

FERN: A Unifying Framework For Name Resolution Across Heterogeneous Architectures

Spencer Sevilla*
spencer@soe.ucsc.edu

* UC Santa Cruz
Santa Cruz, CA

Priya Mahadevan†
priya.mahadevan@parc.com

† Palo Alto Research Center
Palo Alto, CA

J.J. Garcia-Luna-Aceves*†
jj@soe.ucsc.edu

Abstract—A key problem in all name resolution protocols today is that no one protocol performs well across all network architectures. In addition, DNS, the most widespread solution today, depends on a static and connected network layer and cannot support dynamic wireless networks. We introduce FERN (Federated Extensible Resolution of Names), the first framework designed to enable efficient name resolution across heterogeneous name resolution systems operating in dynamic or static networks. FERN organizes nodes into name resolution groups and allows these groups to perform name resolution independently in different ways. FERN arranges these name resolution groups into a hierarchy and allows these groups to communicate efficiently, discover each other’s presence, and resolve each other’s names. We demonstrate the flexibility and interoperability of FERN by deploying and evaluating it across heterogeneous environments, including a MANET, an infrastructure-based wireless network, and the Internet. We show that FERN performs at least as well as DNS, and yet extends name resolution to networks in which DNS is inadequate.

I. INTRODUCTION

Service discovery and name resolution are vital operations in any network. Users and applications use text-strings (e.g., URLs), rather than network addresses, to indicate the content or services they require, and these names must then be mapped to network addresses before communication is possible. This name resolution requirement applies to today’s and future networks and the Internet at large.

Unfortunately, current approaches to name resolution are unable to support this future networking environment because of two constraints. First, no one protocol has been devised that works well across all types of networks. Second, these protocols have not been designed to interoperate with one another. For example, consider the case in which a user accidentally leaves her laptop at home and wishes to access it from the office. The laptop most likely uses multicast DNS (mDNS) [1] to name itself on the home network, but the user has no way of resolving this name outside of that environment, and cannot discover the laptop. As another example, nodes in a MANET may use a distributed protocol to resolve each other’s names, but there is no protocol for them to extend this resolution to the Internet through the domain name system (DNS), despite

the presence of a network-layer gateway bridging the MANET to the Internet.

To address the above limitations of existing name resolution approaches, we introduce the Federated Extensible Resolution of Names (FERN). To the best of our knowledge, FERN is the first system that provides a low-overhead unifying framework for different name-discovery protocols to interoperate. FERN accomplishes this by organizing nodes into name resolution groups and allowing each group to perform internal name resolution in a manner that works best for that group. FERN also defines a technique for requests to propagate across different groups, and organizes these groups into a hierarchy to support deterministic request-forwarding and ensure reachability across these different name resolution groups. Furthermore, FERN performs these tasks while preserving high scalability and backwards compatibility, and can be deployed incrementally alongside existing protocols such as DNS and mDNS.

Section II provides an overview of prior work in name resolution and service discovery for different types of networks. Section III specifies the operation of FERN within a group and across groups. Section IV discusses how to interface FERN with existing name resolution protocols. Section V describes our testbed implementation and provides preliminary results collected with this system. Section VI concludes the paper.

II. RELATED WORK

Prior work in name resolution and service discovery can be loosely divided into client-server systems, peer-to-peer systems, or systems based on overlay networks. Additionally, there exists work on hybrid systems employing more than one of these architectures and on interoperability between these systems.

A. Client-Server Systems

The most prominent approach to name resolution today is the DNS [2]. DNS relies on a hierarchy of servers that must be configured to forward a name-request to the appropriate server, which then resolves that name-request to an IP address. Through the use of this hierarchy, load-balancing “secondary” servers, and caching, the DNS provides name resolution for the entire Internet today. However, this scalability comes at a price. First, the DNS relies completely on these servers: if the

This work was sponsored in part by the Palo Alto Research Center, the Baskin Chair of Computer Engineering at UC Santa Cruz, and the NASA Ames Research Center under Grant No. NAS2-03144.TO.030.10.MD.D.

authoritative DNS server for the subdomain `example.com` is down, overloaded, or configured incorrectly, then all DNS lookups for `*.example.com` fail and `www.example.com` is not reachable, regardless of the state of the web server itself. Refer to [3], [4] for recent examples of DNS outages that affected millions of users. Second, given that the DNS relies on hosts to configure their IP address with their DNS server using out-of-bound communication, it is a static system that cannot support dynamic networks. Dynamic DNS [5] seeks to alleviate these limitations by specifying an UPDATE record type; however, it still requires that (1) the host knows the IP address of its authoritative DNS server *a priori*, and (2) the host successfully sends an update to the authoritative server every single time its IP address changes.

B. Peer-to-Peer Systems

Peer-to-peer systems such as mDNS [1], SSDP [6], and SLP [7] do not require a central server to operate. As a result, these systems require minimal configuration (hence the name “Zeroconf”) and are well suited to dynamic environments where hosts come up, go down, and change IP addresses frequently, such as home networks configured with DHCP [8] or AutoIP [9]. Unfortunately, all peer-to-peer systems currently share a heavy reliance on IP multicast to propagate both name-requests and service announcements through the entire network. As a result they suffer from relatively high latency and cannot scale, which restricts these protocols to LANs, where internal names are denoted by the top level domain (TLD) `.local`.

