



Not Just a Spectator Sport



Identity Management in the Participation Age

By Robin Wilton, Sun Microsystems

We are at a turning point in the evolution of human communication. We have reached the end of the Information Age—the age of passive consumers who view the Internet as the “great database in the sky.” We are now entering the Participation Age—a time of networking where participants aren’t just acquiring information—they are beginning to contribute to it.

Today, weblogs allow anyone with network access to communicate to the whole Internet community in a way that turns the traditional view of mainstream mass-media publishing upside down. New blogs are appearing at a rate of one every second.

Similarly, we’re in a time when wikis—perhaps best exemplified by the Wikipedia (<http://wikipedia.org>)—allow users to cooperatively maintain vast quantities of information. And, open source now allows anyone to contribute to the evolution of an increasingly wide range of software. We have become active participants in an online world—one that transcends peer-to-peer and client-server.

USHERING IN THE PARTICIPATION AGE

Beyond the impact on mass-market publishing, the Participation Age is having a social impact because it allows mass-market communi-

cation technology to enable a similar inversion of the socio-political model.

Two powerful examples from 2005 demonstrate how the public can now be mobilized—and can express opinions—in a way which is visible, easy, and immediate. The rock group **U2** and front-man **Bono** used short message service (SMS) messaging to count supporters as part of a *Make Poverty History* campaign (www.hos.horizon.ie/pressroom/pr280605.html). The *Live8 / Long Walk to Justice* events also used text messaging to secure pledges of support from millions of people (www.live8live.com/).

Every company delivering network services—commerce, financial management, personal ads, photo sharing, gaming, and so on—find their static “information-only” relationships with users changing to become much more interactive. The trend is gaining depth and speed as these services allow users to participate constantly on a wider variety of edge devices.>>





TIPPING THE SCALES

Globally, the Participation Age is about changing the balance between the 1 billion richest inhabitants of this planet and the other 5 billion. Today, only 14 percent of the world's population is online, and yet those of us in the online community are seeing the profound effect that technology can have in terms of enriching our lives, opening up our communities, and stimulating our economies. The Participation Age is

about extending the edge of the network, and capitalizing on the fact that wherever it extends, there's the opportunity to create and empower new participants.

At a recent discussion on the Participation Age with a United Nations panel, **Andrew Zolli**, futurist-in-residence at National Public Radio and *National Geographic*, noted, "In that kind of a world [in which switching service providers is quick and easy], the way in which you hedge against losing everybody in this incredible wave of commoditization,

is by increasing participation.

"So you see this flood of new ideas that have to do with open source. And that idea of open source transcends the idea that people are merely consumers, it transitions them in the view of the company that services them [from] consumers to citizens and participants.

"What's really important here is the exchange of values in both directions. It's not merely that we're empowering in one direction, it's that we're empowering in both directions. Companies get better faster and the greater the degree of participation between people who make things and people who use them, the less likelihood those two people in that relationship are likely to break that bond."

David Kirkpatrick of *Fortune* magazine, the event moderator, said, "There is a change in metaphor. You're hearing a kind of underlying change in metaphor and we talked about it in terms of shifts from the center to the fringe and the empowerment of the edges in a flat-

tened network and people collaborating across great distances inside their network. That is to say people collaborating five, six, seven, eight, nine hops away to create value."

THE NETWORK IS THE COMPUTER

It's possible to create value within a single enterprise, between two enterprises, or even across arbitrary chains of participation by lowering the cost of forging the next link in the chain. Indeed, there are companies that could not exist without this kind of value creation. **Sun Microsystems'** seminal assertion that "The Network is the Computer™" becomes ever more apt as the edge of the network becomes more populated and diverse, and now we are seeing "Everything of Value Connected to the Network" come to life.

Those "things" connected to the network also illustrate the characteristics of the Participation Age. Think back to the '80s, when a network was a rigid hierarchy of logical and physical units. The physical units consisted of computers and terminal devices. Initially, the logical units were just virtual addresses for physical units. Eventually, logical units came to include applications. The relationships between network units were still overwhelmingly hierarchical and usually rigidly defined. And to some extent, assertions of identity were superfluous—a device with a valid network node was implicitly trusted.

