Routing Policies in Named Data Networking

Steve DiBenedetto, Christos Papadopoulos, Dan Massey
Computer Science Dept., Colorado State University
Fort Collins, Colorado, USA
dibenede@cs.colostate.edu, christos@cs.colostate.edu, massey@cs.colostate.edu

ABSTRACT

Modern inter-domain routing with BGP is based on policies rather than finding shortest paths. Network operators devise and implement policies affecting route selection and export independently of others. These policies are realized by tuning a variety of parameters, or knobs, present in BGP. Similarly, NDN, a information-centric future Internet architecture, will utilize a policy-based routing protocol such as BGP. However, NDN allows a finer granularity of policies (content names rather than hosts) and more traffic engineering opportunities.

This work explores what routing policies could look like in an NDN Internet. We describe the knobs available to network operators and their possible settings. Furthermore, we explore the economic incentives present in an NDN Internet and reason how they might drive operators to set their policies.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design

General Terms

Economics, Theory

Keywords

Inter-domain policy, information-centric networking

1. INTRODUCTION

Modern inter-domain routing and packet forwarding is based on policies rather than finding shortest paths. Collections of routers under a common jurisdiction that maintain uniform policies are referred to as Autonomous Systems (AS) [4]. Each AS defines policies independently in order select the best path to to a given prefix and to which neighbors it should be announced. These policies are realized by tuning a variety of parameters or *knobs*, available in the Internet's BGP routing protocol [10].

ICN'11, August 19, 2011, Toronto, Ontario, Canada.

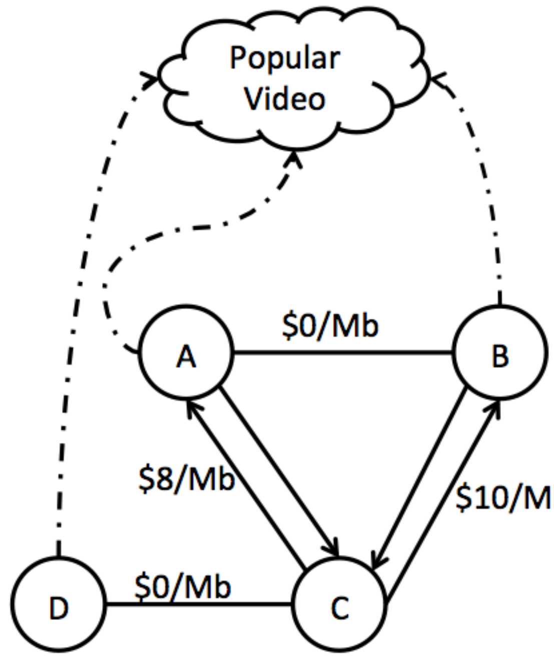

Figure 1: Driving example for our policy investigation. Upward arrows represent customer to provider relationships while lines denote non-settlement peers. Dashed lines represent an arbitrary path to the Popular Video.

The shift to information-centric networks (ICNs) has the potential to dramatically change how inter-domain policies are set. While autonomous organizations and economic driven policies is likely to persist, information-centric networks have the potential to alter both the incentives and the knobs used to the set the policies compared to those found in BGP.

Predicting the possible policies that might be adopted in information-centric networks is very hard given that ICNs offer new incentives. However, we can offer some concreteness by describing the new policy knobs available in ICN and then conjecture on possible policies based on current economic incentives. While such conjectures are bound to be inaccurate by nature, the exercise may offer better insight into the policy space in ICN.

Throughout this paper we will refer to Figure 1 to illustrate new knobs and use scenarios, and describe new incentives. The figure shows an AS-level topology; upward arrows indicate customer-provider relationships and horizontal connections denote settlement-free peering agreements.

In today's Internet a common BGP policy is for customers to not announce routes learned from one provider to another [2] (the valley-free rule). In Figure 1, C learns a route to a popular video website from A, but does not announce this route to provider B or peer D. Economic incentives drive both choices: C pays A for transit and thus incurs a cost for downloading a video from the website. Providing either B or D with a route to the popular video would incur a net cost to C. As a result of the peering agreement, C is willing to provide its peer D with access to C's local content (and that of C's customers, if any), but D does not pay C for access to other content.

