MOBILE INTELLIGENT AGENTS IN ACTIVE VIRTUAL PIPES SUPPORT FOR VIRTUAL ENTERPRISES

Alex Galis <a.galis@ee.ucl.ac.uk>, David Griffin <d.griffin@ee.ucl.ac.uk>, Walter Eaves <weaves@ee.ucl.ac.uk>
University College London
United Kingdom

George Pavlou <g.pavlou@ee.surrey.ac.uk>
University of Surrey
United Kingdom

Stefan Covaci <s.covaci@gmd.fokus.de>
GMD Fokus
Germany

Rudi Broos <rudi.broos@alcatel.be>
Alcatel Bell
Belgium

1. INTRODUCTION

This chapter presents research results achieved in the ACTS AC338 project MIAMI (Mobile Intelligent Agents for Managing the Information Infrastructure). The MIAMI [1], [9] project focuses on the investigation of agent mobility and intelligence and the advantages of mobility in the context of service management and advanced communication services. MIAMI developed a unified Mobile Intelligent Agent framework able to meet the requirements of an emerging open information infrastructure. The key investigation issues were:

- Co-operative management of advanced telecommunications and information services,
- Automated interaction between management systems of different management domains in an Electronic Market context,
- Business-process-driven relationship between the service management layer and the network management layer.

The MIAMI system consists mainly of a two-application testbed: the Active Virtual Pipe (AVP) and the Virtual Enterprise Builder (VEB). The AVP mainly demonstrates the benefit of Mobile Intelligent Agent [20] technology for network and service management. The VEB demonstrates the use of agents for supporting automated negotiation, contracting, and provision of advanced communication and information services.

A Virtual Enterprise (VE) is seen as a temporary federation of autonomous, legally and economically independent companies that come together to collaborate on a common business goal. This type of consortium is formed for a particular task, to share the knowledge, the competence, the resources, and business background of the participants, to contribute to the overall business goal of the VE. The federal character of the VE matches the diversity and complexity of the market offerings and products' structure and increases the flexibility for reacting to the rapidly shifting demands of today's markets.

A VE usually consists of several partners situated at different geographical locations. This physical separation should be as transparent as possible for the members of the VE to enable effective business activities. The Active Virtual Pipe (AVP) provides this location transparency. The AVP is a programmable resource of the information infrastructure that supplies to the VE advanced communication and connectivity services with guaranteed Quality of Service (QoS). AVP is a novel and key example of an Active Service [33]. It provides an abstract view of the dynamically (self-) configurable global connectivity service in charge of the transfer of telecommunication data streams. It is a programmable, dynamic, QoS guarantee Virtual Private Network that can be directly configured according to the demand of the Virtual Enterprise.

2. VIRTUAL ENTERPRISE

The Internet offers any individual the ability to exchange/access information with/from anyone/anywhere in the world as if they were in the same village, the so-called “Global Village”. This makes it possible to create communities of people with common interests where distance is irrelevant. Electronic commerce is part of this scenario since any electronic shop is potentially part of this village. Similarly, workflow can spread across company boundaries anywhere in the world over the Internet. This opens up a completely new way of creating and doing business through Virtual Enterprises.

A Virtual Enterprise (VE) is seen as a temporary federation of autonomous, legally and economically independent companies which come together to collaborate on a common business goal. This type of consortium is formed for a particular task, to share the knowledge, the competence, the resources, and business background of the participants, to contribute to the overall business goal of the VE. The federal character of the VE matches the diversity and complexity of the market offerings and products’ structure and increases the flexibility for reacting to the rapidly shifting demands of today’s markets. A VE is composed of a number of autonomous, real enterprises, or parts of them, which *temporarily* form a new, uniform enterprise which have the main activities of the constituting parts are taking place in the electronic, “virtual” space. This VE is created in order to jointly develop, manufacture and sell products / services on the market. During the life cycle of the VE the various partners will be engaged in adding value to the other’s product or service (build-up a value chain) according to the contract. A major characteristic of a VE is that all partners share the risk and the profit of the joint venture. Each member pursues the strategy of extensive “outsourcing” and concentrates on its core competencies.

Many VEs focus mainly on products which are largely information and communication based, so that the business data have to be transported between the participants over an information transport network. This results in strong requirements to the underlying network infrastructure. The business process data must be transported considering the constraints of each individual business process. Additionally, the stakeholders of a VE are typically not located at the same geographical location and belong to different network management domains. In order to gain benefit of a VE it is required that there is a very close co-operation among all participants like in a single company. This integration of the different parts should not be solved by introducing new levels of local management that results in additional overhead but by directly connecting existing organisational structures. That requires an extensive intra- and inter-organisational use of advanced information, communication, and collaboration services including workflow management services.

For these reasons a highly dynamic and flexible information infrastructure is needed in order to successfully run a VE. This information infrastructure must be able to provide connectivity with guaranteed quality on demand. Conventional solutions like the currently available virtual private networks (VPN) provide mostly few and fixed QoS guarantees (e.g. bandwidth) or no guarantees at all (e.g. VPNs via the internet). The major drawback of the today’s VPNs is their low flexibility to quickly adapt to the changing requirements of the business processes.

Another problem hindering the formation of VEs is the lack of applications that support the automated creation *and administration of VEs*. The creation process includes the finding of adequate business partners, negotiation among them, the definition of business relations and inter-organisational workflow, the collection of configuration information and requirements and the reservation of appropriate resources of the information infrastructure. Because of the spatial distribution and the federal character of the VE and the autonomy of each partner, the administrative management of the VE is a complex task, which requires advanced business support tools.

The MIAMI project provides an extensive VE support environment, which supports the creation, operation and administration phases (including the necessary resources allocation) in an automated way. Figure 1 provides a simplify view of the proposed environment.

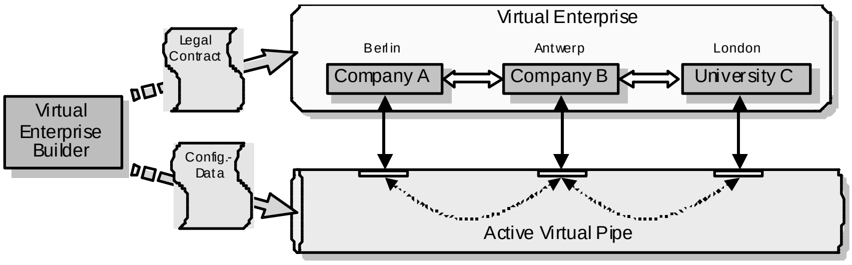

Figure 1: Main Components of a Virtual Enterprise Support Environment

3. VIRTUAL ENTERPRISE BUILDER

The MIAMI Business Model describes the major roles involved in the MIAMI framework and their relations. It is used to provide an overview of the MIAMI framework and to identify interfaces between these entities. In this context, a role is a well-defined business activity, which can not be further subdivided between a number of players. A player can have several roles. Figure 2 shows the MIAMI business model.

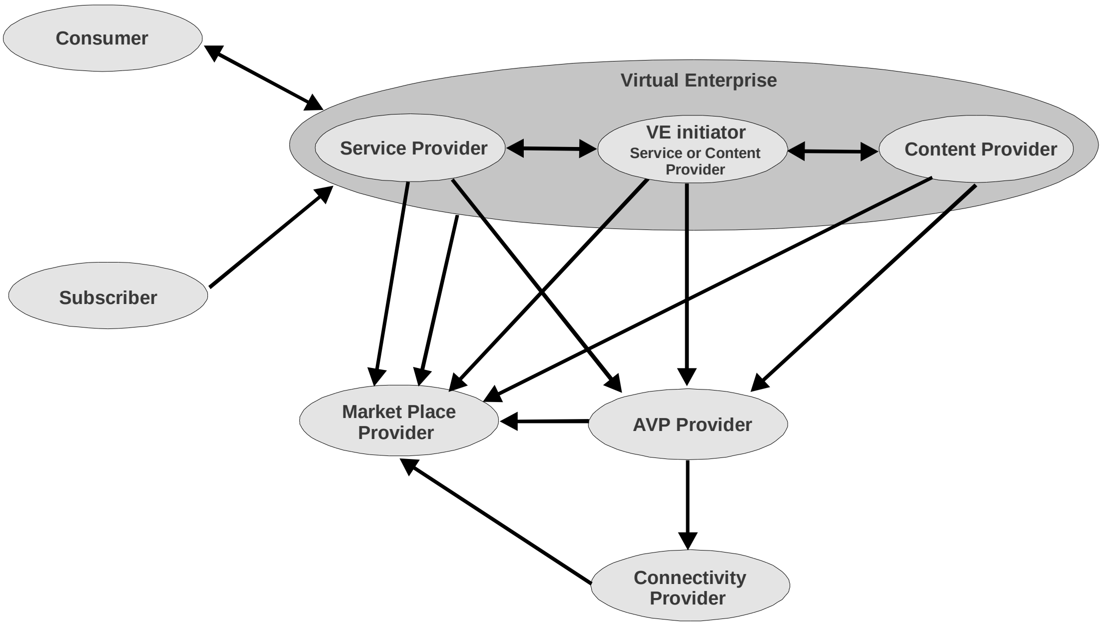

Figure 2: MIAMI Business Model

The MIAMI Business model defines the following roles:

The *SERVICE PROVIDER* provides any kind of service except the content provision covered by the Content Provider role and the connectivity services, which are covered by the Connectivity Provider role.

