

RAIDER: Responsive Architecture for Inter-Domain Economics and Routing

Nirmala Shenoy
Rochester Institute of Technology
Rochester, NY, USA
nxsvks@rit.edu

Murat Yuksele
University of Nevada - Reno
Reno, NV, USA
yuksele@cse.unr.edu

Aparna Gupta, Koushik Kar
Rensselaer Polytechnic Institute
Troy, NY, USA
{guptaa, kark}@rpi.edu

Victor Perotti
Rochester Institute of Technology
Rochester, NY, USA
vperotti@saunders.rit.edu

Manish Karir
Merit Networks
Detroit, MI, USA
mkarir@merit.edu

Abstract— Multi-owner structure shaping inter-domain operations is arguably the most important factor determining the end-to-end performance in the current Internet. Financial sustainability of Internet service provisioning has significantly changed the way Internet grows and all the other business sectors using the Internet as infrastructure. Further, scalability of BGP routing table sizes is becoming a stressing problem. Architectural solutions providing ways to better inter-domain economics and more scalable inter-domain routing protocols are of crucial importance. In this position paper, we present a new inter-domain communications architecture to address routing scalability by leveraging the inherent structure in ISP topologies and using a simplified addressing method. To manage risks in costly backbone business and open the doors for realizing higher quality end-to-end services, our architecture uses protocol-level techniques to increase operational granularity of inter-ISP market with automated service level agreements (SLAs).

Keywords - floating cloud internet network model, contract-switching, inter-domain economics.

I. INTRODUCTION

Among the many challenges facing the current Internet are the tremendous growth in the number of networks and their sizes. The resulting scalability issues have created serious setbacks in terms of security, management and control. Logical address allocation and longest-prefix-matching have led to complex route lookups, increasing the routing table sizes to unmanageable limits [1] with a high imbalance in handling of routing information as noted by the sizes of routing tables in the core routers. In essence, the current inter-domain routing is facing difficulties due to the increasing complexity of routing in policy-based ISP structures and practices.

The ISP industry continues to evolve through this growth of the Internet. The economics and profitability of Internet service provisioning have been critical driving forces in the Internet's great success. However, the commoditization of backbone bandwidth has caused ISPs to suffer financially [17]. Thus, inter-domain economics is a key challenge to address in the future Internet. As it became clear that simply providing access will not constitute a core business by itself, providers created organic ad-hoc solutions through new pricing models.

The current Internet service market's growth and low potential profit pose considerable risk to the ISPs. In this position paper, we suggest there could be two reasons why the

current architecture is not financially sustainable. First, *economics of inter-domain interactions are not flexible* enough, as SLAs take too long to establish and teardown and involve tedious paperwork. As a result, ISPs cannot manage risks associated with QoS investments, and are not sure if their intra-domain investments will produce any returns from applications that need inter-domain end-to-end (E2E) quality. Second, *E2E value is unrecognized* due to the multi-ISP structure of the Internet. Though the multi-ISP structure is key to a healthy market, it currently comes with unrealized collaborative value. The E2E QoS requires all participating ISPs to become more synchronized at lower operational time-scales than allowed by current SLAs.

Motivated by the above architectural issues, we present RAIDER – a Responsive Architecture for Inter Domain Economics and Routing, a clean slate architecture to enable robust Internet performance for users. The architecture supports continued growth, provides better inter-domain economic opportunities and offers more scalable inter-domain routing protocols so that ISPs can collaborate to manage the risks and achieve returns based on their investments. The primary goal of RAIDER is to make the internetworking architecture highly flexible and scalable both in technical and economic dimensions.

The rest of the paper is organized as follows: In section II we highlight some published work closely related to Internet architectures that target routing scalability, contract-switching concepts and inter-domain economics. Section III summarizes RAIDER's vision and key components. In section IV, we discuss the underlying concepts of a Floating Cloud Tiered (FCT) internetworking model and some of its features. Section IV provides details of the contract-switching paradigm, followed by discussions on potential benefits of applying it over the FCT internetwork model. Section V discusses the economics related to contract-switching and the FCT model. We conclude in section VI with a summary of contributions.

