<table>
<tr><td></td><td></td><td colspan="2">Authors' responses to comments received on initial release (Draft 1) of NIST SP 800-189 (draft): "Secure Interdomain Traffic Exchange: BGP Robustness and DDoS Mitigation" (Draft 1 publication date: December 2018). Changes based on these comments/responses are incorporated in Draft 2 of NIST SP 800-189 published on October 17, 2019.</td></tr>
<tr><td></td><td></td><td colspan="2">Note: SR# = Security Recommendation #</td></tr>
<tr><td><b>Lines:</b></td><td><b>SR#</b></td><td><b>Comments</b></td><td><b>Authors' response</b></td></tr>
<tr><td></td><td></td><td><b>Comments set #1</b></td><td></td></tr>
<tr><td>529-530</td><td>3</td><td>This suggestion is pre-ROA. Until more adoption on RPKI, not much movement</td><td>Agree with the observation. Note that about 46% of the announced address space (/24 granularity) in the RIPE region is RPKI registered and covered by ROAs. All regions hopefully ramp up. AT&T and Telia are already doing BGP-OV based on ROAs for filtering peer routes.</td></tr>
<tr><td>531-533</td><td>4</td><td>For lower tiered providers the recommendation might work but for very large global providers, this requirement poses a significant challenge due to the scope and scale as it related to large providers. Most ISPS that are considering RPKI are looking at the HOSTED model. Because the delegated model would require a long time to do their own RPKI. Providing this to customers would be even harder.</td><td>Your points well taken. The intention is to encourage ISPs to provide RPKI registration assistance to customers in any way they can. ISPs at each tier educate their customers and offer support with hosted or delegated model as appropriate.</td></tr>
</table>

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
591-595	5,6	Registering ROAs is not always a simple task. For enterprises who have their own Provider Independent (PI) space this might not be too hard, but for large providers, there are many IP blocks to consider, many of which may be used as Provider Aggregatable (PA) space. Registering ROAs for PA space is a very complex task that could have very negative impacts if not done properly or accurately. Further, doing this for customer prefixes is hard as a public facing, supported front-end needs to be developed and supported. Not sure if this also includes customers PI space, or just PA space. This is particularly hard because all customer prefixes in a block must have accurate ROAs before the overall block ROA is published. This information is not always available to the provider, since customer don't currently need to provide their full routing policy for a BGP session. Customers may have backup arrangements that the carrier is not aware of or they may have backup plans in place that change the prefix length of announcements. All variations must be confirmed with customers, possibly including legal agreements, before the block ROA is published. There are legal challenges to implementing this. To say it should be implemented at this time isn't suggested. Others have publicly commented on this as well as a paper by Penn State. https://pc.nanog.org/static/published/meetings/NANOG75/1900/20190219_Yoo_Rpki_Legal_Barriers_v1.pdf	The security recommendation uses 'should' instead of 'must' in recognition of the deployment time frame and difficulties involved. The community should certainly strive for the objective. Again, the there are regions in the world that have shown rapid progress with ROA registrations (as noted above, RIPE is at 46% ROA coverage). NIST has encouraged and helped support the U. Penn. work that is cited. That effort is aimed to reduce legal hurdles for resource holders in the ARIN region. ARIN is responding.

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
602-603	8	In smaller organizations and enterprises, a single ROA that covers both more specific and less specific prefixes will be more efficient in terms of scarce router resources, so a provider may decide to have a single ROA to cover both sets of routes. If the provider is creating separate ROAs, they do need to make sure that the more specific ROAs are in place before the less specific ROA is published. For large providers, this is particularly difficult because all customer prefixes in a block must have accurate ROAs before the overall block ROA is published. This information is not always available to the provider, since customer don't currently need to provide their full routing policy for a BGP session. Customers may have backup arrangements that the carrier is not aware of, they may have backup plans in place that change the prefix length of announcements, or their prefix has been SWIP'd from a larger provider. All variations must be confirmed with customers, possibly including legal agreements, before the ROA is published. A global recommendation may not work as you suggest in all cases.	(1) The ROAs are not stored on the router. They are stored in a RPKI cache server. The latter provides valid {prefix, maxlength, origin ASN} list to routers on a per prefix basis. So the router memory is not impacted by whether ROAs have single or multiple prefixes. (2) We agree with your other observations. We believe that a lot of efforts for network operator / customer awareness will occur with respect to RPKI and ROAs. And only after the adoption reaches a high mark, BGP-OV will likely be turned on in routers (although AT&T and Telia are already doing BGP-OV based filtering on peers). By then customers are highly likely to be RPKI aware and have ROAs in place, including multi-homing (backup) considerations. We agree this will take time. There is also time lag involved in SP 800 recommendations to enter FISMA considerations and eventually influence procurement requirements.

