

Community-Lab: Architecture of a Community Networking Testbed for the Future Internet

Axel Neumann*, Ivan Vilata*, Xavier León†, Pau Escrich Garcia‡, Leandro Navarro†, Ester López†

*Pangea.org NGO, Barcelona, Spain

neumann@cgws.de, ivan@pangea.org

†Universitat Politècnica de Catalunya, Barcelona, Spain

{xleon, leandro, esterl}@ac.upc.edu

‡guifi.net Community Network, Catalonia, Spain

pau.escrich@guifi.net

Abstract—Community networks are an emerging and successful model for the Future Internet across Europe and far beyond. The CONFINE project complements existing FIRE (Future Internet Research and Experimentation) infrastructures by establishing Community-Lab, a new facility for experimentally-driven research built on the federation of existing community IP networks constituted by more than 20,000 nodes and 20,000 Km of links. In this paper we present the benefits of having such testbed for the research community as well as the improvement and evolution of community networks themselves. This paper presents i) the challenges and requirements for Community-Lab, ii) the resulting testbed architecture, iii) the current state of implementation and iv) the integration of the testbed with existing community networks.

I. INTRODUCTION

In the recent years, the perception of the key role of Internet access for the digital society and the size, growth and scope of the Internet has put under stress its architecture and protocols. The growing Internet penetration at homes, schools, workplaces, and any public or private organization, and the introduction of very demanding and QoS-sensitive applications like distribution of large files, interactive multimedia, real-time audio and video result in an explosive growth in the demand for capacity.

However, the characteristics of the investment needs, cost of ownership, and expected return of investments outside the traditional telecommunication markets in rural areas or in urban areas with low income individuals or special needs not covered by standard products make those areas very unattractive for traditional telecom providers. It has led to the appropriation of these strategic technologies by the communities themselves, by means of self-provision, self-ownership and self-operation. These underserved areas have taken their connectivity needs into their own hands, hence, they developed into “community networks”.

Community networking (also known as bottom-up networking) is an emerging model for the Future Internet across Europe and beyond where communities of citizens build, operate and own open IP-based networks, a key infrastructure for individual and collective digital participation.

These “community networks” (CN) are part of the Internet but differ in several key aspects due to their nature: they are open and transparent as there are no barriers of entry for participation; governance and knowledge is open, with public documentation on every technical and non-technical aspect;

they are self-owned as community members own the network; self-managed with distributed and uncoupled structure that can be open for the participation of community members; self-growing in topology and capacity, based not only on planning but also on immediate demand, as community members can learn and decide to extend the network to reach a new member or provide more capacity. Thus, these networks are not just a path to reach the outside Internet but a wide area network themselves extending the global internet, with internal servers, services, and traffic.

Community-Lab is a testbed being deployed by the CONFINE European Integrated Project¹. This testbed supports experimentally-driven research on community-owned open local IP networks. It integrates and extends three existing community networks: Guifi.net (Catalonia, Spain), FunkFeuer (Vienna, Austria) and AWMN (Athens, Greece); each is in the range of 500 – 20,000 nodes, a greater number of links and even more end-users. These networks are extremely dynamic and diverse. They combine successfully different wireless and wired (optical) link technologies, static, dynamic and ad-hoc routing schemes, and management schemes. They run multiple self-provisioned, experimental and commercial services and applications. This testbed provides researchers with access to these emerging community networks, supporting any stakeholder interested in developing and testing experimental systems and technologies for these open and inter-operable network infrastructures.

The ultimate motivation behind the deployment of such research infrastructure is to support the growth of community networks in terms of their scalability and sustainability by providing the means to conduct experimentally driven research. We expect that this research will directly impact the quality of community networks by improving current shortcomings and providing the basis for a sustainable model of community networking.

Currently, there are several challenges and lines of research to enhance existing community networks. Among others, CNs must consider carefully the impact on the privacy of their users given their openness and degree of decentralization; content distribution should be efficient given the constrained resources by improving network robustness – e.g. routing and cross-layer protocol optimizations; a critical issue is the robustness

¹<http://www.confine-project.eu>

of the nodes participating on the network (e.g. automatic reconfiguration, deployment and failure handling, etc.) because they may sit on remote locations not easily accessible on a day-to-day basis by community owners; and finally, CNs must handle a whole range of heterogeneous physical networks ranging from wireless links to broadband fiber.

