

RouterFarm: Towards a Dynamic, Manageable Network Edge

Mukesh Agrawal*, Susan R. Bailey†, Albert Greenberg‡, Jorge Pastor‡, Panagiotis Sebos‡, Srinivasan Seshan*, Kobus van der Merwe‡, and Jennifer Yates‡

Keywords

Network edge, Availability, Reliability, Manageability

1. INTRODUCTION

Planned maintenance is a fact of life in IP networks. Examples of maintenance activities include updating router software as well as processor upgrades, memory upgrades, installation of additional line cards, and other hardware upgrades. While planned maintenance is clearly necessary, it is also costly. Software upgrades, for example, require rebooting the router. Due to the time required to reboot the router, and then synchronize state (such as BGP routing information) with network neighbors, the upgrade process can yield outages of 10-15 minutes.

ISPs have typically architected their backbones to deal with the unavailability of backbone routers, by dual-homing each edge router to two different backbone routers. The backbones are then designed with sufficient redundancy in the form of alternate paths, available bandwidth, and processing capacity, to accommodate the unavailability of a single router and its associated links. The story for *access routers* – those routers which directly connect customers to the ISP network – however, is quite different. Because most customers maintain only a single connection to an ISP, access routers are a single point of failure for customers. Thus, every minute that an access router is unavailable is a minute of outage time for the customers connecting through (*homed on*) that router.

While planned maintenance is a primary source of downtime (and grief) in ISP networks – accounting for 30% of outages in a recent survey[1] – network operators face other threats to availability as well. Examples include hardware failures and disasters (whether natural or man-made). In

*Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213

†AT&T, One AT&T Way, Bedminster, NJ 07921

‡AT&T Labs – Research, 180 Park Ave, Florham Park, NJ 07932

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'06 Workshops September 11-15, 2006, Pisa, Italy.

Copyright 2006 ACM 1-59593-417-0/06/0009 ...\$5.00.

response to ISP demand, router vendors have offered solutions for some of these problems. For example, redundant processor cards and line cards [2] provide some protection against hardware failures. Similarly, in-service software upgrades, or ISSU [3], reduce the down time during upgrades (provided both old and new versions support ISSU).

We advocate an alternative path to a high available network edge, with two fundamental philosophical differences. First, rather than tackle high-availability by providing point solutions for each threat to availability, we propose an architecture which enables a single solution to the general edge-reliability problem. Second, we address reliability from the outside of routers, using a RAID or server-farm like approach, rather than from the inside, as by making each router highly available. These design choices yield a solution that 1) integrates IP and transport networking at the access edge, 2) is protocol independent, and 3) is vendor independent – thereby reducing the dependence of ISPs on vendor development schedules.

Our solution is an ISP edge architecture that we call *RouterFarm*. The key to our architecture is to remove the static binding between customers and access routers, and treat routers as interchangeable boxes – much like the servers in a server farm. As an example, in the RouterFarm architecture, we mask the unavailability of an access router *A* during planned maintenance as follows: 1) locate alternate router(s) to host the customers currently homed on *A*, 2) re-home the customers of *A* to the selected alternate routers, 3) perform the required maintenance, and finally, 4) re-home customers back to *A*, as desired.

In addition to providing a general edge-reliability solution, the RouterFarm architecture addresses a number of important weaknesses in today's static homing architecture. First, static homing makes it difficult to manage routing resources, such as router CPU and memory. The RouterFarm design allows customers to be reassigned to routers such that load is balanced and resources are used effectively. Second, RouterFarm greatly simplifies the supply chain management of a PoP. When routers fail in the current architecture, replacements must be sent immediately. RouterFarm allows for a more server farm-like approach to management, where all router boxes are viewed as interchangeable. As a result, when a router fails, the affected customers can be re-homed to any available router in any PoP, and the replacement may be dispatched whenever convenient.

