Platform for Active Information Dissemination

Problem Statement

15-413 Software Engineering
Carnegie Mellon University
School of Computer Science
August 25, 1998

Table of Contents

1. The Problem

1.1. As-is Situation

Within Daimler-Benz aftersales information is created and distributed by different departments. The major information sources today are service, parts, and vehicle information. Depending on the type of information the company utilizes a variety of different distribution channels. Table 1 shows the main aftersales informationtypes, the respective end user applications and distribution channels.

Information Type	Application	Distribution Channel
Service Information	WIS	CD-ROM, Paper, Microfiche
Diagnosis Information	STAR DIAGNOSIS	CD-ROM, Paper, Microfiche
Parts Information	EPC	CD-ROM, Paper, Microfiche
Vehicle Information	FDOK	CD-ROM, On-line
Car Configuration Data	MBKS	CD-ROM
Work Units & Operation Texts	ASRA	CD-ROM
Damage Codes	VEGA	On-line

Table 1. Main Aftersales Information Types, Applications, and Distribution Channels

These distribution channels are typically very reliable but also very slow and inefficient. For instance the distribution of service, parts, and vehicle information to the worldwide Mercedes-Benz sales organization is done via a monthly published set of 12 CD-ROMs. This information is already partially outdated when it gets to the dealer.

1.2. The Problem

Today and in the near future Daimler-Benz is extending its business in terms of new product lines (A-Class, M-Class, etc.) and new models of already existing product lines (S-Class, etc.). The amount of aftersales information is increasing due to the introduction of these new products. With the introduction of new aftersales information systems additional information distribution channels are created, which finally lead to a proliferation of distribution channels. All this makes the information management process (creation, publishing, distribution, installation) from a technology and management point of view more complex and expensive. Looking at today's information distribution channels we can see that they are too slow and inflexible to meet these demanding business requirements. The following picture illustrates the current situation:

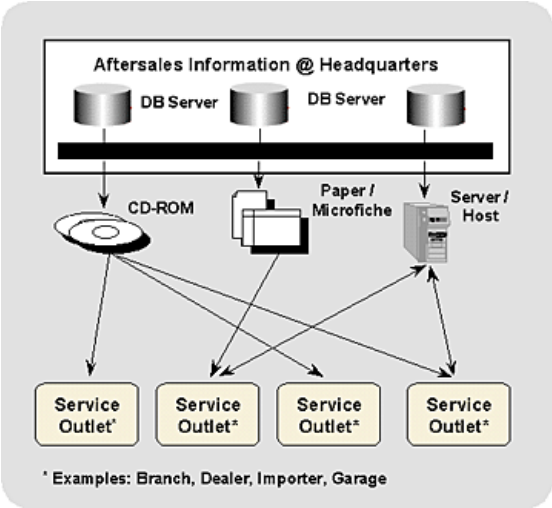

Figure 1: The current distribution process for Daimler-Benz aftersales informations

1.3. The Future IT Platform

STAR NETWORK, as the next generation network based IT platform for Mercedes-Benz Service, needs to provide a common information distribution facility for all applications that run within the framework. Even though STAR NETWORK is conceived as a network based platform giving access to distributed databases, these databases need to be replicated from a few central databases. Without a common information distribution facility, information management within STAR NETWORK and other network based applications continues to be a major concern.