However, these challenges are being addressed currently by means of simulations or on very small controlled lab environments because there exists no real testbed to experiment on, which makes the results hard to translate to real scenarios. The Community-Lab testbed being deployed by CONFINE tries to provide such infrastructure for researchers to evaluate their proposal on an environment with the characteristics (for the best and for the worst) of real community networks. The aim of this paper is presenting the architecture of the Community-Lab testbed. The next section presents related work and background. Section III describes the requirements for the testbed and related challenges. Section IV describes the resulting Community-Lab architecture with its structure, directions, current status integration with community networks, and finally a brief discussion.

II. RELATED WORK AND BACKGROUND

There are several initiatives in Europe beyond the three CNs participating in this initial phase of the testbed –guifi.net (Catalonia, Spain), FunkFeuer (Vienna, Austria) and AWMN (Athens, Greece). The Berlin RoofNet project is an experimental mesh network in Berlin consisting of nearly 50 indoor nodes mainly to provide students with Internet access as well as VoIP [1]. In the USA there are similar initiatives as well. Cities like New York, Seattle, Washington or Austin cooperate with public and non-profit organizations to develop and improve free wireless hot-spots within their cities to provide citizens with Internet access. Also, the MIT Roofnet [2] project is a mesh network consisting of about 50 nodes and covers an area of about four square kilometers allowing for small scale and controlled experimentation thanks to the Click Modular router integrated within their firmware.

However, most of these wireless network deployments are aiming to provide a service (Internet connectivity mainly) to users and do not cover the necessity of researchers to perform experimentation on top of them. Non of them are designed to open the testbed facilities to external researchers.

The first initiative of a large-scale testbed to conduct experimental research was introduced by PlanetLab [3]. They currently manage a planetary-scale experimental testbed composed of hundreds of nodes to deploy and evaluate distributed applications and services in a real environment. The key characteristic of this testbed is the concept of resource sharing (given their scarcity) among different researchers by means of virtualization. PlanetLab introduced a basic abstraction called *slice* which enables the isolation of different experiments. A *slice* is an orthogonal fraction of the testbed, or in other words, a set of nodes on which applications are assigned a fraction of resources by means of a kind of virtual machine called *sliver*. This concept of *slice* has evolved over time to also account for network isolation and other types of resources.

If we look at the wireless research area, there are also several initiatives to create a testbed to conduct experimental

research, but none of them address the problem of high scale deployments nor sharing of resources – two of the main principles behind Community-Lab. For instance, the IBBT w-iLab [4] in Flanders (Belgium), consists of mainly 200 indoor locations spread across three office floors and is mainly composed of sensor networks. In Greece, the NITOS [5] testbed enables researchers to completely control the software on the testbed nodes, which denies the possibility of several simultaneous users. Besides, it consists of 40 WiFi nodes which only enables small scale experiments. Finally, the ORBIT testbed [6] is a large indoor two-dimensional grid of around 400 802.11 radio nodes which can be dynamically interconnected into specified topologies, which makes it the largest indoor wireless lab.

III. CHALLENGES AND REQUIREMENTS

Deploying, managing, and using experimentation facilities in an existing CN, where ownership and administration of network infrastructure is divided over many entities, imposes a set of new challenges and requirements to the design of the architecture for Community-Lab. A selection of most important requirements is briefly discussed by the following subsections.

A. Lightweight, Low Cost, Free, and Open Source

In order to achieve a widespread deployment of CN nodes supporting the Community-Lab testbed, existing installations must be extended with new components which comes with additional equipment acquisition and maintenance cost (e.g. increased energy consumption) for the node owners. Therefore, lightweight and low cost solutions are important to keep expenses as low as possible.

Reusing of generally accepted approaches and free and open source software is mandatory to ease reviewing and trust, lower the burden for contribution, and to aim for sustainability and acceptance in the free networks communities. For the same reasons, the CONFINE project offers the software and documentation needed to set up other testbeds (CONFINE testbeds) as free/open source software.

B. System Stability

To facilitate the preparation and execution of experiments, researchers should be provided with an environment they are familiar with and which allows them to reuse already existing implementations and prototypes. Such environment is given by a Linux OS in combination with root permissions which has evolved to a de-facto standard environment for networking experiments.