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
605-607	9	This recommendation valid but it could very well prevent ISPs from implementing ROAs. It is a chicken/egg problem. We should encourage any adoption at this point from ISPs, not discourage or add roadblocks. If a ROA for the less specific block is added prior to the more specific customer allocations, the customers may end up with invalid routes.	This comment is closely related to your comment for SR#8. It seems we are in agreement. Please also see authors' response immediately above for SR#8.
614-615	10	This is a new concept and could be a very good idea to try and prevent prefix squatting. We'd need to check that the practical validation process functions in this way though. This would require testing of the validation server and all router code. Probably something good to do in a field trial.	Thank you. We made a modification to SR #10 (now SR#11 in the Draft2) consistent with SR #8 and SR #9. With this modification, there should be no issue concerning SR#10 with the validation process.
616-617	11	We are not sure how strictly this is 'enforced' by providers.	Globally, about 0.05% of the unique prefixes have AS_SET in their AS_PATH. Momentum to enforce deprecation of AS_SET and AS_CONFED_SET seems to be picking up. This seems to be important for several reasons including origin validation and route leak prevention. There is an active draft in the IETF that seeks to make this mandatory (it would update the BGP specification [RFC 4271]). See thread: https://mailarchive.ietf.org/arch/msg/idr/bFEht2e-yq4DdCRa6mquUU6xVU0

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
618-619	12	This depends on the decision of the organization to go in this direction. If done, it is important organizations operate more than one cache for resilience reasons. Also, this number should reflect the scale and geographic reach of the organizations network. For instance, a global network with thousands of nodes may require significantly more than two caches.	OK. Yes good observation. Does not call for change in the document.
620	13	There is a potential scale issue for large providers. Large lists of any sort could easily affect memory resources on the router.	The size of the {prefix, maxlength, origin ASN} white list received from the RPKI cache at the router is independent of the size of the provider (default-free zone). Multiple commercial router vendors have implemented origin validation. NCCoE SIDR Project testing did not reveal router memory issues with ROA-based {prefix, maxlength, origin ASN} white list for the full set of Internet routed prefixes.
627-629	14	When organizations start to do incremental updates, this recommendation is valid.	Yes.
671-678	17	In principle this is correct. However yet again the issue of scale comes up. The result of this is large providers could end up with a large number of ROAs per prefix, and if scaled over the entire route table, this could result in a huge number of ROAs with obvious resource and performance considerations on all routers validating the BGP table.	The scale issue is similar what was raised earlier. Please see authors' responses above corresponding to SR #8 and SR #13 above.
709-714	18	Could be promising. Some providers use hard-coded prefix filters based on IANA allocations to prevent announcements of unallocated prefixes. Moving to a more dynamic method with ROAs might be a good idea but it may make sense to have some basic prefix sanity checking should the RPKI ROA services become unavailable. Slow and careful adoption recommended here.	Point well taken. In general, cached RPKI/ROA data will be used when some RPKI/ROA services become temporarily unavailable.

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
719-721	19	It could be argued prefix ranges that should never be announced, should be hard-coded in the router config to ensure they are never announced externally.	OK. Any method that network operator chooses locally for implementation of the SR is fine.
729-731	20	Potential exists of a security product that might preclude this recommendation. More research should be conducted on this point.	OK.
741	21	Support, this is done today.	Good.
753	22	Support, this is done today.	Good.
759-763	23	Support only if the IXP wants the LAN prefix to be globally visible. Member ASNs of an IXP should not originate IXP LAN prefixes, which is sometimes done if members incorrectly redistribute connected prefixes to BGP.	OK.
780-787	24	Support, however, agreed “holes” in prefix blocks the AS originates will need to be allowed for customer mobility and possible security products.	OK. Point well taken. If the AS has suballocated to customers and hence not originating those subprefixes or "holes", then they are not included in the filtering.
788-796	25	Blocking routes learned from other Lateral Peers (via AS_PATH) should be included.	Good catch. This SR (now SR 26) has been updated per suggestion. Old SR #28 (now SR #29) is also updated per this suggestion.
800	26	Support but may need some allowed prefix blocks for customer mobility and security products.	Comment similar to that for SR #24 above. Please see response there.
876-877	34	The size and scale of these prefix lists would almost certainly not be possible for transit providers. We could end up with interface prefix lists of many tens of thousands of lines. These prefixes consume limited resource and therefore are not scalable. There may also be prefixes without ROAs for certain circumstances.	OK. This security recommendation is possibly more applicable to smaller ISPs than larger ISPs. Text or footnote added below the SR to suggest that.
944	35	Support, current common practice.	OK.