The *CONTENT PROVIDER* purely focuses on the provision of content.

The *VE INITIATOR* is the founder of the virtual enterprise. The entity in the VE Initiator role can additionally be a Service or Content Provider.

The *MARKET PLACE PROVIDER* provides an infrastructure for exchanging offers and requests.

The *VIRTUAL ENTERPRISE* is a federation of autonomous service and content providers which acts as one unified enterprise.

The *AVP PROVIDER* provides advanced connectivity services and basic VE management services.

The *CONNECTIVITY PROVIDER* is the manager of the transport network and provides transport services.

The *CONSUMER* utilises the service or the product of the VE.

The *SUBSCRIBER* has ordered the product or the service from the VE and pays for it. It does not consume the service. In many cases the Consumer and the Subscriber are identically.

During the lifetime of a VE three major phases can be distinguished, the VE creation phase, the operational phase and the VE dissolution phase. According to its name, the VEB focuses mainly on the support of the creation of VEs. During the VE creation phase, one future participant of the VE, the VE initiator, has an idea for a profitable business that can not be completely realised by him. He defines the business objectives and identifies basic business roles and their relationships needed to achieve the business objectives. He determines which of resources under his control can be utilised and which external resources are needed. With that information the VE initiator constructs an abstract sketch of the VE by using a visual construction tool - the VEB. The VEB provides pre -defined building blocks or customisable templates that represent possible business roles. These building blocks can be arranged to construct a simplified business model of the future VE. The VE initiator can now start to look for potential business partners that fit into the identified roles. For this he can impose additional constraints to each role.

The search for possible business partners is carried out on the Virtual Market Place. This place represents a real life market place operated (by the Virtual Market Place Provider) in the electronic (virtual) space, on which several entities can offer any kind of goods and other ones can look for specific offers.

The VEB uses the temporary business model to advertise the VE idea on the Virtual Market Place and negotiates automatically with potential future partners based on the defined constraints. According to the results of the negotiation, the temporary business model will be refined and the finally selected partners will be assigned to the identified roles of the business processes. Additionally, the relationships between the building blocks will be defined. These relationships represent the communication aspects in terms of abstract connectivity requirements and the workflow aspects that are connected to the business processes. There is a strong relation between workflow aspects, connectivity aspects and business roles because in some cases the business roles and the workflow description implicitly contain several connectivity requirements. In order to allow for a comfortable creation, the VEB provides different views on the VE definition to visualise these distinct aspects. Moreover, these relations are used to automatically determine the required network/communications resources without an additional explicit specification.

The next step is the legal creation of the VE that makes it to a legal body. For this a contract among all participants of the VE has to be made, including the definition of the business relationships. Parts of this contract may be created automatically using data collected by the VEB from the negotiation results and from the description of roles and their relations. Furthermore, the VEB supports the selection of a suitable AVP Provider from the Virtual Market Place.

Because the description of the VE in the VEB contains also implicit or explicit information about the requirements regarding the connectivity services (e.g. addressing information, used services and extent of usage, schedules, and interfaces), the VEB is able to generate generic configuration data for the AVP. This data is automatically transferred by the VEB to the selected AVP provider. According to that information the AVP provider allocates the required communication resources, configures all AVP components and creates the dynamic interfaces provided to the VE. This instantiated AVP reflects the business process from the viewpoint of the connectivity services. After the internal infrastructure of the VE is installed the entire VE is ready to operate.

During the operational phase of a VE a management and administration application is needed to monitor all activities among the VE participants, to handle exceptions and to allow for adapting the VE to changed business processes. The management application has a strong relation to the VEB because in some circumstances the introduction of new VE partners may be required which results in a complete reconfiguration of the VE. For this reason, the VEB could be an integral part of the VE management application. Since the VE is a federation of autonomous partners that can be distributed around the world, this management application must support distributed VE management.

The last phase of the life cycle is the VE dissolution. In this phase the partners leave the consortium (according to the contract specifications) and all allocated resources are freed by the VEB.

Before the virtual enterprise is created, the VE initiator, who has the idea for a business, contacts the market place to find potential business partners which offer suitable services or contents. The VE initiator itself also contributes the virtual enterprise and can be a service provider or a content provider. Figure 3 shows the corresponding section of the business model.

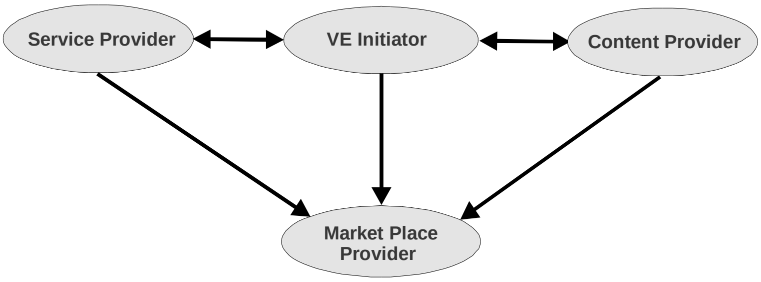

Figure 3: VE Creation Phase

After the negotiations among the business partners, the contract is drawn up and the VE exists legally as a uniform enterprise. In order to obtain the communication infrastructure needed for the operation of the VE, it selects a suitable AVP Provider from the market place, which provides the connectivity and basic VE management services. Figure 4 shows the corresponding section of the business model.

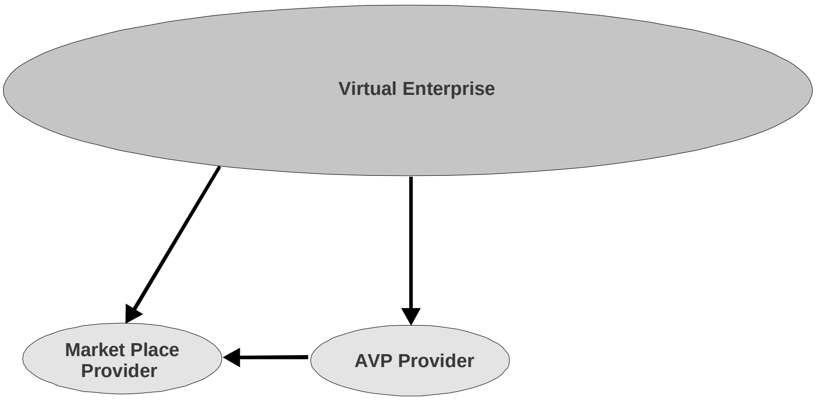

Figure 4: VE Creation Phase

4. ACTIVE VIRTUAL PIPE

Figure 5 shows an overview of the main components of the MIAMI framework.

