A Case of Knowledge-driven Policy Management: Bringing Discipline to Internet Routing

Anduo Wang
Temple University
adw@temple.edu

ABSTRACT

SDN has produced many artifacts that work better than its predecessors: openflow, controller, network programming, just to name a few. But the true advantage of SDN, we argue, is that it also created a discipline — principles that explain why and why not, and realizations that borrow heavily from other fields. Moving beyond the initial context of SDN, namely network state management under a single administrative domain, this poster asks whether the same disciplinary approach can apply to Internet routing with rich policies by multiple independent domains. To answer this question, we identified principles that reveal why BGP was chosen, why extensions and replacements to BGP keep well up, yet BGP remains largely unchallenged. The principles uncovered surprisingly new understanding. Moreover, it leads us to a new policy scheme, which we call knowledge-driven policy routing, that may retain the advantage of BGP while removing the pitfalls. A proof of concept implementation is also sketched by adapting data exchange and knowledge management research.

1 INTRODUCTION

SDN has produced many artifacts that work better than their predecessors: OpenFlow and P4 create unified open interfaces for programming heterogeneous data-plane elements, enabling rapid innovations; Network controllers manage the data-plane state once and for all, freeing control applications from the tedious and complex distributed state management; high-level network programming languages feature a diverse set of control abstractions, drastically simplifying network control and enabling new services such as verification; just to name a few. But what is the true advantage of SDN?

The most essential ingredient to SDN's unprecedented success, in a sense, is not any particular artifact, rather, it is the fact that SDN creates a discipline [7] — principles that explain why, and realizations that borrow from other fields. Open architecture as opposed to vertically integrated architecture, centralized computation rather than distributed protocols are the principles behind SDN that allow rapid innovations, that enable drastic simplification. The new SDN architecture where a network operating system sits between network control and the heterogeneous devices provides an insertion point for new abstractions, allowing one to import a large body of abstractions from the programming languages and distributed system communities.

Moving beyond the initial context of SDN, namely network state management under a single administrative domain, *we ask whether the same disciplinary approach can help Internet routing that features rich policies by multiple administrative domains.*

Current policy based inter-domain routing research, like intra-domain network management before SDN that offer various point solutions, is dominated by ad hoc strategies that address specific limitations of BGP (border gateway protocol) [6], the one single de facto inter-domain routing protocol. Little effort is devoted to fundamental questions like what is missing in BGP that gives rise to the many attempts to replace it? Why was BGP selected for the Internet in the first place? Despite the many new proposals, why BGP seems to be extremely "robust" and resists any change? Popular myth may hold that BGP was chosen merely because the absence of a better alternative [2], or that few extensions and replacements to BGP were ever deployed because — like migration to IP6 — it is simply hard to implement any change on the global Internet. On the contrary, we believe that there are deeper reasons that can be explained by a well-thought discipline.

As a first step towards this discipline, we identify two principles: With what we call the autonomy principle, we argue that, unlike commonly held belief that BGP supports full autonomy [3] — the freedom of individual domains (or autonomous systems, ASes) to pick and re-advertise any route as its best choice, BGP's opaque policy mechanism inherently restricts autonomy. This is because full autonomy demands the freedom to select from the *entire* routing space opaque policy severely limits individual ASes' visibility into the routing space. This explains the mechanism behind many alternatives to make BGP more flexible — they all attempt to improve ASes' visibility in one way or another. On the other hand, with our so called accountability principle, we argue that the long lived BGP was chosen and is hard to change *not* without a good reason, that BGP stands out by including native support for accountability. BGP is capable of verifying or correcting the enforcement of a policy within current Internet architecture. Resources in the Internet are divided and owned by ASes that are interconnected under bi-lateral agreements [1]. In line with this model, BGP policies are either implemented within an AS or along AS borders, in the later case can be factored into the corresponding interconnection agreements.

