

CONTRASTING THE PRESENT AND THE FUTURE

If we try to contrast the present and the future, it is like showing the radical contrast between two different images of the universe. Just imagine the universe before and after Copernicus (visit the web site shown in Figure 2.1 for more information). He was the first to suggest what has since become known as the Copernican theory; namely, that the sun is at the center of the solar system (which was considered the universe at that time), not Earth as everyone had thought up to then.¹ The paradigm shift occurred slowly because Copernicus wasn't immediately sure how to present this revolution in a way that everyone could understand, but eventually the sun became the focal point around which everything was centered. If we use this analogy for today's computing world, we realize see that at present everything is hardware-centric, while the future will be user-centric.

For many people, today's computing possibilities seem infinite, but due to the technologies and paradigms used today, there are limitations. Look, for example, at all the dimensions of the computing infrastructure and the layers that are used to implement and deliver a user experience. Traditionally, networked computing was built on the OSI seven-layer stack that had applications at the seventh layer, with everything else being about networks, connectivity, and communications. The user interface (one of the things at the top layer) was just part of the application, and the user was actually not part of the model at all. This required users of the system to know lots about the application and its user interface.

In contrast, we should define a model that starts with the user at the center, layers various user interface modalities around the user,

¹Nicolaus Copernicus (1543), *De revolutionibus orbium coelestium*, Nuremberg.

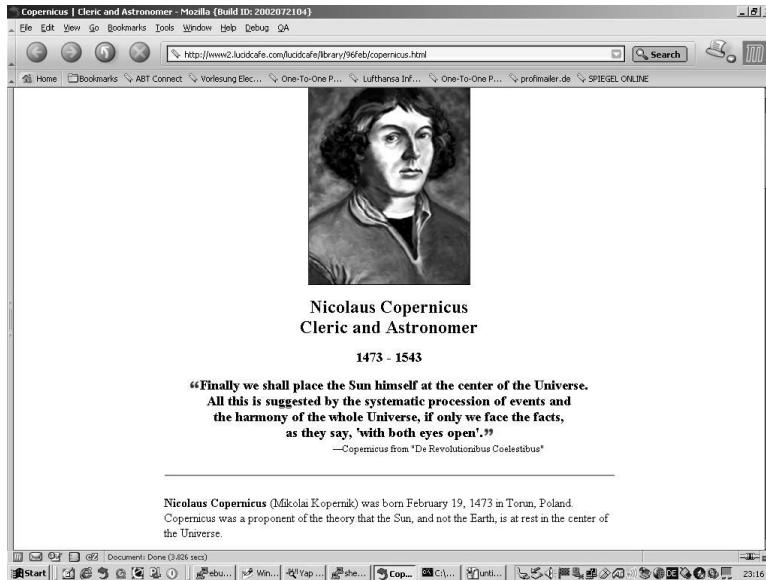


Figure 2.1. Nicolaus Copernicus

and supports these with user models that understand the context and intent of the user. Then we need to deliver value to the user in a manner consistent with concerns for security, privacy, trust, preferences, and other policies. To do that, machines will need to understand how to perform everyday tasks and how to do the master's bidding.

People will be able to tap into this virtual world using any device and have the experience that everything is unified. Neither device differences nor location differences matter anymore. The underlying system remembers the user's activities and other aspects of state; the user can communicate with the system, people, and services whenever convenient and appropriate.

To make this system successful, the user interface experience must be appropriate. Traditional personal computers that are dominating digital communication today will become a minority of access devices in this new system. Computers must understand user intent, most easily achieved using natural modes of communication. In some cases, this will still be the keyboard; in many other cases, it will mean a broad adoption of input devices that utilize speech, text, handwriting, gestures, images, and so forth. Users will adopt input techniques that enable them to convey intent to their appliances most naturally.

This will require efforts to standardize the semantics of processes and how they can be achieved. Thus, starting with XML and Web standards, we will have to fill out the semantics of tasks, people, goals, resources, etc., so that devices from various manufacturers can work with services provided by various service providers to get work done. HP's e-speak concepts, and other research projects in process automation, sketched out much of this (see also Chapter 6).

UDDI and Web services, explained in Chapter 3, are putting down the foundations for service-centric computing, service discovery, and service invocation. Much of what users will notice or care about will be the interface experience; under the covers, everything will be rearranged to delegate work to agents.

What this means is the disposal of current programming concepts and paradigms to make way for new ideas. It's not enough to enhance existing applications; we need to rethink the services behind the applications and rewrite them for the new me-centric environment.

Me-Centric Computing Features

This table provides you with an overview on the key features of a me-centric scenario.

Feature	Description
Me-Centric Appliances	Devices that have knowledge about certain tasks and can do work on behalf of the users
Me-Centric Services	Services that handle information and service requests automatically and autonomously
Me-Centric Environments	Environments that can operate me-centric appliances and provide access to me-centric services

Table 2.1. Me-Centric Computing Features

To give you a better understanding of how the future will look, this chapter offers some example scenarios through which you will see the impact of the new thinking and its radical simplicity. You'll see in each scenario how one segment of life exists in the present, and how it will change in the future. Hopefully you'll glean some input for your own business ideas. For more detailed examples, see *Internet Future Strate-*

gies,² which not only provides a set of scenarios, but also an implementation plan for each of them and a sample ROI calculation.

To put the scenarios in context with me-centric computing, we are providing a table of key features here (see Table 2.1), with a short explanation of what these features mean. At the end of each scenario, you will find a summary of how these features have been implemented in that particular scenario. Later in the book, we will dig into more detail of these features and how they create a me-centric architecture.

Some of the scenarios have been developed by the HP Labs³ and some ideas have been taken from *Internet Future Strategies*.

2.1 Mobile Sales Force

2.1.1 Introduction

Heather works as a salesperson for services in an IT company. Her portfolio includes selling catalog solutions, as well as designing custom services with the help of the client solution design group.

As a salesperson, Heather needs to access information ubiquitously to provide answers to client questions, retrieve order information, up-sell and cross-sell, forecast performance and quotas, and explore opportunities and client contacts. As a designer of custom services, Heather needs to acquire data from the client, understand the client's needs, report these needs to the solution design group, and act as a link between client, sales, and solution design. In general, about 80 percent of a salesperson's time can be divided into four broad activities:

- **Preparation**—These tasks include preparing offers, making travel arrangements, planning for meetings, and reading through the client's history, if available.
- **Travel**—This time is spent on going back and forth to clients and meeting other employees internally, such as the client solution design group.
- **Client Time**—This is the time spent at the client's site, when the salesperson sits down with the client to understand the needs and design a solution for the client.
- **Administrative**—Most of the remaining time is spent estimating forecast performance, regular staff meetings, or completing expense reports.

²Daniel Amor (2001), *Internet Future Strategies*, New York: Prentice Hall.

³<http://www.hpl.hp.com/>

Mobile technology can help reduce the amount of preparation time needed, it can make the client time more productive, and it can speed up the administrative tasks. And in those cases where the salesperson is travelling by means other than driving himself or herself from one client site to another, a mobile solution can provide ways of performing administrative tasks or doing something else productive, such as e-mail or on-line training, during travel time.

Specifically me-centric and mobile wireless solutions will provide her the following advantages on the following topics: sales cycles, trips to the office, preparation time, up-selling and cross-selling, administrative tasks, sales efforts, social networking capabilities, and the delegation of simple tasks. If we look at these points a bit more in detail, you will immediately see the advantages. In the sales cycle, for example, you will find that response time to customer queries can be reduced significantly through the availability of information anytime and anywhere. The salesperson can reduce the number of trips and phone calls back to the office significantly as all information is available on site. The same applies to order entry, which can be done anytime and anywhere. Due to the availability of information, preparation time can be done while travelling and can be delegated to smart agents and intelligent services. Intelligent appliances will be able to provide the right information whenever required.

Up-selling and cross-selling also become easier, as smart agents can find the appropriate products and services that can be used in conjunction with whatever the client is prepared to buy anyway. These personalized proposals can be made while entering the order at the customer site, for example. Many administrative tasks, such as travel expense reports, can be done during travel time and can be automated easily, for example, through smart credit cards that will make sure that all travel-related items will be put into the expense report automatically.

In the following paragraphs, we describe two scenarios in which Heather uses three kinds of connections, today and in the future: managed networks, spontaneous networks, and the Internet. Managed networks are probably the best-known types of network and the most commonly used. These networks include telecommunication networks that allow one to retrieve information from a certain location, by phone or a direct modem connection. Spontaneous networks, on the other hand, allow the spontaneous connections between nearby devices and services. These networks are set up only for specific times and purposes. Wireless LAN and Bluetooth technologies, for example, enable these networks. Although not as reliable as managed networks, they provide a shortcut for exchanging information and services on a very local

basis. The Internet allows one to access information on a world-wide basis. Information and services can be anywhere and can be retrieved easily through the Internet.

2.1.2 The Present

Heather is about to visit a customer. The mission preparation in the office uses the Internet, connected through a managed network. Instead of going back and forth to the office to receive updated information, Heather and Jim, the field engineer, can synchronize through a wireless connection with people and data in the office. In addition, the managed network guides Heather and Jim to the client site.

Heather's primary repository for information is her PDA. She uses it to synchronize data and to communicate wirelessly with scanners, cameras, microphones, and client devices. Before Heather leaves the office, she downloads all information required for this customer to her PDA: mission preparation, updated answers for client, updated client accounts, ongoing actions, sales opportunities, sales forecasts based on account management. During the travel, the PDAs are working in standalone mode until Heather and Jim reach the client site.

Once at the client site, Heather updates Jim while they take the time to establish communications between their appliances. This lets them transfer files and cross-reference devices. During the meeting with the client, they use to-do lists to help increase efficiency in achieving goals and keeping the discussions on track; they use their PDAs to transmit feedback and input for service specifications to Jim's back-end support as well as to Heather's sales office; they can also use their cell phones to make phone calls, if needed. These connections let them share information with people in the room as well as with the back-end office. This information can take various forms, such as pictures, plans, and voice annotations, enabling a more interactive description of the specifications. If needed, they could also have video conferences, scan documents, and share spreadsheets, just to mention a few possibilities. Before parting, Jim and Heather synchronize their data once more and send the information back to the Internet.