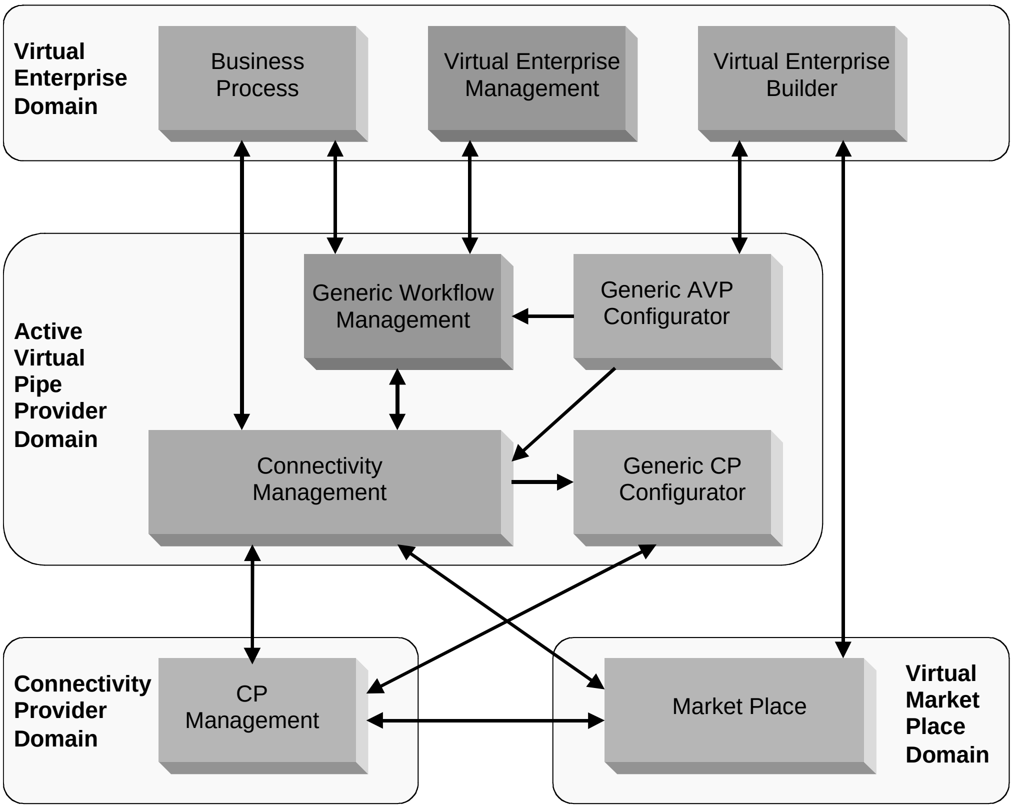

Figure 5: MIAMI Components Overview

A VE usually consists of several partners situated at different geographical locations. This physical separation should be as transparent as possible for the members of the VE to enable effective business activities. This location transparency shall be provided by the AVP. The AVP is a programmable resource of the information infrastructure that supplies to the VE advanced communication and connectivity services with guaranteed QoS. AVP is an example of an Active Service [33]. It provides an abstract view of the dynamically (self) configurable global connectivity service in charge of the transfer of telecommunication data streams or more technically, the AVP is a programmable, dynamic, QoS guarantee VPN that can be directly configured according to the demand of the VE. For instance it is possible that the quality of service parameters of existing connections can be configured dynamically by the VE in accordance to the current business activities.

Since the AVP is not the owner of the transport networks' resources it makes use of the networks of several connectivity providers (CP), which can be dynamically selected and configured in order to adapt to changed market situations. Due to this capability, the AVP is in the position to offer a cost efficient, dynamic, virtual private network that fits better to the needs of VEs than conventional VPNs do.

Additionally, the AVP being the central "node" in the VE that represents the "glue" between all partners includes a generic workflow management facility. This facility can be programmed by the VEB according to the needs of the business processes. It enables an efficient, one-stop-shopping inter-organisational workflow management and provides a generic VE administration service to any VE. Furthermore, in order to reduce the expenditure of running a VE; the AVP allows for outsourcing of common VE management activities. Partners of the VE can focus on their own business logic and outsource the overhead of management of the overall VE to the AVP. Because of its central management function and its role as an advanced connectivity provider, the AVP can establish a business-process-driven relationship between the service management layer and the network management layer. This allows for a business process driven connectivity management. The AVP monitors and controls the inter-organisational workflow and is able to derive appropriate connectivity requirements that are used to configure the underlying network infrastructures. The workflow management facility provides to the VE a connectivity management interface at the level of workflow.

The services of the AVP are provided through dynamically programmable interfaces that can be customised by the VE according to its requirements.

5. VIRTUAL ENTERPRISE DOMAIN

The VE domain contains three functional facilities, the VEB, the VE Management and the Business Process. The VEB is the front end to the VE initiator. It provides support for defining the business processes and connectivity requirements, finding suitable business partners including AVP Providers, negotiating among these partners, creating the VE contract, and configuring the selected AVP according to the contract.

The finding and selection of suitable business partners and AVP Providers is done by using the MIA-technology. After the temporary business model is completed, the VEB creates automatically *MatchMaker Agents (MMA)* which can be configured by the VE initiator with additional constraints and sent to a selected Virtual Market Place. These agents contain the knowledge about the searched roles, the temporary business model and the constraints imposed by the VE initiator. At the market place, the MMAs make announcements by using the Virtual Market Place mechanisms and negotiate with market offers (potential business partners) considering the additional constraints. The potential business partners use a tool similar to the VEB to describe and offer their services on the Virtual Market Place. The MMAs come back with the finally selected partners' offerings (resources). The results of the negotiation can be automatically taken over to become part of the VE contract. The search for the AVP Provider is carried out analogously. The criteria for selecting an appropriate AVP Provider consider also the actual requirements of the VE.

The configuration of the selected AVP Provider is done via the AVP-API. The VEB generates configuration data, which includes the definition of the business process and the connectivity requirements, and programs the AVP accordingly. As a result of the configuration, the Dynamic Workflow Management interface (DWM-IF), the Workflow Invocation interface (WFI-IF) and the Dynamic Extranet Management interface (DXM-IF) are instantiated and provided to the VE. These interfaces are realised by specialised dynamically created static agents. Because of the dynamical instantiation of these interfaces based on configuration data delivered from the VEB, the VE can completely customise and adapt them to its different requirements.

The VE Management facility is a management application within the VE domain, which interacts with the Generic Workflow Management facility of the AVP. It provides the human administrator at the VE premises a graphical user interface to control the VE activities and to visually monitor the business processes. Because the VE Management is working on information defined by the VEB it is closely related to it and has to exchange information about the VE definition with the VEB. In contrast to the VEB, it is a runtime facility, which exists during the whole lifetime of the VE. The VEB can also be an integral part of the VE Management facility. Since the VE is a federation of autonomous partners, a VE Management application can be instantiated at each partners site. Therefore distributed management must be supported. The VE management applications interact with the AVP via DWM-IFs. The DWM-IF allows the VE manager to monitor and control all running workflow instances in the VE.

The Business Process facility comprises the actual VE applications, which make use of the AVP services for communication purposes. The AVP connectivity services can be requested, monitored and managed by AVP capable applications in the VE domain via DXM-interfaces. The WFI interfaces are used by the business processes to accesses workflow data and functions, and to allow the triggering of business activities in the VE domain by the Generic Workflow Management facility.

6. ACTIVE VIRTUAL PIPE PROVIDER DOMAIN

The AVP Provider offers to the VE a programmable information infrastructure. The Generic AVP Configurator receives the configuration data from the VEB and creates and configures all runtime facilities, i.e. a Generic Workflow Management facility and a Connectivity Management facility, and interfaces accordingly. The AVP Provider buys (leases) basic network service resources from different Connectivity Providers and provides high-level connectivity services to the VE. For the selection of the appropriate Connectivity Providers the Virtual Market Place is searched for reasonable offers. The Connectivity Management facility and the Generic CP Configurator manage the Connectivity Providers. The Generic Workflow Management Component and its instance Workflow Management facility provide functionality to support the overall management of the VE. The VEB, the Generic Workflow Management facility of the AVP, and the Virtual Enterprise Management form together a workflow management system which covers, according to the Workflow Reference Model of the Workflow Management Coalition [2] [3], three functional areas of Build-time, Run-time control, and Run-time interaction. The Build-time functions, which are concerned with process design and definition, are covered by the VEB. The Workflow Management and the Virtual Enterprise Management (VEM) facility mainly represent mainly the Run-time control functions concerned with managing the workflow process instances (creation and control). The Run-time interaction functions concerned with the processing of the various activity steps in interaction with human users and IT applications are related to the business objects, which represent entities of the workflow.

The VE initiator using the VEB creates the process definition, which describes the inter-organisational workflow among the VE partners. According to the outcome of the process definition the individual workflow instances are created, executed and controlled by the Workflow Management facility. This facility also provides an interface to the VEM to allow for administration and monitoring of workflow instances from the VE domain. Additionally, the Workflow Management facility instructs the Connectivity Management facility to provide the required connectivity for the execution of the workflow instances. This mechanism establishes the relationship between the business processes and the network management activities.

