TAPAS

IST-2001-34069

Trusted and QoS-Aware Provision of Application Services

TAPAS
Periodic Management Report
PM1 – PMR2

Report Version: Deliverable PMR2

Report Delivery Date: October 2003

Classification: Public Circulation

Contract Start Date: 1 April 2002 **Duration:** 36m

Project Co-ordinator: Newcastle University

Partners: Adesso, Dortmund – Germany; University College London – UK; University of Bologna – Italy; University of Cambridge – UK

Project funded by the European Community under the "Information Society Technology" Programme (1998-2002)

WORK PROGRESS OVERVIEW

The TAPAS project officially started on 1 April 2002. The overall objective of the TAPAS project is to develop novel methods, tools, algorithms and protocols that support the construction and provisioning of Internet application services. The project will achieve the overall objective by developing QoS enabled middleware services capable of meeting Service Level Agreements (SLAs) between application services and will enhance component based middleware technologies such that components can be deployed and interact across organisational boundaries. The project will develop notations for expressing SLAs to enable specification of QoS, such as the availability as well as trust relationships. SLA trust specifications will be used for deriving service invocation primitives enriched with authentication, non-repudiation mechanisms, with or without the involvement of trusted third parties. Middleware services and architectures will be developed using open source application servers and widely used component technologies such as CORBA and Java.

A Second IAB meeting was held on 3 April, in London, ahead of the Annual Review, which was attended by Dr. Tobias C. Kiefer, Dr. Stuart Wheater (Arjuna Technologies), Paul McKee (BT exact Technologies). We presented the results of the first year work, along the lines of the Review.

The plenary workshop (25-26 September) held in Bologna, Italy was organized and carried out.

The primary focus of the work during the reporting period was on technical deliverables due after 24 months:

- D8 – Container for Group Communication
- D9 – Container for Trusted Coordination
- D10 – Container for QoS Monitoring
- D13 – Second Year Evaluation and Assessment Report
- D17 – Updated Dissemination and Use Plan

In addition, 18 month deliverable report, D3, Service composition and analysis, was discussed and reviewed.

Newcastle University

The Reviewers had asked for some clarification on the TAPAS architecture described in the deliverable reports D5 and D7. We worked on preparing a document, an extension of D5, entitled 'An Overview of the TAPAS Architecture' which explained the various subsystems that make up the architecture.

Our main effort involved working on WP3 related deliverables, D8 – Container for Group Communication, D9 – Container for Trusted Coordination and D10 – Container for QoS Monitoring.

D8 – Container for Group Communication:

- In depth study of reliability issues on multicast on existing publish/subscribe systems
- Design of a protocol for QoS adaptive group communication (to be the core of the Event Notification System)
- Simulation of such a protocol
- Analytical approximation of such protcol
- Implementation of such protocol (in progress)

D9 – Container for Trusted Coordination:

Earlier work on non-repudiated information sharing using B2Bobjects, described in D5, has been extended to include transactions. Work is going on to incorporate these ideas within the J2EE component middleware.

D10 – Container for QoS Monitoring:

Overall architecture for monitoring hosted applications, possibly by third parties has been worked out in association with Bologna, and discussed at the September meeting in Bologna. Plans for implementing the monitoring subsystem have been formulated, including how SLA specifications, written in SLAng can be used.

Newcastle and Bologna have worked on the design of QoS enabled application server, with its implementation using JBOSS.

Adesso

WP2 – Design

Discussion regarding the application server decision:

We have been researching and compiling feedback to the JBoss /JOnAS memo, written by the colleagues in Bologna, in form of emails and memos. The discussion of aspects such as JMX and MBean usage, performance and scalability of interceptor concepts and Reflection utilization led to the decision of choosing JBoss during the plenary meeting in Bologna in September, following the suggestion of Giorgia Lodi.

Discussions of architecture and use cases presented in D7

In preparation of the evaluation tasks in Workpackage 4 we have been discussing QoS scenarios with the middleware-related teams. D7 contains as well a first set of use cases that apply to QoS middleware. We have been discussing the use cases as well as the general approach regarding functional and non-functional requirements and it's suitability for research project evaluation with the partners.

As a result we decided to conduct a use case analysis combined with an analysis of non-functional requirements, including QoS violating scenarios.

Feedback on virtual contracts and virtual enterprises

D5 contains some work on virtual contracts and enterprises. We gave feedback from the industrial point of view, based on usual scenarios.