All of the connections and service discoveries made during this scenario are made manually, requiring set-up time and some management operations. The Internet does not provide these low-level services automatically today.

Managed networks offer privacy, trust, and high QoS in a known environment. This enables Heather and Jim's devices to exchange and access company confidential information inside the client's campus during this scenario. Sending trip reports out immediately following the

meeting increases the efficiency and lowers the latency of the answers from the back office.

2.1.3 The Future

While Heather prepares for her customer visit, agents work to carry out certain tasks on her behalf. Specifically, they:

- Revisit the last trip report to verify that action items have been completed (specifically, the answers to client questions)
- Check order status
- Explore potential up-sell and cross-sell opportunities with the company
- Make automatic forecasts based on the client's order history and the goals of the salesperson
- Update all information just prior to the meeting
- Call the client to verify the visit date and location
- Summarize the results of these tasks and download them to Heather's PDA

As Heather and Jim reach the client's campus, their devices automatically discover the local services provided near the campus. They confirm the location of high-quality printing facilities so that Jim and Heather can create updated documents as input to the meeting. At the site, their appliances communicate, synchronize, and interact automatically, taking care of the low-level administrative details such as the day-to-day operations, updates on prices, checking the to-do lists, and so forth. This also enables them to establish a private communication during the meeting, letting them dynamically adapt their strategy based on the client's reactions. They can delegate tasks during the meeting, either to their back-end offices (if any) or to services working on their behalf. For example, they could look for background knowledge on the client's project, or gather evidence on how to reduce the risk on the project by analyzing similar previous contracts and experience. Finally, after the meeting, their agents can fulfill such administrative tasks as trip reports, order forms, and so forth. Heather and Jim can delegate further tasks instantly, resulting in faster answers for the client.

As you can see in this scenario, Heather, Jim, and the client have more time to focus on important issues; their meeting is more efficient

because data is exchanged in the background without human intervention. This is a high competitive advantage to the salesperson, as it reduces cost and time while increasing efficiency. The latency of answers drops considerably, and the salesperson is in a good position to close the deal earlier.

To look at the big picture, this setup is delegated to the devices and infrastructure, providing an automatic service discovery and capability to delegate tasks. This provides a virtual office environment to Heather and Jim for all their meetings, right at the client's site. Instead of storing information on a certain device, such as a laptop computer or a personal agent (a successor to the PDAs), the virtual office will store all the information in a central location, making it possible to access the data from anywhere with any device. Data items are downloaded and cached in anticipation of what might otherwise be long latencies or periods of complete unavailability or low QoS.

2.1.4 Summary

Let's look at the key features of a me-centric future. From an appliance point of view, not much will change. Small, handheld devices will still be used, but with one big difference—they will be connected anytime and anywhere, either through local networks, managed networks, or the Internet. Today, many appliances can connect to only one of these at a time. In the future, it is expected that an intelligent appliance will be able to roam through different networks without losing the connection. Many companies are already providing such solutions. Hewlett-Packard,⁴ for example, is offering its Open Roaming solution. Other companies such as Lucent⁵ and Cisco⁶ offer similar solutions, which solve security and mobility problems for mobile professionals whether they roam within the office using wireless LANs or outside using the varying wired and wireless public networks. Users will maintain the security standards and the IP address no matter where they are or what they do. This is essential to a me-centric environment.

So while the devices do not have to change dramatically, the services that run on these devices do. With agents, processes can be automated to help users resolve issues much faster. As you can see in Figure 2.2, a set of services exists around the core service of mobile sales force support. One of the key services is roaming. The core mobile sales force service contains knowledge about the required processes

⁴<http://www.hp.com/>

⁵<http://www.lucent.com/>

⁶<http://www.cisco.com/>

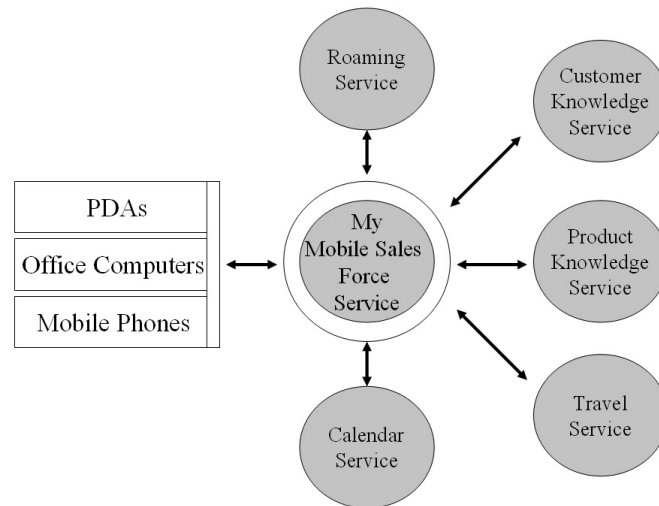


Figure 2.2. Mobile Sales Force Ecosystem

and maintains the connections to all other services. It provides the context on behalf of the salesperson that is required by agents and services to bring back the right information at the right time. By moving the context to this central location, the salesperson is not dependent on a specific device. The adjunct services are not dependent on a particular implementation of context, as the core service will provide translation mechanisms for different technologies.

Besides roaming, knowledge agents are providing an important service to find all relevant information about a customer and the desired products and services he would like to buy. A travel management service will help the sales force to optimize their trips to various clients and make sure that the sales agent's calendar is synchronized with all changes that may occur at the customer site. Therefore, a scheduling agent or service needs to be put in place that does all the timely coordination, based on the feedback from the travel management service and the customer. Additional services may be needed during negotiations, such as status order updates or answers to legal questions. Today these services exist as information pools that can be accessed when required. In the future, this information will be retrieved by agents that will provide contextually relevant information on time.

Through new sales knowledge services, it is possible to coordinate sales efforts towards a target customer group. This makes it possible for salespeople from different regions who work on similar clients to get insight on sales strategies of their colleagues, for example, and show how they can combine offers to create more sales. These collaborative services and agents can also create social networks with other salespeople and clients. These social networks make it possible to reuse existing marketing collateral, proposals, and service brochures more effectively. A salesperson would only see relevant information for that particular kind of deal and it would also allow that simple and standardized tasks can be easily delegated to other employees or to smart agents.

As you can see, both infrastructure and business services must be introduced to make this me-centric scenario a reality. Building up such a scenario would probably be feasible, as the targeted audience of salespeople is enormous and the processes are standardized in many cases. A standardized service with configuration possibilities can help reduce the costs for sales activities drastically.

2.2 Real Estate Agent

2.2.1 Introduction

Henry, a real estate agent is creating a file to describe a house for sale. His goal is to record all information on the house, to present it in its most attractive form, to personalize this presentation for both identified and anonymous potential clients, and finally to diffuse this information in the most appropriate manner.

Collecting information on the house is the first order of business, including a plan of the house, pictures, noise levels in the house, garden, swimming pool, history, maintenance, and conformance to norms. The agent also needs to describe the facilities outside the house, neighborhood safety, risks linked to environment (such as hurricanes, earthquakes, floods), access to facilities (such as highways, train stations and airports), and facilities such as supermarkets and schools.

We develop two scenarios (short-term and long-term) in which our real estate agent uses three kinds of connections for accessing services and information required for the job.

2.2.2 The Present

To-do lists are prepared at the office, taking advantage of the Internet connection and the search engines to gather information on the home

environment. Henry is in constant connection with the office using either a laptop, a PDA, or his cell phone. When he visits the home, all necessary information is recorded and a thumbnail image is sent to a Web site. Potential customers receive an SMS message describing the house. These messages can either be sent to a large base of customers or to a very targeted population, either automatically or handpicked by the real estate agent. More information about potential buyers can be downloaded on the fly while travelling. In addition, the real estate agent can place a key box on the door which sends information to the agency each time somebody visits the home, eliminating the need to go and check these key boxes twice a week as used to be done.

Today, preparation of to-do lists needs to be done at the office and downloaded to the personal agent. During the visit to the house, the personal agent is the primary repository for information and synchronization of data, wirelessly communicating among scanner, camera, and microphone. Communications are established automatically between appliances, and files can be transferred and referred across devices. To-do lists help increase efficiency in achieving goals. On the way to another appointment, the real estate agent needs to pass by a kiosk; the appliances connect to the services offered there and print a generic catalog of the house for sale.

When he reaches the house, the real estate agent connects to the local facilities and can query for information. All appliances can connect to the local network to communicate among themselves and connect to the Web. Therefore, all pictures and information about the home are directly posted on the Web site and are accessible by the clients. Brochures can also be printed out. If there is no network connection at the home, the real estate agent can take advantage of local coverage during his lunch break to synchronize with the office support system and send e-mail to the interested clients.

All tasks are triggered manually, and the end-user needs to verify the result of the action. The real estate agent triggers most operations, the managed architecture provides a reliable network, ensures a quality of service, and allows billing. But, compared to the future, where things do it for us, it is still a lengthy process that requires a lot of time today. Providing a more automated service would enable Henry to increase the quality of service by adding more details to the descriptions or spend more time talking to the prospective clients, for example, for which he does not have enough time today.

2.2.3 The Future

Our real estate agent is in constant connection with the office through his smart device. General preparation work for the house listing is generated automatically based on generic profiles like soccer mom, teenager, or specific profiles of existing clients existing in the office database or on recent statistics. This enables him to show the house to the right clients. The mission preparation is done by a support team. Once the preparation is finished, a precise mission statement, a to-do list, and a software agent are sent to Henry while he drives to the home and downloaded to his smart device.

A guidance system is used to organize visits according to road traffic, other appointments, and opportunities (such as good commission quota or probability of selling the house soon). Five minutes prior to the visit, Henry's personal agent calls the homeowners or clients to reconfirm. While the real estate agent travels to the house, he delegates his agent to conduct research on the neighborhood, safety records, schools, and so forth from police, county, and geographic records.