The Workflow Reference Model of the WfMC defines five interfaces for workflow management systems:

- Process Definition Interface: Bi-directional interface between the modelling and definition tools and the runtime workflow management to exchange process definitions
- Client Application Interface: Interface to support interaction with user interface desktop functions
- Invoked Applications Interface: Interface to support interaction with a variety of IT application types
- Interoperability Interface: Interface to support interaction between different workflow management systems
- Administration and Monitoring: Interface to provide system monitoring and metric functions to facilitate the management of composite workflow application environments

The interfaces of the AVP concerning workflow management aspects can be conceptually mapped to the interfaces specified by the Workflow Management Coalition. This possibly enables an inter-working of the AVP with WfMC compliant applications. Parts of the AVP-APIF cover the functionality of the Process Definition Interface (1). The DWM-IF covers the Administration and Monitoring Interface (5). The WFI-IF covers the Client Application Interface (2) and the parts of the Invoked Application Interface (3). The DXM-IF to the Connectivity Management facility covers parts of the Invoked Application Interface (3). The Interoperability interface is not taken into account.

In contrast to conventional workflow management systems the Workflow Management Facility shall employ Mobile Intelligent Agents in order to investigate their usefulness to manage the business processes. That means, Mobile Intelligent Agents shall represent the base technology of the workflow management infrastructure. In an abstract view, the business process within a VE consists of autonomous entities collaborating to accomplish an overall goal described by process definition. Following this, modelling the business processes and their interactions by using mobile agent technology could provide the following advantages:

- Simple and comprehensive delegation concept: the agents act on behalf of real life entities.
- Decentralised workflow execution and data storage: In contrast to conventional workflow management systems the employment of mobile agent technology leads to a decentralisation because of the distribution of process logic and process data.
- High scalability: Based on the decentralisation through mobile agent technology, it is possible to develop high scalable systems.
- Reduction of network load: The mobile agent technology aims for mostly for local interaction (at the location of the needed resource) which can lead to a reduction of the overall network traffic.
- Reduction of workflow load balancing: Mobile agents make possible the instantiation of an activity at the given place only when this is needed (intelligence on demand). In this way it is eliminated the need of keeping all possible activities up-running all the time in all possible places (reduction of computation resource consumption). Because of the distributed process logic and the negotiation ability and “intelligence” of agents, the agents could perform an automatic load balancing.
- Robust systems: The decentralisation and distribution of logic allows for fault tolerant and robust systems.
- Flexible systems: The negotiation ability and distributed ”intelligence” of agents enable to build systems, which can be dynamically adapted to altered conditions.

A first issue to be considered is the realisation of a traditionally centralised workflow engine with mobile agents. In a first approach, each workflow instance is executed and monitored by at least one mobile agent. This agent can carry its own copy of the process definition. This allows changing a single workflow instance without affecting other running instances. Global changes for all running instances require propagation to all agents.

In contrast to this, a second approach supporting a decentralised storage of workflow definition, e.g. in local manager agents, which contain relevant extracts of process definitions, allows for an uncomplicated changing of all running instances but makes more difficult to change individual.

The second approach represents the idea of shared knowledge and stresses more the negotiation ability and ‘intelligence’ of Mobile Intelligent Agents whereas the first approach represents more a direct implementation of conventional workflow models using mobile software entities. To decide between the two approaches the required performance and demanded flexibility must be weighted against each other.

A second issue is the distributed storage of workflow data. In order to provide a flexible solution, the data transfer should be under the control of the agent. There are two possibilities to approach this issue. The agents can contain workflow data in their bodies, which means the data migrates together with the agent and is stored within the agencies (agent execution environments in a host). This approach requires an enhanced access to the agent transport system in order to make reservations for the transport of the data. The second possibility is to let the agents manage the data transport separately and store the data in external databases. This requires additional transport and storage services resources.

The third issue concerns the involvement of agencies in modelling the business process. Agencies can represent the resources of the execution environment (e.g. free memory or processing capacity). This can be used to build scalable systems, which automatically perform load balancing by forcing agents to move to given agencies. In the same sense it can also be used for fault management purposes, e.g. by forcing agents to migrate to a particular agency. In these cases the agency location shall be transparent for agents, that means they do not require a particular agency to execute their task.

Naturally in the context of mobile agents, agencies represent topographic aspects, which shall be used for the optimisation of network traffic. Agents migrate to other agencies because they simply assume that migration provides a benefit compared to the remote communication/method invocation. If that is always true depends on the distribution of the agencies and the extent of the required communication/method invocation with remote resources. Especially if one computer hosts several agencies or the computers are closely located in a local area network, the migration is not necessarily better than remote communication. Without a real representation of the migration cost and of the remote communication / method invocation cost a well-founded decision cannot be made.

Agencies can also be used to model organisational structures (e.g. a division of a company). This can be used to do more efficient searches by reducing the scope of the search. For instance, if an agent knows he looks for a special service of that division he can migrate to that agency and search locally for an appropriate service. The other solution would be a centralised database, which contains information about all services and their locations. This would contradict the decentralised approach of agent technology.

These three aspects of agencies do not quite fit together. For example, if an agency represents a division it cannot represent topological aspects because another division on the same location would then have to use another agency. The assignment also contradicts with the possibility for load balancing. The use of places to model organisations is also limited to organisations which are not distributed over several locations.

For this reasons the following concept will be used as a base for further study. The agents contain the process definition and the workflow data. The size of workflow data has to be considered during agent transfer. Agencies represent execution environment resources. For agents their location should be transparent. At each location there are manager-agents, which represent organisational structures and physical locations.

The workflow agent contains the process definition and the workflow data. The process definition consists of a list of activity IDs and conditions. Since the workflow agent does not know the current location of an appropriate activity agent, which can perform the searched activity (step 0 – activity agent is triggered by user, step 1 – register for activity), the workflow agents contacts a manager agent (step 2 – find agent for activity). To prevent a bottleneck at the level of manager agents and to enhance the fault tolerance and accessibility to the manager agent, a hierarchical and co-operative structure of all manager agents is proposed. This enables the availability at each location of the knowledge (advertise and acquire mechanism depicted as step 1’) on all topological and organisational aspects. The workflow agent can now migrate to the selected activity agent (step 3 – migrate to activity agent) in order to perform that activity locally (step 4). The activity agent has a relation to a human user (step 5) or to an IT application (step 9) which actually performs the activity. If the activity is completed the workflow agent gets the result and looks for the next activity (step 6 – find agent for next activity) and possibly migrates to another agency (step 7). In case of a split to parallel activities the workflow agent creates subordinated workflow agents which can independently from the main workflow agent perform the activities.

7. CONNECTIVITY PROVIDER DOMAIN

The Connectivity Provider domain comprises one or more transport networks owned by different Connectivity Providers. All Connectivity Providers supply the Connectivity Provider Management facilities, which can be configured by the Connectivity Management facility and the Generic CP Configurator of the AVP in order to provide the connectivity needed by the VE. The Connectivity Providers manage the resources of the transport networks. The AVP Provider (the Generic CP Configurator) can configure the management policies and then according to these policies the CP offers a service management interface to the AVP (Connectivity Management facility). The Connectivity Providers dynamically advertise on the Virtual Market Place temporarily available (unused) service resources, which can be bought by AVP Providers.

8. VIRTUAL MARKET PLACE DOMAIN

The Virtual Market Place domain contains a number of different Virtual Market Places, which can be accessed by the Subscriber, the VE Builder, the CP Management, and the Connectivity Management facility. These market places allow potential business partners, Virtual Enterprises, AVP Providers and Connectivity Providers to advertise their offers. The VE Builder searches for offers from product providers or sends its own product offers. The CPs contacts the market place to offer their services and free service resources. The AVP Provider (Connectivity Management facility) searches the market place for reasonable CP resources, which can be dynamically allocated in order to fulfil requests from the VE.

9. MIAMI NETWORK SCENARIO

The implementation of the MIAMI agent system will support and integrate a variety of technologies and applications. This heterogeneity will be necessary due to the heterogeneous technologies available at the partner sites and due to the requirement of AVP case study.

In the MIAMI project we assume that the network scenario consists of an end-to-end IP network which interconnects the participants of a Virtual Enterprise (VE). In addition to standard best-effort, Internet quality connectivity for general email and web browsing the VE users also require access to higher quality connectivity facilities for real-time services such as high bandwidth video conferencing or for high speed access to large files. The users expect to pay a premium rate for guaranteed quality services, but they also wish to use lower cost, and correspondingly lower quality services for more general-purpose communications.

