

Peer-to-peer Innovations for eBusiness and eWork: A Vision of Emerging Software Service Technologies.

Mícheál Ó FOGHLÚ, Shane DEMPSEY and Eamonn DE LEASTAR
Telecommunications Software & Systems Group
Waterford Institute of Technology
Cork Road, Waterford, IRELAND
Tel: +353 51 302963 Fax: +353 51 302901
Email: {mofoghlu, sdempsey, edeleastar}@tssg.org
Web: <http://www.tssg.org>

Abstract: This paper describes the emerging platform for computing based on a number of developments in converging technological streams. It argues that software developed for this new platform will increasingly be described, designed, constructed and deployed using a service-driven focus. One technology that may well have an interesting impact is that of peer-to-peer (p2p) networking, in contrast with traditional client/server architectures. The most likely application/service domain for such new services will be in the area of on-line presence management, information exchange, and groupware. Some new products and services in this domain are discussed in detail. The analysis of the emerging p2p technology domain is based on current work in designing and developing SIP and SIMPLE services for 3G wireless networks and daily exposure to p2p collaboration services.

1. Introduction

This paper argues that a new paradigm in software systems to support teamwork is evolving. Email has been the dominant group working application since its inception in the 1970s. For large corporations, more powerful messaging solutions have been built around the dominant platforms of the time (from DEC's Notes, through Lotus Notes/Domino and Microsoft's additional functionality in Exchange). Smaller companies on more limited budgets have built low-cost solutions utilising web-based interfaces, often re-using open source components. Some companies have carved out niche markets selling products to this market.

In the next five years these systems will be replaced with more flexible systems enabling dynamic group creation, peer-to-peer (p2p) communication and secure authentication. The first generation of these new tools has already arrived with Groove [1]. As new mobile devices supplant desktop PCs as the standard means of communicating electronic content, new tools like Groove will supplant older email-only solutions, expensive GroupWare and Messaging solutions, and the hybrid lightweight web-based solutions currently in place. These tools will be built upon open and accepted standards for secure p2p communication and data management. These standards are currently emerging as part of a new service-driven platform for development and deployment.

2. A Pervasive New Platform

A profound shift is taking place in the ICT industry. Several technologies, and two major industries, are converging to deliver a new, more pervasive, communications and computing platform. The drivers producing this shift are varied, including technological, social and economic factors. Four technological streams are converging to produce this next generation platform:

- Computing and Communications Platforms
- Telecommunications Software
- Applications Software
- Internet Technologies

The characteristics of this platform combine elements of these four streams, and may embody as yet unforeseen hybrid elements. However it will exhibit certain key capabilities:

- Mobile & Location Sensitive
- Internet-Aware
- Software accessed as Services within the network
- Full telephony support

The platform (summarized in Figure 1) is unlikely to be a single set of closely coupled technology—in the manner of previous platforms—but rather a technological mix incorporating APIs, frameworks, middleware, protocols and wireless infrastructure in highly distributed and dynamic configurations. The challenge is to identify the key features within this mix, and to understand the direction in which these will evolve.

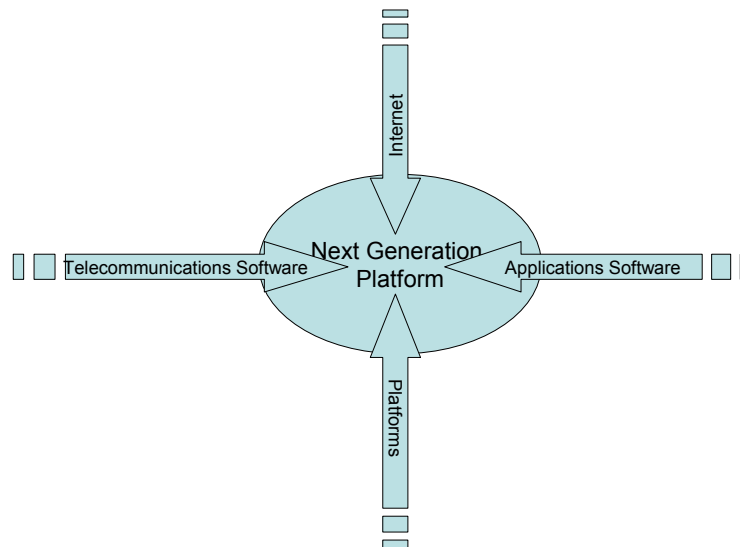


Figure 1 - The Emerging Technology Platform

One key area where new services will be developed is that of teamwork and groupware. Groupware facilitates the communication of ideas between people. These ideas and concepts can be conveyed in a loose or structured manner and with varying degrees of privacy and priority. Groupware includes messaging tools such as MSN Messenger [2] that provide all four of the capabilities listed above. The proliferation of these tools and the anytime, anywhere usage scenarios of mobile services, will strain traditional information delivery systems. New models will be required to meet these demands.