During the visit to the house, the personal agent actively helps Henry carry out the mission; in particular, appliances communicate and interact automatically with those in the home to collect data from them. The camera asks for pictures of the kitchen and gives a description of the important items to capture for all the clients. These pictures are automatically labeled and stored and are ready to be used. The catalogs are automatically composed and are print ready. Video recordings of parties and events casting a good image of the house can be requested uploaded to the house's web site. In addition, the size of each room can be precisely measured and transmitted to the personal agent, which updates the plan of the house and annotates each room's picture. All of this information can be immediately broadcast to potential clients, who can call Henry instantly for additional details.

Meanwhile, Henry follows his to-do list and focuses on taking pictures and collecting information on the house not available elsewhere, or gathering information requested specifically by clients.

The ultimate combination of appliances would enable users to go further and create a virtual home that future clients can navigate and discover at their pace. Interested parties can download the tour or software agents may deliver short lists of houses according to their specific choices. These virtual homes can be tailored to potential client tastes, providing a cozy house to a retiree or a wild party scene for students. Ultimately, a client can create avatars for his family, including the dog and goldfish, to see how these avatars would adapt in the new home, for instance how and when they use the swimming pool (the family, not

the pets). We could also imagine the house twenty years down the road, estimating the maintenance costs and the evolution of the family.

All links between tasks are automatically made for the end-user. All the appliances collaborate in such a way that the user does not need to know where the information is stored, which software is used, or which connections are employed. The end-user is in control of the end results only; the process is completely invisible.

2.2.4 Summary

Let us look again at the key features of this scenario and how they have been implemented. The me-centric infrastructure is built on a high-availability, broadband network connection that links Henry with the office. The reason for a high-speed network is the number of large pictures that need to be transferred back and forth. Unless a new algorithm becomes available that can pack images much better than JPEG without losing quality, the size of the images will be probably around 1 MB. So either a WLAN network or Universal Mobile Telecommunications Service (UMTS) network will be required to make this system work. Roaming does not play an important role here. Speed is more important, and typically you can have either roaming or broadband access. The WLAN network would have a base station in every house that would transmit the information via cable or DSL, for example, back to the office. A UMTS network would require a base station near the house that would transmit the information via satellite to the Internet backbone and from there to the office.

The mobile device that will be used is a combination: mobile phone, PDA, and digital camera. We are seeing the first generation of these devices being introduced into the market today. Nokia⁷, for example, introduced the Nokia 7650, with an integrated digital camera, which takes pictures that then can be sent via MMS (Multimedia Messaging Service) on a high-speed mobile connection. It also provides a photo album for storing pictures and an advanced user interface that runs downloadable personal applications via Java technology. To make it easier to use, it includes joystick navigation and a color display. Although the quality of the pictures is still very poor (the Nokia provides only 640x480 pixels), the bandwidth is not very high, and the PDA functionality is not very good yet, we can expect to see devices that will match our requirements soon.

One of the biggest challenges in this scenario is the services required to run the business. If you look at the ecosystem in Figure 2.3,

⁷<http://www.nokia.com/>

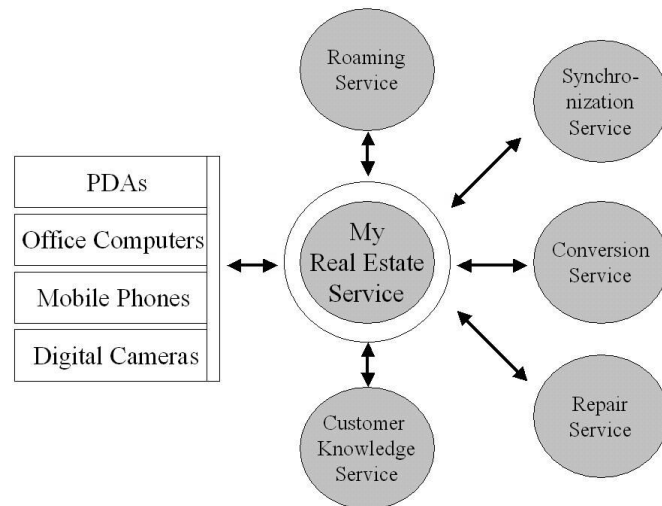


Figure 2.3. Real Estate Ecosystem

you can see that many of them do not exist in an optimal way. A synchronization tool is required to make sure that all data between the office and the device is up to date. While the images are quite large, the synchronization tool needs to check the bandwidth first and update the data only when enough bandwidth is available.

This means that if only little bandwidth is available, it should only synchronize small files, such as text files, and copy the descriptions of the images. If there is an urgent need for a certain image, it would be possible at least to search for it and download it immediately. The device should provide extensive descriptions to the pictures. The real estate agent should speak while taking the pictures and the device should save the spoken words and convert them into text. The device should also note where the pictures have been taken and use the construction plans of the house to create a virtual environment for prospective buyers.

An intelligent agent could be used to collect information about all repair work in the house. While Henry walks through the house and makes comments about repairs that need to be done in order to increase the sale price, the intelligent agent could start searching for the best offers from house repair companies, providing Henry with some time

and cost estimates immediately. Another intelligent agent could collect all the information about the house and create an advertisement for the local newspaper and the local radio/TV stations automatically. Another intelligent agent could start to look through the customer database to find people who may be interested in the house and make appointments with them, if interested. Henry would receive a list of appointments before he even leaves the house.

Obviously, to have all of the services described in this scenario would be expensive and may not be profitable to implement. While most of the scenarios in this chapter are realistic and affordable, we wanted to show you some other ideas that may currently seem otherwise. However, if you look back in history, there have been many technologies and ideas that were considered unrealistic and expensive but have still become commodities today. One of the questions that needs to be answered is what economic model will cause a proliferation of so many diverse services that somehow come to be effectively aggregated into easy-to-use, coherent me-centric customer solutions.

2.3 Family Life

2.3.1 Introduction

Janice, married with three children, works for First Commonwealth Bank (FCB) handling commercial accounts. Her husband, Mike, is a painting contractor. Their oldest, Matt, has started junior college and still lives at home. Claudia is in high school. The youngest child, Sharon, is in middle school.

Janice's desk job comes with a wired phone that is part of the FCB's virtual private network (VPN). She keeps her cell phone on "vibrate" and in her purse during work; it's bad enough being interrupted by one phone when meeting with a client, much less two (one wired and one wireless).

It's a cold, clear winter day, and Janice realizes that she will have to stay an hour later than expected. She would rather not stop at the Bi-Lo or cook dinner. She decides to call Mike to see if he can handle dinner.

2.3.2 The Present

Janice calls Mike on his cell phone to ask if he can stop by the store and pick up barbecued chickens on his way home from work. But, as is often the case, Mike is on the phone. She gets his voicemail and leaves a message:

“Sweetie, could you pick up a couple of barbecued chickens on your way home? I’m going to be late. Let me know if you can’t, and either I’ll call Matt or pick them up myself. Thanks. Bye.”

Half an hour later, she has not gotten a call back from Mike. She checks her cell phone: one unanswered call. She retrieves her voicemail and learns that Mike is actually working on the other side of town and won’t be able to stop by the store. She calls Matt. Again, she gets voicemail:

“Matt, this is mom. On your way home from school, could you buy two barbecued chickens at the Bi-Lo? Thanks. If you can’t, please give me a call at work right away. See you tonight. Bye.”

To make sure that Matt gets the message, she sends him a quick text message asking him to check his voicemail and call her. Janice is finally ready to leave work. Still no call from Matt. “I wonder if he got the chickens?” She calls home: The line is busy. Claudia is no doubt chatting. She calls Matt again: voicemail. As soon as she walks out the door, her office phone rings: Matt leaves a message on his mother’s voicemail:

“Mom, I got the chickens. I put them in the oven to keep warm. Bye.”

The disadvantages of using the present system are pretty clear. It requires too much manual intervention and depends on humans being simultaneously available for synchronous communication requiring their nearly undivided attention.

2.3.3 The Future

It is 6:30 am. Mike and Janice finish their coffee and perform a ritual that’s been going on for almost two years now: they synchronize their organizers. Mike is using the latest Nokia communicator; Janice has an HP iPaq. Lying on the kitchen table, the two devices exchange packets of infrared light.

That afternoon, Janice realizes that she will have to stay an hour later than expected. She would rather not stop at the Bi-Lo or cook dinner. She takes out her iPaq and sees that Mike is working on the other side of town. However, Matt finishes his last class at 5:30. She calls Matt, but gets his voicemail:

“Matt, this is mom. On your way home from school, could you buy two barbecued chickens at the Bi-Lo? Thanks. If you can’t, please give me a call at work right away. See you tonight. Bye.”

Janice picks up her iPaq and selects the shopping page. Touching the Bi-Lo link and then the deli link, she then drags two chicken icons into her “shopping basket.” She then opens the basket and selects a 6

o'clock pick-up. As nobody is at home, automatic delivery would not work without a special storage.

Janice drags the shopping icon into the "message board" space and touches "annotate." After a couple rings, a recorded voice welcomes her to the ATW message board and asks her to speak her full name. "Janice Weaver" is all that it takes for her to authenticate herself and gain access to the message board. Janice says "family" and then when prompted, explains the situation: "Hello everyone. I have to work late tonight so I'd really appreciate it if someone could pick up two barbecued chickens at the Bi-Lo. I've reserved them for 6 o'clock. If you can pick them up, please give me a call. If I don't hear from you, I'll get them myself, but dinner will be even later. Your choice! Thanks. Bye."

Janice goes back to work while the message board attempts to contact and deliver the message to the family. Mike's cell phone is called, but the message board recognizes that the call has been re-directed to voicemail (caller ID has certainly improved over the last few years). After the greeting message goes silent, the message board plays Janice's message and hangs up.

In a similar manner, the home number is called. The fourth call attempt doesn't receive a busy signal, and Sharon answers. She hears her mom's recorded message. Putting down the phone, she calls upstairs to her sister Claudia: "Mom wants you to pick up some chickens. You should call her."

Matt's cell phone doesn't get called: He receives an SMS containing his mom's message, while he's in the middle of an exam. His personal agent refuses to pass on the message to Matt, as the message has only low priority. It will be delivered to him after the exam.