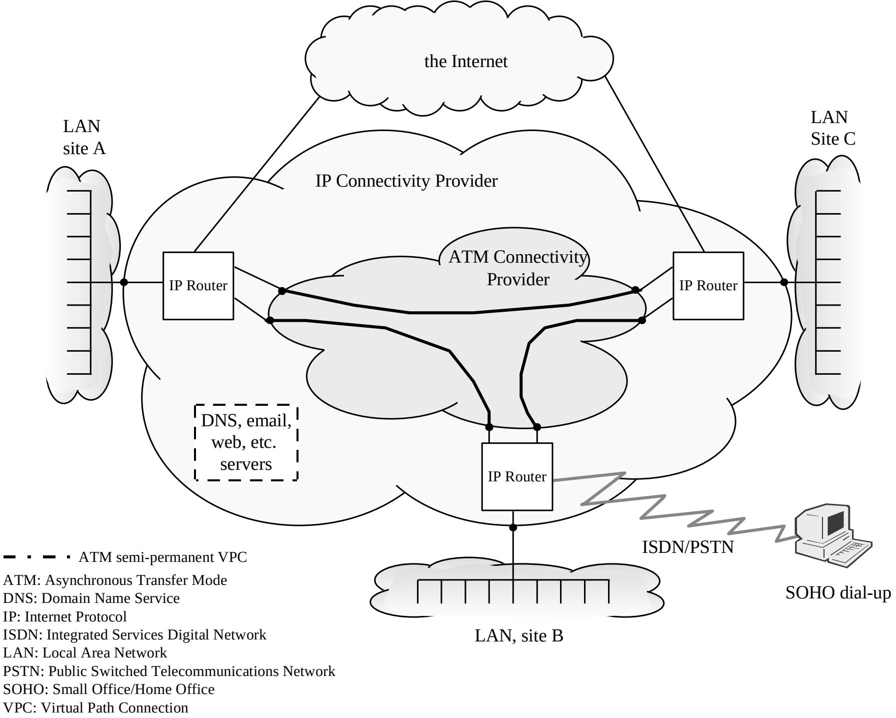

Figure 6 Networking Scenario

In addition to providing access to the Internet, the IP CP also makes use of the services of an underlying CP - offering semi-permanent ATM connections in this case - who is able to provide guaranteed quality leased lines between the IP CP's routers. Figure 6 shows the network scenario we assume in the rest of this paper.

The service provided by the ATM CP is an end-to-end Virtual Path (VP) service offering PVPs (permanent VPs) between specified termination points. Associated with each VP is a number of parameters which defines the capacity of the connection and the level of performance to be provided in terms of end-to-end delay, delay variation and tolerable cell loss ratio. VPs may be created, deleted and modified through client management actions. The clients may monitor VP usage and performance statistics, and initiate fault monitoring activities on their resources.

Issues associated with inter-administration connectivity and federation of management systems are outside the scope of this paper although the dynamic customisable approach presented for a single CP domain could also apply to a multiple CP, inter-domain environment. The remainder of this paper will concentrate on the ATM CP, although the issues discussed are also generally relevant to CPs offering managed services for any network technology.

10. ACTIVE VIRTUAL PIPE APPLICATION

The responsibilities of the IP and ATM connectivity providers in terms of the network resources they manage are showed in Figure 7. The figure shows the layered approach to AVPP. Typically, different organisations will be responsible for the ATM and IP networks. It is also possible that the same organisation be responsible for both network layers, in which case there would be a single CP. The DCM and CP interfaces shown in the figure are introduced in the following section.

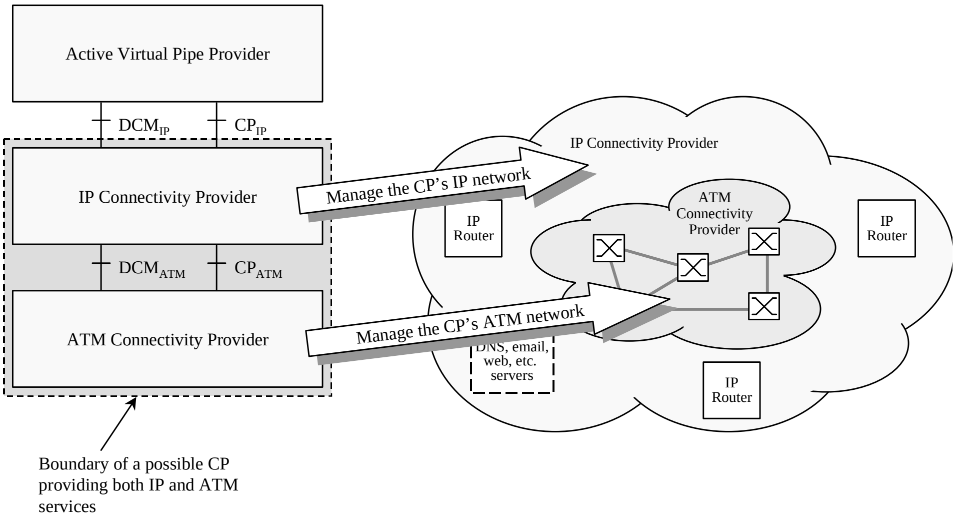

Figure 7 The hierarchy of Connectivity Providers and their responsibilities

10.1. PROVIDER INTERFACES

The ATM CP can be seen as a system with two main classes of interface:

- the interfaces to its clients (either the IP CP or the AVPP), through which it offers a set of services for the management of the connectivity; and
- the interfaces to the underlying Network Elements (NEs) which are ATM VP cross-connects in our scenario. The interfaces to the NEs will be based on the technology offered by the vendors of the switches - either SNMP or CMIP management interfaces; if CORBA-based interfaces are available then this is also an option. It is unlikely that switches will have mobile or intelligent agent enabled NEs in the near future

There are two major interfaces between the ATM CP and its clients: the CP API (Connectivity Provider Application Programming Interface) and the DCM (Dynamic Connectivity Management) interface. Both of these interfaces are based on agent technology. A non-agent based approach to the design of the Connectivity Management is depicted in [15].

Agencies are made visible across each of these interfaces. These are locations where an agent may execute and may be addressed remotely. Within each agency a number of fixed, i.e. permanent, agents will be present together with a number of mobile, or visiting, agents which execute on a temporary basis. In general the fixed agents are created by the host environment- the ATM and the clients of the ATM CP create the mobile agents.

10.2. AGENT-BASED COMMUNICATIONS

We use the term Agent-Based Communication (ABC) to refer to the mechanism by which agents communicate with one another. The term implies that specific protocols and interface definitions are used. These protocols could either be based on general distributed systems techniques for remote method invocation (CORBA or Java RMI, for example) or on higher level semantic/AI languages such as KQML [5] or FIPA's Agent Communication Language (ACL) [21] which support interactions with "semantic heterogeneity”. The specific languages and communications protocols are dependent on the chosen agent platform and are not of direct concern in this paper. Mobile agents may communicate with fixed agents and mobile agents may communicate with other mobile agents using ABC.

The communication between agents may take several forms: to raise asynchronous notifications, to query specific agents to retrieve information and to invoke operations. A mechanism for publishing the facilities offered by an agent operating in the server role is assumed- i.e. a way of formally specifying an agent’s interface. In the design of our system, the Unified Modelling Language (UML) [22] is used to specify formally an agent’s interface, being subsequently mapped to the Java language. It is further assumed that an event/notification service is offered by the host environment for disseminating the events raised by agents based on filtering criteria. Today's agent platforms, including those based on OMG's MASIF specifications, do not currently offer event/notification services. In our environment it has been necessary to implement notifications in a non-generic and fairly inelegant way, on a case-by-case basis. It should be noted that the MIAMI project is currently in the process of extending the Grasshopper agent platform [23] to allow communication between agents using FIPA’s ACL.

ABC extends beyond the local agency to allow communication with agents in remote execution environments. This implies two methods for communications in agent systems:

- either remote operations may be invoked through ABC (in a similar way to traditional distributed systems based on statically located objects); or
- mobile agents may physically travel to the remote agency where they may run in the local environment and invoke the *same* operations through *local* (i.e. intra-node) rather than remote (i.e. inter-node) ABC mechanisms.

By relying on mobile agents, an active and dynamically adaptive management system can be built which is not fixed and limited by initial deployment decisions at system design or build time. The choice of which communications method to use - remote or through mobile agents - is an issue, which may even, be decided dynamically, even at system *runtime.* It is possible to create and deploy a mobile agent when the communications overhead between remote systems rises above a certain threshold, for example. This, however, would be at the cost of physically transferring the agent to the remote execution environment.

