Design Issues in Internet 0 Federation

Karen R. Sollins Ji Li
MIT Computer Science and Artificial Intelligence Laboratory

Abstract— **Internet 0 is proposed as a local area network that supports extremely small network devices with very little capacity for computation, storage, or communication. Internet 0 addresses the issue of connecting very small, inexpensive devices such as lightbulbs and heating vents with their controllers. To achieve this effectively, Internet 0 assumes both that operating between communicating end-nodes should not require third-party support, and that IP will be available all the way to those end-nodes. Several simplifying assumptions are made in Internet 0 to achieve this. The objective of this paper is to explore issues of design in a context where federation of an Internet 0 net either with other Internet 0 nets or the global Internet becomes important. The question we ask is whether the end-node in such an Internet 0 needs to know more or behave differently in such a federated environment, and how one might achieve such federation. We explore three aspects of network design in this study: addressing and routing, traffic collision and congestion control, and security. In each case, based on analysis, we conclude that to reach our goals in a generalizable and extensible fashion, a third party service will be needed to act as an intermediary, and propose that a single service should provide all the required federation services.**

Index Terms— **Internet 0, the Internet, network architecture**

I. Introduction

Internet 0 [1], [2] is a network design intended for low-bandwidth, low-speed network applications. The starting assumption is that there is a class of network end-nodes or devices that are extremely small and inexpensive, such as light bulbs that could valuably be available on a network, especially in a building environment. Upon further thought, Gershenfeld et al. [2] concluded that flexible and extensible local networking for limited communication and limited interaction or functionality would allow for the design of a much simplified network, but that could still provide IP to the end-nodes of the network. Part of the objective for IP to the end-nodes is that it enables applications such as simple web-style interactions. To that end in the project they build tiny web servers that can only report the binary state of a light bulb or switch and change that state.

More generally, the drivers of the Internet 0 project include both energy and cost minimization, ability to deploy in environments such as that of the construction industry, where flexibility and extensibility of many small devices, whether light bulbs, thermostats or components of an alarm system may need both installation after construction and relocation after the fact. In addition, the flexibility of including all such limited capacity devices and their controllers and interfaces on a network allows for improved architectural (in the building sense) expression. To achieve these objectives, Gershenfeld et al. arrived at a number of key design features:

- IP to the end-nodes, to achieve end-to-end IP level service without the need for translation;
- Integrated layer processing, as first described by Clark and Tennenhouse [3] to minimize the cost of end-to-end IP;
- Direct communication and functioning between end-nodes, without the need for a third party node;
- Self-assignment of identities, to avoid the requirement of a third-party for identification assignment;
- Bit-transmission length (time) longer than the length of the network, to allow transmitters to learn about collisions before completing their

This work is funded in part under NSF Grant ANIR-0137403, "Regions: A new architectural capability in networking", and NSF Grant CCR-0122419, "The Center for Bits and Atoms". It is also funded in part by a gift from the Cisco University Research Program.

Karen Sollins is with MIT CSAIL. Email: sollins@csail.mit.edu

Ji Li is with MIT CSAIL. Email: jli@csail.mit.edu

The corresponding address is: MIT CSAIL, 32 Vassar Street, Cambridge, MA 02139.

own transmission;
- Common bit and byte transmission representation independent of the underlying medium using a Manchester encoding scheme. Because the coding scheme transcends all media, no transformation or recoding is needed;
- Adherence to an open standard to enable broad interoperability. One might assume that the Internet protocols provide this, but in Internet 0, a common low-level Manchester style coding scheme is also proposed as a layer of homogeneity.

