

Intel[®] Xeon[®] Processor E5 v4 Product Family

Datasheet Volume 2: Registers

June 2016



Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at Intel.com, or from the OEM or retailer.

No computer system can be absolutely secure. Intel does not assume any liability for lost or stolen data or systems or any damages resulting from such losses.

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.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

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.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

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

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

Copyright © 2016, Intel Corporation. All Rights Reserved.



Contents

1	Registers Overview and Configuration Process	15
1.1	Platform Configuration Structure	15
1.1.1	Processor IIO Devices (CPUBUSNO (0))	15
1.1.2	Processor Uncore Devices (CPUBUSNO (1))	17
1.2	Configuration Register Rules	18
1.2.1	CSR Access	18
1.2.2	MSR Access	21
1.2.3	Memory-Mapped I/O Registers	21
1.3	Register Terminology	22
1.4	Protected Processor Inventory Number	23
2	Integrated Memory Controller (iMC) Configuration Registers	25
2.1	Device 19,22 Function 0	25
2.1.1	pxpcap	26
2.1.2	mcmtr	27
2.1.3	tadwayness_[0:11]	27
2.1.4	mc_init_state_g	28
2.1.5	rcomp_timer	29
2.1.6	mh_sense_500ns_cfg	30
2.1.7	mh_dtycyc_min_asrt_cntr_[0:1]	30
2.1.8	mh_io_500ns_cntr	31
2.1.9	mh_chn_astn	31
2.1.10	mh_ext_stat	32
2.1.11	smb_stat_[0:1]	32
2.1.12	smbcmd_[0:1]	34
2.1.13	smbcntl_[0:1]	35
2.1.14	smb_tsod_poll_rate_cntr_[0:1]	36
2.1.15	smb_period_cfg	36
2.1.16	smb_period_cntr	37
2.1.17	smb_tsod_poll_rate	37
2.2	Device 19,22 Function 1	37
2.2.1	pxpcap	38
2.2.2	spareaddresslo	38
2.2.3	sparectl	39
2.2.4	ssrstatus	40
2.2.5	scrubaddresslo	40
2.2.6	scrubaddresshi	41
2.2.7	scrubctl	41
2.2.8	spareinterval	42
2.2.9	rasenables	42
2.2.10	smisparectl	43
2.2.11	leaky_bucket_cfg	43
2.2.12	leaky_bucket_cntr_lo	45
2.2.13	leaky_bucket_cntr_hi	46
2.3	Device 19,22 Functions 2,3,4,5	46
2.3.1	pxpcap	47
2.3.2	dimmmtr_[0:2]	47
2.3.3	pxpenhcap	49
2.4	Device 20,21,23 Functions 0, 1	49
2.4.1	pxpcap	51
2.4.2	chn_temp_cfg	51
2.4.3	chn_temp_stat	51
2.4.4	dimmm_temp_oem_[0:2]	52



2.4.5	dimmm_temp_th_[0:2].....	52
2.4.6	dimmm_temp_thrt_lmt_[0:2]	53
2.4.7	dimmm_temp_ev_ofst_[0:2]	53
2.4.8	dimmmtempstat_[0:2].....	54
2.4.9	thrt_pwr_dimmm_[0:2]	55
2.5	Device 20,21,23 Functions 2, 3.....	55
2.5.1	correrrcnt_0.....	57
2.5.2	correrrcnt_1.....	58
2.5.3	correrrcnt_2.....	58
2.5.4	correrrcnt_3.....	59
2.5.5	correrrthrshld_0	59
2.5.6	correrrthrshld_1	60
2.5.7	correrrthrshld_2	60
2.5.8	correrrthrshld_3	60
2.5.9	correrrorstatus	61
2.5.10	leaky_bkt_2nd_cntr_reg	61
2.5.11	devtag_cntl_[0:7].....	62
3	Intel® QuickPath Interconnect (Intel® QPI) Agent Registers	65
3.1	Device 8 Function 0	65
3.1.1	QPIMISCSTAT: Intel QPI Misc Status	66
4	Processor Utility Box (UBOX) Registers	67
4.1	Device 16 Function 5.....	67
4.1.1	CPUNODEID.....	68
4.1.2	IntControl.....	68
4.1.3	GIDNIDMAP	69
4.1.4	UBOXErrSts	70
4.2	Device 16 Function 7.....	70
4.2.1	CPUBUSNO	71
4.2.2	SMICtrl	71
5	Power Controller Unit (PCU) Registers	73
5.1	Device 30 Function 0.....	73
5.1.1	MEM_TRML_TEMPERATURE_REPORT	74
5.1.2	MEM_ACCUMULATED_BW_CH_[0:3].....	74
5.1.3	PACKAGE_POWER_SKU	75
5.1.4	PACKAGE_POWER_SKU_UNIT.....	75
5.1.5	PACKAGE_ENERGY_STATUS	76
5.1.6	Package_Temperature.....	76
5.1.7	TEMPERATURE_TARGET	76
5.2	Device 30 Function 1.....	77
5.2.1	SSKPD	77
5.2.2	C2C3TT	78
5.2.3	CSR_DESIRED_CORES	78
5.3	Device 30 Function 2.....	78
5.3.1	PACKAGE_RAPL_PERF_STATUS.....	79
5.3.2	DRAM_POWER_INFO	79
5.3.3	DRAM_ENERGY_STATUS.....	80
5.3.4	DRAM_ENERGY_STATUS_CH[0:3]	80
5.3.5	DRAM_RAPL_PERF_STATUS	81
5.3.6	MCA_ERR_SRC_LOG	81
5.3.7	THERMTRIP_CONFIG	82
5.4	Device 30 Function 3.....	82
5.4.1	CAP_HDR	83
5.4.2	CAPIDO	83
5.4.3	CAPID1	84
5.4.4	CAPID2	85



5.4.5	CAPID3	86
5.4.6	CAPID4	87
5.4.7	CAPID5	89
5.4.8	CAPID6	89
5.4.9	SMT_CONTROL	90
5.4.10	RESOLVED_CORES	90
6	Integrated I/O (IIO) Configuration Registers	91
6.1	Registers Overview	91
6.1.1	Configuration Registers (CSR)	91
6.1.2	BDF: BAR# for Various MMIO BARs in IIO	91
6.1.3	Unimplemented Devices/Functions and Registers	92
6.1.4	PCI Vs. PCIe Device / Function	92
6.2	Device 0 Function 0 DMI, Device 0 Function 0 PCIe, Device 1 Function 0-1, Device 2 Function 0-3 PCIe, Device 3 Function 0-3 PCIe92	
6.2.1	vid	96
6.2.2	did	97
6.2.3	pcicmd	97
6.2.4	pcists	99
6.2.5	rid	100
6.2.6	ccr	101
6.2.7	clsr	101
6.2.8	plat	101
6.2.9	hdr	102
6.2.10	bist	102
6.2.11	pbus	103
6.2.12	secbus	103
6.2.13	subbus	103
6.2.14	iobas	104
6.2.15	iolim	104
6.2.16	secsts	105
6.2.17	mbas	106
6.2.18	mlim	106
6.2.19	pbas	107
6.2.20	plim	107
6.2.21	pbasu	107
6.2.22	plimu	108
6.2.23	capptr	108
6.2.24	intl	109
6.2.25	intpin	109
6.2.26	bctrl	109
6.2.27	scapid	110
6.2.28	snxtptr	110
6.2.29	svid	111
6.2.30	sdid	111
6.2.31	dmircbar	112
6.2.32	msicapid	112
6.2.33	msinxtptr	112
6.2.34	msimsgctl	113
6.2.35	msgadr	113
6.2.36	msgdat	114
6.2.37	msimsk	114
6.2.38	msipending	114
6.2.39	pxpcapid	114
6.2.40	pxpnxtptr	115
6.2.41	pxpcap	115
6.2.42	devcap	116



6.2.43	devctrl	117
6.2.44	devsts	118
6.2.45	lnkcap	119
6.2.46	lnkcon	120
6.2.47	lnksts	122
6.2.48	sltcap	124
6.2.49	sltcon	125
6.2.50	sltsts	127
6.2.51	rootcon	129
6.2.52	rootcap	131
6.2.53	rootsts	131
6.2.54	devcap2	132
6.2.55	devctrl2	133
6.2.56	lnkcap2	134
6.2.57	lnkcon2	134
6.2.58	lnksts2	136
6.2.59	pmcap	137
6.2.60	pmcsr	137
6.2.61	xpreut_hdr_ext	138
6.2.62	xpreut_hdr_cap	139
6.2.63	xpreut_hdr_lef	140
6.2.64	acscaphdr	140
6.2.65	acscap	141
6.2.66	acsctrl	141
6.2.67	apibase	142
6.2.68	apiclimit	143
6.2.69	vsecphdr	143
6.2.70	vshdr	143
6.2.71	errcaphdr	144
6.2.72	uncerrsts	144
6.2.73	uncerrmsk	144
6.2.74	uncerrsev	145
6.2.75	corerrsts	146
6.2.76	corerrmsk	146
6.2.77	errcap	146
6.2.78	hdrlog[0:3]	147
6.2.79	rperrcmd	147
6.2.80	rperrsts	148
6.2.81	errsid	149
6.2.82	perfctrlsts_0	149
6.2.83	perfctrlsts_1	151
6.2.84	miscctrlsts_0	151
6.2.85	miscctrlsts_1	153
6.2.86	pcie_iou_bif_ctrl	156
6.2.87	dmictrl	156
6.2.88	dmists	157
6.2.89	ERRINJCAP	157
6.2.90	ERRINJHDR	157
6.2.91	ERRINJCON	158
6.2.92	ctoctrl	158
6.2.93	xpcorerrsts	159
6.2.94	xpcorerrmsk	159
6.2.95	xpuncerrsts	160
6.2.96	xpuncerrmsk	160
6.2.97	xpuncerrsev	161
6.2.98	xpuncerrptr	161



6.2.99	uncedmask	161
6.2.100	coredmask	162
6.2.101	rpedmask	162
6.2.102	xpuncedmask	163
6.2.103	xpcoredmask	163
6.2.104	xpglberrsts	164
6.2.105	xpglberrptr	164
6.2.106	pxp2cap	165
6.2.107	lnkcon3	165
6.2.108	lnerrsts	166
6.2.109	ln[0:3]eq	166
6.2.110	ln[4:7]eq	168
6.2.111	ln[8:15]eq	169
6.2.112	mcast_cap_hdr	171
6.2.113	mcast_cap_ext	171
6.2.114	mcast_cap	171
6.2.115	mcast_ctrl	172
6.2.116	mcast_base	172
6.2.117	mcast_rcv	172
6.2.118	mcast_blk_all	173
6.2.119	mcast_blk_unt	173
6.2.120	mcast_overlay_bar	173
6.3	Device 0 Function 0 Region DMIRCBAR	174
6.3.1	dmiv0rcap	174
6.3.2	dmiv0rctl	174
6.3.3	dmiv0rst	175
6.3.4	dmiv1rcap	176
6.3.5	dmiv1rctl	176
6.3.6	dmiv1rst	177
6.3.7	dmivpcrcap	178
6.3.8	dmivpcrctl	178
6.3.9	dmivpcrst	179
6.3.10	dmivmrcap	180
6.3.11	dmivmrctl	180
6.3.12	dmivmrst	181
6.3.13	dmiv1cdtthrottle	181
6.3.14	dmivpcdtthrottle	182
6.3.15	dmivmcdtthrottle	182
6.4	Device 4 Function 0-7	183
6.4.1	vid	184
6.4.2	did	184
6.4.3	pcicmd	184
6.4.4	pcists	185
6.4.5	rid	185
6.4.6	ccr	185
6.4.7	clsr	186
6.4.8	hdr	186
6.4.9	cb_bar	186
6.4.10	svid	187
6.4.11	sdid	187
6.4.12	captr	187
6.4.13	intl	187
6.4.14	intpin	188
6.4.15	devcfg	188
6.4.16	msixcapid	188
6.4.17	msixnxtptr	189



6.4.18	msixmsgctl	189
6.4.19	tableoff_bir	189
6.4.20	pbaoff_bir	190
6.4.21	capid	190
6.4.22	nextptr	190
6.4.23	expcap	191
6.4.24	devcap	191
6.4.25	devcon	192
6.4.26	devsts	193
6.4.27	devcap2	193
6.4.28	devcon2	193
6.4.29	pmcap	193
6.4.30	pmcsr	194
6.4.31	dmauncerrsts	195
6.4.32	dmauncerrmsk	196
6.4.33	dmauncerrsev	196
6.4.34	dmauncerrptr	197
6.4.35	dmaglberrptr	197
6.4.36	chanerr_int	197
6.4.37	chanerrmsk_int	199
6.4.38	chanerrsev_int	200
6.4.39	chanerrptr	200
6.5	Device 4 Function 0 - 7 MMIO Region Intel QuickData Technology BARs	200
6.5.1	chancnt	202
6.5.2	xfercap	202
6.5.3	genctrl	202
6.5.4	intrctrl	203
6.5.5	attnstatus	203
6.5.6	cbver	204
6.5.7	intrdelay	204
6.5.8	cs_status	204
6.5.9	dmacapability	205
6.5.10	dcaoffset	206
6.5.11	cbprio	206
6.5.12	chanctrl	207
6.5.13	dma_comp	208
6.5.14	chancmd	208
6.5.15	dmacount	208
6.5.16	chansts_0	209
6.5.17	chansts_1	209
6.5.18	chainaddr_0	210
6.5.19	chainaddr_1	210
6.5.20	chancmp_0	210
6.5.21	chancmp_1	211
6.5.22	chanerr	211
6.5.23	chanerrmsk	213
6.5.24	dcactrl	213
6.5.25	dca_ver	214
6.5.26	dca_reqid_offset	214
6.5.27	csi_capability	214
6.5.28	pcie_capability	214
6.5.29	csi_cap_enable	215
6.5.30	pcie_cap_enable	215
6.5.31	apicid_tag_map	215
6.5.32	dca_reqid[0:1]	216
6.5.33	msgaddr	217



6.5.34	msgupaddr	217
6.5.35	msgdata	218
6.5.36	vecctrl	218
6.5.37	pendingbits	218
6.6	Device 5 Function 0	219
6.6.1	vid	220
6.6.2	did	220
6.6.3	pcicmd	221
6.6.4	pcists	221
6.6.5	rid	222
6.6.6	ccr	222
6.6.7	clsr	223
6.6.8	hdr	223
6.6.9	svid	223
6.6.10	sdid	223
6.6.11	captr	224
6.6.12	intl	224
6.6.13	intpin	224
6.6.14	pxpcapid	224
6.6.15	pxpnxtptr	224
6.6.16	pxpcap	225
6.6.17	hdrtypectrl	225
6.6.18	mmcfg_base	225
6.6.19	mmcfg_limit	226
6.6.20	tommiol_ob	226
6.6.21	tseg	226
6.6.22	genprotrange[1:0]_base	226
6.6.23	genprotrange[1:0]_limit	228
6.6.24	genprotrange2_base	228
6.6.25	genprotrange2_limit	229
6.6.26	tolm	229
6.6.27	tohm	229
6.6.28	tommiol	230
6.6.29	ncmem_base	231
6.6.30	ncmem_limit	231
6.6.31	mencmem_base	231
6.6.32	mencmem_limit	232
6.6.33	cpubusno	232
6.6.34	lmmiol_base	233
6.6.35	lmmiol_limit	233
6.6.36	lmmioh_base	234
6.6.37	lmmioh_limit	234
6.6.38	cipctrl	235
6.6.39	cipsts	236
6.6.40	cipdcasad	236
6.6.41	cipintrc	237
6.6.42	cipintrs	238
6.6.43	vtbar	238
6.6.44	vtgenctrl	239
6.6.45	vtgenctrl2	239
6.6.46	iotlbpartition	240
6.6.47	vtuncerrsts	241
6.6.48	vtuncerrmsk	242
6.6.49	vtuncerrsev	243
6.6.50	vtuncerrptr	243
6.6.51	iiomiscctrl	244



	6.6.52	ltdpr	246
	6.6.53	lcfgbus_base	247
	6.6.54	lcfgbus_limit	248
	6.6.55	csipintrs	248
6.7		Device 5 Function 0 MMIO Region VTBAR	248
	6.7.1	vtd[0:1]_version	250
	6.7.2	vtd[0:1]_cap	250
	6.7.3	vtd[0:1]_ext_cap	251
	6.7.4	vtd[0:1]_glbcmd	252
	6.7.5	vtd[0:1]_glbsts	254
	6.7.6	vtd[0:1]_rootentryadd	255
	6.7.7	vtd[0:1]_ctxcmd	256
	6.7.8	vtd[0:1]_fltsts	257
	6.7.9	nonisoch_fltevtctrl	258
	6.7.10	nonisoch_fltevtdata	258
	6.7.11	vtd[0:1]_fltevtaddr	259
	6.7.12	vtd[0:1]_fltevtupraddr	259
	6.7.13	vtd[0:1]_pmen	259
	6.7.14	vtd[0:1]_prot_low_mem_base	260
	6.7.15	vtd[0:1]_prot_low_mem_limit	260
	6.7.16	vtd[0:1]_prot_high_mem_base	260
	6.7.17	vtd[0:1]_prot_high_mem_limit	261
	6.7.18	vtd[0:1]_inv_queue_head	261
	6.7.19	vtd[0:1]_inv_queue_tail	261
	6.7.20	vtd[0:1]_inv_queue_add	262
	6.7.21	vtd[0:1]_inv_comp_status	262
	6.7.22	nonisoch_inv_cmp_evtctrl	262
	6.7.23	nonisoch_invevtdata	263
	6.7.24	vtd[0:1]_inv_comp_evt_addr	263
	6.7.25	vtd[0:1]_inv_comp_evt_upraddr	263
	6.7.26	vtd[0:1]_intr_remap_table_base	264
	6.7.27	vtd0_fltrec[0:7]_gpa, vtd1_fltrec0_gpa	264
	6.7.28	vtd0_fltrec[0:7]_src, vtd1_fltrec0_src	265
	6.7.29	vtd[0:1]_invaddrreg	265
	6.7.30	vtd[0:1]_iotlbinv	266
6.8		Device 5 Function 2	267
	6.8.1	vid	269
	6.8.2	did	269
	6.8.3	pcicmd	270
	6.8.4	pcists	270
	6.8.5	rid	271
	6.8.6	ccr	271
	6.8.7	clsr	272
	6.8.8	hdr	272
	6.8.9	svid	272
	6.8.10	sdid	272
	6.8.11	capptr	273
	6.8.12	intl	273
	6.8.13	intpin	273
	6.8.14	pxpcapid	273
	6.8.15	pxpnxtptr	273
	6.8.16	pxpcap	274
	6.8.17	irpperrsv	274
	6.8.18	iioerrsv	275
	6.8.19	mierrsv	276
	6.8.20	pcierrsv	276



6.8.21	sysmap	277
6.8.22	vppctl	277
6.8.23	vppsts	278
6.8.24	vppfreq	279
6.8.25	gcerrst	279
6.8.26	gcferrst	280
6.8.27	gcnerrst	280
6.8.28	gnerrst	281
6.8.29	gferrst	282
6.8.30	gerrctl	282
6.8.31	gsysst	283
6.8.32	gsysctl	284
6.8.33	gfferrst, gfnerrst	284
6.8.34	gnferrst, gnnerrst	284
6.8.35	irpp[0:1]errst	285
6.8.36	irpp[0:1]errctl	285
6.8.37	irpp[0:1]fferrst, irpp[0:1]fnerrst	286
6.8.38	irpp[0:1]fferrhd[0:3]	287
6.8.39	irpp[0:1]nferrst, irpp[0:1]nnerrst	287
6.8.40	irpp[0:1]nferrhd[0:3]	288
6.8.41	irpp[0:1]errcntsel	288
6.8.42	irpp[0:1]errcnt	288
6.8.43	iioerrst	289
6.8.44	iioerrctl	290
6.8.45	iiofferrst, iiofnerrst	290
6.8.46	iiofferrhd_[0:3]	290
6.8.47	iionferrst, iionnerrst	291
6.8.48	iionferrhd_[0:3]	291
6.8.49	iioerrcntsel	291
6.8.50	iioerrcnt	292
6.8.51	mierrst	292
6.8.52	mierrctl	292
6.8.53	mifferrst, mifnerrst	293
6.8.54	mifferrhdr_[0:3]	293
6.8.55	minferrst, minnerrst	293
6.8.56	minferrhdr_[0:3]	293
6.8.57	mierrcntsel	294
6.8.58	mierrcnt	294
6.9	Device 5 Function 4	294
6.9.1	vid	295
6.9.2	did	295
6.9.3	pcicmd	295
6.9.4	pcists	296
6.9.5	rid	296
6.9.6	ccr	297
6.9.7	clsr	297
6.9.8	hdr	297
6.9.9	mbar	297
6.9.10	svid	298
6.9.11	sid	298
6.9.12	captr	298
6.9.13	intlin	299
6.9.14	intpin	299
6.9.15	abar	299
6.9.16	pxpcap	300
6.9.17	snapshot_index	300



6.9.18	snapshot_window	300
6.9.19	ioapictetpc	301
6.9.20	pmcap	301
6.9.21	pmcsr	302
6.9.22	loadsels0	303
6.9.23	iointsrc0	303
6.9.24	iointsrc1	304
6.9.25	ioremintcnt	304
6.9.26	ioremgpecnt	305
6.9.27	FauxGV	305
6.10	Device 5 Function 4 I/OxAPIC	305
6.10.1	index	305
6.10.2	window	306
6.10.3	eoi	306
6.11	Device 5 Function 4 Window 0	306
6.12	Device 6-7 Function 0,1,3	310
6.12.1	rx_ctle_peak_gen2	311
6.12.2	rx_ctle_peak_gen2	311
6.12.3	rx_ctle_peak_gen3	311
6.12.4	rx_ctle_peak_gen2	312
6.12.5	rx_ctle_peak_gen3	312
6.13	Non Transparent Bridge Registers	312
6.13.1	Configuration Register Map (NTB Primary Side)	312
6.13.2	Standard PCI Configuration Space - Type 0 Common Configuration Space	316
6.13.3	NTB Port 3A Configured as Primary Endpoint Device	322
6.13.4	PCI Express Configuration Registers (NTB Secondary Side)	357
6.13.5	Configuration Register Map (NTB Secondary Side)	357
6.13.6	NTB Shadowed MMIO Space	386
6.13.7	NTB Primary/Secondary Host MMIO Registers	388
6.13.8	MSI-X MMIO Registers (NTB Primary side)	403
6.13.9	MSI-X MMIO registers (NTB Secondary Side)	405

Figures


1-1	Processor Integrated I/O Device Map	16
1-2	Processor Uncore Devices Map	17

Tables

1-1	Functions Specifically Handled by the Processor	19
1-2	Register Attributes Definitions	22
6-1	BDF:BAR# for Various MMIO BARs in IIO	91
6-2	Function Number of Active Root Ports in Port 1(Dev#1) based on Port Bifurcation	92
6-3	Function Number of Active Root Ports in Port 2(Dev#2) based on Port Bifurcation	92
6-4	Function Number of Active Root Ports in Port 3(Dev#3) based on Port Bifurcation	93
6-5	Device 3 Function 0 (Non-Transparent Bridge) Configuration Map Offset 0x00h - 0xFCh	313
6-6	Device 3 Function 0 (Non-Transparent Bridge) Configuration Map Offset 0x100h - 0x1FCh	313
6-7	Device 3 Function 0 (Non-Transparent Bridge) Configuration Map Offset 0x200h - 0x2FCh	314
6-8	Device 3 Function 0 (Non-Transparent Bridge) Configuration Map Offset 0x300h - 0x3FCh	315
6-9	Device 0 Function 0 (Non -Transparent Bridge) Configuration Map 0x00h - 0xFCh	358
6-10	Device 0 Function 0 (Non -Transparent Bridge) Configuration Map 0x100h - 0x1FCh	358
6-11	NTB MMIO Shadow Registers	386




6-12	NTB MMIO Map	387
6-13	NTB MMIO Map	403
6-14	MSI-X Vector Handling and Processing by IIO on Primary Side.....	404
6-15	NTB MMIO Map	405
6-16	MSI-X Vector Handling and Processing by IIO on Secondary Side.....	407



Revision History

Revision Number	Description	Date
001	<ul style="list-style-type: none">Initial release of the document.	April 2016
002	<ul style="list-style-type: none">Updated Chapters 4 and 6	June 2016

§



1 Registers Overview and Configuration Process

The Intel® Xeon® Processor E5 v4 product family contains one or more PCI devices within each individual functional block. The configuration registers for these devices are mapped as devices residing on the PCI Bus assigned for the processor socket.

Some features are only supported on specific SKUs. In such case, the respective registers would only apply to the specific SKU that contains the feature support.

Refer to the *Intel® Xeon® Processor E5 v4 Product Families Uncore Performance Monitoring Reference Manual* for details on Performance Monitoring Registers.

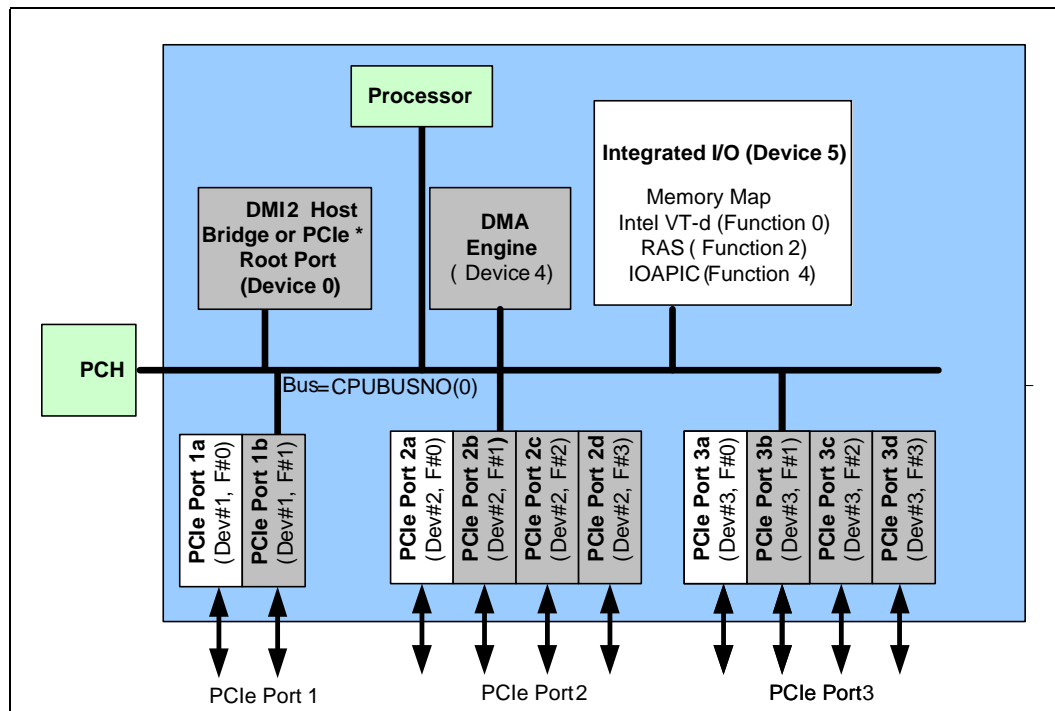
1.1 Platform Configuration Structure

The DMI2 physically connects the processor and the PCH. From a configuration standpoint, the DMI2 is a logical extension of PCI Bus 0. DMI2 and the internal devices in the processor IIO and PCH logically constitute PCI Bus 0 to configuration software. As a result, all devices internal to the processor and the PCH appear to be on PCI Bus 0.


1.1.1 Processor IIO Devices (CPUBUSNO (0))

The processor IIO contains PCI devices within a single, physical component. The configuration registers for the devices are mapped as devices residing on PCI Bus "CPUBUSNO(0)" where CPUBUSNO(0) is programmable by BIOS.

Figure 1-1. Processor Integrated I/O Device Map



- **Device 0:** DMI2 Root Port. Logically this appears as a PCI device residing on PCI Bus 0. Device 0 contains the standard PCI header registers, extended PCI configuration registers and DMI2 device specific configuration registers.
- **Device 1:** PCI Express* Root Port 1a, 1b. Logically this appears as a “virtual” PCI-to-PCI bridge residing on PCI Bus 0 and is compliant with *PCI Express Local Bus Specification Revision 2.0*. Device 1 contains the standard PCI Express/PCI configuration registers including PCI Express Memory Address Mapping registers. It also contains the extended PCI Express configuration space that include PCI Express error status/control registers and Virtual Channel controls.
- **Device 2:** PCI Express* Root Port 2a, 2b, 2c and 2d. Logically this appears as a “virtual” PCI-to-PCI bridge residing on PCI Bus 0 and is compliant with *PCI Express Local Bus Specification Revision 2.0*. Device 2 contains the standard PCI Express/PCI configuration registers including PCI Express Memory Address Mapping registers. It also contains the extended PCI Express configuration space that include PCI Express Link status/control registers and Virtual Channel controls.
- **Device 3:** PCI Express Root Port 3a, 3b, 3c and 3d. Logically this appears as a “virtual” PCI-to-PCI bridge residing on PCI Bus 0 and is compliant with *PCI Express Local Bus Specification Revision 2.0*. Device 3 contains the standard PCI Express/PCI configuration registers including PCI Express Memory Address Mapping registers. It also contains the extended PCI Express configuration space that include PCI Express error status/control registers and Virtual Channel controls.
- **Device 4:** Intel® QuickData Technology DMA. This device contains the Standard PCI registers for each of its functions. This device implements 8 functions for the 8 DMA Channels and also contains Memory Map I/O registers.
- **Device 5:** Integrated I/O Core. This device contains the Standard PCI registers for each of its functions. This device implements three functions; Function 0 contains Address Mapping, Intel® Virtualization Technology (Intel® VT) for Directed I/O

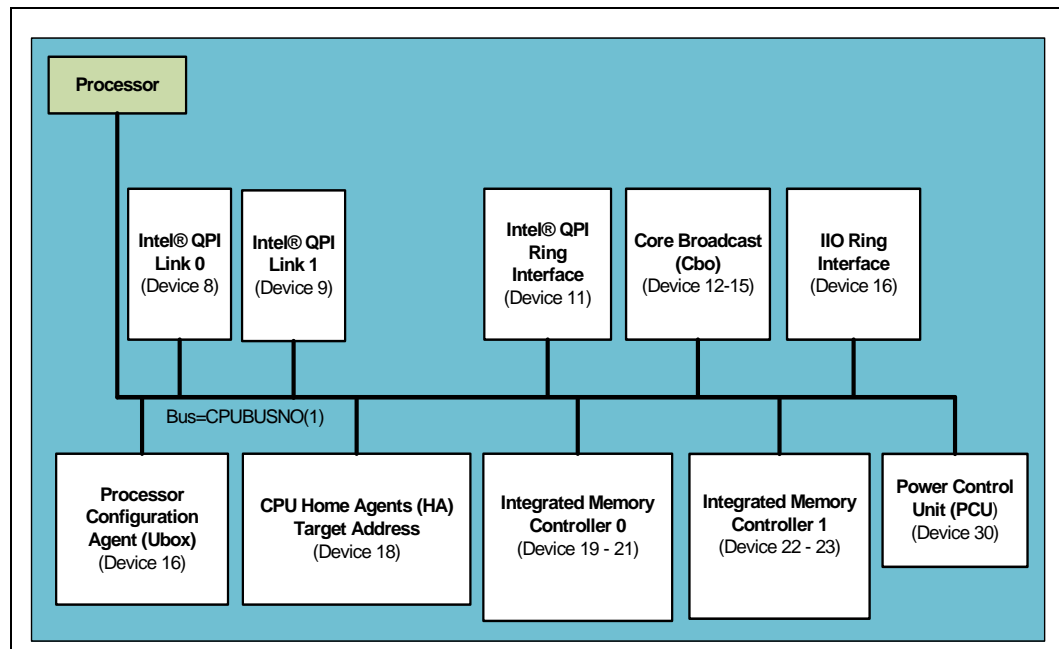


(Intel® VT-d) related registers and other system management registers. Function 1 contains PCIe* Hot-Plug registers. Function 2 contains I/O RAS registers, Function 4 contains System Control/Status registers and miscellaneous control/status registers on power management and throttling.

1.1.2 Processor Uncore Devices (CPUBUSNO (1))

The configuration registers for these devices are mapped as devices residing on the PCI bus assigned for the processor socket. Bus number is derived by the max bus range setting and processor socket number.

Figure 1-2. Processor Uncore Devices Map



- **Device 8:** Intel® QuickPath Interconnect (Intel® QPI) Link 0. Device 8 contains the Intel QPI Link 0 registers.
- **Device 9:** Intel QPI Link 1. Device 9 contains the Intel QPI Link 1 registers.
- **Device 11:** Intel QPI Ring Interface Device. Device 11 contains the processor Ring to Intel QPI registers.
- **Device 12 - 14:** Processor Caching Agent. Device 12 - 14 contain the Cbo Unicast configuration registers.
 - Implemented devices and functions in these devices vary based on SKU.
- **Device 15:** Processor Caching Agent. Device 15 contain the Cbo Broadcast configuration registers.
- **Device 16:** Integrated IO Ring Interface Device. Device 16, Functions 0, 1 contain the processor ring to PCI Express agent registers
- **Device 16:** Processor Configuration Agent. Device 16 contains the Processor Interrupt Event Handling (Ubox) registers.
- **Device 18:** Processor Home Agent(s) (HA). Functions 0-1 contain Home Agent 0 registers. Functions 4-5 contain Home Agent 1 registers. There is one Home Agent per Memory Controller. Not all Intel® Xeon® Processor E5 v4 Product Family



processors support IMC/HA 1. For SKUs with one IMC Functions 4-5 are not implemented as only HA0 is implemented.