On the other hand, from a security point of view, it must be ensured that experiments enabled with root permissions can not affect the general CONFINE Node System used by Community-Lab. This can be achieved by isolating experiments on dedicated hardware or with virtualization techniques. But, despite many mature virtualization solutions exist for UNIX systems, few of them have been successfully ported to embedded hardware which is commonly used for nodes in CNs.

C. Network Stability

Since the Community-Lab testbed aims to enable community integrated network experiments, some (controlled) level of interaction between the experiments and the main CN must be possible. On the other hand it must be ensured that experiments do not interfere seriously with each other nor with the production network. The potential harm of experiments can be classified into two types:

- Resource Consumption: It must be ensured that each experiment gets a limited but fair share of the available link capacities. Traffic Shaping and Distributed Rate Limiting techniques can be used here.
- Conflicting Protocol Behavior: Experimentation with key network services (medium access control, routing, DNS) and other protocol resources (address, port, channel assignments) must be isolated or rigorously controlled. Firewall or Software Defined Networking techniques (like VLAN tagging) or even the complete physical isolation of link resources are considered depending on the lowest accessed network layer of the experiments.

D. Privacy of CN users

While researchers should be conceptually enabled to collect user and network data (for statistical evaluation or as real-time feedback for protocols), privacy regulations of the corresponding CN licenses and peering agreements must be respected. Therefore, the CONFINE system used by Community-Lab must provide convenient mechanisms for traffic filtering and anonymization of networking data before it is provided to researchers.

E. Management Robustness

One of the interesting characteristics for research in real-life CNs is the variety of link setups and stabilities that arise due to the distributed management of CN infrastructure and the objective to provide connectivity even to remote locations at low cost. Often, such locations are abandoned by commercial ISPs because of deployment challenges and poor revenue perspectives.

Therefore, the system (in charge of managing the distributed experiments in CONFINE enabled nodes like those in Community-Lab) must be prepared to deal with connectivity starvation and the spontaneous appearance and outage of links and nodes.

F. IP-Address Provisioning

The objective to allow active experiments on top of existing CN infrastructure as well as interaction with users demands the provisioning of CN public IP addresses for experiments.

However, due to the general scarcity of IPv4 address space, the allocation of unique addresses for experiments is difficult. Also, many CNs are using their own address allocation schemes based on the limited private address ranges for internal routing and end users.

This non-coordinated allocation of IPv4 addresses is raising another challenge for the long-term objective to enable federation of CONFINE testbeds (like Community-Lab) deployed in different CNs.

Relying on IPv6 seems the natural solution to this challenge but we must consider the fact that only few existing CNs yet support native IPv6 routing.

G. Compatibility

New CONFINE nodes (serving as experimentation facilities to host distributed experiments) must be gradually deployable and integratable into existing and productively used CN infrastructure. Therefore, compatibility between CONFINE nodes and existing community nodes (Community Devices) must always be ensured.

Because nodes in CNs are owned and setup individually and may be based on proprietary and closed software systems, no modifications to the original community node firmware should be required and only few assumptions should be made on supported features and protocols or the concrete implementation of the original community nodes.

H. Role-specific APIs and Usability

The operation, deployment and usage of a testbed integrated in existing CNs involves different roles. In particular, the roles of testbed managers, node owners, and researchers have been identified.

In order to simplify the operations related to each role and prevent mis-configurations due to overlapping or missing responsibilities, role-specific APIs and access control should be provided.

Finally, the tools and APIs provided for managing and deploying a CONFINE testbed like Community-Lab should be as simple and intuitive as possible (for each involved role and despite the intrinsic complexity).

IV. COMMUNITY-LAB TESTBED ARCHITECTURE

A CONFINE testbed like Community-Lab consists of at least one CONFINE *server* (or *controller*) and a set of CONFINE-enabled (testbed) *nodes* spread among the existing nodes of a community network. The controller is a normal computer which must be directly reachable from inside the CN (via CN IPs) and usually from the internet. Nodes are integral parts of the CN and form the actual network infrastructure by respecting peering agreements of the CN and relaying other nodes' traffic. A CONFINE testbed node implements additional access control, resource isolation and management capabilities with the objective to grant external researchers/users with a confined access to the node's processing and network resources.