C. Overlay Networks

Several publications [10]–[14] have discussed deploying DNS over an overlay network that uses a distributed hash table (DHT) to reduce the load on individual servers and thus provide higher scalability and better fault tolerance. These papers target the traditional DNS system in the Internet, and thus focus on planet-level scalability. They present mixed results on latency, but note that DHTs serve to decouple the physical location of an entry from its logical location. This architecture helps with load-balancing, removes hot spots and bottlenecks in the hierarchy, and creates a system that is orders of magnitude harder to attack. These benefits are typically achieved by enforcing a flat namespace, where all records in the system are stored as equal objects in one giant DHT. Unfortunately, these approaches rely on a network environment in which the nodes of the overlay are static and available with high uptime, the topology is connected, and links have plenty of bandwidth. The performance of DHTs degrades significantly in dynamic networks as a result of excessive overhead resulting from topology-independent overlay addresses, link failures, and node mobility.

D. Hybrid Systems

There are a number of hybrid approaches to name resolution that attempt to combine the architectures described above. SLP, for example, introduced the concept of an optional “Directory

Agent” (DA) that nodes in a network must contact first if it is present. Kozat et. al. [15], [16] bring this concept to the case of MANETs by proposing a virtual backbone of “Service Broker Nodes” (SBNs) that form a dominating set in a MANET and proactively maintain routes through the MANET to each other. These proposals attempt to increase scalability by only allowing a select subset of nodes to query the entire network, and requiring that other nodes communicate with their closest directory node. However, these approaches all share the same drawback, which is that communication between directory nodes is unstructured and accomplished by flooding a name request to all other directory nodes, which scales as poorly as Section II-B. MDHT [17] addresses this issue by proposing a hierarchy of DHTs, but cannot scale to large numbers of records because it requires the top-level DHT to contain every record in the system.

E. Interoperability

A unifying problem of every approach described above is that every node must be a member of the protocol for name resolution to occur. This lack of interoperability means that the two protocols cannot talk to each other, even though mDNS might be best for home networks and DNS might be best for the Internet. Currently, support for multiple protocols is accomplished by designating some top level domains or TLDs (such as `.local`) for certain protocols and having the node generating a request use the TLD to decide which protocol should be used. A few approaches [18]–[20] have been published on interoperability between multiple resource-discovery protocols, but these works have been limited to developing higher-layer application programming interfaces (APIs) that mask implementation differences between protocols that already share the same basic architecture, such as SSDP and SLP.

Plutarch [21] proposes an architecture for interoperability across different network architectures, both for routing and name resolution. Instead of requiring all networks to use the same protocol, Plutarch divides networks into *contexts* and proposes the use of interstitial functions to translate between contexts, similar to the way network address translation (NAT) is implemented today. Though Plutarch provides a model for interfacing radically different network architectures, it effectively leaves the implementation of these interstitial functions “to the reader.” Plutarch also raises an important question about how the different contexts become aware of each other: it proposes using a gossip protocol to disseminate this information, yet this protocol might cause issues of scalability and coherency if the number of separate contexts becomes too high or if entire contexts exhibit a high degree of mobility.

III. FERN: A UNIFIED FRAMEWORK FOR NAME RESOLUTION

FERN builds on previous name resolution approaches by providing a framework for interoperability among different name resolution protocols, such as the ones described in the previous section. FERN organizes nodes using a common

TABLE I
FERN NAME RESOLUTION GROUP API

| Function Prototype | Comments |
|---|---|
| int (0 = success) joinGroup(args) | args varies as a group-specific parameter |
| int (0 = success) leaveGroup() | Groups must also support ungraceful departures |
| int (0 = success) registerName(name) | name is not fully-qualified (i.e. just "printer") |
| network_address resolveName(name) | name is not fully-qualified (i.e. just "printer") |
| network_address getParent() | assumes the parent group can be reached at this address:udp53 |
| network_address getChild(name) | same as above, but returns (null) if it has no child with this name |
| int (0 = success) registerChild(name) | name is not fully-qualified (i.e. just "lab_3") |
| int (0 = success) deregisterChild(name) | name is not fully-qualified (i.e. just "lab_3") |

name resolution scheme into separate *Name Resolution Groups* (NRGs), and specifies a protocol for NRG intercommunication. FERN then organizes the NRGs into a hierarchy. The primary motivation for organizing nodes into NRGs is to: (a) separate nodes that use different name resolution schemes; and (b) reflect the natural groupings that appear in the underlay network (i.e., subnets), logical hierarchy (i.e., org charts), and users themselves (i.e., social groups). FERN defines a set of operations that a group must support, but explicitly does not define the implementation of these operations. FERN also assumes that all nodes in an NRG are able to exchange messages at the application-layer, and does not assume that all the nodes run a specific network-level protocol.

Given that FERN does not specify the internal mechanics of an NRG, it can be compared to Plutarch, in that FERN's NRGs are equivalent to Plutarch's contexts and intergroup communication is equivalent to Plutarch's interstitial function. However, FERN differs from Plutarch in that it organizes groups into a naming hierarchy instead of a gossip protocol. This hierarchy is to ensure that: (1) the system resolves names deterministically, (2) name requests do not traverse NRGs unnecessarily, and (3) scalability is preserved by enforcing an upper-bound on the number of other NRGs any one group must know.

A. NRG Responsibilities

The first responsibility of a FERN NRG is that every node in a group must be able to resolve names for which the group is responsible. Similar to DNS, NRGs are responsible for names that end in the NRG's fully-qualified name. For example, an NRG named `lab_3` is responsible for queries such as `printer.lab_3` or `johns_pc.lab_3`. To facilitate these responsibilities, NRGs must provide a way for its members to: (1) register names, (2) resolve names, (3) join the NRG, and (4) leave the NRG.