10.3. CP'S MANAGEMENT SYSTEM

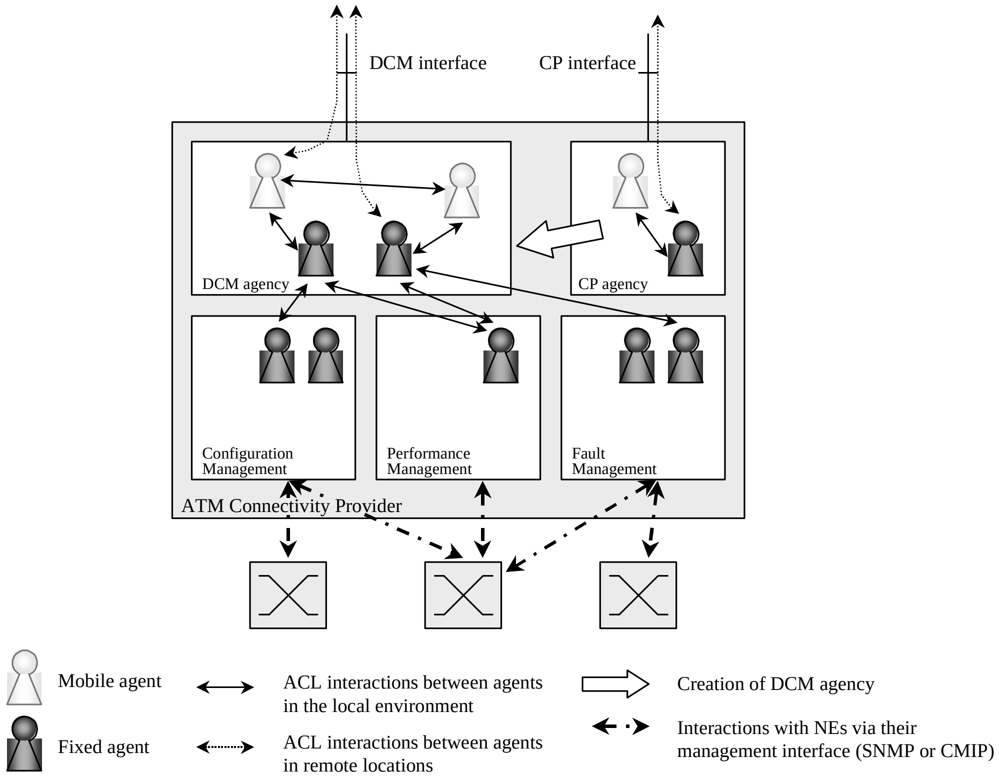

Figure 8 Architecture of an agent-based ATM CP

Figure 8 shows the overall architecture of the CP's management system. The management system consists of three separate agencies for the main management activities of the CP: one each for configuration management, performance management and fault management[1].

Configuration management:

- create connection (ATM VPCs)
- modify connection parameters (e.g. bandwidth) ; tear-down connection
- maintain a database of network resources
- scheduling of resource management and routing management

Performance management:

- QoS monitoring (utilisation of resources, performance of the network)
- Logging of monitored data
- QoS and performance reporting

Fault management:

- Fault detection
- Fault isolation and alarm filtering
- Fault reporting
- Testing

---

[1] Other management functional areas and management services such as accounting and security are also envisioned, but we currently limit ourselves to these three.

In addition there are two agencies, which represent the two classes of interface to the clients of the CP: the CP (Connectivity Provider) agency supports the CP interface, and the DCM (Dynamic Connectivity Management) agency supports the DCM interface. Within these latter two agencies a number of fixed and mobile agents may execute. The fixed agents are provided by the CP, at initialisation time, and form the agent-based interfaces to the basic management services of the CP. The mobile agents belong to the clients of the CP and are dynamically created by remote clients.

Initially, a potential client is unable to invoke CP management services for two reasons: physically it does not have access to a management interface, and legally it does not have a contract with the CP. The first step is to negotiate a contract. A fixed agent in the CP agency offers an interface to allow contract negotiation. This negotiation can be achieved in two ways: either the client creates a mobile agent to move to the CP agency and negotiate locally with the fixed agent; or the client may communicate remotely with the fixed agent.

Following successful contract negotiation, the CP creates an agency and DCM interface for the client. This involves the creation of one or more fixed agents in the DCM agency to offer specific interfaces to the management services, which feature in the contract. The fixed agents tailor (in a *static* sense) the management services of the CP to the requirements of the client and to limit access to the services according to the terms of the contract. For example, not all management services may be made available to all clients, or the geographical coverage of, say, configuration management may be limited to specific locations. In other words, the fixed agents operate as *proxies* to the configuration, performance and fault management services of the CP.

A management service - to create a new VP, for example - could be invoked, either by a locally running mobile agent or by a remote operation from the client. When a management service is invoked the fixed (proxy) agent invokes the corresponding operations on the agents in one of the configuration, performance or fault management agencies within the CP. It is within the latter agencies that the real management work - such as the creation of a VP - is achieved. Through the activity of the CP's configuration, performance or fault management systems, modifications are made to the network elements through their management interfaces (SNMP, CMIP) to reflect the original requests made by the clients at the DCM interface.

10.4. IMPLEMENTATION APPROACH

In the scenario above the interactions between the client and the CP were discussed for contract negotiation, tailoring of offered management services and dynamically invoking specific management services. This section discusses the way in which the configuration, performance and fault management systems *within* the CP are organised.

In general, network management systems are hierarchical with a network-wide view at the top of the hierarchy and an element-specific view at the bottom with zero, one or more intermediate levels according to the needs of the system. This hierarchical approach can be seen in both TMN and TINA architectures. It is assumed that the configuration, fault and performance management systems in the agent-based CP will also follow a hierarchical architecture for many reasons including scalability and compatibility with existing management architectures and information models. There are two ways in which the CP’s management systems could be deployed:

- either through building agent wrappers on existing management software - to represent the highest level of the TMN or TINA compliant system in the agent environment; or
- through building the entire management system from scratch in an agent based way and through building agent wrappers to represent the SNMP or CMIP interfaces of the network elements in the agent environment.

These two approaches are shown in Figure 9. Although the figure shows options for configuration management, the same holds for fault and performance management. Option A shows the approach of wrapping legacy systems with agents at the highest level; option B shows agent wrappers at the NE level and the entire system being built using agent technology.

It is also possible that a hybrid approach (not shown in Figure 9) may be taken. In this case agent wrappers would be provided *at each level* of the hierarchy of a legacy management system. Mobile agents would then be able to visit the hierarchical level that is relevant to their operation and interact with fixed agents representing the legacy software at that level. Mobile agents in such a hybrid environment could relocate by traversing the hierarchy horizontally between subnetworks, or management functional areas, or they could migrate vertically to “zoom-in” or “zoom-out” the level of detail with which they are concerned.

The hierarchical nature of management systems in TMN and TINA (option A, Figure 9) is fixed at system design time and to a certain extent at standardisation time. For each management service to be deployed, the system designers make decisions on the placement of functionality at each hierarchical level and whether to distribute or centralise functionality within a particular hierarchical layer. These decisions are based on many factors, including: the degree of parallelism required; the quantity and complexity of information to be passed between components and whether existing information models need to be modified to support the required information flows; the scalability of the solution; balancing of processing load between management workstations; the complexity of each component.

A promising application of mobile agents for network management is in deploying each management component as a set of co-operating agents (Option B in Figure 9). The way in which these agents are grouped and placed is initially determined by the system designers according to similar criteria as those for the design of static hierarchical TMN or TINA systems. However, now it is possible to revise the grouping and placement decisions during the operation of the management system through the mobility of agents. There are a number reasons for the migration of agents on-the-fly or for spawning new agent instances. These include: to reduce the processing load on an overloaded management workstation, to cater for an expanding set of managed resources or to reduce the quantity of management traffic, or information lag, between remote systems when it crosses a certain unacceptable threshold. This aspect of agent mobility for network management deserves to be studied further but is not the subject of this paper.

