



white paper:
a new platform for social communication

Version: 1.2
Author: Greg Rolan
Date: February, 2008



TABLE OF CONTENTS

| | |
|--|-----------|
| A NEW PLATFORM FOR SOCIAL COMMUNICATION | 1 |
| THE SOCIAL COMMUNICATION LANDSCAPE TODAY | 3 |
| the proliferation of personal identities | 4 |
| the requirement for privacy | 5 |
| islands of information | 8 |
| technological change | 9 |
| changes in social attitudes | 11 |
| REQUIREMENTS OF A NEW APPROACH TO SOCIAL COMMUNICATION | 12 |
| integration | 12 |
| privacy | 13 |
| identity management | 13 |
| call control | 14 |
| peer-to-peer directory architecture | 15 |
| GLYNX: A NEW PLATFORM FOR DIRECTORY SERVICES AND COMMUNICATIONS | 16 |
| introducing the glynx p2p overlay network | 16 |
| glynx information storage and retrieval | 17 |
| special qualities of the glynx 'black pages' | 19 |
| directory publishing and retrieval | 19 |
| glynx identity management | 22 |
| glynx black pages example information flow | 22 |
| Introducing glynx and communications | 24 |
| a-party call control | 26 |
| ubiquity and the glynx communications configuration hypercube | 27 |
| An example call control flow in glynx | 28 |
| CONCLUSION | 30 |
| REFERENCES | 32 |



The proliferation of communications technologies in recent years, coupled with limitations of current internet and telecommunications architectures, has led to a crisis in online identity and information management. In particular, the explosive growth in web-based social networking and other online collaboration and communication platforms has highlighted the weaknesses inherent in these architectures. It puts calling parties in control of their social communications rather than other intermediaries such as directories, networks or websites

A new peer-to-peer architecture promises to solve the problems of today's online identity world.

This paper discusses a new peer-to-peer architecture for social communication (telecommunications and online social collaboration) services. This platform delivers absolute privacy and trust of online relationships; user-centric discovery and control of online identities and information; and a unified mechanism for simplifying the burgeoning online communications world.

the social communication landscape today

The term Social Communication is used here to encompass those technologies and practices employed to establish and maintain relationships via some electronic medium - and to communicate with those with whom we have relationships. It includes traditional technologies such as telephony and electronic address books, but also embraces mechanisms made possible via the ubiquity of the internet E-Mail, Instant messaging, web-site mediated communications etc. Social Communication also includes activities which could be termed "social networking" whether taking place within the confines of services such as FaceBook or LinkedIn etc., or more informally as electronically augmented human activity.

Communications and social networking industries suffer from a chronic inability to adequately serve customer needs in terms of privacy security, and end-user functionality.

The Social Communication industry suffers from a chronic inability to adequately serve customer needs in terms of privacy and security. One barometer of this is the press which delivers us a constant stream of revelations - often either disclosure by service providers that their customer's data has been compromised; or reports of subscribers' ire over changes to terms-of-service which permit further observation, exploitation and/or interference by service providers.

Similarly, there is an enormous unfulfilled customer requirement in even the seemingly simplest telecommunications transactions such as calling someone. Subscribers have no way of unifying all of their communications options across their devices and services. It's now one hundred and thirty years since Bell invented the telephone, and despite ubiquitous broadband connectivity and intelligent devices, we are still playing telephone-tag.

This section describes the problems with existing approaches to Social Communication with respect to:

- the proliferation of personal identities;
- the requirement for privacy;
- the isolationist nature of service provision;
- changes in technology; and
- changes in social attitudes to service provision.

the proliferation of personal identities

It is becoming apparent that, rather than moving toward an oft-touted one-identifier-per-person paradigm (e.g. individuals having a global communications id such as a mobile phone number), personal identities are actually proliferating. As Brough Turner describes it:

“...The advent of the Internet has enabled many new communications services — email, instant messaging, blogging, social networking — and, so far, each new service comes with its own, new, centrally managed, addressing scheme...[It’s] now clear we’re moving to rich identities with dozens of context-dependent identifiers per person...” [1]

Rather than moving toward an oft-touted one-identifier-per-person paradigm personal identities are actually proliferating.

These identifiers include both communications and collaboration identities - identifiers associated with communications channels - (e.g. phone numbers, e-mail addresses, IM handles etc.), identifiers associated with collaboration channels (e.g. social networking handles, internet forum identifiers, etc.) together with identifiers not directly associated with communications addressing (e.g. government IDs, web-site login IDs etc.). Identities are used across a wide spectrum of social communication contexts: work/business; acquaintances; web-sites; friends and family; recreation; federal; state and local government; formal and informal organisations; electronic commerce etc.

Each of these contexts requires disclosure of zero or more identities and the task of administering and processing these context-sensitive identity sets is becoming unmanageable. In addition, the platform underpinning each identity provides information referring to the capabilities of, and disposition towards communication – explicitly or implicitly understood as ‘presence’.

Service providers, thus far, have taken one of two approaches to managing identity proliferation. The first approach is to ignore it, hoping that their service, being so compelling, will surmount all others and that subscribers will need no other identifier than one of theirs. The other approach is to allocate a new (“another”) identifier to the



An individual's identities and communications channels are an expression of self, communicating a 'personal brand'.

People typically have multiple personas which may contain mutually exclusive identity information.

Service providers either ignore identity complexity, or attempt to impose a single 'yet-another-identifier' as a solution,

subscriber and provide some sort of identity mapping or mediation service – for example to allocate a 'global' telephone number from which calls will be redirected to one or more other numbers.

Both of these approaches fail to recognise that an individual's identities and communications channels are an expression of self, communicating a 'personal brand'. Furthermore, individuals typically have multiple personas which may never contain the same identity information. People normally have several mutually exclusive personal brands ('personas') which they seek to maintain separately. No single identifier will suffice as a universal 'personal brand' for all contexts – no matter how compelling the service that is offered along with the identifier. Yet-another-identifier approaches simply add to the problem.

And, in particular, neither of these approaches leads to any insight or facility which would enable individuals to manage their sets of context-sensitive identities in a coherent manner.

Identity proliferation or rather, the dearth of any tool with which to manage user identities across multiple usage contexts, is currently a topic of much discussion. For example, Tim O'Reilly lamenting the lack of Address Book 2.0 [2] which would address exactly these issues recently co-chaired a conference at which these issues were debated [3]. And despite all of the talk, the conference presentations and protestations to the contrary, no incumbent service provider satisfactorily addresses this complexity.

the requirement for privacy

When engaging with others, individuals often seek to limit the disclosure of personal information; such limits based on the context of communication, the parties to the communication, and the impact of violation of privacy. The degree of privacy desired for personal information and, in particular, identities can be considered on a spectrum from totally public (e.g. a name in government registries) to totally private (e.g. private keys held on personal devices in a PKI system).

The requirement for privacy in social communication in a specific instance lies on a spectrum between these two extremes. Social communication services today generally require individuals to delegate custodianship of their identity and relationship information to the service provider – usually to be held in a centralised directory with associated authentication and access controls.