In this paper, we argue that the incentives and the policies can change dramatically in ICN. The reason is that, among others, ICN names content, as opposed to the traditional Internet and BGP, that name hosts (or collections of hosts). The use of named content along and caching can enable new policies in ICN. For example, suppose C has cached a popular large video. Peers C and D may reach an agreement to share not only content present in their networks or their customers' networks (as currently done), but also to their caches. The incentive is that such peering agreements can help reduce the amount of *external content* that C and D download through their providers. In addition, one can imagine more creative policies where C offers provider B access to the cached popular video as well (thus violating the valley-free rule), possibly in exchange for a reduced service charge. In ICN such policies are feasible because content is named, and thus data flow can be enforced at a much finer granularity compared to the current Internet. Interesting new policies are also possible in flavors of ICN that support multi-path routing, such as NDN [6].

In order to keep the discussion concrete we will use NDN as our driving example of ICN. In contrast to the current Internet, which only offers control plane knobs, NDN also offers knobs to manipulate data plane operations such as Interest forwarding, cache access, PIT behavior, and FIB behavior. This enables not only finer policy granularity, but also *new polices* that are based on content names rather than prefixes.

This work is a first step in understanding what inter-domain routing policies might look like in NDN, or ICNs that resemble NDN. We describe the new policy knobs available to network operators and their possible settings. Furthermore, we explore new economic incentives that may be present in an NDN Internet and consider what types of routing policies they may develop. We note that our goal is not to describe how existing BGP policies might translate into NDN (although we find this to be important future work) but what policies might be feasible and reasonable in NDN.

2. BACKGROUND & RELATED WORK

We begin with some Named Data Networking (NDN) [6, 1] background. A node in NDN includes *the content store (CS), the pending interest table (PIT), and the forwarding information base (FIB)*. A consumer requests content by sending an Interest Packet for the desired content. When the Interest Packet arrives at an NDN router, the node (conceptually) consults the CS, PIT, and FIB in that sequence. The router first checks whether the data requested by the Interest Packet is already present in the node's content store (CS). The CS is essentially a packet cache that allows an NDN node to immediately determine if there is an exact match to cached data. If the requested data is not in the CS, the NDN node checks its Pending Interest Table (PIT) if a request for the same data has recently been issued. The PIT maintains records of *faces* (the NDN equivalent of network interfaces) that are currently waiting for this content. If an existing PIT entry is found, the data has already been requested and the Interest Packet is not forwarded. Instead, the face on which the Interest Packet arrived on is appended to the set of faces in the existing PIT entry, so that when the content arrives it is replicated to all requesting faces.

Interest Packets that do not match a CS or PIT entry trigger the creation of a new PIT entry and the Interest Packet is forwarded toward one or more faces that might lead to hosts that have the desired data. To select appropriate faces, the router consults its name-based forwarding table (FIB) to determine where to forward the Interest Packet. Assuming some routing protocol has created correct forwarding state (FIB), the Interest Packet will be forwarded towards a node with the desired data.

The desired content is sent back in Data Packets and follows the reverse path of the Interest Packet that requested it. In other words, the PIT entries effectively create a trail for the Data Packets to follow back to the host(s) that generated the interest. As Data Packets pass through a router, the packet may be cached in the content store before being relayed to the requesting face(s) indicated by the appropriate PIT entry. The PIT entry is subsequently removed. Data Packets that do not have a matching PIT entry are discarded.

To summarize, 1) Interest Packets may be answered using cached data in the Content Store, 2) the system is inherently multi-path and an Interest Packet can be forwarded out multiple faces, 3) the PIT helps ensure that duplicate copies of an Interest Packets are suppressed, 4) Data Packets follow the PIT entries back to the requester(s) and 5) Data Packets not matching a PIT entry are dropped,

In related work, Huston [5] first examined the delicate balance between competitiveness and cooperation ISPs face in order to exist as a business. The author investigates a range of settlement models and business relationships between ISPs. Our work is in a similar spirit as we explore the underlying economic incentives and new possible business relationships that could occur in an information-centric network.

Our work is similar spirit to that of Rajahalme *et. al* [9]. The authors are perhaps the first to analyze the economic incentives, or lack there of, for tier-1 ISPs to deploy information-centric network architectures. Furthermore, the authors also explore novel data-centric policies that may develop. While Rajahalme *et. al* analyze several different ICN architectures, our work focuses on the NDN approach [6] and how its specific features lead to interesting inter-domain routing policies. Additionally, we study ICN policies from the post-deployment perspective in order to focus on the full potential of these novel architectures.

