



Intel® Advanced Vector Extensions 10.2

Architecture Specification

October, 2024

Revision 2.0

Notices & Disclaimers

This document contains information on products in the design phase of development. The information here is subject to change without notice. Do not finalize a design with this information.

Intel technologies may require enabled hardware, software or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

All product plans and roadmaps are subject to change without notice. The products described 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.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Code names are used by Intel to identify products, technologies, or services that are in development and not publicly available. These are not "commercial" names and not intended to function as trademarks.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document, with the sole exception that a) you may publish an unmodified copy and b) code included in this document is licensed subject to the Zero-Clause BSD open source license (0BSD), <https://opensource.org/licenses/0BSD>. You may create software implementations based on this document and in compliance with the foregoing that are intended to execute on the Intel product(s) referenced in this document. No rights are granted to create modifications or derivatives of this document.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Contents

1	CHANGES	28
2	CPUID	29
3	INTRODUCTION	30
3.1	INTEL® AVX10 INTRODUCTION	31
3.1.1	FEATURES AND CAPABILITIES	31
3.1.2	FEATURE ENUMERATION	31
3.1.3	STATE MANAGEMENT	34
3.1.4	VMX FOR INTEL® AVX10 VIRTUAL MACHINES	35
3.1.5	ENCODING EXTENSIONS	36
3.1.6	INTEL® AVX10.2 NEW INSTRUCTIONS	37
3.2	FP8 INTRODUCTION	39
3.2.1	NUMERIC DEFINITION	39
3.2.2	FP8 ROUNDING, DENORMALS, SPECIAL NUMBERS, AND EXCEPTIONS	39
4	EXCEPTION CLASSES	41
4.1	EXCEPTION CLASS INSTRUCTION SUMMARY	42
4.2	EXCEPTION CLASS SUMMARY	52
4.2.1	EXCEPTION CLASS E10NF	52
4.2.2	EXCEPTION CLASS E11	54
4.2.3	EXCEPTION CLASS E2	56
4.2.4	EXCEPTION CLASS E3	58
4.2.5	EXCEPTION CLASS E3NF	60
4.2.6	EXCEPTION CLASS E4	62
4.2.7	EXCEPTION CLASS E4NF	64
4.2.8	EXCEPTION CLASS E9NF	65
5	HELPER FUNCTIONS	66
6	INSTRUCTION TABLE	101
7	INTEL® AVX10 BF16 INSTRUCTIONS	113
7.1	VADDNEPBF16	114
7.1.1	INSTRUCTION OPERAND ENCODING	114
7.1.2	DESCRIPTION	114
7.1.3	OPERATION	115
7.1.4	EXCEPTIONS	115

7.2	VCMPPBF16	116
7.2.1	INSTRUCTION OPERAND ENCODING	116
7.2.2	DESCRIPTION	116
7.2.3	OPERATION	117
7.2.4	EXCEPTIONS	118
7.3	VCOMSBF16	119
7.3.1	INSTRUCTION OPERAND ENCODING	119
7.3.2	DESCRIPTION	119
7.3.3	OPERATION	120
7.3.4	EXCEPTIONS	120
7.4	VDIVNEPBF16	121
7.4.1	INSTRUCTION OPERAND ENCODING	121
7.4.2	DESCRIPTION	121
7.4.3	OPERATION	122
7.4.4	EXCEPTIONS	122
7.5	VF[,N]M[ADD,SUB][132,213,231]NEPBF16	123
7.5.1	INSTRUCTION OPERAND ENCODING	125
7.5.2	DESCRIPTION	125
7.5.3	OPERATION	126
7.5.4	EXCEPTIONS	127
7.6	VFPCLASSPBF16	131
7.6.1	INSTRUCTION OPERAND ENCODING	131
7.6.2	DESCRIPTION	131
7.6.3	OPERATION	132
7.6.4	EXCEPTIONS	133
7.7	VGETEXPPBF16	134
7.7.1	INSTRUCTION OPERAND ENCODING	134
7.7.2	DESCRIPTION	134
7.7.3	OPERATION	135
7.7.4	EXCEPTIONS	135
7.8	VGETMANTPBF16	137
7.8.1	INSTRUCTION OPERAND ENCODING	137
7.8.2	DESCRIPTION	137
7.8.3	OPERATION	139
7.8.4	EXCEPTIONS	140
7.9	VMAXPBF16	141
7.9.1	INSTRUCTION OPERAND ENCODING	141
7.9.2	DESCRIPTION	141
7.9.3	OPERATION	142
7.9.4	EXCEPTIONS	142
7.10	VMINPBF16	144
7.10.1	INSTRUCTION OPERAND ENCODING	144
7.10.2	DESCRIPTION	144
7.10.3	OPERATION	145
7.10.4	EXCEPTIONS	145
7.11	VMULNEPBF16	147
7.11.1	INSTRUCTION OPERAND ENCODING	147
7.11.2	DESCRIPTION	147

7.11.3	OPERATION	148
7.11.4	EXCEPTIONS	148
7.12	VRCPPBF16	149
7.12.1	INSTRUCTION OPERAND ENCODING	149
7.12.2	DESCRIPTION	149
7.12.3	OPERATION	150
7.12.4	EXCEPTIONS	150
7.13	VREDUCENEPBF16	151
7.13.1	INSTRUCTION OPERAND ENCODING	151
7.13.2	DESCRIPTION	151
7.13.3	OPERATION	152
7.13.4	EXCEPTIONS	152
7.14	VRNDSCALENEPBF16	153
7.14.1	INSTRUCTION OPERAND ENCODING	153
7.14.2	DESCRIPTION	153
7.14.3	OPERATION	154
7.14.4	EXCEPTIONS	154
7.15	VRSQRTPBF16	156
7.15.1	INSTRUCTION OPERAND ENCODING	156
7.15.2	DESCRIPTION	156
7.15.3	OPERATION	157
7.15.4	EXCEPTIONS	157
7.16	VSCALEFPBF16	158
7.16.1	INSTRUCTION OPERAND ENCODING	158
7.16.2	DESCRIPTION	158
7.16.3	OPERATION	159
7.16.4	EXCEPTIONS	159
7.17	VSQRTNEPBF16	160
7.17.1	INSTRUCTION OPERAND ENCODING	160
7.17.2	DESCRIPTION	160
7.17.3	OPERATION	161
7.17.4	EXCEPTIONS	161
7.18	VSUBNEPBF16	162
7.18.1	INSTRUCTION OPERAND ENCODING	162
7.18.2	DESCRIPTION	162
7.18.3	OPERATION	163
7.18.4	EXCEPTIONS	163
8	INTEL® AVX10 COMPARE SCALAR FP WITH ENHANCED EFLAGS INSTRUCTIONS	164
8.1	VCOMXSD	165
8.1.1	INSTRUCTION OPERAND ENCODING	165
8.1.2	DESCRIPTION	165
8.1.3	OPERATION	166
8.1.4	SIMD FLOATING-POINT EXCEPTIONS	166
8.1.5	EXCEPTIONS	166
8.2	VCOMXSH	167
8.2.1	INSTRUCTION OPERAND ENCODING	167
8.2.2	DESCRIPTION	167

8.2.3	OPERATION	168
8.2.4	SIMD FLOATING-POINT EXCEPTIONS	168
8.2.5	EXCEPTIONS	168
8.3	VCOMXSS	169
8.3.1	INSTRUCTION OPERAND ENCODING	169
8.3.2	DESCRIPTION	169
8.3.3	OPERATION	170
8.3.4	SIMD FLOATING-POINT EXCEPTIONS	170
8.3.5	EXCEPTIONS	170
8.4	VUCOMXSD	171
8.4.1	INSTRUCTION OPERAND ENCODING	171
8.4.2	DESCRIPTION	171
8.4.3	OPERATION	172
8.4.4	SIMD FLOATING-POINT EXCEPTIONS	172
8.4.5	EXCEPTIONS	172
8.5	VUCOMXSH	173
8.5.1	INSTRUCTION OPERAND ENCODING	173
8.5.2	DESCRIPTION	173
8.5.3	OPERATION	174
8.5.4	SIMD FLOATING-POINT EXCEPTIONS	174
8.5.5	EXCEPTIONS	174
8.6	VUCOMXSS	175
8.6.1	INSTRUCTION OPERAND ENCODING	175
8.6.2	DESCRIPTION	175
8.6.3	OPERATION	176
8.6.4	SIMD FLOATING-POINT EXCEPTIONS	176
8.6.5	EXCEPTIONS	176
9	INTEL® AVX10 CONVERT INSTRUCTIONS	177
9.1	VCVT2PS2PHX	178
9.1.1	INSTRUCTION OPERAND ENCODING	178
9.1.2	DESCRIPTION	178
9.1.3	OPERATION	179
9.1.4	SIMD FLOATING-POINT EXCEPTIONS	179
9.1.5	EXCEPTIONS	179
9.2	VCVTBIASPH2[B,H]F8[S]	180
9.2.1	INSTRUCTION OPERAND ENCODING	180
9.2.2	DESCRIPTION	181
9.2.3	OPERATION	182
9.2.4	SIMD FLOATING-POINT EXCEPTIONS	182
9.2.5	EXCEPTIONS	182
9.3	VCVTHF82PH	184
9.3.1	INSTRUCTION OPERAND ENCODING	184
9.3.2	DESCRIPTION	184
9.3.3	OPERATION	185
9.3.4	SIMD FLOATING-POINT EXCEPTIONS	185
9.3.5	EXCEPTIONS	185
9.4	VCVTNE[,2]PH2[B,H]F8[S]	186

9.4.1	INSTRUCTION OPERAND ENCODING	187
9.4.2	DESCRIPTION	187
9.4.3	OPERATION	189
9.4.4	SIMD FLOATING-POINT EXCEPTIONS	190
9.4.5	EXCEPTIONS	190
10	INTEL® AVX10 INTEGER AND FP16 VNNI, MEDIA NEW INSTRUCTIONS	193
10.1	VDPHPS	194
10.1.1	INSTRUCTION OPERAND ENCODING	194
10.1.2	DESCRIPTION	194
10.1.3	OPERATION	195
10.1.4	EXCEPTIONS	195
10.2	VMPSADBW	197
10.2.1	INSTRUCTION OPERAND ENCODING	197
10.2.2	DESCRIPTION	197
10.2.3	OPERATION	198
10.2.4	EXCEPTIONS	199
10.3	VPDPB[SU,UU,SS]D[,S]	201
10.3.1	INSTRUCTION OPERAND ENCODING	202
10.3.2	DESCRIPTION	202
10.3.3	OPERATION	203
10.3.4	EXCEPTIONS	204
10.4	VPDPW[SU,US,UU]D[,S]	206
10.4.1	INSTRUCTION OPERAND ENCODING	207
10.4.2	DESCRIPTION	207
10.4.3	OPERATION	208
10.4.4	EXCEPTIONS	209
11	INTEL® AVX10 MINMAX INSTRUCTIONS	211
11.1	VMINMAXNEPBF16	212
11.1.1	INSTRUCTION OPERAND ENCODING	212
11.1.2	DESCRIPTION	212
11.1.3	OPERATION	213
11.1.4	EXCEPTIONS	213
11.2	VMINMAX[PH,PS,PD]	214
11.2.1	INSTRUCTION OPERAND ENCODING	214
11.2.2	DESCRIPTION	215
11.2.3	OPERATION	219
11.2.4	SIMD FLOATING-POINT EXCEPTIONS	220
11.2.5	EXCEPTIONS	220
11.3	VMINMAX[SH,SS,SD]	222
11.3.1	INSTRUCTION OPERAND ENCODING	222
11.3.2	DESCRIPTION	222
11.3.3	OPERATION	223
11.3.4	SIMD FLOATING-POINT EXCEPTIONS	224
11.3.5	EXCEPTIONS	224
12	INTEL® AVX10 NEW INSTRUCTION FORMS	225

12.1	VADDPD	226
12.1.1	INSTRUCTION OPERAND ENCODING	226
12.1.2	DESCRIPTION	226
12.1.3	SIMD FLOATING-POINT EXCEPTIONS	226
12.1.4	EXCEPTIONS	226
12.2	VADDPH	227
12.2.1	INSTRUCTION OPERAND ENCODING	227
12.2.2	DESCRIPTION	227
12.2.3	SIMD FLOATING-POINT EXCEPTIONS	227
12.2.4	EXCEPTIONS	227
12.3	VADDP5	228
12.3.1	INSTRUCTION OPERAND ENCODING	228
12.3.2	DESCRIPTION	228
12.3.3	SIMD FLOATING-POINT EXCEPTIONS	228
12.3.4	EXCEPTIONS	228
12.4	VCMPD	229
12.4.1	INSTRUCTION OPERAND ENCODING	229
12.4.2	DESCRIPTION	229
12.4.3	SIMD FLOATING-POINT EXCEPTIONS	229
12.4.4	EXCEPTIONS	229
12.5	VCMPH	230
12.5.1	INSTRUCTION OPERAND ENCODING	230
12.5.2	DESCRIPTION	230
12.5.3	SIMD FLOATING-POINT EXCEPTIONS	230
12.5.4	EXCEPTIONS	230
12.6	VCMP5	231
12.6.1	INSTRUCTION OPERAND ENCODING	231
12.6.2	DESCRIPTION	231
12.6.3	SIMD FLOATING-POINT EXCEPTIONS	231
12.6.4	EXCEPTIONS	231
12.7	VCVTDQ2PH	232
12.7.1	INSTRUCTION OPERAND ENCODING	232
12.7.2	DESCRIPTION	232
12.7.3	SIMD FLOATING-POINT EXCEPTIONS	232
12.7.4	EXCEPTIONS	232
12.8	VCVTDQ2PS	233
12.8.1	INSTRUCTION OPERAND ENCODING	233
12.8.2	DESCRIPTION	233
12.8.3	SIMD FLOATING-POINT EXCEPTIONS	233
12.8.4	EXCEPTIONS	233
12.9	VCVTPD2DQ	234
12.9.1	INSTRUCTION OPERAND ENCODING	234
12.9.2	DESCRIPTION	234
12.9.3	SIMD FLOATING-POINT EXCEPTIONS	234
12.9.4	EXCEPTIONS	234
12.10	VCVTPD2PH	235
12.10.1	INSTRUCTION OPERAND ENCODING	235
12.10.2	DESCRIPTION	235

12.10.3	SIMD FLOATING-POINT EXCEPTIONS	235
12.10.4	EXCEPTIONS	235
12.11	VCVTPD2PS	236
12.11.1	INSTRUCTION OPERAND ENCODING	236
12.11.2	DESCRIPTION	236
12.11.3	SIMD FLOATING-POINT EXCEPTIONS	236
12.11.4	EXCEPTIONS	236
12.12	VCVTPD2QQ	237
12.12.1	INSTRUCTION OPERAND ENCODING	237
12.12.2	DESCRIPTION	237
12.12.3	SIMD FLOATING-POINT EXCEPTIONS	237
12.12.4	EXCEPTIONS	237
12.13	VCVTPD2UDQ	238
12.13.1	INSTRUCTION OPERAND ENCODING	238
12.13.2	DESCRIPTION	238
12.13.3	SIMD FLOATING-POINT EXCEPTIONS	238
12.13.4	EXCEPTIONS	238
12.14	VCVTPD2UQQ	239
12.14.1	INSTRUCTION OPERAND ENCODING	239
12.14.2	DESCRIPTION	239
12.14.3	SIMD FLOATING-POINT EXCEPTIONS	239
12.14.4	EXCEPTIONS	239
12.15	VCVTPH2DQ	240
12.15.1	INSTRUCTION OPERAND ENCODING	240
12.15.2	DESCRIPTION	240
12.15.3	SIMD FLOATING-POINT EXCEPTIONS	240
12.15.4	EXCEPTIONS	240
12.16	VCVTPH2PD	241
12.16.1	INSTRUCTION OPERAND ENCODING	241
12.16.2	DESCRIPTION	241
12.16.3	SIMD FLOATING-POINT EXCEPTIONS	241
12.16.4	EXCEPTIONS	241
12.17	VCVTPH2PS	242
12.17.1	INSTRUCTION OPERAND ENCODING	242
12.17.2	DESCRIPTION	242
12.17.3	SIMD FLOATING-POINT EXCEPTIONS	242
12.17.4	EXCEPTIONS	242
12.18	VCVTPH2PSX	243
12.18.1	INSTRUCTION OPERAND ENCODING	243
12.18.2	DESCRIPTION	243
12.18.3	SIMD FLOATING-POINT EXCEPTIONS	243
12.18.4	EXCEPTIONS	243
12.19	VCVTPH2QQ	244
12.19.1	INSTRUCTION OPERAND ENCODING	244
12.19.2	DESCRIPTION	244
12.19.3	SIMD FLOATING-POINT EXCEPTIONS	244
12.19.4	EXCEPTIONS	244
12.20	VCVTPH2UDQ	245

12.20.1	INSTRUCTION OPERAND ENCODING	245
12.20.2	DESCRIPTION	245
12.20.3	SIMD FLOATING-POINT EXCEPTIONS	245
12.20.4	EXCEPTIONS	245
12.21	VCVTPH2UQQ	246
12.21.1	INSTRUCTION OPERAND ENCODING	246
12.21.2	DESCRIPTION	246
12.21.3	SIMD FLOATING-POINT EXCEPTIONS	246
12.21.4	EXCEPTIONS	246
12.22	VCVTPH2UW	247
12.22.1	INSTRUCTION OPERAND ENCODING	247
12.22.2	DESCRIPTION	247
12.22.3	SIMD FLOATING-POINT EXCEPTIONS	247
12.22.4	EXCEPTIONS	247
12.23	VCVTPH2W	248
12.23.1	INSTRUCTION OPERAND ENCODING	248
12.23.2	DESCRIPTION	248
12.23.3	SIMD FLOATING-POINT EXCEPTIONS	248
12.23.4	EXCEPTIONS	248
12.24	VCVTPS2DQ	249
12.24.1	INSTRUCTION OPERAND ENCODING	249
12.24.2	DESCRIPTION	249
12.24.3	SIMD FLOATING-POINT EXCEPTIONS	249
12.24.4	EXCEPTIONS	249
12.25	VCVTPS2PD	250
12.25.1	INSTRUCTION OPERAND ENCODING	250
12.25.2	DESCRIPTION	250
12.25.3	SIMD FLOATING-POINT EXCEPTIONS	250
12.25.4	EXCEPTIONS	250
12.26	VCVTPS2PH	251
12.26.1	INSTRUCTION OPERAND ENCODING	251
12.26.2	DESCRIPTION	251
12.26.3	SIMD FLOATING-POINT EXCEPTIONS	251
12.26.4	EXCEPTIONS	251
12.27	VCVTPS2PHX	252
12.27.1	INSTRUCTION OPERAND ENCODING	252
12.27.2	DESCRIPTION	252
12.27.3	SIMD FLOATING-POINT EXCEPTIONS	252
12.27.4	EXCEPTIONS	252
12.28	VCVTPS2QQ	253
12.28.1	INSTRUCTION OPERAND ENCODING	253
12.28.2	DESCRIPTION	253
12.28.3	SIMD FLOATING-POINT EXCEPTIONS	253
12.28.4	EXCEPTIONS	253
12.29	VCVTPS2UDQ	254
12.29.1	INSTRUCTION OPERAND ENCODING	254
12.29.2	DESCRIPTION	254
12.29.3	SIMD FLOATING-POINT EXCEPTIONS	254

12.29.4	EXCEPTIONS	254
12.30	VCVTPS2UQQ	255
12.30.1	INSTRUCTION OPERAND ENCODING	255
12.30.2	DESCRIPTION	255
12.30.3	SIMD FLOATING-POINT EXCEPTIONS	255
12.30.4	EXCEPTIONS	255
12.31	VCVTQ2PD	256
12.31.1	INSTRUCTION OPERAND ENCODING	256
12.31.2	DESCRIPTION	256
12.31.3	SIMD FLOATING-POINT EXCEPTIONS	256
12.31.4	EXCEPTIONS	256
12.32	VCVTQ2PH	257
12.32.1	INSTRUCTION OPERAND ENCODING	257
12.32.2	DESCRIPTION	257
12.32.3	SIMD FLOATING-POINT EXCEPTIONS	257
12.32.4	EXCEPTIONS	257
12.33	VCVTQ2PS	258
12.33.1	INSTRUCTION OPERAND ENCODING	258
12.33.2	DESCRIPTION	258
12.33.3	SIMD FLOATING-POINT EXCEPTIONS	258
12.33.4	EXCEPTIONS	258
12.34	VCVTTPD2DQ	259
12.34.1	INSTRUCTION OPERAND ENCODING	259
12.34.2	DESCRIPTION	259
12.34.3	SIMD FLOATING-POINT EXCEPTIONS	259
12.34.4	EXCEPTIONS	259
12.35	VCVTTPD2QQ	260
12.35.1	INSTRUCTION OPERAND ENCODING	260
12.35.2	DESCRIPTION	260
12.35.3	SIMD FLOATING-POINT EXCEPTIONS	260
12.35.4	EXCEPTIONS	260
12.36	VCVTTPD2UDQ	261
12.36.1	INSTRUCTION OPERAND ENCODING	261
12.36.2	DESCRIPTION	261
12.36.3	SIMD FLOATING-POINT EXCEPTIONS	261
12.36.4	EXCEPTIONS	261
12.37	VCVTTPD2UQQ	262
12.37.1	INSTRUCTION OPERAND ENCODING	262
12.37.2	DESCRIPTION	262
12.37.3	SIMD FLOATING-POINT EXCEPTIONS	262
12.37.4	EXCEPTIONS	262
12.38	VCVTTPH2DQ	263
12.38.1	INSTRUCTION OPERAND ENCODING	263
12.38.2	DESCRIPTION	263
12.38.3	SIMD FLOATING-POINT EXCEPTIONS	263
12.38.4	EXCEPTIONS	263
12.39	VCVTTPH2QQ	264
12.39.1	INSTRUCTION OPERAND ENCODING	264

12.39.2	DESCRIPTION	264
12.39.3	SIMD FLOATING-POINT EXCEPTIONS	264
12.39.4	EXCEPTIONS	264
12.40	VCVTTPH2UDQ	265
12.40.1	INSTRUCTION OPERAND ENCODING	265
12.40.2	DESCRIPTION	265
12.40.3	SIMD FLOATING-POINT EXCEPTIONS	265
12.40.4	EXCEPTIONS	265
12.41	VCVTTPH2UQQ	266
12.41.1	INSTRUCTION OPERAND ENCODING	266
12.41.2	DESCRIPTION	266
12.41.3	SIMD FLOATING-POINT EXCEPTIONS	266
12.41.4	EXCEPTIONS	266
12.42	VCVTTPH2UW	267
12.42.1	INSTRUCTION OPERAND ENCODING	267
12.42.2	DESCRIPTION	267
12.42.3	SIMD FLOATING-POINT EXCEPTIONS	267
12.42.4	EXCEPTIONS	267
12.43	VCVTTPH2W	268
12.43.1	INSTRUCTION OPERAND ENCODING	268
12.43.2	DESCRIPTION	268
12.43.3	SIMD FLOATING-POINT EXCEPTIONS	268
12.43.4	EXCEPTIONS	268
12.44	VCVTTPS2DQ	269
12.44.1	INSTRUCTION OPERAND ENCODING	269
12.44.2	DESCRIPTION	269
12.44.3	SIMD FLOATING-POINT EXCEPTIONS	269
12.44.4	EXCEPTIONS	269
12.45	VCVTTPS2QQ	270
12.45.1	INSTRUCTION OPERAND ENCODING	270
12.45.2	DESCRIPTION	270
12.45.3	SIMD FLOATING-POINT EXCEPTIONS	270
12.45.4	EXCEPTIONS	270
12.46	VCVTTPS2UDQ	271
12.46.1	INSTRUCTION OPERAND ENCODING	271
12.46.2	DESCRIPTION	271
12.46.3	SIMD FLOATING-POINT EXCEPTIONS	271
12.46.4	EXCEPTIONS	271
12.47	VCVTTPS2UQQ	272
12.47.1	INSTRUCTION OPERAND ENCODING	272
12.47.2	DESCRIPTION	272
12.47.3	SIMD FLOATING-POINT EXCEPTIONS	272
12.47.4	EXCEPTIONS	272
12.48	VCVTUDQ2PH	273
12.48.1	INSTRUCTION OPERAND ENCODING	273
12.48.2	DESCRIPTION	273
12.48.3	SIMD FLOATING-POINT EXCEPTIONS	273
12.48.4	EXCEPTIONS	273

12.49	VCVTUDQ2PS	274
12.49.1	INSTRUCTION OPERAND ENCODING	274
12.49.2	DESCRIPTION	274
12.49.3	SIMD FLOATING-POINT EXCEPTIONS	274
12.49.4	EXCEPTIONS	274
12.50	VCVTUQQ2PD	275
12.50.1	INSTRUCTION OPERAND ENCODING	275
12.50.2	DESCRIPTION	275
12.50.3	SIMD FLOATING-POINT EXCEPTIONS	275
12.50.4	EXCEPTIONS	275
12.51	VCVTUQQ2PH	276
12.51.1	INSTRUCTION OPERAND ENCODING	276
12.51.2	DESCRIPTION	276
12.51.3	SIMD FLOATING-POINT EXCEPTIONS	276
12.51.4	EXCEPTIONS	276
12.52	VCVTUQQ2PS	277
12.52.1	INSTRUCTION OPERAND ENCODING	277
12.52.2	DESCRIPTION	277
12.52.3	SIMD FLOATING-POINT EXCEPTIONS	277
12.52.4	EXCEPTIONS	277
12.53	VCVTUW2PH	278
12.53.1	INSTRUCTION OPERAND ENCODING	278
12.53.2	DESCRIPTION	278
12.53.3	SIMD FLOATING-POINT EXCEPTIONS	278
12.53.4	EXCEPTIONS	278
12.54	VCVTW2PH	279
12.54.1	INSTRUCTION OPERAND ENCODING	279
12.54.2	DESCRIPTION	279
12.54.3	SIMD FLOATING-POINT EXCEPTIONS	279
12.54.4	EXCEPTIONS	279
12.55	VDIVPD	280
12.55.1	INSTRUCTION OPERAND ENCODING	280
12.55.2	DESCRIPTION	280
12.55.3	SIMD FLOATING-POINT EXCEPTIONS	280
12.55.4	EXCEPTIONS	280
12.56	VDIVPH	281
12.56.1	INSTRUCTION OPERAND ENCODING	281
12.56.2	DESCRIPTION	281
12.56.3	SIMD FLOATING-POINT EXCEPTIONS	281
12.56.4	EXCEPTIONS	281
12.57	VDIVPS	282
12.57.1	INSTRUCTION OPERAND ENCODING	282
12.57.2	DESCRIPTION	282
12.57.3	SIMD FLOATING-POINT EXCEPTIONS	282
12.57.4	EXCEPTIONS	282
12.58	VFCMADDCPH	283
12.58.1	INSTRUCTION OPERAND ENCODING	283
12.58.2	DESCRIPTION	283

12.58.3	SIMD FLOATING-POINT EXCEPTIONS	283
12.58.4	EXCEPTIONS	283
12.59	VFCMULCPH	284
12.59.1	INSTRUCTION OPERAND ENCODING	284
12.59.2	DESCRIPTION	284
12.59.3	SIMD FLOATING-POINT EXCEPTIONS	284
12.59.4	EXCEPTIONS	284
12.60	VFIXUPIMMPD	285
12.60.1	INSTRUCTION OPERAND ENCODING	285
12.60.2	DESCRIPTION	285
12.60.3	SIMD FLOATING-POINT EXCEPTIONS	285
12.60.4	EXCEPTIONS	285
12.61	VFIXUPIMMPS	286
12.61.1	INSTRUCTION OPERAND ENCODING	286
12.61.2	DESCRIPTION	286
12.61.3	SIMD FLOATING-POINT EXCEPTIONS	286
12.61.4	EXCEPTIONS	286
12.62	VFMADD132PD	287
12.62.1	INSTRUCTION OPERAND ENCODING	287
12.62.2	DESCRIPTION	287
12.62.3	SIMD FLOATING-POINT EXCEPTIONS	287
12.62.4	EXCEPTIONS	287
12.63	VFMADD132PH	288
12.63.1	INSTRUCTION OPERAND ENCODING	288
12.63.2	DESCRIPTION	288
12.63.3	SIMD FLOATING-POINT EXCEPTIONS	288
12.63.4	EXCEPTIONS	288
12.64	VFMADD132PS	289
12.64.1	INSTRUCTION OPERAND ENCODING	289
12.64.2	DESCRIPTION	289
12.64.3	SIMD FLOATING-POINT EXCEPTIONS	289
12.64.4	EXCEPTIONS	289
12.65	VFMADD213PD	290
12.65.1	INSTRUCTION OPERAND ENCODING	290
12.65.2	DESCRIPTION	290
12.65.3	SIMD FLOATING-POINT EXCEPTIONS	290
12.65.4	EXCEPTIONS	290
12.66	VFMADD213PH	291
12.66.1	INSTRUCTION OPERAND ENCODING	291
12.66.2	DESCRIPTION	291
12.66.3	SIMD FLOATING-POINT EXCEPTIONS	291
12.66.4	EXCEPTIONS	291
12.67	VFMADD213PS	292
12.67.1	INSTRUCTION OPERAND ENCODING	292
12.67.2	DESCRIPTION	292
12.67.3	SIMD FLOATING-POINT EXCEPTIONS	292
12.67.4	EXCEPTIONS	292
12.68	VFMADD231PD	293

12.68.1	INSTRUCTION OPERAND ENCODING	293
12.68.2	DESCRIPTION	293
12.68.3	SIMD FLOATING-POINT EXCEPTIONS	293
12.68.4	EXCEPTIONS	293
12.69	VFMADD231PH	294
12.69.1	INSTRUCTION OPERAND ENCODING	294
12.69.2	DESCRIPTION	294
12.69.3	SIMD FLOATING-POINT EXCEPTIONS	294
12.69.4	EXCEPTIONS	294
12.70	VFMADD231PS	295
12.70.1	INSTRUCTION OPERAND ENCODING	295
12.70.2	DESCRIPTION	295
12.70.3	SIMD FLOATING-POINT EXCEPTIONS	295
12.70.4	EXCEPTIONS	295
12.71	VFMADDCPH	296
12.71.1	INSTRUCTION OPERAND ENCODING	296
12.71.2	DESCRIPTION	296
12.71.3	SIMD FLOATING-POINT EXCEPTIONS	296
12.71.4	EXCEPTIONS	296
12.72	VFMADDSUB132PD	297
12.72.1	INSTRUCTION OPERAND ENCODING	297
12.72.2	DESCRIPTION	297
12.72.3	SIMD FLOATING-POINT EXCEPTIONS	297
12.72.4	EXCEPTIONS	297
12.73	VFMADDSUB132PH	298
12.73.1	INSTRUCTION OPERAND ENCODING	298
12.73.2	DESCRIPTION	298
12.73.3	SIMD FLOATING-POINT EXCEPTIONS	298
12.73.4	EXCEPTIONS	298
12.74	VFMADDSUB132PS	299
12.74.1	INSTRUCTION OPERAND ENCODING	299
12.74.2	DESCRIPTION	299
12.74.3	SIMD FLOATING-POINT EXCEPTIONS	299
12.74.4	EXCEPTIONS	299
12.75	VFMADDSUB213PD	300
12.75.1	INSTRUCTION OPERAND ENCODING	300
12.75.2	DESCRIPTION	300
12.75.3	SIMD FLOATING-POINT EXCEPTIONS	300
12.75.4	EXCEPTIONS	300
12.76	VFMADDSUB213PH	301
12.76.1	INSTRUCTION OPERAND ENCODING	301
12.76.2	DESCRIPTION	301
12.76.3	SIMD FLOATING-POINT EXCEPTIONS	301
12.76.4	EXCEPTIONS	301
12.77	VFMADDSUB213PS	302
12.77.1	INSTRUCTION OPERAND ENCODING	302
12.77.2	DESCRIPTION	302
12.77.3	SIMD FLOATING-POINT EXCEPTIONS	302

12.77.4	EXCEPTIONS	302
12.78	VFMADDSUB231PD	303
12.78.1	INSTRUCTION OPERAND ENCODING	303
12.78.2	DESCRIPTION	303
12.78.3	SIMD FLOATING-POINT EXCEPTIONS	303
12.78.4	EXCEPTIONS	303
12.79	VFMADDSUB231PH	304
12.79.1	INSTRUCTION OPERAND ENCODING	304
12.79.2	DESCRIPTION	304
12.79.3	SIMD FLOATING-POINT EXCEPTIONS	304
12.79.4	EXCEPTIONS	304
12.80	VFMADDSUB231PS	305
12.80.1	INSTRUCTION OPERAND ENCODING	305
12.80.2	DESCRIPTION	305
12.80.3	SIMD FLOATING-POINT EXCEPTIONS	305
12.80.4	EXCEPTIONS	305
12.81	VFMSUB132PD	306
12.81.1	INSTRUCTION OPERAND ENCODING	306
12.81.2	DESCRIPTION	306
12.81.3	SIMD FLOATING-POINT EXCEPTIONS	306
12.81.4	EXCEPTIONS	306
12.82	VFMSUB132PH	307
12.82.1	INSTRUCTION OPERAND ENCODING	307
12.82.2	DESCRIPTION	307
12.82.3	SIMD FLOATING-POINT EXCEPTIONS	307
12.82.4	EXCEPTIONS	307
12.83	VFMSUB132PS	308
12.83.1	INSTRUCTION OPERAND ENCODING	308
12.83.2	DESCRIPTION	308
12.83.3	SIMD FLOATING-POINT EXCEPTIONS	308
12.83.4	EXCEPTIONS	308
12.84	VFMSUB213PD	309
12.84.1	INSTRUCTION OPERAND ENCODING	309
12.84.2	DESCRIPTION	309
12.84.3	SIMD FLOATING-POINT EXCEPTIONS	309
12.84.4	EXCEPTIONS	309
12.85	VFMSUB213PH	310
12.85.1	INSTRUCTION OPERAND ENCODING	310
12.85.2	DESCRIPTION	310
12.85.3	SIMD FLOATING-POINT EXCEPTIONS	310
12.85.4	EXCEPTIONS	310
12.86	VFMSUB213PS	311
12.86.1	INSTRUCTION OPERAND ENCODING	311
12.86.2	DESCRIPTION	311
12.86.3	SIMD FLOATING-POINT EXCEPTIONS	311
12.86.4	EXCEPTIONS	311
12.87	VFMSUB231PD	312
12.87.1	INSTRUCTION OPERAND ENCODING	312

12.87.2	DESCRIPTION	312
12.87.3	SIMD FLOATING-POINT EXCEPTIONS	312
12.87.4	EXCEPTIONS	312
12.88	VFMSUB231PH	313
12.88.1	INSTRUCTION OPERAND ENCODING	313
12.88.2	DESCRIPTION	313
12.88.3	SIMD FLOATING-POINT EXCEPTIONS	313
12.88.4	EXCEPTIONS	313
12.89	VFMSUB231PS	314
12.89.1	INSTRUCTION OPERAND ENCODING	314
12.89.2	DESCRIPTION	314
12.89.3	SIMD FLOATING-POINT EXCEPTIONS	314
12.89.4	EXCEPTIONS	314
12.90	VFMSUBADD132PD	315
12.90.1	INSTRUCTION OPERAND ENCODING	315
12.90.2	DESCRIPTION	315
12.90.3	SIMD FLOATING-POINT EXCEPTIONS	315
12.90.4	EXCEPTIONS	315
12.91	VFMSUBADD132PH	316
12.91.1	INSTRUCTION OPERAND ENCODING	316
12.91.2	DESCRIPTION	316
12.91.3	SIMD FLOATING-POINT EXCEPTIONS	316
12.91.4	EXCEPTIONS	316
12.92	VFMSUBADD132PS	317
12.92.1	INSTRUCTION OPERAND ENCODING	317
12.92.2	DESCRIPTION	317
12.92.3	SIMD FLOATING-POINT EXCEPTIONS	317
12.92.4	EXCEPTIONS	317
12.93	VFMSUBADD213PD	318
12.93.1	INSTRUCTION OPERAND ENCODING	318
12.93.2	DESCRIPTION	318
12.93.3	SIMD FLOATING-POINT EXCEPTIONS	318
12.93.4	EXCEPTIONS	318
12.94	VFMSUBADD213PH	319
12.94.1	INSTRUCTION OPERAND ENCODING	319
12.94.2	DESCRIPTION	319
12.94.3	SIMD FLOATING-POINT EXCEPTIONS	319
12.94.4	EXCEPTIONS	319
12.95	VFMSUBADD213PS	320
12.95.1	INSTRUCTION OPERAND ENCODING	320
12.95.2	DESCRIPTION	320
12.95.3	SIMD FLOATING-POINT EXCEPTIONS	320
12.95.4	EXCEPTIONS	320
12.96	VFMSUBADD231PD	321
12.96.1	INSTRUCTION OPERAND ENCODING	321
12.96.2	DESCRIPTION	321
12.96.3	SIMD FLOATING-POINT EXCEPTIONS	321
12.96.4	EXCEPTIONS	321

12.97	VFMSUBADD231PH	322
12.97.1	INSTRUCTION OPERAND ENCODING	322
12.97.2	DESCRIPTION	322
12.97.3	SIMD FLOATING-POINT EXCEPTIONS	322
12.97.4	EXCEPTIONS	322
12.98	VFMSUBADD231PS	323
12.98.1	INSTRUCTION OPERAND ENCODING	323
12.98.2	DESCRIPTION	323
12.98.3	SIMD FLOATING-POINT EXCEPTIONS	323
12.98.4	EXCEPTIONS	323
12.99	VFMULCPH	324
12.99.1	INSTRUCTION OPERAND ENCODING	324
12.99.2	DESCRIPTION	324
12.99.3	SIMD FLOATING-POINT EXCEPTIONS	324
12.99.4	EXCEPTIONS	324
12.100	VFNMADD132PD	325
12.100.1	INSTRUCTION OPERAND ENCODING	325
12.100.2	DESCRIPTION	325
12.100.3	SIMD FLOATING-POINT EXCEPTIONS	325
12.100.4	EXCEPTIONS	325
12.101	VFNMADD132PH	326
12.101.1	INSTRUCTION OPERAND ENCODING	326
12.101.2	DESCRIPTION	326
12.101.3	SIMD FLOATING-POINT EXCEPTIONS	326
12.101.4	EXCEPTIONS	326
12.102	VFNMADD132PS	327
12.102.1	INSTRUCTION OPERAND ENCODING	327
12.102.2	DESCRIPTION	327
12.102.3	SIMD FLOATING-POINT EXCEPTIONS	327
12.102.4	EXCEPTIONS	327
12.103	VFNMADD213PD	328
12.103.1	INSTRUCTION OPERAND ENCODING	328
12.103.2	DESCRIPTION	328
12.103.3	SIMD FLOATING-POINT EXCEPTIONS	328
12.103.4	EXCEPTIONS	328
12.104	VFNMADD213PH	329
12.104.1	INSTRUCTION OPERAND ENCODING	329
12.104.2	DESCRIPTION	329
12.104.3	SIMD FLOATING-POINT EXCEPTIONS	329
12.104.4	EXCEPTIONS	329
12.105	VFNMADD213PS	330
12.105.1	INSTRUCTION OPERAND ENCODING	330
12.105.2	DESCRIPTION	330
12.105.3	SIMD FLOATING-POINT EXCEPTIONS	330
12.105.4	EXCEPTIONS	330
12.106	VFNMADD231PD	331
12.106.1	INSTRUCTION OPERAND ENCODING	331
12.106.2	DESCRIPTION	331

12.106.3	SIMD FLOATING-POINT EXCEPTIONS	331
12.106.4	EXCEPTIONS	331
12.107	VFNMADD231PH	332
12.107.1	INSTRUCTION OPERAND ENCODING	332
12.107.2	DESCRIPTION	332
12.107.3	SIMD FLOATING-POINT EXCEPTIONS	332
12.107.4	EXCEPTIONS	332
12.108	VFNMADD231PS	333
12.108.1	INSTRUCTION OPERAND ENCODING	333
12.108.2	DESCRIPTION	333
12.108.3	SIMD FLOATING-POINT EXCEPTIONS	333
12.108.4	EXCEPTIONS	333
12.109	VFNMSUB132PD	334
12.109.1	INSTRUCTION OPERAND ENCODING	334
12.109.2	DESCRIPTION	334
12.109.3	SIMD FLOATING-POINT EXCEPTIONS	334
12.109.4	EXCEPTIONS	334
12.110	VFNMSUB132PH	335
12.110.1	INSTRUCTION OPERAND ENCODING	335
12.110.2	DESCRIPTION	335
12.110.3	SIMD FLOATING-POINT EXCEPTIONS	335
12.110.4	EXCEPTIONS	335
12.111	VFNMSUB132PS	336
12.111.1	INSTRUCTION OPERAND ENCODING	336
12.111.2	DESCRIPTION	336
12.111.3	SIMD FLOATING-POINT EXCEPTIONS	336
12.111.4	EXCEPTIONS	336
12.112	VFNMSUB213PD	337
12.112.1	INSTRUCTION OPERAND ENCODING	337
12.112.2	DESCRIPTION	337
12.112.3	SIMD FLOATING-POINT EXCEPTIONS	337
12.112.4	EXCEPTIONS	337
12.113	VFNMSUB213PH	338
12.113.1	INSTRUCTION OPERAND ENCODING	338
12.113.2	DESCRIPTION	338
12.113.3	SIMD FLOATING-POINT EXCEPTIONS	338
12.113.4	EXCEPTIONS	338
12.114	VFNMSUB213PS	339
12.114.1	INSTRUCTION OPERAND ENCODING	339
12.114.2	DESCRIPTION	339
12.114.3	SIMD FLOATING-POINT EXCEPTIONS	339
12.114.4	EXCEPTIONS	339
12.115	VFNMSUB231PD	340
12.115.1	INSTRUCTION OPERAND ENCODING	340
12.115.2	DESCRIPTION	340
12.115.3	SIMD FLOATING-POINT EXCEPTIONS	340
12.115.4	EXCEPTIONS	340
12.116	VFNMSUB231PH	341

12.116.1	INSTRUCTION OPERAND ENCODING	341
12.116.2	DESCRIPTION	341
12.116.3	SIMD FLOATING-POINT EXCEPTIONS	341
12.116.4	EXCEPTIONS	341
12.117	VFNMSUB231PS	342
12.117.1	INSTRUCTION OPERAND ENCODING	342
12.117.2	DESCRIPTION	342
12.117.3	SIMD FLOATING-POINT EXCEPTIONS	342
12.117.4	EXCEPTIONS	342
12.118	VGETEXPPD	343
12.118.1	INSTRUCTION OPERAND ENCODING	343
12.118.2	DESCRIPTION	343
12.118.3	SIMD FLOATING-POINT EXCEPTIONS	343
12.118.4	EXCEPTIONS	343
12.119	VGETEXPPH	344
12.119.1	INSTRUCTION OPERAND ENCODING	344
12.119.2	DESCRIPTION	344
12.119.3	SIMD FLOATING-POINT EXCEPTIONS	344
12.119.4	EXCEPTIONS	344
12.120	VGETEXPPS	345
12.120.1	INSTRUCTION OPERAND ENCODING	345
12.120.2	DESCRIPTION	345
12.120.3	SIMD FLOATING-POINT EXCEPTIONS	345
12.120.4	EXCEPTIONS	345
12.121	VGETMANTPD	346
12.121.1	INSTRUCTION OPERAND ENCODING	346
12.121.2	DESCRIPTION	346
12.121.3	SIMD FLOATING-POINT EXCEPTIONS	346
12.121.4	EXCEPTIONS	346
12.122	VGETMANTPH	347
12.122.1	INSTRUCTION OPERAND ENCODING	347
12.122.2	DESCRIPTION	347
12.122.3	SIMD FLOATING-POINT EXCEPTIONS	347
12.122.4	EXCEPTIONS	347
12.123	VGETMANTPS	348
12.123.1	INSTRUCTION OPERAND ENCODING	348
12.123.2	DESCRIPTION	348
12.123.3	SIMD FLOATING-POINT EXCEPTIONS	348
12.123.4	EXCEPTIONS	348
12.124	VMAXPD	349
12.124.1	INSTRUCTION OPERAND ENCODING	349
12.124.2	DESCRIPTION	349
12.124.3	SIMD FLOATING-POINT EXCEPTIONS	349
12.124.4	EXCEPTIONS	349
12.125	VMAXPH	350
12.125.1	INSTRUCTION OPERAND ENCODING	350
12.125.2	DESCRIPTION	350
12.125.3	SIMD FLOATING-POINT EXCEPTIONS	350

12.125.4	EXCEPTIONS	350
12.126	VMAXPS	351
12.126.1	INSTRUCTION OPERAND ENCODING	351
12.126.2	DESCRIPTION	351
12.126.3	SIMD FLOATING-POINT EXCEPTIONS	351
12.126.4	EXCEPTIONS	351
12.127	VMINPD	352
12.127.1	INSTRUCTION OPERAND ENCODING	352
12.127.2	DESCRIPTION	352
12.127.3	SIMD FLOATING-POINT EXCEPTIONS	352
12.127.4	EXCEPTIONS	352
12.128	VMINPH	353
12.128.1	INSTRUCTION OPERAND ENCODING	353
12.128.2	DESCRIPTION	353
12.128.3	SIMD FLOATING-POINT EXCEPTIONS	353
12.128.4	EXCEPTIONS	353
12.129	VMINPS	354
12.129.1	INSTRUCTION OPERAND ENCODING	354
12.129.2	DESCRIPTION	354
12.129.3	SIMD FLOATING-POINT EXCEPTIONS	354
12.129.4	EXCEPTIONS	354
12.130	VMULPD	355
12.130.1	INSTRUCTION OPERAND ENCODING	355
12.130.2	DESCRIPTION	355
12.130.3	SIMD FLOATING-POINT EXCEPTIONS	355
12.130.4	EXCEPTIONS	355
12.131	VMULPH	356
12.131.1	INSTRUCTION OPERAND ENCODING	356
12.131.2	DESCRIPTION	356
12.131.3	SIMD FLOATING-POINT EXCEPTIONS	356
12.131.4	EXCEPTIONS	356
12.132	VMULPS	357
12.132.1	INSTRUCTION OPERAND ENCODING	357
12.132.2	DESCRIPTION	357
12.132.3	SIMD FLOATING-POINT EXCEPTIONS	357
12.132.4	EXCEPTIONS	357
12.133	VRANGEPD	358
12.133.1	INSTRUCTION OPERAND ENCODING	358
12.133.2	DESCRIPTION	358
12.133.3	SIMD FLOATING-POINT EXCEPTIONS	358
12.133.4	EXCEPTIONS	358
12.134	VRANGEPS	359
12.134.1	INSTRUCTION OPERAND ENCODING	359
12.134.2	DESCRIPTION	359
12.134.3	SIMD FLOATING-POINT EXCEPTIONS	359
12.134.4	EXCEPTIONS	359
12.135	VREDUCEPD	360
12.135.1	INSTRUCTION OPERAND ENCODING	360

12.135.2	DESCRIPTION	360
12.135.3	SIMD FLOATING-POINT EXCEPTIONS	360
12.135.4	EXCEPTIONS	360
12.136	VREDUCEPH	361
12.136.1	INSTRUCTION OPERAND ENCODING	361
12.136.2	DESCRIPTION	361
12.136.3	SIMD FLOATING-POINT EXCEPTIONS	361
12.136.4	EXCEPTIONS	361
12.137	VREDUCEPS	362
12.137.1	INSTRUCTION OPERAND ENCODING	362
12.137.2	DESCRIPTION	362
12.137.3	SIMD FLOATING-POINT EXCEPTIONS	362
12.137.4	EXCEPTIONS	362
12.138	VRNDSCALEPD	363
12.138.1	INSTRUCTION OPERAND ENCODING	363
12.138.2	DESCRIPTION	363
12.138.3	SIMD FLOATING-POINT EXCEPTIONS	363
12.138.4	EXCEPTIONS	363
12.139	VRNDSCALEPH	364
12.139.1	INSTRUCTION OPERAND ENCODING	364
12.139.2	DESCRIPTION	364
12.139.3	SIMD FLOATING-POINT EXCEPTIONS	364
12.139.4	EXCEPTIONS	364
12.140	VRNDSCALEPS	365
12.140.1	INSTRUCTION OPERAND ENCODING	365
12.140.2	DESCRIPTION	365
12.140.3	SIMD FLOATING-POINT EXCEPTIONS	365
12.140.4	EXCEPTIONS	365
12.141	VSCALEFPD	366
12.141.1	INSTRUCTION OPERAND ENCODING	366
12.141.2	DESCRIPTION	366
12.141.3	SIMD FLOATING-POINT EXCEPTIONS	366
12.141.4	EXCEPTIONS	366
12.142	VSCALEFPH	367
12.142.1	INSTRUCTION OPERAND ENCODING	367
12.142.2	DESCRIPTION	367
12.142.3	SIMD FLOATING-POINT EXCEPTIONS	367
12.142.4	EXCEPTIONS	367
12.143	VSCALEFPS	368
12.143.1	INSTRUCTION OPERAND ENCODING	368
12.143.2	DESCRIPTION	368
12.143.3	SIMD FLOATING-POINT EXCEPTIONS	368
12.143.4	EXCEPTIONS	368
12.144	VSQRTPD	369
12.144.1	INSTRUCTION OPERAND ENCODING	369
12.144.2	DESCRIPTION	369
12.144.3	SIMD FLOATING-POINT EXCEPTIONS	369
12.144.4	EXCEPTIONS	369

12.145	VSQRTPH	370
12.145.1	INSTRUCTION OPERAND ENCODING	370
12.145.2	DESCRIPTION	370
12.145.3	SIMD FLOATING-POINT EXCEPTIONS	370
12.145.4	EXCEPTIONS	370
12.146	VSQRTPS	371
12.146.1	INSTRUCTION OPERAND ENCODING	371
12.146.2	DESCRIPTION	371
12.146.3	SIMD FLOATING-POINT EXCEPTIONS	371
12.146.4	EXCEPTIONS	371
12.147	VSUBPD	372
12.147.1	INSTRUCTION OPERAND ENCODING	372
12.147.2	DESCRIPTION	372
12.147.3	SIMD FLOATING-POINT EXCEPTIONS	372
12.147.4	EXCEPTIONS	372
12.148	VSUBPH	373
12.148.1	INSTRUCTION OPERAND ENCODING	373
12.148.2	DESCRIPTION	373
12.148.3	SIMD FLOATING-POINT EXCEPTIONS	373
12.148.4	EXCEPTIONS	373
12.149	VSUBPS	374
12.149.1	INSTRUCTION OPERAND ENCODING	374
12.149.2	DESCRIPTION	374
12.149.3	SIMD FLOATING-POINT EXCEPTIONS	374
12.149.4	EXCEPTIONS	374
13	INTEL® AVX10 SATURATING CONVERT INSTRUCTIONS	375
13.1	VCVT[T]NEBF162I[U]BS	376
13.1.1	INSTRUCTION OPERAND ENCODING	376
13.1.2	DESCRIPTION	377
13.1.3	OPERATION	378
13.1.4	EXCEPTIONS	378
13.2	VCVTTPD2DQS	380
13.2.1	INSTRUCTION OPERAND ENCODING	380
13.2.2	DESCRIPTION	380
13.2.3	OPERATION	381
13.2.4	SIMD FLOATING-POINT EXCEPTIONS	382
13.2.5	EXCEPTIONS	382
13.3	VCVTTPD2QQS	383
13.3.1	INSTRUCTION OPERAND ENCODING	383
13.3.2	DESCRIPTION	383
13.3.3	OPERATION	384
13.3.4	SIMD FLOATING-POINT EXCEPTIONS	384
13.3.5	EXCEPTIONS	385
13.4	VCVTTPD2UDQS	386
13.4.1	INSTRUCTION OPERAND ENCODING	386
13.4.2	DESCRIPTION	386
13.4.3	OPERATION	387

13.4.4	SIMD FLOATING-POINT EXCEPTIONS	388
13.4.5	EXCEPTIONS	388
13.5	VCVTTPD2UQQS	389
13.5.1	INSTRUCTION OPERAND ENCODING	389
13.5.2	DESCRIPTION	389
13.5.3	OPERATION	390
13.5.4	SIMD FLOATING-POINT EXCEPTIONS	390
13.5.5	EXCEPTIONS	391
13.6	VCVT[T]PH2I[U]BS	392
13.6.1	INSTRUCTION OPERAND ENCODING	392
13.6.2	DESCRIPTION	393
13.6.3	OPERATION	394
13.6.4	SIMD FLOATING-POINT EXCEPTIONS	394
13.6.5	EXCEPTIONS	394
13.7	VCVTTPS2DQS	396
13.7.1	INSTRUCTION OPERAND ENCODING	396
13.7.2	DESCRIPTION	396
13.7.3	OPERATION	397
13.7.4	SIMD FLOATING-POINT EXCEPTIONS	397
13.7.5	EXCEPTIONS	398
13.8	VCVT[T]PS2I[U]BS	399
13.8.1	INSTRUCTION OPERAND ENCODING	399
13.8.2	DESCRIPTION	400
13.8.3	OPERATION	401
13.8.4	SIMD FLOATING-POINT EXCEPTIONS	401
13.8.5	EXCEPTIONS	401
13.9	VCVTTPS2QQS	403
13.9.1	INSTRUCTION OPERAND ENCODING	403
13.9.2	DESCRIPTION	403
13.9.3	OPERATION	404
13.9.4	SIMD FLOATING-POINT EXCEPTIONS	404
13.9.5	EXCEPTIONS	405
13.10	VCVTTPS2UDQS	406
13.10.1	INSTRUCTION OPERAND ENCODING	406
13.10.2	DESCRIPTION	406
13.10.3	OPERATION	407
13.10.4	SIMD FLOATING-POINT EXCEPTIONS	407
13.10.5	EXCEPTIONS	408
13.11	VCVTTPS2UQQS	409
13.11.1	INSTRUCTION OPERAND ENCODING	409
13.11.2	DESCRIPTION	409
13.11.3	OPERATION	410
13.11.4	SIMD FLOATING-POINT EXCEPTIONS	411
13.11.5	EXCEPTIONS	411
13.12	VCVTTS2SIS	412
13.12.1	INSTRUCTION OPERAND ENCODING	412
13.12.2	DESCRIPTION	412
13.12.3	OPERATION	413

13.12.4	SIMD FLOATING-POINT EXCEPTIONS	413
13.12.5	EXCEPTIONS	413
13.13	VCVTTSD2USIS	414
13.13.1	INSTRUCTION OPERAND ENCODING	414
13.13.2	DESCRIPTION	414
13.13.3	OPERATION	415
13.13.4	SIMD FLOATING-POINT EXCEPTIONS	415
13.13.5	EXCEPTIONS	415
13.14	VCVTTSS2SIS	416
13.14.1	INSTRUCTION OPERAND ENCODING	416
13.14.2	DESCRIPTION	416
13.14.3	OPERATION	417
13.14.4	SIMD FLOATING-POINT EXCEPTIONS	417
13.14.5	EXCEPTIONS	417
13.15	VCVTTSS2USIS	418
13.15.1	INSTRUCTION OPERAND ENCODING	418
13.15.2	DESCRIPTION	418
13.15.3	OPERATION	419
13.15.4	SIMD FLOATING-POINT EXCEPTIONS	419
13.15.5	EXCEPTIONS	419
14	INTEL® AVX10 ZERO-EXTENDING PARTIAL VECTOR COPY INSTRUCTIONS	420
14.1	VMOVD	421
14.1.1	INSTRUCTION OPERAND ENCODING	421
14.1.2	DESCRIPTION	421
14.1.3	OPERATION	422
14.1.4	EXCEPTIONS	422
14.2	VMOVW	423
14.2.1	INSTRUCTION OPERAND ENCODING	423
14.2.2	DESCRIPTION	423
14.2.3	OPERATION	424
14.2.4	EXCEPTIONS	424

List of Figures

3.1	Encoding of EVEX instructions	36
5.1	minimum	76
5.2	minimumNumber	76
5.3	minimumMagnitude	77
5.4	minimumMagnitudeNumber	77
5.5	maximum	78

5.6	maximumNumber	78
5.7	maximumMagnitude	79
5.8	maximumMagnitudeNumber	79
5.9	minmax	80

List of Tables

3.1	CPUID Enumeration of Intel® AVX10	32
3.2	Intel® AVX-512 CPUID Feature Flags Included in Intel® AVX10	33
3.3	Feature Differences Between Intel® AVX-512 and Intel® AVX10	34
3.4	XCR0 Feature Bits for Intel® AVX10	35
3.5	EVEX.b and EVEX.u Encoding For Embedded Rounding	37
3.6	FP8 Formats Numeric Definitions	39
3.7	FP8 Downconverts Special Numbers Handling	40
4.1	Exception Class Summary for all Instructions	51
4.2	Type E10NF Class Exception Conditions	53
4.3	Type E11 Class Exception Conditions	55
4.4	Type E2 Class Exception Conditions	57
4.5	Type E3 Class Exception Conditions	59
4.6	Type E3NF Class Exception Conditions	61
4.7	Type E4 Class Exception Conditions	63
4.8	Type E4NF Class Exception Conditions	64
4.9	Type E9NF Class Exception Conditions	65
7.1	VF[,N]M[ADD,SUB][132,213,231]NEPBF16 Notation for Operands	126
7.2	VRCPPBF16 Special Cases	150
7.3	VRSQRTPBF16 Special Cases	157
8.1	Valid Comparison Tests	166
8.2	Valid Comparison Tests	168
8.3	Valid Comparison Tests	170
8.4	Valid Comparison Tests	172
8.5	Valid Comparison Tests	174
8.6	Valid Comparison Tests	176
11.1	MINMAX operation selection according to imm8[4] and imm8[1:0].	216
11.2	MINMAX sign control according to imm8[3:2].	216
11.3	NaN propagation of one or more NaN input values and effect of imm8[3:2] for minimum, minimumMagnitude, maximum, maximumMagnitude MINMAX operations.	217

11.4 NaN propagation of one or more NaN input values and effect of imm8[3:2] for minimumNumber, minimumMagnitudeNumber, maximumNumber, maximumMagnitudeNumber MINMAX operations. 217

11.5 MINMAX Operation Behavior With Signed Comparison of Opposite-Signed Zeros (src1=-0 and src2=+0, or src1=+0 and src2=-0) 217

11.6 MINMAX Operation Behavior With Equal Magnitude Comparisons (src1=a and src2=b, or src1=b and src2=a, where |a|=|b| and a>0 and b<0) 218

Chapter 1

CHANGES

Revision Number	Description	Date
1.0	1. Initial document release	July 11, 2024
2.0	<ol style="list-style-type: none">1. Updated introduction with AVX10 CPUID enumeration clarification on VL128 bit reserved-at-1 behavior2. Updated introduction with AVX10 state management and VMX extensions3. Updated VFPCCLASSPBF16 DAZ behavior (does not consult MXCSR)4. Updated encodings of VCOM*/VUCOM* ops (swizzling of required prefixes, F2 and F3)5. Updated encoding for VGETEXPPBF16 (updating map and required prefixes)	October 1, 2024

Chapter 2

CPUID

This section summarizes the CPUID names and leaf mappings referenced in this document.

