Distributed Universal Number Discovery (DUNDi™) and the General Peering Agreement (GPA™)

Mark Spencer, Digium, Inc.

October 10, 2004 – Version 0.4.1

Abstract:

Voice over IP has certainly made tremendous progress over the past several years, and many companies, including Digium™, have been exploring exciting new technologies for communication, such as instant messaging, video and speech recognition. However, the basic expectation that many people have of VoIP, that is, the ability to call someone else they know for free, still has not yet been fully realized. This paper explores some of the reasons for that and explains a new peer to peer (p2p) system developed by Digium to solve these problems.

Problems:

1. If you are sitting on the Internet and want to call someone, how do you know where to deliver the call (i.e. which server, which protocol, and which URI)?
2. If you make your number freely callable over the Internet, how can you set policies to prevent VoIP spam and undesired unsolicited calling?
3. Within an enterprise, how can the maintenance of a diverse range of communication servers be made easier and more robust

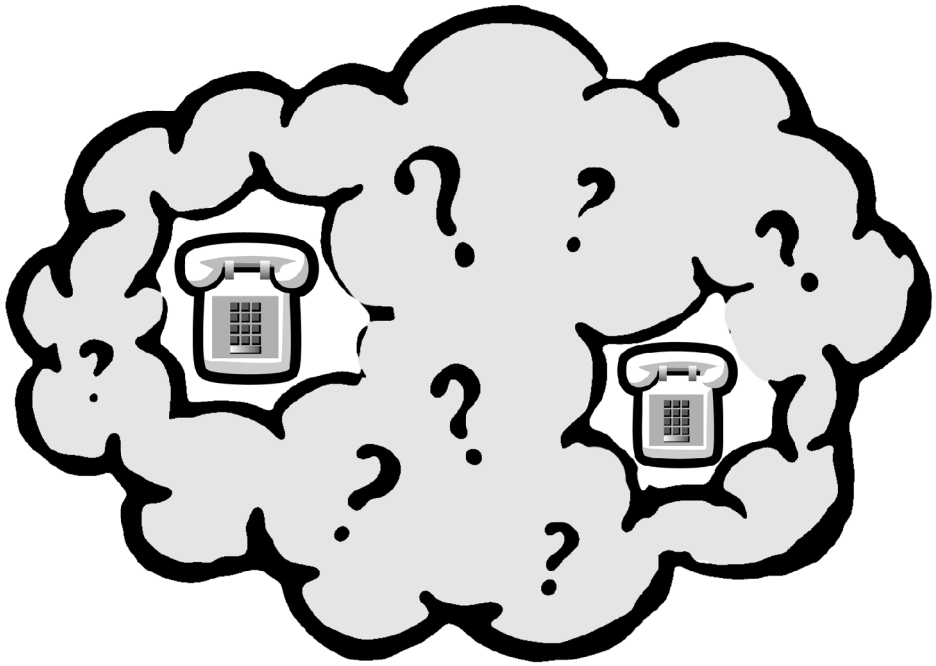

Existing Solution:

Russell: “I hit that pothole outside and lost a tire. Hey, do you know anywhere I can get a new tire for my car?”
Sara: “Check the yellow pages, why don't you.”
Russell: “Oh, good idea. You know, I wonder how much it costs to have an ad in the yellow pages.”
Tom: “You won't believe this, but I hit a gigantic pothole outside and now I need a new tire for my car. Do you know where I can go?”
Sara: “Ugh, Russell just asked me that. Check the yellow pages.”
Tom: “Oh yah, I'll try that. Man, I wonder how much those ads cost.”

ENUM is considered the traditional method for solving this problem. By utilizing the existing infrastructure of DNS, ENUM provides a framework for a strictly hierarchical method for locating services. However, ENUM presents several implementation issues:

