1. Publishable summary

The vision of the PrimeLife project is to enable individuals in the information society to protect their privacy and retain control over their personal information, irrespective of the activities they are performing. Today, individuals are put at risk as information technologies hardly consider those privacy requirements of the citizens, customers, and consumers.

Today, individuals and businesses are increasingly using social networking, online collaboration applications, mesh-ups of different services, and the Internet in general, for both private and business purposes. This clearly emphasizes the increasing need to easily use, share, and manage identities and associated trust data in an organized way. Despite all of this, the new information technologies hardly consider the privacy requirements of the individuals. This is apparent in applications such as today's electronic social networks and user-created content applications such as blogs and Wikis. This lack of respect for privacy is particularly unacceptable as most people need to use such information technologies, at least occasionally.

PrimeLife's Approach

PrimeLife's Approach is threefold. First, PrimeLife will pick up the results and technologies from the PRIME project and work towards their adoption in the real world by providing materials for standardization and education. Second, PrimeLife will eliminate the remaining hurdles for large-scale adoption as identified by PRIME including user interfaces, policy languages for the technologies' integration into currently-used access control schemes, and infrastructural components. Third, and certainly not least, PrimeLife aims to solve the privacy, identity, and trust management challenges for the cases where to the paradigm of privacy protection by data minimization fails. This includes Web 2.0 applications, i.e., where users provide content and, in contrast to the paradigm of privacy protection by data minimization, want and need to reveal substantial amounts of personal data. It also includes the management of personal data throughout our entire lives.

Privacy for Life – Beyond Data Minimization

The objective of PrimeLife in this area is to 1) bring better privacy and identity management to selected emerging scenarios in today's Internet; 2) investigate how to maintain life-long privacy; and 3) progress the foundation of privacy technology in today's interactions over electronic communication networks such as the Internet or cellular networks. The first goal is to find means to achieve privacy and security in specific scenarios including trustworthy user-generated content, electronic social networks, or collaborative workspaces, the second one is more generic, operates in a much less known space, and is geared towards long-term aspects of privacy protection and identity management. PrimeLife carries out foundational research to understand the underlying issues and to develop solutions for these challenges. The viability of our solutions is ensured by building and evaluating targeted prototypes. The third goal is generic in that it tackles technology and usability issues of various key problem areas in the privacy and identity space, particularly while supporting the first two challenges with foundational work.

Our work on *trusted content* is concerned with injecting trust in (user-generated) content and allowing other readers an assessment of the trustworthiness of the content. We have built our first *demonstrator* showing how trust metadata can be attached to user-generated data, such as blogs or Wiki entries by a content producer or reviewer and assessed by a content consumer. The work has focussed on a blog scenario in the first year and developed some foundations of technologies and an associated demonstrator to show the practical application of the approach. The results are encouraging in that they show the technical feasibility of achieving more trustworthy (user-generated) content. Though, the evaluation has also shown the need for a system to give people incentives for contributing to the establishment of the trustworthiness of content, e.g., when reviewing and rating content, and to take into consideration the trustworthiness of rating users. These evaluation results have resulted in our roadmap for continuing the work in the trusted content area. Based on this, in the second year of PrimeLife, we have focussed research and prototype development on privacy-enhanced incentive schemes and reputation systems to facilitate trusted content. We envision obtaining a practical toolbox of research results and prototypes that will be made available to the community.

The area of on-line social interactions and collaborations such associal networks has gained significant attention during the past years. Unfortunately, privacy has not been accounted for seriously so far. In the first project year, we have studied the functioning and aims of social networks and collaborative spaces and havecome up with a number of approaches to, on the one hand, offer (better) protection of the personal data in such application and, on the other hand, to support users in assessing privacy consequences of their actions in these applications. In year 2

of the project, we have built prototypes that demonstrate the feasibility of our approaches to protect privacy in such applications. Our social network prototype *Clique* shows how users can be assisted in defining the target audiences for data in their social network profile as well as how to make users aware of the audiences. Our second social network prototype *Scramble!* shows how access control to data in one's social network profile can be enforced in any existing social network employing encryption. Furthermore, our demonstrator in the area of collaborative workspaces shows how access control to data is done in an attribute-based way, using anonymous credential technology for authentication of parties.

The area of life-time sustainable privacy and identity management has gained more attention during the second year of PrimeLife. We have investigated what it means to protect privacy and to manage our identities throughout our lives. In particular, we have come up with a number of scenarios that exemplify these aspects and provide requirements. Based on this, we have picked a suitable scenario for our prototype work for the third year. It will demonstrate our concepts for the life-time sustainability of a user's data collection, particularly backup and access of her data. A focus is put on how this will be done in a privacy-friendly way.

