



Intel[®] Intelligent Storage Acceleration Library (Intel[®] ISA-L) Open Source Version

API Reference Manual - Version 2.10

April 8, 2014

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S LICENSE AGREEMENT FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2011 - 2014 Intel Corporation. All rights reserved.

Contents

1	Storage Library	1
1.1	About This Document	1
1.2	Overview	1
1.3	Erasur e Code Functions	1
1.4	System Requirements	2
2	Function Version Numbers	3
2.1	Function Version Numbers	3
2.2	Function Version Numbers Tables	4
3	Instruction Set Requirements	5
4	File Index	7
4.1	File List	7
5	File Documentation	8
5.1	erasure_code.h File Reference	8
5.1.1	Detailed Description	10
5.1.2	Function Documentation	10
5.1.2.1	ec_encode_data	10
5.1.2.2	ec_encode_data_base	11
5.1.2.3	ec_encode_data_sse	11
5.1.2.4	ec_init_tables	12
5.1.2.5	gf_2vect_dot_prod_avx	12
5.1.2.6	gf_2vect_dot_prod_avx2	13
5.1.2.7	gf_2vect_dot_prod_sse	13
5.1.2.8	gf_3vect_dot_prod_avx	14
5.1.2.9	gf_3vect_dot_prod_avx2	14
5.1.2.10	gf_3vect_dot_prod_sse	15
5.1.2.11	gf_4vect_dot_prod_avx	16
5.1.2.12	gf_4vect_dot_prod_avx2	16
5.1.2.13	gf_4vect_dot_prod_sse	17
5.1.2.14	gf_5vect_dot_prod_avx	17
5.1.2.15	gf_5vect_dot_prod_avx2	18
5.1.2.16	gf_5vect_dot_prod_sse	19
5.1.2.17	gf_6vect_dot_prod_avx	19

5.1.2.18	gf_6vect_dot_prod_avx2	20
5.1.2.19	gf_6vect_dot_prod_sse	20
5.1.2.20	gf_gen_cauchy1_matrix	21
5.1.2.21	gf_gen_rs_matrix	21
5.1.2.22	gf_inv	22
5.1.2.23	gf_invert_matrix	22
5.1.2.24	gf_mul	22
5.1.2.25	gf_vect_dot_prod	23
5.1.2.26	gf_vect_dot_prod_avx	23
5.1.2.27	gf_vect_dot_prod_avx2	24
5.1.2.28	gf_vect_dot_prod_base	24
5.1.2.29	gf_vect_dot_prod_sse	25
5.2	gf_vect_mul.h File Reference	25
5.2.1	Detailed Description	26
5.2.2	Function Documentation	26
5.2.2.1	gf_vect_mul	26
5.2.2.2	gf_vect_mul_avx	27
5.2.2.3	gf_vect_mul_base	27
5.2.2.4	gf_vect_mul_init	27
5.2.2.5	gf_vect_mul_sse	28
5.3	types.h File Reference	28
5.3.1	Detailed Description	28

1.1 About This Document

This document describes the software programming interface and operation of functions in the library. Sections in this document are grouped by the functions found in individual header files that define the function prototypes. Subsections include function parameters, description and type.

This document refers to the open release version of the library. A separate, general release called the Intel® Intelligent Storage Acceleration Library (Intel® ISA-L) is also available and contains an extended set of functions.

1.2 Overview

The Intel® Intelligent Storage Acceleration Library (Intel® ISA-L) Open Source Version is a collection of functions used in storage applications optimized for Intel architecture Intel® 64. In some cases, multiple versions of the same function are available that are optimized for a particular Intel architecture and instruction set. This software takes advantage of new instructions and users should ensure that the chosen function is compatible with hardware it will run on.

1.3 Erasure Code Functions

Functions pertaining to erasure codes implement a general Reed-Solomon type encoding for blocks of data to protect against erasure of whole blocks. Individual operations can be described in terms of arithmetic in the Galois finite field $GF(2^8)$ with the particular field-defining primitive or reducing polynomial $x^8 + x^4 + x^3 + x^2 + 1$ (0x1d).

For example, the function `ec_encode_data()` will generate a set of parity blocks P_i from the set of k source blocks D_i and arbitrary encoding coefficients $a_{i,j}$ where each byte in P is calculated from sources as:

$$P_i = \sum_{j=1}^k a_{i,j} \cdot D_j$$

where addition and multiplication \cdot is defined in $GF(2^8)$. Since any arbitrary set of coefficients $a_{i,j}$ can be supplied, the same fundamental function can be used for encoding blocks or decoding from blocks in erasure.

As noted in the above document, the scheduler routines do not enforce atomic access to the context structure. If a single scheduler state structure is being used by multiple threads, then the application must take care that calls are not made from different threads at the same time, i.e. thread-safety should be implemented at a level higher than these routines. This could be implemented by employing a separate context structure for each worker thread.

1.4 System Requirements

Individual functions may have various run-time requirements such as the minimum version of SSE as described in [Instruction Set Requirements](#). General requirements are listed below.

Recommended Hardware:

- em64t: A system based on the Intel® Xeon® processor with Intel® 64 architecture.
- IA32: When available for 32-bit functions; A system based on the Intel® Xeon® processor or subsequent IA-32 architecture based processor.