Our design is conceptually simple. It does, however, reflect a fundamental change in the architecture of ISP-edge networks. Accordingly, there are significant challenges to

its realization, as we discuss in Section 3. In many scenarios, possibly the most pressing issue is the service interruption that the customer experiences during re-homing in the Router farm architecture. In planned maintenance scenarios, for example, if customers cannot be re-homed relatively rapidly, then it is unclear what availability gains this design will actually provide. Accordingly, the bulk of this paper focuses on an experimental evaluation of re-homing performance (Section 4).

2. ROUTERFARM OVERVIEW

Figure 1 illustrates a typical Tier 1 ISP network, and the re-homing process. The upper cloud in the figure represents the ISP’s IP network. In the IP network, customers connect to access routers, which in turn connect to *core* or *backbone* routers. Core routers direct the traffic from one access router to another. The transport network, underneath the IP network, provides the links that connect IP routers. The transport network may be implemented using an intelligent mesh optical network, SONET/SDH rings or a combination. For simplicity, we illustrate the transport network with rings.

In Figure 1, the customer to be re-homed is initially homed at a router within its metropolitan area. We refer to this router as its *initial router*. The customer’s link to its home router, its *access link*, runs over transport layer links 1, 2, and 3. After re-homing, the customer is homed at a spare router, which we refer to as the *target router*. The target router may be co-located in the same office as the initial router, or may be in a completely different region of the network. The IP-layer link now runs over transport links 1, 4, 5, 6, and 7. Note that, while this example illustrates a customer being re-homed to another region, in general, it is also possible to re-home a customer to a router within the same region.

Figure 2 details the re-homing process at the IP layer. The target router is assumed to be permanently connected to the ISP’s backbone network. Re-homing a customer is achieved by copying the customer-specific configuration from the initial router to the target router and moving the customer circuit (as detailed above). Once physical layer connectivity is established to the target router, routing protocols on the target router learn routes from the customer router, and re-advertise those routes to the ISP backbone. After this process completes for all customers, the initial router may be serviced without customer impact. When servicing is complete, customers may be returned to the initial router by repeating the procedure.

3. ROUTERFARM CHALLENGES

A number of challenges must be resolved to make RouterFarms practical. These include network design, resource mapping, service diversity, and re-homing time. We elaborate on these challenges in turn.

Network Design The RouterFarm architecture introduces an important new dimension to access network design: the placement of backup resources. As a single, fully-populated, access router costs millions of dollars, ISPs may be able to significantly reduce capital expenditure by sharing backup routers across PoPs. Such cost considerations argue for a single, centralized backup facility. However, such a design increases packet latency for re-homed customers,

1. Copy customer configuration from initial to target router
2. Reset BGP session between initial router and customer router
3. Disable customer interface on initial router
4. Reconfigure transport network
5. Enable customer interface on target router
6. Routing protocols learn and re-advertise customer routes

2: Re-homing Procedure

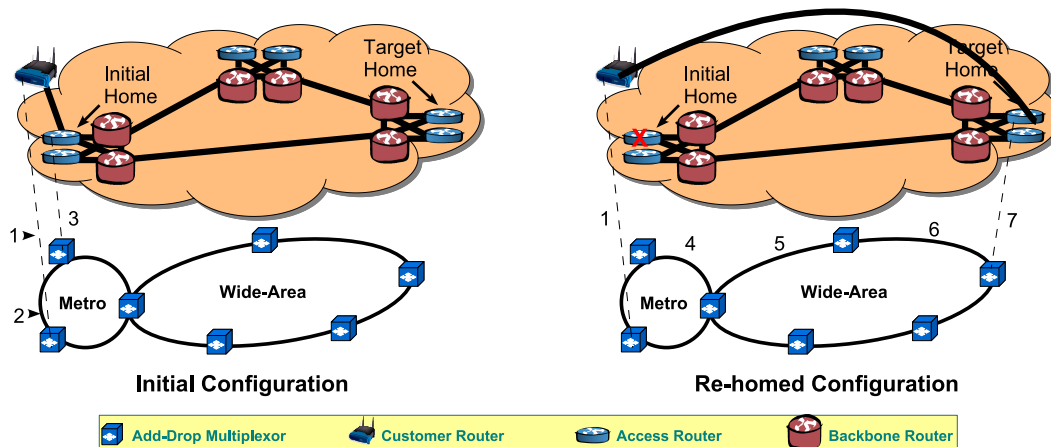
and requires large amounts of spare capacity in the transport network. In general, spare routers must be placed in a manner that balances performance goals, capital costs, and transport network utilization. Solving this optimization problem will require cost and performance models which capture these trade-offs.