The result is an extremely simple protocol stack, based on a model of broadcast at the MAC layer and an assumption that the transmission of a bit is longer (slower) than the length of the network, with encoding at that layer in their simple Manchester coding scheme, self-selected IP addresses, unique at least within the scope of individual Internet 0 net, and IP running over everything, in order to support extremely simple HTTP. In light of this significantly simplified design for a network, this paper explores the question of how to allow that limited network environment to continue to exist while considering the implications of attaching such an individual Internet 0 network to a larger composite of multiple Internet 0 networks or a metanet [4]. In the larger Internet, there has been ongoing thought put into the end-to-end arguments [5] and which functions and services of the network should be provided end-to-end or not. In this paper, we consider three such functions and examine the alternatives for provision of them either end-to-end or not. The three functions are identity, collision control leading to congestion avoidance, and security. With respect to identity, Internet 0 takes the position that a node should be able to decide its identity for use in packet transmission to and from its peers independently of any third party. In the domain of security, under discussion is a proposal for encryption capabilities to provide privacy and authentication between peers. [6] presents the scalable encryption algorithm that makes this possible. Finally, the position taken in Internet 0 with respect to congestion is to design it out of existence. What this means is that when contention occurs, it will not be found in the network, but cause a transmitting node to discover collisions during transmission and then back off, if necessary. To the extent that the design of Internet 0 is stable and fixed, we will demonstrate that a common approach of providing an intermediary service will be necessary in all three cases. We expect this to generalize well.

The remainder of the paper is organized as follows. First, we will analyze three scenarios, the individual Internet 0 net, a federation of Internet 0 nets, and a federation including the global Internet. We will then evaluate how best to provide the basic functions of addressing, collision control, and security in these three contexts without requiring changes to an individual Internet 0 node.. We conclude that a third-party gateway service will be necessary in each case.

II. The Problems: Internet 0 Organization and Federation

We begin with more detail about Internet 0 organization and mechanism, in order to explore the issues of federating Internet 0 networks with each other and the global Internet itself.

A. Internet 0 organization

First, let us consider a single Internet 0. At its lowest layer it is a broadcast medium. In fact, it may be a set of broadcast media with translators between them. At this level, the design criteria include that the propagation time of a bit from one end to the other (between the most distant devices on the net) is longer than twice the transmission time of that bit from its source. This must include any medium translation. Thus, the length of such a network will have an upper limit determined by the combination of the length of a transmission and its propagation time. Such a scheme allows for discovery of collisions of bits prior to completion of the transmission of the bit, and hence in upper layers will allow for sharing of the resource without need for a great deal of congestion control. It also provides a higher probability that the transmitted bit will have been received correctly, without need of acknowledgment. It is valuable to note here that there is no theoretical or algorithm upper limit on the number of nodes attached to an Internet 0. The number of nodes is only limited by the number that can be physically connected to the network. On the other hand, as the number increases, bit collisions will increase. In fact, as was studied for Ethernet [7] performance will degrade exponentially as attempted use of the network increases. On the other hand, the expectation here is that the traffic between devices will be

limited to low volume and slow activities, for example (1) polling for state such as whether the light is off or on, (2) measurements such as temperature and humidity, and (3) setting behavior or state of a device, such as turning it off or on, or setting the temperature at which it turns itself off or on. Because Internet 0 is not intended for general purpose computing, but rather for low-bandwidth, low-traffic situations, it is likely that an Internet 0 net could support many more devices than a typical Ethernet of the same caliber.

Above this, Internet 0 provides a common coding scheme for bits, based on Manchester coding. This simplifies the design and hopefully increases performance in transmitting across different media, rather than requiring coding translation at such boundaries. It also provides a layer of homogeneity, not typically provided in other device networks, which often only operate over small numbers of meters to a device that is more powerful and on a different sort of network, whether 802.11, Ethernet, or something else.

Above the bit transmission level, Internet 0 supports IP, the Internet Protocol [8]. The interesting feature of this layer is that although IP is used for communicating among nodes on the network, some of the ancillary functions that are traditionally part of the "IP layer" are not present in Internet 0, but the functions provided differently. In particular, the two that we want to address here are routing and addressing. In its simplest form, the Internet assumes that the addresses of source and destination are globally unique and can be used directly for transmission. [1] Typically address assignment, whether behind a NAT or not is done either manually or in an automated manner as with DHCP. Both of these approaches involve avoiding address conflicts. Internet 0 makes a similar assumption, but only within the bounds of a single Internet 0. Because an Internet 0 is a small broadcast environment, there is no need for coordination or a remote service to allocate non-conflicting addresses. A new node can simply pick an address and test it. If it is unused, then it is available. It is important to note that this approach implies that no third party service need be available in order to allow peer nodes to communicate.

[1]Things are not quite this simple in the Internet, because we have allowed for NATs or Network Address Translators. NATs allow for addresses behind them to be reused elsewhere and provide address translation between the global Internet and the networks behind the NATs.