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
948-950	36	This is a good idea in many circumstances, especially for leaf networks. For transit service providers there are corner cases where this may not be a good idea.	Sounds good. There was a NANOG list discussion and several AS operators shared that they perform the ingress tagging and use it to ensure route leaks are prevented. https://mailman.nanog.org/pipermail/nanog/2016-June/thread.html#86348
1108-1112	37	Support in these limited circumstances.	OK.
1117-1125	38	Suggest enterprise networks should announce all their IP space to all providers unless there are specific reasons not to. For instance, a /16 prefix could be announced to two upstream providers, then announce specific /17s of the /16 to upstream providers to balance inbound traffic. This recommendation limits certain legitimate load balancing and backup configurations for enterprise customers in order to support uRPF. Engineering decisions like this should be made by enterprises.	We have modified and reworded the SR (now SR 41 in Draft2) keeping your observations in mind.
1126-1131	39	As per SR 38, in these circumstances, the operator should provide covering supernet announcements. An Enterprise cannot always rely on AS_PATH prepending to affect routing across their transit ISPs. The ISP could simply override that via Local Preference. There will be cases where an enterprise will have to stop advertising prefixes on one ISP. That does not preclude the Enterprise from sending traffic to that ISP, however.	We have modified and reworded SR 39 (now SR 42 in Draft2) keeping your observations in mind. Concerning your comment "There will be cases where an enterprise will have to stop advertising prefixes on one ISP ...", interestingly there is an IETF draft in progress (soon to be an RFC/BCP) that proposes an enhanced feasible path uRPF (EFP-uRPF) to effectively address that scenario https://datatracker.ietf.org/doc/draft-ietf-opsec-urpf-improvements/ . But since it is work in progress, there is no SR in SP 800-189 that is based on the IETF draft at this time.

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
1139-1147	41/42	Due to issues with vendor inter-operational support the use of Feasible Path uRPF is not globally adopted and may not be for some time. Loose is recommended for large providers at this time.	We have modified and reworded the SRs (now SR 45, SR 46 in Draft 2) keeping your observations in mind. Yes, large providers would more likely use loose uRPF. The goal here is not global deployment of feasible path uRPF (FP-uRPF). Smaller ISPs (those closer to the edge of the Internet) can use FP-uRPF or the EFP-uRPF (soon-to-be RFC as noted above) when permitted by their specific scenario. FP-uRPF or EFP-uRPF can be deployed independently on a per edge-router basis. So it is not clear that router interoperability is an issue here.
1148-1149	43	Support, current best practice. Customers should have the option of overriding this recommended practice if needed to support their engineering goals.	OK.
1153-1156	44	See reasoning in SR 41/42.	Please see responses above re: SR 41/42. We have modified and reworded the SR (now SR 48 in Draft 2) keeping your observations in mind.
1157-1160	45	Support except large ACL lists consume router resources and could cause a network to become unstable. This may not scale for large providers.	Yes, the recommendation calls for loose uRPF or ACLs. So an ISP can choose whichever is more feasible in their scenario.
1161-1164	46	Support though due to resource limitations these will be simple blocking ACLs where the ACL is likely to be a standard one applied to all interfaces. This will likely block obvious martians only.	Yes, the intent is to use simple blocking ACLs.
1170-1171	47	The construction of BGP prefix-lists is already complex enough and not dynamic enough in nature. Adding more dynamic content (published ROAs) to this process is not likely to have a positive impact. Generating large ACLs based on ROAs is not a good idea for the resource reasons already discussed. Again, scale and scope are different for large providers.	We have reworded the SR. Yes, this SR is not for larger ISPs. Smaller ISPs (those closer to the edge of the Internet) can have simple ACLs or RPF lists based on announced prefixes (by customers) augmented by the relevant ROAs that pertain to their customer cone ASes. The customer cone size would be typically small for the participating small ISP.

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
1198	48	While true port 0 is reserved, non-initial fragments have no ports assigned and will show up as port 0. DNS EDNS0 and other non-initial fragments will be blocked by such a filter. This could render DNSSEC inoperable.	Thank you for the information. We have studied the issues you and other reviewers have raised, and deleted the previous SR 48 related to port 0. There is no mention of dropping port 0 traffic in the revised document.
1200-1203	49	Support, generally done already with control plane access lists where possible.	OK.
1204	50	Support, current best practice.	OK.
1206	51	Support, current best practice though care needs to be taken not break tools such as traceroute where still required. Also consider ICMP (and other) traffic.	OK.
1212-1219	52	Allowing FlowSpec from customers is not considered safe at this time and is not recommended. Some research has been conducted that will investigate using FlowSpec between peers that may offer the improvements you suggest. See the talk at NANOG 71, https://pc.nanog.org/static/published/meetings/NANOG71/1447/20171003_Levy_Operationalizing_Isp_v2.pdf. Per-source monitoring can be 'very' resource intensive and take down the resolver if done locally. This is especially true if spoofed packets are used.	The Flowspec technology is maturing and can be used with customers just like with peers. It is just a matter of having a proper contract/agreement in place that helps the customer appreciate the benefits of Flowspec and its risks (if misused). The potential for misuse can be minimized by training. (We have viewed the video of the NANOG 71 presentation and corresponded with the presenter. There is also a more recent NANOG 75 presentation by Charter: https://youtu.be/rKEz8mXcC7o . We will keep a tab on the issues you've raised for future revisions of the document.)
1232	53	Support, current best practice.	OK
1235-1239	54	Applicable to smaller providers, will not work for global DNS providers. What is the definition of network in this context (subnet/ASN/etc.)? This may be unimplementable and unrealistic at an ISP level with a global footprint. The anycast infrastructure should provide similar protection assuming the associated unicast address doesn't answer outside known friendlies. Perhaps an exception for anycast resolvers and a stronger definition of network?	The SR (now SR 57 in Draft2) has been modified to reflect your comments.