Resource Mapping In order to re-home customers, we must assign them to target routers. Due to the diversity of router hardware configurations in IP networks, it is impractical to statically map customers from initial routers to target routers. Doing so would require large numbers of target routers, increasing capital expenditures. Resource mapping algorithms will thus need to cope with a variety of constraints, such as equipment compatibility (whether or not a router line-card supports the link technology of a particular customer), security requirements (such as segregation of VPN customers and Internet customers on separate routers), prioritization of customers (when there are insufficient resources to aid all customers), and available transport network capacity (which may vary due to failures at the optical layer).

Service Diversity There is considerable diversity in service offerings. For example, a customer may use static or dynamic (typically BGP) routing; may use bundled links, such as multi-link PPP; may or may not subscribe to class of service offerings; may or may not use VPN services, etc. One implication of such diversity is that the re-homing system must understand the configuration of varying services and service combinations and be able to copy the relevant configuration to the backup router. In commercial router operating systems, such as Cisco’s IOS, individual, customer-related configuration is sprinkled throughout the configuration file. It thus becomes a non-trivial task to identify all of the configuration components defining a customer’s service so as to copy the relevant configuration to the target router. Beyond this, some of the customer-specific configuration is shared across multiple customers. Thus, careful configuration modeling is required to ensure that customer configuration is complete on the backup router as each customer is re-homed and reverted. For example, when a VPN customer is re-homed, the target router must already be configured for VPN service, or the relevant configuration must be identified and established on the target router as part of re-homing.

Re-homing Time

Re-homing time is affected by factors such as the time taken to physically re-home the customer’s transport cir-



1: Tier 1 ISP network architecture, and the re-homing process. Solid lines indicate links within a layer, while dashed lines indicate links between IP and transport network equipment.

cuit, the time to establish routing sessions between a customer router and an access router, and the time required to propagate routing information throughout the network. Quantifying these factors, and proposing methods for reducing their impact, is a key goal in this paper, and the focus of Section 4.

4. EXPERIMENTAL VALIDATION

As noted in Section 1, for many scenarios, the utility of the RouterFarm architecture depends on the length of the service interruption that customers experience during re-homing. In particular, RouterFarm will be of limited value for improving planned maintenance outages if re-homing requires extended outages. Furthermore, to be useful for existing networks, RouterFarm must be feasible with stock commercial routers.

Accordingly, in this Section, we experimentally investigate the time required to re-home a customer, using a test-bed environment consisting of ISP-grade commercial routers. The goals of our evaluation are 1) to establish that re-homing is possible with unmodified commercial routers 2) to determine the length of the service interruption caused by re-homing, and 3) to identify possibilities for improvement.

While we have implemented re-homing for statically routed Internet service, BGP-routed Internet service, and MPLS VPN service, we limit ourselves here to results for BGP-routed Internet service, due to space constraints. To re-home a customer, we execute the re-homing procedure described in Figure 2.

Note that we focus here on re-homing the customer from an initial router to a target router. If it is necessary to return a customer to the initial router, as will typically be the case for planned maintenance, the re-homing procedure would need to be executed twice. Hence, the total outage observed by the customer would be twice the time reported herein.

4.1 Testbed Environment

Our testbed, illustrated in Figure 3, emulates the portions of an ISP network most relevant to the re-homing process. It consists of two customer routers (CE1 and CE2), and an ISP

network consisting of three access routers.¹ One customer router, denoted CE1, is connected to its initial home router, denoted PAR (primary access router) in Figure 3, over a simple reconfigurable optical network, implemented using a single optical cross-connect. This optical cross-connect is also connected to the target router, denoted BAR (backup access router), and is used to emulate transport layer re-homing.