CPUID	Allocation
AVX10.2	CPUID.(0x7.0x1).EDX[19] and CPUID.(0x24.0x0).EBX[7:0] >= 2

Chapter 3

INTRODUCTION

3.1 INTEL® AVX10 INTRODUCTION

Intel® Advanced Vector Extensions 10 (Intel® AVX10) represents the first major new vector ISA since the introduction of Intel® Advanced Vector Extensions 512 (Intel® AVX-512) in 2013. This ISA will establish a common, converged vector instruction set across all Intel® architectures, incorporating the modern vectorization aspects of Intel® AVX-512. This ISA will be supported on all future processors, including Performance cores (P-cores) and Efficient cores (E-cores).

The Intel® AVX10 ISA represents the latest in ISA innovations, instructions, and features moving forward. Based on the Intel® AVX-512 ISA feature set and including all Intel® AVX-512 instructions introduced with future Intel® Xeon® processors based on the Granite Rapids microarchitecture, it will support all instruction vector lengths (128, 256, and 512), as well as scalar and opmask instructions. A “converged” version of Intel® AVX10 with vector lengths of at least 256 bits will be supported across all Intel® processors.

3.1.1 FEATURES AND CAPABILITIES

The Intel® AVX10 architecture introduces several features and capabilities beyond the Intel® AVX2 ISA:

- Version-based instruction set enumeration.
- Intel® AVX10/256 – Converged implementation support on all Intel® processors to include all the existing Intel® AVX-512 capabilities such as EVEX encoding, 32 vector registers and eight mask registers at a maximum vector length of 256 bits.
- Embedded rounding and Suppress All Exceptions (SAE) control for YMM (256-bit) versions of the instructions.
- VMX capability to create Intel® AVX10/256 virtual machines that provide a hardware enforced Intel® AVX10/256 execution environment on an Intel® AVX10/512 capable processor.

3.1.2 FEATURE ENUMERATION

Intel® AVX10 introduces a versioned approach for enumeration that is monotonically increasing, inclusive, and supporting all vector lengths. This is introduced to simplify application development by ensuring that all Intel® processors support the same features and instructions at a given Intel® AVX10 version number, as well as reduce the number of CPUID feature flags required to be checked by an application to determine feature support. In this enumeration paradigm, the application developer will only need to check three fields:

1. A CPUID feature bit indicating that the Intel® AVX10 ISA is supported.
2. A version number to ensure that the supported version is greater than or equal to the desired version.
3. A vector length (VL) bit indicating the maximum supported vector length.

The “AVX10 Converged Vector ISA Enable” bit will indicate processor support for the ISA and the presence of an “AVX10 Converged Vector ISA” leaf containing fields for the version number and the supported vector bit lengths. See Table 3.1 for details.

CPUID Bit	Description	Type
CPUID.(EAX=07H, ECX=01H):EDX[bit 19]	If 1, the Intel® AVX10 Converged Vector ISA is supported.	Bit (0/1)
CPUID.(EAX=24H, ECX=00H):EAX[bits 31:0]	Reports the maximum supported sub-leaf.	Integer
CPUID.(EAX=24H, ECX=00H):EBX[bits 7:0]	Reports the Intel® AVX10 Converged Vector ISA version.	Integer (≥ 1)
CPUID.(EAX=24H, ECX=00H):EBX[bits 15:8]	Reserved	N/A
CPUID.(EAX=24H, ECX=00H):EBX[bit 16]	VL128 Support – Reserved at 1, indicates that 128-bit vector support is present	Bit (1)
CPUID.(EAX=24H, ECX=00H):EBX[bit 17]	VL256 Support – If 1, indicates that 256-bit vector support is present.	Bit (0/1)
CPUID.(EAX=24H, ECX=00H):EBX[bit 18]	VL512 Support – If 1, indicates that 512-bit vector support is present.	Bit (0/1)
CPUID.(EAX=24H, ECX=00H):EBX[bits 31:19]	Reserved	N/A
CPUID.(EAX=24H, ECX=00H):ECX[bits 31:0]	Reserved	N/A
CPUID.(EAX=24H, ECX=00H):EDX[bits 31:0]	Reserved	N/A
CPUID.(EAX=24H, ECX=01H):EAX[bits 31:0]	Reserved for discrete feature bits.	N/A
CPUID.(EAX=24H, ECX=01H):EBX[bits 31:0]	Reserved for discrete feature bits.	N/A
CPUID.(EAX=24H, ECX=01H):ECX[bits 31:0]	Reserved for discrete feature bits.	N/A
CPUID.(EAX=24H, ECX=01H):EDX[bits 31:0]	Reserved for discrete feature bits.	N/A

Table 3.1: CPUID Enumeration of Intel® AVX10

The versioned approach to ISA enumeration is expected to adhere to the following rules when incrementing from version N to N+1:

- All contemporary processor families¹ support Intel® AVX10 Version N+1.
- Intel® AVX10 Version N+1 delivers significant value over version N to justify the associated software enabling efforts.

In the rare case of a feature needing to be introduced in-between versions, a discrete CPUID feature bit of

¹Contemporary processor families supporting Intel® AVX10 begin with future Intel® Xeon processors based on Granite Rapids microarchitecture

the form “AVX10-XXXX” may be allocated and enumerated in sub-leaf 1 of CPUID leaf 24H, i.e., CPUID.(EAX=24H, ECX=01H). However, this is expected to be the exception rather than the norm due to the necessity for entrenched legacy support.

Several other important tenets regarding Intel® AVX10 enumeration are as follows:

- Versions are expected to be inclusive such that version N+1 is a superset of version N. Once an instruction is introduced in Intel® AVX10.x, it is expected to be carried forward in all subsequent Intel® AVX10 versions, allowing a developer to check only for a version greater than or equal to the desired version.
- Any processor that enumerates support for Intel® AVX10 will also enumerate support for Intel® AVX and Intel® AVX2.
- Developers can assume that the highest supported vector length for a processor implies that all lesser vector lengths are also supported. Scalar Intel® AVX-512 instructions will be supported independent of the maximum vector width.

The initial, fully-featured version of Intel® AVX10 will be enumerated as Version 2 (denoted as Intel® AVX10.2). This will include the new YMM-form of embedded rounding and Suppress All Exceptions (SAE), the new enumeration scheme, and several new sets of instructions.

An early version of Intel® AVX10 (Version 1, or Intel® AVX10.1) that only enumerates the Intel® AVX-512 instruction set at 128, 256, and 512 bits will be enabled on the Granite Rapids microarchitecture for software pre-enabling. Applications written to Intel® AVX10.1 will run on any future Intel® processor (P-core or E-core) that enumerates Intel® AVX10.1 or higher at the matching desired vector lengths. Intel® AVX-512 instruction families included in Intel® AVX10.1 are shown in Table 3.2.

Feature Introduction	Intel® AVX-512 CPUID Feature Flags Included in Intel® AVX10
Intel® Xeon® Scalable Processor Family based on Skylake microarchitecture	AVX512F, AVX512CD, AVX512BW, AVX512DQ
Intel® Core™ processors based on Cannon Lake microarchitecture	AVX512-VBMI, AVX512-IFMA
2nd generation Intel® Xeon® Scalable Processor Family based on Cascade Lake product	AVX512-VNNI
3rd generation Intel® Xeon® Scalable Processor Family based on Cooper Lake product	AVX512-BF16
3rd generation Intel® Xeon® Scalable Processor Family based on Ice Lake microarchitecture	AVX512-VPOPCNTDQ, AVX512-VBMI2, VAES, GFNI, VPCLMULQDQ, AVX512-BITALG
4th generation Intel® Xeon® Scalable Processor Family based on Sapphire Rapids microarchitecture	AVX512-FP16

Table 3.2: Intel® AVX-512 CPUID Feature Flags Included in Intel® AVX10

Note

VAES, VPCLMULQDQ, and GFNI EVEX instructions will be supported on Intel® AVX10.1 machines but will continue to be enumerated by their existing discrete CPUID feature flags. This requires the developer to check for both the feature and Intel® AVX10, e.g., {AVX10.1 AND VAES}.

Intel® AVX-512 will continue to be supported on P-core-only processors for the foreseeable future to support legacy applications. However, new vector ISA features will only be added to the Intel® AVX10 ISA moving forward. While Intel® AVX10/512 includes all Intel® AVX-512 instructions, it is important to note that applications compiled to Intel® AVX-512 with vector length limited to 256 bits are not guaranteed to be compatible on an Intel® AVX10/256 processor. Intel® will develop tools to enable developers to validate their code prior to deployment.

Feature	Intel® AVX-512	Intel® AVX10.1/256	Intel® AVX10.2/256	Intel® AVX10.1/512	Intel® AVX10.2/512
128-bit vector (XMM) register support	Yes	Yes	Yes	Yes	Yes
256-bit vector (YMM) register support	Yes	Yes	Yes	Yes	Yes
512-bit vector (ZMM) register support	Yes	No	No	Yes	Yes
YMM embedded rounding	No	No	Yes	No	Yes
ZMM embedded rounding	Yes	No	No	Yes	Yes
New AVX10.2 Instructions	No	No	Yes	No	Yes

Table 3.3: Feature Differences Between Intel® AVX-512 and Intel® AVX10

3.1.3 STATE MANAGEMENT

Intel® AVX10 state enumeration in CPUID leaf 0xD and enabling in XCR0 register are identical to Intel® AVX-512 regardless of the maximum vector length supported. The CPUID leaf D enumeration will enumerate the state components and its sizes exactly as Intel® AVX-512 state. Intel® AVX10 state will be enabled in XCR0 register exactly as Intel® AVX-512 state. Table 3.4 highlights the Intel® AVX10 state components and their corresponding XCR0 bits. Please refer to Section 13.1, Volume 1 of Intel Software Developer Manual for the complete definition of these state components.

State Component	XCR0 Index Bit
SSE State	1
AVX State	2
Opmask State	5
ZMM_Hi256 state	6
Hi16_ZMM state	7

Table 3.4: XCR0 Feature Bits for Intel® AVX10

XSAVE*/XRSTOR* instructions on Intel® AVX10/512 processors behave exactly as in Intel® AVX-512 processors.

XSAVE*/XRSTOR* instructions on Intel® AVX10/256 processors operate only on the lower 256 bits of each ZMM register.

- There are no changes to SSE, AVX and Opmask state save/restore and INIT tracking.
- There are no changes to Hi16_ZMM state INIT tracking. However, ZMM_Hi256 state will always be tracked as INIT.
- XRSTOR* when state component's XSTATE_BV header is 1:
 - Restore the lower 256 bit of each ZMM register in Hi16_ZMM state and ignore the contents of the memory corresponding to upper 256 bits.
 - Ignore the contents of the memory corresponding to ZMM_Hi256 state and the corresponding XSTATE_BV bit. ZMM_Hi256 state INIT tracker remains 0.
 - Note that XRSTOR* instructions are allowed to access memory corresponding to the upper 256 bits of ZMM registers.
- XSAVEOPT/C/S when state component is non-INIT or XSAVE
 - Save the lower 256 bits of each ZMM registers in Hi16_ZMM state and zero the memory corresponding to upper 256 bits.
 - Zero the memory corresponding ZMM_Hi256 (This applies only to XSAVE. Other XSAVE* instructions incorporate INIT optimization. Hence when the state component is INIT, they only update the XSTATE_BV to 0 and don't save the state.)

3.1.4 VMX FOR INTEL® AVX10 VIRTUAL MACHINES

Processors that support the Intel® AVX10 instruction set will provide a new processor-based VM execution control bit (tertiary control bit 13) in virtual machine control structure (VMCS) to create Intel® AVX10/256 virtual machines (VMs).

On Intel® AVX10/512 processors, when the execution control bit is set and the guest operating system (OS) enabled AVX10 by setting XCR0[7:5] = 111, the following behavior will be observed during VMX non-root operation.

- Intel® AVX10 vector instructions with vector length (VL) = 512 encoded as EVEX.512 will produce an undefined opcode (#UD) exception to provide #UD compatibility with Intel® AVX10/256 processors.
- XRSTOR* instructions will behave as follows.
 - Restore the lower 256 bit of each ZMM register in Hi16_ZMM state and zero the upper 256 bits of each ZMM register. Ignore the contents of the memory corresponding to upper 256 bits.
 - Zero the upper 256 bits of each ZMM register in Hi16_ZMM state irrespective of the corresponding XSTATE_BV value. Ignore the contents of the memory corresponding to Hi16_ZMM state.
 - Note that XRSTOR* instructions are allowed to access memory corresponding to the upper 256 bits ZMM registers.
- XSAVE* instructions will behave exactly as native Intel® AVX10/256 CPU as described in section 3.1.3

Intel® AVX10/256 processors will also provide the above-mentioned processor-based execution control bit for software compatibility but does not modify any execution behavior. This allows VMs to set execution control bits whenever it creates an Intel® AVX10/256 VM in a way that enables migration of the VM between Intel® AVX10/256 processors and Intel® AVX10/512 processors without the need to modify execution control configuration.

3.1.5 ENCODING EXTENSIONS

As Intel® AVX10 is expected to closely follow the Intel® AVX-512 ISA specification, EVEX encoding changes will be minimal. In the current EVEX encoding specification, the EVEX.b bit encodes multiple functionality which can override the vector length L'L bits when combined with modrm.reg = 0x11 (register-register) to enable static/embedded rounding control. The L'L bits define the rounding mode definition, overriding the MXCSR.RC field while also implying Suppress All Exceptions (SAE). Because the vector length bits (L'L) have been repurposed, the vector length of instructions utilizing embedded rounding is constrained to 512-bits or scalar.

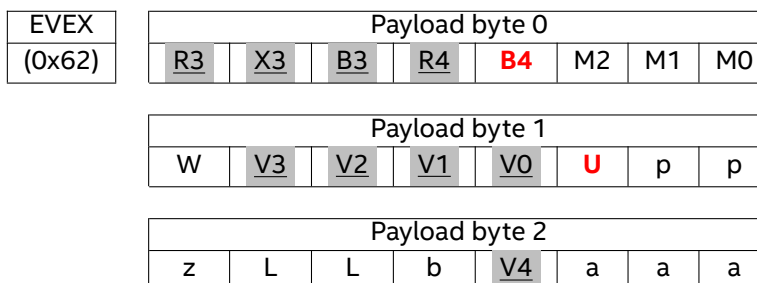


Figure 3.1: Encoding of EVEX instructions

In order to support embedded rounding capabilities for YMM/256 bit Intel® AVX10.2 instructions, the EVEX.U bit is repurposed. It's current assignment is set to always be "1". For Intel® AVX10.2, an EVEX.U bit

value of "0" combined with an EVEX.b value of "1" will enable embedded rounding for supported YMM/256 bit instructions in any Intel® AVX10 ISA version 2 or greater. Embedded rounding for scalars will utilize the same flow as the current design (using the pp bits) to minimize duplication and validation.

EVEX.b	EVEX.U	MODRM.MOD	Semantics
0	1	11	Legacy. Reg-reg VL based on EVELX.LL. No embedded rounding
1	1	11	Legacy. 512 bit reg-reg supporting embedded rounding
1	0	11	New YMM/256 bit reg-reg supporting embedded rounding
0	0	11	Reserved
0/1	0	00/10/01	Reserved

Table 3.5: EVEX.b and EVEX.u Encoding For Embedded Rounding

The following pseudocode will be added for each Intel® AVX10.2 instruction to support embedded rounding at 256 bit and 512 bit vector lengths:

```

If ModR/M=1 and EVEX.b=1 and {VL=512 or {VL=256 and EVEX.u=0}}
    SET_ROUNDING_MODE{EVEX.RC} \\ Embedded rounding control enabled
Else
    SET_ROUNDING_MODE{MXCSR.RC} \\ Rounding control determined from MXCSR
End

```

3.1.6 INTEL® AVX10.2 NEW INSTRUCTIONS

Intel® AVX10 Version 2 (Intel® AVX10.2) includes a suite of new instructions delivering new AI features and performance, accelerated media processing, expanded Web Assembly, and Cryptography support, along with enhancements to existing legacy instructions for completeness and efficiency. All new instructions will be enumerated via the Intel® AVX10.2 feature flags and can be considered to conform to the taxonomy below.

- **AI Datatypes, Conversions, and post-Convolution Instructions:** introduces a suite of AI instructions including FP8 datatypes, convert instructions supporting single, half, and quarter precision (FP32, FP/BF16, E5M2/E4M3 FP8), and post-convolution targeted instructions including arithmetic, scale, min/max, and transcendentals. Note: the FP8 data types are defined as in the Open Compute Project 'OCP 8-bit Floating Point Specification (OFP8)'. The two FP8 (OFP8) formats are E5M2 and E4M3. In instruction mnemonics and pseudo-code, these are denoted throughout this document by BF8 (or bf8) and HF8 (or hf8), respectively.
- **Media Acceleration:** hardware support for codecs through two new media-targeted instructions. VMPSADBW extends the existing MPSADBW instructions to 512 bits, accelerating motion estimation. Also, 16-bit VNNI now supports all sign combinations further accelerating 16-bit video processing.
- **IEEE-754-2019 Minimum and Maximum Support:** introduces min/max instructions supporting NAN behavior as specified in IEEE-754-2019, making it compatible for WebAssembly application development. Also applies to other numeric codes to indicate invalid results.

- **Saturating Conversions:** saturating conversion instructions to support languages such as RUST and WASM. Also applicable for machine learning.
- **Zero-extending Partial Vector Copies:** aligns the zero-extending partial vector copies with existing memory move instructions. The instructions will clear the destination registers irrespective of the load from memory or register.
- **FP Scalar Comparison:** extensions to test all common FP relationships directly. Previously not all flags were capable of being set directly with a single instruction. These new instructions remove previous limitations by more comprehensively setting the appropriate flags.

3.2 FP8 INTRODUCTION

FP8 consists of new datatypes of Floating Point numbers consisting of 8 bits. They are aimed to speedup both training and inference AI workloads. The two new FP8 formats, E5M2 and E4M3, are important optimizations for memory footprint and core AI compute density as well as power and performance efficiency. These formats are introduced in the Open Compute Project 'OCP 8-bit Floating Point Specification (OFP8)' and conversion to and from the two formats will be supported as part of Intel® AVX10.2 ISA.

3.2.1 NUMERIC DEFINITION

Two different FP8 formats are supported: E5M2 FP8 which has 1 sign bit, 5 exponent bits and 2 mantissa bits and E4M3 FP8 which has 1 sign bit, 4 exponent bits and 3 mantissa bits. Due to the very small range and precision of these datatypes, both formats are needed to converge and reach the required accuracy across a wide range of AI topologies. Table 3.6 describes the numerics of each format. While E5M2 follows standard floating point representations, the E4M3 format has a non-standard definition, including the same representation for Infinity and NaN in order to increase its range.

Number	E5M2	E4M3
Exponent Bias	15	7
Max Normal	S.11110.11 = 57344.0 ($1.75 * 2^{15}$)	S.1111.110 = 448.0 ($1.75 * 2^8$)
Min Normal	S.00001.00 = 6.10e-05 (2^{-14})	S.0001.000 = 1.56e-02 (2^{-6})
Max Denormal	S.00000.11 = 4.57e-05 ($0.75 * 2^{-14}$)	S.0000.111 = 1.36e-02 ($0.875 * 2^{-6}$)
Min Denormal	S.00000.01 = 1.52e-05 ($0.25 * 2^{-14}$)	S.0000.001 = 1.95e-03 ($0.125 * 2^{-6}$)
NaNs	S.11111.[01, 10, 11]	S.1111.111
Infinity	S.11111.00	NA

Table 3.6: FP8 Formats Numeric Definitions

3.2.2 FP8 ROUNDING, DENORMALS, SPECIAL NUMBERS, AND EXCEPTIONS

Intel® AVX10.2 dot product instructions, upconverts and downconverts are supported. The down converts have two flavors: *NE* and *BIAS*, which indicate the rounding modes.

- **FP rounding:** Excluding the *BIAS* downconverts, all instructions use RNE (Round to nearest tie to even) rounding mode. The *BIAS* downconverts use RNE in case the input is denormal and truncate (round towards zero) for normal input.
- **Denormal Handling:** All instructions function as if FP exceptions are masked. For any type of input, instructions behave as if MXCSR.DAZ is not set. For FP8 output type, instructions behave as if MXCSR.FTZ is not set. For any other type of output, instructions behave as if MXCSR.FTZ is set.

- Special numbers - Nan/inf/overflow handling: Excluding downconvert instructions, all instructions behave as expected. Infinity is bypassed, NaN is bypassed as QNaN, and Overflow returns Infinity. The downconverts have saturation and non-saturation flavors. Table 3.7 describes the downconverts behavior in these cases:

Flavor	Scenario	E5M2	E4M3
Regular	NaN at input	S.111111.[10, 11]	S.11111.111
	+/- infinity at input	S.111111.00	S.11111.111
	Overflow due to conversion/rounding	S.111111.00	S.11111.111
Saturated	NaN at input	S.111111.[10, 11]	S.11111.111
	+/- infinity at input	S.111111.00	S.11111.111
	Overflow due to conversion/rounding	S.111110.11	S.11111.110

Table 3.7: FP8 Downconverts Special Numbers Handling

- FP exceptions:
 - No instructions consult MXCSR.
 - No instructions raise exceptions.
 - Special rules:
 - * Upconverts and downconverts between different datatypes always update MXCSR.
 - * All other instructions never update MXCSR.

Chapter 4

EXCEPTION CLASSES

4.1 EXCEPTION CLASS INSTRUCTION SUMMARY

Type E10NF			
VCOMSBF16	xmm1, xmm2/m16	AVX10.2	EVEX
Type E11			
VCVTPH2PS	ymm1, xmm2/m128	AVX10.2	EVEX
VCVTPS2PH	xmm1, ymm2, imm8	AVX10.2	EVEX
Type E2			
VADDPD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VADDPH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VADDPS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VCMPPD	k1, ymm2, ymm3/m256, imm8	AVX10.2	EVEX
VCMPPH	k1, ymm2, ymm3/m256, imm8	AVX10.2	EVEX
VCMPPS	k1, ymm2, ymm3/m256, imm8	AVX10.2	EVEX
VCVTDQ2PH	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTDQ2PS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTHF82PH	xmm1, xmm2/m64	AVX10.2	EVEX
VCVTHF82PH	ymm1, xmm2/m128	AVX10.2	EVEX
VCVTHF82PH	zmm1, ymm2/m256	AVX10.2	EVEX
VCVTPD2DQ	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTPD2PH	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTPD2PS	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTPD2QQ	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTPD2UDQ	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTPD2UQQ	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTPH2DQ	ymm1, xmm2/m128	AVX10.2	EVEX
VCVTPH2IBS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTPH2IBS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTPH2IBS	zmm1, zmm2/m512	AVX10.2	EVEX
VCVTPH2IUBS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTPH2IUBS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTPH2IUBS	zmm1, zmm2/m512	AVX10.2	EVEX
VCVTPH2PD	ymm1, xmm2/m64	AVX10.2	EVEX
VCVTPH2PSX	ymm1, xmm2/m128	AVX10.2	EVEX
VCVTPH2QQ	ymm1, xmm2/m64	AVX10.2	EVEX
VCVTPH2UDQ	ymm1, xmm2/m128	AVX10.2	EVEX
VCVTPH2UQQ	ymm1, xmm2/m64	AVX10.2	EVEX
VCVTPH2UW	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTPH2W	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTPS2DQ	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTPS2IBS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTPS2IBS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTPS2IBS	zmm1, zmm2/m512	AVX10.2	EVEX
VCVTPS2IUBS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTPS2IUBS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTPS2IUBS	zmm1, zmm2/m512	AVX10.2	EVEX
VCVTPS2PD	ymm1, xmm2/m128	AVX10.2	EVEX

VCVTPS2PHX	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTPS2QQ	ymm1, xmm2/m128	AVX10.2	EVEX
VCVTPS2UDQ	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTPS2UQQ	ymm1, xmm2/m128	AVX10.2	EVEX
VCVTQQ2PD	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTQQ2PH	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTQQ2PS	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTTPD2DQ	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTTPD2DQS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTTPD2DQS	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTTPD2DQS	ymm1, zmm2/m512	AVX10.2	EVEX
VCVTTPD2QQ	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTPD2QQS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTTPD2QQS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTPD2QQS	zmm1, zmm2/m512	AVX10.2	EVEX
VCVTTPD2UDQ	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTTPD2UDQS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTTPD2UDQS	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTTPD2UDQS	ymm1, zmm2/m512	AVX10.2	EVEX
VCVTTPD2UQQ	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTPD2UQQS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTTPD2UQQS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTPD2UQQS	zmm1, zmm2/m512	AVX10.2	EVEX
VCVTTPH2DQ	ymm1, xmm2/m128	AVX10.2	EVEX
VCVTTPH2IBS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTTPH2IBS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTPH2IBS	zmm1, zmm2/m512	AVX10.2	EVEX
VCVTTPH2IUBS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTTPH2IUBS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTPH2IUBS	zmm1, zmm2/m512	AVX10.2	EVEX
VCVTTPH2QQ	ymm1, xmm2/m64	AVX10.2	EVEX
VCVTTPH2UDQ	ymm1, xmm2/m128	AVX10.2	EVEX
VCVTTPH2UQQ	ymm1, xmm2/m64	AVX10.2	EVEX
VCVTTPH2UW	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTPH2W	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTPS2DQ	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTPS2DQS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTTPS2DQS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTPS2DQS	zmm1, zmm2/m512	AVX10.2	EVEX
VCVTTPS2IBS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTTPS2IBS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTPS2IBS	zmm1, zmm2/m512	AVX10.2	EVEX
VCVTTPS2IUBS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTTPS2IUBS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTPS2IUBS	zmm1, zmm2/m512	AVX10.2	EVEX
VCVTTPS2QQ	ymm1, xmm2/m128	AVX10.2	EVEX
VCVTTPS2QQS	xmm1, xmm2/m64	AVX10.2	EVEX

VCVTTPS2QQS	ymm1, xmm2/m128	AVX10.2	EVEX
VCVTTPS2QQS	zmm1, ymm2/m256	AVX10.2	EVEX
VCVTTPS2UDQ	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTPS2UDQS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTTPS2UDQS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTPS2UDQS	zmm1, zmm2/m512	AVX10.2	EVEX
VCVTTPS2UQQ	ymm1, xmm2/m128	AVX10.2	EVEX
VCVTTPS2UQQS	xmm1, xmm2/m64	AVX10.2	EVEX
VCVTTPS2UQQS	ymm1, xmm2/m128	AVX10.2	EVEX
VCVTTPS2UQQS	zmm1, ymm2/m256	AVX10.2	EVEX
VCVTUDQ2PH	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTUDQ2PS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTUQQ2PD	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTUQQ2PH	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTUQQ2PS	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTUW2PH	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTW2PH	ymm1, ymm2/m256	AVX10.2	EVEX
VDIVPD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VDIVPH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VDIVPS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFIXUPIMMPD	ymm1, ymm2, ymm3/m256, imm8	AVX10.2	EVEX
VFIXUPIMMPS	ymm1, ymm2, ymm3/m256, imm8	AVX10.2	EVEX
VFMADD132PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADD132PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADD132PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADD213PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADD213PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADD213PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADD231PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADD231PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADD231PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADDSUB132PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADDSUB132PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADDSUB132PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADDSUB213PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADDSUB213PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADDSUB213PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADDSUB231PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADDSUB231PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADDSUB231PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUB132PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUB132PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUB132PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUB213PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUB213PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUB213PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUB231PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX

VFMSUB231PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUB231PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUBADD132PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUBADD132PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUBADD132PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUBADD213PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUBADD213PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUBADD213PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUBADD231PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUBADD231PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUBADD231PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMADD132PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMADD132PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMADD132PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMADD213PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMADD213PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMADD213PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMADD231PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMADD231PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMADD231PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMSUB132PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMSUB132PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMSUB132PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMSUB213PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMSUB213PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMSUB213PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMSUB231PD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMSUB231PH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMSUB231PS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VGETEXPPD	ymm1, ymm2/m256	AVX10.2	EVEX
VGETEXPPH	ymm1, ymm2/m256	AVX10.2	EVEX
VGETEXPPS	ymm1, ymm2/m256	AVX10.2	EVEX
VGETMANTPD	ymm1, ymm2/m256, imm8	AVX10.2	EVEX
VGETMANTPH	ymm1, ymm2/m256, imm8	AVX10.2	EVEX
VGETMANTPS	ymm1, ymm2/m256, imm8	AVX10.2	EVEX
VMAXPD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VMAXPH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VMAXPS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VMINMAXPD	xmm1, xmm2, xmm3/m128, imm8	AVX10.2	EVEX
VMINMAXPH	ymm1, ymm2, ymm3/m256, imm8	AVX10.2	EVEX
VMINMAXPD	zmm1, zmm2, zmm3/m512, imm8	AVX10.2	EVEX
VMINMAXPH	xmm1, xmm2, xmm3/m128, imm8	AVX10.2	EVEX
VMINMAXPH	ymm1, ymm2, ymm3/m256, imm8	AVX10.2	EVEX
VMINMAXPH	zmm1, zmm2, zmm3/m512, imm8	AVX10.2	EVEX
VMINMAXPS	xmm1, xmm2, xmm3/m128, imm8	AVX10.2	EVEX
VMINMAXPS	ymm1, ymm2, ymm3/m256, imm8	AVX10.2	EVEX
VMINMAXPS	zmm1, zmm2, zmm3/m512, imm8	AVX10.2	EVEX

VMINPD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VMINPH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VMINPS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VMULPD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VMULPH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VMULPS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VRANGEPD	ymm1, ymm2, ymm3/m256, imm8	AVX10.2	EVEX
VRANGEPS	ymm1, ymm2, ymm3/m256, imm8	AVX10.2	EVEX
VREDUCEPD	ymm1, ymm2/m256, imm8	AVX10.2	EVEX
VREDUCEPH	ymm1, ymm2/m256, imm8	AVX10.2	EVEX
VREDUCEPS	ymm1, ymm2/m256, imm8	AVX10.2	EVEX
VRNDSCALEPD	ymm1, ymm2/m256, imm8	AVX10.2	EVEX
VRNDSCALEPH	ymm1, ymm2/m256, imm8	AVX10.2	EVEX
VRNDSCALEPS	ymm1, ymm2/m256, imm8	AVX10.2	EVEX
VSCALEFPD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VSCALEFPH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VSCALEFPS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VSQRTPD	ymm1, ymm2/m256	AVX10.2	EVEX
VSQRTPH	ymm1, ymm2/m256	AVX10.2	EVEX
VSQRTPS	ymm1, ymm2/m256	AVX10.2	EVEX
VSUBPD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VSUBPH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VSUBPS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
Type E3			
VMINMAXSD	xmm1, xmm2, xmm3/m64, imm8	AVX10.2	EVEX
VMINMAXSH	xmm1, xmm2, xmm3/m16, imm8	AVX10.2	EVEX
VMINMAXSS	xmm1, xmm2, xmm3/m32, imm8	AVX10.2	EVEX
Type E3NF			
VCOMXSD	xmm1, xmm2/m64	AVX10.2	EVEX
VCOMXSH	xmm1, xmm2/m16	AVX10.2	EVEX
VCOMXSS	xmm1, xmm2/m32	AVX10.2	EVEX
VCVTSD2SIS	r32, xmm1/m64	AVX10.2	EVEX
VCVTSD2SIS	r64, xmm1/m64	AVX10.2	EVEX
VCVTSD2USIS	r32, xmm1/m64	AVX10.2	EVEX
VCVTSD2USIS	r64, xmm1/m64	AVX10.2	EVEX
VCVTSS2SIS	r32, xmm1/m32	AVX10.2	EVEX
VCVTSS2SIS	r64, xmm1/m32	AVX10.2	EVEX
VCVTSS2USIS	r32, xmm1/m32	AVX10.2	EVEX
VCVTSS2USIS	r64, xmm1/m32	AVX10.2	EVEX
VUCOMXSD	xmm1, xmm2/m64	AVX10.2	EVEX
VUCOMXSH	xmm1, xmm2/m16	AVX10.2	EVEX
VUCOMXSS	xmm1, xmm2/m32	AVX10.2	EVEX
Type E4			
VADDNEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VADDNEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VADDNEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VCMPPBF16	k1, xmm2, xmm3/m128, imm8	AVX10.2	EVEX

VCMPPBF16	k1, ymm2, ymm3/m256, imm8	AVX10.2	EVEX
VCMPPBF16	k1, zmm2, zmm3/m512, imm8	AVX10.2	EVEX
VCVTNEBF162IBS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTNEBF162IBS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTNEBF162IBS	zmm1, zmm2/m512	AVX10.2	EVEX
VCVTNEBF162IUBS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTNEBF162IUBS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTNEBF162IUBS	zmm1, zmm2/m512	AVX10.2	EVEX
VCVTTNEBF162IBS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTTNEBF162IBS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTNEBF162IBS	zmm1, zmm2/m512	AVX10.2	EVEX
VCVTTNEBF162IUBS	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTTNEBF162IUBS	ymm1, ymm2/m256	AVX10.2	EVEX
VCVTTNEBF162IUBS	zmm1, zmm2/m512	AVX10.2	EVEX
VDIVNEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VDIVNEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VDIVNEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VDPPHPS	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VDPPHPS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VDPPHPS	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VFCMADDCPH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFCMULCPH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADD132NEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VFMADD132NEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADD132NEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VFMADD213NEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VFMADD213NEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADD213NEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VFMADD231NEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VFMADD231NEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADD231NEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VFMADDCPH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUB132NEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VFMSUB132NEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUB132NEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VFMSUB213NEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VFMSUB213NEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUB213NEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VFMSUB231NEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VFMSUB231NEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMSUB231NEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VFMULCPH	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMADD132NEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VFNMADD132NEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMADD132NEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VFNMADD213NEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VFNMADD213NEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX

VFMADD213NEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VFMADD231NEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VFMADD231NEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFMADD231NEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VFNMSUB132NEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VFNMSUB132NEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMSUB132NEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VFNMSUB213NEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VFNMSUB213NEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMSUB213NEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VFNMSUB231NEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VFNMSUB231NEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VFNMSUB231NEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VFPCLASSPBF16	k1, xmm2/m128, imm8	AVX10.2	EVEX
VFPCLASSPBF16	k1, ymm2/m256, imm8	AVX10.2	EVEX
VFPCLASSPBF16	k1, zmm2/m512, imm8	AVX10.2	EVEX
VGETEXPPBF16	xmm1, xmm2/m128	AVX10.2	EVEX
VGETEXPPBF16	ymm1, ymm2/m256	AVX10.2	EVEX
VGETEXPPBF16	zmm1, zmm2/m512	AVX10.2	EVEX
VGETMANTPBF16	xmm1, xmm2/m128, imm8	AVX10.2	EVEX
VGETMANTPBF16	ymm1, ymm2/m256, imm8	AVX10.2	EVEX
VGETMANTPBF16	zmm1, zmm2/m512, imm8	AVX10.2	EVEX
VMAXPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VMAXPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VMAXPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VMINMAXNEPBF16	xmm1, xmm2, xmm3/m128, imm8	AVX10.2	EVEX
VMINMAXNEPBF16	ymm1, ymm2, ymm3/m256, imm8	AVX10.2	EVEX
VMINMAXNEPBF16	zmm1, zmm2, zmm3/m512, imm8	AVX10.2	EVEX
VMINPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VMINPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VMINPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VMULNEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VMULNEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VMULNEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VPDPBSSD	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VPDPBSSD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VPDPBSSD	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VPDPBSSDS	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VPDPBSSDS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VPDPBSSDS	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VPDPBSUD	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VPDPBSUD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VPDPBSUD	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VPDPBSUDS	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VPDPBSUDS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VPDPBSUDS	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VPDPBUUD	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX

VPDPBUUD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VPDPBUUD	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VPDPBUUDS	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VPDPBUUDS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VPDPBUUDS	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VPDPWSUD	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VPDPWSUD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VPDPWSUD	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VPDPWSUDS	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VPDPWSUDS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VPDPWSUDS	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VPDPWUSD	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VPDPWUSD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VPDPWUSD	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VPDPWUSDS	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VPDPWUSDS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VPDPWUSDS	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VPDPWUUD	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VPDPWUUD	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VPDPWUUD	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VPDPWUUDS	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VPDPWUUDS	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VPDPWUUDS	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VRCPPBF16	xmm1, xmm2/m128	AVX10.2	EVEX
VRCPPBF16	ymm1, ymm2/m256	AVX10.2	EVEX
VRCPPBF16	zmm1, zmm2/m512	AVX10.2	EVEX
VREDUCENEPBF16	xmm1, xmm2/m128, imm8	AVX10.2	EVEX
VREDUCENEPBF16	ymm1, ymm2/m256, imm8	AVX10.2	EVEX
VREDUCENEPBF16	zmm1, zmm2/m512, imm8	AVX10.2	EVEX
VRNDSCALENEPBF16	xmm1, xmm2/m128, imm8	AVX10.2	EVEX
VRNDSCALENEPBF16	ymm1, ymm2/m256, imm8	AVX10.2	EVEX
VRNDSCALENEPBF16	zmm1, zmm2/m512, imm8	AVX10.2	EVEX
VRSQRTPBF16	xmm1, xmm2/m128	AVX10.2	EVEX
VRSQRTPBF16	ymm1, ymm2/m256	AVX10.2	EVEX
VRSQRTPBF16	zmm1, zmm2/m512	AVX10.2	EVEX
VSCALEFPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VSCALEFPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VSCALEFPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VSQRTNEPBF16	xmm1, xmm2/m128	AVX10.2	EVEX
VSQRTNEPBF16	ymm1, ymm2/m256	AVX10.2	EVEX
VSQRTNEPBF16	zmm1, zmm2/m512	AVX10.2	EVEX
VSUBNEPBF16	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VSUBNEPBF16	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VSUBNEPBF16	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
Type E4NF			
VCVT2PS2PHX	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VCVT2PS2PHX	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX

VCVT2PS2PHX	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VCVTBIASPH2BF8	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VCVTBIASPH2BF8	xmm1, ymm2, ymm3/m256	AVX10.2	EVEX
VCVTBIASPH2BF8	ymm1, zmm2, zmm3/m512	AVX10.2	EVEX
VCVTBIASPH2BF8S	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VCVTBIASPH2BF8S	xmm1, ymm2, ymm3/m256	AVX10.2	EVEX
VCVTBIASPH2BF8S	ymm1, zmm2, zmm3/m512	AVX10.2	EVEX
VCVTBIASPH2HF8	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VCVTBIASPH2HF8	xmm1, ymm2, ymm3/m256	AVX10.2	EVEX
VCVTBIASPH2HF8	ymm1, zmm2, zmm3/m512	AVX10.2	EVEX
VCVTBIASPH2HF8S	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VCVTBIASPH2HF8S	xmm1, ymm2, ymm3/m256	AVX10.2	EVEX
VCVTBIASPH2HF8S	ymm1, zmm2, zmm3/m512	AVX10.2	EVEX
VCVTNE2PH2BF8	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VCVTNE2PH2BF8	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VCVTNE2PH2BF8	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VCVTNE2PH2BF8S	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VCVTNE2PH2BF8S	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VCVTNE2PH2BF8S	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VCVTNE2PH2HF8	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VCVTNE2PH2HF8	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VCVTNE2PH2HF8	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VCVTNE2PH2HF8S	xmm1, xmm2, xmm3/m128	AVX10.2	EVEX
VCVTNE2PH2HF8S	ymm1, ymm2, ymm3/m256	AVX10.2	EVEX
VCVTNE2PH2HF8S	zmm1, zmm2, zmm3/m512	AVX10.2	EVEX
VCVTNEPH2BF8	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTNEPH2BF8	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTNEPH2BF8	ymm1, zmm2/m512	AVX10.2	EVEX
VCVTNEPH2BF8S	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTNEPH2BF8S	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTNEPH2BF8S	ymm1, zmm2/m512	AVX10.2	EVEX
VCVTNEPH2HF8	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTNEPH2HF8	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTNEPH2HF8	ymm1, zmm2/m512	AVX10.2	EVEX
VCVTNEPH2HF8S	xmm1, xmm2/m128	AVX10.2	EVEX
VCVTNEPH2HF8S	xmm1, ymm2/m256	AVX10.2	EVEX
VCVTNEPH2HF8S	ymm1, zmm2/m512	AVX10.2	EVEX
VMPSADBW	xmm1, xmm2, xmm3/m128, imm8	AVX10.2	EVEX
VMPSADBW	ymm1, ymm2, ymm3/m256, imm8	AVX10.2	EVEX
VMPSADBW	zmm1, zmm2, zmm3/m512, imm8	AVX10.2	EVEX
Type E9NF			
VMOVD	xmm1, xmm2/m32	AVX10.2	EVEX
VMOVD	xmm1/m32, xmm2	AVX10.2	EVEX
VMOVW	xmm1, xmm2/m16	AVX10.2	EVEX
VMOVW	xmm1/m16, xmm2	AVX10.2	EVEX

Table 4.1: Exception Class Summary for all Instructions

4.2 EXCEPTION CLASS SUMMARY

The following descriptions contain tabular summaries of the behavior of each exception class across the operating modes of Intel® processors.

Notations and abbreviations include:

- CM: Compatibility Mode
- PM: Protected Mode

4.2.1 EXCEPTION CLASS E10NF

EVEX-encoded scalar instructions that ignore EVEX.L'L vector length encoding, that cause no SIMD FP exceptions, and do not support memory fault suppression follow exception class E10NF.

Exception	Real	Virtual 8086	PM & CM	64-bit	Cause of exception
Invalid Opcode, #UD	X	X			EVEX prefix is present.
			X	X	If CR4.OSXSAVE[bit 18] = 0. If any of the following conditions applies <ul style="list-style-type: none"> • State requirements not met. • Opcode independent #UD conditions not met. • Operand encoding #UD conditions not met. • Opmask encoding #UD conditions not met. • EVEX.b encoding #UD conditions not met and in E4.nb sub-class.
	X	X	X	X	If preceded by a LOCK prefix (F0H)
			X	X	If any REX, F2, F3 or 66 prefixes precede a VEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, #SS(0)			X		For an illegal address in the SS segment.
				X	If a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		For an illegal memory operand effective address in the CS, DS, ES, FS or GS segments. If the DS, ES, FS, or GS register is used to access memory and it contains a NULL segment selector.
				X	If the memory address is in a non-canonical form.
	X	X			If any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault, #PF(faultcode)		X	X	X	For a page fault.
Alignment Check, #AC(0)		X	X	X	For 2, 4, or 8 byte memory access if alignment checking is enabled and an unaligned memory access is made while CPL=3

Table 4.2: Type E10NF Class Exception Conditions

4.2.2 EXCEPTION CLASS E11

EVEX-encoded instructions that can cause SIMD FP exception, memory operand fault suppression, but do not cause #AC follow exception class E11.

Exception	Real	Virtual 8086	PM & CM	64-bit	Cause of exception
Invalid Opcode, #UD	X	X			EVEX prefix is present.
			X	X	If CR4.OSXSAVE[bit 18] = 0. If any of the following conditions applies <ul style="list-style-type: none"> • State requirements not met. • Opcode independent #UD conditions not met. • Operand encoding #UD conditions not met. • Opmask encoding #UD conditions not met. • EVEX.b encoding #UD conditions not met. • Instruction specific EVEX.L'L restriction not met.
	X	X	X	X	If preceded by a LOCK prefix (F0H)
			X	X	If any REX, F2, F3 or 66 prefixes precede a VEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, #SS(0)			X		If fault suppression not set, and an illegal address in the SS segment.
				X	If fault suppression not set, and a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		If fault suppression not set, and an illegal memory operand effective address in the CS, DS, ES, FS or GS segments. If fault suppression not set, and the DS, ES, FS, or GS register is used to access memory and it contains a NULL segment selector.
				X	If fault suppression not set, and the memory address is in a non-canonical form.
	X	X			If fault suppression not set, and any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault, #PF(faultcode)		X	X	X	If fault suppression not set, and a page fault.
SIMD Floating-point Exception, #XM	X	X	X	X	If an unmasked SIMD floating-point exception, sae not set, and CR4.OSXMMEXCPT[bit 10] = 1

Table 4.3: Type E11 Class Exception Conditions

4.2.3 EXCEPTION CLASS E2

EVEX-encoded vector instructions with arithmetic semantics follow exception class E2.

Exception	Real	Virtual 8086	PM & CM	64-bit	Cause of exception
Invalid Opcode, #UD	X	X			EVEX prefix is present.
	X	X	X	X	If an unmasked SIMD floating-point exception and CR4.OSXMMEXCPT[bit 10] = 0
			X	X	If CR4.OSXSAVE[bit 18] = 0. If any of the following conditions applies <ul style="list-style-type: none"> • State requirements not met. • Opcode independent #UD conditions not met. • Operand encoding #UD conditions not met. • Opmask encoding #UD conditions not met. • EVEX.b encoding #UD conditions not met. • Instruction specific EVEX.L'L restriction not met.
	X	X	X	X	If preceded by a LOCK prefix (F0H)
			X	X	If any REX, F2, F3 or 66 prefixes precede a VEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, #SS(0)			X		If fault suppression not set, and an illegal address in the SS segment.
				X	If fault suppression not set, and a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		If fault suppression not set, and an illegal memory operand effective address in the CS, DS, ES, FS or GS segments. If fault suppression not set, and the DS, ES, FS, or GS register is used to access memory and it contains a NULL segment selector.
				X	If fault suppression not set, and the memory address is in a non-canonical form.
	X	X			If fault suppression not set, and any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault, #PF(faultcode)		X	X	X	If fault suppression not set, and a page fault.
Alignment Check, #AC(0)		X	X	X	For 2, 4, or 8 byte memory access if alignment checking is enabled and an unaligned memory access is made while CPL=3
SIMD Floating-point Exception, #XM	X	X	X	X	If an unmasked SIMD floating-point exception and CR4.OSXMMEXCPT[bit 10] = 1

Table 4.4: Type E2 Class Exception Conditions

4.2.4 EXCEPTION CLASS E3

EVEX-encoded scalar instructions with arithmetic semantics that support memory fault suppression follow exception class E3.

