

Next Generation of Digital Identity

FULUP AR FOLL, JASON BARAGRY



Fulup Ar Foll is Master Architect in the global software practice of Sun Microsystems, Inc.



Jason Baragry is Lead Architect in SOA/Business Integration in Sun Microsystems, Inc.

While traditional paper based identity and digital identity share the same fundamentals and aim towards similar goals, they nevertheless have some significant differences. This paper provides a high level introduction to the fundamentals of Identity – its technical constituents, its requirements, and its risks. It closes with an explanation of how Project Liberty provides solutions to many of the issues raised by next generation digital identities.

The concept of “Identity” and having to prove your identity is nothing new. A national ID card was introduced in France around 1920, and while we had to wait until 1940 to see its generalisation, it became mandatory for every citizen older than 16 only around 2000. Obviously comparing traditional and electronic identification is a risky business, nevertheless there are some invariants that remain valid whether your ID is paper based or digitally based.

Most identity documents include more information than your basic identification. The first set of attributes is usually information allowing others to determine that you are effectively the person you claim to be – whether this information is a picture, a login/password, a certificate or any other credentials – the final goal is to authenticate you. The second set of information is usually used to apply authorization, e.g. you can drive but only with glasses, you’re a French citizen which allows you to move freely within the Schengen area, your name is highlighted in red and thus you have access to confidential documents, etc. The last set of information is dedicated to controlling the validity of the document in the digital world. This will typically be an electronic signature attached to some revocation list. In the paper world, this would typically be the place and date of issuance plus some serial numbers.

Quite surprisingly, the main target of most identity documents is not to authenticate you, but to authorize you. The general mindset is to first perform authentication; this is mostly because traditionally, in order to access needed personal attributes for a given authorization, you first need to answer the question “who is he?” This is not because attributes necessary for the authorization would not be valid outside of an authentication context, but more because the information flow is such that without first doing an authentication you cannot find revealing attributes.

In order to authorize you someone must be in position of controlling the validity of your claim. While information attached to most identity documents is not very visible and thus not very accessible for most

users, the way you control the information forces a direct interaction between the user and the controller, thus transforming this operation in something very accessible. The perception of what is acceptable or not depends on the context, for instance, young people versus old, Europe versus North America, etc. English people have trouble accepting any identity document and most drivers will find it perfectly normal not to have any ID while driving in their neighbourhood. On the other hand they find it perfectly acceptable to have video cameras on each street corner or to apply DNA tests to immigrants. In France, the DNA test is unacceptable and was recently refused by the chamber of senators, video cameras are never welcome, but on the other hand most people find it normal for the police to ask for an ID card whenever they want. Furthermore, if you cross the Atlantic, you have to provide your fingerprints to enter the country, show your ID card to buy alcohol, etc. Therefore, identity is something personal that may have ramifications for your private life. Depending on culture and history, some control or cross connection can be either very natural or completely unacceptable. This issue is very often outside of any serious requirement or risk analysis.

Until now identity control was mostly a paper and manual process. For this reason collusion between authorization and authentication was not a serious issue. For instance, for most of us, voting is the only time where we expect real anonymity; we do not mind our doctor knowing who we are, or our car insurance company knowing our home address – but what we don’t want is our insurance to know what our doctor knows. In fact, most of us have only very few secrets that we don’t want people to know about. Nevertheless, in order to protect ourselves, we want to make sure that personal information about us will not be cross-border from one sector to another. Until very recently those exchanges were only done manually. They were slow, complex, and could not happen on a large scale. Moving from manual to computer-assisted processes completely changed this paradigm. Technology has moved very fast in the past year and what was impossible a few years ago is

now something a 1000 Euro computer can do. In fact, when looking back we realise that until very recently limits were more due to technology than to legal regulation, and without doubt, had DNA been available during World War II it would have been used. Today technological limits are pushed further and further every single day. It looks more like a revolution than a simple evolution and both mentality and legal frameworks have a lot of difficulty in keeping up. We are moving from a long period where we were limited by “what is possible” to a new period where we want to limit by “what is acceptable”.