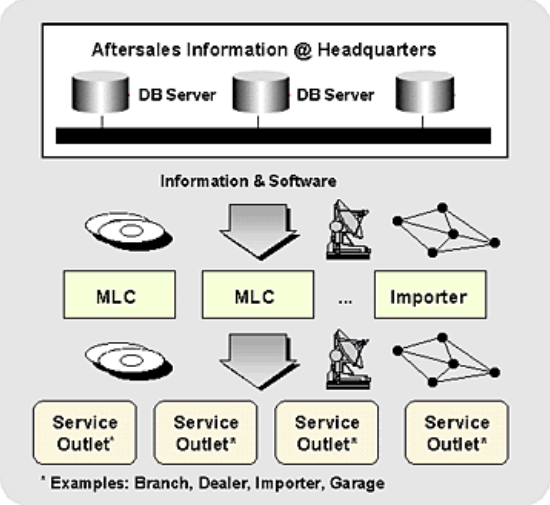

Figure 2: PAID Information Distribution Process

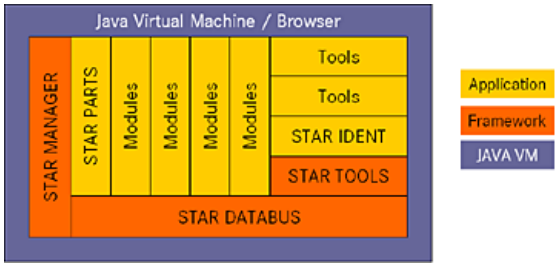

Figure 3: The STAR NETWORK Framework

The STAR NETWORK framework is based on a service based architecture. The architecture subdivides an application in two parts: an application service running on the server implementing the business logic resp. data access and a client part running on the client implementing the GUI and parts of the business logic. Within the client the framework offers a public databus used by the client parts for local communication. There are two classes of client parts: modules and tools. Modules are typical applications. Tools are common services used by several modules. STAR Parts is the application formerly known as EPC (Electronic Parts Catalog). STAR IDENT replaces the legacy system FDOK.

2. Objectives

2.1. Business Vision

Mercedes-Benz is extending its business in terms of new products and market share. In conjunction with the extension of the business the amount of aftersales information grows. In order to meet these business requirements Mercedes-Benz needs to improve its service applications or systems. The timely distribution of aftersales information is a success factor for the Mercedes-Benz Sales and the PAID architecture can be seen as an enabling technology.

PAID as a standardized architecture for the distribution of aftersales information will allow:

- Distribution of all kinds of information (content & software) to the sales organization worldwide through the same facility or technology,

- Low entry and administration cost, scalable (cost proportional to added applications or systems) and...
- Easy development of new applications or systems concerning the distribution of data

PAID is in line with other IT integration efforts, in particular STAR NETWORK©. To leverage the potential of STAR NETWORK© as the next generation Internet-based IT platform for Mercedes-Benz Service, a service like PAID is mandatory. For a business application like STAR NETWORK© to run in various locations worldwide, with the required performance (e.g. response time) and stability (24 hours every day), an information or software distribution facility based on a PAID-like architecture is needed. The combination of online access and the intelligent distribution of most used information to the worldwide sales organization offers the best trade-off between up-to-date content and reliable and fast access.

The scenarios in section three illustrate the distribution processes of three important aftersales information sources: service information, parts information and vehicle information.

•Service information encompasses among others the following document types: Service Instructions, Maintenance procedures, Fault trees, Service Bulletins, Data sheets and Wiring Diagrams.

•Parts information encompasses all the MB spare parts with their assigned parts numbers. Parts Information is one part of the Mercedes-Benz Aftersales Database (MAD). In order to find the appropriate documents to service a certain car another information source needs to be accessed:

•Vehicle information is stored in the vehicle documentation database FDOK. It holds the complete documentation of all cars produced since 1986.

Even though the scenarios focus on the information sources mentioned above, other information sources should be handled in the same way. In general the information distribution takes place on the DB Intranet (Daimler-Benz corporate network) and its extension to the worldwide dealers and service outlets, the DB Extranet. This network is heterogeneous in terms of bandwidth and quality. Other storage technologies like DVD and communication technologies such as satellite-based broadcasting should be considered as additional distribution vehicles.

2.2. Expected Achievements

The publication of the PAID architecture should initiate and lead to a common architecture for information distribution within Daimler-Benz. In a first step, the architecture should enable or support effective distribution of up-to-date information to the worldwide sales or service units. A common architecture should reduce the cost of implementation of new information systems in the future.

From a Daimler-Benz point of view, PAID offers the opportunity of increased consistency of information in the sales organization leading to better service and fewer customer complaints. At the same time, it should minimize distribution costs.

From a dealer point of view, PAID provides instant access to the right aftersales information without the limitations of today's mainly CD-ROM based distribution.