Exception	Real	Virtual 8086	PM & CM	64-bit	Cause of exception
Invalid Opcode, #UD	X	X			EVEX prefix is present.
	X	X	X	X	If an unmasked SIMD floating-point exception and CR4.OSXMMEXCPT[bit 10] = 0
			X	X	If CR4.OSXSAVE[bit 18] = 0. If any of the following conditions applies <ul style="list-style-type: none"> • State requirements not met. • Opcode independent #UD conditions not met. • Operand encoding #UD conditions not met. • Opmask encoding #UD conditions not met. • EVEX.b encoding #UD conditions not met.
	X	X	X	X	If preceded by a LOCK prefix (F0H)
			X	X	If any REX, F2, F3 or 66 prefixes precede a VEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, #SS(0)			X		If fault suppression not set, and an illegal address in the SS segment.
				X	If fault suppression not set, and a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		If fault suppression not set, and an illegal memory operand effective address in the CS, DS, ES, FS or GS segments. If fault suppression not set, and the DS, ES, FS, or GS register is used to access memory and it contains a NULL segment selector.
				X	If fault suppression not set, and the memory address is in a non-canonical form.
	X	X			If fault suppression not set, and any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault, #PF(faultcode)		X	X	X	If fault suppression not set, and a page fault.
Alignment Check, #AC(0)		X	X	X	For 2, 4, or 8 byte memory access if alignment checking is enabled and an unaligned memory access is made while CPL=3
SIMD Floating-point Exception, #XM	X	X	X	X	If an unmasked SIMD floating-point exception, sae or er not set, and CR4.OSXMMEXCPT[bit 10] = 1

Table 4.5: Type E3 Class Exception Conditions

4.2.5 EXCEPTION CLASS E3NF

EVEX-encoded scalar instructions with arithmetic semantics that do not support memory fault suppression follow exception class E3NF.

Exception	Real	Virtual 8086	PM & CM	64-bit	Cause of exception
Invalid Opcode, #UD	X	X			EVEX prefix is present.
	X	X	X	X	If an unmasked SIMD floating-point exception and CR4.OSXMMEXCPT[bit 10] = 0
			X	X	If CR4.OSXSAVE[bit 18] = 0. If any of the following conditions applies <ul style="list-style-type: none"> • State requirements not met. • Opcode independent #UD conditions not met. • Operand encoding #UD conditions not met. • Opmask encoding #UD conditions not met. • EVEX.b encoding #UD conditions not met.
	X	X	X	X	If preceded by a LOCK prefix (F0H)
			X	X	If any REX, F2, F3 or 66 prefixes precede a VEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, #SS(0)			X		For an illegal address in the SS segment.
				X	If a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		For an illegal memory operand effective address in the CS, DS, ES, FS or GS segments. If the DS, ES, FS, or GS register is used to access memory and it contains a NULL segment selector.
				X	If the memory address is in a non-canonical form.
	X	X			If any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault, #PF(faultcode)		X	X	X	For a page fault.
Alignment Check, #AC(0)		X	X	X	For 2, 4, or 8 byte memory access if alignment checking is enabled and an unaligned memory access is made while CPL=3
SIMD Floating-point Exception, #XM	X	X	X	X	If an unmasked SIMD floating-point exception, sae or er not set, and CR4.OSXMMEXCPT[bit 10] = 1

Table 4.6: Type E3NF Class Exception Conditions

4.2.6 EXCEPTION CLASS E4

EVEX-encoded scalar instructions that cause no SIMD FP exceptions, and that support memory fault suppression follow exception class E4.

Exception	Real	Virtual 8086	PM & CM	64-bit	Cause of exception
Invalid Opcode, #UD	X	X			EVEX prefix is present.
			X	X	If CR4.OSXSAVE[bit 18] = 0. If any of the following conditions applies <ul style="list-style-type: none"> • State requirements not met. • Opcode independent #UD conditions not met. • Operand encoding #UD conditions not met. • Opmask encoding #UD conditions not met. • EVEX.b encoding #UD conditions not met and in E4.nb subclass. • Instruction specific EVEX.L'L restriction not met
	X	X	X	X	If preceded by a LOCK prefix (F0H)
			X	X	If any REX, F2, F3 or 66 prefixes precede a VEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, #SS(0)			X		If fault suppression not set, and an illegal address in the SS segment.
				X	If fault suppression not set, and a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		If fault suppression not set, and an illegal memory operand effective address in the CS, DS, ES, FS or GS segments. If fault suppression not set, and the DS, ES, FS, or GS register is used to access memory and it contains a NULL segment selector.
				X	If fault suppression not set, and the memory address is in a non-canonical form.
	X	X			If fault suppression not set, and any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault, #PF(faultcode)		X	X	X	If fault suppression not set, and a page fault.
Alignment Check, #AC(0)		X	X	X	For 2, 4, or 8 byte memory access if alignment checking is enabled and an unaligned memory access is made while CPL=3

Table 4.7: Type E4 Class Exception Conditions

4.2.7 EXCEPTION CLASS E4NF

EVEX-encoded scalar instructions that cause no SIMD FP exceptions, and that do not support memory fault suppression follow exception class E4NF.

Exception	Real	Virtual 8086	PM & CM	64-bit	Cause of exception
Invalid Opcode, #UD	X	X			EVEX prefix is present.
			X	X	If CR4.OSXSAVE[bit 18] = 0. If any of the following conditions applies <ul style="list-style-type: none"> State requirements not met. Opcode independent #UD conditions not met. Operand encoding #UD conditions not met. Opmask encoding #UD conditions not met. EVEX.b encoding #UD conditions not met and in E4.nb sub-class. Instruction specific EVEX.L'L restriction not met
	X	X	X	X	If preceded by a LOCK prefix (F0H)
			X	X	If any REX, F2, F3 or 66 prefixes precede a VEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CR0.TS[bit 3]=1.
Stack, #SS(0)			X		For an illegal address in the SS segment.
				X	If a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		For an illegal memory operand effective address in the CS, DS, ES, FS or GS segments. If the DS, ES, FS, or GS register is used to access memory and it contains a NULL segment selector.
				X	If the memory address is in a non-canonical form.
	X	X			If any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault, #PF(faultcode)		X	X	X	For a page fault.
Alignment Check, #AC(0)		X	X	X	For 2, 4, or 8 byte memory access if alignment checking is enabled and an unaligned memory access is made while CPL=3

Table 4.8: Type E4NF Class Exception Conditions

4.2.8 EXCEPTION CLASS E9NF

EVEX-encoded vector or partial-vector instructions that cause no SIMD FP exceptions, and do not support memory fault suppression follow exception class E9NF.

Exception	Real	Virtual 8086	PM & CM	64-bit	Cause of exception
Invalid Opcode, #UD	X	X			EVEX prefix is present.
			X	X	If CR4.OSXSAVE[bit 18] = 0. If any of the following conditions applies <ul style="list-style-type: none"> • State requirements not met. • Opcode independent #UD conditions not met. • Operand encoding #UD conditions not met. • Opmask encoding #UD conditions not met. • EVEX.b encoding #UD conditions not met and in E4.nb sub-class. • Instruction specific EVEX.L'L restriction not met
	X	X	X	X	If preceded by a LOCK prefix (F0H)
			X	X	If any REX, F2, F3 or 66 prefixes precede a VEX prefix.
	X	X	X	X	If any corresponding CPUID feature flag is '0'.
Device Not Available, #NM	X	X	X	X	If CRO.TS[bit 3]=1.
Stack, #SS(0)			X		For an illegal address in the SS segment.
				X	If a memory address referencing the SS segment is in a non-canonical form.
General Protection, #GP(0)			X		For an illegal memory operand effective address in the CS, DS, ES, FS or GS segments. If the DS, ES, FS, or GS register is used to access memory and it contains a NULL segment selector.
				X	If the memory address is in a non-canonical form.
	X	X			If any part of the operand lies outside the effective address space from 0 to FFFFH.
Page Fault, #PF(faultcode)		X	X	X	For a page fault.
Alignment Check, #AC(0)		X	X	X	For 2, 4, or 8 byte memory access if alignment checking is enabled and an unaligned memory access is made while CPL=3

Table 4.9: Type E9NF Class Exception Conditions

Chapter 5

HELPER FUNCTIONS

```
1 define convert_hf8_to_fp16( in ):
2     fp16_bias = 15
3     hf8_bias = 7
4     s = ( in & 0x80 ) << 8
5     e = ( in & 0x78 ) >> 3
6     m = ( in & 0x07 )
7     e_norm = e + (fp16_bias - hf8_bias)
8
9     /* convert denormal hf8 number into a normal fp16 number */
10    if ( (e == 0) && (m != 0) ):
11        lz_cnt = 2
12        lz_cnt = ( m > 0x1 ) ? 1 : lz_cnt
13        lz_cnt = ( m > 0x3 ) ? 0 : lz_cnt
14        e_norm -= lz_cnt
15        m = (m << (lz_cnt+1)) & 0x07
16    else if ( (e == 0) && (m == 0) ):
17        e_norm = 0
18    else if ( (e == 0xf) && (m == 0x7) ):
19        e_norm = 0x1f
20
21    /* set result */
22    res = 0x0
23    res |= (e_norm << 10)
24    res |= (m << 7)
25    res |= s
26
27    return res
```

```

1  define convert_fp16_to_bf8(x,s):
2  // The x parameter is the data
3  // The s parameter indicates whether we saturate in case of
4  // overflow due to conversion or rounding
5
6  if *x is infinity*
7      if (s==0x1):                // Max Value
8          dest[7] := x[15]
9          dest[6:0] := 0x7B
10         else:                    // INF
11             dest[7:0] := x[15:8]
12
13     else if *x is nan*:
14         dest[7:0] := x[15:8]      // truncate and set QNaN
15         dest[1] := 1
16
17     else // normal, zero or denormal number apply RNE
18         lsb := x[8]
19         rounding_bias[15:0] := 0x007F + lsb
20         temp[15:0] := x[15:0] + rounding_bias[15:0]      // temp is rounded data (RNE)
21
22         if temp [14:8] == 0x7C && s==1 // saturating to E5M2_MAX due to infinity result
23             dest[7] := temp[15]
24             dest[6:0] := 0x7B
25         else
26             dest[7:0] := temp[15:8]
27     return dest
28
29

```

```

1  define convert_fp16_to_hf8(x, s):
2  // The x parameter is the data
3  // The s parameter indicates whether we saturate in case of
4  // overflow due to conversion or rounding
5  // Round mode RNE
6  fp16_bias := 15
7  hf8_bias := 7
8  fp16_to_hf8_exp_rebias = fp16_bias - hf8_bias
9
10
11  sign := ( x & 0x8000 ) >> 8
12  e_fp16 := ( x & 0x7c00 ) >> 10
13  m_fp16 := ( x & 0x03ff )
14
15  if *x is infinity*
16      e := 0xF

```

```

17         if (s==0x1):                // Max Value
18             m := 0x6
19         else:                        // NaN
20             m := 0x7
21
22     else if *x is nan*:
23         e := 0xF
24         m := 0x7
25
26     /* overflow --> make it NaN or Saturate to E4M3_MAX*/
27     else if ((e_fp16 > (fp16_to_hf8_exp_rebias + 15)) ||
28             ((e_fp16 == (fp16_to_hf8_exp_rebias + 15)) && (m_fp16 > 0x0340))):
29         e := 0xf
30         if ( s == 0x1 ):
31             m := 0x6
32         else:
33             m := 0x7
34
35     /* Zero */
36     else if ((e_fp16 ==0) & (m_fp16==0)):
37         e := 0
38         m := 0
39
40     /* underflow */
41     else if ( e_fp16 <= fp16_to_hf8_exp_rebias ):
42     /* denormalized mantissa */
43         /* set Jbit */
44         m := m_fp16 | 0x0400
45
46         /* additionally subnormal shift */
47         m = m >> ((fp16_to_hf8_exp_rebias) + 1 - e_fp16)
48
49         /* preserve sticky bit (some sticky bits are lost when denormalizing) */
50         ShiftOutSticky = (((m_fp16 & 0x007f) + 0x007f) >> 7)
51         m = m | ShiftOutSticky           // OR m.lsb with sticky
52
53         /* RNE Round */
54         fixup := (m >> 7) & 0x1;        // get Lbit
55         m := m + 0x003f + fixup;        // RNE
56
57         // in case of round overflow, m>>10 declares the carry into exponent
58         e := m>>10
59
60         m := (m >> 7) & 0x7 // Truncate Round and ignore carry
61
62
63     /* normal */
64     else:

```

```

65     /* RNE round */
66     fixup := (m_fp16 >> 7) & 0x1;
67     RneX := x + 0x003f + fixup;
68     e := ((RneX & 0x7c00) >> 10);
69     m := ( RneX & 0x03ff );
70     e = e - (fp16_to_hf8_exp_rebias);
71     m := m >> 7;
72
73 res = 0x0
74 res |= e << 3 // set exp
75 res |= m      //set mant
76 res |= sign   // set sign
77
78 return res
79

```

```

1  define convert_fp16_to_fp8(x,y,s):
2  // The x parameter is the data
3  // The y parameter is the destination format indicating bfloat8 or hfloat8.
4  // The s parameter indicates whether we saturate in case of overflow due to
5  // conversion or rounding
6  if *y is bf8*:
7      return convert_fp16_to_bf8(x,s)
8  else:
9      return convert_fp16_to_hf8(x,s)

```

```

1  define convert_fp16_to_hf8_bias(x, b, s):
2  // The x parameter is the data
3  // The b paramemter is the bias 8b integer to be added to the data for RS
4  // The s parameter indicates whether we saturate in case of overflow due to
5  // conversion or rounding
6
7  fp16_bias := 15
8  hf8_bias := 7
9  fp16_to_hf8_exp_rebias = fp16_bias- hf8_bias
10
11 // extract original sign, mantissa and exponent
12 sign := ( x & 0x8000 ) >> 8
13 e_fp16 := ( x & 0x7c00 ) >> 10
14 m_fp16 := ( x & 0x03ff )
15
16 // extract biased mantissa and biased exponent
17 x_bias := x + ( b >> 1 )
18 e_fp16_bias := ( x_bias & 0x7c00 ) >> 10
19 m_fp16_bias := ( x_bias & 0x03ff )
20

```

```

21     if *x is infinity*
22         e := 0xF
23         if (s==0x1):                // Max Value
24             m := 0x6
25         else:                        // NaN
26             m := 0x7
27
28     else if *x is nan*:
29         e := 0xF
30         m := 0x7
31
32     /* overflow --> make it HF8 NaN/Inf == s.1111.111 */
33     /* or Saturate to E4M3_MAX == s.1111.110 */
34     else if ((e_fp16_bias > (fp16_to_hf8_exp_rebias + 15)) ||
35             ((e_fp16_bias == (fp16_to_hf8_exp_rebias + 15)) &&
36             ( m_fp16_bias >= 0x0380))) then:
37         e = 0xf
38         if ( s == 0x1 ) then:
39             m = 0x6
40         else:
41             m = 0x7
42
43     // input denormal (or zero)
44     // in this case the bias rounding will end up with min-denormal or 0
45
46     else if (e_fp16 == 0x0):
47         m = m_fp16+(b<<7)           // align bias
48         m = m >> (fp16_to_hf8_exp_rebias) + 7 // align mantissa 8+7
49         e = 0x0;                   // set exp to zero
50
51     /* underflow*/
52     // underflow happens if e_fp16_bias is too small for representing in
53     // the destination format in this case the Jbit should be set and bias
54     // rounding should be done after aligning to the destination format
55     else if ( e_fp16_bias <= fp16_to_hf8_exp_rebias ):
56         /* set Jbit */
57         m = m_fp16 | 0x0400
58
59         // m += aligned b to mantissa
60         m = m +( b<< (fp16_to_hf8_exp_rebias - e_fp16))
61
62         // now mantissa is aligned to destination exponent + subnormal shift
63         m = m >> ((fp16_to_hf8_exp_rebias) + 1 - e_fp16)
64
65         // in case of round overflow, m>>10 declares the carry into exponent
66         e = m >> 10;              // e=1 in case of overflow
67
68         /* Truncate Round and ignore carry*/

```

```

69         m = ( m >> 7 ) & 0x7;
70
71     /* normal */
72     else then:
73         /* Stochastic Round by truncating */
74         e = ( x_bias & 0x7c00 ) >> 10
75         e -= (fp16_to_hf8_exp_rebias)
76         m = ( x_bias & 0x03ff )
77         m = m >> 7
78
79     res = 0x0
80     res |= e << 3 // set exp
81     res |= m      //set mant
82     res |= sign  // set sign
83
84 return dest

```

```

1  define convert_fp16_to_bf8_bias(x, b, s):
2  // The x parameter is the data
3  // The b parameter is the bias 8b integer to be added to
4  // the data before the downconvert
5  // The s parameter indicates whether we saturate in case of
6  // overflow due to conversion or rounding
7
8  if *x is infinity*
9      if (s==0x1):                // Max Value
10         dest[7] := x[15]
11         dest[6:0] := 0x7B
12     else:                        // INF
13         dest[7:0] := x[15:8]
14
15 else if *x is nan*:
16     dest[7:0] := x[15:8]        // truncate and set QNaN
17     dest[1] := 1
18
19 else // normal, denormal or zero input operand apply RS
20     rounding_bias[15:8] := 0
21     rounding_bias[7:0] := b[7:0]
22     temp[15:0] := x[15:0] + rounding_bias // temp is rounded data (RS)
23     // saturating to E5M2_MAX in case of infinity result
24     if temp [14:8] == 0x7C && s==1
25         dest[7] := temp[15]
26         dest[6:0] := 0x7B
27     else
28         dest[7:0] := temp[15:8]
29 return dest
30

```


31

```

1  define convert_fp16_to_fp8_bias(x, b, y, s):
2  // The x parameter is the data
3  // The b parameter is the bias
4  // The y parameter is the destination format indicating bfloat8 or hfloat8
5  // The s parameter indicating saturation
6  if *y is bf8*:
7      return convert_fp16_to_bf8_bias(x, b, s)
8  else:
9      return convert_fp16_to_hf8_bias(x, b, s)

```

```

1  define convert_fp32_to_fp16(x):
2
3  // The x parameter is the data
4  // rm=MXCSR.RC or embedded.RC
5  Fp32_bias := 127
6  Fp16_bias := 15
7
8  sign := ( x & 0x80000000 ) >> 16
9  e_fp32 := ( x & 0x7F800000 ) >> 23
10 m_fp32 := ( x & 0x007FFFFF)
11
12 m_fp16 = m_fp32 >> 13
13 e_fp16 = e_fp32 - fp32_bias + fp16_bias
14
15 if (e_fp32==0xFF) && (m_fp32==0x000000) // *x is infinity*
16     e_fp16 := 0x1F
17     m_fp16 := 0x000
18
19 else if (e_fp32==0xFF) && (m_fp32 != 0x000000) // *x is nan*:
20     e_fp16 := 0x1F
21     m_fp16 := (m_fp32 | 0x400000) >> 13
22
23 /* Zero */
24 else if (e_fp32==0x00) && (DAZ || (m_fp32 == 0x000000)) // *x is zero*:
25     e_fp16 := 0x00
26     m_fp16 := 0x000
27
28 else if (e_fp32==0x00) && (m_fp32 != 0x000000) // *x is denormal*:
29     // check if result is zero or min-denormal
30     RoundAdd1 := (sign) ? (rm==RDN) : (rm==RUP)
31     e_fp16 := 0x00
32     m_fp16 := 0x000 + RoundAdd1
33
34 /* underflow */

```

```

35 else if ( e_fp32 <= fp32_bias - fp16_bias ):
36 /* denormalized mantissa */
37     /* set Jbit */
38     m := (m_fp32 | 0x800000);
39
40     /* Calculate shift out sticky */
41     ShiftOutGbitMask = (0x1 << ((fp32_bias - fp16_bias) - e_fp32));
42     ShiftOutStickyMask = ShiftOutGbitMask - 1;
43
44     ShiftOutSticky = OR (m & ShiftOutStickyMask);
45     ShiftOutGbit = OR (m & ShiftOutGbitMask);
46
47     /* denormalization */
48     m := m >> (fp32_bias - fp16_bias) - e_fp32 + 1;
49     Lbit = m & 1;
50
51     /* Rounding */
52     RoundAdd1 = (~sign && (rm==RUP) && (ShiftOutGbit | ShiftOutSticky)) ||
53                (sign && (rm==RDN) && (ShiftOutGbit | ShiftOutSticky)) ||
54                (rm==RNE && ShiftOutGbit &(Lbit | ShiftOutSticky));
55     m = m + RoundAdd1;
56
57     m_fp16 = m & 0x3FF;
58     e_fp16 = (m>>10) & 1;
59
60 else
61     /* normal round parameters */
62     Lbit1 = (m_fp32 & 0x2000) >> 13
63     Gbit1 = (m_fp32 & 0x1000) >> 12
64     Stky1 = OR(m_fp32 & 0x0FFF)
65     RoundAdd1 := (~sign && (rm==RUP) && (Gbit1 | Stky1)) ||
66                (sign && (rm==RDN) && (Gbit1 | Stky1)) ||
67                (rm==RNE && Gbit1 &(Lbit1 | Stky1));
68
69     /* overflow --> make it INF (depending on round mode rm) */
70     if (e_fp32  (fp32_bias - fp16_bias + 31) ||
71         (e_fp32 == (fp32_bias - fp16_bias + 30)) && (m_fp32 > 0x7FE000) && RoundAdd1:
72         e_fp16 := 0x1F
73         m_fp16 := 0x000
74
75     else /* normal */
76         m = (m_fp32 & 0x7FE000) >> 13;
77         m = m + RoundAdd1;
78         m_fp16 = m & 0x3FF;
79         e0 = (m>>10) & 1;
80         e_fp16 = e_fp32 - fp32_bias + fp16_bias + e0;
81
82

```

```
83  /* set result */
84  res := 0x0
85  res |= e_fp16 << 10
86  res |= m_fp16
87  res |= sign
88  return res
```

Figure 5.1: minimum

```
1 def minimum(a,b):
2     if a is SNAN or (a is QNAN and b is not SNAN):
3         return QNAN(a)
4     else if b is NAN:
5         return QNAN(b)
6     else if (a == +0.0 and b == -0.0) or (a == -0.0 and b == +0.0):
7         return -0.0
8     else if a <= b:
9         return a
10    else:
11        return b
```

Figure 5.2: minimumNumber

```
1 def minimum_number(a,b):
2     if a is NAN and B is NAN:
3         if a is SNAN or (a is QNAN and b is QNAN):
4             return QNAN(a)
5         else: // a is QNAN and b is SNAN
6             return QNAN(b)
7     else if a is NAN:
8         return b
9     else if b is NAN:
10        return a
11    else if (a == +0.0 and b == -0.0) or (a == -0.0 and b == +0.0):
12        return -0.0
13    else if a <= b:
14        return a
15    else:
16        return b
```

Figure 5.3: minimumMagnitude

```
1 def minimum_magnitude(a,b):
2     if a is SNAN or (a is QNAN and b is not SNAN):
3         return QNAN(a)
4     else if b is NAN:
5         return QNAN(b)
6     else if abs(a) < abs(b):
7         return a
8     else if abs(b) < abs(a):
9         return b
10    else:
11        return minimum(a,b)
```

Figure 5.4: minimumMagnitudeNumber

```
1 def minimum_magnitude_number(a,b):
2     if a is NAN and B is NAN:
3         if a is SNAN or (a is QNAN and b is QNAN):
4             return QNAN(a)
5         else: // a is QNAN and b is SNAN
6             return QNAN(b)
7     else if a is NAN:
8         return b
9     else if b is NAN:
10        return a
11    else if abs(a) < abs(b):
12        return a
13    else if abs(b) < abs(a):
14        return b
15    else:
16        return minimum_number(a,b)
```

Figure 5.5: maximum

```
1 def maximum(a,b):
2     if a is SNAN or (a is QNAN and b is not SNAN):
3         return QNAN(a)
4     else if b is NAN:
5         return QNAN(b)
6     else if (a == +0.0 and b == -0.0) or (a == -0.0 and b == +0.0):
7         return +0.0
8     else if a >= b:
9         return a
10    else:
11        return b
```

Figure 5.6: maximumNumber

```
1 def maximum_number(a,b):
2     if a is NAN and B is NAN:
3         if a is SNAN or (a is QNAN and b is QNAN):
4             return QNAN(a)
5         else: // a is QNAN and b is SNAN
6             return QNAN(b)
7     else if a is NAN:
8         return b
9     else if b is NAN:
10        return a
11    else if (a == +0.0 and b == -0.0) or (a == -0.0 and b == +0.0):
12        return +0.0
13    else if a >= b:
14        return a
15    else:
16        return b
```

Figure 5.7: maximumMagnitude

```
1 def maximum_magnitude(a,b):
2     if a is SNAN or (a is QNAN and b is not SNAN):
3         return QNAN(a)
4     else if b is NAN:
5         return QNAN(b)
6     else if abs(a) > abs(b):
7         return a
8     else if abs(b) > abs(a):
9         return b
10    else:
11        return maximum(a,b)
```

Figure 5.8: maximumMagnitudeNumber

```
1 def maximum_magnitude_number(a,b):
2     if a is NAN and B is NAN:
3         if a is SNAN or (a is QNAN and b is QNAN):
4             return QNAN(a)
5         else: // a is QNAN and b is SNAN
6             return QNAN(b)
7     else if a is NAN:
8         return b
9     else if b is NAN:
10        return a
11    else if abs(a) > abs(b):
12        return a
13    else if abs(b) > abs(a):
14        return b
15    else:
16        return maximum_number(a,b)
```

Figure 5.9: minmax

```

1 def minmax(a,b,imm,daz,except):
2     op_select := imm[1:0]
3     sign_control := imm[3:2]
4     nan_prop_select := imm[4]
5
6     if daz == true:
7         if a is denormal:
8             a.fraction := 0
9         if b is denormal:
10            b.fraction := 0
11
12    if except == true:
13        if a is SNAN or b is SNAN:
14            set_MXCSR(IE)
15        else if a is QNAN or b is QNAN:
16            // QNAN prevents lower-priority exceptions (SDM Vol.3A Table 6-8)
17        else if a is denormal or b is denormal:
18            set_MXCSR(DE)
19
20    if nan_prop_select == 0: //propagate NaNs
21        if op_select == 0:
22            tmp := minimum(a,b)
23        else if op_select == 1:
24            tmp := maximum(a,b)
25        else if op_select == 2:
26            tmp := minimum_magnitude(a,b)
27        else: //op_select == 3
28            tmp := maximum_magnitude(a,b)
29    else: //do not propagate NaNs
30        if op_select == 0:
31            tmp := minimum_number(a,b)
32        else if op_select == 1:
33            tmp := maximum_number(a,b)
34        else if op_select == 2:
35            tmp := minimum_magnitude_number(a,b)
36        else: //op_select == 3
37            tmp := maximum_magnitude_number(a,b)
38    if tmp is not NAN:
39        if (sign_control == 3):
40            tmp.sign := 1
41        else if (sign_control == 2):
42            tmp.sign := 0
43        else if (sign_control == 1) or (a is NAN):
44            // Keep sign of comparison result
45            // i.e. tmp.sign is un-changed
46        else: // sign_control == 0
47            tmp.sign := a.sign
48    return tmp

```



```
1 define convert_bf16_to_signed_byte_rne_saturate(src.bf16):
2     /* VCVTNEBF162IBS converts brain-float16 floating point elements into
3     signed byte integer elements. When a conversion is inexact, the rounding mode
4     is RNE. If a converted result cannot be represented in the destination format
5     then: In case value is too big, the INT_MAX value (2^(w-1)-1, where w
6     represents the number of bits in the destination format) is returned. In case
7     value is too small, the INT_MIN value -(2^(w-1)) is returned. In case of NaN,
8     (0) is returned. */
9
10    Dest;
11    TMP = 0;
12
13    IF (src.bf16==NaN):
14        TMP[31:0]=0x00000000; // return zero in case of NaN
15    ELSE IF (src.bf16 > 127) || (src.bf16 == +INF):
16        TMP[31:0]=0x0000007F; // saturate to max signed value
17    ELSE IF (src.bf16 < -128) || (src.bf16 == -INF):
18        TMP[31:0]=0x00000080; // saturate to min signed value
19    ELSE:
20        // make it fp32 and then convert to INT using RNE
21        TMP[31:0] = src.bf16 << 16
22        TMP[31:0] = Convert_Single_Precision_Floating_Point_To_Integer_RNE(TMP);
23
24    Dest.b = TMP[7:0];
25
26    return Dest;
```

```
1 define convert_bf16_to_signed_byte_truncate_saturate(src.bf16):
2
3     /* VCVTNEBF162IBS converts brain-float16 floating point elements into
4     signed byte integer elements. When a conversion is inexact, a truncated (round
5     toward zero) result is returned. If a converted result cannot be represented in
6     the destination format then: In case the value is too big, the INT_MAX value
7     (2^(w-1)-1, where w represents the number of bits in the destination format) is
8     returned. In case the value is too small, the INT_MIN value -(2^(w-1)) is returned.
9     In case of NaN, (0) is returned. */
10
11     Dest;
12     TMP = 0;
13
14     IF (src.bf16==NaN):
15         TMP[31:0]=0x00000000;    // return zero in case of NaN
16     ELSE IF (src.bf16 > 127) || (src.bf16 == +INF):
17         TMP[31:0]=0x0000007F;    // saturate to max signed value
18     ELSE IF (src.bf16 < -128) || (src.bf16 == -INF):
19         TMP[31:0]=0x00000080;    // saturate to min signed value
20     ELSE:
21         // make it fp32 and then convert to INT using RTZ (Truncate)
22         TMP[31:0] = src.bf16 << 16
23         TMP[31:0] = Convert_Single_Precision_Floating_Point_To_Integer_RTZ(TMP);
24
25     Dest.b = TMP[7:0];
26     return Dest;
27
```

```
1 define convert_bf16_to_unsigned_byte_rne_saturate(src.bf16):
2     /* VCVTNEBF162IUBS converts brain-float16 floating point elements into
3     un-signed byte integer elements. When a conversion is inexact, the rounding
4     mode is RNE. If a converted result cannot be represented in the destination
5     format then: In case value is too big, the UINT_MAX value ( $2^w-1$ , where w
6     represents the number of bits in the destination format) is returned. In case
7     value is too small, the UINT_MIN value (0) is returned. In case of NaN, (0) is
8     returned. */
9
10    Dest;
11    TMP = 0;
12
13    IF (src.bf16==NaN):
14        TMP[31:0]=0x00000000; // return zero in case of NaN
15    ELSE IF (src.bf16 > 255) || (src.bf16 == +INF):
16        TMP[31:0]=0x000000FF; // saturate to max unsigned value
17    ELSE IF (src.bf16 < 0) || (src.bf16 == -INF):
18        TMP[31:0]=0x00000000; // saturate to min unsigned value
19    ELSE:
20        // make it fp32 and then convert to INT using RNE
21        TMP[31:0] = src.bf16 << 16
22        TMP[31:0] = Convert_Single_Precision_Floating_Point_To_Integer_RNE(TMP);
23
24    Dest.b = TMP[7:0];
25    return Dest;
26
```

```
1 define convert_bf16_to_unsigned_byte_truncate_saturate(src.bf16):
2
3     /* VCVTTNEBF162IUBS converts brain-float16 floating point elements into
4     un-signed byte integer elements. When a conversion is inexact, a truncated
5     (round toward zero) result is returned. If a converted result cannot be
6     represented in the destination format then: In case value is too big, the
7     UINT_MAX value (2w-1, where w represents the number of bits in the destination
8     format) is returned. In case value is too small, the UINT_MIN value (0) is
9     returned. In case of NaN, (0) is returned. */
10
11     Dest;
12     TMP = 0;
13
14     IF (src.bf16==NaN):
15         TMP[31:0]=0x00000000; // return zero in case of NaN
16     ELSE IF (src.bf16 > 255) || (src.bf16 == +INF):
17         TMP[31:0]=0x000000FF; // saturate to max unsigned value
18     ELSE IF (src.bf16 < 0) || (src.bf16 == -INF):
19         TMP[31:0]=0x00000000; // saturate to min unsigned value
20     ELSE:
21         // make it fp32 and then convert to INT using RTZ (Truncate)
22         TMP[31:0] = src.bf16 << 16
23         TMP[31:0] = Convert_Single_Precision_Floating_Point_To_Integer_RTZ(TMP);
24
25     Dest.b = TMP[7:0];
26     return Dest;
```

```

1  define convert_fp16_to_signed_byte_saturate(src.fp16):
2
3      /* VCVTPH2IBS converts half-precision floating point elements into
4      signed byte integer elements. When a conversion is inexact, floating-point
5      precision exception is raised and the value returned is rounded according to
6      the rounding control bits in the MXCSR register or the embedded rounding
7      control bits. If a converted result cannot be represented in the destination
8      format, the floating-point invalid exception is raised, and if this exception
9      is masked then: In case the value is too big, the INT_MAX value (2^(w-1)-1, where
10     represents the number of bits in the destination format) is returned. In case the
11     value is too small, the INT_MIN value -(2^(w-1)) is returned. In case of NaN, (0)
12     is returned.  * /
13     W=8;          // Byte destination size
14     RC = MXCSR.RC
15     EXP = 2^(W-1)
16     OutOfDestRepresentation = Case (RC) :
17         RUP:      ((src.fp16 <= -(EXP + 1)) || (src.fp16 > (EXP - 1))),
18         RDN:      ((src.fp16 < -(EXP)) || (src.fp16 >= (EXP))),
19         RTZ:      ((src.fp16 <= -(EXP + 1)) || (src.fp16 >= (EXP))),
20         RNE:      ((src.fp16 < -(EXP + 1/2)) || (src.fp16 >= (EXP - 1/2)))
21
22     Check IF (src.fp16 ==NaN) || (src.fp16 == +/-INF) || OutOfDestRepresentation;
23                                                     // IE=1
24
25     Dest;
26     TMP = 0;
27
28     IF (src.fp16==NaN):
29         TMP[15:0]=0x0000; // return zero in case of NaN
30
31     ELSE IF (src.fp16 >= 127) || (src.fp16 == +INF):
32         TMP[15:0]=0x007F; // saturate to max signed value
33
34     ELSE IF (src.fp16 <=-128) || (src.fp16 == -INF):
35         TMP[15:0]=0x0080; // saturate to min signed value
36
37     ELSE:
38         // convert to INT using MXCSR.RC
39         TMP[15:0] = Convert_fp16_to_integer16(src.fp16);
40         // MXCSR.PE is updated according to result
41     IF (src.fp16 !=half(TMP)) && (NOT OutOfDestRepresentation ) : PE=1
42
43     Dest.b = TMP[7:0];
44     return Dest;

```

```

1  define convert_fp16_to_signed_byte_truncate_saturate(src.fp16):
2
3      /* VCVTTPH2IBS converts half-precision floating point elements into
4      signed byte integer elements. When a conversion is inexact, floating-point
5      precision exception is raised and a truncated (round toward zero) result is
6      returned. If a converted result cannot be represented in the destination
7      format, the floating-point invalid exception is raised, and if this exception
8      is masked then: In case the value is too big, the INT_MAX value ( $2^{(w-1)}-1$ , where
9      represents the number of bits in the destination format) is returned. In case the
10     value is too small, the INT_MIN value ( $-(2^{(w-1)})$ ) is returned. In case of NaN, (0)
11     is returned. * /
12
13     W=8;          // Byte destination size
14     RC = RTZ
15     EXP = 2^(W-1)
16     OutOfDestRepresentation = ((src.fp16 <= -(EXP + 1)) || (src.fp16 >= (EXP)));
17     Check IF (src.fp16 ==NaN) || (src.fp16 == +/-INF) ||OutOfDestRepresentation;
18     // IE=1
19
20     Dest;
21     TMP = 0;
22     IF (src.fp16==NaN):
23         TMP[15:0]=0x0000;    // return zero in case of NaN
24
25     ELSE IF (src.fp16 >= 127) || (src.fp16 == +INF):
26         TMP[15:0]=0x007F; // saturate to max signed value
27         // PE=1
28     ELSE IF (src.fp16 <=-128) || (src.fp16 == -INF):
29         TMP[15:0]=0x0080; // saturate to min signed value
30         // PE=1
31     ELSE:
32         // convert to INT using RTZ (Truncate)
33         TMP[15:0] = Convert_fp16_to_integer16_truncate(src.fp16<<16);
34         // MXCSR.PE is updated according to result
35
36     IF (src.fp16 !=half(TMP)) && (NOT OutOfDestRepresentation ) : PE=1
37
38     Dest.b = TMP[7:0];
39     return Dest;

```

```

1  define convert_fp16_to_unsigned_byte_saturate(src.fp16):
2
3      /* VCVTPH2IUBS converts half-precision floating point elements into
4      un-signed byte integer elements. When a conversion is inexact, floating-point
5      precision exception is raised and the value returned is rounded according to
6      the rounding control bits in the MXCSR register or the embedded rounding
7      control bits. If a converted result cannot be represented in the destination
8      format, the floating-point invalid exception is raised, and if this exception
9      is masked then: In case the value is too big, the UINT_MAX value (2^(w-1)), where w
10     represents the number of bits in the destination format) is returned. In case the
11     value is too small, the UINT_MIN value (0) is returned. In case of NaN, (0)
12     is returned. * /
13
14     W=8;          // Byte Destination Size
15     RC = MXCSR.RC
16     EXP = 2^(W)
17     OutOfDestRepresentation = Case (RC) :
18         RUP:      ((src.fp16 <= -1) || (src.fp16 > (EXP - 1))),
19         RDN:      ((src.fp16 < 0) || (src.fp16 >= (EXP))),
20         RTZ:      ((src.fp16 <= -1) || (src.fp16 >= (EXP))),
21         RNE:      ((src.fp16 < -1/2) || (src.fp16 >= (EXP - 1/2)))
22
23     Check IF (src.fp16 ==NaN) || (src.fp16 == +/-INF) || OutOfDestRepresentation
24                                                    // IE=1
25
26     Dest;
27     TMP = 0;
28     IF (src.fp16==NaN):
29         TMP[15:0]=0x0000; // return zero in case of NaN
30
31     ELSE IF (src.fp16 >= 255) || (src.fp16 == +INF):
32         TMP[15:0]=0x00FF; // saturate to max unsigned value
33
34     ELSE IF (src.fp16 <=0 ) || (src.fp16 == -INF):
35         TMP[15:0]=0x0000; // saturate to min unsigned value
36
37     ELSE:
38         // convert to INT using MXCSR.RC
39         TMP[15:0] = Convert_fp16_to_unsigned_integer16(src.fp16);
40         // MXCSR.PE is updated according to result
41         IF (src.fp16 !=half(TMP)) && (NOT OutOfDestRepresentation ) : PE=1
42
43     Dest.b = TMP[7:0];
44     return Dest;

```

```

1  define convert_fp16_to_unsigned_byte_truncate_saturate(src.fp16):
2
3      /* VCVTTPH2IUBS converts half-precision floating point elements into
4      un-signed byte integer elements. When a conversion is inexact, floating-point
5      precision exception is raised and a truncated (round toward zero) result is
6      returned. If a converted result cannot be represented in the destination
7      format, the floating-point invalid exception is raised, and if this exception
8      is masked then: In case the value is too big, the UINT_MAX value (2w-1, where w
9      represents the number of bits in the destination format) is returned. In case the
10     value is too small, the UINT_MIN value (0) is returned. In case of NaN, (0) is
11     returned. */
12
13     W=8;          // Byte Destination Size
14     RC = RTZ
15     EXP = 2(W)
16     OutOfDestRepresentation = ((src.fp16 <= -1) || (src.fp16 >= (EXP)));
17
18     Check IF (src.fp16 ==NaN) || (src.fp16 == +/-INF) || OutOfDestRepresentation
19                                     // IE=1
20
21     Dest;
22     TMP = 0;
23     IF (src.fp16==NaN):
24         TMP[15:0]=0x0000; // return zero in case of NaN
25
26     ELSE IF (src.fp16 >= 255) || (src.fp16 == +INF):
27         TMP[15:0]=0x00FF; // saturate to max unsigned value
28
29     ELSE IF (src.bf16 <=0) || (src.bf16 == -INF):
30         TMP[15:0]=0x0000; // saturate to min unsigned value
31
32     ELSE:
33         // make it fp32 and then convert to INT using RTZ (Truncate)
34         TMP[15:0] = Convert_fp16_to_unsigned_integer16_truncate(src.fp16);
35         // MXCSR.PE is updated according to result
36         IF (src.fp16 !=half(TMP)) && (NOT OutOfDestRepresentation ) : PE=1
37         Dest.b = TMP[7:0];
38     return Dest;

```



```

1  define convert_fp32_to_signed_byte_saturate(src.fp32):
2
3
4      /* VCVTPS2IBS converts single-precision floating point elements into
5      signed byte integer elements. When a conversion is inexact, floating-point
6      precision exception is raised and the value returned is rounded according to
7      the rounding control bits in the MXCSR register or the embedded rounding
8      control bits. If a converted result cannot be represented in the destination
9      format, the floating-point invalid exception is raised, and if this exception
10     is masked then: If value is too big, the INT_MAX value ( $2^{(w-1)}-1$ , where w
11     represents the number of bits in the destination format) is returned. If value
12     is too small, the INT_MIN value ( $-(2^{(w-1)})$ ) is returned. For NaN, (0) is
13     returned. */
14
15     W=8;          // Byte destination size
16     RC = MXCSR.RC
17     EXP = 2^(W-1)
18     OutOfDestRepresentation = Case (RC) :
19         RUP:      ((src.fp32 <= -(EXP + 1)) || (src.fp32 > (EXP - 1))),
20         RDN:      ((src.fp32 < -(EXP)) || (src.fp32 >= (EXP))),
21         RTZ:      ((src.fp32 <= -(EXP + 1)) || (src.fp32 >= (EXP))),
22         RNE:      ((src.fp32 < -(EXP + 1/2)) || (src.fp32 >= (EXP - 1/2)))
23
24     Check IF (src.fp32 ==NaN) || (src.fp32 == +/-INF) || OutOfDestRepresentation;
25                                                    // IE=1
26
27     Dest;
28     TMP = 0;
29
30     IF (src.fp32==NaN):
31         TMP[31:0]=0x00000000; // return zero in case of NaN
32     ELSE IF (src.fp32 >= 127) || (src.fp32 == +INF):
33         TMP[31:0]=0x0000007F; // saturate to max signed value
34     ELSE IF (src.fp32 <=-128) || (src.fp32 == -INF):
35         TMP[31:0]=0x00000080; // saturate to min signed value
36     ELSE:
37         // (-128<x<127) make it fp32 and then convert to INT using MXCSR.RC
38         TMP[31:0] = Convert_Single_Precision_Floating_Point_To_Integer(src.fp32);
39         // MXCSR.PE is updated according to result
40
41     IF (src.fp32 != float(TMP)) && (NOT OutOfDestRepresentation) : PE=1
42     Dest.b = TMP[7:0];
43     return Dest;

```

```

1  define convert_fp32_to_signed_byte_truncate_saturate(src.fp32):
2
3      /* VCVTTPS2IBS converts single-precision floating point elements into
4      signed byte integer elements. When a conversion is inexact, floating-point
5      precision exception is raised and a truncated (round toward zero) result is
6      returned. If a converted result cannot be represented in the destination
7      format, the floating-point invalid exception is raised, and if this exception
8      is masked then: If value is too big, the INT_MAX value ( $2^{(w-1)}-1$ , where w
9      represents the number of bits in the destination format) is returned. If value
10     is too small, the INT_MIN value ( $-(2^{(w-1)})$ ) is returned. For NaN, (0) is
11     returned. */
12
13     Dest;
14     TMP = 0;
15
16     W=8; // Byte destination size
17     RC = RTZ
18     EXP =  $2^{(W-1)}$ 
19     OutOfDestRepresentation = ((src.fp32 <= -(EXP + 1)) || (src.fp32 >= EXP));
20
21     IF ((src.fp32 ==NaN) ||
22         (src.fp32 == +INF ||src.fp32 == -INF) ||
23         OutOfDestRepresentation):
24         // signal IE=1
25
26     IF (src.fp32 ==NaN):
27         TMP[31:0]=0x00000000; // return zero in case of NaN
28     ELIF (src.fp32 >= 127) || (src.fp32 == +INF):
29         TMP[31:0]=0x0000007F; // saturate to max signed value
30     ELIF (src.fp32 <= -128) || (src.fp32 == -INF):
31         TMP[31:0]=0x00000080; // saturate to min signed value
32     ELSE:
33         // make it fp32 and then convert to INT using RTZ (Truncate)
34         TMP[31:0] = Convert_Single_Precision_Floating_Point_To_Integer_RTZ(src.fp32);
35         // set PE if inexact conversion.
36     IF (src.fp32 != float(TMP)) && (NOT OutOfDestRepresentation) : PE=1
37     Dest.b = TMP[7:0];
38
39     return Dest;

```

```

1  define convert_fp32_to_unsigned_byte_saturate(src.fp32):
2
3      /* VCVTNEBF162IUBS converts brain-float16 floating point elements into
4      un-signed byte integer elements. When a conversion is inexact, the rounding
5      mode is RNE. If a converted result cannot be represented in the destination
6      format then: In case value is too big, the UINT_MAX value ( $2^w-1$ , where w
7      represents the number of bits in the destination format) is returned. In case
8      value is too small, the UINT_MIN value (0) is returned. In case of NaN, (0) is
9      returned. */
10
11     TMP = 0;
12     W=8;          // Byte Destination Size
13     RC = MXCSR.RC
14     OutOfDestRepresentation = ((src.fp32 <= -1) || (src.fp32 >= (2w)))
15
16     IF ((src.fp32 ==NaN) || (src.fp32 == +/-INF) || OutOfDestRepresentation):
17         // signal IE=1
18
19     Dest;
20     TMP = 0;
21     IF (src.fp32==NaN):
22         TMP[31:0]=0x00000000;    // return zero in case of NaN
23
24     ELSE IF (src.fp32 > 255) || (src.fp32 == +INF):
25         TMP[31:0]=0x000000FF;    // saturate to max unsigned value
26
27     ELSE IF (src.fp32 < 0) || (src.fp32 == -INF):
28         TMP[31:0]=0x00000000;    // saturate to min unsigned value
29     ELSE:
30         // convert to INT using MXCSR or embedded rounding mode
31         TMP[31:0] = Convert_Single_Precision_Floating_Point_To_UInteger(src.fp32);
32     IF (src.fp32 != float(TMP)) && (NOT OutOfDestRepresentation) : PE=1
33     Dest.b = TMP[7:0];
34     return Dest;

```

```

1 define convert_fp32_to_unsigned_byte_truncate_saturate(src.fp32):
2
3     /* VCVTTPS2IUBS converts single-precision floating point elements into
4     un-signed byte integer elements. When a conversion is inexact, floating-point
5     precision exception is raised and a truncated (round toward zero) result is
6     returned. If a converted result cannot be represented in the destination
7     format, the floating-point invalid exception is raised, and if this exception
8     is masked then: If value is too big, the UINT_MAX value ( $2^w-1$ , where w
9     represents the number of bits in the destination format) is returned. If value
10    is too small, the UINT_MIN value (0) is returned. For NaN, (0) is returned. */
11
12    Dest;
13    TMP = 0;
14    W=8;          // Byte Destination Size
15    RC = RTZ
16    OutOfDestRepresentation = ((src.fp32 <= -1) || (src.fp32 >= (2w)))
17
18    IF (src.fp32 ==NaN) || (src.fp32 == +/-INF) || OutOfDestRepresentation:
19        // signal IE=1
20
21    IF (src.fp32==NaN):
22        TMP[31:0]=0x00000000; // return zero in case of NaN
23
24        ELSE IF (src.fp32 > 255) || (src.fp32 == +INF):
25            TMP[31:0]=0x000000FF; // saturate to max unsigned value
26
27    ELSE IF (src.fp32 < 0) || (src.fp32 == -INF):
28        TMP[31:0]=0x00000000; // saturate to min unsigned value
29
30    ELSE:
31        // convert to INT using RTZ (Truncate)
32        TMP = src.fp32
33        TMP[31:0] = Convert_Single_Precision_Floating_Point_To_UInteger_RTZ(TMP);
34        // MXCSR.PE is updated according to result
35    IF (src.fp32 != float(TMP)) && (NOT OutOfDestRepresentation) : PE=1
36    Dest.b = TMP[7:0];
37    return Dest;

```

```

1  define convert_SP_to_DW_SignedInteger_TruncateSaturate(src.fp32):
2
3      /* VCVTTPS2DQS converts single-precision floating point elements into signed
4      double word Integer elements. When a conversion is inexact, floating-point
5      precision exception is raised and a truncated (round toward zero) result is
6      returned. If a converted result cannot be represented in the destination
7      format, the floating-point invalid exception is raised, and if this exception
8      is masked then: If value is too big, the UINT_MAX value ( $2^{w-1}$ , where w
9      represents the number of bits in the destination format) is returned. If value
10     is too small, the INT_MIN value ( $-(2^{w-1})$ ) is returned. For NaN, (0) is returned. */
11
12     Dest;
13     TMP = 0;
14
15     W=32;          // DW destination size
16     RC = RTZ
17     OutOfDestRepresentation = ((src.fp32 <=  $-(2^{w-1} + 1)$ ) ||
18                               (src.fp32 >=  $(2^{w-1})$ ))
19                               );
20
21     IF (src.fp32 ==NaN) || (src.fp32 == +/-INF) || OutOfDestRepresentation:
22         // signal IE=1
23
24     IF (src.fp32==NaN):
25         TMP[31:0]=0x00000000; // return zero in case of NaN
26     ELSE IF (src.fp32 >=  $+2^{31} - 1$ ) || (src.fp32 == +INF):
27         TMP[31:0]=0x7FFFFFFF; // saturate to max signed value
28     ELSE IF (src.fp32 <=  $-2^{31}$ ) || (src.fp32 == -INF):
29         TMP[31:0]=0x80000000; // saturate to min signed value
30     ELSE:
31         // make it fp32 and then convert to INT using RTZ (Truncate)
32         //set PE if inexact conversion
33         TMP[31:0] = Convert_SP_TO_DW_SignedInteger_RTZ(src.fp32);
34     IF (src.fp32 != float(TMP)) && (NOT OutOfDestRepresentation) : PE=1
35     Dest.dw = TMP[31:0];
36     return Dest;

```

```

1  define convert_SP_to_DW_UnSignedInteger_TruncateSaturate(src.fp32):
2
3      /* VCVTTPS2UDQS converts single-precision floating point elements into
4      unsigned double word Integer elements. When a conversion is inexact,
5      floating-point precision exception is raised and a truncated (round
6      toward zero) result is returned. If a converted result cannot be
7      represented in the destination format, the floating-point invalid
8      exception is raised, and if this exception is masked then: If value is
9      too big, the UINT_MAX value ( $2^w-1$ , where w represents the number of
10     bits in the destination format) is returned. If value is too small, the
11     UINT_MIN value (0) is returned. For NaN, (0) is returned. */
12
13     Dest;
14     TMP = 0;
15
16     W=32;          // DW destination size
17     RC = RTZ
18     EXP = 2^(W)
19     OutOfDestRepresentation = ((src.fp32 <=-1) || (src.fp32 >= EXP));
20
21     IF ((src.fp32 ==NaN) ||
22         (src.fp32 == +INF ||src.fp32 == -INF) ||
23         OutOfDestRepresentation):
24         // signal IE=1
25
26     IF (src.fp32 == NaN):
27         TMP[31:0]=0x00000000;          // return zero in case of NaN
28     ELIF (src.fp32 >= 2^32 - 1) || (src.fp32 == +INF):
29         TMP[31:0]=0xFFFFFFFF;        // saturate to max unsigned value
30     ELIF (src.fp32 <= 0) || (src.fp32 == -INF):
31         TMP[31:0]=0x00000000;        // saturate to min unsigned value
32     ELSE:
33         // Convert to fp32 and then convert to INT using RTZ (Truncate)
34         TMP[31:0] = Convert_Single_Precision_Floating_Point_To_Integer_RTZ(src.fp32);
35         // set PE if inexact conversion.
36     IF (src.fp32 != float(TMP)) && (NOT OutOfDestRepresentation) : PE=1
37     Dest.Dword = TMP[31:0];
38
39     return Dest;

```

```

1 define convert_SP_to_QW_SignedInteger_TruncateSaturate(src.fp32):
2
3     /* VCVTTPS2QQS converts single-precision floating point elements into
4     packed signed quadword Integer elements. When a conversion is inexact,
5     floating-point precision exception is raised and a truncated (round
6     toward zero) result is returned. If a converted result cannot be
7     represented in the destination format, the floating-point invalid
8     exception is raised, and if this exception is masked then: If value is
9     too big, the UINT_MAX value ( $2^w-1$ , where w represents the number of
10    bits in the destination format) is returned. If value is too small, the
11    INT_MIN value ( $-(2^{(w-1)})$ ) is returned. For NaN, (0) is returned. */
12
13    Dest;
14    TMP = 0;
15    W=64;          // DW destination size
16    RC = RTZ
17    EXP =  $2^{(W-1)}$ 
18    OutOfDestRepresentation = ((src.fp32 <= -(EXP + 1)) || (src.fp32 >= (EXP)));
19
20    IF (src.fp32 ==NaN) || (src.fp32 == +/-INF) || OutOfDestRepresentation):
21        // signal IE=1
22
23    IF (src.fp32==NaN):
24        TMP[63:0]=0x00000000.00000000; // return zero in case of NaN
25    ELSE IF (src.fp32 >= + $2^{31}$  - 1) || (src.fp32 == +INF):
26        TMP[63:0]=0x7FFFFFFF.FFFFFFFF; // saturate to max signed value
27    ELSE IF (src.fp32 <= - $2^{31}$ ) || (src.fp32 == -INF):
28        TMP[63:0]=0x80000000.00000000; // saturate to min signed value
29    ELSE:
30        // convert to INT using RTZ (Truncate)//set PE if inexact conversion
31        TMP[63:0] = Convert_SP_TO_QW_SignedInteger_RTZ(src.fp32);
32    IF (src.fp32 != float(TMP)) && (NOT OutOfDestRepresentation) : PE=1
33    Dest.qw = TMP[63:0];