The big difficulty with what is acceptable is that, depending on whom you ask the question, you get a different answer. And obviously whatever technologists, politicians, etc may believe, some people will always break legal limits, whether officially or not. Furthermore, the impact of bypassing some of those limits is not visible within a normal human understandable time frame. For instance, many teenagers write very personal things on their blog without realizing that in twenty years from now, when they apply for a new job or a political investiture, that information will almost certainly re-surface. People within rich democracies tend to ignore that our world may still change for better or for worse. They are often ready to sell their privacy for only a few per cent extra discount, allowing airlines, banks, supermarkets, etc to know everything about what they eat, the clothes they like, and the music they listen to. Last but not least, what is acceptable within the current context (i.e. European Nordic countries that use a unique national ID to index almost every single identity record a citizen owns) might not be such a good idea if one day the bad guys take ownership of the country. Implementing privacy aware mechanisms increases complexity, and it is very hard to justify when your immediate context does not allow you to contemplate the consequences of not doing so.

Regardless of what technologists may have you believe, today the strongest limit to digital identity is neither technological cost nor availability. In fact, the necessary technology already exists in the mass market, and as of today almost no single commercial company could exist without some form of electronic identity transaction – the cost of ignoring this technology would be so high that it would bankrupt anyone trying this path. Governments do not have the exact same financial constraints that the private sector has; nevertheless they cannot ignore technology anymore and will have to leverage it in order to both reduce cost and increase the quality of services to citizens. In fact the only remaining constraint is complexity, and as of today this is the strongest limitation of digital ID penetration. While digital ID may lever-

age the same general concepts as traditional identity, it remains very abstract. Indeed, even if basic daily usage is understandable for a significant part of the educated population, as soon something goes wrong, you would very quickly feel like everything is getting out of control. Most of you have already had a refused credit card, usually in an unfriendly environment: for instance a foreign country, the middle of the night, during a train strike, etc. And very rarely can someone explain what actually happened. In fact, most people understand digital ID and electronic transactions like they understand a TV remote control – when it does not work they simply want to change the batteries and if the problem persists they will simply change TV!

Implementing privacy has never been simple, and even in a traditional “paper based” world privacy is most of the time required due to loss of information and not the fact that information is not created. Let’s take an example: imagine that you want to buy a book without anyone knowing that you bought it. If you buy it in a shop, even paying with cash, the salesman knows what you have bought. If you ask someone to buy it for you, then that person knows what you bought. In order to break the information chain, you would need something like the following:

- Take an envelope, place a piece of paper into it with the name of the book you are looking for and enough cash.
- Ask someone you trust to go to the shop for you with the envelope.
- The salesman opens the envelope, reads what you want, and takes the money.
- The salesman puts your book in the envelope, closes it, and hands it back to your trusted courier.
- Your partner sends back the envelope to you.

In fact most of the time this is not necessary, and if you buy a book far enough from your home and in a busy place like a central railway station, it might be more than enough for the information to get lost. Nevertheless this shows that in a traditional world privacy is mostly due to loss of information and not the fact that information is not created. This situation has had a significant consequence, it has allowed the mass population to have an acceptable level of privacy except in specific cases, e.g. criminal investigations, which permit the implementation of special processes to ensure that the loss of information is avoided. Obviously the fact that the information is created has limits, and some governments like the

old East German one have proved that collecting and manually cross connecting all those small pieces of information was possible.

One of the strongest differences between paper and digital ID is that in the digital world it is very easy to automate processes, and what was once only possible for huge organizations like the Stasi in the recent past, is now possible for almost any student in a garage. Today technology provides almost unlimited capabilities for both storage and processing, making research of correlation between small pieces of information far too simple to guarantee privacy. It is a real risk that regardless of whatever legal frameworks are implemented to protect citizens, that privacy as we know it today, will just disappear. It is obvious that in only a few years from now, we will have the technological capability to determine from our mobile phone camera who everyone is, just by taking a picture and asking the question to whatever will be the next generation Google search engine. The Government cannot prevent rain; it is a good question to ask ourselves if they can limit technological capabilities and, if they can, should they?

Knowing that the cross connection of information will be possible, easy, and unlimited, the only option to protect privacy is not to create the information in the first place. In fact, the next generation of digital ID architecture should do more to limit the creation of information than anything else. Limiting the creation of information is hard but not impossible unless you own the system. It is the only way to guarantee that if the ownership of the system does change, then even if the new owner changes all the rules, they will not get the information they want. This is because it will not exist. We should not forget that the only information no-one can steal is that which does not exist.

As we've seen before, identity information can be spliced in three classes:

- Indexes to search and find the information;
- Attributes containing part of the information about the user;
- Control mechanisms to guarantee the authenticity of the information.

