More Instant Messaging Interoperability (mimi)
Internet-Draft
Intended status: Informational
Expires: 23 December 2025

G. Hogben
F. Olumofin
Google
J. Peterson
TransUnion
J. Rosenberg
Five9

21 June 2025

User Discovery Requirements
draft-interop-mimi-discovery-requirements-03

Abstract

This document defines requirements for the user discovery problem within the More Instant Messaging Interoperability (MIMI) working group. User discovery is essential for interoperability, allowing message senders to locate recipients across diverse platforms using globally unique, cross-service identifiers (e.g., email addresses, phone numbers). The core challenge involves reliably mapping these identifiers to messaging service providers and determining the reachability of a recipient’s identifier across multiple providers.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at https://datatracker.ietf.org/doc/draft-interop-mimi-discovery-requirements/. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-interop-mimi-discovery-requirements/.

Discussion of this document takes place on the mimi Working Group mailing list (mailto:mimi@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/mimi/. Subscribe at https://www.ietf.org/mailman/listinfo/mimi/.

Source for this draft and an issue tracker can be found at https://github.com/ietf-wg-mimi/draft-ietf-mimi-user-discovery-reqs/.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Copyright Notice

Table of Contents

1. Introduction

MIMI user discovery enables a message sender to locate messaging service providers on which a particular recipient can be reached. Currently, users often need to ask contacts what service they are on out-of-band or try multiple services, which creates friction. Specifically, discovery helps a user on one messaging service (Alice on Service A) to find another user on a potentially different or the same service (Bob on Service B) without prior knowledge of Bob's provider, and in a provider-neutral manner.

Discovery is necessary because the identifiers we commonly have for contacts (phone numbers, email addresses, etc.) do not necessarily tell us which messaging service they are using. Someone with the email alice@gmail.com might use iMessage, Signal, or another service entirely. Thus, the core problem is how to take one of these cross-service identifiers and learn the messaging service provider that the user is using and how to communicate with them on that service.

2. Conventions and Definitions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1. _Cross-Service Identifier (CSI)_: A globally unique identifier for a user across multiple services (e.g., E.164 phone number, email address).

2. _Cross-Service Identifier Provider (CSIP)_: An entity that issues and manages CSIs (e.g., telecom providers, email providers).

3. _Messaging Service Provider (MSP)_: An entity offering messaging services to end users (e.g., WhatsApp, Signal, iMessage).

4. _Service Specific Identifier (SSI)_: A unique identifier for a user within a single MSP (e.g., a Twitter handle).

5. _Discovery Provider (DP)_: An entity that stores and facilitates the discovery of MSPs for a given CSI. A DP may be operated by an MSP or a third party.

6. _Verifiable Mapping_: A representation of the cryptographic binding of a CSI and the set of MSPs where the CSI is reachable.

7. _Recipient_: A user who possesses a CSI that has been assigned by a CSIP and authorizes verifiable mapping of that CSI to an MSP set.

8. _Sender_: A user who queries a DP to discover the MSPs for a given CSI.

3. User Discovery Problem Statement

User discovery involves two key aspects:

1. _Authorization_: Recipients must be able to authorize verifiable mappings between their CSIs and their chosen MSPs.

2. _Lookup_: Senders must be able to query these mappings to determine where a recipient can be reached.

To authorize verifiable mappings, the recipient must:

* Possess a valid, CSIP-issued CSI.

* Have active accounts with each MSP to be mapped.

* Associate the CSI with each of those accounts.

Discovered mappings must be verifiable to ensure they are accurate. Crucially, discovery must prioritize user privacy, allowing users to control their discoverability, and it must integrate well with end-to-end encryption and other MIMI protocols.

Note: This document focuses on discovering _which_ MSPs a user is on. Retrieving SSIs and cryptographic keys can be a separate, subsequent step.

The rest of this document describes a series of requirements for the discovery problem.