Now, the "things" connected to the network include people, Web services, and devices of increasing cleverness such as phones, PDAs, vehicles, and radio frequency identification (RFID) tags. Few of these things expect to be treated as mere dumb physical units. They expect to participate in bilateral application-level dialogues, and to share their own processing capacity with that of the other things to which they are connected. Grid computing, sensor networks, and mesh networks of tiny wireless devices all illustrate the principle of networks as flat, dynamically configured, loose structures of collaborating devices.

The network has become the basis for applications that no longer consist of monolithic, server-centric executable code. Rather, an application becomes an assembly of users, devices, and services which addresses a specific functional need at any given time.

Consider **HousingMaps** ([Regardless of whether it's a commercial or public sector service, providing value usually implies some form of liability. Without the notion of identity, it's hard to make any kind of liability stick.](http://www.hous-</p></div>
<div data-bbox=)



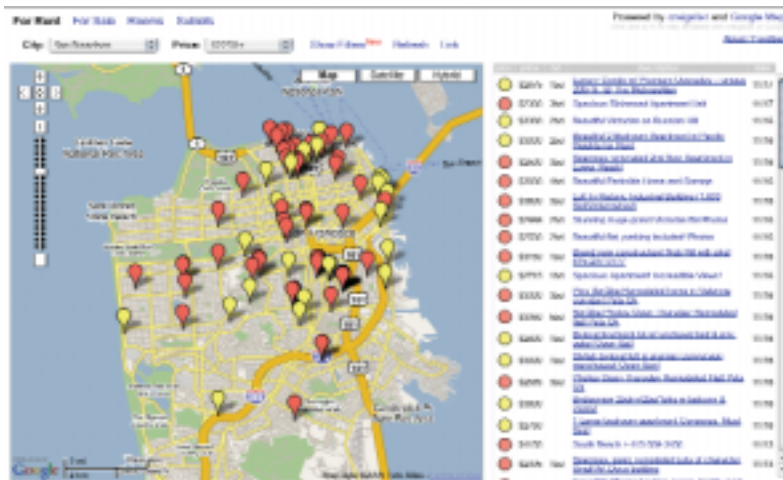
ingmaps.com) (see Figure 1). This Web application—what is called a “mash-up”—integrates listings from **Craigslist** (www.craigslist.org) with mapping data from **Google** (http://maps.google.com) to produce a geographical view of available living accommodations in a given area. Remarkably, HousingMaps is not affiliated in any way with either Craigslist or Google; it merely combines their public interfaces to synthesize a whole that is more than the sum of its parts.

If we look to the enterprise data center, we can clearly see the hallmark of the Participation Age in service oriented architectures (SOAs)—loosely coupled services enabling an agile IT response to rapidly changing business requirements. However, such an architecture does not exist in pristine isolation. As **EDS** Executive Vice President **Charlie Feld** put it recently, “So many enterprises have built up decades of complexity in their legacy systems. In addition to all the security and privacy issues, it’s really quite a feat to get these infrastructures to support SOAs and the explosion of edge devices we’re seeing today.”

In response to this challenge, we see IT departments re-factoring mainframe systems away from dumb terminals and screen-scraping interfaces, and providing business functionality as Web services in an SOA. This move to loosely coupled, course-grained interfaces, accessible via standard protocols, has resulted in some surprising new integrations along the lines of the HousingMaps example above.

For example, Sun Microsystems has long maintained an enterprise directory, accessible via LDAP, and calendar services, accessible via HTTP. The fact that these services are universally available throughout Sun via standard protocols made it easy for a system engineer, in his own time, to develop a white pages application called Namefinder (see Figure 2), that allows Sun employees to search for any Sun employee and retrieve his or her contact details, photo, and daily schedule. Despite originally being developed as a “skunkworks” project outside Sun’s IT department, Namefinder is now deployed throughout Sun, and even ships as part of Sun’s Directory Server Resource Kit.

Figure 1



“WE ID”

As we focus on the impact of the Participation Age on the enterprise and its IT department, we see that identity becomes a key issue. To unlock the “value chain” it’s essential to have a flexible, portable, secure notion of identity so that the full range of customization, access control, and business models is available.

Mainframe applications previously required a login when they were accessed via dumb terminals. Even with the move to Web interfaces via screen-scraping, the user still had to provide credentials to access the application. In turn, the application itself maintained its own list of authorized users—though over time, this function was often delegated to enterprise-level access control applications.