2.3. Design Goals

The PAID system will consist of a set of services distributed over the enterprise system platforms. The software architecture of PAID is based on the following design goals:

•Extensibility: New applications can be added later. These applications subscribe to the services offered by PAID.

•Scalability: PAID needs to be scalable in terms of number of clients, subscribing applications, protocols, and bandwidth.

•Location Transparency: PAID provides location transparent access to data.

Actuality: A dealer always accesses up-to-date information.

To cope with the complexity of these objectives and goals, a software prototype must be developed in the PAID project starting on August 27, 1998. The software prototype will be evaluated by the customer on November 5, 1998 to provide feedback to the developers. This might result in a refinement of the software architecture and the class library. A revised prototype must be delivered in February, 1999.

3. Scenarios

3.1. The scenario-infrastructure

All scenarios within this document are based upon the following fictitious configuration of the Daimler-Benz Intra-respective Extranet. To simplify the description, the network is reduced to the central database of Daimler-Benz at Stuttgart (DBHQ), a local database Mercedes-Benz North America (MBNA), and three different dealers: two of them in the USA and the other one in Germany.

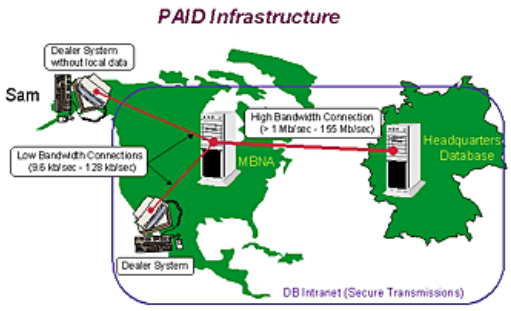

The network driven by PAID consists of two types of computers:

1.Servers: They represent the backbone of the Daimler-Benz Intranet and are connected via high-bandwidth network connections. To make the scenarios easier to understand, we concentrate on two servers: the headquarters database of Daimler-Benz in Stuttgart (Germany) and the database of Daimler-Benz North America in the United States.

2.Dealer systems: The clients of Daimler-Benz use these systems to provide aftersales support to their customers. Dealer systems normally have only low-bandwidth connections (e.g. via modem) to the next server.

3.2. Scenario 1: Adding a dealer

Two dealers (Bratt and Klaus) are planning to integrate their computers into the DB Extranet. Bratt (residing in Beverly Hills) has a high performance network connection to MBNA. Klaus (Munich) has to pay for his 64KBit ISDN connection. Both dealers want a local copy of some important information according to their business requirements.

Because of his good network capabilities, Bratt decides to install the software and data for STAR PARTS (EPC) and STAR IDENT (FDOK). He retrieves the information from the remote MBNA server. After this process has finished, he owns a copy of an up-to-date subset of the complete DB after-sales database.

Because Klaus has to pay for the ISDN, he decides to install the bulk of the software and data from a CD he received from Daimler 2 weeks ago. After the installation, the system automatically contacts the closest server (the central server of DB in Stuttgart) and retrieves all updates necessary to bring the local data up-to-date.

3.3. Scenario 2: No service due to poor network performance

A customer needs a new clutch. Sam tries to access the aftersales database at MBNA. Using the STAR NETWORK standard confi-guration without local data, he experiences a network problem: The transmission rate from MBNA to his computer is very bad, so Sam cannot get the right part number for the clutch, and he does not remember it either. As a consequence, he can not guarantee that he orders the right clutch.

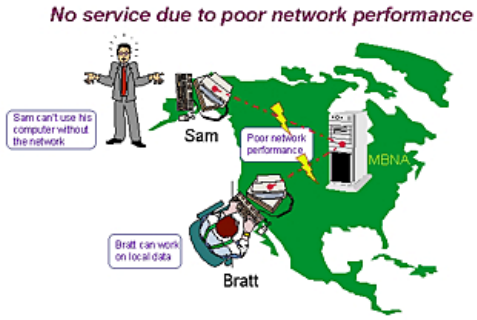

Bratt has the same network problems as Sam, but he has access to the price for the clutch on his machine. He can serve the customer because he does not need to access the aftersales database at MBNA.