By doing so, absolute privacy of information is traded for utility in communication or dissemination. Generally, services allow subscribers to manage their information (e.g. social networking web



The privacy of social communication subscribers is being compromised either deliberately or inadvertently

The discovery, setup & maintenance of social communication relationships together with the exchange of personal information must be able to take place privately – without observation, interference or exploitation by any third party

sites holding subscribers' details and those of others), while the service providers retain visibility of all information and relationships.

This may be exactly what is intended – such as with Yellow Pages listings or Social Networking public profiles where the objective is to publicly disseminate information to as wide an audience as possible.

In other social communication applications, the privacy of social communication subscribers is being compromised either deliberately or inadvertently seemingly with increasing frequency. And, as importantly the marginal cost of complexity in managing such profiles is still significant, so such compromises are costly – subscribers cannot simply receive and publish a new set of identity claims

Service providers may deliberately compromise the privacy of their subscribers in various ways;

- Via exploitation of information or relationships inherent in the Business model of the service provider – for example targeted advertising;
- Via data aggregation and syndication which allows correlation of personal information for purposes beyond the original scope of disclosure; and
- By simply selling personal information to some other party – often to be used for unsolicited communication.

Of course there is a major problem with the White or Yellow types of public directories. In traditional communications networks, the privacy of directory information is not particularly problematic. This is because 'brute force' searching - i.e. contacting large numbers of subscribers to find an individual subscriber or for mass unsolicited communications- is normally expensive, for example because of the cost of making many calls can be significant.

With the communication mechanisms such as e-mail, Instant Messaging (IM) and Voice over Internet Protocol (VoIP) telephony there is little or no cost to transmit information to the destination resulting in the proliferation of unsolicited communications (SPAM, SPIM and SPIT). The SPAM (et al) problem has prevented the publishing of public e-mail address directories as these are too readily harvestable by spammers. Similar problems exist in relation to IM and VoIP directories as they rely on communication between IP addresses. Publishing IP addresses has the potential not only to lead to unwanted communication but also to other forms of attack on specific IP addresses.

The effect of these threats is that service subscribers concerned with unsolicited communication or identity fraud, where possible, tend not



to include complete personal and/or communications information in their public directory listings such as. Instant Messaging directories.

There are also inadvertent violations of trust, made possible by holding all information centrally:

- The inappropriate exploitation of entrusted information by unauthorised employees; and
- Compromising of entrusted information by hackers, government agencies and other external parties.

More fundamentally, however, there is nothing inherent in a centralised architecture which prevents such privacy violations. A service provider using a centralised architecture can only make assurances regarding the fidelity of its business practices, employees and infrastructure security – there is nothing architectural which guarantees it.

Storing identity information centrally presents a systemic security risk.

There are whole demographics which cannot be satisfied using a centralised approach

Storing identity information centrally presents a systemic security risk in that if the central database (or communications with a central database) is compromised (either through a technical attack, legal method, personnel misuse or other means) all contact information may become accessible without the knowledge or consent of the subscriber, contact or possible central agency. Also by observing the information flow occurring over the Internet to and from the server an observer may be able to deduce private information about subscribers and contacts without directly compromising the server or central agency.

At best, such services can only satisfy that portion of its target demographic for which the impact of such violations is not great.

At worst there are whole demographics which cannot be satisfied using a centralised approach: those for whom personal privacy is paramount (e.g. those in government, in entertainment; the wealthy and/or famous, and those needing to act clandestinely etc.); those who have private information entrusted to them and cannot disclose it to any third party without consent (e.g. those in the legal and medical professions etc.); and those for whom privacy is a fundamental right, the violation of which is a social concern

As social communication becomes more ubiquitous, dissatisfaction in the privacy aspects of service provision is growing - both with subscribers who's privacy is continually being eroded (for example, the recent outcry among Face Book users over the introduction of a 'feature' which was, in fact, a relationship data-mining and advertising exploitation) [4]; and those for whom it is apparent that their needs are not being met in any way by service providers (e.g. the boycott of the UK medical records system by British doctors due to the blatant



interlinking of it's data with other systems with far-reaching but poorly considered implications);

The discovery, setup & maintenance of social communication relationships together with the exchange of personal information must be able to take place privately – without observation, interference or exploitation by any third party.

islands of information

Today, the social communication setting resembles 'islands' of information floating in the 'seas' of the Internet and corporate or private networks.

For example: your Face Book profile and friends, your Linked In social network; your Outlook contacts; your Skype ID, buddy list and presence states; your MSN ID, buddy list and presence states; your Mobile phone details and network state (e.g. in-call/off-network etc.); the number of your geographically nearest fixed phone; your geographical location suggested by network address, mobile phone cell, wireless beacon, GPS etc. to name but a few.

These islands can be considered provider-centric directories inasmuch as they contain information useful to subscriber and operator alike, but are optimised and constrained to facilitate a particular service provider's business.

At best, service providers only share information which furthers their own business objectives. At worst they prevent subscribers or anyone else from freely using useful subscriber-related information (e.g. Telco's not disclosing mobile phone state – in-call/off-network etc.) forcing others to pay for access. through additional call connections and voicemail .

More than this, any value to subscribers to be gained from interconnecting these directories is almost never realised due to cannibalisation which limits a network's incentive to "play ball".

For example, the telecommunications, VoIP, E-Mail and Instance Messaging industries have formed some uneasy alliances - having largely agreed on interconnect of basic carriage services (in some cases) but have not addressed interconnect of directory services in any meaningful way.

Of course one industry initiative within scope of this objective is ENUM which provides:

- A directory for multiple communications identities associated with a subscriber

Today, social communication resembles 'islands' of information based upon provider-centric directories

Any value to subscribers to be gained from interconnecting these directories is almost never realised

Much of the latent innovation in identity management has been stifled by the constraints imposed by service providers



- A unified numbering plan (eg a new country code for VoIP numbers)
- Interconnect arrangements between disparate network (eg PSTN) to another (eg VoIP)

However, as described below, ENUM, 10 years on, still falls short of true directory interconnect.

Most telling though is that, as the marginal cost of carriage drops, and the barriers to churn are removed, service providers, looking to wring value from their directory assets, are not expected to willingly share this information with others. For example, a PSTN, mobile, or conventional VoIP carrier directory has no direct interest in informing a caller that the B-Party is available on a free VoIP service as well as their own network. The ENUM project designed to address these concerns from a carrier perspective has been underway for ten or so years, but telcos have been so concerned about cannibalisation that this initiative has not moved much beyond an engineering curiosity. And, more recently, Social Networks are still discussing how to deal with the “Open Social” initiative.

Service providers have custody of subscriber's information and relationships and use this position as the basis for their respective business models

Thus, service providers have custody of subscriber's information and relationships and use this position as the basis for their respective business models. They have limited interest in fostering either interoperability with other services, or unrestricted private interaction between subscribers. Their business models necessitate striving to be the only ‘island in the sea’, mediating subscriber relationships for exploitation.

Much of the latent innovation in identity management has been stifled by the constraints imposed by service providers: by restricting the directory information available to subscribers; by confining control of communications in the networks and not devolving it to subscriber-controlled devices; and by only largely addressing carriage interconnect but not directory services.

technological change

Possibly the biggest issue with current approaches concerns technological change.

