TRUST AND ON DEMAND: ENABLING PRIVACY, SECURITY, TRANSPARENCY, AND ACCOUNTABILITY IN DISTRIBUTED SYSTEMS

Michael R. Nelson, IBM mrn@us.ibm.com (Washington, DC)
Matthias Schunter, IBM Research, mts@zurich.ibm.com (Zurich, Switzerland)
Michael R. McCullough, IBM mccumich@us.ibm.com (Washington, DC)
John S. Bliss, IBM jblisslv@us.ibm.com (Las Vegas, NV)

Abstract

The deployment of gigabit-per-second networks; the development of Grid computing, peer-to-peer applications, and Web services; the spread of wireless networks and pervasive computing; the development of autonomic, self-monitoring systems; and the adoption of open standards that enable tighter integration of IT systems and business processes within companies and across companies are ushering in the new paradigm of distributed computing. Companies will be able to integrate their corporate IT systems into grids, making it far easier to move data within the company and fostering collaboration between employees. In time, it is likely that companies and organizations will link their networked systems with those of business partners and with "utility grids" and data centers run by third parties. One barrier to the adoption of this model is that customers may doubt that their personal data can be protected when it is moving between multiple servers and storage systems scattered across the country or even around the world. In Canada and other countries this is becoming a serious policy issue and there have been proposals to ban out-of-country outsourcing of government IT services.

In this changing landscape, today's privacy solutions — whether technologies, services, or regulations — will be unable to meet the needs of either individuals or corporations. Entirely new classes of products and services are badly needed. The paper explores new types of privacy-enhancing technologies to address employees' and customers' privacy concerns. Some questions to be examined by this paper include:

> How will privacy and expectations of privacy change when distributed computing is commonplace?
>
> How can privacy and transparency be built into distributed computing systems?
>
> What can governments do to protect privacy without inhibiting the growth of distributed computing (e.g. virtualization, Grid computing, peer-to-peer applications)?

Introduction

In recent years, a number of information technology companies, including IBM, have embraced a vision of a new way of providing computing resources to companies, governments, and other organizations. The names used vary—distributed computing, On Demand Business, the Adaptive Enterprise, or utility computing—and so do some of the details, but there are several common themes: (1) Companies will have almost instant access to the computing resources (computing cycles, software, and storage) they need when and where they need them — rather than having to wait to install hardware and software, (2) More and more functionality will migrate from the desktop and the laptop onto distributed systems linked together over the Internet or intranets, (3) Employees will have far better and faster access to the crucial data they need to do their job, (4) Business processes will be integrated across business units and even across different companies (accelerating the outsourcing trend), (5) Computing systems will be more reliable, more secure, and easier to manage, and (6) Companies will pay for only the computing resources they use—rather than buying computer systems.This will trigger a shift from long-term strategic partnerships to short-term ad-hoc value networks – from contract-backed trust towards virtual enterprises with far higher trust requirements.

The implications of this shift will be incredibly disruptive. To quote Nicholas G. Carr, this shift "will overturn strategic and operating assumptions, alter industrial economics, upset markets and pose daunting challenges to every user and vendor. The history of the commercial application of information technology has been characterized by astounding leaps, but nothing that has come before—not even the introduction of the personal computer or the opening of the Internet—will match the upheaval that lies just over the horizon."[1]

In addition to the huge economic implications of the new distributed computing paradigm, it also poses a number of thorny policy challenges, which if not properly addressed could hinder the growth of distributed computing and prevent companies and individuals from enjoying the benefits of this new, cheaper, more reliable, and more flexible approach to computing. Privacy and security are likely to be the most difficult — and most emotional and political—issues to be addressed. In a world where organizations are not defined by buildings and employees but instead consist of loose networks of services provided by a dynamically changing group of contractors and subcontractors and those services are running on a network of machines scattered around the country and around the world, customers will need to be convinced that the privacy and security of their personal data can be guaranteed. Already, the Canadian government and other governments are investigating the privacy implications of the offshoring of personal data. A recent scandal involving CardSystems Solutions, an American credit card processing company, that retained data from more than forty million credit card users in violation of its contract with MasterCard, Visa, and other credit card companies highlighted the difficulty of securing private, personal data when it is passed from one company to

[1] Nicholas G. Carr, *The End of Corporate Computing,* MIT Sloan Management Review, Spring 2005, pp. 67-73

another. Clearly, if the potential of distributed computing can be realized, it will enhance collaboration and provide users with access to far more computing resources than they have today. However, that requires a higher level of trust among the collaborating partners than is normally present today (see Figure 1).

Used effectively, distributed computing could enable companies and governments to better respond to consumer demand, increasing globalization, unpredictable competition, and emerging security issues. It could also allow new competitors to challenge industry leaders.

This paper describes the technologies and standards that are enabling the rise of distributed computing and describes some of the ways in which new privacy-enhancing technologies and business practices might help address the privacy concerns it raises. In addition, we examine how perceptions of privacy may change and how companies will have to find new approaches to winning the trust of their customers.

The Technological Foundation of Distributed Computing[2]

A number of new technologies and standards provide the foundation upon which distributing computing systems will be built in the future. Those systems in turn provide the underpinning for On Demand Business and the new business models made possible when companies can respond more quickly to new opportunities and customer needs. The key technological building blocks include:

Service-Oriented Architecture (SOA) enables a company's IT and business functions to be broken into a series of interconnected services. Activities ranging from employee benefits and travel that serve a company's employees to sales, delivery, and help desk support that serve the company's customers can be standardized, made more flexible, and more adaptable to changing business conditions and challenges.