```

```

1  define convert_SP_to_QW_UnSignedInteger_TruncateSaturate(src.fp32):
2
3      /* VCVTTPS2UQQS converts single-precision floating point elements into
4      packed unsigned quadword Integer elements. When a conversion is
5      inexact, floating-point precision exception is raised and a truncated
6      (round toward zero) result is returned. If a converted result cannot
7      be represented in the destination format, the floating-point invalid
8      exception is raised, and if this exception is masked then: If value is
9      too big, the UINT_MAX value ( $2^w-1$ , where w represents the number of
10     bits in the destination format) is returned. If value is too small, the
11     UINT_MIN value (0) is returned. For NaN, (0) is returned. */
12
13     Dest;
14     TMP = 0;
15
16     W=64;          // QW destination size
17     EXP = 2^(W)
18     OutOfDestRepresentation = ((src.fp32 <=-1) || (src.fp32 >= EXP));
19
20     IF ((src.fp32 == NaN) ||
21         (src.fp32 == +INF ||src.fp32 == -INF) ||
22         OutOfDestRepresentation):
23         // Signal IE=1
24
25     IF (src.fp32 == NaN):
26         TMP[63:0]=0x00000000.00000000;          // return zero in case of NaN
27     ELIF (src.fp32 >= EXP - 1) || (src.fp32 == +INF):
28         TMP[63:0]=0xFFFFFFFF.FFFFFFFF;      // saturate to max unsigned value
29     ELIF (src.fp32 <= 0) || (src.fp32 == -INF):
30         TMP[63:0]=0x00000000.00000000;      // saturate to min unsigned value
31     ELSE:
32         // make it fp32 and then convert to INT using RTZ (Truncate)
33         TMP[63:0] = cvt_SP_FP_To_QW_Integer_RTZ(src.fp32);
34         // set PE if inexact conversion.
35     IF (src.fp32!= float(TMP)) && (NOT OutOfDestRepresentation) : PE=1
36     Dest.Qword = TMP[63:0];
37
38     return Dest;

```



```

1  define convert_DP_to_DW_SignedInteger_TruncateSaturate(src.fp64):
2
3      /* VCVTTPD2DQS converts double-precision floating point elements into
4      signed double word Integer elements. When a conversion is inexact,
5      floating-point precision exception is raised and a truncated (round
6      toward zero) result is returned. If a converted result cannot be
7      represented in the destination format, the floating-point invalid
8      exception is raised, and if this exception is masked then: If value is
9      too big, the UINT_MAX value (2^(w-1)-1), where w represents the number of
10     bits in the destination format) is returned. If value is too small, the
11     INT_MIN value -(2^(w-1)) is returned. For NaN, (0) is returned. */
12
13     Dest;
14     TMP = 0;
15     W=32;
16     EXP = 2^(W-1)
17     RC = RTZ
18     OutOfDestRepresentation = ((src.fp64 <= -(EXP + 1)) || (src.fp64 >= (EXP)));
19
20     IF (src.fp64 ==NaN) || (src.fp64 == +/-INF) || OutOfDestRepresentation:
21         // signal IE=1
22
23     IF (src.fp64==NaN):
24         TMP[31:0]=0x00000000; // return zero in case of NaN
25         ELSE IF (src.fp64 >= +2^31 - 1) || (src.fp64 == +INF):
26             TMP[31:0]=0x7FFFFFFF; // saturate to max signed value
27     ELSE IF (src.fp64 <=-2^31) || (src.fp64 == -INF):
28         TMP[31:0]=0x80000000; // saturate to min signed value
29     ELSE:
30         //make it fp64 and convert to INT using RTZ (Truncate)
31         //set PE if inexact conversion
32         TMP[31:0] = Convert_DP_TO_DW_SignedInteger_RTZ(src.fp64);
33     IF (src.fp64 != double(TMP)) & (NOT OutOfDestRepresentation) : PE=1
34     Dest.dw = TMP[31:0];
35     return Dest;

```

```

1 define convert_DP_to_DW_UnSignedInteger_TruncateSaturate(src.fp64):
2
3     /* VCVTTPD2UDQS converts double-precision floating point elements into
4     unsigned double word Integer elements. When a conversion is inexact,
5     floating-point precision exception is raised and a truncated (round
6     toward zero) result is returned. If a converted result cannot be
7     represented in the destination format, the floating-point invalid
8     exception is raised, and if this exception is masked then: If value is
9     too big, the UINT_MAX value ( $2^w-1$ , where w represents the number of
10    bits in the destination format) is returned. If value is too small, the
11    UINT_MIN value (0) is returned. For NaN, (0) is returned. */
12
13    Dest;
14    TMP = 0;
15    W=32;          // QW destination size
16    EXP = 2^(W)
17    OutOfDestRepresentation = ((src.fp64 <=-1) || (src.fp64 >= EXP));
18
19    IF (src.fp64 == NaN) ||
20        (src.fp64 == +INF || src.fp64 == -INF) ||
21        OutOfDestRepresentation:
22        // Signal IE=1
23
24    IF (src.fp64 < 0) || (src.fp64 > EXP - 1):
25        // Signal PE=1
26
27    IF (src.fp64==NaN):
28        TMP[31:0]=0x00000000; // return zero in case of NaN
29    ELSE IF (src.fp64 >= +2^32 - 1) || (src.fp64 == +INF):
30        TMP[31:0]=0xFFFFFFFF; // saturate to max signed value
31    ELSE IF (src.fp64 <=0) || (src.fp64 == -INF):
32        TMP[31:0]=0x00000000; // saturate to min signed value
33    ELSE:
34        // make it fp64 and then convert to INT using RTZ (Truncate)
35        //set PE if inexact conversion
36        TMP[31:0] = Convert_DP_TO_DW_UnSignedInteger_RTZ(src.fp64);
37    IF (src.fp64 != double(TMP)) & (NOT OutOfDestRepresentation) : PE=1
38    Dest.dw = TMP[31:0];
39    return Dest;

```

```

1  define convert_DP_to_QW_SignedInteger_TruncateSaturate(src.fp64):
2
3      /* VCVTTPD2QQS converts double-precision floating point elements into
4      packed signed quadword Integer elements. When a conversion is inexact,
5      floating-point precision exception is raised and a truncated (round
6      toward zero) result is returned. If a converted result cannot be
7      represented in the destination format, the floating-point invalid
8      exception is raised, and if this exception is masked then: If value is
9      too big, the UINT_MAX value ( $2^w-1$ , where w represents the number of
10     bits in the destination format) is returned. If value is too small, the
11     INT_MIN value ( $-(2^{(w-1)})$ ) is returned. For NaN, (0) is returned. */
12
13     Dest = 0;
14     TMP = 0;
15     W=64;          // QW destination size
16     EXP=2^(W-1)
17     RC = RTZ
18     OutOfDestRepresentation = ((src.fp64 <= -(EXP + 1)) || (src.fp64 >= (EXP)));
19
20     IF ((src.fp64 == NAN) || (x == +-INF) || OutOfDestRepresentation):
21         // signal IE=1
22
23     IF (src.fp64==NaN):
24         TMP[63:0]=0x00000000.00000000; // return zero in case of NaN
25     ELSE IF (src.fp64 >= EXP - 1) || (src.fp64 == +INF):
26         TMP[63:0]=7FFFFFFF.FFFFFFFF; // saturate to max signed value
27     ELSE IF (src.fp64 <=EXP) || (src.fp64 == -INF):
28         TMP[63:0]=0x80000000.00000000; // saturate to min signed value
29     ELSE:
30         // convert to INT using RTZ (Truncate)
31         //set PE if inexact conversion
32         TMP[63:0] = Convert_DP_To_QW_SignedInteger_RTZ(src.fp64);
33     IF (src.fp64 != double(TMP)) & (NOT OutOfDestRepresentation) : PE=1
34     Dest.qw = TMP[63:0];
35     return Dest;

```

```

1  define convert_DP_to_QW_UnSignedInteger_TruncateSaturate(src.fp64):
2
3      /* VCVTTPD2UQQS converts double-precision floating point elements into
4      packed unsigned quadword Integer elements. When a conversion is
5      inexact, floating-point precision exception is raised and a truncated
6      (round toward zero) result is returned. If a converted result cannot
7      be represented in the destination format, the floating-point invalid
8      exception is raised, and if this exception is masked then: If value is
9      too big, the UINT_MAX value ( $2^w-1$ , where w represents the number of
10     bits in the destination format) is returned. If value is too small, the
11     UINT_MIN value (0) is returned. For NaN, (0) is returned. */
12
13     Dest;
14     TMP = 0;
15     EXP = 2^(W)
16     OutOfDestRepresentation = ((src.fp64 <=-1) || (src.fp64 >= EXP))
17
18     IF (src.fp64 ==NaN) || (src.fp64 == +/-INF) || OutOfDestRepresentation:
19         // signal IE=1
20
21     IF (src.fp64==NaN):
22         TMP[63:0]=0x00000000.00000000; // return zero in case of NaN
23     ELSE IF (src.fp64 >= EXP) || (src.fp64 == +INF):
24         TMP[63:0]=FFFFFFFF.FFFFFFFF; // saturate to max signed value
25     ELSE IF (src.fp64 <=0 ) || (src.fp64 == -INF):
26         TMP[63:0]=0x00000000.00000000; // saturate to min signed value
27     ELSE:
28         // make it fp32 and then convert to INT using RTZ (Truncate)
29         //set PE if inexact conversion
30         TMP[63:0] = Convert_DP_To_QW_UnSignedInteger_RTZ(src.fp64);
31     IF (src.fp64 != double(TMP)) & (NOT OutOfDestRepresentation) : PE=1
32     Dest.qw = TMP[63:0];
33     return Dest;

```

Chapter 6

INSTRUCTION TABLE

CPUID: AVX10.2	OPERANDS	ENCSPACE
VADDNEPBF16	xmm1, xmm2, xmm3/m128	EVEX
VADDNEPBF16	ymm1, ymm2, ymm3/m256	EVEX
VADDNEPBF16	zmm1, zmm2, zmm3/m512	EVEX
VADDDP	ymm1, ymm2, ymm3/m256	EVEX
VADDPH	ymm1, ymm2, ymm3/m256	EVEX
VADDP	ymm1, ymm2, ymm3/m256	EVEX
VCMPPBF16	k1, xmm2, xmm3/m128, imm8	EVEX
VCMPPBF16	k1, ymm2, ymm3/m256, imm8	EVEX
VCMPPBF16	k1, zmm2, zmm3/m512, imm8	EVEX
VCMPPD	k1, ymm2, ymm3/m256, imm8	EVEX
VCMPPH	k1, ymm2, ymm3/m256, imm8	EVEX
VCMPPS	k1, ymm2, ymm3/m256, imm8	EVEX
VCOMSBF16	xmm1, xmm2/m16	EVEX
VCOMXSD	xmm1, xmm2/m64	EVEX
VCOMXSH	xmm1, xmm2/m16	EVEX
VCOMXSS	xmm1, xmm2/m32	EVEX
VCVT2PS2PHX	xmm1, xmm2, xmm3/m128	EVEX
VCVT2PS2PHX	ymm1, ymm2, ymm3/m256	EVEX
VCVT2PS2PHX	zmm1, zmm2, zmm3/m512	EVEX
VCVTBIASPH2BF8	xmm1, xmm2, xmm3/m128	EVEX
VCVTBIASPH2BF8	xmm1, ymm2, ymm3/m256	EVEX
VCVTBIASPH2BF8	ymm1, zmm2, zmm3/m512	EVEX
VCVTBIASPH2BF8S	xmm1, xmm2, xmm3/m128	EVEX
VCVTBIASPH2BF8S	xmm1, ymm2, ymm3/m256	EVEX
VCVTBIASPH2BF8S	ymm1, zmm2, zmm3/m512	EVEX
VCVTBIASPH2HF8	xmm1, xmm2, xmm3/m128	EVEX
VCVTBIASPH2HF8	xmm1, ymm2, ymm3/m256	EVEX
VCVTBIASPH2HF8	ymm1, zmm2, zmm3/m512	EVEX
VCVTBIASPH2HF8S	xmm1, xmm2, xmm3/m128	EVEX
VCVTBIASPH2HF8S	xmm1, ymm2, ymm3/m256	EVEX
VCVTBIASPH2HF8S	ymm1, zmm2, zmm3/m512	EVEX
VCVTDQ2PH	xmm1, ymm2/m256	EVEX
VCVTDQ2PS	ymm1, ymm2/m256	EVEX
VCVTHF82PH	xmm1, xmm2/m64	EVEX
VCVTHF82PH	ymm1, xmm2/m128	EVEX
VCVTHF82PH	zmm1, ymm2/m256	EVEX
VCVTNE2PH2BF8	xmm1, xmm2, xmm3/m128	EVEX
VCVTNE2PH2BF8	ymm1, ymm2, ymm3/m256	EVEX
VCVTNE2PH2BF8	zmm1, zmm2, zmm3/m512	EVEX
VCVTNE2PH2BF8S	xmm1, xmm2, xmm3/m128	EVEX

Table continued on next page...

CPUID: AVX10.2	OPERANDS	ENCSPACE (contd.)
VCVTNE2PH2BF8S	ymm1, ymm2, ymm3/m256	EVEX
VCVTNE2PH2BF8S	zmm1, zmm2, zmm3/m512	EVEX
VCVTNE2PH2HF8	xmm1, xmm2, xmm3/m128	EVEX
VCVTNE2PH2HF8	ymm1, ymm2, ymm3/m256	EVEX
VCVTNE2PH2HF8	zmm1, zmm2, zmm3/m512	EVEX
VCVTNE2PH2HF8S	xmm1, xmm2, xmm3/m128	EVEX
VCVTNE2PH2HF8S	ymm1, ymm2, ymm3/m256	EVEX
VCVTNE2PH2HF8S	zmm1, zmm2, zmm3/m512	EVEX
VCVTNEBF162IBS	xmm1, xmm2/m128	EVEX
VCVTNEBF162IBS	ymm1, ymm2/m256	EVEX
VCVTNEBF162IBS	zmm1, zmm2/m512	EVEX
VCVTNEBF162IUBS	xmm1, xmm2/m128	EVEX
VCVTNEBF162IUBS	ymm1, ymm2/m256	EVEX
VCVTNEBF162IUBS	zmm1, zmm2/m512	EVEX
VCVTNEPH2BF8	xmm1, xmm2/m128	EVEX
VCVTNEPH2BF8	xmm1, ymm2/m256	EVEX
VCVTNEPH2BF8	ymm1, zmm2/m512	EVEX
VCVTNEPH2BF8S	xmm1, xmm2/m128	EVEX
VCVTNEPH2BF8S	xmm1, ymm2/m256	EVEX
VCVTNEPH2BF8S	ymm1, zmm2/m512	EVEX
VCVTNEPH2HF8	xmm1, xmm2/m128	EVEX
VCVTNEPH2HF8	xmm1, ymm2/m256	EVEX
VCVTNEPH2HF8	ymm1, zmm2/m512	EVEX
VCVTNEPH2HF8S	xmm1, xmm2/m128	EVEX
VCVTNEPH2HF8S	xmm1, ymm2/m256	EVEX
VCVTNEPH2HF8S	ymm1, zmm2/m512	EVEX
VCVTPD2DQ	xmm1, ymm2/m256	EVEX
VCVTPD2PH	xmm1, ymm2/m256	EVEX
VCVTPD2PS	xmm1, ymm2/m256	EVEX
VCVTPD2QQ	ymm1, ymm2/m256	EVEX
VCVTPD2UDQ	xmm1, ymm2/m256	EVEX
VCVTPD2UQQ	ymm1, ymm2/m256	EVEX
VCVTPH2DQ	ymm1, xmm2/m128	EVEX
VCVTPH2IBS	xmm1, xmm2/m128	EVEX
VCVTPH2IBS	ymm1, ymm2/m256	EVEX
VCVTPH2IBS	zmm1, zmm2/m512	EVEX
VCVTPH2IUBS	xmm1, xmm2/m128	EVEX
VCVTPH2IUBS	ymm1, ymm2/m256	EVEX
VCVTPH2IUBS	zmm1, zmm2/m512	EVEX
VCVTPH2PD	ymm1, xmm2/m64	EVEX
VCVTPH2PS	ymm1, xmm2/m128	EVEX

Table continued on next page...

CPUID: AVX10.2	OPERANDS	ENCSPACE (contd.)
VCVTPH2PSX	ymm1, xmm2/m128	EVEX
VCVTPH2QQ	ymm1, xmm2/m64	EVEX
VCVTPH2UDQ	ymm1, xmm2/m128	EVEX
VCVTPH2UQQ	ymm1, xmm2/m64	EVEX
VCVTPH2UW	ymm1, ymm2/m256	EVEX
VCVTPH2W	ymm1, ymm2/m256	EVEX
VCVTPS2DQ	ymm1, ymm2/m256	EVEX
VCVTPS2IBS	xmm1, xmm2/m128	EVEX
VCVTPS2IBS	ymm1, ymm2/m256	EVEX
VCVTPS2IBS	zmm1, zmm2/m512	EVEX
VCVTPS2IUBS	xmm1, xmm2/m128	EVEX
VCVTPS2IUBS	ymm1, ymm2/m256	EVEX
VCVTPS2IUBS	zmm1, zmm2/m512	EVEX
VCVTPS2PD	ymm1, xmm2/m128	EVEX
VCVTPS2PH	xmm1, ymm2, imm8	EVEX
VCVTPS2PHX	xmm1, ymm2/m256	EVEX
VCVTPS2QQ	ymm1, xmm2/m128	EVEX
VCVTPS2UDQ	ymm1, ymm2/m256	EVEX
VCVTPS2UQQ	ymm1, xmm2/m128	EVEX
VCVTQQ2PD	ymm1, ymm2/m256	EVEX
VCVTQQ2PH	xmm1, ymm2/m256	EVEX
VCVTQQ2PS	xmm1, ymm2/m256	EVEX
VCVTTNEBF162IBS	xmm1, xmm2/m128	EVEX
VCVTTNEBF162IBS	ymm1, ymm2/m256	EVEX
VCVTTNEBF162IBS	zmm1, zmm2/m512	EVEX
VCVTTNEBF162IUBS	xmm1, xmm2/m128	EVEX
VCVTTNEBF162IUBS	ymm1, ymm2/m256	EVEX
VCVTTNEBF162IUBS	zmm1, zmm2/m512	EVEX
VCVTTPD2DQ	xmm1, ymm2/m256	EVEX
VCVTTPD2DQS	xmm1, xmm2/m128	EVEX
VCVTTPD2DQS	xmm1, ymm2/m256	EVEX
VCVTTPD2DQS	ymm1, zmm2/m512	EVEX
VCVTTPD2QQ	ymm1, ymm2/m256	EVEX
VCVTTPD2QQS	xmm1, xmm2/m128	EVEX
VCVTTPD2QQS	ymm1, ymm2/m256	EVEX
VCVTTPD2QQS	zmm1, zmm2/m512	EVEX
VCVTTPD2UDQ	xmm1, ymm2/m256	EVEX
VCVTTPD2UDQS	xmm1, xmm2/m128	EVEX
VCVTTPD2UDQS	xmm1, ymm2/m256	EVEX
VCVTTPD2UDQS	ymm1, zmm2/m512	EVEX
VCVTTPD2UQQ	ymm1, ymm2/m256	EVEX

Table continued on next page...

CPUID: AVX10.2	OPERANDS	ENCSPACE (contd.)
VCVTPPD2UQQS	xmm1, xmm2/m128	EVEX
VCVTPPD2UQQS	ymm1, ymm2/m256	EVEX
VCVTPPD2UQQS	zmm1, zmm2/m512	EVEX
VCVTPPH2DQ	ymm1, xmm2/m128	EVEX
VCVTPPH2IBS	xmm1, xmm2/m128	EVEX
VCVTPPH2IBS	ymm1, ymm2/m256	EVEX
VCVTPPH2IBS	zmm1, zmm2/m512	EVEX
VCVTPPH2IUBS	xmm1, xmm2/m128	EVEX
VCVTPPH2IUBS	ymm1, ymm2/m256	EVEX
VCVTPPH2IUBS	zmm1, zmm2/m512	EVEX
VCVTPPH2QQ	ymm1, xmm2/m64	EVEX
VCVTPPH2UDQ	ymm1, xmm2/m128	EVEX
VCVTPPH2UQQ	ymm1, xmm2/m64	EVEX
VCVTPPH2UW	ymm1, ymm2/m256	EVEX
VCVTPPH2W	ymm1, ymm2/m256	EVEX
VCVTPPS2DQ	ymm1, ymm2/m256	EVEX
VCVTPPS2DQS	xmm1, xmm2/m128	EVEX
VCVTPPS2DQS	ymm1, ymm2/m256	EVEX
VCVTPPS2DQS	zmm1, zmm2/m512	EVEX
VCVTPPS2IBS	xmm1, xmm2/m128	EVEX
VCVTPPS2IBS	ymm1, ymm2/m256	EVEX
VCVTPPS2IBS	zmm1, zmm2/m512	EVEX
VCVTPPS2IUBS	xmm1, xmm2/m128	EVEX
VCVTPPS2IUBS	ymm1, ymm2/m256	EVEX
VCVTPPS2IUBS	zmm1, zmm2/m512	EVEX
VCVTPPS2QQ	ymm1, xmm2/m128	EVEX
VCVTPPS2QQS	xmm1, xmm2/m64	EVEX
VCVTPPS2QQS	ymm1, xmm2/m128	EVEX
VCVTPPS2QQS	zmm1, ymm2/m256	EVEX
VCVTPPS2UDQ	ymm1, ymm2/m256	EVEX
VCVTPPS2UDQS	xmm1, xmm2/m128	EVEX
VCVTPPS2UDQS	ymm1, ymm2/m256	EVEX
VCVTPPS2UDQS	zmm1, zmm2/m512	EVEX
VCVTPPS2UQQ	ymm1, xmm2/m128	EVEX
VCVTPPS2UQQS	xmm1, xmm2/m64	EVEX
VCVTPPS2UQQS	ymm1, xmm2/m128	EVEX
VCVTPPS2UQQS	zmm1, ymm2/m256	EVEX
VCVTTSD2SIS	r32, xmm1/m64	EVEX
VCVTTSD2SIS	r64, xmm1/m64	EVEX
VCVTTSD2USIS	r32, xmm1/m64	EVEX
VCVTTSD2USIS	r64, xmm1/m64	EVEX

Table continued on next page...

CPUID: AVX10.2	OPERANDS	ENCSPACE (contd.)
VCVTTSS2SIS	r32, xmm1/m32	EVEX
VCVTTSS2SIS	r64, xmm1/m32	EVEX
VCVTTSS2USIS	r32, xmm1/m32	EVEX
VCVTTSS2USIS	r64, xmm1/m32	EVEX
VCVTUDQ2PH	xmm1, ymm2/m256	EVEX
VCVTUDQ2PS	ymm1, ymm2/m256	EVEX
VCVTUQQ2PD	ymm1, ymm2/m256	EVEX
VCVTUQQ2PH	xmm1, ymm2/m256	EVEX
VCVTUQQ2PS	xmm1, ymm2/m256	EVEX
VCVTUW2PH	ymm1, ymm2/m256	EVEX
VCVTW2PH	ymm1, ymm2/m256	EVEX
VDIVNEPBF16	xmm1, xmm2, xmm3/m128	EVEX
VDIVNEPBF16	ymm1, ymm2, ymm3/m256	EVEX
VDIVNEPBF16	zmm1, zmm2, zmm3/m512	EVEX
VDIVPD	ymm1, ymm2, ymm3/m256	EVEX
VDIVPH	ymm1, ymm2, ymm3/m256	EVEX
VDIVPS	ymm1, ymm2, ymm3/m256	EVEX
VDPHPS	xmm1, xmm2, xmm3/m128	EVEX
VDPHPS	ymm1, ymm2, ymm3/m256	EVEX
VDPHPS	zmm1, zmm2, zmm3/m512	EVEX
VFCMADDCPH	ymm1, ymm2, ymm3/m256	EVEX
VFCMULCPH	ymm1, ymm2, ymm3/m256	EVEX
VFIXUPIMMPD	ymm1, ymm2, ymm3/m256, imm8	EVEX
VFIXUPIMMPS	ymm1, ymm2, ymm3/m256, imm8	EVEX
VFMADD132NEPBF16	xmm1, xmm2, xmm3/m128	EVEX
VFMADD132NEPBF16	ymm1, ymm2, ymm3/m256	EVEX
VFMADD132NEPBF16	zmm1, zmm2, zmm3/m512	EVEX
VFMADD132PD	ymm1, ymm2, ymm3/m256	EVEX
VFMADD132PH	ymm1, ymm2, ymm3/m256	EVEX
VFMADD132PS	ymm1, ymm2, ymm3/m256	EVEX
VFMADD213NEPBF16	xmm1, xmm2, xmm3/m128	EVEX
VFMADD213NEPBF16	ymm1, ymm2, ymm3/m256	EVEX
VFMADD213NEPBF16	zmm1, zmm2, zmm3/m512	EVEX
VFMADD213PD	ymm1, ymm2, ymm3/m256	EVEX
VFMADD213PH	ymm1, ymm2, ymm3/m256	EVEX
VFMADD213PS	ymm1, ymm2, ymm3/m256	EVEX
VFMADD231NEPBF16	xmm1, xmm2, xmm3/m128	EVEX
VFMADD231NEPBF16	ymm1, ymm2, ymm3/m256	EVEX
VFMADD231NEPBF16	zmm1, zmm2, zmm3/m512	EVEX
VFMADD231PD	ymm1, ymm2, ymm3/m256	EVEX
VFMADD231PH	ymm1, ymm2, ymm3/m256	EVEX

Table continued on next page...

CPUID: AVX10.2	OPERANDS	ENCSPACE (contd.)
VFNMADD132NEPBF16	zmm1, zmm2, zmm3/m512	EVEX
VFNMADD132PD	ymm1, ymm2, ymm3/m256	EVEX
VFNMADD132PH	ymm1, ymm2, ymm3/m256	EVEX
VFNMADD132PS	ymm1, ymm2, ymm3/m256	EVEX
VFNMADD213NEPBF16	xmm1, xmm2, xmm3/m128	EVEX
VFNMADD213NEPBF16	ymm1, ymm2, ymm3/m256	EVEX
VFNMADD213NEPBF16	zmm1, zmm2, zmm3/m512	EVEX
VFNMADD213PD	ymm1, ymm2, ymm3/m256	EVEX
VFNMADD213PH	ymm1, ymm2, ymm3/m256	EVEX
VFNMADD213PS	ymm1, ymm2, ymm3/m256	EVEX
VFNMADD231NEPBF16	xmm1, xmm2, xmm3/m128	EVEX
VFNMADD231NEPBF16	ymm1, ymm2, ymm3/m256	EVEX
VFNMADD231NEPBF16	zmm1, zmm2, zmm3/m512	EVEX
VFNMADD231PD	ymm1, ymm2, ymm3/m256	EVEX
VFNMADD231PH	ymm1, ymm2, ymm3/m256	EVEX
VFNMADD231PS	ymm1, ymm2, ymm3/m256	EVEX
VFNMSUB132NEPBF16	xmm1, xmm2, xmm3/m128	EVEX
VFNMSUB132NEPBF16	ymm1, ymm2, ymm3/m256	EVEX
VFNMSUB132NEPBF16	zmm1, zmm2, zmm3/m512	EVEX
VFNMSUB132PD	ymm1, ymm2, ymm3/m256	EVEX
VFNMSUB132PH	ymm1, ymm2, ymm3/m256	EVEX
VFNMSUB132PS	ymm1, ymm2, ymm3/m256	EVEX
VFNMSUB213NEPBF16	xmm1, xmm2, xmm3/m128	EVEX
VFNMSUB213NEPBF16	ymm1, ymm2, ymm3/m256	EVEX
VFNMSUB213NEPBF16	zmm1, zmm2, zmm3/m512	EVEX
VFNMSUB213PD	ymm1, ymm2, ymm3/m256	EVEX
VFNMSUB213PH	ymm1, ymm2, ymm3/m256	EVEX
VFNMSUB213PS	ymm1, ymm2, ymm3/m256	EVEX
VFNMSUB231NEPBF16	xmm1, xmm2, xmm3/m128	EVEX
VFNMSUB231NEPBF16	ymm1, ymm2, ymm3/m256	EVEX
VFNMSUB231NEPBF16	zmm1, zmm2, zmm3/m512	EVEX
VFNMSUB231PD	ymm1, ymm2, ymm3/m256	EVEX
VFNMSUB231PH	ymm1, ymm2, ymm3/m256	EVEX
VFNMSUB231PS	ymm1, ymm2, ymm3/m256	EVEX
VFPCLASSPBF16	k1, xmm2/m128, imm8	EVEX
VFPCLASSPBF16	k1, ymm2/m256, imm8	EVEX
VFPCLASSPBF16	k1, zmm2/m512, imm8	EVEX
VGETEXPPBF16	xmm1, xmm2/m128	EVEX
VGETEXPPBF16	ymm1, ymm2/m256	EVEX
VGETEXPPBF16	zmm1, zmm2/m512	EVEX
VGETEXPPD	ymm1, ymm2/m256	EVEX

Table continued on next page...

CPUID: AVX10.2	OPERANDS	ENCSPACE (contd.)
VGETEXPPH	ymm1, ymm2/m256	EVEX
VGETEXPPS	ymm1, ymm2/m256	EVEX
VGETMANTPBF16	xmm1, xmm2/m128, imm8	EVEX
VGETMANTPBF16	ymm1, ymm2/m256, imm8	EVEX
VGETMANTPBF16	zmm1, zmm2/m512, imm8	EVEX
VGETMANTPD	ymm1, ymm2/m256, imm8	EVEX
VGETMANTPH	ymm1, ymm2/m256, imm8	EVEX
VGETMANTPS	ymm1, ymm2/m256, imm8	EVEX
VMAXPBF16	xmm1, xmm2, xmm3/m128	EVEX
VMAXPBF16	ymm1, ymm2, ymm3/m256	EVEX
VMAXPBF16	zmm1, zmm2, zmm3/m512	EVEX
VMAXPD	ymm1, ymm2, ymm3/m256	EVEX
VMAXPH	ymm1, ymm2, ymm3/m256	EVEX
VMAXPS	ymm1, ymm2, ymm3/m256	EVEX
VMINMAXNEPBF16	xmm1, xmm2, xmm3/m128, imm8	EVEX
VMINMAXNEPBF16	ymm1, ymm2, ymm3/m256, imm8	EVEX
VMINMAXNEPBF16	zmm1, zmm2, zmm3/m512, imm8	EVEX
VMINMAXPD	xmm1, xmm2, xmm3/m128, imm8	EVEX
VMINMAXPD	ymm1, ymm2, ymm3/m256, imm8	EVEX
VMINMAXPD	zmm1, zmm2, zmm3/m512, imm8	EVEX
VMINMAXPH	xmm1, xmm2, xmm3/m128, imm8	EVEX
VMINMAXPH	ymm1, ymm2, ymm3/m256, imm8	EVEX
VMINMAXPH	zmm1, zmm2, zmm3/m512, imm8	EVEX
VMINMAXPS	xmm1, xmm2, xmm3/m128, imm8	EVEX
VMINMAXPS	ymm1, ymm2, ymm3/m256, imm8	EVEX
VMINMAXPS	zmm1, zmm2, zmm3/m512, imm8	EVEX
VMINMAXSD	xmm1, xmm2, xmm3/m64, imm8	EVEX
VMINMAXSH	xmm1, xmm2, xmm3/m16, imm8	EVEX
VMINMAXSS	xmm1, xmm2, xmm3/m32, imm8	EVEX
VMINPBF16	xmm1, xmm2, xmm3/m128	EVEX
VMINPBF16	ymm1, ymm2, ymm3/m256	EVEX
VMINPBF16	zmm1, zmm2, zmm3/m512	EVEX
VMINPD	ymm1, ymm2, ymm3/m256	EVEX
VMINPH	ymm1, ymm2, ymm3/m256	EVEX
VMINPS	ymm1, ymm2, ymm3/m256	EVEX
VMOVD	xmm1, xmm2/m32	EVEX
VMOVD	xmm1/m32, xmm2	EVEX
VMOVW	xmm1, xmm2/m16	EVEX
VMOVW	xmm1/m16, xmm2	EVEX
VMPSADBW	xmm1, xmm2, xmm3/m128, imm8	EVEX
VMPSADBW	ymm1, ymm2, ymm3/m256, imm8	EVEX

Table continued on next page...

CPUID: AVX10.2	OPERANDS	ENCSPACE (contd.)
VMPSADBW	zmm1, zmm2, zmm3/m512, imm8	EVEX
VMULNEPBF16	xmm1, xmm2, xmm3/m128	EVEX
VMULNEPBF16	ymm1, ymm2, ymm3/m256	EVEX
VMULNEPBF16	zmm1, zmm2, zmm3/m512	EVEX
VMULPD	ymm1, ymm2, ymm3/m256	EVEX
VMULPH	ymm1, ymm2, ymm3/m256	EVEX
VMULPS	ymm1, ymm2, ymm3/m256	EVEX
VPDPBSSD	xmm1, xmm2, xmm3/m128	EVEX
VPDPBSSD	ymm1, ymm2, ymm3/m256	EVEX
VPDPBSSD	zmm1, zmm2, zmm3/m512	EVEX
VPDPBSSDS	xmm1, xmm2, xmm3/m128	EVEX
VPDPBSSDS	ymm1, ymm2, ymm3/m256	EVEX
VPDPBSSDS	zmm1, zmm2, zmm3/m512	EVEX
VPDPBSUD	xmm1, xmm2, xmm3/m128	EVEX
VPDPBSUD	ymm1, ymm2, ymm3/m256	EVEX
VPDPBSUD	zmm1, zmm2, zmm3/m512	EVEX
VPDPBSUDS	xmm1, xmm2, xmm3/m128	EVEX
VPDPBSUDS	ymm1, ymm2, ymm3/m256	EVEX
VPDPBSUDS	zmm1, zmm2, zmm3/m512	EVEX
VPDPBUUD	xmm1, xmm2, xmm3/m128	EVEX
VPDPBUUD	ymm1, ymm2, ymm3/m256	EVEX
VPDPBUUD	zmm1, zmm2, zmm3/m512	EVEX
VPDPBUUDS	xmm1, xmm2, xmm3/m128	EVEX
VPDPBUUDS	ymm1, ymm2, ymm3/m256	EVEX
VPDPBUUDS	zmm1, zmm2, zmm3/m512	EVEX
VPDPWSUD	xmm1, xmm2, xmm3/m128	EVEX
VPDPWSUD	ymm1, ymm2, ymm3/m256	EVEX
VPDPWSUD	zmm1, zmm2, zmm3/m512	EVEX
VPDPWSUDS	xmm1, xmm2, xmm3/m128	EVEX
VPDPWSUDS	ymm1, ymm2, ymm3/m256	EVEX
VPDPWSUDS	zmm1, zmm2, zmm3/m512	EVEX
VPDPWUSD	xmm1, xmm2, xmm3/m128	EVEX
VPDPWUSD	ymm1, ymm2, ymm3/m256	EVEX
VPDPWUSD	zmm1, zmm2, zmm3/m512	EVEX
VPDPWUSDS	xmm1, xmm2, xmm3/m128	EVEX
VPDPWUSDS	ymm1, ymm2, ymm3/m256	EVEX
VPDPWUSDS	zmm1, zmm2, zmm3/m512	EVEX
VPDPWUUD	xmm1, xmm2, xmm3/m128	EVEX
VPDPWUUD	ymm1, ymm2, ymm3/m256	EVEX
VPDPWUUD	zmm1, zmm2, zmm3/m512	EVEX
VPDPWUUDS	xmm1, xmm2, xmm3/m128	EVEX

Table continued on next page...

CPUID: AVX10.2	OPERANDS	ENCSPACE (contd.)
VPDPWUUDS	ymm1, ymm2, ymm3/m256	EVEX
VPDPWUUDS	zmm1, zmm2, zmm3/m512	EVEX
VRANGEPD	ymm1, ymm2, ymm3/m256, imm8	EVEX
VRANGEPS	ymm1, ymm2, ymm3/m256, imm8	EVEX
VRCPPBF16	xmm1, xmm2/m128	EVEX
VRCPPBF16	ymm1, ymm2/m256	EVEX
VRCPPBF16	zmm1, zmm2/m512	EVEX
VREDUCENEPBF16	xmm1, xmm2/m128, imm8	EVEX
VREDUCENEPBF16	ymm1, ymm2/m256, imm8	EVEX
VREDUCENEPBF16	zmm1, zmm2/m512, imm8	EVEX
VREDUCEPD	ymm1, ymm2/m256, imm8	EVEX
VREDUCEPH	ymm1, ymm2/m256, imm8	EVEX
VREDUCEPS	ymm1, ymm2/m256, imm8	EVEX
VRNDSCALENEPBF16	xmm1, xmm2/m128, imm8	EVEX
VRNDSCALENEPBF16	ymm1, ymm2/m256, imm8	EVEX
VRNDSCALENEPBF16	zmm1, zmm2/m512, imm8	EVEX
VRNDSCALEPD	ymm1, ymm2/m256, imm8	EVEX
VRNDSCALEPH	ymm1, ymm2/m256, imm8	EVEX
VRNDSCALEPS	ymm1, ymm2/m256, imm8	EVEX
VRSQRTPBF16	xmm1, xmm2/m128	EVEX
VRSQRTPBF16	ymm1, ymm2/m256	EVEX
VRSQRTPBF16	zmm1, zmm2/m512	EVEX
VSCALEFPBF16	xmm1, xmm2, xmm3/m128	EVEX
VSCALEFPBF16	ymm1, ymm2, ymm3/m256	EVEX
VSCALEFPBF16	zmm1, zmm2, zmm3/m512	EVEX
VSCALEFPD	ymm1, ymm2, ymm3/m256	EVEX
VSCALEFPH	ymm1, ymm2, ymm3/m256	EVEX
VSCALEFPS	ymm1, ymm2, ymm3/m256	EVEX
VSQRTNEPBF16	xmm1, xmm2/m128	EVEX
VSQRTNEPBF16	ymm1, ymm2/m256	EVEX
VSQRTNEPBF16	zmm1, zmm2/m512	EVEX
VSQRTPD	ymm1, ymm2/m256	EVEX
VSQRTPH	ymm1, ymm2/m256	EVEX
VSQRTPS	ymm1, ymm2/m256	EVEX
VSUBNEPBF16	xmm1, xmm2, xmm3/m128	EVEX
VSUBNEPBF16	ymm1, ymm2, ymm3/m256	EVEX
VSUBNEPBF16	zmm1, zmm2, zmm3/m512	EVEX
VSUBPD	ymm1, ymm2, ymm3/m256	EVEX
VSUBPH	ymm1, ymm2, ymm3/m256	EVEX
VSUBPS	ymm1, ymm2, ymm3/m256	EVEX
VUCOMXSD	xmm1, xmm2/m64	EVEX

Table continued on next page...

CPUID: AVX10.2	OPERANDS	ENCSPACE (<i>contd.</i>)
VUCOMXSH	xmm1, xmm2/m16	EVEX
VUCOMXSS	xmm1, xmm2/m32	EVEX

Chapter 7

INTEL® AVX10 BF16 INSTRUCTIONS

7.1 VADDNEPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.66.MAP5.W0 58 /r VADDNEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.66.MAP5.W0 58 /r VADDNEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.66.MAP5.W0 58 /r VADDNEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2

7.1.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

7.1.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

This instruction adds packed BF16 values from source operands and stores the packed BF16 result in the destination operand. The destination elements are updated according to the writemask. This instruction does not generate floating point exceptions and does not consult or update MXCSR. Denormal BF16 input operands are treated as zeros (DAZ) and denormal BF16 outputs are flushed to zero (FTZ). Rounding Mode is always RNE.

7.1.3 OPERATION

```

1 VADDNEPBF16 (EVEX encoded versions) when src2 operand is a register
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         DEST.bf16[j] := SRC1.bf16[j] + SRC2.bf16[j] //DAZ, FTZ, RNE, SAE
7     ELSE IF *zeroing*:
8         DEST.bf16[j] := 0
9         // else dest.bf16[j] remains unchanged
10
11 DEST[MAX_VL-1:VL] := 0

```

```

1 VADDNEPBF16 (EVEX encoded versions) when src2 operand is a memory source
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         IF EVEX.b == 1:
7             DEST.bf16[j] := SRC1.bf16[j] + SRC2.bf16[0] //DAZ, FTZ, RNE, SAE
8         ELSE:
9             DEST.bf16[j] := SRC1.bf16[j] + SRC2.bf16[j] //DAZ, FTZ, RNE, SAE
10
11     ELSE IF *zeroing*:
12         DEST.bf16[j] := 0
13         // else dest.bf16[j] remains unchanged
14
15 DEST[MAX_VL-1:VL] := 0

```

7.1.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VADDNEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VADDNEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VADDNEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2

7.2 VCMPPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.F2.0F3A.W0 C2 /r /ib VCMPPBF16 k1{k2}, xmm2, xmm3/m128/m16bcst, imm8	A	V/V	AVX10.2
EVEX.256.F2.0F3A.W0 C2 /r /ib VCMPPBF16 k1{k2}, ymm2, ymm3/m256/m16bcst, imm8	A	V/V	AVX10.2
EVEX.512.F2.0F3A.W0 C2 /r /ib VCMPPBF16 k1{k2}, zmm2, zmm3/m512/m16bcst, imm8	A	V/V	AVX10.2

7.2.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	IMM8(r)

7.2.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

This instruction compares packed BF16 values from source operands and stores the result in the destination mask operand. The comparison predicate operand (immediate byte bits 4:0) specifies the type of comparison performed on each of the pairs of packed values. The destination elements are updated according to the writemask. This instruction does not generate floating point exceptions and does not consult or update MXCSR. Denormal BF16 input operands are treated as zeros (DAZ)

7.2.3 OPERATION

```
1 CASE (imm8 & 0x1F) OF
2 0: CMP_OPERATOR := EQ_OQ;
3 1: CMP_OPERATOR := LT_OS;
4 2: CMP_OPERATOR := LE_OS;
5 3: CMP_OPERATOR := UNORD_Q;
6 4: CMP_OPERATOR := NEQ_UQ;
7 5: CMP_OPERATOR := NLT_US;
8 6: CMP_OPERATOR := NLE_US;
9 7: CMP_OPERATOR := ORD_Q;
10 8: CMP_OPERATOR := EQ_UQ;
11 9: CMP_OPERATOR := NGE_US;
12 10: CMP_OPERATOR := NGT_US;
13 11: CMP_OPERATOR := FALSE_OQ;
14 12: CMP_OPERATOR := NEQ_OQ;
15 13: CMP_OPERATOR := GE_OS;
16 14: CMP_OPERATOR := GT_OS;
17 15: CMP_OPERATOR := TRUE_UQ;
18 16: CMP_OPERATOR := EQ_OS;
19 17: CMP_OPERATOR := LT_OQ;
20 18: CMP_OPERATOR := LE_OQ;
21 19: CMP_OPERATOR := UNORD_S;
22 20: CMP_OPERATOR := NEQ_US;
23 21: CMP_OPERATOR := NLT_UQ;
24 22: CMP_OPERATOR := NLE_UQ;
25 23: CMP_OPERATOR := ORD_S;
26 24: CMP_OPERATOR := EQ_US;
27 25: CMP_OPERATOR := NGE_UQ;
28 26: CMP_OPERATOR := NGT_UQ;
29 27: CMP_OPERATOR := FALSE_OS;
30 28: CMP_OPERATOR := NEQ_OS;
31 29: CMP_OPERATOR := GE_OQ;
32 30: CMP_OPERATOR := GT_OQ;
33 31: CMP_OPERATOR := TRUE_US;
34 ESAC
```

```

1  VCMPPBF16 (EVEX encoded versions)
2  (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4  FOR j := 0 TO KL-1:
5      IF k2[j] OR *no writemask*:
6          IF EVEX.b == 1:
7              tsrc2 := SRC2.bf16[0]
8          ELSE:
9              tsrc2 := SRC2.bf16[j]
10             DEST.bit[j] := SRC1.bf16[j] CMP_OPERATOR tsrc2 //DAZ, SAE
11         ELSE *zero masking only*:
12             DEST.bit[j] := 0
13
14     DEST[MAX_KL-1:KL] := 0

```

7.2.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCMPPBF16 k1, xmm2, xmm3/m128, imm8	E4	N/A	AVX10.2
VCMPPBF16 k1, ymm2, ymm3/m256, imm8	E4	N/A	AVX10.2
VCMPPBF16 k1, zmm2, zmm3/m512, imm8	E4	N/A	AVX10.2

7.3 VCOMSBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.LLIG.66.MAP5.W0 2F /r VCOMSBF16 xmm1, xmm2/m16	A	V/V	AVX10.2

7.3.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	SCALAR	MODRM.REG(r)	MODRM.R/M(r)	N/A	N/A

7.3.2 DESCRIPTION

Note:

For SCALAR instructions with a CPUID feature flag specifying Intel® AVX10, the programmer can skip the check for the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and usually determines the set of instructions available to the programmer listed in the above opcode table. However, scalar instructions are architected without the notion of vector length, so this step does not strictly apply.

Compares the half-precision floating-point values in the low word of operand 1 (first operand) and operand 2 (second operand), and sets the ZF, PF, and CF flags in the EFLAGS register according to the result (unordered, greater than, less than, or equal). The OF, SF and AF flags in the EFLAGS register are set to 0. The unordered result is returned if either source operand is a NaN (QNaN or SNaN). This instruction does not generate floating point exceptions and does not consult or update MXCSR. Denormal BF16 input operands are treated as zeros (DAZ)

Operand 1 is an XMM register; operand 2 can be an XMM register or a 16-bit memory location.

EVEX.vvvv are reserved and must be 1111b, otherwise instructions will #UD

7.3.3 OPERATION

```

1 VCOMSBF16
2
3 RESULT := UnorderedCompare(SRC1.bf16[0],SRC2.bf16[0])
4 IF RESULT is UNORDERED:
5     ZF, PF, CF := 1, 1, 1
6 ELIF RESULT is GREATER_THAN:
7     ZF, PF, CF := 0, 0, 0
8 ELIF RESULT is LESS_THAN:
9     ZF, PF, CF := 0, 0, 1
10 ELSE: // RESULT is EQUALS
11     ZF, PF, CF := 1, 0, 0
12
13 OF, AF, SF := 0, 0, 0

```

7.3.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCOMSBF16 xmm1, xmm2/m16	E10NF	N/A	AVX10.2

7.4 VDIVNEPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.66.MAP5.W0 5E /r VDIVNEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.66.MAP5.W0 5E /r VDIVNEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.66.MAP5.W0 5E /r VDIVNEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2

7.4.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

7.4.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

This instruction divides packed BF16 values from the first source operand by the corresponding elements in the second source operand, storing the packed BF16 result in the destination operand. The destination elements are updated according to the writemask. This instruction does not generate floating point exceptions and does not consult or update MXCSR. Denormal BF16 input operands are treated as zeros (DAZ) and denormal BF16 outputs are flushed to zero (FTZ). Rounding Mode is always RNE.

7.4.3 OPERATION

```

1 VDIVNEPBF16 (EVEX encoded versions) when src2 operand is a register
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         DEST.bf16[j] := SRC1.bf16[j] / SRC2.bf16[j] //DAZ, FTZ, RNE, SAE
7     ELSE IF *zeroing*:
8         DEST.bf16[j] := 0
9         // else dest.bf16[j] remains unchanged
10
11 DEST[MAX_VL-1:VL] := 0

```

```

1 VDIVNEPBF16 (EVEX encoded versions) when src2 operand is a memory source
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         IF EVEX.b == 1:
7             DEST.bf16[j] := SRC1.bf16[j] / SRC2.bf16[0] //DAZ, FTZ, RNE, SAE
8         ELSE:
9             DEST.bf16[j] := SRC1.bf16[j] / SRC2.bf16[j] //DAZ, FTZ, RNE, SAE
10
11     ELSE IF *zeroing*:
12         DEST.bf16[j] := 0
13         // else dest.bf16[j] remains unchanged
14
15 DEST[MAX_VL-1:VL] := 0

```

7.4.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VDIVNEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VDIVNEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VDIVNEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2

7.5 VF[,N]M[ADD,SUB][132,213,231]NEPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.NP.MAP6.W0 98 /r VFMADD132NEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP6.W0 98 /r VFMADD132NEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.W0 98 /r VFMADD132NEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.NP.MAP6.W0 A8 /r VFMADD213NEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP6.W0 A8 /r VFMADD213NEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.W0 A8 /r VFMADD213NEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.NP.MAP6.W0 B8 /r VFMADD231NEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP6.W0 B8 /r VFMADD231NEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.W0 B8 /r VFMADD231NEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.NP.MAP6.W0 9A /r VFMSUB132NEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP6.W0 9A /r VFMSUB132NEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.W0 9A /r VFMSUB132NEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.NP.MAP6.W0 AA /r VFMSUB213NEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP6.W0 AA /r VFMSUB213NEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.W0 AA /r VFMSUB213NEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.NP.MAP6.W0 BA /r VFMSUB231NEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2

Continued on next page...

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.MAP6.W0 BA /r VFMSUB231NEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.W0 BA /r VFMSUB231NEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.NP.MAP6.W0 9C /r VFNMADD132NEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP6.W0 9C /r VFNMADD132NEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.W0 9C /r VFNMADD132NEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.NP.MAP6.W0 AC /r VFNMADD213NEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP6.W0 AC /r VFNMADD213NEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.W0 AC /r VFNMADD213NEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.NP.MAP6.W0 BC /r VFNMADD231NEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP6.W0 BC /r VFNMADD231NEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.W0 BC /r VFNMADD231NEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.NP.MAP6.W0 9E /r VFNMSUB132NEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP6.W0 9E /r VFNMSUB132NEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.W0 9E /r VFNMSUB132NEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.NP.MAP6.W0 AE /r VFNMSUB213NEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP6.W0 AE /r VFNMSUB213NEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.W0 AE /r VFNMSUB213NEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2

Continued on next page...

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.NP.MAP6.W0 BE /r VFNMSUB231NEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP6.W0 BE /r VFNMSUB231NEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.W0 BE /r VFNMSUB231NEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2

7.5.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

7.5.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

This instruction performs a packed multiply-add, multiply-subtract, negated multiply-add or negated multiply-subtract computation on BF16 values using three source operands and writes the results in the destination operand. The destination operand is also the first source operand. The “N” (negated) forms of this instruction add/subtract the remaining operand to/from the negated infinite precision intermediate product. The notation “132”, “213” and “231” indicate the use of the operands in $\pm A * B - C$, where each digit corresponds to the operand number, with the destination being operand 1. The destination elements are updated according to the writemask. This instruction does not generate floating point exceptions and does not consult or update MXCSR. Denormal BF16 input operands are treated as zeros (DAZ) and denormal BF16 outputs are flushed to zero (FTZ). Rounding Mode is always RNE.