Key to the development of a new platform is the emerging role of software as a service. This influence comes both from the applications software development stream, which has evolved distributed systems and component architectures such as CORBA [3] and COM+ [4] through to the new 'web services' [5] concept. Coupled with this is the strong history of service architectures in the traditional telecommunications world. These are now being

re-engineered to fit with such services being delivered over IP. Witness the use of the Session Initiation Protocol (SIP) [6] as a core part of 3G mobile network service architectures [7]. It is important to recognise the importance of protocols in the evolution of the Internet. The Internet has evolved based on a set of open protocols, defined in Request For Comments' (RFCs) and facilitating a range of communication scenarios e.g. Email, web browsing. These RFCs provide individual strands of Internet DNA or pieces in the larger communications puzzle shown in Figure 2.

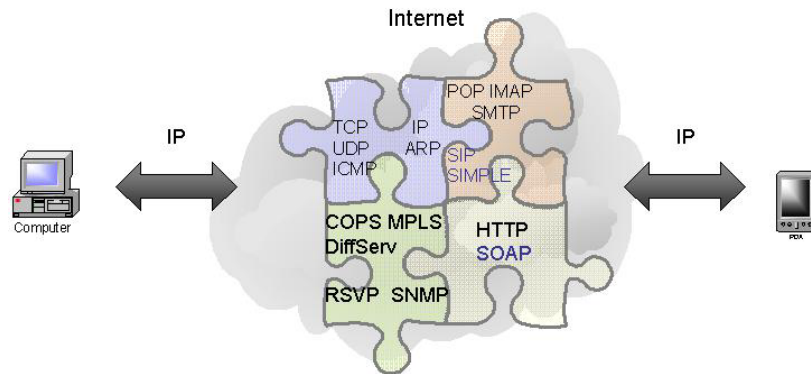


Figure 2 - Internet Protocol Puzzle

Together with commonly used Internet protocols Figure 2 shows SIP and its Instant Messaging (IM) counterpart, the Sip Multimedia and Presence Leveraging Extensions (SIMPLE) [8]. These messaging protocols have received widespread acceptance [9] among major software and telecommunications industry players (e.g. Microsoft, Motorola) and standards bodies (e.g. IETF, 3GPP). Through the work of the SIMPLE and Instant Messaging and Presence Protocol (IMPP) [10] working groups a globally accepted instant messaging protocol looks possible.




Figure 3 - Services Spectrum

Figure 3 attempts to show a complete service spectrum, taking existing telecommunications and Internet services and showing the requirements of elements within the distributed service network. This categorisation is subjective as each category is quite broad with many design and implementation possibilities, potentially affecting the spectrum. However, the diagram successfully illustrates how widely used applications will place very different demands on both terminals and access. An assessment of current Quality of Service (QoS) capabilities [11] promotes the view that a range of successful

person-to-person and peer-to-peer applications can be built using services requiring medium to low bandwidth and intelligence from network and terminal.

3. Groupware Technology Context

In delivering these services there is a **dynamic tension** between large, and sometimes proprietary, corporate solutions and lower cost open source web-based alternatives. This applies to many areas of Internet computing such as eLearning, Intranets, Knowledge-based systems and so on. This tension increases where mobile devices are used rather than traditional PCs.

However, it is in software encapsulating communications between ad-hoc groups where this dynamic tension is at its most strained (potentially its most productive), i.e. in Groupware. Jon Udell has pioneered both the analysis of groupware evolution and the implementation of cheap web-based utilities leveraging the power of open standards and lightweight scripting languages. His on-line report on Internet Groupware [12] and his book on the practicalities of Internet-based groupware [13] based on his own experience had a large influence on the way many industry watchers perceived the issue.