3.4. Scenario 3: Dealer's workshop at 8:00 am

It's eight o'clock in the morning. A long line of customers are waiting to get served at Sam´s workshop. The aftersales database at MBNA is under heavy load.

•The response time is so slow that Sam has to ask his customers to be patient. The customers get angry.

Bratt is opening his dealership at the same time. He also has a long line of customers impatiently waiting to get service.

•When he starts up the PAID enabled STAR NETWORK system, he is asked whether he wants to update his local database. He doesn't want to add any load on the network while his customers are waiting, so he decides to delay the update and serves his customers right away. Later in the morning, when his business has calmed down, he is asked again, and now he decides to update his database.

3.5. Scenario 4: Introduction of M-Class in Germany

Klaus, who doesn't have the data for the new M-Class in his local database, has to access the data remotely.

Serving customers, he has to access the same data over and over again, and notices that the response becomes faster each time because more parts of the data are cached locally.

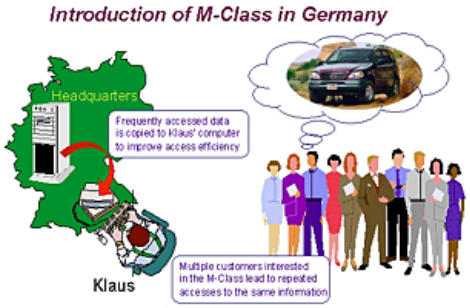

After some days of repeated accesses, the PAID enabled STAR NETWORK system asks him whether he wants to copy the data for the new M-class into his local database.

Klaus says yes, and is happy to learn that this also guarantees that all necessary updates in the future are done automatically.

After this initial update, only incremental updates for the M-Class are sent to his machine.

3.6. Scenario 5: Minimizing Connection Costs.

Sam has a free of charge connection to the internet using his Bell Atlantic local phone number. He decides to be on-line all day long, because he does not pay anything to access all the data remotely.

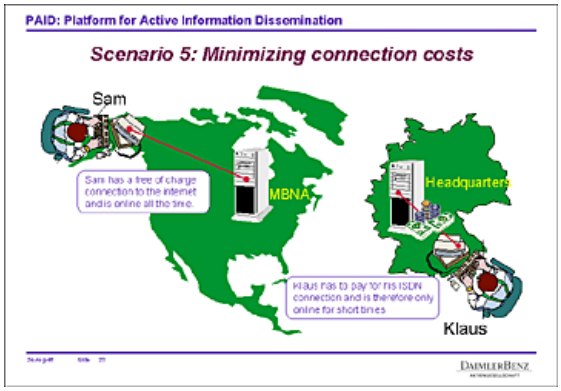

Klaus cannot afford to be on-line all the time because he has to pay for each second he uses his 64KBit ISDN dial-up line. With the PAID enabled STAR NETWORK system, he connects to the internet only for short times, synchronizes his data, or accesses data and infrequently needed information for which he doesn't want to waste space in his local database.

3.7. Scenario 6: Mobile Garage

Bratt receives a phone call from Frank Farmer. Frank lives far outside of Beverly Hills and his truck is broken down. He needs the truck for his harvest and the weather forecast for the next week is bad, so he asks Bratt for a fast repair.

Bratt activates his laptop and connects it to the computer in the workshop. He selects the service, parts, and diagnosis information belonging to

Frank's specific truck and gets an extract of the database installed on his laptop. Due to the limited storage capacity of the laptop, Bratt is asked either to remove some other, yet not necessary, information manually or to let this be done by the system automatically. He selects the "automatic" option.