A second way that the IP environment of Internet 0 is different from the global Internet is in routing. In the Internet, routing is achieved using a hierarchical approach. The network is divided in interconnected Autonomous Systems (ASs), between which the BGP protocol is used, to find paths among ASs, and separate local routing protocols, designed for the smaller and more homogeneous environments are used inside ASs. The assumption is that something outside the end node will determine the path taken to by packets to arrive at the destination. In Internet 0, this is more or less a moot point, because it is simply a broadcast environment, so all packets arrive everywhere and a node simply needs to know its own address and then determine whether an incoming packet is destined for it or not, with no routers involved, thus again preserving the goal that peer nodes be able to communicate without the assistance of a third party service.

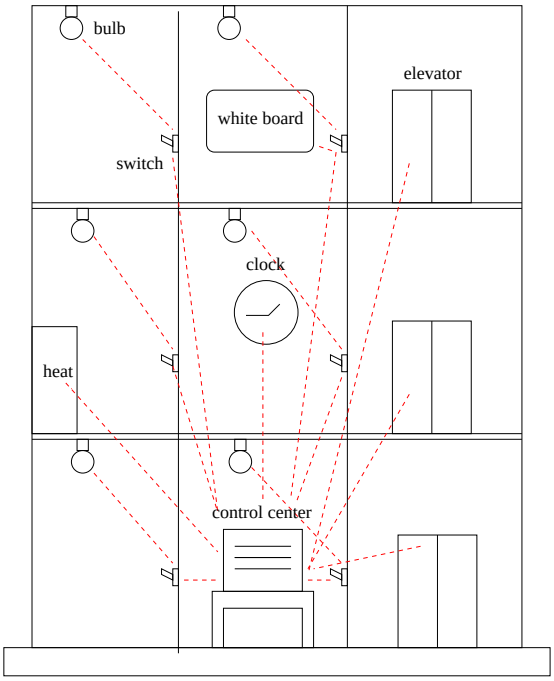

Fig. 1. A single Internet 0s within a building. The dashed lines represent the connections in the Internet 0. In the implementation, the connections may be some shared media.

Figure 1 depicts a single Internet 0 arrangement in a building. As a proof of concept, Krikorian[1] describes implementations of both UDP/IP and TCP/IP stacks on their microchip of choice. Further implementations have also been done, but are not published at this time.

B. A federation of Internet 0 nets

It is important to recognize that because of the limitations of a single Internet 0 net, we will often find that we need more than one. Consider the simple example of a large office building. One can imagine that one or a small number of floors may be served by a single Internet 0, but not the whole building. Then we find that there are functions that must cross the boundaries, so that the Internet 0 nets must be interconnected. One such example may be the sensors and controllers for elevators. Another may be overall building heating, lighting, power management, and emergency systems. These may have reasonably centralized facilities but need to perform functions across all the Internet 0 nets of the building. Another example is a campus containing multiple buildings under a single set of management controls, in terms of the facilities. Again, even if each building is supported by a single Internet 0, there may be need for the management, located in only one of those buildings to be able to monitor, query and control devices in all of them, either individually or as a group. In Figure 2 we depict such a federation. The question we must ask ourselves is whether and how a device can be in an Internet 0, yet participate in a federation of Internet 0 nets. Part of the question here is whether the node in an Internet 0 needs more information and capability, or whether it can continue to operate as though it were in only its private Internet 0. Below, we will explore both the questions of addressing and routing in such a federation

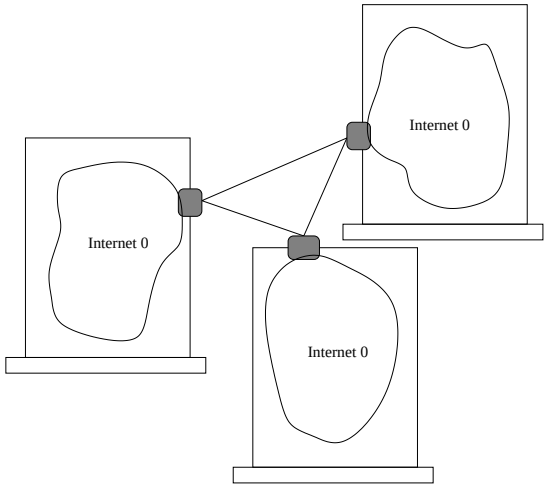

Fig. 2. Multiple Internet 0 between buildings. The shaded blocks represent the translation between the Internet 0s.