Claudia is the first to phone back, but she calls her mother's cell phone instead of the office number. Nobody picks up the phone. Therefore she decides to ask the message service agent for more information on the details. Fortunately, Janice linked her shopping agent to the message service agent to form a workflow, making it easy for Claudia to see the details of the order. Claudia confirms taking on the task to the message service which would know how to close the loop with Janice and also close the task out, keeping others from taking it on or being bothered with the message.

Janice is finally ready to leave work. Still no call from Matt, but an automatic message telling her that he will get back after the exam. She calls home: busy. Claudia is no doubt chatting. As soon as she walks out the door, her office phone rings: Claudia's agent reminds her to leave a message on her mother's voicemail in addition to the automatic notification of the message service agent:

“Mom, I got the chickens. I put them in the oven to keep warm. See you later. Bye.”

2.3.4 Summary

If we look at the key features of the me-centric implementation in this scenario, we can see that the community is the most important feature. There is no need for high-tech devices or high-speed networks. Today’s PDAs and mobile phones can be enhanced to make this scenario reality. The difficulty lies in coordinating the communication within the family, which is sometimes already difficult enough, even when everyone is in the same room.

The future scenario is an example of the use of a standardized template for tossing a task out to a small community of potential volunteers and a way for the task creator to be notified of successful delegation and a way for other volunteers to be freed from further communication about the task and prevented from unnecessarily taking on an already spoken for task redundantly.

In Europe, many children age six and up already have their own mobile phones. Parents give them phones to call home (for the younger ones, it is the only number they are able to dial; all other numbers are disabled). One advantage of having a mobile phone is that it is possible to locate the person easily. Although GSM⁸ provides only rough location capabilities, newer phones have GPS built in, enabling exact positioning. This means that parents can check where their children are without having to call them. An agent could be introduced to make sure children stick to a routine path when walking home from school, for example. The agent could alert parents if their child starts moving in the wrong direction, or begins to move too fast as though they had entered a vehicle.

In our scenario, an intelligent agent could coordinate the locations and the supermarkets and check in advance if a certain market has the desired chicken in stock. Another agent could coordinate and synchronize the agendas of all family members. GPS could also be used to locate the person closest to the market for convenience.

In Figure 2.4, you can see the basic infrastructure requires a similar set up like a company, where each family member has certain privileges and rights in the system. A communication agent should track where each family member is and how that family member can be reached best. If Mike, for example, tells his mobile that he is in a meeting, this information could be relayed back to somebody who may be try-

⁸Global system for mobile communication

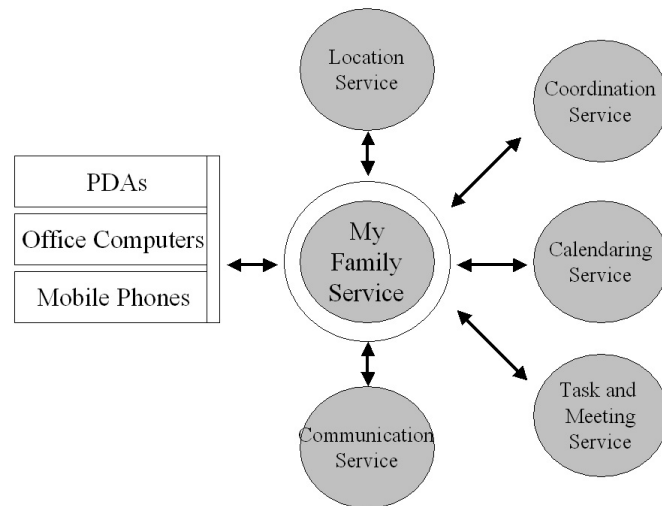


Figure 2.4. Family Life Ecosystem

ing to call him. Additionally, important information could be relayed to Mike in text form while he's in his meeting, so Janice could speak to the voicemail system and the agent could evaluate its importance. Today's voicemail systems allow the caller to specify the importance of the message, which is no good, as callers always think that their calls are of utmost importance. The agent acts on behalf of its user and not on behalf of the caller, which makes the agent trustworthy.

The ecosystem should also support task distribution among family members. In this case, Janice could send out the task to be completed to everyone in the family, and a task agent would coordinate the task with everyone's agenda and capabilities to find who would be ideally suited for the task. The task agent would also coordinate the family with the outside world, e.g., with friends and other relatives, and would connect to other agents that may be of interest to the family, such as shopping agents or intelligent vacation agents.

2.4 Work-Life Balance

2.4.1 Introduction

Chris is a working parent and is involved in two entertainment activities right now, one for work and one at home, which he needs to coordinate in order to maintain a good work-life balance.

At work, Chris is in sales and has a need to keep in touch with his customers and develop new ones. Chris sells a line of teen-oriented clothing. A number of his customers are into a future-shock, cyberpunk role-playing game, called “bunk,” that connects them through phone, pager, Web, and e-mail. Chris has chosen to play this extended online game so that he can meet potential customers, interact with his existing customers, keep tabs on the latest lingo and fads, and just to have fun.

At home, Chris likes to relax in his media room and is very into music, music videos, and photographs. He has a collection of old college buddies who are into this too, but they are spread all over the world. He and his friends like to get together to share music, talk about the bands, concert dates, new videos, and so forth. He also likes to have access to his music when he’s in the car on the way to his next appointment or just when he’s ferrying his daughter to her horseback riding lessons. Ozzy Osbourne is a great accompaniment to dressage.

2.4.2 The Present

In the near future, we have serious bandwidth issues for our wide area wireless providers and very spotty coverage for 802.11. The appliances Chris uses speak both. Chris likes to travel light, so he has a very light headset he uses for both music and phone work, a palm-sized computer that doubles as a phone dialer/communication hub, and a wristwatch that shows simple paging info and doubles as a control panel for playing music, etc.

On his way to the mall to talk to his next client, Chris is notified through his watch that a bunk player is nearby who has some info he needs. He uses his PDA to IM and makes a stop at the Jamba Juice to chat and trade bunk secrets, PDA to PDA. On his way to his customer, he passes a print kiosk where he prints a small booklet about the next bunk challenge and some info he needs.

His new-music-notifier service also sends him a page, based on his music interests, showing that there is a new Ozzy track. He vows to track that down later. At his customer’s place of business, he finds out what the needs are, beams over a catalog customized to his customer,

picks up some information about what's selling (the local store inventory has been told to allow him to get aggregate info on his products only), checks on the status of some open orders, and swaps bunk stories with the manager. Because his watch has swapped interests with the manager, he finds they are both Ozzy fans, and the manager tells him about a new single that is out. They share bunk IDs so they can help each other later in the week. A new custom catalog will be printed and shipped out that same day, and an online catalog will be uploaded that the customer can peruse.

On the way out of the mall, Chris stops at the music store to get a demo track of the Ozzy song, adding it to his playlist. He will listen to it now and then share it with his buds tonight when he gets to his media room. On the road again, Chris fires up his playlist on the car stereo that was synchronized with his home system overnight when the car was in the garage. The Ozzy single will be waiting for him at home. He can listen to it and share it with his friends later tonight.

2.4.3 The Future

In the farther future, we have better bandwidth for our wide area wireless providers and less spotty coverage for 802.11 WLAN technology. The appliances Chris uses can connect to both types of networks. Chris likes to travel light, so he has a very light headset he uses for both music and phone work, a palm-sized computer that doubles as a phone dialer/communication hub, and a wristwatch that shows simple paging info and doubles as a control panel for playing music.

The personal agent informs Chris about new music from Ozzy, which is automatically downloaded to his music repository and which will be charged only after he has listened to it twice. The first listening is free and lets Chris decide if he likes the song enough to purchase it. The song is encrypted and cannot be copied so the record label does not suffer from pirate copies of the song. Once he decides that he likes it, he receives not only the audio file, but automatically also the video file with it. Depending on the device Chris is using, the quality of the video is recalculated on the spot or only the audio track is heard.

While driving to the customer site, Chris can already start talking to the customer about his needs, because his car has a built in camera and displays the incoming video stream on the windshield, so that Chris is not distracted from driving. At the customer site, the print-on-demand service has delivered the new catalogues, and Chris can use the time to have some discussions with the manager about the future.

2.4.4 Summary

If we analyze the me-centric implementation in this scenario, we see that mobility plays an important role, as Chris moves around and requires lots of bandwidth. The major issue is the design of the device that Chris will be using. The light headset, the wristwatch, and the PDA need to be connected wirelessly and need to be able to operate on their own, if the other components are not reachable. While the wireless connection can be achieved easily with today's technology, such as Bluetooth for local connection and WLAN for external connection, it can be difficult to create self-sufficient devices.

Each of the devices should be able to take over the functionality of the other. The headphone should be able to provide spoken page information, while the PDA should be able to play music, if required. Therefore, the devices themselves should not hold any information or services, but rely on the connection to a service provider that can provide the required information and services to each of the devices. The only job of the device is to convert the information to its output interface. An agent needs to configure the devices that Chris uses. The devices should work together automatically and respond to the service provider the configuration in which they work. In case a device is switched off, the service provider may need to reconfigure the service to make it useful for Chris. This means that when he uses his earphones, the music will be streamed to that device; if he is at home, the music will be played on his personal stereo. The device should have a fat cache to operate well with downloaded data and programs, and continuously when offline.

In Figure 2.5, you can see the networking infrastructure, which provides a seamless network to Chris and his music. Depending on the type of network, the quality of the music is automatically adjusted. On low-bandwidth networks, the quality is automatically dropped, and as soon as Chris connects to a higher bandwidth network, the quality is enhanced again. This requires seamless roaming just like in the first scenario.

Some bandwidth needs to be reserved to support the agents that help Chris with his work. Chris's car can be used for video-conferencing and can provide all the relevant information that Chris needs for his work. An agent collects the past orders of a certain dealer and compares the information with other dealers and their recent orders. The agent prepares a proposal for that dealer based on this information. Another agent would check the stock level at that dealer and restock the bestsellers automatically.