Mobile devices today have the bandwidth and capacity to solve many of the issues outlined above. The marginal cost of managing complexity is declining rapidly – both in terms of the cost of bandwidth and that of computing power itself. A major potential of technological change is to unleash possibilities for a paradigm shift in the delivery of online services.



And yet, service providers and consumers are locked within a paradigm of distinct in-network service provision and edge-device consumption of services under network control. This architecture prevents substantial innovation in the provision of services which subscribers want.

Despite accelerating advances in personal computing and telecommunications, service providers would have us believe we are still back in the nineteen seventies with 'dumb' devices being serviced by highly efficient and leading edge centralised servers.

Service providers would have us believe we are still back in the nineteen seventies with 'dumb' devices being serviced by highly efficient and leading edge centralised servers

Consider a simple example from the telecommunications world. Why isn't an 'answering machine' standard functionality on a mobile phone? It is not a problem of circuit miniaturisation or of a phone's ability to handle audio stream reception or recording. It's that service providers only allow and promote phone features that stimulate revenue such as cameras, music players and radios (where the sending of MMS generates data traffic and possibly revenues) - rather than features such as answering machine that could optimise the full customer experience. And such voicemail is necessarily mediated through a set of centralised directory-based services - resulting in a net service cost to the subscriber.

However, the centralised architecture is a two-edged sword for service providers. On the one hand, as the marginal cost of carriage decreases, this paradigm dictates that the onus shifts to directory-based services to provide subscriber value. On the other, there are limits to subscriber functionality which centralised business models allow, as well as limits to the profitable scalability of centralised services (not every organisation can become a Google-size island).

Thus service providers find themselves with business models that do not scale as easily with carriage provision, and so they will not deliver against subscriber needs. At the same time, subscribers are in possession of computing power and bandwidth which they cannot use to satisfy these requirements under the delivered paradigm (or at least within the prevailing terms-of-service).

While service providers operating with centralised directories are the sole gatekeepers of social communication services, the true potential of consumer devices to satisfy subscriber requirements will not be realised through this model. An alternative approach, leveraging subscriber computing power and bandwidth would appear to be the way forward.



changes in social attitudes

Individuals are moving from being passive consumers of services and information toward being active participants in the global communications milieu. Fuelled by the internet, mobile and related technologies, social communication is transitioning from geographically based delivery to a global marketplace for services and dissemination of information.

There is a definite shifting of attitude with respect to service provider control of the subscriber experience

Social communication consumers are increasingly unwilling to accept structural and deliberately imposed inadequacies in privacy and functionality, and are demanding real solutions to their problems.

Social communication consumers are aware of the issues described above, and the capabilities of devices to solve them. In particular they are aware that their devices are increasingly capable of managing this complexity and so are demanding that their own devices manage these issues – rather than accept partial solutions offered by a patchwork of competing, uncooperative service providers.

Whereas, in the past, consumers had little choice but to accept what was offered to them, they are now well connected in the global marketplace of services. Philosophically, subscribers are increasingly unlikely to merely accept providers' assurances of the necessity of constraints. Dissemination of information via the internet is fast and widespread – consumers can become remarkably well-informed about service deficiencies and competitor differences with little effort. Consumers expect responsiveness as their right. Service providers ignore unfavourable blogs and forum postings at their peril.

The internet generation is more aware than ever before of the possibilities that technology can bring. Service subscribers expect 'cool' solutions to their social communication needs and are perplexed when technology itself becomes part of the problem. They are 'savvy' consumers and become irate when service providers ignore requests for new services or improvements to existing ones. They are able to quickly change providers if a clear differentiator emerges.

In the telecommunications industry, for example, there is a definite shifting of attitude with respect to service provider control of the subscriber experience. Mobile phones need no longer be necessarily 'locked' to any network (though operators still proffer such subscription plans) and consumers are quite happy to purchase phones outright and churn providers based on support or features. Device and platform suppliers – who are supplanting service providers as the source of innovation – are, in some cases dictating terms to operators.

The launch of the iPhone, the announcement of Google's Android platform, and Nokia's divestiture of their in-network equipment business in favour of an exclusive focus on handsets are all indicators



of a major shift if power from service providers to device manufacturers and consumers.

Thus, not only are there large unmet requirements for delivery of social communication services, subscribers are aware of these gaps in service delivery. In much the same way that they have embraced the PC revolution to incorporate computing at a personal level in their businesses and homes, social communication consumers choose and understand the technological capabilities of their devices, and are willing to take responsibility for managing their social communication environments.

Social communication consumers are increasingly unwilling to accept structural and deliberately imposed inadequacies in privacy and functionality, and are demanding real solutions to their problems.

requirements of a new approach to social communication

The setting is right for a new approach to social communication; one which breaks this nexus of the centralised directory while linking the islands of information, and enables the power of consumer devices to deliver true privacy and subscriber functionality in social communication.

Requirements comprise

- *Integration with existing systems and business*
- *True privacy*
- *Identity Management*
- *User-centric communications*
- *A peer-to-peer architecture*

This approach requires a new architecture based on control of social communication by the subscriber at the subscriber's device(s). Control may be delegated to centralised services where appropriate, but done so at the discretion of the subscriber – not as a requirement of the service provider.

True privacy can only be delivered within an environment which is under the end-users' direct control and communication can only be free of interference from service providers if is unmediated. In this way privacy and independence can be guaranteed at an architectural level.

This section describes the requirements of a new approach to social communications along four dimensions and concludes that a peer-to-peer approach must be adopted. The four dimensions are: integration, privacy, identity management and call control.

integration

It almost goes without saying that any new approach to Social Communication cannot exist in a vacuum. It cannot be a replacement for all which has gone before, and, despite being disruptive, must form a commercial basis for Social Communication.

- The private directory must be able to integrate with existing public directories and services – in effect to provide a private



channel for social communication alongside existing public channels.

- Any new approach must allow opportunity for commercialisation without compromising privacy or service independence.

privacy

The overarching requirement for privacy is that the discovery, setup and maintenance of social communication relationships together with the exchange of personal information must be able to take place privately – without observation, interference or exploitation by any third party.

- The subscriber must be able to choose which identifiers and other personal information should be publicly visible and which should be treated privately.
- The private information – the private portion of the directory – must be under the direct control of the subscriber. The only meaningful way to achieve this is for the private directory to be located on the subscriber's device(s). In this way, the subscriber is in full control of their information (and entrusted information) and can choose whether or not to delegate control or visibility of it at arbitrary levels of granularity.
- Personal information must be able to be aggregated at arbitrary levels of granularity and such aggregates must be able to be differentially disclosed to different audiences – ultimately resulting in the potential for custom views of personal information to multiple audiences-of-one.
- Any exchange of personal information or information regarding relationships must be private at subscriber's discretion.
- The directory must not be able to be harvested for candidate addresses to be used for unsolicited communication. It should be searchable (i.e. provides utility for locating target subscribers) but non-browsable (i.e. cannot be used to browse for entries about which there was no prior knowledge).

identity management

A major goal of any new approach is to allow expression of identity as a 'personal brand'.

- The approach must facilitate the management of many personal identifiers together with multiple aggregates at arbitrary levels of granularity which may be disclosed to or



used by audiences of arbitrary size. Ultimately, management should be able to be extended to audiences-of-one as every relationship is unique.