We believe that these principles can serve as a proof that an SDN-like disciplined understanding for policy management in the Internet is possible, explaining how and why we arrive at BGP. More importantly, we envision that the principles can guide the design of future policy infrastructure to address BGP's limitations — autonomy, as well as preserve its strength — accountability. Specifically, we propose a novel policy scheme that aims for both autonomy and accountability by separating policy distribution and route computation. The idea is that, unlike BGP that keeps policies private, that implicitly embeds policies in route computation and distribution, we make policy explicit "knowledge" that can be exchanged and reassembled by individual ASes. The main feature of the new scheme is not to support policy-based route computation while hiding policies, but to help determine what is the right policies to disclose to which neighbor.

As a realization of this vision, we develop an implementation by borrowing from knowledge management community. We adapt concepts from data exchange and integrity constraints (IC) research [4]: We adopt the IC representation in the form of answer set programming as the knowledge representation for AS polices, and develop new techniques to address unique networking challenges. First, we extend the standard notion of ICs — a set of "equally" important logic statements — to "prioritized" networking policies. We introduce the notion of prioritized logic statements, and develop a semantic transformation method that can incorporate priority into standard logic statements. Second, we extend the standard data exchange method [5] to integrity constraints to support policy exchange across AS borders. We develop a new transformation process that can map a policy of one AS — knowledge constraints over that AS's routing data — to policy over its neighboring AS — constraints over that neighbor's data — by using dependencies among the two ASes' routing data.

2 PRELIMINARY RESULT

This poster presents a policy scheme as a realization of the autonomy and accountability principles. Note that we do not claim our scheme to be the final solution to all policy routing problems, nor do we intend to exclude other designs. Rather, out goal is to illustrate that it is possible to build a policy architecture from well thought principles by borrowing from other fields.

The goal of our policy scheme is to achieve fully autonomous policy at each individual AS by improving its visibility into the routing space, and to preserve the accountability of BGP-like policies by leveraging an AS's interconnection agreements with its neighbors. The key idea is to discard BGP's policy-based route computation and distribution. Instead, we split routes and policies: Routes are propagated similar to distance-vector protocol except that they are now computed (selected) by the policies at each AS; The policies, on the other hand, are explicitly exchanged between ASes — a policy announced by an upstream (downstream) AS represents a service request (restriction, respectively) according to the inter-AS agreements. Upon receiving a policy, an AS attempts to realize it: when a policy-compliant route exists, the receiving AS will simply select that route; otherwise, the AS will further propagate the policy (request or constraint) to other neighbors.

To illustrate this vision, our poster will present a policy distribution mechanism by borrowing from data exchange research and knowledge (integrity constraints) management: Policies are knowledge represented as data integrity constraints — logic statements about preference over routing information (data); Policies are disclosed between neighbors ASes according to their relations in the form of knowledge exchange — the logic statements over one AS's data are mapped to that over a neighboring AS's data via semantic transformation; Finally, policies with varying importance — priority — are integrated via a normalization process that incorporates the notion of priority into standard logic.

REFERENCES

[1] M. F. B. Woodcock. 2016 survey of internet carrier interconnection agreements.

[2] L. Breslau and D. Estrin. Design of inter-administrative domain routing protocols. In *Proceedings of the ACM Symposium on Communications Architectures &Amp; Protocols*, SIGCOMM '90, pages 231–241, New York, NY, USA, 1990. ACM.

[3] N. Feamster, R. Johari, and H. Balakrishnan. Implications of autonomy for the expressiveness of policy routing. In *ACM SIGCOMM*, 2005.

[4] P. Godfrey, J. Grant, J. Gryz, and J. Minker. Integrity constraints: Semantics and applications. In *Logics for Databases and Information Systems*, 1998.

[5] P. G. Kolaitis. Schema mappings, data exchange, and metadata management. In *Proceedings of the Twenty-fourth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, PODS '05, pages 61–75, New York, NY, USA, 2005. ACM.

[6] Y. Rekhter, T. Li, and S. Hares. A Border Gateway Protocol 4 (BGP-4). RFC 4271, RFC Editor, 2006.

[7] The Future of Networking, and the Past of Protocols. http://www.opennetsummit.org/archives/apr12/site/talks/shenker-tue.pdf.