Legacy applications have subsequently been “Web-ified” by a variety of means, though in doing so, identity was often overlooked in the rush to get Web services online. Architects often made the assumption that confidentiality and authentication at the wire level—typically via SSL/TLS—would suffice. They were wrong. The old paradigm, based on remote execution models such as remote procedure calls (RPC) or object-oriented remote method >>






invocation (RMI), tended towards the view that if the request that arrived was correctly formatted, that was proof enough of authenticity.

In cases where the requester is a human being, that's not a very good match for the way trust works in real-world transactions. In the growing number of instances where the requester is another application, it becomes a very risky way to expose your services online. Regardless of whether it's a commercial or public sector service, providing value usually implies some form of liability. Without the notion of identity, it's hard to make any kind of liabil-

Figure 2



ally true of almost every valuable transaction, since typically some kind of auditing and identity tracking needs to be involved for billing, security, access control, or regulatory compliance purposes.

If we look at the wider system, the user typically establishes identity via a desktop or portal login. As services cluster together into applications, the user's identity must be transported between services. A real-world use case illustrates the requirement.

The **Business Industry Political Action Committee** (BIPAC, www.bipac.org) provides businesses with political research and analysis through publications, conferences, and consulting. BIPAC provides an online service where employees of member organizations can find information relating to political participation, such as online voting, voting records for their political representatives, and so on. The information legally available to an employee depends on their position in the enterprise, shareholder status, and other factors. In order to give users accurate, legally appropriate information, BIPAC must know several pieces of data (attributes) relating to the user, including their home zip code (to assign the user to electoral districts) and shareholder status.

Each BIPAC member organization—the user's employer—has all of this information, but needs to be able to share it with BIPAC without divulging the user's name, since BIPAC does not necessarily need to know this to provide their services.

As employees access the BIPAC site from within the enterprise, both BIPAC and the member organization need to ensure that only valid users gain access—but want to do this without exposing users' corporate user IDs and passwords, or requiring them to log on twice with different IDs and passwords. Additionally, BIPAC is not in a position to dictate the IT infrastructure that each of its partners must use for all this attribute-sharing and mutual authentication. The answer to their requirements came in the form of a standard protocol, namely the Liberty Alliance's Identity Federation Framework (ID-FF).

ity stick. Whether we're looking at auditability of commercial transactions, governance of public services, or the security of borders and citizens, identity turns out to be indispensable.

Consider one of the most widely used examples of Web services: the stock ticker. It's a great example because it's concise—send a ticker symbol, receive a quote—and it's almost universally understood. But, if we look at real-world stock quote services, the situation is a little more complex. For example, anonymous requesters get delayed data, while authenticated and paid-up users are allowed to access real-time quotes. So, in order to return appropriate data, the service must require every stock quote request to be made in the context of an identity. This is gener-



STANDARDIZING IDENTITY

The BIPAC example shows us that a key enabler of the Participation Age is standardization. TCP, HTTP, HTML, SOAP, GSM, SMS—if you want to give the widest opportunity for participation, you need them all. Standards enable us to provide applications spanning whole classes of systems, such as a user requesting a share deal via a Web page and receiving a confirmation message to be digitally signed on their cell phone. Standards mean that it doesn't matter who makes the server hardware, the Web server software, the cell phone base station, or even the cell phone itself—it all needs to work together to deliver an experience that is more than the sum of its parts.

Let's take a closer look, then, at modern identity standards. The classic identity problem, now well addressed in the standards world, is single sign-on: logging in to one application and then being able to use additional applications without being required to re-authenticate.

Enterprises commonly provide Web-based single sign-on internally via products such as Sun Java™ System Access Manager. These products typically use proprietary mechanisms based on browser cookies to allow and control user access to intranet resources based on a single login. The scope of a user's access is typically limited to a single domain.

A FEDERATED APPROACH

Federated Web single sign-on extends this process so that employees can log in to the company portal, click on a link to their 401(k) retirement account provider, and manage their retirement account without logging in again, or providing any other identifying data.

The benefits of this are profound. Employees gain the convenience and added security of single sign-on, and have one fewer username/password combinations to remember, while the employer and the 401(k) provider achieve cost savings through a reduction in help-desk calls requesting password resets. It also avoids the requirement for the 401(k) provider to keep and maintain a list of the valid user IDs and passwords of client employees. Importantly, federation is an opt-in process. Employees who would rather keep >>

The Impact of IT

"In the last 25 years, we have been living through the 'Information Age,' so named because of the impact information technologies have had on our lives. It's a valid label, as the commerce of information today represents a huge percentage of all economic activity in the world. Millions upon millions of people produce information, refine it, store it, and distribute it; billions consume it in the same way we consume air, food, and water. Some of us even suffer withdrawal symptoms if denied access to it.