Technical Agenda and Progress

PrimeLife has identified the key technical areas to be addressed so that privacy-enhancing identity and trust management can be realized in practice: Human-computer interaction, policy languages, access control technology, user-support tools for identity management, privacy-enhanced processing of data sets, infrastructural aspects such as trusted hardware tokens, service compositions, or mobile infrastructures, and privacy-enhancing cryptography.

Within the first two years of the project, the teams in the different activities of work have collaborated extremely well and produced a considerable amount of research results in the targeted areas. The work has been driven by the foundation of requirements and use cases worked out early in the project, and continuously updated as well as the joint goals across the different activities of work. Our excellent progress on the main demonstrators which are located in the area of social networks and collaborative spaces deserves being mentioned separately.

The *basic research* on privacy technology has progressed the state of the art in the areas we have been doing research in. In *cryptography*, multiple extensions to anonymous credentials systems have been researched resulting in a more complete anonymous credential protocol toolbox and thus better viability of deployments of such systems in the real world. This line of research is aligned with our strategy of making our Identity Mixer anonymous credential system open source and thus available to the privacy community. One result to be highlighted here is a scheme for delegatable anonymous credentials solving a long-standing open problem in cryptography, as well as various efficiency improvements and improved credential revocation mechanisms. On the more practical frontier, we have been working on concepts and designs around a secure wallet implementation for better protecting of a user's private keys and credentials as well as protecting the integrity of computations and credential selection user interface. Considering our results in the field of *user supporting mechanisms* for privacy-enhancing identity management, requirements have been elicited and results have been obtained for transparency support, establishment of collaborative groups, trust management and user awareness. Multiple tools have been made available showing how transparency for the user can be increased for their on-line lives. Transparency-enhancing mechanisms play an important role in the life-long data protection for the end user. In the area of *privacy of data* in large data collections, such as the ones held and being capitalized on by social network providers, we have progressed our research. This work becomes particularly relevant in the light of recent results that claim re-identifiability of "anonymized" data sets from social network providers. It is expected that our results in this area will become highly relevant in practice as the number of privacy incidents increases and re-identification capabilities get stronger. Our work covers privacy metrics for large-scale data collections, as well as the enforcement of data privacy in such data collections based on specified confidentiality constraints. In the latter area, we have built a tool for privacy-aware design and fragmentation of relational data schemas which has been made open source. Our research on *protection of user-generated data* has targeted the limitation of disclosure of users' data, that is, dissemination control and secondary use restrictions, as well as efficient access control to users' data hosted on third-party servers. For dissemination control, we have obtained results on dissemination limitation based on semantics of constraints captured by ontologies and protection of users' location information by preferences-defined obfuscation. For access control to 3rd-party-hosted data we have defined a method and system for access control through encryption and key management to prevent the server from learning the stored data and keys. A prototype of the system has been implemented and made open source.

One major research and development package of PrimeLife comprises *policies*, more specifically, access control and data handling policies for a wide range of uses.. Within the second year we have finalized our comprehensive analysis of *requirements* for privacy-enhancing access control and data handling policy languages that has been

serving as foundation for our policy efforts. *Fundamental research* in the policy area has been continued to provide the basic foundation of our work in terms of policy models and theory. Based on the requirements and results in the model area, we have finalized the implementation of a first policy engine based partially also on work taken over from PRIME. Our final policy language has been designed, capturing our main policy requirements, thus being widely applicable in practice covering both access control and data handling functionality. The design and implementation of our second policy engine for this language has started in the second year. We align its development closely with a widespread standard in order to foster future adoption. We extend a well-maintained open source policy engine for our implementation. The policy engine is planned to be utilized within one of our demonstrators. Overall, the efforts in the first two project years have resulted in a solid set of *requirements* for policy languages as well as the design of the *policy model, architecture, and language* for PrimeLife, and foundational results.