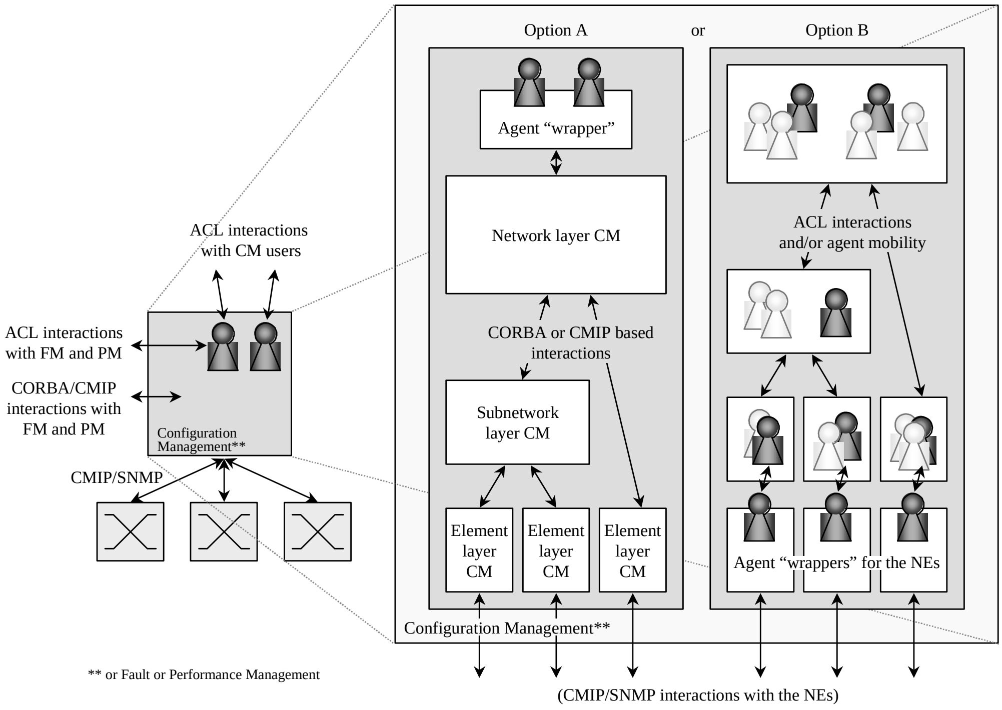

Figure 9- Implementation options for agent-based management functions

10.5. INTEGRATED SCENARIO FOR THE MIAMI MANAGEMENT APPLICATION

So far the basic operation of the CP's configuration, performance and fault management systems have been presented together with the means by which clients in the IP CP or AVPP may access them. This section presents the mechanism by which the management services may be customised and the following section presents three integrated scenarios to demonstrate the power of the agent-based approach to management.

With the basic operation described in the previous section there is very little apparent advantage in adopting an agent based management system - very similar facilities are available in a traditional system based on distributed systems: TMN or TINA for example. However, there was one distinguishing advantage in the system as presented above, that was the way in which the DCM interface could be *customised* - with client-owned agents - to *tailor* the services offered to specific clients - this is one area where agents gain an advantage over traditional software systems.

In the scenarios below mobile agents are able to autonomously interact with one or more management functional area in the server to add value to the original management services. In traditional client-server[2] distributed systems for network management (e.g. SNMP, TMN, TINA) the client is limited to working with the in-built facilities of the server. Any further manipulation of management information, beyond that which was generically provided by the relevant standards or by the developer of the server, must be performed in the client application code.

If customers wish to add value to the basic management services offered by the provider they must have the capability of running a management platform in their premises which supports the protocols and information models offered by the server. In addition, the customers must deploy suitable applications running on their local platform to house the required logic. The client applications must interact with the remote server to receive notifications and to initiate management operations. The quantity of information to be exchanged between client and server is a function of the management activities being undertaken and on the efficiency of the protocols and information models supported by the server for the task in hand. The delay and cost associated with each remote operation is a function of the network interconnecting the client and server. Given the particular protocols and information models supported by the server and the characteristics of the network interconnecting them, it may be not be cost effective or even possible to perform certain management tasks remotely in the client. For example, the cost of communication may outweigh the benefits of performing some management tasks (such as fine grain, real-time monitoring of performance parameters) or the information may be out of date by the time an appropriate course of action has been determined by the client application, if the network delays are too large.

With mobile agents it is possible for a client to program a mobile agent with specific functionality which may then be deployed *at run-time* in the server to add value above and beyond the server's basic facilities. Following deployment, the mobile agent may act autonomously to interact with the local environment and make *local* decisions which may then be implemented as *local* management actions without needing to interact with the remote client every time a decision is required. Through this approach it is possible for a remote client to deploy management behaviour and algorithms *inside* the remote server.

To illustrate the use of dynamically deployed agents we have identified the following three integrated scenarios which are being studied further in the MIAMI project:

- Intelligent reporting

  Mobile agents may respond to reports from both the performance and fault management systems. According to their programmed policies, rather than relaying *all* fault and performance reports back to the remote client, they will only report when certain conditions have been fulfilled. An example might be performance degradation on one connection following the failure of a connection in a remote part of the network, which forms an alternate route. Only the correlation of these two events might be relevant to the client. Alternatively, observed performance degradations might cause the agent to initiate tests to verify that unreported failures have not taken place. This scenario integrates the facilities of fault and performance management.

- Fault repair

  A mobile agent is programmed to listen to fault reports from the fault management system when connections have been interrupted by network failures. According to a pre-programmed policy it can initiate new connection requests between the same end points as the failed connection to restore connectivity. This scenario integrates fault and configuration management.

---

[2] The terms "client" and "server" are used rather than "manager" and "agent" to avoid confusing mobile or intelligent *agents* with OSI management or SNMP *agents*.

- Bandwidth management

  Assuming that performance monitoring agents have been deployed (either by the client or the bandwidth management agent itself) to monitor and report on the utilisation of connections, a bandwidth management agent will be programmed to listen to utilisation reports for certain connections. Depending on the policy for a specific connection the bandwidth management agent may decide to request increased bandwidth on highly utilised connections or to reduce the capacity of a lightly utilised connection. The decision may depend on the cost of changing the bandwidth and so a negotiation between the configuration management agents and the bandwidth management agent may take place. This scenario integrates performance and configuration management.

The integrated scenarios introduced above combine the basic facilities of the configuration, fault and performance management services offered by the server with additional customer specific logic. In other words the customer is able to program the offered management service to a certain degree. This concept has its parallel in traditional management systems through the use of the OSI management Systems Management Functions (SMFs) [24] for event forwarding and logging, resource monitoring [25] [26], and testing [27], albeit in a more limited way. Previous research work [28] [29] has demonstrated how clients can take advantage of these generic facilities to simplify the construction of intelligent clients.

We now consider how these basic facilities could be implemented through the use of mobile agents. Rather than being restricted to standardised capabilities such as the SMFs it is now possible to build entirely arbitrary and powerful behaviour into mobile agents which will be physically located in the managed system’s environment. This embedded intelligence not only allows event reports tailored to the client's requirements to be emitted but it enables the migration of the client's logic and decision making algorithms to the server. This has an obvious impact on reducing the quantity of management traffic between remote systems and achieves more timely access to information generated by the remote server.

In OSI management, SMFs were standardised by international organisations and encapsulated in the compiled functions of OSI agents. In CORBA-based management systems the SMF-like facilities could be determined by the designers of the systems and embedded at design and system-build time; with mobile agents and intelligent reporting in agent-based management systems the SMF-like facilities can be enhanced and extended almost infinitely and deployed *at run time*.

As seen in the examples above, it would be very difficult to capture such behaviour in traditional TMN or TINA systems without standardising such a bandwidth management service at the Xuser or ConS interface. If such a service was to be standardised it would be difficult to capture all possible potential behaviours that clients may request without making a comprehensive and therefore complex specification of the service in GDMO or IDL. However, through the use of programmable, intelligent agents based on mobile code for dynamic and customisable network management this is achievable and deployable on the fly and at the whim of the client. This is clearly a very powerful application of mobile agents for telecommunications management.

11. ANALYSIS

Traditional and emerging frameworks for network management such as TMN and TINA allow customers electronic access to management services. These services, however, are fixed in the sense that new features can only be added after a lengthy research-standardisation-deployment cycle. In this paper we have discussed the advent of mobile agent technologies and how they may enhance traditional connectivity management services making them dynamically customisable by clients. We presented three examples which add value to the offered management services as perceived by and required by the *client* rather than by the researchers, standardisation bodies, equipment vendors or service providers. Clients may introduce their own value-added logic during service operation to cater for the dynamics of their environment and to enforce their own policies.

This prompts for a new paradigm for building network management services. Instead of providers building services that attempt to encapsulate the requirements of all clients, they build the necessary hooks and let the clients apply their logic. Customisation and programmability of management services was always possible in traditional systems based on client-server paradigms through the development of client applications in the customers' premises management platform to capture the required logic and intelligence. However, the cost and efficiency associated with such remote operations compared to the proposed agent-based approach should be considered.