Gao [2], which is widely regarded as the seminal work on routing policy inference techniques for today's Internet, was the first to use economic incentives and Huston's peering suggestions to infer AS relationships. Using this approach, an AS-level topology can be annotated with business relationships according to observed routes, the *valley-free* rule, and node degree. Our work uses Gao's business relationships and the valley-free rule as a starting point to explore possible NDN policies.

3. NDN POLICY KNOBS

Table 1 provides an overview of the policy knobs available to NDN network operators. As with BGP, there are control plane knobs; but in addition, NDN offers several data plane knobs. These operate on components that are new in NDN, such as content store (CS), the Pending Interest Table (PIT) and the Forwarding Information Base (FIB). Note that the FIB in NDN has significant differences than the FIB in BGP: in NDN the FIB is used to forward interest packets only, and may contain multiple entries for the same name prefix. Data packets follow the reverse path marked by PIT entries back to the originator of an interest packet.

3.1 Control Plane Knobs

In the current Internet AS policies are implemented by the BGP routing control plane [3] and a number of *prefix attributes*. BGP selects a *single best route* to a prefix, which may then be exported to a select set of neighbors. While similar knobs are available in NDN there are a couple of important differences. First, in NDN route selection influences forwarding of interest packets, not data. Data packets follow the state left in the PIT by interest packets, so packet flow is symmetric. Note that this is in stark contrast to the current Internet, where packet flow is often *asymmetric*. The second difference is that NDN can select *multiple feasible interfaces* in an attempt to find content nearby.[1]

The table shows the possible settings of the route selection knob at the control plane. An advertised route (or name prefix) can be ignored, selected, or selected and propagated to neighbors. Furthermore, NDN may maintain multiple routes to a given content. Which routes to use when searching for the content is determined by FIB policies, which are described below.

3.2 Content Store

The Content Store (CS) acts as a cache for an NDN router and adds a new set of policy knobs at the data plane, and thus not present in the current Internet. The CS Access knobs are shown in Table 1) and provide the operator an opportunity to determine whether or not an incoming Interest Packet is allowed to access cached content. The first setting *Allow* determines if an arriving interest packet is allowed to search the CS for the desired content. An Interest may be denied access to CS if, for example, the interest came from a peer and there is no agreement to allow the peer unfettered access to the CS.

3.3 Forwarding Information Base

The control planes in both BGP and NDN utilize a routing information base (RIB) to record available routes. BGP will select a single, best route for each prefix from the RIB and install it into a forwarding information base (FIB), typically on the router's line card. The data plane decision is simple. If a FIB entry exists, it indicates which interface *will be used* to forward a matching IP packet.

NDN allows for multi-path routing and thus multiple faces can be associated with a FIB entry. The control plane only determines which faces *could be used* by the data plane to forward an Interest Packet. The data planes face selection may be controlled through the FIB Usage knob. For example, Interest Packets from a customer may be forwarded over all possible faces selected by the control plane policy. However, an Interest Packet (for the same name) from a settlement free peer may only be forwarded over subset of the possible faces.

[1]Note the PIT helps protect against loops in the multi-path selection.

3.4 Data Packet Policies

Note our knobs have focused on policies for Interest Packets. Data Packets follow the trail left by Interest Packets and will be processed only if there is a matching PIT entry. Undesired Data Packets are blocked preemptively by not forwarding Interest Packets. Conversely, if a node created a PIT entry and forwarded an Interest Packet, it is expected to also relay the corresponding Data Packet. As such, the only policy related question for a Data Packet is whether to cache the content. Caching strategies are an interesting challenge, but beyond the scope of this paper.

4. POLICIES & INCENTIVES

Policies in the current Internet tend to be driven by economic incentives [5, 2, 7, 8]. We contend economic incentives will continue to drive policy in an information centric network, but there are a number of new policy options. This section introduces several example policies, argues the policies can't be easily adopted in the current Internet, and shows how our policy knobs can be used to implement the policies. All our examples use the network shown in Figure 1, unless stated otherwise.