Software Requirements:

Most functions in the library use the 64-bit embedded and Unix standard for calling convention http://refspecs.linuxfoundation.org/elf/x86_64-abi-0.95.pdf. When available, 32-bit versions use cdecl. Individual functions are written to be statically linked with an application.

Building Library Functions:

- Yasm Assembler: version at least v1.2.0.

Building Examples and Tests:

Examples and test source follow simple command line POSIX standards and should be portable to any mostly POSIX-compliant OS.

Note

Please note that the library assumes 1MB = 1,000,000 bytes in reported performance figures.

CHAPTER 2

FUNCTION VERSION NUMBERS

2.1 Function Version Numbers

Individual functions are given version numbers with the format mm-vv-ssss.

- mm = Two hex digits indicating the processor a function was optimized for.

- 00 = Nehalem/Jasper Forest/Multibinary
- 01 = Westmere
- 02 = Sandybridge
- 03 = Ivy Bridge
- 04 = Haswell
- 05 = Silvermont

- vv = function version number

- ssss = function serial number

2.2 Function Version Numbers Tables

Function	Version
gf_vect_mul_sse	00-02-0034
gf_vect_mul_init	00-02-0035
gf_vect_mul_avx	01-02-0036
gf_vect_dot_prod_sse	00-03-0060
gf_vect_dot_prod_avx	02-03-0061
gf_2vect_dot_prod_sse	00-02-0062
gf_3vect_dot_prod_sse	00-03-0063
gf_4vect_dot_prod_avx2	04-03-0064
gf_4vect_dot_prod_avx	00-02-0064
gf_4vect_dot_prod_sse	00-03-0064
gf_5vect_dot_prod_sse	00-03-0065
gf_6vect_dot_prod_sse	00-03-0066
ec_init_tables	00-01-0068
ec_encode_data_sse	00-02-0069
ec_encode_data	00-02-0133
gf_vect_mul	00-02-0134
ec_encode_data_base	00-01-0135
gf_vect_mul_base	00-01-0136
gf_vect_dot_prod_base	00-01-0137
gf_vect_dot_prod	00-01-0138
gf_vect_dot_prod_avx2	04-03-0190
gf_2vect_dot_prod_avx	02-03-0191
gf_3vect_dot_prod_avx	02-03-0192
gf_5vect_dot_prod_avx	02-03-0194
gf_6vect_dot_prod_avx	02-03-0195
gf_2vect_dot_prod_avx2	04-03-0196
gf_3vect_dot_prod_avx2	04-03-0197
gf_5vect_dot_prod_avx2	04-03-0199
gf_6vect_dot_prod_avx2	04-03-019a

CHAPTER 3

INSTRUCTION SET REQUIREMENTS

`ec_encode_data_sse` (int len, int k, int rows, unsigned char *gftbls, unsigned char **data, unsigned char **coding)
SSE4.1

`gf_2vect_dot_prod_avx` (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char **dest)
AVX

`gf_2vect_dot_prod_avx2` (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char **dest)
AVX2

`gf_2vect_dot_prod_sse` (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char **dest)
SSE4.1

`gf_3vect_dot_prod_avx` (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char **dest)
AVX

`gf_3vect_dot_prod_avx2` (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char **dest)
AVX2

`gf_3vect_dot_prod_sse` (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char **dest)
SSE4.1

`gf_4vect_dot_prod_avx` (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char **dest)
AVX

`gf_4vect_dot_prod_avx2` (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char **dest)
AVX2

`gf_4vect_dot_prod_sse` (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char **dest)
SSE4.1

`gf_5vect_dot_prod_avx` (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char **dest)
AVX

`gf_5vect_dot_prod_avx2` (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char **dest)
AVX2

`gf_5vect_dot_prod_sse` (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char **dest)
SSE4.1

`gf_6vect_dot_prod_avx` (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char **dest)
AVX

`gf_6vect_dot_prod_avx2` (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char **dest)
AVX2

`gf_6vect_dot_prod_sse` (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char **dest)
SSE4.1

[gf_vect_dot_prod_avx](#) (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char *dest)
AVX

[gf_vect_dot_prod_avx2](#) (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char *dest)
AVX2

[gf_vect_dot_prod_sse](#) (int len, int vlen, unsigned char *gftbls, unsigned char **src, unsigned char *dest)
SSE4.1

[gf_vect_mul_avx](#) (int len, unsigned char *gftbl, void *src, void *dest)
AVX

[gf_vect_mul_sse](#) (int len, unsigned char *gftbl, void *src, void *dest)
SSE4.1

4.1 File List

Here is a list of all documented files with brief descriptions:

erasure_code.h	Interface to functions supporting erasure code encode and decode	8
gf_vect_mul.h	Interface to functions for vector (block) multiplication in GF(2 ⁸)	25
types.h	Defines standard width types	28

5.1 erasure_code.h File Reference

Interface to functions supporting erasure code encode and decode.

```
#include "gf_vect_mul.h"
```