The design described in this section intends to facilitate the addition of testbed nodes in any location of a CN by connecting additional devices to existing CN nodes with minimum or no testbed-specific changes to their configuration. At the same time, it provides researchers with familiar Linux environments with root access and rich connectivity options (see IV-A4) including:

- Simple NAT access (like that of most home computers).
- Public CN addresses (to serve incoming connections).
- Full L2 access to VLAN-tagged networks (to implement routing experiments).
- CN traffic sniffing (with anonymization).
- Raw L1 access (under certain conditions).




Fig. 1. The architecture of a Community-Lab node.

A. Long-term vision

As shown in Figure 1, a **CONFINE node** (node) consists of two or three devices: the **community device (CD)**, the **research device (RD)** and an optional **recovery device**. This separation tries to preserve the stability of the CD (and thus of the CN) and is most compatible with any firmware it may run. All devices in a node are linked by a wired **local network (LN)**, possibly shared with other non-testbed devices like CN clients using the CD as a gateway.

1) *The community device:* The CD is a *completely normal community network device* with at least two different network interfaces: a wired or wireless one connecting to the CN (a **community interface** with a **community address**) and a wired one to the LN (the **gateway interface** with its **gateway address**). It can act as a simple gateway for hosts connected to the LN, some of them with fixed addresses not assigned to other devices via DHCP or similar. This allows the use of a uniform address scheme for testbed elements in the node.

The node is accessed by testbed servers, administrators and researchers via the community interfaces of the CD, although node administration may occasionally proceed directly through the LN.

2) *The research device:* The RD is a relatively powerful device (compared to the CD) running a *customized OpenWrt firmware* provided by CONFINE which allows simultaneously running several *slivers* implemented as Linux containers [7]. Slivers have limited access to the device's resources and to one another, thus ensuring slice and network *isolation*. This is guaranteed by the **control software** run in the RD through tools like `ebtables`, `arptables`, `iptables`, `tc`, Open vSwitch. . .

The RD implements an **internal bridge** with an **internal address** that is the same in all testbed nodes and which belongs to a private network which does not clash with CN or LN addresses. The RD offers some basic **sliver services** (see

IV-A5) on the internal address, including NAT gateway access to the CN.

The RD also implements a **local bridge** which connects to the LN through a wired interface (the **local interface**). The bridge is used for simple network layer access to the CN through the CD's gateway address, and the **local address** of the RD in it is fixed and used for testbed management and remote administration. For easy RD setup and local administration, the local interface may also sport a **recovery address** that is easily predictable or the same in all testbed nodes and that belongs to a private network which does not clash with CN or LN addresses (nor those of the internal bridge). A **debug address** which is also private and easily predictable can be used to access different RDs in the same local network for debugging purposes.

The RD may have additional **direct interfaces**, each one connected to its own **direct bridge**. These interfaces may be connected to the CN at the link layer and used for experiments below the network level (see IV-A4).

All the aforementioned bridges are managed by the control software in order to ensure network isolation between slices (i.e. between slivers running in the RD) as mentioned above, and to keep CN stability and privacy.

3) *The recovery device:* The node may include some simple recovery device whose purpose is to *force reboot* of the RD in case of malfunction using some direct hardware mechanism (like a GPIO port connected to the power supply of the RD), thus avoiding the need for physical presence for rebooting devices in places with difficult access.

The recovery device may get remote instructions from the CN (via the LN) or via different sensors, preferably based on wide-range technologies suffering from low interference and differing from those used by the CN (like ham radio signals, GSM calls or SMS). It may also receive a *heartbeat signal* kept by control software via some direct link like a serial line; when the recovery device misses a number of heartbeats, it reboots the RD.

A more advanced version of this device may help the RD boot some recovery environment (e.g. via PXE) or collaborate in some techniques for safe device upgrade and recovery (see [8]) to allow restoring its firmware to a known state.

4) *Node and sliver connectivity:* The research and recovery devices have fixed addresses in the LN and they need no routing protocol software to reach the CN: they simply use the CD's gateway address as a default gateway. Static configuration or DHCP suffice as long as the addresses they get are fixed.

In contrast, the connectivity of a sliver is determined by the network interfaces it sports, which are requested by the researcher at sliver definition time and depend on the interfaces provided by the RD and their features. Default routing configuration is explicitly controlled by the researcher to avoid traffic unexpectedly flowing through unwanted interfaces.