The glue holding these components together is a set of software components built around open standards and designed to connect disparate processes and systems. This lowers the cost of designing and deploying new systems, while giving a company maximum flexibility to redesign, reconfigure and outsource its business functions as necessary.[3]

Virtualization. Today's IT infrastructures are built out of disparate resources—legacy and new equipment, running different operating systems and different applications. Virtualization connects these assets—not just the computers themselves, but storage systems and networks—to form a single pool of resources that can be accessed and managed across an entire organization, within a single location or around the world.

[2] Cf. IBM's On Demand Web page http://www-306.ibm.com/e-business/ondemand/us/teamperformance/infrastructure/infrastructure_guide.shtml

[3] Two of the best papers on SOA are John Hagel, III, and John Seely Brown, *Your Next IT Strategy*, Harvard Business Review, October, 2001, and John Hagel, III, and John Seely Brown, Service Grids: The Missing Link in Web Services (available at http://www.johnhagel.com/blog20030716.html)

The result is a less complicated, more easily and efficiently managed network. This lowers management costs and empowers employees, while helping organizations serve more people and extract more value from their older systems.

Autonomic computing systems have the ability to monitor, troubleshoot and manage themselves. They allow companies to realign their infrastructure automatically, using pre-established rules based on an organization's business policies and objectives. These systems are more resilient, responsive and secure—and because they can perform their own routine maintenance, they can free IT staff to work on higher-level projects to improve productivity and provide employees with new tools and information resources.

Grid computing takes open-standards integration a step further. It joins disparate systems into a single, tightly-knit resource that acts as a single, powerful computer. Within a single company or organization, grids allow the "harvesting" of unused computer power—from idle desktop computers to underutilized servers—to more quickly and cheaply tackle massive computer processing problems and meet storage needs. In the future, third party grid providers will use global grids to provide computing resources to their clients in locations around the world.

Many of the concepts related to distributed computing have been around for years. However, in just the past few years there has been a huge increase in interest in actually applying them for routine corporate computing functions (and not just niche research and supercomputing projects). This is happening because of two important developments:

Open standards—like the WS standards for Web Services and the Globus middleware for Grid computing—which provide the glue that holds the infrastructure together. They allow heterogeneous systems and software to speak a common language, forming seamless networks that can bridge technological and organizational boundaries. Various efforts at the Global Grid Forum and other standards groups are developing the tools needed to manage distributed systems.[4]

Faster, more reliable, and more affordable networks—The dramatic reduction in the cost of megabit-per-second network connections makes it much easier to provide offices with instant access to remote resources such as the Grid while enabling the sale of unused resources.

Privacy and Security Implications of Distributed Computing

Bringing together disparate, separate, disconnected computer systems and databases and combining them into a single "virtual supercomputer" or Grid will make crucial data more available more quickly to the employees who need it. It will also raise serious privacy and security concerns that must be addressed. How can a customer be assured that their data will be secure in a virtualized, distributed system? How will national

[4] E.g., http://www-dse.doc.ic.ac.uk/research.distsystmgmt.html and http://www.gridforum.org

privacy and consumer protection laws apply to a global grid that spans national boundaries? How can consumers be assured that their data is not being misused when a single transaction with a single vendor might actually involve three, four, or more different companies, each handling one piece of the transaction? How can individuals ensure that data remains in a trustworthy jurisdiction? What methods of redress exist for a customer's private data when it is under foreign jurisdiction?

And perhaps most important of all, as distributed computing makes it easier to collect and exploit information about individuals, how much more personal data will be collected and how many new ways will be it used—or abused?

Distributed computing means that it will be even more difficult for individuals to understand who has data about them, where it is stored, and how it is being used. Fortunately, work is being done on both the technological and business practices needed to enhance security, privacy, and transparency.

Potential Futures[6]

If current technological trends continue, it is certain that more and more personal data will be collected and used—for good or ill. The Transparent Society portrayed by David Brin is just one potential future. In his provocative book, *The Transparent Society*, David Brin[7] concludes that in just a few years so much data will be routinely collected about each of us—credit card records, data about online purchases, surveillance videos, government records, electronic turnpike toll records, cell phone logs, and much more—that traditional concepts of privacy will have to change in fundamental ways. He argues that instead of trying to restrict how much data others know about us, we should recognize and accept that we can no longer have much control over the kinds of data that is collected about us. Instead, he argues that individuals and governments should look for ways to achieve "reciprocal transparency," where instead of trying to restrict the flow of information, we try to increase the **two-way** flow of information between individuals and their government and between individuals and companies with which they do business. Brin argues that individuals will be willing to provide more information about themselves—provided they also get information back about how that information will be protected, who accesses it, and how it is used. Brin argues that reciprocal transparency could foster trust between individuals and the institutions they rely upon—and that if that could be done, it would alleviate fears about "the end of privacy" due to the explosion in collection of personal data. George Orwell's classic, *1984,* paints a far more ominous future where privacy and transparency are non-existent. Figure 2 shows that these two potential futures are end cases that bracket a whole of range of possibilities. The vertical axis of the figure represents how much information "they" know about an individual,

[6] This section is based in part on material first presented by Harriet Pearson, IBM's Chief Privacy Officer, at IBM's Almaden Institute Symposium on Privacy in April, 2003

[7] David Brin, *The Transparent Society*, Addison-Wesley, 1998, 378 pp.

where "they" can be: (1) the government and large corporations, (2) the government, large corporation, and a single individual, or (3) everyone. The horizontal axis represents transparency, specifically which institutions know what "they" know about me.