Functions

- void [ec_init_tables](#) (int k, int rows, unsigned char *a, unsigned char *gftbbs)
Initialize tables for fast Erasure Code encode and decode.
- void [ec_encode_data_sse](#) (int len, int k, int rows, unsigned char *gftbbs, unsigned char **data, unsigned char **coding)
Generate or decode erasure codes on blocks of data.
- void [ec_encode_data](#) (int len, int k, int rows, unsigned char *gftbbs, unsigned char **data, unsigned char **coding)
Generate or decode erasure codes on blocks of data, runs appropriate version.
- void [ec_encode_data_base](#) (int len, int srcs, int dests, unsigned char *v, unsigned char **src, unsigned char **dest)
Generate or decode erasure codes on blocks of data, runs baseline version.
- void [gf_vect_dot_prod_sse](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char *dest)
GF(2⁸) vector dot product.
- void [gf_vect_dot_prod_avx](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char *dest)
GF(2⁸) vector dot product.
- void [gf_vect_dot_prod_avx2](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char *dest)
GF(2⁸) vector dot product.
- void [gf_2vect_dot_prod_sse](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char **dest)
GF(2⁸) vector dot product with two outputs.
- void [gf_2vect_dot_prod_avx](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char **dest)
GF(2⁸) vector dot product with two outputs.
- void [gf_2vect_dot_prod_avx2](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char **dest)
GF(2⁸) vector dot product with two outputs.
- void [gf_3vect_dot_prod_sse](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char **dest)

GF(2⁸) vector dot product with three outputs.

- void [gf_3vect_dot_prod_avx](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char **dest)

GF(2⁸) vector dot product with three outputs.

- void [gf_3vect_dot_prod_avx2](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char **dest)

GF(2⁸) vector dot product with three outputs.

- void [gf_4vect_dot_prod_sse](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char **dest)

GF(2⁸) vector dot product with four outputs.

- void [gf_4vect_dot_prod_avx](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char **dest)

GF(2⁸) vector dot product with four outputs.

- void [gf_4vect_dot_prod_avx2](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char **dest)

GF(2⁸) vector dot product with four outputs.

- void [gf_5vect_dot_prod_sse](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char **dest)

GF(2⁸) vector dot product with five outputs.

- void [gf_5vect_dot_prod_avx](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char **dest)

GF(2⁸) vector dot product with five outputs.

- void [gf_5vect_dot_prod_avx2](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char **dest)

GF(2⁸) vector dot product with five outputs.

- void [gf_6vect_dot_prod_sse](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char **dest)

GF(2⁸) vector dot product with six outputs.

- void [gf_6vect_dot_prod_avx](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char **dest)

GF(2⁸) vector dot product with six outputs.

- void [gf_6vect_dot_prod_avx2](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char **dest)

GF(2⁸) vector dot product with six outputs.

- void [gf_vect_dot_prod_base](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char *dest)

GF(2⁸) vector dot product, runs baseline version.

- void [gf_vect_dot_prod](#) (int len, int vlen, unsigned char *gftbbs, unsigned char **src, unsigned char *dest)

GF(2⁸) vector dot product, runs appropriate version.

- unsigned char [gf_mul](#) (unsigned char a, unsigned char b)

Single element GF(2⁸) multiply.

- unsigned char [gf_inv](#) (unsigned char a)

Single element GF(2⁸) inverse.

- void `gf_gen_rs_matrix` (unsigned char *a, int m, int k)
Generate a matrix of coefficients to be used for encoding.
- void `gf_gen_cauchy_matrix` (unsigned char *a, int m, int k)
Generate a Cauchy matrix of coefficients to be used for encoding.
- int `gf_invert_matrix` (unsigned char *in, unsigned char *out, const int n)
Invert a matrix in GF(2⁸)

5.1.1 Detailed Description

Interface to functions supporting erasure code encode and decode. This file defines the interface to optimized functions used in erasure codes. Encode and decode of erasures in GF(2⁸) are made by calculating the dot product of the symbols (bytes in GF(2⁸)) across a set of buffers and a set of coefficients. Values for the coefficients are determined by the type of erasure code. Using a general dot product means that any sequence of coefficients may be used including erasure codes based on random coefficients. Multiple versions of dot product are supplied to calculate 1-6 output vectors in one pass. Base GF multiply and divide functions can be sped up by defining GF_LARGE_TABLES at the expense of memory size.

5.1.2 Function Documentation

5.1.2.1 void `ec_encode_data` (int *len*, int *k*, int *rows*, unsigned char * *gftbls*, unsigned char ** *data*, unsigned char ** *coding*)

Generate or decode erasure codes on blocks of data, runs appropriate version.

Given a list of source data blocks, generate one or multiple blocks of encoded data as specified by a matrix of GF(2⁸) coefficients. When given a suitable set of coefficients, this function will perform the fast generation or decoding of Reed-Solomon type erasure codes.

This function determines what instruction sets are enabled and selects the appropriate version at runtime.

Parameters

<i>len</i>	Length of each block of data (vector) of source or dest data.
<i>k</i>	The number of vector sources or rows in the generator matrix for coding.
<i>rows</i>	The number of output vectors to concurrently encode/decode.
<i>gftbls</i>	Pointer to array of input tables generated from coding coefficients in <code>ec_init_tables()</code> . Must be of size 32*k*rows
<i>data</i>	Array of pointers to source input buffers.
<i>coding</i>	Array of pointers to coded output buffers.