- The approach must not require additional identifiers. The directory service must be able to use existing identities – without compromising privacy or facilitating unsolicited communication
- The approach must shift control from issuers of identities to recipients. Recipients must be free to use identifiers in any way (subject to commercial terms) which enhances their social communication experience – irrespective of the business wishes of the issuer.
- Identity claims must be verifiable by identity providers and others.

call control

As discussed above, meaningful integration of service directories is unlikely to happen, therefore aggregation of communications identities and associated information must occur on the subscriber's device(s). Only a subscriber's device(s) have full access to communications options and state at any point in time – no other directory may be queried for this information.

- Control of communications (e.g. call set-up, channel etc. negotiation) must take place device-to-device and not via in-service switching.
- Communications must be able to handle non-predictable calling situations – i.e. in real time at the time of call set-up, not as a result of some pre-defined calling sequence.
- Control of communication must satisfy both the calling party and called party preferences – irrespective of the business preferences of the carriage providers.
- Communications must be optimised on a end-to-end basis – not just at the calling party end (c.f. a BluePhone device which chooses between Cell and VoIP carriage depending on the proximity of the calling device to services – irrespective of the capabilities or status of the called party. With BluePhone if, for example, the calling party is within range of a VoIP service, VoIP will be chosen – even if a Cellular call may be cheaper and a higher quality option from an end-to-end perspective).
- Where possible, deliberate interference by carriage providers should be avoided – for example, the practice of carriage



providers to terminate an un-answered call at an in-network voice-mail service.

- Call control should be easily integrated into other aspects of social communication – such as public/private directory services

peer-to-peer directory architecture

The requirement for the private directory residing on the subscriber's device(s) leads to a major shift in the architectural paradigm for directory services. In order for the public and private components of the directory to work together in a cohesive and deterministic manner, they must be linked via a decentralised, secure and private network - a peer to-peer (P2P) directory network with the following characteristics:

- The network must ensure the efficient, timely and deterministic storage and retrieval of data.
- The information retrieved from the directory must be verifiably correct.
- The directory must ensure consistency of data and response time despite participant subscriber nodes of varying computing capability and bandwidth of connection. It must also operate under any anticipated levels of churn – i.e. nodes leaving and entering the network.
- The directory must be able to store an arbitrarily large number of pieces of information via unique key/data mapping. This is particularly important as common algorithms employed by various implementations of P2P networks may map different data to identical keys.
- The directory must be resistant to attacks whereby a party is able to take over a large number of nodes in a predictable segment of the network which would undermine the ability to guarantee service.
- The network must be able to ensure connectivity where possible between nodes with interposed NAT and firewall devices. Of course firewalls which are configured to block such P2P traffic may not be traversed; however the directory network must be able to work with all devices where possible.
- The directory network must facilitate anonymity at the IP level in order to properly protect against observers compromising subscriber privacy.,.



Glynx: a new platform for directory services and communications

Glynx is the result of 5 years R&D and which satisfies the requirements for Social communication.

Glynx is the result of over 5 years research and development and is the only known platform which protects users' online identities and gives users full control of their communications and identities. It allows subscribers to locate people in real time and securely share private information used in call control and other purposes.

This section introduces Glynx. It is divided into two parts. The first part describes the Glynx P2P directory overlay with the following parts,

- Introducing the Glynx P2P overlay network,
- Glynx information storage and retrieval,
- the special qualities of this directory - Glynx Black Pages,
- directory listing publishing and retrieval,
- Glynx identity management and a
- Glynx black pages example information flow.

The second part covers communications control in the Glynx environment including:

- Introducing Glynx and communications
- A-Party call control
- Ubeity and the Glynx communications configuration hypercube
- An example call control flow in Glynx

introducing the glynx p2p overlay network

The core of the Glynx architecture is a structured P2P network which provides the foundation of the private directory components.

A P2P network is typically implemented as an application layer overlay to the existing IP-based internet. Overlay networks fall into two major architecture groups – un-structured and structured

In an **un-structured** P2P network, nodes are interconnected haphazardly – i.e. links between nodes are formed as they are discovered or referred to by other nodes. There is no deterministic way to locate and/or communicate with an arbitrary node in the network and messages are passed on from node to node with an expectation of a reasonable probability that the message will reach the desired destination node. Un-structured P2P networks may be thought of as analogous to of human social networks. Examples of un-structured P2P networks are the various file downloading networks, Freenet and the JXTA project.

Glynx is a structured P2P overlay network which works across NAT & firewall devices.

It scales logarithmically with network size keeping query times within an acceptable range.



In a **structured** P2P network, by contrast every node has a unique identifier (node-id) and assumes a 'location' in a mathematically derived network topology.

The topological algorithms ensure that each node holds the addresses of a number of other nodes such that (a) every node in the network can be contacted deterministically and (b) the maximum number of nodes that need to be contacted in order to locate any arbitrary node varies logarithmically with the number of nodes in the network. Logarithmic behaviour is important to ensure that query response times remain within an acceptable range even though the network may grow by many orders of magnitude.

In order to communicate with devices behind firewalls and NAT devices, the Glynx P2P networking layer contains the following facilities:

- A UPnP client which allows inbound Glynx traffic to traverse compatible Internet Gateway devices (e.g. domestic and SOHO modem/router devices).
- A 'supernode' configuration which, while not performing any special role in data storage or retrieval, acts as a relay for those nodes which lie behind a NAT or firewall device and which cannot therefore accept inbound connections.

It is expected that, over time, Glynx will adopt other NAT and Firewall traversal strategies.

In addition, the Glynx P2P networking layer provide relay functionality to obscure the IP addresses of source and destination nodes of a query.

glynx information storage and retrieval

A structured P2P networks is usually implemented as a Distributed Hash Tables or **DHT** allowing the timely and deterministic storage and retrieval of data from nodes in the network. A DHT typically derives a digest or hash of the descriptor of data to be stored, and subsequently uses that hash as the identifier of the node to which the data is stored. Examples of structured P2P algorithms/implementations are CHORD, PASTry, Ocean Store/Tapestry, Bamboo, and Kademia.

The Glynx DHT stores keys which represent identity claims of subscribers together with their corresponding values that contain network addressing information and certification keys. This information enables the retriever to transmit a query request to the private directory associated with the published identity.

More specifically, the publish/retrieve mechanism works as follows:

Individual nodes within the Glynx overlay are allocated a unique

The Glynx directory is a Distributed Hash Table with keys generated from identities via a one-way algorithm.

It provides true privacy and is resistant to attack.



Node-Id. The node-id is not chosen by the node joining the Glynx DHT but is allocated randomly by Glynx central services (and certified by the Glynx Certificate Authority (CA)). In this way the Glynx P2P network is resistant to attacks where a party takes over a large number of nodes in a predictable segment of the DHT network.

A subscriber node applies the universal Glynx one-way digest algorithm to an individual identity claim (e.g. an e-mail address, a telephone number, an instant messaging handle, a web site ID etc.) resulting in an Identity-Key. It then publishes the Identity-Key in the DHT together with an IP address and an identity certificate through which the publisher's private directory may be contacted (often the same device as the publishing node) and other identifying information explained below (see Figure 1).

| | |
|------------------|--|
| DHT Key | The Identity-Key |
| DHT Value | <ul style="list-style-type: none">• The IP address of a communication path to private directory• Identity certificate• Collision discriminator• Endorsement certificate(s) – if any |

Figure 1 – DHT Listing Structure

The DHT layer stores this key/value pair by mapping the Identity-Key to a node in the network and requesting the node to store the key/value pair.