C. Federation with the global Internet

In fact, the problem is more challenging than that because there will be occasions in which one wants to support communication between a device on an Internet 0 and a device that is only in the larger Internet. Consider, for example, the office worker. Clearly Internet 0, as provided in an office will be inadequate and inappropriate for the worker's workstation. For that low latency and high bandwidth are increasingly important, both for remote work and for real-time applications. At the same time, from the desk, it would be valuable to be able to poll and control the devices in the office, without requiring that the workstation also be on the Internet 0 provided for the lighting and thermostat. At a greater distance, consider the traveler who is arriving home from a long trip. From the airport, perhaps using a wireless hand-carried device that is on the Internet, the traveler would like to turn up the heat and turn on the lights, prior to arriving home. In this case, access to the Internet 0 devices will be from a more remote part of the Internet. Again, we must consider whether and how such a federation of Internet 0 and the global Internet can occur, and whether that implies that the device on the Internet 0 will now need to support greater functionality. What we hope is that we can provide end-to-end IP connectivity, but not require that the Internet 0 device need to be a full participant in the whole Internet story. In particular, again we will ask the questions in the context of addressing and routing. Figure 3 depicts such a federation of Internet 0 and Internet 1.

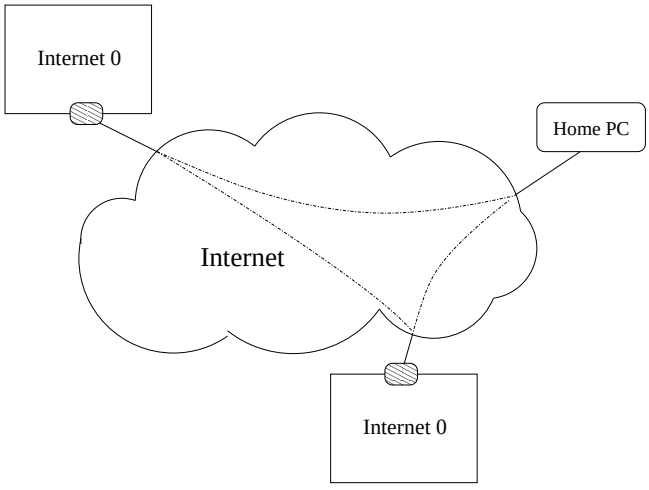

Fig. 3. Internet 0 nets and the Internet. The shaded blocks represent the entry and exit points of the Internet 0 nets.

D. Security

Finally, there is a proposal on the table to use the scalable encryption algorithm by Standaert et al. [6] as the basis for provision of authentication and privacy in an Internet 0. With the use of the functionality of an algorithm such as the Diffie Hellman algorithm[9] to provide unsupervised private key exchange and a scalable encryption algorithm that provides guarantees even in computationally limited situations, encryption can be achieved that will allow for authentication and privacy. Such functionality can be shared among any nodes that share a single key, so it allows for authenticated and private group communication, in the Internet 0 broadcast environment. Although this is not an IP layer question necessarily, we ask the same questions that we ask about addressing and routing, namely, can this function be provided effectively across a federation of either multiple Internet 0 nets or the global Internet without the Internet 0 nodes needing to be modified in order to know about the federation.

III. Analysis and approaches

In this section, we will consider in more detail the questions of addressing, routing and the Internet 0 security functions in the context of federation. Each will be considered in turn.

A. Addressing

As mentioned above, the approach taken in Internet 0 is to create a locally unique address in IP format. This is described in detail in Krikorian [1]. In that work, three categories of identity are explored, MAC address, IP address, and user level naming. We actually use at least one additional one, the Domain Name. Krikorian rejects both MAC layer and upper layer identification. In addition, he could have rejected Domain Names. The reason for choosing IP addresses is to meet the objective of providing IP to the leaf nodes. For that, the packets must retain the same form as normal IP. The scheme is the following.