Returns

none

5.1.2.2 void `ec_encode_data_base` (int *len*, int *srcs*, int *dsts*, unsigned char * *v*, unsigned char ** *src*, unsigned char ** *dest*)

Generate or decode erasure codes on blocks of data, runs baseline version.

Given a list of source data blocks, generate one or multiple blocks of encoded data as specified by a matrix of GF(2⁸) coefficients. When given a suitable set of coefficients, this function will perform the fast generation or decoding of Reed-Solomon type erasure codes.

Parameters

<i>len</i>	Length of each block of data (vector) of source or dest data.
<i>srcs</i>	The number of vector sources or rows in the generator matrix for coding.
<i>dsts</i>	The number of output vectors to concurrently encode/decode.
<i>v</i>	Pointer to array of input tables generated from coding coefficients in <code>ec_init_tables()</code> . Must be of size 32*k*rows
<i>src</i>	Array of pointers to source input buffers.
<i>dest</i>	Array of pointers to coded output buffers.

Returns

none

5.1.2.3 void `ec_encode_data_sse` (int *len*, int *k*, int *rows*, unsigned char * *gftbls*, unsigned char ** *data*, unsigned char ** *coding*)

Generate or decode erasure codes on blocks of data.

Given a list of source data blocks, generate one or multiple blocks of encoded data as specified by a matrix of GF(2⁸) coefficients. When given a suitable set of coefficients, this function will perform the fast generation or decoding of Reed-Solomon type erasure codes.

Requires SSE4.1

Parameters

<i>len</i>	Length of each block of data (vector) of source or dest data.
<i>k</i>	The number of vector sources or rows in the generator matrix for coding.
<i>rows</i>	The number of output vectors to concurrently encode/decode.
<i>gftbls</i>	Pointer to array of input tables generated from coding coefficients in <code>ec_init_tables()</code> . Must be of size 32*k*rows
<i>data</i>	Array of pointers to source input buffers.
<i>coding</i>	Array of pointers to coded output buffers.

Returns

none

5.1.2.4 void ec_init_tables (int *k*, int *rows*, unsigned char * *a*, unsigned char * *gftbls*)

Initialize tables for fast Erasure Code encode and decode.

Generates the expanded tables needed for fast encode or decode for erasure codes on blocks of data. 32bytes is generated for each input coefficient.

Parameters

<i>k</i>	The number of vector sources or rows in the generator matrix for coding.
<i>rows</i>	The number of output vectors to concurrently encode/decode.
<i>a</i>	Pointer to sets of arrays of input coefficients used to encode or decode data.
<i>gftbls</i>	Pointer to start of space for concatenated output tables generated from input coefficients. Must be of size 32*k*rows.

Returns

none

5.1.2.5 void gf_2vect_dot_prod_avx (int *len*, int *vlen*, unsigned char * *gftbls*, unsigned char ** *src*, unsigned char ** *dest*)

GF(2⁸) vector dot product with two outputs.

Vector dot product optimized to calculate two outputs at a time. Does two GF(2⁸) dot products across each byte of the input array and two constant sets of coefficients to produce each byte of the outputs. Can be used for erasure coding encode and decode. Function requires pre-calculation of a 2*32*vlen byte constant array based on the two sets of input coefficients.

Requires AVX

Parameters

<i>len</i>	Length of each vector in bytes. Must be >= 16.
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to 2*32*vlen byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Array of pointers to destination data buffers.

Returns

none

5.1.2.6 `void gf_2vect_dot_prod_avx2 (int len, int vlen, unsigned char * gftbls, unsigned char ** src, unsigned char ** dest)`

GF(2⁸) vector dot product with two outputs.

Vector dot product optimized to calculate two outputs at a time. Does two GF(2⁸) dot products across each byte of the input array and two constant sets of coefficients to produce each byte of the outputs. Can be used for erasure coding encode and decode. Function requires pre-calculation of a 2*32*vlen byte constant array based on the two sets of input coefficients.

Requires AVX2

Parameters

<i>len</i>	Length of each vector in bytes. Must be ≥ 32 .
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to 2*32*vlen byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Array of pointers to destination data buffers.

Returns

none

5.1.2.7 `void gf_2vect_dot_prod_sse (int len, int vlen, unsigned char * gftbls, unsigned char ** src, unsigned char ** dest)`

GF(2⁸) vector dot product with two outputs.

Vector dot product optimized to calculate two outputs at a time. Does two GF(2⁸) dot products across each byte of the input array and two constant sets of coefficients to produce each byte of the outputs. Can be used for erasure coding encode and decode. Function requires pre-calculation of a 2*32*vlen byte constant array based on the two sets of input coefficients.

Requires SSE4.1

Parameters

<i>len</i>	Length of each vector in bytes. Must be ≥ 16 .
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to $2*32*vlen$ byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Array of pointers to destination data buffers.

Returns

none

5.1.2.8 `void gf_3vect_dot_prod_avx (int len, int vlen, unsigned char * gftbls, unsigned char ** src, unsigned char ** dest)`

GF(2^8) vector dot product with three outputs.