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
1240-1242	55	We believe you were referring to DNS servers for these recommendations. IPv4 anycasting isn't always apparent to anyone outside of the AS. Was this intended for IPv6? The use of RRL in large deployments may cause performance issues and end up slowing down DNS.	We have taken your comments for SR 55 and SR 56 into consideration. Old SR 55 and SR 56 had overlap in their objective. They are now merged into a single SR (now SR 58 in Draft2).
1243-1244	56	This seems too vague to be implemented. Can you elaborate your intent?	As stated above, old SR 55 and SR 56 are now merged into a new SR (now SR 58 in Draft2). Also see the reasoning provided below this SR 58 in Draft2).
1246-1261		There has been some discussion on this topic in the Network Operating groups. We would like to hear some feedback on limiting the types used in this type of communication. Refer to the above listed NANOG 71 presentation. FlowSpec may be not be supported or supported well, the DBHF or SBHF may be the correct method in many cases.	OK. We have viewed the video of the NANOG 71 presentation and corresponded with the presenter. Also viewed the NANOG 75 presentation mentioned above (see response for SR 52 above). We will keep a tab on this issue.
1271-1273	57	Support DBHF, SBHF could be a possibility.	OK
1274-1276	58	Support where capable. Start with small well controlled architecture and discover issues first. Implementing on EDGE for all customers is not recommended at this time.	OK. Point well taken.
1277	59	Support, current best practice for DBHF. Adding ROA for the future is a nice concept and suggest any/all DBHF prefixes are checked against IRR based prefix-lists.	SR 59 (now SR 62 in Draft2) modified accordingly.
1285	60	Support, current best practice.	OK
1288-1295	61/62	Support but must be used with caution.	OK
		Comments set #2	

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
477-488		It may be worth adding that RIPE NCC, APNIC and AfriNIC each run Internet Routing Registries (IRRs) that are integrated with Regional Internet Registry (RIR) allocation data that facilitates stronger authentication schemes. These are documented in RFC2725 Routing Policy System Security [4]. However, while the IRR-related recommendations are important practices in line with the current operational reality, it is also important that these recommendations do not discourage RPKI deployment. RPKI provides an even stronger authentication and validation framework for network operators.	We've incorporated these suggestions in Section 4.1 in the revised draft.
485-487		Along with ARIN, LACNIC also runs a Shared Whois Project (SWIP). However, unlike ARIN, LACNIC does not provide an IRR of their own.	Updated Section 4.1 per your suggestion.
594-595	6	About Security Recommendation 6: Transit providers cannot provide this service for the address space they do not hold. Instead, Security Recommendation 6 could read “Transit providers should provide a service where the customers that use the space sub-allocated from their providers can create, publish, and maintain ROAs for their prefixes.”	We have reworded the security recommendation based on your suggestion.
620-626	13	It seems the recommendation specifies how BGP-OV should be implemented. A BGP router should validate received routes through a local RPKI cache server, and base the routing decisions on RPKI validity. BGP-OV is implemented by the majority of major router vendors.	This is explained in the text preceding SR 13 (now SR 14 in Draft 2) in Section 4.3 and also new text is added immediately following SR 14 in Draft2. However, in accordance with your observations, we have also updated the text in the SR 13 (now SR 14 in Draft 2).