Notation	Operands
132	dest = ± dest*src3-src2
231	dest = ± src2*src3-dest
213	dest = ± src2*dest-src3

Table 7.1: VF[,N]M[ADD,SUB][132,213,231]NEPBF16 Notation for Operands

7.5.3 OPERATION

```

1 VF[,N]M[ADD,SUB][132,213,231]NEPBF16 (EVEX encoded versions) when src3 operand is
2 a register
3 (KL, VL) = (8, 128), (16, 256), (32, 512)
4
5 IF *132 form*:
6     a := DEST
7     b := SRC3
8     c := SRC2
9 ELIF *213 form*:
10    a := SRC2
11    b := DEST
12    c := SRC3
13 ELIF *231 form*:
14    a := SRC2
15    b := SRC3
16    c := DEST
17
18 IF *negative form*:
19     a := -a
20
21 IF *add form*:
22     OP := +
23 ELIF *sub form*:
24     OP := -
25
26 FOR j := 0 TO KL-1:
27     IF k1[j] OR *no writemask*:
28         //DAZ, FTZ, SAE
29         DEST.bf16[j] := RoundFPControl_RNE(a.bf16[j]*b.bf16[j] OP c.bf16[j])
30     ELSE IF *zeroing*:
31         DEST.bf16[j] := 0
32     // else dest.bf16[j] remains unchanged
33
34 DEST[MAX_VL-1:VL] := 0

```

```

1 VF[,N]M[ADD,SUB][132,213,231]NEPBF16 (EVEX encoded versions) when src3 operand
2 is a memory source
3 (KL, VL) = (8, 128), (16, 256), (32, 512)
4
5 IF *132 form*:
6     a := DEST
7     b := SRC3
8     c := SRC2
9 ELIF *213 form*:
10    a := SRC2
11    b := DEST
12    c := SRC3
13 ELIF *231 form*:
14    a := SRC2
15    b := SRC3
16    c := DEST
17
18 IF *negative form*:
19     a := -a
20
21 IF *add form*:
22     OP := +
23 ELIF *sub form*:
24     OP := -
25
26 FOR j := 0 TO KL-1:
27     IF k1[j] OR *no writemask*:
28         IF EVEX.b == 1:
29             //DAZ, FTZ, SAE
30             DEST.bf16[j] := RoundFPControl_RNE(a.bf16[j]*b.bf16[j] OP c.bf16[0])
31         ELSE:
32             //DAZ, FTZ, SAE
33             DEST.bf16[j] := RoundFPControl_RNE(a.bf16[j]*b.bf16[j] OP c.bf16[j])
34         ELSE IF *zeroing*:
35             DEST.bf16[j] := 0
36         // else dest.bf16[j] remains unchanged
37
38 DEST[MAX_VL-1:VL] := 0

```

7.5.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD132NEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VFMADD132NEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VFMADD132NEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VFMADD213NEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VFMADD213NEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VFMADD213NEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VFMADD231NEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VFMADD231NEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VFMADD231NEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VFMSUB132NEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VFMSUB132NEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VFMSUB132NEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VFMSUB213NEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VFMSUB213NEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VFMSUB213NEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2

Continued on next page...

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUB231NEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VFMSUB231NEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VFMSUB231NEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VFNMADD132NEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VFNMADD132NEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VFNMADD132NEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VFNMADD213NEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VFNMADD213NEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VFNMADD213NEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VFNMADD231NEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VFNMADD231NEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VFNMADD231NEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VFNMSUB132NEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VFNMSUB132NEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VFNMSUB132NEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2

Continued on next page...

Instruction	Exception Type	Arithmetic Flags	CPUID
VFNMSUB213NEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VFNMSUB213NEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VFNMSUB213NEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VFNMSUB231NEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VFNMSUB231NEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VFNMSUB231NEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2

7.6 VFPCLASSPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.F2.0F3A.W0 66 /r /ib VFPCLASSPBF16 k1{k2}, xmm2/m128/m16bcst, imm8	A	V/V	AVX10.2
EVEX.256.F2.0F3A.W0 66 /r /ib VFPCLASSPBF16 k1{k2}, ymm2/m256/m16bcst, imm8	A	V/V	AVX10.2
EVEX.512.F2.0F3A.W0 66 /r /ib VFPCLASSPBF16 k1{k2}, zmm2/m512/m16bcst, imm8	A	V/V	AVX10.2

7.6.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	IMM8(r)	N/A

7.6.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

The VFPCLASSPBF16 instruction checks the packed bfloat16 floating point values for special categories, specified by the set bits in the imm8 byte. Each set bit in imm8 specifies a category of floating-point values that the input data element is classified against. The classified results of all specified categories of an input value are ORed together to form the final boolean result for the input element. The result of each element is written to the corresponding bit in a mask register k2 according to the writemask k1. Bits [MAX_KL-1:32/16/8] of the destination are cleared. This instruction does not generate floating point exceptions and does not consult or update MXCSR. Denormal BF16 input operands are treated as zeros (DAZ).

The source operand is a ZMM/YMM/XMM register, a 512/256/128-bit memory location, or a 512/256/128-bit vector broadcasted from a 16-bit memory location.

EVEX.vvvv is reserved and must be 1111b otherwise instructions will #UD.

7.6.3 OPERATION

```

1 def check_fp_class_bf16(src, imm8):
2     negative := src[15]
3     exponent_all_ones := (src[14:7] == 0xFF)
4     exponent_all_zeros := (src[14:7] == 0)
5     IF(exponent_all_zeros):
6         mantissa_all_zeros:=1
7     ELSEIF (src[6:0] == 0 OR exponent_all_zeros):
8         mantissa_all_zeros := 1
9     zero := exponent_all_zeros and mantissa_all_zeros
10    signaling_bit := src[6]
11
12    snan := exponent_all_ones and not(mantissa_all_zeros) and not(signaling_bit)
13    qnan := exponent_all_ones and not(mantissa_all_zeros) and signaling_bit
14    positive_zero := not(negative) and exponent_all_zeros and mantissa_all_zeros
15    negative_zero := negative and exponent_all_zeros and mantissa_all_zeros
16    positive_infinity := not(negative) and exponent_all_ones and mantissa_all_zeros
17    negative_infinity := negative and exponent_all_ones and mantissa_all_zeros
18    denormal := exponent_all_zeros and not(mantissa_all_zeros)
19    finite_negative := negative and not(exponent_all_ones) and not(zero)
20
21    return (imm8[0] and qnan) OR
22           (imm8[1] and positive_zero) OR
23           (imm8[2] and negative_zero) OR
24           (imm8[3] and positive_infinity) OR
25           (imm8[4] and negative_infinity) OR
26           (denormal) OR
27           (imm8[6] and finite_negative) OR
28           (imm8[7] and snan)

```

```

1 VFPCLASSPBF16 destk2k1, src, imm8
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         IF SRC is memory and (EVEX.b == 1):
7             tsrc := SRC.bf16[0]
8         ELSE:
9             tsrc := SRC.bf16[j]
10        DEST.bit[j] := check_fp_class_bf16_(tsrc, imm8)
11    ELSE // zero masking only
12        DEST.bit[j] := 0
13
14 DEST[MAX_KL-1:KL] := 0

```

7.6.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFPCLASSPBF16 k1, xmm2/m128, imm8	E4	N/A	AVX10.2
VFPCLASSPBF16 k1, ymm2/m256, imm8	E4	N/A	AVX10.2
VFPCLASSPBF16 k1, zmm2/m512, imm8	E4	N/A	AVX10.2

7.7 VGETEXPPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.NP.MAP6.W0 42 /r VGETEXPPBF16 xmm1{k1}{z}, xmm2/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP6.W0 42 /r VGETEXPPBF16 ymm1{k1}{z}, ymm2/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.W0 42 /r VGETEXPPBF16 zmm1{k1}{z}, zmm2/m512/m16bcst	A	V/V	AVX10.2

7.7.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

7.7.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

Extracts the biased exponents from the each bfloat16 data element of the source operand (the second operand) as unbiased signed integer value. Each integer value of the unbiased exponent is converted to bfloat16 FP value and written to the corresponding bfloat16 elements of the destination operand (the first operand) as bfloat16 FP numbers.

The destination operand is a ZMM/YMM/XMM register and updated under the writemask. The source operand can be a ZMM/YMM/XMM register, a 512/256/128-bit memory location, or a 512/256/128-bit vector broadcasted from a 16-bit memory location.

EVEX.vvvv is reserved and must be 1111b, otherwise instructions will #UD.

Each GETEXP operation converts the exponent value into a bfloat16 number.

The formula is: $GETEXP(x) = \text{floor}(\log_2(|x|))$

Notation $\text{floor}(x)$ stands for maximal integer not exceeding real number x .

This instruction does not generate floating point exceptions and does not consult or update MXCSR. Denormal BF16 input operands are treated as zeros (DAZ)

7.7.3 OPERATION

```

1 def getexp_bf16(src):
2     IF (*src is nan*):
3         return QNAN(src)
4     ELIF (*src is positive infinity*):
5         return INF
6     ELIF (*src is denormal or zero*):
7         return -INF
8     ELSE:
9         tmp := SAR(src, 7) //shift arithmetic right
10        tmp := tmp - 127 //subtract bias
11        return convert_integer_to_bf16(tmp)

```

```

1 VGETEXPPBF16 destk1, src
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         IF SRC is memory and (EVEX.b == 1):
7             tsrc := SRC.bf16[0]
8         ELSE:
9             tsrc := SRC.bf16[j]
10        DEST.bf16[j] := getexp_bf16(tsrc)
11    ELSE IF *zeroing*:
12        DEST.bf16[j] := 0
13    // else dest.bf16[j] remains unchanged
14
15 DEST[MAX_VL-1:VL] := 0

```

7.7.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VGETEXPPBF16 xmm1, xmm2/m128	E4	N/A	AVX10.2
VGETEXPPBF16 ymm1, ymm2/m256	E4	N/A	AVX10.2

Continued on next page...

Instruction	Exception Type	Arithmetic Flags	CPUID
VGETEXPPBF16 zmm1, zmm2/m512	E4	N/A	AVX10.2

7.8 VGETMANTPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.F2.0F3A.W0 26 /r /ib VGETMANTPBF16 xmm1{k1}{z}, xmm2/m128/m16bcst, imm8	A	V/V	AVX10.2
EVEX.256.F2.0F3A.W0 26 /r /ib VGETMANTPBF16 ymm1{k1}{z}, ymm2/m256/m16bcst, imm8	A	V/V	AVX10.2
EVEX.512.F2.0F3A.W0 26 /r /ib VGETMANTPBF16 zmm1{k1}{z}, zmm2/m512/m16bcst, imm8	A	V/V	AVX10.2

7.8.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	IMM8(r)	N/A

7.8.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

Convert bfloat16 floating values in the source operand (the second operand) to bfloat16 FP values with the mantissa normalization specified by the imm8. The converted results are written to the destination operand (the first operand) using writemask k1. The normalized mantissa is specified by interv (imm8[1:0]) and the sign control (sc) is specified by bits 3:2 of the immediate byte.

The destination operand is a ZMM/YMM/XMM register updated under the writemask. The source operand can be a ZMM/YMM/XMM register, a 512/256/128-bit memory location, or a 512/256/128-bit vector broadcasted from a 16-bit memory location.

For each input SP FP value x, The conversion operation is: $\text{GetMant}(x) = \pm 2^k |x.\text{significand}|$ where:

$$1 \leq |x.\text{significand}| < 2$$

Unbiased exponent k depends on the interval range defined by `interv` and whether the exponent of the source is even or odd. The sign of the final result is determined by `sc` and the source sign.

If `interv` \neq 0 then $k = -1$, otherwise $k = 0$.

Each converted bfloat16 FP result is encoded according to the sign control, the unbiased exponent k (adding bias) and a mantissa normalized to the range specified by `interv`.

This instruction is writemasked, so only those elements with the corresponding bit set in vector mask register `k1` are computed and stored into `zmm1`.

This instruction does not generate floating point exceptions and does not consult or update MXCSR. Denormal BF16 input operands are treated as zeros (DAZ)

Note: `EVEX.vvvv` is reserved and must be 1111b; otherwise instructions will #UD.

7.8.3 OPERATION

```
1 def getmant_bf16(src, sign_control, normalization_interval):
2     dst.sign := sign_control[0] ? 0 : src.sign
3     signed_one := sign_control[0] ? +1.0 : -1.0
4     dst.exp := src.exp
5     dst.fraction := src.fraction
6     bias := 127
7
8     IF (*src is nan*):
9         return QNaN(src)
10    ELIF (*src is positive zero or positive infinity*):
11        return 1.0
12    ELIF (*src is negative*):
13        IF (*src is zero*):
14            return signed_one
15        ELIF (*src is infinity*):
16            IF (sign_control[1]):
17                return QNaN_Indefinite
18            ELSE:
19                return signed_one
20        ELIF (sign_control[1]):
21            return QNaN_Indefinite
22    IF (*src is denormal*):
23        dst.fraction := 0
24
25    unbiased_exp := dst.exp - bias
26    odd_exp := unbiased_exp[0]
27    signaling_bit := dst.fraction[6]
28
29    IF (normalization_interval := 0b00):
30        dst.exp := bias
31    ELIF (normalization_interval := 0b01):
32        dst.exp := odd_exp ? bias-1 : bias
33    ELIF (normalization_interval := 0b10):
34        dst.exp := bias-1
35    ELIF (normalization_interval := 0b11):
36        dst.exp := signaling_bit ? bias-1 : bias
37
38    return dst
```

```

1 VGETMANTPBF16 destk1, src, imm8
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 sign_control := imm8[3:2]
5 normalization_interval := imm8[1:0]
6
7 FOR j := 0 TO KL-1:
8     IF k1[j] OR *no writemask*:
9         IF SRC is memory and (EVEX.b == 1):
10            tsrc := SRC.bf16[0]
11        ELSE:
12            tsrc := SRC.bf16[j]
13        DEST.bf16[j] := getmant_bf16(tsrc)
14    ELSE IF *zeroing*:
15        DEST.bf16[j] := 0
16    // else dest.bf16[j] remains unchanged
17
18 DEST[MAX_VL-1:VL] := 0

```

7.8.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VGETMANTPBF16 xmm1, xmm2/m128, imm8	E4	N/A	AVX10.2
VGETMANTPBF16 ymm1, ymm2/m256, imm8	E4	N/A	AVX10.2
VGETMANTPBF16 zmm1, zmm2/m512, imm8	E4	N/A	AVX10.2

7.9 VMAXPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.66.MAP5.W0 5F /r VMAXPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.66.MAP5.W0 5F /r VMAXPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.66.MAP5.W0 5F /r VMAXPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2

7.9.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

7.9.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

Performs a SIMD compare of the packed half-precision floating-point values in the first source operand and the second source operand and returns the maximum value for each pair of values to the destination operand.

If the values being compared are both 0.0s (of either sign), the value in the second operand (source operand) is returned. If a value in the second operand is an SNaN, then SNaN is forwarded unchanged to the destination (that is, a QNaN version of the SNaN is not returned).

If only one value is a NaN (SNaN or QNaN) for this instruction, the second operand (source operand), either a NaN or a valid floating-point value, is written to the result. If instead of this behavior, it is required that the NaN source operand (from either the first or second operand) be returned, the action of VMAXNEPBF16 can be emulated using a sequence of instructions, such as, a comparison followed by AND, ANDN and OR. This instruction does not generate floating point exceptions and does not consult or update MXCSR. Denormal BF16 input operands are treated as zeros (DAZ)

EVEX encoded versions: The first source operand (the second operand) is a ZMM/YMM/XMM register. The second source operand can be a ZMM/YMM/XMM register, a 512/256/128-bit memory location or a 512/256/128-bit vector broadcasted from a 16-bit memory location. The destination operand is a ZMM/YMM/XMM register conditionally updated with writemask k1.

7.9.3 OPERATION

```

1  DEFINE MAX(SRC1, SRC2):
2      IF (SRC1 = 0.0) and (SRC2 = 0.0):
3          DEST := SRC2
4      ELSE IF (SRC1 = NaN):
5          DEST := SRC2
6      ELSE IF (SRC2 = NaN):
7          DEST := SRC2
8      ELSE IF (SRC1 > SRC2):
9          DEST := SRC1
10     ELSE:
11         DEST := SRC2
12
13
14  VMAXPFB16 (EVEX encoded versions)
15  (KL, VL) = (8, 128), (16, 256), (32, 512)
16
17  FOR j := 0 TO KL-1:
18      IF k1[j] OR *no writemask*:
19          IF EVEX.b == 1:
20              tsrc2 := SRC2.bf16[0]
21          ELSE:
22              tsrc2 := SRC2.bf16[j]
23          DEST.bf16[j] := MAX(SRC1.bf16[j], tsrc2) //DAZ, SAE
24      ELSE IF *zeroing*:
25          DEST.bf16[j] := 0
26      // else dest.bf16[j] remains unchanged
27
28  DEST[MAX_VL-1:VL] := 0

```

7.9.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMAXPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VMAXPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2

Continued on next page...

Instruction	Exception Type	Arithmetic Flags	CPUID
VMAXPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2

7.10 VMINPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.66.MAP5.W0 5D /r VMINPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.66.MAP5.W0 5D /r VMINPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.66.MAP5.W0 5D /r VMINPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2

7.10.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

7.10.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

Performs a SIMD compare of the packed half-precision floating-point values in the first source operand and the second source operand and returns the minimum value for each pair of values to the destination operand.

If the values being compared are both 0.0s (of either sign), the value in the second operand (source operand) is returned. If a value in the second operand is an SNaN, then SNaN is forwarded unchanged to the destination (that is, a QNaN version of the SNaN is not returned).

If only one value is a NaN (SNaN or QNaN) for this instruction, the second operand (source operand), either a NaN or a valid floating-point value, is written to the result. If instead of this behavior, it is required that the NaN source operand (from either the first or second operand) be returned, the action of VMINPBF16 can be emulated using a sequence of instructions, such as, a comparison followed by AND, ANDN and OR. This instruction does not generate floating point exceptions and does not consult or update MXCSR. Denormal BF16 input operands are treated as zeros (DAZ)

EVEX encoded versions: The first source operand (the second operand) is a ZMM/YMM/XMM register. The second source operand can be a ZMM/YMM/XMM register, a 512/256/128-bit memory location or a 512/256/128-bit vector broadcasted from a 16-bit memory location. The destination operand is a ZMM/YMM/XMM register conditionally updated with writemask k1.

7.10.3 OPERATION

```

1  DEFINE MIN(SRC1, SRC2):
2      IF (SRC1 = 0.0) and (SRC2 = 0.0):
3          DEST := SRC2
4      ELSE IF (SRC1 = NaN):
5          DEST := SRC2
6      ELSE IF (SRC2 = NaN):
7          DEST := SRC2
8      ELSE IF (SRC1 < SRC2):
9          DEST := SRC1
10     ELSE:
11         DEST := SRC2
12
13
14  VMINPFB16 (EVEX encoded versions)
15  (KL, VL) = (8, 128), (16, 256), (32, 512)
16
17  FOR j := 0 TO KL-1:
18      IF k1[j] OR *no writemask*:
19          IF EVEX.b == 1:
20              tsrc2 := SRC2.bf16[0]
21          ELSE:
22              tsrc2 := SRC2.bf16[j]
23          DEST.bf16[j] := MIN(SRC1.bf16[j], tsrc2) //DAZ, SAE
24      ELSE IF *zeroing*:
25          DEST.bf16[j] := 0
26      // else dest.bf16[j] remains unchanged
27
28  DEST[MAX_VL-1:VL] := 0

```

7.10.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMINPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VMINPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2

Continued on next page...

Instruction	Exception Type	Arithmetic Flags	CPUID
VMINPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2

7.11 VMULNEPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.66.MAP5.W0 59 /r VMULNEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.66.MAP5.W0 59 /r VMULNEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.66.MAP5.W0 59 /r VMULNEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2

7.11.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

7.11.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

This instruction multiplies packed BF16 values from source operands and stores the packed BF16 result in the destination operand. The destination elements are updated according to the writemask. This instruction does not generate floating point exceptions and does not consult or update MXCSR. Denormal BF16 input operands are treated as zeros (DAZ) and denormal BF16 outputs are flushed to zero (FTZ). Rounding Mode is always RNE.

7.11.3 OPERATION

```

1 VMULNEPBF16 (EVEX encoded versions) when src2 operand is a register
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5   IF k1[j] OR *no writemask*:
6     DEST.bf16[j] := SRC1.bf16[j] * SRC2.bf16[j] //DAZ, FTZ, RNE, SAE
7   ELSE IF *zeroing*:
8     DEST.bf16[j] := 0
9     // else dest.bf16[j] remains unchanged
10
11 DEST[MAX_VL-1:VL] := 0

```

```

1 VMULNEPBF16 (EVEX encoded versions) when src2 operand is a memory source
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5   IF k1[j] OR *no writemask*:
6     IF EVEX.b == 1:
7       DEST.bf16[j] := SRC1.bf16[j] * SRC2.bf16[0] //DAZ, FTZ, RNE, SAE
8     ELSE:
9       DEST.bf16[j] := SRC1.bf16[j] * SRC2.bf16[j] //DAZ, FTZ, RNE, SAE
10
11   ELSE IF *zeroing*:
12     DEST.bf16[j] := 0
13     // else dest.bf16[j] remains unchanged
14
15 DEST[MAX_VL-1:VL] := 0

```

7.11.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMULNEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VMULNEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VMULNEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2

7.12 VRCPPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.NP.MAP6.W0 4C /r VRCPPBF16 xmm1{k1}{z}, xmm2/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP6.W0 4C /r VRCPPBF16 ymm1{k1}{z}, ymm2/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.W0 4C /r VRCPPBF16 zmm1{k1}{z}, zmm2/m512/m16bcst	A	V/V	AVX10.2

7.12.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

7.12.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

This instruction performs a SIMD computation of the approximate reciprocals of 8/16/32 packed BF16 values in the source operand (the second operand) and stores the packed BF16 results in the destination operand. The maximum relative error for this approximation is less than $2^{-8} + 2^{-14}$. For special cases, see Table 7.2.

Input Value	Result Value	Comments
$0 \leq X < 2^{-126}$	+INF	DAZ
$-2^{-126} < X \leq 0$	-INF	DAZ
$X = +\text{INF}$	+0	
$X = -\text{INF}$	-0	
$X = 2^{-n}$	2^n	
$X = -2^{-n}$	-2^n	

Table 7.2: VRCPPBF16 Special Cases

7.12.3 OPERATION

```

1 VRCPPBF16 destk1, src
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         IF SRC is memory and (EVEX.b == 1):
7             tsrc := SRC.bf16[0]
8         ELSE:
9             tsrc := SRC.bf16[j]
10            DEST.bf16[j] := APPROXIMATE(1.0 / tsrc) //DAZ, FTZ, RNE, SAE
11        ELSE IF *zeroing*:
12            DEST.bf16[j] := 0
13        // else dest.bf16[j] remains unchanged
14
15 DEST[MAX_VL-1:VL] := 0

```

7.12.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VRCPPBF16 xmm1, xmm2/m128	E4	N/A	AVX10.2
VRCPPBF16 ymm1, ymm2/m256	E4	N/A	AVX10.2
VRCPPBF16 zmm1, zmm2/m512	E4	N/A	AVX10.2

7.13 VREDUCENEPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.F2.0F3A.W0 56 /r /ib VREDUCENEPBF16 xmm1{k1}{z}, xmm2/m128/m16bcst, imm8	A	V/V	AVX10.2
EVEX.256.F2.0F3A.W0 56 /r /ib VREDUCENEPBF16 ymm1{k1}{z}, ymm2/m256/m16bcst, imm8	A	V/V	AVX10.2
EVEX.512.F2.0F3A.W0 56 /r /ib VREDUCENEPBF16 zmm1{k1}{z}, zmm2/m512/m16bcst, imm8	A	V/V	AVX10.2

7.13.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	IMM8(r)	N/A

7.13.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

Extracts the reduced argument of bfloat16 Floating-Point values in the first source operand by the number of bits specified in the immediate operand (imm8) and places the result in the destination operand.

The reduced argument extraction is formulated below:

$$zmm1 := zmm2 - (\text{ROUND}(2^M * zmm2)) * 2^{-M};$$

$$\text{Given } zmm2 = 2^{\text{exp}2} * \text{man}2$$

$$\text{Then } 0 \leq |zmm1| < 2^{\text{exp}2 - M - 1}$$

The scaling value M is determined by the imm8[7:4].

The operation is write masked.

7.13.3 OPERATION

```

1 def reduce_bf16_ne(src,imm8):
2   IF (*src is nan*):
3     return QNAN(src)
4   m := imm8[7:4]
5   tmp := 2~m * ROUND(2m * src, RNE) //DAZ, SAE
6   tmp := src - tmp //FTZ, RNE, SAE
7   return tmp

```

```

1 VREDUCENEPBF16 destk1, src, imm8
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5   IF k1[j] OR *no writemask*:
6     IF SRC is memory and (EVEX.b == 1):
7       tsrc := SRC.bf16[0]
8     ELSE:
9       tsrc := SRC.bf16[j]
10    DEST.bf16[j] := reduce_bf16_ne(tsrc, imm8)
11  ELSE IF *zeroing*:
12    DEST.bf16[j] := 0
13  // else dest.bf16[j] remains unchanged
14
15 DEST[MAX_VL-1:VL] := 0

```

7.13.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VREDUCENEPBF16 xmm1, xmm2/m128, imm8	E4	N/A	AVX10.2
VREDUCENEPBF16 ymm1, ymm2/m256, imm8	E4	N/A	AVX10.2
VREDUCENEPBF16 zmm1, zmm2/m512, imm8	E4	N/A	AVX10.2

7.14 VRNDSCALENEPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.F2.0F3A.W0 08 /r /ib VRNDSCALENEPBF16 xmm1{k1}{z}, xmm2/m128/m16bcst, imm8	A	V/V	AVX10.2
EVEX.256.F2.0F3A.W0 08 /r /ib VRNDSCALENEPBF16 ymm1{k1}{z}, ymm2/m256/m16bcst, imm8	A	V/V	AVX10.2
EVEX.512.F2.0F3A.W0 08 /r /ib VRNDSCALENEPBF16 zmm1{k1}{z}, zmm2/m512/m16bcst, imm8	A	V/V	AVX10.2

7.14.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	IMM8(r)	N/A

7.14.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

Round the bfloat16 floating-point values in the source operand by the rounding mode specified in the immediate operand and places the result in the destination operand.

The destination operand (the first operand) is a ZMM/YMM/XMM register conditionally updated according to the writemask. The source operand (the second operand) can be a ZMM/YMM/XMM register, a 512/256/128-bit memory location, or a 512/256/128-bit vector broadcasted from a 16-bit memory location.

The rounding process rounds the input to an integral value, plus number bits of fraction that are specified by imm8[7:4] (to be included in the result) and returns the result as a bfloat16 floating-point value. RNE rounding mode is used. the imm8[7:4] value).

If any source operand is an SNaN then it will be converted to a QNaN. Denormals will be converted to zero

before rounding. The sign of the result of this instruction is preserved, including the sign of zero. The formula of the operation on each data element for VRNDSCALENEPBF16 is

$$\text{ROUND}(x) = 2^{-M} \cdot \text{Round_to_INT}(x \cdot 2^M, \text{RNE}), M = \text{imm}[7:4];$$

The operation of $x \cdot 2^M$ is computed as if the exponent range is unlimited (i.e. no overflow ever occurs).

Note: EVEX.vvvv is reserved and must be 1111b, otherwise instructions will #UD.

7.14.3 OPERATION

```

1 def round_bf16_to_integer_ne(src,imm8):
2     IF (*src is nan*):
3         return QNAN(src)
4     m := imm8[7:4]
5     tmp := ROUND_TO_NEAREST_EVEN_INTEGER(2^m * src) //DAZ, SAE
6     tmp := 2^(-m) * tmp //FTZ, RNE, SAE
7     return tmp

```

```

1 VRNDSCALENEPBF16 destk1, src, imm8
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         IF SRC is memory and (EVEX.b == 1):
7             tsrc := SRC.bf16[0]
8         ELSE:
9             tsrc := SRC.bf16[j]
10            DEST.bf16[j] := round_bf16_to_integer_ne(tsrc, imm8)
11        ELSE IF *zeroing*:
12            DEST.bf16[j] := 0
13        // else dest.bf16[j] remains unchanged
14
15 DEST[MAX_VL-1:VL] := 0

```

7.14.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VRNDSCALENEPBF16 xmm1, xmm2/m128, imm8	E4	N/A	AVX10.2

Continued on next page...

Instruction	Exception Type	Arithmetic Flags	CPUID
VRNDSCALENEPBF16 ymm1, ymm2/m256, imm8	E4	N/A	AVX10.2
VRNDSCALENEPBF16 zmm1, zmm2/m512, imm8	E4	N/A	AVX10.2

7.15 VRSQRTPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.NP.MAP6.W0 4E /r VRSQRTPBF16 xmm1{k1}{z}, xmm2/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP6.W0 4E /r VRSQRTPBF16 ymm1{k1}{z}, ymm2/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.W0 4E /r VRSQRTPBF16 zmm1{k1}{z}, zmm2/m512/m16bcst	A	V/V	AVX10.2

7.15.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

7.15.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

This instruction performs a SIMD computation of the approximate reciprocals square-root of 8/16/32 packed BF16 floating-point values in the source operand (the second operand) and stores the packed BF16 floating-point results in the destination operand. The maximum relative error for this approximation is less than $2^{-8} + 2^{-14}$. For special cases, see Table 7.3. The destination elements are updated according to the writemask.

Input Value	Result Value	Comments
$0 \leq X < 2^{-126}$	+INF	DAZ
$-2^{-126} < X \leq 0$	-INF	DAZ
$X = 2^{-2n}$	2^n	
$X < 0$	QNaN In-definite	Including -INF

X = +INF	+0	
----------	----	--

Table 7.3: VRSQRTPBF16 Special Cases

7.15.3 OPERATION

```

1 VRSQRTPBF16 destk1, src
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         IF SRC is memory and (EVEX.b == 1):
7             tsrc := SRC.bf16[0]
8         ELSE:
9             tsrc := SRC.bf16[j]
10            DEST.bf16[j] := APPROXIMATE(1.0 / SQRT(tsrc)) //DAZ, FTZ, RNE, SAE
11        ELSE IF *zeroing*:
12            DEST.bf16[j] := 0
13        // else dest.bf16[j] remains unchanged
14
15 DEST[MAX_VL-1:VL] := 0

```

7.15.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VRSQRTPBF16 xmm1, xmm2/m128	E4	N/A	AVX10.2
VRSQRTPBF16 ymm1, ymm2/m256	E4	N/A	AVX10.2
VRSQRTPBF16 zmm1, zmm2/m512	E4	N/A	AVX10.2

7.16 VSCALEFPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.NP.MAP6.WO 2C /r VSCALEFPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP6.WO 2C /r VSCALEFPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP6.WO 2C /r VSCALEFPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2

7.16.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

7.16.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

Performs a floating-point scale of the packed bfloat16 floating-point values in the first source operand by multiplying it by 2 power of the bfloat16 values in second source operand.

The equation of this operation is given by:

$$zmm1 := zmm2 * 2^{\text{floor}(zmm3)}$$

Floor(zmm3) means maximum integer value \leq zmm3.

Denormal BF16 input operands are treated as zeros (DAZ) and denormal BF16 outputs are flushed to zero (FTZ).

7.16.3 OPERATION

```

1 def scale_bf16(src1,src2):
2     tmp1 := src1
3     tmp2 := src2
4     IF (src1 is denormal):
5         tmp1 := 0
6     IF (src2 is denormal):
7         tmp2 := 0
8     return tmp1 * POW(2, FLOOR(tmp2)) //FTZ, SAE

```

```

1 VSCALEFPBF16 destk1, src1, src2
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         IF SRC2 is memory and (EVEX.b == 1):
7             tsrc2 := SRC2.bf16[0]
8         ELSE:
9             tsrc2 := SRC2.bf16[j]
10            DEST.bf16[j] := scale_bf16(src1, tsrc2)
11    ELSE IF *zeroing*:
12        DEST.bf16[j] := 0
13    // else dest.bf16[j] remains unchanged
14
15 DEST[MAX_VL-1:VL] := 0

```

7.16.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VSCALEFPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VSCALEFPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VSCALEFPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2

7.17 VSQRTNEPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.66.MAP5.W0 51 /r VSQRTNEPBF16 xmm1{k1}{z}, xmm2/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.66.MAP5.W0 51 /r VSQRTNEPBF16 ymm1{k1}{z}, ymm2/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.66.MAP5.W0 51 /r VSQRTNEPBF16 zmm1{k1}{z}, zmm2/m512/m16bcst	A	V/V	AVX10.2

7.17.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

7.17.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

This instruction computes the SIMD square root of 8/16/32 packed BF16 floating-point values in the source operand (second operand), rounded to nearest. Outputs for special cases follow the IEEE specification for this operation.

This instruction does not generate floating point exceptions and does not consult or update MXCSR. Denormal bfloat16 input operands are treated as zeros (DAZ) and denormal bfloat16 outputs are flushed to zero (FTZ).

7.17.3 OPERATION

```

1 VSQRTNEPBF16 dest{k1}, src
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         IF SRC is memory and (EVEX.b == 1):
7             tsrc := SRC.bf16[0]
8         ELSE:
9             tsrc := SRC.bf16[j]
10        DEST.bf16[j] := SQRT(tsrc) //DAZ, FTZ, RNE, SAE
11    ELSE IF *zeroing*:
12        DEST.bf16[j] := 0
13    // else dest.bf16[j] remains unchanged
14
15 DEST[MAX_VL-1:VL] := 0

```

7.17.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VSQRTNEPBF16 xmm1, xmm2/m128	E4	N/A	AVX10.2
VSQRTNEPBF16 ymm1, ymm2/m256	E4	N/A	AVX10.2
VSQRTNEPBF16 zmm1, zmm2/m512	E4	N/A	AVX10.2

7.18 VSUBNEPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.66.MAP5.W0 5C /r VSUBNEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.66.MAP5.W0 5C /r VSUBNEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.66.MAP5.W0 5C /r VSUBNEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2

7.18.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

7.18.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

This instruction subtracts packed BF16 values from second source operand from the corresponding elements in the first source operand, storing the packed BF16 result in the destination operand. The destination elements are updated according to the writemask This instruction does not generate floating point exceptions and does not consult or update MXCSR. Denormal BF16 input operands are treated as zeros (DAZ) and denormal BF16 outputs are flushed to zero (FTZ). Rounding Mode is always RNE.

7.18.3 OPERATION

```

1 VSUBNEPBF16 (EVEX encoded versions) when src2 operand is a register
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         DEST.bf16[j] := SRC1.bf16[j] - SRC2.bf16[j] //DAZ, FTZ, RNE, SAE
7     ELSE IF *zeroing*:
8         DEST.bf16[j] := 0
9         // else dest.bf16[j] remains unchanged
10
11 DEST[MAX_VL-1:VL] := 0

```

```

1 VSUBNEPBF16 (EVEX encoded versions) when src2 operand is a memory source
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         IF EVEX.b == 1:
7             DEST.bf16[j] := SRC1.bf16[j] - SRC2.bf16[0] //DAZ, FTZ, RNE, SAE
8         ELSE:
9             DEST.bf16[j] := SRC1.bf16[j] - SRC2.bf16[j] //DAZ, FTZ, RNE, SAE
10
11     ELSE IF *zeroing*:
12         DEST.bf16[j] := 0
13         // else dest.bf16[j] remains unchanged
14
15 DEST[MAX_VL-1:VL] := 0

```

7.18.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VSUBNEPBF16 xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VSUBNEPBF16 ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VSUBNEPBF16 zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2

Chapter 8

INTEL® AVX10 COMPARE SCALAR FP WITH ENHANCED EFLAGS INSTRUCTIONS

8.1. VCOMXSD

8.1 VCOMXSD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.LLIG.F2.OF.W1 2F /r VCOMXSD xmm1, xmm2/m64 {sae}	A	V/V	AVX10.2

8.1.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	SCALAR	MODRM.REG(r)	MODRM.R/M(r)	N/A	N/A

8.1.2 DESCRIPTION

Note:

For SCALAR instructions with a CPUID feature flag specifying Intel® AVX10, the programmer can skip the check for the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and usually determines the set of instructions available to the programmer listed in the above opcode table. However, scalar instructions are architected without the notion of vector length, so this step does not strictly apply.

VCOMXSD: Compares the double-precision floating-point values in the low quadwords of operand 1 (first operand) and operand 2 (second operand), and sets the OF, SF, ZF, PF, and CF flags in the EFLAGS register according to the result (unordered, greater than, less than, or equal). The AF flag in the EFLAGS register is set to 0. The unordered result is returned if either source operand is a NaN (QNaN or SNaN). Operand 1 is an XMM register; operand 2 can be an XMM register or a 64-bit memory location. The VCOMXSD instruction differs from the VCOMISD instruction in that the EFLAGS register is set differently to allow more relationships to be tested directly, with unsigned integer tests used for testing greater than conditions and signed integer tests used for testing less than. The VCOMXSD instruction differs from the VUCOMXSD instruction in that it signals a SIMD floating-point invalid operation exception (#XM, specifically #I) when a source operand is either a QNaN or SNaN. The VUCOMXSD instruction signals an invalid operation exception only if a source operand is an SNaN. The EFLAGS register is not updated if an unmasked SIMD floating-point exception is generated. EVEX.vvvv are reserved and must be 1111b, otherwise instructions will #UD. Software should ensure VCOMXSD is encoded with EVEX.LL=00b, otherwise instructions will #UD.

VCOMXSD allows any combination of unordered, greater than, less than, and equal to be tested for using a single instruction.

Additionally: #UD if EVEX.vvvv ≠ 1111B, EVEX.vvvv ≠ 1111B, EVEX.L ≠ 0, or EVEX.LL ≠ 00B

8.1. VCOMXSD

Meaning	True	Inverse
Greater than	JA, CMOVA, SETA	JNA, CMOVNA, SETNA
Greater than or equal	JAE, CMOVAE, SETAE	JNAE, CMOVNAE, SETNAE
Equal	JE, CMOVE, SETE	JNE, CMOVNE, SETNE
Less than or equal	JLE, CMOVLE, SETLE	JNLE, CMOVNLE, SETNLE
Less than	JL, CMOVL, SETL	JNL, CMOVNL, SETNL
Ordered	JNP, CMOVNP, SETNP	JP, CMOVNP, SETP

Table 8.1: Valid Comparison Tests

8.1.3 OPERATION

```

1
2 VCOMXSD (all versions)
3 RESULT := OrderedCompare(DEST[63:0] != SRC[63:0]) {
4 (* Set EFLAGS *) CASE (RESULT) OF
5     UNORDERED: OF,SF,ZF,PF,CF := 11011;
6     GREATER_THAN: OF,SF,ZF,PF,CF := 00000;
7     LESS_THAN: OF,SF,ZF,PF,CF := 10001;
8     EQUAL: OF,SF,ZF,PF,CF := 11100;
9 ESAC;
10 AF := 0; }
11
12

```

8.1.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

8.1.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCOMXSD xmm1, xmm2/m64	E3NF	ID	AVX10.2

8.2. VCOMXSH

8.2 VCOMXSH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.LLIG.F3.MAP5.W0 2F /r VCOMXSH xmm1, xmm2/m16 {sae}	A	V/V	AVX10.2

8.2.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	SCALAR	MODRM.REG(r)	MODRM.R/M(r)	N/A	N/A

8.2.2 DESCRIPTION

Note:

For SCALAR instructions with a CPUID feature flag specifying Intel® AVX10, the programmer can skip the check for the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and usually determines the set of instructions available to the programmer listed in the above opcode table. However, scalar instructions are architected without the notion of vector length, so this step does not strictly apply.

VCOMXSH: Compares the 16 bit floating point values in the low words of operand 1 (first operand) and operand 2 (second operand), and sets the OF, SF, ZF, PF, and CF flags in the EFLAGS register according to the result (unordered, greater than, less than, or equal). The AF flag in the EFLAGS register is set to 0. The unordered result is returned if either source operand is a NaN (QNaN or SNaN). Operand 1 is an XMM register; operand 2 can be an XMM register or a 16-bit memory location. The VCOMXSH instruction differs from the VUCOMISH instruction in that the EFLAGS register is set differently to allow more relationships to be tested directly, with unsigned integer tests used for testing greater than conditions and signed integer tests used for testing less than. The EFLAGS register is not updated if an unmasked SIMD floating-point exception is generated. Note: EVEX.vvvv are reserved and must be 1111b, otherwise instructions will #UD. Software should ensure VCOMXSH is encoded with EVEX.LL=00b, otherwise instructions will #UD.

Additionally: #UD If EVEX.vvvv ≠ 1111B, EVEX.vvvv ≠ 1111B, EVEX.L ≠ 0, or EVEX.LL ≠ 00B

8.2. VCOMXSH

Meaning	True	Inverse
Greater than	JA, CMOVA, SETA	JNA, CMOVNA, SETNA
Greater than or equal	JAE, CMOVAE, SETAE	JNAE, CMOVNAE, SETNAE
Equal	JE, CMOVE, SETE	JNE, CMOVNE, SETNE
Less than or equal	JLE, CMOVLE, SETLE	JNLE, CMOVNLE, SETNLE
Less than	JL, CMOVL, SETL	JNL, CMOVNL, SETNL
Ordered	JNP, CMOVNP, SETNP	JP, CMOVNP, SETP

Table 8.2: Valid Comparison Tests

8.2.3 OPERATION

```

1
2 VCOMXSH (all versions)
3 RESULT := OrderedCompare(DEST[16:0] != SRC[16:0]) {
4 (* Set EFLAGS *) CASE (RESULT) OF
5     UNORDERED: OF,SF,ZF,PF,CF := 11011;
6     GREATER_THAN: OF,SF,ZF,PF,CF := 00000;
7     LESS_THAN: OF,SF,ZF,PF,CF := 10001;
8     EQUAL: OF,SF,ZF,PF,CF := 11100;
9 ESAC;
10 AF := 0; }
11

```

8.2.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

8.2.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCOMXSH xmm1, xmm2/m16	E3NF	ID	AVX10.2

8.3. VCOMXSS

8.3 VCOMXSS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.LLIG.F3.OF.W0 2F /r VCOMXSS xmm1, xmm2/m32 {sae}	A	V/V	AVX10.2

8.3.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	SCALAR	MODRM.REG(r)	MODRM.R/M(r)	N/A	N/A

8.3.2 DESCRIPTION

Note:

For SCALAR instructions with a CPUID feature flag specifying Intel® AVX10, the programmer can skip the check for the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and usually determines the set of instructions available to the programmer listed in the above opcode table. However, scalar instructions are architected without the notion of vector length, so this step does not strictly apply.

VCOMXSS: Compares the single-precision floating-point values in the low doublewords of operand 1 (first operand) and operand 2 (second operand), and sets the OF, SF, ZF, PF, and CF flags in the EFLAGS register according to the result (unordered, greater than, less than, or equal). The AF flag in the EFLAGS register is set to 0. The unordered result is returned if either source operand is a NaN (QNaN or SNaN). Operand 1 is an XMM register; operand 2 can be an XMM register or a 32-bit memory location. The VCOMXSS instruction differs from the VCOMISS instruction in that the EFLAGS register is set differently to allow more relationships to be tested directly, with unsigned integer tests used for testing greater than conditions and signed integer tests used for testing less than. The VCOMXSS instruction differs from the VUCOMXSS instruction in that it signals a SIMD floating-point invalid operation exception (#XM, specifically #I) when a source operand is either a QNaN or SNaN. The VUCOMXSS instruction signals an invalid operation exception only if a source operand is an SNaN. The EFLAGS register is not updated if an unmasked SIMD floating-point exception is generated. EVEX.vvvv are reserved and must be 1111b, otherwise instructions will #UD. Software should ensure VCOMXSS is encoded with EVEX.LL=00b, otherwise instructions will #UD.

Additionally: #UD If EVEX.vvvv ≠ 1111B, EVEX.vvvv ≠ 1111B, EVEX.L ≠ 0, or EVEX.LL ≠ 00B

8.3. VCOMXSS

Meaning	True	Inverse
Greater than	JA, CMOVA, SETA	JNA, CMOVNA, SETNA
Greater than or equal	JAE, CMOVAE, SETAE	JNAE, CMOVNAE, SETNAE
Equal	JE, CMOVE, SETE	JNE, CMOVNE, SETNE
Less than or equal	JLE, CMOVLE, SETLE	JNLE, CMOVNLE, SETNLE
Less than	JL, CMOVL, SETL	JNL, CMOVNL, SETNL
Ordered	JNP, CMOVNP, SETNP	JP, CMOVNP, SETP

Table 8.3: Valid Comparison Tests

8.3.3 OPERATION

```

1
2 VCOMXSS (all versions)
3 RESULT := OrderedCompare(DEST[31:0] != SRC[31:0]) {
4 (* Set EFLAGS *) CASE (RESULT) OF
5     UNORDERED: OF,SF,ZF,PF,CF := 11011;
6     GREATER_THAN: OF,SF,ZF,PF,CF := 00000;
7     LESS_THAN: OF,SF,ZF,PF,CF := 10001;
8     EQUAL: OF,SF,ZF,PF,CF := 11100;
9 ESAC;
10 AF := 0; }
11

```

8.3.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

8.3.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCOMXSS xmm1, xmm2/m32	E3NF	ID	AVX10.2

8.4. VUCOMXSD

8.4 VUCOMXSD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.LLIG.F2.0F.W1 2E /r VUCOMXSD xmm1, xmm2/m64 {sae}	A	V/V	AVX10.2

8.4.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	SCALAR	MODRM.REG(r)	MODRM.R/M(r)	N/A	N/A

8.4.2 DESCRIPTION

Note:

For SCALAR instructions with a CPUID feature flag specifying Intel® AVX10, the programmer can skip the check for the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and usually determines the set of instructions available to the programmer listed in the above opcode table. However, scalar instructions are architected without the notion of vector length, so this step does not strictly apply.

VUCOMXSD: Performs an unordered compare of the double-precision floating-point values in the low quadwords of operand 1 (first operand) and operand 2 (second operand), and sets the OF, SF, ZF, PF, and CF flags in the EFLAGS register according to the result (unordered, greater than, less than, or equal). The AF flag in the EFLAGS register is set to 0. The unordered result is returned if either source operand is a NaN (QNaN or SNaN). Operand 1 is an XMM register; operand 2 can be an XMM register or a 64-bit memory location. The VUCOMXSD instruction differs from the VUCOMISD instruction in that the EFLAGS register is set differently to allow more relationships to be tested directly, with unsigned integer tests used for testing greater than conditions and signed integer tests used for testing less than. The VUCOMXSD instruction differs from the VCOMXSD instruction in that it signals a SIMD floating-point invalid operation exception (#XM, specifically #I) only when a source operand is an SNaN. The EFLAGS register is not updated if an unmasked SIMD floating-point exception is generated. Note: EVEX.vvvv are reserved and must be 1111b, otherwise instructions will #UD. Software should ensure VUCOMXSD is encoded with EVEX.LL=00b, otherwise instructions will #UD.

Additionally: #UD If EVEX.vvvv ≠ 1111B, EVEX.vvvv ≠ 1111B, EVEX.L ≠ 0, or EVEX.LL ≠ 00B

8.4. VUCOMXSD

Meaning	True	Inverse
Greater than	JA, CMOVA, SETA	JNA, CMOVNA, SETNA
Greater than or equal	JAE, CMOVAE, SETAE	JNAE, CMOVNAE, SETNAE
Equal	JE, CMOVE, SETE	JNE, CMOVNE, SETNE
Less than or equal	JLE, CMOVLE, SETLE	JNLE, CMOVNLE, SETNLE
Less than	JL, CMOVL, SETL	JNL, CMOVNL, SETNL
Ordered	JNP, CMOVNP, SETNP	JP, CMOVNP, SETP

Table 8.4: Valid Comparison Tests

8.4.3 OPERATION

```

1
2 VUCOMXSD (all versions)
3 RESULT := OrderedCompare(DEST[63:0] != SRC[63:0]) {
4 (* Set EFLAGS *) CASE (RESULT) OF
5     UNORDERED: OF,SF,ZF,PF,CF := 11011;
6     GREATER_THAN: OF,SF,ZF,PF,CF := 00000;
7     LESS_THAN: OF,SF,ZF,PF,CF := 10001;
8     EQUAL: OF,SF,ZF,PF,CF := 11100;
9 ESAC;
10 AF := 0; }
11

```

8.4.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

8.4.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VUCOMXSD xmm1, xmm2/m64	E3NF	ID	AVX10.2

8.5. VUCOMXSH

8.5 VUCOMXSH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.LLIG.F3.MAP5.W0 2E /r VUCOMXSH xmm1, xmm2/m16 {sae}	A	V/V	AVX10.2

8.5.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	SCALAR	MODRM.REG(r)	MODRM.R/M(r)	N/A	N/A

8.5.2 DESCRIPTION

Note:

For SCALAR instructions with a CPUID feature flag specifying Intel® AVX10, the programmer can skip the check for the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and usually determines the set of instructions available to the programmer listed in the above opcode table. However, scalar instructions are architected without the notion of vector length, so this step does not strictly apply.

VUCOMXSH: Performs an unordered compare of the 16 bit floating point values in the low words of operand 1 (first operand) and operand 2 (second operand), and sets the OF, SF, ZF, PF, and CF flags in the EFLAGS register according to the result (unordered, greater than, less than, or equal). The AF flag in the EFLAGS register is set to 0. The unordered result is returned if either source operand is a NaN (QNaN or SNaN). Operand 1 is an XMM register; operand 2 can be an XMM register or a 16-bit memory location. The VUCOMXSH instruction differs from the VCOMISH instruction in that the EFLAGS register is set differently to allow more relationships to be tested directly, with unsigned integer tests used for testing greater than conditions and signed integer tests used for testing less than. The VUCOMXSH instruction differs from the VCOMXSH instruction in that it signals a SIMD floating-point invalid operation exception (#XM, specifically #I) only if a source operand is an SNaN. The VCOMXSH instruction signals an invalid operation exception when a source operand is either a QNaN or SNaN. The EFLAGS register is not updated if an unmasked SIMD floating-point exception is generated. Note: EVEX.vvvv are reserved and must be 1111b, otherwise instructions will #UD. Software should ensure VUCOMXSH is encoded with EVEX.LL=00b, otherwise instructions will #UD.

Additionally: #UD If EVEX.vvvv ≠ 1111B, EVEX.vvvv ≠ 1111B, EVEX.L ≠ 0, or EVEX.LL ≠ 00B

8.5. VUCOMXSH

Meaning	True	Inverse
Greater than	JA, CMOVA, SETA	JNA, CMOVNA, SETNA
Greater than or equal	JAE, CMOVAE, SETAE	JNAE, CMOVNAE, SETNAE
Equal	JE, CMOVE, SETE	JNE, CMOVNE, SETNE
Less than or equal	JLE, CMOVLE, SETLE	JNLE, CMOVNLE, SETNLE
Less than	JL, CMOVL, SETL	JNL, CMOVNL, SETNL
Ordered	JNP, CMOVNP, SETNP	JP, CMOVNP, SETP

Table 8.5: Valid Comparison Tests

8.5.3 OPERATION

```

1
2 VUCOMXSH (all versions)
3 RESULT := OrderedCompare(DEST[16:0] != SRC[16:0]) {
4 (* Set EFLAGS *) CASE (RESULT) OF
5     UNORDERED: OF,SF,ZF,PF,CF := 11011;
6     GREATER_THAN: OF,SF,ZF,PF,CF := 00000;
7     LESS_THAN: OF,SF,ZF,PF,CF := 10001;
8     EQUAL: OF,SF,ZF,PF,CF := 11100;
9 ESAC;
10 AF := 0; }
11

```

8.5.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

8.5.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VUCOMXSH xmm1, xmm2/m16	E3NF	ID	AVX10.2

8.6 VUCOMXSS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.LLIG.F3.0F.W0 2E /r VUCOMXSS xmm1, xmm2/m32 {sae}	A	V/V	AVX10.2

8.6.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	SCALAR	MODRM.REG(r)	MODRM.R/M(r)	N/A	N/A

8.6.2 DESCRIPTION

Note:

For SCALAR instructions with a CPUID feature flag specifying Intel® AVX10, the programmer can skip the check for the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and usually determines the set of instructions available to the programmer listed in the above opcode table. However, scalar instructions are architected without the notion of vector length, so this step does not strictly apply.

VUCOMXSS: Compares the single-precision floating-point values in the low doublewords of operand 1 (first operand) and operand 2 (second operand), and sets the OF, SF, ZF, PF, and CF flags in the EFLAGS register according to the result (unordered, greater than, less than, or equal). The AF flag in the EFLAGS register is set to 0. The unordered result is returned if either source operand is a NaN (QNaN or SNaN). Operand 1 is an XMM register; operand 2 can be an XMM register or a 32-bit memory location. The VUCOMXSS instruction differs from the VUCOMISS instruction in that the EFLAGS register is set differently to allow more relationships to be tested directly, with unsigned integer tests used for testing greater than conditions and signed integer tests used for testing less than. The VUCOMXSS instruction differs from the VCOMXSS instruction in that it signals a SIMD floating-point invalid operation exception (#XM, specifically #I) only if a source operand is an SNaN. The VCOMXSS instruction signals an invalid operation exception when a source operand is either a QNaN or SNaN. The EFLAGS register is not updated if an unmasked SIMD floating-point exception is generated. Note: EVEX.vvvv are reserved and must be 1111b, otherwise instructions will #UD. Software should ensure VUCOMXSS is encoded with EVEX.LL=00b, otherwise instructions will #UD.