Unfortunately, though, there's one thing wrong with this world view: The Information Age is so last millennium.

Get past it!

Welcome instead to the 'Participation Age.' Advances in technology have made it possible for more and more people to connect with each other to participate and to share work flows, to compete for jobs, to purchase goods and services, to learn and create.

Information Age thinking says, 'Control the creation and distribution of information and you dominate markets.' Participation Age is the antithesis of that. It's all about access. That access allows for value to be created through networked human beings who share, interact and solve problems. Because of participation, meaningful content, connections, and relationships are created like never before.

In the Participation Age, there are no arbitrary distinctions between passengers and crew, actors and audience. Be one, be both, be everything in between.

Welcome to the revolution."

—Scott McNealy
CEO, Sun Microsystems





their standalone 401(k) account access are free to do so.

Similar use cases abound within enterprises, where federated Web single sign-on is used to provide access across divisions, eliminating the need for a single, global “master directory.” In fact, analyst studies show that the majority of federation deployments are intra-enterprise. The typical “IT junk yard” contains a variety of third-party and homegrown identity directories and databases, and the problem can grow worse with every merger and acquisition.

Standards are crucial for federated Web single sign-on, given the difficulty and expense of tightly coupling systems across departmental and organizational boundaries. Without standards, the choice would be between imposing a vendor decision on every participant, or expensive point-to-point integrations between disparate proprietary technologies. With standards, you only need to impose the creation of a front end that conforms to a “contract” for taking part in a well-documented, widely supported protocol.

Web single sign-on is an area well-plumbed by standards efforts. “Efforts” here might sound ominous, given the cliché: “What’s so good about standards is that there are so many to choose from.” But in fact the news is good: Because of a massive convergence project completed this year among three players (OASIS, the Liberty Alliance, and the Shibboleth initiative), there is wide agreement about the right standard to use.

The **Organization for the Advancement of Structured Information Standards** (OASIS, www.oasis-open.org) works to drive the development, convergence, and adoption of e-business standards. The technical committee on security services within OASIS originally defined the Security Assertion Markup Language (SAML) in 2002 to provide a standard mechanism for conveying security assertions between systems (for example, “User ‘Bob’ has the ‘Manager’ role” or “User ‘Mary’ logged in using her smart card at 9:07 this morning”), with a focus on single sign-on as the first interoperability use case.

The Liberty Alliance (www.projectliberty.org) is a consortium representing over 150 diverse organizations from around the world, including telecommunications, financial, computer, and health care companies, government

agencies, and identity vendors. It was created in 2001 to address the technical, business, and policy challenges around identity and identity-based Web services. For example, it offers business and legal guidelines on how to enable privacy, the management of identities, and the complex technical needs of environments where anonymity must be preserved.

Building on the foundation of SAML, the Liberty Alliance developed a full-fledged standard framework for Web single sign-on and identity federation as well as an identity Web services framework called ID-WSF. Liberty ultimately donated ID-FF back into OASIS as the basis for future SAML work, while it continues to refine and enhance its Web services solution. OASIS went on to produce SAML Version 2.0 in March 2005, reflecting not only Liberty’s input but also the requirements of Shibboleth (<http://shibboleth.internet2.edu>), a software solution for identity federation provided by the higher-education Internet2 initiative. Today both Liberty and Shibboleth are in the process of moving to a SAML V2.0 basis.

However, SAML is not the only game in town for federated Web single sign-on. WS-Federation is a similar specification originating from a group of vendors including Microsoft and IBM. Microsoft plans to support WS-Federation for federated Web single sign-on in the upcoming **Active Directory Federation Services** (ADFS).

Recognizing that many enterprises are deploying Liberty and SAML right now and will require interoperability with WS-Federation when ADFS arrives, Sun and Microsoft worked together to define a pair of specifications providing interoperability between disparate federated Web single sign-on schemes. The twin results, Web Single Sign-On Metadata Exchange Protocol and Web Single Sign-On Interoperability Profile, were published in May 2005 (<http://developers.sun.com/techtopics/identity/interop/index.html>).

The companies jointly demonstrated the technology in action at a press event, with Sun and Microsoft engineers each logging in to an identity provider on their “native” platforms and then gaining single sign-on access to applications (service providers) based on the other’s platform.