4. Prior Efforts

Discovery services are far from new on the Internet. The whois protocol [RFC3912], largely focused on mapping domain names to associated services, was one of the earliest discovery services deployed on the Internet. DNS SRV records, specified in [RFC2782], allow a similar process - given a domain name, a user can discover available services, such as VoIP based on the Session Initiation Protocol (SIP) [RFC3261] [RFC3263]. SRV records were adapted specifically for messaging in [RFC3861]. However, both whois and DNS SRV records rely on domain names as lookup keys, making them unsuitable for identifiers like mobile phone numbers, which don’t have inherent domain associations.

ENUM [RFC6117] addressed this limitation. It used DNS to look up phone numbers by reversing the digits and adding the "e164.arpa" suffix. This allowed delegation of portions of the namespace to telecom providers who owned specific number prefixes. While technically straightforward, public deployment of ENUM was hampered by challenges in establishing authority for prefixes. However, private ENUM [RFC6116] services have become relatively common, facilitating functions like MMS routing within messaging.

Another attempt was ViPR (Verification Involving PSTN Reachability) [I-D.rosenberg-dispatch-vipr-overview] [I-D.petithuguenin-vipr-pvp]. ViPR utilized a peer-to-peer network based on RELOAD (Resource Location and Discovery) [RFC6940] to operate between enterprises. It addressed the authority problem by authorizing records based on proof of forward routability but faced the same network effects issue as ENUM. ViPR attempted to address incentive problems by focusing on enterprises seeking cost savings by bypassing the phone network. Ultimately, network effects challenges (among other protocol-unrelated issues) prevented widespread deployment.

Discovery and lookup services are now commonplace on the Internet but are largely scoped within large providers such as Facebook, Twitter, WhatsApp, and others.

The MIMI discovery service requires a solution that works across providers.

5. Summary of Requirements

This section provides a summary of the requirements for MIMI user discovery.

#	Requirement	Mandatory	Optional
	Recipients		
R1	Only the recipient MUST be able to authorize verifiable mapping of the recipient's CSI to an MSP set	x	
R2	A mapping MUST allow for the inclusion of tags or similar constructs to indicate the recipient's preferences for using each included MSP	x	
	Senders		
R3	Senders MUST be able to retrieve all verifiable mappings for a CSI	x	
R4	Senders MUST be able to verify the authenticity of mappings	x	
	Discovery Provider		
R5	Discovery MUST support results with zero, one, or multiple mappings	x	
R6	DPs MUST NOT be able to learn both the sender's identity and the recipient's CSI	x	
R7	All data exchanged in the processing of a discovery request MUST be encrypted in transit	x	
R8	DPs SHOULD be protected against enumeration attacks		x
R9	DPs MUST provide a way to remove mappings	x	

R10	Only the recipient or the CSIP SHOULD be able to remove mappings		x	

Table 1

6. Recipient Authorization of Mappings

R1: Only the recipient MUST be able to authorize verifiable mapping of the recipient's CSI to an MSP set.

This requirement prevents unauthorized mapping and potential impersonation attacks. The integrity of mappings requires the recipient to be the sole party to authorize mappings of the recipient's CSI to the designated MSP set. Allowing any other user besides the recipient to authorize the creation of verifiable mappings could open the recipient to impersonation attacks, where the MSP set will include MSPs where the attacker controls the accounts.

It is expected that the recipient has active accounts with each MSP in the set and has already linked the CSI with each of those MSP accounts. If a recipient includes MSPs that do not meet these preconditions, the recipient will not be reachable on the incorrect MSPs. While this problem can be resolved if each MSP participates in creating verifiable mappings to confirm the recipient's account exists, mapping creation will be more involved. Hence, this issue is deferred as a post-discovery problem to address.

7. Recipient Preferences in Mappings

R2: A mapping MUST allow for the inclusion of tags or similar constructs to indicate the recipient's preferences for using each included MSP.