Today, we are still in the lower left corner of the diagram, where the total amount of information that any given company or government agency has about a single individual is still limited (certainly compared with what will be the case in the future if current trends continue). Many privacy advocates worry that we are moving inexorably to the Orwellian *1984* future, where governments, with the help of large corporations, are able to collect and analyze huge quantities of information about every aspect of every citizen's purchases, medical history, movements, relationship, and phone conversations—and citizens have no ability to find out what information has been collected. The growing concern about homeland security and terrorism could be a driver for the development of new government systems for collecting information on citizens—and provide arguments for keeping that data secret.

In the far upper right corner is David Brin's Transparent Society, where everyone is able to find out what everyone else is doing. There are video cameras on every street corner, but the video is available not just to the police and intelligence agencies but also to any citizen who wants to see what their neighbors are doing—or where the police patrols are going and whom they are arresting. Brin argues that this could be a type of utopia, where citizens watch out for each other, where familiarity breeds tolerance, and where criminal activity is almost non-existent. This future may be inevitable if technology trends continue; it would certainly be accelerated by a growing desire in the United States and other countries to rebuild community ties and the widespread tendency of Internet users in their teens and twenties to share the most intimate secrets of their lives with the world via Web logs, cell phone photos, and Web pages. This might be an indication of a fundamental change in attitude about privacy.

In between these two extreme scenarios is what might be called the On Demand Privacy future. In this future, individuals are ready and willing to provide and share a great deal of personal information about themselves, their families, their purchases, and other aspects of their lives. But in return, they get detailed information on what data has been collected, how that data will be stored, and how it is used. And they will be able to be confident that the data will not be shared beyond the companies and institutions that they provided it to. And most importantly, they will be provided with a huge range of customized services that save them time and money. Of course, all this assumes that individuals are willing to go to the time and trouble to see that their data is properly protected (or hire trusted third party to do it for them). This also assumes that companies will start to compete with regard to how much privacy and how many customized services they provide their customers.

Besides today's status quo, there are two other possible futures in which current trends reverse and the steady increase in the amount of personal data collected about individuals slows or stops. One reason for this could that malicious hackers and criminals start to find and exploit more and more vulnerabilities in corporate and government networks and computer systems. If the criminals start to win the ongoing arms race between malicious hackers and IT software and hardware companies (and the customers that use their products), faith in e-commerce, e-government, and almost all types of online activities will be undermined. More and more people will decide doing transactions online is just too unsafe. In such a world, personal data would be routinely stolen from corporate and government databases and end up posted on hackers' Web pages. No one's data would be safe. In such a world, companies and governments would have little incentive to develop systems for collecting more personal data and for using that data to provide better and more customized services. New regulations or court decisions that hold companies liable for security breaches could make sure a future more likely.

The last future to consider is one in which anonymizing technologies, digital cash, pseudonyms, and other privacy-enhancing techniques are used by individuals to do e-business transactions online while leaving only minimal traces in cyberspace. This future is named for David Chaum, a cryptographer who has been an outspoken promoter of anonymizing technologies. Unfortunately for Chaum and others who have tried to market strong anonymizing technologies, very few users and companies seem willing to go to the expense and trouble needed to minimize the amount of personal data they leave behind when doing online transactions. For this reason, solutions following this paradigm are limited to selected applications[8] while a complete implementation of this future seems the least likely of the six end cases in Figure 2.

"Building in" Security, Privacy, Accountability, and Transparency

As the technologies and standards for distributed computing are developed and deployed it is essential that they are designed in ways that foster trust—that customers are given more control over their personal data and more information on how that data is used. If that is not done, the full potential of distributed computing will not be realized because individuals and the companies that serve them will not be willing to make the shift to this new paradigm of computing.

In order to examine in detail how trust can be fostered in a world where distributed computing is commonplace, it is useful to divide up the processes by which personal information is collected and used into a number of discrete steps. These steps describe what might be called the "data lifecycle." At each step, measures can be taken to provide more privacy and more transparency. For each of the steps illustrated in Figure 3, we can examine how new technologies, standards, and business practices might foster consumer confidence and trust.

Step 1 – The Privacy Notice

When a consumer visits an e-business Web site or fills out an application for a loyalty card at a grocery store, he or she is usually provided with a privacy policy describing what kind of data the vendor will collect, how it will be protected, how it will be used,

[8] Many enterprises use basic data minimization to minimize risk. Specific cryptographic protocols are, e.g., used for chip authentication using the DAA protocols of the Trusted Computing Group.

and whether it will be shared with other parties. Unfortunately, few Web users take the time to read the details of these privacy policies and would be surprised or confused if they did. The P3P (Platform for Privacy Preferences) standard developed by the World Wide Web Consortium was designed to communicate a privacy notice to a user in a machine-readable form. It is currently mainly used adjust his or her browser so that it will automatically block access to any Web site that tries to collect data that the user is unwilling to share. The browser, not the user, reads the privacy policy. P3P has been incorporated in the most popular Web browsers and is now in use by more than one-third of the one hundred most popular American e-commerce sites.[9]

Step 2 – Authentication

Once a Web user has decided to accept the privacy promise provided by an e-commerce or e-government Web site, the user must authenticate himself or herself. Today, too often, e-commerce Web sites require users to provide ten or even twenty different pieces of personal information (address, phone number, credit card information, birthplace, etc.) in order to verify the identify of the user. This personal information can end up being misused or even stolen. Stronger authentication such as user tokens can prevent identity misuse. In addition, privacy-enhancing authentication technologies exist that can minimize the collection of personal data while still providing strong authentication[10]. Better yet, sites could allow users to use authentication that accepts multiple pseudonyms for a single individual and thus breaks the link between the individual's name and their personal information.

