

Infrastructure - Story #8227

ExceptionHandler regurgitates long html pages into the wrapping ServiceFailure message, which goes into CN log files.

2017-12-13 21:19 - Rob Nahf

Status:	In Progress	Start date:	2017-12-13
Priority:	Normal	Due date:	
Assignee:	Rob Nahf	% Done:	30%
Category:	d1_common_java	Estimated time:	0.00 hour
Target version:	CCI-2.4.0		
Story Points:			
Description			
<p>While useful to know what was returned in the error response when it was not the correct response, HTML pages can be verbose and include excessive markup that's not useful. Especially when a GMN MN is in debugging mode and there is a systematic error being returned (like during an authentication issue), these logged html pages can end up being 75% of the log files, and cause meaningful log lines from scrolling off the end of the log rotation.</p> <p>An option should be provided to limit the amount of characters being returned in the ServiceFailure.</p> <p>Options are to:</p> <ol style="list-style-type: none">1. eliminate the message body altogether2. truncate the message body3. only print the visible parts of the HTML (remove and elements)4. combination of 2 & 3 <p>since a new feature, develop in trunk.</p>			

Associated revisions

Revision 19047 - 2017-12-13 21:38 - Rob Nahf

refs #8227: added an HTML condenser for situations where the returned response is an exception written in HTML. Defined the configuration property: "dataone.exceptionHandler.limitOriginalReponseInMessage", with a default to true.

Revision 19047 - 2017-12-13 21:38 - Rob Nahf

refs #8227: added an HTML condenser for situations where the returned response is an exception written in HTML. Defined the configuration property: "dataone.exceptionHandler.limitOriginalReponseInMessage", with a default to true.

History

#1 - 2017-12-13 21:26 - Rob Nahf

a Pattern-based routine has been established to limit the HTML to just printable characters and also removes blank lines. On a typical GMN debugging page for a NotFound, it reduced the length of the message body from 8200 characters to 2500, and only adds 7ms to the process.

Interestingly, adding a conditional to check a Settings property from the parsing routine (to see if it should be limited) costs 150ms, so it would be much more costly to allow on-the-fly changes to the way in which things are logged.

The choice of verbosity style is put in a static block, so a restart would be needed to change styles.

#2 - 2017-12-26 17:41 - Dave Vieglaiss

- Sprint set to Infrastructure backlog