The second responsibility of a FERN NRG is that it must forward queries for which it is not responsible. This requirement is accomplished by organizing the NRGs into a naming hierarchy and allowing NRGs themselves to be

members of another NRG. In this situation, the *children* of an NRG are its members, and the *parent* of an NRG is the NRG of which it is a member. This relationship is denoted using the same dot-notation as in DNS. For example, an NRG with the name `lab_3.parc.usa` is a member of the NRG `parc.usa`, which is itself a member of `usa`. For clarification and brevity, in the remainder of this paper, the *shortname* of this NRG is `lab_3`, whereas its *fullname* is `lab_3.parc.usa`. This child-parent relationship between NRGs creates a *name resolution tree* (NRT) as in the DNS, with the *root NRG* "*P*" at the top, and this tree powers the forwarding of requests among NRGs. Though the NRG `lab_3.parc.usa` in our example is not responsible for the query `x.ccrq.ucsc.usa`, it can forward the request up the NRT to `usa`, and then down to `ccrg`. Furthermore, `lab_3` can perform this task without knowing the network address of `ccrg` itself; all NRG `lab_3` needs to know is how to contact its children NRGs and parent NRG. We express FERN programmatically as a pseudocode API in Table I, and formally define the set of rules in List 1.

List 1: FERN NRG Rules

- 1) NRG *X* has at most one parent NRG *Y* in the FERN NRT, and $fullname_X = shortname_X.fullname_Y$.
 - 2) NRG *X* can have several child NRGs in the FERN NRT, and each of these children has the full name $childname.fullname_X$.
 - 3) NRG *X* must be able to communicate with its parent and children NRGs in the FERN NRT.
 - 4) NRG *X* must know the addresses of all its ancestor NRGs in the FERN NRT.
 - 5) NRG *X* is responsible for directly answering all queries that end in $fullname_X$.
 - 6) NRG *X* must forward queries to the best match of NRG possible, adhering to the caching rules in List 2.
 - 7) NRG *X* must return an error for a query that it cannot answer or forward.
-

B. NRG Communication

For the sake of interoperability with DNS, we have chosen to use the traditional DNS record format (A, CNAME, ...) and port (UDP 53). This choice means that to support request-forwarding along a branch in the NRT, all an NRG has to store is the network address of the other NRG. This results in an exceedingly simple interstitial function, and means that intergroup resolution through the entire hierarchy can be supported by simple recursion.

Figure 1 shows an example of request forwarding in FERN. Here, the NRG `example` uses a server, `subgroup1` uses request-flooding, `subgroup2` uses a DHT, and a smartphone in `subgroup1` wants the address of the printer in `subgroup2`. First, the request is flooded through `subgroup1` until it reaches a node that can communicate with `example`. Next, the interstitial function of FERN is used to forward the request up to `example`, and then again to forward the request down to `subgroup2`. Lastly, `subgroup2`




Fig. 1. Request-Forwarding Across Groups

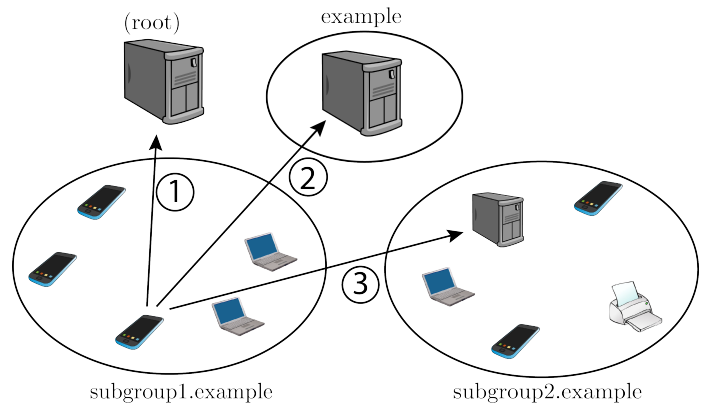


Fig. 2. Request-Forwarding In DNS

uses its DHT to resolve the address of the printer. This behavior is contrasted with Figure 2, which shows the same name resolved iteratively in DNS. As Figure 2 illustrates, DNS requires that (1) each name group be supported by an authoritative name server, (2) resolution starts at the root server and descends the NRT, and (3) servers support iterative resolution, where the resolver communicates with each name-server in turn.

C. Internal Group Policies

FERN places no constraints on the number of services or names an individual node may register, the nature of these services, or the number of NRGs of which a node may be a member simultaneously. It is left to individual NRGs to implement and enforce rules such as restricting group membership to certain nodes or restricting the names that a particular node may register. NRGs may choose to adopt and enforce certain naming conventions (similar to the mDNS service registry), and these conventions may even standardize across different NRGs; however, this standardization is outside the scope of FERN.

FERN treats group security the same way. NRGs may choose to use encryption, MAC addresses, or other out-of-bound information to authenticate, authorize, and verify their members and names. They may also decide to use name resolution to enforce other security policies, such as only allowing certain nodes to resolve the address of certain services. However, the administration and implementation of these policies are left to the individual NRG, not the entire framework.