- Every sliver has a **private interface and address** whose host side *veth* interface is placed in the internal bridge. The address is automatically assigned from the RD's private network, thus allowing access to the RD's internal address and sliver services (see IV-A5). The researcher may choose to use the latter address as the default

gateway, in which case traffic is routed by the RD through the local bridge to the CD's gateway address after performing NAT. This allows *client access to the CN* but not connections from the CN to the sliver (similar to a home computer behind a NAT gateway on the Internet) *nor traffic between slivers*. Thus the sliver is ensured that there will be no incoming connections on that interface, obviating the need for firewalls or access control.

Sliver container 1 in Figure 1 only has a private interface in the internal bridge.

- If the RD has been allocated some public addresses, the researcher may request a **public interface and address** whose host side *veth* interface will be placed in the local bridge, thus allowing access to the LN. The address is automatically assigned from the RD's pool of public addresses, and the researcher may choose to use the CD's gateway address as the default gateway. This allows *connections from the CN to the sliver* (similar to a computer directly connected to the Internet through a normal gateway).

Sliver container 2 in Figure 1 has a public interface in the local bridge.

- If the RD has a direct interface, the researcher may request a **passive interface** (with no network address) whose host side *veth* interface is placed in the associated direct bridge, thus allowing direct access to the CN. Permission is granted only for *traffic capture* on the passive interface, which is anonymized by control software (e.g. an OpenFlow controller on an Open vSwitch-based bridge [9]). This allows *CN traffic analysis* while respecting privacy.

Sliver containers 3 and 5 in Figure 1 have passive interfaces in direct bridge X.

- If the RD has a direct interface, the researcher may request an **isolated interface** (with no pre-assigned network address), i.e. a VLAN interface on the associated direct bridge using one of the VLAN tags allocated to the slice at creation time. *Any kind of traffic can be transmitted and captured* on such an interface at the cost of being isolated from the CN at the link layer and delivered only to neighboring slivers of the same slice. This allows experimentation with network layer protocols and addresses (eg *routing experiments*) to operate safely on groups of nearby located testbed nodes.

Sliver container 4 in Figure 1 has an isolated interface on direct bridge X.

- If the RD has a direct interface, the researcher may request *raw access* to the interface's network device. The **raw interface** is *moved* into the sliver container and the associated direct bridge is disabled while the sliver is running. Since the sliver has *full physical control* on the network device, network isolation can not be guaranteed, so only that sliver is allowed to interact with that interface (exclusive access). Moreover, this access can disrupt CN operation and privacy, so it should only be allowed *under very particular circumstances* (e.g. out of reach of the CN).

Sliver container 5 in Figure 1 *owns* the RD's direct interface Y.

Besides the included private interface, a sliver may be granted several interfaces using the local bridge or any direct bridge. Conversely, the local bridge or a direct bridge may provide several slivers with shared access to the same interface.

The setup of a sliver's networking is accomplished by populating its image with appropriate configuration files. Static configuration allows testbed servers to know the addresses of a sliver *a priori* (so they can be propagated to other slivers in the same slice) and facilitate the configuration of routing, filtering and traffic control rules in the RD.

5) *Sliver services*: The RD offers on its internal address some basic services which become available to slivers through their private interfaces. This setup intends to relieve researchers from configuring those services themselves while providing some features tailored for slivers (some of which are not feasible at the sliver level), all in a trivial and reliable manner by accessing an address which is always the same on an interface which is always available and closed to external access. Factoring services out of the slivers also saves resources and minimizes the changes of breakage by accidental misconfiguration on the sliver. Some examples of these (optional) services are:

- DNS: The RD acts as a name server, possibly with customized testbed- or slice-wide domains.
- SMTP: The RD acts as a mail gateway, possibly rewriting source mail addresses for easy sliver identification.
- A NAT gateway for slivers needing basic client-only connectivity to the network.
- A DHCP service which can be used by a sliver as a last resort to regain minimum (NAT) connectivity in case of lost network configuration.
- DLEP [10] can offer link quality and speed information from network devices to slivers (useful for routing experiments using cross-layer information).
- Other services specific to the testbed for retrieving testbed or slice information, or running privileged operations in the RD (like PlanetLab's vSys [11]).

6) *Out-of-band access to slivers*: Regardless of whether a slice has network access at a given moment, a researcher should be able to access any of its slivers through the RD's local address in a way similar to PlanetLab, allowing for remote command execution and file transfer for collecting experiment results.