II. RELATED WORK

The scalability problems faced by the current Internet routing protocols has been of increasing concern as the routing tables in the core routers experience unmanageable growth rates. Several research efforts were directed towards solving this problem. However, most proposed solutions were constrained by the fact that they had to operate in the existing internetwork

architecture and with the well-established Internet Protocol, its forwarding and routing mechanisms and logical addresses. *Hybrid Link State Protocol* [2] availed the natural hierarchy in the AS structures to provide a solution to excessive route churns through route information aggregation within an AS hierarchy. The *New Inter-Domain Routing Architecture* [3] used a provider-rooted hierarchy and showed improvements in the number of forwarding entries and convergence times. A routing research group at *Internet Engineering Task Force* [4] proposed ‘core-edge separation’ to temporarily solve the routing table size problem by ‘address indirection’ or ‘Map-and-Encap’ to keep the de-aggregated IP prefixes out of the global routing table. *Routing Architecture for Next Generation Internet* [5] uses locator/identifier split ideas with ‘routing scalability’ as one goal. *Routing on Flat Labels* [6] uses flat routing to separate location and identity for both inter and intra-domain routing. Enhanced *Mobility and Multi-homing Supporting Identifier Locator Split Architecture* [7] is based on a hybrid design of ID/locator split and core-edge separation to address routing scalability among others.

The European Future Internet Portal (FP7) funded projects focuses on several key areas, and routing scalability is one of them. Likewise, the National Science Foundation in the United States initiated the clean slate Future Internet Design (FIND) program [8] and funded several projects [9], some of which tackled the routing scalability issue.

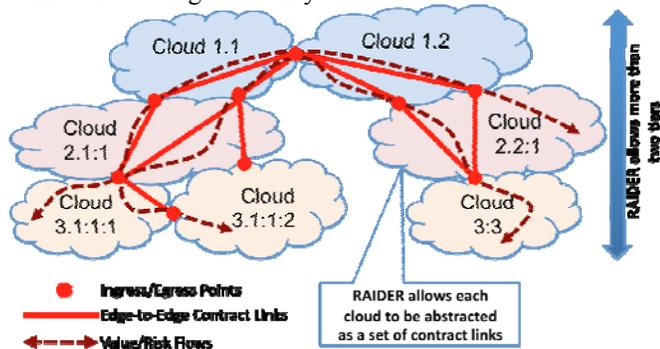


Figure 1. RAIDER supports Contract Routing over an FCT internetwork Model: Multiple E2E contracts can be automatically established via concatenation of several single-hop contracts.

The literature on price discovery, contract design and implementation for Internet services is vast. We focus here on our past, related work, which will be utilized in the Contract Switching paradigm. We have developed pricing framework for expected bandwidth and loss-guarantees based contingent claims for intra-domain contracts [12]. The spot pricing framework for intra-domain contracts is designed for cost recovery while maximizing the total surplus. The framework also accounts for the competitiveness structure of the market. The loss-guarantee contingent claim pricing framework utilizes the concept of state-price density for pricing the risk inherent in the service guarantee for loss of customer-data through the contract duration. Utilizing the intra-domain frameworks, we have developed solutions for end-to-end contracts under the link-state and path-vector approach to end-to-end contract creation [13] For temporal risk management, we have also developed a pricing framework for a specific kind of forward contract for the risk in bandwidth demand [12]. A bailout clause embedded in the forward contracts allows the provider

to be relieved of its obligation should the network be significantly congested at the time of maturity of the contract.

III. RAIDER

RAIDER’s key goal is to make “*the internetworking architecture highly flexible and scalable both technically and economically, so that it can be responsive to future needs of the network’s users and providers.*” We believe that being adaptive and responsive to emerging needs is a crucial characteristic of the internetwork architecture. Further, it is generally not possible to predict future technical and economic innovations, and thus, “flexibility and scalability” should be key ingredients of any architecture. As shown in Figure 1, RAIDER aims to provide this flexibility and scalability via two novel components:

Technical Responsiveness – Floating Cloud Tiered (FCT) Model: RAIDER introduces FCT internetwork model that allows portions of the Internet to be abstracted as tiered autonomous network clouds that can easily associate and disassociate with other network clouds at different tiers to achieve rapid configurability.

Economic Responsiveness – Contract-Switching Paradigm: RAIDER uses “edge-to-edge contracts” as the building block that allows ISPs to define their own local policies and embed “values” as contracts expressing service and price levels they are willing to accept for their network.

In order to fully utilize the advantages offered by Contract-Switching it is imperative that interconnecting networks have the ability to rapidly connect or disconnect with other networks. FCT internetwork model allows precisely this ability to rapidly attach and detach to different network clouds. Thus, these two components complement each other to achieve a highly flexible and responsive architecture.