WP3 – Implementation

Feedback regarding J2EE QoS implementation issues and feasibility

We have been in contact with the different teams to discuss and give feedback about diverse issues. These issues included the component replication approach proposed by Graham Morgan, extensions to J2EE servers and clients to support QoS features and network QoS details.

WP5 – Dissemination

Our current dissemination activities are focussing on competing attention within the IT business community in Germany. In order to address as many decision makers as possible we contributed in different forms to articles in German magazines (see as well Publications).

WP6 – project management

After evaluation of the first year review in April further project planning helped to get a first, more detailed picture of workpackage 4.

In accordance to the review results the project team was intensifying the interactions among the different partners. Driven by these integration approach we joined the executive board meeting on 12 June in London at UCL, where the most important aspects were discussed.

UniBo

During the last six months of the TAPAS project, our activity has concentrated principally on WP3, of which we are responsible. In addition, we have provided contribution to WP4. Specifically, as part of our activity in WP3, we have investigated the detailed design of the TAPAS middleware architecture, in order to proceed with the prototype implementation of the QoS-aware TAPAS middleware core services, as planned in WP3. To this end, we have carried out a careful examination of two candidates, open-source, middleware platforms, namely JBoss and JOnAS, and compared and contrasted these two platforms in order to assess which suits better the TAPAS project implementation needs. (This work is described in the TAPAS project report mentioned below). This work has stimulated a number of relevant comments from our partners that have led us to propose the use of the JBoss application server, in our TAPAS middleware implementation, as the latest version of this application server appears to be better documented, more stable, and incorporating a larger number of services than JOnAS. Then, we have started some preliminary implementation exercise, using the JBoss application server. Moreover, we have contributed to the WP 3-Task 3; our contribution is summarized in [3]. As part of WP 4, we have investigated the design and development of a scalable architecture for responsive auction services over the Internet [1]. The scope of this exercise has been that of providing our project partners involved in WP4, with an optional architecture to use for the purpose of validating the TAPAS middleware.

UCL

The approach of the TAPAS project to guaranteeing QoS for internet services is to introduce Service Level Agreements (SLAs) to form part of the legal contracts binding suppliers and clients. In a marketplace for QoS-aware services it is contingent upon service providers to reason about their capacity to deliver services within the constraints place upon them by their offered SLAs, and on service clients to determine acceptable QoS levels based on their requirements. Although much research has been completed in the field on non-functional analysis, this reasoning still lies beyond the capabilities of most software development organisation. My work has focussed on delivering this reasoning through tool integration.

We have chosen the Unified Modelling Language (UML), and by extension the emerging Model-Driven Architecture (MDA) as the technical and theoretical foundation for my work. Reasoning in this domain generally requires modelling, and the UML is the most widely applied and supported modelling language for software architectures.

The work is broadly divided into three themes.

1. The automated derivation of analysis models from designs. The paper `A Model-Driven Approach to the Non-Functional Analysis of Software Architectures' by Skene and Emmerich describes an approach for encoding transformations from design models to formal mathematical models using logical constraints.
2. The definition of the semantics for SLAs and the integration of SLA information into designs. The paper `Precise Service Level Agreements' by Skene, Lamanna and Emmerich, defines the semantics of SLAs using a constraints over a reference model of service behaviour. The reference model also serves as the semantic basis for design annotations.

3. The development of a coherent approach to managing the semantics of designs which include multiple annotations, for example pertaining to performance or service level agreements. The paper `Modelling Electronic Services Using UML' by Skene, Piccinelli and Stearns, discusses the use of a reference model to define the semantics of design annotation pertaining to electronic services.

Cambridge

The main task in this phase of the project has been to identify the performance limitations experienced by Internet users and the available techniques for enhancing performance as well as the timescales in which these techniques operate. In this context, Internet users are enterprises, which interact with remote sites scattered around the Internet and they are primarily interested either in receiving or transmitting content to these sites.

We have described a framework for endpoint-control techniques, which was presented at the Bologna project meeting in September 2003. In particular we have identified Provider Selection, Route Selection and Transmission Control Behaviour selection as performance enhancement techniques which operate across the entire range of meaningful timescales and which are orthogonal to traditional performance enhancement techniques such as caching and load-balancing.

