).

Security and access control

There has been a lot of work on security and access control in general. For XML, we already mention the emerging XACML proposal [20] and some of

its limitations. XACML has a very restricted notion of request context. In GUP^{ster}, the quality of the privacy shield will depend on the context information the user can use to define her requirements. Moreover, XACML policies are expressed using some simple rules (no forward chaining) which are not suitable for complex access control policies. Offering an expressive framework with good enough performance is clearly a challenge. Another key issue is the provisioning interface that end users will use in order to express their preferences.

An elegant solution for conditional access to XML (based on encryption) data has been presented in [19]. Some efficient algorithms for XML access control have been proposed in [31].

In any case, GUP^{ster} will have to be compliant with whatever standard gets accepted for network identity (*e.g.*, Liberty Alliance [2]).

7 Conclusion

This paper identifies a new and challenging direction in data management, namely to support easy (but controlled) access and sharing of profile data in support of so-called converged services. These services often involve real-time, interactive communication between a user and other users and/or multiple network-hosted applications, and the associated profile data is typically spread across multiple networks and organizations. Additional aspects and requirements in this application area include the expectation that a single data schema will emerge for (most) profile data, that privacy and access controls are highly relevant, and the current lack of clarity on the business models for managing profile data will be resolved. In addition to describing these directions in general terms, this paper describes the key kinds of profile data that need to be accessed in support of converged services, and also how and where that data is currently stored in the various networks.

The paper has proposed a high-level framework for supporting access and sharing of profile data, which satisfies the requirements agreed upon to date by the 3GPP GUP standards body, and that also follows the spirit of the Liberty Alliance consortium. This GUP^{ster} framework is fundamentally a combination of two things: the federated database paradigm and a peer-to-peer, Napster-inspired ap-

proach for organizing meta-data information. The fundamental technology challenges of GUP^{ster} are how to manage and control access to profile meta-data and the associated data, and how to provide very efficient, scalable, and reliable support for sharing of the profile data.

While the GUP^{ster} proposal provides a sound basis for developing technologies for sharing profile data in support of converged services, it raises many challenges. We mention three of the most fundamental challenges here.

The core challenge is to continue work on meta-data management itself — what is the right conceptual model for meta-data (in the context of managing profile data), what are the efficient ways to store and access meta-data, what integrity constraints are relevant and how can they be enforced, is XPath sufficient for expressing the partitioning of meta-data and/or data that will be managed by disparate organizations, how should XACML [20] be adapted or extended, how should the Schema Adjunct Framework [26] be applied to capture these aspects, and what is a systematic framework for supporting the extension of the global profile schema (for both local and global extensions)?

A second challenge concerns efficient, scalable, reliable implementation. Although we have presented some approaches in this direction, the area requires much more detailed consideration, including an analysis of performance requirements, the development of testbeds and benchmarks, and ultimately the development of self-adapting implementations for data access, transformation, distribution, caching, etc.

The third core challenge involves data provenance [8], that is, the tracking of where data (and meta-data) have come from, and where they have been used. In e-commerce, when a user buys something, she gives her credit card number and the merchant gives a confirmation number for that specific purchase. The user trusts that the merchant won't use the credit card number beyond the purchases that the user authorizes. This illustrates just one example of the many kinds of tracking mechanisms that will be needed around access to profile data and meta-data. A special case of the data provenance challenge, in a different direction, concerns management of overlapping sets of profile data, e.g., a user's personal and enterprise address books might be held by different organizations but hold over-

lapping data. What are systematic ways to support data reconciliation, to identify a single data source that holds all the data needed for a specific application, and to avoid distribution of data from one source that violates access controls given for another source?

Acknowledgements: the authors would like to thank Vinod Anupam, Bharat Kumar and Prasan Roy for comments and discussions about ideas presented in this paper.