IV. THE FLOATING CLOUD TIERED (FCT) MODEL

The proposed Floating Cloud Tiered (FCT) internetwork model is different from any prior work in this area as for the first time the tiered ISP topological structure is being leveraged to solve the routing scalability problem. We introduce a new tier based addressing scheme to support tiered structure based packet forwarding, which introduces a new form of tier based address aggregation, which can have a very significant impact on routing scalability.

The tiered topological connections existing among ISPs are well known. In this tiered structure there are several tier 1 ISPs that connect the tier 2 ISPs, which connect tier 3 ISPs and so on. In Figure 2, we capture one such simplified ISP topology, where ISPs A and B are tier 1 ISPs, while ISP C is a tier 2 ISP and ISP E, a tier 3 ISP. In the figure, each ISP is shown as a network cloud. The broad arrows are indicative of multiple connections between ISPs. We use the relative positional properties among ISP clouds across the tiers to introduce a structured packet forwarding paradigm through the concept of a ‘tiered cloud address’ (CloudAddr) assigned to each ISP cloud. The CloudAddr is function of the tier in which the cloud resides and the clouds that it is associated with. Packets will be forwarded to the appropriate cloud based on the CloudAddr. Clouds can associate to or disassociate from the network topology via acquisition or release of CloudAddrs.

The forwarding and routing within the cloud can adopt

either the tiered approach, or any other mechanism. We would thus decouple the inter- and intra-cloud dynamics, such that a change in CloudAddr will not impact the internal structure or addresses within a cloud. This decoupling will allow for easy movement (*floating*) of network clouds across tiers offering dynamic flexibility in the network topology.

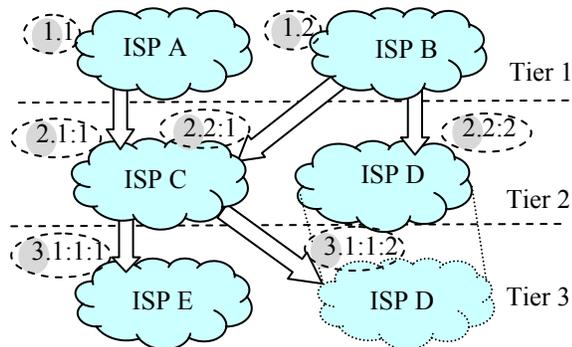


Figure 2. The Concept of Floating Clouds across Tiers

To explain the CloudAddr and its significance we assign CloudAddrs to the different network clouds in Figure 2. ISP A, the first cloud in tier 1 has CloudAddr = 1.1 following a format TierValue.MyCloudID. ISP B similarly has a CloudAddr 1.2. ISP C is at tier 2 and connected to ISP A (via CloudAddr 2.1:1) and ISP B (CloudAddr 2.2:1) simultaneously, where the CloudAddr format is TierValue.ParentCloudID:MyCloudID. ISP D is connected to ISP B through CloudAddr 2.2:2; it can however decide to change its service provider to ISP C by moving to tier 3, or remain simultaneously connected to ISPs B and C, by using two CloudAddrs at different tiers.

Packet Forwarding: The first field in the CloudAddr is the TierValue, shown inside a shaded circle, will be used to forward the packets across clouds. The decision to forward in a particular direction, up the tiers, down or sideways, depends on the relative positions of the source (SRC) and destination (DST) clouds in the tier structure and the links between sibling clouds in a tier. To illustrate packet forwarding we use a simple example from Figure 2; if the SRC cloud is 3.1:1:1 and DST cloud is 2.2:2; the source compares the two addresses to determine the tier of a common parent (or grandparent) cloud for the SRC and DST. In this case, it will be tier '1' as there are no address components after the TierValue, common in the SRC and DST CloudAddrs. The remaining fields in the DST address (after the common part) are then appended to the TierValue to provide the forwarding address; in this case it will be 1.2:2. All intermediate clouds between 3.1:1:1 and 1.2 will forward the packet upwards, using the tier value until it reaches cloud 1.2. This cloud then identifies that the destination is at tier 2 because of the two address fields following the tier value, replaces the TierValue with 2 and forwards the packet down to the DST cloud. However if there were a link between ISPs C and D), the routers in ISP C could be made aware of the sibling cloud connection and can then directly forward the packet to cloud 2.2:2.