D. Bootstrapping Group Membership

For a node to join an NRG with the `joinGroup(args)` operation in Table I, it must already know (1) the group architecture, (2) any `args` the group requires, and (3) to whom to send this information. Though the mechanics and specifics of joining an NRG should be handled by the NRG itself, the process of group discovery and acquiring the information listed above must be standardized, because it is a process that exists outside of any individual NRG and may interact with other

protocols and systems. There are several protocols (e.g. DHCP [8] and AutoIP [9]) used to help nodes join a network by supporting discovery, authentication, and address acquisition. They also bootstrap DNS resolution by providing hosts with the address of a local DNS server to be used. Thus, FERN can extend these existing protocols by defining an extra FERN record to be passed to a node when it joins the network. This record should contain: (1) the fullname of the NRG, (2) the structure of the NRG, (3) any group-specific arguments, and (4) a fallback network address to be used as a local DNS server if the node does not recognize the value in (2) or is FERN-unaware.

E. Caching

Caching name responses and intermediate name referrals significantly reduces latency and overall network load. It changes the system performance and may even result in different behavior. In the DNS, caching benefits stem primarily from reducing the number of round-trips a query takes. In FERN, benefits come from “short-circuiting” the group hierarchy. For instance, in the same example of Figure 1, if `subgroup1` has a cached network address for `subgroup2`, it may skip the group `example` entirely. Caching in FERN is enabled by allowing groups to append an A record for their group itself when they answer a query or recursively return the answer to a query. Thus, if a request originates at group *A* and traverses groups *B*, *C*, and *D* before finishing at *E*, the requesting node could end up caching the network address of groups *B* through *E* if these groups elect to append their network address to the response. Additionally, intermediate groups may also read these records, so in this example group *C* could also learn the network address of groups *D* and *E*.

1) *Hybrid Behavior*: Caching leads to behavior that closely resembles a hybrid system. In our example above, the bottom groups use architectures better-suited for dynamic networks. The first time a node in one of these groups attempts to resolve a name outside of the group, it must first call `getParent` and use the group to resolve the address of its parent. However, the resolving node may then cache this address and send all future requests directly to its parent group without needing

to re-resolve its address. This behavior is remarkably similar to the hybrid approaches described in Section II, where local requests stay local and system-wide requests get forwarded to the appropriate SBN or DA, yet FERN enables this behavior without the added protocol complexity of specifying how it should be done, figuring out what constitutes a local request, or forcing that system on all network scenarios. This behavior can also be compared to current systems, where requests are either multicasted over mDNS or sent to a local DNS server based on the TLD of the name-request. FERN exhibits very similar behavior, yet accomplishes this without fragmenting the namespace.

2) *Caching Up The Name Resolution Tree*: In DNS, caching can only occur *down* the tree, and this improves performance by reducing the load on top-level name-servers and the number of referrals. However, in FERN, caching can also occur *up* the tree. While this is a feature aimed at improving performance, it could make FERN perform much like DNS if nodes were to use the address of the NRT root to resolve names. For example, consider the case where a node `node1.subgroup1.example.usa` needs to resolve the name `node2.subgroup2.example2.uk`, and caching is enabled for any NRG in the FERN NRT. Since the root of the NRT is the closest common ancestor between the node and the name that must be resolved, the node caches the network address of the NRT root once resolution is complete. After that, anytime the same node needs to resolve a name outside of `usa`, the closest-matching group will always be the NRT root, and the node will contact it directly. The FERN caching rules stated in List 2 eliminate this problem.

List 2: FERN Caching Rules

- 1) A node in NRG X may cache the address of nodes of NRGs for which NRG X has a branch in the FERN NRT.
 - 2) A node in NRG X *may not* cache addresses of nodes in NRGs that are closer to the root of the FERN NRT.
 - 3) A node in NRG X may cache the addresses of nodes in NRGs that are at the same level of NRG X in the NRT, or further down the NRT.
-

The FERN caching rules significantly reduce the load on nodes in NRGs that are higher in the hierarchy and serve to create a much more distributed system. To revisit the previous example, once `node1.subgroup1.example.usa` resolves `node2.subgroup2.example2.uk`, the group `uk` is cached only by two groups: the root and `usa`. Not only does this help reduce traffic on the top name-servers, it also helps provide cached information to other nodes. Now, if the same node wishes to resolve a name in the TLD `china`, rather than query the root directly (and get a direct response) it must go up the tree through the group `usa`. This behavior ensures that now `usa` is on the return-path and has the opportunity to cache the network address for `china`, which further reduces traffic on the root group, since all subsequent requests from nodes in `usa` for nodes in `china` would be able to take advantage

of the cache-hit in `usa`.

Together, the FERN rules in Lists 1 and 2 provide interoperability across different architectures while limiting the amount of information that any one NRG must maintain.

F. Fault Tolerance and Resilience

The FERN process of forwarding requests up and then down the NRT also affects the fault tolerance and resilience of the system. In DNS, if a node is unable to contact the root server it is unable to perform any resolution, as shown in Figure 2. This behavior makes the root server an attractive target for attackers, and also restricts the usefulness of DNS to nodes that can access a root name server, as opposed to nodes in a private network or MANET. Conversely, FERN requests only travel up the NRT as far as necessary. Thus, in Figure 3, the only queries that would reach the root NRG are requests from NRG `ucsc` to NRG `parc` or vice-versa. All other traffic stays within either NRG, and thus would function normally independently of their ability to access the root NRG.

1) *Internal Resolution*: By forwarding queries in the manner described above, FERN reduces the reliance on the top-level NRGs of the NRT and improves resilience among lower-level NRGs. If a root or TLD server fails, or if a NRG is cut off from these servers due to a network partition, internal resolution is unaffected. Thus, name resolution in FERN is much more distributed. Ideally, if there exists an active route between two hosts, they should be able to resolve each other's names and communicate. Conversely, if no route exists between the hosts, then name resolution is unimportant because even in the event of successful resolution, no communication can occur.