7) *The community container*: This architecture also allows a low-cost version of a Community-Lab node where the CD runs virtualized inside the RD as a **community container** (CC) thus saving on hardware at the price of stability. In this case each community interface is considered a direct interface in the RD, and the CC has a *veth* interface placed in its associated bridge. The gateway interface of the CC is a *veth* interface placed in the local bridge, and the RD's local interface can be kept for wired node administration and for other non-testbed devices. The CC has few restrictions on the local and direct interfaces while slivers can still access them via passive or isolated interfaces (but not as raw interfaces). The CC may manage several direct interfaces in this way.




Fig. 2. The IPv6 overlay spanning two community networks.

B. Current state

At the time of writing this document (end of May 2012) we are just finishing our first release of the CONFINE Node System (CNS) for RDs [12]. The main objective of this first version is to provide a proof-of concept for selected technologies and be able to gather experience from the execution of first experiments on real deployments.

CNS is built on top of the upcoming “Attitude Adjustment” release from OpenWrt [13]. The decision to base our work on the current developer branch (trunk) of OpenWrt was a trade off between the requirement to provide up-to-date implementations for quite a number of relatively new Linux networking and virtualization technologies and the lack of maturity that comes with development branches.

The current implementation provides role-specific APIs and user interfaces for node owners, testbed management, and researchers.

By default, the task of integrating a new testbed RD into an existing CN is under the control of the node owner. After the initial flashing of the RD with our firmware image (implementing the CONFINE Node System for a pre-selection of supported x86 hardware like Alix or Com-mell), a node system must be customized for its specific deployment. Therefore, a small set of testbed- and node-specific attributes (node ID, public keys...) must be exchanged between the node owner and the testbed operators via a central web portal and used to complete the configuration template of the RD. The node-management functions `confine_node_enable/disable` can be used at any time to cleanly activate or deactivate the participation of the RD in the testbed.

The abstract concept of a sliver (as introduced in Section I) is realized via virtualization. To provide researchers with a common experimentation environment we decided to use Linux Container (LXC [7]) as a very light-weight virtualization solution. This allows us to grant researchers root access to an OpenWrt or Debian guest system (the sliver), but running completely under control of the CONFINE RD host system. LXC user-space tools have been ported to OpenWrt as a new package and are used by the CNS to control the LXC container internally. This sliver environment given by the container OS can be further extended by the researcher with additional tools, data, and init scripts to automatically execute the processes that

define experiments. Interactive access to slivers is provided at any time via SSH and the testbed management network (*tinc* IPv6 overlay, see next point).

The role of the testbed operators is to manage the deployment and life cycle of experiments (via the concepts of slices and slivers). The following set of functions were implemented as Bash scripts for remote invocation (via SSH and the CONFINE management network) from a central server of the testbed.

- `confine_sliver_allocate` is used to allocate the resources (like local namespaces, IP addresses, storage) on a RD that are necessary to execute an experiment. The function requires a sliver description in UCI syntax specifying the required interfaces, addresses and OS type. On success, it returns precise information about the allocated sliver attributes. The attributes of all slivers of a slice must be collected by the server for later summarization and provisioning as slice attributes during the deployment of each sliver. Thereby, each sliver has access to the sliver attributes (e.g. IP addresses) of all other slivers in its slice before it is started.
- `confine_sliver_deploy` is used to setup the sliver environment in a RD. Therefore a linux container and root file system is created and customized according to RD, slice, and sliver specific attributes. The function requires the slice attributes in UCI syntax as input.
- `confine_sliver_start/stop/remove` can be used to start, stop, and remove the LXC container related to a given sliver (in the RD where the command is invoked). The last function does also release all allocated resources reserved for that sliver.

1) *The IPv6 overlay*: Creating testbeds that span several community networks (CNs) is drastically hampered by their lack of widespread IPv6 support. They normally use private IPv4 addressing that can not be trusted to be compatible across CNs. Also, the diversity of CN devices using IPv4 mean that research devices may sit behind NAT boxes or firewalls further limiting their reachability. IPv4 address scarcity also keeps the testbed from using a clean and uniform address scheme for their elements.

IPv6 migration solutions like 6to4, 6in4, or Teredo either have problems with protocol 41 handling in NAT boxes or firewalls, or they use host IPv4-based addresses which can

also clash between CNs. Together with most VPN solutions they get complicated to setup in a mesh-like cloud, or use centralized architectures which alter the resulting topology and can turn the VPN server into a bottleneck –e.g. OpenSSH.