A retriever then, knowing an identity of the target subscriber, constructs an Identity-Key by applying the Glynx one-way digest algorithm to the identity and requests the DHT to return the value associated with this Identity-Key.

Again, the DHT maps the Identity-Key to a node in the network and requests the node to retrieve the value corresponding to the Identity-Key if it exists.

The Glynx DHT layer ensures that there are sufficient replicas of published information, located at topologically optimal points throughout the DHT, to ensure integrity of the directory in the face of node churn and to load balance high traffic nodes. Not only does the DHT re-balance itself in the event of nodes joining or leaving the DHT, but the originating nodes themselves periodically ensure that their listings have integrity and can be retrieved in a timely and deterministic manner.



*The Glynx directory is called **Black Pages** as it cannot be browsed and cannot be harvested for use in unsolicited communication*

An enquirer must have a priori knowledge of a published identity in order to search Black Pages.

special qualities of the glynx 'black pages'

Thus if a enquirer has a priori knowledge of a published identity claim, target subscribers can be located and queried for additional information, in a way that does not leak any additional information about the published identity, such as other identities or claims, real-time contact information, or any other private information. The target subscriber's private directory can evaluate the query and return zero or more pieces information depending on the target subscriber's relationship with the enquirer based on credential presented by the enquirer in the query.

However, because identity-keys are derived via a one-way algorithm (i.e. there is no algorithmic way of obtaining the original identity given an Identity-Key), the Glynx directory as a whole or even listings resident on any particular node cannot be browsed for identity claims and cannot be harvested for identity claims to be used for unsolicited communication. In addition, the use of relays prevents any observation of traffic to/from a particular node from yielding useful information about resident identity claims or their relationships.

The Glynx directory is termed **Black Pages** (as compared to 'White pages' or 'Yellow pages') due to its property of being browsable but non-searchable. It is opaque to enquirers who do not have a priori knowledge of subscribers' identities or who do not have sufficient credentials to yield a response from queried entries.

Identity claims within Black Pages are defined with reference to an arbitrarily extensible schema and Glynx uses a plug-in architecture to handle integration of emerging or unforeseen directories and services.

It should also be noted that the IP address associated with an Identity-Key may not be the direct IP address of the publishing node, but may be the address of an intermediate node in a relay-scheme designed to protect the anonymity of publishers' devices from enquirers of the directory.

directory publishing and retrieval

Glynx subscribers publish and search for any identity claim within Black Pages without reference to a 'master identity'. This enables subscribers to manage their identity claims natively as they naturally would outside of the electronic context. Using identities as query keys, subscribers can organise their identities and other personal information consistently into 'personas' and present these to audiences as they see fit, independently of the constraints of any identity-issuing entity or its associated services.

Every identity claim is associated with a PKI public/private **key-pair** and a **certificate** referencing the Identity claim and containing its



public-key. These certificates are published alongside the IP address within the Black Pages listing (see Figure 1. above). All request and response messages between peer nodes are protected by PKI encryption to ensure integrity and non-repudiation – i.e. all messages are signed by the sender's private key and encrypted with the receiver's public key.

Most P2P implementations either assume nodes are 'friendly' and will not corrupt information (e.g. those that are academic in nature and usually structured) or do not care if a node returns corrupted or false information (e.g. those that are public, un-structured and predominantly used for file-sharing). Glynx, on the other hand, does not assume anything about the integrity of participant nodes, and provides mechanisms to ensure that information held in the directory is not corrupted or false.

Identities within Glynx are protected by strong encryption and may be verified by a trusted third party.

Glynx can store approximately 10^{77} distinct identities.

In order to ensure that Blackpages information can be trusted, subscribers can choose to have their identity claims **verified**. Verification enables enquirers to know (to the degree of confidence specified by the verification process) that an identity claim in Blackpages represents the true publishing subscriber.

Once an identity claim has been verified, an **endorsement** certificate is issued by the verifier and returned to the subscriber that owns the verified identity. This certificate, which references the Identity-Key and is signed by the verifier, is also published alongside the IP address within the Black Pages listing (see Figure 1. above).

The Glynx architecture allows for five levels of verification of identity claims, which are:

- Not verified

Unverified identity claims do not have any associated endorsement certificates.

- Verified by the publishing subscriber

Self-verified identity claims have an endorsement certificate signed by another identity of the same subscriber. For example, if there is an E-Mail identity endorsed by a phone number certificate, then "...if I trust that this is your phone number, and I trust you, then I can also trust that this is your e-mail address ..."

- Verified by another subscriber (peer verification)

With this mechanism, endorsement certificates are issued by peer subscribers to Black Pages – usually as the result of some sort of out-of-band verification transaction.



It is anticipated that peer-verification will be used by those who do not wish to use third-party or issuer verification for privacy reasons.

- Verified by a third party verifier

Third party verification is performed by a trusted authority via a non P2P transaction from a Glynx Client. It is achieved by sending an out-of-band message to the Identity (using the Identity's communications medium) asking for confirmation that the subscriber wishes to be listed in the directory. In this way subscribers must authorise for their identities to be published in the directory using information only obtained from them privately over means outside of the system.

Verification processing and the issuing of the verification certificates is currently performed by Glynx central services using a generalised verification engine together with the Glynx Certificate Authority. Note there is no architectural constraint requiring Glynx to uniquely perform this service.

- Verified by the Identity Issuer

In this case, the endorsement certificate is issued along with the identity by the service provider concerned – e.g. a telco for phone numbers, ISP for email addresses etc. as with other Black pages listings, the subscriber publishes the endorsement certificate along with the identity claim..

The various levels of verification enable enquirers to choose the level of endorsement they will trust when querying or retrieving Black Pages listings. For example “only trust listings verified by subscribers I know, by third party verifiers I trust, or Identity Issuers.

Currently, only Glynx provides verification services. In the future it is expected that other organizations will wish to become verifiers and that Glynx will expose an API to allow other parties (such as third-parties and issuer verifiers) to verify identity claims.

The Glynx one-way digest algorithm has been selected to minimise the possibility of key collision and can accommodate 2^{256} (approx 10^{77}) unique Identity-Keys. However as all digest algorithms lose information there is always the possibility of a collision –i.e. two or more Identical Identity-Keys which refer to different identities. To resolve this issue, Glynx also stores a public-key-encrypted version of the identity within the Black Pages listing to enable enquirers, in the event of a collision, to discern the correct listing.

It should be noted that during verification is the only time in the core Glynx process that Glynx is aware of a true identity of a subscriber. At all other times the Glynx software only deals with Identity-Keys.

Glynx provides five levels of trust for identity endorsers.



Glynx does not keep records of the subscribers once the verification process is complete.

glynx identity management

A major goal of the Glynx Black Pages directory is to enable subscribers to maintain sets of Identity information that may be disclosed to particular Audiences. These sets of Identities are termed **Personas** and may contain any information which is relevant to the application under consideration. Of particular interest are those Persona elements which relate to social communication. For example, a subscriber may have a business Persona which contains different Identity information to a Friends-and-Family Persona, or a Tennis-Club Persona.