This requirement allows recipients to guide how senders contact them on different MSPs. For example, a preference tag may be formulated as closed or open-ended strings (e.g., "Business", "Personal", "BasketballFriends", "WhatsApp") or a list (e.g., "Business, WhatsApp"). A recipient might want senders to utilize a specific "Business" mapping for business messaging and a different one for other types of messaging.

This requirement prioritizes recipient preferences because senders already have flexibility in choosing their MSP (and potentially DP) when initiating requests. Further considerations regarding sender and DP preferences, as well as more complex recipient preferences, are deemed to be implementation issues. See Appendix B: Recipient's Critical User Journeys for some more examples for recipients.

8. Sender Retrieval of Mappings

R3: Senders MUST be able to retrieve all verifiable mappings for a CSI.

This requirement ensures senders have complete information for reaching a recipient. When a sender queries for mappings established for a given recipient's CSI, all available verifiable mappings that exist must be returned as a response.

9. Mapping Authenticity Verification

R4: Senders MUST be able to verify the authenticity of mappings.

This requirement enables senders to trust the mapping information. On receipt of the mappings that exist for a recipient's CSI, the sender should have a means to verify the authenticity of the mappings by learning that the recipient authorized that mapping. The verification process may leverage all or parts of the mapping content and may involve network calls to an oracle to help with the verification, provided the oracle query process is consistent with Requirement 6.

10. Support for Varying Mapping Results

R5: Discovery MUST support results with zero, one, or multiple mappings.

This requirement covers cases where a user has no mappings, has opted out of discovery, or uses multiple MSPs. Regardless of how the verifiable mapping structure is represented (e.g., one CSI to a set of MSPs, or one CSI to one MSP repeated for each MSP), senders must be able to determine the following from a user discovery response for a given CSI:

* No mapping exists, e.g., the recipient has not created any mapping.

* One or more mappings exist, but without listed MSPs, e.g., the recipient has intentionally configured discovery to return no mappings.

* One mapping exists with one listed MSP.

* One mapping exists with multiple listed MSPs.

* Multiple mappings exist, each potentially with a different number of MSPs.

In situations where the recipient who authorized mappings using a CSI, such as an email address, becomes unreachable (e.g., due to the user's death), implementations can rely on the MSPs included in the mapping to resolve non-reachability issues outside user discovery.

11. Protecting Sender and Recipient Privacy

R6: DPs MUST NOT be able to learn both the sender's identity and the recipient's CSI.

This requirement prevents the DP from building a complete social graph of users across MSPs. This protects user privacy by limiting the information the DP can gather about relationships between senders and recipients.

User discovery lookups inherently create a connection point (an edge) on the social graph between the sender and the recipient. To protect the privacy of both parties, the DP must be prevented from learning this connection consisting of the sender's identity and the recipient's CSI/mapping. Specifically, the DP can learn at most one of the following:

1. Identifying information about the sender, such as their source IP address, username, etc.

2. The recipient's CSI and any associated response mappings.

While the DP might be able to infer this edge or connection later if both users communicate through the same MSP, this requirement focuses on preventing the DP from directly learning this during discovery. To understand the importance of this, it's helpful to distinguish between two types of social graphs:

* Messaging social graph: Reflects actual communication between users.

* Discovery social graph: Encompasses all attempted user discovery lookups, which can be significantly larger than the messaging social graph, as it includes searches for users the sender may not ultimately contact.

Protecting user privacy during discovery is crucial because, without it, a sender's entire discovery social graph could be revealed during bulk discovery that requires the lookup of all contacts on the sender's address book, which exposes a much broader range of potential connections than their actual messaging activity. Appendix C Protecting Sender-Recipient Social Graph Edge has some recommendations for implementation.

12. In-Transit Encryption

R7: All data exchanged in the processing of a discovery request MUST be encrypted in transit.

This requirement ensures the confidentiality of discovery requests and responses.

13. Protection Against Enumeration Attacks

R8: DPs SHOULD be protected against enumeration attacks.

This requirement ensures there is a security defense against attacks aimed at scraping mapping data. Discovery providers should implement mechanisms to defend against large-scale scraping of mappings from their database, which can be used to compile spam targeting lists.