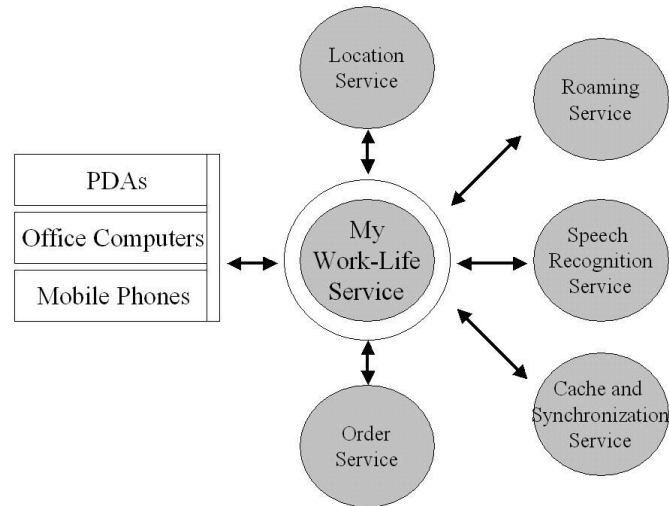


Figure 2.5. Work-Life Balance Ecosystem

2.5 Construction Site

2.5.1 Introduction

On a construction site, a new set of technologies is going to be introduced. Every worker will receive an online connection that allows him to communicate with the other workers and that monitors progress. This is done, for example, by monitoring the number of bricks that are already used for a wall. The computer in the back can verify that everything is going according to the project plan. It can also make projects, monitor when the next step is going to be complete, and order electricians just-in-time, as needed. A satellite will check the construction from space and make sure that it is built to its specifications. The architect can always have the full picture. He is able to see a 3D model of the construction on the screen, both the desired solution and then the actual solution, and make modifications on the fly. As every worker is connected, it is possible to make changes while building, because everyone will be able to receive the update immediately.

2.5.2 The Present

Although very ambitious, it would be possible to start with such a project today. First, a framework must be created that enables all parties involved to work together. A collaboration tool needs to be introduced that can be used by anyone on the construction site. As the workers won't be able to use a laptop while working, other devices are required to allow them to work in a group and exchange information electronically. Although very expensive, it is possible to build intelligent helmets for workers and display the content and the services on a small monitor that can be built into the worker's glasses. But it may not really be necessary for all workers. Information and services can be provided by mobile phone, for example, and it may not be required that the workers be online all the time. Therefore, it is necessary to think about the processes involved before setting up the infrastructure.

Although it may seem that a construction site is organized in a decentralized manner, there is a management of the whole site that makes sure that everything is done in the right order. Therefore, it would be important to support the management functionality of the site first. This would mean, for example, that materials are delivered at the right time, in the right amount. If the architect's drawings are the basis for all further planning, the software should be extended to support a project plan and procure the materials required for building.

The architect should just enter the drawings into a CAD program and provide a project plan based on the drawings. From this input, the system should be able to generate a list of items that are required to build the house. The materials list is sent to one or several market-places where companies can provide prices for the required building blocks of the house. In addition to the quantity, the delivery date can also be sent to the supplier. This reduces the space required on the construction site for storing materials. So far, it is not so difficult. It allows the architect to tell the sponsor of the project how much it would cost in total and if it is actually possible to build the house in time, as the architect now has a commitment from the suppliers to deliver the materials on time.

The service for the planning phase can be extended into the development phase on different levels. First, communication can be enhanced through mobile devices that can transmit voice and data. Simple instructions could be sent via voice, and more complex issues could be sent as data, making it possible to visualize a part of the house. While it may be impractical for many to have anything larger than a PDA, it actually may be sufficient to display on a small screen the parts of the house that the worker is actually working on.

Although security is not a major concern, it still needs to be taken into account. If a hacker could invade the construction network, he could create confusion and damage. Therefore, the wireless network needs to be protected. End-to-end security is a must again. The server, the transmission and the clients need to be part of the overall security model. Security is further enhanced by using devices that contain smart cards or require fingerprint verification, for example. This ensures that nobody can steal passwords without someone noticing it. If a hardware device gets lost, it can be excluded. Through this procedure, it is easier to secure the network, as only devices that are used by authorized persons are allowed on the network, and any device that is not owned by the construction company is automatically denied access.

The communication service could be enhanced by providing feedback from the workers about their status. Once they have finished their task, they could send a status update back to the server, which allows the construction management to track the progress. To make it easy for the workers, there should be a button labeled “I have finished this task” on the mobile device. By clicking on this button, and confirming the data that has been collected during the task, the worker sends the information back to the server. As each mobile device is attached to a certain worker and a certain worker has a certain task, this concept could work quite well and would not require the workers to have a deep understanding of the technology.

The planning service consists of components that exist today. CAD applications, project management tools, and marketplaces are common, but the integration is missing. This means that the main focus for the planning part of the service is integration of different software solutions that are not necessarily Internet-enabled. Most construction companies will have CAD applications and project management tools in place. In order to create this new service, it needs to be determined whether an extension of the current products makes sense or a replacement with newer software is required. Both will require some investment. Both applications require an Internet-enabled interface that allows them to share information with the marketplace and other services, such as communication services on-site. Introducing a messaging backbone and a central database for this service is not cheap, but still within the range that a company would like to pay for.

More expensive are the mobile devices and the infrastructure required to make them work in a collaborative environment. As we are talking about a system with special buttons to make use as easy as possible, special hardware needs to be introduced. Although it is an expanding market, companies will charge a premium for these devices

in the beginning. Today it may not be feasible to introduce such a new communication infrastructure.

Although possible to implement today, only the planning phase would make business sense; introducing the collaboration part of the solution will be probably be too expensive on the hardware side, even though the software and the infrastructure can be set up at a low cost.

2.5.3 The Future

In the future the solution overview that was laid out in the preceding section can become reality more easily. The major reason is that mobile devices will become more common and prices will drop significantly. Not only will the prices for the hardware drop, the software components in the solution will become service offerings, reducing the need to buy a certain piece of software. It can be easily rented instead. Due to the new e-services infrastructure that will be in place, exchanging information between the services will become much easier and cheaper. First, there will be a connection to a security service that authenticates and authorizes the devices and workers who want to use the service.

Other than in a national security marketplace, it is not necessary that a governmental organization run the authentication service, but governmental organizations may want to enforce proper staffing on the construction site. Governments may actually require their service to be used on the site. Illegal workers will have a hard time and companies that try to save money through clandestine employment will have an even harder time with this type of authentication in place.

All construction-related data is stored in the construction management system, with contributions from the architect, the marketplace, and the workers on the construction site. Information can be pushed to the parties involved, but it is not stored there. This reduces the need for memory on the required devices. In the near future, it will also be easier to integrate new devices into such a network. Through wireless LAN technologies, any device can be made part of a network easily. Thus, creating new devices for a particular cause, such as providing feedback about the status of a workstream, is easy. Imagine a mason building a wall with a scoop. Instead of giving the mason a new device, it would be possible to enhance the scoop with wireless LAN access. The mason presses a button when he starts to build the wall and another button once he is finished. These simple functions can also be included in other tools the workers require anyway. Thus, the information technology is hidden away from them and does not require additional learning to improve the performance of their job.

To make the CAD application as effective as possible, it needs to be based on a client-server architecture. Just like all other applications, it should be provided as a service to the client in order to make it easy to exchange information with other services that are required for this meta-service.

In addition to the CAD application, a special visualization tool is required. It not only visualizes the current state, but also the complete plan and the delta between the two. It shows this not only from an architectural point of view, but also from a project management point of view. Ideally, the tool would also allow the user to create projections and reports for management. Very important also is the feature of visualizing information not only on high-resolution screens, but also on small displays of personal agents and cellular phones, which require different levels of detail.

The same applies to communication. The system should be able to transform messages from one format to another on the fly. Voice, text, e-mail, and fax, just to mention a few, need to be converted and stored in a single format, to make sure that everyone on the construction site is able to access the information in the way that best suits a particular moment. Unified messaging will be a key service for collaboration on the construction site. It will allow people to receive voice mails that are transformed automatically into e-mails and SMS that will be forwarded as fax, for example.

2.5.4 Summary

Let's look at the key features of a me-centric future in this scenario (see also Figure 2.6). While the workers do not require highly advanced technology to communicate with each other, the architect will require some highly sophisticated devices and services to support him in constructing the building. The worker communication is similar to the family from the scenario in Section 2.3.

The system needs to know the location of the workers and their status, i.e., what they are doing and what they should be doing according to the project plan. The device could be an enhanced mobile phone integrated into the helmet of the workers to make sure that they are not distracted from their work. The mobile phone needs to take lots of context into account, such as the location and the process, to send back to the central service provider.

In the construction site ecosystem, you can clearly see that tasks could be done more effectively with a little support from this central service, although the workers would probably not like it. While this advanced communication system would improve the performance of the

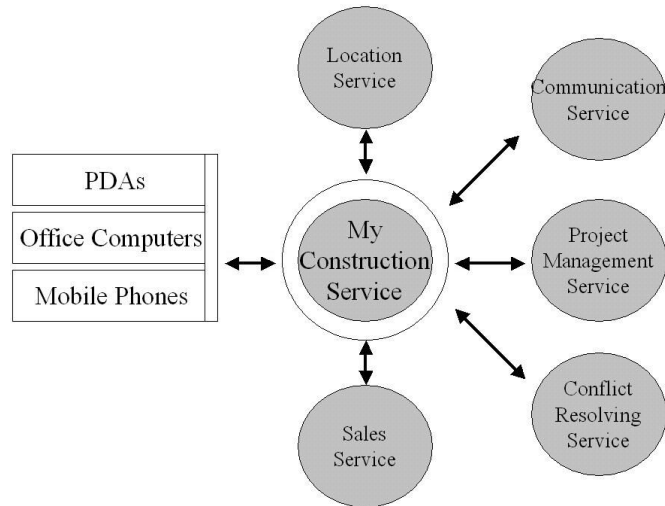


Figure 2.6. Construction Site Ecosystem

workers, only additional technology and services would help the architect to ensure the quality and the timing of the project. Satellites are able to shoot high-resolution images of any spot on Earth. Using these high-resolution images and the original construction plan can help the architect to verify that the plan is executed correctly. An agent would compare the plan with the images every day and would try to correct issues itself; only if this is not possible would it inform the architect about the issue. Another agent would compare the images with the project plan to see when which wall should have been built, for example. Again, the agent would try to resolve conflicts on its own where possible and would inform the architect only in cases where it cannot resolve the issue.