- The Intel® Xeon® Processor E5-1600 v4 product family and some SKUs of the Intel® Xeon® Processor E5-2600 v4 product family implement 1 IMC/HA.
- **Device 19 - 21:** Integrated Memory Controller 0 configuration registers. For SKUs with one IMC, this IMC supports up to 4 channels (0-3) off of IMC 0. For SKUs with two IMC, this IMC supports 2 channels (0,1) and device 19 Functions 4, 5 (channel 2,3) and device 21 are not used and should be ignored as they apply to channels 2 and 3.
 - The Intel® Xeon® Processor E5-1600 v4 product family and some SKUs of the Intel® Xeon® Processor E5-2600 v4 product family implement 1 IMC/HA. For SKUs which implement 1 IMC/HA channels 0, 1, 2 and 3 reside off of these devices. For SKUs which implement 2 IMC/HA channels 0 and 1 reside off of these devices.
- **Device 22 - 23:** Integrated Memory Controller 1 configuration registers. For SKUs with one IMC, these devices are not used. For SKUs with two IMC, this IMC supports 2 channels (2,3) and device 22 Functions 4, 5 (channel 2,3) are not used and should be ignored as they apply to unimplemented channels.
 - The Intel® Xeon® Processor E5-1600 v4 product family and some SKUs of the Intel® Xeon® Processor E5-2600 v4 product family implement 1 IMC/HA. For SKUs which implement 1 IMC/HA these devices are not used. For SKUs which implement 2 IMC/HA channels 2 and 3 reside off of these devices.
- **Device 30:** Processor Power Control Unit. Device 30 contain the PCU registers.

1.2 Configuration Register Rules

The Intel® Xeon® Processor E5 v4 product family supports the following configuration register types:


- PCI Configuration Registers (CSRs): CSRs are chipset specific registers that are located at PCI defined address space.
- Machine Specific Registers (MSRs): MSRs are machine specific registers that can be accessed by specific read and write instructions. MSRs are OS ring 0 and BIOS accessible, though some can only be accessed in certain modes (that is, SMM mode).
- Memory-mapped I/O registers: These registers are mapped into the system memory map as MMIO low or MMIO high. They are accessed by any code typically an OS driver running on the platform. This register space is introduced with the integration of some of the chipset functionality.

1.2.1 CSR Access

Configuration space registers are accessed via the well known configuration transaction mechanism defined in the PCI specification and this uses the bus:device:function number concept to address a specific device's configuration space. If initiated by a remote CPU, accesses to PCI configuration registers are achieved via NcCfgRd/Wr transactions on Intel QPI.

All configuration register accesses are accessed over Message Channel through the Ubox but might come from a variety of different sources:

- Local cores
- Remote cores (over Intel QPI)



Configuration registers can be read or written in Byte, WORD (16-bit), or DWORD (32-bit) quantities. *Accesses larger than a DWORD to PCI Express configuration space will result in unexpected behavior.* All multi-byte numeric fields use “little-endian” ordering (that is, lower addresses contain the least significant parts of the field).

1.2.1.1 PCI Bus Number

In the tables shown for IIO devices (0 - 7), the PCI Bus numbers are all marked as “Bus 0”. This means that the actual bus number is variable depending on which socket is used. The specific bus number for all PCIe devices in the Intel® Xeon® Processor E5 v4 product family is specified in the CPUBUSNO register which exists in the I/O module’s configuration space. Bus number is derived by the max bus range setting and processor socket number.

1.2.1.2 Uncore Bus Number

In the tables shown for Uncore devices (8 - 31), the PCI Bus numbers are all marked as “bus 1”. This means that the actual bus number is CPUBUSNO(1) where CPUBUSNO(1) is programmable by BIOS depending on which socket is used. The specific bus number for all PCIe devices in the Intel® Xeon® Processor E5 v4 product family is specified in the CPUBUSNO register.

1.2.1.3 Device Mapping

Each component in the processor is uniquely identified by a PCI bus address consisting of Bus Number, Device Number and Function Number. Device configuration is based on the PCI Type 0 configuration conventions. All processor registers appear on the PCI bus assigned for the processor socket. Bus number is derived by the max bus range setting and processor socket number.

Table 1-1. Functions Specifically Handled by the Processor (Sheet 1 of 3)

Register Group	DID	Device	Function	Comment
DMI2	2F00h	0	0	x4 link from Processor to PCH
PCI Express* Root Port in DMI2 Mode	2F01h	0	0	Device 0 operating as a x4 PCI Express Port instead of a link to the PCH
PCI Express Root Port 2	2F04h, 2F05h, 2F06h, 2F07h	2	0-3	PCIe* Device 2 Root Ports x16, x8 or x4 max link width
PCI Express Root Port 3	2F08, 2F09h, 2FOAh, 2F0Bh	3	0-3	PCIe Device 3 Root Ports x16, x8 or x4 max link width
IIO	2F28h	5	0	Address Map, Intel VT-d, System Management
IIO	2F2Ah	5	2	RAS, Control Status and Global Errors
IIO	2F2Ch	5	4	I/O APIC
Intel QuickData Technology	2F20h, 2F21h, 2F22h, 2F23h, 2F24h, 2F25h, 2F26h, 2F27h	4	0-7	DMA Channel 0 to Channel 7
Intel QPI Link	2F80h	8	0	Intel QPI Link 0





Table 1-1. Functions Specifically Handled by the Processor (Sheet 2 of 3)

Register Group	DID	Device	Function	Comment
Intel QPI Link	2F90h	9	0	Intel QPI Link 1
PCU	2F98h, 2F99h, 2F9Ah 2FC0h 2F9Ch	30	0-4	Power Control Unit
UBOX	2F1Eh	16	5	Scratchpad and Semaphores
UBOX	2F7Dh	16	6	Scratchpad and Semaphores
UBOX	2F1F	16	7	Scratchpad and Semaphores
Integrated Memory Controller 0	2FA8h	19	0	IMC Main
Integrated Memory Controller 0	2F71h	19	1	IMC RAS Registers
Integrated Memory Controller 0	2FAAh, 2FABh,	19	2-3	IMC Channel 0-1 Target Address Decoder Registers
Integrated Memory Controller 0	2FACh, 2FADh	19	4-5	IMC Channel 2-3 Target Address Decoder Registers
Integrated Memory Controller 0	2FB4, 2FB5	20	0,1	IMC Channel 0-1 Registers
Integrated Memory Controller 0	2FBAh, 2FBBh, 2FBEh, 2FBFh	17	2,3,6,7	For 1 HA: DDRIO 0,1,2,3 Multicast
Integrated Memory Controller 0	2FBCh, 2FBDh	17	4,5	For 2 HA: DDRIO 0 & 1
Integrated Memory Controller 0	2FBEh, 2FBFh	17	6,7	For 2 HA: DDRIO 0,1,Multicast
Integrated Memory Controller 0	2FD8h, 2FD9h, 2FDAh, 2FDBh,	31	0-3	For 2 HA only: DDRIO 2 & 3
Integrated Memory Controller 0	2FB6, 2FB7	20	2,3	IMC Channel 0-1 Registers
Integrated Memory Controller 1	2FB0, 2FB1,	21	0,1	IMC Channel 2-3 Registers
Integrated Memory Controller 1	2FB2, 2FB3,	21	2,3	IMC Channel 2-3 Registers
Integrated Memory Controller 1	2F68h	22	0	IMC Main
Integrated Memory Controller 1	2F79h	22	1	IMC RAS Registers
Integrated Memory Controller 1	2F6Ah, 2F6Bh,	22	2-3	IMC Channel 0-1 Target Address Decoder Registers
Integrated Memory Controller 1	2FD4, 2FD5	23	0,1	IMC Channel 0-1 Registers
Integrated Memory Controller 0	2FBAh, 2FBBh, 2FBEh, 2FBFh	17	2,3,6,7	For 1 HA: DDRIO 0,1,2,3 Multicast
Integrated Memory Controller 0	2FBCh, 2FBDh	17	4,5	For 2 HA: DDRIO 0 & 1
Integrated Memory Controller 0	2FBEh, 2FBFh	17	6,7	For 2 HA: DDRIO 0,1,Multicast
Integrated Memory Controller 0	2FD8h, 2FD9h, 2FDAh, 2FDBh,	31	0-3	For 2 HA only: DDRIO 2 & 3


Table 1-1. Functions Specifically Handled by the Processor (Sheet 3 of 3)

Register Group	DID	Device	Function	Comment
Integrated Memory Controller 1	2FD6, 2FD7,	23	0,1	IMC Channel 0-1 Registers
R2PCIe	2F1Dh	16	0	Integrated IO Ring Interface
R2PCIe	2F34h	16	1	PCI Express Ring Performance Monitoring
R3QPI	2F81h,	11	0	Intel QPI Ring Interface
R3QPI	2F36h, 2F37h	11	1,2	Intel QPI Ring Performance Monitoring

1.2.1.4 Unimplemented Devices/Functions and Registers

Configuration reads to unimplemented functions and devices will return all ones emulating a master abort response. Note that there is no asynchronous error reporting that happens when a configuration read master aborts. Configuration writes to unimplemented functions and devices will return a normal response.

Software should not attempt or rely on reads or writes to unimplemented registers or register bits. Unimplemented registers should return all zeros when read. Writes to unimplemented registers are ignored. For configuration writes to these register (require a completion), the completion is returned with a normal completion status (not master-aborted).

1.2.1.5 Device Hiding

The Intel® Xeon® Processor E5 v4 product family provides a mechanism by which various PCI devices or functions within the unit can be hidden from the host configuration software; that is, all PCI configuration accesses to the devices' configuration space from Intel QPI will be master aborted. This mechanism is needed in cases where a device or function is not used or is available for use, because either the device is turned off or the device is not serving any meaningful purpose in a given platform configuration. This hiding mechanism is implemented via the DEVHIDE register.

1.2.2 MSR Access

Machine specific registers are architectural and only accessed by using specific ReadMSR/WriteMSR instructions. MSRs are always accessed as a naturally aligned 4 or 8 byte quantity.

For common IA-32 architectural MSRs, please refer to the *Intel® 64 and IA-32 Software Developer's Manual*.

1.2.3 Memory-Mapped I/O Registers

The PCI standard provides not only configuration space registers but also registers which reside in memory-mapped space. For PCI devices, this is typically where the majority of the driver programming occurs and the specific register definitions and characteristics are provided by the device manufacturer. Access to these registers are typically accomplished via CPU reads and writes to non-coherent (UC) or write-combining (WC) space.

Reads and writes to memory-mapped registers can be accomplished with 1, 2, 4 or 8 byte transactions.




1.3 Register Terminology

The bits in configuration register descriptions will have an assigned attribute from the following table. Bits without a Sticky attribute are set to their default value by a hard reset.

Table 1-2. Register Attributes Definitions

Attr	Description
RO	Read Only: These bits can only be read by software, writes have no effect. The value of the bits is determined by the hardware only.
RW	Read / Write: These bits can be read and written by software.
RC	Read Clear Variant: These bits can be read by software, and the act of reading them automatically clears them. HW is responsible for writing these bits, and therefore the -V modifier is implied.
W1S	Write 1 to Set : Writing a 1 to these bits will set them to 1. Writing 0 will have no effect. Reading will return indeterminate values and read ports are not required on the register.
WO	Write Only: These bits can only be written, reads return indeterminate values.
RW-O	Read / Write Once: These bits can be read by software. After reset, these bits can only be written by software once, after which the bits becomes 'Read Only'.
RW-L	Read / Write Lock: These bits can be read and written by software. Hardware can make these bits 'Read Only' via a separate configuration bit or other logic.
RW1C	Read / Write 1 to Clear: These bits can be read and cleared by software. Writing a '1' to a bit clears it, while writing a '0' to a bit has no effect.
ROS	RO Sticky: These bits can only be read by software, writes have no effect. The value of the bits is determined by the hardware only. These bits are only re-initialized to their default value by a PWRGOOD reset.
RW1S	Read, Write 1 to Set: These bits can be read. Writing a 1 to a given bit will set it to 1. Writing a 0 to a given bit will have no effect. It is not possible for software to set a bit to "0". The 1->0 transition can only be performed by hardware. These registers are implicitly -V.
RWS	R / W Sticky: These bits can be read and written by software. These bits are only re-initialized to their default value by a PWRGOOD reset.
RW1CS	R / W1C Sticky: These bits can be read and cleared by software. Writing a '1' to a bit clears it, while writing a '0' to a bit has no effect. These bits are only re-initialized to their default value by a PWRGOOD reset.
RW-LB	Read/Write Lock Bypass: Similar to RWL, these bits can be read and written by software. HW can make these bits "Read Only" via a separate configuration bit or other logic. However, RW-LB is a special case where the locking is controlled by the lock-bypass capability that is controlled by the lock-bypass enable bits. Each lock-bypass enable bit enables a set of config request sources that can bypass the lock. The requests sourced from the corresponding bypass enable bits will be lock-bypassed (i.e. RW).
RO-FW	Read Only Forced Write: These bits are read only from the perspective of the cores.
RWS-O	R / W Sticky Once: If a register is both sticky and "once" then the sticky value applies to both the register value and the "once" characteristic. Only a PWRGOOD reset will reset both the value and the "once" so that the register can be written to again.
RW-V	R / W Volatile: These bits may be modified by hardware. Typically, this occurs based on values from hardware configuration straps for functions such as DMI2 and PCIe I/O configuration. They also could be changed based on status or modes within internal state machines. Software cannot expect the values to stay unchanged.
RWS-L	R / W Sticky Locked: If a register is both sticky and locked, then the sticky behavior only applies to the value. The sticky behavior of the lock is determined by the register that controls the lock.
RV, RSVD	Reserved: These bits are reserved for future expansion and their value must not be modified by software. When writing these bits, software must preserve the value read.




1.4 Protected Processor Inventory Number

Protected Processor Inventory Number (PPIN) is a solution for inventory management available on Intel® Xeon® processor E5 v4 product family for use in server platforms.







2 Integrated Memory Controller (iMC) Configuration Registers

The Integrated Memory Controller registers are listed below and are specific to the Intel® Xeon® Processor E5 v4 product family.

- The Intel® Xeon® Processor E5-1600 v4 product family implement 1 Memory Controller. The Intel® Xeon® Processor E5-2600 v4 product family implement either 1 Memory Controller or two. The Intel® Xeon® Processor E5-4600 v4 product family implement 2 Memory Controllers.
 - For SKUs with one IMC Device 22 and Device 23 (IMC 1) are not used and should be ignored, as there are 4 channels on Device IMC 0.
 - For SKUs with two IMC, each IMC has 2 channels. Device 19,22 Functions 4, 5 (channel 2,3) and device 21 are not used and should be ignored.

For Device 19 and 22 Functions 0-5 for offsets >= 256, PCIe extended configuration space are not designed for direct usage by OS or device drivers, and may not be accessible directly by OS components such as device drivers. The PCI Capability Pointer Register (CAPPTR) is set to a value of 00h. BIOS/firmware and/or BMC can access these registers, combine the information obtained with system implementation specifics, and if required, make it available to the OS through firmware and/or BMC interfaces.

2.1 Device 19,22 Function 0

	100h	SMB_STAT_0	180h
MH_MAINCNTL	104h	SMBCMD_0	184h
	108h	SMBCntL_0	188h
MH_SENSE_500NS_CFG	10Ch	SMB_TSOD_POLL_RATE_CNTR_0	18Ch
MH_DTYCYC_MIN_ASRT_CNTR_0	110h	SMB_STAT_1	190h
MH_DTYCYC_MIN_ASRT_CNTR_1	114h	SMBCMD_1	194h
MH_IO_500NS_CNTR	118h	SMBCntL_1	198h
MH_CHN_ASTN	11Ch	SMB_TSOD_POLL_RATE_CNTR_1	19Ch
	120h	SMB_PERIOD_CFG	1A0h
MH_EXT_STAT	124h	SMB_PERIOD_CNTR	1A4h
	128h	SMB_TSOD_POLL_RATE	1A8h
	12Ch		1ACh
	130h		1B0h
	134h		1B4h
	138h		1B8h
	13Ch		1BCh
	140h		1C0h
	144h		1C4h
	148h		1C8h
	14Ch		1CCh



Integrated Memory Controller (iMC) Configuration Registers

150h	1D0h
154h	1D4h
158h	1D8h
15Ch	1DCh
160h	1E0h
164h	1E4h
168h	1E8h
16Ch	1ECh
170h	1F0h
174h	1F4h
178h	1F8h
17Ch	1FCh

2.1.1 pxpcap

PCI Express Capability.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x40		Function: 0	
Bit	Attr	Default	Description
29:25	RO	0x0	Interrupt Message Number (interrupt_message_number): N/A for this device
24:24	RO	0x0	Slot Implemented (slot_implemented): N/A for integrated endpoints
23:20	RO	0x9	Device/Port Type (device_port_type): Device type is Root Complex Integrated Endpoint
19:16	RO	0x1	Capability Version (capability_version): PCI Express Capability is Compliant with Version 1.0 of the PCI Express Spec. Note: This capability structure is not compliant with Versions beyond 1.0, since they require additional capability registers to be reserved. The only purpose for this capability structure is to make enhanced configuration space available. Minimizing the size of this structure is accomplished by reporting version 1.0 compliance and reporting that this is an integrated root port device. As such, only three Dwords of configuration space are required for this structure.
15:8	RO	0x0	Next Capability Pointer (next_ptr): Pointer to the next capability. Set to 0 to indicate there are no more capability structures.
7:0	RO	0x10	Capability ID (capability_id): Provides the PCI Express capability ID assigned by PCI-SIG.



2.1.2 mcmtr

Memory Technology

Bit	Attr	Default	Description
Type: CFG PortID: N/A Bus: 1 Device: 19,22 Function: 0 Offset: 0x7c			
21:18	RW_LB	0x0	CHN_DISABLE(chn_disable): Channel disable control. When set, the corresponding channel is disabled.
17:16	RW_LB	0x0	pass76(pass76): 00: do not alter ChnAdd calculation 01: replace ChnAdd[6] with SysAdd[6] 10: Reserved 11: replace ChnAdd[7:6] with SysAdd[7:6]
14	RW_LB	0x0	ddr4 (ddr4): DDR4 mode
13:12	RW_LB	0x0	IMC_MODE (imc_mode): Memory mode: 00: Native DDR All others reserved.
8:8	RW_LB	0x0	NORMAL (normal): 0: Training mode 1: Normal Mode
3:3	RW_LBV	0x0	DIR_EN (dir_en): If the directory disabled in SKU, this register bit is set to Read-Only (RO) with 0 value, that is, the directory is disabled. When this bit is set to zero, IMC ECC code uses the non-directory CRC-16. If the SKU supports directory and enabled, that is, the directory is not disabled, the DIR_EN bit can be set by BIOS, MC ECC uses CRC-15 in the first 32B code word to yield one directory bit. It is important to know that changing this bit will require BIOS to re-initialize the memory.
2:2	RW_LBV	0x0	ECC_EN (ecc_en): ECC enable. DISECC will force override this bit to 0.
1:1	RW_LBV	0x0	LS_EN (ls_en): Use lock-step channel mode if set; otherwise, independent channel mode. This field should only be set for native DDR lockstep.
0:0	RW_LB	0x0	CLOSE_PG (close_pg): Use close page address mapping if set; otherwise, open page.

2.1.3 tadwayness_[0:11]

TAD Range Wayness, Limit and Target.


There are total of 12 TAD ranges ($N + P + 1 =$ number of TAD ranges; $P =$ how many times channel interleave changes within the SAD ranges.).

Note for mirroring configuration:

For 1-way interleave, channel 0-2 mirror pair: target list <0,2,x,x>, TAD ways = "00"

For 1-way interleave, channel 1-3 mirror pair: target list <1,3,x,x>, TAD ways = "00"

For 2-way interleave, 0-2 mirror pair and 1-3 mirror pair: target list <0,1,2,3>, TAD ways = "01"



Integrated Memory Controller (iMC) Configuration Registers

For 1-way interleave, lockstep mirroring, target list <0,2,x,x>, TAD ways = "00"

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x80, 0x84, 0x88, 0x8c, 0x90, 0x94, 0x98, 0x9c, 0xa0, 0xa4, 0xa8, 0xac		Function: 0	
Bit	Attr	Default	Description
31:12	RW_LB	0x0	TAD_LIMIT (tad_limit): Highest address of the range in system address space, 64MB granularity, i.e. TADRANGLIMIT[45:26].
11:10	RW_LB	0x0	TAD_SKT_WAY (tad_skt_way): socket interleave wayness 00 = 1 way, 01 = 2 way, 10 = 4 way, 11 = 8 way.
9:8	RW_LB	0x0	TAD_CH_WAY (tad_ch_way): Channel interleave wayness 00 - interleave across 1 channel or mirror pair 01 - interleave across 2 channels or mirror pairs 10 - interleave across 3 channels 11 - interleave across 4 channels This parameter effectively tells iMC how much to divide the system address by when adjusting for the channel interleave. Since both channels in a pair store every line of data, divide by 1 when interleaving across one pair and 2 when interleaving across two pairs. For HA, it tells how many channels to distribute the read requests across. When interleaving across 1 pair, this distributes the reads to two channels, when interleaving across 2 pairs, this distributes the reads across 4 pairs. Writes always go to both channels in the pair when the read target is either channel.
7:6	RW_LB	0x0	TAD_CH_TGT3 (tad_ch_tgt3): Target channel for channel interleave 3 (used for 4-way TAD interleaving). This register is used in the iMC only for reverse address translation for logging sparepatrol errors, converting a rank address back to a system address.
5:4	RW_LB	0x0	TAD_CH_TGT2 (tad_ch_tgt2): Target channel for channel interleave 2 (used for 3/4-way TAD interleaving).
3:2	RW_LB	0x0	TAD_CH_TGT1 (tad_ch_tgt1): Target channel for channel interleave 1 (used for 2/3/4-way TAD interleaving).
1:0	RW_LB	0x0	TAD_CH_TGT0 (tad_ch_tgt0): Target channel for channel interleave 0 (used for 1/2/3/4-way TAD interleaving).

2.1.4 mc_init_state_g

Initialization state for boot and training.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0xb4		Function: 0	
Bit	Attr	Default	Description
12:9	RWS_L	0x0	cs_oe_en:



Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0xb4		Function: 0	
Bit	Attr	Default	Description
8:8	RWS_L	0x1	MC is in SR (safe_sr): This bit indicates if it is safe to keep the MC in self refresh (SR) during MC-reset. If it is clear when reset occurs, it means that the reset is without warning and the DDR-reset should be asserted. If set when reset occurs, it indicates that DDR is already in SR and it can keep it this way. This bit can also indicate MRC if reset without warning has occurred, and if it has, cold-reset flow should be selected. BIOS need to clear this bit at MRC entry.
7:7	RW_L	0x0	MRC_DONE (mrc_done): This bit indicates the PCU that the MRC is done, IMC is in normal mode, ready to serve. MRC should set this bit when MRC is done, but it doesn't need to wait until training results are saved in BIOS flash.
5:5	RW_L	0x1	DDRIO Reset (reset_io): Training Reset for DDRIO. Make sure this bit is cleared before enabling DDRIO.
3:3	RW_L	0x0	Refresh Enable (refresh_enable): If cold reset, this bit should be set by BIOS after: 1) Initializing the refresh timing parameters 2) Running DDR through reset ad init sequence. If warm reset or S3 exit, this bit should be set immediately after SR exit.
2:2	RW_L	0x0	DCLK Enable (for all channels) (dclk_enable):
1:1	RW_L	0x1	DDR_RESET (ddr_reset): DIMM reset. Controls all channels.

2.1.5 rcomp_timer

RCOMP wait timer. Defines the time from IO starting to run RCOMP evaluation until RCOMP results are definitely ready. This counter is added in order to keep determinism of the process if operated in different mode. This register also indicates that first RCOMP has been done.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0xc0		Function: 0	
Bit	Attr	Default	Description
31:31	RW_V	0x0	rcomp_in_progress: RCOMP in progress status bit
30:30	RW	0x0	rcomp: RCOMP start via message channel control for BIOS. RCOMP start only triggered when the register bit output is changing from 0 -> 1. iMC is not be responsible for clearing this bit. When Rcomp is done via first_rcomp_done bit field.
21:21	RW	0x0	ignore_mdll_locked_bit Ignore DDRIO MDLL lock status during rcomp when set.
20:20	RW	0x0	no_mdll_fsm_override: Do not force DDRIO MDLL on during rcomp when set.



Integrated Memory Controller (iMC) Configuration Registers

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0xc0		Function: 0	
Bit	Attr	Default	Description
16:16	RW_LV	0x0	First RCOMP has been done in DDRIO (first_rcomp_done): This is a status bit that indicates the first RCOMP has been completed. It is cleared on reset, and set by IMC HW when the first RCOMP is completed. BIOS should wait until this bit is set before executing any DDR command.
15:0	RW	0xc00	COUNT (count): DCLK cycle count that IMC needs to wait from the point it has triggered RCOMP evaluation until it can trigger the load to registers.

2.1.6 mh_sense_500ns_cfg


MEMHOT Sense and 500 ns Config.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x10c		Function: 0	
Bit	Attr	Default	Description
25:16	RW	0xc8	MH_SENSE_PERIOD (mh_sense_period): MEMHOT Input Sense Period in number of CNTR_500_NANOSEC. BIOS calculate number of CNTR_500_NANOSEC for 50 micro-sec / 100 micro-sec / 200 micro-sec / 400 micro-sec.
15:13	RW	0x2	MH_IN_SENSE_ASSERT (mh_in_sense_assert): MEMHOT Input Sense Assertion Time in number of CNTR_500_NANOSEC. BIOS calculate number of CNFG_500_NANOSEC for 1 micro-sec / 2 micro-sec inputsense duration. MH_IN_SENSE_ASSERT ranges: 0 or 1: Reserved 2 - 7: 1 micro-sec - 3.5 micro-sec sense assertion time in 500 ns increment.
9:0	RW-LS	0x190	CNFG_500_NANOSEC (cnfg_500_nanosec): 500 ns equivalent in DCLK. BIOS calculate number of DCLK to be equivalent to 500 ns. This value is loaded into CNTR_500_NANOSEC when it is decremented to zero.

2.1.7 mh_dtycyc_min_asrt_cntr_[0:1]

MEMHOT Duty Cycle Period and Min Assertion Counter.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x110, 0x114		Function: 0	
Bit	Attr	Default	Description
31:20	RO_V	0x0	MH_MIN_ASRTN_CNTR (mh_min_asrtn_cntr): MEM_HOT[1:0]# Minimum Assertion Time Current Count in number of CNTR_500_NANOSEC decrement by 1 every CNTR_500_NANOSEC. When the counter is zero, the counter is remain at zero and it is only loaded with MH_MIN_ASRTN only when MH_DUTY_CYC_PRD_CNTR is reloaded.
19:0	RW_LV	0x0	MH_DUTY_CYC_PRD_CNTR (mh_duty_cyc_prd_cntr): MEM_HOT[1:0]# DUTY Cycle Period Current Count in number of CNTR_500_NANOSEC decrement by 1 every CNTR_500_NANOSEC. When the counter is zero, the next cycle is loaded with MH_DUTY_CYC_PRD.



2.1.8 mh_io_500ns_cntr


MEMHOT Input Output and 500 ns Counter.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x118		Function: 0	
Bit	Attr	Default	Description
31:22	RW_LV	0x0	MH1_IO_CNTR (mh1_io_cntr): MEM_HOT[1:0]# Input Output Counter in number of CNTR_500_NANOSEC. When MHO_IO_CNTR is zero, the counter is loaded with MH_SENSE_PERIOD in the next CNTR_500_NANOSEC. When count is greater than MH_IN_SENSE_ASSERT, the MEM_HOT1# output driver may be turn on if the corresponding MEM_HOT#event is asserted. The receiver is turned off during this time. When count is equal or less than MH_IN_SENSE_ASSERT, MEM_HOT[1:0]# output is disabled and receiver is turned on. Hardware decrements this counter by 1 every time CNTR_500_NANOSEC is decremented to zero. When the counter is zero, the next CNTR_500_NANOSEC count is loaded with MH_IN_SENSE_ASSERT.
21:12	RW_LV	0x0	MHO_IO_CNTR (mh0_io_cntr): MEM_HOT[1:0]# Input Output Counter in number of CNTR_500_NANOSEC. When MH_IO_CNTR is zero, the counter is loaded with MH_SENSE_PERIOD in the next CNTR_500_NANOSEC. When count is greater than MH_IN_SENSE_ASSERT, the MEM_HOT[1:0]# output driver may be turn on if the corresponding MEM_HOT#event is asserted. The receiver is turned off during this time. When count is equal or less than MH_IN_SENSE_ASSERT, MEM_HOT[1:0]# output is disabled and receiver is turned on. BIOS calculate number of CNTR_500_NANOSEC hardware will decrement this register by 1 every CNTR_500_NANOSEC. When the counter is zero, the next CNTR_500_NANOSEC count is loaded with MH_IN_SENSE_ASSERT.
9:0	RW_LV	0x0	CNTR_500_NANOSEC (cntr_500_nanosec): 500 ns base counters used for the MEMHOT counters and the SMBus counters. BIOS calculate number of DCLK to be equivalent to 500 nanoseconds. CNTR_500_NANOSEC hardware will decrement this register by 1 every CNTR_500_NANOSEC. When the counter is zero, the next CNTR_500_NANOSEC count is loaded with CNTR_500_NANOSEC.

2.1.9 mh_chn_astn

MEMHOT Domain Channel Association.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x11c		Function: 0	
Bit	Attr	Default	Description
23:20	RO	0xb	MH1_2ND_CHN_ASTN (mh1_2nd_chn_astn): MemHot[1]# 2nd Channel Association bit 23: is valid bit. Note: Valid bit means the association is valid and it does not implies the channel is populated. Bit 22-20: 2nd channel ID within this MEMHOT domain.
19:16	RO	0xa	MH1_1ST_CHN_ASTN (mh1_1st_chn_astn): MemHot[1]# 1st Channel Association bit 19: is valid bit. Note: Valid bit means the association is valid and it does not implies the channel is populated. Bit 18-16: 1st channel ID within this MEMHOT domain.



Integrated Memory Controller (iMC) Configuration Registers

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x11c		Function: 0	
Bit	Attr	Default	Description
7:4	RO	0x9	MHO_2ND_CHN_ASTN (mh0_2nd_chn_astn): MemHot[0]# 2nd Channel Association bit 7: is valid bit. Note: Valid bit means the association is valid and it does not implies the channel is populated. Bit 6-4: 2nd channel ID within this MEMHOT domain.
3:0	RO	0x8	MHO_1ST_CHN_ASTN (mh0_1st_chn_astn): MemHot[0]# 1st Channel Association bit 3: is valid bit. Note: Valid bit means the association is valid and it does not implies the channel is populated or exist. Bit 2-0: 1st channel ID within this MEMHOT domain.

2.1.10 mh_ext_stat

Capture externally asserted MEM_HOT[1:0]# assertion detection.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x124		Function: 0	
Bit	Attr	Default	Description
1:1	RW1C	0x0	MH_EXT_STAT_1 (mh_ext_stat_1): MEM_HOT[1]# assertion status at this sense period. Set if MEM_HOT[1]# is asserted externally for this sense period, this running status bit will automatically updated with the next sensed value in the next MEMHOT input sense phase.
0:0	RW1C	0x0	MH_EXT_STAT_0 (mh_ext_stat_0): MEM_HOT[0]# assertion status at this sense period. Set if MEM_HOT[0]# is asserted externally for this sense period, this running status bit will automatically updated with the next sensed value in the next MEMHOT input sense phase.

2.1.11 smb_stat_[0:1]

SMBus Status. This register provides the interface to the SMBus/I2C* SCL and SDA signals that is used to access the Serial Presence Detect EEPROM (SPD) or Thermal Sensor on DIMM (TSOD) that defines the technology, configuration, and speed of the DIMMs controlled by iMC.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x180		Function: 0	
Bit	Attr	Default	Description
31:31	RO_V	0x0	SMB_RDO (smb_rdo): Read Data Valid This bit is set by iMC when the Data field of this register receives read data from the SPD/TSOD after completion of an SMBus read command. It is cleared by iMC when a subsequent SMBus read command is issued.



Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x180		Function: 0	
Bit	Attr	Default	Description
30:30	RO_V	0x0	<p>SMB_WOD (smb_wod): Write Operation Done</p> <p>This bit is set by iMC when a SMBus Write command has been completed on the SMBus. It is cleared by iMC when a subsequent SMBus Write command is issued.</p>
29:29	RO_V	0x0	<p>SMB_SBE (smb_sbe): SMBus Error</p> <p>This bit is set by iMC if an SMBus transaction (including the TSOD polling or message channel initiated SMBus access) that does not complete successfully (non-Ack has been received from slave at expected Ack slot of the transfer). If a slave device is asserting clock stretching, iMC does not have logic to detect this condition to set the SBE bit directly; however, the SMBus master will detect the error at the corresponding transaction's expected ACK slot.</p> <p>Once SMBUS_SBE bit is set, iMC stops issuing hardware initiated TSOD polling SMBUS transactions until the SMB_SBE is cleared. iMC will not increment the SMB_STAT_x.TSOD_SA until the SMB_SBE is cleared. Manual SMBus command interface is not affected, that is, new command issue will clear the SMB_SBE like AO silicon behavior.</p>
28:28	ROS_V	0x0	<p>SMB_BUSY (smb_busy): SMBus Busy state. This bit is set by iMC while an SMBus/I2C command (including TSOD command issued from iMC hardware) is executing. Any transaction that is completed normally or gracefully will clear this bit automatically. By setting the SMB_SOFT_RST will also clear this bit.</p> <p>This register bit is sticky across reset so any surprise reset during pending SMBus operation will sustain the bit assertion across surprised warm-reset. BIOS reset handler can read this bit before issuing any SMBus transaction to determine whether a slave device may need special care to force the slave to idle state (for example, via clock override toggling SMB_CKOV RD and/or via induced time-out by asserting SMB_CKOV RD for 25-35 ms).</p>
27:24	RO_V	0x7	<p>Last Issued TSOD Slave Address (tsod_sa): This field captures the last issued TSOD slave address. Here is the slave address and the DDR CHN and DIMM slot mapping:</p> <p>Slave Address: 0 -- Channel: Even Chn; Slot #: 0 Slave Address: 1 -- Channel: Even Chn; Slot #: 1 Slave Address: 2 -- Channel: Even Chn; Slot #: 2 Slave Address: 3 -- Channel: Even Chn; Slot #: 3 (reserved) Slave Address: 4 -- Channel: Odd Chn; Slot #: 0 Slave Address: 5 -- Channel: Odd Chn; Slot #: 1 Slave Address: 6 -- Channel: Odd Chn; Slot #: 2 Slave Address: 7 -- Channel: Odd Chn; Slot #: 3 (reserved)</p> <p>Since this field only captures the TSOD polling slave address. During SMB error handling, software should check the hung SMB_TSOD_POLL_EN state before disabling the SMB_TSOD_POLL_EN in order to qualify whether this field is valid.</p>
15:0	RO_V	0x0	<p>SMB_RDATA (smb_rdata): Read DataHolds data read from SMBus Read commands.</p> <p>Since TSOD/EEPROM are I2C* devices and the byte order is MSByte first in a word read, reading of I2C using word read should return SMB_RDATA[15:8] = I2C_MSB and SMB_RDATA[7:0] = I2C_LSB. If reading of I2C using byte read, the SMB_RDATA[15:8] = dont care; SMB_RDATA[7:0] = readbyte.</p> <p>If there is a SMB slave connected on the bus, reading of the SMBus slave using word read returns SMB_RDATA[15:8] = SMB_LSB and SMB_RDATA[7:0] = SMB_MSB.</p> <p>If the software is not sure whether the target is I2C or SMBus slave, please use byte access.</p>



2.1.12 smbcmd_[0:1]

A write to this register initiates a DIMM EEPROM access through the SMBus/I2C.


Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x184,		Function: 0	
Bit	Attr	Default	Description
31:31	RW_V	0x0	SMB_CMD_TRIGGER (smb_cmd_trigger): CMD trigger: After setting this bit to 1, the SMBus master will issue the SMBus command using the other fields written in SMBCMD_[0:1] and SMBCntI_[0:1]. Note: The '-V' in the attribute implies the hardware will reset this bit when the SMBus command is being started.
30:30	RWS	0x0	SMB_PNTR_SEL (smb_pntr_sel): Pointer Selection: SMBus/I2C present pointer-based access enable when set; otherwise, use random access protocol. Hardware based TSOD polling will also use this bit to enable the pointer word read. Important Note: CPU hardware-based TSOD polling can be configured with pointer based access. If software manually issue SMBus transaction to other address, i.e. changing the pointer in the slave device, it is software's responsibility to restore the pointer in each TSOD before returning to hardware based TSOD polling while keeping the SMB_PNTR_SEL = 1.
29:29	RWS	0x0	SMB_WORD_ACCESS (smb_word_access): Word access: SMBus/I2C word 2B access when set; otherwise, it is a byte access.
28:28	RWS	0x0	SMB_WRT_PNTR (smb_wrt_pntr): Bit[28:27] = 00: SMBus Read Bit[28:27] = 01: SMBus Write Bit[28:27] = 10: illegal combination Bit[28:27] = 11: Write to pointer register SMBus/I2C pointer update (byte). bit 30, and 29 are ignored. Note: SMBCntI_[0:1] [26] will NOT disable WrtPntr update command.
27:27	RWS	0x0	SMB_WRT_CMD (smb_wrt_cmd): When '0', it's a read command When '1', it's a write command
26:24	RWS	0x0	SMB_SA (smb_sa): Slave Address: This field identifies the DIMM SPD/TSOD to be accessed.
23:16	RWS	0x0	SMB_BA (smb_ba): Bus Txn Address: This field identifies the bus transaction address to be accessed. Note: In WORD access, 23:16 specifies 2B access address. In Byte access, 23:16 specified 1B access address.
15:0	RWS	0x0	SMB_WDATA (smb_wdata): Write Data: Holds data to be written by SPDW commands. Since TSOD/EEPROM are I2C devices and the byte order is MSByte first in a word write, writing of I2C using word write should use SMB_WDATA[15:8] = I2C_MSB and SMB_WDATA[7:0] = I2C_LSB. If writing of I2C using byte write, the SMB_WDATA[15:8] = dont care; SMB_WDATA[7:0] = writebyte. If we have a SMB slave connected on the bus, writing of the SMBus slave using word write should use SMB_WDATA[15:8] = SMB_LSB and SMB_WDATA[7:0] = SMB_MSB. It is software responsibility to figure out the byte order of the slave access.



2.1.13 smbcntl_[0:1]

SMBus Control.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x188,		Function: 0	
Bit	Attr	Default	Description
31:28	RWS	0xa	<p>SMB_DTI (smb_dti):</p> <p>Device Type Identifier: This field specifies the device type identifier. Only devices with this device-type will respond to commands.</p> <p>'0011' specifies TSOD.</p> <p>'1010' specifies EEPROM's.</p> <p>'0110' specifies a write-protect operation for an EEPROM.</p> <p>Other identifiers can be specified to target non-EEPROM devices on the SMBus.</p> <p>Note: iMC based hardware TSOD polling uses hardcoded DTI. Changing this field has no effect on the hardware based TSOD polling.</p>
27:27	RWS_V	0x1	<p>SMB_CKOVRD (smb_ckovrd):</p> <p>Clock Override</p> <p>'0' Clock signal is driven low, overriding writing a '1' to CMD.</p> <p>'1' Clock signal is released high, allowing normal operation of CMD.</p> <p>Toggleing this bit can be used to 'budge' the port out of a 'stuck' state.</p> <p>Software can write this bit to 0 and the SMB_SOFT_RST to 1 to force hung SMBus controller and the SMB slaves to idle state without using power good reset or warm reset.</p> <p>Note: Software need to set the SMB_CKOVRD back to 1 after 35ms in order to force slave devices to time-out in case there is any pending transaction. The corresponding SMB_STAT_x.SMB_SBE error status bit may be set if there was such pending transaction time-out (non-graceful termination). If the pending transaction was a write operation, the slave device content may be corrupted by this clock override operation. A subsequent SMB command will automatically cleared the SMB_SBE.</p> <p>iMC added SMBus time-out control timer in B0. When the time-out control timer expired, the SMBCKOVRD# will "de-assert", i.e. return to 1 value and clear the SMBSBE0.</p>
26:26	RW_LB	0x1	<p>SMB_DIS_WRT (smb_dis_wrt):</p> <p>Disable SMBus Write</p> <p>Writing a '0' to this bit enables CMD to be set to 1; Writing a 1 to force CMD bit to be always 0, i.e. disabling SMBus write. This bit can only be written in SMMode. SMBus Read is not affected. I2C Write Pointer Update Command is not affected.</p> <p>Important Note to BIOS: Since BIOS is the source to update SMBCNTL_x register initially after reset, it is important to determine whether the SMBus can have write capability before writing any upper bits (bit24-31) via byte-enable config write (or writing any bit within this register via 32b config write) within the SMBCNTL register.</p>
10:10	RW	0x0	<p>SMB_SOFT_RST (smb_soft_rst):</p> <p>SMBus software reset strobe to graceful terminate pending transaction after ACK and keep the SMB from issuing any transaction until this bit is cleared. If slave device is hung, software can write this bit to 1 and the SMB_CKOVRD to 0 (for more than 35ms) to force hung the SMB slaves to time-out and put it in idle state without using power good reset or warm reset.</p> <p>Note: Software need to set the SMB_CKOVRD back to 1 after 35ms in order to force slave devices to time-out in case there is any pending transaction. The corresponding SMB_STAT_x.SMB_SBE error status bit may be set if there was such pending transaction time-out (non-graceful termination). If the pending transaction was a write operation, the slave device content may be corrupted by this clock override operation. A subsequent SMB command will automatically cleared the SMB_SBE.</p> <p>If the iMC HW perform SMB time-out with the SMB_SBE_EN = 1. Software should simply clear the SMB_SBE and SMB_SOFT_RST sequentially after writing the SMB_CKOVRD = 0 and SMB_SOFT_RST = 1 asserting clock override and perform graceful txn termination. Hardware will automatically de-assert the SMB_CKOVRD update to 1 after the pre-configured 35ms/65ms time-out.</p>



Integrated Memory Controller (iMC) Configuration Registers

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x188,		Function: 0	
Bit	Attr	Default	Description
9:9	RW_LB	0x0	start_tsod_poll: This bit starts the reading of all enabled devices. Note that the hardware will reset this bit when the SMBus polling has started.
8:8	RW_LB	0x0	SMB_TSOD_POLL_EN (smb_tsod_poll_en): TSOD polling enable '0': disable TSOD polling and enable SPDCMD accesses. '1': disable SPDCMD access and enable TSOD polling. It is important to make sure no pending SMBus transaction and the TSOD polling must be disabled (and pending TSOD polling must be drained) before changing the TSOD_POLL_EN.
7:0	RW_LB	0x0	TSOD_PRESENT for the lower and upper channels (tsod_present): DIMM slot mask to indicate whether the DIMM is equipped with TSOD sensor. Bit 7: must be programmed to zero. Upper channel slot #3 is not supported Bit 6: TSOD PRESENT at upper channel (ch 1 or ch 3) slot #2 Bit 5: TSOD PRESENT at upper channel (ch 1 or ch 3) slot #1 Bit 4: TSOD PRESENT at upper channel (ch 1 or ch 3) slot #0 Bit 3: must be programmed to zero. Lower channel slot #3 is not supported Bit 2: TSOD PRESENT at lower channel (ch 0 or ch 2) slot #2 Bit 1: TSOD PRESENT at lower channel (ch 0 or ch 2) slot #1 Bit 0: TSOD PRESENT at lower channel (ch 0 or ch 2) slot #0

2.1.14 smb_tsod_poll_rate_cntr_[0:1]

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x18c,		Function: 0	
Bit	Attr	Default	Description
17:0	RW_LV	0x0	SMB_TSOD_POLL_RATE_CNTR (smb_tsod_poll_rate_cntr): TSOD poll rate counter. When it is decremented to zero, reset to zero or written to zero, SMB_TSOD_POLL_RATE value is loaded into this counter and appear the updated value in the next DCLK.

2.1.15 smb_period_cfg

SMBus Clock Period Config.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x1a0		Function: 0	
Bit	Attr	Default	Description
31:16	RWS	0x445c	Reserved
15:0	RWS	0xfa0	SMB_CLK_PRD (smb_clk_prd): This field specifies both SMBus Clock in number of DCLK. Note: In order to generate a 50% duty cycle SCL, half of the SMB_CLK_PRD is used to generate SCL high. SCL must stay low for at least another half of the SMB_CLK_PRD before pulling high. It is recommend to program an even value in this field since the hardware is simply doing a right shift for the divided by 2 operation. Note the 100 KHz SMB_CLK_PRD default value is calculated based on 800 MTs (400 MHz) DCLK.



2.1.16 smb_period_cntr

SMBus Clock Period Counter.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x1a4		Function: 0	
Bit	Attr	Default	Description
31:16	RO_V	0x0	SMB1_CLK_PRD_CNTR (smb1_clk_prd_cntr): SMBus #1 Clock Period Counter for Ch 23. This field is the current SMBus Clock Period Counter Value.
15:0	RO_V	0x0	SMB0_CLK_PRD_CNTR (smb0_clk_prd_cntr): SMBus #0 Clock Period Counter for Ch 01. This field is the current SMBus Clock Period Counter Value.

2.1.17 smb_tsod_poll_rate

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x1a8		Function: 0	
Bit	Attr	Default	Description
17:0	RWS	0x3e800	SMB_TSOD_POLL_RATE (smb_tsod_poll_rate): TSOD poll rate configuration between consecutive TSOD accesses to the TSOD devices on the same SMBus segment. This field specifies the TSOD poll rate in number of 500 ns per CNFG_500_NANOSEC register field definition.

2.2 Device 19,22 Function 1

DID		VID		0h	SPAREADDRESSLO	80h
PCISTS		PCICMD		4h		84h
CCR		RID		8h		88h
BIST	HDR	PLAT	CLSR	Ch		8Ch
				10h	SPARECTL	90h
				14h	SSRSTATUS	94h
				18h	SCRUBADDRESSLO	98h
				1Ch	SCRUBADDRESSHI	9Ch
				20h	SCRUBCTL	A0h
				24h		A4h
				28h	SPAREINTERVAL	A8h
SDID		SVID		2Ch	RASENABLES	ACh
				30h		B0h
				CAPPTR	34h	SMISPARECTL
				38h	LEAKY_BUCKET_CFG	B8h
MAXLAT	MINGNT	INTPIN	INTL	3Ch		BCh
PXPCAP				40h	LEAKY_BUCKET_CNTR_LO	C0h
				44h	LEAKY_BUCKET_CNTR_HI	C4h



Integrated Memory Controller (iMC) Configuration Registers

	48h	C8h
	4Ch	CCh
	50h	D0h
	54h	D4h
	58h	D8h
	5Ch	DCh
	60h	E0h
	64h	E4h
	68h	E8h
	6Ch	ECh
	70h	F0h
	74h	F4h
	78h	F8h
	7Ch	FCh

2.2.1 pxpcap

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x40		Function: 1	
Bit	Attr	Default	Description
29:25	RO	0x0	Interrupt Message Number (interrupt_message_number): NA for this device
24:24	RO	0x0	Slot Implemented (slot_implemented): NA for integrated endpoints
23:20	RO	0x9	Device/Port Type (device_port_type): Device type is Root Complex Integrated Endpoint
19:16	RO	0x1	Capability Version (capability_version): PCI Express Capability is Compliant with Version 1.0 of the PCI Express Spec. Note: This capability structure is not compliant with Versions beyond 1.0, since they require additional capability registers to be reserved. The only purpose for this capability structure is to make enhanced configuration space available. Minimizing the size of this structure is accomplished by reporting version 1.0 compliance and reporting that this is an integrated root port device. As such, only three Dwords of configuration space are required for this structure.
15:8	RO	0x0	Next Capability Pointer (next_ptr): Pointer to the next capability. Set to 0 to indicate there are no more capability structures.
7:0	RO	0x10	Capability ID (capability_id): Provides the PCI Express capability ID assigned by PCI-SIG.

2.2.2 spareaddresslo

Spare Address Low


Always points to the lower address for the next sparing operation. This register is not affected by the HA access to the spare source rank during the HA window.



Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x80		Function: 1	
Bit	Attr	Default	Description
30:0	RW_LV	0x0	RANKADD (rankadd): Always points to the lower address for the next sparing operation. This register will not be affected by the HA access to the spare source rank during the HA window.

2.2.3 sparectl

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x90		Function: 1	
Bit	Attr	Default	Description
29:29	RW_LB	0x0	DisWPQWM (diswpqwm): Disable WPQ level based water mark, so that sparing wm is only based on HaFifoWM. If DisWPQWM is clear, the spare window is started when the number of hits to the failed DIMM exceed max (# of credits in WPQ not yet returned to the HA, HaFifoWM). If DisWPQWM is set, the spare window starts when the number of hits to the failed DIMM exceed HaFifoWM. In either case, if the number of hits to the failed DIMM do not hit the WM, the spare window will still start after SPAREINTERVAL.NORMOPDUR timer expiration.
28:24	RW_LB	0x0	HaFifoWM (hafifowm): minimum water mark for HA writes to failed rank. Actual wm is max of WPQ credit level and HaFifoWM. When wm is hit the HA is backpressured and a sparing window is started. If DisWPQWM is clear, the spare window is started when the number of hits to the failed DIMM exceed max (# of credits in WPQ not yet returned to the HA, HaFifoWM). If DisWPQWM is set, the spare window starts when the number of hits to the failed DIMM exceed HaFifoWM.
23:16	RW	0x0	SCRATCH_PAD (scratch_pad): This field is available as a scratch pad.
10:8	RW_LB	0x0	DST_RANK (dst_rank): Destination logical rank used for the memory copy.
6:4	RW_LB	0x0	SRC_RANK (src_rank): Source logical rank that provides the data to be copied.
3:2	RW_LB	0x0	CHANNEL SELECT FOR THE SPARE COPY (chn_sel): Since there is only one spare-copy logic for all channels, this field selects the channel or channel-pair for the spare-copy operation. For independent channel operation: 00 = channel 0 is selected for the spare-copy operation 01 = channel 1 is selected for the spare-copy operation 10 = channel 2 is selected for the spare-copy operation 11 = channel 3 is selected for the spare-copy operation For lock-step channel operation: 0x = channel 0 and channel 1 are selected for the spare-copy operation 1x = channel 2 and channel 3 are selected for the spare-copy operation
0:0	RW_LBV	0x0	SPARE_ENABLE (spare_enable): Spare enable when set to 1. Hardware clear after the sparing completion.



2.2.4 ssrstatus

Provides the status of a spare-copy memory Init operation.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x94		Function: 1	
Bit	Attr	Default	Description
2:2	RW1C	0x0	PATCMPLT (patcmplt): All memory has been scrubbed. Hardware sets this bit each time the patrol engine steps through all memory locations. If software wants to monitor 0 --> 1 transition after the bit has been set, the software will need to clear the bit by writing a one to clear this bit in order to distinguish the next patrol scrub completion. Clearing the bit will not affect the patrol scrub operation.
1:1	RO_V	0x0	SPRCMPLT (sprcmplt): Spare Operation Complete. Set by hardware once operation is complete. Bit is cleared by hardware when a new operation is enabled. Note: just before MC release the HA block prior to the completion of the sparing operation, iMC logic will automatically update the corresponding RIR_RNK_TGT target to reflect new DST_RANK.
0:0	RO_V	0x0	SPRINPROGRESS (sprinprogress): Spare Operation in progress. This bit is set by hardware once operation has started. It is cleared once operation is complete or fails.

2.2.5 scrubaddresslo

Scrub Address Low.

This register contains part of the address of the last patrol scrub request issued. When running memtest, the failing address is logged in this register on memtest errors. Software can write the next address to be scrubbed into this register. The STARTSCRUB bit will then trigger the specified address to be scrubbed. Patrol scrubs must be disabled to reliably write this register.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x98		Function: 1	
Bit	Attr	Default	Description
30:0	RW_LB V	0x0	RANKADD (rankadd): Contains the rank address of the last scrub issued. Can be written to specify the next scrub address with STARTSCRUB. Patrol Scrubs must be disabled when writing to this field.



2.2.6 scrubaddresshi

Scrub Address High.

This register pair contains part of the address of the last patrol scrub request issued. Software can write the next address into this register. Scrubbing must be disabled to reliably read and write this register. The STARTSCRUB bit will then trigger the specified address to be scrubbed.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x9c		Function: 1	
Bit	Attr	Default	Description
11:10	RW_LBV	0x0	CHNL (chnl): Can be written to specify the next scrub address with STARTSCRUB. This register is updated with channel address of the last scrub address issued. Patrol Scrubs must be disabled when writing to this field.
7:4	RW_LBV	0x0	RANK (rank): Contains the physical rank ID of the last scrub issued. Can be written to specify the next scrub address with STARTSCRUB. Patrol Scrubs must be disabled when writing to this field.

2.2.7 scrubctl

This register contains the Scrub control parameters and status.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0xa0		Function: 1	
Bit	Attr	Default	Description
31:31	RW_L	0x0	Scrub Enable (scrub_en): Scrub Enable when set.
30:30	RW_LB	0x0	Stop on complete (stop_on_cmpl): Stop patrol scrub at end of memory range. This mode is meant to be used as part of memory migration flow. Intel® Scalable Memory Interconnect (Intel® SMI) is signaled by default.
29:29	RW_LBV	0x0	patrol range complete (ptl_cmpl): When stop_on_cmpl is enabled, patrol will stop at the end of the address range and set this bit. Patrol will resume from beginning of address range when this bit or stop_on_cmpl is cleared by BIOS and patrol scrub is still enabled by scrub_en.
28:28	RW_LB	0x0	Stop on error (stop_on_err): Stop patrol scrub on poison or uncorrectable. On poison, patrol will log error then stop. On uncorr, patrol will convert to poison if enabled then stop. This mode is meant to be used as part of memory migration flow. Intel SMI is signaled by default.
27:27	RW_LBV	0x0	patrol stopped (ptl_stopped): When stop_on_err is set, patrol will stop on error and set this bit. Patrol will resume at the next address when this bit or stop_on_err is cleared by BIOS and patrol scrub is still enabled by scrub_en.
26:26	RW_LBV	0x0	SCRUBISSUED (scrubissued): When Set, the scrub address registers contain the last scrub address issued.
25:25	RW_LB	0x0	ISSUEONCE (issueonce): When Set, the patrol scrub engine will issue the address in the scrub address registers only once and stop.



Integrated Memory Controller (iMC) Configuration Registers

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0xa0		Function: 1	
Bit	Attr	Default	Description
24:24	RW_LBV	0x0	STARTSCRUB (startscrub): When Set, the Patrol scrub engine will start from the address in the scrub address registers. Once the scrub is issued this bit is reset.
23:0	RW_LB	0x0	SCRUBINTERVAL (scrubinterval): Defines the interval in DCLKS between patrol scrub requests. The calculation for this register to get a scrub to every line in 24 hours is: $((86400)/(\text{memory capacity}/64))/\text{cycle time of DCLK}$. RESTRICTIONS: Can only be changed when patrol scrubs are disabled. Set to a minimum value of 1500.

2.2.8 spareinterval

Defines the interval between normal and sparing operations. Interval is defined in dclk.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0xa8		Function: 1	
Bit	Attr	Default	Description
28:16	RW-LB	0x320	NUMSPARE (numspare): Sparing operation duration. System requests will be blocked during this interval and only sparing copy operations will be serviced.
15:0	RW-LB	0xc80	NORMAL OPERATION DURATION (normopdur): Normal operation duration. System requests will be serviced during this interval.

2.2.9 rasenables

RAS Enables Register

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0xac		Function: 1	
Bit	Attr	Default	Description
0:0	RW_LB	0x0	MIRROREN (mirroren): Mirror mode enable. The channel mapping must be set up before this bit will have an effect on iMC operation. This changes the error policy.



2.2.10 smisparectl

System Management Interrupt and Spare control register.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0xb4		Function: 1	
Bit	Attr	Default	Description
17:17	RW-LB	0x0	INTRPT_SEL_PIN (intrpt_sel_pin): Enable pin signaling. When set the interrupt is signaled via the ERROR_N[0] pin to get the attention of a BMC.
16:16	RW-LB	0x0	INTRPT_SEL_CMCI (intrpt_sel_cmci): (CMCI used as a proxy for NMI signaling). Set to enable NMI signaling. Clear to disable NMI signaling. If both NMI and Intel SMI enable bits are set then only Intel SMI is sent.
15:15	RW-LB	0x0	INTRPT_SEL_SMI (intrpt_sel_smi): Intel SMI enable. Set to enable Intel SMI signaling. Clear to disable Intel SMI signaling.

2.2.11 leaky_bucket_cfg

The leaky bucket is implemented as a 53-bit DCLK counter. The upper 42-bit of the 53-bit counter is captured in LEAKY_BUCKET_CNTR_LO and LEAKY_BUCKET_CNTR_HI registers. The carry "strobe" from the not-shown least significant 11-bit counter will trigger this 42-bit counter-pair to count. LEAKY_BUCKET_CFG contains two hot encoding thresholds LEAKY_BKT_CFG_HI and LEAKY_BKT_CFG_LO. The 42-bit counter-pair is compared with the two thresholds pair specified by LEAKY_BKT_CFG_HI and LEAKY_BKT_CFG_LO.



Integrated Memory Controller (iMC) Configuration Registers

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0xb8		Function: 1	
Bit	Attr	Default	Description
11:6	RW	0x0	<p>LEAKY_BKT_CFG_HI (leaky_bkt_cfg_hi):</p> <p>This is the higher order bit select mask of the two hot encoding threshold. The value of this field specify the bit position of the mask:</p> <p>00h: reserved 01h: LEAKY_BUCKET_CNTR_LO bit 1, i.e. bit 12 of the full 53b counter ... 1Fh: LEAKY_BUCKET_CNTR_LO bit 31, i.e. bit 42 of the full 53b counter 20h: LEAKY_BUCKET_CNTR_HI bit 0, i.e. bit 43 of the full 53b counter ... 29h: LEAKY_BUCKET_CNTR_HI bit 9, i.e. bit 52 of the full 53b counter 2Ah - 3F: reserved</p> <p>When both counter bits selected by the LEAKY_BKT_CFG_HI and LEAKY_BKT_CFG_LO are set, the 53b leaky bucket counter will be reset and the logic will generate a primary leak Strobe which is used by a 2-bit LEAKY_BKT_2ND_CNTR. LEAKY_BKT_2ND_CNTR_LIMIT specifies the value to generate LEAK pulse which is used to decrement the correctable error counter by 1 as shown below:</p> <p>LEAKY_BKT_2ND_CNTR_LIMIT LEAK pulse to decrement CE counter by 1</p> <p>00b (default): 4 x Primary leak strobe (four times the value programmed by the LEAKY_BKT_CFG_HI and LEAKY_BKT_CFG_LO)</p> <p>01b: 1x Primary leak strobe (same as the value programmed by the LEAKY_BKT_CFG_HI and LEAKY_BKT_CFG_LO)</p> <p>10b: 2x Primary leak strobe (two times the value programmed by the LEAKY_BKT_CFG_HI and LEAKY_BKT_CFG_LO)</p> <p>11b: 3x Primary leak strobe (two times the value programmed by the LEAKY_BKT_CFG_HI and LEAKY_BKT_CFG_LO)</p> <p>Note: A value of all zeros in LEAKY_BUCKET_CFG register is equivalent to no leaky bucketing.</p> <p>BIOS must program this register to any non-zero value before switching to NORMAL mode.</p>



Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0xb8		Function: 1	
Bit	Attr	Default	Description
5:0	RW	0x0	<p>LEAKY_BKT_CFG_LO (leaky_bkt_cfg_lo):</p> <p>This is the lower order bit select mask of the two hot encoding threshold. The value of this field specify the bit position of the mask:</p> <p>00h: reserved</p> <p>01h: LEAKY_BUCKET_CNTR_LO bit 1, i.e. bit 12 of the full 53b counter</p> <p>...</p> <p>1Fh: LEAKY_BUCKET_CNTR_LO bit 31, i.e. bit 42 of the full 53b counter</p> <p>20h: LEAKY_BUCKET_CNTR_HI bit 0, i.e. bit 43 of the full 53b counter</p> <p>...</p> <p>29h: LEAKY_BUCKET_CNTR_HI bit 9, i.e. bit 52 of the full 53b counter</p> <p>2Ah - 3F: reserved</p> <p>When both counter bits selected by the LEAKY_BKT_CFG_HI and LEAKY_BKT_CFG_LO are set, the 53b leaky bucket counter will be reset and the logic will generate a primary leak Strobe which is used by a 2-bit LEAKY_BKT_2ND_CNTR. LEAKY_BKT_2ND_CNTR_LIMIT specifies the value to generate LEAK pulse which is used to decrement the correctable error counter by 1 as shown below:</p> <p>LEAKY_BKT_2ND_CNTR_LIMIT LEAK pulse to decrement CE counter by 1</p> <p>00b (default): 4 x Primary leak strobe (four times the value programmed by the LEAKY_BKT_CFG_HI and LEAKY_BKT_CFG_LO)</p> <p>01b: 1x Primary leak strobe (same as the value programmed by the LEAKY_BKT_CFG_HI and LEAKY_BKT_CFG_LO)</p> <p>10b: 2x Primary leak strobe (two times the value programmed by the LEAKY_BKT_CFG_HI and LEAKY_BKT_CFG_LO)</p> <p>11b: 3x Primary leak strobe (two times the value programmed by the LEAKY_BKT_CFG_HI and LEAKY_BKT_CFG_LO)</p> <p>Note: A value of all zeros in LEAKY_BUCKET_CFG register is equivalent to no leaky bucketing.</p> <p>MRC BIOS must program this register to any non-zero value before switching to NORMAL mode.</p>

2.2.12 leaky_bucket_cntr_lo

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0xc0		Function: 1	
Bit	Attr	Default	Description
31:0	RW_V	0x0	<p>Leaky Bucket Counter Low (leaky_bkt_cntr_lo):</p> <p>This is the lower half of the leaky bucket counter. The full counter is actually a 53b "DCLK" counter. There is a least significant 11b of the 53b counter is not captured in CSR. The carry "strobe" from the not-shown least significant 11b counter will trigger this 42b counter pair to count. The 42b counter-pair is compared with the two-hot encoding threshold specified by the LEAKY_BUCKET_CFG_HI and LEAKY_BUCKET_CFG_LO pair. When the counter bits specified by the LEAKY_BUCKET_CFG_HI and LEAKY_BUCKET_CFG_LO are both set, the 53b counter is reset and the leaky bucket logic will generate a LEAK strobe last for 1 DCLK.</p>



2.2.13 leaky_bucket_cntr_hi

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0xc4		Function: 1	
Bit	Attr	Default	Description
9:0	RW_V	0x0	Leaky Bucket Counter High Limit (leaky_bkt_cntr_hi): This is the upper half of the leaky bucket counter. The full counter is actually a 53b "DCLK" counter. There is a least significant 11b of the 53b counter is not captured in CSR. The carry "strobe" from the not-shown least significant 11b counter will trigger this 42b counter pair to count. The 42b counter-pair is compared with the two-hot encoding threshold specified by the LEAKY_BUCKET_CFG_HI and LEAKY_BUCKET_CFG_LO pair. When the counter bits specified by the LEAKY_BUCKET_CFG_HI and LEAKY_BUCKET_CFG_LO are both set, the 53b counter is reset and the leaky bucket logic will generate a LEAK strobe last for 1 DCLK.

2.3 Device 19,22 Functions 2,3,4,5

DID		VID		0h	DIMMMTR_0		80h	
PCISTS		PCICMD		4h	DIMMMTR_1		84h	
CCR			RID	8h	DIMMMTR_2		88h	
BIST	HDR	PLAT	CLSR	Ch			8Ch	
				10h			90h	
				14h			94h	
				18h			98h	
				1Ch			9Ch	
				20h			A0h	
				24h			A4h	
				28h			A8h	
				2Ch			ACH	
SDID		SVID		30h				
				CAPPTR	34h	B4h		
				38h	B8h			
MAXLAT	MINGNT	INTPIN	INTL	3Ch				
PXPCAP				40h			C0h	
				44h			C4h	
				48h			C8h	
				4Ch			CCh	
				50h			D0h	
				54h			D4h	
				58h			D8h	
				5Ch			DCh	
				60h			E0h	
64h	E4h							
68h	E8h							



	6Ch	ECh
	70h	F0h
	74h	F4h
	78h	F8h
	7Ch	FCh

2.3.1 pxpcap

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x40		Function: 2,3,4,5	
Bit	Attr	Default	Description
29:25	RO	0x0	Interrupt Message Number (interrupt_message_number): NA for this device
24:24	RO	0x0	Slot Implemented (slot_implemented): NA for integrated endpoints
23:20	RO	0x9	Device/Port Type (device_port_type): Device type is Root Complex Integrated Endpoint
19:16	RO	0x1	Capability Version (capability_version): PCI Express Capability is Compliant with Version 1.0 of the PCI Express Spec. Note: This capability structure is not compliant with Versions beyond 1.0, since they require additional capability registers to be reserved. The only purpose for this capability structure is to make enhanced configuration space available. Minimizing the size of this structure is accomplished by reporting version 1.0 compliance and reporting that this is an integrated root port device. As such, only three Dwords of configuration space are required for this structure.
15:8	RO	0x0	Next Capability Pointer (next_ptr): Pointer to the next capability. Set to 0 to indicate there are no more capability structures.
7:0	RO	0x10	Capability ID (capability_id): Provides the PCI Express capability ID assigned by PCI-SIG.

2.3.2 dimmmtr_[0:2]

DIMM Memory Technology.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x80, 0x84, 0x88		Function: 2,3,4,5	
Bit	Attr	Default	Description
22:22	RW_LB	0x0	hdr1_parity: When set, will enable parity calculation to include address bits 17:16 which are sent on chip select lines 7:6 and 3:2.
21:21	RW_LB	0x0	hdr1: When set, will enable High Density Reduced Load mode which will transmit Row address bits 17:16 on chip select lines 7:6 and 3:2.
20:20	RW_LB	0x0	ddr4_mode: When set, indicating DDR4 DIMM type is used. Channel 0 and 1, and channel 2 and 3 must have matching values even if both ddr channels are not populated.



2.3.3 pxpenhcap

This field points to the next Capability in extended configuration space.

Type: CFG		PortID: N/A	
Bus: 1		Device: 19,22	
Offset: 0x100		Function: 2,3,4,5	
Bit	Attr	Default	Description
31:20	RO	0x0	Next Capability Offset (next_capability_offset):
19:16	RO	0x0	Capability Version (capability_version): Indicates there are no capability structures in the enhanced configuration space.
15:0	RO	0x0	Capability ID (capability_id): Indicates there are no capability structures in the enhanced configuration space.

2.4 Device 20,21,23 Functions 0, 1


DID		VID		0h		80h
PCISTS		PCICMD		4h		84h
CCR			RID	8h		88h
BIST	HDR	PLAT	CLSR	Ch		8Ch
				10h		90h
				14h		94h
				18h		98h
				1Ch		9Ch
				20h		
				24h		A0h
				28h		A4h
				SDID		SVID
				30h		ACh
				CAPPTR		34h
				38h		B4h
				MAXLAT		MINGNT
PXPCAP				40h		BCh
				44h		C0h
				48h		C4h
				4Ch		
				50h		
				54h		D0h
				58h		D4h
				5Ch		D8h
				60h		DCh
				64h		E0h
				68h		E4h
				6Ch		E8h



Integrated Memory Controller (iMC) Configuration Registers

	70h		F0h
	74h		F4h
	78h		F8h
	7Ch		FCh

	100h			180h
	104h			184h
CHN_TEMP_CFG	108h			188h
CHN_TEMP_STAT	10Ch			18Ch
DIMM_TEMP_OEM_0	110h	THRT_PWR_DIMM_1	THRT_PWR_DIMM_0	190h
DIMM_TEMP_OEM_1	114h		THRT_PWR_DIMM_2	194h
DIMM_TEMP_OEM_2	118h			198h
	11Ch			19Ch
DIMM_TEMP_TH_0	120h			1A0h
DIMM_TEMP_TH_1	124h			1A4h
DIMM_TEMP_TH_2	128h			1A8h
	12Ch			1ACh
DIMM_TEMP_THRT_LMT_0	130h			1B0h
DIMM_TEMP_THRT_LMT_1	134h			1B4h
DIMM_TEMP_THRT_LMT_2	138h			1B8h
	13Ch			1BCh
DIMM_TEMP_EV_OFST_0	140h			1C0h
DIMM_TEMP_EV_OFST_1	144h			1C4h
DIMM_TEMP_EV_OFST_2	148h			1C8h
	14Ch			1CCh
DIMMTEMPSTAT_0	150h			1D0h
DIMMTEMPSTAT_1	154h			1D4h
DIMMTEMPSTAT_2	158h			1D8h
	15Ch			1DCh
	160h			1E0h
	164h			1E4h
	168h			1E8h
	16Ch			1ECh
	170h			1F0h
	174h			1F4h
	178h			1F8h
	17Ch			1FCh



2.4.1 pxpcap

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x40		Function: 0,1	
Bit	Attr	Default	Description
7:0	RO	0x10	Capability ID (capability_id): Provides the PCI Express capability ID assigned by PCI-SIG.