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
627-629	14	Security Recommendation 14 could read “In partial/incremental deployment state of the RPKI, BGP-OV should be augmented by using the prefix filters generated from the IRR data, and customer contracts.”	BGP-OV and "prefix filtering" (see Sections 4.4 and 4.5 or RFC 7454) have somewhat different connotations. But we have slightly changed the wording in the SR in question based on your suggestion.
709-714	18	If "whitelist" filtering (based on the IRR+RPKI) is used, then this recommendation is no longer needed. In general, different approaches (and types of filters) are used for different types of peers. For example, building "whitelist" filters for transit providers is rare, while for customers it is quite common – as seen in MANRS. One possible approach is to provide descriptions of various types of filters and technologies in 4.3/4.4. and move relevant recommendations to section 4.5., as is done in RFC7454 BGP Operations and Security.	We cannot get rid of traditional prefix filtering (Sections 4.4 and 4.5) as long as RPKI/BGP-OV adoption is not complete with nearly 100% ROA coverage. For example, when the BGP-OV result is NotFound, the router needs to reject the route if the prefix is unallocated.
780-781	24	Security Recommendation 24 could be further strengthened by advising that providers explicitly whitelist filtering of peers and their customer cones, as implemented by the members of the MANRS IXP Program.	We feel that this security recommendation is good as is for now.
788-796	25	The “customer cone” is mentioned in the text of Security Recommendation 25, but omitted from the following list. Similarly, a definition of the "Customer cone" prefix filter may be helpful to readers in section 4.4.	The 'following list' does not mention customer cone ASes or prefixes since those are allowed (not filtered). Added that prefixes received from the AS's transit providers and other lateral peers should not be sent to the lateral peer in question (this is basically avoidance of route leaks). We have now included defections of customer cone, lateral peer, etc. earlier in the document in Section 2.3.
		Comments set #3	

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
		[Our organization] supports NIST’s goals and highlights parallel industry efforts to tackle difficult security issues. [Our organization] appreciates that NIST has drawn heavily from industry-driven work and urges NIST to continue to do so.	Thank you.
		SP 800-189 builds upon years of work by the private sector, often in tandem with the government. SP 800-189 cites numerous products created by private sector entities, including Cisco, Comcast, Juniper Networks, and Symantec. Many of the government products cited in SP 800-189 were created with private sector partners, such as the CSRIC documents and the Botnet Road Map. SP 800-189 also notes that the MANRS Implementation Guide—developed by the Internet Society—can “be thought as complementary to [SP 800-189] since it provides implementation guidance for some of the solution technologies described in [various sections of SP 800-189].” [Our organization] applauds NIST’s use of collaborative public and private sector work.	Thank you for these observations and appreciation. We (NIST in general) have many past and ongoing collaborations with many industry partners in the Internet infrastructure / cyber security areas, exemplified by many joint contributions (IETF RFCs, Internet Drafts, Botnet Road Map, Cybersecurity Framework, etc.).

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
		NIST should incorporate the recently adopted Communications Security, Reliability and Interoperability Council ("CSRIC") report on Best Practices and Recommendations to Mitigate Security Risks to Current IP-based Protocols ("CSRIC Report"). Given the closeness in timing of the CSRIC Report adoption and this Draft's comment cycle, it is not surprising that there are gaps between what NIST is recommending and what is in the CSRIC Report. ....While the CSRIC Report puts forward significantly less prescriptive recommendations and includes a more robust discussion of current limitations and future developments in this area, there is considerable overlap between the documents. .... [Our organization] would be happy to work with NIST to ensure a close mapping between the CSRIC Report and SP 800-189. It may be especially prudent for SP 800-189 to note the areas in which CSRIC is aspirational or identifies recommendations that require refinement.	The CSRIC report (CISRIC VI, Group 3) was published (Mar 2019) after this NIST draft SP 800-189 was put out for public comments (Dec 2019). We have since read the CSRIC report in full. As you have noted there is significant commonality in terms of objectives for routing security and DDoS mitigation between the two documents. In Sections 4.11 and 6.5 of the CSRIC report, a somewhat detailed and complimentary view of the NIST 800-189 draft is offered. They even go on to say that the NIST 800-189 "recommendations should be reviewed by future CSRICs for inclusion in future DNS/BGP reports". As we have revised the NIST draft, we have kept in view the desired alignment with the CSRIC report and other efforts such as MANRS. We have added the following wording in the Introduction and in Section 4: "This document addresses many of the same concerns as highlighted in [CSRIC4-WG6] regarding BGP vulnerabilities and DoS/DDoS attacks, but goes into greater technical depth in describing the standards-based security mechanisms and in providing specific security recommendations."
		[Our organization] applauds SP 800-189's incorporation of several international programs, including the Mutually Agreed Norms for Routing Security ("MANRS") initiative and the Resource Public Key Infrastructure (RPKI) system. [Our organization] urges NIST to further highlight this work and emphasize that "global collective action" is necessary to address routing security threats.	Thank you. Yes, your observations are well taken. As mentioned in the response immediately above, we are committed to maintaining coordination with other interested groups that have shared interest in promoting security practices related to Internet routing, and ensuring that our document is in alignment with other efforts such as MANRS and the CSRIC report.