References

- [1] Enosys. <http://www.enosysmarkets.com/>.
- [2] Liberty Alliance. <http://www.projectliberty.org>.
- [3] Napster. <http://www.napster.com>.
- [4] SyncML. <http://www.syncml.org>.
- [5] Universal Description, Discovery and Integration (UDDI) project. <http://www.uddi.org>.
- [6] 3GPP. Generic User Profile, 2001. <http://www.3gpp.org>.
- [7] Peter Buneman, Susan Davidson, Wenfei Fan, Carmem Hara, and Wang-Chiew Tan. Keys for XML. In *WWW10*, May 2001.
- [8] Peter Buneman, Alin Deutsch, and Wang-Chiew Tan. A deterministic model for semistructured data. In *Workshop on Query Processing for Semistructured Data and Non-Standard Data Formats*, 1998. Available from <http://db.cis.upenn.edu/DL/icdt.ps.gz>.
- [9] Michael J. Carey, Jerry Kiernan, Jayavel Shanmugasundaram, Eugene J. Shekita, and Subbu N. Subramanian. XPERANTO: Middleware for publishing object-relational data as XML documents. In *VLDB 2000*, pages 646–648, 2000.
- [10] DEN home page, 2002. <http://www.dmtf.org/spec/denh.html>.
- [11] Alin Deutsch and Val Tannen. Containment and Integrity Constraints for XPath Fragments. In *KRDB 2001*, 2001.
- [12] Kazutaka Murakami et. al. Global Roaming and Personal Mobility with COPS architecture in SuperHLR. Submitted for publication.
- [13] M. F. Fernandez, W.-C. Tan, and D. Suciu. “SilkRoute: Trading between Relations and XML”. In *Int’l World Wide Web Conf. (WWW)*, Amsterdam, Netherlands, May 2000.
- [14] Mary Fernandez, Atsuyuki Morishima, Dan Suciu, and WangChiew Tan. Publishing Relational Data in XML: the SilkRoute Approach. *IEEE Data Engineering Bulletin*, 24(2), 2001.
- [15] The Parlay Group. Parlay 3.0 Policy Management Specification, 2002. Available from http://www.parlay.org/docs/Spec3_ParlayPolicyManagement.html.
- [16] Dennis Heimbigner and Dennis McLeod. A federated architecture for information management. *TOIS*, 1985.
- [17] Richard Hull, Bharat Kumar, Arnaud Sahuguet, and Ming Xiong. Have It Your Way: Personalization of Network-Hosted Services. In *Advances in Databases, 19th British National Conference on Databases*, 2002.
- [18] IETF. Open Pluggable Edge Services (OPES). <http://www.ietf-opes.org>.
- [19] Gerome Miklau and Dan Suciu. Cryptographically Enforced Conditional Access for XML. In *Proceedings of WebDB*, 2002.
- [20] OASIS. eXtensible Access Control Markup Language. Available from <http://www.oasis-open.org/committees/xacml/>.
- [21] Yannis Papakonstantinou and Vasilis Vassalos. Architecture and Implementation of an XQuery-based Information Integration Platform. *IEEE Data Engineering Bulletin*, 25(1):18–26, 2002.
- [22] Arnaud Sahuguet. *ubQL: a Distributed Query Language to Program Distributed Query Systems*. PhD thesis, University of Pennsylvania, Department of Computer and Information Science, 2002.
- [23] Jayavel Shanmugasundaram, Jerry Kiernan, Eugene J. Shekita, and Catalina Fanand John Funderburk. Querying XML Views of Relational Data. In *VLDB*, 2001.
- [24] Amit P. Sheth. Federated database systems for managing distributed, heterogeneous, and autonomous databases. In Guy M. Lohman, Amílcar Sernadas, and Rafael Camps, editors, *VLDB*, page 489, 1991.
- [25] Kristin Tufte and David Maier. Aggregation and Accumulation of XML Data. *IEEE Data Engineering Bulletin*, 24(2):34–39, 2001.
- [26] Scott Vorthmann and Lee Buck. Schema Adjunct Framework. Draft Specification 24 February 2000, Feb 2000.
- [27] W3C. Composite Capability/Preference Profiles (CC/PP). <http://www.w3.org/TR/NOTE-CCPP/>.
- [28] W3C. XML Path Language (XPath). Available from <http://www.w3.org/TR/xpath>.
- [29] Gio Wiederhold. Mediators in the architecture of future information systems. *IEEE Computer*, 1992.
- [30] Joe Wilcox. Survey: Passport required—not appealing. CNET, April 2002. <http://news.com.com/2100-1001-884730.html>.
- [31] Ting Yu, Divesh Srivastava, Laks V. S. Lakshmanan, and H. V. Jagadish. Compressed Access Control for XML. In *Advances in Databases, 19th British National Conference on Databases*, 2002.