Step 3 – Authorization

Once a user has been authenticated by an e-commerce site, he or she is authorized to use a particular service. As with authentication, this can be done in a way that minimizes the collection and storage of personal information. This is particularly important when a single Web site might call upon several other companies' applications (using Web Services and other technologies.) The Shibboleth authentication and authorization software developed by Internet2 was designed to work across institutions (such as colleges and their suppliers) and is growing increasingly popular, in part because privacy protection is built in.[11]

Step 4 – Data Collection

The best way to prevent disclosure of sensitive person data is not to collect it in the first place. Unfortunately, too many online retailers and e-government services collect more data than is necessary for a given transaction, in part because they hope to use the data to better tailor their service to users' needs or because they hope to identify new opportunities for subsequent sales. Data minimization can protect customers and make them feel more secure about doing business online. On the other hand, many Web users are savvy enough to know that once they have provided their home address, which is

[9] Personal communication, Lorrie Cranor, August, 2005.

[10] Cf. Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. Lecture Notes in Computer Science, 2045:93-118, 2001.

[11] Cf. www.shibboleth.internet2.edu

needed for products to be delivered or a credit card to be billed, the Web merchant can access third party databases that provide a host of information about the user.

A key technology that could provide additional privacy and foster trust is data policy metadata that would be embedded with the personal data and specify the level of privacy protection that the individual or company providing the data expects for the different records and data elements. To quote Kevin Kelly, the former editor of *Wired* magazine, "The answer to the whole privacy question is more knowledge. More knowledge about who's watching you. More knowledge about the information flows between us — particularly the meta-information about who knows what and where it's going."[12] IBM has announced an effort to develop a WS-Privacy for exchanging such metadata as part of the Web-services security roadmap. Once completed, such a standard should specify whether certain data can be shared, how long it should be retained, and what kind of logging must be provided to ensure that there is a complete and accurate record of who has accessed and used the data. Hopefully, as the volume of personal data being collected grows, as the challenges posed by distributed computing become more clear, and as consumers grow increasingly vocal in expressing their concerns about privacy, there will be increased interest in such standards.

Step 5 – Use

Once a company or organization has information on a user's identity, it can then take the next step and process their orders, provide information, news, or music that the user is looking for, link the user to other users, or provide some other service. Proper security and privacy procedures and enforcement mechanisms are needed to protect the user's personal information (e.g. credit card information) at this step.

Step 6 – Storage

This is a critical stage in the data cycle. Before the user's personal data (and the results of the applications) are stored, it is essential that it be minimized and properly encrypted. Encryption does add some complexity and overhead to computer systems, especially distributed systems. Fortunately, recently-announced systems, such as IBM's z9 mainframe, will soon provide the computing power and the hardware-based encryption technology needed to provide top-level encryption for data stored in servers, storage systems, and tape drives connected to it. The recent rash of cases where corporate back-up tapes containing personal data from hundreds of thousands or even millions individuals were lost or stolen has highlighted the need for the widespread use of such end-to-end encryption architectures and devices, a need that is even more critical in distributed systems.

Step 7 – Knowledge Discovery

After personal data has been collected, processed, and stored, there will often be occasions where companies will want to extract knowledge from the data for various purposes—researching their customer base, searching for additional marketing opportunities, trying to identify fraud, or responding to law enforcement requests for information. A pattern-based, behavior-driven form of knowledge discovery known as

[12] Josh Quittner, *Invasion of Privacy, Time,* August 25, 1997

"data mining" has received increasing attention on Capitol Hill in Washington and in other capitals, in part because of the potential for privacy and civil liberties abuses associated with that particular technique. In 2003, for example, there was a serious flap when it was revealed that several US airlines had provided millions of airline passenger records to a US government contractor that was developing data mining software to spot terrorists.

One form of knowledge discovery is known as "subject-based query link analysis". Subject-based queries (which as a technique are far preferable to pattern-based queries because they are consistent with Fourth Amendment "reasonable particularity" tests) can now be performed anonymously. Major advances in analytics inside the anonymized data space now offer the possibility of "knowledge discovery without disclosure" – that is, resolving identities and ascertaining relationships between identities while all identity data remains in non-human readable and non-reversable form. This could, for example, enable law enforcement agencies to seek records of two or three suspected terrorists from a data provider's files containing millions of records of innocent people without any of the records ever being read by either the government or the database provider until a match was found, and even then, only pursuant to a lawful judicial order.

Step 8 – Partnering

As explained in the first section on distributed computing, outsourcing, Web services, the Grid, and other developments mean that more and more data is being shared between companies. Protecting data as it moves between institutions is perhaps the largest privacy challenge posed by the shift to distributed computing. This is one reason why standards for cross-organizational privacy are so badly needed. If the level of privacy protection needed for a given database is embedded with the data, it will be much easier for a company to ensure that its partners abide by those requirements. Such metadata can be generated through P3P-enabled browsers and transferred using machine readable privacy policies[13]. Another criticial aspect is appropriate tethering that ensures that updates to data and rights are appropriately managed along the disclosure chain.

Step 9 – Disclose to a Partner

Once a partner has agreed to abide by the data policies specified for a given set of data, the owner of that data will share it. Properly security procedures (encryption, audits, etc.) can minimize the chances that privacy will be compromised in the process.