The Nesting Concept: Inside of an ISP there are several Point of Presence (POPs). Each POP comprises of backbone, distribution and access routers. The backbone routers are

primarily for connecting to other backbone routers in other POPs. An interesting observation to be made at this point is, that a tiered structure can be noticed inside of an ISP POP. The set of backbone routers can be associated to tier 1 within the POP. The distribution routers connected to a backbone router can be associated to tier 2. Lastly the access routers can be associated to tier 3.

Nested tiered address inside an ISP POP: We illustrate the tiered address allocation inside a POP in the AT&T network in the US. We can assign a Cloud Addr 1.1 to the Seattle POP, 1.2 to the NY POP and 1.3 to the Chicago POP and so on. There are 17 distribution clouds in Seattle POP. The CloudAddrs for the Seattle distribution clouds hence could be assigned as 2.1:1, 2.1:2 and so on with the last CloudAddr being 2.1:17. The access routers connected to the first distribution cloud would have CloudAddr that started with 3.1:1:1. The packet forwarding across POPs in the AT&T network thus can follow a similar process outlined under Figure 2. If the packet has to be forwarded to another ISP, then the globally visible CloudAddr for AT&T will be used. This would require the internal address to be nested behind the globally visible address. Note, that unlike tunneling, in this approach we are not encapsulating one address inside another, but an outer cloud address will be prepended when the packets have to be forwarded outside of a given cloud. This type of tiered address allocation can be extend a stub AS or network, where it is common to have border routers that connect to distribution routers which may connect the sub-networks directly or via access routers.

Summarizing, in the FCT model, an ISP can be considered to be a cloud. The ISP cloud further has several POP clouds. Each POP cloud can have a 'backbone cloud' and several 'distribution clouds' and 'access clouds'. The routers in the 'access cloud' of a POP could be part of the 'backbone cloud' in a stub network, which can have its own distribution and access clouds. A 'cloud' thus is a set of routers that have a specific purpose i.e. for backbone routing and distribution routing amongst others – leading to an abstract 'cloud' concept that can be defined to any level of granularity. We call the property of defining a cloud(s) inside another cloud as 'nesting'. The above explanation is intended to highlight some attributes of the FCT internetwork model that demonstrate flexibility in cloud sizing and connections to network clouds to overcome the structural rigidity in the current Internet.

V. CONTRACT-BASED INTER-DOMAIN ROUTING

The current SLA establishment process in the Internet is too rigid and coarse in the *temporal*, *spatial*, and *economic* dimensions. The SLAs require tedious paperwork and are established over long time periods, i.e. months to years, which results in inefficient resource usage as demands vary dynamically at much finer time granularities. Such temporal granularity has been successfully implemented in other markets such as power grid [18]. Another technical limitation of current SLAs is their point-to-anywhere (pt-any) design as managing E2E QoS becomes impractical with such pt-any SLAs among ISPs. The other extreme, i.e. point-to-point (pt-pt) design between end-systems or access routers/providers, is known to have scalability limitations as was observed in Asynchronous Transfer Mode (ATM). Lastly, current SLAs

lack economic tools (e.g., insurances, money-back guarantees, and bailouts) that have been widely used in many other markets. With limited economic tools in SLAs, ISPs cannot manage risks involved in their high-cost investments.

We believe that the solution to these issues lies in allowing dynamic, flexible pt-pt *contracting* between the *edge/peering points of ISPs* [19]. Such *edge-to-edge* “contracts” would allow as much flexibility in spatial composition as an E2E pt-pt design would provide, while preserving scalability of implementation. In the simplest form, a contract represents an SLA between the seller and the buyer of Internet access, transit or delivery services. To address the temporal and economic limitations of traditional SLAs, our notion of contracts offer more flexibility in terms of the following components: (i) *performance* component (ii) *time* component, (iii) *financial* component. The performance component relates to the specification of the QoS metrics of the contract. For example, a contract will usually be associated with an agreement on the baseline bandwidth, possibly with additional metrics like delay, loss and jitter. The time-scale over which these agreements should be met must also be specified; this time-scale specification is also useful for contract monitoring, as discussed later. The time component relates to the time duration for which the contract is negotiated. It can be at the current time (*spot* contract), or for the future but negotiated in advance (*forward* contract) [14]. Finally, the financial component relates to issues of how the contract is priced and terms of payback in case the contract is violated.