4.1 Multi-Path Means Multi-Pay

In today's Internet node C from Figure 1 learns two possible IP routes to the popular video, one route from provider A one from provider B. C must select a single best route to the popular video location. Assuming other factors such as bandwidth and connection speed are comparable, C is likely to prefer the route from A due to its price of \$8/Mb of traffic. The same path selection knobs also apply in NDN, but NDN routing can select *multiple feasible faces* rather than a *single best exit.*

In NDN, the control plane can select *both* A and B to reach the popular video. C can simultaneously send Interest Packets to both A and B to retrieve the popular video content. The resulting Data Packets forwarded to the consumer can even be an interleaved combination of Data Packets from A and B for maximum robustness and performance. The location of the data is irrelevant so the consumer does not know or care if the Data Packet arrived from A or B. This is a new potential in NDN, but is it in C's best interest to use multiple paths?

Both providers will charge for the transit service. If the policy is to send all Interest Packets to both A and B, instead of the previous \$8/Mb pricing under BGP, C is now effectively paying \$18/Mb for the same content it would have received had it only used a single path. C could use the control plane policy knob to select a single path and then change paths in the event of problems. This is essentially today's BGP routing choice.

Instead, C's Control Plane policy knob could be used to select the set $\{A, B\}$ and then C applies the FIB Usage knob for finer grained traffic engineering. For example, initial Interest Packets could be forwarded on both faces and later Interest Packets forwarded only along the better face. Periodically, an Interest Packet can be forwarded over both links to see if performance has changed.

For clarity, our examples use a simple cost per MB model in both Internet and NDN. Other models yield similar results. For example, one could also argue that provider's A

Component	Name	Controls	Setting	Description
Control Plane	Routing/RIB	Routes	Not Selected	Route cannot be used to forward Interest Packets.
			Selected	Route may be used for forwarding Interest Packets.
			Select and Announce	Route may be used for forwarding Interest Packets and may be announced to peer groups.
Content Store	Cache Access	Interests	Allow	Check interest against cache and proceed to PIT if no matches are found.
			Cache Only	Check interest against cache and drop if no matches are found.
			Deny	Drop interest without checking cache.
FIB	FIB Usage		Full Usage	New matching faces may be found and added to the existing face list.
			Limited Usage	No new faces may be added to the existing face list.

Table 1: Examples of NDN policy knobs (not exhaustive).

and B charge C a flat rate per month. In this case, C does not pay more by requesting the data from both providers. But C does double the bandwidth used, causing C to pay for higher capacity lines from both A and B.

The relatively small size of Interest Packets, in conjunction with intermediate PITs and content stores, would reduce the number of packets that need to be forwarded and possibly cause significant decreases in connectivity pricing. However, we argue these changes would still favor a single, or relatively small number of, selected paths to content. The operator that pays $X for connectivity under BGP is unlikely to pick sufficient reduced price NDN paths to continue paying $X. Instead, the economic incentives favor picking a single, reduced price, path. That said, we expect a price point where operators will want to pick up a small number of additional paths to improve robustness and performance. The total price is expected to be less than the original amount.

4.2 Cache Sharing Peers

The Content Store adds perhaps the most interesting new policy options. In Figure 1, nodes C and D are non-settlement peers. Neighboring ASes that are of roughly equal size and tend to exchange similar volumes of traffic often form non-settlement peering relationships. In this scenario, the neighbors negotiate a limited scope of connectivity and no money is exchanged. The scope of connectivity is usually limited to each ASes' own destinations and those of its customers. C does not offer D a route to the popular video since providing this transit service would cost C. Similarly, D does not offer C a route to the popular video.

In NDN, node D may have cached the popular video in its Content Store. At this point, offering C access to the cached popular video would not impose a new cost on D. If D shared its cached data with C, C can obtain some cached data over the peering link instead paying provider A or B for transit service. In return C could offer D access to its cached data and reduce D's transit costs. Provided there is some available bandwidth on the peering link, both C and D have an economic incentive to share cached content.

Cache sharing could be implemented using the Control Plane knob. When the popular video is added to its Content Store, D announces the popular video's content name to C. If the video is later flushed from the content store, D should send C a withdraw for the popular video content name. While feasible, this could result in considerable churn to the routing protocol.