In our setup, the PAR and BAR routers are Cisco GSRs (12406) whilst the customer and remote PE router (PE) are Cisco 7200s. The routers use default values for link layer and IP layer timers (such as those controlling neighbor failure detection). The optical cross-connect is a 2-D MEMS switch.

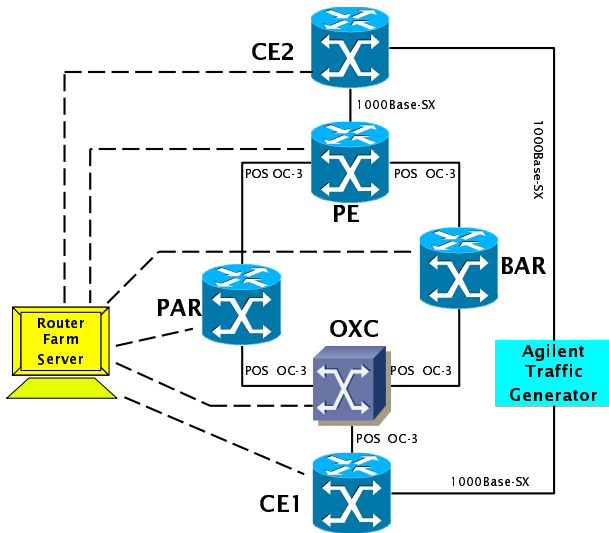
The customer re-homing is managed by our RouterFarm server, which is implemented on a Linux server. We implemented the re-homing software using three basic modules: (1) RouterFarm controller, which is responsible for making decisions and managing the overall procedure; (2) Router configuration module, which is responsible for communicating with the routers, reading the configuration from the initial router and copying over the relevant customer-related configuration over to the target router; and (3) Cross-connect controller, which communicates with our cross-connect to execute connection establishment and deletion. The controller is also used to capture experimental measurements.

Finally, we use an Agilent RouterTester as a traffic generator and protocol emulator. Specifically, the RouterTester was connected to the two customer routers, and was used to generate a large number of BGP routes. Customer routes are advertised by the RouterTester to CE1, while the “Internet” routes are emulated using advertisements to CE2. The impact of our re-homing procedure is measured using traffic generated by the RouterTester and then transmitted between CE1 and CE2. The outage time is calculated using the number of lost packets observed by the RouterTester.

4.2 Experimental Results

Herein, we evaluate the performance of the fast fail-over re-homing procedure. Our metric is the outage observed

¹Backbone routers are not incorporated here as they are expected to have negligible impact on re-homing performance.



3: Testbed


by the customer being re-homed, as measured by the time from which the first customer packet (either inbound or outbound) is dropped, until bi-directional traffic flow is restored (i.e., packets are delivered successfully both inbound and outbound).

Figure 4 depicts measured outage times as a function of number of routes advertised by the customer (CE1). The figure illustrates that with a small number of routes (10), we have a total average outage time of approximately 60 seconds, as measured over 10 runs. We scale this up to 5000 routes, which represents approximately 3% of all the routes in the Internet today — a massive number of routes for any single customer. Even at this extreme value, we observe that the outage time has not increased dramatically, with an average measurement of approximately 75 seconds.

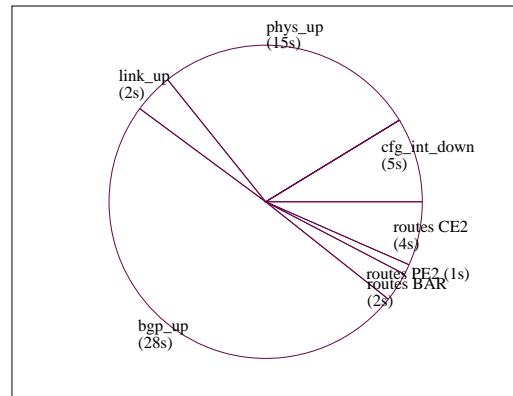
The total outage time consists of a number of different contributors. Re-homing the customer involves shutting down the interface on the primary router, moving the physical layer circuit from the PAR to the BAR, establishing a PPP connection up once the SONET layer is up, establishing a BGP session between CE1 and BAR, advertising customer routes from CE1 to BAR and then propagating those routes throughout the network, and advertising network routes from BAR to CE1. The route exchanges to and from the re-homed customer occur concurrently, with the routes being distributed from the customer to the Internet taking longer than the routes being learned by the customer. All other times are cumulative.