The initial approach for the configuration management domain in MIAMI was to base it on an existing TMN system for ATM PVP set-up originating from the ACTS MISA project [30], to which an agent interface would be added (option A in Figure 9). For logistical reasons, the final implementation is based on *static* agents internally (Figure 9, option B), which communicate using remote method calls. This implementation could be also based on distributed object technology e.g. CORBA. Agents were chosen for two reasons: first for uniformity, since there is no need for an adaptation agent-based interface; and second for evaluating mobile agent platforms in the same role as distributed object frameworks. It should be finally noted that we do not see any immediate benefits from applying mobile agent technology to configuration management.

On the other hand, the performance and fault management systems use agents internally (Figure 9, option B) in a way that mobility is exercised and exploited. In the performance management domain, customised agents replace, augment and allow customisation of the functionality of the TMN/OSI-SM metric monitoring and summarisation objects [25][26], while in the fault management domain customised agents replace, augment and allow customisation of the TMN/OSI-SM testing objects [27]. In both domains, mobile agents are instantiated at the “network management level” of a management hierarchy according to requests originating from the DCM domain, migrate to network elements and perform relevant tasks locally, enjoying minimal latency and reducing network traffic. Details of the performance and fault management approaches are described in [31] and [32] respectively.

In summary, agent mobility in the presented network management architecture is used in a fashion, which we would term “weak mobility”. Mobile agents are instantiated at a control point by a master static agent and then move to another point (i.e. network node) where they stay until their task is accomplished, this can be considered as an intelligent software deployment activity. The key benefit of this approach is *programmability*, allowing clients to “push” functionality to a point offering elementary hooks, which can be accessed to provide derived, higher-level services. In a similar fashion, we could term “strong mobility” as a situation in which a mobile agent moves from point to point using its built-in logic, adapting to changing situations in the problem domain where it is involved. We have not yet found convincing cases in our network management research where strong mobility could offer substantial benefits.

12. CONCLUSIONS

A dynamic framework for creating and realising business in the information infrastructure environment has been proposed. The framework capitalises on the benefits of the Mobile Intelligent Agents technology to offer an easy to use automated and efficient solution to the operation of a VE across all its lifecycle phases. The framework supports customisation and as such enabling innovation in the area of future infoware products. The in this paper proposed model for the VE Domain can be applied equally to the other identified domains, as these can be at their turn individual VEs.

A key mediator role of the VE framework, Active Virtual Pipe (AVP) between the business and the service and network management levels has been developed. This mediator role is a novel way to integrated legacy network and service management solutions into a business driven virtual enterprise environment using Mobile and Intelligent Agent technology. AVP is a novel and key example of an Active Service. In this paper we have discussed the advent of mobile agent technologies and how they may enhance traditional connectivity management services in a form of AVP services making them dynamically customisable by consumers.

13. ACKNOWLEDGEMENTS

This paper describes work undertaken in the context of the ACTS MIAMI (AC338) project. The ACTS programme is partially funded by the Commission of the European Union.

REFERENCES

[1] MIAMI Project Public Web Site, http://www.fokus.gmd.de/research/cc/ima/miami/

[2] Workflow Management Coalition, The Workflow Reference Model (1994)

[3] Workflow Management Coalition Web Site, http://www.wfmc.org/

[4] http://drogo.cselt.stet.it/fipa/

[5] Labrou, Y., and Finin, T., A semantics approach for KQML - a general-purpose communication language for software agents, Third International Conference on Information and Knowledge Management (CIKM'94), November 1994.

[6] ITU-T M.3010, Principles for a Telecommunications Management Network.

[7] TINA consortium, Overall Concepts and Principles of TINA, Document label TB_MDC.018_1.0_94, TINA-C, February 1995.

[8] Mobile Agent System Interoperability Facilities Specification OMG TC Document orbos/97-10-05, November 10, 1997. ftp://ftp.omg.org/pub/docs/orbos/97-10-05.pdf

[9] MIAMI project deliverables - , http://www.fokus.gmd.de/research/cc/ima/miami/

[10] CORBAservices: Common Object Service Specification", Object Management Group, 1996

[11] "Common Object Request Broker: Architecture and Specification", r2.0, Object Management Group, 1995

[12] Finin T et. al. (1994) KQLM as an Agent Communication Language, Proc CIKM'94, ACM Press

[13] Galis A, Brianza C, Leone C, Salvatori C, Gantenbein D, Covaci, Mykoniatis G, Karayannis F - Towards Integrated Network Management for ATM and SDH Networks supporting a Global Broadband Connectivity Management Service, in "Intelligence in Services and Networks: Technology for Co-operative Competition", Mullery A (Ed.)- Springer- Verlag, Berlin, 1997, ISBN 3-540-63135-6

[14] Griffin, D., Pavlou, G., Tin, T., "Implementing TMN-like Management Services in a TINA Compliant Architecture - A Case Study on Resource Configuration Management," Intelligence in Services and Networks: Technology for Co-operative Competition, ed. A. Mullery, M. Besson, M. Campolargo, R. Gobbi, R. Reed, pp. 263-274, Springer Verlag, 1997.

[15] ACTS MISA AC080- Consortium and Deliverables- Management of Integrated SDH and ATM Networks, http://misa.zurich.ibm.com/

[16] "Network Resource Architecture," The TINA Consortium, Version 3.0, Feb. 1997.

[17] "Network Resource Information Model," The TINA Consortium, Document No. TB_LRS.011_2.0_94, Dec. 1994.

[18] Ranc, D., Pavlou, G., Griffin, D., Galis, A., "A Staged Approach for TMN to TINA Migration", proceedings of the TINA'97 Conference, Santiago, Chile, November 1997, Proceedings to be published by IEEE

[20] Kurt Rothermel, Radu Popescu Zeletin (Eds.) - Mobile Agents, First International Workshop, MA'97, Berlin, April 1997, IBSN 3-540-62803-7 Springer Verlag

[21] Agent Communication Language, FIPA 97 Specification, Version 2.0, October 1998, http://www.fipa.org/spec/FIPA97.html

[22] Booch, G., Rumbaugh, J., Jacobson, I., The Unified Modelling Language User Guide, Addison-Wesley, 1999.

[23] http://www.ikv.de/products/grasshopper/index.html

[24] ITU-T Recommendations X.730-750, Information Technology - Open Systems Interconnection - Systems Management Functions.

[25] ITU-T X.738, Information Technology - Open Systems Interconnection - Systems Management: Metric Objects and Attributes, 1994.

[26] ITU-T X.739, Information Technology - Open Systems Interconnection - Systems Management: Summarisation Function, 1994.

[27] ITU-T X.745, Information Technology - Open Systems Interconnection - Systems Management: Test management function, 1994.

[28] Georgatsos, P., Griffin, D., Management Services for Performance Verification in Broadband Multi-Service Networks, in - Bringing Telecommunication Services to the People - IS&N'95, Clarke, Campolargo and Karatzas, eds., pp. 275-289, Springer, 1995.

[29] Pavlou, G., Mykoniatis, G., Sanchez, J., Distributed Intelligent Monitoring and Reporting Facilities, IEE Distributed Systems Engineering Journal, Vol. 3, No. 2, pp. 124-135, IOP Publishing, 1996.

[30] Karayannis, F., Berdekas, K., Diaz, R., Serrat, J., A Telecommunication Operators Inter-domain Interface Enabling Multi-domain, Multi-technology Network Management, Interoperable Communication Networks Journal, Vol. 2, No. 1, pp. 1-10, Baltzer Science Publishers, March 1999.

[31] Bohoris, C., Pavlou, G., Cruickshank, H., *Using Mobile Agents for Network Performance Management*, to appear in the Proc. of the IFIP/IEEE Network Operations and Management Symposium (NOMS'00), Hawaii, USA, IEEE, April 2000.

[32] Sugauchi, K., Miyazaki, S., Covaci, S., Zhang, T., Efficiency Evaluation of a Mobile Agent Based Network Management System, in Intelligence in Services and Networks - Paving the Way for an Open Service Market, Zuidweg, Campolargo, Delgado, Mullery, eds., pp. 527-535, Springer, 1999.

[33] Stefan Covaci (Editor) - "Active Networks"- Proceedings First International Working Conference, IWAN'99 - Berlin June 1999, Springer ISBN 3-540-66238-3