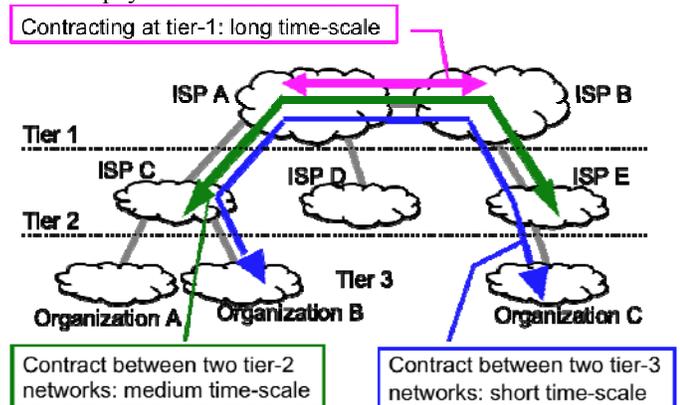


Figure 3. Contract Routing over an FCT internetwork Model: Multiple E2E contracts can be automatically established via concatenation of several single-hop contracts.

A. Contract-Switching Paradigm

With contracts as basic inter-domain building blocks, we introduce a contract-based inter-domain routing (or *contract routing*) framework over the highly flexible floating clouds described earlier. The key goal of contract routing is to provide the tools to bundle intra-domain services to compose E2E services appreciable by users. Such an inter-domain routing architecture requires a paradigm shift to “contract-switching” [20] by generalizing the existing packet-switching paradigm. The essence of contract-switching is to use contracts as the key building block for inter-domain networking, where each contract in its basic form refers to a service agreement between the neighboring ISPs made at their peering points. At a broad level, both packet-switching and

(virtual) circuit-switching can be viewed as special cases of contract switching. For example, packet size can be viewed as a short-term bandwidth contract (between a router with its previous hop) whose duration is proportional to the transmission time of a packet. Virtual circuit-switching on the other hand, is a contract between the network and the end-user (end-system) for bandwidth and other QoS parameters, whose duration is proportional to the lifetime of the virtual circuits. There are however two key facts that are worth noting. Firstly, the contracts that we consider are made between peering ISPs or between the end-systems and their access provider ISPs. Secondly, contract-switching associates “value” to each bit in the network and focuses on managing “value flows”, rather than treating each bit as it is and focusing only on managing “data flows”, as traditionally done in packet-switching and virtual circuit switching

B. Contract Routing and Automated E2E SLA Establishment

The composition of end-to-end inter-domain contracts poses a major research problem that we formulate as a “contract routing” problem by using single-domain contracts as “contract links”. We abstract the point-to-point QoS services provided by each ISP as a set of “overlay contracts” each being defined between peering points, i.e., ingress/egress points of the ISP. Given contracts between peering points of ISPs, the “contract-routing” problem involves discovering and composing end-to-end QoS-enabled contracts from per-ISP contracts. As we describe below, it is possible to attain this with BGP-style path-vector routing protocols that compose end-to-end “contract paths”. To provide enough flexibility capturing more dynamic technical and economic behaviors in the network, it is possible to design contract routing that operates at short time-scales, i.e., *tens of minutes*. This time-scale is reasonable as current inter-domain BGP routing operates with prefix changes and route updates occurring at the order of a few minutes [15]. Further, an ISP might want to advertise a spot price for an edge-to-edge contract to a subset of other ISPs instead of flooding it to all. Similarly, a user might want to query a specific contracting capability for short-term and involving various policy factors.

Just like BGP composes paths, E2E contract paths can be calculated in an on-demand lazy manner. In our design, each ISP has the option of initiating contract path calculations by advertising its contract links to its neighbors. Depending on various technical, financial, or policy factors, those neighbors may or may not use these contracts in composing a two-hop contract path. If they do, then they advertise a two-hop contract path to their neighbors. This path-vector composition process continues as long as there are participating ISPs into the contract paths. Users or ISPs receiving these contract paths will have the choice of using them or leaving them to invalidation by the time the contract path term expires.

The QoS metric appropriate for a contract and the time duration (timescale) over which a contract is negotiated (monitored) will depend on the tier of the cloud in the FCT model at which the contracting occurs (Figure 3). In general, contracting at the lower tiers (i.e. higher tier values) of the FCT network cloud will involve shorter time-scales and durations than those at upper tiers (i.e. lower tier values). This also becomes necessary to preserve scalability of establishing and managing the contracts. For example, the appropriate QoS metric for an end-user may be file transfer delay or latency,

measured over a file or over data transfers over a short period, say an hour. However, the relevant QoS metric for the contract established by a Tier-2 ISP with its corresponding provider, a Tier-1 ISP, can be average bandwidth over a one-month period.