A CONFINE testbed like Community-Lab works around these problems by creating an **IPv6 overlay** based on the *tinc* [14] VPN software. *Tinc* allows setting up a mesh network where data is exchanged directly between endpoints with the VPN software taking care of routing traffic and propagating endpoint information, thus needing minimal configuration.

As shown in Figure 2, a *tinc* daemon in a testbed node only connects to the testbed server or a CN testbed gateway. Gateways connect between themselves through another network like the Internet. Together they form a single routed IPv6 network where individual hosts (servers, gateways and clients) and subnetworks (node devices and slivers) become available to one another.

Since the overlay makes all elements in the testbed easily reachable through predictable IPv6 addresses with automatic encryption and authentication between endpoints, the IPv6 overlay is used as a **management network** while CNs gain widespread IPv6 support.

C. Community-Lab integration in existing community networks

What makes the Community-Lab testbed different from existing ones is the integration with actual community networks (CNs). The nodes belonging to the testbed are deployed among the community network to be able to perform experiments in a real environment with real users. However these networks are indeed production networks, so the experiments must have some limitations or provide some degree of protection to community network owners. For instance, nodes placed in the core of the network cannot, as a general rule, execute link layer experiments because it can easily disrupt the correct operation of the production environment.

Community networks have multiple owners, i.e. they conform a *distributed property*, which means that each user is the owner of his own device. In a city, where people live in buildings, the users normally share a single device placed on the roof – a good practice though not always possible.

A typical community network deployment looks like Figure 3a. The green lines represent WiFi links. Most of these networks are using this technology because it is cheap, easy to use and can be easily obtained. However it might be any other kind of technology such as WiMax, Ethernet or Optical Fiber depending on the owner’s needs.

The integration of the Community-Lab devices inside existing CN requires the collaboration of at least one part of the community. To make it possible, the CONFINE project provides additional resources in exchange of covering some of the costs like new links or new nodes. We envision three different levels of integration.

1) Single research devices connected to existing nodes:

As shown in Figure 3b (left), a research device (RD) without antennas is adopted by a community user not necessary related with the project. The RD will be connected to the community device (or CD, which is owned by the person who adopts it) through an Ethernet connection. The research device will

obtain a valid community network IP – depending on the availability of a DHCP server on the community device – and it will be reachable from the rest of the network and from the Community-Lab testbed server. Experiments between research devices can be performed, however restricted to application layer experiments because of the lack of direct access to the physical layer. If the owner allows it, a WiFi antenna might be added to the research device to allow researchers to obtain valuable link layer information, which could be useful to perform long term studies on the real usage of such links.

In this case the CONFINE project only needs to install a single device (the RD) because it uses a existing CD owned by the user.

2) *Nodes to extend the current network*: One of the explicit objectives of the CONFINE project is to support the growth and usage of this *distributed property* network model. In some cases where the community network needs one or several extra links to improve performance or availability, the project will contribute to provide extra resources. These resources will be placed in some community member’s location, and so the adoption of a research device (following the integration model described before) will be a requirement. In this case the CONFINE project will add a new CD and the common RD.

3) *Multiple nodes (cloud) deployed in a specific scenario*: In some specific scenarios controlled by a CONFINE partner – i.e. university campus – an entire cloud may be deployed. As shown in Figure 3b (right), it implies the installation of several nodes. The CD must be a node compatible with the existing CN to be able to connect to it. The RD may contain one or several WiFi antennas and radio devices depending on the characteristics of the experiments to perform. Having several radio devices will allow researchers to execute any kind of link layer experiments given that the access to the physical layer will be exclusive.

If the set of links (between the nodes) are interesting for the community, the users may make use of them. Moreover, they must be aware that such a cloud is an experimental network and its availability and stability cannot be guaranteed. However, this bilateral open access between community network users and researchers can provide highly valuable information to evaluate new protocols and services.

In this case the CONFINE project will add the new CD, the common RD and any other hardware needed (like wires, metallic mast, electricity box, etc.).