The agent would also provide a completion date, passing on the information to an intelligent sales agent that would send out the information about the building to some prospects. They would be able to wander through the 3D construction plans, which would be completed by the images of the satellite. The sales agent could be owned by the real estate agent from the scenario in Section 2.2.

Based on the progress of the construction, another agent could check with the project plan and start negotiation processes with suppliers for

bricks or windows, for example. The bricks or windows would be delivered just-in-time to the construction site, reducing the need for storage and its associated costs.

2.6 Holiday Break-In

2.6.1 Introduction

Dan and his family leave for a long-needed holiday. They plan to spend two weeks at the Cannes Film Festival. On the third day of their trip, burglars break into their home and steal most of their valuables, including their computer system, home theater system, silverware, and several pieces of priceless jewelry.

During their raid, they also deliberately break several items in the home, including their cabinet of fine china and an antique curio holding Dan's wife's Lladro collection, the value of which the burglars do not realize. A neighbor notifies Dan at the hotel and the family cuts their trip short to return home and assess their losses.

2.6.2 The Present

Unfortunately, this scenario is all too common. Today, there are a few options to protect against such a scenario, but the biggest problem would be the assets in the home. Dan did what he could to prepare for such a scenario.

A year before this burglary took place, Dan's neighbor suggested that he videotape the contents of his home for insurance purposes. Dan takes this a step further and describes each item he photographs to the best of his ability, including his idea of each item's approximate value. Dan stores the videotape in a safe deposit box at his bank.

He also has a motion-sensor alarm system installed, which notifies a central alarm office when activated. The central office first telephones the home to determine whether the alarm was tripped accidentally or it's a real security breach. If they receive no answer within a few minutes, they then notify the police, who are dispatched to investigate.

In this case, the police are notified within ten minutes of the break-in and arrive at Dan's house within twenty minutes. By this time, the burglars are long gone and the damage is done. Dan files his insurance claim using the videotape he created for reference. The problem is that the tape is over a year old, and many of the stolen items are not recorded. To compensate for items he may not remember, Dan claims several items that were not stolen, and fudges the age and approximate value of many of the known items. The insurance company settles

his claim without incident and he is satisfied with the reimbursement. However, a month later, he discovers that his insurance premium is increased by fifteen percent as a result of both his claim and the continuous problem of insurance fraud that plagues the industry.

2.6.3 The Future

As we have seen in this chapter, tomorrow's world will be much different. More things will have chips built into them that will enable them to communicate with other devices nearby, as well as with the Internet. These chips provide information such as the status of the object and its location.

Instead of Dan's videotape, which contains ad hoc and unreliable descriptions of his home's contents and their value, every item in his home - including furniture, cupboards, even walls and artwork, in addition to electronic devices - will register itself into an electronic catalog. The system will check each item's status and record the appropriate information, including exact descriptions of each item's age, condition, and value. Some items, such as the priceless jewelry, will be keyed in manually because they will not contain chips to identify themselves. But high-quality photographs will accompany all items, facilitating claim procedures and assessments.

After an online registration, all of this information will transfer over a WLAN into a database on the home server, which in turn is linked to the insurance company over the Internet, as well as a replacement service that will monitor the status of each item and process replacements should that status change.

Before Dan leaves for his vacation, he tells his alarm system to automatically notify the police if triggered. The system thus knows that Dan isn't home without him having to disclose this sensitive information to another human being. Once Dan activates the service, the objects with chips built in begin to monitor their own status and that of neighboring objects. An XML description is used to communicate the status of each object. Either at regular intervals or in real time, changes in the status are filtered through the database and reported to the replacement service, the insurance company, and to Dan.

When the break-in occurs, the police are dispatched immediately and arrive at the home within five minutes. Dan is automatically notified via his cell phone that a burglary has been reported and that the status of stolen and broken items has changed. The police aren't in time to prevent the burglary, but the thieves only manage to get away with Dan's stereo equipment and some jewelry, and in their scramble to get out of the house, they topple the curio housing the Lladro collec-

tion. As soon as the pieces shatter, the replacement service is notified and replacement pieces are processed immediately. They use an XML description to report the shipping status of each piece.

Within a half-hour, the police use a GPS system to track the chips in Dan's stereo equipment. A van is pulled over minutes later, the burglars are arrested, and the equipment is returned. Dan is reimbursed for the jewelry and the Lladro collection. He claims that his receiver was broken during the burglary, but what he doesn't know is that his son broke the receiver just before they left for their holiday. The database informs the insurance company that the receiver was in fact broken before the burglary took place, so he is not reimbursed accordingly. A month later, Dan discovers that his insurance premium has been lowered by two percent because of the efficiency of the system he has in place.

2.6.4 Summary

This implementation combines several of the me-centric features that we've seen already in this chapter, with several advantages to these features easily recognized in this scenario.

In the present situation, items in the home are too diverse to be treated by a single system such that we will see in the future. Although many household devices today have chips built in, they lack an Internet connection or even a local connection to exchange information. No system knows about the existence of any others, and there is no central system that knows about all the items in the home. To realize such a service, it would be necessary to create an online service where homeowners can register all the items they have in the house. There are several problems with this approach. First, there is no standard way to describe the objects. Either one builds a catalogue of all products that have been built and sold over the last hundred years, or one lets the homeowners describe their goods themselves. With the catalogue, one could be sure that every item is described in a very precise way, but the catalogue would be too big to be handled by the homeowners. Every category would contain thousands of different items, and often the owner does not know how to describe the item best.

The other option of free text does not really help either, because it makes it difficult to find a suitable replacement. Probably the best way would be to create generic categories (such as cabinet), define some attributes (such as size, material, and color), and add a free text field where additional information can be entered to make sure that the item is described as clearly as possible. This still means lots of work without immediate reward, and is still not always accurate.

To make it work, the system could focus only on consumer electronics and household devices. This means that the television set, the DVD player, the stereo, and the refrigerator would be registered at the service. These devices are easy to register because they have code names and numbers and are all standardized. Having their model names and numbers will not only make sure that a certain product is correctly described, but also the age and the approximate cost for the object. This makes it easier to find an equivalent replacement. It also helps insurance companies to evaluate the damage.

The number of objects to be entered into the system will be probably be pretty low—around ten—meaning that the effort is manageable and will ensure that many people will participate. However, the majority of household items would remain unprotected. Additionally, a change in each item's status would be impossible to monitor. We would still rely on unreliable user descriptions of broken or stolen items.

Using today's technologies it is possible to implement the basic idea of this scenario, but due to the complexity of it, we doubt that there is a business plan to support the effort. It would not be possible to implement the technology with reasonable costs. Too many unknowns are in the project to make it reality.

From an infrastructure point of view, the future scenario does not require much upgrade of today's networks. In-house communication can be based on WLAN, an agent will make sure that all affected parties are notified, and stolen items will be tracked via GPS.

To make this service work, it would be necessary for every object in the house to be scanned into the catalog in some way. It will be impracticable for the next few years to introduce chips to all objects, but putting a bar code on each object would work. In this case, broken items would be re-scanned for replacement purposes. Soon, though, we can expect to buy furniture and other items that include a small chip that can transmit the existence and state of each object back to a central service hub automatically.

In fact, each item in the home would have its own agent that would work on behalf of that item and make sure that everything is in order with it. The central service hub would know when to repair or replace the item. The service's ecosystem, shown in Figure 2.7, shows the required infrastructure services and the application services for the different parties involved.

Another advantage to this system is that insurance fraud would be virtually eliminated. Because items report their own status to a central system, insurance customers will no longer be able to make inflated or even false claims, thus keeping insurance premiums for all customers

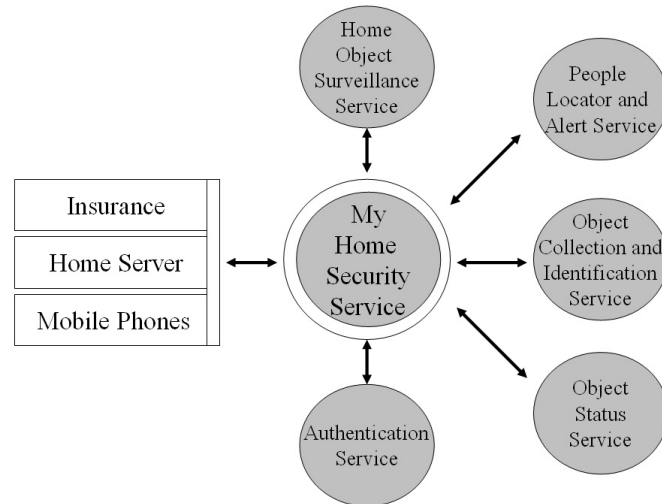


Figure 2.7. Home Security Ecosystem

at reasonable levels. Likewise, customers can rest easy knowing that accurate information is recorded for all of their valuables, so insurance companies cannot underestimate the value of any given item. Finally, stolen goods are easily tracked, not only ensuring their return to their rightful owners, but also increasing the likelihood that thieves will be caught red-handed.

2.7 Extended Manufacturing

2.7.1 Introduction

A manufacturer wants to automate its design process by moving it to the Internet. Instead of creating large-scale designs of its products on paper, it wants to use the Internet in order to reduce the complexity of the task and the number of errors in the designs. Internal designers work on the projects, but also some external companies and freelancers. Today, the designs are sent back and forth by courier. A tracer had to take the paper designs and put them into CAD applications in the end. The idea is to use the Internet as the collaboration platform for all parties involved in the design process.

2.7.2 The Present

In this case also, the technology is available today to move the whole design process to the Internet. To implement such a solution, several steps can be envisioned. It may be a pure communication medium, which allows files to be exchanged through a secure channel, or a full-blown collaboration services that enables people to work together online on a certain design.