Our contribution has been to areas of Transmission Control Behaviour and Provider Selection. In particular we developed the idea of Transport Options (discussed in detail in [2]), which was presented at the RIPQoS ACM workshop in August 2003 in Karlsruhe, Germany. It has also been accepted for presentation at the research track of the NANOG Meeting (North America Network Operators Group) in October 2003 in Chicago. Our Transport Options work is a direct contribution to WP2; it involves the design and deployment plan of a QoS-aware infrastructure for application hosting (here QoS refers to relative service differentiation with statistical performance guarantees). It involves the specification of mechanisms for QoS enhancement and negotiation, which will allow the component containers in the TAPAS architecture to become QoS-enabled. Transport-level SLAs is a relatively new approach and therefore we expect this work to provide input to WP1 which develops the notation for expressing QoS support in SLAs.

There has also been work in the area of Provider Selection where we have now finalised the basic model and described a methodology, which allows an Internet site (e.g. an ASP) to search for optimal multihoming solutions.

We consider the Provider Selection work as a contribution to WP1 as it involves the specification of the performance characteristics of Internet Service Providers and the development of analysis techniques, which allow a multihomed ASP to make optimal network provider selection.

Over the summer, we have also collaborated with our colleagues at the University of Newcastle in the initial phase of their work on QoS adaptive multi-party communication scheme (part of WP3) – transferring our research and operational experiences with IP multicast related protocols and technologies.

Overview

The main theme in this phase of the project has been to identify the performance limitations and the available techniques for improving the performance experienced by Internet users as well as the timescales in which these techniques operate.

We distinguish between users which are primarily *sources* of traffic (ASPs, content providers, web farms which are mainly interested in the traffic they transmit) – these users are said to be of *server-type*, and those other users which are interested in the traffic they *receive* from sites of interest (remote/target endpoints) (said to be of *client-type*). In both cases we assume that the user (endpoint) is a stub network, i.e. traffic either originates or terminates at it (*no transit traffic*). The techniques are different in each case.

We make the pragmatic assumption of a best effort Internet *without QoS mechanisms in the routers*. Following that we make two basic observations

1. The performance to different parts of the network does vary significantly depending on the provider networks between the communicating endpoints.
2. The performance of reliable transport depends heavily on the adaptive transmission control behaviour of the endpoints (aka congestion control).

From the first observation we realize the strategic importance of multihoming while the second one suggests an opportunity for introducing differentiated services or just enhancing end-to-end performance by simple modifications to the adaptive transmission control mechanisms which operating at the endpoints as part of their transport protocol.

We described a framework for *endpoint control techniques* and the associated *timescales* and presented the ideas in September 2003 in the Bologna TAPAS Project meeting. We have identified the following techniques

1. Multihoming and Provider Selection, operating in the months/weeks timescale,
2. Route Selection (assuming multihoming) operating in days/hours and
3. Transport Selection (which is orthogonal to multihoming) operating at the sec/msec timescale.

We believe that this classification will influence the design of next generation network Service Level Agreements (SLAs).

There are already commercial products that offer performance oriented route selection to multihomed, stub-networks (see for example www.sockeye.com, www.routescience.com]. Our main contribution is at points (1) and (3).

In particular, different ISPs offer SLAs with different pricing structures and delay / bandwidth commitments however there has been no prior-work known to the authors on provider selection methodologies that involves performance vs. cost tradeoffs. One of the reasons is that data on ISP tariffs are hard to obtain.

Provider Selection

In our model we assume that there is a number of providers active in the area. We focus on a particular user (aka *site* or *reference endpoint*), which is either of the *server* or the *client* type. Using historical data the user in question knows

1. the *remote endpoints* of interest (target endpoints) (AS numbers, network numbers),
2. the expected traffic *volume* to be sent (received) to (from) each of the remote endpoints over a period of time,
3. the *cost* incurred per Mbyte for communicating with each of the endpoints of interest and
4. the (average) *delay* to (from) each of these remote endpoints of interest.

Using this information the reference endpoint computes the total cost of its traffic and assigns a *weight* to each remote endpoint of interest according to its relative contribution to the total cost (or traffic volume or importance – according to business criteria). The optimisation objectives now become

i. Minimize total costs
   s.t. maintaining the current characteristics of the delay distribution

ii. Optimise the delay distribution (mean/variance/median or other metric)
   s.t maintaining (or reducing) the total cost of traffic

Normally there will be no single "best" provider for reaching all the remote endpoints of interest, with the lowest delay at the lowest price, so the endpoint will naturally seek a multihoming solution, which involves a certain number of bilateral SLAs with different providers. It is assumed that the fixed interconnection costs (cost of line cards, cost of setting up the link etc.) are negligible compared to traffic costs (this may not be true if we consider multihoming to a large number of providers).