- A node determines that it needs a new address.
- The node sends out a request for an address using DHCP. If it gets a response, it uses the IP address provided, trusting that the DHCP server will guarantee uniqueness of the addresses.
- If it gets no answer, it selects a random address from the address space 169.254.0.0/16 (excepting 169.254.0.0/24 and 169.254.254.0/24).[2]
- It then sends out a "who-has" ARP for that address.
- If it receives a response, the address is already taken and it tries again with another address.
- if it receives no ARP response to its ARP request, it uses the address and defends it in the future by responding appropriate to other ARP requests.

It is valuable to notice that although in the global Internet, a node must change its IP address if it moves any distance, as long as the Internet 0 node stays anywhere within the Internet 0, it can continue to use the same address. The almost 2^{16} addresses both allows for quite populous nets as well as being large enough that the probability of choosing a pre-assigned address is low.

Now let us consider what happens when two pre-existing Internet 0 nets are connected. If the two nets merge in such a way that broadcast still holds, then they can be considered one larger Internet 0. There will be a transition problem, because the same IP address may have been assigned in both, so, in order to avoid an N^2 problem, every device may need to acquire a new address in this new, larger net. Depending on the basis on which associations may have been made, either for interaction or security, more work may need to be done to re-establish those relationships using the new identities.

A lower overhead approach is to put a gateway between the two nets. This node will have several tasks. First, it will act as an address translator. Each address on one side of the gateway will have a substitute address on the other side of the gateway. This way, from the perspective of one side of the gateway, every node in both nets will have a unique address locally. As packets travel across the gateway, both source and destination addresses will need to be translated appropriately. Note that on one side of the gateway, it will claim the addresses of all the nodes on the other side of it, in order to process and transmit them, and respond the ARPs for any of those addresses to protect them from reuse.

Figure 4 and the accompanying Table I provide a simple example of connecting several Internet 0 nets

[2]The first of these two regions is reserved for "link local" addresses. The second is generally used for "self-assigned" addresses.

and an bit of a routing table to achieve this.

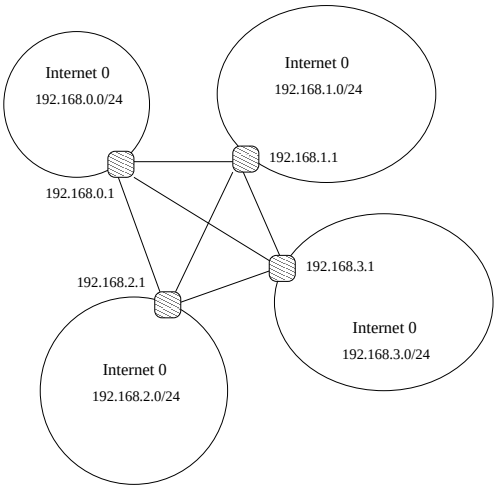

Fig. 4. Internet 0s within an organization can be connected into a complete graph due to its small scale. The shaded blocks refer to the servers, one for each Internet 0.

IP Prefix	Server Address
192.168.0.0/24	192.168.0.1
192.168.1.0/24	192.168.1.1
192.168.2.0/24	192.168.2.1
192.168.3.0/24	192.168.3.1

TABLE I
ROUTING TABLE FOR FIGURE 4

Thus, we conclude that providing a gateway, although it violates the dictum of peer-to-peer communication with the need for a third party, provides a smoother transition to a federation. We note here that a number of Internet 0 nodes can be tied together with gateways, always with the proviso that the total number of nodes in the whole federation not exceed (and probably not approach for efficiency reasons) the total number of addresses available in one Internet 0 address space. From any point, every other node in the federation must be addressable.

The question of federation with the global Internet has a similar flavor. In this case, there is no option for re-identifying all the potential members of the federation, because renumbering the whole Internet is not possible, nor is it possible to send an ARP to the whole interned to search for duplicate addresses. Thus, the only alternative is the gateway. In this case, from the perspective of the Internet 0 world, only a small selection of nodes from the global Internet world will have local Internet 0 addresses. There will need to be some higher level protocol for discovering new nodes in the global Internet that are of interest and giving them identities in the Internet 0. With that in place, things should work more or less the same as with federating Internet 0 nets.

Consider the example shown in Figures 5 and Table II. In this case identity translation is provided across the boundary with the global Internet. Thus, with respect to addressing in our federation situations, we propose a gateway as the solution to the problems of identity in the case of federation.