To understand the possibilities for improvement, we delve into the details of the average outage time for a customer advertising 10 routes. We divide the outage time into components, and present the results in Figure 5. We define these components depicted in the figure as follows:

cfg_int_down Time to shut down PAR’s customer-facing interface and confirm success. This is measured as the time between the first packet dropped and positive acknowledgment that the interface is disabled.



4: Outage time as a function of the number of routes advertised by the re-homed customer, for Internet service. Each data point gives the mean and 95%-ile confidence interval from 10 runs.



5: Breakdown of outage time for Internet service, focusing on a customer with a small number of routes (10).

phys_up Time taken to reconfigure the physical layer from the PAR to the BAR. This includes the time taken to cross-connect and for the routers to identify that the physical layer is up and operational.

link_up Time for the customer router and the target access router to establish a PPP connection once the physical layer is established between the routers.

bgp_up Time for the customer router (CE1) and target access router to establish an E-BGP session.

routes BAR Time for the target access router (BAR) to receive route advertisements from the customer router (CE1).

routes PE2 Time for the PE router to learn customer CE1’s route advertisements from the BAR.

routes CE2 Time for CE2 to learn the customer’s route advertisements from PE.

Figure 5 illustrates that the major contributors to our customer outage time are, in order of impact:

- the time required to establish an E-BGP session between the customer and BAR (**bgp_up**)

- the time involved in bringing up the physical layer between the routers (`phys_up`)
- time for router controller to configure interface down and confirm status (`cfg_int_down`)
- time to exchange routes between the remote PE and customer router (CE2) (`routes CE2`)

By analyzing the system log messages generated by the routers, we have found that the long `bgp_up` time is due to the BGP establishment time, which we define as the time between when the interface is up and operational (BGP protocol IDLE state), and when the routers have established a BGP session (BGP ESTABLISHED state). In our experiments, this time varied between 15 and 45 seconds; a result of a random timer used by the Cisco routers before initiating BGP session establishment. If this timer were configurable at the access router, tuning it down during the re-homing process would eliminate 28 seconds.

The approximately 15 second `phys_up` delay is caused by two factors: the time to disconnect and then re-connect the optical cross-connect (5 seconds) followed by the time for the routers to declare the physical layer up once it is fully cross-connected. This second time relates to the Cisco routers which wait 10 seconds following receipt of valid SONET signal before they declare the physical layer as available. Unlike the timer controlling BGP session establishment, however, eliminating this cost would require control at both, the ISP, and customer ends.

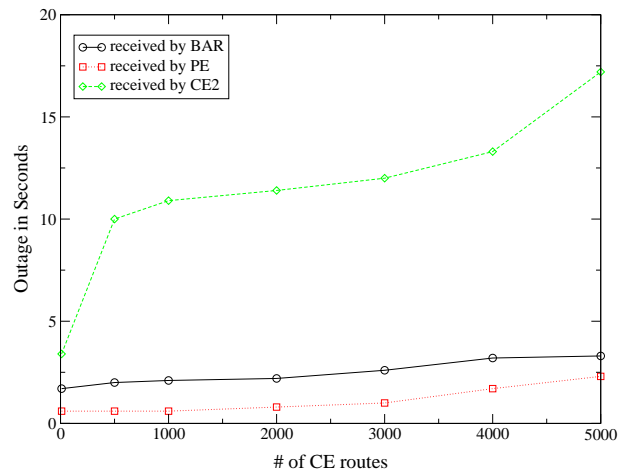
The `cfg_int_down` delay of 5 seconds is, to a large extent, due to our conservative implementation of our re-homing procedure. Specifically, our current implementation waits for confirmation that the interface has been successfully shut down on the PAR before proceeding with physical layer re-homing. Optimizing this process so that confirmation occurs in parallel with physical layer re-homing would reduce the `cfg_int_down` time, saving up to 4 seconds.