D. Discussion

There are three major characteristics which make the Community-Lab testbed of the CONFINE project unique compared to the deployments presented in Section II: i) the larger scale of the testbed with an expected size of thousands of nodes by the end of the project compared to the few tens of nodes of current deployments; ii) the integration within already existing production community networks which provides a more realistic scenario for experimentation opposed to controlled lab environments; iii) the support for long term studies of new network protocols and parallel experimental services which share the same wireless infrastructure – both from the testbed and from the participating wireless community networks – by means of node and network virtualization




Fig. 3. Community-Lab integration with existing community networks.

opposed to other testbeds which provide dedicated nodes to testbed users. Of course, the development of current deployments presented in this section can and should provide a solid foundation for building this unique infrastructure.

V. CONCLUSION

Community networks are an emerging field to provide citizens with connectivity in a sustainable and distributed manner in which the owners of the networks are the users themselves. Research on this field is necessary to support CNs growth and scope, and improve their operation and quality. However, a few initiatives have arose to provide researchers with a platform to conduct such novel experimentation which must push forward this new networking model. The CONFINE project aims to fill this gap by deploying Community-Lab, a testbed for community networks inside existing community networks.

Throughout this paper, we have presented the motivation, requirements and challenges that pose the deployment of such testbed. We have described our envisioned architecture and specified the decisions made during the design process taking into account the challenges and peculiarities of CNs. Our next step is to make this architecture a reality for researchers to experiment with novel protocols and services which should provide community network owners, operators and citizens overall a better experience.

ACKNOWLEDGMENT

This work is supported by the European Community Framework Programme 7 within the Future Internet Research and Experimentation Initiative (FIRE), Community Networks Testbed for the Future Internet (CONFINE), contract FP7-288535.

REFERENCES

- [1] R. Sombrutski, A. Zubow, M. Kurth, and J.-P. Redlich, "Self-organization in community mesh networks the berlin rooftop," in *Operator-Assisted (Wireless Mesh) Community Networks, 2006 1st Workshop on*, sept. 2006, pp. 1–11.
- [2] D. Aguayo, J. Bicket, S. Biswas, R. Morris, B. Chambers, and D. De Couto, "Mit rooftop," in *Proceedings of the International Conference on Mobile Computing and Networking*, 2003.
- [3] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman, "Planetlab: an overlay testbed for broad-coverage services," *SIGCOMM Comput. Commun. Rev.*, vol. 33, no. 3, pp. 3–12, Jul. 2003. [Online]. Available: <http://doi.acm.org/10.1145/956993.956995>
- [4] S. Bouckaert, W. Vandenberghe, B. Jooris, I. Moerman, and P. Demeester, "The w-ilab testbed," in *Proceedings of TRIDENTCOM*, 2010.
- [5] A. Anadiotis, A. Apostolaras, D. Syrivelis, T. Korakis, L. Tassioulas, L. Rodriguez, and M. Ott, "A new slicing scheme for efficient use of wireless testbeds," in *WINTech'09: Proceedings of the 4th ACM international workshop on Experimental evaluation and characterization*, 2009, pp. 83–84.
- [6] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Krems, R. Siracusa, H. Liu, and M. Singh, "Overview of the orbit radio grid testbed for evaluation of next-generation wireless network protocols," in *Wireless Communications and Networking Conference, 2005 IEEE*, vol. 3. IEEE, 2005, pp. 1664–1669.
- [7] "Linux containers." [Online]. Available: <http://lxc.sourceforge.net>
- [8] "Safe upgrade of embedded systems." [Online]. Available: https://archive.fosdem.org/2012/schedule/event/safe_upgrade_of_embedded_systems.html
- [9] B. Pfaff, J. Pettit, T. Koponen, K. Amidon, M. Casado, and S. Shenker, "Extending networking into the virtualization layer," *Proc. HotNets (October 2009)*, 2009.
- [10] S. Ratliff, D. Satterwhite, S. Jury, G. Harrison, and C. Cisco, "Dynamic link exchange protocol (dlep)," 2012.
- [11] S. Bhatia, G. Di Stasi, T. Haddow, A. Bavier, S. Muir, and L. Peterson, "Vsys: A programmable sudo," in *Proceedings of the 2011 USENIX conference on USENIX annual technical conference*. USENIX Association, 2011, pp. 22–22.
- [12] "A hack: 1st milestone of the confine project," May 2012. [Online]. Available: <http://wiki.confine-project.eu/milestones:a-hack>
- [13] "The openwrt embedded development framework," 2012. [Online]. Available: <http://www.openwrt.org>
- [14] "Tinc vpn," 2012. [Online]. Available: <http://www.tinc-vpn.org>