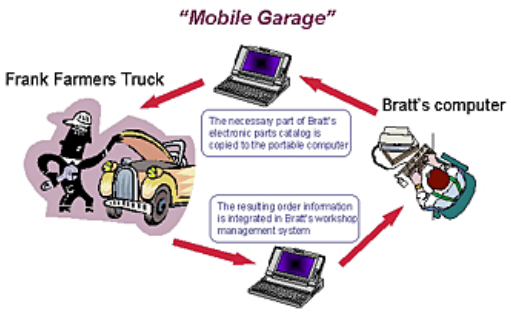

Bratt asks his mechanic Steve to go to Frank Farmer to diagnose the problem. Steve takes the laptop and after several tests with the diagnosis system it becomes clear that one of the ECUs is broken.

Steve searches for the appropriate part number in the electronic parts catalog on his laptop and suggests that Frank should install a new ECU.

Back in the garage, Steve reconnects the laptop with the computer in the garage and the ordering information for the ECU is integrated into the dealer management system.

3.8. Scenario 7: Billing
Problem: Not every dealer is an affiliated partner of Mercedes-Benz.

Independent partners like Sam pay for access to the aftersales database.

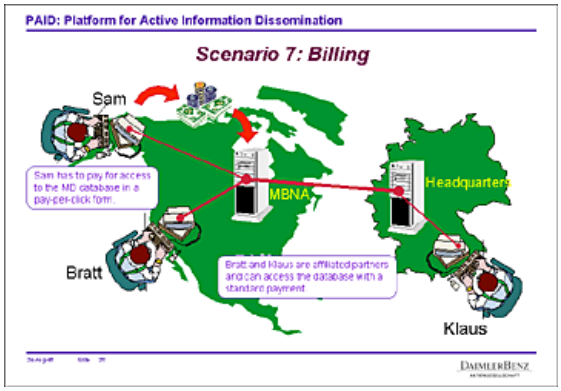

PAID provides accounting mechanisms that allow different forms of billing such as pay-per-click or flat rates.

3.9. Scenario 8: Security
Problem: Information within the aftersales database belongs to different security classes.

Problem: Not every dealer is directly connected to the Daimler-Benz Intra- or Extranet.

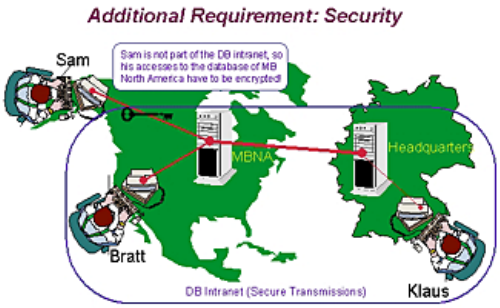

Certain information can only be accessed by certain user groups.

Transmissions to and from third party dealers have to be encrypted.

Different security mechanisms such as encryption and authentication have to be supported depending on the accessed information.

4. Requirements and Constraints

4.1. Requirements

The PAID system will consist of a set of applications and servers distributed over enterprise system platforms. The system realizes an adaptive, selective multicast of information. The PAID enhanced STAR NETWORK is:

- Adaptive because the network is able to adapt itself to the behavior of the users. A typical example for this functionality is given in scenario 5, where frequent accesses to the same data lead to a reconfiguration of the users local database.

- Selective because only the information a user is interested in is transmitted to his computer. No

user receives data that he will not use. The same behavior also applies to the server machines. MBNA for example does not receive information for vehicles which are sold only in Asia.

• A Multicast-Network because PAID does not simply broadcast new information but it instead sends it explicitly to those people who can use it.

The PAID system should also provide the following functionality:

Hierarchical Caching
The informations within the PAID driven STAR NETWORK are cached hierarchically. This means, that on the different servers of the DB Intranet different subsets of the headquarters database are cached.

Push/Pull-Mode Support
The PAID driven STAR NETWORK supports both standard access modes:

• New information from headquarters is pushed through the network to the dealers.

• Any information that is not present at the dealer's local database is pulled from the next server that stores this information.

Security
Data transmissions over the PAID driven network can be encrypted if necessary. The necessary encryption level can either be explicitly specified by the author of the information (e.g. DB headquarters) or based upon the profile of the user (e.g. affiliated dealer or not?) who is accessing the information.

The system provides an authentication mechanism to identify every user.