2) *Intermediate Failure Points*: Consider the case where the NRG `parc` fails. In DNS, all nodes (including nodes within the domain `parc`) would be unable to resolve any names below domain `parc` in the tree, but are able to resolve all other names. Conversely, in FERN, requests that stay inside NRGs `cs1` or `is1` would still succeed, but none of the nodes in these NRGs would be able to resolve any names outside of `parc`, unless the NRT is modified to reflect the failure that took place.

FERN addresses this problem by allowing nodes to cache the network address of other nodes in their ancestor NRGs all the way up to the root of the NRT. Note that, in accordance with List 2, these network addresses cannot be used for the forwarding of requests. The addresses are used solely for fault tolerance. An NRG may use these addresses to forward requests to its grandparent if and only if its parent is unresponsive. With this rule in place, FERN may often do better than DNS (by preserving internal resolution when possible) but it never does worse, since it effectively reduces to DNS when intermediate NRGs fail.

To reduce the risk of node failure, NRGs may also choose to replicate records across $K > 1$ separate nodes. Choosing a proper value for K depends heavily on the underlying network. In the Internet case, the DNS itself shows that small values of K are sufficient. For example, over 80% of DNS entries were supported by just one or two name servers in

2004 [11]. In other network scenarios as MANETs, $K = 1$ might be completely acceptable if the only node bridging name-requests is also the only node able to perform network address translation (in which case its failure also partitions the network). Ideally, K should be sufficiently large so that name resolution reflects network connectivity.

G. FERN Correctness

The FERN rules in Lists 1 and 2 allow us to formally prove that requests processed in FERN deterministically terminate, do not loop, and are resolved correctly. The following proof focuses solely on loops resulting from the misconfiguration of FERN groups, and does not consider underlay network errors or malicious behavior.

Theorem 1: FERN request forwarding is loop-free.

Proof: Assume that there exists a request-forwarding loop among nodes using FERN. Given that NRGs are organized as a tree, which is acyclic, the existence of a request-forwarding loop necessarily implies that an NRG i must forward the request to another NRG k that is not its parent or child in the FERN naming tree. However, according to Rule 6 in List 1 and List 2, an NRG that forwards a request must do so to either (a) its parent or (b) one of its children. Given that the NRT is acyclic, the request-forwarding loop must occur as a result of misconfiguring either the parent NRG or one of its children.

According to Rules 1 and 3, an NRG cannot mistake its parent and hence NRG i cannot consider NRG k to be its parent NRG mistakenly. This means that the request-forwarding loop must result from the misconfiguration of a child NRG, i.e., assuming that NRG k is a child NRG when in fact it is not. For this to be the case, NRG i must know how to contact NRG k , and NRG i can acquire this knowledge only through the registration process. Since the registration process is always initiated by the child NRG, it is impossible for NRG i to mistakenly assume that NRG k is its child when it is in fact not the case. This completes the proof.

Though Theorem 1 may appear trivial, referral loops are possible in the DNS, and these loops significantly impact network performance. Jung et. al. [22] observe that a very small portion (3%) of requests to misconfigured name-servers result in referral loops, and these requests generate over 12% of all DNS packets, on average retrying each query over ten times before giving up. Theorem 2 below serves as a general proof of correctness for FERN requests.

Proposition 1: Any two NRGs have a common ancestor.

Proof: The proof is immediate from the fact that the root group “/” is the parent of all TLD groups, hence it is a common ancestor of every group.

Theorem 2: FERN name resolution is provably correct

Proof: Without loss of generality, assume that some node in group X wishes to resolve the name of a node in group

Y . From Proposition 1, it follows that groups X and Y must have a common ancestor; call that ancestor group Z . By Rule 3 and Theorem 1, the request originating in group X can be forwarded up the tree until it reaches group Z . Again by Rule 3 and Theorem 1, once the request reaches group Z , it is forwarded down the tree until it reaches group Y , which resolves the request. Hence, any node in any group is always able to resolve the name of any node in any other group.

IV. DEPLOYMENT

The previous sections define the FERN framework and explain why FERN explicitly does not specify internal group considerations. However, these considerations still have a large impact on the overall system performance, and thus merit discussion.

A. Interfacing with Existing Protocols

Since internal group resolution can take many forms, FERN can be used to bridge all existing name resolution protocols today without modifying them. Supporting mDNS is trivial, and can be accomplished by simply appending `.local` to the end of a name-request before sending it to a mDNS daemon. DNS integration is equally straightforward, though it comes with one caveat: given that the resolution of DNS queries start at the root, if the DNS is used to power a FERN NRG, the NRG must be the highest group in the FERN NRT; otherwise, unnecessary request-forwarding and group traversal can occur. However, because the DNS is already well-established, we believe that FERN NRGs could exist “underneath” the current DNS hierarchy, using the DNS for Internet resolution, while still supporting other networks where the DNS is not appropriate.

B. Internal Group Communication

The best choice for internal group communication depends on both the underlying network topology and the number of nodes in the NRG. Though an NRG may specify that only a certain number of nodes may join, the number of nodes in an NRG is determined primarily by external factors, which in turn determine group communication. These external factors could be logical (the number of people in an organization), hierarchical (an org chart), or based on the underlying network topology (e.g., nodes in a MANET).

In the case of the Internet, a connected underlay network with static addresses, the client-server architecture has been shown by the DNS to be efficient, scalable, and provides an attractive first choice. For fully-connected networks with dynamic network addresses (such as an internal subnet or home network), a DHT may be a better choice for both robustness and dynamic updating.