2.4.2 chn_temp_cfg

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x108		Function: 0,1	
Bit	Attr	Default	Description
31:31	RW	0x1	OLTT_EN (oltt_en): Enable OLTT temperature tracking.
29:29	RW	0x0	CLTT_OR_PCODE_TEMP_MUX_SEL (cltt_or_pcode_temp_mux_sel): The TEMP_STAT byte update mux select control to direct the source to update DIMMTEMPSTAT_[0:3][7:0]: 0: Corresponding to the DIMM TEMP_STAT byte from PCODE_TEMP_OUTPUT. 1: TSOD temperature reading from CLTT logic.
28:28	RW_O	0x1	CLTT_DEBUG_DISABLE_LOCK (cltt_debug_disable_lock): Lock bit of DIMMTEMPSTAT_[0:3][7:0]: Set this lock bit to disable configuration write to DIMMTEMPSTAT_[0:3][7:0].
27:27	RW	0x1	Enables thermal bandwidth throttling limit (bw_limit_thrt_en):
23:16	RW	0x0	THRT_EXT (thrt_ext): Max number of throttled transactions to be issued during BWLIMITTF due to externally asserted MEMHOT#.
15:15	RW	0x0	THRT_ALLOW_ISOCH (thrt_allow_isoch): When this bit is zero, MC will lower CKE during Thermal Throttling, and ISOCH is blocked. When this bit is one, MC will NOT lower CKE during Thermal Throttling, and ISOCH will be allowed base on bandwidth throttling setting. However, setting this bit would mean more power consumption due to CKE is asserted during thermal or power throttling.
10:0	RW	0x3ff	BW_LIMIT_TF (bw_limit_tf): BW Throttle Window Size in DCLK. Note: This value is left shifted 3 bits before being used.

2.4.3 chn_temp_stat

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x10c		Function: 0,1	
Bit	Attr	Default	Description
2:2	RW1C	0x0	Event Asserted on DIMM ID 2 (ev_asrt_dimm2): Event Asserted on DIMM ID 2
1:1	RW1C	0x0	Event Asserted on DIMM ID 1 (ev_asrt_dimm1): Event Asserted on DIMM ID 1



Integrated Memory Controller (iMC) Configuration Registers

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x10c		Function: 0,1	
Bit	Attr	Default	Description
0:0	RW1C	0x0	Event Asserted on DIMM ID 0 (ev_asrt_dimm0): Event Asserted on DIMM ID 0

2.4.4 dimm_temp_oem_[0:2]

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x110, 0x114, 0x118		Function: 0,1	
Bit	Attr	Default	Description
26:24	RW	0x0	TEMP_OEM_HI_HYST (temp_oem_hi_hyst): Positive going Threshold Hysteresis Value. This value is subtracted from TEMPOEMHI to determine the point where the asserted status for that threshold will clear. Set to 00h if sensor does not support positive-going threshold hysteresis
18:16	RW	0x0	TEMP_OEM_LO_HYST (temp_oem_lo_hyst): Negative going Threshold Hysteresis Value. This value is added to TEMPOEMLO to determine the point where the asserted status for that threshold will clear. Set to 00h if sensor does not support negative-going threshold hysteresis.
15:8	RW	0x50	TEMP_OEM_HI (temp_oem_hi): Upper Threshold value - TCase threshold at which to Initiate System Interrupt (Intel SMI or MEMHOT#) at a+ going rate. Note: The default value is listed in decimal. Valid range: 32 - 127 in degree (C). Others: reserved.
7:0	RW	0x4b	TEMP_OEM_LO (temp_oem_lo): Lower Threshold Value - TCase threshold at which to Initiate System Interrupt (Intel SMI or MEMHOT#) at a - going rate. Note: the default value is listed in decimal. Valid range: 32 - 127 in degree (C). Others: reserved.

2.4.5 dimm_temp_th_[0:2]

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x120, 0x124, 0x128		Function: 0,1	
Bit	Attr	Default	Description
26:24	RW-LB	0x0	TEMP_THRT_HYST (temp_thrt_hyst): Positive going Threshold Hysteresis Value. Set to 00h if sensor does not support positive-going threshold hysteresis. This value is subtracted from TEMP_THRT_XX to determine the point where the asserted status for that threshold will clear.
23:16	RW-LB	0x5f	TEMP_HI (temp_hi): TCase threshold at which to Initiate THRTCRIT and assert THERMTRIP# valid range: 32 - 127 in degree (C). Note: the default value is listed in decimal. FF: Disabled Others: reserved. TEMP_HI should be programmed so it is greater than TEMP_MID.



Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x120, 0x124, 0x128		Function: 0,1	
Bit	Attr	Default	Description
15:8	RW	0x5a	TEMP_MID (temp_mid): TCase threshold at which to Initiate THRTHI and assert valid range: 32 - 127 in degree (C). Note: The default value is listed in decimal. FF: Disabled Others: reserved. TEMP_MID should be programmed so it is less than TEMP_HI.
7:0	RW	0x55	TEMP_LO (temp_lo): TCase threshold at which to Initiate 2x refresh andor THRTMID and initiate Interrupt (MEMHOT#). Note: The default value is listed in decimal.valid range: 32 - 127 in degree (C). FF: Disabled Others: reserved. TEMP_LO should be programmed so it is less than TEMP_MID

2.4.6 dimm_temp_thrt_lmt_[0:2]

All three THRT_CRIT, THRT_HI and THRT_MID are per DIMM BW limit, i.e. all activities (ACT, READ, WRITE) from all ranks within a DIMM are tracked together in one DIMM activity counter.

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x130, 0x134, 0x138		Function: 0,1	
Bit	Attr	Default	Description
23:16	RW-LB	0x0	THRT_CRIT (thrt_crit): Max number of throttled transactions (ACT, READ, WRITE) to be issued during BWLIMITTF.
15:8	RW-LB	0xf	THRT_HI (thrt_hi): Max number of throttled transactions (ACT, READ, WRITE) to be issued during BWLIMITTF.
7:0	RW	0xff	THRT_MID (thrt_mid): Max number of throttled transactions (ACT, READ, WRITE) to be issued during BWLIMITTF.

2.4.7 dimm_temp_ev_ofst_[0:2]

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x140, 0x144, 0x148		Function: 0,1	
Bit	Attr	Default	Description
31:24	RO	0x0	TEMP_AVG_INTRVL (temp_avg_intrvl): Temperature data is averaged over this period. At the end of averaging period (ms), averaging process starts again. 0x1 - 0xFF Averaging data is read via TEMPDIMM STATUSREGISTER (Byte 1/2) as well as used for generating hysteresis based interrupts. 00 Instantaneous Data (non-averaged) is read via TEMPDIMM STATUSREGISTER (Byte 1/2) as well as used for generating hysteresis based interrupts. Note: CPU does not support temp averaging.



Integrated Memory Controller (iMC) Configuration Registers

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x140, 0x144, 0x148		Function: 0,1	
Bit	Attr	Default	Description
14:14	RW	0x0	Initiate THRTMID on TEMPLO (ev_thrtmid_templo): Initiate THRTMID on TEMPLO
13:13	RW	0x1	Initiate 2X refresh on TEMPLO (ev_2x_ref_templo_en): Initiate 2X refresh on TEMPLO DIMM with extended temperature range capability will need double refresh rate in order to avoid data lost when DIMM temperature is above 85C but below 95C. Warning: If the 2x refresh is disable with extended temperature range DIMM configuration, system cooling and power thermal throttling scheme must guarantee the DIMM temperature will not exceed 85C.
12:12	RW	0x0	Assert MEMHOT Event on TEMPHI (ev_mh_temphi_en): Assert MEMHOT# Event on TEMPHI
11:11	RW	0x0	Assert MEMHOT Event on TEMPMID (ev_mh_tempmid_en): Assert MEMHOT# Event on TEMPMID
10:10	RW	0x0	Assert MEMHOT Event on TEMPLO (ev_mh_templo_en): Assert MEMHOT# Event on TEMPLO
9:9	RW	0x0	Assert MEMHOT Event on TEMPOEMHI (ev_mh_tempoemhi_en): Assert MEMHOT# Event on TEMPOEMHI
8:8	RW	0x0	Assert MEMHOT Event on TEMPOEMLO (ev_mh_tempoemlo_en): Assert MEMHOT# Event on TEMPOEMLO
3:0	RW	0x0	DIMM_TEMP_OFFSET (dimm_temp_offset): Temperature Offset Register.

2.4.8 dimmtempstat_[0:2]

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x150, 0x154, 0x158		Function: 0,1	
Bit	Attr	Default	Description
28:28	RW1C	0x0	Event Asserted on TEMPHI going HIGH (ev_asrt_temphi): Event Asserted on TEMPHI going HIGH It is assumed that each of the event assertion is going to trigger Configurable interrupt (Either MEMHOT# only or both Intel SMI and MEMHOT#) defined in bit 30 of CHN_TEMP_CFG.
27:27	RW1C	0x0	Event Asserted on TEMPMID going High (ev_asrt_tempmid): Event Asserted on TEMPMID going High It is assumed that each of the event assertion is going to trigger configurable interrupt (Either MEMHOT# only or both Intel SMI and MEMHOT#) defined in bit 30 of CHN_TEMP_CFG.
26:26	RW1C	0x0	Event Asserted on TEMPLO Going High (ev_asrt_templo): Event Asserted on TEMPLO Going High It is assumed that each of the event assertion is going to trigger Configurable interrupt (Either MEMHOT# only or both Intel SMI and MEMHOT#) defined in bit 30 of CHN_TEMP_CFG.
25:25	RW1C	0x0	Event Asserted on TEMPOEMLO Going Low (ev_asrt_tempoemlo): Event Asserted on TEMPOEMLO Going Low It is assumed that each of the event assertion is going to trigger Configurable interrupt (Either MEMHOT# only or both Intel SMI and MEMHOT#) defined in bit 30 of CHN_TEMP_CFG.



Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x150, 0x154, 0x158		Function: 0,1	
Bit	Attr	Default	Description
24:24	RW1C	0x0	Event Asserted on TEMPOEMHI Going High (ev_asrt_tempoemhi): Event Asserted on TEMPOEMHI Going High It is assumed that each of the event assertion is going to trigger Configurable interrupt (Either MEMHOT# only or both Intel SMI and MEMHOT#) defined in bit 30 of CHN_TEMP_CFG.
7:0	RW_LV	0x55	DIMM_TEMP (dimm_temp): Current DIMM Temperature for thermal throttling. Lock by CLTT_DEBUG_DISABLE_LOCK. When the CLTT_DEBUG_DISABLE_LOCK is set, this field becomes read-only, i.e. configuration write to this byte is aborted. This byte is updated from internal logic from a 2:1 Mux which can be selected from either CLTT temperature or from the corresponding temperature registers output (PCODE_TEMP_OUTPUT) updated from pcode. The mux select is controlled by CLTT_OR_PCODE_TEMP_MUX_SEL defined in CHN_TEMP_CFG register. Valid range from 0 to 127 i.e. 0C to +127C. Any negative value read from TSOD is forced to 0. TSOD decimal point value is also truncated to integer value.


2.4.9 thrt_pwr_dimm_[0:2]

bit[10:0]: Max number of transactions (ACT, READ, WRITE) to be allowed during the 1 usec throttling timeframe per power throttling.

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x190, 0x192, 0x194		Function: 0,1	
Bit	Attr	Default	Description
15:15	RW	0x1	THRT_PWR_EN (thrt_pwr_en): bit[15]: set to one to enable the power throttling for the DIMM.
11:0	RW	0xffff	Power Throttling Control (thrt_pwr): bit[11:0]: Max number of transactions (ACT, READ, WRITE) to be allowed (per DIMM) during the 1 micro-sec throttling timeframe per power throttling.

2.5 Device 20,21,23 Functions 2, 3


DID		VID		0h		80h	
PCISTS		PCICMD		4h		84h	
CCR			RID	8h		88h	
BIST	HDR	PLAT	CLSR	Ch		8Ch	
				10h		90h	
				14h		94h	
				18h		98h	
				1Ch		9Ch	
				20h		A0h	
				24h		A4h	
				28h		A8h	
				2Ch		ACh	
SDID		SVID		2Ch			ACh



Integrated Memory Controller (IMC) Configuration Registers

				30h		B0h
				CAPPTR	34h	B4h
					38h	B8h
MAXLAT	MINGNT	INTPIN	INTL		3Ch	BCh
PXPCAP					40h	C0h
					44h	C4h
					48h	C8h
					4Ch	CCh
					50h	D0h
					54h	D4h
					58h	D8h
					5Ch	DCh
					60h	E0h
					64h	E4h
					68h	E8h
					6Ch	ECh
					70h	F0h
					74h	F4h
					78h	F8h
					7Ch	FCh

					100h		180h
CORRERRCNT_0					104h		184h
CORRERRCNT_1					108h		188h
CORRERRCNT_2					10Ch		18Ch
CORRERRCNT_3					110h		190h
					114h		194h
					118h		198h
CORRERRTHRSILD_0					11Ch		19Ch
CORRERRTHRSILD_1					120h		1A0h
CORRERRTHRSILD_2					124h		1A4h
CORRERRTHRSILD_3					128h		1A8h
					12Ch		1ACh
					130h		1B0h
CORRERRORSTATUS					134h		1B4h
LEAKY_BKT_2ND_CNTR_REG					138h		1B8h
					13Ch		1BCh
DEVTAG_C NTL_3	DEVTAG_C NTL_2	DEVTAG_C NTL_1	DEVTAG_C NTL_0		140h		1C0h
DEVTAG_C NTL_7	DEVTAG_C NTL_6	DEVTAG_C NTL_5	DEVTAG_C NTL_4		144h		1C4h
					148h		1C8h
					14Ch		1CCh
					150h		1D0h



154h	1D4h
158h	1D8h
15Ch	1DCh
160h	1E0h
164h	1E4h
168h	1E8h
16Ch	1ECh
170h	1F0h
174h	1F4h
178h	1F8h
17Ch	1FCh

2.5.1 correrrcnt_0

Per Rank corrected error counters.

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x104		Function: 2,3	
Bit	Attr	Default	Description
31:31	RW1CS	0x0	RANK 1 OVERFLOW (overflow_1): The corrected error count for this rank has been overflowed. Once set it can only be cleared via a write from BIOS.
30:16	RWS_LV	0x0	RANK 1 CORRECTABLE ERROR COUNT (cor_err_cnt_1): The corrected error count for this rank. Hardware automatically clears this field when the corresponding OVERFLOW_x bit is changing from 0 to 1.
15:15	RW1CS	0x0	RANK 0 OVERFLOW (overflow_0): The corrected error count for this rank has been overflowed. Once set it can only be cleared via a write from BIOS.
14:0	RWS_LV	0x0	RANK 0 CORRECTABLE ERROR COUNT (cor_err_cnt_0): The corrected error count for this rank. Hardware automatically clear this field when the corresponding OVERFLOW_x bit is changing from 0 to 1.



2.5.2 correrrcnt_1


Per Rank corrected error counters.

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x108		Function: 2,3	
Bit	Attr	Default	Description
31:31	RW1CS	0x0	RANK 3 OVERFLOW (overflow_3): The corrected error count has crested over the limit for this rank. Once set it can only be cleared via a write from BIOS.
30:16	RWS_LV	0x0	RANK 3 COR_ERR_CNT (cor_err_cnt_3): The corrected error count for this rank.
15:15	RW1CS	0x0	RANK 2 OVERFLOW (overflow_2): The corrected error count has crested over the limit for this rank. Once set it can only be cleared via a write from BIOS.
14:0	RWS_LV	0x0	RANK 2 COR_ERR_CNT (cor_err_cnt_2): The corrected error count for this rank.

2.5.3 correrrcnt_2

Per Rank corrected error counters.

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x10c		Function: 2,3	
Bit	Attr	Default	Description
31:31	RW1CS	0x0	RANK 5 OVERFLOW (overflow_5): The corrected error count has crested over the limit for this rank. Once set it can only be cleared via a write from BIOS.
30:16	RWS_LV	0x0	RANK 5 COR_ERR_CNT (cor_err_cnt_5): The corrected error count for this rank.
15:15	RW1CS	0x0	RANK 4 OVERFLOW (overflow_4): The corrected error count has crested over the limit for this rank. Once set it can only be cleared via a write from BIOS.
14:0	RWS_LV	0x0	RANK 4 COR_ERR_CNT (cor_err_cnt_4): The corrected error count for this rank.



2.5.4 correrrcnt_3


Per Rank corrected error counters.

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x110		Function: 2,3	
Bit	Attr	Default	Description
31:31	RW1CS	0x0	RANK 7 OVERFLOW (overflow_7): The corrected error count for this rank.
30:16	RWS_LV	0x0	RANK 7 COR_ERR_CNT_7 (cor_err_cnt_7): The corrected error count for this rank.
15:15	RW1CS	0x0	RANK 6 OVERFLOW (overflow_6): The corrected error count has crested over the limit for this rank. Once set it can only be cleared via a write from BIOS.
14:0	RWS_LV	0x0	RANK 6 COR_ERR_CNT (cor_err_cnt_6): The corrected error count for this rank.

2.5.5 correrrthrshld_0

This register holds the per rank corrected error thresholding value.

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x11c		Function: 2,3	
Bit	Attr	Default	Description
30:16	RW-LB	0x7fff	RANK 1 COR_ERR_TH (cor_err_th_1): The corrected error threshold for this rank that will be compared to the per rank corrected error counter.
14:0	RW-LB	0x7fff	RANK 0 COR_ERR_TH (cor_err_th_0): The corrected error threshold for this rank that will be compared to the per rank corrected error counter.



2.5.6 **correrrthrshld_1**

This register holds the per rank corrected error thresholding value.

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x120		Function: 2,3	
Bit	Attr	Default	Description
30:16	RW-LB	0x7fff	RANK 3 COR_ERR_TH (cor_err_th_3): The corrected error threshold for this rank that will be compared to the per rank corrected error counter.
14:0	RW-LB	0x7fff	RANK 2 COR_ERR_TH (cor_err_th_2): The corrected error threshold for this rank that will be compared to the per rank corrected error counter.

2.5.7 **correrrthrshld_2**

This register holds the per rank corrected error thresholding value.

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x124		Function: 2,3	
Bit	Attr	Default	Description
30:16	RW-LB	0x7fff	RANK 5 COR_ERR_TH (cor_err_th_5): The corrected error threshold for this rank that will be compared to the per rank corrected error counter.
14:0	RW-LB	0x7fff	RANK 4 COR_ERR_TH (cor_err_th_4): The corrected error threshold for this rank that will be compared to the per rank corrected error counter.

2.5.8 **correrrthrshld_3**

This register holds the per rank corrected error thresholding value.

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x128		Function: 2,3	
Bit	Attr	Default	Description
30:16	RW-LB	0x7fff	RANK 7 COR_ERR_TH (cor_err_th_7): The corrected error threshold for this rank that will be compared to the per rank corrected error counter.
14:0	RW-LB	0x7fff	RANK 6 COR_ERR_TH (cor_err_th_6): The corrected error threshold for this rank that will be compared to the per rank corrected error counter.



2.5.9 corrrerrorstatus

Per rank corrected error status. These bits are reset by bios.

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x134		Function: 2,3	
Bit	Attr	Default	Description
31:24	RW_V	0x0	ddr4crc_rank_log: This field get set with 1'b1 if the corresponding rank detected DDR4 CRC in one of its write data. This will be cleared by BIOS.
7:0	RW1C	0x0	ERR_OVERFLOW_STAT (err_overflow_stat): This 8 bit field is the per rank error over-threshold status bits. The organization is as follows: Bit 0 : Rank 0 Bit 1 : Rank 1 Bit 2 : Rank 2 Bit 3 : Rank 3 Bit 4 : Rank 4 Bit 5 : Rank 5 Bit 6 : Rank 6 Bit 7 : Rank 7 Note: The register tracks which rank has reached or exceeded the corresponding CORRERTRSHLD threshold settings.

2.5.10 leaky_bkt_2nd_cntr_reg

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x138		Function: 2,3	
Bit	Attr	Default	Description
31:16	RW	0x0	LEAKY_BKT_2ND_CNTR_LIMIT(leaky_bkt_2nd_cntr_limit): Secondary Leaky Bucket Counter Limit (2b per DIMM). This register defines secondary leaky bucket counter limit for all 8 logical ranks within channel. The counter logic will generate the secondary LEAK pulse to decrement the rank's correctable error counter by 1 when the corresponding rank leaky bucket rank counter roll over at the predefined counter limit. The counter increment at the primary leak pulse from the LEAKY_BUCKET_CNTR_LO and LEAKY_BUCKET_CNTR_HI logic. Bit[31:30]: Rank 7 Secondary Leaky Bucket Counter Limit Bit[29:28]: Rank 6 Secondary Leaky Bucket Counter Limit Bit[27:26]: Rank 5 Secondary Leaky Bucket Counter Limit Bit[25:24]: Rank 4 Secondary Leaky Bucket Counter Limit Bit[23:22]: Rank 3 Secondary Leaky Bucket Counter Limit Bit[21:20]: Rank 2 Secondary Leaky Bucket Counter Limit Bit[19:18]: Rank 1 Secondary Leaky Bucket Counter Limit Bit[17:16]: Rank 0 Secondary Leaky Bucket Counter Limit The value of the limit is defined as the following: 0: The LEAK pulse is generated one DCLK after the primary LEAK pulse is asserted. 1: the LEAK pulse is generated one DCLK after the counter roll over at 1. 2: the LEAK pulse is generated one DCLK after the counter roll over at 2. 3: the LEAK pulse is generated one DCLK after the counter roll over at 3.



Integrated Memory Controller (iMC) Configuration Registers

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x138		Function: 2,3	
Bit	Attr	Default	Description
15:0	RW_V	0x0	<p>LEAKY_BKT_2ND_CNTR (leaky_bkt_2nd_cntr):</p> <p>Per rank secondary leaky bucket counter (2b per rank)</p> <p>bit [15:14]: rank 7 secondary leaky bucket counter</p> <p>bit [13:12]: rank 6 secondary leaky bucket counter</p> <p>bit [11:10]: rank 5 secondary leaky bucket counter</p> <p>bit [9:8]: rank 4 secondary leaky bucket counter</p> <p>bit [7:6]: rank 3 secondary leaky bucket counter</p> <p>bit [5:4]: rank 2 secondary leaky bucket counter</p> <p>bit [3:2]: rank 1 secondary leaky bucket counter</p> <p>bit [1:0]: rank 0 secondary leaky bucket counter</p>

2.5.11 devtag_cntl_[0:7]

SDDC Usage model

When the number of correctable errors (CORRERRCNT_x) from a particular rank exceeds the corresponding threshold (CORRERRTHRSHLD_y), hardware will generate a Intel SMI interrupt and log and preserve the failing device in the FailDevice field. SMM software will read the failing device on the particular rank. Software then set the EN bit to enable substitution of the failing device/rank with the parity from the rest of the devices in line.

For independent channel configuration, each rank can tag once. Up to 8 ranks can be tagged.

For lock-step channel configuration, only one x8 device can be tagged per rank-pair. SMM software must identify which channel should be tagged for this rank and only set the valid bit for the channel from the channel-pair.

There is no hardware logic to report incorrect programming error. Unpredictable error and or silent data corruption will be the consequence of such programming error.

If the rank-sparing is enabled, it is recommended to prioritize the rank-sparing before triggering the device tagging due to the nature of the device tagging would drop the correction capability and any subsequent ECC error from this rank would cause uncorrectable error.

Type: CFG		PortID: N/A	
Bus: 1		Device: 20,21,23	
Offset: 0x140, 0x141, 0x142, 0x143, 0x144, 0x145, 0x146, 0x147		Function: 2,3	
Bit	Attr	Default	Description
7:7	RWS_L	0x0	<p>Device tagging enable for this rank (en):</p> <p>Device tagging SDDC enable for this rank. Once set, the parity device of the rank is used for the replacement device content. After tagging, the rank will no longer have the "correction" capability. ECC error "detection" capability will not degrade after setting this bit.</p> <p>For lock-step channel configuration, only one x8 device can be tagged per rank-pair. SMM software must identify which channel should be tagged for this rank and only set the corresponding DEVTAG_CNTL_x.EN bit for the channel contains the fail device. The DEVTAG_CNTL_x.EN on the other channel of the corresponding rank must not be set.</p>



Integrated Memory Controller (iMC) Configuration Registers



3 Intel® QuickPath Interconnect (Intel® QPI) Agent Registers

The Intel® QuickPath Interconnect (Intel® QPI) Agent is a is the coherent communication interface between processors.

- The Intel® Xeon® Processor E5-1600 v4 does not implement Intel QPI links as it is single socket.
- The Intel® Xeon® Processor E5-2600 and E5-4600 v4 implement 2 Intel QPI links (0,1).

3.1 Device 8 Function 0

DID		VID		0h		80h	
PCISTS		PCICMD		4h		84h	
CCR			RID	8h		88h	
BIST	HDR	PLAT	CLSR	Ch		8Ch	
				10h		90h	
				14h		94h	
				18h		98h	
				1Ch		9Ch	
				20h		A0h	
				24h		A4h	
				28h		A8h	
SDID		SVID		2Ch		ACh	
				30h		B0h	
				CAPPTR		34h	B4h
				38h		B8h	
				MAXLAT		MINGNT	INTPIN
				40h		C0h	
				44h		C4h	
				48h		C8h	
				4Ch		CCh	
				50h		D0h	
				54h		QPIMISCSTAT	D4h
				58h			D8h
				5Ch			DCh
				60h			E0h
				64h			E4h
				68h			E8h
6Ch	ECh						
70h	F0h						




	74h		F4h
	78h		F8h
	7Ch		FCh

3.1.1 QPIMISCSTAT: Intel QPI Misc Status

This is a status register for Common logic in Intel QPI. It is shared between Intel QPI 0 and Intel QPI 1.

QPIMISCSTAT			
Bus: 1		Device: 8	Function: 0
		Offset: D4	
Bit	Attr	Default	Description
2:0	RO-V	011b	Intel QPI Rate This reflects the current Intel QPI rate setting into the PLL. 011 - 6.4 GT/s 101 - 8 GT/s 111 - 9.6 GT/s other - Reserved

§




4 Processor Utility Box (UBOX) Registers

The UBOX is the piece of processor logic that deals with the non mainstream flows in the system. This includes transactions like the register accesses, interrupt flows, lock flows and events. In addition, the UBOX houses coordination for the performance architecture, and also houses scratchpad and semaphore registers.

4.1 Device 16 Function 5

DID		VID		0h		80h
PCISTS		PCICMD		4h		84h
CCR			RID	8h		88h
BIST	HDR	PLAT	CLSR	Ch		8Ch
				10h		90h
				14h		94h
				18h		98h
				1Ch		9Ch
				20h		A0h
				24h		A4h
				28h		A8h
SDID		SVID		2Ch		ACH
				30h		B0h
				CAPPTR		34h
				38h		B8h
				MAXLAT		MINGNT
CPUNODEID				40h		C0h
				44h		C4h
				IntControl		48h
				4Ch		CCh
				50h		D0h
GIDNIDMAP				54h		D4h
				58h		D8h
				5Ch		DCh
				60h		E0h
UBOXErrSts				64h		E4h
				68h		E8h
				6Ch		ECh
				70h		F0h



	74h		F4h
	78h		F8h
	7Ch		FCh

4.1.1 CPUNODEID

Node ID Configuration Register

Type: CFG		Port ID: N/A	
Bus: 1		Device: 16	
Offset: 0x40		Function: 5	
Bit	Attr	Default	Description
12:10	RW_LB	0x0	NodeID of the legacy socket(LgcNodeID): NodeID of the legacy socket.
7:5	RW_LB	0x0	NodeID of the lock master(LockNodeID): NodeID of the lock master.
2:0	RW_LB	0x0	NodeID of the local register(LclNodeID): NodeID of the local Socket.

4.1.2 IntControl

Interrupt Configuration Register

Type: CFG		Port ID: N/A	
Bus: 1		Device: 16	
Offset: 0x48		Function: 5	
Bit	Attr	Default	Description
18:18	RW_LB	0x0	IA32 Logical Flat or Cluster Mode Override Enable(LogFlatClustOvrEn): 0: IA32 Logical Flat or Cluster Mode bit is locked as Read only bit. 1: IA32 Logical Flat or Cluster Mode bit may be written by SW, values written by xTPR update are ignored. For one time override of the IA-32 Logical Flat or Cluster Mode value, return this bit to it's default state after the bit is changed. Leaving this bit as '1' will prevent automatic update of the filter.
17:17	RW_LBV	0x0	IA32 Logical Flat or Cluster Mode(LogFltClustMod): Set by BIOS to indicate if the OS is running logical flat or logical cluster mode. This bit can also be updated by IntPrioUpd messages. This bit reflects the setup of the filter at any given time. 0 - flat, 1 - cluster.
16:16	RW_LB	0x0	Cluster Check Sampling Mode(ClastChkSmpMod): 0: Disable checking for Logical_APICID[31:0] being non-zero when sampling flat cluster mode bit in the IntPrioUpd message as part of setting bit 1 in this register 1: Enable the above checking



Type: CFG		Port ID: N/A	
Bus: 1		Device: 16	
Offset: 0x48		Function: 5	
Bit	Attr	Default	Description
10:8	RW_LB	0x0	Vecor Based Hashe Mode Control(HashModCtr): Indicates the hash mode control for the interrupt control. Select the hush function for the Vector based Hash Mode interrupt redirection control: 000 select bits 7: 4 / 5: 4 for vector cluster / flat algorithm 001 select bits 6: 3 / 4: 3 010 select bits 4: 1 / 2: 1 011 select bits 3: 0 / 1: 0 Other - reserved
6:4	RW_LB	0x0	Redirection Mode Select for Logical Interrupts(RdrModSel): Selects the redirection mode used for MSI interrupts with lowest-priority delivery mode. The following schemes are used: 000: Fixed Priority 001: Round-robin 010: Interrupt vector hash
1:1	RW_LB	0x0	Force to X2 APIC Mode(ForceX2APIC): Write: 1: Forces the system to move into X2APIC Mode. 0: No affect
0:0	RW_LB	0x1	Extended APIC Enable(xApicEn): 1: Extended XAPIC configuration. This bit can be written directly, and can also be updated using XTPR messages.

4.1.3 GIDNIDMAP

Node ID Mapping Register. Mapping between group id and NodeID

Type: CFG		Port ID: N/A	
Bus: 1		Device: 16	
Offset: 0x54		Function: 5	
Bit	Attr	Default	Description
23:21	RW_LB	0x0	NodeID 7(NodeId7): NodeID for group id 7
20:18	RW_LB	0x0	NodeID 6(NodeId6): NodeID for group 6
17:15	RW_LB	0x0	NodeID 5(NodeId5): NodeID for group 5
14:12	RW_LB	0x0	NodeID 4(NodeId4): NodeID for group id 4
11:9	RW_LB	0x0	NodeID 3(NodeId3): NodeID for group 3
8:6	RW_LB	0x0	NodeID 2(NodeId2): NodeID for group Id 2
5:3	RW_LB	0x0	NodeID 1(NodeId1): NodeID for group Id 1
2:0	RW_LB	0x0	NodeID 0(NodeId0): NodeID for group 0



4.1.4 UBOXErrSts

This is error status register in the UBOX and covers most of the interrupt related errors.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 16	
Offset: 0x64		Function: 5	
Bit	Attr	Default	Description
16:16	RW_V	0x0	Intel SMI delivery valid(SMI_delivery_valid): Intel SMI interrupt delivery status valid, write 1'b1 to clear valid status
7:7	RWS_V	0x0	MasterLock Timeout received by UBOX(MasterLockTimeOut): Master Lock Timeout received by UBOX
6:6	RWS_V	0x0	Intel SMI Timeout received by UBOX(SMITimeOut): Intel SMI Timeout received by UBOX
5:5	RWS_V	0x0	MMCFG Write Address Misalignment received by UBOX(CFGWrAddrMisAligned): MMCFG Write Address Misalignment received by UBOX
4:4	RWS_V	0x0	MMCFG Read Address Misalignment received by UBOX(CFGRdAddrMisAligned): MMCFG Read Address Misalignment received by UBOX
3:3	RWS_V	0x0	Unsupported Opcode received by UBOX(UnsupportedOpcode): Unsupported opcode received by UBOX
2:2	RWS_V	0x0	Poison was received by UBOX(PoisonRsvd): UBOX received a poisoned transaction
1:1	RWS_V	0x0	Intel SMI source iMC(SMISrciMC): Intel SMI is caused due to an indication from the iMC
0:0	RWS_V	0x0	Intel SMI is caused due to a locally generated UMC(SMISrcUMC): This is a bit that indicates that an Intel SMI was caused due to a locally generated UMC.