If we look at the communication needs, a simple Web site that hosts a set of newsgroups and file exchange would basically do it. In basic technology, there are some challenges that need to be considered when implementing such a solution. As the competition would be very interested in knowing which designs are in progress, the site must be well-secured. To achieve this, security needs to be enforced on the server, during transmission, and on the clients. Only with end-to-end security implemented properly is it possible to use the Internet as a communication medium. This requires a set of tools for the server security, such as a firewall, a secure operating system, and a security procedure. This also requires application security. In many cases, hackers get onto a server by using design flaws within the application, so your applications must stick to the same standards as your secure server. This means that you can make the server as secure as you want, but if your application does not stick to the same standards, the whole system will break at its weakest link.

You must also ensure that data cannot be intercepted during transmission. Today, 128-bit SSL encryption can be used to ensure that nobody can eavesdrop on the connection and steal, alter, or remove information that is flowing between parties. To ensure the security on the client-side, data should only be accessible by password on an encrypted filesystem, making it impossible to hack the client's computer and steal the information. In many cases, the client is the weakest link, as the operating systems are either not secure or not securely configured. Special applications for most client operating systems exist that can create data vaults, making it impossible for hackers to steal the data. They still can hack the system but are locked out of private company information.

Security is further enhanced by introducing smart cards. This ensures that nobody can steal passwords without someone noticing it. An additional card-reader and software needs to be installed on the client's computer as it is still not common to use smart cards in a Web environment.

Once the security infrastructure has been set up, it is possible to use a communication portal Web site to exchange information and data. In-

formation can be exchanged either through Web e-mail or newsgroups that are hosted on the system. The advantage of using secure Web e-mail is that no special e-mail client is required and no additional software needs to be installed by the client. Through Web e-mail and newsgroups, files could be also exchanged and stored on the server.

This simple application would be the first step towards an online collaboration tool. To make this tool happen, several other technologies need to be introduced. First, a design repository is required, which allows all designers involved to contribute. A project management tool and a revision control system are also required to support the collaborative private network. Through this set of tools, designers can share their work online and create project groups online. The revision control system ensures that people can check files in and out, thus making clear who is working on which part. The information about the file usage can flow back into the project management system, if files are grouped into projects and the work can be assigned to the files. Several technologies are available to handle this part of the collaboration.

At this stage, the collaboration process is online, but the actual work is still done offline. In order to move the work online as well, it is necessary to have some visualization tools and CAD applications online. The first part can be realized quite easily today, but online CAD tools are not yet available. Most CAD applications can export DXF files, which means that this should be the base format for all online activity, if designers come from different companies and use different technologies. By allowing only DXF on the server, the system is independent of the CAD system used. DXF can also be easily converted to X3D, which is a new format based on VRML (Virtual Reality Modeling Language) that has been available for several years now. X3D can be visualized in a Web browser, making it possible that work groups can view a certain part of the design and use the newsgroups or e-mail to exchange information about it. This could even be done in realtime via an online chat. Modifying the design online would require a CAD application running. Via VNC⁹, for example, it would be possible to export the display of a client to the others, but this is not really a convenient or good way to work.

Building a basic communication portal is not very costly, since there is free software that could be used for this task. A Web server, Web e-mail server, and a newsgroups server are required. The costly part of the communication portal is the security infrastructure required. The costs are not fixed for the system, but depend on the number of people that are going to use the system. For every user, a digital certificate

⁹<http://www.uk.research.att.com/vnc/>

needs to be created. This requires a public key infrastructure to be built up. Every client computer that is going to use the service will require a data vault. Although the cost per user and client computer is decreasing with the larger number, it still means that a large investment needs to be made today.

The collaboration portal would cost significantly more. The project management software, the revision control system and the visualization tools can be quite expensive and the integration of these tools with the communication platform will create some additional costs. To save costs, an offline integration could be done that will exchange information between the different applications in a batch-mode. This means that the information would be updated every few hours only. Although the information would not be always up to date, it would reduce the integration costs significantly. Once the system proves to be successful, a tighter integration could be envisioned. Still, it will be expensive, as each application has its own set of interfaces and no common exchange format exists.

Designers would create a design, upload it to the server, and have an online discussion about it. Another designer could check it out, modify it, and upload it back to the server. The revision control system would track all changes, making sure that the pre-defined steps in the project plan are followed. In case of major changes in the design or failure to follow the project plan, the project manager would need to revisit the project plan and make adjustments to it.

To many unknowns are in the project to make it reality today, as it is highly complicated to build up such a solution using today's technologies. We doubt that there is a business plan to support the effort or implement the technology with reasonable costs.

2.7.3 The Future

The solution tomorrow will look very similar from an architectural point of view, but the integration will be much easier. Several companies are working on online services that allow the tight integration of the different modules and provide a common interface to the designers. Over time, the need for paper printouts of CAD designs becomes less relevant. More designers will be able to integrate their drawings and ideas into an overall project and can see immediately the effect of their designs in the overall context.

From a technical point of view, things become easier as the integration between the services is standardized. The core service will be connected to various other services. There will be a connection to a security service that authenticates and authorizes the person who wants

to access the extended manufacturing portal. In the ideal future, this service is run by a governmental agency, offering all companies an identification service, which can be combined with a role model that is flexible and can be adapted for each business case. The current problem of proper user identification can thus be easily outsourced, and if run by the government, the element of trust would be very high. Alternatively this service could be run by banks, as their trust factor is also very high, but it would require people to register several times, depending on which bank a particular web service is using. Another model has the governmental agency holding the master database for identification, while the service is operated by the banks. The master database would be required to hold several types of data to provide different authentication mechanisms. Personal data, iris and retina scans, fingerprints, and a digital certificate need to be part of the digital identity of every single person.

Many companies try to set up their own trust centers, but they will not work efficiently if they are not trusted by the companies and people wanting to use the service. Microsoft and Sun are already trying to compete with rival technologies, but neither will win, as they are not considered to be a trusted third-party by the majority of people and because they are too technology-driven.

As the web services infrastructure provides an encrypted exchange of information, it is not necessary to implement something for this particular project. A data vault is not required on the client computers anymore, as the CAD editing service provides the possibility to edit the designs online. Therefore, a local copy of the data is not required, reducing the risk significantly. The CAD editing service is basically the client-server version of the standard CAD application that used to run on a single computer. This service could be hosted at an application service provider that also offers other applications online. This is probably the most complex part of the solution, as redesigning a CAD application to run in client-server mode will be expensive and difficult to achieve. Broadband Internet access and high-end computers will be required to make this service possible. High-speed access is required to download the application on demand and high-end computers to actually execute the application, as it most likely will be written in Java, which is not necessarily the fastest programming language. A part of the CAD application could be the visualization tool, but typically it should be another application, because it provides a different functionality. While the CAD application focuses on the editor, the visualization tool includes features to display the model in different views and qualities and provides a chat feature for online meetings.

E-mail, chat, and newsgroup services can be bought from a community provider and easily integrated. The flow of information between these services is limited, but a tight integration into the project management service and the visualization and CAD tool needs to be provided. The total solution will consist of a Web front-end that will allow designers to do all their work through a unified interface. A separate interface needs to be built for the project managers where they can create their project plans and provide updates to the designers. The project manager can also register designers and assign work packages to them.

Multiple access for the designers to the portal is nice to have, but not a requirement. Designers will require a computer to make changes in the designs. This is not possible through a mobile phone, for example, but in order to participate in discussion groups, alternative access devices should be supported.

2.7.4 Summary

Let's now look at the me-centric features of this implementation. A broadband networking infrastructure is required to support the needs of the designers in order to keep the data flow up. Security is an important issue in this scenario, as the designs could range from product designs for new cars to design plans for a space station. Everyone connected to this virtual community network requires a security clearance, and this applies to the devices that are connected.

In this scenario, no fancy devices or WLAN are required (although they do not harm the scenario). Traditional personal computers are mainly used for this kind of work, but the applications and services need to be made more intelligently. This means the community needs to employ agents to regulate the communication and the workflow and make it more efficient. These agents are especially important if the work is done in a distributed environment. They could ensure that the work is well coordinated.

If you look at the extended manufacturing ecosystem (see Figure 2.8), you can see the importance of the security aspect in this scenario. The security is important to make sure that nobody breaks in, but also to make sure that only people within a design team can add, modify, or delete the portions of the design they are responsible for. The data vault that contains all the information also needs special security features. Intelligent agents would be in place that act as security guards to protect the information. Another agent would be responsible for backup services and high availability of the solution. These agents would work autonomously and try to resolve all the issues related to

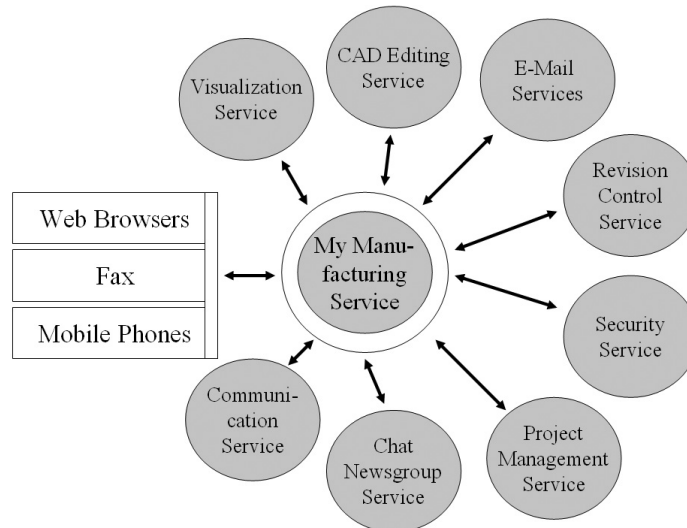


Figure 2.8. Extended Manufacturing Ecosystem

their tasks themselves. This would mean a huge advance in IT management. Today an administrator needs to monitor the whole IT environment; using these intelligent agents, an administrator is required only once in the beginning to set up the agents and tell them what to do.

The communication agents make sure that every party involved in the design knows about updates in the process and changes of decisions. They also ensure that people are working according to the project plan, making automatic static updates to the project plan, and informing other agents about the progress. It is important to note that there are two types of intelligent agents: one that is working on behalf of a person and one that is working on behalf of a process. The agents that work for people help to create the me-centric experience for the user, while the process agents help to get on with the predefined process.