Instead, the Cache Access policy knob can accomplish the same policy. D announces a default route to C and then applies the *Cache Only* setting to Interest Packets received from C. Interest Packets that can be satisfied using cached data are answered. Cache misses are dropped; D is not willing to pay its providers to fetch the data.

Finally note Cache Sharing impacts our previous discussion on multi-path routing. With Cache Sharing enabled, the Control Plane knob at C selects $\{A, B, D\}$ as possible faces. The FIB Usage knob may initially try all three, preferring the free path via D if it answers over the paid path via A or B.

4.3 Routing Rebates

Clearly, there is an economic incentive for C and D to share their cached content. There also may be an incentive for C to share cached content with its provider B. In today's Internet, C would not offer transit service to its provider. This is the basic premise of BGP policies [5, 2].

But suppose C has cached the popular video. C has already paid the cost of downloading the content. When the link to B is lightly loaded, C may offer B access to its cached data in exchange for reduction in charges. This is somewhat analogous to solar power in the U.S. A customers install solar panels on their home. When home's energy load is low and the power companies load is high, excess power generated by the home is automatically sold back to the power company and appears as a credit on the home's bill. Similarly, when C's traffic load is low and provider B's load is high, B may choose to fetch cached data from C and provide C with a rebate.

From C's perspective, implementing the policy is identical to implement cache sharing with peer D. As long as the added load does cause C to drop packets, it is clearly in C incentive to provide B with data in exchange for a rebate. Provider B must make a more complex decision that compares the cost of the rebate paid to C with the gains it obtains by using C's cached data, but [7, 8] already observe that ISP polices are increasingly complex as one moves from edge networks to providers.

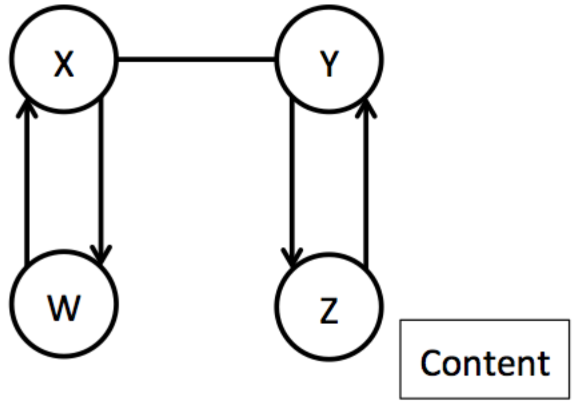

Figure 2: Driving example for the economic incentives of caching. A is attempting to retrieve content stored at D.

4.4 Policy in the Content Store

Caching data is analogous to providing redundancy and load balancing for the content's publisher. However, caching also provides faster access for an organization and its customers as well as free content retrieval. Consider the simplified example figure 2 where W and Z are customers of X and Y, respectively, while X and Y are peers. Here, W wishes to retrieve a piece of content that can only be found at Z. When W expresses an interest in Z's content, X and Y will provide transit services and carry back Z's response as each are commissioned by their respective customer. However, as the response travels back from Z to W, each node must make a decision of whether or not they wish to keep a local copy.[2]

Traditional BGP business relationships only provide the ground work for the transit aspect of this example, which is fulfilled. Our first observation is that as Z's provider, Y is able to charge Z for traffic sent to it. As such, Y's has an incentive to not cache the response data. This gives Y leverage to charge Z more for caching in addition to transit services.

W's caching policy is very different from that of Y. Here, W will attempt to cache as much content as possible to avoid paying X for subsequent requests. In this example, X's economic interests are closer to W's than Y's as X would prefer to avoid sending traffic (Interests) to Y in order to preserve its non-settlement peering. Furthermore, if X is able to cache sufficient data from a range of Y's customers, relative to what Y is able to cache from X, then X could conceivably become a paid provider of Y.

4.5 Policy in the PIT

Finally, we provide one example of policy decisions for the PIT. Suppose node C receives an Interest Packet from node D. If the content is not present in the cache, C may forward the Interest Packet to customers E and F. C will create a PIT entry for this content but will not forward the Interest Packet to Providers A or B. Note because the Interest Packet was not forwarded to Providers A and B, it may not elicit a response. But this is the desired policy result since D is not paying C for transit service.

