

Gen-Z Pattern Operations

October 2017

This presentation covers Gen-Z Pattern operations. Pattern operations can be used to simplify content management of large buffers, or to accelerate big data analytics operations.

Disclaimer

This document is provided 'as is' with no warranties whatsoever, including any warranty of merchantability, noninfringement, fitness for any particular purpose, or any warranty otherwise arising out of any proposal, specification, or sample. Gen-Z Consortium disclaims all liability for infringement of proprietary rights, relating to use of information in this document. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Gen-Z is a trademark or registered trademark of the Gen-Z Consortium.

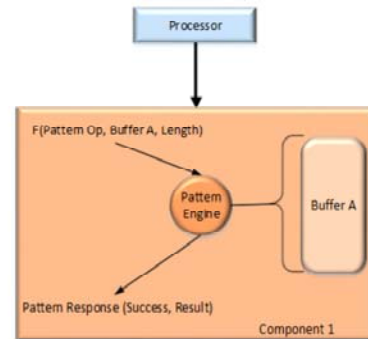
All other product names are trademarks, registered trademarks, or servicemarks of their respective owners.

All material is subject to change at any time at the discretion of the Gen-Z Consortium

<http://genzconsortium.org/>

Pattern Operations

- Pattern requests enable a pattern to be applied to or against a resource range within a Responder
 - Pattern operations provide a standardized set of data-centric computation operations applicable to a variety of workloads
- The Pattern Engine is a processing element
 - General-purpose or specialized processor (discrete or integrated)
 - FPGA
 - Custom or augmented ASIC
 - Etc.
- The Pattern Engine is co-located near or integrated within the memory or storage media controller
 - Provides direct access to the media
 - For example, a small embedded processor within the media controller. Upon receipt of pattern or other data-centric requests, the Pattern Engine directly accesses the media to execute the request
- Pattern Engine may support a regular expression table to enable applications to apply any configured rule
- Pattern Engine may support multiple execution threads to enable multiple comparisons / rules to be simultaneously applied across all media



© Copyright 2016 by Gen-Z. All rights reserved.

GENZ

Pattern operations are executed by a pattern processing engine. The processing engine is co-located with memory or storage media, for example, it can be integrated into the memory or storage media controller, or it can be a discrete element that works through the media controller to access the underlying media. Multiple implementation options are possible.

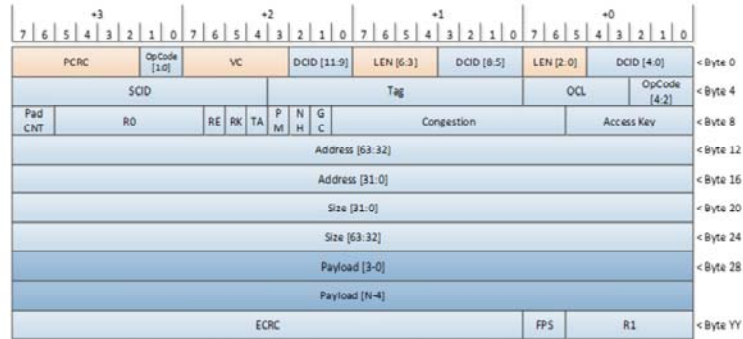
Pattern Request Types

- Pattern Set
 - Request iteratively copies the Pattern at Pattern size strides starting at Address
 - For example, can be used to set the addressed buffer to zero or to a fixed string
- Pattern Count
 - Request iteratively compares the Pattern with memory starting at Address and returns the total count upon completion
- Pattern Match
 - Request iteratively compares the Pattern or applies the supplied regular expression with the memory starting at Address and returns zero or more 64-bit addresses corresponding to successful comparisons.

The following pattern requests are specified:

- Pattern Set operation repeatedly copies the pattern within the request packet's payload to targeted address range. For example, a Pattern Set operation can set the address range to 0x0 or to a fixed pattern, e.g., 0x12345678.
 - Note: If a component supports the Component Media structure, then the primary or secondary media can be set to zero without using a Pattern Set operation.
- Pattern Count returns the number of times the pattern appears within the targeted address range. For example, count the number of times a keyword appears.
- Pattern Match returns one or more addresses that correspond to data locations that matched the supplied pattern or supplied regular expression. The regular expression needs to be comprehended by the Responder.

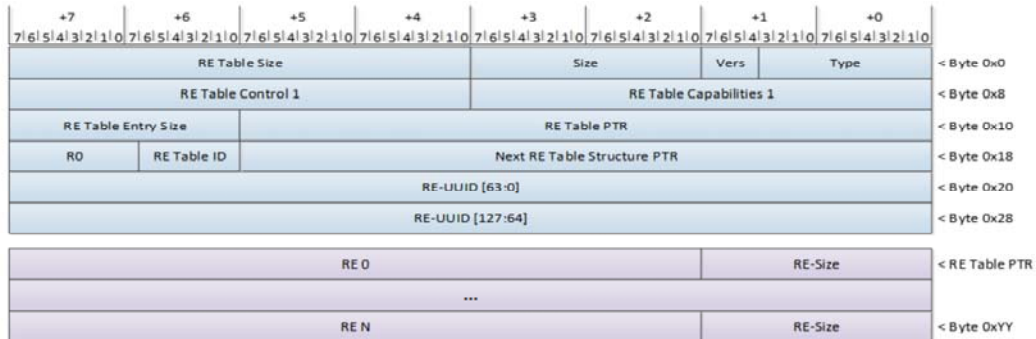
Pattern Request Packet Format



- Address—starting location of the data impacted by the pattern operation
- Size—length of the data buffer
- Payload—pattern to apply or identifies regular expression (if RE == 1b)

Pattern operations are included in the Advanced 1 OpClass. For the most part, the request packets contains the standard Explicit OpClass protocol fields.

Regular Expressions



- Regular expression rules stored in the Component RE Table
 - Each table instance may hold up to 2^{32} regular expressions, aka rules
- Each Component RE Table may support a unique Regular Expression syntax
 - POSIX Based Regular Syntax (BRE), Extended Regular Syntax (ERE), Perl 6 Syntax
- Pattern request identifies which table and which rule within the table to apply to the addressed buffer

Thank you

This concludes this presentation. Thank you.