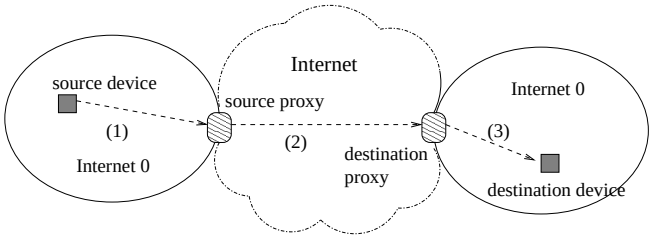

Fig. 5. Two devices talk to each other through their servers and the Internet.

B. Collision control

In considering collision control, we find a similar story to that of addressing. First, consider the federation of two Internet 0 nets. Internet 0 assumes that bit collisions can be discovered. If a collision occurs, the sender will need to retry. Again, we can consider two possibilities, either simply building a single larger net or including a gateway. If the two merged nets combined are small enough for the bit collision detection to continue to work, then there is no problem. But, it is more likely that the purpose of two nets was because of distances, so we must consider a gateway. The gateway causes a problem, because collisions will not be discovered across it. So, the gateway will need to provide an additional function of receiving bits, buffering them, and then, on the other side retransmitting them until they are sent without collision. By doing this, the gateway will preserve for the sending node the appearance that all collisions are detected.

When we consider a federation including the global Internet, the story must change a little, for two reasons. First, the uniform coding scheme of Internet 0 is not a standard and will not be used throughout the whole Internet. Therefore, the gateway must perform a coding translation between the

IP	name	Type	Location	Status	...
192.168.0.3	Light	G806	On	...	
192.168.0.5	Thermostat	G401	Off	...	
192.168.0.11	Elevator	Building G	On	...	
192.168.0.24	Fire Alarm	G8	On	...	

TABLE II
DATABASE ON A SERVER.

two worlds. Second, the large bit detection will no longer be viable. The gateway will now need to collect the whole packet, transform it, and send it out into the Internet. For traffic flowing into an Internet 0, again recoding Will be required and now the gateway will also use the bit collision detection scheme,, with one minor issue with respect to bit collisions. Since the global Internet does not use the Internet 0 Manchester encoding and IP is an unreliable protocol, both translation to other coding schemes will be need on the global Internet side of the gateway and there will be no collision detection or back-off mechanisms. Since IP is best-effort anyway, any protocol that is based on IP must assume that some packets will not arrive at their destination, and provide recovery at a higher level to the extent it is important. Since part of the definition of the IP protocol is that it is only best-effort, although the collision detection may have avoided certain problems, any layer sitting on top of IP must be prepared to deal with some of the packets not arriving, so this approach adds no new problems.

C. Security

The situation with respect to security is a little different than addressing and congestion. In this case, the encryption/decryption provided by SEA is not an issue. The proposition is that payload is encrypted using SEA, parameterized to be computationally feasible for Internet 0 nodes. SEA was designed as a parameterized algorithm to make it adaptable to a range of computation and memory constraints. Hence the strength of the algorithm can be tuned to the best capabilities of an end-node. Thus, when SEA is used, it will need to be tuned to be usable by the weakest node. It may mean that some negotiation is needed between the source and destination to determine this. But, that said, in terms of encryption of the payload of a packet, there is no issue with respect to federation. As discussed above the payload will be transmitted across a gateway as needed, but needs nothing further done to it. The encryption can be end-to-end, as long as it only includes the payload.

There are two further central problems, key distribution and authentication. We will consider them separately, beginning with key distribution. We can begin with Krikorian's idea of an introduction mechanism. If one assumes that physical control can be maintained for the Internet 0 devices, then one can postulate (and Gershenfeld has demonstrated [10]) a small device used for this purpose. If two Internet 0 nodes are to be introduced to each other, the small device is brought up to one to acquire its identifier (IP address). The device is then physically moved near the other Internet 0 device to "give" it the address of the first one. At this point the second one can contact the first. If bi-directional introduction is needed, that can just be added. This approach also allows for transfer of secret keys. It is important to notice that this is a very small third-party device, used to enable both pairwise identification and pairwise secret-key sharing.