It is widely acknowledged that *usability* is a key factor in the deployment of any user-facing technology. Many of our PETs are user facing and usability is therefore receiving considerable attention. PrimeLife has started its usability work by building on the results available from the predecessor project PRIME while opening up new areas at the same time. From a content perspective, our usability efforts not only support the other activities, where required but also focuses on the following key usability areas: *Identity selection*, *trust and assurance*, and *policy display and management*. In the area of secure and easy-to-use identity selection, we have been exploring the intrinsic properties of the identity card metaphor in the context of anonymous credential systems as underlying identity federation system. Conveying the ideas of data minimization to the user has been found to be a real challenge and has lead to an iterative improvement of our approach for presenting an identity selection user interface to the end user. Both card-based and other approaches have been thought of and tested so far. For the mediation of trust and assurance, a graphical trust evaluation interface has been built and tested in multiple iterations. Work on a data tracking tool allowing users to keep track of their data disclosures and exercise their rights has been improved, with transparency tools in social networks being a specific sub-topic. Throughout the user interface work, for evaluation of results we employ the methodology of cycling through small user tests on variants of an interface and new such tests on resulting updates as this has been shown to be most economic yet effective. In the area of policy display and administration, work on a simplified policy definition has been continued based on the developments on our policy language. Within the second year, a close synchronization with the upcoming policy language has been started to jointly leverage our competence in both those areas. The definition of user preferences “on the fly” has been tested iteratively and work on a policy editor has started. Legal aspects for users defining policies are being considered as important aspects. We researched and tested user-friendly representations of privacy policies. A particular effort was a collaboration with Stanford University on work on policy icons for handling of e-mail content. Data protection requires cross-cutting consideration across technology layers, thereby particularly also appropriate support by the underlying *infrastructures* which must be aligned in their capabilities with the privacy protection goals. PrimeLife drives a research and development effort in the area of infrastructures, split into three different areas. First, we investigate lightweight identity management, adoption models for identity management, economic considerations, and interoperability. Results have been achieved in the area of an identity management interoperability framework as well as a comparison of different identity management approaches. In the second year, particularly the so-called “enabler” concept is discussed in the light of privacy-enhancing identity management. Second, trusted end user devices and hardware tokens, such as smart cards, or a combination thereof, that act as trust anchors in an identity management system, are investigated. Such trust anchors are required in a practical system in order to meet certain security as well as privacy requirements. The project has started to elaborate on both of those. A design for a demonstrator of a trusted mobile device has started, based on the upcoming Android platform and a removable secure element. This captures the user-facing and trusted-hardware-related side of our infrastructure work. Thirdly, the ever more important landscape of dynamic service compositions is investigated in terms of privacy, particularly in the context of the peculiarities of policy composition in such a setting. The research challenge comes from the fact that the usual policy issues get aggravated due to a multiplicity of players being involved in providing a service to a user in such a setting. This effort is closely related to our policy language efforts. A joint demonstrator related to our infrastructure work will be delivered in the next project period, comprising on the one hand a service composition back end, integrating our latest policy engine, as well as the secure mobile front-end on an Android-based device.

Within the first two years of the project, PrimeLife has published a considerable number of papers at international conferences, workshops, and in journals to disseminate its results to the scientific community. The project has made its key decisions towards its prototype strategy to lay the foundation for the remaining period of the project. The decisions on the prototypes to implement comprise a major part of the project strategy, due to our focus on

open source. The focal prototypes for demonstrating trusted user-generated content and on-line collaborations have been delivered. A first prototype demonstrating an integration of advanced privacy-enhancing identity management functionality with a mainstream open-source Wiki application as well as a prototype for demonstrating the feasibility trusted content has been built. A set of prototypes and code components has been made available to the public as open source in order to reach out to external stakeholders with practical results as early as possible. In the second year, the tradition of publishing quality results at major conferences has been continued. On the practical side, tools and demonstrators have been built to show various aspects of privacy in today's on-line interactions.

Making Privacy Live

It is understood that doing research and showing the viability of the results through prototypes is not enough for impacting privacy in already-deployed or upcoming information systems. For this reason, the project is executing a variety of activities to *make privacy live*, that is, to reach out to relevant external parties for influencing existing and future product offerings. Visibility of our results by decision makers in industry and by policy makers is another bold goal of PrimeLife. PrimeLife's outreach activities comprise interaction with and contribution to open source initiatives and standards bodies, interaction with and support of partner projects, and organizing external workshops to act as multipliers for supporting our goals in technology transfer. The project is furthermore organizing two summer schools as a service to the privacy research community and for knowledge transfer. PrimeLife is a member of Europe's Future Internet Assembly and active in the area of privacy and identity management. Overall, the PrimeLife thinking is intended to influence many more than the mentioned activities of our partners and thus facilitate our ideas to get to practice. This will particularly also comprise advertising privacy in conference talks or conversations with relevant stakeholders or decision makers. As a further opportunity PrimeLife has collaborations with a variety of other projects funded by the European Commission, mainly those in the area of privacy, trust, identity management, and formal methods, in order to align parts of our efforts if this will increase the value delivered to Europe.