Additionally: #UD If EVEX.vvvv ≠ 1111B, EVEX.vvvv ≠ 1111B, EVEX.L ≠ 0, or EVEX.LL ≠ 00B

8.6. VUCOMXSS

Meaning	True	Inverse
Greater than	JA, CMOVA, SETA	JNA, CMOVNA, SETNA
Greater than or equal	JAE, CMOVAE, SETAE	JNAE, CMOVNAE, SETNAE
Equal	JE, CMOVE, SETE	JNE, CMOVNE, SETNE
Less than or equal	JLE, CMOVLE, SETLE	JNLE, CMOVNLE, SETNLE
Less than	JL, CMOVL, SETL	JNL, CMOVNL, SETNL
Ordered	JNP, CMOVNP, SETNP	JP, CMOVNP, SETP

Table 8.6: Valid Comparison Tests

8.6.3 OPERATION

```

1
2 VUCOMXSS (all versions)
3 RESULT := OrderedCompare(DEST[31:0] != SRC[31:0]) {
4 (* Set EFLAGS *) CASE (RESULT) OF
5     UNORDERED: OF,SF,ZF,PF,CF := 11011;
6     GREATER_THAN: OF,SF,ZF,PF,CF := 00000;
7     LESS_THAN: OF,SF,ZF,PF,CF := 10001;
8     EQUAL: OF,SF,ZF,PF,CF := 11100;
9 ESAC;
10 AF := 0; }
11

```

8.6.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

8.6.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VUCOMXSS xmm1, xmm2/m32	E3NF	ID	AVX10.2

Chapter 9

INTEL® AVX10 CONVERT INSTRUCTIONS

9.1 VCVT2PS2PHX

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.66.0F38.W0 67 /r VCVT2PS2PHX xmm1{k1}{z}, xmm2, xmm3/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.66.0F38.W0 67 /r VCVT2PS2PHX ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2
EVEX.512.66.0F38.W0 67 /r VCVT2PS2PHX zmm1{k1}{z}, zmm2, zmm3/m512/m32bcst {er}	A	V/V	AVX10.2

9.1.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

9.1.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

This instruction converts two SIMD registers of Packed Single data into a single register of packed FP16 data. This instruction does not support memory fault suppression.

This instruction updates MXCSR as if all MXCSR numerical exceptions flags are masked and does not generate floating point exceptions. MXCSR (and EVEX embedded rounding) determine the rounding mode. Input FP32 denormals are affected by MXCSR.DAZ but output FP16 denormals are not affected by MXCSR.FTZ and not flushed to zero.

9.1.3 OPERATION

```

1 VCVT2PS2PH dest, src1, src2, k1    // EVEX encoded version
2 VL = (128,256,512)
3 KL := VL/16
4
5 for i := 0 to KL-1:
6     if (k1[i] or *no writemask*):
7         if i < KL/2:
8             if src2 is memory and evex.b == 1:
9                 t := src2.fp32[0]
10            else:
11                t := src2.fp32[i]
12            else:
13                t := src1.fp32[i-KL/2]
14            dest.word[i] := convert_fp32_to_fp16(t)
15
16        else:
17            if *merging-masking*:
18                dest.word[i] remains unchanged
19            else: // zero masking
20                dest.word[i] := 0
21
22 dest[MAX_VL-1:VL] := 0

```

9.1.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

9.1.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVT2PS2PHX xmm1, xmm2, xmm3/m128	E4NF	IOUPD	AVX10.2
VCVT2PS2PHX ymm1, ymm2, ymm3/m256	E4NF	IOUPD	AVX10.2
VCVT2PS2PHX zmm1, zmm2, zmm3/m512	E4NF	IOUPD	AVX10.2

9.2 VCVTBIASPH2[B,H]F8[S]

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.NP.OF38.W0 74 /r VCVTBIASPH2BF8 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.OF38.W0 74 /r VCVTBIASPH2BF8 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.OF38.W0 74 /r VCVTBIASPH2BF8 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.NP.MAP5.W0 74 /r VCVTBIASPH2BF8S xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP5.W0 74 /r VCVTBIASPH2BF8S ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP5.W0 74 /r VCVTBIASPH2BF8S zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.NP.MAP5.W0 18 /r VCVTBIASPH2HF8 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP5.W0 18 /r VCVTBIASPH2HF8 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP5.W0 18 /r VCVTBIASPH2HF8 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.NP.MAP5.W0 1B /r VCVTBIASPH2HF8S xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP5.W0 1B /r VCVTBIASPH2HF8S ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.NP.MAP5.W0 1B /r VCVTBIASPH2HF8S zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2

9.2.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

9.2.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

The instructions add a packed 8-bit unsigned INT number to the 8 bits below the LSB of the rounded number and then convert the result into a single register of E5M2 FP8 values or E4M3 FP8 values by normalizing and truncation to align with E5M2/E4M3 dynamic range. Infinity is preserved for the non-saturated versions. Nan at input is propagated as qnan. If the result is infinity or too big to be represented then for the saturated version, E5M2_MAX or E4M3_MAX is returned, for the non-saturated versions, infinity is returned (PH2BF8) or NaN (PH2HF8). In the 2 source versions the upper bits of the destination register beyond the downconverted E5M2/E4M3 elements are zeroed. DAZ is not obeyed and always assumed DAZ==0. FTZ is not obeyed and always assumed FTZ==0. Execution occurs as if all MXCSR exceptions are masked. All MXCSR mask bits DM, IM, OM, PM, UM are implicitly set for these instructions.

VCVTBIASPH2BF8 integer adds eight, sixteen or thirty-two packed INT8 in the lower half of the first source operand with eight, sixteen or thirty-two packed FP16 values in the second source operand. Addition is aligned to the FP16 LSB. The result is converted to eight, sixteen or thirty-two packed E5M2 FP8 values in the dest operand. The result is truncated. The UE (underflow) flag is set when the result is both denormal and inexact.

VCVTBIASPH2HF8 integer adds eight, sixteen or thirty-two packed INT8 in the lower half of the first source operand with eight, sixteen or thirty-two packed FP16 values in the second source operand. First the first source operand is shifted right by '1 and then the addition is aligned to the FP16 LSB. The result is converted to eight, sixteen or thirty-two packed E4M3 FP8 values in the dest operand. The result is truncated.

VCVTBIASPH2BF8S integer adds eight, sixteen or thirty-two packed INT8 in the lower half of the first source operand with eight, sixteen or thirty-two packed FP16 values in the second source operand. Addition is aligned to the FP16 LSB. The result is converted to eight, sixteen or thirty-two packed E5M2 FP8 values in the dest operand. The result is truncated. If the result is too big to be represented E5M2_MAX is returned. The UE (underflow) flag is set when the result is both denormal and inexact.

VCVTBIASPH2HF8S integer adds eight, sixteen or thirty-two packed INT8 in the lower half of the first source operand with eight, sixteen or thirty-two packed FP16 values in the second source operand. First the first source operand is shifted right by '1 and then the addition is aligned to the FP16 LSB. The result is converted to eight, sixteen or thirty-two packed E4M3 FP8 values in the dest operand. The result is truncated. If the result is too big to be represented E4M3_MAX is returned.

9.2.3 OPERATION

```

1  VCVTBIASPH2[B,H]F8[S] dest {k1}, src1, src2
2  VL = (128,256,512)
3  KL = VL/8
4  origdest := dest
5
6  if *dest is bfloat8*:
7     type := bf8
8  else:
9     type := hf8
10
11 if *saturation*:
12    s := 1
13 else:
14    s := 0
15
16 for i := 0 to KL/2-1:
17    if k1[i] or *no writemask*:
18       if src is memory and evex.b == 1:
19          t := src2.fp16[0]
20       else:
21          t := src2.fp16[i]
22       dest.fp8[i] := convert_fp16_to_fp8_bias(t, src1.byte[2i], type, s)
23    else if *zeroing*:
24       dest.fp8[i] := 0
25    else: // merge masking, dest element unchanged
26       dest.fp8[i] := origdest.fp8[i]
27 dest[max_vl-1:vl/2] := 0

```

9.2.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

9.2.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTBIASPH2BF8 xmm1, xmm2, xmm3/m128	E4NF	IOUPD	AVX10.2

Continued on next page...

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTBIASPH2BF8 xmm1, ymm2, ymm3/m256	E4NF	IOUPD	AVX10.2
VCVTBIASPH2BF8 ymm1, zmm2, zmm3/m512	E4NF	IOUPD	AVX10.2
VCVTBIASPH2BF8S xmm1, xmm2, xmm3/m128	E4NF	IOUPD	AVX10.2
VCVTBIASPH2BF8S xmm1, ymm2, ymm3/m256	E4NF	IOUPD	AVX10.2
VCVTBIASPH2BF8S ymm1, zmm2, zmm3/m512	E4NF	IOUPD	AVX10.2
VCVTBIASPH2HF8 xmm1, xmm2, xmm3/m128	E4NF	IOUPD	AVX10.2
VCVTBIASPH2HF8 xmm1, ymm2, ymm3/m256	E4NF	IOUPD	AVX10.2
VCVTBIASPH2HF8 ymm1, zmm2, zmm3/m512	E4NF	IOUPD	AVX10.2
VCVTBIASPH2HF8S xmm1, xmm2, xmm3/m128	E4NF	IOUPD	AVX10.2
VCVTBIASPH2HF8S xmm1, ymm2, ymm3/m256	E4NF	IOUPD	AVX10.2
VCVTBIASPH2HF8S ymm1, zmm2, zmm3/m512	E4NF	IOUPD	AVX10.2

9.3 VCVTHF82PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.F2.MAP5.W0 1E /r VCVTHF82PH xmm1{k1}{z}, xmm2/m64	A	V/V	AVX10.2
EVEX.256.F2.MAP5.W0 1E /r VCVTHF82PH ymm1{k1}{z}, xmm2/m128	A	V/V	AVX10.2
EVEX.512.F2.MAP5.W0 1E /r VCVTHF82PH zmm1{k1}{z}, ymm2/m256	A	V/V	AVX10.2

9.3.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	HALF	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

9.3.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

VCVTHF82PH converts E4M3 FP8 value datatype elements into FP16 elements. DAZ is not obeyed and always assumed DAZ==0. Since any E4M3 FP8 number can be represented in FP16, the conversion result is exact and no rounding is needed. This instruction only updates DE in MXCSR.

9.3.3 OPERATION

```

1 VCVTHF82PH dest k1, src
2 VL = (128,256,512)
3 KL := VL/16
4 origdest := dest
5 for i := 0 to kl-1:
6     if k1[i] OR *no writemask*:
7         tsrc := src.hf8[i]
8         dest.fp16[i] := convert_hf8_to_fp16(tsrc)
9     else if *zeroing*:
10        dest.fp16[i] := 0
11    else: // merge masking, dest element unchanged
12        dest.fp16[i] := origdest.fp16[i]
13 dest[max_vl-1:vl] := 0

```

9.3.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal In abbreviated form, this includes: DE

9.3.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTHF82PH xmm1, xmm2/m64	E2	D	AVX10.2
VCVTHF82PH ymm1, xmm2/m128	E2	D	AVX10.2
VCVTHF82PH zmm1, ymm2/m256	E2	D	AVX10.2

9.4 VCVTNE[2]PH2[B,H]F8[S]

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.F2.0F38.W0 74 /r VCVTNE2PH2BF8 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.F2.0F38.W0 74 /r VCVTNE2PH2BF8 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.F2.0F38.W0 74 /r VCVTNE2PH2BF8 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.F2.MAP5.W0 74 /r VCVTNE2PH2BF8S xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.F2.MAP5.W0 74 /r VCVTNE2PH2BF8S ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.F2.MAP5.W0 74 /r VCVTNE2PH2BF8S zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.F2.MAP5.W0 18 /r VCVTNE2PH2HF8 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.F2.MAP5.W0 18 /r VCVTNE2PH2HF8 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.F2.MAP5.W0 18 /r VCVTNE2PH2HF8 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.F2.MAP5.W0 1B /r VCVTNE2PH2HF8S xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.F2.MAP5.W0 1B /r VCVTNE2PH2HF8S ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.F2.MAP5.W0 1B /r VCVTNE2PH2HF8S zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.F3.0F38.W0 74 /r VCVTNEPH2BF8 xmm1{k1}{z}, xmm2/m128/m16bcst	B	V/V	AVX10.2
EVEX.256.F3.0F38.W0 74 /r VCVTNEPH2BF8 xmm1{k1}{z}, ymm2/m256/m16bcst	B	V/V	AVX10.2
EVEX.512.F3.0F38.W0 74 /r VCVTNEPH2BF8 ymm1{k1}{z}, zmm2/m512/m16bcst	B	V/V	AVX10.2
EVEX.128.F3.MAP5.W0 74 /r VCVTNEPH2BF8S xmm1{k1}{z}, xmm2/m128/m16bcst	B	V/V	AVX10.2

Continued on next page...

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F3.MAP5.W0 74 /r VCVTNEPH2BF8S xmm1{k1}{z}, ymm2/m256/m16bcst	B	V/V	AVX10.2
EVEX.512.F3.MAP5.W0 74 /r VCVTNEPH2BF8S ymm1{k1}{z}, zmm2/m512/m16bcst	B	V/V	AVX10.2
EVEX.128.F3.MAP5.W0 18 /r VCVTNEPH2HF8S xmm1{k1}{z}, xmm2/m128/m16bcst	B	V/V	AVX10.2
EVEX.256.F3.MAP5.W0 18 /r VCVTNEPH2HF8S xmm1{k1}{z}, ymm2/m256/m16bcst	B	V/V	AVX10.2
EVEX.512.F3.MAP5.W0 18 /r VCVTNEPH2HF8S ymm1{k1}{z}, zmm2/m512/m16bcst	B	V/V	AVX10.2
EVEX.128.F3.MAP5.W0 1B /r VCVTNEPH2HF8S xmm1{k1}{z}, xmm2/m128/m16bcst	B	V/V	AVX10.2
EVEX.256.F3.MAP5.W0 1B /r VCVTNEPH2HF8S xmm1{k1}{z}, ymm2/m256/m16bcst	B	V/V	AVX10.2
EVEX.512.F3.MAP5.W0 1B /r VCVTNEPH2HF8S ymm1{k1}{z}, zmm2/m512/m16bcst	B	V/V	AVX10.2

9.4.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A
B	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

9.4.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

These instructions convert one or two SIMD registers of packed FP16 data into a single register of packed E5M2 FP8 values or E4M3 FP8 values. The upper bits of the destination register beyond the

downconverted E5M2/E4M3 elements are zeroed. Rounding Mode is always RNE (Round To Nearest Ties to Even). DAZ is not obeyed and always assumed DAZ=0. FTZ is not obeyed and always assumed FTZ=0. These instructions do not raise exceptions but do set status. For saturated instructions, infinity at input is saturated to maximal normal values. For non-saturated versions, infinity at input is preserved (PH2BF8) or converted to NaN (PH2HF8). In case of overflow due to conversion or rounding, the Saturated instructions (with 'S' suffix) returns the corresponding E5M2_MAX or E4M3_MAX, the non-saturated versions return infinity (PH2BF8) or NaN (PH2HF8). All MXCSR mask bits DM, IM, OM, PM, UM are implicitly set for these instructions.

VCVTNEPH2BF8 converts eight, sixteen or thirty-two packed FP16 values in the source operand to eight, sixteen or thirty-two packed E5M2 FP8 values in the destination operand. The UE (underflow) flag is set when the result is both denormal and inexact.

VCVTNEPH2HF8 converts eight, sixteen or thirty-two packed FP16 values in the source operand to eight, sixteen or thirty-two packed E4M3 FP8 values in the destination operand.

VCVTNEPH2BF8S converts eight, sixteen or thirty-two packed FP16 values in the source operand to eight, sixteen or thirty-two packed E5M2 FP8 values in the destination operand. Returns Saturated values in case of overflow due to conversion or rounding. The UE (underflow) flag is set when the result is both denormal and inexact.

VCVTNEPH2HF8S converts eight, sixteen or thirty-two packed FP16 values in the source operand to eight, sixteen or thirty-two packed E4M3 FP8 values in the destination operand. Returns Saturated values in case of overflow due to conversion or rounding.

VCVTNE2PH2BF8 converts sixteen, thirty-two or sixty-four packed FP16 values in the 2 source operands to sixteen, thirty-two or sixty-four packed E5M2 FP8 values in the destination operand. The UE (underflow) flag is set when the result is both denormal and inexact.

VCVTNE2PH2HF8 converts sixteen, thirty-two or sixty-four packed FP16 values in the 2 source operands to sixteen, thirty-two or sixty-four packed E4M3 FP8 values in the destination operand.

VCVTNE2PH2BF8S converts sixteen, thirty-two or sixty-four packed FP16 values in the 2 source operands to sixteen, thirty-two or sixty-four packed E5M2 FP8 values in the destination operand. Returns Saturated values in case of overflow due to conversion or rounding. The UE (underflow) flag is set when the result is both denormal and inexact.

VCVTNE2PH2HF8S converts sixteen, thirty-two or sixty-four packed FP16 values in the 2 source operands to sixteen, thirty-two or sixty-four packed E4M3 FP8 values in the destination operand. Returns Saturated values in case of overflow due to conversion or rounding.

9.4.3 OPERATION

```
1 VCVTNEPH2[B,H]F8[S] dest {k1}, src
2 VL = (128,256,512)
3 KL = VL/8
4 origdest := dest
5 if *dest is bfloat8*:
6     y := bf8
7 else:
8     y := hf8
9 if *saturation*:
10    s := 1
11 else:
12    s := 0
13
14 for i := 0 to KL/2-1:
15     if k1[i] or *no writemask*:
16         if src is memory and evex.b == 1:
17             t := src.fp16[0]
18         else:
19             t := src.fp16[i]
20
21         dest.fp8[i] := convert_fp16_to_fp8(t, y, s)
22
23     else if *zeroing*:
24         dest.fp8[i] := 0
25     else: // merge masking, dest element unchanged
26         dest.fp8[i] := origdest.fp8[i]
27 dest[max_vl-1:vl/2] := 0
```

```

1 VCVTNE2PH2[B,H]F8[,S] dest, src1, src2, k1
2 VL = (128,256,512)
3 KL = VL/8
4 if *dest is bfloat8*:
5     y := bf8
6 else:
7     y := hf8
8 if *saturation*:
9     s := 1
10 else:
11     s := 0
12
13 for i := 0 to KL-1:
14     if (k1[i] or *no writemask*):
15         if i < KL/2:
16             if src2 is memory and evex.b == 1:
17                 t := src2.fp16[0]
18             else:
19                 t := src2.fp16[i]
20         else:
21             t := src1.fp16[i-KL/2]
22
23         dest.fp8[i] := convert_fp16_to_fp8(t, y, s)
24     else:
25         if *merging-masking*:
26             dest.fp8[i] remains unchanged
27         else:
28             dest.fp8[i] := 0
29
30 dest[max_vl-1:vl] := 0

```

9.4.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

9.4.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTNE2PH2BF8 xmm1, xmm2, xmm3/m128	E4NF	IOUPD	AVX10.2
VCVTNE2PH2BF8 ymm1, ymm2, ymm3/m256	E4NF	IOUPD	AVX10.2
VCVTNE2PH2BF8 zmm1, zmm2, zmm3/m512	E4NF	IOUPD	AVX10.2
VCVTNE2PH2BF8S xmm1, xmm2, xmm3/m128	E4NF	IOUPD	AVX10.2
VCVTNE2PH2BF8S ymm1, ymm2, ymm3/m256	E4NF	IOUPD	AVX10.2
VCVTNE2PH2BF8S zmm1, zmm2, zmm3/m512	E4NF	IOUPD	AVX10.2
VCVTNE2PH2HF8 xmm1, xmm2, xmm3/m128	E4NF	IOUPD	AVX10.2
VCVTNE2PH2HF8 ymm1, ymm2, ymm3/m256	E4NF	IOUPD	AVX10.2
VCVTNE2PH2HF8 zmm1, zmm2, zmm3/m512	E4NF	IOUPD	AVX10.2
VCVTNE2PH2HF8S xmm1, xmm2, xmm3/m128	E4NF	IOUPD	AVX10.2
VCVTNE2PH2HF8S ymm1, ymm2, ymm3/m256	E4NF	IOUPD	AVX10.2
VCVTNE2PH2HF8S zmm1, zmm2, zmm3/m512	E4NF	IOUPD	AVX10.2
VCVTNEPH2BF8 xmm1, xmm2/m128	E4NF	IOUPD	AVX10.2
VCVTNEPH2BF8 xmm1, ymm2/m256	E4NF	IOUPD	AVX10.2
VCVTNEPH2BF8 ymm1, zmm2/m512	E4NF	IOUPD	AVX10.2
VCVTNEPH2BF8S xmm1, xmm2/m128	E4NF	IOUPD	AVX10.2

Continued on next page...

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTNEPH2BF8S xmm1, ymm2/m256	E4NF	IOUPD	AVX10.2
VCVTNEPH2BF8S ymm1, zmm2/m512	E4NF	IOUPD	AVX10.2
VCVTNEPH2HF8 xmm1, xmm2/m128	E4NF	IOUPD	AVX10.2
VCVTNEPH2HF8 xmm1, ymm2/m256	E4NF	IOUPD	AVX10.2
VCVTNEPH2HF8 ymm1, zmm2/m512	E4NF	IOUPD	AVX10.2
VCVTNEPH2HF8S xmm1, xmm2/m128	E4NF	IOUPD	AVX10.2
VCVTNEPH2HF8S xmm1, ymm2/m256	E4NF	IOUPD	AVX10.2
VCVTNEPH2HF8S ymm1, zmm2/m512	E4NF	IOUPD	AVX10.2

Chapter 10

INTEL® AVX10 INTEGER AND FP16 VNNI, MEDIA NEW INSTRUCTIONS

10.1 VDPPHPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.NP.OF38.W0 52 /r VDPPHPS xmm1{k1}{z}, xmm2, xmm3/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.NP.OF38.W0 52 /r VDPPHPS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst	A	V/V	AVX10.2
EVEX.512.NP.OF38.W0 52 /r VDPPHPS zmm1{k1}{z}, zmm2, zmm3/m512/m32bcst	A	V/V	AVX10.2

10.1.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

10.1.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

This instruction performs a SIMD dot-product of two FP16 pairs and accumulates into a packed single precision register.

“Round to nearest even” rounding mode is used when doing each accumulation of the FMA. Output denormals are always flushed to zero. DAZ==0 is assumed for FP16 input. DAZ==1 is assumed for FP32 input. MXCSR is not consulted nor updated.

When more than one input is NaN, the NaN that is propagated to the result (after being quietized) is the first NaN that occurs in the expression Src2low*Src3low + Src2high*Src3high + SrcDest.

10.1. VDPPHPS

10.1.3 OPERATION

```

1  VDPPHPS srcdest, src1, src2 // EVEX ENCODED VERSION
2  VL = (128,256,512)
3
4  kl := vl/32
5  origdest := srcdest
6  for i := 0 to kl-1:
7      if kl[i] or *no writemask*:
8          if src2 is memory and evex.b == 1:
9              t := src2.dword[0]
10             else:
11                 t := src2.dword[i]
12
13             s1o = convert_fp16_to_fp32(src1.fp16[2*i+1])
14             s2o = convert_fp16_to_fp32(t.fp16[1])
15             s1e = convert_fp16_to_fp32(src1.fp16[2*i+0])
16             s2e = convert_fp16_to_fp32(t.fp16[0])
17
18             // MXCSR is neither consulted nor updated.
19             // Masked response for all floating point exceptional conditions.
20             // MXCSR.DAZ=0 is assumed for FP16 input.
21             // MXCSR.DAZ=1 is assumed for FP32 input.
22             // MXCSR.FTZ=1 is assumed for FP32 outputs.
23             // Uses RNE rounding.
24
25             srcdest.fp32[i] = fma32(srcdest.fp32[i], s1o, s2o, daz=1, ftz=1, sae=1, rc=RNE)
26             srcdest.fp32[i] = fma32(srcdest.fp32[i], s1e, s2e, daz=1, ftz=1, sae=1, rc=RNE)
27
28         else if *zeroing*:1
29             srcdest.dword[i] := 0
30
31         else: // merge masking, dest element unchanged
32             srcdest.dword[i] := origdest.dword[i]
33
34 srcdest[max_vl-1:vl] := 0

```

10.1.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VDPPHPS xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VDPPHPS ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2

Continued on next page...

10.1. VDPPHPS

Instruction	Exception Type	Arithmetic Flags	CPUID
VDPPHPS zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2

10.2 VMPSADBW

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.F3.0F3A.W0 42 /r /ib VMPSADBW xmm1{k1}{z}, xmm2, xmm3/m128, imm8	A	V/V	AVX10.2
EVEX.256.F3.0F3A.W0 42 /r /ib VMPSADBW ymm1{k1}{z}, ymm2, ymm3/m256, imm8	A	V/V	AVX10.2
EVEX.512.F3.0F3A.W0 42 /r /ib VMPSADBW zmm1{k1}{z}, zmm2, zmm3/m512, imm8	A	V/V	AVX10.2

10.2.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULLMEM	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	IMM8(r)

10.2.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

This is the EVEX encoded version of the VMPSADBW instruction. There are AVX and AVX2 versions of this instruction which were introduced on Sandybridge and Haswell respectively.

The EVEX VMPSADBW calculates packed word results of sum-absolute-difference (SAD) of unsigned bytes from two blocks of 32-bit dword elements, using two select fields in the immediate byte to select the offsets of the two blocks within the first source operand and the second operand. Packed SAD word results are calculated within each 128-bit lane. Each SAD word result is calculated between a stationary block_2 (whose offset within the second source operand is selected by a two bit select control, multiplied by 32 bits) and a sliding block_1 at consecutive byte-granular position within the first source operand. The offset of the first 32-bit block of block_1 is selectable using a one bit select control, multiplied by 32 bits. For the 512-bit version of this instruction the control bits for the lower two lanes are replicated to the upper two lanes.

10.2.3 OPERATION

```
1 def zero_to_max_vl(dst, vl):  
2     for i in range(vl, MAXVL):  
3         dst.bit[i] = 0
```

```

1 def emulate_vmepsadbw(inst, dst, msk, src1, src2, control):
2     blk2_pos = (control&3) * 4
3     blk1_pos = ((control>>2) & 1) * 4
4
5     # range(x,y) is x through y-1 inclusive
6     for i in range(0, 11):
7         blk1.byte[i] = src1.byte[i+blk1_pos]
8
9     for i in range(0, 4):
10        blk2.byte[i] = src2.byte[i+blk2_pos]
11
12    for i in range(0, 8):
13        for j in range(0, 4):
14            x = blk1.byte[j+i] - blk2.byte[j] # 16b signed arithmetic
15            tmp.word[j] = abs(x)
16
17        if inst.write_masking == 0 or msk[i]:
18            s = 0
19            for j in range(0, 4):
20                s += tmp.word[j]
21            dst.word[i] = s
22        elif inst.zeroing:
23            dst.word[i] = 0
24        #else dst.word[i] remains unchanged
25
26
27 def vmepsadbw_128(inst, dst, msk, src1, src2, imm8):
28     emulate_vmepsadbw(inst, dst.xmm[0], msk>>0, src1.xmm[0], src2.xmm[0], imm8)
29     zero_to_max_vl(dst, inst.VL)
30
31 def vmepsadbw_256(inst, dst, msk, src1, src2, imm8):
32     emulate_vmepsadbw(inst, dst.xmm[0], msk>>0, src1.xmm[0], src2.xmm[0], imm8)
33     emulate_vmepsadbw(inst, dst.xmm[1], msk>>8, src1.xmm[1], src2.xmm[1], imm8>>3)
34     zero_to_max_vl(dst, inst.VL)
35
36 def vmepsadbw_512(inst, dst, msk, src1, src2, imm8):
37     emulate_vmepsadbw(inst, dst.xmm[0], msk>>0, src1.xmm[0], src2.xmm[0], imm8)
38     emulate_vmepsadbw(inst, dst.xmm[1], msk>>8, src1.xmm[1], src2.xmm[1], imm8>>3)
39     emulate_vmepsadbw(inst, dst.xmm[2], msk>>16, src1.xmm[2], src2.xmm[2], imm8)
40     emulate_vmepsadbw(inst, dst.xmm[3], msk>>24, src1.xmm[3], src2.xmm[3], imm8>>3)
41     zero_to_max_vl(dst, inst.VL)

```

10.2.4 EXCEPTIONS

10.2. VMPSADBW

Instruction	Exception Type	Arithmetic Flags	CPUID
VMPSADBW xmm1, xmm2, xmm3/m128, imm8	E4NF	N/A	AVX10.2
VMPSADBW ymm1, ymm2, ymm3/m256, imm8	E4NF	N/A	AVX10.2
VMPSADBW zmm1, zmm2, zmm3/m512, imm8	E4NF	N/A	AVX10.2

10.3. VPDPB[SU,UU,SS]D[,S]

10.3 VPDPB[SU,UU,SS]D[,S]

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.F2.0F38.W0 50 /r VPDPBSSD xmm1{k1}{z}, xmm2, xmm3/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.F2.0F38.W0 50 /r VPDPBSSD ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst	A	V/V	AVX10.2
EVEX.512.F2.0F38.W0 50 /r VPDPBSSD zmm1{k1}{z}, zmm2, zmm3/m512/m32bcst	A	V/V	AVX10.2
EVEX.128.F2.0F38.W0 51 /r VPDPBSSDS xmm1{k1}{z}, xmm2, xmm3/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.F2.0F38.W0 51 /r VPDPBSSDS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst	A	V/V	AVX10.2
EVEX.512.F2.0F38.W0 51 /r VPDPBSSDS zmm1{k1}{z}, zmm2, zmm3/m512/m32bcst	A	V/V	AVX10.2
EVEX.128.F3.0F38.W0 50 /r VPDPBSUD xmm1{k1}{z}, xmm2, xmm3/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.F3.0F38.W0 50 /r VPDPBSUD ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst	A	V/V	AVX10.2
EVEX.512.F3.0F38.W0 50 /r VPDPBSUD zmm1{k1}{z}, zmm2, zmm3/m512/m32bcst	A	V/V	AVX10.2
EVEX.128.F3.0F38.W0 51 /r VPDPBSUDS xmm1{k1}{z}, xmm2, xmm3/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.F3.0F38.W0 51 /r VPDPBSUDS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst	A	V/V	AVX10.2
EVEX.512.F3.0F38.W0 51 /r VPDPBSUDS zmm1{k1}{z}, zmm2, zmm3/m512/m32bcst	A	V/V	AVX10.2
EVEX.128.NP.0F38.W0 50 /r VPDPBUUD xmm1{k1}{z}, xmm2, xmm3/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.NP.0F38.W0 50 /r VPDPBUUD ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst	A	V/V	AVX10.2
EVEX.512.NP.0F38.W0 50 /r VPDPBUUD zmm1{k1}{z}, zmm2, zmm3/m512/m32bcst	A	V/V	AVX10.2
EVEX.128.NP.0F38.W0 51 /r VPDPBUUDS xmm1{k1}{z}, xmm2, xmm3/m128/m32bcst	A	V/V	AVX10.2

Continued on next page...

10.3. VPDPB[SU,UU,SS]D[,S]

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.0F38.W0 51 /r VPDPBUUDS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst	A	V/V	AVX10.2
EVEX.512.NP.0F38.W0 51 /r VPDPBUUDS zmm1{k1}{z}, zmm2, zmm3/m512/m32bcst	A	V/V	AVX10.2

10.3.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

10.3.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

Multiplies the individual bytes of the first source operand by the corresponding bytes of the second source operand, producing intermediate word results. The word results are then summed and accumulated in the destination dword element size operand.

For unsigned saturation, when an individual result value is beyond the range of an unsigned doubleword (that is, greater than FFFF_FFFFH), the saturated unsigned doubleword integer value of FFFF_FFFFH stored in the doubleword destination.

For signed saturation, when an individual result is beyond the range of a signed doubleword integer (that is, greater than 7FFF_FFFFH or less than 8000_0000H), the saturated value of 7FFF_FFFFH or 8000_0000H, respectively, is written to the destination operand.

The EVEX encoded form of this instruction supports memory fault suppression.

10.3. VPDPB[SU,UU,SS]D[S]

10.3.3 OPERATION

```

1  VPDPB[UU,SS,SU]D[S] dest, src1, src2    // EVEX encoded version
2  VL = (128,256,512)
3
4  KL := VL/32
5  ORIGDEST := DEST
6  FOR i := 0 TO KL-1:
7      IF k1[i] or *no writemask*:
8          // Elements of SRC1 are zero-extended to 16b and
9          // byte elements of SRC2 are sign extended to 16b before multiplication.
10         if SRC2 is memory and EVEX.b == 1:
11             t := SRC2.dword[0]
12         ELSE:
13             t := SRC2.dword[i]
14
15         if *src1 is signed*:
16             src1extend := SIGN_EXTEND    // SU, SS
17         else:
18             src1extend := ZERO_EXTEND    // UU
19         if *src2 is signed*:
20             src2extend := SIGN_EXTEND    // SS
21         else:
22             src2extend := ZERO_EXTEND    // UU, SU
23
24         p1word := src1extend(SRC1.byte[4*i+0]) * src2extend(t.byte[0])
25         p2word := src1extend(SRC1.byte[4*i+1]) * src2extend(t.byte[1])
26         p3word := src1extend(SRC1.byte[4*i+2]) * src2extend(t.byte[2])
27         p4word := src1extend(SRC1.byte[4*i+3]) * src2extend(t.byte[3])
28
29         IF *saturating*:
30             IF *UU instruction version*:
31                 DEST.dword[i] := UNSIGNED_DWORD_SATURATE(ORIGDEST.dword[i] +
32                     p1word + p2word + p3word + p4word)
33             ELSE:
34                 DEST.dword[i] := SIGNED_DWORD_SATURATE(ORIGDEST.dword[i] +
35                     p1word + p2word + p3word + p4word)
36         ELSE:
37             DEST.dword[i] := ORIGDEST.dword[i] + p1word + p2word + p3word + p4word
38
39         ELSE IF *zeroing*:
40             DEST.dword[i] := 0
41         ELSE: // Merge masking, dest element unchanged
42             DEST.dword[i] := ORIGDEST.dword[i]
43     DEST[MAX_VL-1:VL] := 0

```

10.3. VPDPB[SU,UU,SS]D[,S]

```

1  VPDPBUU,SS,SUD,S dest, src1, src2    // VEX encoded version
2  VL = (128,256)
3
4  KL := VL/32
5  ORIGDEST := DEST
6  FOR i := 0 TO KL-1:
7    // Elements of SRC1 are zero-extended to 16b and
8    // byte elements of SRC2 are sign extended to 16b before multiplication.
9
10   if *src1 is signed*:
11     src1extend := SIGN_EXTEND    // SU, SS
12   else:
13     src1extend := ZERO_EXTEND    // UU
14   if *src2 is signed*:
15     src2extend := SIGN_EXTEND    // SS
16   else:
17     src2extend := ZERO_EXTEND    // UU, SU
18
19   p1word := src1extend(SRC1.byte[4*i+0]) * src2extend(SRC2.dword[4*i+0])
20   p2word := src1extend(SRC1.byte[4*i+1]) * src2extend(SRC2.dword[4*i+1])
21   p3word := src1extend(SRC1.byte[4*i+2]) * src2extend(SRC2.dword[4*i+2])
22   p4word := src1extend(SRC1.byte[4*i+3]) * src2extend(SRC2.dword[4*i+3])
23
24   IF *saturating*:
25     IF *UU instruction version*:
26       DEST.dword[i] := UNSIGNED_DWORD_SATURATE(ORIGDEST.dword[i] +
27         p1word + p2word + p3word + p4word)
28     ELSE:
29       DEST.dword[i] := SIGNED_DWORD_SATURATE(ORIGDEST.dword[i] +
30         p1word + p2word + p3word + p4word)
31   ELSE:
32     DEST.dword[i] := ORIGDEST.dword[i] + p1word + p2word + p3word + p4word
33
34  DEST[MAX_VL-1:VL] := 0

```

10.3.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VPDPBSSD xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VPDPBSSD ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2

Continued on next page...

10.3. VPDPB[SU,UU,SS]D[,S]

Instruction	Exception Type	Arithmetic Flags	CPUID
VPDPBSSD zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VPDPBSSDS xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VPDPBSSDS ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VPDPBSSDS zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VPDPBSUD xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VPDPBSUD ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VPDPBSUD zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VPDPBSUDS xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VPDPBSUDS ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VPDPBSUDS zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VPDPBUUD xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VPDPBUUD ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VPDPBUUD zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VPDPBUUDS xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VPDPBUUDS ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VPDPBUUDS zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2

10.4. VPDPW[SU,US,UU]D[,S]

10.4 VPDPW[SU,US,UU]D[,S]

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.F3.0F38.W0 D2 /r VPDPWSUD xmm1{k1}{z}, xmm2, xmm3/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.F3.0F38.W0 D2 /r VPDPWSUD ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst	A	V/V	AVX10.2
EVEX.512.F3.0F38.W0 D2 /r VPDPWSUD zmm1{k1}{z}, zmm2, zmm3/m512/m32bcst	A	V/V	AVX10.2
EVEX.128.F3.0F38.W0 D3 /r VPDPWSUDS xmm1{k1}{z}, xmm2, xmm3/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.F3.0F38.W0 D3 /r VPDPWSUDS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst	A	V/V	AVX10.2
EVEX.512.F3.0F38.W0 D3 /r VPDPWSUDS zmm1{k1}{z}, zmm2, zmm3/m512/m32bcst	A	V/V	AVX10.2
EVEX.128.66.0F38.W0 D2 /r VPDPWUSD xmm1{k1}{z}, xmm2, xmm3/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.66.0F38.W0 D2 /r VPDPWUSD ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst	A	V/V	AVX10.2
EVEX.512.66.0F38.W0 D2 /r VPDPWUSD zmm1{k1}{z}, zmm2, zmm3/m512/m32bcst	A	V/V	AVX10.2
EVEX.128.66.0F38.W0 D3 /r VPDPWUSDs xmm1{k1}{z}, xmm2, xmm3/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.66.0F38.W0 D3 /r VPDPWUSDs ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst	A	V/V	AVX10.2
EVEX.512.66.0F38.W0 D3 /r VPDPWUSDs zmm1{k1}{z}, zmm2, zmm3/m512/m32bcst	A	V/V	AVX10.2
EVEX.128.NP.0F38.W0 D2 /r VPDPWUUD xmm1{k1}{z}, xmm2, xmm3/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.NP.0F38.W0 D2 /r VPDPWUUD ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst	A	V/V	AVX10.2
EVEX.512.NP.0F38.W0 D2 /r VPDPWUUD zmm1{k1}{z}, zmm2, zmm3/m512/m32bcst	A	V/V	AVX10.2
EVEX.128.NP.0F38.W0 D3 /r VPDPWUUDS xmm1{k1}{z}, xmm2, xmm3/m128/m32bcst	A	V/V	AVX10.2

Continued on next page...

10.4. VPDPW[SU,US,UU]D[,S]

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.0F38.W0 D3 /r VPDPWUUDS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst	A	V/V	AVX10.2
EVEX.512.NP.0F38.W0 D3 /r VPDPWUUDS zmm1{k1}{z}, zmm2, zmm3/m512/m32bcst	A	V/V	AVX10.2

10.4.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

10.4.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

Multiplies the individual words of the first source operand by the corresponding words of the second source operand, producing intermediate dword results. The dword results are then summed and accumulated in the destination dword element-size operand.

For unsigned saturation, when an individual result value is beyond the range of an unsigned doubleword (that is, greater than FFFF_FFFFH), the saturated unsigned doubleword integer value of FFFF_FFFFH stored in the doubleword destination.

For signed saturation, when an individual result is beyond the range of a signed doubleword integer (that is, greater than 7FFF_FFFFH or less than 8000_0000H), the saturated value of 7FFF_FFFFH or 8000_0000H, respectively, is written to the destination operand.

The EVEX encoded form of this instruction supports memory fault suppression.

The EVEX version of VPDPWSSD{,S} was previously introduced with AVX512-VNNI. The VEX version of VPDPWSSD{,S} was previously introduced with AVX-VNNI.

10.4. VPDPW[SU,US,UU]D[,S]

10.4.3 OPERATION

```

1  VPDPW[UU,SU,US]D[,S] dest, src1, src2  // VEX version
2  VL = (128, 256)
3  KL = VL/32
4  ORIGDEST := DEST
5
6  IF *src1 is signed*:  // SU
7     src1extend := SIGN_EXTEND
8  ELSE:                // UU,US
9     src1extend := ZERO_EXTEND
10 IF *src2 is signed*: // US
11    src2extend := SIGN_EXTEND
12 ELSE:              // UU, SU
13    src2extend := ZERO_EXTEND
14
15 FOR i := 0 TO KL-1:
16    p1dword := src1extend(SRC1.word[2*i+0]) * src2extend(SRC2.word[2*i+0])
17    p2dword := src1extend(SRC1.word[2*i+1]) * src2extend(SRC2.word[2*i+1])
18    IF *saturating version*:
19        IF *UU instruction version*:
20            DEST.dword[i] := UNSIGNED_DWORD_SATURATE(ORIGDEST.dword[i] + p1dword + p2dword)
21        ELSE:
22            DEST.dword[i] := SIGNED_DWORD_SATURATE(ORIGDEST.dword[i] + p1dword + p2dword)
23    ELSE:
24        DEST.dword[i] := ORIGDEST.dword[i] + p1dword + p2dword
25 DEST[MAX_VL-1:VL] := 0

```


10.4. VPDPW[SU,US,UU]D[,S]

```

1  VPDPW[UU,US,SU]D[,S] dest, src1, src2      // EVEX version
2  VL = (128,256,512)
3  KL = VL/32
4  ORIGDEST := DEST
5
6  IF *src1 is signed*: // SU
7     src1extend := SIGN_EXTEND
8  ELSE: // UU,US
9     src1extend := ZERO_EXTEND
10 IF *src2 is signed*: // US
11    src2extend := SIGN_EXTEND
12 ELSE: // UU, SU
13    src2extend := ZERO_EXTEND
14
15 FOR i := 0 TO KL-1:
16    IF k1[i] or *no writemask*:
17        IF SRC2 is memory and EVEX.b == 1:
18            t := SRC2.dword[0]
19        ELSE:
20            t := SRC2.dword[i]
21
22        p1dword := src1extend(SRC1.word[2*i+0]) * src2extend(t.word[0])
23        p2dword := src1extend(SRC1.word[2*i+1]) * src2extend(t.word[1])
24
25        IF *saturating*:
26            IF *UU instruction version*:
27                DEST.dword[i] := UNSIGNED_DWORD_SATURATE(ORIGDEST.dword[i] +
28                    p1dword + p2dword)
29            ELSE:
30                DEST.dword[i] := SIGNED_DWORD_SATURATE(ORIGDEST.dword[i] +
31                    p1dword + p2dword)
32        ELSE:
33            DEST.dword[i] := ORIGDEST.dword[i] + p1dword + p2dword
34
35    ELSE IF *zeroing*:
36        DEST.dword[i] := 0
37    ELSE: // Merge masking, dest element unchanged
38        DEST.dword[i] := ORIGDEST.dword[i]
39 DEST[MAX_VL-1:VL] := 0

```

10.4.4 EXCEPTIONS

10.4. VPDPW[SU,US,UU]D[,S]

Instruction	Exception Type	Arithmetic Flags	CPUID
VPDPWSUD xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VPDPWSUD ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VPDPWSUD zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VPDPWSUDS xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VPDPWSUDS ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VPDPWSUDS zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VPDPWUSD xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VPDPWUSD ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VPDPWUSD zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VPDPWUSDS xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VPDPWUSDS ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VPDPWUSDS zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VPDPWUUD xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VPDPWUUD ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VPDPWUUD zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2
VPDPWUUDS xmm1, xmm2, xmm3/m128	E4	N/A	AVX10.2
VPDPWUUDS ymm1, ymm2, ymm3/m256	E4	N/A	AVX10.2
VPDPWUUDS zmm1, zmm2, zmm3/m512	E4	N/A	AVX10.2

Chapter 11

INTEL® AVX10 MINMAX INSTRUCTIONS

11.1 VMINMAXNEPBF16

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.F2.0F3A.W0 52 /r /ib VMINMAXNEPBF16 xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst, imm8	A	V/V	AVX10.2
EVEX.256.F2.0F3A.W0 52 /r /ib VMINMAXNEPBF16 ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst, imm8	A	V/V	AVX10.2
EVEX.512.F2.0F3A.W0 52 /r /ib VMINMAXNEPBF16 zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst, imm8	A	V/V	AVX10.2

11.1.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	IMM8(r)

11.1.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

This instruction perform min/max comparison between two SIMD registers. An immediate control selects which comparison operation is performed, and also allows to override the sign of the comparison result.

The numeric behavior of the comparison operations are slightly different than VRANGE* and FP MIN/MAX instructions, and follows the "Minimum and Maximum operations" definitions in IEEE-754-2019 standard section 9.6.

The immediate bytes controls the comparison operation according to Table 11.1 and the sign control according to Table 11.2.

The sign control indication ignores NaN signs: it does not manipulate the sign if the result is a NaN, and does not copy the sign of SRC1 (for sign control = 0b00) if SRC1 is a NaN. NaN propagation and sign control behavior in case of NaN operands are described in Table 11.3 and Table 11.4.

This instruction does not generate floating point exceptions and does not consult or update MXCSR. Denormal bfloat16 input operands are treated as zeros (DAZ) and denormal bfloat16 outputs are flushed to zero (FTZ).

11.1.3 OPERATION

```

1  VMINMAXNEPBF16 dest {k1}, src1, src2, imm8
2  VL = 128, 256 or 512
3  KL := VL / 16
4
5  for i := 0 to KL-1:
6      if k1[i] or *no writemask*:
7          if src2 is memory and (EVEX.b == 1):
8              dest.bf16[i] := minmax(src1.bf16[i], src2.bf16[0],
9                                  imm8, daz=true, except=false)
10         else:
11             dest.bf16[i] := minmax(src1.bf16[i], src2.bf16[i],
12                                 imm8, daz=true, except=false)
13         else if *zeroing*:
14             dest.bf16[i] := 0
15         //else dest.bf16[i] remains unchanged
16
17 dest[MAX_VL-1:VL] := 0

```

11.1.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMINMAXNEPBF16 xmm1, xmm2, xmm3/m128, imm8	E4	N/A	AVX10.2
VMINMAXNEPBF16 ymm1, ymm2, ymm3/m256, imm8	E4	N/A	AVX10.2
VMINMAXNEPBF16 zmm1, zmm2, zmm3/m512, imm8	E4	N/A	AVX10.2

11.2 VMINMAX[PH,PS,PD]

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.66.0F3A.W1 52 /r /ib VMINMAXPD xmm1{k1}{z}, xmm2, xmm3/m128/m64bcst, imm8	A	V/V	AVX10.2
EVEX.256.66.0F3A.W1 52 /r /ib VMINMAXPD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {sae}, imm8	A	V/V	AVX10.2
EVEX.512.66.0F3A.W1 52 /r /ib VMINMAXPD zmm1{k1}{z}, zmm2, zmm3/m512/m64bcst {sae}, imm8	A	V/V	AVX10.2
EVEX.128.NP.0F3A.W0 52 /r /ib VMINMAXPH xmm1{k1}{z}, xmm2, xmm3/m128/m16bcst, imm8	A	V/V	AVX10.2
EVEX.256.NP.0F3A.W0 52 /r /ib VMINMAXPH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {sae}, imm8	A	V/V	AVX10.2
EVEX.512.NP.0F3A.W0 52 /r /ib VMINMAXPH zmm1{k1}{z}, zmm2, zmm3/m512/m16bcst {sae}, imm8	A	V/V	AVX10.2
EVEX.128.66.0F3A.W0 52 /r /ib VMINMAXPS xmm1{k1}{z}, xmm2, xmm3/m128/m32bcst, imm8	A	V/V	AVX10.2
EVEX.256.66.0F3A.W0 52 /r /ib VMINMAXPS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {sae}, imm8	A	V/V	AVX10.2
EVEX.512.66.0F3A.W0 52 /r /ib VMINMAXPS zmm1{k1}{z}, zmm2, zmm3/m512/m32bcst {sae}, imm8	A	V/V	AVX10.2

11.2.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	IMM8(r)

11.2.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

This instruction perform min/max comparison between two SIMD registers. An immediate control selects which comparison operation is performed, and also allows to override the sign of the comparison result.

The numeric behavior of the comparison operations are slightly different than VRANGE* and FP MIN/MAX instructions, and follows the "Minimum and Maximum operations" definitions in IEEE-754-2019 standard section 9.6.

Table 11.1 describes the encoding of the imm8 operand. Bits [7:5] are reserved. #IE signalling behavior differs between operations which do not have a "Number" suffix¹ and those that are suffixed with "Number"².

This instruction raises Invalid exception (#IE) if either of the operands is an SNAN, and a denormal exception (#DE) if either operand is a denormal and none of the operands are NAN.

The sign control indication ignores NAN signs: it does not manipulate the sign if the result is a NAN, and does not copy the sign of SRC1 (for sign control = 0b00) if SRC1 is a NAN. NaN propagation and sign control behavior in case of NaN operands are described in Table 11.3 and Table 11.4.

¹-0 < +0, #IE is reported in case of any sNaN source operand

²-0 < +0, #IE if both are NaN, and either one is sNaN

imm8[4] min/max	imm8[1:0] Op-select	Operation	Description
0b0	0b00	minimum ¹	x if $x \leq y$, y if $y < x$, and a quiet NaN if either operand is a NaN.
0b0	0b01	maximum ¹	x if $x \geq y$, y if $y > x$, and a quiet NaN if either operand is a NaN.
0b0	0b10	minimumMagnitude ¹	x if $ x < y $, y if $ y < x $, otherwise minimum(x, y).
0b0	0b11	maximumMagnitude ¹	x if $ x > y $, y if $ y > x $, otherwise maximum(x, y).
0b1	0b00	minimumNumber ²	x if $x \leq y$, y if $y < x$, and the number if one operand is a number and the other is a NaN (including signaling NaN which is ignored and not converted to a quiet NaN). If both operands are NaNs, a quiet NaN is returned. If either operand is a signaling NaN, an invalid operation exception is signaled.
0b1	0b01	maximumNumber ²	x if $x \geq y$, y if $y > x$, and the number if one operand is a number and the other is a NaN (including signaling NaN which is ignored and not converted to a quiet NaN). If both operands are NaNs, a quiet NaN is returned. If either operand is a signaling NaN, an invalid operation exception is signaled.
0b1	0b10	minimumMagnitudeNumber ²	x if $ x < y $, y if $ y < x $, otherwise minimumNumber(x, y).
0b1	0b11	maximumMagnitudeNumber ²	x if $ x > y $, y if $ y > x $, otherwise maximumNumber(x, y).