We envisage the process to be as follows; the endpoint submits its requirements to the *candidate* providers stating a *bound* for a performance metric of choice (e.g. average delay or a combination of delay-based statistical metrics) as this applies to each of the remote endpoints of interest. Each candidate provider will respond with a *bid* for the contract specifying the tariff ($/Mbyte), the supported traffic volume and the performance offering. The parameters in the provider's bid for the contract may be specific to individual target endpoints – or not. They will almost certainly be target specific in the case of the performance metrics. In the most general case the offer for each remote endpoint will have a tariff, a volume and a performance parameter associated with each remote endpoint of interest; the traffic volume parameter simply states the maximum amount of data to (from) the remote endpoint for the duration of the contract and requires further bandwidth (rate) specification (e.g. peak, sustained etc).

We assume that the providers declare truthfully the performance metric (delay) to different parts of the network and that at the end of the contract the customer will be able to verify how reliable the provider is. Moreover because the interaction between the customer and the provider is not one-off and there is always the prospect of future business it is in the best interests of the provider to be truthful in the performance aspects of its bid in order to be considered by the customer for future contracts.

Assuming that the reference endpoint considers a certain number of providers for multihoming (a number which will normally be significantly smaller than the number of the remote endpoints of interest).

There may be a provider which offers minimum delay at a minimum cost for a certain number of remote sites but not to the rest of the sites, should this provider be selected to be part of the multihoming solution or not and on what criteria should such decision be based.

Obviously user *preferences* (attitudes towards the expected delay, delay variance and/or cost) will play an important role at this point.

Next we need to define an *objective function,* which involves quantifiable aspects of these preferences. Our approach involves

1. rating the different providers according to different factors (e.g. performance and cost offerings) for each target endpoint, using the weights of the target endpoints,
2. combine these ratings using appropriate weights for the different factors in order to derive a final rating

Approaches which involve searching through all possible combinations of providers can be time consuming and inefficient. The precise description of our approach is outside the scope of this document and more details will be published in a separate document.

Transport Options

In today's Internet the vast majority of the traffic is being transmitted using the standard TCP congestion control procedures [3]. There is no doubt that this one-size-fits-all approach to Internet transport with a "single" reliable transport protocol and a "single", "standard", adaptive transmission control behaviour can be limiting to performance in such diverse environments (ranging from thin wireless links to Gigabit networks) and traffic conditions. Clearly different transport protocols and/or adaptive transmission control behaviours should be used, but this is far from straightforward.

Standardisation, implementation by vendors, wide-spread deployment, backwards compatibility with older versions, and "airness ("TCP-friendliness") considerations are only some of the issues involved and they take years to resolve.

TCP itself changed at a steady pace since congestion control was first introduced in TCP Tahoe in 1988 and there has been a lot of research in transport protocols for streaming media applications, congestion control for aggregates. Currently TCP New Reno and TCP SACK are the most widely used versions of the protocol. Some more recent proposals like High Speed TCP allows the standard congestion control rules to be overridden once the congestion window grows past a certain threshold.

However, the communicating endpoints still have no way of selecting what is the appropriate transport, *leveraging information* they might have about the environment in which they operate.

We refer to the selection of transport protocol and/or adaptive transmission control behaviour within the transport protocol, simply as choice of "transport selection".

Granting transport selection to the endpoints has interesting game-theoretic implications because the endpoints are greedy and there are no guarantees that they will behave in a "socially responsible" manner.

However our hypothesis is that the stability of the network now depends more in the longterm provisioning decisions of network operators (ISPs) and less on the behaviour of the endpoints and that congestion occurs at inter-provider links while their internal networks are relatively well-provisioned.

We envisage a framework in which the "transport" is determined on-the-fly (late binding) based on *policy criteria* and *negotiation* between the endpoints.

The design choices here involve

1. transport protocol code that can be loaded dynamically if it is not already loaded (it may be already available at a host or it may need to be downloaded from the network i.e. mobile code) and executed, or
2. the transport protocol that is already built into the kernel could be appropriately parametrised at run-time.

In [2] we presented an architecture for addressing these issues, mainly (2) since it allows for rapid prototyping and deployment. We called our approach Transport Options because it allows the endpoints to select and negotiate a suitable mode of transport. Transport Options can be used in two ways:

1. for enabling high performance transfers in networks which are known to be well-provisioned and lightly utilised and
2. in a more commercial setting, for introducing service differentiation when they become part of a customer-provider Service Level Agreement,

The approach naturally raises a number of questions about the risk from increased transmission aggressiveness. However the proposed changes can be made to affect transmission behaviour only when the network is not congested and in the opposite case the transport will fall back to what is considered standard transmission behaviour.