Vector dot product optimized to calculate three outputs at a time. Does three GF(2^8) dot products across each byte of the input array and three constant sets of coefficients to produce each byte of the outputs. Can be used for erasure coding encode and decode. Function requires pre-calculation of a $3*32*vlen$ byte constant array based on the three sets of input coefficients.

Requires AVX

Parameters

<i>len</i>	Length of each vector in bytes. Must be ≥ 16 .
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to $3*32*vlen$ byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Array of pointers to destination data buffers.

Returns

none

5.1.2.9 `void gf_3vect_dot_prod_avx2 (int len, int vlen, unsigned char * gftbls, unsigned char ** src, unsigned char ** dest)`

GF(2^8) vector dot product with three outputs.

Vector dot product optimized to calculate three outputs at a time. Does three GF(2⁸) dot products across each byte of the input array and three constant sets of coefficients to produce each byte of the outputs. Can be used for erasure coding encode and decode. Function requires pre-calculation of a 3*32*vlen byte constant array based on the three sets of input coefficients.

Requires AVX2

Parameters

<i>len</i>	Length of each vector in bytes. Must be ≥ 32 .
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to 3*32*vlen byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Array of pointers to destination data buffers.

Returns

none

5.1.2.10 `void gf_3vect_dot_prod_sse (int len, int vlen, unsigned char * gftbls, unsigned char ** src, unsigned char ** dest)`

GF(2⁸) vector dot product with three outputs.

Vector dot product optimized to calculate three outputs at a time. Does three GF(2⁸) dot products across each byte of the input array and three constant sets of coefficients to produce each byte of the outputs. Can be used for erasure coding encode and decode. Function requires pre-calculation of a 3*32*vlen byte constant array based on the three sets of input coefficients.

Requires SSE4.1

Parameters

<i>len</i>	Length of each vector in bytes. Must be ≥ 16 .
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to 3*32*vlen byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Array of pointers to destination data buffers.

Returns

none

5.1.2.11 `void gf_4vect_dot_prod_avx (int len, int vlen, unsigned char * gftbls, unsigned char ** src, unsigned char ** dest)`

GF(2⁸) vector dot product with four outputs.

Vector dot product optimized to calculate four outputs at a time. Does four GF(2⁸) dot products across each byte of the input array and four constant sets of coefficients to produce each byte of the outputs. Can be used for erasure coding encode and decode. Function requires pre-calculation of a 4*32*vlen byte constant array based on the four sets of input coefficients.

Requires AVX

Parameters

<i>len</i>	Length of each vector in bytes. Must be ≥ 16 .
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to 4*32*vlen byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Array of pointers to destination data buffers.

Returns

none

5.1.2.12 `void gf_4vect_dot_prod_avx2 (int len, int vlen, unsigned char * gftbls, unsigned char ** src, unsigned char ** dest)`

GF(2⁸) vector dot product with four outputs.

Vector dot product optimized to calculate four outputs at a time. Does four GF(2⁸) dot products across each byte of the input array and four constant sets of coefficients to produce each byte of the outputs. Can be used for erasure coding encode and decode. Function requires pre-calculation of a 4*32*vlen byte constant array based on the four sets of input coefficients.

Requires AVX2

Parameters

<i>len</i>	Length of each vector in bytes. Must be ≥ 32 .
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to $4*32*vlen$ byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Array of pointers to destination data buffers.

Returns

none

5.1.2.13 `void gf_4vect_dot_prod_sse (int len, int vlen, unsigned char * gftbls, unsigned char ** src, unsigned char ** dest)`

GF(2^8) vector dot product with four outputs.

Vector dot product optimized to calculate four outputs at a time. Does four GF(2^8) dot products across each byte of the input array and four constant sets of coefficients to produce each byte of the outputs. Can be used for erasure coding encode and decode. Function requires pre-calculation of a $4*32*vlen$ byte constant array based on the four sets of input coefficients.

Requires SSE4.1

Parameters

<i>len</i>	Length of each vector in bytes. Must be ≥ 16 .
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to $4*32*vlen$ byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Array of pointers to destination data buffers.

Returns

none

5.1.2.14 `void gf_5vect_dot_prod_avx (int len, int vlen, unsigned char * gftbls, unsigned char ** src, unsigned char ** dest)`

GF(2^8) vector dot product with five outputs.

Vector dot product optimized to calculate five outputs at a time. Does five GF(2⁸) dot products across each byte of the input array and five constant sets of coefficients to produce each byte of the outputs. Can be used for erasure coding encode and decode. Function requires pre-calculation of a 5*32*vlen byte constant array based on the five sets of input coefficients.

Requires AVX

Parameters

<i>len</i>	Length of each vector in bytes. Must ≥ 16 .
<i>vlen</i>	Number of vector sources.
<i>gftbbs</i>	Pointer to 5*32*vlen byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Array of pointers to destination data buffers.

Returns

none

5.1.2.15 `void gf_5vect_dot_prod_avx2 (int len, int vlen, unsigned char * gftbbs, unsigned char ** src, unsigned char ** dest)`

GF(2⁸) vector dot product with five outputs.