The tiered structure of our proposed FCT network model can be leveraged in composing the E2E contract paths. Naturally, higher tier clouds will connect higher capacity links, though with potentially higher number of E2E contract paths traversing them. A key issue in composing E2E contract paths with QoS options has been the enormous number of E2E path possibilities. We can simplify the contract path composition problem by forcing rules based on the tiered structure of the FCT model. However, imposing such rules will result in a significant reduction in the number of E2E paths. The open research problem is to find the right set of rules such that the set of available E2E high-quality paths is not undesirably reduced. The other research question of interest is how E2E contracts with a QoS metric that is of interest to the end-user can be composed using the inter-cloud contracts which can differ in their SLA agreements in terms of their QoS metrics and contracting time-scales, as discussed earlier.

C. Contract Monitoring, Reporting, and Verification

An important architectural part of the future Internet with contract-switching will be monitoring, reporting, and verification of ongoing E2E contracts. While SLA verification is a major issue in the Internet as well, it will play a bigger role in a contract-switched future Internet. Therefore, we envision that technologies at each tier of the FCT model must actively involve in SLA verification. Depending on the contracting time-scale and financial responsibilities of each contracting entity, there needs to be various software and hardware support for contract verification and monitoring at each floating cloud involved in contracting. Under-performing contracting entities will need to be penalized according to these monitoring results and the terms of the contract. These monitoring agents will operate at time-scales in correlation with the contracting time-scales, and necessary reporting functions will be incorporated into the associated signaling protocols. For example, contract monitoring at higher tiers can be done more appropriately over pre-agreed, fixed, relatively longer durations of time like weeks or months. Contracts with money-back guarantees and significant costs, typically at macro-level with long-term commitments, will certainly require such support. For the current Internet, legacy SLA technologies address this issue by placing hardware agents/verifiers at the ISP peering points.

While there has been significant body of work in E2E QoS monitoring [21][22], majority of the focus has been on per-packet and single-contract schemes with no explicit consideration of multiple contracts involved in an E2E service. Our contract-based architecture introduces new problems in E2E QoS verification such as: (i) identification of intermediate clouds failing the E2E contract's QoS targets, (ii) appropriate accounting such that failing clouds can be charged the cost of any money-back guarantees to the end-customer, and (iii) measurements robust to gaming by intermediate clouds.

VI. INTER-DOMAIN ECONOMICS

The consensus emerging from the ACM SIGCOMM workshop, RIPQoS (e.g., [10][11]) leads us to conclude that the future problems for QoS in the wired domain would transcend

the technical domain. However, the economic solutions have to rely on facilitating technical advancements. Economic flexibilities and possibilities for creation of new revenue streams will require the mechanisms designed to be sufficiently dynamic and scalable. Moreover these solutions, both economic and technical, have to be available in intra-domain settings, before they become useful means for inter-domain interaction and collaborations.

Transparency: In the FCT model, through prices charged for services, some of an ISP's internal network information can be reverse-engineered by the outside world, which can be beneficial for the overall health of the Internet service ecosystem. For instance, managing and reducing harmful information asymmetry is a desirable goal for enhancing social welfare and optimal functioning of the Internet service ecosystem. Towards this goal, the extent of inclusion of floating clouds in an ISP structure may be controlled to make a certain level of information regarding an ISP become available to the outside world, including the end-users, and thereby limit the potential of excessive market power of certain dominating ISPs. Provision of floating clouds in the FCT model creates options for how an ISP constructs pt-pt contracts for its domain in order to collaborate with neighboring and other ISPs. The real-options methodology [24][25] can be utilized to incorporate this flexibility in the intra-domain contract definition and pricing. Dynamically configurable clouds combined with contract switching create enhanced collaborative opportunities with other ISPs, where the goal is to take advantage of these opportunities without hitting the scalability limits.

Depending on the performance guarantee, contracting time-scale and financial responsibilities of the provider, the software and hardware support for verification and monitoring of the contractual terms adequately empower the end-user and establish best practices for collaboration between providers.