One potential implementation approach is to utilize time-bound blind signatures. This method limits the number of user discovery lookups a sender can perform within a given timeframe. Each lookup request must include a unique, unblinded signature that cannot be linked to the sender's identity. To facilitate rate-limiting across different server entities (e.g., DPs and MSPs), this unique signature should be passed along during communication.

14. Mapping Removal

R9: DPs MUST provide a way to remove mappings.

R10: Only the recipient or the CSIP SHOULD be able to remove mappings

These requirements allow recipients to manage their discoverability and remove outdated mappings. Authorization to remove verifiable mappings for a given CSI should be limited to only the recipient who authorized the mapping or the CSIP managing that CSI (for cases where a CSI is no longer assigned to any recipient).

Recipients can change their minds and decide to make a prior mapping not discoverable; they may want to update mappings or start all over. Thus, a DP should provide a means to remove verifiable mappings. Once removed, a mapping should no longer be returned for senders' requests.

A recipient that authorized the creation of a verifiable mapping should also be able to authorize its removal from the backend of DPs. Similarly, a CSIP may request the removal of mappings for a CSI that has become unassigned.

A possible implementation approach is for verifiable mapping authorization to include a bit that the recipient can use to indicate if a mapping is new. When that bit is set, a DP may proactively de-list any existing mapping for that CSI (after asking the user to re-confirm). Note that the CSIP does not need to be aware of or be involved with the de-listing of mappings with such an approach. There are other approaches to ensure mappings are fresh and are not impacted when the same CSI is transferred between two recipients.

15. Security Considerations

Security considerations are addressed throughout the document, particularly in requirements R1, R4, R6, R7, and R8. These requirements focus on preventing unauthorized mapping, ensuring the authenticity of mappings, protecting user privacy, and securing data in transit and at rest.

16. IANA Considerations

This document has no IANA actions.

17. Normative References

[I-D.petithuguenin-vipr-pvp]
Petit-Huguenin, M., Rosenberg, J., and C. F. Jennings, "The Public Switched Telephone Network (PSTN) Validation Protocol (PVP)", Work in Progress, Internet-Draft, draft-petithuguenin-vipr-pvp-04, 12 March 2012, <https://datatracker.ietf.org/doc/html/draft-petithuguenin-vipr-pvp-04>.

[I-D.rosenberg-dispatch-vipr-overview]
Rosenberg, J., Jennings, C. F., and M. Petit-Huguenin, "Verification Involving PSTN Reachability: Requirements and Architecture Overview", Work in Progress, Internet-Draft, draft-rosenberg-dispatch-vipr-overview-04, 25 October 2010, <https://datatracker.ietf.org/doc/html/draft-rosenberg-dispatch-vipr-overview-04>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/rfc/rfc2119>.

[RFC2782] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, DOI 10.17487/RFC2782, February 2000, <https://www.rfc-editor.org/rfc/rfc2782>.

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, DOI 10.17487/RFC3261, June 2002, <https://www.rfc-editor.org/rfc/rfc3261>.

[RFC3263] Rosenberg, J. and H. Schulzrinne, "Session Initiation Protocol (SIP): Locating SIP Servers", RFC 3263, DOI 10.17487/RFC3263, June 2002, <https://www.rfc-editor.org/rfc/rfc3263>.

[RFC3861] Peterson, J., "Address Resolution for Instant Messaging and Presence", RFC 3861, DOI 10.17487/RFC3861, August 2004, <https://www.rfc-editor.org/rfc/rfc3861>.

[RFC3912] Daigle, L., "WHOIS Protocol Specification", RFC 3912, DOI 10.17487/RFC3912, September 2004, <https://www.rfc-editor.org/rfc/rfc3912>.

[RFC5222] Hardie, T., Newton, A., Schulzrinne, H., and H. Tschofenig, "LoST: A Location-to-Service Translation Protocol", RFC 5222, DOI 10.17487/RFC5222, August 2008, <https://www.rfc-editor.org/rfc/rfc5222>.