The final major contributor to the re-homing times is `routes CE2`, which is the propagation of routes from the remote access router (PE) to the remote customer, CE2. As illustrated in Figure 6, this is the dominant component of the overall scaling of outage time seen in Figure 4.

Investigation of the data reveal that the scaling is due to batching of route advertisements, and the `MinRouteAdver` timer. In particular, as the number of routes advertised by the customer increase, the probability that the routes are transmitted in a single batch decreases. Each batch of advertisements is separated from the previous batch by a delay of 30 seconds, caused by the `MinRouteAdver` timer. While this timer is tunable, we choose not to adjust it, because the adjustment would have to be made on every access router in the network, and leaving it in place during normal operation would increase BGP processing load during network transients.

4.3 Discussion

Our testbed evaluation demonstrates that re-homing is feasible on stock commercial equipment, and that a single customer can be re-homed in approximately 60 to 75 seconds (with default router timers), depending on the number of routes advertised by that customer. This time is an order of magnitude improvement over the commonly used approach of router reboots.



6: Component-wise scaling of outage time as a function of number of routes. Components which exhibit no scaling are omitted. Scaling effects are dominated by the time taken for routes to be received by the CE2 router.

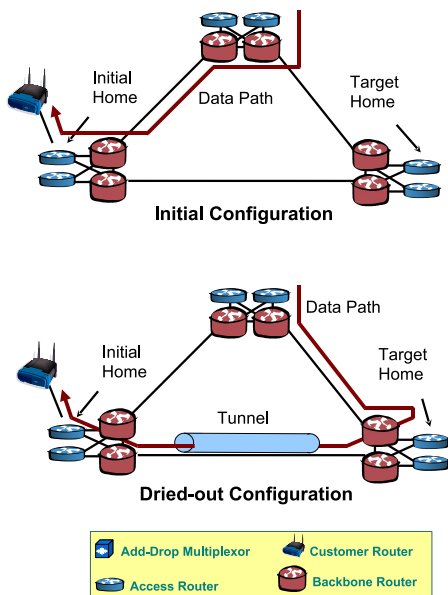
Our analysis of the components of re-homing times reveals that much of the delay is due to timers which are intended to protect the network from harm due to hazards such as link flapping. On commercial hardware today, these timers are non-configurable (such as for BGP session establishment and SONET link availability), and/or re-configuration at an access router alone cannot eliminate their costs (SONET link availability and BGP `MinRouteAdver`). This is not a significant concern in a paradigm where links go down quickly and come up slowly. The RouterFarm architecture, however, needs the ability to re-home customers with minimal outage times.

To meet this challenge, we propose an optimized re-homing procedure, which we call *Virtual Multi-Homing* (illustrated in Figure 7). Virtual Multi-Homing makes the re-homing outage time independent of IP-layer behavior by bringing up the link to the target router, before taking down the link between the customer and the initial router. More specifically, Virtual Multi-Homing avoids the `bgp_up` and `routes` costs by establishing the BGP session between the customer and the target router before tearing down the session between the customer and initial router. As the target router learns the customer routes, traffic for the customer is diverted through the target router, *drying-out* the paths through the current router [4]. After dry-out is complete, the customer access link is reconfigured, connecting the customer router to the target router. Note that, in addition to eliminating the effects of the IP-layer timers, Virtual Multi-Homing eliminates the dependence of outage times on route propagation times. Thus, we expect Virtual Multi-Homing to yield an outage time of 25 seconds (when used with SONET physical layer links), independent of the number of routes advertised by the customer.²

5. RELATED WORK

While the RouterFarm architecture can be used to improve resource management in a number of ways, this paper focuses on the availability improvements it provides. Here,

²We have not yet measured the outage time for this procedure experimentally.



7: Virtual Multi-homing Procedure

we briefly describe other efforts at improving availability of routers as well as other systems that have made use of transport network re-homing.