Table 11.1: MINMAX operation selection according to imm8[4] and imm8[1:0].

imm8[3:2] Sign control	Sign
0b00	Select sign (src1)
0b01	Select sign (compare result)
0b10	Set sign to 0
0b11	Set sign to 1

Table 11.2: MINMAX sign control according to imm8[3:2].

src1	src2	Result	IE Signaling Due To Comparison	Imm8[3:2] effect to range output
sNaN1	sNaN2	Quiet(sNaN1)	Yes	Ignored
sNaN1	qNaN2	Quiet(sNaN1)	Yes	Ignored
sNaN1	Norm2	Quiet(sNaN1)	Yes	Ignored
qNaN1	sNaN2	Quiet(sNaN2)	Yes	Ignored
qNaN1	qNaN2	qNaN1	No	Ignored
qNaN1	Norm2	qNaN1	No	Ignored
Norm1	sNaN2	Quiet(sNaN2)	Yes	Ignored
Norm1	qNaN2	qNaN2	No	Ignored

Table 11.3: NaN propagation of one or more NaN input values and effect of imm8[3:2] for minimum, minimumMagnitude, maximum, maximumMagnitude MINMAX operations.

src1	src2	Result	IE Signaling Due To Comparison	Imm8[3:2] effect to range output
sNaN1	sNaN2	Quiet(sNaN1)	Yes	Ignored
sNaN1	qNaN2	Quiet(sNaN1)	Yes	Ignored
sNaN1	Norm2	Norm2	Yes	Imm8[3:2] == 0b00 ignored other values applicable
qNaN1	sNaN2	Quiet(sNaN2)	Yes	Ignored
qNaN1	qNaN2	qNaN1	No	Ignored
qNaN1	Norm2	Norm2	No	Imm8[3:2] == 0b00 ignored other values applicable
Norm1	sNaN2	Norm1	Yes	Applicable
Norm1	qNaN2	Norm1	No	Applicable

Table 11.4: NaN propagation of one or more NaN input values and effect of imm8[3:2] for minimumNumber, minimumMagnitudeNumber, maximumNumber, maximumMagnitudeNumber MINMAX operations.

imm8[4] min/max	imm8[1:0] Op-select	Operation	Comparison Result for Opposite-Signed Zero
0	00	Minimum	-0
1	00	MinimumNumber	-0
0	10	MinimumMagnitude	-0
1	10	MinimumMagnitudeNumber	-0
0	01	Maximum	+0
1	01	MaximumNumber	+0
0	11	MaximumMagnitude	+0
1	11	MaximumMagnitudeNumber	+0

Table 11.5: MINMAX Operation Behavior With Signed Comparison of Opposite-Signed Zeros (src1=-0 and src2=+0, or src1=+0 and src2=-0)

imm8[4] min/max	imm8[1:0] Op-select	Operation	Comparison Result for Equal Magnitude Inputs
0	00	Minimum	b
1	00	MinimumNumber	b
0	10	MinimumMagnitude	b
1	10	MinimumMagnitudeNumber	b
0	01	Maximum	a
1	01	MaximumNumber	a
0	11	MaximumMagnitude	a
1	11	MaximumMagnitudeNumber	a

Table 11.6: MINMAX Operation Behavior With Equal Magnitude Comparisons (src1=a and src2=b, or src1=b and src2=a, where $|a|=|b|$ and $a>0$ and $b<0$)

11.2.3 OPERATION

```

1  VMINMAXPD dest {k1}, src1, src2, imm8
2  VL = 128, 256 or 512
3  KL := VL / 64
4
5  for i := 0 to KL-1:
6      if k1[i] or *no writemask*:
7          if src2 is memory and (EVEX.b == 1):
8              dest.f64[i] := minmax(src1.f64[i], src2.f64[0],
9                                  imm8, daz=MXCSR.DAZ, except=true)
10         else:
11             dest.f64[i] := minmax(src1.f64[i], src2.f64[i],
12                                 imm8, daz=MXCSR.DAZ, except=true)
13         else if *zeroing*:
14             dest.f64[i] := 0
15         //else dest.f64[i] remains unchanged
16
17 dest[MAX_VL-1:VL] := 0

```

```

1  VMINMAXPS dest {k1}, src1, src2, imm8
2  VL = 128, 256 or 512
3  KL := VL / 32
4
5  for i := 0 to KL-1:
6      if k1[i] or *no writemask*:
7          if src2 is memory and (EVEX.b == 1):
8              dest.f32[i] := minmax(src1.f32[i], src2.f32[0],
9                                  imm8, daz=MXCSR.DAZ, except=true)
10         else:
11             dest.f32[i] := minmax(src1.f32[i], src2.f32[i],
12                                 imm8, daz=MXCSR.DAZ, except=true)
13         else if *zeroing*:
14             dest.f32[i] := 0
15         //else dest.f32[i] remains unchanged
16
17 dest[MAX_VL-1:VL] := 0

```

```

1  VMINMAXPH dest {k1}, src1, src2, imm8
2  VL = 128, 256 or 512
3  KL := VL / 16
4
5  for i := 0 to KL-1:
6      if k1[i] or *no writemask*:
7          if src2 is memory and (EVEX.b == 1):
8              dest.f16[i] := minmax(src1.f16[i], src2.f16[0],
9                                  imm8, daz=false, except=true)
10         else:
11             dest.f16[i] := minmax(src1.f16[i], src2.f16[i],
12                                 imm8, daz=false, except=true)
13     else if *zeroing*:
14         dest.f16[i] := 0
15     //else dest.f16[i] remains unchanged
16
17 dest[MAX_VL-1:VL] := 0

```

11.2.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

11.2.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMINMAXPD xmm1, xmm2, xmm3/m128, imm8	E2	ID	AVX10.2
VMINMAXPD ymm1, ymm2, ymm3/m256, imm8	E2	ID	AVX10.2
VMINMAXPD zmm1, zmm2, zmm3/m512, imm8	E2	ID	AVX10.2
VMINMAXPH xmm1, xmm2, xmm3/m128, imm8	E2	ID	AVX10.2
VMINMAXPH ymm1, ymm2, ymm3/m256, imm8	E2	ID	AVX10.2

Continued on next page...

Instruction	Exception Type	Arithmetic Flags	CPUID
VMINMAXPH zmm1, zmm2, zmm3/m512, imm8	E2	ID	AVX10.2
VMINMAXPS xmm1, xmm2, xmm3/m128, imm8	E2	ID	AVX10.2
VMINMAXPS ymm1, ymm2, ymm3/m256, imm8	E2	ID	AVX10.2
VMINMAXPS zmm1, zmm2, zmm3/m512, imm8	E2	ID	AVX10.2

11.3 VMINMAX[SH,SS,SD]

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.LLIG.66.0F3A.W1 53 /r /ib VMINMAXSD xmm1{k1}{z}, xmm2, xmm3/m64 {sae}, imm8	A	V/V	AVX10.2
EVEX.LLIG.NP.0F3A.W0 53 /r /ib VMINMAXSH xmm1{k1}{z}, xmm2, xmm3/m16 {sae}, imm8	A	V/V	AVX10.2
EVEX.LLIG.66.0F3A.W0 53 /r /ib VMINMAXSS xmm1{k1}{z}, xmm2, xmm3/m32 {sae}, imm8	A	V/V	AVX10.2

11.3.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	SCALAR	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	IMM8(r)

11.3.2 DESCRIPTION

Note:

For SCALAR instructions with a CPUID feature flag specifying Intel® AVX10, the programmer can skip the check for the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and usually determines the set of instructions available to the programmer listed in the above opcode table. However, scalar instructions are architected without the notion of vector length, so this step does not strictly apply.

This instruction perform min/max comparison between two elements in SIMD registers. An immediate control selects which comparison operation is performed, and also allows to override the sign of the comparison result.

The numeric behavior of the comparison operations are slightly different than VRANGE* and FP MIN/MAX instructions, and follows the "Minimum and Maximum operations" definitions in IEEE-754-2019 standard section 9.6.

The immediate bytes controls the comparison operation according to Table 11.1 and the sign control according to Table 11.2.

This instruction raises Invalid exception (#IE) if either of the operands is an SNAN, and a denormal exception (#DE) if either operand is a denormal and none of the operands are NAN.

The sign control indication ignores NaN signs: it does not manipulate the sign if the result is a NaN, and does not copy the sign of SRC1 (for sign control = 0b00) if SRC1 is a NaN. NaN propagation and sign control behavior in case of NaN operands are described in Table 11.3 and Table 11.4.

Bits 127:16/32/64 are copied to the destination from the respective elements of the first operand

11.3.3 OPERATION

```
1 VMINMAXSD dest {k1}, src1, src2, imm8
2 VL = 128, 256 or 512
3
4 if k1[0] or *no writemask*:
5     dest.f64[0] := minmax(src1.f64[0], src2.f64[0],
6                          imm8, daz=MXCSR.DAZ, except=true)
7 else if *zeroing*:
8     dest.f64[0] := 0
9 //else dest.f64[0] remains unchanged
10
11 dest[127:64] := src1[127:64]
12 dest[MAX_VL-1:VL] := 0
```

```
1 VMINMAXSS dest {k1}, src1, src2, imm8
2 VL = 128, 256 or 512
3
4 if k1[0] or *no writemask*:
5     dest.f32[0] := minmax(src1.f32[0], src2.f32[0],
6                          imm8, daz=MXCSR.DAZ, except=true)
7 else if *zeroing*:
8     dest.f32[0] := 0
9 //else dest.f32[0] remains unchanged
10
11 dest[127:32] := src1[127:32]
12 dest[MAX_VL-1:VL] := 0
```

```

1 VMINMAXSH dest {k1}, src1, src2, imm8
2 VL = 128, 256 or 512
3
4 if k1[0] or *no writemask*:
5     dest.f16[0] := minmax(src1.f16[0], src2.f16[0],
6                           imm8, daz=false, except=true)
7 else if *zeroing*:
8     dest.f16[0] := 0
9 //else dest.f16[0] remains unchanged
10
11 dest[127:16] := src1[127:16]
12 dest[MAX_VL-1:VL] := 0

```

11.3.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

11.3.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMINMAXSD xmm1, xmm2, xmm3/m64, imm8	E3	ID	AVX10.2
VMINMAXSH xmm1, xmm2, xmm3/m16, imm8	E3	ID	AVX10.2
VMINMAXSS xmm1, xmm2, xmm3/m32, imm8	E3	ID	AVX10.2

Chapter 12

INTEL® AVX10 NEW INSTRUCTION FORMS

12.1 VADDPD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W1 58 /r VADDPD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.1.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.1.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.1.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.1.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VADDPD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.2 VADDPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.MAP5.W0 58 /r VADDPH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.2.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.2.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.2.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.2.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VADDPH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.3 VADDPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF.W0 58 /r VADDPS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.3.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.3.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.3.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.3.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VADDPS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.4 VCMPPD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W1 C2 /r /ib VCMPPD k1{k2}, ymm2, ymm3/m256/m64bcst {sae}, imm8	A	V/V	AVX10.2

12.4.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	IMM8(r)

12.4.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.4.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.4.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCMPPD k1, ymm2, ymm3/m256, imm8	E2	ID	AVX10.2

12.5 VCMPPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF3A.W0 C2 /r /ib VCMPPH k1{k2}, ymm2, ymm3/m256/m16bcst {sae}, imm8	A	V/V	AVX10.2

12.5.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	IMM8(r)

12.5.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.5.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.5.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCMPPH k1, ymm2, ymm3/m256, imm8	E2	ID	AVX10.2

12.6 VCMPPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF.WO C2 /r /ib VCMPPS k1{k2}, ymm2, ymm3/m256/m32bcst {sae}, imm8	A	V/V	AVX10.2

12.6.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	IMM8(r)

12.6.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.6.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.6.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCMPPS k1, ymm2, ymm3/m256, imm8	E2	ID	AVX10.2

12.7 VCVTDQ2PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.MAP5.W0 5B /r VCVTDQ2PH xmm1{k1}{z}, ymm2/m256/m32bcst {er}	A	V/V	AVX10.2

12.7.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.7.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.7.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Overflow, Precision In abbreviated form, this includes: OE, PE

12.7.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTDQ2PH xmm1, ymm2/m256	E2	PO	AVX10.2

12.8 VCVTDQ2PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF.W0 5B /r VCVTDQ2PS ymm1{k1}{z}, ymm2/m256/m32bcst {er}	A	V/V	AVX10.2

12.8.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.8.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.8.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Precision In abbreviated form, this includes: PE

12.8.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTDQ2PS ymm1, ymm2/m256	E2	P	AVX10.2

12.9 VCVTPD2DQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F2.0F.W1 E6 /r VCVTPD2DQ xmm1{k1}{z}, ymm2/m256/m64bcst {er}	A	V/V	AVX10.2

12.9.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.9.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.9.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.9.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPD2DQ xmm1, ymm2/m256	E2	IP	AVX10.2

12.10 VCVTPD2PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP5.W1 5A /r VCVTPD2PH xmm1{k1}{z}, ymm2/m256/m64bcst {er}	A	V/V	AVX10.2

12.10.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.10.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.10.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.10.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPD2PH xmm1, ymm2/m256	E2	IOUPD	AVX10.2

12.11 VCVTPD2PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W1 5A /r VCVTPD2PS xmm1{k1}{z}, ymm2/m256/m64bcst {er}	A	V/V	AVX10.2

12.11.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.11.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.11.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.11.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPD2PS xmm1, ymm2/m256	E2	IOUPD	AVX10.2

12.12 VCVTPD2QQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W1 7B /r VCVTPD2QQ ymm1{k1}{z}, ymm2/m256/m64bcst {er}	A	V/V	AVX10.2

12.12.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.12.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.12.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.12.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPD2QQ ymm1, ymm2/m256	E2	IP	AVX10.2

12.13 VCVTPD2UDQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF.W1 79 /r VCVTPD2UDQ xmm1{k1}{z}, ymm2/m256/m64bcst {er}	A	V/V	AVX10.2

12.13.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.13.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.13.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.13.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPD2UDQ xmm1, ymm2/m256	E2	IP	AVX10.2

12.14 VCVTPD2UQQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W1 79 /r VCVTPD2UQQ ymm1{k1}{z}, ymm2/m256/m64bcst {er}	A	V/V	AVX10.2

12.14.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.14.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.14.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.14.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPD2UQQ ymm1, ymm2/m256	E2	IP	AVX10.2

12.15 VCVTPH2DQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP5.W0 5B /r VCVTPH2DQ ymm1{k1}{z}, xmm2/m128/m16bcst {er}	A	V/V	AVX10.2

12.15.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	HALF	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.15.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.15.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.15.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPH2DQ ymm1, xmm2/m128	E2	IP	AVX10.2

12.16 VCVTPH2PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.MAP5.W0 5A /r VCVTPH2PD ymm1{k1}{z}, xmm2/m64/m16bcst {sae}	A	V/V	AVX10.2

12.16.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	QUARTER	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.16.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.16.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.16.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPH2PD ymm1, xmm2/m64	E2	ID	AVX10.2

12.17 VCVTPH2PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 13 /r VCVTPH2PS ymm1{k1}{z}, xmm2/m128 {sae}	A	V/V	AVX10.2

12.17.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	HALFMEM	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.17.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.17.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid In abbreviated form, this includes: IE

12.17.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPH2PS ymm1, xmm2/m128	E11	I	AVX10.2

12.18 VCVTPH2PSX

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 13 /r VCVTPH2PSX ymm1{k1}{z}, xmm2/m128/m16bcst {sae}	A	V/V	AVX10.2

12.18.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	HALF	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.18.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.18.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.18.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPH2PSX ymm1, xmm2/m128	E2	ID	AVX10.2

12.19 VCVTPH2QQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP5.W0 7B /r VCVTPH2QQ ymm1{k1}{z}, xmm2/m64/m16bcst {er}	A	V/V	AVX10.2

12.19.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	QUARTER	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.19.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.19.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.19.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPH2QQ ymm1, xmm2/m64	E2	IP	AVX10.2

12.20 VCVTPH2UDQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.MAP5.W0 79 /r VCVTPH2UDQ ymm1{k1}{z}, xmm2/m128/m16bcst {er}	A	V/V	AVX10.2

12.20.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	HALF	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.20.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.20.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.20.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPH2UDQ ymm1, xmm2/m128	E2	IP	AVX10.2

12.21 VCVTPH2UQQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP5.W0 79 /r VCVTPH2UQQ ymm1{k1}{z}, xmm2/m64/m16bcst {er}	A	V/V	AVX10.2

12.21.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	QUARTER	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.21.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.21.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.21.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPH2UQQ ymm1, xmm2/m64	E2	IP	AVX10.2

12.22 VCVTPH2UW

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.MAP5.W0 7D /r VCVTPH2UW ymm1{k1}{z}, ymm2/m256/m16bcst {er}	A	V/V	AVX10.2

12.22.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.22.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.22.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.22.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPH2UW ymm1, ymm2/m256	E2	IP	AVX10.2

12.23 VCVTPH2W

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP5.W0 7D /r VCVTPH2W ymm1{k1}{z}, ymm2/m256/m16bcst {er}	A	V/V	AVX10.2

12.23.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.23.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.23.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.23.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPH2W ymm1, ymm2/m256	E2	IP	AVX10.2

12.24 VCVTPS2DQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W0 5B /r VCVTPS2DQ ymm1{k1}{z}, ymm2/m256/m32bcst {er}	A	V/V	AVX10.2

12.24.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.24.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.24.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.24.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPS2DQ ymm1, ymm2/m256	E2	IP	AVX10.2

12.25 VCVTSP2PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF.W0 5A /r VCVTSP2PD ymm1{k1}{z}, xmm2/m128/m32bcst {sae}	A	V/V	AVX10.2

12.25.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	HALF	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.25.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.25.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.25.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTSP2PD ymm1, xmm2/m128	E2	ID	AVX10.2

12.26 VCVTSP2PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F3A.W0 1D 11:rrr:bbb /ib VCVTSP2PH xmm1{k1}{z}, ymm2 {sae}, imm8	A	V/V	AVX10.2

12.26.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	HALFMEM	MODRM.R/M(w)	MODRM.REG(r)	IMM8(r)	N/A

12.26.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.26.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.26.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTSP2PH xmm1, ymm2, imm8	E11	IOUPD	AVX10.2

12.27 VCVTSP2PHX

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP5.W0 1D /r VCVTSP2PHX xmm1{k1}{z}, ymm2/m256/m32bcst {er}	A	V/V	AVX10.2

12.27.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.27.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.27.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.27.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTSP2PHX xmm1, ymm2/m256	E2	IOUPD	AVX10.2

12.28 VCVTPS2QQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W0 7B /r VCVTPS2QQ ymm1{k1}{z}, xmm2/m128/m32bcst {er}	A	V/V	AVX10.2

12.28.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	HALF	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.28.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.28.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.28.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPS2QQ ymm1, xmm2/m128	E2	IP	AVX10.2

12.29 VCVTPS2UDQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF.WO 79 /r VCVTPS2UDQ ymm1{k1}{z}, ymm2/m256/m32bcst {er}	A	V/V	AVX10.2

12.29.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.29.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.29.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.29.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPS2UDQ ymm1, ymm2/m256	E2	IP	AVX10.2

12.30 VCVTPS2UQQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W0 79 /r VCVTPS2UQQ ymm1{k1}{z}, xmm2/m128/m32bcst {er}	A	V/V	AVX10.2

12.30.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	HALF	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.30.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.30.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.30.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPS2UQQ ymm1, xmm2/m128	E2	IP	AVX10.2

12.31 VCVTQQ2PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F3.0F.W1 E6 /r VCVTQQ2PD ymm1{k1}{z}, ymm2/m256/m64bcst {er}	A	V/V	AVX10.2

12.31.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.31.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.31.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Precision In abbreviated form, this includes: PE

12.31.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTQQ2PD ymm1, ymm2/m256	E2	P	AVX10.2

12.32 VCVTQQ2PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.MAP5.W1 5B /r VCVTQQ2PH xmm1{k1}{z}, ymm2/m256/m64bcst {er}	A	V/V	AVX10.2

12.32.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.32.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.32.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Overflow, Precision In abbreviated form, this includes: OE, PE

12.32.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTQQ2PH xmm1, ymm2/m256	E2	PO	AVX10.2

12.33 VCVTQQ2PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF.W1 5B /r VCVTQQ2PS xmm1{k1}{z}, ymm2/m256/m64bcst {er}	A	V/V	AVX10.2

12.33.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.33.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.33.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Precision In abbreviated form, this includes: PE

12.33.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTQQ2PS xmm1, ymm2/m256	E2	P	AVX10.2

12.34 VCVTTPD2DQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W1 E6 /r VCVTTPD2DQ xmm1{k1}{z}, ymm2/m256/m64bcst {sae}	A	V/V	AVX10.2

12.34.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.34.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.34.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.34.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPD2DQ xmm1, ymm2/m256	E2	IP	AVX10.2

12.35 VCVTTPD2QQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W1 7A /r VCVTTPD2QQ ymm1{k1}{z}, ymm2/m256/m64bcst {sae}	A	V/V	AVX10.2

12.35.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.35.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.35.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.35.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPD2QQ ymm1, ymm2/m256	E2	IP	AVX10.2

12.36 VCVTTPD2UDQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF.W1 78 /r VCVTTPD2UDQ xmm1{k1}{z}, ymm2/m256/m64bcst {sae}	A	V/V	AVX10.2

12.36.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.36.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.36.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.36.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPD2UDQ xmm1, ymm2/m256	E2	IP	AVX10.2

12.37 VCVTTPD2UQQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W1 78 /r VCVTTPD2UQQ ymm1{k1}{z}, ymm2/m256/m64bcst {sae}	A	V/V	AVX10.2

12.37.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.37.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.37.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.37.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPD2UQQ ymm1, ymm2/m256	E2	IP	AVX10.2

12.38 VCVTTPH2DQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F3.MAP5.W0 5B /r VCVTTPH2DQ ymm1{k1}{z}, xmm2/m128/m16bcst {sae}	A	V/V	AVX10.2

12.38.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	HALF	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.38.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.38.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.38.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPH2DQ ymm1, xmm2/m128	E2	IP	AVX10.2

12.39 VCVTTPH2QQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP5.W0 7A /r VCVTTPH2QQ ymm1{k1}{z}, xmm2/m64/m16bcst {sae}	A	V/V	AVX10.2

12.39.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	QUARTER	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.39.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.39.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.39.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPH2QQ ymm1, xmm2/m64	E2	IP	AVX10.2

12.40 VCVTTPH2UDQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.MAP5.W0 78 /r VCVTTPH2UDQ ymm1{k1}{z}, xmm2/m128/m16bcst {sae}	A	V/V	AVX10.2

12.40.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	HALF	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.40.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.40.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.40.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPH2UDQ ymm1, xmm2/m128	E2	IP	AVX10.2

12.41 VCVTTPH2UQQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP5.W0 78 /r VCVTTPH2UQQ ymm1{k1}{z}, xmm2/m64/m16bcst {sae}	A	V/V	AVX10.2

12.41.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	QUARTER	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.41.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.41.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.41.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPH2UQQ ymm1, xmm2/m64	E2	IP	AVX10.2

12.42 VCVTTPH2UW

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.MAP5.W0 7C /r VCVTTPH2UW ymm1{k1}{z}, ymm2/m256/m16bcst {sae}	A	V/V	AVX10.2

12.42.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.42.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.42.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.42.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPH2UW ymm1, ymm2/m256	E2	IP	AVX10.2

12.43 VCVTTPH2W

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP5.W0 7C /r VCVTTPH2W ymm1{k1}{z}, ymm2/m256/m16bcst {sae}	A	V/V	AVX10.2

12.43.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.43.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.43.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.43.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPH2W ymm1, ymm2/m256	E2	IP	AVX10.2

12.44 VCVTTPS2DQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F3.0F.W0 5B /r VCVTTPS2DQ ymm1{k1}{z}, ymm2/m256/m32bcst {sae}	A	V/V	AVX10.2

12.44.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.44.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.44.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.44.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPS2DQ ymm1, ymm2/m256	E2	IP	AVX10.2

12.45 VCVTTPS2QQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W0 7A /r VCVTTPS2QQ ymm1{k1}{z}, xmm2/m128/m32bcst {sae}	A	V/V	AVX10.2

12.45.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	HALF	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.45.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.45.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.45.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPS2QQ ymm1, xmm2/m128	E2	IP	AVX10.2

12.46 VCVTTPS2UDQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF.W0 78 /r VCVTTPS2UDQ ymm1{k1}{z}, ymm2/m256/m32bcst {sae}	A	V/V	AVX10.2

12.46.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.46.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.46.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.46.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPS2UDQ ymm1, ymm2/m256	E2	IP	AVX10.2

12.47 VCVTTPS2UQQ

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W0 78 /r VCVTTPS2UQQ ymm1{k1}{z}, xmm2/m128/m32bcst {sae}	A	V/V	AVX10.2

12.47.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	HALF	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.47.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.47.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.47.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPS2UQQ ymm1, xmm2/m128	E2	IP	AVX10.2

12.48 VCVTUDQ2PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F2.MAP5.W0 7A /r VCVTUDQ2PH xmm1{k1}{z}, ymm2/m256/m32bcst {er}	A	V/V	AVX10.2

12.48.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.48.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.48.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Overflow, Precision In abbreviated form, this includes: OE, PE

12.48.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTUDQ2PH xmm1, ymm2/m256	E2	PO	AVX10.2

12.49 VCVTUDQ2PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F2.0F.W0 7A /r VCVTUDQ2PS ymm1{k1}{z}, ymm2/m256/m32bcst {er}	A	V/V	AVX10.2

12.49.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.49.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.49.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Precision In abbreviated form, this includes: PE

12.49.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTUDQ2PS ymm1, ymm2/m256	E2	P	AVX10.2

12.50 VCVTUQQ2PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F3.0F.W1 7A /r VCVTUQQ2PD ymm1{k1}{z}, ymm2/m256/m64bcst {er}	A	V/V	AVX10.2

12.50.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.50.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.50.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Precision In abbreviated form, this includes: PE

12.50.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTUQQ2PD ymm1, ymm2/m256	E2	P	AVX10.2

12.51 VCVTUQQ2PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F2.MAP5.W1 7A /r VCVTUQQ2PH xmm1{k1}{z}, ymm2/m256/m64bcst {er}	A	V/V	AVX10.2

12.51.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.51.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.51.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Overflow, Precision In abbreviated form, this includes: OE, PE

12.51.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTUQQ2PH xmm1, ymm2/m256	E2	PO	AVX10.2

12.52 VCVTUQQ2PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F2.0F.W1 7A /r VCVTUQQ2PS xmm1{k1}{z}, ymm2/m256/m64bcst {er}	A	V/V	AVX10.2

12.52.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.52.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.52.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Precision In abbreviated form, this includes: PE

12.52.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTUQQ2PS xmm1, ymm2/m256	E2	P	AVX10.2

12.53 VCVT UW2PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F2.MAP5.W0 7D /r VCVTUW2PH ymm1{k1}{z}, ymm2/m256/m16bcst {er}	A	V/V	AVX10.2

12.53.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.53.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.53.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Overflow, Precision In abbreviated form, this includes: OE, PE

12.53.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTUW2PH ymm1, ymm2/m256	E2	PO	AVX10.2

12.54 VCVTW2PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F3.MAP5.W0 7D /r VCVTW2PH ymm1{k1}{z}, ymm2/m256/m16bcst {er}	A	V/V	AVX10.2

12.54.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.54.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.54.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Precision In abbreviated form, this includes: PE

12.54.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTW2PH ymm1, ymm2/m256	E2	P	AVX10.2

12.55 VDIVPD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W1 5E /r VDIVPD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.55.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.55.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.55.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Divide-by-Zero, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE, ZE

12.55.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VDIVPD ymm1, ymm2, ymm3/m256	E2	IOUPDZ	AVX10.2

12.56 VDIVPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.MAP5.W0 5E /r VDIVPH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.56.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.56.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.56.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Divide-by-Zero, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE, ZE

12.56.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VDIVPH ymm1, ymm2, ymm3/m256	E2	IOUPDZ	AVX10.2

12.57 VDIVPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF.WO 5E /r VDIVPS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.57.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.57.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.57.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Divide-by-Zero, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE, ZE

12.57.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VDIVPS ymm1, ymm2, ymm3/m256	E2	IOUPDZ	AVX10.2

12.58 VFCMADDCPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F2.MAP6.W0 56 /r VFCMADDCPH ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.58.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.58.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.58.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.58.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFCMADDCPH ymm1, ymm2, ymm3/m256	E4	IOUPD	AVX10.2

12.59 VFCMULCPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F2.MAP6.W0 D6 /r VFCMULCPH ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.59.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.59.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.59.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.59.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFCMULCPH ymm1, ymm2, ymm3/m256	E4	IOUPD	AVX10.2

12.60 VFIXUPIMMPD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F3A.W1 54 /r /ib VFIXUPIMMPD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {sae}, imm8	A	V/V	AVX10.2

12.60.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	IMM8(r)

12.60.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.60.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Divide-by-Zero, Invalid In abbreviated form, this includes: IE, ZE

12.60.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFIXUPIMMPD ymm1, ymm2, ymm3/m256, imm8	E2	IZ	AVX10.2

12.61 VFIXUPIMMPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F3A.W0 54 /r /ib VFIXUPIMMPS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {sae}, imm8	A	V/V	AVX10.2

12.61.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	IMM8(r)

12.61.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.61.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Divide-by-Zero, Invalid In abbreviated form, this includes: IE, ZE

12.61.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFIXUPIMMPS ymm1, ymm2, ymm3/m256, imm8	E2	IZ	AVX10.2

12.62 VFMADD132PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 98 /r VFMADD132PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.62.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.62.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.62.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.62.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD132PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.63 VFMADD132PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 98 /r VFMADD132PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.63.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.63.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.63.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.63.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD132PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.64 VFMADD132PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 98 /r VFMADD132PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.64.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.64.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.64.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.64.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD132PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.65 VFMADD213PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 A8 /r VFMADD213PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.65.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.65.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.65.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.65.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD213PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.66 VFMADD213PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 A8 /r VFMADD213PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.66.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.66.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.66.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.66.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD213PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.67 VFMADD213PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 A8 /r VFMADD213PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.67.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.67.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.67.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.67.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD213PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.68 VFMADD231PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 B8 /r VFMADD231PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.68.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.68.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.68.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.68.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD231PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.69 VFMADD231PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 B8 /r VFMADD231PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.69.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.69.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.69.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.69.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD231PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.70 VFMADD231PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 B8 /r VFMADD231PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.70.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.70.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.70.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.70.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD231PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.71 VFMADDCPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F3.MAP6.W0 56 /r VFMADDCPH ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.71.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.71.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.71.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.71.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADDCPH ymm1, ymm2, ymm3/m256	E4	IOUPD	AVX10.2

12.72 VFMADDSUB132PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 96 /r VFMADDSUB132PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.72.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.72.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.72.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.72.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADDSUB132PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.73 VFMADDSUB132PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 96 /r VFMADDSUB132PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.73.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.73.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.73.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.73.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADDSUB132PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.74 VFMADDSUB132PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 96 /r VFMADDSUB132PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.74.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.74.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.74.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.74.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADDSUB132PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.75 VFMADDSUB213PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 A6 /r VFMADDSUB213PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.75.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.75.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.75.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.75.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADDSUB213PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.76 VFMADDSUB213PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 A6 /r VFMADDSUB213PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.76.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.76.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.76.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.76.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADDSUB213PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.77 VFMADDSUB213PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 A6 /r VFMADDSUB213PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.77.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.77.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.77.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.77.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADDSUB213PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.78 VFMADDSUB231PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 B6 /r VFMADDSUB231PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.78.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.78.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.78.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.78.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADDSUB231PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.79 VFMADDSUB231PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 B6 /r VFMADDSUB231PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.79.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.79.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.79.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.79.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADDSUB231PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.80 VFMADDSUB231PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 B6 /r VFMADDSUB231PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.80.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.80.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.80.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.80.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADDSUB231PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.81 VFMSUB132PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 9A /r VFMSUB132PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.81.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.81.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.81.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.81.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUB132PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.82 VFMSUB132PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 9A /r VFMSUB132PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.82.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.82.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.82.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.82.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUB132PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.83 VFMSUB132PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 9A /r VFMSUB132PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.83.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.83.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.83.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.83.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUB132PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.84 VFMSUB213PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 AA /r VFMSUB213PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.84.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.84.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.84.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.84.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUB213PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.85 VFMSUB213PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 AA /r VFMSUB213PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.85.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.85.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.85.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.85.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUB213PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.86 VFMSUB213PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 AA /r VFMSUB213PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.86.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.86.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.86.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.86.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUB213PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.87 VFMSUB231PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 BA /r VFMSUB231PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.87.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.87.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.87.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.87.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUB231PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.88 VFMSUB231PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 BA /r VFMSUB231PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.88.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.88.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.88.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.88.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUB231PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.89 VFMSUB231PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 BA /r VFMSUB231PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.89.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.89.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.89.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.89.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUB231PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.90 VFMSUBADD132PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 97 /r VFMSUBADD132PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.90.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.90.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.90.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.90.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUBADD132PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.91 VFMSUBADD132PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 97 /r VFMSUBADD132PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.91.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.91.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.91.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.91.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUBADD132PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.92 VFMSUBADD132PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 97 /r VFMSUBADD132PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.92.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.92.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.92.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.92.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUBADD132PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.93 VFMSUBADD213PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 A7 /r VFMSUBADD213PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.93.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.93.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.93.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.93.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUBADD213PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.94 VFMSUBADD213PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 A7 /r VFMSUBADD213PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.94.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.94.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.94.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.94.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUBADD213PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.95 VFMSUBADD213PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 A7 /r VFMSUBADD213PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.95.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.95.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.95.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.95.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUBADD213PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.96 VFMSUBADD231PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 B7 /r VFMSUBADD231PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.96.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.96.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.96.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.96.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUBADD231PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.97 VFMSUBADD231PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 B7 /r VFMSUBADD231PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.97.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.97.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.97.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.97.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUBADD231PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.98 VFMSUBADD231PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 B7 /r VFMSUBADD231PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.98.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.98.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.98.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.98.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMSUBADD231PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.99 VFMULCPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.F3.MAP6.W0 D6 /r VFMULCPH ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.99.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.99.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.99.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.99.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMULCPH ymm1, ymm2, ymm3/m256	E4	IOUPD	AVX10.2

12.100 VFMADD132PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 9C /r VFMADD132PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.100.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.100.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.100.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.100.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD132PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.101 VFMADD132PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 9C /r VFMADD132PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.101.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.101.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.101.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.101.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD132PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.102 VFMADD132PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 9C /r VFMADD132PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.102.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.102.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.102.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.102.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD132PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.103 VFMADD213PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 AC /r VFMADD213PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.103.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.103.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.103.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.103.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD213PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.104 VFMADD213PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 AC /r VFMADD213PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.104.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.104.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.104.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.104.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD213PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.105 VFMADD213PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 AC /r VFMADD213PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.105.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.105.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.105.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.105.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD213PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.106 VFMADD231PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 BC /r VFMADD231PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.106.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.106.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.106.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.106.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD231PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.107 VFMADD231PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 BC /r VFMADD231PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.107.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.107.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.107.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.107.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD231PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.108 VFMADD231PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 BC /r VFMADD231PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.108.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.108.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.108.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.108.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFMADD231PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.109 VFNMSUB132PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 9E /r VFNMSUB132PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.109.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.109.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.109.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.109.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFNMSUB132PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.110 VFNMSUB132PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 9E /r VFNMSUB132PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.110.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.110.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.110.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.110.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFNMSUB132PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.111 VFNMSUB132PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 9E /r VFNMSUB132PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.111.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.111.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.111.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.111.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFNMSUB132PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.112 VFNMSUB213PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 AE /r VFNMSUB213PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.112.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.112.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.112.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.112.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFNMSUB213PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.113 VFNMSUB213PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 AE /r VFNMSUB213PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.113.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.113.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.113.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.113.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFNMSUB213PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.114 VFNMSUB213PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 AE /r VFNMSUB213PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.114.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.114.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.114.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.114.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFNMSUB213PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.115 VFNMSUB231PD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 BE /r VFNMSUB231PD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.115.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.115.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.115.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.115.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFNMSUB231PD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.116 VFNMSUB231PH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 BE /r VFNMSUB231PH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.116.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.116.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.116.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.116.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFNMSUB231PH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.117 VFNMSUB231PS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 BE /r VFNMSUB231PS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.117.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(rw)	VVVV(r)	MODRM.R/M(r)	N/A

12.117.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.117.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.117.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VFNMSUB231PS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.118 VGETEXPPD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 42 /r VGETEXPPD ymm1{k1}{z}, ymm2/m256/m64bcst {sae}	A	V/V	AVX10.2

12.118.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.118.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.118.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.118.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VGETEXPPD ymm1, ymm2/m256	E2	ID	AVX10.2

12.119 VGETEXPPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 42 /r VGETEXPPH ymm1{k1}{z}, ymm2/m256/m16bcst {sae}	A	V/V	AVX10.2

12.119.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.119.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.119.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.119.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VGETEXPPH ymm1, ymm2/m256	E2	ID	AVX10.2

12.120 VGETEXPPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 42 /r VGETEXPPS ymm1{k1}{z}, ymm2/m256/m32bcst {sae}	A	V/V	AVX10.2

12.120.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.120.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.120.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.120.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VGETEXPPS ymm1, ymm2/m256	E2	ID	AVX10.2

12.121 VGETMANTPD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F3A.W1 26 /r /ib VGETMANTPD ymm1{k1}{z}, ymm2/m256/m64bcst {sae}, imm8	A	V/V	AVX10.2

12.121.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	IMM8(r)	N/A

12.121.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.121.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.121.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VGETMANTPD ymm1, ymm2/m256, imm8	E2	ID	AVX10.2

12.122 VGETMANTPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF3A.W0 26 /r /ib VGETMANTPH ymm1{k1}{z}, ymm2/m256/m16bcst {sae}, imm8	A	V/V	AVX10.2

12.122.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	IMM8(r)	N/A

12.122.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.122.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.122.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VGETMANTPH ymm1, ymm2/m256, imm8	E2	ID	AVX10.2

12.123 VGETMANTPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F3A.W0 26 /r /ib VGETMANTPS ymm1{k1}{z}, ymm2/m256/m32bcst {sae}, imm8	A	V/V	AVX10.2

12.123.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	IMM8(r)	N/A

12.123.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.123.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.123.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VGETMANTPS ymm1, ymm2/m256, imm8	E2	ID	AVX10.2

12.124 VMAXPD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W1 5F /r VMAXPD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {sae}	A	V/V	AVX10.2

12.124.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.124.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.124.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.124.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMAXPD ymm1, ymm2, ymm3/m256	E2	ID	AVX10.2

12.125 VMAXPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.MAP5.W0 5F /r VMAXPH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {sae}	A	V/V	AVX10.2

12.125.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.125.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.125.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.125.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMAXPH ymm1, ymm2, ymm3/m256	E2	ID	AVX10.2

12.126 VMAXPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF.WO 5F /r VMAXPS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {sae}	A	V/V	AVX10.2

12.126.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.126.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.126.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.126.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMAXPS ymm1, ymm2, ymm3/m256	E2	ID	AVX10.2

12.127 VMINPD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W1 5D /r VMINPD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {sae}	A	V/V	AVX10.2

12.127.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.127.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.127.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.127.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMINPD ymm1, ymm2, ymm3/m256	E2	ID	AVX10.2

12.128 VMINPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.MAP5.W0 5D /r VMINPH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {sae}	A	V/V	AVX10.2

12.128.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.128.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.128.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.128.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMINPH ymm1, ymm2, ymm3/m256	E2	ID	AVX10.2

12.129 VMINPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF.W0 5D /r VMINPS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {sae}	A	V/V	AVX10.2

12.129.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.129.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.129.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.129.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMINPS ymm1, ymm2, ymm3/m256	E2	ID	AVX10.2

12.130 VMULPD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W1 59 /r VMULPD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.130.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.130.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.130.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.130.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMULPD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.131 VMULPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.MAP5.W0 59 /r VMULPH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.131.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.131.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.131.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.131.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMULPH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.132 VMULPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF.W0 59 /r VMULPS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.132.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.132.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.132.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.132.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMULPS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.133 VRANGEPD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F3A.W1 50 /r /ib VRANGEPD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {sae}, imm8	A	V/V	AVX10.2

12.133.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	IMM8(r)

12.133.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.133.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.133.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VRANGEPD ymm1, ymm2, ymm3/m256, imm8	E2	ID	AVX10.2

12.134 VRANGEPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F3A.W0 50 /r /ib VRANGEPS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {sae}, imm8	A	V/V	AVX10.2

12.134.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	IMM8(r)

12.134.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.134.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid In abbreviated form, this includes: DE, IE

12.134.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VRANGEPS ymm1, ymm2, ymm3/m256, imm8	E2	ID	AVX10.2

12.135 VREDUCEPD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F3A.W1 56 /r /ib VREDUCEPD ymm1{k1}{z}, ymm2/m256/m64bcst {sae}, imm8	A	V/V	AVX10.2

12.135.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	IMM8(r)	N/A

12.135.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.135.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.135.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VREDUCEPD ymm1, ymm2/m256, imm8	E2	IP	AVX10.2

12.136 VREDUCEPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF3A.W0 56 /r /ib VREDUCEPH ymm1{k1}{z}, ymm2/m256/m16bcst {sae}, imm8	A	V/V	AVX10.2

12.136.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	IMM8(r)	N/A

12.136.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.136.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.136.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VREDUCEPH ymm1, ymm2/m256, imm8	E2	IP	AVX10.2

12.137 VREDUCEPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F3A.W0 56 /r /ib VREDUCEPS ymm1{k1}{z}, ymm2/m256/m32bcst {sae}, imm8	A	V/V	AVX10.2

12.137.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	IMM8(r)	N/A

12.137.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.137.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.137.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VREDUCEPS ymm1, ymm2/m256, imm8	E2	IP	AVX10.2

12.138 VRNDSCALEPD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F3A.W1 09 /r /ib VRNDSCALEPD ymm1{k1}{z}, ymm2/m256/m64bcst {sae}, imm8	A	V/V	AVX10.2

12.138.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	IMM8(r)	N/A

12.138.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.138.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.138.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VRNDSCALEPD ymm1, ymm2/m256, imm8	E2	IP	AVX10.2

12.139 VRNDSCALEPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF3A.W0 08 /r /ib VRNDSCALEPH ymm1{k1}{z}, ymm2/m256/m16bcst {sae}, imm8	A	V/V	AVX10.2

12.139.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	IMM8(r)	N/A

12.139.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.139.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision, Underflow In abbreviated form, this includes: IE, PE, UE

12.139.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VRNDSCALEPH ymm1, ymm2/m256, imm8	E2	IPU	AVX10.2

12.140 VRNDSCALEPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F3A.W0 08 /r /ib VRNDSCALEPS ymm1{k1}{z}, ymm2/m256/m32bcst {sae}, imm8	A	V/V	AVX10.2

12.140.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	IMM8(r)	N/A

12.140.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.140.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

12.140.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VRNDSCALEPS ymm1, ymm2/m256, imm8	E2	IP	AVX10.2

12.141 VSCALEFPD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W1 2C /r VSCALEFPD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.141.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.141.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.141.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.141.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VSCALEFPD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.142 VSCALEFPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.MAP6.W0 2C /r VSCALEFPH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.142.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.142.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.142.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.142.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VSCALEFPH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.143 VSCALEFPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F38.W0 2C /r VSCALEFPS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.143.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.143.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.143.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.143.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VSCALEFPS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.144 VSQRTPD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W1 51 /r VSQRTPD ymm1{k1}{z}, ymm2/m256/m64bcst {er}	A	V/V	AVX10.2

12.144.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.144.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.144.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Precision In abbreviated form, this includes: DE, IE, PE

12.144.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VSQRTPD ymm1, ymm2/m256	E2	IPD	AVX10.2

12.145 VSQRTPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.MAP5.W0 51 /r VSQRTPH ymm1{k1}{z}, ymm2/m256/m16bcst {er}	A	V/V	AVX10.2

12.145.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.145.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.145.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Precision In abbreviated form, this includes: DE, IE, PE

12.145.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VSQRTPH ymm1, ymm2/m256	E2	IPD	AVX10.2

12.146 VSQRTPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF.WO 51 /r VSQRTPS ymm1{k1}{z}, ymm2/m256/m32bcst {er}	A	V/V	AVX10.2

12.146.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

12.146.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.146.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Precision In abbreviated form, this includes: DE, IE, PE

12.146.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VSQRTPS ymm1, ymm2/m256	E2	IPD	AVX10.2

12.147 VSUBPD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.66.0F.W1 5C /r VSUBPD ymm1{k1}{z}, ymm2, ymm3/m256/m64bcst {er}	A	V/V	AVX10.2

12.147.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.147.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.147.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.147.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VSUBPD ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.148 VSUBPH

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.MAP5.W0 5C /r VSUBPH ymm1{k1}{z}, ymm2, ymm3/m256/m16bcst {er}	A	V/V	AVX10.2

12.148.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.148.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.148.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.148.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VSUBPH ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

12.149 VSUBPS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.256.NP.OF.W0 5C /r VSUBPS ymm1{k1}{z}, ymm2, ymm3/m256/m32bcst {er}	A	V/V	AVX10.2

12.149.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	VVVV(r)	MODRM.R/M(r)	N/A

12.149.2 DESCRIPTION

This instruction's description remains substantially the same as that found in Volume 2A of the Intel® 64 and IA-32 Architectures Software Developer's Manual, except being suitably modified by new Intel® AVX10 functionalities as explained in Section 3.1 of this document.

12.149.3 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Denormal, Invalid, Overflow, Precision, Underflow In abbreviated form, this includes: DE, IE, OE, PE, UE

12.149.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VSUBPS ymm1, ymm2, ymm3/m256	E2	IOUPD	AVX10.2

Chapter 13

INTEL® AVX10 SATURATING CONVERT INSTRUCTIONS

13.1 VCVT[,T]NEBF162I[,U]BS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.F2.MAP5.W0 69 /r VCVTNEBF162IBS xmm1{k1}{z}, xmm2/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.F2.MAP5.W0 69 /r VCVTNEBF162IBS ymm1{k1}{z}, ymm2/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.F2.MAP5.W0 69 /r VCVTNEBF162IBS zmm1{k1}{z}, zmm2/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.F2.MAP5.W0 6B /r VCVTNEBF162IUBS xmm1{k1}{z}, xmm2/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.F2.MAP5.W0 6B /r VCVTNEBF162IUBS ymm1{k1}{z}, ymm2/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.F2.MAP5.W0 6B /r VCVTNEBF162IUBS zmm1{k1}{z}, zmm2/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.F2.MAP5.W0 68 /r VCVTTNEBF162IBS xmm1{k1}{z}, xmm2/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.F2.MAP5.W0 68 /r VCVTTNEBF162IBS ymm1{k1}{z}, ymm2/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.F2.MAP5.W0 68 /r VCVTTNEBF162IBS zmm1{k1}{z}, zmm2/m512/m16bcst	A	V/V	AVX10.2
EVEX.128.F2.MAP5.W0 6A /r VCVTTNEBF162IUBS xmm1{k1}{z}, xmm2/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.F2.MAP5.W0 6A /r VCVTTNEBF162IUBS ymm1{k1}{z}, ymm2/m256/m16bcst	A	V/V	AVX10.2
EVEX.512.F2.MAP5.W0 6A /r VCVTTNEBF162IUBS zmm1{k1}{z}, zmm2/m512/m16bcst	A	V/V	AVX10.2

13.1.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

13.1.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

These instructions convert eight, sixteen or thirty-two packed brain-float16 values (aka bf16) in the source operand to eight, sixteen or thirty-two signed or un-signed byte integers in the destination operand

The downconverted 8-bit result is written inplace at the lower 8-bit of the corresponding 16-bit element. The upper byte is zeroed.

These instructions does not generate floating point exceptions and does not consult or update MXCSR. Denormal bfloat16 input operands are treated as zeros (DAZ)

VCVTNEBF162IBS converts brain-float16 floating point elements into signed byte integer elements. When a conversion is inexact, the rounding mode is RNE. If a converted result cannot be represented in the destination format then: If value is too big, then INT_MAX value ($2^{(w-1)}-1$, where w represents the number of bits in the destination format) is returned. If value is too small, then INT_MIN value ($2^{(w-1)}$) is returned. For NaN, (0) is returned.

VCVTNEBF162IUBS converts brain-float16 floating point elements into un-signed byte integer elements. When a conversion is inexact, the rounding mode is RNE. If a converted result cannot be represented in the destination format then: If value is too big, the UINT_MAX value (2^w-1 , where w represents the number of bits in the destination format) is returned. If value is too small, the UINT_MIN value (0) is returned. For NaN, (0) is returned.

VCVTTNEBF162IBS converts brain-float16 floating point elements into signed byte integer elements. When a conversion is inexact, a truncated (round toward zero) result is returned. If a converted result cannot be represented in the destination format then: If value is too big, the INT_MAX value ($2^{(w-1)}-1$, where w represents the number of bits in the destination format) is returned. If value is too small, the INT_MIN value ($2^{(w-1)}$) is returned. For NaN, (0) is returned.

VCVTTNEBF162IUBS converts brain-float16 floating point elements into un-signed byte integer elements. When a conversion is inexact, a truncated (round toward zero) result is returned. If a converted result cannot be represented in the destination format then: If value is too big, the UINT_MAX value (2^w-1 , where w represents the number of bits in the destination format) is returned. If value is too small, the UINT_MIN value (0) is returned. For NaN, (0) is returned.

13.1. VCVT[,T]NEBF162I[,U]BS

13.1.3 OPERATION

```

1 VCVT[,T]NEBF162[,U]BS dest {k1}, src
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         IF src is memory and (EVEX.b == 1):
7             tsrc := src.bf16[0]
8         ELSE:
9             tsrc := src.bf16[j]
10        tmp := 0
11        IF *TRUNCATED*:
12            IF *dest is signed*:
13                tmp.byte[0] := convert_bf16_to_signed_byte_truncate_saturate(tsrc)
14            ELSE:
15                tmp.byte[0] := convert_bf16_to_unsigned_byte_truncate_saturate(tsrc)
16        ELSE:
17            IF *dest is signed*:
18                tmp.byte[0] := convert_bf16_to_signed_byte_rne_saturate(tsrc)
19            ELSE:
20                tmp.byte[0] := convert_bf16_to_unsigned_byte_rne_saturate(tsrc)
21        dest.word[j] := tmp
22    ELSE IF *zeroing*:
23        dest.word[j] := 0
24    // else dest.word[j] remains unchanged
25
26 dest[MAX_VL-1:VL] := 0

```

13.1.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTNEBF162IBS xmm1, xmm2/m128	E4	N/A	AVX10.2
VCVTNEBF162IBS ymm1, ymm2/m256	E4	N/A	AVX10.2
VCVTNEBF162IBS zmm1, zmm2/m512	E4	N/A	AVX10.2
VCVTNEBF162IUBS xmm1, xmm2/m128	E4	N/A	AVX10.2
VCVTNEBF162IUBS ymm1, ymm2/m256	E4	N/A	AVX10.2
VCVTNEBF162IUBS zmm1, zmm2/m512	E4	N/A	AVX10.2

Continued on next page...

13.1. VCVT[,T]NEBF162I[,U]BS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTNEBF162IBS xmm1, xmm2/m128	E4	N/A	AVX10.2
VCVTTNEBF162IBS ymm1, ymm2/m256	E4	N/A	AVX10.2
VCVTTNEBF162IBS zmm1, zmm2/m512	E4	N/A	AVX10.2
VCVTTNEBF162IUBS xmm1, xmm2/m128	E4	N/A	AVX10.2
VCVTTNEBF162IUBS ymm1, ymm2/m256	E4	N/A	AVX10.2
VCVTTNEBF162IUBS zmm1, zmm2/m512	E4	N/A	AVX10.2

13.2 VCVTTPD2DQS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.NP.MAP5.W1 6D /r VCVTTPD2DQS xmm1{k1}{z}, xmm2/m128/m64bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP5.W1 6D /r VCVTTPD2DQS xmm1{k1}{z}, ymm2/m256/m64bcst {sae}	A	V/V	AVX10.2
EVEX.512.NP.MAP5.W1 6D /r VCVTTPD2DQS ymm1{k1}{z}, zmm2/m512/m64bcst {sae}	A	V/V	AVX10.2

13.2.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

13.2.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

VCVTTPD2DQS: Converts packed double-precision floating-point values in the source operand (the second operand) with truncation and saturation to packed doubleword integers in the destination operand (the first operand). The source operand is a ZMM/YMM/XMM register or a 512/256/128-bit memory location. The destination operand is a YMM/XMM register conditionally updated with writemask k1. When a conversion is inexact, the source value is first truncated (round toward zero) and then checked against the range of the destination type. If the truncated value is representable in the result type that value is returned, otherwise the floating-point invalid exception is raised, and if this exception is masked, the representable value closest to the truncated value is returned. If the source value is NaN, zero is returned. If some rounding mode other than truncation is desired, it can be applied using the VROUNDPD or VRNDSCALEPD instruction before the conversion. Note: EVEX.vvvv is reserved and must be 1111b, otherwise instructions will #UD.

13.2.3 OPERATION

```

1
2 VCVTTPD2DQS (EVEX encoded versions) when src2 operand is a register
3 (KL, VL) = (2, 128), (4, 256), (8, 512)
4 FOR j := 0 TO KL-1
5   i := j * 32
6   k := j * 64
7   IF k1[j] OR *no writemask*
8   THEN
9     DEST[i+31:i] := convert_DP_to_DW_SignedInteger_TruncateSaturate(SRC[k+63:k])
10  ELSE
11    IF *merging-masking* ; merging-masking
12    THEN *DEST[i+31:i] remains unchanged*
13    ELSE
14    DEST[i+31:i] := 0 ; zeroing-masking
15    FI
16  FI;
17 ENDFOR
18 DEST[MAXVL-1:VL/2] := 0
19 VCVTTPD2DQS (EVEX encoded versions) when src operand is a memory source
20 (KL, VL) = (2, 128), (4, 256), (8, 512)
21 FOR j := 0 TO KL-1
22   i := j * 32
23   k := j * 64
24   IF k1[j] OR *no writemask*
25   THEN
26     IF (EVEX.b = 1)
27     THEN
28       DEST[i+31:i] := convert_DP_to_DW_SignedInteger_TruncateSaturate(SRC[63:0])
29     ELSE
30       DEST[i+31:i] := convert_DP_to_DW_SignedInteger_TruncateSaturate(SRC[k+63:k])
31     FI;
32   ELSE
33     IF *merging-masking* ; merging-masking
34     THEN *DEST[i+31:i] remains unchanged*
35     ELSE
36     DEST[i+31:i] := 0 ; zeroing-masking
37     FI
38   FI;
39 ENDFOR
40 DEST[MAXVL-1:VL/2] := 0
41

```

13.2. VCVTTPD2DQS

13.2.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

13.2.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPD2DQS xmm1, xmm2/m128	E2	IP	AVX10.2
VCVTTPD2DQS xmm1, ymm2/m256	E2	IP	AVX10.2
VCVTTPD2DQS ymm1, zmm2/m512	E2	IP	AVX10.2

13.3 VCVTTPD2QQS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.66.MAP5.W1 6D /r VCVTTPD2QQS xmm1{k1}{z}, xmm2/m128/m64bcst	A	V/V	AVX10.2
EVEX.256.66.MAP5.W1 6D /r VCVTTPD2QQS ymm1{k1}{z}, ymm2/m256/m64bcst {sae}	A	V/V	AVX10.2
EVEX.512.66.MAP5.W1 6D /r VCVTTPD2QQS zmm1{k1}{z}, zmm2/m512/m64bcst {sae}	A	V/V	AVX10.2

13.3.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

13.3.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

VCVTTPD2QQS: Converts packed double-precision floating-point values in the source operand (second operand) to packed quadword integers in the destination operand (first operand) with truncation and saturation. The source operand is a ZMM/YMM/XMM register or a 512/256/128-bit memory location. The destination operand is a ZMM/YMM/XMM register conditionally updated with writemask k1. When a conversion is inexact, the source value is first truncated (round toward zero) and then checked against the range of the destination type. If the truncated value is representable in the result type that value is returned, otherwise the floating-point invalid exception is raised, and if this exception is masked, the representable value closest to the truncated value is returned. If the source value is NaN, zero is returned. If some rounding mode other than truncation is desired, it can be applied using the VROUNDPD or VRNDSCALEPD instruction before the conversion. Note: EVEX.vvvv is reserved and must be 1111b, otherwise instructions will #UD.