Glynx provides direct support for managing multiple identities, personas, and audiences.

Glynx allows subscribers to maintain their contacts in groups of arbitrary **Audiences** to which one or more Personas may be disclosed. For example members of a work Audience may receive only a Business Persona but members of a Family Audience may receive a Business and a Friends-and-Family Persona. A role of the private directory component of Black Pages is to ascertain the Audience to which an enquirer belongs and to disclose the correct Persona(s) (if any) to the enquirer.

Once Personas have been exchanged the two parties accept each other's Identities with some level of confidence, negotiate a level of trust based on their respective Audiences and are now considered to have formed a Glynx relationship. They typically then link whenever they come online and update each other with changed information relevant to the Audience categorisation – both the information which changes infrequently or that which changes in real time.

glynx black pages example information flow

The following is a typical information flow showing the relationship between directory publisher, enquirer, Glynx verification and the DHT:

1. The publisher chooses an identity and generates an identity key, a PKI key-pair and a self-signed certificate referencing this Identity-Key and the generated public-key.

For example, the identity Email-Address/publisher@host.com is transformed into Identity-Key 6789.....

2. The publisher initiates a Verification Request containing the Identity-Key; in this case Email-Address/publisher@host.com and the generated certificate.

The Glynx Verification Service sends (in this case) an e-mail containing a verification token to the identity. The publisher completes the verification request by supplying the token to the Glynx Verification Service.



3. The Glynx Verification Service returns an endorsement certificate for the Identity-Key to the publisher
4. The publisher issues a PUT request to the DHT comprising:
 - a. The Identity-Key
 - b. The self-signed certificate
 - c. The endorsement certificate
 - d. A collision discriminator obtained by encrypting the Identity by the public-key generated in step 1 above.
 - e. An IP address which may be used to communicate with the publisher.
5. An enquirer wishing to communicate with the publisher, and knowing the publisher's e-mail address, generates an identity key using the same algorithm as the publisher and issues a GET request to the DHT for Identity-Key=6789....
6. The DHT locates the nearest node before node-id = 6789... (or one of its replicas), the 'listing node' and obtains the listing published by the publisher and returns it to the enquirer.
7. The enquirer formulates and issues to the publisher a cryptographically protected application-level request for Persona information (note that in order to protect the IP anonymity of the publisher the path may not be direct).
8. The publisher, after checking the credentials of the enquirer and ascertaining the relevant Audience, responds with cryptographically protected Persona information.

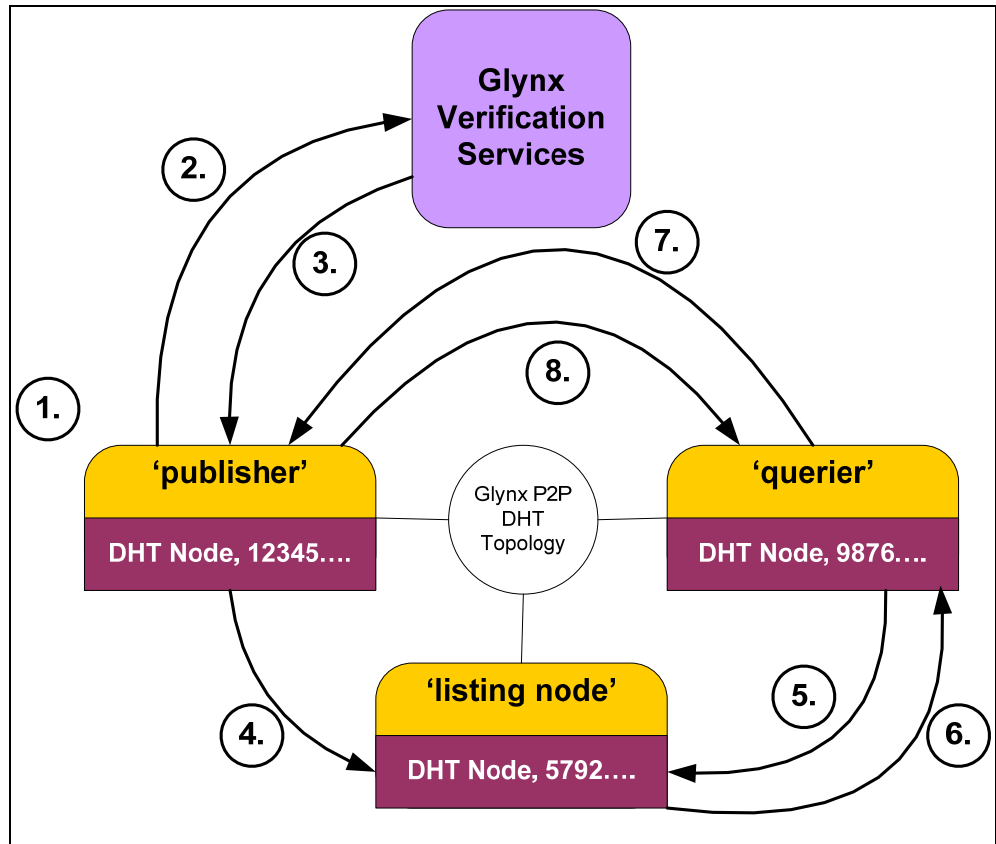


Figure 2 - Glynx Black Pages information flow.

This flow is shown in Figure 2. It can be seen that at no time is Glynx aware of the relationships between subscribers nor does Glynx have any visibility of information transferred.

Introducing glynx and communications

The Glynx approach to communications leverages the Black Pages directory to provide real-time control of communications options from subscribers' devices at the time of calling.

Glynx uses the Black Pages directory to provides real-time control of communications options from subscribers' devices at the time of calling

This section covers key issues with current communications control approaches, A-party call control, ubiquity and the Glynx communications configuration hypercube, an example call control flow in the Glynx environment,

Glynx considers communications to occur between two or more parties over Transports and Networks. A **Transport** is the physical infrastructure over which communication occurs and a **Network** is a separate logical group of Parties who can communicate together directly (i.e. without relaying communications through another party or a gateway). Each Network has a unique Network name and clusters of Networks may also be known by a single name. Networks and Transports may be connected to each other through gateways or



means to form larger groupings. For example the fixed line TDM (Time Division Multiplex) voice network in each country is typically connected to other country's fixed line TDM voice networks through gateways forming the international PSTN (Public Switch Telecommunications Network) for voice calls. Networks run over Transports and may span several Transports.

For example mobile networks may transit calls over wireless and/or wired networks and end mobile devices may be attached wirelessly or through a wired fixed connection to the network. Several Networks may run over the same Transport but each Network will have independent relations with its Parties.

A **Communications Identity** is the unique address by which a service subscriber is known on a specific Network at a specific time. The Communications Identity is the combination of both the Network address and the Network name. As described above, people may have many Communications Identities spanning a wide variety of Networks.

Glynx allows the user to control all aspects of communications based on identities, device capabilities, available networks, and modes of communication – and exercises this control differentially based on personas and audiences.