Step 10 – Track and Audit

One of the most important ways to foster trust between customers and vendors (and to ensure that a vendor's partners are following proper data protection procedures is to ensure proper audits that can determine who has been accessing and using a particular data set. Embedded metadata (policies, consent, and audit requirements) can make it simpler to check procedures against the privacy protection practices specified for a given data set. Another powerful tool under development is an immutable audit capability that

[13] Cf. M. Backes, B. Pfitzmann, and M. Schunter. *A toolkit for managing enterprise privacy policies*. In Proc. 8th European Symposium on Research in Computer Security (ESORICS), LNCS 2808, pages 162--180. Springer, 2003.

would generate unalterable electronic logs of who has accessed a particular data set.[14] Immutable audits are still a research topic and a great deal of work is needed before practical, cost-effective systems are marketed.

Policy Implications

The explosion of data made possible by the plummeting cost of computer cycles and storage, the spread of high-speed networks, the growth of e-business, and the development of distributed computing (e.g. Web Services, peer-to-peer applications, and the Grid) raise serious privacy and security concerns. Unless properly addressed, these issues could stymie the growth of new applications of e-business and distributed computing. Unfortunately, in several countries concerns about privacy has led some politicians and regulators to propose strengthening existing privacy regulations — which focus on limiting the collection and sharing of information.. This could prove to be counterproductive. As explained above, in the future, what will be needed to foster trust is not LESS sharing of data but MORE sharing of data between consumers and providers — data on the consumers' privacy preferences, data on what data is stored, metadata on how the data is to be stored, and data on who is using the data and how. Provided the proper infrastructure is in place, in the future more data will be shared, which could increase trust, leading to a virtuous circle where even more data is shared and even better and more customized services are available to consumers.

One of the most serious barriers to widespread use of distributed computing could be varying and often conflicting national privacy regulations. It is hard to see how an international grid spanning ten countries could be used to process a company's payroll if national regulations prohibit export of personal data. Such national laws will need to be updated to reflect technological change or else end up being ignored and unenforceable.

Rather than trying to regulate data flows, governments should focus first on leading by example. By finding ways to promote transparency in their own data systems, governments agencies can make customers more comfortable with sharing data with government. Particularly because many agencies have strict requirements for privacy, security, and confidentiality, governments can set a good example for the commercial sector when it comes to privacy, security, and transparency. At the same time, government agencies can learn from the private sector. Why shouldn't individual U.S. citizens have the ability to view their income tax records in the same way they can access the details of the books and records they have ordered at Amazon.com? Why shouldn't government IT systems for collecting and managing citizens' data have audit systems built in that enable citizens to see their data and learn who's used it and for what purposes. Immutable audit records could certainly help here. When developing data mining and knowledge discovery systems to hunt for terrorists and other threats, law enforcement and intelligence agencies should be required to use anonymous data mining, anonymous

[14] See the work of Doug Tygar and his team at the University of California Berkeley, http://trust.eecs.berkeley.edu/socialScience.htm

resolution, and similar technologies to avoid the unintended disclosure of the personal records of average citizens.

Furthermore, governments can be an early adopter of technologies such as P3P and Direct Anonymous Attestation (DAA) and support development and adoption of WS-Policy (and other similar standards for policy metadata). Even more importantly, governments could invest more money in privacy-enhancing technologies, just as they have, in recent years, spent more money on the development of IT security technologies.

If IT companies and their customers work together to develop the standards needed for new privacy-enhancing technologies and then work quickly to implement and deploy those standards in a variety of products and services, it is quite likely that concerns about privacy and security will not delay the shift to distributed computing and the development of on demand business. On the other hand, if government regulations designed for the era of unconnected mainframes are not relaxed or eliminated, and if companies fail to develop and deploy privacy-enhancing technologies, the huge cost savings and new services that distributed computing could make possible will not be realized.