Vector dot product optimized to calculate five outputs at a time. Does five GF(2⁸) dot products across each byte of the input array and five constant sets of coefficients to produce each byte of the outputs. Can be used for erasure coding encode and decode. Function requires pre-calculation of a 5*32*vlen byte constant array based on the five sets of input coefficients.

Requires AVX2

Parameters

<i>len</i>	Length of each vector in bytes. Must ≥ 32 .
<i>vlen</i>	Number of vector sources.
<i>gftbbs</i>	Pointer to 5*32*vlen byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Array of pointers to destination data buffers.

Returns

none

5.1.2.16 `void gf_5vect_dot_prod_sse (int len, int vlen, unsigned char * gftbls, unsigned char ** src, unsigned char ** dest)`

GF(2⁸) vector dot product with five outputs.

Vector dot product optimized to calculate five outputs at a time. Does five GF(2⁸) dot products across each byte of the input array and five constant sets of coefficients to produce each byte of the outputs. Can be used for erasure coding encode and decode. Function requires pre-calculation of a 5*32*vlen byte constant array based on the five sets of input coefficients.

Requires SSE4.1

Parameters

<i>len</i>	Length of each vector in bytes. Must ≥ 16 .
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to 5*32*vlen byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Array of pointers to destination data buffers.

Returns

none

5.1.2.17 `void gf_6vect_dot_prod_avx (int len, int vlen, unsigned char * gftbls, unsigned char ** src, unsigned char ** dest)`

GF(2⁸) vector dot product with six outputs.

Vector dot product optimized to calculate six outputs at a time. Does six GF(2⁸) dot products across each byte of the input array and six constant sets of coefficients to produce each byte of the outputs. Can be used for erasure coding encode and decode. Function requires pre-calculation of a 6*32*vlen byte constant array based on the six sets of input coefficients.

Requires AVX

Parameters

<i>len</i>	Length of each vector in bytes. Must be ≥ 16 .
<i>vlen</i>	Number of vector sources.
<i>gftbbs</i>	Pointer to $6 \cdot 32 \cdot vlen$ byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Array of pointers to destination data buffers.

Returns

none

5.1.2.18 `void gf_6vect_dot_prod_avx2 (int len, int vlen, unsigned char * gftbbs, unsigned char ** src, unsigned char ** dest)`

GF(2^8) vector dot product with six outputs.

Vector dot product optimized to calculate six outputs at a time. Does six GF(2^8) dot products across each byte of the input array and six constant sets of coefficients to produce each byte of the outputs. Can be used for erasure coding encode and decode. Function requires pre-calculation of a $6 \cdot 32 \cdot vlen$ byte constant array based on the six sets of input coefficients.

Requires AVX2

Parameters

<i>len</i>	Length of each vector in bytes. Must be ≥ 32 .
<i>vlen</i>	Number of vector sources.
<i>gftbbs</i>	Pointer to $6 \cdot 32 \cdot vlen$ byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Array of pointers to destination data buffers.

Returns

none

5.1.2.19 `void gf_6vect_dot_prod_sse (int len, int vlen, unsigned char * gftbbs, unsigned char ** src, unsigned char ** dest)`

GF(2^8) vector dot product with six outputs.

Vector dot product optimized to calculate six outputs at a time. Does six GF(2⁸) dot products across each byte of the input array and six constant sets of coefficients to produce each byte of the outputs. Can be used for erasure coding encode and decode. Function requires pre-calculation of a 6*32*vlen byte constant array based on the six sets of input coefficients.

Requires SSE4.1

Parameters

<i>len</i>	Length of each vector in bytes. Must be ≥ 16 .
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to 6*32*vlen byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Array of pointers to destination data buffers.

Returns

none

5.1.2.20 void gf_gen_cauchy1_matrix (unsigned char * a, int m, int k)

Generate a Cauchy matrix of coefficients to be used for encoding.

Cauchy matrix example of encoding coefficients where high portion of matrix is identity matrix I and lower portion is constructed as $1/(i+j) \mid i \neq j, i: \{0, k-1\} j: \{k, m-1\}$. Any sub-matrix of a Cauchy matrix should be invertable.

Parameters

<i>a</i>	[mxk] array to hold coefficients
<i>m</i>	number of rows in matrix corresponding to srcs + parity.
<i>k</i>	number of columns in matrix corresponding to srcs.

Returns

none

5.1.2.21 void gf_gen_rs_matrix (unsigned char * a, int m, int k)

Generate a matrix of coefficients to be used for encoding.

Vandermonde matrix example of encoding coefficients where high portion of matrix is identity matrix I and lower portion is constructed as $2^{i*(j-k+1)} \mid i: \{0, k-1\} j: \{k, m-1\}$. Commonly used method for choosing coefficients in

erasure encoding but does not guarantee invertable for every sub matrix. For large k it is possible to find cases where the decode matrix chosen from sources and parity not in erasure are not invertable. Users may want to adjust for $k > 5$.

Parameters

<i>a</i>	[mxk] array to hold coefficients
<i>m</i>	number of rows in matrix corresponding to srcs + parity.
<i>k</i>	number of columns in matrix corresponding to srcs.

Returns

none

5.1.2.22 unsigned char gf_inv (unsigned char *a*)

Single element $GF(2^8)$ inverse.