[2]There may also be more complex decisions, such as how long the content should be cached, but for these are beyond the scope of this work.

Now suppose C's customer F also sends an Interest Packet for the same content name. First, consider what happens if there was no PIT entry. C would have forwarded the Interest Packet to customer E and also forwarded the Interest Packet to Providers A and B. C is obligated to pay the cost of using A and/or B since F is a customer and the customer is paying for transit service.

Unfortunately, a PIT entry was created when D requested the content. Basic NDN would simply note that the content has already been requested. The fact F also wants the data would be recorded in the existing PIT entry, but no new Interest Packets would be forwarded. Thus no Interest Packets are sent to provider A and B and the Interest Packet may not get response. Thus customer may F may not get the transit service it thought it had purchased. Our solution to this problem is to modify PIT entries to also record the used outgoing faces. This allows for a simple check against the second to longest matching name and sending an interest to faces belonging to the unused faces.[3]

4.5.1 Discount Multi-path through Caching

On some occasions, an NDN node may only wish to provide access to its content store and a limited set of names to others. More specifically, the only offered names would belong to itself or its customers. It is necessary to offer these names to neighbors, who otherwise only have access to the content store, as the node's customers are paying to ensure they are reachable to the world. This policy could also be modified to allow usage of the PIT. In practice, the difference between an entry existing in the PIT and the content store should be roughly a round trip time. Such policies would allow operators to offer a form of discounted connectivity that would primarily serve popular content. If this practice lead to significant discounts in connectivity, it may present a cost effective realization of multi-path routing.

5. CONCLUSIONS & FUTURE WORK

This work explores the economic incentives an potential routing policies in a full deployment of the NDN architecture. We explored the primary policy knobs available to NDN network operators and noted that NDN provides operators with a number of new data plane knobs that can be used to set policies. However, this paper does not aim to identify every possible policy. Rather we argue economic incentives will drive operators to use the policy knobs to achieve their goals. Furthermore, we show examples of how economic incentives lead to new policies that are not feasible in the current Internet.

We note that while multi-path routing may be considered one of NDN's greatest strengths, economic incentives could also encourage the system to favor fewer paths in order to reduce costs. Furthermore, we show, like Rajahalme *et. al* [9], that the economic incentives of content caching offer a number of new possibilities such as Cache Sharing between peers and Routing Rebates between customers and providers. Our work also shows that policy can impact the PIT as well the traditional forwarding table (FIB) and new cache (Content Store). We are also working on extending our understanding of NDN and BGP policies to to help network operators

[3]We do not consider the peer receiving a response that would only have been yielded from a customer's interest to be incorrect. The peer is stumbling into better service than it would have normally receive and such events are expected to be reciprocated.

develop a more complete understanding of the NDN routing environment.

6. REFERENCES

[1] Named data networking. http://named-data.org.

[2] L. Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM TRANSACTIONS ON NETWORKING*, 9(6), dec 2001.

[3] T. Griffin, F. Shepherd, and G. Wilfong. The stable paths problem and interdomain routing. *IEEE/ACM Transactions on Networking (TON)*, 10(2):232–243, 2002.

[4] J. Hawkinson and T. Bates. Guidelines for creation, selection, and registration of an autonomous system (as), March 1996.

[5] G. Huston. Interconnection, peering, and settlements. *Proc. INET*, 1999.

[6] V. Jacobson, D. Smetters, J. Thornton, M. Plass, N. Briggs, and R. Braynard. Networking named content. *Proceedings of the 5th international conference on Emerging networking experiments and technologies*, pages 1–12, 2009.

[7] W. Mühlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig. Building an AS-topology model that captures route diversity. *ACM SIGCOMM Computer Communication Review*, 36(4):206, 2006.

[8] W. Mühlbauer, S. Uhlig, B. Fu, M. Meulle, and O. Maennel. In search for an appropriate granularity to model routing policies. *ACM SIGCOMM Computer Communication Review*, 37(4):145–156, 2007.

[9] J. Rajahalme, M. Särelä, P. Nikander, and S. Tarkoma. Incentive-compatible caching and peering in data-oriented networks. *ReArch '08 - Re-Architecting the Internet*, 2008.

[10] Y. Rekhter, T. Li, and S. Hares. A border gateway protocol 4 (bgp-4), January 2006.