Early in the first project year, PrimeLife has launched its public Web site which has been one of the pillars of our dissemination strategy. The project will make most of its results publicly available on our Web site. This comprises project deliverables, research papers, and slides of presentations. A comprehensive overview of our open source has been be made available on our Web site to attract potential users or contributors of our technologies. In addition to this, PrimeLife will organize events such as workshops or summer schools targeted at specific stakeholder groups to maximize our impact in dissemination. For making certain project results practically viable, we are working together with selected standardization bodies. This is aligned with our open source strategy to have impact on multiple dimensions. Involvement in ISO activities as well in W3C's Policy Language Interest Group (PLING) have been carried out in the first two project years. Collaborations have been started with other EC-funded projects to leverage complementary competences. In this context, two cluster workshops have been organized jointly by the European Commission and PrimeLife for deepening project collaborations.

Expected Results and Overall Impact

The key overall project results of PrimeLife will comprise relevant research results, prototypes, and open source components. Parts of the source code being built will be made open source, while other parts will remain under the control of the creators and potentially utilized otherwise, or both. The research results are targeted at driving the community in the user-centric identity management space and leading to a take-up of our ideas in the research community. Particularly, our expected results in the field of life-time aspects of identity and privacy have the potential for opening up a new development in the identity management space taking a broader view at identity management. The project results will demonstrate the feasibility of our ideas and will ideally influence product and service offerings of companies and the open source space and support decision makers. This goal is supported by our demonstrator activities in prominent Web 2.0 scenarios such as social networks and online collaboration systems. The social network demonstrators finished in the second year have a potential of driving the adoption of the presented privacy features into real-world social network services offerings or diffusing into open source social network offerings.

The composition of the PrimeLife consortium creates potential for outreach which is unique in Europe and the world by comprising both leading players from industry and academia. We will contribute parts of our source code to selected open source initiatives to directly have an impact in the development community in the open source area and as a next step find our results or derivations thereof in commercial offerings based on open source. Our intention is clearly to trigger a wide-spread take-up of our privacy technologies and implementations thereof. At a larger scale, our efforts are intended to create traction in the privacy space within Europe and the rest

of the world, with Europe providing great conditions in terms of legislation and policy support. Among other benefits, this will help strengthen Europe's thought and technology leadership in the area of privacy. The wider societal benefit of the work carried out by PrimeLife is intended to encompass a substantial improvement of the citizens' privacy in their electronic interactions over the Internet.

Following Successful Traditions

PrimeLife has emerged from the award-winning PRIME project which has marked the beginning of Europe's targeted research in user-centric privacy and identity management. The overall vision of PRIME and PrimeLife are the same: *Allowing people to act in a secure way in the information society while retaining control over their private sphere*. PRIME has mainly been focused on research related to solutions generally applicable to user-centric privacy-enhancing identity management with a particular focus on data minimization while PrimeLife is going beyond this. PrimeLife is targeted at life-time aspects of identity management as well as privacy-enhancing identity management for complex real-life scenarios. PrimeLife builds on technical results produced and experience gained in PRIME as well as the overall traction in Europe and beyond created by PRIME. Both PRIME and PrimeLife help Europe to come closer to offering citizens a protection of their rights of privacy in electronic interactions. This makes Europe unique in the world, particularly when comparing with the United States or Asia. PrimeLife's definition of its research agenda as well as its open approach facilitates the European vision and will facilitate the adoption of our technologies in practice.

Please visit our Web site for obtaining project deliverables, selected publications, presentation slides, project news, and further up-to-date information. Contact the project coordinator for requests and any further details.

Consortium

The PrimeLife Consortium comprises partners from seven EU Member States, Switzerland, and the United States of America. Each partner's unique competence complements the PrimeLife Consortium's joint competence that makes it unique in tackling today's most exciting privacy and identity management challenges. The following organisations are members of the PrimeLife Consortium: IBM Research GmbH (Switzerland; project coordination and scientific and technical lead), Unabhängiges Landeszentrum für Datenschutz (Germany), Technische Universität Dresden (Germany), Karlstads Universitet (Sweden), Università degli Studi di Milano (Italy), Johann Wolfgang Goethe-Universität Frankfurt am Main (Germany), Stichting Katholieke Universiteit Brabant (the Netherlands), GEIE ERCIM (France), Katholieke Universiteit Leuven (Belgium), Università degli Studi di Bergamo (Italy), Giesecke & Devrient GmbH (Germany), Center for Usability Research & Engineering (Austria), Europäisches Microsoft Innovations Center GmbH (Germany), SAP AG (Germany), Brown University (United States of America).

Project Manager

Dieter M. Sommer

IBM Research GmbH, Switzerland

primelife@zurich.ibm.com

http://www.primelife.eu/