In the past, one of the authors of this paper has argued that the open source technology platform has matured to the stage where technologically-aware developers can utilise large amounts of basic open source components (e.g. operating systems, web servers, databases, scripting language etc.) to build powerful applications targeted at the eBusiness needs of SMEs [14]. It seems probable that one of the first areas where open source (or at the very least open standards-compliant) end-user applications may become available to SMEs is in the area of p2p communications. There is some precedence for this.

Email is essentially a p2p solution mediated by a network of servers. Most machines are configured to have mail delivered to a central server, but in principle every machine permanently connected to the Internet, or with a stable identity, could elect to manage its own email. Email is currently dominated by three open standards; SMTP, IMAP and POP3. SMTP controls delivery to the server. IMAP allows a client to access a server that stores all email folders centrally; POP3 requires users to download and manage their messages locally. As the number of email users grows, the business community is beginning to discover the problems with this form of communication:

- Email is useful for person-to-person or small coherent group communication;
- Large email lists become unmanageable;
- Centralised storage and management of email messages for individuals and groups tends to be done in an ad-hoc manner. This makes it difficult for users to search old content or establish context for discussions in newly joined lists;
- Email is inherently inefficient as much content (e.g. attachments) is duplicated.

The original Internet solution to this problem was Usenet News, now accessible though web interfaces such as Google's [15]. This was a highly distributed common discussion group system that worked across servers that had intermittent Internet access (much like email itself). Every message could be directly referenced, allowing sophisticated display of threaded views. In practice, most discussion groups are now implemented using stand-alone web servers so the concept of a global discussion community is gone (except for some specialist areas which are still quite active), though the Usenet News network service still operates, new Internet users almost never use it.

Within corporations, solutions such as DEC Notes, Lotus Notes, Novell's Groupwise and Microsoft's Exchange have evolved from email systems towards messaging systems and groupware systems with much more functionality than the simple email list. Here content can be structured into thematic areas and made available to appropriate groups of users.

Some of these systems also include high levels of security (necessary for commercially sensitive material). Lotus Notes has come to dominate this market but groupware has not become part of the standard desktop environment for all Internet users in same way as a web browser or an email client.

Therefore, if you currently want to reach a large audience the best way is to utilise the web browser as a client for some form of web discussion forum. There are a myriad of such systems available now, from free open source options such as PHP-Nuke [16] to web front ends for the big corporate solutions such as Lotus Domino. Facilities such as this are often built-in to other larger tools such as Learning Management Systems (LMS) like WebCT.

Instant Messaging is in its infancy but its popularity is growing rapidly [17]. It offers dynamic presence, you can see whether your contacts are on-line, and they can indicate their state (e.g. busy, happy to chat). IM would appear useful for work or recreation but many companies are unsure about the benefits of this technology. IM services have been blocked on many corporate firewalls to prevent workers wasting productive time in non-work related IM activity. Other companies have seized this mode of communication as being core to the evolving information systems environment in industry. Popular IM solutions include ICQ [18], AIM [19], Microsoft's MSN Messenger [2] and Yahoo! IM [20]. There is a battle over who can control the largest subscription bases to such services (most of these big players claim over 100million registered users). As one illustration of IM's growing popularity, it has been reported that "Time spent chatting over the Internet via IM also rose sharply, increasing 110 percent to 4.9 billion minutes at work—and up 48 percent to 13.6 billion minutes at home—mostly using the three major IM providers: AOL (NYSE: AOL), Microsoft (Nasdaq: MSFT) and Yahoo! (Nasdaq: YHOO)." [17]

There has been a great buzz in the open source world around the Jabber [21] project that has IM wrappers available for contacting most of the popular IM services (there are ongoing debates about the legality of reverse engineering the APIs for these services) thus undermining the proprietary nature of the subscription wars. Jabber's XML-based message exchange mechanism could also be used to build any number of services (similar to the web services paradigm). One extensible and flexible messaging and presence protocol is probably enough, however, and the aforementioned SIMPLE could unite the IM world.

