

BIEN 3 - Task #561

make VegBIEN ID fields plain-text instead of numeric

02/14/2013 04:47 AM - Aaron Marcuse-Kubitza

Status:	New	Start date:	02/14/2013
Priority:	Normal	Due date:	
Assignee:	Aaron Marcuse-Kubitza	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:			
Activity type:			
Description			
Rationale			
<ul style="list-style-type: none">this makes it possible to append data from multiple sources without having pkey collisions			
Forming the ID value			
<ul style="list-style-type: none">for datasource-specific data, use the datasource name + <i>intrinsic</i> identifiers such as a barcode or author plot code<ul style="list-style-type: none">avoid using the datasource's pkey if it's autogenerated, because this is random and not related to the data itselfuse a standard path syntax:<ul style="list-style-type: none">LDAP dnURNXPathfor hierarchical names, create a path from the names<ul style="list-style-type: none">include only as many parent nodes as are needed to guarantee global uniqueness<ul style="list-style-type: none">e.g. for taxonomic names, the family is globally unique, so higher taxa do not need to be includednote that the ID field does not need to be the <i>only</i> way to uniquely identify a record, it just needs to be a way that is a <i>single text value</i>the value should be a human-readable text string rather than the cryptographic hash of such a string, to make it visually obvious what the ID refers to			
Performance considerations			
<ul style="list-style-type: none">to avoid affecting row lookup performance when using arbitrary-length strings as pkeys, a hash index should be added on the ID field<ul style="list-style-type: none">this is in addition to the usual btree index, which is still needed for merge joins and retrieving the data in sorted ordera hashed string is roughly equivalent to using just the hash value because the hash is usually cacheda hash value is roughly equivalent to an autogenerated integer, because it is a similar number of bytes long			

History

#1 - 02/14/2013 05:45 AM - Aaron Marcuse-Kubitza

- Description updated

#2 - 02/14/2013 05:58 AM - Aaron Marcuse-Kubitza

- Description updated

#3 - 02/14/2013 06:21 AM - Aaron Marcuse-Kubitza

- Description updated

#4 - 02/14/2013 04:40 PM - Aaron Marcuse-Kubitza

- Description updated

#5 - 12/03/2013 05:54 AM - Aaron Marcuse-Kubitza

- Subject changed from make VegBIEN ID fields globally unique instead of autogenerated to make VegBIEN ID fields plain-text instead of numeric