

Path Aware Networking: Obstacles to Deployment
(A Bestiary of Roads Not Taken)
(**draft-irtf-panrg**-what-not-to-do-00)

Note the name change after research group adoption!

Spencer Dawkins

Things to remember about this draft

- This draft was an outcome of the first PANRG meeting at IETF 99
 - "We keep encountering the same problems. We should write them down"
- Don't describe every idea, but capture every lesson
 - There is no shame in describing your learning experience!
 - This stuff was research, even if we **were** trying to engineer it in the IETF
- People contributed subsections from their own experience
 - Requested format for contributions is in Section 3 of the draft
- This is now an RG draft, so you all own it (make good choices!)
 - Please send comments on draft contents to the mailing list

New Summary of "Summary of Lessons Learned"

- Benefits of adoption must be big enough to overcome inertia
- **Providing benefits for early adopters is key**
- Benefits must be big enough for operators to justify deployment
- Current operational practices can prevent deployment of good ideas
- Per-connection state in intermediate devices is an impediment
- **Many routers don't process inband/HBH signals in hardware**
- **Middleboxes and end systems each have to trust the other**
- Increasing distance from sources makes path information less useful
- Many applications don't know things they need to tell PAN transports

Status of draft-irtf-panrg-what-not-to-do-00

- Version 00 submitted prior to IETF 103
- Added pointer to IAB ITAT workshop report on protocol design and adoption
 - Not specific to Path-aware Networking, but worth reading
- Section 2 (Summary of Lessons Learned) builds on new contributions
- Now includes contributions on
 - Integrated Services (IntServ), provided by Ron Bonica
 - Quick-Start TCP, provided by Michael Scharf
 - Triggers for Transport (TRIGTRAN), mostly unchanged from -00
 - Shim6, based on input from Erik Nordmark
 - NSIS, provided by Roland Bless and Martin Stiemerling
- I do want to get these descriptions right. Please send corrections!

"Please send corrections" - well, Wes Eddy did ...

- *(1) Section 4.1.2 on RSVP says: "BGP uses Flow Specs to distribute firewall filters."*

That's not quite right. BGP flowspecs (RFC 5575) are not related to RSVP flowspecs. The same term is used for different things, unfortunately.

- Well, that's embarrassing. RFC 5575 doesn't reference RFC 2210, but it does explicitly define 4 "BGP Flow Specs". This should be rephrased as "**BGP reused the concept of flowspecs when distributing firewall filters**".

"Please send corrections" - well, Wes Eddy did ...

- *(2) Would it make sense to say anything about XCP, PCN, or CONEX?*
- **High-order bit - what other contributions would provide lessons that we don't already have in Section 2, "Lessons Learned"? Those are the additional contributions we should solicit.**
- Given that filter, what lessons would contributions about XCP, PCN, or CONEX provide?
- Given that filter, what lessons would other contributions provide, that we don't already have?

"Please send corrections" - well, Wes Eddy did ...

- *(3) In the shim6 discussion in Section 4.4.1, I'm not sure what is meant by "file descriptions with associated HTTP state" ... something is missing in the description of the problem, because HTTP is mentioned out of nowhere.*
- Should be "The technology issues centered around scaling concerns that Shim6 relied on the host to track all the **transport connections** ~~TCP connections and the file descriptions with associated HTTP~~ **state**, while also tracking Identity/Locator mappings in the kernel, and tracking failures to recognize that a backup path has failed."

"Please send corrections" - well, Wes Eddy did ...

- *(4) Section 4.4.2 talks about SCTP and MPTCP somehow, but I don't think those are really related to shim6 as coherently as the text might imply (they're transport layer versus network layer), nor do I think those are really "path aware" in any meaningful way. They know about multiple interfaces locally, which lets them possibly use multiple paths, but they're not really aware of anything about the innards of those paths in the way that other technologies mentioned are.*
- My point here is that operators were concerned about shim6 because hosts shift traffic on paths operators traffic-engineer, and SCTP/MPTCP hosts can also shift traffic, but **now**, no one cares???

But that comment makes a larger point

- (only) Spencer's opinion
 - **One goal for this draft is that we can use it to spot proposals that might trip over known obstacles to deployment in Lessons Learned**
 - **Let's emphasize that IRTF RGs don't police the Internet, or the IETF**
 - **But the truth is always appropriate**
- If that's true
 - I'll remove observations about how shim6 hosts moving traffic around is different from MPTCP and QUIC hosts doing the same thing in this draft
 - I'll ask the chairs and the Research Group if it's time to start the next draft

Spencer's understanding of where the goalposts are

- I'd like for this document to be useful advice from PANRG to IETF
 - IETF developed these protocols that were not implemented and deployed
 - There's no reason to think IETF will not "learn these lessons" again
 - So, what other guidance can we give to the IETF?
- I'd like for this document to be useful advice to PANRG
 - [Open Questions in Path Aware Networking](#) to guide PANRG research
 - This document could help identify other open questions to guide PANRG
 - So, what other guidance can we give ourselves? :-)

What Next? Questions

- Are there more Path Aware Networking lessons to document?
- How many more contributions do we need? Which ones? From who?
 - CONEX?
 - Are there notable lessons from outside the IETF to guide us?

What Next? What seems reasonable to Spencer

- We could reasonably look at current path-aware proposals
 - Will current proposals encounter same obstacles? Why or why not?
- We could reasonably compare this draft to panrg-questions
 - Are there obstacles that we don't have questions for yet?
- We could reasonably finish this draft and publish it
 - IRTF RGs don't **have** to publish RFCs, but certainly **can** publish RFCs