Parameters

<i>a</i>	Input element
----------	---------------

Returns

Field element b such that $a \times b = \{1\}$

5.1.2.23 int gf_invert_matrix (unsigned char * *in*, unsigned char * *out*, const int *n*)

Invert a matrix in $GF(2^8)$

Parameters

<i>in</i>	input matrix
<i>out</i>	output matrix such that $[in] \times [out] = [I]$ - identity matrix
<i>n</i>	size of matrix [nxn]

Returns

0 successful, other fail on singular input matrix

5.1.2.24 unsigned char gf_mul (unsigned char *a*, unsigned char *b*)

Single element $GF(2^8)$ multiply.

Parameters

<i>a</i>	Multiplicand a
<i>b</i>	Multiplicand b

Returns

Product of a and b in GF(2⁸)

5.1.2.25 void `gf_vect_dot_prod (int len, int vlen, unsigned char * gftbls, unsigned char ** src, unsigned char * dest)`

GF(2⁸) vector dot product, runs appropriate version.

Does a GF(2⁸) dot product across each byte of the input array and a constant set of coefficients to produce each byte of the output. Can be used for erasure coding encode and decode. Function requires pre-calculation of a 32*vlen byte constant array based on the input coefficients.

This function determines what instruction sets are enabled and selects the appropriate version at runtime.

Parameters

<i>len</i>	Length of each vector in bytes. Must be >= 32.
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to 32*vlen byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Pointer to destination data array.

Returns

none

5.1.2.26 void `gf_vect_dot_prod_avx (int len, int vlen, unsigned char * gftbls, unsigned char ** src, unsigned char * dest)`

GF(2⁸) vector dot product.

Does a GF(2⁸) dot product across each byte of the input array and a constant set of coefficients to produce each byte of the output. Can be used for erasure coding encode and decode. Function requires pre-calculation of a 32*vlen byte constant array based on the input coefficients.

Requires AVX

Parameters

<i>len</i>	Length of each vector in bytes. Must be ≥ 16 .
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to $32*vlen$ byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Pointer to destination data array.

Returns

none

5.1.2.27 `void gf_vect_dot_prod_avx2 (int len, int vlen, unsigned char * gftbls, unsigned char ** src, unsigned char * dest)`

GF(2^8) vector dot product.

Does a GF(2^8) dot product across each byte of the input array and a constant set of coefficients to produce each byte of the output. Can be used for erasure coding encode and decode. Function requires pre-calculation of a $32*vlen$ byte constant array based on the input coefficients.

Requires AVX2

Parameters

<i>len</i>	Length of each vector in bytes. Must be ≥ 32 .
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to $32*vlen$ byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Pointer to destination data array.

Returns

none

5.1.2.28 `void gf_vect_dot_prod_base (int len, int vlen, unsigned char * gftbls, unsigned char ** src, unsigned char * dest)`

GF(2^8) vector dot product, runs baseline version.

Does a GF(2^8) dot product across each byte of the input array and a constant set of coefficients to produce each byte of the output. Can be used for erasure coding encode and decode. Function requires pre-calculation of a $32*vlen$ byte constant array based on the input coefficients.

Parameters

<i>len</i>	Length of each vector in bytes. Must be ≥ 16 .
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to $32*vlen$ byte array of pre-calculated constants based on the array of input coefficients. Only elements $32*CONST*j + 1$ of this array are used, where $j = (0, 1, 2\dots)$ and $CONST$ is the number of elements in the array of input coefficients. The elements used correspond to the original input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Pointer to destination data array.

Returns

none

5.1.2.29 void gf_vect_dot_prod_sse (int *len*, int *vlen*, unsigned char * *gftbls*, unsigned char ** *src*, unsigned char * *dest*)

GF(2^8) vector dot product.

Does a GF(2^8) dot product across each byte of the input array and a constant set of coefficients to produce each byte of the output. Can be used for erasure coding encode and decode. Function requires pre-calculation of a $32*vlen$ byte constant array based on the input coefficients.

Requires SSE4.1

Parameters

<i>len</i>	Length of each vector in bytes. Must be ≥ 16 .
<i>vlen</i>	Number of vector sources.
<i>gftbls</i>	Pointer to $32*vlen$ byte array of pre-calculated constants based on the array of input coefficients.
<i>src</i>	Array of pointers to source inputs.
<i>dest</i>	Pointer to destination data array.

Returns

none

5.2 gf_vect_mul.h File Reference

Interface to functions for vector (block) multiplication in GF(2^8).

Functions

- int [gf_vect_mul_sse](#) (int len, unsigned char *gftbl, void *src, void *dest)
GF(2⁸) vector multiply by constant.
- int [gf_vect_mul_avx](#) (int len, unsigned char *gftbl, void *src, void *dest)
GF(2⁸) vector multiply by constant.
- int [gf_vect_mul](#) (int len, unsigned char *gftbl, void *src, void *dest)
GF(2⁸) vector multiply by constant, runs appropriate version.
- void [gf_vect_mul_init](#) (unsigned char c, unsigned char *gftbl)
Initialize 32-byte constant array for GF(2⁸) vector multiply.
- void [gf_vect_mul_base](#) (int len, unsigned char *a, unsigned char *src, unsigned char *dest)
GF(2⁸) vector multiply by constant, runs baseline version.