Incremental Updates:
Unlike in the currently used system (this is valid for the CD based delivery as well as for the standard STAR NETWORK), only new or modified data is transmitted. This feature reduces the amount of data which has to be submitted.

Disconnected Mode
This allows Mobile Computing with laptops or other personal digital assistants. This does not mean, that any application can be installed on a mobile system

Disconnected mode makes the system insensible to unreliable network connections.

Transparency of Medium
All types of media like networks, CD-ROM or satellite-based broadcast are handled in a transparent way by PAID. The end-user can select the medium he prefers.

The software architecture of PAID must also meet the following requirements:

PAID is extensible: New applications and data sources can be added after PAID is launched. This requires an open architecture for PAID and all subsystems of it.

The system is customizable to the needs of the users:

• Every user can select automatic or manually initiated updates.

• Every user can customize his local database according to his business requirements.

PAID is scalable in terms of number of clients, number of applications, used protocols and bandwidth.

Internal management aspects such as data location and replication are transparent for the end-user and only visible to people who are engaged in the management of the system.

4.2. Constraints
The reference applications for the proof-of-concept prototype are applications from StarNetwork. These applications (StarIdent and StarParts) have been selected in collaboration with the customer.

A web site will be established, that allows the client to participate in the project remotely from Stuttgart.

All Applications must be written in 100% Java

The PAID software must be published in Javadoc

The software architecture and class library must be placed under configuration control

All project documents must be published in HTML or PDF

5. System Design

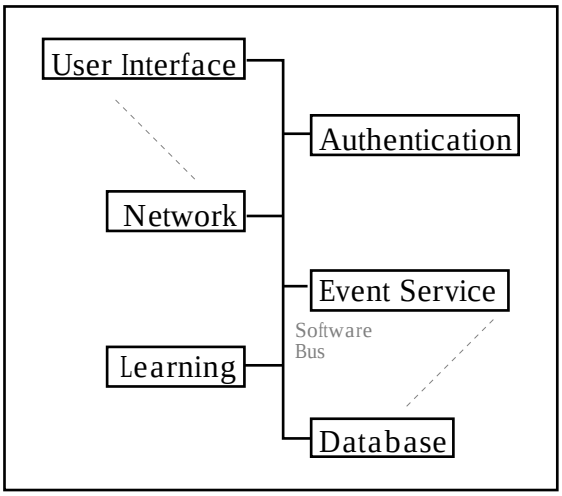

Figure 4: Top level design for the PAID system

1. Authentication & User Management Subsystem
Provides an authentication mechanism to be used within the PAID system. Models all types of users that use the system. Defines and handles access rights as well as user profiles which allow, for example, to create a partial copy of a user's database on one's laptop. The specific tasks of the authentication subsystems are:

• Definition of a secure authentication scheme using a Java card.

• Definition of a user model that identifies all users within PAID and defines their access rights to information.

• Provision of this user model for all other subsystems in PAID; that is, access rights are propagated into the other subsystems.

• Provide the notion of a session and allow sessions to migrate from one computer to another. Sessions allow people to start working on a computer, in the garage for example, and resume at another computer,perhaps a mobile laptop.

2. User Interface Subsystem
Provides a set of graphical user interfaces that deal with the scenarios. The user interface must support computers used in the server administration, as well as the computers at the dealer reception area, in the dealer garage, and the mobile laptop. The user interface subsystem must be able to deal with dialog boxes from other PAID Subsystems, as well as user interfaces from not yet identified subsystems within the STAR network system. Specific tasks for the user interface subsystem are:

• Definition of a style guide to be used for all user interfaces within PAID

• Design of a user interface that is independent from a particularscreen size.

• Provision of a standard for communication with external subsystems that have their own idiosyncratic user interfaces.

3. Event Service Subsystem
This subsystem provides a model based on publish/subscribe, which allows creators of information to reach only nodes such as dealers or services within the Daimler Benz intranet that have expressed their interest, explicitly or implicitly, in receiving this information. Specific tasks for the event service subsystem include:

• Definition of the event model

• Definition of a naming service for publishers, subscribers and Daimler Benz products

• Definition of a set of standard events (Information has changed, node is down, new model has been released, etc) to be used for communication between the publishers and subscribers.