[RFC6116] Bradner, S., Conroy, L., and K. Fujiwara, "The E.164 to Uniform Resource Identifiers (URI) Dynamic Delegation Discovery System (DDDS) Application (ENUM)", RFC 6116, DOI 10.17487/RFC6116, March 2011, <https://www.rfc-editor.org/rfc/rfc6116>.

[RFC6117] Hoeneisen, B., Mayrhofer, A., and J. Livingood, "IANA Registration of Enumservices: Guide, Template, and IANA Considerations", RFC 6117, DOI 10.17487/RFC6117, March 2011, <https://www.rfc-editor.org/rfc/rfc6117>.

[RFC6376] Crocker, D., Ed., Hansen, T., Ed., and M. Kucherawy, Ed., "DomainKeys Identified Mail (DKIM) Signatures", STD 76, RFC 6376, DOI 10.17487/RFC6376, September 2011, <https://www.rfc-editor.org/rfc/rfc6376>.

[RFC6940] Jennings, C., Lowekamp, B., Ed., Rescorla, E., Baset, S., and H. Schulzrinne, "REsource LOcation And Discovery (RELOAD) Base Protocol", RFC 6940, DOI 10.17487/RFC6940, January 2014, <https://www.rfc-editor.org/rfc/rfc6940>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/rfc/rfc8174>.

[RFC8616] Levine, J., "Email Authentication for Internationalized Mail", RFC 8616, DOI 10.17487/RFC8616, June 2019, <https://www.rfc-editor.org/rfc/rfc8616>.

Appendix A. Architectural Models

This appendix explores various architectural models for MIMI DP to provide an overview and address practical implementation considerations. The working group observed these requirements are similar among these models and opted to maintain architectural neutrality for the discovery protocol. However, we will outline requirements for the roles of DPs, how they interact with each other, MSPs in a federated model, and how DPs accommodate queries from both MSPs and users.

A.1. Centralized DP (Monolithic Service)

A globally accessible, authoritative database (potentially sharded/replicated) stores all authenticated CSI mappings. This monolithic service, implemented across synchronized global nodes, handles all CSI queries from messaging platforms and acts as the single source of truth for mapping data, even if certain mappings may be restricted in specific regions due to geopolitical considerations.

A.1.1. Advantages

* Standardization and uniform control over mapping creation, updates, and data formats.

* Promotes fair and unbiased CSI mapping discovery.

* Single source of truth for mapping simplifies data management and ensures consistency.

* Sharding/replication can address geographical distribution and performance needs.

A.1.2. Drawbacks

* Centralization of sensitive user data raises privacy risks.

* May conflict with data localization regulations.

* Single point of failure vulnerability from outages affecting the entire system.

* Potential difficulty with immediate global updates for rapidly changing mappings.

A.2. Hierarchical DPs (Discovery Resolvers)

A global root Discovery Provider (DP) directs mapping requests to authoritative DPs based on CSI structure (e.g., country codes for E.164 phone numbers) or sharding mechanisms. The root DP, similar to hierarchical DNS, acts as a directory service, holding pointers to authoritative DPs rather than mappings themselves. Alternatives to hierarchical resolution, like the LoST protocol [RFC5222], or distributed hash tables (DHTs), can achieve similar outcomes.

A.2.1. Advantages

* Scalability from distributed load and mapping management across multiple DPs.

* Flexibility that allows different DPs to specialize in specific CSI ranges for regions.

* Better alignment with data localization requirements.

A.2.2. Drawbacks

* Requires coordination and maintenance of hierarchy.

* Root DP failure could disrupt the entire system.

* Potential delays due to additional hops in the discovery process.

A.3. Federated DPs (Distributed Peer Service)

In a federated model, multiple independent DPs collaborate to provide CSI mapping discovery. Each DP holds a subset of mappings and pointers to other DPs, with no central authority dictating mapping locations. Discovering all mappings for a CSI may involve querying multiple DPs. DPs can exchange CSI information or recursively query each other. The specifics of DP federation are determined by business agreements, not technical requirements. A variation of this model involves messaging platforms acting as their own DPs, managing mappings for their users.