4.2 Device 16 Function 7

DID		VID		0h		80h
PCISTS		PCICMD		4h		84h
CCR			RID	8h		88h
BIST	HDR	PLAT	CLSR	Ch		8Ch
				10h		90h
				14h		94h
				18h		98h
				1Ch		9Ch
				20h		A0h
				24h		A4h
				28h		A8h
SDID		SVID		2Ch		ACh
				30h		B0h
				CAPPTR		34h
				38h		B8h
				MAXLAT		MINGNT
				40h		C0h



	44h		C4h
	48h		C8h
	4Ch		CCh
	50h	CPUBUSNO	D0h
	54h		D4h
	58h	SMICtrl	D8h
	5Ch		DCh
	60h		E0h
	64h		E4h
	68h		E8h
	6Ch		ECh
	70h		F0h
	74h		F4h
	78h		F8h
	7Ch		FCh

4.2.1 CPUBUSNO

Bus Number Configuration for the Intel® Xeon® processor E5 v4 product family.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 16	
Offset: 0xd0		Function: 7	
Bit	Attr	Default	Description
31:31	RW_LB	0x0	Valid: Indicates whether the bus numbers have been initialized or not
15:8	RW_LB	0x0	CPU Bus Number 1(CPUBUSNO1): Bus Number for non IIO devices in the uncore in the processor.
7:0	RW_LB	0x0	CPU Bus Number 0(CPUBUSNO0): Bus Number for IIO devices in the processor.

4.2.2 SMICtrl

Intel SMI generation control


Type: CFG		Port ID: N/A	
Bus: 1		Device: 16	
Offset: 0xd8		Function: 7	
Bit	Attr	Default	Description
28:28	RW_LB	0x0	Disable Generation of Intel SMI from CSMI from MsgCh(SMIDis4): Disable Generation of Intel SMI from CSMI from MsgCh
27:27	RW_LB	0x0	Disable Generation of Intel SMI for new Ubox erros(SMIDis3): Disable generation of Intel SMI from message channel
26:26	RW_LB	0x1	Disable Generation of Intel SMI for new Ubox erros(SMIDis2): Disable generation of Intel SMI for lock timeout, cfg write mis-align access, and cfg read mis-align access.



Processor Utility Box (UBOX) Registers

Type: CFG		Port ID: N/A	
Bus: 1		Device: 16	
Offset: 0xd8		Function: 7	
Bit	Attr	Default	Description
25:25	RW_LB	0x0	Disable Generation of Intel SMI (all)(SMIDis): Disable generation of Intel SMI
24:24	RW_LB	0x0	UMC Intel SMI Enable (UMCSMIEn): This is the enable bit that enables Intel SMI generation due to a UMC 1 - Generate Intel SMI after the threshold counter expires. 0 - Disable generation of Intel SMI
19:0	RW_LB	0x0	Intel SMI generation threshold (Threshold): This is the countdown that happens in the hardware before an Intel SMI is generated due to a UMC.

§



5 Power Controller Unit (PCU) Registers

The Power Controller Unit (PCU) is a dedicated controller that provides power and thermal management for the processor.

All Intel® Xeon® Processor E5 v4 Processors implement a PCU.

- The Intel® Xeon® Processor E5-1600/2600 and E5-4600 v4 implements 4 channels for DIMMs.

5.1 Device 30 Function 0

DID		VID		0h		80h	
PCISTS		PCICMD		4h	PACKAGE_POWER_SKU	84h	
CCR			RID	8h		88h	
BIST	HDR	PLAT	CLSR	Ch	PACKAGE_POWER_SKU_UNIT	8Ch	
				10h	PACKAGE_ENERGY_STATUS	90h	
				14h		94h	
				18h		98h	
				1Ch		9Ch	
				20h		A0h	
				24h		A4h	
				28h		A8h	
SDID		SVID		2Ch		ACh	
				30h		B0h	
				CAPPTR		34h	B4h
						38h	B8h
MAXLAT	MINGNT	INTPIN	INTL	3Ch		BCh	
				40h		C0h	
				44h		C4h	
				48h	Package_Temperature	C8h	
				4Ch		CCh	
				50h		D0h	
				54h	PCU_REFERENCE_CLOCK	D4h	
				58h		D8h	
5Ch		DCh					
MEM_TRML_TEMPERATURE_REPORT				60h		E0h	
MEM_ACCUMULATED_BW_CH_0				64h	TEMPERATURE_TARGET	E4h	
MEM_ACCUMULATED_BW_CH_1				68h		E8h	
MEM_ACCUMULATED_BW_CH_2				6Ch		ECh	
MEM_ACCUMULATED_BW_CH_3				70h		F0h	



		74h			F4h
		78h			F8h
		7Ch			FCh
DID	VID	0h			80h

5.1.1 MEM_TRML_TEMPERATURE_REPORT

This register is used to report the thermal status of the memory.


The channel max temperature field is used to report the maximal temperature of all ranks.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x60		Function: 0	
Bit	Attr	Default	Description
31:24	RO_V	0x0	Channel 3 Maximum Temperature(Channel3_Max_Temperature): Temperature in Degrees (C).
23:16	RO_V	0x0	Channel 2 Maximum Temperature(Channel2_Max_Temperature): Temperature in Degrees (C).
15:8	RO_V	0x0	Channel 1 Maximum Temperature(Channel1_Max_Temperature): Temperature in Degrees (C).
7:0	RO_V	0x0	Channel 0 Maximum Temperature(Channel0_Max_Temperature): Temperature in Degrees (C).

5.1.2 MEM_ACCUMULATED_BW_CH_[0:3]

This register contains a measurement proportional to the weighted DRAM BW for the channel including all ranks. The weights are configured in the memory controller channel register PM_CMD_PWR.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x64, 0x68, 0x6c, 0x70		Function: 0	
Bit	Attr	Default	Description
31:0	RO_V	0x0	Data(DATA): The weighted BW value is calculated by the memory controller based on the following formula: $\text{NumPrecharge} * \text{PM_CMD_PWR}[\text{PWR_RAS_PRE}] +$ $\text{NumReads} * \text{PM_CMD_PWR}[\text{PWR_CAS_R}] +$ $\text{NumWrites} * \text{PM_CMD_PWR}[\text{PWR_CAS_W}]$



5.1.3 PACKAGE_POWER_SKU


Defines allowed SKU power and timing parameters.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x84		Function: 0	
Bit	Attr	Default	Description
54:48	RO_V	0x12	Maximal Time Window(PKG_MAX_WIN): The maximal time window allowed for the SKU. Higher values will be clamped to this value. $x = \text{PKG_MAX_WIN}[54:53]$ $y = \text{PKG_MAX_WIN}[52:48]$ The timing interval window is Floating Point number given by $1.x * \text{power}(2,y)$. The unit of measurement is defined in PACKAGE_POWER_SKU_UNIT_MSR[TIME_UNIT].
46:32	RO_V	0x240	Maximal Package Power(PKG_MAX_PWR): The maximal package power setting allowed for the SKU. Higher values will be clamped to this value. The maximum setting is typical not guaranteed. The units for this value are defined in PACKAGE_POWER_SKU_UNIT_MSR[PWR_UNIT].
30:16	RO_V	0x60	Minimal Package Power(PKG_MIN_PWR): The minimal package power setting allowed for the SKU. Lower values will be clamped to this value. The minimum setting is typical not guaranteed. The units for this value are defined in PACKAGE_POWER_SKU_UNIT_MSR[PWR_UNIT].
14:0	RO_V	0x118	TDP Package Power(PKG_TDP): The TDP package power setting allowed for the SKU. The TDP setting is typical not guaranteed. The units for this value are defined in PACKAGE_POWER_SKU_UNIT_MSR[PWR_UNIT].

5.1.4 PACKAGE_POWER_SKU_UNIT

Defines units for calculating SKU power and timing parameters.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x8c		Function: 0	
Bit	Attr	Default	Description
19:16	RO_V	0xa	Time Unit(TIME_UNIT): Time Units used for power control registers. The actual unit value is calculated by $1 / \text{Power}(2, \text{TIME_UNIT})$ second. The default value of 0Ah corresponds to 976 usec.
12:8	RO_V	0xe	Energy Units(ENERGY_UNIT): Energy Units used for power control registers. The actual unit value is calculated by $1 / \text{Power}(2, \text{ENERGY_UNIT})$ J.
3:0	RO_V	0x3	Power Units(PWR_UNIT): Power Units used for power control registers. The actual unit value is calculated by $1 / \text{Power}(2, \text{PWR_UNIT})$ W. The default value of 0011b corresponds to 18 W.



5.1.5 PACKAGE_ENERGY_STATUS

Package energy consumed by the core and uncore. The counter will wrap around and continue counting when it reaches its limit.

The energy status is reported in units which are defined in PACKAGE_POWER_SKU_UNIT_MSR[ENERGY_UNIT].

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x90		Function: 0	
Bit	Attr	Default	Description
31:0	RO_V	0x0	Energy Value(DATA): Energy Value

5.1.6 Package_Temperature

Package temperature in degrees (C). This field is updated by FW.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0xc8		Function: 0	
Bit	Attr	Default	Description
7:0	RO_V	0x0	Temperature(DATA): Package temperature in degrees (C).

5.1.7 TEMPERATURE_TARGET

Legacy register holding temperature related constants for Platform use. This register is updated by FW.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0xe4		Function: 0	
Bit	Attr	Default	Description
27:24	RO_V	0x0	Max TCC Offset (MAX_TCC_OFFSET): Temperature offset in degrees (C) from the Processor Hot. Used for throttling temperature. Will not impact temperature reading. If offset is allowed and set, the throttle will occur and reported at lower than Processor Hot.
23:16	RO_V	0x0	Thermal Monitor Reference Temperature(REF_TEMP): This field indicates the maximum junction temperature, also referred to as the throttle temperature, TCC activation temperature or prohot temperature. This is the temperature at which the Thermal Monitor is activated.
15:8	RO_V	0x0	Fan Temperature target offset(FAN_TEMP_TARGET_OFST): Fan Temperature target offset also known as T-Control. Indicates the relative offset from the Thermal Monitor Trip Temperature at which fans should be engaged.



5.2 Device 30 Function 1

DID		VID		0h		80h
PCISTS		PCICMD		4h		84h
CCR			RID	8h		88h
BIST	HDR	PLAT	CLSR	Ch		8Ch
				10h		90h
				14h		94h
				18h		98h
				1Ch		9Ch
				20h		A0h
				24h		CSR_DESIRED_CORES
				28h	A8h	
SDID		SVID		2Ch	ACh	
				30h	B0h	
				CAPPTR	34h	B4h
				38h	B8h	
MAXLAT	MINGNT	INTPIN	INTL	3Ch	BCh	
				40h	C0h	
				44h	C4h	
				48h	C8h	
				4Ch	CCh	
				50h	D0h	
				54h	D4h	
				58h	D8h	
				5Ch	DCh	
				60h	E0h	
				64h	E4h	
SSKPD				68h	E8h	
				6Ch	ECh	
				70h	F0h	
C2C3TT				74h	F4h	
				78h	F8h	
				7Ch	FCh	

5.2.1 SSKPD

Sticky Scratchpad Data.

This register holds 64 writable bits with no functionality behind them.



Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x68		Function: 1	
Bit	Attr	Default	Description
63:0	RWS	0x0	Scratchpad Data(SKPD): 4 WORDs of data storage.

5.2.2 C2C3TT

C2 to C3 Transition Timer. BIOS can update this value during run-time. Unit for this register is usec with a range of 0-4095 us.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x74		Function: 1	
Bit	Attr	Default	Description
11:0	RW	0x32	Pop Down Initialization Value(PPDN_INIT): Value in micro-seconds.

5.2.3 CSR_DESIRED_CORES

Number of cores/threads BIOS wants to exist on the next reset. A processor reset must be used for this register to take effect. Note, programming this register to a value higher than the product has cores should not be done.

This register is reset only by PWRGOOD.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0xa4		Function: 1	
Bit	Attr	Default	Description
31:31	RWS_KL	0x0	Lock(LOCK): Once written to a '1', changes to this register cannot be done. Cleared only by a power-on reset
30:30	RWS_L	0x0	SMT Disable (SMT_DISABLE): Disable simultaneous multithreading in all cores if this bit is set to '1'.
23:0	RWS_L	0x0	Cores Off Mask (CORE_OFF_MASK): BIOS will set this bit to request that the matching core should not be activated coming out of reset. The default value of this registers means that all cores are enabled. Restrictions: At least one core needs to be left active. Otherwise, FW will ignore the setting altogether.

5.3 Device 30 Function 2

DID	VID	0h		80h
PCISTS	PCICMD	4h		84h
CCR	RID	8h	PACKAGE_RAPL_PERF_STATUS	88h



BIST	HDR	PLAT	CLSR	Ch		8Ch	
				10h	DRAM_POWER_INFO	90h	
				14h		94h	
				18h			98h
				1Ch			9Ch
				20h	DRAM_ENERGY_STATUS	A0h	
				24h		A4h	
				28h	DRAM_ENERGY_STATUS_CH0	A8h	
SDID		SVID		2Ch		ACh	
				30h	DRAM_ENERGY_STATUS_CH1	B0h	
				CAPPTR		34h	B4h
				38h	DRAM_ENERGY_STATUS_CH2	B8h	
MAXLAT	MINGNT	INTPIN	INTL	3Ch		BCh	
				40h	DRAM_ENERGY_STATUS_CH3	C0h	
				44h		C4h	
				48h			C8h
				4Ch			CCh
				50h			D0h
				54h			D4h
				58h	DRAM_RAPL_PERF_STATUS	D8h	
				5Ch		DCh	
				60h			E0h
				64h			E4h
				68h			E8h
				6Ch	MCA_ERR_SRC_LOG	ECh	
				70h			F0h
				74h			F4h
				78h	THERMTRIP_CONFIG	F8h	
7Ch			FCh				


5.3.1 PACKAGE_RAPL_PERF_STATUS

This register is used to report Package Power limit violations.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x88		Function: 2	
Bit	Attr	Default	Description
31:0	RO_V	0x0	Power Limit Throttle Counter (PWR_LIMIT_THROTTLE_CTR): Reports the number of times the Power limiting algorithm had to clip the power limit due to hitting the lowest power state available. Accumulated PACKAGE throttled time.

5.3.2 DRAM_POWER_INFO

Defines allowed DRAM power and timing parameters.



Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x90		Function: 2	
Bit	Attr	Default	Description
63:63	RW_KL	0x0	Lock: Lock bit to lock the Register
54:48	RW_L	0x28	Maximal Time Window (DRAM_MAX_WIN): The maximal time window allowed for the DRAM. Higher values will be clamped to this value. $x = \text{PKG_MAX_WIN}[54:53]$ $y = \text{PKG_MAX_WIN}[52:48]$ The timing interval window is Floating Point number given by $1.x * \text{power}(2,y)$. ENERGY_UNIT for DRAM domain is 15.3uJ.
46:32	RW_L	0x258	Maximal Package Power (DRAM_MAX_PWR): The maximal power setting allowed for DRAM. Higher values will be clamped to this value. The maximum setting is typical (not guaranteed). ENERGY_UNIT for DRAM domain is 15.3uJ.
30:16	RW_L	0x78	Minimal DRAM Power (DRAM_MIN_PWR): The minimal power setting allowed for DRAM. Lower values will be clamped to this value. The minimum setting is typical (not guaranteed). ENERGY_UNIT for DRAM domain is 15.3uJ.
14:0	RW_L	0x118	Spec DRAM Power (DRAM_TDP): The Spec power allowed for DRAM. The TDP setting is typical (not guaranteed). ENERGY_UNIT for DRAM domain is 15.3uJ.

5.3.3 DRAM_ENERGY_STATUS

DRAM energy consumed by all the DIMMS in all the Channels. The counter will wrap around and continue counting when it reaches its limit.

ENERGY_UNIT for DRAM domain is 15.3uJ.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0xa0		Function: 2	
Bit	Attr	Default	Description
31:0	RO_V	0x0	Energy Value(DATA): Energy of the DDR plane. This counter rolls over upon an overflow and continues counting. To determine the power consumed by the DDR, BIOS/SW can read the counter at a specific interval and divide the difference by the interval time. $\text{Power} = [\text{Value}(t + x) - \text{Value}(t)]/x$

5.3.4 DRAM_ENERGY_STATUS_CH[0:3]

DRAM energy consumed by all the DIMMS in ChannelX (X = 0, 1, 2, 3). The counter will wrap around and continue counting when it reaches its limit.



Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0xa8, 0xb0, 0xb8, 0xc0		Function: 2	
Bit	Attr	Default	Description
31:0	RO_V	0x0	Energy Value(DATA): Energy Value

5.3.5 DRAM_RAPL_PERF_STATUS

This register is used to report DRAM Plane Power limit violations.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0xd8		Function: 2	
Bit	Attr	Default	Description
31:0	RO_V	0x0	Power Limit Throttle Counter (PWR_LIMIT_THROTTLE_CTR): Reports the number of times the Power limiting algorithm had to clip the power limit due to hitting the lowest power state available. Accumulated DRAM throttled time.

5.3.6 MCA_ERR_SRC_LOG

MCA Error Source Log.

MC Source Log is used by the PCU to log the error sources. This register is initialized to zeros during reset. The PCU will set the relevant bits when the condition they represent appears. The PCU never clears the registers-the UBox or off-die entities should clear them when they are consumed, unless their processing involves taking down the platform.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0xec		Function: 2	
Bit	Attr	Default	Description
31:31	RWS_V	0x0	CATERR: External error: The package asserted CATERR# for any reason. It is or(bit 30, bit29); functions as a Valid bit for the other two package conditions. It has no effect when a local core is associated with the error.
30:30	RWS_V	0x0	IERR: External error: The package asserted IERR.
29:29	RWS_V	0x0	MCERR: External error: The package asserted MCERR.
23:23	RWS_V	0x0	MSMI: External error: The package observed MSMI# (for any reason). It is or(bit 22, bit21); functions as a Valid bit for the other two package conditions. It has no effect when a local core is associated with the error.
22:22	RWS_V	0x0	MSMI_IERR: External error: The package observed MSMI_IERR.
21:21	RWS_V	0x0	MSMI_MCERR: External error: The package observed MSMI_MCERR.




5.3.7 THERMTRIP_CONFIG

This register is used to configure whether the Thermtrip signal only carries the processor Trip information, or does it carry the Mem trip information as well. The register will be used by HW to enable ORing of the memtrip info into the thermtrip OR tree.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0xf8		Function: 2	
Bit	Attr	Default	Description
0:0	RW	0x0	Enable MEM Trip(EN_MEMTRIP): If set to 1, PCU will OR in the MEMtrip information into the ThermTrip OR Tree. If set to 0, PCU will ignore the MEMtrip information and ThermTrip will just have the processor indication.

5.4 Device 30 Function 3

DID		VID		0h	CAP_HDR	80h				
PCISTS		PCICMD		4h	CAPID0	84h				
CCR			RID	8h	CAPID1	88h				
BIST	HDR	PLAT	CLSR	Ch	CAPID2	8Ch				
				10h	CAPID3	90h				
				14h	CAPID4	94h				
				18h	CAPID5	98h				
				1Ch	CAPID6	9Ch				
				20h		A0h				
				24h		A4h				
				28h		A8h				
SDID		SVID		2Ch		ACh				
				30h	SMT_CONTROL	B0h				
				CAPPTR		34h	RESOLVED_CORES	B4h		
				38h		B8h				
				MAXLAT		MINGNT	INTPIN	INTL	3Ch	BCh
				40h		C0h				
				44h		C4h				
				48h		C8h				
				4Ch		CCh				
				50h		D0h				
				54h		D4h				
				58h		D8h				
				5Ch		DCh				
				60h		E0h				
				64h		E4h				
				68h		E8h				
				6Ch		ECh				
				70h		F0h				



	74h		F4h
	78h		F8h
	7Ch		FCh

Note: The CSR located at offset in Device 30, Function 3, Offset 0x10 is not a Configuration Space Header and SW should not treat it as such.

5.4.1 CAP_HDR


This register is a Capability Header. It enumerates the CAPID registers available, and points to the next CAP_PTR.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x80		Function: 3	
Bit	Attr	Default	Description
27:24	RO_FW	0x1	CAPID_Version: This field has the value 0001b to identify the first revision of the CAPID register definition.
23:16	RO_FW	0x18	CAPID_Length: This field indicates the structure length including the header in Bytes.
15:8	RO_FW	0x0	Next_Cap_Ptr: This field is hardwired to 00h indicating the end of the capabilities linked list.
7:0	RO_FW	0x9	CAP_ID: This field has the value 1001b to identify the CAPID assigned by the PCI SIG for vendor dependent capability pointers.

5.4.2 CAPIDO

This register is a Capability Register used to expose feature support for BIOS use. Default value varies base on SKU.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x84		Function: 3	
Bit	Attr	Default	Description
31:31	RO_FW	0x0	PCLMULQ_DIS: PCLMULQ instruction disabled.
29:29	RO_FW	0x0	PECI_EN: PECI to the Processor enabled.
26:26	RO_FW	0x0	GSSE256_DIS: GSSE instructions disabled.
23:23	RO_FW	0x0	AES_DIS: AES (Advanced Encryption Standard) disabled.
20:20	RO_FW	0x0	LT_SX_EN: Intel TXT and FIT-boot enabled.
19:19	RO_FW	0x0	LT_PRODUCTION: Intel TXT enabled.
18:18	RO_FW	0x0	SMX_DIS: Intel TXT enabled.




Power Controller Unit (PCU) Registers

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x84		Function: 3	
Bit	Attr	Default	Description
17:17	RO_FW	0x0	VMX_DIS: VMX (Virtual-Machine Extensions) disabled.
15:15	RO_FW	0x0	VT_X3_EN: VT-x3 (Intel® Virtualization Technology) enabled.
12:12	RO_FW	0x0	HT_DIS: Multithreading disabled.
8:8	RO_FW	0x0	PRG_TDP_LIM_EN: Usage of TURBO_POWER_LIMIT MSRs enabled.
4:4	RO_FW	0x0	DE_SKTR1_EX: Set to 0 for Intel Xeon processor E5 v4 Product Family.
3:3	RO_FW	0x0	DE_SKTR_EP4S: Set to 0 for Intel Xeon Processor E5-2600 v4 Product Families. Set to 1 for Intel Xeon Processor E5-4600 v4 Product Families.
2:2	RO_FW	0x0	DE_SKTR_EP2S: Indicates that device is a 2S SKU, independent of package.
0:0	RO_FW	0x0	DE_SKT2_UP: Indicates that device is a UP SKU, independent of package.

5.4.3 CAPID1

This register is a Capability Register used to expose feature support for BIOS use. Default value varies base on SKU.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x88		Function: 3	
Bit	Attr	Default	Description
31:31	RO_FW	0x0	DIS_MEM_MIRROR: Disable memory channel mirroring mode. In the mirroring mode, the server maintains two identical copies of all data in memory. The contents of branch 0 (containing channel 0/1) is duplicated in the DIMMs of branch 1 (containing channel 2/3). In the event of an uncorrectable error in one of the copies, the system can retrieve the mirrored copy of the data. The use of memory mirroring means that only half of the installed memory is available to the operating system.
30:30	RO_FW	0x0	DIS_MEM_LT_SUPPORT: Intel TXT support disabled.
29:26	RO_FW	0x0	DMFC: This field controls which values may be written to the Memory Frequency Select field 6:4 of the Clocking Configuration registers. Any attempt to write an unsupported value will be ignored. [3:3] - If set, over-clocking is supported and bits 2:0 are ignored. [2:0] - Maximum allowed memory frequency. 3b110 - up to DDR-1333 (5 x 266) 3b101 - up to DDR-1600 (6 x 266) 3b100 - up to DDR-1866 (7 x 266) 3b011 - up to DDR-2133 (8 x 266) 3b010 - up to DDR-2400 (9 x 266) All others reserved



Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x88		Function: 3	
Bit	Attr	Default	Description
25:23	RO_FW	0x0	MEM_PA_SIZE: Physical address size supported in the core low two bits (uncore is 44 by default) 000: 46 010: 44 101: 36 110: 40 111: 39 reserved
8:8	RO_FW	0x0	rsvd
7:7	RO_FW	0x0	X2APIC_EN: Extended APIC support enabled. When set the enables the support of x2APIC (Extended APIC) in the core and uncore.
5:5	RO_FW	0x0	PWRBITS_DIS: 0b Power features activated during reset. 1b Power features (i.e. clock gating) are not activated.
4:4	RO_FW	0x0	GV3_DIS: Intel SpeedStep® Technology disabled. Does not allow for the writing of the IA32_PERF_CTL register in order to change ratios.
1:1	RO_FW	0x0	CORE_RAS_EN: Data Poisoning, MCA recovery enabled.
0:0	RO_FW	0x0	DCA_EN: DCA (Direct Cache Access) enabled.

5.4.4 CAPID2

This register is a Capability Register used to expose feature support for BIOS use. Default value varies base on SKU. Default value varies base on SKU.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x8c		Function: 3	
Bit	Attr	Default	Description
29:25	RO_FW	0x0	QPI_ALLOWED_CFCLK_RATIO_DIS: Allowed Intel QPI link speeds. bit 8 = 6.4GT/s bit 10 = 8.0GT/s bit 12 = 9.6GT/s
24:24	RO_FW	0x0	QPI_LINK1_DIS: Intel QPI link 1 disabled.
23:23	RO_FW	0x0	QPI_LINK0_DIS: Intel QPI link 0 disabled.
19:19	RO_FW	0x0	PCIE_DISNTB: NTB (Non Transparent Bridge) support disabled.
18:18	RO_FW	0x0	PCIE_DISROL: Raid-on-load disabled.
17:17	RO_FW	0x0	PCIE_DISLTSX: Intel TXT disabled.



Power Controller Unit (PCU) Registers

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x8c		Function: 3	
Bit	Attr	Default	Description
16:16	RO_FW	0x0	PCIE_DISLT: Intel TXT disabled.
15:15	RO_FW	0x0	PCIE_DISPCIEG3: PCIe Gen 3 disabled.
14:14	RO_FW	0x0	PCIE_DISDMA: DMA engine and supporting functionality disabled.
13:13	RO_FW	0x0	PCIE_DISDMI: DMI2 interface disabled.
12:3	RO_FW	0x0	PCIE_DISXPDEV: Specific PCIe port disabled.
2:1	RO_FW	0x0	PCIE_DISx16: PCIe x16 ports disabled (limit to x8's only).
0:0	RO_FW	0x0	PCIE_DISWS: WS features such as graphics cards in PCIe slots disabled.

5.4.5 CAPID3

This register is a Capability Register used to expose feature support for BIOS use. Default value varies base on SKU.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x90		Function: 3	
Bit	Attr	Default	Description
30:30	RO_FW	0x0	DISABLE_MEM_DDR4: DDR4 disabled.
29:24	RO_FW	0x0	MC2GD: Bit0: 1.35V DDR3L LVDDR disable
22:22	RO_FW	0x0	DISABLE_SMBUS_WRT: SMBUS write capability disable control. When set, SMBus write is disabled.
21:21	RO_FW	0x0	DISABLE_ROL_OR_ADR: RAID-On-LOAD disable control. When set, memory ignores ADR event. Download may change the default value after reset de-assertion.
20:20	RO_FW	0x0	DISABLE_EXTENDED_ADDR_DIMM: Extended addressing DIMM disable control. When set, DIMM with extended addressing (MA[17/16] is forced to be zero when driving MA[17:16]).
19:19	RO_FW	0x0	DISABLE_EXTENDED_LATENCY_DIMM: Extended latency DIMM disable control. When set, DIMM with extended latency is forced to CAS to be less than or equal to 14.
18:18	RO_FW	0x0	DISABLE_PATROL_SCRUB: Patrol scrub disable control. When set, rank patrol scrub is disabled.
17:17	RO_FW	0x0	DISABLE_SPARING: Sparing disable control. When set, rank sparing is disabled.
16:16	RO_FW	0x0	DISABLE_LOCKSTEP: LOCKSTEP disable control. When set, channel lockstep operation is disabled.



Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x90		Function: 3	
Bit	Attr	Default	Description
15:15	RO_FW	0x0	DISABLE_CLTT: CLTT disable control. When set, CLTT support is disabled by disabling TSOD polling.
14:14	RO_FW	0x0	DISABLE_UDIMM: UDIMM disable control. When set, UDIMM support is disabled by disabling address bit swizzling.
13:13	RO_FW	0x0	DISABLE_RDIMM: RDIMM disable control. When set, RDIMM support is disabled.
12:12	RO_FW	0x0	DISABLE_3N: 3N disable control. When set, 3N mode under normal operation (excluding MRS) is disabled.
11:11	RO_FW	0x0	DISABLE_DIR: DIR disable control. When set, directory is disabled.
10:10	RO_FW	0x0	DISABLE_ECC: ECC disable control. When set, ECC is disabled.
9:9	RO_FW	0x0	DISABLE_OR_DIMM: QR DIMM disable control. When set, CS signals for QR-DIMM in slot 0-1 is disabled.
8:8	RO_FW	0x0	DISABLE_4GBIT_DDR3: 4 GB disable control. When set, the address decode to the corresponding 4 Gb mapping is disabled. Note: LR-DIMM's logical device density is also limited to 4 Gb when this is set.
7:7	RO_FW	0x0	DISABLE_8GBIT_DDR3: 8 Gb or higher disable control. When set, the address decode to the corresponding 8 Gb or higher mapping is disabled. Note: LR-DIMM's logical device density is also limited to 8 Gb when this is set.
5:5	RO_FW	0x0	DISABLE_3_DPC: 3 DPC disable control. When set, CS signals for DIMM slot 2 are disabled.
4:4	RO_FW	0x0	DISABLE_2_DPC: 2 DPC disable control. When set, CS signals for DIMM slot 1-2 (i.e. slots 0 is not disabled) are disabled.
3:0	RO_FW	0x0	CHN_DISABLE: Channel disable control. When set, the corresponding memory channel is disabled. 1 IMC and HA on SKU: • 0000 = Channels 0/1/2/3 enabled 2 IMC and HA on SKU: • 1100 = Channels 0/1/2/3 enabled (2 channels per IMC/HA)

5.4.6 CAPID4

This register is a Capability Register used to expose feature support for BIOS use. Default value varies base on SKU.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x94		Function: 3	
Bit	Attr	Default	Description
31:31	RO_FW	0x0	Disable DRAM Power Meter (DRAM_POWER_METER_DISABLE)



Power Controller Unit (PCU) Registers

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x94		Function: 3	
Bit	Attr	Default	Description
30:30	RO_FW	0x0	Disable DRAM RAPL(DRAM_RAPL_DISABLE)
26:26	RO_FW	0x0	EET_ENABLE: Energy efficient turbo enabled.
25:25	RO_FW	0x0	PCPS_DISABLE: Per-core P-state disabled.
24:24	RO_FW	0x0	UFS_DISABLE: UFS (Uncore Frequency Scaling) disabled.
19:19	RO_FW	0x0	ENHANCED_MCA_DIS: Enhanced MCA disabled
14:14	RO_FW	0x0	FMA_DIS: FMA (Floating point Multiple Add) instructions disabled.
8:6	RO_FW	0x0	PHYSICAL: Physical configuration of processor. 10:configuration 2; 01:configuration 1; 00:configuration 0;
5:4	RO_FW	0x0	PROD_TYPE — Product type 00 = Intel® Xeon® Processor E5 v4 Product Family



5.4.7 CAPID5

This register is a Capability Register used to expose feature support for BIOS use. Default value varies base on SKU.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x98		Function: 3	
Bit	Attr	Default	Description
31:31	RO_FW	0x0	COD_ENABLE: COD (Cluster on die) support enabled where the processor clusters cores to the near memory controller.
30:30	RO_FW	0x0	HITME_ENABLE : Directory Cache enabled.
27:27	RO_FW	0x0	QD_DIS Intel QuickData Technology disabled
26:26	RO_FW	0x0	ACSTATE_ENABLED Autonomous C-state control enabled.
25:25	RO_FW	0x0	HWP_ENABLE Hardware-Controlled Performance States (HWP) enabled
24:24	RO_FW	0x0	HSW_NI_DIS New instructions except LZCNT, TZCNT, MOVBE disabled which Intel® Xeon® Processor E5 v4 Product Family disabled.
23:0	RO_FW	0x0	LLC_SLICE_EN: Enabled Cbo slices (Cbo with enabled LLC slice).

5.4.8 CAPID6

This register is a Capability Register used to expose feature support for BIOS use. Default value varies base on SKU.

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0x9C		Function: 3	
Bit	Attr	Default	Description
30:30	RO_FW	0x0	IIO_LLCCONFIG_EN: IIO to allocate in LLC enabled.
29:29	RO_FW	0x0	DE_SKT_SECONDHA: Indicates when second Home Agent and Memory Controller is enabled.
23:0	RO_FW	0x0	LLC_IA_CORE_EN: Cores enabled on SKU of the Intel® Xeon® Processor E5 v4 Product Family.



5.4.9 SMT_CONTROL

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0xb0		Function: 3	
Bit	Attr	Default	Description
24:24	RO_V	0x0	SMT Capability: Enabled threads in the package. 0b 1 thread 1b 2 threads
9:8	RO_V	0x0	Thread Mask (THREAD_MASK): Thread Mask indicates which threads are enabled in the core. The LSB is the enable bit for Thread 0, whereas the MSB is the enable bit for Thread 1. This field is determined by FW based on CSR_DESIRED_CORES[SMT_DISABLE] and SKU capability.

5.4.10 RESOLVED_CORES

Type: CFG		Port ID: N/A	
Bus: 1		Device: 30	
Offset: 0xb4		Function: 3	
Bit	Attr	Default	Description
23:0	RO_V	0x0	CORE_MASK — The resolved IA core mask contains the functional (enabled in SKU) and non-defeatured IA cores. The mask is indexed by logical ID. It is normally contiguous, unless BIOS defeature is activated on a particular core. BSP and APIC IDs will be set by the processor based on this value. This field is determined by FW based on CSR_DESIRED_CORES[CORE_OFF_MASK].

§



6 Integrated I/O (IIO) Configuration Registers

The Integrated I/O (IIO) contains the DMI2 link, PCI Express* link, Intel QuickData Technology, IOAPIC, Intel VT-d and other related logic.