C. NRT Height

The current DNS hierarchy is relatively shallow, with a typical height of three or four levels, but is almost exclusively limited to naming Internet servers. We believe that a full FERN NRT would be allowed to have more levels, because part of




Fig. 3. Testbed FERN Name Resolution Tree

the intent of FERN is to expand name resolution to devices in different network environments. As described above, the addition of NRGs in the NRT could be the result of several logical or organizational factors, as well as underlay network concerns (such as bridging resolution across two MANETs). In Section V below, we investigate the performance overhead of adding extra groups to the NRT.

V. TESTBED DEVELOPMENT AND EVALUATION

The main contribution of FERN is a low-overhead modular framework that enables different systems to interconnect. This modularity makes it hard to compare its performance to existing systems, because any results collected are heavily dependent on the individual groups and their structure. Additionally, all existing approaches today, such as the DNS or UPnP, can function as a FERN group, so they can be considered a subset of FERN functionality.

To confirm the arguments we made in the above sections, and to collect more information about FERN performance, we built a FERN daemon in Java and used it to support three different internal NRG protocols. These protocols, Chord, Server, and Flood, are detailed below, though it is important to note that FERN can support many other forms of communication. We then deployed this daemon on eleven separate nodes located at UCSC, PARC, and UCSB. At each campus we deployed one server connected to the Internet to handle inter-campus queries, at PARC we also used three laptops connected to the campus wireless network, and at UCSC we set up a MANET of four smartphones (and one laptop) running OLSR, with the laptop also connected to the LAN via ethernet.

Figure 4 shows how we configured the nodes into the NRT shown in Figure 3. The NRT, network topology, and choice of name resolution protocols was designed to emulate what we envision a wide scale FERN system might look like, with significantly more mobility and dynamic behavior in the lower layers of the NRT than in the upper layers. We ran several name resolution tests using this testbed to evaluate our system and summarize our results below.

Chord: The Chord NRG requires that every node in the group be a member of the Chord DHT [23]. Once a member of Chord, nodes use key-value pairings to map their IP address to any names (nodes or groups) they are responsible for. We use Chord at two levels, shown in Figures 3 and 4, to illustrate its use in different contexts: in the group *ccrg*, we use Chord on top of OLSR to simulate a MANET environment where routing and discovery are done through a DHT [24], [25]. We also use Chord at the highest level, *global*, to emulate




Fig. 4. Testbed Implementation

an Internet-wide DHT of static name-servers, as discussed in Section II-C.

Server: The Server NRG very closely resembles the current DNS. The group consists of one node acting as a name server, and all other nodes in the group are clients of that server. Names and services are registered with this server and resolved by querying the server.

Flood: The Flood NRG is completely decentralized and emulates mDNS. When nodes join a Flood NRG, they subscribe to a multicast address and port combination. Rather than publishing or announcing any names or services, they listen to the multicast address and respond in unicast to any requests for a name or service that they can provide. Conversely, nodes resolve a name or service by simply sending a request to that multicast address. This approach is used at the lowest levels of the hierarchy for two reasons: first, to emulate a MANET running a reactive routing protocol, such as AODV. Second, to illustrate the current role of mDNS and UPnP in today's home networking environment, where discovery is limited to the reach of the multicast tree.

A. Latency

We evaluate latency in our testbed by disabling caching and using `tcpdump` to time 50 requests internal to each NRG, and 50 requests (each way) between every pair of NRGs. In analyzing these results, we find that the latency of an inter-group request can be expressed as the sum of two main components: the latency of internal group resolution to determine the address of the next hop, and the latency of sending a request or response from one group to another.

1) *Internal Group Latency:* Internal group latency varies tremendously by group, and can dominate the end-to-end latency. As Table II shows, in server architectures (including DNS) where the address of the next hop is already known and exists in a table, this time is typically under 5ms, whereas other groups based on DHTs or multicast take significantly longer. Additionally, the groups *csl* and *global* merit additional discussion. In *csl*, requests are supported by IP multicast, which results in both a high mean and a significant variance ($\mu = 451$ $\sigma = 173$), which are problems also seen with mDNS. The performance of IP multicast varies

significantly with the network topology, traffic load, and even the implementation of the multicast tree. Thus, the results we present in Table II do not necessarily reflect performance in other LANs, though they do highlight the performance problems with using multicast for name resolution.

In our topology, the Chord used for `global` contains three nodes, located at UCSC, UCSB, and PARC. When querying this Chord, we observed latencies ranging from 80 to several hundred milliseconds, depending on the physical location of the data. Though this serves as a proof-of-concept, the low number of nodes and their geographical proximity does not accurately reflect the topology of a large-scale Chord. To better reflect a large-scale geographically diverse Chord deployment, we choose to present and compare numbers from the original Chord paper [23] instead, where the authors measure a large 190-node Chord deployed across the entire USA. This topology is closer to what we would expect to see, and therefore their results are a better indicator of latency in this scenario.

2) *Inter-Group Latency*: Inter-group latency is determined by two relatively static factors: the physical distance a message must travel (from one node in one group to another node in another group), and the number of times it travels between groups. We compare inter-group latency in FERN to DNS values to get a more accurate understanding of how group hierarchy and structure affects latency. We compare this particular metric to DNS latency because DNS does not have an equivalent internal group component.