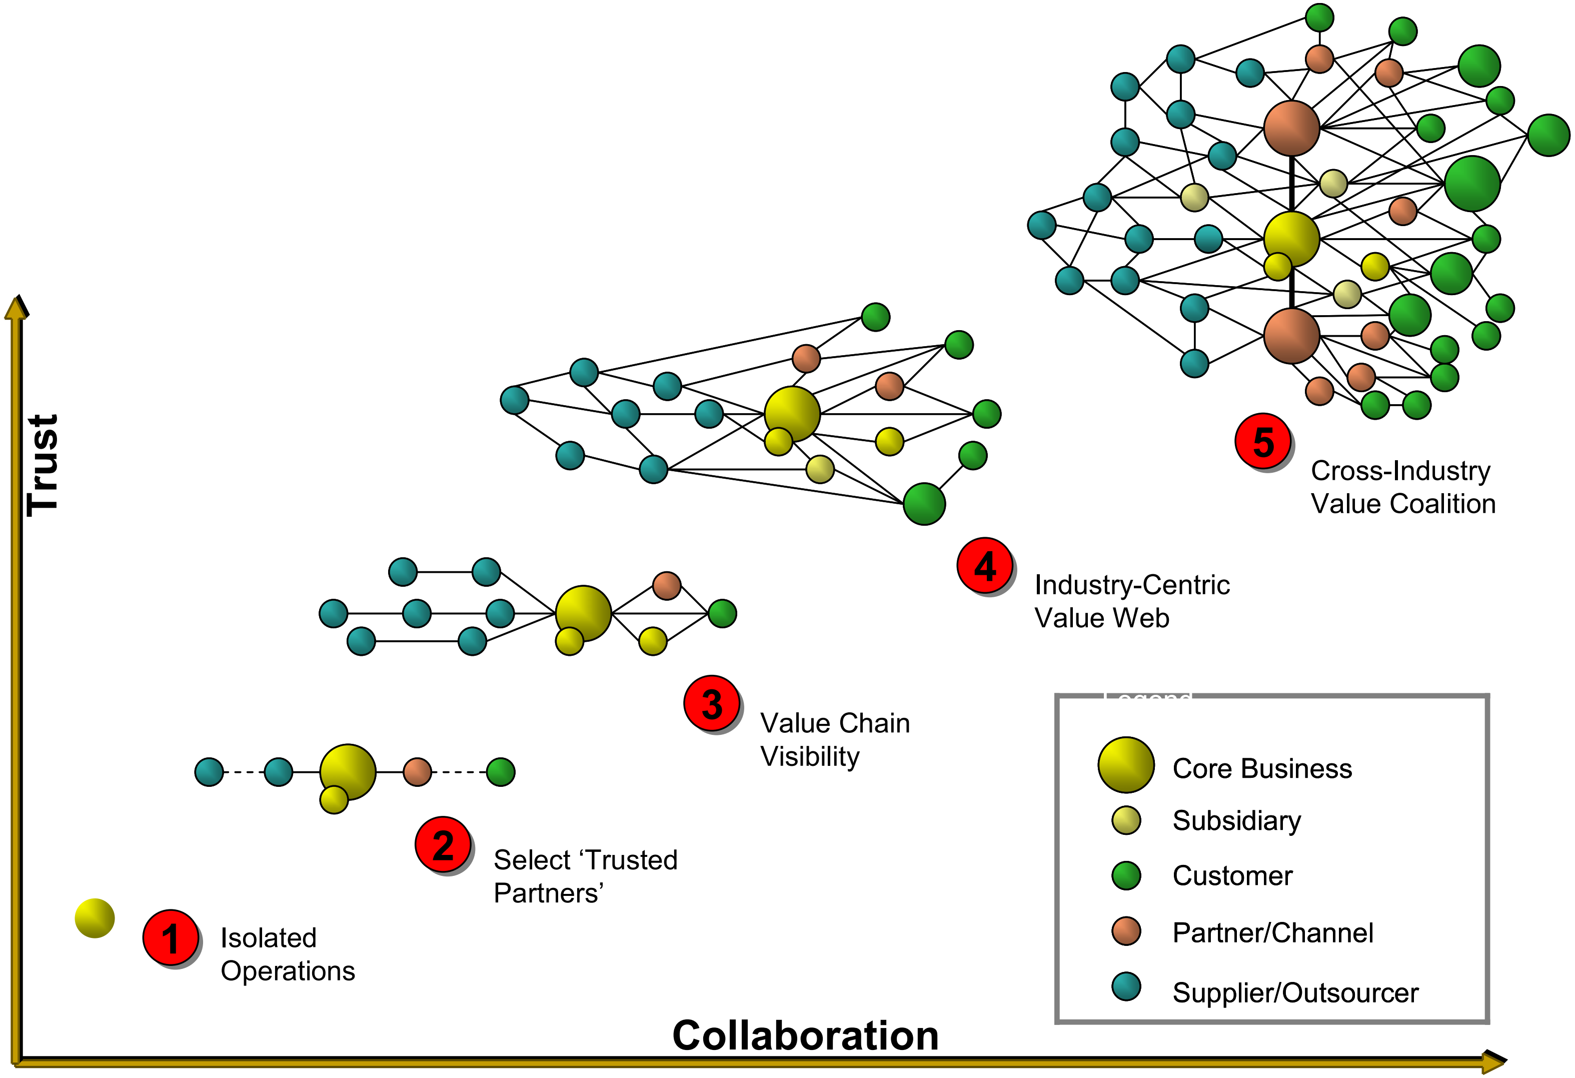

1 Isolated Operations

2 Select 'Trusted Partners'

3 Value Chain Visibility

4 Industry-Centric Value Web

5 Cross-Industry Value Coalition

● Core Business
● Subsidiary
● Customer