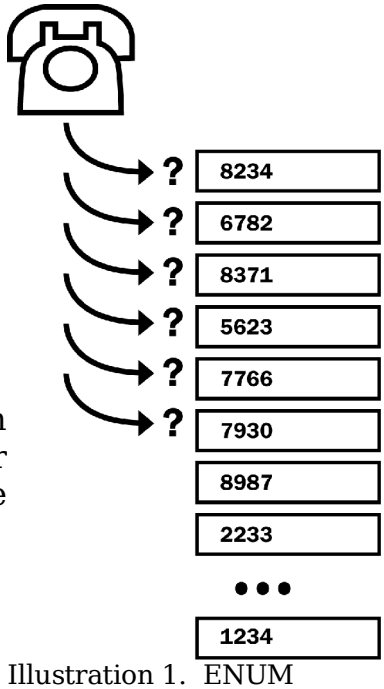

Illustration 1. ENUM

1. ENUM does not have any built-in access control for services – that must be handled separately (e.g. by the SIP security model).
2. ENUM requires some entity to provide the service of managing the root ENUM domain (e.g. e164.arpa) and requires authorities for each delegation (e.g. a single entity must be responsible for 6.5.2.1.e164.arpa and provide any delegations thereof). This entity is free to impose a charge, tax, etc for this service and everyone is dependent upon how well this entity runs the service.
3. ENUM requires making a complete list of all numbers or patterns that a provider terminates available to the party managing the ENUM database.
4. ENUM does not provide security or privacy for the numbers being requested. Anyone can observe which numbers are being looked up.

5. ENUM cannot be used to indicate that a number is partially correct (e.g. “1256428” is not a valid number, but there are valid numbers beginning with that pattern), so it can only be used for handling whole numbers, and is thus incapable of handling a complex dial plan for traditional analog phone service.
6. ENUM provides no information regarding the preference of the user with respect to unsolicited calling.
7. Within an enterprise, ENUM requires work to provide a failsafe solution against a server failure.

New Solution:

Russell: “I hit that pothole outside and lost a tire. Hey, do you know anywhere I can get a new tire for my car?”
Sara: “I'm not sure. Hey Andy, do you know where Russell can get new tires for his car?”
Andy: “Not off the top of my head. Hey Beth, do you know where Russell can get new tires for his car?”
Beth: “Yah, there's a place at the corner of Main and Wall that usually has good deals.”
Andy: “Beth says there is a place at Main and Wall.”
Sara: “Beth says there's a place at Main and Wall.”
Russell: “Thanks!”
Tom: “You won't believe this, but I hit a gigantic pothole outside and now I need a new tire for my car. Do you know where I can go?”
Sara: “Beth just said there's a place at Main and Wall.”
Tom: “Thanks!”

The new proposed solution to the stated problems is a combination of two pieces, one technical and one non-technical. This paper will begin with the technical portion and continue to the non-technical portion.

The Technical Portion: DUNDi

As a substitute or supplement to ENUM, Digium has developed Distributed Universal Number Discovery or DUNDi (pronounced “dun-

dee”). It is easiest to begin by considering DUNDi within an enterprise. In a traditional model, an enterprise would have a series of servers, potentially at different geographic locations, sharing a centralized ENUM repository to be able to locate other resources. The ENUM server may be made redundant through the use of specialized layer four switching. When a client requests a number or extension it does not know how to terminate, it would query the ENUM repository for a list of egress gateways for that service.

In the DUNDi model, there is no central repository. Instead, nodes (uniquely identified by a 6 byte Endpoint Identifier, typically the MAC address of an ethernet interface) participate in a trust system in which each node has a trust relationship with at least one other node in the system (often two or more for redundancy). When the client requests a number or extension it does not know how to terminate, it queries the nodes directly connected to it. Those nodes in turn will query nodes which are directly connected to them and so on (care is taken in the protocol to minimize the actual number of transactions and nodes that are queried, while maintaining the same effective answer as querying the entire system). The resulting answers are collated, cached appropriately at each involved node, and passed along to the original requesting party.