In order to limit the risk of correlation, the first thing to do is to decouple the indexes that connect the principal from the attributes that define himself. If we take as an example your medical records, this is a very private document and in many cases you would not want your employer, nor your insurance com-

pany, nor your wife/husband, etc to see it. On the other hand, if we remove your name from this document, then most of us would be more than happy for all the data to be given to a university for research. This is a very good example of what endangers privacy is not the information in itself, but the connection with what we call the principal, e.g. the name and address with the data.

The second action is to make sure that we limit the information itself. The first and most obvious action is to build a process where we only provide necessary information. Most of today's processes request far too many pieces of information. Why does a driving licence require a place of birth? Why would your telephone company need to know your age? Why a hotel your home address? ... Very often when you ask people why they want this type of information they simply cannot provide an answer. In the manual world, collecting the information was very complex and it was somehow understandable that people would ask for more information than needed, just in case they would need that information at a later time. Even if we could argue this, due to the cost of manual correlation being generally unacceptable, the risk to your privacy was probably acceptable. A modern digital ID architecture should allow you to provide only the minimum level of information mandatory for a given action. Furthermore, in many cases the information should even be reduced. For instance, not providing your full date of birth, but just the fact you're an adult. While providing only required attributes for a given service is complex, it leaves the owner of the attribute in a position to choose whether to provide the information or not. Unfortunately, as soon as the attribute is given, you have no idea and no control over how or what for the receiving party will use it. Nevertheless, in most cases you can somehow trust the receiving party to handle the information as you wish, for example they claim not to store it on disk, it should be deleted after six months, etc. This obviously only applies if you have a means of communicating your wishes in a manner that the receiving party will understand. Last but not least, all of this should be done under user control, that is, not by asking a user to renounce any of his rights as is too often the case, but by providing the end-user with real control over what he wants and what he does not want. In fact, a new generation of digital ID architecture would require:

- Separation of identity principal from attributes. Needs a mechanism to get attributes about someone you don't know (i.e. anonymous access to the "adult" attribute only);

- Detach authentication from authorization and attributes exchange. By coupling all those pieces of information together it creates a new big brother.
- Provide mechanisms to prove an attribute is authentic. Many services create copies of information only because they have no way of trusting other parties.
- Allow users to control the overall system, within an adequate level of complexity.
 - Build a distributed system to both make it scalable to country/continent level and to guarantee that if an element is compromised only a limited number of information pieces/users will be affected.
- A way to handle exceptions, tracking of bad behaviour, terrorism prevention, etc.

Last but not least, what about risk and security? The world is not perfect, no system is fully secure, the world is not as gentle as we would like it to be. Identity attributes are critical pieces of information for many actors: government, intelligence, commercial companies, etc. While the number of risk factors is probably unlimited, the following ones seem to be the biggest.

External attack: Quite surprisingly this is probably the easiest one to deal with. Mail spam and virus are probably the best example of external attack. They create useless trouble and cost money, but at the end of the day they do not kill us. I tend to consider external attack as street graffiti; they create useless damage, without comprising national security. In fact the biggest issue with external attack is that this is the most visible one. It is also the most understood one by IT people and for this reason we tend to allocate too much resources to it, forgetting the other ones.

Internal risk: In my opinion the biggest one. This for the simple reason that in order to make operations possible you need at least some of the employees to get access to some of your critical information. There are two classes of risk. First, human error, where something is leaked unintentionally. Second, criminal, where some bad guy finds a cheaper and faster way to buy someone to get the information from the inside rather than trying to get it by hacking your system from the outside.

Change factor: One risk that today is hardly ever addressed. While we may hope that modern democracies will never be replaced by the bad guys, should it not be taken into consideration? More commonly,

when a commercial company changes ownership, how can a given user prevent his personal information being transferred without his consent?

Project Liberty is a technology framework that has been designed to address this class of problems. It might not be perfect, but it is a solution that has been proven to work. Digging deep inside Liberty or SAML technologies is out of scope for this paper, but hopefully even while keeping a high level of abstraction we may explain how it addresses some of these key issues.

Federation: Technically implemented through SAML2 protocols. It is a weak link between different identities, or a given principal. Federation allows us to keep the level of complexity simple enough for end-users to browse seamlessly from one service to another without re-authenticating. Federation is the cornerstone that allows us to decouple the identity of the end-user (principal) from its attributes, allowing a given principal identity to be slotted in many different places in a transparent manner for the end-user.

