Paper for Consideration by S100TSM9

Exchange Sets containing only external resources

Submitted by:	PRIMAR
Executive Summary:	Several topics regarding external resources, and in particular support files, are discussed in this paper. The topics are related to: a) S-100 allowing external resources (support files) to be standalone entities in the Exchange set. b) S-100 allowing external resources (support files) to be shared by multiple products. c) S-100 allowing support files to be cancelled without the accompanying support file itself included in the data delivery (fileless cancellation).
Related Documents:	Additional TSM9 paper "Management of external resources extended descriptions" Additional TSM9 paper "SupportFileDiscoveryMetadata attributes supportedResource and resourcePurpose" S-100 5.0.0.
Related Projects:	

Introduction / Background

a) S-100 allows for external resources to be standalone entities in the Exchange Catalogue without an accompanying dataset file:
S-100 17-4.3:
"…External resources can support either datasets or Catalogues or can be standalone entities in the Exchange Set".

Technically this can be done by including only external resources (support files) in the Exchange Set, accompanied by the appropriate S100_SupportFileDiscoveryMetadata.

b) S-100 is allowing external resources (support files) to be shared by multiple products:
S-100 17-4.3:
"…As long as the mapping from the external resource metadata to the dataset metadata is unique it is valid, so multiple datasets are able to "share" common external resources within an Exchange Catalogue without ambiguity".

c) S-100 does not specifically define fileless cancellation of support files – but it is assumed the delete instruction could be transferred within the S100_SupportFileDiscoveryMetadata without having an accompanying support file in the exchange set.

The above described options in S-100 (a,b,c) pose new challenges for service providers currently moving from S-57 distribution towards the future of S-100 products distribution. This paper seeks to describe some of these challenges and suggested solutions – in addition to posting the question to reverse the option for one support file to be shared between multiple dataset files.

Analysis/Discussion

1. Consider reversing decision of support files being used by multiple products in S-100.

To make sure we use resources efficiently we want to question the use case for 1 support file being used by multiple products. Are we sure this approach is beneficial from a data producer and service provider viewpoint? We can see the potential positive benefits from the producer side, where support file management may require

less resources. However, the anticipated benefits should be well documented, and seen against the added complexity in the data delivery chain before implementing this solution in S-100.

The complexity we face is to ensure that all subscribers of a product (dataset) receive updated information whenever a supported resource is updated. When an updated version of a supported resource is received, it would be essential to know all the products that are using the resource, and we will have to build automation ensuring all subscribers to existing products get the new support file. Further complexity is also described throughout this document.

The complexity also brings in added risk that must be considered. As a service provider we see the challenges, but this will also have to be addressed further down the value chain by distributors packing their own Exchange Sets, and by OEMs implementing support for this flexibility in the end user systems. Potentially this can lead to invalid/missing data in end user systems. Is the benefits for the data producers worth the added complexity and the risk it brings in?

Conclusion: Discuss the benefits from supported resources referencing multiple dataset files, seen against the added complexity and risk this puts on the data delivery services. Consider reversing the option to use 1 support file by multiple datasets.

2. Mandating introduction of a new support file to be distributed with a dataset file.

If we look at the following scenario:
One support file is being referred to by several datasets, used by several customers. The support file is updated.

Publishing a new edition of the support file must trigger delivery of the new updated support file to all datasets referring to it. For a service provider this means an update has to be made available for all those subscribing to the involved datasets. This again means PRIMAR and other service providers will have to develop support for issuing updated support files to customers without accompanying data set file updates.

To build consistent support for this S-100 clearly must state that a dataset support file in its 1st edition must have as a minimum one accompanying dataset file. This is sensible as an external resource (at least the type of support file) will not be issued unless there is a dataset file referencing it. This will help service providers produce sensible checks ensuring that new support files are not introduced unless they are referred to by a dataset.

So, the rule should be that support files of 1st edition must be distributed together with a minimum of one accompanying dataset file. Support files of edition number 2 or higher could be distributed as standalone entities. (The specific use of the term “1st edition” should also cater for the possibility where a support file in its first release is > edition 1).

Conclusion: S-100 clearly must state that an S-100 dataset support file in its 1st edition must have as a minimum one accompanying dataset file.

3. Mandating that support files cannot be shared across products from different product specifications.

It is not possible for multiple S-100 products to share a common support file since the support file naming convention XXXYYYY….. requires XXX to reference the product code (for example 101 for S-101) and YYYY is the producer code. If the need should occur to share support files across product specifications, a product independent support file must be duplicated once for every product type. S-100 should clearly state that one support file cannot be shared between different S-100 products.

Conclusion: S-100 clearly must specify that an S-100 support file cannot be shared between different S-100 product specifications.

4. Mandate updating of S100_SupportFileDiscoveryMetadata attribute supportedResource under the specific conditions that can arise due to multiple datasets referencing the same external resource.

The situation where a number of dataset files are referencing the same external resource, and an update to that resource does not apply to all those dataset files, is further described in the TSM9 paper “Management of external resources extended descriptions”.

As that paper explains, we are facing a situation where a new support file is created for some of the datasets previously using the older support file.

This means the datasets that are referencing the new support file will have to be updated with new references from the dataset attribution to the new support file.