2.8 Fleet Management

2.8.1 Introduction

A large corporation wants to track all its vehicles in use. It therefore sends out an RFP to fleet management companies to introduce a new

service for automatic vehicle detection systems, in order to make sure that employees do not misuse the cars. Part of the new service requires a tight integration into a satellite navigation system, connection to all gas stations, and a fleet management system that is able to track the use of the cars.

In the past, employees used company cars for private excursions, and used the company gas card to refill their private cars. It was never clear which car was used by whom and for how long. The new system is designed to cut down on employee abuse. Additionally, in case of emergencies, accidents, and breakdowns, fleet management would be able to react quickly and efficiently.

2.8.2 The Present

If we look at the possibilities today, it becomes quite clear that the required technology is available. GPS, mobile Internet, and fleet management software packages can track the location and the use of the car. Due to the different technologies used for all these components, it is quite expensive to integrate them and make them work efficiently. The problem becomes apparent if you look at the number of vehicles that need to be tracked. If we talk about five or ten cars, it would be almost possible to track them manually, but fleet management service companies often have thousands of vehicles in operation and rented out to many different companies by a fleet management service provider.

In order to make this system work, the information would need to flow from the car via the satellite to the fleet management service provider, which in turn offers an extranet that allows the company to check the status of each car. Although not very common yet, telemetric data could also be sent with the location to make sure that the fleet management service provider is able to track repairs and mileage of the cars out in the field. This would require adding lots of technology to the cars, as very few support this type of integrated information exchange today.

To realize such a fleet management service, it would be necessary to create an online service where companies can register all cars they use. There are several problems with this approach. First, there is no standard way to describe vehicles. Every manufacturer uses a similar, but not identical way of providing a description of a car. Therefore, the company needs to create its own vocabulary to describe the vehicles in use. This vocabulary can be used to build a catalogue, where one could be sure that every item is described in a very precise way.

Such a system could be very expensive, even for a large corporation, which may outsource the whole operation to reduce the cost signifi-

cantly. Once all cars are registered, employees can book vehicles based on a set of specifications provided by the company. This could range from car transportation between two offices to doing a roadshow with large trucks. The system would check which vehicles are available and their proximity to the desired location. It does not really help if a car in South Africa is available and you need one at the airport in Milan, Italy. The Web site could provide not only a set of vehicle specifications, but add a different view by providing use cases, meaning that an employee does not select the type of car, but the type of situation he or she is in, and the Web site will select the vehicle accordingly.

Once the employee has decided what he needs the vehicle for, the exact route can be specified by an online route planner. This would enable the employee to calculate gas costs and provide the duration of the travel. Today, the information can be stored on a PDA or printed out, for the convenience of the driver. If the vehicle has a built-in navigation system, the employee can re-enter the data and be guided by that system. Today, it is not possible to transfer the data between an online system and a car navigation system automatically.

The employee picks up the car and sets off on the journey. With GPS tracking enabled, the vehicle can be tracked. This way, fleet management could offer additional help if the driver hits a traffic jam or has a breakdown.

En route, drivers would be able to refill the fuel tank using a gas card, with the money being deducted automatically from the account of the company. In theory, this works very well, because it is limited to the car that the employee is using. In reality, this means that the drivers may use the card to buy fuel for something other than the company car. Although the cashier at the gas station should check the license plate, this is often not done thoroughly; hence, the high rate of fraud. It is difficult to overcome this problem today.

Once the vehicle has arrived at the destination, telemetric data could be sent back to the fleet management for statistical reasons. Imagine lots of employees travel from Milan to Florence every day for business reasons. You would like to know which time is the best to get there, how long it takes at a certain time of day, and how much it costs to get there. By collecting this telemetric data and creating the statistics, it would be possible to determine that it is not a good idea to leave Milan at nine o'clock in winter, because the motorway is often very foggy. In summer, it is okay to leave at nine, because most people are already at work and the motorway is free. Many other statistics can be created, to make the use as efficient as possible. Another example could be that ten people are driving from Milan to Florence in the

morning, but all use their own cars. If you can find three people that leave at the same time and return at the same time, you could have them car pool, thus reducing costs heavily.

The complexity of this scenario makes it impossible to introduce such a solution today, although most components are available and could be interconnected.

2.8.3 The Future

The solution in the near future will look quite different. Several companies are working on introducing new technologies to make life easier for fleet management. New cars will all include an online connection either to the Internet or to a dedicated wireless connection to a server. The technology to do so is becoming cheaper every day.

All vehicles need to be equipped with mobile connectivity that connects them securely to a telematics service. This service collects data and provides the information to the fleet management service. The telematics service is responsible for maintaining the connection between the vehicles and the fleet management within the company. This includes operating the satellites and transforming the data that comes from the car. As companies often have cars from different manufacturers, and no standard has been set on how the information should be transmitted, the data is structured differently for every car type or manufacturer. The role of the telematics service is to convert the data from the car and present it in a unified way to the fleet management.

The gasoline payment service is an automated gas station that detects which car is filling up. This can be done via a camera that takes a picture of the number plate and compares it to the database of the fleet management service and makes sure that only the cars that are eligible can pay through the corporate gasoline card. More advanced systems will include a Bluetooth service that enables the gas station to communicate with the car, meaning that the car can tell what type of gas it requires and send the company credit card information electronically. The Fraunhofer Gesellschaft¹⁰ in Germany is even a step ahead by providing robots that fill up the gas automatically. By combining these new technologies, vehicles could be served faster and more efficiently.

The central fleet management service consists of an employee Web front-end, which allows employees to perform the required transaction for booking a vehicle and for getting more information about the service. Employees can register for the service or can be automatically

¹⁰<http://www.fhg.de>

eligible for the service because of the role the employee has in the company. More importantly, the fleet management service provides many administrative functions for managing the fleet and providing information to the users. Tightly connected to the fleet management service is the route planning service that will not only help employees find their way easily, it is also the basis for statistics. The statistics service takes all the data, puts it into a data warehouse, and provides management statistics to the fleet management service.

The travel service provides the fleet management service with information about delays, traffic jams, important news, and other related information. This would mean a huge reduction in effort and lower the entry barrier significantly, if the fleet management service portal ensures that the information is securely stored and only authorized systems and people are allowed to see it.

Multiple access for the employees to the fleet management portal needs to be guaranteed. This means that employees can use the service via the Web, the phone, fax, or any other means of communication. The same is true for the information that comes from the fleet management system. Depending on where the employee is, the information should be sent as e-mail, SMS, voicemail, fax, or any other means of communication.

2.8.4 Summary

Let's take a look at the me-centric features of this scenario. Broadband access is not a requirement here. Although everything needs to be connected, the amount of data that is flowing back and forth is minimal compared to the other scenarios. Telematics data can be easily transmitted in XML files over low-bandwidth networks.

The biggest issue to solve is coordinating a large number of vehicles and their users (see also Figure 2.9). While the company does not need a car for each employee, it can reduce the number of vehicles drastically with good fleet management. Therefore, agents need to collaborate and find out who should use which car when. Each employee who needs a car should also have an agent that makes car pool reservations on his or her behalf based on the schedule of the employee.

To make sure that employees use the cars only for the intended business travel, the route can be planned by the coordination agent, which allows the user to predict the time of the journey and when the next person can use the car. By using the GPS to track the car, it is also possible to see if the employee is taking the shortest route to the business meeting or if the employee is going some other routes to do some other things. This can also be used for filling up the car with gas.

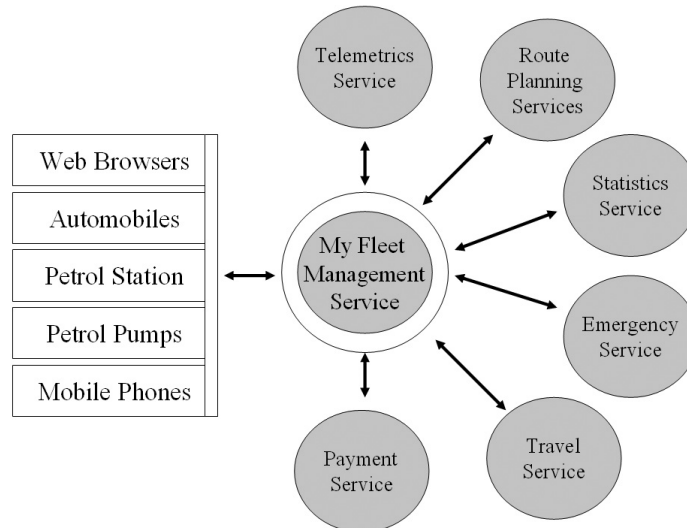


Figure 2.9. Fleet Management Ecosystem

The telemetrics system sends back information about gas usage and allows an intelligent agent to look out for gas stations in the vicinity. The prices are prenegotiated, and the driver can stop at the gas station and fill up the car. The driver does not have to pay personally. As soon as he leaves, the company pays for the amount of fuel that was bought. This ensures that only company cars are filled up.

2.9 Conclusion

As you can see quite clearly, these scenarios are built upon the trends we talked about in Chapter 1. If you look at these scenarios, you will find the kinds of interactions that “agents” would have to have and the kind of protocols that would be efficient or inefficient. You will find more detailed descriptions of agents throughout this book. These agents, their interactions, and their protocols are the basis for me-centric computing architecture—they make it possible to receive and perform delegated work.

The ecosystem of me-centric appliances is glued together by infrastructure that provides enabling communications and servers/services, and to the extent the appliances are mobile, this is about wireless mo-

bility. The key elements of this me-centric architecture are the smart infrastructure, the smart services, and the smart devices. These scenarios should help you to understand what the future can look like and help to validate the general proposition that there are countless arenas to be addressed, with many similarities but with a wealth of diverse functionalities, as appropriate to the domain.

From these examples, it has become quite clear that me-centric computing will be everywhere, changing the nature of every occupation. Me-centric computing consists of a set of capabilities/roles operating in an ecosystem glued together by pervasive computing and communications. In the next chapters, we will go into more detail how these futuristic scenarios can become reality soon.