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
		These multifaceted efforts overlap. Publishing SP 800-189 is a “task” under the Botnet Road Map’s Workstream 1: Improvements to Routing Security, the goal of which is to “advance deployment of longstanding anti-spoofing techniques and newer technologies to protect against route hijacks and leaks.” But this workstream includes additional tasks such as (1) “Remov[ing] Legal and Policy Barriers to Resource Public Key Infrastructure (RPKI) Adoption” (contributors are Academia, Internet engineers, NIST, NTIA, DOD, and regional and local Internet registries); and (2) “Extend[ing] Adoption, Awareness and Application of Anti-Spoofing Mechanisms” (contributors are Internet infrastructure owners and operators, civil society, NIST, NTIA, and DHS). SP 800-189 is one critical step among many.	Thank you for the observations; also for your comment "SP 800-189 is one critical step among many." As you've noted, NIST is involved closely (along with many other stake holders) in multiple efforts which are all important to achieving overall success.

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
		[Our organization] applauds the voluntary nature of SP 800-189 as applied to industry. [Our organization's] understanding is that NIST intends for its guidance to be implemented in “federal enterprise networks” and “the service agreements for federal contracts for hosted application services and Internet transit services,” but not on industry more broadly. However, NIST should explicitly incorporate this understanding into the final draft of SP 800-189 because the document creates ambiguity in two ways. First, it notes “[t]he guidance will also be useful for enterprise and transit network operators and equipment vendors in general.” The reference to these parties “in general” could create the impression that SP 800-189 is targeted at private-sector entities. Second, the publication’s uses some broad language that could be interpreted as applying to the private sector. For example, Security Recommendation 1 says: “All Internet Number Resources (e.g., address blocks and ASNs) should be properly registered . . .” Additionally, many Security Recommendations use broad terminology, such as “enterprise,” “ISP,” and “transit provider,” which are not limited to the government or its contractors. Given these ambiguities, NIST should disclaim any binding application of its recommendations on industry.	Based on your suggestion, we have replaced "will also be useful" with "may also be useful" in this sentence in the audience section: "The guidance may also be useful for enterprise and transit network operators and equipment vendors in general." The Draft SP800-189 states upfront, "This publication may be used by nongovernmental organizations on a voluntary basis and is not subject to copyright in the United States." Nothing elsewhere in the Draft is intended to imply otherwise. The security recommendations say "should" rather than use stronger language such as "must". Please note that the MANRS and CSRIC documents reach out to a broader audience and make similar recommendations as in Draft SP800-189. Both BGP and DNS are global distributed protocols and hence voluntary participation in security practices by as many entities (ISPs, enterprises) as possible is helpful to protect all users of the Internet from the impacts of BGP hijacks, DDoS, etc.

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
		Moreover, it may be premature to incorporate all aspects of SP 800-189 into federal contracts and network management. In particular, the CSRIC Report highlights a paper from the University of Pennsylvania that discussed "legal barriers that may be hindering RPKI adoption in North America." One such legal barrier is "the North American RIR's (Regional Internet Registry) requirement for RPKI users to enter a Relying Party Agreement and certain terms in that agreement." NIST should recognize these issues as it revises SP 800-189.	NIST has been actively involved in fostering and facilitating support for the University of Pennsylvania work. The North American (ARIN) region is lagging behind while about 46% address space in the RIPE region is already RPKI registered. In response to the UPenn work, ARIN has expressed its commitment (at NANOG meetings) to work with the community to resolve the perceived legal barriers. Having said that, we have followed your advise to incorporate wording in the Draft SP 800-189 to recognize these issues.
		Comments set #4	
		1) Microsoft LDAP servers (Active Directory) support LDAP over UDP (also referred to as "CLDAP"). Reflection attacks against these servers are now common, so LDAP should probably be listed in Table 1 in Section 5.4. Some additional info can be found here: https://www.akamai.com/kr/ko/multimedia/documents/state-of-the-internet/cldap-threat-advisory.pdf	Yes, LDAP is now included in Table 1. We have also cited the reference you've provided.
		2) Line 427 lists the acronym as DoS, but it should be DDoS	Correction made.

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
		3) Line 449 says “query and response are contained in a single packet”, which makes it sound like one packet contains both a query and response. This should probably be changed to “query and response are each contained in a single packet” to make it clearer that there is one packet for the query plus one packet for the response.	Yes. Suggested rewording is incorporated.
		Comments set #5	
		In this document on page 31, line 1192 there is a table showing common DDoS amplification ports. I have some comments regarding it: Looking at 25 million subscribers, I have never seen DNS DDoS attacks on port 853 or 953. I have never seen RPC DDoS attacks on port 369. I do see these attacks on port UDP 111. I have never seen any RIPng attacks on port 521. I have seen RIPv1 attacks on UDP port 520. LDAP on UDP port 389 is not mentioned. That is the second most common DDoS attack port that I see. I have never seen any RTSP DDoS attacks on 554 or 1755.	Thank you for sharing. We have included some of this information in the document; we've updated Table 1.
		DNS, LDAP and other DDoS amplification protocols generate a lot of UDP fragment traffic. We do policing / rate-limiting of UDP fragments at our peering edge to reduce the impact of DDoS amplification traffic. It might be worthwhile to include this.	Thank you for the suggestion. A new SR 58 in Draft2 has been added. We have added text just above the SR to explain the motivation.
		We also do policing on LDAP, SNMP, and RPC to reduce the impact of DDoS attack using these vectors.	Thank you for sharing.