At the same time, the S100_SupportFileDiscoveryMetadata attribute supportedResource must be updated so that *references to the previous support file are removed*. This because S-100 has 2 mechanisms for support file referencing:

A. Datasets hold a reference to the supported resource as an attribute value (e.g. the attribute “file reference” in the S-101 data model).
B. The encoding of the S100_SupportFileDiscoveryMetadata attribute supportedResource, where identifiers of the datasets using the support file are encoded:

Attribute	supportedResource	Identifier of the resource supported by this support file	0..*	CharacterString	Conventions for identifiers are still to be developed and will be defined later

If the second mechanism mentioned above (B) is also being used, there must be rules for updating the supportedResource information to reflect the changes in the datasets no longer referencing this support file.

This probably also means S-100 has to state the relation between those two mechanisms and ensure rules are implemented to ensure the information is consistent if/when both mechanisms are used in the data exchange.

If the mechanisms are not intended to coexist within one product specification, this must be clearly stated in S-100. As of now implementing the second mechanism is not mandatory ref multiplicity [0..*]. However, as proposed in TSM9 paper: “05_S100_SupportFileDiscoveryMetadata attributes supportedResource and resourcePurpose” there may be a good use case for making supportedResource a mandatory metadata encoding requirement for all S-100 based product specifications.

Conclusion: S-100 has to describe the relation between the two mechanisms for supported resources referencing and have defined rules implemented to ensure the information is consistent if/when both mechanisms are used in the data exchange. If the mechanisms are not intended to coexist within one product specification, this must be clearly stated.

Conclusion: S-100 must specify that the S100_SupportFileDiscoveryMetadata attribute supportedResource must be updated when a new support file replaces the old support file.

5. Clarify how a support file can be deleted from the end user system without having an accompanying support file in the data delivery exchange set.

Another use case is the possibility for cancelling a support file without the accompanying support file itself (fileless cancellation). Technically this can be done by including the support file information in the exchange catalogue metadata and encode the S100_SupportFileDiscoveryMetadata attribute “revisionStatus” (Type = S100_SupportFileRevisionStatus) with the value 3 (deletion):

Attribute	revisionStatus	The purpose for which the support file has been issued	1	S100_SupportFileRevisionStatus	For example new, replacement, etc

S100_SupportFileRevisionStatus

Role Name	Name	Description	Code	Remarks
Enumeration	S100_SupportFileRevisionStatus	The reason for inclusion of the support file in this Exchange Set	-	-
Value	new	A file which is new	1	Signifies a new file
Value	replacement	A file which replaces an existing file	2	Signifies a replacement for a file of the same name
Value	deletion	Deletes an existing file	3	Signifies deletion of a file of that name

If this is the intended mechanism to delete support files a clarifying note should be added to Part 17 S100_SupportFileRevisionStatus.

Conclusion: Add clarifying note to Part 17 S100_SupportFileRevisionStatus stating that: "When a support file is to be deleted from the end user system, this is done by encoding value = 3(deletion). Such a delete instruction is conveyed by the information within this exchange catalogue without any accompanying support file".

6. Be aware of time differences when multiple datasets reference to the same support file.

PRIMAR has a service where producers are able to upload their data to a database for validation before being flagged as released when QA procedures are accepted and the producer decides it can go to the market. In this process there may be time differences to consider amongst when the involved datasets referencing the same support file enters the database and further is given the release flag.

This adds complexity to our system, where we probably have to implement a structure where we have to make sure a support file is made available before dataset files referencing it are being updated. This to ensure that the support file is available when the datasets referencing it are released. The consequence of not doing this would be datasets in the market referencing support files that are not yet available for the customers.

A consequence of this is that an end-user can receive a support file where none of his datasets are referencing it. A new update dataset is received later with an attribute update that includes a file reference to the support file. End-user systems should be aware of this use case situation in their implementations.

Conclusion: End-user can receive a support file where none of his datasets are referencing it - End-user systems should be aware of this use case situation in their implementations

Conclusions

- Discuss the benefits from supported resources referencing multiple dataset files, seen against the added complexity and risk this puts on the data delivery services. Consider reversing the option to use 1 support file by multiple datasets.
- S-100 clearly must state that an S-100 support file in its 1st edition must have as a minimum one accompanying dataset file.
- S-100 clearly must specify that an S-100 support file cannot be shared between different S-100 product specifications.
- S-100 has to describe the relation between the two mechanisms for supported resources referencing and have defined rules implemented to ensure the information is consistent if/when both mechanisms are used in the data exchange. If the mechanisms are not intended to coexist within one product specification, this must be clearly stated.
- S-100 must specify that the S100_SupportFileDiscoveryMetadata attribute supportedResource must be updated when a new support file replaces the old support file.
- Add clarifying note to Part 17 S100_SupportFileRevisionStatus stating that: "When a support file is to be deleted from the end user system, this is done by encoding value = 3(deletion). Such a delete instruction is conveyed by the information within this exchange catalogue without any accompanying support file".

- End-user can receive a support file where none of his datasets are referencing it - End-user systems should be aware of this use case situation in their implementations.

Action Required of S100TSM
The S100TSM is invited to:
Note the paper and discuss proposed changes.