Although physical introduction is very appealing and is perhaps the best approach in a small environment, it has drawbacks both in that it does not scale (to either multiple federated Internet 0 nets nor the global Internet, where physically proximity cannot be assumed). Thus, for example, to provide campus wide control of electrical appliances, in order to do power management, one cannot depend on pairwise introduction of a central service to every component in every office. Thus, other tools will be needed for management structure, introduction, and key exchange. Again, one can imagine that a service will be needed. In this case, it will need to be trusted perhaps on a longer term basis than the small introduction device, with an appropriate ini-

tialization scheme. For this key and certificate services have been designed for both public and shared-secret key systems. The Kerberos system is one example in the shared-secret universe. (See Peterson and Davie [11], Schneier [12], and Kaufman et al.[13] for examples of this, with further discussion of the limitations of such schemes as well.) Another alternative for secret-key sharing is the Diffie-Hellman approach [9]. In this case, assuming that two nodes have addresses for each other, they can agree on a shared secret without actually exchanging that secret, using algorithms that are known to be extremely difficult to reverse, but that are composable. Notice, that although a trusted key or certificate distribution service is not needed for such an exchange, for the pair to find each other originally, there must still be some trusted third party. In a public-key environment, this trusted service may be extremely wide public distribution of the non-private part of the Diffie-Hellman exchange, in order to make it extremely difficult for someone to corrupt the service.

A final problem, which argues further against the manual approach to introduction is that one needs a mechanism for future key distribution. It is well known that the longer a key is used the higher the risk of compromise. First, keys need to be replaced before compromise. Then, if there is a compromise, a further mechanism may be needed to reinitialize on a larger basis. Again, it is unlikely that a manual approach will be viable, and that a third party will be needed to provide the required functionality. Again Kerberos or a Diffie-Hellman exchange with the appropriate servers allow for more limited use keys with replacement as part of the scheme.

Authentication needs to be provided by something more than simply using the encryption algorithm, in this case SEA parameterized appropriately with a shared secret key. It is important to understand that even knowing that there exists an uncompromised shared secret is not enough for authentication. If node B receives an encrypted message from node A, how is it to know that the message really came from A and is not being replayed by some malicious intruder? The typical approach is a paired challenge, or "three-way handshake". It involves three messages as follows:

1) A sends an encrypted challenge to B, which B decrypts;
2) B responds by transforming the challenging, including its own challenge and encrypting both, A decrypts the paired response and challenge;
3) A transforms B's challenge, encrypts it, sends it back to B, and B decrypts it.

Assuming the challenge changes each time and the key has not been compromised, each node now knows that it is communicating with the only other node that knows the secret. Without this, if A simply sent an encrypted message to B, B could not distinguish it from a duplicate of an old message, and A would not really know that B received it. Clearly, as one goes deeper into such protocols, there is a longer string of problems one can attempt to address, with increasingly complex protocols. An alternative approach to this is to use certificates, which may provide short-lived keys tied to particular identities, so that theft and replay problems are reduced. The certificate approach involves a third-party, the trusted certificate authority, or in many situations a hierarchy of them.

As the reader has seen, a simple third party as suggested by Krikorian and Gershenfeld may be adequate for the extremely local Internet 0 case, but scaling, mobility, and longer distances will require at least a single, shared and trusted third party, and often a larger structure of such servers to make adequate identity, key or certificate discovery feasible.

IV. Related work

Related work on this subject is broad and can only be suggested here. The problem of federation of networks initially designed or instantiated independently is not a new problem Examples can be found as far back as the early days of Ethernet[14] and IBM's SNA architecture. More recently, the problem arises repeatedly in the Internet Engineering Task Force, for example with respect to the existence and effective usage of Network Address Translators[15]. In fact, if we consider routing at the IP layer, the actual routing for a packet across the Internet Firewalls are another context in which federation is a key element. Echelon[16], in conjunction with Cisco, recently released the iLON-1000 Internet Server, which seamlessly links IP networks with ANSI 709.1 (LonWorks) based control networks. Cisco NetWorks certified, the device works as both a web server and a tunneling

router, allowing low cost sensors and actuators running the ANSI 709.1 protocol to communicate peer-to-peer over Ethernet networks. This design approach uses the ANSI 709.1 protocol and industrial-grade wiring/connectors to communicate with sensors and actuators, and TCP/IP Ethernet to link together manufacturing pods, facilities, campuses etc. It is important to realize that these few references are only an extremely limited sampling.