● Partner/Channel
● Supplier/Outsourcer

Possible Futures

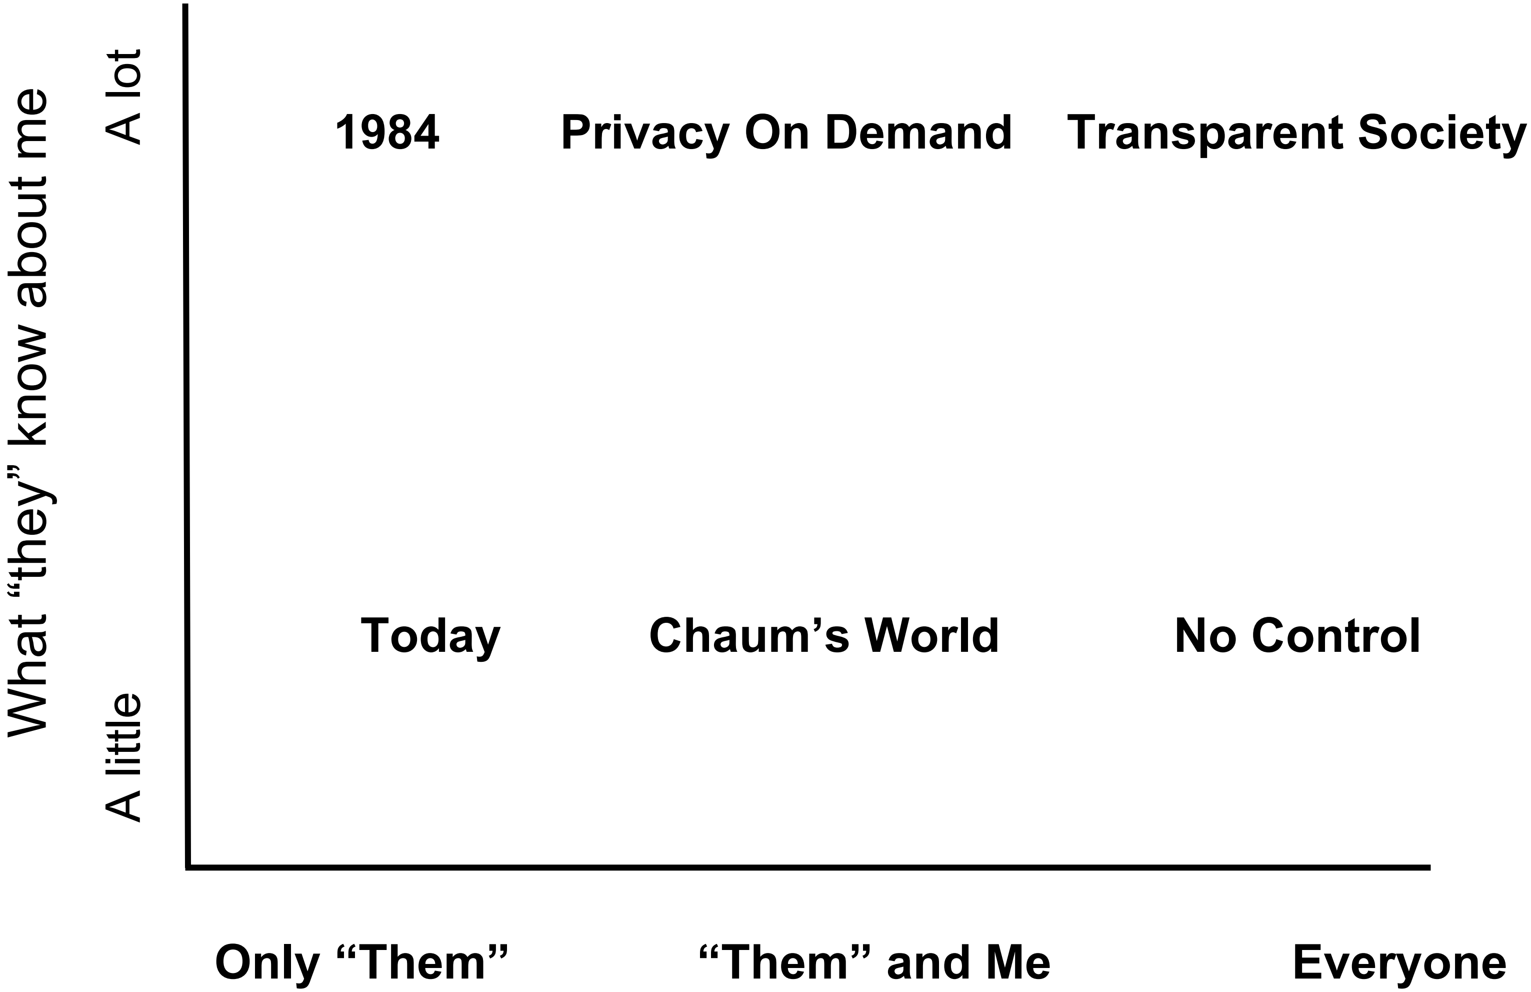

What "they" know about me

A lot

A little

1984

Privacy On Demand

Transparent Society

Today

Chaum's World

No Control

Only "Them"

