

1 Comments on EXI Profile, W3C Draft 31 July 2012

2 Rumen Kyusakov

3 2012-09-28

4

5 DISCLAIMER: The views and opinions expressed here are solely based on the author's experience and knowledge of the EXI specification and the EXI Profile. Many of the conclusions and comments might be biased by misunderstanding of these specifications or the author's lack of broader knowledge in the area.

6

7 1)

8 Subject: Evaluations of EXI in application areas with memory restrictions

9 Section: 1. Introduction

10 Paragraph: "Certain evaluations of EXI in the context of such areas..."

11

12 Comments:

13 It is good to have access to these evaluations. From my own perspective and use cases this profile will not bring much benefits. I am working with very small messages over low-bandwidth wireless links (up to 250 kb/s). Of course, memory is an issue but the compression is even more important. It would be good to have some statistics on what are the trade-offs (memory usage vs compression) by using the Profile in different use-cases (size and structure of the messages).

14

15 2)

16 Subject: Limiting the number of build-in grammars and Limiting the evolutions of build-in grammars

17 Section: 2. Grammar Capping

18 Paragraph: Whole section

19

20 Comments:

21 In the grammar capping section it is given a single mechanism to limit the number of build-in grammars by using the xsi:type attribute (possibly with xsd:anyType). There is no mentioning of the second very distinct mechanism described later on: limiting the evolutions of build-in grammars by restricting the number of dynamic productions inserted. Mixing up between these two mechanisms in the spec is the main source of confusion for me.

22

23 Example:

24 In the second paragraph:

25 "Note that the EXI profile can only limit grammar learning for schema-informed EXI streams but not for schema-less EXI streams. In the case where no schema is available, the "schemaId" element may be set to the empty value, so that all grammars derived from the built-in XML Schema types become available through xsi:type grammar switching, in particular the grammar corresponding to the xsd:anyType complex type."

26

27 The EXI Profile can also limit the "grammar learning" even in schema-less streams with the second mechanism (limiting dynamic productions). Without the information for the second mechanism the paragraph should read 'the "schemaId" element MUST be set to the empty value' and not 'may'.

28

29 (Note also the spelling error of "gramar" in the paragraph)

30

31 3)

32 Subject: Event code length for xsi:type attribute event

33 Section: 2.1 Grammar Learning Disabling Mechanism

34 Paragraph: "The xsi:type attribute event MUST always be represented by the AT(*) production whose event code length is 2."

35

36 Comments:

37 I am sure there is a good reason to require this but I cannot deduct it. Maybe a hint in the spec will make the job of the implementers easier. My guess is that this is connected to the separation of the "original" xsi:type attributes and the xsi:type attributes added because of the Profile.

38

39 (Note also: "For a given element E, the disabling of grammar learning E ..."
40 should be "For a given element E, the disabling of grammar learning for E ...";
41 "..., the following rule happens:" it is better "..., the following rules apply:")

42

43 4)

44 Subject: Limiting the evolutions of build-in grammars

45 Section: 2.1 Grammar Learning Disabling Mechanism

46 Paragraph: "If grammar learning is disabled in the case of a production insertion in a given grammar, named G, the following rule happens ..."

47

48 Comments:

49 This is the first time the second mechanism (Limiting the evolutions of build-in grammars) is mentioned in the spec. It is almost impossible to make the connection with the rest of the spec without reading it few/several times. In my opinion this mechanism should be clearly described and differentiated with the first one (using xsi:type).

50 A proper way of doing this is by mentioning it in the 2. Grammar Capping sections and then describing it as a separate subsection. Giving it a proper name would also help a lot as "grammar learning is disabled in the case of a production insertion" is definitely not a good one.