Service subscribers use **Devices** to communicate and may be human individuals using Devices, organisations with human representatives using Devices, or may be systems that can communicate, for example programmed computer systems such as Interactive Voice Response Systems. Devices connect subscribers to Networks. Examples of Devices are telephones and computers running communications software.

A Device may connect to several Networks. Each Network will have at least one associated Device that provides the functionality to enable a service subscriber to connect to the Network using the Device and provides functionality to facilitate communications. For example some Devices may allow video communications and others only voice.

Typically Networks are not directly aware of the state or identity of subscribers using each Device. So, for example, a Network can generally not direct calls to an subscriber, it can only direct calls to a Communications Identity associated with a Device at a particular address and the calling party makes the assumption that the Communications Identity is associated with the subscriber of interest. A caller will select a device & network with which to place a call, the choice typically being made on the basis of the caller's assumptions about availability, cost and quality of the Device and/or Network.

However, Communication may occur via qualitatively different mechanisms or **Channels** which have different characteristics in terms of cost, fidelity, usability access, and real-time responsiveness, - for example audio (conventional telephony equipment), video (video-phones, video-conferencing equipment, computer-computer



applications), messaging (e-mail, Instant-Messaging, facsimile) etc. Particular Networks and Devices will support different sets of Channels, which may be associated with a given Communications Identity, Device, or Network. For example, a Communications Identity associated with a Device which supports video and audio will have different Communications Capabilities than if the Device supported audio only.

Communications Capabilities may change over time in response to environmental or other changes. For example, a Network may change from full video to audio-only in the event of a network degradation – while maintaining the Communications Identity, Network And Devices of the current communications.

a-party call control

Glynx enables A-parties (calling parties) to use the Black Pages directory, along with the elements below, to exercise full control of a call and so optimise their communications rather than relying on Networks' or B-parties' (called parties') optimisation rules. This optimises A-party productivity and minimises full end-to-end network costs and quality of a call.

Glynx uses Ring Timeouts to avoid network or third-party interference during call set-up or maintenance.

Call Control is exercised by the Glynx-enabled Device by allowing the Device to obtain and use cost, quality or other factors of communication and, in particular, Ring Timeout Parameters (RTPs). RTPs establish the minimum time at which the Network or called party Device is expected to intervene with their prevailing call control rules and are usually expressed as a time interval, for example a number of seconds.

RTPs are published to Audiences along with their corresponding Communications Identities using the methods described above.

The RTP is used by the Glynx dialler so as to (a) avoid interference or termination of the communication by other network elements (for example, the Network, PBXs or B-Party Device), or (b) seek to improve the continuing cost, quality or other factor of the communications. The RTP may be employed to attempt to establish several alternative communications methods with or without calling party user intervention or presumption of information regarding the called party in parallel or series while avoiding interference by networks or called party devices.

The calling party Device only initiates calls to Communication Identities with a viable (i.e. non-zero) RTPs and then only for a duration up to, but not including, the duration of the RTP. In this way, barring unanticipated Network or called party intervention, the calling party Device remains in control of communications to Communication Identities for which it has an RTP. The initiation of communications which remain unanswered at the expiration of the RTP may be



abandoned by the calling party Device before Network or called party Call Control Device intervene.

*Glynx unifies the disposition for communications and capabilities into a complex 'Rich Presence' service termed **Ubeity**.*

Ubeity allows Glynx subscribers to optimize communications with each other.

ubeity and the glynx communications configuration hypercube

Glynx employs the concept of **Ubiety** to determine the optimum communications channel to be used between parties. Ubiety provides a mechanism for a subscriber to make representations to particular Audiences regarding the collective availability and capabilities of its Communication Identities. Ubiety is a function of Device capability, subscriber **Location**, the subscribers **Disposition** or receptiveness to receiving communications, and the real-time availability of Devices, Transports, Networks and Channels.

For example, the Ubiety as a result of a subscriber being at the location "at office", disposition "working", Voice Channel "engaged" can imply a different set of Transports and Communication Identity Presence States to an Audience, than Location "at airport", Disposition "receptive to public communications", Voice Channel "not engaged", Transport WIFI "unavailable"; or Location "home", Disposition "not working", Voice Channel "engaged". In fact, Location and Disposition are, themselves generalisations and the specificity or granularity of these elements is dependent at least partly upon the subscriber's ability to specify them or the Device's capability to determine them automatically.

The Presence State of a given Communication Identity may be associated with an RTP where, for example, a value of zero may mean 'do not attempt communication', and a non-zero value may represent the timeout duration after which an attempted call should be abandoned. The Presence State may also be associated with other attributes which A-Party Call Control may take into account. These may be quantitative metrics, for example bandwidth, latency, or priority as well as qualitative assessments such as "high quality" or "preferred". The RTP for a given Communication Identity may, at any given point in time, be dependent on not only the Entity's Ubiety (derived from, say, location and disposition), but also on any Device, Transport, Network, Channel or other override information currently in effect.

Glynx models Ubeity information as a five dimensional hypercube.

The complete Communications state is conceptually modelled as a five-dimensional "hypercube", the dimensions being Identity, Device, Transport, Channel and Ubiety, as the communications parameters may vary depending upon variation in any of these dimensions. At any point in time, a given Presence State is represented by a locus of co-ordinate points in this hypercube space, as the set of table rows identified by a given set of Identity, Device, Transport, Channel and Ubiety values.



Glynx uses the Ubeity hypercube information exchanged between subscribers to optimize each and every communication.

An example call control flow in glynx

The Glynx dialler uses the hypercube information received via Persona exchange to control communications using the complete capabilities and preferences of the calling and called parties.

Parties who wish to communicate typically form a relationship as described in the previous sections regarding the Glynx Black Pages. However once a Glynx relationship has been formed, they each inform each other with Ubeity and hypercube update information relevant to their respective Audiences. This information takes the form of fairly static Hypercube information together with Ubeity and override information which changes in real time.

A Glynx subscriber defines their communications information hypercube by specifying Communications Identities (organised into personae) together with information regarding their available devices and communications priorities and preferences etc.

1. At some point prior to initiating or at initiation of communication, the calling party locates the called party in the Glynx Black Pages directory and seeks to establish a relationship with the called party. The calling party supplies relevant Persona information, as credentials, to the Audience to which the called party has been associated.
2. The called party, after checking the credentials of the calling party, responds with its Persona and Hypercube information. Each party is now aware of the static set of communications options and preferences of the other party.
3. At the time of communication, the calling party queries the called party for its current Ubeity (while Glynx allows subscribers to act pro-actively by informing others, in real-time, of changes in Ubeity, a just-in-time check is necessary to avoid the mobile-phone-just-stepped-into-the-elevator effect.).
4. The called party responds with its real-time Ubeity together with any overrides which may affect the status of communication options defined in the hypercube (e.g. Network unavailable, VoIP client not running etc.). Overrides may also be discerned by the calling party from sources other than the called party device e.g. network or service statuses etc.
5. The calling party uses the supplied Ubeity and override information to parse the hypercube obtained previously from the caller and determine the Presence State; the communications options and preferences currently applicable to the called party. This is then merged with the calling party's communications capabilities and preferences to obtain a prioritised list of communications channels applicable to this communication instance – at the time of calling.