- The Intel® Xeon® Processor E5-1600/2600 and E5-4600 v4 includes a single x4 DMI2 link and 40 lanes of PCI Express 3.0. Device 0 is the DMI2 link, which can also operate as a PCI Express 2.0 x4 Root Port if not connected to a PCH. Device 1 is a x8 PCIe 3.0 Root Port. Device 2 is a x16 PCIe 3.0 Root Port. Device 3 is a x16 PCIe 3.0 Root Port.

6.1 Registers Overview

6.1.1 Configuration Registers (CSR)

There are two distinct CSR register spaces supported by the IIO Module.

The first one is the traditional PCI-defined configuration registers. These registers are accessed via the well known configuration transaction mechanism defined in the PCI specification and this uses the bus:device:function number concept to address a specific device's configuration space.


The second is via MMIO space for Intel® QuickData Technology, Intel VT-d, and I/OxAPIC runtime registers.

6.1.2 BDF: BAR# for Various MMIO BARs in IIO

This is needed for any entity trying to access MMIO registers in the IIO module over message channel.

Table 6-1. BDF:BAR# for Various MMIO BARs in IIO

BAR Name	B	D	F	BAR#
DMIRCBAR	DC	0	0	0
CB-BAR0	DC	4	0	0
CB-BAR1	DC	4	1	0
CB-BAR2	DC	4	2	0
CB-BAR3	DC	4	3	0
CB-BAR4	DC	4	4	0
CB-BAR5	DC	4	5	0
CB-BAR6	DC	4	6	0
CB-BAR7	DC	4	7	0
VT-d VTBAR	DC	5	0	0
I/OxAPIC-MBAR	DC	5	4	0
I/OxAPIC-ABAR	DC	5	4	1



6.1.3 Unimplemented Devices/Functions and Registers

If the IIO module receives a configuration access over message channel or directly via the JTAG mini-port, to a device/function or BAR# that does not exist in the IIO module, the IIO module will abort these accesses. Software should not attempt or rely on reads or writes to unimplemented registers or register bits.

6.1.4 PCI Vs. PCIe Device / Function

PCI devices/functions do NOT have a PCIe capability register set and do not decode offsets 100h and beyond. Accesses to 100h and beyond are master aborted by these devices. I/OxAPIC functions are PCI functions. All other functions in the IIO module are PCIe functions and these have a PCIe capability register set and also decode address offsets 100h and beyond.

6.2 Device 0 Function 0 DMI , Device 0 Function 0 PCIe, Device 1 Function 0-1, Device 2 Function 0-3 PCIe, Device 3 Function 0-3 PCIe

Device 0 Function 0 PCIe Mode - Port 0 (X4)

Device 1 - Port 1 (X8)

Device 2 - Port 2 (X16)

Device 3 - Port 3 (X16)

Table 6-2. Function Number of Active Root Ports in Port 1 (Dev#1) based on Port Bifurcation

Port Bifurcation	Function# of Active Root Port	
	7:4	3:0
x8	0	
x4x4	1	0

Table 6-3. Function Number of Active Root Ports in Port 2 (Dev#2) based on Port Bifurcation

Port Bifurcation	Function# of Active Root Port			
	15:12	11:8	7:4	3:0
x16	0			
x8x8	2		0	
x8x4x4	2		1	0
x4x4x8	3	2	0	
x4x4x4x4	3	2	1	0


Table 6-4. Function Number of Active Root Ports in Port 3(Dev#3) based on Port Bifurcation

Port Bifurcation	Function# of Active Root Port			
	15:12	11:8	7:4	3:0
x16	0			
x8x8	2		0	
x8x4x4	2		1	0
x4x4x8	3	2	0	
x4x4x4x4	3	2	1	0

Register Name	Offset	Size	Device 0 Function	Device 1 Function	Device 2 Function	Device 3 Function
vid	0x0	16	0	0-1	0 - 3	0 - 3
did	0x2	16	0	0-1	0 - 3	0 - 3
pcicmd	0x4	16	0	0-1	0 - 3	0 - 3
pcists	0x6	16	0	0-1	0 - 3	0 - 3
rid	0x8	8	0	0-1	0 - 3	0 - 3
ccr	0x9	24	0	0-1	0 - 3	0 - 3
clsr	0xc	8	0	0-1	0 - 3	0 - 3
plat	0xd	8	0	0-1	0 - 3	0 - 3
hdr	0xe	8	0	0-1	0 - 3	0 - 3
bist	0xf	8	0	0-1	0 - 3	0 - 3
pbus	0x18	8	0 (PCIe)	0-1	0 - 3	0 - 3
secbus	0x19	8	0 (PCIe)	0-1	0 - 3	0 - 3
subbus	0x1a	8	0 (PCIe)	0-1	0 - 3	0 - 3
iobas	0x1c	8	0 (PCIe)	0-1	0 - 3	0 - 3
iolim	0x1d	8	0 (PCIe)	0-1	0 - 3	0 - 3
secsts	0x1e	16	0 (PCIe)	0-1	0 - 3	0 - 3
mbas	0x20	16	0 (PCIe)	0-1	0 - 3	0 - 3
mlim	0x22	16	0 (PCIe)	0-1	0 - 3	0 - 3
pbas	0x24	16	0 (PCIe)	0-1	0 - 3	0 - 3
plim	0x26	16	0 (PCIe)	0-1	0 - 3	0 - 3
pbasu	0x28	32	0 (PCIe)	0-1	0 - 3	0 - 3
plimu	0x2c	32	0 (PCIe)	0-1	0 - 3	0 - 3
capptr	0x34	8	0	0-1	0 - 3	0 - 3
intl	0x3c	8	0	0-1	0 - 3	0 - 3
intpin	0x3d	8	0	0-1	0 - 3	0 - 3
bctrl	0x3e	16	0 (PCIe)	0-1	0 - 3	0 - 3
scapid	0x40	8	0 (PCIe)	0-1	0 - 3	0 - 3
snxptr	0x41	8	0 (PCIe)	0-1	0 - 3	0 - 3
svid	0x2c	16	0 (DMI2)			
svid	0x44	16	0 (PCIe)	0-1	0 - 3	0 - 3
sdid	0x2e	16	0 (DMI2)			



Register Name	Offset	Size	Device 0 Function	Device 1 Function	Device 2 Function	Device 3 Function
sdid	0x46	16	0 (PCIe)	0-1	0 - 3	0 - 3
dmircbar	0x50	32	0			
msicapid	0x60	8	0	0-1	0 - 3	0 - 3
msinxtptr	0x61	8	0	0-1	0 - 3	0 - 3
msimsgctl	0x62	16	0	0-1	0 - 3	0 - 3
msgadr	0x64	32	0	0-1	0 - 3	0 - 3
msgdat	0x68	32	0	0-1	0 - 3	0 - 3
msimsk	0x6c	32	0	0-1	0 - 3	0 - 3
msipending	0x70	32	0	0-1	0 - 3	0 - 3
pxpcapid	0x90	8	0	0-1	0 - 3	0 - 3
pxpnxtptr	0x91	8	0	0-1	0 - 3	0 - 3
pxpcap	0x92	16	0	0-1	0 - 3	0 - 3
devcap	0x94	32	0	0-1	0 - 3	0 - 3
devctrl	0xf0	16	0 (DMI2)			
devctrl	0x98	16	0 (PCIe)	0-1	0 - 3	0 - 3
devsts	0xf2	16	0 (DMI2)			
devsts	0x9a	16	0 (PCIe)	0-1	0 - 3	0 - 3
lnkcap	0x9c	32	0	0-1	0 - 3	0 - 3
lnkcon	0x1b0	16	0 (DMI2)			
lnkcon	0xa0	16	0 (PCIe)	0-1	0 - 3	0 - 3
lnksts	0x1b2	16	0 (DMI2)			
lnksts	0xa2	16	0 (PCIe)	0-1	0 - 3	0 - 3
sltcap	0xa4	32	0 (PCIe)	0-1	0 - 3	0 - 3
sltcon	0xa8	16	0 (PCIe)	0-1	0 - 3	0 - 3
sltsts	0xaa	16	0 (PCIe)	0-1	0 - 3	0 - 3
rootcon	0xac	16	0	0-1	0 - 3	0 - 3
rootcap	0xae	16	0	0-1	0 - 3	0 - 3
rootsts	0xb0	32	0 (PCIe)	0-1	0 - 3	0 - 3
devcap2	0xb4	32	0	0-1	0 - 3	0 - 3
devctrl2	0xf8	16	0 (DMI2)			
devctrl2	0xb8	16	0 (PCIe)	0-1	0 - 3	0 - 3
lnkcap2	0xbc	32	0	0-1	0 - 3	0 - 3
lnkcon2	0x1c0	16	0 (DMI2)			
lnkcon2	0xc0	16	0 (PCIe)	0-1	0 - 3	0 - 3
lnksts2	0x1c2	16	0 (DMI2)			
lnksts2	0xc2	16	0 (PCIe)	0-1	0 - 3	0 - 3
pmcap	0xe0	32	0	0-1	0 - 3	0 - 3
pmcsr	0xe4	32	0	0-1	0 - 3	0 - 3
xpreut_hdr_ext	0x100	32	0	0-1	0 - 3	0 - 3
xpreut_hdr_cap	0x104	32	0	0-1	0 - 3	0 - 3
xpreut_hdr_lef	0x108	32	0	0-1	0 - 3	0 - 3
acscaphdr	0x110	32	0 (PCIe)	0-1	0 - 3	0 - 3



Register Name	Offset	Size	Device 0 Function	Device 1 Function	Device 2 Function	Device 3 Function
acscap	0x114	16	0 (PCIe)	0-1	0 - 3	0 - 3
acsctrl	0x116	16	0 (PCIe)	0-1	0 - 3	0 - 3
apibase	0x140	16	0	0-1	0 - 3	0 - 3
apiclimit	0x142	16	0	0-1	0 - 3	0 - 3
vsecphdr	0x144	32	0 (DMI2)			
vshdr	0x148	32	0 (DMI2)			
errcaphdr	0x148	32	0 (PCIe)	0-1	0 - 3	0 - 3
uncerrsts	0x14c	32	0	0-1	0 - 3	0 - 3
uncerrmsk	0x150	32	0	0-1	0 - 3	0 - 3
uncerrsev	0x154	32	0	0-1	0 - 3	0 - 3
corerrsts	0x158	32	0	0-1	0 - 3	0 - 3
corerrmsk	0x15c	32	0	0-1	0 - 3	0 - 3
errcap	0x160	32	0	0-1	0 - 3	0 - 3
hdrlog0	0x164	32	0	0-1	0 - 3	0 - 3
hdrlog1	0x168	32	0	0-1	0 - 3	0 - 3
hdrlog2	0x16c	32	0	0-1	0 - 3	0 - 3
hdrlog3	0x170	32	0	0-1	0 - 3	0 - 3
rperrcmd	0x174	32	0	0-1	0 - 3	0 - 3
rperrsts	0x178	32	0	0-1	0 - 3	0 - 3
errsid	0x17c	32	0	0-1	0 - 3	0 - 3
perfctrlsts_0	0x180	32	0	0-1	0 - 3	0 - 3
perfctrlsts_1	0x184	32	0	0-1	0 - 3	0 - 3
miscctrlsts_0	0x188	32	0	0-1	0 - 3	0 - 3
miscctrlsts_1	0x18c	32	0	0-1	0 - 3	0 - 3
pcie_iou_bif_ctrl	0x190	16	0		0	0
dmictrl	0x1a0	64	0 (DMI2)			
dmists	0x1a8	32	0 (DMI2)			
ERRINJCAP	0x1d0	32	0	0-1	0 - 3	0 - 3
ERRINJHDR	0x1d4	32	0	0-1	0 - 3	0 - 3
ERRINJCON	0x1d8	16	0	0-1	0 - 3	0 - 3
ctoctrl	0x1e0	32	0	0-1	0 - 3	0 - 3
xpcorerrsts	0x200	32	0	0-1	0 - 3	0 - 3
xpcorerrmsk	0x204	32	0	0-1	0 - 3	0 - 3
xpuncerrsts	0x208	32	0	0-1	0 - 3	0 - 3
xpuncerrmsk	0x20c	32	0	0-1	0 - 3	0 - 3
xpuncerrsev	0x210	32	0	0-1	0 - 3	0 - 3
xpuncerrptr	0x214	8	0	0-1	0 - 3	0 - 3
uncedmask	0x218	32	0	0-1	0 - 3	0 - 3
coredmask	0x21c	32	0	0-1	0 - 3	0 - 3
rpcedmask	0x220	32	0	0-1	0 - 3	0 - 3
xpuncedmask	0x224	32	0	0-1	0 - 3	0 - 3
xpcoredmask	0x228	32	0	0-1	0 - 3	0 - 3



Register Name	Offset	Size	Device 0 Function	Device 1 Function	Device 2 Function	Device 3 Function
xpglberrsts	0x230	16	0	0-1	0 - 3	0 - 3
xpglberrptr	0x232	16	0	0-1	0 - 3	0 - 3
pxp2cap	0x250	32		0-1	0 - 3	0 - 3
lnkcon3	0x254	32		0-1	0 - 3	0 - 3
lnerrsts	0x258	32		0-1	0 - 3	0 - 3
ln0eq	0x25c	16		0-1	0 - 3	0 - 3
ln1eq	0x25e	16		0-1	0 - 3	0 - 3
ln2eq	0x260	16		0-1	0 - 3	0 - 3
ln3eq	0x262	16		0-1	0 - 3	0 - 3
ln4eq	0x264	16		0-1	0, 2	0, 2
ln5eq	0x266	16		0-1	0, 2	0, 2
ln6eq	0x268	16		0-1	0, 2	0, 2
ln7eq	0x26a	16		0-1	0, 2	0, 2
ln8eq	0x26c	16			0	0
ln9eq	0x26e	16			0	0
ln10eq	0x270	16			0	0
ln11eq	0x272	16			0	0
ln12eq	0x274	16			0	0
ln13eq	0x276	16			0	0
ln14eq	0x278	16			0	0
ln15eq	0x27a	16			0	0
mcast_cap_hdr	0x300	32	0	0-1	0 - 3	0 - 3
mcast_cap_ext	0x304	32	0	0-1	0 - 3	0 - 3
mcast_cap	0x30c	16	0	0-1	0 - 3	0 - 3
mcast_ctrl	0x30e	16	0	0-1	0 - 3	0 - 3
mcast_base	0x310	64	0	0-1	0 - 3	0 - 3
mcast_rcv	0x318	64	0	0-1	0 - 3	0 - 3
mcast_blk_all	0x320	64	0	0-1	0 - 3	0 - 3
mcast_blk_unt	0x328	64	0	0-1	0 - 3	0 - 3
mcast_overlay_bar	0x330	64	0	0-1	0 - 3	0 - 3

6.2.1 vid

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0x0			
Bit	Attr	Default	Description
15:0	RO	0x8086	vendor_identification_number: The value is assigned by PCI-SIG to Intel.



6.2.2 did

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x2				

Bit	Attr	Default	Description
15:0	RO RO_V (Device 0 and 3 Function 0)	For Device 0 Function 0: 0x2f00 (DMI2 Mode) 0x2f01 (PCIe Mode) For Device 2: 0x2f04 (Function 0) 0x2f05 (Function 1) 0x2f06 (Function 2) 0x2f07 (Function 3) For Device 3: 0x2f08 (Function 0) 0x2f09 (Function 1) 0x2f0a (Function 2) 0x2f0b (Function 3)	device_identification_number: Device ID values vary from function to function.


6.2.3 pcicmd

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x4				

Bit	Attr	Default	Description
10:10	RW	0x0	interrupt_disable: Interrupt Disable. Controls the ability of the PCI Express port to generate INTx messages. This bit does not affect the ability of the processor to route interrupt messages received at the PCI Express port. However, this bit controls the generation of legacy interrupts to the DMI for PCI Express errors detected internally in this port (for example, Malformed TLP, CRC error, completion time out, and so forth) or when receiving RP error messages or interrupts due to Hot Plug/Power Management events generated in legacy mode within the processor. 1: Legacy Interrupt mode is disabled 0: Legacy Interrupt mode is enabled
9:9	RO	0x0	fast_back_to_back_enable: Fast Back-to-Back Enable Not applicable to PCI Express must be hardwired to 0.



Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0x4		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
8:8	RW	0x0	serre: SERR Enable For PCI Express/DMI ports, this field enables notifying the internal core error logic of occurrence of an uncorrectable error (fatal or non-fatal) at the port. The internal core error logic of the IIO module then decides if/how to escalate the error further (pins/message, and so forth). This bit also controls the propagation of PCI Express ERR_FATAL and ERR_NONFATAL messages received from the port to the internal IIO core error logic. 1: Fatal and Non-fatal error generation and Fatal and Non-fatal error message forwarding is enabled 0: Fatal and Non-fatal error generation and Fatal and Non-fatal error message forwarding is disabled		
7:7	RO	0x0	idssel_stepping_wait_cycle_control: IDSEL Stepping/Wait Cycle Control Not applicable to PCI Express must be hardwired to 0.		
6:6	RW	0x0	perre: Parity Error Response For PCI Express/DMI ports, the IIO module ignores this bit and always does ECC/parity checking and signaling for data/address of transactions both to and from IIO. This bit though affects the setting of bit 8 in the PCISTS register.		
5:5	RO	0x0	vga_palette_snoop_enable: Not applicable to PCI Express must be hardwired to 0.		
4:4	RO	0x0	mwie: Not applicable to PCI Express must be hardwired to 0.		
3:3	RO	0x0	sce: Not applicable to PCI Express must be hardwired to 0.		
2:2	RW RW_L (Device 0 Function 0)	0x0	bme:		
1:1	RW RW_L (Device 0 Function 0)	0x0	mse: Memory Space Enable 1: Enables a PCI Express port's memory range registers to be decoded as valid target addresses for transactions from secondary side. 0: Disables a PCI Express port's memory range registers (including the Configuration Registers range registers) to be decoded as valid target addresses for transactions from secondary side. All memory accesses received from secondary side are UR'ed.		
0:0	RW RW_L (Device 0 and 3 Function 0)	0x0	iose: IO Space Enable Controls a device's response to I/O Space accesses. A value of 0 disables the device response. A value of 1 allows the device to respond to I/O Space accesses. State after RST# is 0.		



6.2.4 pcists


Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0x6		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
15:15	RW1C	0x0	dpe: Detected Parity Error This bit is set by a root port when it receives a packet on the primary side with an uncorrectable data error (including a packet with poison bit set) or an uncorrectable address/control parity error. The setting of this bit is regardless of the Parity Error Response bit (PERRE) in the PCICMD register.		
14:14	RW1C	0x0	sse: Signaled System Error 1: The root port reported fatal/non-fatal (and not correctable) errors it detected on its PCI Express interface to the IIO core error logic (which might eventually escalate the error through the ERR[2:0] pins or message to cpu core or message to PCH). Note that the SERRE bit in the PCICMD register must be set for a device to report the error the IIO core error logic. Software clears this bit by writing a '1' to it. This bit is also set (when SERR enable bit is set) when a FATAL/NON-FATAL message is forwarded to the IIO core error logic. Note that the IIO internal 'core' errors (like parity error in the internal queues) are not reported via this bit. 0: The root port did not report a fatal/non-fatal error		
13:13	RW1C	0x0	rma: Received Master Abort This bit is set when a root port experiences a master abort condition on a transaction it mastered on the primary interface (uncore internal bus). Note that certain errors might be detected right at the PCI Express interface and those transactions might not 'propagate' to the primary interface before the error is detected (for example, accesses to memory above TOCM in cases where the PCIe interface logic itself might have visibility into TOCM). Such errors do not cause this bit to be set, and are reported via the PCI Express interface error bits (secondary status register). Conditions that cause bit 13 to be set, include: Device receives a completion on the primary interface (internal bus of uncore) with Unsupported Request or master abort completion Status. This includes UR status received on the primary side of a PCI Express port on peer-to-peer completions also.		
12:12	RW1C	0x0	rta: Received Target Abort This bit is set when a device experiences a completer abort condition on a transaction it mastered on the primary interface (uncore internal bus). Note that certain errors might be detected right at the PCI Express interface and those transactions might not 'propagate' to the primary interface before the error is detected (for example, accesses to memory above VTBAR). Such errors do not cause this bit to be set, and are reported via the PCI Express interface error bits (secondary status register). Conditions that cause bit 12 to be set, include: Device receives a completion on the primary interface (internal bus of uncore) with completer abort completion Status. This includes CA status received on the primary side of a PCI Express port on peer-to-peer completions also.		
11:11	RW1C	0x0	sta: Signaled Target Abort This bit is set when a root port signals a completer abort completion status on the primary side (internal bus of uncore). This condition includes a PCI Express port forwarding a completer abort status received on a completion from the secondary.		



Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0x6		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
10:9	RO	0x0	devsel_timing: Not applicable to PCI Express. Hardwired to 0.		
8:8	RW1C	0x0	mdpe: Master Data Parity Error This bit is set by a root port if the Parity Error Response bit in the PCI Command register is set and it either receives a completion with poisoned data from the primary side or it forwards a packet with data (including MSI writes) to the primary side with poison.		
7:7	RO	0x0	fast_back_to_back: Not applicable to PCI Express. Hardwired to 0.		
5:5	RO	0x0	pci66mhz_capable: Not applicable to PCI Express. Hardwired to 0.		
4:4	RO	0x1	capabilities_list: Not applicable to PCI Express. Hardwired to 0.		
3:3	RO_V	0x0	intx_status: This Read-only bit reflects the state of the interrupt in the PCI Express Root Port. Only when the Interrupt Disable bit in the command register is a 0 and this Interrupt Status bit is a 1, will this device generate INTx interrupt. Setting the Interrupt Disable bit to a 1 has no effect on the state of this bit. This bit does not get set for interrupts forwarded to the root port from downstream devices in the hierarchy. When MSI are enabled, Interrupt status should not be set.		

6.2.5 rid

Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0x8		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
7:0	RO_V	0x0	revision_id: Reflects the Uncore Revision ID after reset. Reflects the Compatibility Revision ID after BIOS writes 0x69 to any RID register in any Intel® Xeon® Processor E5 v4 product family function.		



6.2.6 ccr

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x9				

Bit	Attr	Default	Description
23:16	RO_V	0x6	base_class: Generic Device
15:8	RO_V	0x4 0x80 (Device 3 Function 0 only)	sub_class: Generic Device
7:0	RO_V	0x0	interface: This field is hardwired to 00h for PCI Express port.

6.2.7 clsr

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0xc				

Bit	Attr	Default	Description
7:0	RW	0x0	cacheline_size: This register is set as RW for compatibility reasons only. Cacheline size is always 64B. IIO hardware ignores this setting.

6.2.8 plat

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0xd				

Bit	Attr	Default	Description
7:0	RO	0x0	primary_latency_timer: Not applicable to PCI Express. Hardwired to 00h.



6.2.9 hdr

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0xe				
Bit	Attr	Default	Description		
7:7	RO_V RO (Device 0 Function 0)	0x1 0x0 (Device 0 Function 0)	<p>mf: Multi-function Device</p> <p>This bit defaults to 0 for Device 0.</p> <p>This bit defaults to 1 for Devices 2-3.</p> <p>BIOS can individually control the value of this bit in Function 0 of these devices, based on HDRTPCTRL register. BIOS will write to that register to change this field to 0 in Function 0 of these devices, if it exposes only Function 0 in the device to OS.</p> <p>Note: In product SKUs where only Function 0 of the device is exposed to any software (BIOS/OS), BIOS would have to still set the control bits mentioned above to set the this bit in this register to be compliant per PCI rules.</p>		
6:0	RO RO_V (Device 0 Function 0)	0x1 0x0 (Device 0 Function 0)	<p>cl: Configuration Layout</p> <p>This field identifies the format of the configuration header layout.</p> <p>In DMI mode, default is 00h indicating a conventional type 00h PCI header.</p> <p>In PCIe mode, the default is 01h, corresponding to Type 1 for a PCIe root port.</p>		

6.2.10 bist

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0xf				
Bit	Attr	Default	Description		
7:0	RO	0x0	<p>bist_tests:</p> <p>Not Supported. Hardwire to 00h.</p>		



6.2.11 pbus

Primary Bus Number Register.

Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0x18		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 (PCIe Mode) Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
7:0	RW	0x0	pbn: Configuration software programs this field with the number of the bus on the primary side of the bridge. This register has to be kept consistent with the Internal Bus Number 0 in the CPUBUSNO01 register. BIOS (and OS if internal bus number gets moved) must program this register to the correct value since IIO hardware would depend on this register for inbound configuration cycle decode purposes.		

6.2.12 secbus

Secondary Bus Number Register.

Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0x19		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 (PCIe Mode) Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
7:0	RW	0x0	sbn: This field is programmed by configuration software to assign a bus number to the secondary bus of the virtual P2P bridge. IIO uses this register to either forward a configuration transaction as a Type 1 or Type 0 to PCI Express.		

6.2.13 subbus

Subordinate Bus Number Register.

Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0x1a		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 (PCIe Mode) Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
7:0	RW	0x0	subordinate_bus_number: This register is programmed by configuration software with the number of the highest subordinate bus that is behind the PCI Express port. Any transaction that falls between the secondary and subordinate bus number (both inclusive) of an Express port is forwarded to the express port.		



6.2.14 iobas

I/O Base Register.

Type:	CFG	PortID:	N/A	Function:	0 (PCIe Mode)
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x1c				

Bit	Attr	Default	Description
7:4	RW	0xf	i_o_base_address: Corresponds to A[15:12] of the IO base address of the PCI Express port. See also the IOLIM register description.
3:2	RW_L	0x0	more_i_o_base_address: When EN1K is set in the IOMISCCTRL register, these bits become RW and allow for 1K granularity of I/O addressing, otherwise these are RO.
1:0	RO	0x0	i_o_address_capability: IIO supports only 16 bit addressing

6.2.15 iolim

I/O Limit Register.

Type:	CFG	PortID:	N/A	Function:	0 (PCIe Mode)
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x1d				

Bit	Attr	Default	Description
7:4	RW	0x0	i_o_address_limit: Corresponds to A[15:12] of the I/O limit address of the PCI Express port. The I/O Base and I/O Limit registers define an address range that is used by the PCI Express port to determine when to forward I/O transactions from one interface to the other using the following formula: $IO_BASE \leq A[15:12] \leq IO_LIMIT$ The bottom of the defined I/O address range will be aligned to a 4KB boundary (1KB if EN1K bit is set. Refer to the IOMISCCTRL register for definition of EN1K bit) while the top of the region specified by IO_LIMIT will be one less than a 4 KB (1KB if EN1K bit is set) multiple. Notes: Setting the I/O limit less than I/O base disables the I/O range altogether. General the I/O base and limit registers won't be programmed by software without clearing the IOSE bit first.
3:2	RW_L	0x0	more_i_o_address_limit: When EN1K is set in the IOMISCCTRL register, these bits become RW and allow for 1K granularity of I/O addressing, otherwise these are RO.
1:0	RO	0x0	i_o_address_limit_capability: IIO only supports 16 bit addressing



6.2.16 secsts

Secondary Status Register.

Type: CFG		PortID: N/A		Function: 0 (PCIe Mode)	
Bus: 0		Device: 0		Function: 0-1	
Bus: 0		Device: 1		Function: 0-3	
Bus: 0		Device: 2		Function: 0-3	
Bus: 0		Device: 3		Function: 0-3	
Offset: 0x1e					
Bit	Attr	Default	Description		
15:15	RW1C	0x0	dpe: Detected Parity Error This bit is set by the root port whenever it receives a poisoned TLP in the PCI Express port. This bit is set regardless of the state the Parity Error Response Enable bit in the Bridge Control register.		
14:14	RW1C	0x0	rse: Received System Error This bit is set by the root port when it receives a ERR_FATAL or ERR_NONFATAL message from PCI Express. Note this does not include the virtual ERR* messages that are internally generated from the root port when it detects an error on its own.		
13:13	RW1C	0x0	rma: Received Master Abort Status This bit is set when the root port receives a Completion with 'Unsupported Request Completion' Status or when the root port master aborts a Type0 configuration packet that has a non-zero device number.		
12:12	RW1C	0x0	rta: Received Target Abort Status This bit is set when the root port receives a Completion with 'Completer Abort' Status.		
11:11	RW1C	0x0	sta: Signaled Target Abort This bit is set when the root port sends a completion packet with a 'Completer Abort' Status (including peer-to-peer completions that are forwarded from one port to another).		
10:9	RO	0x0	devsel_timing: Not applicable to PCI Express. Hardwired to 0.		
8:8	RW1C	0x0	mdpe: Master Data Parity Error This bit is set by the root port on the secondary side (PCI Express link) if the Parity Error Response Enable bit (PERRE) is set in Bridge Control register and either of the following two conditions occurs: The PCI Express port receives a Completion from PCI Express marked poisoned. The PCI Express port poisons an outgoing packet with data. If the Parity Error Response Enable bit in Bridge Control Register is cleared, this bit is never set.		
7:7	RO	0x0	fast_back_to_back_transactions_capable: Not applicable to PCI Express. Hardwired to 0.		
5:5	RO	0x0	pci66_mhz_capability: Not applicable to PCI Express. Hardwired to 0.		



6.2.17 mbas

Memory Base.

Type:	CFG	PortID:	N/A	Function:	0 (PCIe Mode)
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x20				

Bit	Attr	Default	Description
15:4	RW	0xffff	memory_base_address: Corresponds to A[31:20] of the 32 bit memory window's base address of the PCI Express port. See also the MLIM register description.

6.2.18 mlim

Memory Limit Register.

Type:	CFG	PortID:	N/A	Function:	0 (PCIe Mode)
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x22				

Bit	Attr	Default	Description
15:4	RW	0x0	memory_limit_address: Corresponds to A[31:20] of the 32 bit memory window's limit address that corresponds to the upper limit of the range of memory accesses that will be passed by the PCI Express bridge. The Memory Base and Memory Limit registers define a memory mapped IO non-prefetchable address range (32-bit addresses) and the IIO directs accesses in this range to the PCI Express port based on the following formula: MEMORY_BASE <= A[31:20] <= MEMORY_LIMIT The upper 12 bits of both the Memory Base and Memory Limit registers are read/write and corresponds to the upper 12 address bits, A[31:20] of 32-bit addresses. Thus, the bottom of the defined memory address range will be aligned to a 1 MB boundary and the top of the defined memory address range will be one less than a 1 MB boundary. Notes: Setting the memory limit less than memory base disables the 32-bit memory range altogether. Note that in general the memory base and limit registers won't be programmed by software without clearing the MSE bit first.



6.2.19 pbas

Prefetchable Memory Base Register.

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0 (PCIe Mode)
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0x24			
Bit	Attr	Default	Description
15:4	RW	0xfff	prefetchable_memory_base_address: Corresponds to A[31:20] of the prefetchable memory address range's base address of the PCI Express port. See also the PLIMU register description.
3:0	RO	0x1	prefetchable_memory_base_address_capability: IIO sets this bit to 01h to indicate 64bit capability.

6.2.20 plim


Prefetchable Memory Limit Register.

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0 (PCIe Mode)
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0x26			
Bit	Attr	Default	Description
15:4	RW	0x0	prefetchable_memory_limit_address: Corresponds to A[31:20] of the prefetchable memory address range's limit address of the PCI Express port. See also the PLIMU register description.
3:0	RO	0x1	prefetchable_memory_limit_address_capability: IIO sets this field to 01h to indicate 64bit capability.

6.2.21 pbasu

Prefetchable Memory Base Upper 32 bits.

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0 (PCIe Mode)
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0x28			
Bit	Attr	Default	Description
31:0	RW	0xffffffff	prefetchable_upper_32_bit_memory_base_address: Corresponds to A[63:32] of the prefetchable memory address range's base address of the PCI Express port. See also the PLIMU register description.



6.2.22 plimu

Prefetchable Memory Limit Upper 32 bits.

Type:	CFG	PortID:	N/A	Function:	0 (PCIe Mode)
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x2c				

Bit	Attr	Default	Description
31:0	RW	0x0	<p>prefetchable_upper_32_bit_memory_limit_address: Corresponds to A[63:32] of the prefetchable memory address range's limit address of the PCI Express port. The Prefetchable Memory Base and Memory Limit registers define a memory mapped I/O prefetchable address range (64-bit addresses) which is used by the PCI Express bridge to determine when to forward memory transactions based on the following formula: $\text{PREFETCH_MEMORY_BASE_UPPER} :: \text{PREFETCH_MEMORY_BASE} \leq A[63:20] \leq \text{PREFETCH_MEMORY_LIMIT_UPPER} :: \text{PREFETCH_MEMORY_LIMIT}$ The upper 12 bits of both the Prefetchable Memory Base and Memory Limit registers are read/write and corresponds to the upper 12 address bits, A[31:20] of 32-bit addresses. The bottom of the defined memory address range will be aligned to a 1 MB boundary and the top of the defined memory address range will be one less than a 1 MB boundary. The bottom 4 bits of both the Prefetchable Memory Base and Prefetchable Memory Limit registers are read-only, contain the same value, and encode whether or not the bridge supports 64-bit addresses. If these four bits have the value 0h, then the bridge supports only 32 bit addresses. If these four bits have the value 1h, then the bridge supports 64-bit addresses and the Prefetchable Base Upper 32 Bits and Prefetchable Limit Upper 32 Bits registers hold the rest of the 64-bit prefetchable base and limit addresses respectively. Setting the prefetchable memory limit less than prefetchable memory base disables the 64-bit prefetchable memory range altogether. Notes: In general the memory base and limit registers won't be programmed by software without clearing the MSE bit first.</p>

6.2.23 capptr

Capability Pointer.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x34				

Bit	Attr	Default	Description
7:0	RO_V (Device 0 Function 0, Device 2 Function 0-3) RW_V (Device 3 Function 0) RO (Device 3 Function 1-3)	0x40 0x60 (Device 3 Function 0) 0x90 (Device 0 Function 0)	<p>capability_pointer: Points to the first capability structure for the device which is the PCIe capability.</p>



6.2.24 intl

Interrupt Line Register.

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0x3c			

Bit	Attr	Default	Description
7:0	RW RO (Device 0 Function 0)	0x0	interrupt_line: N/A for these devices

6.2.25 intpin

Interrupt Pin Register.