4. Data Management Subsystem
This subsystem is responsible for storing and retrieving all types of data from relational databases as well as from files. Defines a data model for all the information available within PAID such as technical drawings, sales brochures, parts, error codes, marketing information, standard procedures for repairs, and diagnosis procedures. Specific tasks for the data management subsystems are:

• Definition of a recovery mechanism that allows to deal with all types of failures (missing files, electricity blackouts, computer crashes).

• Expansion and reduction of local databases as a response to new data usage patterns.

• Provide access to a dealer management system.

5. Learning Subsystem
This subsystem monitors the network activity and provides statistical information about the activities. The subsystem must be able to aggregate and filter this information in an appropriate format useful for other PAID subsystems as the basis for decision making. Specific tasks for the learning subsystem include:

• Be able to deal with a large number of users (more than 6000 nodes).

• Investigate mechanisms such as machine learning algorithms and agent technologies.

• Investigate existing monitoring and diagnosis tools.

• Provision of reports detailing the network activities with a graphical user interface.

6. Network Subsystem
Provides an efficient transport mechanism that deals with all kinds of information within the PAID system. Specific tasks for the network subsystem include:

• compression on the fly

• integrity checks on the information

• clever routing of information

• Investigation of existing solutions, such as the Marimba Castanet mechanism, to provide push/pull communication modes.

6. Development Environment

The development of PAID is split into two projects at Carnegie Mellon University (CMU) and Technische Universität München (TUM).

At CMU, development makes use of the facilities offered in the clusters on CMU Campus as well as the software engineering lab in Building D, Room 154. This lab can only be used by 15-413 students.

At TUM, the computer laboratory in Room 3175 will be used for PAID. For additional work, students can also use the laboratory in Room 3219 and the computers of the Faculty for Computer Science at TUM.

The following development environments are provided:

Director
A tool that allows fast prototyping of graphical user interfaces.

Together/J (CASE Tool)
A CASE tool supporting the object-oriented model based development of systems written in Java. It provides modeling based on UML and OMT and supports round-trip re-engineering between models and source code as well as generation of documentation.

Database Management System (DBMS)
A not yet defined database management system

to mirror the Daimler Benz databases to be used in the project. The particular system will be selected during the development phase.

CodeWarrior Pro (Interactive Development Environment)
A software development environment with a powerful symbolic debugger for Java 1.1. CodeWarrior runs on the platforms Windows 95, Windows NT and MacOS 8. It allows cross compilation.

Visibroker for Java (Object Request Broker)
Middleware following the OMG CORBA standard. It provides remote method invocation across heterogenous platforms.

CVS
A configuration management system.

7. Target Environment

The development environment at CMU and TUM will be used for demonstrating the prototype. The target environment for the field-test demonstration of PAID will be selected during the development phase.

8. Client Acceptance

The client considers this problem statement to be a broad definition and does not expect that all the functionality mentioned in this document will be demonstrated at the client acceptance test on December 10, 1998. However, the analysis and design should be extensible to include this functionality in a future version of the system.

During the requirements analysis phase of the project the client will negotiate with the software engineers an acceptable prototype for delivery. After the negotiation phase the specific requirements for the client acceptance tests will be baselined. The client expects to sign off on the negotiated deliverables within 4-6 weeks of the client presentation on August 27, 1998.

For a demonstration of the system on client acceptance day, appropriate scenarios will be developed in cooperative work between the client and the project leader respective the software developers.

9. Deliverables

The client expects a successful demonstration of the PAID prototype on December 10, 1998 in the Software Engineering Lab at CMU with participation of remote observers at Daimler Benz Headquarters, Stuttgart and other viewers around the world viewing the demonstration remotely in real time over the Internet.

A set of documents on a CD-ROM describing the requirements analysis (RAD), the system design (SDD), object design (ODD), testing procedures (TM) and user manual of the PAID system should accompany the final demonstration in Munich.