The Glynx dialler then either presents this list to the caller for communication to be invoked manually, or automatically initiates communication using RTP information to 'pull back' unanswered calls and progress to the next channel in the list as required to complete the call.

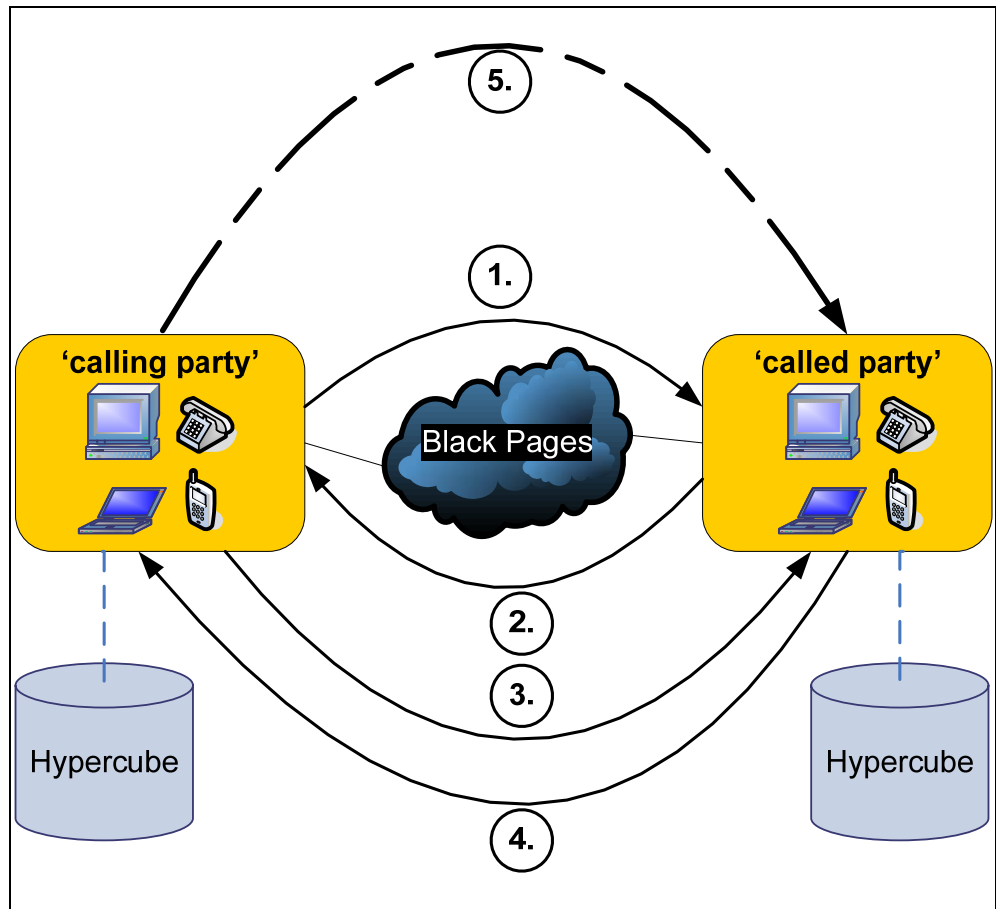


Figure 3 - Call control flow

At all times control of communication is in the hands of the calling and called parties and is totally managed from their respective devices.

Ubeity and Hypercube information are updated in real time between parties and may result in changes to the real time optimum communications method between the parties. Either party may initiate a change in communications methods using a new Ubeity or Hypercube update to ensure communications are continually optimised for changes in Ubeity or Hypercube information. The Device handles synchronising the call during changes.

The effect of this is to optimise each and every communication instance between Glynx subscribers and greatly reduce the incidence of 'telephone-tag'. Subscribers place a call to the person – not a series of phone numbers, VoIP address or Messaging IDs – and the Glynx dialler determines the best mechanism for connection. And all this is done without reference to a centralised directory



For example, if the called party is at home, and has previously shared a Friends-and-family Persona with the caller, then the called party's home number may be at the top of the caller's dialler list (and the first number to be automatically called by the Glynx dialler if so configured). If the caller has received only a business Persona, then the called-party may be 'unavailable'.

If they are both active on a computer and have a common VoIP client running, then a VoIP channel may have priority. If the called party is currently engaged on a call, then voice channels may be removed from the list or, alternatively, a message or chat may be initiated.

Most importantly, if a called party is not available for communication, the caller will be informed of this and can choose to leave a message for the called party using whatever mechanism they mutually prefer (email, SMS, IM, etc) preventing unnecessary network voicemail. Glynx also provides the capability to record multimedia messages and deliver them via the Glynx private and secure channel.

At all times control of communication is in the hands of the calling and called parties and is totally managed from their respective Devices.

conclusion

We have seen that Glynx delivers a new approach to social communication which uses the power of consumer devices to deliver true privacy and subscriber functionality in social communication.

- Glynx facilitates true privacy and security. With Glynx, discovery, setup & maintenance of social communication relationships together with the exchange of personal information takes place privately – without observation, interference or exploitation by any third party. Glynx is the private channel used in conjunction with other social communication services
- The Glynx Black Pages directory is searchable and non-browsable. It is the first viable directory to be used for E-Mail, VoIP, and IM addresses without fear of harvesting for SPAM.
- Glynx provides mechanisms for subscribers to manage identity proliferation in terms of Personas and arbitrarily granular Audiences, enabling expression of Identity as a 'personal brand'.
- The Glynx call control regime and dialler enables subscribers to reduce costs and lost productivity incurred due to calling at an inappropriate Ubeity (telephone tag) or calling a sub-optimum channel - greatly increasing communications productivity

Glynx delivers a new approach to social communication which uses the power of consumer devices to deliver true privacy and subscriber functionality in social communication



Glynx breaks the hold by service providers on directory information allowing innovation in social communication which has been previously impossible. For example:

- Optimisation of communications across networks which will enable innovative calling options – e.g. “call me back wherever I am”, “call any member of an audience” etc.
- Reduced inbound call centre queues and costs through call me back requests when I am in an appropriate Ubeity about my query
- Industrial strength decentralised information repositories - e.g. Medical records
- Private location ‘mash-ups’
- Etc.

Glynx is the new peer-to-peer architecture for social communication services

Thus it can be seen that Glynx is the new peer-to-peer architecture for social communication services which delivers absolute privacy and trust of online relationships; user-centric control of online identities and information; and a unified mechanism for simplifying the burgeoning online communications world.



References

1. Turner Brough, "Phone Numbers and Our Evolving Communications Identity" Internet Telephony Magazine September 2007 / Volume 10 / Number 9
<<http://www.tmcnet.com/voip/0907/the-next-wave-redux-phone-numbers-and-our-evolving-communications-identity.htm>>
2. O'Reilly Tim, "Social Network Fatigue and the Missing Web 2.0 Address Book", O'Reilly Radar, Sun 02.11.07
<http://radar.oreilly.com/archives/2007/02/social_network_1.html>
3. <http://www.web2summit.com>
4. Perez Juan Carlos, "Facebook's Beacon More Intrusive Than Previously Thought", PC World Friday, November 30, 2007 4:10 PM PST,
<<http://www.pcworld.com/article/id,140182-c,onlineprivacy/article.html>>