"Them" and Me

Everyone

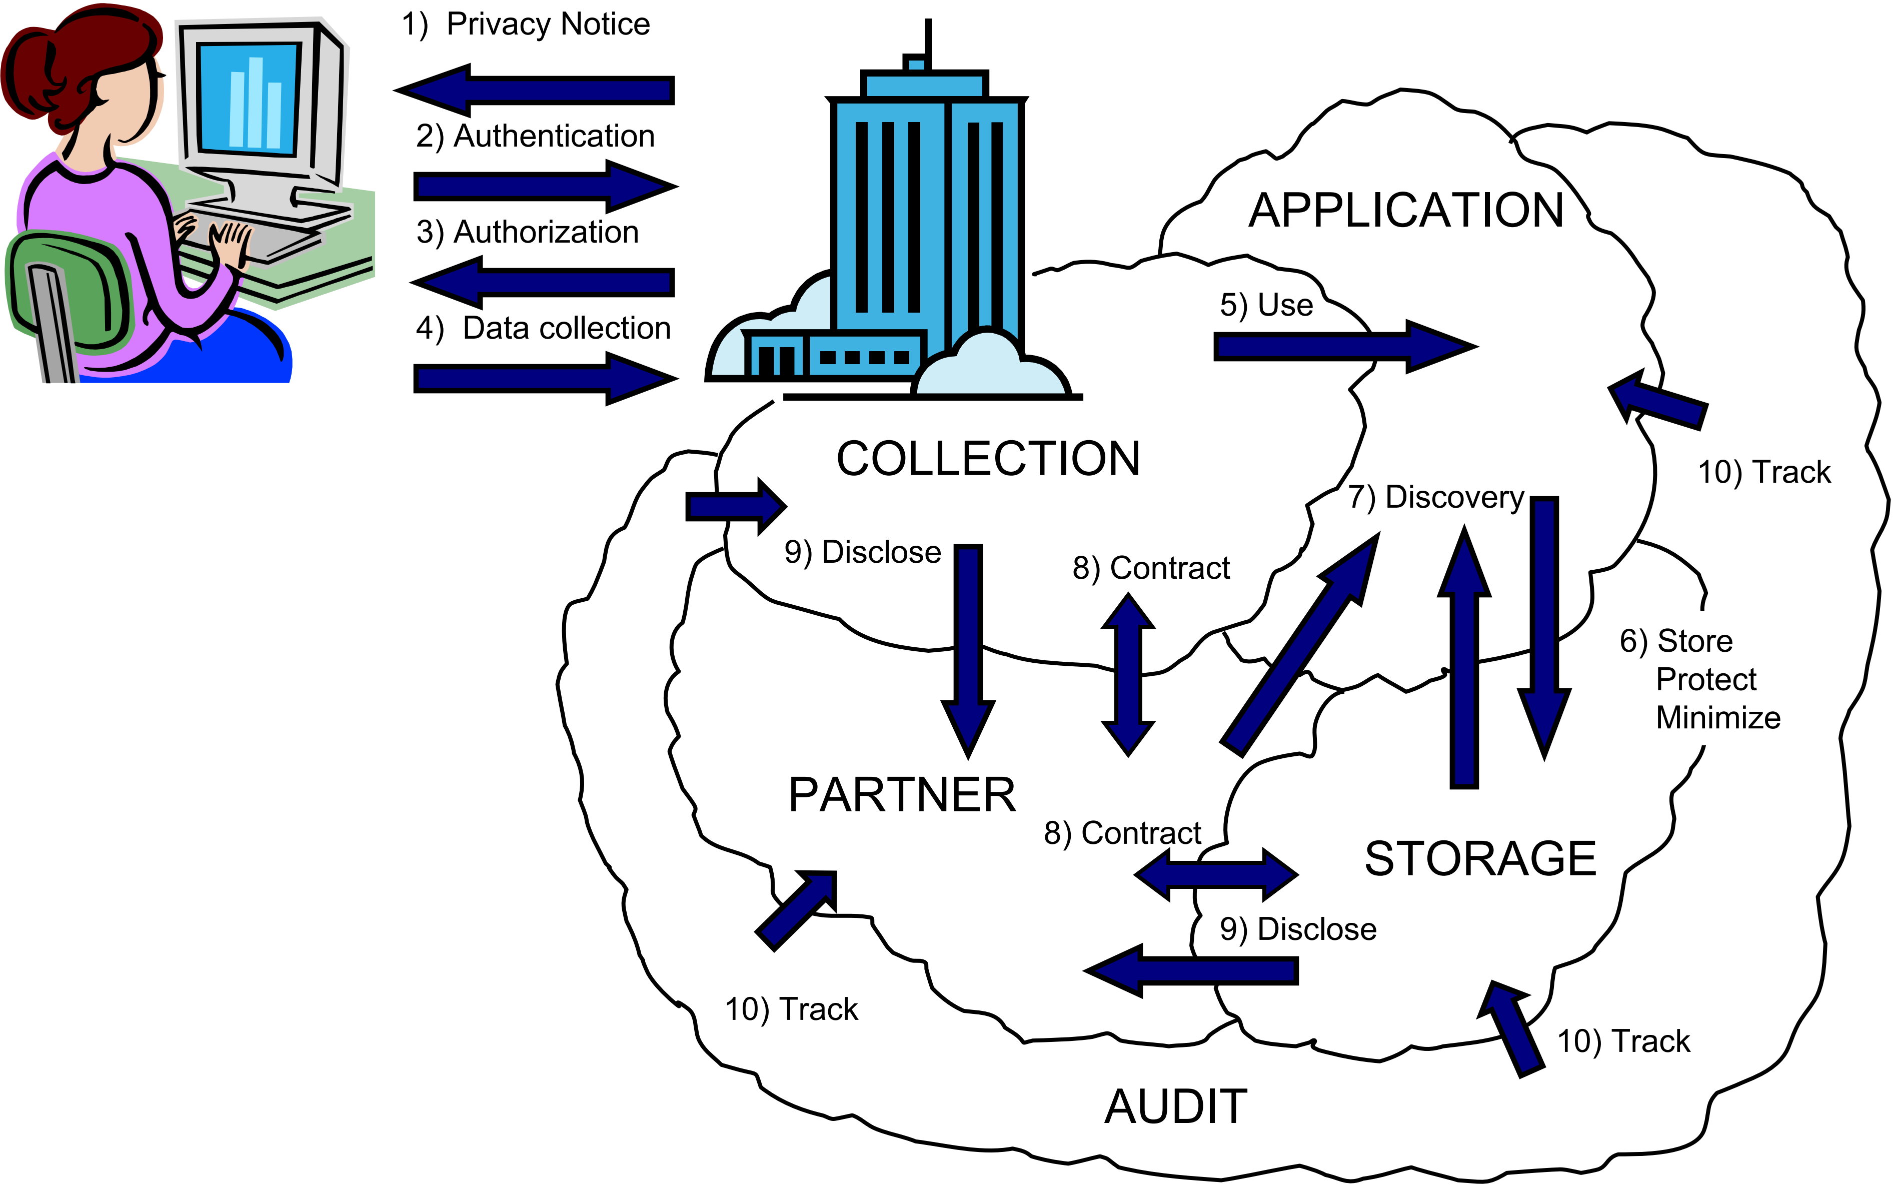

1) Privacy Notice

2) Authentication

3) Authorization

4) Data collection

APPLICATION

COLLECTION

5) Use

10) Track

7) Discovery

9) Disclose

8) Contract

6) Store
Protect
Minimize

PARTNER

8) Contract

STORAGE

9) Disclose

10) Track

10) Track

AUDIT