Let’s take a closer look at how this works using the following example: Bob is a manager



in the IT department at Harrington Insights, a computer industry analyst firm. Trey Research supplies Harrington with computer hardware. Bob needs to access Trey's online order application to purchase supplies. Harrington's employee portal is built on Sun's Java Enterprise System, which implements Liberty's Identity Federation Framework (ID-FF) for federated single sign-on. Trey has a Microsoft environment; the online order system is a .NET application that uses Active Directory Federation Services' WS-Federation implementation for the same purpose. The Sun and Microsoft environments must agree on a common protocol to interoperate. Here's how that happens:

Step One: After Bob logs on to the Harrington employee portal, he clicks on a link from the Harrington (identity provider) employee portal to the online order system at Trey (service provider). The online order system can't allow Bob access (yet), since it has no record of his identity.

Step Two: At this point Harrington and Trey use the Web Single Sign-On Metadata Exchange Protocol to agree on a protocol to use for single sign-on. (The Web Single Sign-On Interoperability Profile lists ID-FF and WS-Federation as alternative protocols.)

Step Three: Harrington then creates an identity assertion, digitally signs it, and sends it to Bob's browser. Bob's browser repeats the request from step 1, but this time, it includes the assertion in an acceptable form. Since the assertion states that Bob is a Harrington employee and Trey trusts the Harrington identity provider, Trey's system can allow him access to the online order application.

This process is completely transparent to Bob; he clicks on a link and gains access to the online order system. The protocol negotiations take place in the background, and an assertion of his identity, including his role as a manager vs. a regular employee—but not his password—is transmitted silently and securely between the two systems to allow the appropriate level of access.

Since releasing these specifications, Sun and Microsoft have established an online feedback forum for interested participants to collaborate on refining and ultimately standardizing them (<http://groups.yahoo.com/group/WebSSO-Workshops>).

CIRCLES OF TRUST

Web single sign-on is a fairly basic example of adding value by making the exchange of identity information both open and interoperable. When it's possible for identity to pervade a Participation Age network, Web applications and services themselves become participants in what the Liberty Alliance calls circles of trust—federations of service providers in which technical, business, and legal agreements provide a whole new level of customization, access control, and privacy. In your work-life persona you might participate in an enterprise circle of trust to use customized employment-related services, such as the BIPAC example mentioned earlier, or department purchasing programs. By contrast, in your home-life persona you might participate in a personal circle of trust in which you can use your online wallet service to securely (and even anonymously) pay for personal items bought online from other services. Within each circle, information about you becomes an online phenomenon—a network identity.

In many realms, whether business-to-consumer, across enterprises, or even within a single enterprise, federated identity is showing itself to be the best basis for overcoming users' issues with secure identity management. It's also proving to be a better solution than centralized approaches to large-scale identity management, which tend to fail at providing workable online analogs for real-world trust relationships.

There are valid reasons why the early models for request/response authentication are not good ways to map real-world trust onto the online services of the future. For instance, the older models tend to assume that authorization is binary. Once you have authenticated yourself (that is, proved your identity), you're in, and you can do what you want. The argument to the contrary has best been expressed by the Jericho Forum (www.open-group.org/jericho), which points out that today's enterprise security model, rather than resembling a fortified castle, is an open structure where the perimeter is permeable, and users' actions must be controlled once they are inside. >>





Second, previous models for authentication and authorization tend to assume that all authentications carry equal weight. Biometrics and digital certificates aside, logging on to your airline’s loyalty program, your local video store Web site, and your online bank account all use the same technology (user ID and password), so why not just have single sign-on between them and be done with it? Of course, the issue is that there’s an implicit “gradient of trust” between those service providers. If you log on to your bank first and then link over to the video store, you might expect them to trust that it’s you. By contrast, if you log on to the video store’s site first and then link over to your bank, you wouldn’t expect the bank to let you start transferring funds without re-authenticating yourself.

Finally, here’s a handy test you can apply to your own IT infrastructure: Does it allow you to manage your employees, customers, and partners as people, or does it force you to manage them as poorly linked clusters of computer accounts? Managing them as people—with rights, responsibilities, roles, entitlements, etc.—requires a more sophisticated data model than the one described above.

SINGLE SIGN-ON OR INTEROPERABLE AUTHENTICATION?