Highly Available Routers. As the reliability demands for IP networks have increased, router vendors have responded with various features for improving the availability of their routers. Examples include redundant control planes, redundant data planes [2], modular software [1], and in-service software upgrades [3]. Rather than tackle individual causes of failure, RouterFarm provides a general solution to the availability problem.

An alternative general purpose solution is redundancy. Examples include 1+1 redundancy at the ISP end [2], and multi-homing to multiple access routers. Either approach, however, requires substantial capital and operational expenditure, as it dedicates redundant hardware to each customer. In contrast, RouterFarm provides a lower-cost solution by sharing redundant/spare hardware across the network.

Re-homing. In [5], Sebos et al. propose re-homing for recovery from access router failures, and demonstrate feasibility using Linux-based software routers. In this paper, we propose a network architecture with re-homing as a first-class primitive. We demonstrate the suitability of re-homing in planned maintenance scenarios, and show that it is practical with existing commercial routers. Our experiments yield insights into opportunities for further optimizations.

Cisco’s Broadband Access Center provides a provisioning solution for broadband customers, such as cable modem users. As such, it could be used to support re-homing, via de-provisioning and re-provisioning customers. However, broadband and BGP/enterprise differ in important ways. First, ISPs have significantly more control over home users than enterprises. For example, an ISP might simply re-assign IP addresses when moving a home user. However, it must make the re-homing transparent to an enterprise and propagate routing updates to reflect re-homing changes. Second, enterprises are likely to employ a broader range of services than home users. Third, enterprises generally de-

mand greater reliability than home users.

6. CONCLUSION

In this paper, we tackle the problem of customer outages resulting from the planned maintenance of network access routers. We propose the *RouterFarm architecture* that moves customers from one access router to another and, thus, minimizes customer outages resulting from access router unavailability. In addition to planned maintenance, this design can also reduce the customer impact of other important tasks such as recovering from router failure and balancing load across access routers. The key observation that enables this architecture is that modern reconfigurable transport networks (e.g., intelligent optical networks) can dynamically reconfigure connectivity in short periods of time.

Our experiments demonstrate that customers can be re-homed with customer outages of 60-75 seconds (depending on the number of routes advertised by the customer), with stock ISP-grade commercial routers. This outage time is an order of magnitude reduction less than the commonly used approach of rebooting router for these tasks. Our results indicate that re-homing performance scales well with the number of routes announced by a customer.

To identify opportunities for improvement, we carefully break down the re-homing times. We find that over half of the re-homing time is due to timers, but that it is difficult to adjust these timers. Instead, we propose to mask the effects of IP-layer timers by *virtually* multi-homing customers during re-homing. In addition to masking the effects of these timers, Virtual Multi-Homing eliminates the scaling of outage times with number of customer routes. Thus, we expect that Virtual Multi-Homing will reduce outages to about 25 seconds (for SONET access links). While there remain a number of issues to resolve to make the RouterFarm architecture the truly flexible infrastructure we envision, the work in this paper illustrates the feasibility and attractiveness of our approach.

7. ACKNOWLEDGEMENTS

We thank Zihui Ge and Gary Flack for code used to configure routers during re-homing; Dan Pei, Jennifer Rexford, Olaf Maennel, and Anja Feldmann for answering our many questions about BGP; and Meeyoung Cha for comments on drafts of this paper.

8. REFERENCES

- [1] M. Reardon, “IP reliability,” *Light Reading*, March 2003.
- [2] Cisco Systems, “A brief overview of packet over SONET APS.” Cisco website, Document ID 13566, July 2004.
- [3] Cisco Systems, “Cisco IOS software: guide to performing in-service software upgrades,” March 2006.
- [4] J. V. der Merwe et.al., “Dynamic Connectivity Management with an Intelligent Route Service Control Point.” Sigcomm INM Workshop, September 2006.
- [5] P. Sebos, J. Yates, G. Li, D. Rubenstein, and M. Lazer, “An integrated IP/optical approach for efficient access router failure recovery,” in *Optical Fiber Communications Conference*, IEEE, 2003.