The well-defined intra-domain contracts are the building block for inter-domain business models and financial settlement methods (i.e. monetary flows to compensate for QoS traffic flows), and for flexible risk management mechanisms (including contingent claims contracting, insurance, money-back-guarantees). End-to-end contracts are compositions of per-ISP contracts (i.e. composition over space), and may be specified at longer time-scales (i.e. composition over time), and may be a derivative contract (i.e. functions of underlying end-to-end spot contracts). Both link-state and path-vector style contract routing protocols can be applied, the former for long-term proactive investment decisions by providers and the latter for short-term on-demand investment decisions by providers and users. Hybrid contract routing schemes can also be used by users wanting to compose an end-to-end contract path with some portion of the end-to-end capacity being composed over long-term contracts and the rest being composed by short-term more dynamic contracts.

Risk Management: Tools for risk management for intra-domain and inter-domain risks are crucial. Risk management helps reduce volatility in earnings and gives the provider stability in its cash flows providing assurance for being open to future investments opportunities. Risks pertain to both performance characteristics of the contracts and their temporal properties. Contingent claims or derivative contracts are

designed in the financial markets as a mechanism to transfer appropriately defined risk to the party that is most capable/willing to bear the risk. The major risk underlying a performance guarantee in a contract is that there is a non-zero probability that the provider will not be able to deliver on the guarantee. A contingent claim pricing methodology assigns a price to this risk, or in other words, assigns a price to the scenario that the provider will deliver on the performance guarantee. The guarantee given to the customer comes at the price of risk of the favorable scenario. If the price of all the scenarios is 1, and Z defines the favorable scenario for the performance guarantee, the price of the guarantee is $E_Q[Z] < 1$, where 'Q' is an appropriate probability measure. ΩZ is the unfavorable scenario, the risk of which the provider would bear under this performance guarantee, since the price of the guarantee corresponds only to the favorable scenario. Hence contingent claims help the provider and the customer share the underlying risk of contracts.

Risks underlying network service provision are significant, arising from stochastic demand, network failures, moral hazard between ISPs, and malicious attacks. ISPs need to provide network services, as well as safeguard their own, and their end-users, interests in presence of these risks. Advanced contracting terms, such as bailout forward contracts [14], can be developed to manage some of these risks. Similar and additional risk sharing, risk transfer and risk management mechanisms can be developed for intra-domain risks under the FCT model. For instance, a money-back guarantee essentially serves as insurance for a service, for as long as the risk of service failure is not realized, the customer or end-user pays the fee (premiums, in insurance terminology) for the service. However, if the risks defined under the money-back guarantee are realized, the provider must compensate for the loss. All the outstanding contracts, with or without money-back guarantees, leave a certain level of aggregate risk exposure of a provider. The provider may seek to protect itself through insurance contracts for this aggregated risk. Appropriate contract design is required based on an analysis of the aggregate risk profile for its insurability. Risk sharing and transfer mechanisms for intra-domain risks can be utilized for E2E risk management.

The new revenue streams arising from a greater variety of contract types, the more precise and flexible management of risks, and the response to users' growing and evolving needs will create incentives for the management of a provider's infrastructure and its upgrades. Such incentives created for all network stakeholders can encourage industrial innovation while improving Quality-of-Experience (QoE) for end-users.

VII. SUMMARY AND FUTURE DIRECTIONS

We outlined a future architecture, RAIDER, to address the routing scalability issues and inter-domain economics in the current Internet. The proposed architecture focuses on making the internetwork design more responsive to (i) inter-ISP structuring via a tiered addressing scheme that allows ISPs to swiftly form their domains as "floating clouds", and (ii) economic incentives via "contract routing" protocols to automatically compose end-to-end inter-cloud contracts. RAIDER follows the "contract-switching" paradigm as a generalization of the packet-switching paradigm, and promotes usage of edge-to-edge contracts as inter-ISP tussle

units that promote economic incentives and risk sharing. RAIDER's capability to break the rigidity in the current Internet architecture will aid in flexible service provisioning. Future directions may include extending the proposed inter-domain contracting mechanisms to enhance the market power of consumers, thereby enabling more industrial innovation.

ACKNOWLEDGMENTS

This work is supported in part by the NSF under awards 0832008, 0831957, and 0831830.