13.3.3 OPERATION

```

1  VCVTTPD2QQS (EVEX encoded version) when src operand is a register
2  (KL, VL) = (2, 128), (4, 256), (8, 512)
3  FOR j := 0 TO KL-1
4  i := j * 64
5  IF k1[j] OR *no writemask*
6      THEN DEST[i+63:i] := convert_DP_to_QW_SignedInteger_TruncateSaturate(SRC[i+63:i])
7  ELSE
8      IF *merging-masking*      ; merging-masking
9          THEN *DEST[i+63:i] remains unchanged*
10     ELSE
11         DEST[i+63:i] := 0      ; zeroing-masking
12     FI
13 FI;
14 ENDFOR
15 DEST[MAXVL-1:VL] := 0
16
17 VCVTTPD2QQS (EVEX encoded version) when src operand is a memory source
18 (KL, VL) = (2, 128), (4, 256), (8, 512)
19 FOR j := 0 TO KL-1
20 i := j * 64
21 IF k1[j] OR *no writemask*
22     THEN
23         IF (EVEX.b == 1)
24             THEN
25                 DEST[i+63:i] := convert_DP_to_QW_SignedInteger_TruncateSaturate(SRC[63:0])
26             ELSE
27                 DEST[i+63:i] := convert_DP_to_QW_SignedInteger_TruncateSaturate(SRC[i+63:i])
28             FI;
29     ELSE
30         IF *merging-masking*      ; merging-masking
31             THEN *DEST[i+63:i] remains unchanged*
32         ELSE
33             DEST[i+63:i] := 0      ; zeroing-masking
34         FI
35     FI;
36 ENDFOR
37 DEST[MAXVL-1:VL] := 0

```

13.3.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

13.3. VCVTTPD2QQS

13.3.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPD2QQS xmm1, xmm2/m128	E2	IP	AVX10.2
VCVTTPD2QQS ymm1, ymm2/m256	E2	IP	AVX10.2
VCVTTPD2QQS zmm1, zmm2/m512	E2	IP	AVX10.2

13.4 VCVTTPD2UDQS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.NP.MAP5.W1 6C /r VCVTTPD2UDQS xmm1{k1}{z}, xmm2/m128/m64bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP5.W1 6C /r VCVTTPD2UDQS xmm1{k1}{z}, ymm2/m256/m64bcst {sae}	A	V/V	AVX10.2
EVEX.512.NP.MAP5.W1 6C /r VCVTTPD2UDQS ymm1{k1}{z}, zmm2/m512/m64bcst {sae}	A	V/V	AVX10.2

13.4.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

13.4.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

VCVTTPD2UDQS: Converts packed double-precision floating-point values in the source operand (the second operand) with truncation and saturation to packed unsigned doubleword integers in the destination operand (the first operand). The source operand is a ZMM/YMM/XMM register or a 512/256/128-bit memory location. The destination operand is a YMM/XMM register conditionally updated with writemask k1. When a conversion is inexact, the source value is first truncated (round toward zero) and then checked against the range of the destination type. If the truncated value is representable in the result type that value is returned, otherwise the floating-point invalid exception is raised, and if this exception is masked, the representable value closest to the truncated value is returned. If the source value is NaN, zero is returned. If some rounding mode other than truncation is desired, it can be applied using the VROUNDPD or VRNDSCALEPD instruction before the conversion. Note: EVEX.vvvv is reserved and must be 1111b, otherwise instructions will #UD.

13.4. VCVTTPD2UDQS

13.4.3 OPERATION

```

1 VCVTTPD2UDQS (EVEX encoded versions) when src2 operand is a register
2 (KL, VL) = (2, 128), (4, 256), (8, 512)
3 FOR j := 0 TO KL-1
4   i := j * 32
5   k := j * 64
6   IF k1[j] OR *no writemask*
7     THEN
8       DEST[i+31:i] := convert_DP_to_DW_UnSignedInteger_TruncateSaturate(SRC[k+63:k])
9     ELSE
10      IF *merging-masking*      ; merging-masking
11      THEN *DEST[i+31:i] remains unchanged*
12      ELSE
13      DEST[i+31:i] := 0      ; zeroing-masking
14      FI
15    FI;
16  ENDFOR
17  DEST[MAXVL-1:VL/2] := 0
18
19 VCVTTPD2UDQS (EVEX encoded versions) when src operand is a memory source
20 (KL, VL) = (2, 128), (4, 256), (8, 512)
21 FOR j := 0 TO KL-1
22   i := j * 32
23   k := j * 64
24   IF k1[j] OR *no writemask*
25     THEN
26       IF (EVEX.b = 1)
27         THEN
28           DEST[i+31:i] := convert_DP_to_DW_UnSignedInteger_TruncateSaturate(SRC[63:0])
29         ELSE
30           DEST[i+31:i] := convert_DP_to_DW_UnSignedInteger_TruncateSaturate(SRC[k+63:k])
31         FI;
32     ELSE
33       IF *merging-masking*      ; merging-masking
34       THEN *DEST[i+31:i] remains unchanged*
35       ELSE
36       DEST[i+31:i] := 0      ; zeroing-masking
37       FI
38     FI;
39  ENDFOR
40  DEST[MAXVL-1:VL/2] := 0

```

13.4. VCVTTPD2UDQS

13.4.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

13.4.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPD2UDQS xmm1, xmm2/m128	E2	IP	AVX10.2
VCVTTPD2UDQS xmm1, ymm2/m256	E2	IP	AVX10.2
VCVTTPD2UDQS ymm1, zmm2/m512	E2	IP	AVX10.2

13.5 VCVTTPD2UQQS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.66.MAP5.W1 6C /r VCVTTPD2UQQS xmm1{k1}{z}, xmm2/m128/m64bcst	A	V/V	AVX10.2
EVEX.256.66.MAP5.W1 6C /r VCVTTPD2UQQS ymm1{k1}{z}, ymm2/m256/m64bcst {sae}	A	V/V	AVX10.2
EVEX.512.66.MAP5.W1 6C /r VCVTTPD2UQQS zmm1{k1}{z}, zmm2/m512/m64bcst {sae}	A	V/V	AVX10.2

13.5.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

13.5.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

VCVTTPD2UQQS: Converts packed double-precision floating-point values in the source operand (second operand) to packed unsigned quadword integers in the destination operand (first operand) with truncation and saturation. The source operand is a ZMM/YMM/XMM register or a 512/256/128-bit memory location. The destination operand is a ZMM/YMM/XMM register conditionally updated with writemask k1. When a conversion is inexact, the source value is first truncated (round toward zero) and then checked against the range of the destination type. If the truncated value is representable in the result type that value is returned, otherwise the floating-point invalid exception is raised, and if this exception is masked, the representable value closest to the truncated value is returned. If the source value is NaN, zero is returned. If some rounding mode other than truncation is desired, it can be applied using the VROUNDPD or VRNDSCALEPD instruction before the conversion. Note: EVEX.vvvv is reserved and must be 1111b, otherwise instructions will #UD.

13.5.3 OPERATION

```

1  VCVTTPD2UQQS (EVEX encoded version) when src operand is a register
2  (KL, VL) = (2, 128), (4, 256), (8, 512)
3  FOR j := 0 TO KL-1
4  i := j * 64
5  IF k1[j] OR *no writemask*
6      THEN DEST[i+63:i] := convert_DP_to_QW_UnSignedInteger_TruncateSaturate(SRC[i+63:i])
7  ELSE
8      IF *merging-masking*      ; merging-masking
9          THEN *DEST[i+63:i] remains unchanged*
10     ELSE
11         DEST[i+63:i] := 0      ; zeroing-masking
12     FI
13 FI;
14 ENDFOR
15 DEST[MAXVL-1:VL] := 0
16
17 VCVTTPD2UQQS (EVEX encoded version) when src operand is a memory source
18 (KL, VL) = (2, 128), (4, 256), (8, 512)
19 FOR j := 0 TO KL-1
20 i := j * 64
21 IF k1[j] OR *no writemask*
22     THEN
23         IF (EVEX.b == 1)
24             THEN
25                 DEST[i+63:i] :=
26                 ELSE    convert_DP_to_QW_UnSignedInteger_TruncateSaturate(SRC[63:0])
27                 DEST[i+63:i] := convert_DP_to_QW_UnSignedInteger_TruncateSaturate(SRC[i+63:i])
28             FI;
29     ELSE
30         IF *merging-masking*      ; merging-masking
31             THEN *DEST[i+63:i] remains unchanged*
32         ELSE
33             DEST[i+63:i] := 0      ; zeroing-masking
34         FI
35     FI;
36 ENDFOR
37 DEST[MAXVL-1:VL] := 0

```

13.5.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

13.5. VCVTTPD2UQQS

13.5.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPD2UQQS xmm1, xmm2/m128	E2	IP	AVX10.2
VCVTTPD2UQQS ymm1, ymm2/m256	E2	IP	AVX10.2
VCVTTPD2UQQS zmm1, zmm2/m512	E2	IP	AVX10.2

13.6 VCVT_[,T]PH2I_[,U]BS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.NP.MAP5.W0 69 /r VCVTPH2IBS xmm1{k1}{z}, xmm2/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP5.W0 69 /r VCVTPH2IBS ymm1{k1}{z}, ymm2/m256/m16bcst {er}	A	V/V	AVX10.2
EVEX.512.NP.MAP5.W0 69 /r VCVTPH2IBS zmm1{k1}{z}, zmm2/m512/m16bcst {er}	A	V/V	AVX10.2
EVEX.128.NP.MAP5.W0 6B /r VCVTPH2IUBS xmm1{k1}{z}, xmm2/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP5.W0 6B /r VCVTPH2IUBS ymm1{k1}{z}, ymm2/m256/m16bcst {er}	A	V/V	AVX10.2
EVEX.512.NP.MAP5.W0 6B /r VCVTPH2IUBS zmm1{k1}{z}, zmm2/m512/m16bcst {er}	A	V/V	AVX10.2
EVEX.128.NP.MAP5.W0 68 /r VCVTTPH2IBS xmm1{k1}{z}, xmm2/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP5.W0 68 /r VCVTTPH2IBS ymm1{k1}{z}, ymm2/m256/m16bcst {sae}	A	V/V	AVX10.2
EVEX.512.NP.MAP5.W0 68 /r VCVTTPH2IBS zmm1{k1}{z}, zmm2/m512/m16bcst {sae}	A	V/V	AVX10.2
EVEX.128.NP.MAP5.W0 6A /r VCVTTPH2IUBS xmm1{k1}{z}, xmm2/m128/m16bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP5.W0 6A /r VCVTTPH2IUBS ymm1{k1}{z}, ymm2/m256/m16bcst {sae}	A	V/V	AVX10.2
EVEX.512.NP.MAP5.W0 6A /r VCVTTPH2IUBS zmm1{k1}{z}, zmm2/m512/m16bcst {sae}	A	V/V	AVX10.2

13.6.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

13.6.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

These instructions convert eight, sixteen or thirty-two packed half-precision floating-point values (aka fp16) in the source operand to eight, sixteen or thirty-two signed or un-signed byte integers in the destination operand

The downconverted 8-bit result is written inplace at the lower 8-bit of the corresponding 16-bit element. The upper byte is zeroed.

VCVTPH2IBS converts half-precision floating point elements into signed byte integer elements. When a conversion is inexact, floating-point precision exception is raised and the value returned is rounded according to the rounding control bits in the MXCSR register or the embedded rounding control bits. If a converted result cannot be represented in the destination format, the floating-point invalid exception is raised, and if this exception is masked then: If value is too big, the INT_MAX value ($2^{(w-1)}-1$, where w represents the number of bits in the destination format) is returned. If value is too small, the INT_MIN value ($2^{-(w-1)}$) is returned. For NaN, (0) is returned.

VCVTPH2IUBS converts half-precision floating point elements into un-signed byte integer elements. When a conversion is inexact, floating-point precision exception is raised and the value returned is rounded according to the rounding control bits in the MXCSR register or the embedded rounding control bits. If a converted result cannot be represented in the destination format, the floating-point invalid exception is raised, and if this exception is masked then: If value is too big, the UINT_MAX value (2^w-1 , where w represents the number of bits in the destination format) is returned. If value is too small, the UINT_MIN value (0) is returned. For NaN, (0) is returned.

VCVTTPH2IBS converts half-precision floating point elements into signed byte integer elements. When a conversion is inexact, floating-point precision exception is raised and a truncated (round toward zero) result is returned. If a converted result cannot be represented in the destination format, the floating-point invalid exception is raised, and if this exception is masked then: If value is too big, the INT_MAX value ($2^{(w-1)}-1$, where w represents the number of bits in the destination format) is returned. If value is too small, the INT_MIN value ($2^{-(w-1)}$) is returned. For NaN, (0) is returned.

VCVTTPH2IUBS converts half-precision floating point elements into un-signed byte integer elements. When a conversion is inexact, floating-point precision exception is raised and a truncated (round toward zero) result is returned. If a converted result cannot be represented in the destination format, the floating-point invalid exception is raised, and if this exception is masked then: If value is too big, the UINT_MAX value (2^w-1 , where w represents the number of bits in the destination format) is returned. If value is too small, the UINT_MIN value (0) is returned. For NaN, (0) is returned.

13.6. VCVT[,T]PH2I[,U]BS

13.6.3 OPERATION

```

1 VCVT[,T]PH2I[,U]BS dest {k1}, src
2 (KL, VL) = (8, 128), (16, 256), (32, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         IF src is memory and (EVEX.b == 1):
7             tsrc := src.fp16[0]
8         ELSE:
9             tsrc := src.fp16[j]
10        tmp := 0
11        IF *TRUNCATED*:
12            IF *dest is signed*:
13                tmp.byte[0] := convert_fp16_to_signed_byte_truncate_saturate(tsrc)
14            ELSE:
15                tmp.byte[0] := convert_fp16_to_unsigned_byte_truncate_saturate(tsrc)
16        ELSE:
17            IF *dest is signed*:
18                tmp.byte[0] := convert_fp16_to_signed_byte_saturate(tsrc)
19            ELSE:
20                tmp.byte[0] := convert_fp16_to_unsigned_byte_saturate(tsrc)
21        dest.word[j] := tmp
22    ELSE IF *zeroing*:
23        dest.word[j] := 0
24    // else dest.word[j] remains unchanged
25
26 dest[MAX_VL-1:VL] := 0

```

13.6.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

13.6.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPH2IBS xmm1, xmm2/m128	E2	IP	AVX10.2
VCVTPH2IBS ymm1, ymm2/m256	E2	IP	AVX10.2

Continued on next page...

13.6. VCVT[*T*]PH2I[*U*]BS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPH2IBS zmm1, zmm2/m512	E2	IP	AVX10.2
VCVTPH2IUBS xmm1, xmm2/m128	E2	IP	AVX10.2
VCVTPH2IUBS ymm1, ymm2/m256	E2	IP	AVX10.2
VCVTPH2IUBS zmm1, zmm2/m512	E2	IP	AVX10.2
VCVTTPH2IBS xmm1, xmm2/m128	E2	IP	AVX10.2
VCVTTPH2IBS ymm1, ymm2/m256	E2	IP	AVX10.2
VCVTTPH2IBS zmm1, zmm2/m512	E2	IP	AVX10.2
VCVTTPH2IUBS xmm1, xmm2/m128	E2	IP	AVX10.2
VCVTTPH2IUBS ymm1, ymm2/m256	E2	IP	AVX10.2
VCVTTPH2IUBS zmm1, zmm2/m512	E2	IP	AVX10.2

13.7 VCVTTPS2DQS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.NP.MAP5.W0 6D /r VCVTTPS2DQS xmm1{k1}{z}, xmm2/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP5.W0 6D /r VCVTTPS2DQS ymm1{k1}{z}, ymm2/m256/m32bcst {sae}	A	V/V	AVX10.2
EVEX.512.NP.MAP5.W0 6D /r VCVTTPS2DQS zmm1{k1}{z}, zmm2/m512/m32bcst {sae}	A	V/V	AVX10.2

13.7.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

13.7.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

VCVTTPS2DQS: Converts packed single-precision floating-point values in the source operand to doubleword integers in the destination operand with truncation and saturation. The source operand is a ZMM/YMM/XMM register or a 512/256/128-bit memory location. The destination operand is a ZMM/YMM/XMM register conditionally updated with writemask k1. When a conversion is inexact, the source value is first truncated (round toward zero) and then checked against the range of the destination type. If the truncated value is representable in the result type that value is returned, otherwise the floating-point invalid exception is raised, and if this exception is masked, the representable value closest to the truncated value is returned. If the source value is NaN, zero is returned. If some rounding mode other than truncation is desired, it can be applied using the VROUNDPS or VRNDSCALEPS instruction before the conversion. Note: EVEX.vvvv is reserved and must be 1111b otherwise instructions will #UD

13.7.3 OPERATION

```

1  VCVTTPS2DQS (EVEX encoded versions) when src2 operand is a register
2  (KL, VL) = (2, 128), (4, 256), (8, 512)
3  FOR j := 0 TO KL-1
4      i := j * 32
5      IF k1[j] OR *no writemask*
6          THEN
7              DEST[i+31:i] := convert_SP_to_DW_SignedInteger_TruncateSaturate(SRC[k+31:k])
8          ELSE
9              IF *merging-masking*      ; merging-masking
10             THEN *DEST[i+31:i] remains unchanged*
11             ELSE
12                 DEST[i+31:i] := 0      ; zeroing-masking
13             FI
14         FI;
15     ENDFOR
16     DEST[MAXVL-1:VL/2] := 0
17
18     VCVTTPS2DQS (EVEX encoded versions) when src operand is a memory source
19     (KL, VL) = (2, 128), (4, 256), (8, 512)
20     FOR j := 0 TO KL-1
21         i := j * 32
22         IF k1[j] OR *no writemask*
23             THEN
24                 IF (EVEX.b = 1)
25                     THEN
26                         DEST[i+31:i] := convert_SP_to_DW_SignedInteger_TruncateSaturate(SRC[31:0])
27                     ELSE
28                         DEST[i+31:i] := convert_SP_to_DW_SignedInteger_TruncateSaturate(SRC[k+31:k])
29                     FI;
30             ELSE
31                 IF *merging-masking*      ; merging-masking
32                 THEN *DEST[i+31:i] remains unchanged*
33                 ELSE
34                     DEST[i+31:i] := 0      ; zeroing-masking
35                 FI
36             FI;
37     ENDFOR
38     DEST[MAXVL-1:VL/2] := 0

```

13.7.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

13.7. VCVTTPS2DQS

13.7.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPS2DQS xmm1, xmm2/m128	E2	IP	AVX10.2
VCVTTPS2DQS ymm1, ymm2/m256	E2	IP	AVX10.2
VCVTTPS2DQS zmm1, zmm2/m512	E2	IP	AVX10.2

13.8 VCVT_[,T]PS2I_[,U]BS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.66.MAP5.W0 69 /r VCVTPS2IBS xmm1{k1}{z}, xmm2/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.66.MAP5.W0 69 /r VCVTPS2IBS ymm1{k1}{z}, ymm2/m256/m32bcst {er}	A	V/V	AVX10.2
EVEX.512.66.MAP5.W0 69 /r VCVTPS2IBS zmm1{k1}{z}, zmm2/m512/m32bcst {er}	A	V/V	AVX10.2
EVEX.128.66.MAP5.W0 6B /r VCVTPS2IUBS xmm1{k1}{z}, xmm2/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.66.MAP5.W0 6B /r VCVTPS2IUBS ymm1{k1}{z}, ymm2/m256/m32bcst {er}	A	V/V	AVX10.2
EVEX.512.66.MAP5.W0 6B /r VCVTPS2IUBS zmm1{k1}{z}, zmm2/m512/m32bcst {er}	A	V/V	AVX10.2
EVEX.128.66.MAP5.W0 68 /r VCVTTPS2IBS xmm1{k1}{z}, xmm2/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.66.MAP5.W0 68 /r VCVTTPS2IBS ymm1{k1}{z}, ymm2/m256/m32bcst {sae}	A	V/V	AVX10.2
EVEX.512.66.MAP5.W0 68 /r VCVTTPS2IBS zmm1{k1}{z}, zmm2/m512/m32bcst {sae}	A	V/V	AVX10.2
EVEX.128.66.MAP5.W0 6A /r VCVTTPS2IUBS xmm1{k1}{z}, xmm2/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.66.MAP5.W0 6A /r VCVTTPS2IUBS ymm1{k1}{z}, ymm2/m256/m32bcst {sae}	A	V/V	AVX10.2
EVEX.512.66.MAP5.W0 6A /r VCVTTPS2IUBS zmm1{k1}{z}, zmm2/m512/m32bcst {sae}	A	V/V	AVX10.2

13.8.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

13.8.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

These instructions convert four, eight or sixteen packed single-precision floating-point values in the source operand to four, eight or sixteen signed or unsigned byte integers in the destination operand.

The downconverted 8-bit result is written in-place at the lower 8-bit of the corresponding 32-bit element. The upper 3 bytes are zeroed.

VCVTPS2IBS converts single-precision floating point elements into signed byte integer elements. When a conversion is inexact, floating-point precision exception is raised and the value returned is rounded according to the rounding control bits in the MXCSR register or the embedded rounding control bits. If a converted result cannot be represented in the destination format, the floating-point invalid exception is raised, and if this exception is masked then: If value is too big, the INT_MAX value ($2^{(w-1)}-1$, where w represents the number of bits in the destination format) is returned. If value is too small, the INT_MIN value ($2^{(w-1)}$) is returned. For NaN, (0) is returned.

VCVTPS2IUBS converts single-precision floating point elements into un-signed byte integer elements. When a conversion is inexact, floating-point precision exception is raised and the value returned is rounded according to the rounding control bits in the MXCSR register or the embedded rounding control bits. If a converted result cannot be represented in the destination format, the floating-point invalid exception is raised, and if this exception is masked then: If value is too big, the UINT_MAX value (2^w-1 , where w represents the number of bits in the destination format) is returned. If value is too small, the UINT_MIN value (0) is returned. For NaN, (0) is returned.

VCVTTPS2IBS converts single-precision floating point elements into signed byte integer elements. When a conversion is inexact, floating-point precision exception is raised and a truncated (round toward zero) result is returned. If a converted result cannot be represented in the destination format, the floating-point invalid exception is raised, and if this exception is masked then: If value is too big, the INT_MAX value ($2^{(w-1)}-1$, where w represents the number of bits in the destination format) is returned. If value is too small, the INT_MIN value ($2^{(w-1)}$) is returned. For NaN, (0) is returned.

VCVTTPS2IUBS converts single-precision floating point elements into un-signed byte integer elements. When a conversion is inexact, floating-point precision exception is raised and a truncated (round toward zero) result is returned. If a converted result cannot be represented in the destination format, the floating-point invalid exception is raised, and if this exception is masked then: If value is too big, the UINT_MAX value (2^w-1 , where w represents the number of bits in the destination format) is returned. If value is too small, the UINT_MIN value (0) is returned. For NaN, (0) is returned.

13.8. VCVT[,T]PS2I[,U]BS

13.8.3 OPERATION

```

1 VCVT[,T]PS2I[,U]BS dest {k1}, src
2 (KL, VL) = (4, 128), (8, 256), (16, 512)
3
4 FOR j := 0 TO KL-1:
5     IF k1[j] OR *no writemask*:
6         IF src is memory and (EVEX.b == 1):
7             tsrc := src.fp32[0]
8         ELSE:
9             tsrc := src.fp32[j]
10        tmp := 0
11        IF *TRUNCATED*:
12            IF *dest is signed*:
13                tmp.byte[0] := convert_fp32_to_signed_byte_truncate_saturate(tsrc)
14            ELSE:
15                tmp.byte[0] := convert_fp32_to_unsigned_byte_truncate_saturate(tsrc)
16        ELSE:
17            IF *dest is signed*:
18                tmp.byte[0] := convert_fp32_to_signed_byte_saturate(tsrc)
19            ELSE:
20                tmp.byte[0] := convert_fp32_to_unsigned_byte_saturate(tsrc)
21        dest.dword[j] := tmp
22    ELSE IF *zeroing*:
23        dest.dword[j] := 0
24    // else dest.dword[j] remains unchanged
25
26 dest[MAX_VL-1:VL] := 0

```

13.8.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

13.8.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPS2IBS xmm1, xmm2/m128	E2	IP	AVX10.2
VCVTPS2IBS ymm1, ymm2/m256	E2	IP	AVX10.2

Continued on next page...

13.8. VCVT[,T]PS2I[,U]BS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTPS2IBS zmm1, zmm2/m512	E2	IP	AVX10.2
VCVTPS2IUBS xmm1, xmm2/m128	E2	IP	AVX10.2
VCVTPS2IUBS ymm1, ymm2/m256	E2	IP	AVX10.2
VCVTPS2IUBS zmm1, zmm2/m512	E2	IP	AVX10.2
VCVTTPS2IBS xmm1, xmm2/m128	E2	IP	AVX10.2
VCVTTPS2IBS ymm1, ymm2/m256	E2	IP	AVX10.2
VCVTTPS2IBS zmm1, zmm2/m512	E2	IP	AVX10.2
VCVTTPS2IUBS xmm1, xmm2/m128	E2	IP	AVX10.2
VCVTTPS2IUBS ymm1, ymm2/m256	E2	IP	AVX10.2
VCVTTPS2IUBS zmm1, zmm2/m512	E2	IP	AVX10.2

13.9 VCVTTPS2QQS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.66.MAP5.W0 6D /r VCVTTPS2QQS xmm1{k1}{z}, xmm2/m64/m32bcst	A	V/V	AVX10.2
EVEX.256.66.MAP5.W0 6D /r VCVTTPS2QQS ymm1{k1}{z}, xmm2/m128/m32bcst {sae}	A	V/V	AVX10.2
EVEX.512.66.MAP5.W0 6D /r VCVTTPS2QQS zmm1{k1}{z}, ymm2/m256/m32bcst {sae}	A	V/V	AVX10.2

13.9.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	HALF	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

13.9.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

VCVTTPS2QQS: Converts packed single-precision floating-point values in the source operand to quadword integers in the destination operand with truncation and saturation. The source operand is a YMM/XMM register or a 256/128/64-bit memory location. The destination operand is a ZMM/YMM/XMM register conditionally updated with writemask k1. When a conversion is inexact, the source value is first truncated (round toward zero) and then checked against the range of the destination type. If the truncated value is representable in the result type that value is returned, otherwise the floating-point invalid exception is raised, and if this exception is masked, the representable value closest to the truncated value is returned. If the source value is NaN, zero is returned. If some rounding mode other than truncation is desired, it can be applied using the VROUNDPS or VRNDSCALEPS instruction before the conversion. Note: EVEX.vvvv is reserved and must be 1111b otherwise instructions will #UD.

13.9.3 OPERATION

```

1  VCVTTPS2QQS (EVEX encoded versions) when src operand is a register
2  (KL, VL) = (2, 128), (4, 256), (8, 512)
3  FOR j := 0 TO KL-1
4      i := j * 64
5      k := j * 32
6      IF k1[j] OR *no writemask*
7      THEN DEST[i+63:i] := convert_SP_to_QW_SignedInteger_TruncateSaturate(SRC[k+31:k])
8      ELSE
9          IF *merging-masking*      ; merging-masking
10         THEN *DEST[i+63:i] remains unchanged*
11         ELSE
12             DEST[i+63:i] := 0      ; zeroing-masking
13         FI
14     FI;
15 ENDFOR
16 DEST[MAXVL-1:VL] := 0
17
18 VCVTTPS2QQS (EVEX encoded versions) when src operand is a memory source
19 (KL, VL) = (2, 128), (4, 256), (8, 512)
20 FOR j := 0 TO KL-1
21     i := j * 64
22     k := j * 32
23     IF k1[j] OR *no writemask*
24     THEN
25         IF (EVEX.b == 1)
26         THEN DEST[i+63:i] := convert_SP_to_QW_SignedInteger_TruncateSaturate(SRC[31:0])
27         ELSE
28             DEST[i+63:i] := convert_SP_to_QW_SignedInteger_TruncateSaturate(SRC[k+31:k])
29         FI;
30     ELSE
31         IF *merging-masking*      ; merging-masking
32         THEN *DEST[i+63:i] remains unchanged*
33         ELSE
34             DEST[i+63:i] := 0      ; zeroing-masking
35         FI
36     FI;
37 ENDFOR
38 DEST[MAXVL-1:VL] := 0

```

13.9.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

13.9.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPS2QQS xmm1, xmm2/m64	E2	IP	AVX10.2
VCVTTPS2QQS ymm1, xmm2/m128	E2	IP	AVX10.2
VCVTTPS2QQS zmm1, ymm2/m256	E2	IP	AVX10.2

13.10 VCVTTTPS2UDQS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.NP.MAP5.W0 6C /r VCVTTTPS2UDQS xmm1{k1}{z}, xmm2/m128/m32bcst	A	V/V	AVX10.2
EVEX.256.NP.MAP5.W0 6C /r VCVTTTPS2UDQS ymm1{k1}{z}, ymm2/m256/m32bcst {sae}	A	V/V	AVX10.2
EVEX.512.NP.MAP5.W0 6C /r VCVTTTPS2UDQS zmm1{k1}{z}, zmm2/m512/m32bcst {sae}	A	V/V	AVX10.2

13.10.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	FULL	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

13.10.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

VCVTTTPS2UDQS: Converts packed single-precision floating-point values in the source operand to unsigned doubleword integers in the destination operand with truncation and saturation. The source operand is a ZMM/YMM/XMM register or a 512/256/128-bit memory location. The destination operand is a ZMM/YMM/XMM register conditionally updated with writemask k1. When a conversion is inexact, the source value is first truncated (round toward zero) and then checked against the range of the destination type. If the truncated value is representable in the result type that value is returned, otherwise the floating-point invalid exception is raised, and if this exception is masked, the representable value closest to the truncated value is returned. If the source value is NaN, zero is returned. If some rounding mode other than truncation is desired, it can be applied using the VROUNDPS or VRNDSCALEPS instruction before the conversion. Note: EVEX.vvvv is reserved and must be 1111b otherwise instructions will #UD

13.10. VCVTTTPS2UDQS

13.10.3 OPERATION

```

1  VCVTTTPS2UDQS (EVEX encoded versions) when src2 operand is a register
2  (KL, VL) = (2, 128), (4, 256), (8, 512)
3  S := 31
4  FOR j := 0 TO KL-1
5      i := j * 32
6      IF k1[j] OR *no writemask*
7          THEN
8              DST[i+S:i]:=Convert_Single_Precision_Floating_Point_To_Uinteger_Saturate(SRC[k+S:k])
9          ELSE
10             IF *merging-masking*      ; merging-masking
11                 THEN *DST[i+S:i] remains unchanged*
12             ELSE
13                 DST[i+S:i] := 0      ; zeroing-masking
14             FI
15         FI;
16     ENDFOR
17     DST[MAXVL-1:VL/2] := 0
18
19     VCVTTTPS2UDQS (EVEX encoded versions) when src operand is a memory source
20     (KL, VL) = (2, 128), (4, 256), (8, 512)
21     S := 31
22     FOR j := 0 TO KL-1
23         i := j * 32
24         IF k1[j] OR *no writemask*:
25             IF (EVEX.b = 1):
26                 DST[i+S:i]:=Convert_Single_Precision_Floating_Point_To_Uinteger_Saturate(SRC[S:0])
27             ELSE:
28                 DST[i+S:i]:=Convert_Single_Precision_Floating_Point_To_Uinteger_Saturate(SRC[k+S:k])
29             FI;
30         ELSE
31             IF *merging-masking*      ; merging-masking
32                 THEN *DST[i+S:i] remains unchanged*
33             ELSE
34                 DST[i+S:i] := 0      ; zeroing-masking
35             FI
36         FI;
37     ENDFOR
38     DST[MAXVL-1:VL/2] := 0

```

13.10.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

13.10. VCVTTPS2UDQS

13.10.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPS2UDQS xmm1, xmm2/m128	E2	IP	AVX10.2
VCVTTPS2UDQS ymm1, ymm2/m256	E2	IP	AVX10.2
VCVTTPS2UDQS zmm1, zmm2/m512	E2	IP	AVX10.2

13.11 VCVTTPS2UQQS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.66.MAP5.W0 6C /r VCVTTPS2UQQS xmm1{k1}{z}, xmm2/m64/m32bcst	A	V/V	AVX10.2
EVEX.256.66.MAP5.W0 6C /r VCVTTPS2UQQS ymm1{k1}{z}, xmm2/m128/m32bcst {sae}	A	V/V	AVX10.2
EVEX.512.66.MAP5.W0 6C /r VCVTTPS2UQQS zmm1{k1}{z}, ymm2/m256/m32bcst {sae}	A	V/V	AVX10.2

13.11.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	HALF	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

13.11.2 DESCRIPTION

Note:

For instructions with a CPUID feature flag specifying Intel® AVX10, the programmer must check the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and as such will determine the set of instructions available to the programmer listed in the above opcode table.

VCVTTPS2UQQS: Converts packed single-precision floating-point values in the source operand to unsigned quadword integers in the destination operand with truncation and saturation. The source operand is a YMM/XMM register or a 256/128/64-bit memory location. The destination operand is a ZMM/YMM/XMM register conditionally updated with writemask k1. When a conversion is inexact, the source value is first truncated (round toward zero) and then checked against the range of the destination type. If the truncated value is representable in the result type that value is returned, otherwise the floating-point invalid exception is raised, and if this exception is masked, the representable value closest to the truncated value is returned. If the source value is NaN, zero is returned. If some rounding mode other than truncation is desired, it can be applied using the VROUNDPS or VRNDSCALEPS instruction before the conversion. Note: EVEX.vvvv is reserved and must be 1111b otherwise instructions will #UD.

13.11. VCVTTTPS2UQQS

13.11.3 OPERATION

```

1  VCVTTTPS2UQQS (EVEX encoded versions) when src operand is a register
2  (KL, VL) = (2, 128), (4, 256), (8, 512)
3  FOR j := 0 TO KL-1
4      i := j * 64
5      k := j * 32
6      IF k1[j] OR *no writemask*
7      THEN DEST[i+63:i] := convert_SP_to_QW_UnSignedInteger_TruncateSaturate(SRC[k+31:k])
8      ELSE
9          IF *merging-masking*      ; merging-masking
10         THEN *DEST[i+63:i] remains unchanged*
11         ELSE
12             DEST[i+63:i] := 0      ; zeroing-masking
13         FI
14     FI;
15 ENDFOR
16 DEST[MAXVL-1:VL] := 0
17
18 VCVTTTPS2UQQS (EVEX encoded versions) when src operand is a memory source
19 (KL, VL) = (2, 128), (4, 256), (8, 512)
20 FOR j := 0 TO KL-1
21     i := j * 64
22     k := j * 32
23     IF k1[j] OR *no writemask*
24     THEN
25         IF (EVEX.b == 1)
26         THEN
27             DEST[i+63:i] := convert_SP_to_QW_UnSignedInteger_TruncateSaturate(SRC[31:0])
28         ELSE
29             DEST[i+63:i] := convert_SP_to_QW_UnSignedInteger_TruncateSaturate(SRC[k+31:k])
30         FI;
31     ELSE
32         IF *merging-masking*      ; merging-masking
33         THEN *DEST[i+63:i] remains unchanged*
34         ELSE
35             DEST[i+63:i] := 0      ; zeroing-masking
36         FI
37     FI;
38 ENDFOR
39 DEST[MAXVL-1:VL] := 0

```

13.11. VCVTTPS2UQQS

13.11.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

13.11.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTPS2UQQS xmm1, xmm2/m64	E2	IP	AVX10.2
VCVTTPS2UQQS ymm1, xmm2/m128	E2	IP	AVX10.2
VCVTTPS2UQQS zmm1, ymm2/m256	E2	IP	AVX10.2

13.12 VCVTTSD2SIS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.LLIG.F2.MAP5.W0 6D /r VCVTTSD2SIS r32, xmm1/m64 {sae}	A	V/V ¹	AVX10.2
EVEX.LLIG.F2.MAP5.W1 6D /r VCVTTSD2SIS r64, xmm1/m64 {sae}	A	V/N.E.	AVX10.2
Notes:			
1: Outside of 64b mode, the EVEX.W field is ignored. The instruction behaves as if W=0 were used.			

13.12.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	SCALAR	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

13.12.2 DESCRIPTION

Note:

For SCALAR instructions with a CPUID feature flag specifying Intel® AVX10, the programmer can skip the check for the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and usually determines the set of instructions available to the programmer listed in the above opcode table. However, scalar instructions are architected without the notion of vector length, so this step does not strictly apply.

VCVTTSD2SIS: Converts a double-precision floating-point value in the source operand (the second operand) to a doubleword integer (or quadword integer if operand size is 64 bits) in the destination operand (the first operand) with truncation and saturation. The source operand can be an XMM register or a 64-bit memory location. The destination operand is a general-purpose register. When the source operand is an XMM register, the double-precision floating-point value is contained in the low quadword of the register. When a conversion is inexact, the source value is first truncated (round toward zero) and then checked against the range of the destination type. If the truncated value is representable in the result type that value is returned, otherwise the floating-point invalid exception is raised, and if this exception is masked, the representable value closest to the truncated value is returned. If the source value is NaN, zero is returned. If some rounding mode other than truncation is desired, it can be applied using the VROUNDSD or VRNDSCALESD instruction before the conversion. EVEX.W1 version: promotes the instruction to produce 64-bit data in 64-bit mode.

13.12. VCVTTSD2SIS

13.12.3 OPERATION

```

1
2 VCVTTSD2SIS (EVEX encoded version)
3 IF 64-bit Mode and OperandSize = 64
4 THEN
5     DEST[63:0] := convert_DP_to_QW_SignedInteger_TruncateSaturate(SRC[63:0]);
6 ELSE
7     DEST[31:0] := convert_DP_to_DW_SignedInteger_TruncateSaturate(SRC[63:0]);
8 FI;
9
10 SIMD Floating-Point Exceptions
11 -Invalid, Precision
12
13 Other Exceptions:
14 -EVEX-encoded instructions, see Table 2-48, "Type E3NF Class Exception
15 Conditions".
16
17 NOTES:
18
19 1. For this specific instruction, EVEX.W in non-64 bit is ignored; the
20 instructions behaves as if the W0 version is used.
21

```

13.12.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

13.12.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTSD2SIS r32, xmm1/m64	E3NF	IP	AVX10.2
VCVTTSD2SIS r64, xmm1/m64	E3NF	IP	AVX10.2

13.13 VCVTTSD2USIS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.LLIG.F2.MAP5.W0 6C /r VCVTTSD2USIS r32, xmm1/m64 {sae}	A	V/V ¹	AVX10.2
EVEX.LLIG.F2.MAP5.W1 6C /r VCVTTSD2USIS r64, xmm1/m64 {sae}	A	V/N.E.	AVX10.2
Notes:			
1: Outside of 64b mode, the EVEX.W field is ignored. The instruction behaves as if W=0 were used.			

13.13.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	SCALAR	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

13.13.2 DESCRIPTION

Note:

For SCALAR instructions with a CPUID feature flag specifying Intel® AVX10, the programmer can skip the check for the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and usually determines the set of instructions available to the programmer listed in the above opcode table. However, scalar instructions are architected without the notion of vector length, so this step does not strictly apply.

VCVTTSD2USIS: Converts a double-precision floating-point value in the source operand (the second operand) to an unsigned doubleword integer (or an unsigned quadword integer if operand size is 64 bits) in the destination operand (the first operand) with truncation and saturation. The source operand can be an XMM register or a 64-bit memory location. The destination operand is a general-purpose register. When the source operand is an XMM register, the double-precision floating-point value is contained in the low quadword of the register. When a conversion is inexact, the source value is first truncated (round toward zero) and then checked against the range of the destination type. If the truncated value is representable in the result type that value is returned, otherwise the floating-point invalid exception is raised, and if this exception is masked, the representable value closest to the truncated value is returned. If the source value is NaN, zero is returned. If some rounding mode other than truncation is desired, it can be applied using the VROUNDSD or VRNDSCALESD instruction before the conversion. EVEX.W1 version: promotes the instruction to produce 64-bit data in 64-bit mode.

13.13. VCVTTSD2USIS

13.13.3 OPERATION

```

1 VCVTTSD2USIS (EVEX encoded version)
2 IF 64-bit Mode and OperandSize = 64
3 THEN
4     DEST[63:0] := convert_DP_to_QW_UnSignedInteger_TruncateSaturate(SRC[63:0]);
5 ELSE
6     DEST[31:0] := convert_DP_to_DW_SignedInteger_TruncateSaturate(SRC[63:0]);
7 FI;
8
9 SIMD Floating-Point Exceptions
10 -Invalid, Precision
11 Other Exceptions
12 -EVEX-encoded instructions, see Table 2-48, "Type E3NF Class Exception Conditions".
13
14 NOTES:
15 1. For this specific instruction, EVEX.W in non-64 bit is ignored; the
16 instructions behaves as if the W0 version is used.
17

```

13.13.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

13.13.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTSD2USIS r32, xmm1/m64	E3NF	IP	AVX10.2
VCVTTSD2USIS r64, xmm1/m64	E3NF	IP	AVX10.2

13.14 VCVTTSS2SIS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.LLIG.F3.MAP5.W0 6D /r VCVTTSS2SIS r32, xmm1/m32 {sae}	A	V/V ¹	AVX10.2
EVEX.LLIG.F3.MAP5.W1 6D /r VCVTTSS2SIS r64, xmm1/m32 {sae}	A	V/N.E.	AVX10.2
Notes:			
1: Outside of 64b mode, the EVEX.W field is ignored. The instruction behaves as if W=0 were used.			

13.14.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	SCALAR	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

13.14.2 DESCRIPTION

Note:

For SCALAR instructions with a CPUID feature flag specifying Intel® AVX10, the programmer can skip the check for the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and usually determines the set of instructions available to the programmer listed in the above opcode table. However, scalar instructions are architected without the notion of vector length, so this step does not strictly apply.

VCVTTSS2SIS: Converts a single-precision floating-point value in the source operand (the second operand) to a doubleword integer (or quadword integer if operand size is 64 bits) in the destination operand (the first operand) with truncation and saturation. The source operand can be an XMM register or a memory location. The destination operand is a general-purpose register. When the source operand is an XMM register, the single-precision floating-point value is contained in the low doubleword of the register. When a conversion is inexact, the source value is first truncated (round toward zero) and then checked against the range of the destination type. If the truncated value is representable in the result type that value is returned, otherwise the floating-point invalid exception is raised, and if this exception is masked, the representable value closest to the truncated value is returned. If the source value is NaN, zero is returned. If some rounding mode other than truncation is desired, it can be applied using the VROUNDSS or VRNDSCALESS instruction before the conversion. EVEX.W1 version: promotes the instruction to produce 64-bit data in 64-bit mode. Note: EVEX.vvvv is reserved and must be 1111b, otherwise instructions will #UD.

13.14. VCVTTSS2SIS

13.14.3 OPERATION

```

1
2 VCVTTSS2SIS (EVEX encoded version)
3 IF 64-bit Mode and OperandSize = 64
4 THEN
5     DEST[63:0] := convert_SP_to_QW_SignedInteger_TruncateSaturate(SRC[31:0]);
6 ELSE
7     DEST[31:0] := convert_SP_to_DW_SignedInteger_TruncateSaturate(SRC[31:0]);
8 FI;
9
10 SIMD Floating-Point Exceptions
11 -Invalid, Precision
12
13 Other Exceptions
14 -EVEX-encoded instructions, see Table 2-48, "Type E3NF Class Exception Conditions".
15
16 NOTES:
17
18 1. For this specific instruction, EVEX.W in non-64 bit is ignored; the
19 instructions behaves as if the W0 version is used.
20

```

13.14.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

13.14.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTSS2SIS r32, xmm1/m32	E3NF	IP	AVX10.2
VCVTTSS2SIS r64, xmm1/m32	E3NF	IP	AVX10.2

13.15 VCVTTSS2USIS

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.LLIG.F3.MAP5.W0 6C /r VCVTTSS2USIS r32, xmm1/m32 {sae}	A	V/V ¹	AVX10.2
EVEX.LLIG.F3.MAP5.W1 6C /r VCVTTSS2USIS r64, xmm1/m32 {sae}	A	V/N.E.	AVX10.2
Notes:			
1: Outside of 64b mode, the EVEX.W field is ignored. The instruction behaves as if W=0 were used.			

13.15.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	SCALAR	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A

13.15.2 DESCRIPTION

Note:

For SCALAR instructions with a CPUID feature flag specifying Intel® AVX10, the programmer can skip the check for the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and usually determines the set of instructions available to the programmer listed in the above opcode table. However, scalar instructions are architected without the notion of vector length, so this step does not strictly apply.

VCVTTSS2USIS: Converts a single-precision floating-point value in the source operand (the second operand) to a doubleword unsigned integer (or unsigned quadword integer if operand size is 64 bits) in the destination operand (the first operand) with truncation and saturation. The source operand can be an XMM register or a memory location. The destination operand is a general-purpose register. When the source operand is an XMM register, the single-precision floating-point value is contained in the low doubleword of the register. When a conversion is inexact, the source value is first truncated (round toward zero) and then checked against the range of the destination type. If the truncated value is representable in the result type that value is returned, otherwise the floating-point invalid exception is raised, and if this exception is masked, the representable value closest to the truncated value is returned. If the source value is NaN, zero is returned. If some rounding mode other than truncation is desired, it can be applied using the VROUNDSS or VRNDSCALESS instruction before the conversion. EVEX.W1 version: promotes the instruction to produce 64-bit data in 64-bit mode. Note: EVEX.vvvv is reserved and must be 1111b, otherwise instructions will #UD.

13.15. VCVTTSS2USIS

13.15.3 OPERATION

```

1 VCVTTSS2USIS (EVEX encoded version)
2 IF 64-bit Mode and OperandSize = 64
3 THEN
4     DEST[63:0] := convert_SP_to_QW_UnSignedInteger_TruncateSaturate(SRC[31:0]);
5 ELSE
6     DEST[31:0] := convert_SP_to_DW_UnSignedInteger_TruncateSaturate(SRC[31:0]);
7 FI;
8
9 SIMD Floating-Point Exceptions
10 -Invalid, Precision
11
12 Other Exceptions
13 -EVEX-encoded instructions, see Table 2-48, "Type E3NF Class Exception Conditions".
14
15 NOTES:
16 1. For this specific instruction, EVEX.W in non-64 bit is ignored; the
17 instructions behaves as if the W0 version is used.
18

```

13.15.4 SIMD FLOATING-POINT EXCEPTIONS

This instruction can set the following flags in MXCSR: Invalid, Precision In abbreviated form, this includes: IE, PE

13.15.5 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VCVTTSS2USIS r32, xmm1/m32	E3NF	IP	AVX10.2
VCVTTSS2USIS r64, xmm1/m32	E3NF	IP	AVX10.2

Chapter 14

INTEL® AVX10 ZERO-EXTENDING PARTIAL VECTOR COPY INSTRUCTIONS

14.1. VMOVD

14.1 VMOVD

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.F3.0F.W0 7E /r VMOVD xmm1, xmm2/m32	A	V/V	AVX10.2
EVEX.128.66.0F.W0 D6 /r VMOVD xmm1/m32, xmm2	B	V/V	AVX10.2

14.1.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	TUPLE1	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A
B	TUPLE1	MODRM.R/M(w)	MODRM.REG(r)	N/A	N/A

14.1.2 DESCRIPTION

Note:

For SCALAR instructions with a CPUID feature flag specifying Intel® AVX10, the programmer can skip the check for the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and usually determines the set of instructions available to the programmer listed in the above opcode table. However, scalar instructions are architected without the notion of vector length, so this step does not strictly apply.

VMOVD: Copies a doubleword from the source operand (second operand) to the destination operand (first operand). The source and destination operands can be XMM registers, or 32-bit memory locations. This instruction can be used to move a doubleword between two XMM registers or between an XMM register and a 32-bit memory location. The instruction cannot be used to transfer data between memory locations. When the source operand is an XMM register, the low doubleword is moved; when the destination operand is an XMM register, the doubleword is stored to the low doubleword of the register, and the remaining bits are cleared to all 0s. Note: EVEX.vvvv is reserved and must be 1111b, otherwise instructions will #UD. EVEX.LL must be 00b, otherwise instructions will #UD.

14.1. VMOVD

14.1.3 OPERATION

```

1
2 VMOVD (7E) with XMM register source and destination
3 DEST[31:0] := SRC[31:0]
4 DEST[MAXVL-1:32] := 0
5
6 VMOVD (D6) with XMM register source and destination
7 DEST[31:0] := SRC[31:0]
8 DEST[MAXVL-1:32] := 0
9
10 VMOVD (7E) with memory source
11 DEST[31:0] := SRC[31:0]
12 DEST[:MAXVL-1:32] := 0
13
14 VMOVD (D6) with memory dest
15 DEST[31:0] := SRC2[31:0]
16
17 Flags Affected
18 None.
19
20

```

14.1.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMOVD xmm1, xmm2/m32	E9NF	N/A	AVX10.2
VMOVD xmm1/m32, xmm2	E9NF	N/A	AVX10.2

14.2 VMOVW

Encoding / Instruction	Op/En	64/32-bit mode	CPUID
EVEX.128.F3.MAP5.W0 6E /r VMOVW xmm1, xmm2/m16	A	V/V	AVX10.2
EVEX.128.F3.MAP5.W0 7E /r VMOVW xmm1/m16, xmm2	B	V/V	AVX10.2

14.2.1 INSTRUCTION OPERAND ENCODING

Op/En	Tuple	Operand 1	Operand 2	Operand 3	Operand 4
A	TUPLE1	MODRM.REG(w)	MODRM.R/M(r)	N/A	N/A
B	TUPLE1	MODRM.R/M(w)	MODRM.REG(r)	N/A	N/A

14.2.2 DESCRIPTION

Note:

For SCALAR instructions with a CPUID feature flag specifying Intel® AVX10, the programmer can skip the check for the available vector options on the processor at run-time via the CPU_SUPPORTED_VECTOR_LENGTHS field in Converged Vector ISA Leaf 0x24. This field enumerates the maximum supported vector width and usually determines the set of instructions available to the programmer listed in the above opcode table. However, scalar instructions are architected without the notion of vector length, so this step does not strictly apply.

VMOVW: Copies a word from the source operand (second operand) to the destination operand (first operand). The source and destination operands can be XMM registers, or 16-bit memory locations. This instruction can be used to move a word between two XMM registers or between an XMM register and a 16-bit memory location. The instruction cannot be used to transfer data between memory locations. When the source operand is an XMM register, the low word is moved; when the destination operand is an XMM register, the word is stored to the low word of the register, and the remaining bits are cleared to all 0s. Note: EVEX.vvvv is reserved and must be 1111b, otherwise instructions will #UD. EVEX.LL must be 00b, otherwise instructions will #UD.

14.2. VMOVW

14.2.3 OPERATION

```

1
2 VMOVW (xx) with XMM register source and destination
3 DEST[15:0] := SRC[15:0]
4 DEST[MAXVL-1:16] := 0
5
6 VMOVW (xx) with XMM register source and destination
7 DEST[15:0] := SRC[15:0]
8 DEST[MAXVL-1:16] := 0
9
10 VMOVW (xx) with memory source
11 DEST[15:0] := SRC[15:0]
12 DEST[:MAXVL-1:16] := 0
13
14 VMOVW (xx) with memory dest
15 DEST[15:0] := SRC2[15:0]
16

```

14.2.4 EXCEPTIONS

Instruction	Exception Type	Arithmetic Flags	CPUID
VMOVW xmm1, xmm2/m16	E9NF	N/A	AVX10.2
VMOVW xmm1/m16, xmm2	E9NF	N/A	AVX10.2