Jung et. al. [22] conclude that the number of DNS referrals has a strong effect on the latency of a DNS request. Unfortunately, a DNS request with N referrals could potentially result in $2N$ FERN referrals, since it must both climb and descend the tree. However, DNS requests (especially to root and TLD servers) are generally performed iteratively. This distinction is important because the latency of an iterative request with two referrals corresponds to the latency of two complete round-trips from a local DNS server to an authoritative name-server that may or may not be close. This problem is highlighted by their KAIST dataset analysis, where they identify a latency “bump” that they correlate with round-trips traversing the Pacific Ocean. Iterative resolution makes this problem worse, since it results in potentially N transpacific round-trips.

In contrast, FERN minimizes inter-group latency by (1) requiring groups to resolve name-requests recursively and (2) organizing nodes in a hierarchy that reflects physical proximity (i.e. assigning countries or physical regions to TLDs). These two concepts combine to forward these requests to their destination and ensure that requests only traverse a particular long-haul link (i.e. the Pacific Ocean) once. With this feature in place, we find that the latency overhead of adding another logical group to the hierarchy is minimal: although it is unlikely that the network address of the group is directly on the route to the target group, with these rules it should be relatively close, and in our tests we find this overhead to typically be under 10ms.

TABLE II
MEAN LATENCY OF INTERNAL GROUP RESOLUTION

| name | ccrg | ucsc | global | parc | csl |
|---------|--------|--------|--------|--------|--------------|
| NRG | Chord | Server | Chord | Server | Flood |
| network | OLSR | LAN | WAN | LAN | IP Multicast |
| mean | 106ms | 2ms | 180ms* | 3.5ms | 451ms* |
| stdev | 21.8ms | 0.7ms | 60ms* | 0.8ms | 173ms* |

TABLE III
SUCCESSFUL NAME-REQUESTS WITH A FAILURE IN THE TOPOLOGY

| Group: | global | ucsc | ccrg | Link: | parc-global | csl-parc |
|--------|--------|--------|--------|-------|-------------|----------|
| DNS | 0/110 | 66/110 | 77/110 | DNS | 30/110 | 42/110 |
| FERN | 40/110 | 92/110 | 77/110 | FERN | 50/110 | 54/110 |

B. Fault Tolerance

Given that the topology in Figure 4 consists of 11 nodes, there is a total of 110 different source-destination pairs for a name-request, and in a fully-connected topology they all succeed. However, it is the case that sometimes individual nodes or network links in the system will fail and partition the network into smaller sub-trees, shown in Figure 5 by the dotted lines. We measure fault tolerance in FERN by introducing failures into the system and observing how many name-requests succeed; we present these numbers in Table III. For both DNS and FERN, we assume that a link-failure occurs in the underlay network (meaning that the entire network is partitioned) and a group-failure occurs in the NRT (meaning that network-layer connectivity still exists).

As discussed in Section III-F2, since FERN groups can contact higher-up groups if necessary, FERN always performs at least as well as DNS with regards to fault tolerance. However, Table III shows that FERN usually outperforms DNS when failures occur, especially when the failures occur in higher levels of the hierarchy. This benefit is mainly due to preserving internal connectivity when higher-level groups and links fail, whereas in DNS no-one can resolve names below a failure in the hierarchy.

C. Caching

Because FERN allows for nodes to be a part of multiple NRGs simultaneously, we augmented the NRT in Figure 3 by registering eight “fake” TLDs of the form $\{\text{fake1, fake2, } \dots\}$ and adding one entry, `test`, under each fake TLD. We then ran an experiment where five nodes in `ccrg` each resolved one name under each TLD (for a total of 11 names and 55 requests) in our hierarchy three separate times: once with caching disabled, once with all caching enabled, and once with the caching rules in Section III-E. Table IV shows our results, represented by three metrics: the number of requests sent to the root group, the number of cache-hits that occurred, and the number of cached entries in the system. In addition to the total values of these metrics, we breakdown the total by node, numbering the nodes 1-5 to show the order they issued their requests in.

The results with no caching serve primarily as a baseline for comparison. Turning on all caching reduces the number of root




Fig. 5. Hierarchy Partitioning Due to Link-Failure

TABLE IV
COMPARISON OF AGGREGATED CACHING RULES

| Metric | Rules | 1 | 2 | 3 | 4 | 5 | Total |
|---------------|-------------|----|----|----|----|----|-------|
| Root Requests | No Caching | 10 | 10 | 10 | 10 | 10 | 50 |
| | All Caching | 10 | 9 | 8 | 7 | 6 | 40 |
| | Our Rules | 10 | 0 | 0 | 0 | 0 | 10 |
| Cache Hits | No Caching | 0 | 0 | 0 | 0 | 0 | 0 |
| | All Caching | 0 | 1 | 2 | 3 | 4 | 10 |
| | Our Rules | 0 | 10 | 10 | 10 | 10 | 40 |
| Cache Entries | No Caching | 0 | 0 | 0 | 0 | 0 | 0 |
| | All Caching | 14 | 13 | 13 | 13 | 13 | 66 |
| | Our Rules | 12 | 1 | 1 | 1 | 1 | 16 |

requests, but only slightly: this is because once a node issues a request that results in a root group query, that node caches the address of the root and then sends all subsequent queries to the root group directly. When an individual node queries the root group directly, it is able to cache the responses, but it cannot share this information with the other nodes. Correspondingly, when we turn on our caching rules we see a tremendous decrease in the number of root queries: the group `ucsc` is the only one that can query the root, and once `ucsc` has a cache entry for every other TLD in the system, there is no need for any further root queries. This benefit is also reflected in the total number of cache entries for the system, which is substantially lower because the TLD entries only exist in `ucsc` instead of being duplicated at each node.