So is single sign-on a mirage? No, but it needs to be acknowledged as a carefully qualified first step. Let’s start by calling it something less catchy but more accurate: interoperable authentication. It’s a way of exchanging identity assertions between organizations—nothing more, nothing less—and each service’s idea of proving your identity may differ from that of the site you just left. The point is that once authentication has been dealt with, we can move on to authorization.

Authorization says, “Now that you have done something to prove your identity, what am I going to let you do?” That decision may be based on a number of things. What user profile does your authentication unlock? What history of previous transactions do I hold, and which can I associate with you? What information about you can I get from other sources?

A useful analogy would be the distinction between a passport and a visa. A passport is one form of assertion of identity. When you present



it at the immigration desk, the officer can compare the photo with your real appearance, and judge whether the passport itself looks genuine. You still might not be admitted to the country, though. Your assertion of entitlement comes in the form of your visa, which effectively says, “The holder of this passport has also met the criteria for entry into the country.” Then there’s other information about you, perhaps held by your own country, perhaps held by the one you’re visiting, which allows them to refine that judgment. Do you have a criminal record? Are you a known money-launderer? None of that would appear in either your passport or your visa, but is relevant to the decision to grant or deny you access.

Taken as a layered data model, based on open standard specifications, and backed up by contractual agreements, this gives us a much better basis for trusted online interaction.

This is the province of Liberty-style circles of trust. The whole concept of a network identity is based on the assumption that information about you will be held by multiple online entities. This contrasts with many current architectures, which try to avoid the issue by simply pretending that this is not the case now, or will not be in the future. A recognition of the realities of network identity ushers in two linked concepts which will define the identity technologies of the next few years.

The first is user consent: If you are to be “defined” to multiple online service providers in terms of the data they hold about you, then you should have the right to say which data they





may or may not share with each other. Note that this is not the same as insisting that you (and only you) should be allowed to hold that data. By analogy, think of the way you use your bank: It's a secure and convenient way for you to store credit, and to transfer it to service providers whenever you need. There is no need for you to keep all your cash in a sack under your mattress. The important thing is that the bank should preserve the integrity of your deposit, and ensure that only you are able to transfer credit to anyone else.

The second concept is respect for user privacy. Again, if third parties are to be the trusted custodians of data on your behalf (and there are good reasons why they should continue to hold this role), then you need to be able to trust them to keep your private data private. In some cases, this may even extend to letting you make anonymous transactions. To use the banking analogy again: If you make an online payment, there is not necessarily any more need for the bank to tell the merchant your identity than there is if you make a cash payment in a store. With a cash payment, the merchant knows they are getting fair compensation for the goods, and has no need to know your identity. The same principle can apply to online transactions. A bank should be able to make an assertion of your creditworthiness without having to reveal your identity at the same time.

One of the most compelling motivations for these circles of trust may come from the government sector, given the strong need for e-government initiatives to "get it right" on con-

sent, privacy, and citizen trust issues.

A recent article in *Information Age* (www.infoeconomy.com) sums up these related issues with the following penetrating argument in favor of federation, taking U.K. national identity cards as a case in point. Andrew Lawrence wrote, "In the case of ID cards, the grand, centralized, technological solution being proposed has already drawn much criticism. Aside from the obvious civil liberties objections, many are puzzled at the [U.K.] government's determination to take a centralized approach, rather than a distributed one based on legacy investments. With dozens of ID schemes in existence (bank cards, driving licenses, passports, etc.), why not legislate to encourage or allow greater harmonization of these, rather than start all over again?"

Open, standard federation opens the way to interoperable authentication within and between organizations. A federated system makes it possible to give users control over how and when their identity-related information is exchanged. Federation allows real-world concepts such as consent, privacy, and trust to be reflected in technological implementations. In short, an open, standards-based federated approach makes it possible for enterprises to manage their employees and customers as people, not computer accounts, and for governments to treat their citizens as people, not identification numbers.

Technically, it might be possible to sacrifice one or more of these and still have a working system. However, there is more to the Participation Age than just the technology of identity management. Socially, economically, and globally, the technology is necessary but not sufficient. Without it, we can't evolve from the Information Age. With pieces of it, the developed world can benefit from participation. With a federation based on a fully open set of standards and an understanding of the underlying principles, the Participation Age can become a global reality. [s]

About the Author: Robin Wilton is corporate architect of Federated Identity for Sun Microsystems. Sun's Pat Patterson and Eve Maler contributed to this article.