Identity attributes: Implemented through IDWSF. We have seen that it is very important to separate attributes from the principal identity. Liberty Identity Web Services allow us to discover, request, and retrieve updates of attributes about a given user without needing access to the principal identity. It is designed in such a way that different users may choose different services to hold a given attribute, or even to allow a given user to store a given attribute in two different places with eventually two different values.

Social network: Implemented through people services. There are many cases where you need to group identities, this is either to give special access right to a given group (e.g. allow all parents from a given class to see pictures from a school sport event) or to allow someone to act on behalf of someone else (e.g. an accountant acting on behalf of a company).

User consent: Implemented through Liberty Interaction Service. Allows a given service to request user consent independently of the level of imbrication a given request has.

Identity governance: A Liberty ongoing effort known as IGF (Identity Governance Framework). It is an XACML based layer that, first, allows a service owning an attribute to take a decision on whether or not it can release the attribute to a requesting party, and second, allows a request to be sent attached with the attributes and some metadata for the receiving parties to know how they should/can handle that informa-

tion. In many government cases, the fact that you can release or not an attribute is based on a legal framework, and IGF rules should be a flat translation of legal constraints. On the other hand, the receiving party agrees to respect the constraints attached to requested information (e.g. not allow to write on disk, but be deleted after six months,).

Obviously technology cannot do everything. Nevertheless it should handle privacy as a first class citizen and propose a framework that handles enough of the problem automatically, thus leaving the remaining part manageable outside of technology by manual and legal mechanisms. Obviously nothing is perfect, but not doing anything because we cannot do everything would be criminal. Not only should a given user own the right to verify and change most of the information held about him, but in many cases he should also own the right to simply forget. How a digital ID will allow someone to be a very bad teenager and then years later to be a very nice respectful father, employee, etc., is only one example of what a modern digital ID framework should implement.

Project Liberty was designed to handle this global problem. Some people complain that it is too complex or too restrictive. However, we had to make it complex enough to support the complexity of our world and we chose to make it restrictive to keep the cost of adoption acceptable.

About Project Liberty

The Liberty Alliance is a global identity consortium formed in 2001 by approximately 30 organizations with the goal of developing open technical, business and privacy standards for federated identity management. Liberty Alliance achieved this goal in 2002 with the release of Liberty Federation and in 2003 released Liberty Web Services, an open framework for deploying and managing a variety of identity-enabled Web Services. Having grown to nearly 150 members from around the world, the Liberty Alliance is currently working toward developing ID-SAFE, the industry's first open framework for deploying and managing interoperable strong authentication. The Liberty Alliance is the only global identity organization approaching identity issues from a holistic perspective, addressing the technology, business and privacy aspects of identity management in order to build a more trusted global Internet for consumers and organizations worldwide. Liberty Alliance background information including a listing of timelines and industry milestones is available for download from <http://www.projectliberty.org>.

About Sun Microsystems

Since its inception in 1982, a singular vision – “The Network Is The Computer” – has propelled Sun Microsystems, Inc. (Nasdaq: SUNW) to its position as a leading provider of industrial strength hardware, software and services that make the Net work. Sun can be found in more than 100 countries and on the World Wide Web at <http://www.sun.com>.

Fulup Ar Foll holds a Master degree in Computer Science from the French Military School. Before joining Sun he was a research engineer for ten years on distributed technologies for the French Department of Defence and he taught Internet and Java technologies for six years at South Brittany University. For Sun he has been the lead Internet Architect for many projects related to European Telecom operators (France Telecom, Orange, Vodafone, Telefonica, Turkcell, T-Systems, ...), as well as for other strong identity and security infrastructure users such as banks and governments. In the recent past he helped France and Norway to move toward the Liberty Alliance Federated model. He is currently Master Architect inside Sun global software practice and focuses on high scale federated identity issues. He represents Sun software customer service group inside Liberty Technology Expert Group standardization committee, and inside OMA (Open Mobile Alliance) MWS (Mobile Web Services) group and works as Lead Architect for major identity projects on a world-wide level. He has also been speaker at many international conferences.

email: fulup@sun.com

Jason Baragry holds a PhD in Software Engineering from La Trobe University, Australia. He is based in Norway and is the Lead Architect in SOA / Business Integration for Sun in Central and Northern Europe. He has been an Architecture Advisor for many government and telco projects both in Norway and in the wider Nordic area.

email: Jason.Baragry@sun.com