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
		NTP Monlist traffic can be mitigated with an ACL that blocks the monlist reply traffic with the maximum number of IPs defined in a packet (6). These packets are 468 bytes for IPv4 and 488 for IPv6 excluding the ethernet frame. So filtering on UDP with a source port of 123 and a packet length of 468 will pretty much stop NTP amplification attacks.	Thank you for the information.
		On line 1224 it is stated to do a RTBH using Flowspec. A RTBH can be done without Flowspec. The advantage of Flowspec is that the filtering can be much more surgical blocking just the attack traffic and permitting all other traffic. I would call this more a filter than a RTBH.	Thanks for the observations. Yes, Flowspec facilitates a more precise and automated way of specification of IP addresses that must be blocked.
		Comments set #6	
		Security recommendation 1: ARIN allows for setting the Origin AS in the RIR database, though this is optional, it is a stronger attestation than IRR data at present. Example: https://whois.arin.net/rest/net/NET-128-3-0-0-1/pft?s=128.3.0.0 . I would recommend that the contact information be up to date, and also that the resources be covered by an appropriate registration services agreement. (this is required for #5)	Wording in SR 1 updated per your suggestion. New SR 2 added per your suggestion. Also new text added in Section 4.1 corresponding to this new SR.
		Security recommendation 6: Is this actually reasonable? My understanding of the software ecosystem to accomplish this is that it is immature at best. There has also not been any true testing of this approach at scale (hundreds to thousands of ISP's, for example) to my knowledge.	As you've observed, SR 6 is one way to facilitate SR 9. (Note: These are old SR numbers; the corresponding new SR number are 7 and 10 in Draft2.)

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
		Security Recommendation 9: Is there a way around this? Waiting for the last customer to deploy would lead to the case where large blocks would be unlikely to ever get covered. (Or I guess this helps explain why #6 needs to be viable)	Please see response above.
		Security Recommendation 12: the box for Enterprise should be checked 'X'	Done.
		Security Recommendation 13: Is "white list" a standard term in this context? If the router has this list, how is it to be applied?	ROA payloads contain authorized {prefix, maxlength, origin ASN} information. It seems appropriate to call this white list (as used in the router). The RPKI cache typically passes on the validated {prefix, maxlength, origin ASN} tuples to the router.
		Security Recommendation 15: Only Drop-Invalid works in practice. Do not recommend prefer. Some IXP's and a handful of networks are now running drop-invalid.	We plan to monitor operator experience that will be reported over time and then refine this recommendation in a future release. Yes, the goal should be to drop all invalid routes.
		Security Recommendation 16/17: Is this language strong enough? See "MaxLength Considered Harmful to the RPKI" doi:10.1145/3143361.3143363 Best practice I would believe is that the length MUST match reality, not SHOULD.	We have changed the wording now from "The maxlength in the ROA should preferably not exceed …" to "The maxlength in the ROA should not exceed …". 'MUST' would make sense but none of the RFCs or drafts (RFC 7115, [maxlength]) have used the 'MUST' language yet.
		Security Recommendation 18: Why is this only an IPv6 recommendation (what about IPv4, or I guess #19 is adequate for v4)? Also, due to sparse allocation practices from RIR's and ISP's, is this even realistic in practice? https://www.team-cymru.org/Services/Bogons/fullbogons-ipv6.txt has 100k lines at this point and may exceed the FIB of low-cost hardware. An example would be that an enterprise that only takes a default route from their ISP may not choose hardware with a large FIB.	Prefix filtering is performed in the control plane. So, the permissible IPv6 prefix list will not be stored in the FIB. But we need to keep an eye on this from a performance (look up delay) point of view in the future when the list possibly grows much bigger.