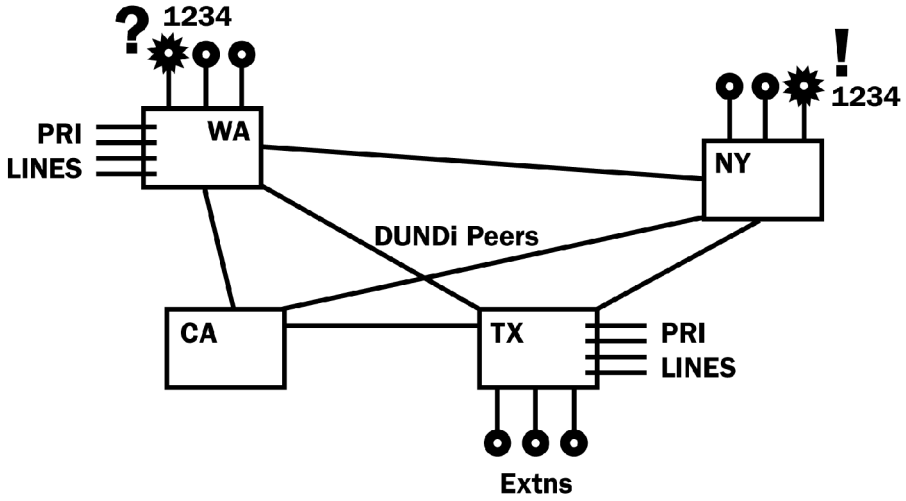

Illustration 2. DUNDi within an Enterprise

Within an enterprise, DUNDi allows an administrator to create a federation of entirely independent communication servers, each providing resources to, or demanding resources from, the trust network at large, using an efficient, binary encoded, AES and RSA encrypted and authenticated protocol both to keep the bandwidth utilization minimal and to ensure the privacy of the queries.

Care is taken at each node to minimize the number of queries, trim large portions of the graph or tree of nodes, and maximize the efficiency of caches.

DUNDi includes a weight system, similar to the way DNS provides a weight in MX records. This weight can be used to indicate a recommended priority to the party originating the query.

DUNDi is agnostic to the actual protocol used to set up and carry the call. IAX™ is the most frequently used protocol, followed by SIP.

DUNDi contains flags which indicate the availability of a number in more detail, for example whether a number is partially correct, whether a route may be used for placing unsolicited commercial or non-commercial calls, etc.

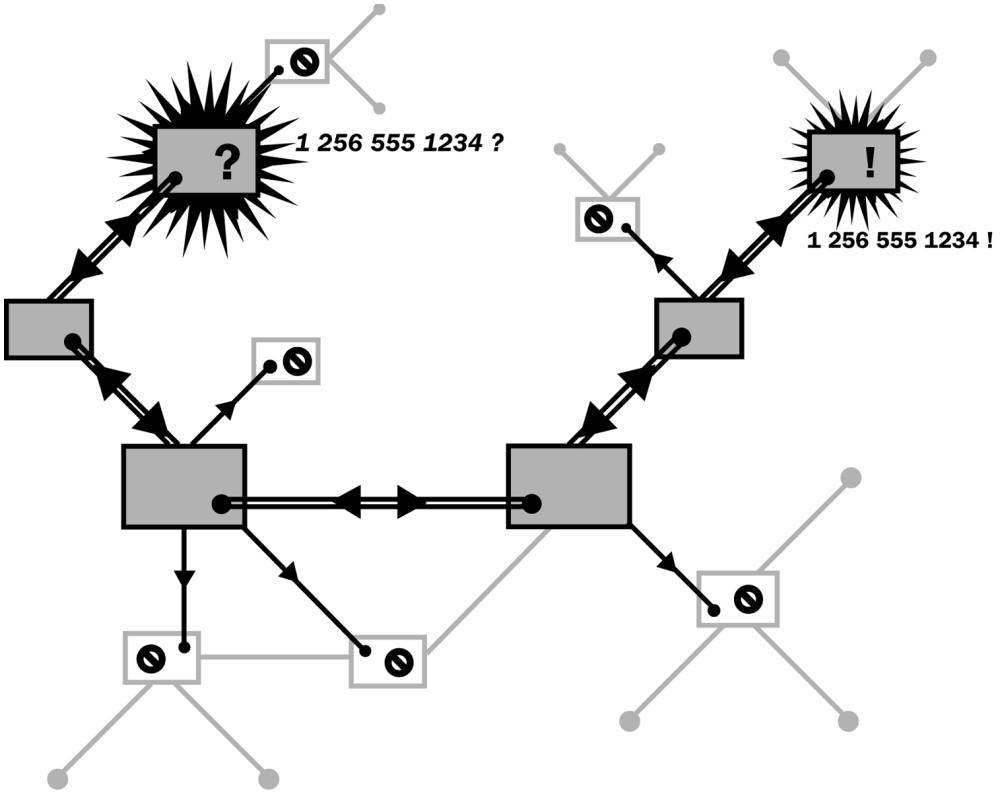

Illustration 3. DUNDi within E.164

DUNDi, like Asterisk™, includes the concept of extension contexts. Extension contexts are collections of numbers or extensions. A system may support more than one context, for example the “mycompany” context might contain 4 digit local extensions within an enterprise PBX, while the “e164” context would be reserved for true E.164 numbers outside enterprise. In this way, nodes can limit access for queries to

specific extensions or resources from certain other nodes.