5.2.1 Detailed Description

Interface to functions for vector (block) multiplication in GF(2⁸). This file defines the interface to routines used in fast RAID rebuild and erasure codes.

5.2.2 Function Documentation

5.2.2.1 int gf_vect_mul (int len, unsigned char * gftbl, void * src, void * dest)

GF(2⁸) vector multiply by constant, runs appropriate version.

Does a GF(2⁸) vector multiply $b = Ca$ where a and b are arrays and C is a single field element in GF(2⁸). Can be used for RAID6 rebuild and partial write functions. Function requires pre-calculation of a 32-element constant array based on constant C . $gftbl(C) = \{C\{00\}, C\{01\}, C\{02\}, \dots, C\{0f\}\}, \{C\{00\}, C\{10\}, C\{20\}, \dots, C\{f0\}\}$. len and src must be aligned to 32B.

This function determines what instruction sets are enabled and selects the appropriate version at runtime.

Parameters

<i>len</i>	Length of vector in bytes. Must be aligned to 32B.
<i>gftbl</i>	Pointer to 32-byte array of pre-calculated constants based on C .
<i>src</i>	Pointer to src data array. Must be aligned to 32B.
<i>dest</i>	Pointer to destination data array. Must be aligned to 32B.

Returns

0 pass, other fail

5.2.2.2 int gf_vect_mul_avx (int *len*, unsigned char * *gftbl*, void * *src*, void * *dest*)

GF(2⁸) vector multiply by constant.

Does a GF(2⁸) vector multiply $b = Ca$ where a and b are arrays and C is a single field element in GF(2⁸). Can be used for RAID6 rebuild and partial write functions. Function requires pre-calculation of a 32-element constant array based on constant C . $gftbl(C) = \{C\{00\}, C\{01\}, C\{02\}, \dots, C\{0f\}\}, \{C\{00\}, C\{10\}, C\{20\}, \dots, C\{f0\}\}$. len and src must be aligned to 32B.

Requires AVX

Parameters

<i>len</i>	Length of vector in bytes. Must be aligned to 32B.
<i>gftbl</i>	Pointer to 32-byte array of pre-calculated constants based on C .
<i>src</i>	Pointer to src data array. Must be aligned to 32B.
<i>dest</i>	Pointer to destination data array. Must be aligned to 32B.

Returns

0 pass, other fail

5.2.2.3 void gf_vect_mul_base (int *len*, unsigned char * *a*, unsigned char * *src*, unsigned char * *dest*)

GF(2⁸) vector multiply by constant, runs baseline version.

Does a GF(2⁸) vector multiply $b = Ca$ where a and b are arrays and C is a single field element in GF(2⁸). Can be used for RAID6 rebuild and partial write functions. Function requires pre-calculation of a 32-element constant array based on constant C . $gftbl(C) = \{C\{00\}, C\{01\}, C\{02\}, \dots, C\{0f\}\}, \{C\{00\}, C\{10\}, C\{20\}, \dots, C\{f0\}\}$. len and src must be aligned to 32B.

Parameters

<i>len</i>	Length of vector in bytes. Must be aligned to 32B.
<i>a</i>	Pointer to 32-byte array of pre-calculated constants based on C . only use 2nd element is used.
<i>src</i>	Pointer to src data array. Must be aligned to 32B.
<i>dest</i>	Pointer to destination data array. Must be aligned to 32B.

5.2.2.4 void gf_vect_mul_init (unsigned char *c*, unsigned char * *gftbl*)

Initialize 32-byte constant array for GF(2⁸) vector multiply.

Calculates array $\{C\{00\}, C\{01\}, C\{02\}, \dots, C\{0f\}\}, \{C\{00\}, C\{10\}, C\{20\}, \dots, C\{f0\}\}$ as required by other fast vector multiply functions.

Parameters

<i>c</i>	Constant input.
<i>gftbl</i>	Table output.

5.2.2.5 int gf_vect_mul_sse (int *len*, unsigned char * *gftbl*, void * *src*, void * *dest*)

GF(2⁸) vector multiply by constant.

Does a GF(2⁸) vector multiply $b = Ca$ where a and b are arrays and C is a single field element in GF(2⁸). Can be used for RAID6 rebuild and partial write functions. Function requires pre-calculation of a 32-element constant array based on constant C . $gftbl(C) = \{C\{00\}, C\{01\}, C\{02\}, \dots, C\{0f\}\}, \{C\{00\}, C\{10\}, C\{20\}, \dots, C\{f0\}\}$. len and src must be aligned to 32B.

Requires SSE4.1

Parameters

<i>len</i>	Length of vector in bytes. Must be aligned to 32B.
<i>gftbl</i>	Pointer to 32-byte array of pre-calculated constants based on C .
<i>src</i>	Pointer to src data array. Must be aligned to 32B.
<i>dest</i>	Pointer to destination data array. Must be aligned to 32B.

Returns

0 pass, other fail

5.3 types.h File Reference

Defines standard width types.

5.3.1 Detailed Description

Defines standard width types.