REFERENCES

- [1] <http://www.internetworldstats.com/stats.htm>
- [2] L. Subramanian, M. Caesar, C. T. Ee, M. Handley, M. Mao, S. Shenker, and I. Stoica, "HLP: A next generation inter-domain routing protocol," in *Proceedings of ACM SIGCOMM*, pp. 13–24, 2005.
- [3] X. Yang, "NIRA: A new Internet routing architecture," *Proceedings of the ACM SIGCOMM Workshop on Future Directions in Network Architecture*, pp. 301–312, Karlsruhe, Germany, August 2003.
- [4] Internet Research Task Force Routing Research Group Wiki page, 2008. <http://trac.tools.ietf.org/group/irtf/trac/wiki/RoutingResearchGroup>
- [5] X. Xu and R. Jain, "Routing Architecture for the Next Generation Internet (RANGI)," <http://tools.ietf.org/id/draft-xu-rangi>, Mar. 2009.
- [6] M. Caesar, T. Condie, J. Kannan, K. Lakshminarayanan, and I. Stoica, "ROFL: Routing on flat labels," *SIGCOMM CCR* vol. 36, no. 4 pp. 363–374, August 2006.
- [7] J. L. Pan, R. Jain, S. Paul, M. Bowman, X. Xu, and S. Chen, "Enhanced MILSA Architecture for Naming, Addressing, Routing and Security Issues in the Next Generation Internet," *Proceedings of IEEE ICC*, Dresden, Germany, June 2009.
- [8] NSF NeTS Future Internet Design Initiative, <http://www.nets-find.net>
- [9] <http://cse.wustl.edu/Research/Lists/Technical%20Reports/Attachments/891/3SURVEY.pdf>
- [10] B. Teitelbaum and S. Shalunov, "What QoS research hasn't understood about risk," *P. of ACM SIGCOMM Workshops*, pp. 148–150, Aug 2003.
- [11] L. Burgstahler, K. Dohler, C. Hauser, J. Jahnert, S. Junghans, C. Macian, and W. Payer, "Beyond technology: The missing pieces for QoS success," *P. of ACM SIGCOMM Workshops*, pp. 121–130, Aug 2003.
- [12] A. Gupta, S. Kalyanaraman, and L. Zhang, "Pricing of risk for loss guaranteed intra-domain internet service contracts," *Computer Networks*, vol. 50, pp. 2787–2804, 2006.
- [13] A. Gupta and L. Zhang, "Pricing for end-to-end assured bandwidth services," *International Journal of Information Technology & Decision Making*, vol. 7, no. 2, pp. 361–389, Jun 2008.
- [14] W. Liu, H.T. Karaoglu, A. Gupta, M. Yuksel, and K. Kar, "Edge-to-edge bailout forward contracts for single-domain internet services," *Proc. of IEEE International Workshop on Quality of Service*, 2008.
- [15] R. Teixeira, S. Agarwal, and J. Rexford, "BGP routing changes: Merging views from two ISPs," *ACM SIGCOMM Computer Communication Review (CCR)*, October 2005.
- [16] L. G. Roberts, "A radical new router," *IEEE Spectrum*, vol. 46, no. 7, pp. 34–39, July 2009.
- [17] A. M. Odlyzko, "Internet pricing and the history of communications," *Computer Networks*, vol. 36, pp. 493–517, 2001.
- [18] S. Stoft, *Market Operations in Electric Power Systems*, Wiley, 2002.
- [19] M. Yuksel and S. Kalyanaraman, "Distributed Dynamic Capacity Contracting: An overlay congestion pricing framework," *Computer Communications*, vol. 26, no. 13, pp. 1484–1503, January 2003.
- [20] M. Yuksel, A. Gupta, and S. Kalyanaraman, "Contract-Switching Paradigm for Internet Value Flows and Risk Management", *Proceedings of IEEE Global Internet Symposium*, April 2008.
- [21] C. M. Leung and J. A. Schormans, "Measurement-based end to end latency performance prediction for SLA verification," *Proceedings of IEEE ICCCN*, pp. 412–417, 2002.
- [22] "Advanced architecture for INTER-domain quality of service MONitoring, modelling and visualisation," <http://www.ist-intermon.org>.
- [23] B. Stiller, D. Hausheer, T. Hoffeld, A. Liotta, S. Spirou, M. Waldburger, "A Discussion of Socio-economic Management and Incentives for the Future Internet," *Proc. of GLOBECOM Workshops*, pp. 33–38, 2009.
- [24] J. Mun, *Real Options Analysis: Tools and Techniques for Valuing Strategic Investment and Decisions*, Wiley Finance, November 2005.
- [25] F. P. Boer, *The Real Options Solution: Finding Total Value in a High-Risk World*, Wiley Finance, February 2002.