V. Conclusion

In this brief paper, we began with Internet 0 as a premise and explored the question of how to combine either multiple pre-existing Internet 0 nets or an Internet 0 net with the larger global Internet. The question really revolves around choices and assumptions of operating in isolation and how those might conflict when crossing boundaries. Since a key assumption of Internet 0 was to provide IP to extremely tiny and low-capability devices, the question was whether that pervasive IP would actually allow for wider communication services, beyond the single Internet 0.

We considered three aspects of the system, identification, congestion, and security. In each of the three cases, we concluded that the only real option is to provide an intermediary component or service. This allows for a certain degree of shielding a local Internet 0 world from the rest of the world, making it possible to avoid disruption of the simple leaf node, while allowing for more distant communication. As suggested in the discussion, there will be further requirements on the upper layer protocols, whether naming (e.g. DNS or user friendly naming) decisions about authentication and authorization of access among devices, and management of unreliability, especially in the global Internet.

In terms of systems design, we propose that a generalization can be made that a gateway or proxy is the correct approach to providing federation, not only to solve the problems of identity, congestion, and security, but more generally. This suggests that our prior work on Regions [17], [18] and Wroclawski original Metanet proposal [4] provide a valuable architectural model for designing such a federated network environment.

References

[1] Raffi Chant Krikorian, "Internet 0: Inter-device internetworking," M.S. thesis, MIT, Aug. 2004.

[2] Neil Gershenfeld, Raffi Krikorian, and Danny Cohen, "Internet 0: Interdevice internetworking," in Scientific American, 2004.

[3] David D. Clark and David L. Tennenhouse, "Architectural considerations for a new generation of protocols," in Proc. ACM SIGCOMM 1990, Philadelphia, PA, 1990.

[4] John T. Wroclawski, "The metanet," in Proc. Workshop on Research Challenges for the Next Generation Internet, Computing Research Association, Ed., Washington, DC, 1997.

[5] Jerome H. Saltzer, David P. Reed, and David D. Clark, "End-to-end arguments in system design," ACM Transactions on Computer Systems, vol. 2, no. 4, pp. 277–288, Nov. 1984.

[6] François-Xavier Standaert, Gilles Piret, Neil Gershenfeld, and Jean-Jacques Quisquater, "Sea: A scalable encryption algorithm for small embedded applications," in Workshop on RFIP and Lightweight Crypto, Graz, Austria, 2005, ECRYPT.

[7] J. Hastad, T. Leighton, and B. Rogoff, "Analysis of backoff protocols for multiple access channels," in Symposium on Theory of Computing, New York, NY, 1987, ACM, pp. 274 – 284.

[8] Jon Postel, "Internet protocol," Tech. Rep. STD005/RFC 791, IETF, 1981.

[9] Whitfield Diffie and Martin E. Hellman, "New directions in cryptography," IEEE Transactions on Information Theory, vol. IT-22, no. 6, pp. 644–654, 1976.

[10] Neil Gershenfeld, "http://cba.mit.edu/projects/i0/i0.mpg," 2003.

[11] Larry L. Peterson and Bruce S. Davie, Computer Networks A Systems Approach Edition 3, Morgan Kaufman Publishers, San Francisco, CA, USA, 2003.

[12] Bruce Schneier, Secrets and Lies Digital Security in a Networked World, Wiley Computer Publishing, John Wiley and Sons, Inc., New York, NY, USA, 2000.

[13] Charlie Kaufman, Radia Perlman, and Mike Speciner, Network Security Private Communication in a Public World Second Edition, Prentice Hall PTR, Upper Saddle River, NJ, USA, 2002.

[14] Robert M. Metcalfe and David R. Boggs, "Ethernet: Distributed packet switching for local computer networks," CACM, vol. 19, no. 5, pp. 395 – 404, 1976.

[15] G. Tsirtsis and P. Srisuresh, "Network address translation - protocol translation (nat-pt)," Tech. Rep. RFC 2766, IETF, 2000.

[16] Echelon, "http://www.echelon.com/," .

[17] Karen R. Sollins, "Designing for scale and differentiation," in ACM SIGCOMM Workshop on Future Directions in Network Architecture, Karlsruhe, Germany, 2003.

[18] Ji Li, "Improving application-level network services with regions," Tech. Rep. LCS TR 897, MIT, 2003.