A.3.1. Advantages

* Decentralization ensures there is no single point of control or failure. The system can continue functioning even if some DPs are unavailable.

* Mappings can be distributed according to local regulations.

A.3.2. Drawbacks

* Discovery process may involve querying multiple DPs, increasing network load.

* Potential for bias as DPs may prioritize their own mappings, leading to uneven results.

* Requires robust mechanisms to prevent CSI impersonation and ensure trust between DPs.

* Pairwise relationships in a federated DP model could create a barrier to entry for smaller MSP/DPs, similar to the trust requirements in the DKIM [RFC6376] and DMARC [RFC8616] ecosystems.

A.4. Additional Considerations on Architectural Models

A.4.1. Telephone Number CSIs

Telephone number portability is complex due to its reliance on real-time queries to proprietary and legacy systems. Overall, the authenticated mappings proposed for MIMI may necessitate additional platform measures to assess mapping freshness and ensure up-to-date reachability responses.

A.4.2. Bias Mitigation

Bias occurs when a DP prioritizes mappings to its affiliated MSP without consideration of what is best for end users. Mitigating bias is essential to ensure fair and equitable discovery of authenticated mappings across different services. The working group has decided to defer such mitigation to policies and regulations, excluding it from the discovery protocol.

Appendix B. Recipient’s Critical User Journeys

Here are some Critical User Journeys (CUJs) that are the most important to discovery recipients.

In the CUJs below, Bob is the recipient that creates verifiable mappings, and Alice is the sender or user performing discovery:

1. Sender mapping preferences: Bob only wants to be found by Alice and other users on WhatsApp, not his other messaging apps.

2. Same-app preferences: Bob prefers that Alice can find him on the same messaging service that she is using. In other words, Bob does not want cross-app discovery and messaging.

3. No-random mapping preferences: Bob does not want to go through multiple apps to find a message from Alice when discovery returns one of 10 mappings that Bob has established with discovery providers.

4. No-duplication preferences: Bob does not want Alice’s messages to be broadcasted to all or a subset of his apps based on the result of discovery.

5. Per-sender preferences: Bob wants to control which app messages from Alice go to and do the same for other users (e.g., Carol’s messages may go to a different app than Alice’s).

6. Closed group preferences: Bob only wants his soccer parents to discover and contact him on WhatsApp, not his Wire app. That is, a group of senders has the same mapping results based on Bob’s preferences.

7. Open-ended group preferences: Bob wants his business contacts to discover and reach him on Wire, not WhatsApp. That is, an open-ended list of senders (i.e., including leads) are provided with a designated mapping.

Appendix C. Protecting Sender-Recipient Social Graph Edge

To protect user privacy, implementations should consider:

1. Sender Identity Protection: This approach focuses on concealing the sender's identity from the Discovery Provider (DP). Techniques like IP blinding (e.g., using a Private Relay) can be employed to achieve this, ensuring the DP only learns the recipient's CSI.

2. Recipient Identity Protection: This approach aims to hide the recipient's CSI from the DP. Methods like Private Information Retrieval (PIR) or Private Set Membership (PSM) allow the sender to perform the lookup without revealing the recipient's information to the DP, effectively limiting the DP's knowledge to the sender's identity.

Acknowledgments

We gratefully acknowledge the valuable feedback and constructive discussions received within the working group, in individual conversations, and during the MIMI interim meetings, as well as IETF 119, 120, and 121.

Authors' Addresses

Giles Hogben
Google
United States of America
Email: gih@google.com

Femi Olumofin
Google
United States of America
Email: fgolu@google.com

Jon Peterson
TransUnion
United States of America
Email: jon.peterson@transunion.com

Jonathan Rosenberg
Five9
United States of America
Email: jdrosen@jdrosen.net