51

52 5)

53 Subject: Confusing description
54 Section: 2.2 Grammar Learning Disabling Parameters
55 Paragraph: "Grammar learning is disabled...
56 ...
57 Grammar learning is disabled in the case of a production ..."
58
59 Comments:
60 There should really be a better name for "Grammar learning is disabled in the case of a
61 production ..." A suggestion: "Disabling the evolution of build-in grammars"?
62 6)
63 Subject: The Profile parameters
64 Section: 4. Parameters representation
65 Paragraph: "In such a case, the actual profile parameters (or fine-grained capping strategies)
66 should be defined by an out-of-bound mechanism."
67 Comments:
68 The maximumNumberOfBuiltInElementGrammars and localValuePartitions parameters should not be
69 mandatory to be set/communicated-out-of-bound in my opinion. For the number of build in grammar
70 the encoder knows the number of the grammars so no need to be included in the header. On the
71 decoder side if the implementation is "smart" it won't create a new build-in element grammar
72 before examining the next event. If it is a xsi:type="SOME_EXISTING_TYPE", there is no need to
73 create a new build-in element grammar.
74 For the localValuePartitions again - no issues for the encoder if the parameter is not included in
75 the header. The decoder on the other hand can create local value tables only if it has enough
76 memory to do so. If it cannot support local value tables, it will only expect string values
77 represented literally in the stream. Receiving a string as a compact identifier will produce error
78 during decoding in any way if the decoder does not have enough memory.
79 7)
80 Subject: Unclear
81 Section: C Prefix Workarounds (Non-Normative)
82 Paragraph: All
83 Comments:
84 All the paragraph could be summarize by saying that the application should use only ONE prefix per
85 namespace. Then the two requirements are fulfilled.
86 8)
87 Subject: Striping of Profile xsi:type attributes
88 Section: E.2 Grammar Restriction Decoder Considerations
89 Paragraph: "An EXI profile decoder SHOULD strip any xsi:type attribute with the xsd:anyType value
90 from the infoset that corresponds to the grammar learning disabling mechanism."
91 Comments:
92 It is strange to have "SHOULD" in a (Non-Normative) section. In any case a non-Profile EXI decoder
93 will include the Profile xsi:type attributes.
94 8)
95 Subject: Summary
96 Section: All spec
97 Paragraph: All
98 Comments:
99 The conformance with the Profile does not guarantee in any way a bounded memory for processing
100 (for example not setting valuePartitionCapacity to 0 is enough to make a Profile complying
101 implementation consuming huge amount of memory). Moreover a non-conforming to the Profile
102 implementation can still use the Limiting the number of build-in grammars mechanism without
103 declaring that in the EXI header.
104 Without being an expert in the field my personal experience is that the EXI spec and even more the
105 EXI Profile spec are very much requiring the users to be aware of some implementation issues in
106 order to use the format in an optimal way.
107 For example, I was involved to some small extend in the Smart Energy Profile 2.0 standard/draft
108 that considers using EXI. Main goal is compactness but also lean processing. The knowledge on what
109 is the optimal way to set all the parameters (schemaId, preserve, selfContained, valueMaxLength,
110 valuePartitionCapacity and now the Profile parameters) is simply not there.
111 The great flexibility provided by having so many things to tune is working against the wide spread
112 adoption. The users should not be burdened by knowing what all these parameters mean especially
113 the Profile ones that are purely connected to the implementation.
114 Again, without claiming to be an expert and have a wide range of use cases in mind, I think a more
115 restrictive Profile that actually provides some guarantees and not so much freedom of setting
116 implementation parameters will be much more useful. In any case the actual use of the EXI Profile
117 will require some "Application-specific Profile" that defines the fixed values for all these
118 parameters. So why not preset all these parameters and provide some guarantees with a good memory
119 usage vs compression trade-off? I know one size does not fit all, but then why defining an EXI
120 Profile in the first place?
121 In any case, if every single application sets these parameters in a different way then the
122 implementations won't be able to interoperate due to some memory issues or size constrains. What I

see as a need is a Profile that provides some guarantees for resource constrained devices. For example, a conforming implementations should use only:

- 99 - maximumNumberOfBuiltInElementGrammars = 0
- 100 - valuePartitionCapacity = 0
- 101 - maximumNumberOfBuiltInProductions = 0
- 102 - localValuePartitions = 0
- 103 - selfContained = false
- 104 - preserve = all false
- 105 - fragment = false
- 106 - either strict schema mode OR "schemaId" set to the empty value schema-less mode
- 107
- 108

Note that by fixing the values of these parameters will not only make the RAM consumption predictive but also lower the programming memory which is also an issue in some cases. The implementations of the Profile will be much simpler as a codebase as compared to implementing the generic version of the Profiles as it is now. In fact, the Profile as it is now will just extend the code and make it more complex which I believe is the opposite of what is needed for embedded applications.