3.1. Practical P2P Examples

P2P isn't just about **messaging**; it is also about **file sharing** and **collaboration/groupware**. The most famous p2p file-sharing service is Napster [22] that provided peer-to-peer file sharing capabilities across the Internet, mediated by a central server bank, and was primarily used for sharing digital music in MP3 format. The music industry shut down these servers after protracted legal battles, and Napster's new commercial service has yet to emerge from the legal wrangling. However, the p2p file-sharing gauntlet has been taken up by a number of challengers, primarily KaZaA [23] (in fact the two portals Morpheus and KaZaA both shared the same FastTrack network, but have recently fallen out) and Gnutella [24]. The growing popularity of file-sharing clients and services has sparked an interest in the protocols and mechanisms used to search and download files. Taking Gnutella as an example, the degree to which the search space is distributed and the mechanism used to search and retrieve files can adversely affect performance [25]. This is not confined to file-sharing activities, as all operations based on one-to-many communications should be optimised to reduce the risk of exponentially generating messages with every hop in the chain. Clustering of p2p resources [26] based

on personal or corporate relationships may provide the solutions to more efficient groupware services.

Perhaps the most promising really usable p2p application that crosses into the commercial world is Groove [1]. A team led by Ray Ozzie, (ex Lotus Notes developer), has recently released version 2 of this product which combines p2p features such as IM with some centrally managed security features to create a flexible groupware product promising open standards interfaces and extensibility. Groove exposes SOAP interfaces so its internal components can be accessed and extended as cross-platform web services.

Groove facilitates both off-line and on-line activity. Each time a client is connected to the Internet, it re-synchronises itself with the latest activity using peer-to-peer communication with close peers. All Groove network activity takes place over what is effectively a VPN (i.e. secure private tunnels over the Internet using a combination of public and private key cryptography). The company is working on a two-pronged approach to push Groove out from standard PC clients (the current client software requires a subset of Microsoft's COM+ component architecture) to the emerging pervasive computing/ubiquitous computing platforms:

- Development of Groove Edge Services that use a web services architecture (i.e. SOAP) to provide mobile access to a local access point.
- Design of the core Groove architecture so that it can be ported to alternative platforms (a promise yet to be proven).

3.2. P2P is SIMPLE JuXTApositions

Key to the success of any new technology is a set of software tools and libraries that enable developers to create applications using the concepts, protocols and APIs of that technology with the minimum of effort. P2P is already popular with the open source community as p2p is inherently community-based software. There are many examples of open source projects and open protocols such as Gnutella, Freenet and Jabber. O'Reilly Net's 'Open P2P' website [27] provides a useful resource for tracking new developments in open p2p technologies.

With JXTA, SUN Microsystems have taken a bold step in providing developers with a complete framework for p2p development. JXTA is not limited to java developers as SUN have focussed on the protocols necessary to solve real-world p2p issues such as discovering groups of resources and establishing secure communications with them when desired. JXTA's XML-based protocols for peer authentication, discovery and usage are encapsulated within C++ and Java implementations available at the JXTA website. [28]. JXTA may not be the ideal solution to all p2p problems but it provides developers with valuable experience with complex p2p application scenarios.

Microsoft's support of SIP and SIMPLE protocols in its MSN Messenger tools is good news for proponents of open P2P. Using their Real-Time Communications (RTC) Client API [29] Windows developers can add SIP-based messaging and IP telephony capabilities to their own applications. As Jonathan Rosen, the primary architect of the SIMPLE protocol, points out [30] IM requires a protocol that can: (i) Identify users independent of their location; (ii) Route messages between those users; (iii) Scale from 2 to 2 million users and beyond.

SIP was designed to meet these criteria for Internet Telephony. SIMPLE adds a MESSAGE request to send messages and presence event subscription via SIP's SUBSCRIBE/NOTIFY messages. XML messages can be sent in a body of a NOTIFY message to provide presence information about a 'watched' resource. This is covered in [31].

