

Applicability Statement for DNS MIB Extensions

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

This document explains why, after more than six years as proposed standards, the DNS Server and Resolver MIB extensions were never deployed, and recommends retiring these MIB extensions by moving them to Historical status.

1. History

The road to the DNS MIB extensions was paved with good intentions.

In retrospect, it's obvious that the working group never had much agreement on what belonged in the MIB extensions, just that we should have some. This happened during the height of the craze for MIB extensions in virtually every protocol that the IETF was working on at the time, so the question of why we were doing this in the first place never got a lot of scrutiny. Very late in the development cycle we discovered that much of the support for writing the MIB extensions in the first place had come from people who wanted to use SNMP SET operations to update DNS zones on the fly. Examination of the security model involved, however, led us to conclude that this was not a good way to do dynamic update and that a separate DNS Dynamic Update protocol would be necessary.

The MIB extensions started out being fairly specific to one particular DNS implementation (BIND-4.8.3); as work progressed, the BIND-specific portions were rewritten to be as implementation-neutral as we knew how to make them, but somehow every revision of the MIB extensions managed to create new counters that just happened to closely match statistics kept by some version of BIND. As a result, the MIB extensions ended up being much too big, which raised a number

of concerns with the network management directorate, but the WG resisted every attempt to remove any of these variables. In the end, large portions of the MIB extensions were moved into optional groups in an attempt to get the required subset down to a manageable size.

The DNS Server and Resolver MIB extensions were one of the first attempts to write MIB extensions for a protocol usually considered to be at the application layer. Fairly early on it became clear that, while it was certainly possible to write MIB extensions for DNS, the SMI was not really designed with this sort of thing in mind. A case in point was the attempt to provide direct indexing into the caches in the resolver MIB extensions: while arguably the only sane way to do this for a large cache, this required much more complex indexing clauses than is usual, and ended up running into known length limits for object identifiers in some SNMP implementations.

Furthermore, the lack of either real proxy MIB support in SNMP managers or a standard subagent protocol meant that there was no reasonable way to implement the MIB extensions in the dominant implementation (BIND). When the AgentX subagent protocol was developed a few years later, we initially hoped that this would finally clear the way for an implementation of the DNS MIB extensions, but by the time AgentX was a viable protocol it had become clear that nobody really wanted to implement these MIB extensions.

Finally, the MIB extensions took much too long to produce. In retrospect, this should have been a clear warning sign, particularly when the WG had clearly become so tired of the project that the authors found it impossible to elicit any comments whatsoever on the documents.

2. Lessons

Observations based on the preceding list of mistakes, for the benefit of anyone else who ever attempts to write DNS MIB extensions again:

- Define a clear set of goals before writing any MIB extensions. Know who the constituency is and make sure that what you write solves their problem.
- Keep the MIB extensions short, and don't add variables just because somebody in the WG thinks they'd be a cool thing to measure.
- If some portion of the task seems to be very hard to do within the SMI, that's a strong hint that SNMP is not the right tool for whatever it is that you're trying to do.

- If the entire project is taking too long, perhaps that's a hint too.

3. Recommendation

In view of the community's apparent total lack of interest in deploying these MIB extensions, we recommend that RFCs 1611 and 1612 be reclassified as Historical documents.

4. Security Considerations

Re-classifying an existing MIB document from Proposed Standard to Historic should not have any negative impact on security for the Internet.

5. IANA Considerations

Getting rid of the DNS MIB extensions should not impose any new work on IANA.

6. Acknowledgments

The author would like to thank all the people who were involved in this project over the years for their optimism and patience, misguided though it may have been.

7. References

- [DNS-SERVER-MIB] Austein, R. and J. Saperia, "DNS Server MIB Extensions", RFC 1611, May 1994.
- [DNS-RESOLVER-MIB] Austein, R. and J. Saperia, "DNS Resolver MIB Extensions", RFC 1612, May 1994.
- [DNS-DYNAMIC-UPDATE] Vixie, P., Thomson, S., Rekhter, Y. and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, April 1997.
- [AGENTX] Daniele, M., Wijnen, B., Ellison, M., and D. Francisco, "Agent Extensibility (AgentX) Protocol Version 1", RFC 2741, January 2000.

8. Author's Address

Rob Austein
InterNetShare, Incorporated
325M Sharon Park Drive, Suite 308
Menlo Park, CA 94025
USA

EMail: sra@hactrn.net

9. Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