VI. CONCLUSIONS

FERN is novel in its ability to interface radically different name resolution architectures. By providing a unifying framework for these protocols, we have laid a foundation for interoperability between future name resolution protocols that are highly specialized for a particular network environment. Furthermore, we show how to seamlessly extend the current DNS to support FERN-style name resolution.

We have examined and highlighted the differences between FERN and DNS. Our results show that the extra group traversals in FERN do not significantly impact latency, and FERN's forcing recursive queries significantly improves performance. We have discussed the effect of caching and confirmed FERN's fault tolerance and ability to handle network partitions. We have also illustrated FERN's scalability and proved that FERN is deterministic and loop-free.

FERN provides a robust framework for name resolution and service discovery. It provides one global namespace and supports both global and local name resolution, yet does

so without the previous constraints on both namespaces. By supporting different name resolution architectures, FERN paves the way for optimization of name resolution protocols for their corresponding networks and serves as an important stepping-stone for interoperability between heterogeneous networks, such as wireless sensor networks and MANETs, home "Internet-of-Things" networks, and the general Internet.

REFERENCES

- [1] S. Cheshire and D. Steinberg, *Zero configuration networking: The definitive guide*. O'Reilly Media, Inc., 2005.
- [2] P. Mockapetris, "RFC 1035: Domain Names - Implementation and Specification," *IETF Standard*, 1987.
- [3] "Godaddy outage takes down millions of sites," <http://techcrunch.com/2012/09/11/godaddy-says-it-wasnt-anonymous-it-wasnt-a-hack-it-wasnt-a-ddos-it-was-internal-network-issues/>.
- [4] "Facebook outage blamed on dns issue," <http://www.itproportal.com/2012/12/12/facebook-outage-blamed-dns-issue/>.
- [5] P. Vixie et al., "RFC 2136: Dynamic Updates in the Domain Name System," pp. 1–26, Mar. 2002.
- [6] Y. Goland et al., "IETF Draft: Simple Service Discovery Protocol," pp. 1–19, Jul. 2009.
- [7] E. Guttman and J. Veizades, "RFC 2608 - Service Location Protocol, Version 2," *IETF standard*, 1999.
- [8] R. Droms, "RFC 2131: Dynamic Host Configuration Protocol," *IETF Standard*, 1997.
- [9] S. Cheshire, B. Aboba, and E. Guttman, "RFC 3927: Dynamic Configuration of IPv4 Link-Local Addresses," *IETF Standard*, 2005.
- [10] R. Cox, A. Muthitacharoen, and R. Morris, "Serving DNS using a peer-to-peer lookup service," *Peer-to-Peer Systems*, pp. 155–165, 2002.
- [11] V. Ramasubramanian and E. Sirer, "The design and implementation of a next generation name service for the Internet," in *Proc. ACM SIGCOMM*, 2004.
- [12] V. Pappas et al., "A comparative study of the DNS design with DHT-based alternatives," *Proc. IEEE INFOCOM*, 2006.
- [13] Y. Song and K. Koyanagi, "Study on a hybrid P2P based DNS," *Proc. IEEE CSAE*, vol. 4, pp. 152–155, 2011.
- [14] T. Vu et al., "Dmap: A shared hosting scheme for dynamic identifier to locator mappings in the global internet," pp. 698–707, 2012.
- [15] U. Kozat and L. Tassiulas, "Network layer support for service discovery in mobile ad hoc networks," in *Proc. IEEE INFOCOM*, 2003.
- [16] —, "Service discovery in mobile ad hoc networks: an overall perspective on architectural choices and network layer support issues," *Ad Hoc Networks*, vol. 2, no. 1, pp. 23–44, Jan. 2004.
- [17] M. D'Ambrosio et al., "MDHT: a hierarchical name resolution service for information-centric networks," pp. 7–12, 2011.
- [18] P. Grace, G. Blair, and S. Samuel, "ReMMoC: A reflective middleware to support mobile client interoperability," in *Proc. DOA*, 2003.
- [19] A. Friday, N. Davies, N. Wallbank, E. Catterall, and S. Pink, "Supporting service discovery, querying and interaction in ubiquitous computing environments," *Wireless Networks*, vol. 10, no. 6, pp. 631–641, 2004.
- [20] J. Allard, V. Chinta, S. Gundala, and G. Richard III, "Jini meets UPnP: an architecture for Jini/UPnP interoperability," in *Proc. Symposium on Applications and the Internet*. IEEE, 2003, pp. 268–275.
- [21] J. Crowcroft et al., "Plutarch: an argument for network pluralism," in *Proc. ACM FDNA*, 2003.
- [22] J. Jung, E. Sit, H. Balakrishnan, and R. Morris, "DNS Performance and the Effectiveness of Caching," *IEEE/ACM Trans. Networking*, vol. 10, no. 5, pp. 589–603, 2002.
- [23] I. Stoica et al., "Chord: a scalable peer-to-peer lookup protocol for Internet applications," *IEEE/ACM Trans. Networking*, vol. 11, no. 1, pp. 17–32, 2003.
- [24] M. Caesar, M. Castro, E. B. Nightingale, G. O'Shea, and A. Rowstron, "Virtual ring routing: network routing inspired by DHTs," *Proc. ACM SIGCOMM*, vol. 36, no. 4, pp. 351–362, 2006.
- [25] D. Sampath and J. Garcia-Luna-Aceves, "Scalable integrated routing using prefix labels and distributed hash tables for MANETs," *Proc. IEEE MASS*, 2009.