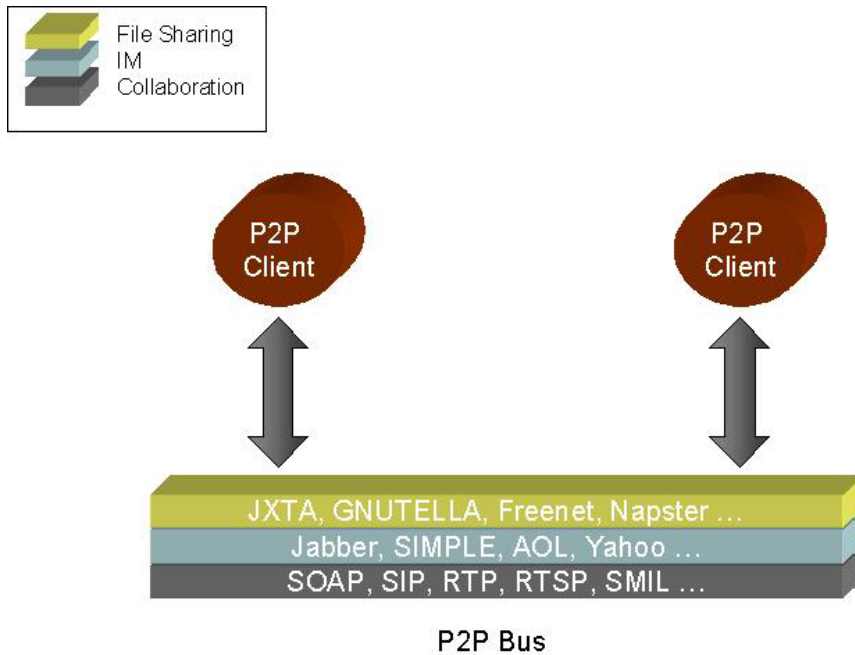


Figure 4 - P2P Communications Bus

Figure 4 illustrates the state-of-the-art in P2P communications, in particular the protocols that facilitate 3 broad categories of group interaction. At every layer there is at least one open alternative to the proprietary protocols of the large P2P providers. For instance, collaboration services may build upon open standardised protocols for multimedia session control and streaming (SIP, RTP, RTSP). Widely accepted protocols such as SOAP provide access to collaboration web services. Synchronous Multimedia Integration Language (SMIL) [32] is an XML-based presentation language that defines the spatial and temporal layout of media data in a presentation. SMIL enables p2p users to dynamically create and share media presentations, even over 3G wireless networks [33].

4. Conclusions

The area of groupware crosses into a number of key emerging technology areas: pervasive/ubiquitous computing on mobile devices, knowledge and information management (capturing the information lost in individual emails by imposing a more structured form), web services as a development and deployment methodology for thin client services, XML as a mark-up language for content and as an exchange protocol for web services, wireless networks linking mobile devices. This paper has outlined the technologies converging to influence the emergence of a new technology platform. This platform will be a combination of elements from the areas outlined here, with input from the telecommunications world where the concept of a service is key. Software development will become the development of services for deployment on this new platform, rather than the development of applications for deployment on Windows PCs.

It is clear that p2p concepts have a role to play in these emerging services. Standardisation of protocols such as SIMPLE is an important step. Whilst it may be unrealistic to expect large bandwidth for heavy p2p communication, some form of ad hoc networking, and direct communication, within a managed service architecture may well become the norm.

The TSSG is involved in a number of EU IST projects applying this technology set to practical problems, in particular the development of 3G mobile services: Opium [34], AlbatrOSS [35].