Expanding the system beyond the enterprise, we reserve a special context ("e164") for the purpose of creating a global Peering System in which anyone can exchange access to E.164 numbers.

Because DUNDi does not advertise routes but instead simply answers queries, it is not necessary to explicitly publish changes in the peering system with respect to new peering agreements, additions, removals or relocation of answers within the system (i.e. number portability does not require any additional work to implement).

An Internet Draft for the DUNDi specification is available at the official DUNDi web site at http://www.dundi.com.

Russell's PC: "Hey, do you know where I can find 12565551212?"
Sara's PC: "I'm not sure. Hey Andy's PC, do you know where Russell's PC can find 12565551212?"
Andy's PC: "Not off the top of my head. Hey Beth's PC, do you know where Russell's PC can find 12565551212?"
Beth's PC: "Yes, you can find it at sip:12565551212@hsv.com."
Andy's PC: "Beth's PC says you can find it at sip:12565551212@hsv.com."
Sara's PC: "Beth's PC says you can find it at sip:12565551212@hsv.com."
Russell PC: "Thanks!"
Tom's PC: "Hey, do you know where I can find 12565551212?"
Sara's PC: "Beth's PC just said you can find it at sip:12565551212@hsv.com."
Tom's PC: "Thanks!"

The Non Technical Portion: The GPA

The obvious problems with trying to build a global system with which people can call one another via their real E.164 phone numbers for free

are how to preserve the accuracy of the data, and how to prevent the use of free routes for VoIP spam and other undesired abuses. With ENUM, these problems are solved by the creation of an entity to oversee the routes, and through a SIP security model, also with an entity in charge of issuing certificates and authenticity.

With DUNDi, we replace both these entities with a document, the General Peering Agreement, or GPA. That is to say, rather than seek accuracy and accountability through a company or government organization, we create a document, establishing the rules of peering within the “e164” context of the DUNDi protocol with respect to the publishing and use of routes. The General Peering Agreement, executed by all members of the Trust System, places the following responsibilities on its participants:

1. Participants will only publish numbers for which they are in good faith authorized to represent, in accordance with a specific policy.
2. Participants and their customers may only use the routes in accordance with a (brief) acceptable use policy, including those implied by the flags provided on a route. For example, if a route says it may not be used for unsolicited calls, a Participant or its customer must not use that route for making any kind of unsolicited call.
3. Participants must only populate the weight field in accordance with the agreement, so that the priority of answers may be established (i.e. 0-99 is subscriber-specified only, 100-199 is only used by the subscriber's direct service provider, etc.).
4. Participants must only provide routes which do not require a fee to connect to the advertised service. Other fees are generally permissible as outlined in the agreement text.
5. Participants must provide valid or empty caller identification only (i.e caller identification is optional, but if supplied must be accurate).
6. Participants agree to give all other members of the Trust System legal standing with respect to the enforcement of the GPA with any other member.
7. Participants agree to a specific dispute resolution strategy.
8. Participants must provide valid contact information via the DUNDi protocol.

The critical advantage of distribution in a nontechnical environment is that the GPA replaces a monopoly entity (which is required under the ENUM architecture) with a static document, where the terms and

conditions for peering are well understood and not subject to the control of a single entity. Further, because the GPA places all participants on an equal standing, there is no artificial lower bound to the cost of service. In contrast, under the current DNS system for example, the cost that Verisign® is permitted to charge for their registration service through third parties forms a lower limit to the charge assessed to anyone seeking to register a domain name. Similarly, the charges that incumbent local exchange carriers are permitted to levy on long distance carriers for termination into their networks form a lower bound for the cost of a traditional long distance PSTN call.

The DUNDi protocol and the GPA are designed to facilitate rapid location of invalid routes and enforcement of the agreement. The source identifier of the node which originally provided the route is included in the response as well as the node which finally supplied the response to the entity making the original request. Contact information can be obtained on any node on the system, and the GPA gives any member of the Peering System legal standing to enforce the agreement, against any other member, regardless of whether those two particular members had directly executed the GPA.

The full terms and conditions of the GPA are available at the official DUNDi web site at http://www.dundi.com . In order to join the Peering System it is only necessary to execute the GPA with another member of the Peering System and then perform the technical requirements of building the trust relationship.

DUNDi and Asterisk:

Although DUNDi is a public specification intended to be implemented by anyone, Asterisk already contains an implementation which may be freely downloaded at http://www.asterisk.org.