In order to address this issue we intend to design an experiment where the same workload (i.e. amount of data consisting of a number of TCP flows) will be transferred over a link over a period of time so that the average utilisation is reasonably low. (e.g. 20-30%). Then we repeat the experiment with same workload but now the TCP connections will use a more aggressive transmission strategy. The goal of this experiment is to show that in well-provisioned environments it is possible increased aggressiveness allows faster completion times for the TCP flows without increasing loss rates.

Another issue is whether an ISP should care about how server clusters in their customers' sites or neighbouring ISPs operate or whether simply monitoring the traffic levels at the interconnection points is sufficient. This issue is critical for the success of the scheme because requiring the ISP to police or manage the transmission behaviours of its neighbours may turn out to be prohibitively expensive.

REFERENCES

[1] J. Crowcroft and P. Oechslin "Differentiated end-to-end Internet services using a weighted proportional fair sharing TCP" ACM Computer Communication Review 28(3), July 1998

[2] P. Gevros "Internet Service Differentiation Using Transport Options: the case for policy-aware congestion control" ACM Workshop on Revisiting IP Quality of Service, August 2003.

[3] M. Allman, V. Paxson, W. Stevens RFC 2581 "TCP Congestion Control", April 1999

PROJECT MANAGEMENT AND CO-ORDINATION:

During the period, plans for progressing all work packages were agreed by the Executive Board, with particular attention being given to deliverables due after 18 and 24 months:

- D3 - Method for Service Composition and Analysis
- D8 - Container for Group Communication
- D9 - Container fro Trusted Coordination
- D10 - Container for QoS Monitoring
- D13 - Second Year Evaluation and Assessment Report
- D17 - Updated Dissemination and Use Plan

Further meetings have been arranged as follows:

- Late Oct: Ncl/UniBo/UCL 5-day meeting: Monitoring, QoS enabled application server
- mid Nov. The same group 2-3 day visit to Adesso, Dortmund
- 9-10th Dec. Technical Meeting, Bologna, to discuss Auction demonstration implementation plan
- W/b 2nd Feb. IAB Meeting, UCL
- 16-27th Feb.04 Plenary Workshop sometime during this period

CONFERENCE MEETINGS:

14th Database and Expert Systems Applications (DEXA'03) attended by Nicola Mezzetti.

Test and Analysis of Component Based Systems, 13th April, overall European Joint Conferences on Theory and Practice of Software (ETAPS) in Warsaw 5 – 13 April ’03 attended by J Skene.

Future Trends of Distributed Computing Systems (FTDCS) May 28 30, 2003, Puerto Rico attended by D. Lamanna

4th IEEE International Workshop on Policies for Distributed Systems and Networks, Lake of Como, 4-6 June 2003 attended by A di Ferdinando

Middleware 2003 16-20 June 2003, Brazil attended by D Lamanna

Workshop on Revisiting IP QoS: Why do we care, what have we learned? (RIPQOS) Karlsruhe, Germany, August 27, 2003 attended by Jon Crowcroft and Panos Gevros

Bertinoro Summer School on Formal Methods for Software, Bertinoro, Italy, 22-24 September, 2003 attended by W. Emmerich

Workshop on Service Based Software Engineering in Pisa (FM2003-SBSE), 8th September with FME, Pisa, 8 – 14 Sept. 03 attended by J Skene

IEEE International Conference on Electronic Commerce, Newport Beach, CA, June 2003, attended by Carlos Molina

Seventeenth Annual IFIP WG 11.3 Working Conference on Data and Applications Security, Estes Park, Colorado, August 2003, attended by Santosh Shrivastava

PUBLICATIONS:

Amoroso, A. and Panzieri, F., "*A scalable architecture for responsive auction services over the Internet*", TR UBLCS-2003-09, Dept. of Computer Science, University of Bologna, June 2003.

Cook, N., Shrivastava, S. and Wheater, S. "*Middleware Support for Non-repudiable Transactional Information Sharing between Enterprises*", 4th IFIP International Conf. on Distributed Applications and Interoperable Systems, DAIS 03, November 2003, Paris

Crowcroft, J., Hand, S. Mortier, R., Roscoe, T., Warfield, A "*QoS`s Downfall: At the bottom, or not at all!*" In Proceedings of the ACM Workshop on Revisitng IP Quality of Service (RIPQoS), pp. 109-114, August 2003, Karlsruhe, Germany.