References

- [1] Groove Networks Inc, Platform for Inter-enterprise communication and collaboration, <http://www.groove.net> 2002 (Viewed 2002-04-29).
- [2] .NET Messenger Service, <http://messenger.msn.com/>, 2002 (Viewed 2002-04-30).
- [3] The Common Object Request Broker: Architecture and Specification Revision 2.5, Object Management Group, <http://cgi.omg.org/docs/formal/01-09-34.pdf> Sep 2002 (Viewed 2002-06-22)
- [4] COM+ Papers, Presentations, Media Coverage and Resources for COM+, Microsoft Corp., <http://www.microsoft.com/com/tech/COMPlus.asp> 2002 (Viewed 2002-06-22)
- [5] W3C Web Services, W3C, <http://www.w3.org/2002/ws/> 2002 (Viewed 2002-06-22)
- [6] Session Initiation Protocol - IETF RFC 2543, Schulzrinne H [et al.], 1999 <http://www.ietf.org/rfc/rfc2543.txt?number=2543>
- [7] 3GPP TS 24.229: IP Multimedia Call Control Protocol based on SIP and SDP; stage 3, 3rd Generation Partnership Programme (3GPP), http://www.3gpp.org/ftp/Specs/archive/24_series/24.229/ 2002, (Viewed 2002-06-22)
- [8] SIP Instant Messaging and Presence Leveraging Extensions (SIMPLE) charter, IETF, <http://www.ietf.org/html.charters/simple-charter.html> 2002 (Viewed 2002-06-22).
- [9] Message Vendors Rally Around SIMPLE Protocol, Network World Fusion, http://www.nwfusion.com/archive/2001/123534_08-06-2001.html 06-08-2001, (Viewed 2002-06-22)
- [10] Instant Messaging and Presence Protocol Charter (IMPP), IETF, <http://www.ietf.org/html.charters/impp-charter.html> 2002 (Visited 2002-
- [11] NEC GlobalNet – Whitepaper on IP QoS Control, NEC, <http://www.nec-globalnet.com/qos/> 2002 (Viewed 2002-06-22)
- [12] Jon Udell, <http://udell.roninhouse.com> 2002 (Visited 2002-04-29).
- [13] Jon Udell, *Practical Internet Groupware*. ISBN-1565925378 O'Reilly & Associates, Sebastopol CA, 1999.
- [14] Micheál Ó Foghlú, ConnectKey.com: The Use of Open Source Technologies to Provide Cost Effective Services for eBusiness to SMEs. In: E-Work and E-Commerce Vol 2 Brian Stanford-Smith and Enrica Chiozza (Eds.), ISBN 1 58603 205 4. IOS Press, Amsterdam, 2001.
- [15] Google News, <http://news.google.com> 2002 (Visited 2002-04-30).
- [16] PHP-Nuke, The Future of the Web, <http://www.phpnuke.org> 2002 (Visited 2002-04-30).
- [17] Jay Lyman, Report: Instant Messenger Use Exploding, E-Commerce Times, <http://www.ecommercetimes.com/perl/story/14793.html> Nov 2001 (Visited 2002-04-30).
- [18] ICQ.com, <http://www.icq.com> 2002 (Visited 2002-04-30).
- [19] AOL Instant Messenger, <http://www.aim.com/index.adp> 2002 (Visited 2002-04-30).
- [20] Yahoo! Messenger, <http://messenger.yahoo.com/> 2002 (Visited 2002-04-30).
- [21] Jabber Software Foundation, <http://jabber.org> & <http://www.jabber.com> 2002 (Visited 2002-04-30).
- [22] Napster, <http://www.napster.com/> 2002 (Viewed 2002-04-30).
- [23] KaZaA, <http://www.kazaa.com/> 2002 (Viewed 2002-04-30).
- [24] Gnutella.com, <http://www.gnutella.com/> 2002 (Viewed 2002-04-30).
- [25] The Popularity of Gnutella and its Implications for Scalability, Sripanidkulchai K., Carnegie Mellon University, <http://www-2.cs.cmu.edu/~kunwadee/research/p2p/gnutella.html> 2001 (Viewed 2002-06-22)
- [26] Early Measurement of a Cluster-based Architecture for P2P System, Krishnamurthy B [et al.] ACM Sigcomm Internet Measurement Workshop 2001, <http://www.icir.org/vern/imw-2001/imw2001-papers/29.pdf> 2001 (Viewed 2002-06-23).
- [27] P2P Development, Open Source Development, O'Reilly Net, <http://www.openp2p.com> 2002, (Viewed 2002-06-23)
- [28] Jxta.org, Project JXTA, <http://www.jxta.org> 2002, (Viewed 2002-06-23)
- [29] Real Time Communication (RTC) Client API, Microsoft Corp, http://msdn.microsoft.com/library/default.asp?url=/library/en-us/rtclnt/rtcover_91ih.asp 2002, (Viewed 2002-06-23)
- [30] SIP Presence and Instant Messaging,, Rosen J., O' Reilly Open P2P developer conference, 2001.
- [31] A Session Initiation Protocol (SIP) Event Template-Package for Watcher Information (Internet Draft), Rosenberg J., <http://www.ietf.org/internet-drafts/draft-ietf-simple-winfo-package-02.txt> May 2002
- [32] Synchronous Multimedia Integration Language (SMIL) 1.0 Specification, Ayars J. [et al.], W3C, <http://www.w3.org/TR/smil20/> 08/2001, (Viewed 2002-06-26)
- [33] Streaming Technology in 3G Mobile Communications Systems, Elson I [et al.], IEEE Computer Magazine, 9/2001, pgs. 46-52
- [34] Opium IST Project <http://www.ist-opium.org> (Viewed 2002-06-23)
- [35] Albatross IST Project <http://www.ist-albatross.org> (Viewed 2002-06-23)