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
		Security Recommendation 21: One can also filter covering prefixes as well. For example, ESnet will not accept IPv4 < /8 or IPv6 < /11	We've added new text in Section 4.4.3: "It may be noted that some operators may choose to reject prefix announcements that are less specific than /8 and /11 for IPv4 and IPv6, respectively."
		Security Recommendation 23: The IXP should use an ROA for the lan prefix with AS0, purposefully making it invalid, taking specific care not to override it, (RFC6483 section 4).	SR 23 (now SR 24 in Draft2) states that RS's LAN prefix should be announced to the RS's member ASes, and that a member AS should reject any more specifics prefixes (of the IXP announced prefix) from any of its eBGP peers. It is expected that the ISP would create a normal ROA for the LAN prefix (with maxlength equal to the prefix length). That would make any more specific prefix announcements in consideration Invalid. So, it seems not necessary to create an AS 0 ROA in this case.
		Security Recommendation 48: 0 is not a reserved source port. RFC8085 states "A UDP sender SHOULD NOT use a source port value of zero." rather than MUST NOT. Notably any application written before 8085 may also still be following the guidance from RFC768 "Source Port is an optional field, when meaningful, it indicates the port of the sending process, and may be assumed to be the port to which a reply should be addressed in the absence of any other information. If not used, a value of zero is inserted". Services such as interdomain multicast are an example still following the old convention. There are others. Secondly, filtering UDP 0 on some platforms is known to be problematic due to syntactic issues in the filter language leading operators to mistakenly drop all packet fragments. See https://kb.juniper.net/InfoCenter/index?page=content&id=KB31437	Thank you for the information. We have studied the issues you and other reviewers have raised, and deleted the previous SR 48 related to port 0. There is no mention of dropping port 0 traffic in the revised document.

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
		Security Recommendation 49: This recommendation or an additional one should require use of TTL checking, (GTSM).	See newly added Section 4.10 on GTSM and Security Recommendation 39. Thank you.
		Comments set #7	
		Security Recommendation 15 provides a brief discussion on typical policy choices. We would like to see further details on this recommendation such as comparing and contrasting these typical policies and implications on the overall security posture. In future releases of the document a potential example may provide additional insight.	Yes, we can plan to do this in a future release when further network operator experience with BGP-OV is possibly reported.
		SRs 24, 25, 26, 31, 32 provide suggested prefix filters that should be used to enhance security. Is it implied that all of the filters suggested for a SR should be applied as a group? How is the security strength affected if one or more of the filters in the list are not implemented? Further clarification would be helpful in this context.	These recommendations are provided based on the peering relationship (transit provider, lateral peer, or customer facing) and also based on the direction (inbound or outbound). An operator should follow the relevant recommendations based on the relationship and the direction on the interface in consideration. Each operator accrues benefit locally at their AS by implementing the relevant recommendations.

		Note: SR# = Security Recommendation #	
Lines:	**SR#**	**Comments**	**Authors' response**
		In our opinion, AS Path Validation will further enhance security and reduce the attack surface, which is described in section 4.7 as an Emerging/Future capability. There has been substantial concentrated IETF community effort to standardize BGPsec. Additionally, BGPsec, both from an architectural and operational perspective, has been described comprehensively in RFCs and peer-reviewed publications. Functional and high-performance prototype implementations of BGPsec/BGP-PV are currently available. Future versions of the document should provide security recommendations on deploying BPGsec.	Yes, we can plan to do this in a future release when BGPsec (path validation) possibly begins to gain traction.
		Comments set #8	
		[Should you] keep SAV separate from BGP security? The deployment, policies, and scope of work addresses different people inside of an organization.	The focus of the guidance is on the services between an enterprise and their ISPs, not just BGP. So, the DDoS issues are just another aspect of the main focus. We have stated upfront in the document, “This document provides technical guidance and recommendations for technologies that improve the security and robustness of interdomain traffic exchange. The primary focus of these recommendations are the points of interconnection between enterprise networks, or hosted-service providers, and the public Internet.” Our expectation is that Federal CIOs and IT security folks who help write contracts with ISPs would find this comprehensive approach more useful.
		[Seem to be] doing a survey of existing uRPF capabilities? VRF mode is not mentioned while much of the larger Cisco equipment can deploy it.	We didn’t see it as a survey. The idea was to provide a brief overview of the underlying technologies before listing a set of recommendations in each category (BGP origin validation, prefix filtering, SAV, etc.). VRF mode is included and discussed in the revised version.

		Note: SR# = Security Recommendation #	
Lines:	SR#	Comments	Authors' response
		IETF “future” work should not be mentioned in a NIST document until there is deployed working code in the industry.	There is a difference between NIST 800 SP recommendations and FISMA requirements. With the latter, we would strongly stick with only available technologies. Traditionally, NIST 800 SPs provide brief technology overview including what is in the pipeline, i.e., evolving technologies that address the gaps in security coverage. As you may have noticed, where we provide promising pointers to evolving technologies, we also carefully state, “ … this section briefly describes the technology and standards effort but does not make a security recommendation concerning use of …..”. The idea is that Federal CIOs and IT security folks should not only be aware that the existing standards and technologies have limitations, but also that there is evolving work that is addressing the gaps. This is done very briefly in each category. We devoted about 2 pages (out of about 70 pages in the whole document) to evolving technologies.