Di Ferdinando, A., McKee, P. and Amoroso, A.: *A Policy Based Approach for Automated Topology Management of Peer To Peer Networks and a Prototype Implementation*

Ezhilchelvan, P.D. and Shrivastava, S.K. "*Systematic Development of a Family of Fair Exchange Protocols*", Seventeenth Annual IFIP WG 11.3 Working Conference on Data and Applications Security, Estes Park, Colorado, August 2003.

Gevros, P. "*Internet Service Differentiation using Transport Options: the case for policy-aware congestion control*" In Proceedings of the ACM Workshop on Revisitng IP Quality of Service (RIPQoS), pp. 151-157, August 2003, Karlsruhe, Germany.

Kaveh, N and Emmerich, W. *Validating Distributed Object and Component Designs* in Formal Methods for Software Architecture, Springer Verlag, Lecture Notes in Computer Science, vol. 2804, 2003, pages 63-91, Edited by M. Bernardo and P. Inverardi KE03: Validating)

Kistijantoro, A.I., Morgan, G, Shrivastava, S.K. and Little, M.C. "*Component Replication in Distributed Systems: a Case study using Enterprise Java Beans*", 22nd IEEE/IFIP Symposium on Reliable Distrinbuted Systems (SRDS2003), Florence, October 2003, pp. 89-98, ISBN: 0-7695-1955-5.

Lamanna, D.D., Skene, J and Emmerich, W *SLAng: A Language for Service Level Agreements* In Proceeedings of the 9th IEEE Workshop on Future Trends in Distributed Computing Systems (LSE03: SLAng), 2003, (pages 100-106) IEEE Computer Society Press

Lodi, G. and Panzieri, F. "*JBoss vs. JOnAS*", TAPAS Project Internal Report, June 2003.

Mezzetti, N.: "*Towards a Model for Trust Relationships in Virtual Enterprises*" In Proceedings of 14th Database and Expert Systems Applications (DEXA'03) Workshop, 1 - 5 September 2003, Prague (Czech Republic).

Molina-Jimenez, C., Shrivastava, S.K., Solaiman, E. and Warne, J. "*Contract Representation for Run-time Monitoring and Enforcement*", IEEE Conference on Electronic Commerce (CEC'03), Newport Beach, CA, June 2003, pp. 103-110.

Rodrigues G N, Roberts, G, Emmerich, W and Skene, J *Reliability Support for the Model Driven Architecture* In Proceedings of the ICSE Workshop on Software Architecture for Dependable Systems 2003 (RRES03: Reliability Support), ICSE 2003.

Skene, J. and Emmerich, W. *Model Driven Performance Analysis of Enterprise Information Systems* in Electronic Notes in Theoretical Computer Science, April 2003, vol. 82, number 6 (SE03: Performance)

Skene, J and Emmerich, W *A Model Driven Architecture Approach to Analysis of Non-Functional Properties of Software Architectures* In Proceedings of the 18th IEEE Conference on Automated Software Engineering (SE03: AModel), October 2003, Montreal, Canada (pages 236-239), 2003. IEEE Computer Society Press

Skene, J., Piccinelli, G. and Stearns, M *Modelling Electronic Service Systems Using UML* in Workshop on Service Based Software Engineering, FM2003-SBSE, Pisa, Italy, 2003, September, "Technische Universitaet Muenchen", pages15—30, url: http://www.cs.ucl.ac.uk/staff/J.Skene/phd/sbse2.pdf (SPS03: Modelling)

Solaiman, E., Molina-Jimenez, C. and Shrivastava, S.K. "*Model Checking Correctness Properties of Electronic Contracts*", International Conference on Service Oriented Computing, Trento, November, 2003

„Die zweite ASP-Welle ist auf dem Weg“ (The second wave of ASP is on it's way) in „Computerwoche“, pp 34-35, IDG Business Verlag GmbH, 26/2003, München, Germany

"TAPAS macht Appetit auf ASP" (TAPAS whets one's appetite for ASP), eCommerce Magazin 06-07 /2003, IWT Magazin Verlags-GmbH, Vaterstetten, Germany

"An Answer to the JBoss vs. JOnAS Comparison", W. Beckmann, M. Koßmann, adesso AG, 30 June 2003

PROBLEMS AND (POSSIBLE) REMEDIAL ACTIONS:

Regrettably, in July this year, Adesso lost a member of staff, which reduced the effort being deployed at Adesso.