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0x3d			

Bit	Attr	Default	Description
7:0	RW_O	0x1	intp: N/A since these devices do not generate any interrupt on their own

6.2.26 bctrl

Bridge Control Register.

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0 (PCIe Mode)
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0x3e			

Bit	Attr	Default	Description
6:6	RW	0x0	sbr: 1: Setting this bit triggers a hot reset on the link for the corresponding PCI Express port and the PCI Express hierarchy domain subordinate to the port. This sends the LTSSM into the Training (or Link) Control Reset state, which necessarily implies a reset to the downstream device and all subordinate devices. The transaction layer corresponding to port will be emptied by virtue of the link going down when this bit is set. This means that in the outbound direction, all posted transactions are dropped and non-posted transactions are sent a UR response. In the inbound direction, completions for inbound NP requests are dropped when they arrive. Inbound posted writes are retired normally. Note also that a secondary bus reset will not reset the virtual PCI-to-PCI bridge configuration registers of the targeted PCI Express port. 0: No reset happens on the PCI Express port.



Type:	CFG	PortID:	N/A	Function:	0 (PCIe Mode)
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x3e				
Bit	Attr	Default	Description		
4:4	RW	0x0	vga16b: This bit enables the virtual PCI-to-PCI bridge to provide 16-bit decoding of VGA I/O address precluding the decoding of alias addresses every 1 KB. 0: execute 10-bit address decodes on VGA I/O accesses. 1: execute 16-bit address decodes on VGA I/O accesses. Note: This bit only has meaning if bit 3 of this register is also set to 1, enabling VGA IO decoding and forwarding by the bridge.		
3:3	RW	0x0	vgaen: Controls the routing of CPU initiated transactions targeting VGA compatible IO and memory address ranges. This bit must only be set for one p2p port in the entire system.		
2:2	RW	0x0	isaen: Modifies the response by the root port to an I/O access issued by the core that target ISA I/O addresses. This applies only to I/O addresses that are enabled by the IOBASE and IOLIM registers. 1: The root port will not forward to PCI Express any IO transactions addressing the last 768 bytes in each 1KB block even if the addresses are within the range defined by the IOBASE and IOLIM registers. 0: All addresses defined by the IOBASE and IOLIM for core issued IO transactions will be mapped to PCI Express.		
1:1	RW	0x0	serre: SERR Response Enable This bit controls forwarding of ERR_COR, ERR_NONFATAL and ERR_FATAL messages from the PCI Express port to the primary side. 1: Enables forwarding of ERR_COR, ERR_NONFATAL and ERR_FATAL messages. 0: Disables forwarding of ERR_COR, ERR_NONFATAL and ERR_FATAL		
0:0	RW	0x0	perre: Parity Error Response Enable This only effect this bit has is on the setting of bit 8 in the SECSTS register		

6.2.27 scapid

Subsystem Capability Identity.

Type:	CFG	PortID:	N/A	Function:	0 (PCIe Mode)
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x40				
Bit	Attr	Default	Description		
7:0	RO RW_0 (Device 0 Function 0)	0xd	capability_id: Assigned by PCI-SIG for subsystem capability ID		

6.2.28 snxtptr

Subsystem ID Next Pointer.



6.2.31 dmircbar

DMI Root Complex Register Block Base Address.

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	
Offset: 0x50		Function: 0	
Bit	Attr	Default	Description
31:12	RW_LB	0x0	dmircbar: This field corresponds to bits 32 to 12 of the base address DMI Root Complex register space. BIOS will program this register resulting in a base address for a 4KB block of contiguous memory address space. This register ensures that a naturally aligned 4KB space is allocated within the first 64GB of addressable memory space. System Software uses this base address to program the DMI Root Complex register set. All the Bits in this register are locked in Intel TXT mode.
0:0	RW_LB	0x0	dmircbaren: 0: DMIRCBAR is disabled and does not claim any memory 1: DMIRCBAR memory mapped accesses are claimed and decoded Note: Accesses to registers pointed to by the DMIRCBAR, via message channel or JTAG mini-port are not gated by this enable bit i.e. accesses these registers are honored regardless of the setting of this bit.

6.2.32 msicapid

MSI Capability ID.

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	
Bus: 0		Device: 1	
Bus: 0		Device: 2	
Bus: 0		Device: 3	
Offset: 0x60		Function: 0-3	
Bit	Attr	Default	Description
7:0	RO	0x5	capability_id: Assigned by PCI-SIG for MSI root ports.

6.2.33 msinxtptr

MSI Next Pointer.

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	
Bus: 0		Device: 1	
Bus: 0		Device: 2	
Bus: 0		Device: 3	
Offset: 0x61		Function: 0-3	
Bit	Attr	Default	Description
7:0	RW_O	0x90	next_ptr: This field is set to 90h for the next capability list PCI Express capability structure in the chain.



6.2.34 msimsgctl

MSI Control.

Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0x62		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
8:8	RO	0x1	pvmc: This bit indicates that PCI Express ports support MSI per-vector masking.		
7:7	RO	0x0	b64ac: This field is hardwired to 0h since the message addresses are only 32-bit addresses (for example, FEEx_xxxxh).		
6:4	RW	0x0	mme: Multiple Message Enable. Applicable only to PCI Express ports. Software writes to this field to indicate the number of allocated messages which is aligned to a power of two. When MSI is enabled, the software will allocate at least one message to the device. A value of 000 indicates 1 message. Any value greater than or equal to 001 indicates a message of 2.		
3:1	RO	0x1	mmc: Multiple Message Capable. Intel® Xeon® Processor E5 v4 product family's Express ports support two messages for all their internal events.		
0:0	RW	0x0	msien: Software sets this bit to select INTx style interrupt or MSI interrupt for root port generated interrupts. 0: INTx interrupt mechanism is used for root port interrupts, provided the override bits in MISCCTRLSTS allow it 1: MSI interrupt mechanism is used for root port interrupts, provided the override bits in MISCCTRLSTS allow it Note there bits 4:2 and bit 2 MISCCTRLSTS can disable both MSI and INTx interrupt from being generated on root port interrupt events.		

6.2.35 msgadr

The MSI Address Register (MSIAR) contains the system specific address information to route MSI interrupts from the root ports and is broken into its constituent fields.

Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0x64		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
31:2	RW	0x0	address_id:		



6.2.36 msgdat

MSI Data Register.

Type:	CFG	PortID:	N/A		
Bus:	0	Device:	0	Function:	0
Bus:	0	Device:	1	Function:	0-1
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x68				
Bit	Attr	Default	Description		
15:0	RW	0x0	data:		

6.2.37 msimsk

MSI Mask Bit.

Type:	CFG	PortID:	N/A		
Bus:	0	Device:	0	Function:	0
Bus:	0	Device:	1	Function:	0-1
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x6c				
Bit	Attr	Default	Description		
1:0	RW	0x0	mask_bits: Relevant only when MSI is enabled and used for interrupts generated by the root port. For each Mask bit that is set, the PCI Express port is prohibited from sending the associated message. When only one message is allocated to the root port by software, only mask bit 0 is relevant and used by hardware.		


6.2.38 msipending

MSI Pending Bit.

Type:	CFG	PortID:	N/A		
Bus:	0	Device:	0	Function:	0
Bus:	0	Device:	1	Function:	0-1
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x70				
Bit	Attr	Default	Description		
1:0	RO_V	0x0	pending_bits: Relevant only when MSI is enabled and used for interrupts generated by the root port. When MSI is not enabled or used by the root port, this register always reads a value 0. For each Pending bit that is set, the PCI Express port has a pending associated message. When only one message is allocated to the root port by software, only pending bit 0 is set/cleared by hardware and pending bit 1 always reads 0. Hardware sets this bit whenever it has an interrupt pending to be sent. This bit remains set till either the interrupt is sent by hardware or the status bits associated with the interrupt condition are cleared by software.		

6.2.39 pxpcapid

PCI Express Capability Identity.



Type: CFG PortID: N/A Bus: 0 Device: 0 Function: 0 Bus: 0 Device: 1 Function: 0-1 Bus: 0 Device: 2 Function: 0-3 Bus: 0 Device: 3 Function: 0-3 Offset: 0x90			
Bit	Attr	Default	Description
7:0	RO	0x10	capability_id: Provides the PCI Express capability ID assigned by PCI-SIG.

6.2.40 pxpnxtptr

PCI Express Next Pointer.

Type: CFG PortID: N/A Bus: 0 Device: 0 Function: 0 Bus: 0 Device: 1 Function: 0-1 Bus: 0 Device: 2 Function: 0-3 Bus: 0 Device: 3 Function: 0-3 Offset: 0x91			
Bit	Attr	Default	Description
7:0	RO	0xe0	next_ptr: This field is set to the PCI Power Management capability.

6.2.41 pxpcap

PCI Express Capabilities Register.

Type: CFG PortID: N/A Bus: 0 Device: 0 Function: 0 Bus: 0 Device: 1 Function: 0-1 Bus: 0 Device: 2 Function: 0-3 Bus: 0 Device: 3 Function: 0-3 Offset: 0x92			
Bit	Attr	Default	Description
13:9	RO	0x0	interrupt_message_number: Applies to root ports. This field indicates the interrupt message number that is generated for Power Management/Hot Plug/BW-change events. When there are more than one MSI interrupt Number allocated for the root port MSI interrupts, this register field is required to contain the offset between the base Message Data and the MSI Message that is generated when there are Power Management/Hot Plug/BW-change interrupts. IIO assigns the first vector for Power Management/Hot Plug/BW-change events and so this field is set to 0.
8:8	RW_O	0x0	slot_implemented: Applies only to the root ports. 1: indicates that the PCI Express link associated with the port is connected to a slot. 0: indicates no slot is connected to this port.
7:4	RO_V	0x4	device_port_type: This field identifies the type of device. It is set to 0x4 for all the Express ports.



Type:	CFG	PortID:	N/A	
Bus:	0	Device:	0	Function: 0
Bus:	0	Device:	1	Function: 0-1
Bus:	0	Device:	2	Function: 0-3
Bus:	0	Device:	3	Function: 0-3
Offset:	0x92			

Bit	Attr	Default	Description
3:0	RW_O	0x2	<p>capability_version:</p> <p>This field identifies the version of the PCI Express capability structure, which is 2h as of now. This register field is left as RW-O to cover any unknowns with Gen3.</p>

6.2.42 devcap

The PCI Express Device Capabilities register identifies device specific information for the device.

Type:	CFG	PortID:	N/A	
Bus:	0	Device:	0	Function: 0
Bus:	0	Device:	1	Function: 0-1
Bus:	0	Device:	2	Function: 0-3
Bus:	0	Device:	3	Function: 0-3
Offset:	0x94			

Bit	Attr	Default	Description
27:26	RO	0x0	<p>captured_slot_power_limit_scale:</p> <p>Does not apply to root ports or integrated devices.</p>
25:18	RO	0x0	<p>captured_slot_power_limit_value:</p> <p>Does not apply to root ports or integrated devices.</p>
15:15	RO	0x1	<p>role_based_error_reporting:</p> <p>IIO is 1.1 compliant and so supports this feature</p>
14:14	RO	0x0	<p>power_indicator_present_on_device:</p> <p>Does not apply to root ports or integrated devices.</p>
13:13	RO	0x0	<p>attention_indicator_present:</p> <p>Does not apply to root ports or integrated devices.</p>
12:12	RO	0x0	<p>attention_button_present:</p> <p>Does not apply to root ports or integrated devices.</p>
11:9	RO	0x0	<p>endpoint_l1_acceptable_latency:</p> <p>N/A</p>
8:6	RO	0x0	<p>endpoint_l0s_acceptable_latency:</p> <p>N/A</p>
5:5	RW_O	0x0 0x1 (Device 3 Function 0)	<p>extended_tag_field_supported:</p>
4:3	RO	0x0	<p>phantom_functions_supported:</p> <p>IIO does not support phantom functions.</p>
2:0	RO	0x1 0x0 (Device 0 Function 0)	<p>max_payload_size_supported:</p> <p>Max payload is 128B on the DMI/PCIe port corresponding to Port 0.</p>



6.2.43 devctrl

PCI Express Device Control.

Type: CFG Bus: 0 Offset: 0xf0		PortID: N/A Device: 0		Function: 0 (DMI 2 Mode)
Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0x98		Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 (PCIe Mode) Function: 0-1 Function: 0-3 Function: 0-3
Bit	Attr	Default	Description	
14:12	RO	0x0	max_read_request_size: PCI Express/DMI ports in Processor do not generate requests greater than 64B and this field is RO.	
11:11	RO	0x0	enable_no_snoop: Not applicable to DMI or PCIe root ports since they never set the 'No Snoop' bit for transactions they originate (not forwarded from peer) to PCI Express/DMI. This bit has no impact on forwarding of NoSnoop attribute on peer requests.	
10:10	RO	0x0	auxiliary_power_management_enable: Not applicable to Processor	
9:9	RO	0x0	phantom_functions_enable: Not applicable to IIO since it never uses phantom functions as a requester.	
8:8	RW RO (Device 0 Function 0)	0x0	extended_tag_field_enable: N/A since IIO it never generates any requests on its own that uses tags 7:5. Note though that on peer to peer writes, IIO forwards the tag field along without modification and tag fields 7:5 could be set and that is not impacted by this bit.	
7:5	RW_LV RW (Device 0 Function 0)	0x0	max_payload_size: 000: 128B max payload size 001: 256B max payload size others: alias to 128B IIO can receive packets equal to the size set by this field. IIO generate read completions as large as the value set by this field. IIO generates memory writes of max 64B.	
4:4	RO	0x0	enable_relaxed_ordering: Not applicable to root/DMI ports since they never set relaxed ordering bit as a requester (this does not include tx forwarded from peer devices). This bit has no impact on forwarding of relaxed ordering attribute on peer requests.	
3:3	RW	0x0	unsupported_request_reporting_enable: This bit controls the reporting of unsupported requests that IIO itself detects on requests its receives from a PCI Express/DMI port. 0: Reporting of unsupported requests is disabled 1: Reporting of unsupported requests is enabled.	
2:2	RW	0x0	fatal_error_reporting_enable: Controls the reporting of fatal errors that IIO detects on the PCI Express/DMI interface. 0 = Reporting of Fatal error detected by device is disabled 1 = Reporting of Fatal error detected by device is enabled	
1:1	RW	0x0	non_fatal_error_reporting_enable: Controls the reporting of non-fatal errors that IIO detects on the PCI Express/DMI interface. 0 = Reporting of Non Fatal error detected by device is disabled 1 = Reporting of Non Fatal error detected by device is enabled	




Type:	CFG	PortID:	N/A	Function:	0 (DMI2 Mode)
Bus:	0	Device:	0		
Offset:	0xf0				
Bus:	0	Device:	0	Function:	0 (PCIe Mode)
Bus:	0	Device:	1	Function:	0-1
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x98				
Bit	Attr	Default	Description		
0:0	RW	0x0	correctable_error_reporting_enable: Controls the reporting of correctable errors that IIO detects on the PCI Express/DMI interface 0 = Reporting of link Correctable error detected by the port is disabled 1 = Reporting of link Correctable error detected by port is enabled		

6.2.44 devsts

PCI Express Device Status.

Type:	CFG	PortID:	N/A	Function:	0 (DMI2 Mode)
Bus:	0	Device:	0		
Offset:	0xf2				
Bus:	0	Device:	0	Function:	0 (PCIe Mode)
Bus:	0	Device:	1	Function:	0-1
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x9a				
Bit	Attr	Default	Description		
5:5	RO	0x0	transactions_pending: Does not apply to Root/DMI ports, that is, bit hardwired to 0 for these devices.		
4:4	RO	0x0	aux_power_detected: Does not apply to the processor		
3:3	RW1C	0x0	unsupported_request_detected: This bit indicates that the root port or DMI port detected an Unsupported Request. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control Register. 1: Unsupported Request detected at the device/port. These unsupported requests are NP requests inbound that the root port or DMI port received and it detected them as unsupported requests (for example, address decoding failures that the root port detected on a packet, receiving inbound lock reads, BME bit is clear and so forth). 0: No unsupported request detected by the root or DMI port Note: This bit is not set on peer-to-peer completions with UR status that are forwarded by the root port or DMI port to the PCIe/DMI link.		
2:2	RW1C	0x0	fatal_error_detected: This bit indicates that a fatal (uncorrectable) error is detected by the root or DMI port. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. 1: Fatal errors detected 0: No Fatal errors detected		
1:1	RW1C	0x0	non_fatal_error_detected: This bit gets set if a non-fatal uncorrectable error is detected by the root or DMI port. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. 1: Non Fatal errors detected 0: No non-Fatal Errors detected		



Type:	CFG	PortID:	N/A	Function:	0 (DMI2 Mode)
Bus:	0	Device:	0		
Offset:	0xf2				
Bus:	0	Device:	0	Function:	0 (PCIe Mode)
Bus:	0	Device:	1	Function:	0-1
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x9a				
Bit	Attr	Default	Description		
0:0	RW1C	0x0	correctable_error_detected: This bit gets set if a correctable error is detected by the root or DMI port. Errors are logged in this register regardless of whether error reporting is enabled or not in the PCI Express Device Control register. 1: correctable errors detected 0: No correctable errors detected		

6.2.45 Inkcab

PCI Express Link Capabilities

The Link Capabilities register identifies the PCI Express specific link capabilities. The link capabilities register needs some default values setup by the local host.


Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x9c				
Bit	Attr	Default	Description		
31:24	RW_O	0x0	port_number: This field indicates the PCI Express port number for the link and is initialized by software/BIOS. IIO hardware does nothing with this bit.		
22:22	RW_O	0x1	aspm_optionality_compliance:		
21:21	RO_V	0x1	link_bandwidth_notification_capability: A value of 1b indicates support for the Link Bandwidth Notification status and interrupt mechanisms.		
20:20	RO	0x1	data_link_layer_link_active_reporting_capable: IIO supports reporting status of the data link layer so software knows when it can enumerate a device on the link or otherwise know the status of the link.		
19:19	RO	0x1	surprise_down_error_reporting_capable: IIO supports reporting a surprise down error condition		
18:18	RO	0x0	clock_power_management: Does not apply to processor		



Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0x9c			
Bit	Attr	Default	Description
17:15	RW_O	0x2	l1_exit_latency: This field indicates the L1 exit latency for the given PCI Express port. It indicates the length of time this port requires to complete transition from L1 to L0. 000: Less than 1us 001: 1 us to less than 2 us 010: 2 us to less than 4 us 011: 4 us to less than 8 us 100: 8 us to less than 16 us 101: 16 us to less than 32 us 110: 32 us to 64 us 111: More than 64us
14:12	RW_O	0x3	l0s_exit_latency: This field indicates the L0s exit latency (i.e L0s to L0) for the PCI Express port. 000: Less than 64 ns 001: 64 ns to less than 128 ns 010: 128 ns to less than 256 ns 011: 256 ns to less than 512 ns 100: 512 ns to less than 1 us 101: 1 is to less than 2 us 110: 2 is to 4 us 111: More than 4 us
11:10	RW_O	0x3	active_state_link_pm_support: This field indicates the level of active state power management supported on the given PCI Express port. 00: Disabled 01: L0s Entry Supported 10: Reserved 11: L0s and L1 Supported
9:4	RW_O	0x4	maximum_link_width: This field indicates the maximum width of the given PCI Express Link attached to the port. 000001: x1 000010: x2 000100: x4 001000: x8 010000: x16 Others: Reserved This is left as a RW-O register for bios to update based on the platform usage of the links.
3:0	RW_O	0x3 0x1 (Device 0 Function 0)	maxlnkspd: This field indicates the maximum link speed of this Port. The encoding is the binary value of the bit location in the Supported Link Speeds Vector in LNKCAP2 that corresponds to the maximum link speed.

6.2.46 Inkcon

PCI Express Link Control



The PCI Express Link Control register controls the PCI Express Link specific parameters. The link control register needs some default values setup by the local host.

Type: CFG Bus: 0 Offset: 0x1b0		PortID: N/A Device: 0		Function: 0 (DMI2 Mode)	
Bus: 0 Device: 0 Function: 0 (PCIe Mode)		Device: 1 Function: 0-1		Function: 0-3	
Bus: 0 Device: 0 Function: 0-3		Device: 2 Function: 0-3		Function: 0-3	
Bus: 0 Device: 3 Function: 0-3		Device: 3 Function: 0-3		Function: 0-3	
Offset: 0xa0					
Bit	Attr	Default	Description		
11:11	RW	0x0	link_autonomous_bandwidth_interrupt_enable: For root ports, when set to 1b this bit enables the generation of an interrupt to indicate that the Link Autonomous Bandwidth Status bit has been set. For DMI mode on Dev#0, interrupt is not supported and hence this bit is not useful. Expectation is that BIOS will set bit 27 in MISCCTRLSTS to notify the system of autonomous BW change event on that port.		
10:10	RW	0x0	link_bandwidth_management_interrupt_enable: For root ports, when set to 1b this bit enables the generation of an interrupt to indicate that the Link Bandwidth Management Status bit has been set. For DMI mode on Dev#0, interrupt is not supported and hence this bit is not useful. Expectation is that BIOS will set bit 27 in MISCCTRLSTS to notify the system of autonomous BW change event on that port.		
9:9	RW	0x0	hardware_autonomous_width_disable: When Set, this bit disables hardware from changing the Link width for reasons other than attempting to correct unreliable Link operation by reducing Link width. Note that IIO does not by itself change width for any reason other than reliability. So this bit only disables such a width change as initiated by the device on the other end of the link.		
8:8	RO	0x0	enable_clock_power_management:		
7:7	RW	0x0	extended_synch: This bit when set forces the transmission of additional ordered sets when exiting L0s and when in recovery.		
6:6	RW_V (Function 0) RW (Function 1-3)	0x0	common_clock_configuration: Software sets this bit to indicate that this component and the component at the opposite end of the Link are operating with a common clock source. A value of 0b indicates that this component and the component at the opposite end of the Link are operating with separate reference clock sources. Default value of this bit is 0b. Components utilize this common clock configuration information to report the correct L0s and L1 Exit Latencies in the NFTS. The values used come from these registers depending on the value of this bit: 0: Use NFTS values from CLSPHYCTL3 1: Use NFTS values from CLSPHYCTL4		
5:5	WO	0x0	retrain_link: A write of 1 to this bit initiates link retraining in the given PCI Express/DMI port by directing the LTSSM to the recovery state if the current state is [L0, L0s or L1]. If the current state is anything other than L0, L0s, L1 then a write to this bit does nothing. This bit always returns 0 when read. It is permitted to write 1b to this bit while simultaneously writing modified values to other fields in this register. If the LTSSM is not already in Recovery or Configuration, the resulting Link training must use the modified values. If the LTSSM is already in Recovery or Configuration, the modified values are not required to affect the Link training that's already in progress.		



Type:	CFG	PortID:	N/A	Function:	0 (DMI2 Mode)
Bus:	0	Device:	0		
Offset:	0x1b0				
Bus:	0	Device:	0	Function:	0 (PCIe Mode)
Bus:	0	Device:	1	Function:	0-1
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0xa0				
Bit	Attr	Default	Description		
4:4	RW	0x0	link_disable: This field controls whether the link associated with the PCI Express/DMI port is enabled or disabled. When this bit is a 1, a previously configured link would return to the 'disabled' state as defined in the PCI Express Base Specification, Revision 2.0. When this bit is clear, an LTSSM in the 'disabled' state goes back to the detect state. 0: Enables the link associated with the PCI Express port 1: Disables the link associated with the PCI Express port		
3:3	RO	0x0	read_completion_boundary: Set to zero to indicate IIO could return read completions at 64B boundaries		
1:0	RW_V (Function 0) RW (Function 1-3)	0x0	active_state_link_pm_control: When 01b or 11b, L0s on transmitter is enabled, otherwise it is disabled. 10 and 11 enables L1 ASPM.		

6.2.47 Inksts

PCI Express Link Status

The PCI Express Link Status register provides information on the status of the PCI Express Link such as negotiated width, training, and so forth. The link status register needs some default values setup by the local host.

Type:	CFG	PortID:	N/A	Function:	0 (DMI2 Mode)
Bus:	0	Device:	0		
Offset:	0x1b2				
Bus:	0	Device:	0	Function:	0 (PCIe Mode)
Bus:	0	Device:	1	Function:	0-1
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0xa2				
Bit	Attr	Default	Description		
15:15	RW1C	0x0	link_autonomous_bandwidth_status: This bit is set to 1b by hardware to indicate that hardware has autonomously changed link speed or width, without the port transitioning through DL_Down status, for reasons other than to attempt to correct unreliable link operation. IIO does not, on its own, change speed or width autonomously for non-reliability reasons. IIO only sets this bit when it receives a width or speed change indication from downstream component that is not for link reliability reasons.		
14:14	RW1C	0x0	link_bandwidth_management_status: This bit is set to 1b by hardware to indicate that either of the following has occurred without the port transitioning through DL_Down status: a) A link retraining initiated by a write of 1b to the Retrain Link bit has completed b) Hardware has autonomously changed link speed or width to attempt to correct unreliable link operation Note IIO also sets this bit when it receives a width or speed change indication from downstream component that is for link reliability reasons.		



6.2.48 sltcap

PCI Express Slot Capabilities

The Slot Capabilities register identifies the PCI Express specific slot capabilities.

Type:	CFG	PortID:	N/A
Bus:	0	Device:	0
Bus:	0	Device:	1
Bus:	0	Device:	2
Bus:	0	Device:	3
Offset:	0xa4	Function:	0 (PCIe Mode)
		Function:	0-1
		Function:	0-3
		Function:	0-3

Bit	Attr	Default	Description
31:19	RW_O	0x0	physical_slot_number: This field indicates the physical slot number of the slot connected to the PCI Express port and is initialized by BIOS.
18:18	RO	0x0	command_complete_not_capable: Intel® Xeon® Processor E5 v4 product family is capable of command complete interrupt.
17:17	RW_O	0x0	electromechanical_interlock_present: This bit when set indicates that an Electromechanical Interlock is implemented on the chassis for this slot and that lock is controlled by bit 11 in Slot Control register. This field is initialized by BIOS based on the system architecture. BIOS note: this capability is not set if the Electromechanical Interlock control is connected to main slot power control. This is expected to be used only for Express Module hot-pluggable slots.
6:6	RW_O	0x0	hot_plug_capable: This field defines hot-plug support capabilities for the PCI Express port. 0: indicates that this slot is not capable of supporting Hot-plug operations. 1: indicates that this slot is capable of supporting Hot-plug operations This bit is programmed by BIOS based on the system design. This bit must be programmed by bios to be consistent with the VPP enable bit for the port.
5:5	RW_O	0x0	hot_plug_surprise: This field indicates that a device in this slot may be removed from the system without prior notification. This field is initialized by BIOS. 0: indicates that hot-plug surprise is not supported 1: indicates that hot-plug surprise is supported Generally this bit is not expected to be set because the only know usage case for this is the ExpressCard FF. But that is not really expected usage in Intel® Xeon® Processor E5 v4 product family context. But this bit is present regardless to allow a usage if it arises. This bit is used by IIO hardware to determine if a transition from DL_active to DL_Inactive is to be treated as a surprise down error or not. If a port is associated with a hot-pluggable slot and the hot-plug surprise bit is set, then any transition to DLInactive is not considered an error.
4:4	RW_O	0x0	power_indicator_present: This bit indicates that a Power Indicator is implemented for this slot and is electrically controlled by the chassis. 0: indicates that a Power Indicator that is electrically controlled by the chassis is not present. 1: indicates that Power Indicator that is electrically controlled by the chassis is present. BIOS programs this field with a 1 for CEMExpress Module FFs, if the slot is hot-plug capable.



Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0xa4		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 (PCIe Mode) Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
3:3	RW_O	0x0	attention_indicator_present: This bit indicates that an Attention Indicator is implemented for this slot and is electrically controlled by the chassis 0: indicates that an Attention Indicator that is electrically controlled by the chassis is not present 1: indicates that an Attention Indicator that is electrically controlled by the chassis is present BIOS programs this field with a 1 for CEMExpress Module FFs, if the slot is hot-plug capable.		
2:2	RW_O	0x0	mrl_sensor_present: This bit indicates that an MRL Sensor is implemented on the chassis for this slot. 0: indicates that an MRL Sensor is not present 1: indicates that an MRL Sensor is present BIOS programs this field with a 0 for Express Module FF always. If CEM slot is hot-plug capable, BIOS programs this field with either 0 or 1 depending on system design.		
1:1	RW_O	0x0	power_controller_present: This bit indicates that a software controllable power controller is implemented on the chassis for this slot. 0: indicates that a software controllable power controller is not present 1: indicates that a software controllable power controller is present BIOS programs this field with a 1 for CEMExpress Module FFs, if the slot is hot-plug capable.		
0:0	RW_O	0x0	attention_button_present: This bit indicates that the Attention Button event signal is routed from slot or on-board in the chassis to the IIO's hot-plug controller. 0: indicates that an Attention Button signal is not routed to IIO 1: indicates that an Attention Button is routed to IIO BIOS programs this field with a 1 for CEMExpress Module FFs, if the slot is hot-plug capable.		

6.2.49 sltcon

PCI Express Slot Control.

Any write to this register will set the Command Completed bit in the SLTSTS register, only if the VPP enable bit for the port is set. If the port's VPP enable bit is set (i.e. hot-plug for that slot is enabled) then the required actions on VPP are completed before the Command Completed bit is set in the SLTSTS register. If the VPP enable bit for the port is clear, then the write simply updates this register see individual bit definitions for details but the Command Completed bit in the SLTSTS register is not set.



Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0xa8		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 (PCIe Mode) Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
12:12	RWS	0x0	data_link_layer_state_changed_enable: When set to 1, this field enables software notification when Data Link Layer Link Active bit in the LNKSTS register changes state		
11:11	RW	0x0	electromechanical_interlock_control: When software writes either a 1 to this bit, IIO pulses the EMIL pin per It;Bluegt;PCI Express ServerWorkstation Module Electromechanical Spec Rev 1.0. Write of 0 has no effect. This bit always returns a 0 when read. If electromechanical lock is not implemented, then either a write of 1 or 0 to this register has no effect.		
10:10	RWS	0x1	power_controller_control: If a power controller is implemented, when writes to this field will set the power state of the slot per the defined encodings. Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not executed yet at the VPP, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. 0: Power On 1: Power Off Note: If the link experiences an unexpected DL_Down condition that is not the result of a Hot Plug removal, the processor follows the PCI Express specification for logging Surprise Link Down. SW is required to set SLTCON[10] to 0 (Power On) in all devices that do not connect to a slot that supports Hot-Plug to enable logging of this error in that device. For devices connected to slots supporting Hot-Plug operations, SLTCON[10] usage to control PWREN# assertion is as described elsewhere.		
9:8	RW	0x3	power_indicator_control: If a Power Indicator is implemented, writes to this field will set the Power Indicator to the written state. Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not executed yet at the VPP, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. 00: Reserved. 01: On 10: Blink (IIO drives 1 Hz square wave for Chassis mounted LEDs) 11: Off IIO does not generated the Power_Indicator_On/Off/Blink messages on PCI Express when this field is written to by software.		
7:6	RW	0x3	attention_indicator_control: If an Attention Indicator is implemented, writes to this field will set the Attention Indicator to the written state. Reads of this field reflect the value from the latest write, even if the corresponding hot-plug command is not executed yet at the VPP, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined. 00: Reserved. 01: On 10: Blink (Processor drives 1 Hz square wave) 11: Off IIO does not generated the Attention_Indicator_On/Off/Blink messages on PCI Express when this field is written to by software.		
5:5	RW	0x0	hot_plug_interrupt_enable: When set to 1b, this bit enables generation of Hot-Plug interrupt MSI or INTx interrupt depending on the setting of the MSI enable bit in MSICTRL on enabled Hot-Plug events, provided ACPI mode for hot-plug is disabled. 0: disables interrupt generation on Hot-plug events 1: enables interrupt generation on Hot-plug events		



Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0xa8		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 (PCIe Mode) Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
4:4	RW	0x0	command_completed_interrupt_enable: This field enables software notification Interrupt - MSIINTx or WAKE when a command is completed by the Hot-plug controller connected to the PCI Express port 0 = disables hot-plug interrupts on a command completion by a hot-plug Controller 1 = Enables hot-plug interrupts on a command completion by a hot-plug Controller		
3:3	RW	0x0	presence_detect_changed_enable: This bit enables the generation of hot-plug interrupts or wake messages via a presence detect changed event. 0 = Disables generation of hot-plug interrupts or wake messages when a presence detect changed event happens. 1 = Enables generation of hot-plug interrupts or wake messages when a presence detect changed event happens.		
2:2	RW	0x0	mrl_sensor_changed_enable: This bit enables the generation of hot-plug interrupts or wake messages via a MRL Sensor changed event. 0: disables generation of hot-plug interrupts or wake messages when an MRL Sensor changed event happens. 1: Enables generation of hot-plug interrupts or wake messages when an MRL Sensor changed event happens.		
1:1	RW	0x0	power_fault_detected_enable: This bit enables the generation of hot-plug interrupts or wake messages via a power fault event. 0 = Disables generation of hot-plug interrupts or wake messages when a power fault event happens. 1 = Enables generation of hot-plug interrupts or wake messages when a power fault event happens.		
0:0	RW	0x0	attention_button_pressed_enable: This bit enables the generation of hot-plug interrupts or wake messages via an attention button pressed event. 0 = Disables generation of hot-plug interrupts or wake messages when the attention button is pressed. 1 = Enables generation of hot-plug interrupts or wake messages when the attention button is pressed.		

6.2.50 sltsts

PCI Express Slot Status

The PCI Express Slot Status register defines important status information for operations such as hot-plug and Power Management.



Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0 (PCIe Mode)
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0xaa			
Bit	Attr	Default	Description
8:8	RW1C	0x0	<p>data_link_layer_state_changed:</p> <p>This bit is set (if it is not already set) when the state of the Data Link Layer Link Active bit in the Link Status register changes. Software must read Data Link Layer Active field to determine the link state before initiating configuration cycles to the hot plugged device.</p>
7:7	RO_V	0x0	<p>electromechanical_latch_status:</p> <p>When read this register returns the current state of the Electromechanical Interlock (the EMILS pin) which has the defined encodings as:</p> <p>0 = Electromechanical Interlock Disengaged 1 = Electromechanical Interlock Engaged</p>
6:6	RO_V	0x0	<p>presence_detect_state:</p> <p>For ports with slots (where the Slot Implemented bit of the PCI Express Capabilities Registers is 1b), this field is the logical OR of the Presence Detect status determined via an in-band mechanism and sideband Present Detect pins.</p> <p>0 = Card/Module slot empty 1 = Card/module Present in slot (powered or unpowered)</p> <p>For ports with no slots, IIO hardwires this bit to 1b.</p> <p>Note: OS could get confused when it sees an empty PCI Express root port i.e. 'no slots + no presence', since this is now disallowed in the spec. So bios must hide all unused root ports devices in IIO config space, via the DEVHIDE register.</p>
5:5	RO_V	0x0	<p>mrl_sensor_state:</p> <p>This bit reports the status of an MRL sensor if it is implemented.</p> <p>0 = MRL Closed 1 = MRL Open</p>
4:4	RW1C	0x0	<p>command_completed:</p> <p>This bit is set by IIO when the hot-plug command has completed and the hot-plug controller is ready to accept a subsequent command. It is subsequently cleared by software after the field has been read and processed. This bit provides no guarantee that the action corresponding to the command is complete. Any write to SLTCON (regardless of the port is capable or enabled for hot-plug) is considered a 'hot-plug' command.</p> <p>If the port is not hot-plug capable or hot-plug enabled, then the hot-plug command does not trigger any action on the VPP port but the command is still completed via this bit.</p>
3:3	RW1C	0x0	<p>presence_detect_changed:</p> <p>This bit is set by IIO when the value reported in bit 6 is changes. It is subsequently cleared by software after the field has been read and processed.</p>
2:2	RW1C	0x0	<p>mrl_sensor_changed:</p> <p>This bit is set if the value reported in bit 5 changes. It is subsequently cleared by software after the field has been read and processed.</p>
1:1	RW1C	0x0	<p>power_fault_detected:</p> <p>This bit is set by IIO when a power fault event is detected by the power controller (which is reported via the VPP bit stream). It is subsequently cleared by software after the field has been read and processed.</p>
0:0	RW1C	0x0	<p>attention_button_pressed:</p> <p>This bit is set by IIO when the attention button is pressed. It is subsequently cleared by software after the field has been read and processed.</p> <p>IIO silently discards the AttentionButtonPressed message if received from PCI Express link without updating this bit.</p>



6.2.51 rootcon

PCI Express Root Control.

Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0xac		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
4:4	RW	0x0	crsswwisen: CRS software visibility Enable This bit, when set, enables the Root Port to return Configuration Request Retry Status (CRS) Completion Status to software. If 0, retry status cannot be returned to software.		
3:3	RW RW_L (Device 3 Function 0 only)	0x0	pmeinten: This field controls the generation of MSI interrupts INTx interrupts for PME messages. 1 = Enables interrupt generation upon receipt of a PME message 0 = Disables interrupt generation for PME messages		
2:2	RW	0x0	sefeen: System Error on Fatal Error Enable This field enables notifying the internal IIO core error logic of occurrence of an uncorrectable fatal error at the port or below its hierarchy. The internal core error logic of IIO then decides if/how to escalate the error further (pins/message etc). 1: indicates that an internal IIO core error logic notification should be generated if a fatal error (ERR_FATAL) is reported by any of the devices in the hierarchy associated with and including this port. 0: No internal IIO core error logic notification should be generated on a fatal error (ERR_FATAL) reported by any of the devices in the hierarchy associated with and including this port. Note that generation of system notification on a PCI Express fatal error is orthogonal to generation of an MSI/INTx interrupt for the same error. Both a system error and MSI/INTx can be generated on a fatal error or software can chose one of the two. Note that since this register is defined only in PCIe mode for Device#0, this bit will read a 0 in DMI mode. So, to enable core error logic notification on DMI mode fatal errors, BIOS must set bit 35 of MISCCTRLSTS to a 1 (to override this bit) on Device#0 in DMI mode.		
1:1	RW	0x0	senfeen: System Error on Non-Fatal Error Enable This field enables notifying the internal IIO core error logic of occurrence of an uncorrectable non-fatal error at the port or below its hierarchy. The internal IIO core error logic then decides if/how to escalate the error further (pins/message etc). 1: indicates that a internal IIO core error logic notification should be generated if a non-fatal error (ERR_NONFATAL) is reported by any of the devices in the hierarchy associated with and including this port. 0: No internal core error logic notification should be generated on a non-fatal error (ERR_NONFATAL) reported by any of the devices in the hierarchy associated with and including this port. Note that generation of system notification on a PCI Express non-fatal error is orthogonal to generation of an MSI/INTx interrupt for the same error. Both a system error and MSI/INTx can be generated on a non-fatal error or software can chose one of the two. Note that since this register is defined only in PCIe mode for Device#0, this bit will read a 0 in DMI mode. So, to enable core error logic notification on DMI mode non-fatal errors, BIOS must set bit 34 of MISCCTRLSTS to a 1 (to override this bit) on Device#0 in DMI mode.		



Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0xac				

Bit	Attr	Default	Description
0:0	RW	0x0	<p>seceen:</p> <p>System Error on Correctable Error Enable</p> <p>This field controls notifying the internal IIO core error logic of the occurrence of a correctable error in the device or below its hierarchy. The internal core error logic of IIO then decides if/how to escalate the error further (pins/ message etc).</p> <p>1: indicates that an internal core error logic notification should be generated if a correctable error (ERR_COR) is reported by any of the devices in the hierarchy associated with and including this port.</p> <p>0: No internal core error logic notification should be generated on a correctable error (ERR_COR) reported by any of the devices in the hierarchy associated with and including this port.</p> <p>Note that generation of system notification on a PCI Express correctable error is orthogonal to generation of an MSI/INTx interrupt for the same error. Both a system error and MSI/INTx can be generated on a correctable error or software can chose one of the two.</p> <p>Note that since this register is defined only in PCIe mode for Device#0, this bit will read a 0 in DMI mode. So, to enable core error logic notification on DMI mode correctable errors, BIOS must set bit 33 of MISCCTRLSTS to a 1 (to override this bit) on Device#0 in DMI mode.</p>



6.2.52 rootcap

PCI Express Root Capabilities.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0xae				

Bit	Attr	Default	Description
0:0	RO RW_O (Device 0 Function 0)	0x1 0x0 (Device 0 Function 0, DMI2 mode)	crs_software_visibility: This bit, when set, indicates that the Root Port is capable of returning Configuration Request Retry Status (CRS) Completion Status to software. Intel® Xeon® Processor E5 v4 product family supports this capability.

6.2.53 rootsts

PCI Express Root Status.

Type:	CFG	PortID:	N/A	Function:	0 (PCIe Mode)
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0xb0				

Bit	Attr	Default	Description
17:17	RO_V	0x0	pme_pending: This field indicates that another PME is pending when the PME Status bit is set. When the PME Status bit is cleared by software; the pending PME is delivered by hardware by setting the PME Status bit again and updating the Requestor ID appropriately. The PME pending bit is cleared by hardware if no more PMEs are pending.
16:16	RW1C	0x0	pme_status: This field indicates a PM_PME message (either from the link or internally from within that root port) was received at the port. 1: PME was asserted by a requester as indicated by the PME Requester ID field This bit is cleared by software by writing a '1'. Note that the root port itself could be the source of a PME event when a hot-plug event is observed when the port is in D3hot state.
15:0	RO_V	0x0	pme_requester_id: This field indicates the PCI requester ID of the last PME requester. If the root port itself was the source of the (virtual) PME message, then a RequesterID of CPUBUSNO0:DevNo:FunctionNo is logged in this field.



6.2.54 devcap2

PCI Express Device Capabilities 2 Register.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0xb4				

Bit	Attr	Default	Description
13:12	RW_O	0x1	tph_completer_supported: Indicates the support for TLP Processing Hints. Processor does not support the extended TPH header. 00: TPH and Extended TPH Completer not supported. 01: TPH Completer supported; Extended TPH Completer not supported. 10: Reserved. 11: Both TPH and Extended TPH Completer supported.
9:9	RO	0x1	atomic128bcascompsup:
8:8	RO	0x1	atomic64bccompsup:
7:7	RO	0x1	atomic32bccompsup:
6:6	RO	0x0	atomicroutsup:
5:5	RW_O	0x1	ari_en: Alternative RID InterpretationCapable This bit is set to 1b indicating Root Port supports this capability.
4:4	RO	0x1	cmpltodissup: Completion Timeout Disable Supported IIO supports disabling completion timeout
3:0	RO	0xe	cmpltovalsup: Completion Timeout Values Supported This field indicates device support for the optional Completion Timeout programmability mechanism. This mechanism allows system software to modify the Completion Timeout range. Bits are one-hot encoded and set according to the table below to show timeout value ranges supported. A device that supports the optional capability of Completion Timeout Programmability must set at least two bits. Four time values ranges are defined: Range A: 50 us to 10 ms Range B: 10 ms to 250 ms Range C: 250 ms to 4 s Range D: 4 s to 64 s Bits are set according to table below to show timeout value ranges supported. 0000b: Completions Timeout programming not supported – values is fixed by implementation in the range 50 us to 50 ms. 0001b: Range A 0010b: Range B 0011b: Range A & B 0110b: Range B & C 0111b: Range A, B, & C 1110b: Range B, C & D 1111b: Range A, B, C & D All other values are reserved. IIO supports timeout values up to 10 ms-64 s



6.2.55 devctrl2

PCI Express Device Control Register 2.

Type: CFG Bus: 0 Offset: 0xf8		PortID: N/A Device: 0		Function: 0 (DMI 2 Mode)
Bus: 0 Device: 0 Function: 0 Offset: 0xb8		Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 (PCIe Mode) Function: 0-1 Function: 0-3 Function: 0-3
Bit	Attr	Default	Description	
7:7	RO	0x0	atomicgressblock:	
6:6	RO	0x0	atomicregen:	
5:5	RW_L	0x0	ari: Alternative RID InterpretationEnable Applies only to root ports. When set to 1b, ARI is enabled for the Root Port. For Device#0 in DMI mode, this bit is ignored	
4:4	RW_V (Device 2 and 3 Function 0) RW (Device 0 Function0, Device 2 and 3 Function 1-3)	0x0 0x1 (Device 0 Function 0)	compltodis: Completion Timeout Disable When set to 1b, this bit disables the Completion Timeout mechanism for all NP tx that IIO issues on the PCIe/DMI link. When 0b, completion timeout is enabled. Software can change this field while there is active traffic in the root/ DMI port.	
3:0	RW_V (Device 2 and 3 Function 0) RW (Device 0 Function0, Device 2 and 3 Function 1-3)	0x0	compltovall: Completion Timeout Value on NP Tx that IIO issues on PCIe/DMI In Devices that support Completion Timeout programmability, this field allows system software to modify the Completion Timeout range. The following encodings and corresponding timeout ranges are defined: 0000b = 10 ms to 50 ms 0001b = Reserved (IIO aliases to 0000b) 0010b = Reserved (IIO aliases to 0000b) 0101b = 16 ms to 55 ms 0110b = 65 ms to 210 ms 1001b = 260 ms to 900 ms 1010b = 1 s to 3.5 s 1101b = 4 s to 13 s 1110b = 17 s to 64 s When software selects 17 s to 64 s range, CTOCTRL further controls the timeout value within that range. For all other ranges selected by OS, the timeout value within that range is fixed in IIO hardware. Software can change this field while there is active traffic in the root port. This value will also be used to control PME_TO_ACK Timeout. That is this field sets the timeout value for receiving a PME_TO_ACK message after a PME_TURN_OFF message has been transmitted. The PME_TO_ACK Timeout has meaning only if bit 6 of MISCCTRLSTS register is set to a 1b.	



6.2.56 Inkcap2

PCI Express Link Capabilities 2.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0xbc				

Bit	Attr	Default	Description
7:1	RW_0	0x7 0x3 (Device 0 Function 0)	Inkspdvec: Supported Link Speeds Vector - This field indicates the supported Link speeds of the associated Port. For each bit, a value of 1b indicates that the corresponding Link speed is supported; otherwise, the Link speed is not supported. Bit definitions are: Bit 1 2.5 GTs set in CPU Bit 2 5.0 GTs set in CPU Bit 3 8.0 GTs set in CPU Bits 7:4 reserved

6.2.57 Inkcon2

Type:	CFG	PortID:	N/A	Function:	0 (DMI2 Mode)
Bus:	0	Device:	0	Function:	0 (PCIe Mode)
Offset:	0x1c0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0xc0				

Bit	Attr	Default	Description
15:12 12:12 (Device 0 Function 0)	RWS	0x0	compliance_de_emphasis: For 8 GT/s Data Rate: This bit sets the Transmitter Preset level in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b. The Encodings are defined as follows: 0000b: -6 dB for de-emphasis, 0 dB for preshoot 0001b: -3.5 dB for de-emphasis, 0 dB for preshoot 0010b: -4.5 dB for de-emphasis, 0 dB for preshoot 0011b: -2.5 dB for de-emphasis, 0 dB for preshoot 0100b: 0 dB for de-emphasis, 0 dB for preshoot 0101b: 0 dB for de-emphasis, 2 dB for preshoot 0110b: 0 dB for de-emphasis, 2.5 dB for preshoot 0111b: -6 dB for de-emphasis, 3.5 dB for preshoot 1000b: -3.5 dB for de-emphasis, 3.5 dB for preshoot 1001b: 0 dB for de-emphasis, 3.5 dB for preshoot Others: reserved For 5 GT/s Data Rate: This bit sets the de-emphasis level in Polling.Compliance state if the entry occurred due to the Enter Compliance bit being 1b. Encodings: 0001b: -3.5 dB 0000b: -6 dB For 2.5 GT/s Data Rate: The setting of this field has no effect. Components that support only 2.5 GT/s speed are permitted to hardwire this field to 0h. Notes: This bit is intended for debug, compliance testing purposes. System firmware and software is allowed to modify this bit only during debug or compliance testing.




Type: CFG Bus: 0 Offset: 0x1c0		PortID: N/A Device: 0		Function: 0 (DMI2 Mode)	
Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0xc0		Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 (PCIe Mode) Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
11:11	RWS	0x0	compliance_sos: When set to 1b, the LTSSM is required to send SKP Ordered Sets periodically in between the (modified) compliance patterns.		
10:10	RWS	0x0	enter_modified_compliance: When this bit is set to 1b, the device transmits Modified Compliance Pattern if the LTSSM enters Polling.Compliance substate.		
9:7	RWS_V	0x0	transmit_margin: This field controls the value of the nondeemphasized voltage level at the Transmitter pins.		
6:6	RW_O	0x0	selectable_de_emphasis: When the Link is operating at 5.0 GT/s speed, this bit selects the level of de-emphasis for an Upstream component.Encodings: 1b -3.5 dB 0b -6 dB When the Link is operating at 2.5 GT/s speed, the setting of this bit has no effect.		
5:5	RWS	0x0	hardware_autonomous_speed_disable: When Set, this bit disables hardware from changing the Link speed for device specific reasons other than attempting to correct unreliable Link operation by reducing Link speed.		
4:4	RWS_V	0x0	enter_compliance: Software is permitted to force a link to enter Compliance mode at the speed indicated in the Target Link Speed field by setting this bit to 1b in both components on a link and then initiating a hot reset on the link.		
3:0	RWS_V	0x3 0x2 (Device 0 Function 0)	target_link_speed: This field sets an upper limit on link operational speed by restricting the values advertised by the upstream component in its training sequences. Defined encodings are: 0001b 2.5Gb/s Target Link Speed 0010b 5Gb/s Target Link Speed 0011b 8Gb/s Target Link Speed (Reserved for Device 0 Function 0) All other encodings are reserved. If a value is written to this field that does not correspond to a speed included in the Supported Link Speeds field, IIO will default to Gen1 speed. This field is also used to set the target compliance mode speed when software is using the Enter Compliance bit to force a link into compliance mode.		



6.2.58 Inksts2

PCI Express Link Status Register 2.

Type: CFG Bus: 0 Offset: 0x1c2		PortID: N/A Device: 0		Function: 0 (DMI2 Mode)
Bus: 0 Device: 0 Function: 0 (PCIe Mode)		Device: 0 Device: 1 Device: 2 Device: 3		Function: 0-1 Function: 0-3 Function: 0-3
Bus: 0 Device: 0 Function: 0-3 Offset: 0xc2				
Bit	Attr	Default	Description	
5:5	RW1CS	0x0	Inkeqreq: This bit is Set by hardware to request Link equalization process to be performed on the link. Reserved for Device 0 Function 0.	
4:4	RO_V	0x0	eqph3_succ: When set to 1b, this indicates that Phase 3 of the Transmitter Equalization procedure has successfully completed. Reserved for Device 0 Function 0.	
3:3	RO_V	0x0	eqph2_succ: When set to 1b, this indicates that Phase 2 of the Transmitter Equalization procedure has successfully completed. Reserved for Device 0 Function 0.	
2:2	RO_V	0x0	eqph1_succ: When set to 1b, this indicates that Phase 1 of the Transmitter Equalization procedure has successfully completed. Reserved for Device 0 Function 0.	
1:1	RO_V	0x0	eqcmp: When set to 1b, this indicates that the Transmitter Equalization procedure has completed. Reserved for Device 0 Function 0.	
0:0	RO_V	0x0	current_de_emphasis_level: When operating at Gen2 speed, this reports the current de-emphasis level. This field is Unused for Gen1 speeds 1b: -3.5 dB 0b: -6 dB	



6.2.59 pmcap

Power Management Capabilities

The Power Management Capabilities Register defines the capability ID, next pointer and other power management related support. The following Power Management registers/capabilities are added for software compliance.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0xe0				
Bit	Attr	Default	Description		
31:27	RO_V	0x19	pme_support: For DMI it should be 0, 0x19 for the PCIe ports. Bits 31, 30 and 27 must be set to q1q for PCI-PCI bridge structures representing ports on root complexes.		
26:26	RO	0x0	d2_support: Does not support power management state D2.		
25:25	RO	0x0	d1_support: Does not support power management state D1.		
24:22	RO	0x0	aux_current:		
21:21	RO	0x0	device_specific_initialization:		
19:19	RO	0x0	pme_clock: This field is hardwired to 0h as it does not apply to PCI Express.		
18:16	RO	0x3	version: This field is set to 3h Power Management 1.2 compliant as version number. Bit is RW-O to make the version 2h incase legacy OS'es have any issues.		
15:8	RO	0x0	next_capability_pointer: This is the last capability in the chain and hence set to 0.		
7:0	RO	0x1	capability_id: Provides the Power Management capability ID assigned by PCI-SIG.		

6.2.60 pmcsr

Power Management Control and Status Register

This register provides status and control information for Power Management events in the PCI Express port of the IIO.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0xe4				
Bit	Attr	Default	Description		
31:24	RO	0x0	data: N/A		
23:23	RO	0x0	bus_power_clock_control_enable: N/A		



Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0xe4				
Bit	Attr	Default	Description		
22:22	RO	0x0	b2_b3_support: N/A		
15:15	RW1CS	0x0	pme_status: N/A		
14:13	RO	0x0	data_scale: N/A		
12:9	RO	0x0	data_select: N/A		
8:8	RWS RWS_L (Device 3 Function 0)	0x0	pme_enable: N/A		
3:3	RW_O	0x1	no_soft_reset: Indicates does not reset its registers when transitioning from D3hot to D0.		
1:0	RW RW_L (Device 0 Function 0)	0x0	power_state: This 2-bit field is used to determine the current power state of the function and to set a new power state as well. 00: D0 01: D1 (not supported by IIO) 10: D2 (not supported by IIO) 11: D3hot If Software tries to write 01 or 10 to this field, the power state does not change from the existing power state which is either D0 or D3hot and nor do these bits1:0 change value. When in D3hot state, IOxAPIC will a) respond to only Type 0 configuration transactions targeted at the device's configuration space, when in D3hot state c) will not respond to memory i.e. D3hot state is equivalent to MSE , accesses to MBAR region note: ABAR region access still go through in D3hot state, if it enabled d) will not generate any MSI writes		

6.2.61 xpreut_hdr_ext

REUT PCIe Header Extended.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x100				
Bit	Attr	Default	Description		
31:20	RO RO_V (Device 0 Function 0)	0x110	pcienextptr: Next Capability Pointer This field contains the offset to the next PCI capability structure or 00h if no other items exist in the linked list of capabilities. In DMI Mode, it points to the Vendor Specific Error Capability. In PCIe Mode, it points to the ACS Capability.		



6.2.63 xpreut_hdr_lef

REUT Header Leaf Capability.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x108				
Bit	Attr	Default	Description		
15:8	RO_V	0x38 0x30 (Device 0 Function 0)	leafreutdevnum: This field identifies the PCI Device/Function # where the REUT engine associated with this link resides. Device6 = 00110b & function0 = 000b = 30h		
7:0	RO_V	0x7	leafreutengid: This field identifies the REUT engine associated with the link (same as the REUT ID).		

6.2.64 accscaphdr

Access Control Services Extended Capability Header.

Type:	CFG	PortID:	N/A	Function:	0 (PCIe Mode)
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x110				
Bit	Attr	Default	Description		
31:20	RO_V	0x148	next_capability_offset: This field points to the next Capability in extended configuration space. In PCIe Mode, it points to the Advanced Error Capability.		
19:16	RO	0x1	capability_version: Set to 1h for this version of the PCI Express logic		
15:0	RO	0xd	pci_express_extended_cap_id: Assigned for Access Control Services capability by PCISIG.		



6.2.65 acscap

Access Control Services Capability Register.

Type:	CFG	PortID:	N/A	Function:	0 (PCIe mode)
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x114				

Bit	Attr	Default	Description
15:8	RO	0x0	egress_control_vector_size: N/A for IIO
6:6	RO	0x0	t: Applies only to root ports. Indicates that the component does not implement ACS Direct Translated P2P.
5:5	RO	0x0	e: Applies only to root ports. Indicates that the component does not implement ACS P2P Egress Control.
4:4	RO_V (Device 2 and 3 Function 0) RO (Device 0 Function 0, Device 2 and 3 Function 1-3)	0x1	u: Applies only to root ports. Indicates that the component implements ACS Upstream Forwarding.
3:3	RO_V (Device 2 and 3 Function 0) RO (Device 0 Function 0, Device 2 and 3 Function 1-3)	0x1	c: Applies only to root ports. Indicates that the component implements ACS P2P Completion Redirect.
2:2	RO_V (Device 2 and 3 Function 0) RO (Device 0 Function 0, Device 2 and 3 Function 1-3)	0x1	r: Applies only to root ports. Indicates that the component implements ACS P2P Request Redirect.
1:1	RO_V (Device 2 and 3 Function 0) RO (Device 0 Function 0, Device 2 and 3 Function 1-3)	0x1	b: Applies only to root ports. Indicates that the component implements ACS Translation Blocking.
0:0	RO_V (Device 2 and 3 Function 0) RO (Device 0 Function 0, Device 2 and 3 Function 1-3)	0x1	v: Applies only to root ports. Indicates that the component implements ACS Source Validation.

6.2.66 acsctrl

Access Control Services Control Register.

Type:	CFG	PortID:	N/A	Function:	0 (PCIe Mode)
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x116				

Bit	Attr	Default	Description
6:6	RO	0x0	t: Applies only to root ports. This is hardwired to 0b as the component does not implement ACS Direct Translated P2P.




Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0 (PCIe Mode)
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0x116			
Bit	Attr	Default	Description
5:5	RO	0x0	e: Applies only to root ports. The component does not implement ACS P2P Egress Control and hence this bit should not be used by SW.
4:4	RW_L (Device 2 and 3 Function 0) RW (Device 0 Function 0, Device 2 and 3 Function 1-3)	0x0	u: When this bit is set, transactions arriving from a root port that target the same port back down, will be forwarded. Normally such traffic would be aborted. Applies only to root ports.
3:3	RW_L (Device 2 and 3 Function 0) RW (Device 0 Function 0, Device 2 and 3 Function 1-3)	0x0	c: Applies only to root ports. Determines when the component redirects peer-to-peer Completions upstream; applicable only to Read Completions whose Relaxed Ordering Attribute is clear.
2:2	RW_L (Device 2 and 3 Function 0) RW (Device 0 Function 0, Device 2 and 3 Function 1-3)	0x0	r: When this bit is set, transactions arriving from a root port that target the same port back down, will be forwarded. Normally such traffic would be aborted. Applies only to root ports.
1:1	RW_L (Device 2 and 3 Function 0) RW (Device 0 Function 0, Device 2 and 3 Function 1-3)	0x0	b: Applies only to root ports. When set, the component blocks all upstream Memory Requests whose Address Translation AT field is not set to the default value.
0:0	RW_L (Device 2 and 3 Function 0) RW (Device 0 Function 0, Device 2 and 3 Function 1-3)	0x0	v: Applies only to root ports. When set, the component validates the Bus Number from the Requester ID of upstream Requests against the secondary subordinate Bus Numbers.

6.2.67 apicbase

ACPI Base Register.

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0x140			
Bit	Attr	Default	Description
11:1	RW	0x0	addr: Bits 31:20 are assumed to be 0xFECh. Bits 8:0 are a don't care for address decode. Address decoding to the APIC range is done as APICBASE.ADDR[31:8] <= A[31:8] <= APICLIMIT.ADDR[31:8]. Outbound accesses to the APIC range are claimed by the root port and forwarded to PCIe, if bit 0 is set, even if the MSE bit of the root port is clear or the root port itself is in D3hot state.
0:0	RW	0x0	en: enables the decode of the APIC window



6.2.68 apiclimit

ACPI Limit Register.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x142				

Bit	Attr	Default	Description
11:1	RW	0x0	addr: Applies only to root ports. Bits 31:20 are assumed to be 0xFECh. Bits 8:0 are a don't care for address decode. Address decoding to the APIC range is done as $APICBASE.ADDR[31:8] \leq A[31:8] \leq APICLIMIT.ADDR[31:8]$. Outbound accesses to the APIC range are claimed by the root port and forwarded to PCIe, if the range is enabled, even if the MSE bit of the root port is clear or the root port itself is in D3hot state.

6.2.69 vsecphdr

PCI Express Enhanced Capability Header - DMI2 Mode.

Type:	CFG	PortID:	N/A	Function:	0 (DMI2 Mode)
Bus:	0	Device:	0		
Offset:	0x144				

Bit	Attr	Default	Description
31:20	RO	0x1d0	next_capability_offset: This field points to the next Capability in extended configuration space or is 0 if it is that last capability.
19:16	RO	0x1	capability_version: Set to 1h for this version of the PCI Express logic.
15:0	RO	0xb	pci_express_extended_cap_id: Assigned for Vendor Specific Capability.

6.2.70 vshdr

Vendor Specific Header - DMI2 Mode.

Type:	CFG	PortID:	N/A	Function:	0 (DMI2 Mode)
Bus:	0	Device:	3		
Offset:	0x148				

Bit	Attr	Default	Description
31:20	RO	0x3c	vsec_length: This field points to the next Capability in extended configuration space which is the ACS capability at 150h.
19:16	RO	0x1	vsec_version: Set to 1h for this version of the PCI Express logic
15:0	RO	0x4	vsec_id: Identifies Intel Vendor Specific Capability for AER on DMI



6.2.71 errcaphdr

PCI Express Enhanced Capability Header - Root Ports.

Type:	CFG	PortID:	N/A	Function:	0 (PCIe Mode)
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x148				

Bit	Attr	Default	Description
31:20	RO	0x1d0	next_capability_offset: This field points to the next Capability in extended configuration space or is 0 if it is that last capability.
19:16	RO	0x1	capability_version: Set to 1h for this version of the PCI Express logic
15:0	RO	0x1	pci_express_extended_cap_id: Assigned for advanced error reporting

6.2.72 uncrrsts

Uncorrectable Error Status.

This register identifies uncorrectable errors detected for PCI Express/DMI port.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x14c				

Bit	Attr	Default	Description
21:21	RW1CS	0x0	acs_violation_status:
20:20	RW1CS	0x0	received_an_unsupported_request:
19:19	RW1CS	0x0	ecrc_error_status:
18:18	RW1CS	0x0	malformed_tlp_status:
17:17	RW1CS	0x0	receiver_buffer_overflow_status:
16:16	RW1CS	0x0	unexpected_completion_status:
15:15	RW1CS	0x0	completer_abort_status:
14:14	RW1CS	0x0	completion_time_out_status:
13:13	RW1CS	0x0	flow_control_protocol_error_status:
12:12	RW1CS	0x0	poisoned_tlp_status:
5:5	RW1CS	0x0	surprise_down_error_status:
4:4	RW1CS	0x0	data_link_protocol_error_status:

6.2.73 uncermsk

Uncorrectable Error Mask.

This register masks uncorrectable errors from being signaled.



Type: CFG		PortID: N/A			
Bus: 0		Device: 0		Function: 0	
Bus: 0		Device: 1		Function: 0-1	
Bus: 0		Device: 2		Function: 0-3	
Bus: 0		Device: 3		Function: 0-3	
Offset: 0x150					
Bit	Attr	Default	Description		
21:21	RWS	0x0	acs_violation_mask:		
20:20	RWS	0x0	unsupported_request_error_mask:		
18:18	RWS	0x0	malformed_tlp_mask:		
19:19	RWS	0x0	ecrc_error_mask:		
17:17	RWS	0x0	receiver_buffer_overflow_mask:		
16:16	RWS	0x0	unexpected_completion_mask:		
15:15	RWS	0x0	completer_abort_mask:		
14:14	RWS	0x0	completion_time_out_mask:		
13:13	RWS	0x0	flow_control_protocol_error_mask:		
12:12	RWS	0x0	poisoned_tlp_mask:		
5:5	RWS	0x0	surprise_down_error_mask:		
4:4	RWS	0x0	data_link_layer_protocol_error_mask:		

6.2.74 uncerrsev

Uncorrectable Error Severity.

This register indicates the severity of the uncorrectable errors.

Type: CFG		PortID: N/A			
Bus: 0		Device: 0		Function: 0	
Bus: 0		Device: 1		Function: 0-1	
Bus: 0		Device: 2		Function: 0-3	
Bus: 0		Device: 3		Function: 0-3	
Offset: 0x154					
Bit	Attr	Default	Description		
21:21	RWS	0x0	acs_violation_severity:		
20:20	RWS	0x0	unsupported_request_error_severity:		
19:19	RWS	0x1	ecrc_error_severity:		
18:18	RWS	0x1	malformed_tlp_severity:		
17:17	RWS	0x1	receiver_buffer_overflow_severity:		
16:16	RWS	0x0	unexpected_completion_severity:		
15:15	RWS	0x0	completer_abort_severity:		
14:14	RWS	0x0	completion_time_out_severity:		
13:13	RWS	0x1	flow_control_protocol_error_severity:		
12:12	RWS	0x0	poisoned_tlp_severity:		
5:5	RWS	0x1	surprise_down_error_severity:		
4:4	RWS	0x1	data_link_protocol_error_severity:		



6.2.75 corerrsts

Correctable Error Status.

This register identifies the status of the correctable errors that have been detected by the PCI Express port.

Type:	CFG	PortID:	N/A		
Bus:	0	Device:	0	Function:	0
Bus:	0	Device:	1	Function:	0-1
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x158				
Bit	Attr	Default	Description		
13:13	RW1CS	0x0	advisory_non_fatal_error_status:		
12:12	RW1CS	0x0	replay_timer_time_out_status:		
8:8	RW1CS	0x0	replay_num_rollover_status:		
7:7	RW1CS	0x0	bad_dllp_status:		
6:6	RW1CS	0x0	bad_tlp_status:		
0:0	RW1CS	0x0	receiver_error_status:		

6.2.76 corerrmsk

Correctable Error Mask.

This register masks correctable errors from being signaled.

Type:	CFG	PortID:	N/A		
Bus:	0	Device:	0	Function:	0
Bus:	0	Device:	1	Function:	0-1
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x15c				
Bit	Attr	Default	Description		
13:13	RWS	0x1	advisory_non_fatal_error_mask:		
12:12	RWS	0x0	replay_timer_time_out_mask:		
8:8	RWS	0x0	replay_num_rollover_mask:		
7:7	RWS	0x0	bad_dllp_mask:		
6:6	RWS	0x0	bad_tlp_mask:		
0:0	RWS	0x0	receiver_error_mask:		

6.2.77 errcap

Advanced Error capabilities and Control Register.



Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0x160			
Bit	Attr	Default	Description
8:8	RWS	0x0	ecrc_check_enable: PCIe ECRC enable.
7:7	RW_O	0x1	ecrc_check_capable: PCIe ECRC capable.
6:6	RWS	0x0	ecrc_generation_enable: PCIe ECRC generation enable.
5:5	RW_O	0x1	ecrc_generation_capable: PCIe ECRC generation capable.
4:0	ROS_V	0x0	first_error_pointer: The First Error Pointer is a read-only register that identifies the bit position of the first unmasked error reported in the Uncorrectable Error register. In case of two errors happening at the same time, fatal error gets precedence over non-fatal, in terms of being reported as first error. This field is rearmed to capture new errors when the status bit indicated by this field is cleared by software.

6.2.78 hdrlog[0:3]

Header Log 0-3.

This register contains the header log when the first error occurs. Headers of the subsequent errors are not logged.


Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0x164, 0x168, 0x16c, 0x170			
Bit	Attr	Default	Description
31:0	ROS_V	0x0	hdr: Logs the first DWORD of the header on an error condition.

6.2.79 rperrcmd

Root Port Error Command.

This register controls behavior upon detection of errors.

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0x174			
Bit	Attr	Default	Description



2:2	RW	0x0	fatal_error_reporting_enable: Applies to root ports onlyEnable MSIINTx interrupt on fatal errors when set.
1:1	RW	0x0	non_fatal_error_reporting_enable: Applies to root ports onlyEnable interrupt on a non-fatal error when set.
0:0	RW	0x0	correctable_error_reporting_enable: Applies to root ports onlyEnable interrupt on correctable errors when set.

6.2.80 rperrsts

Root Port Error Status.

The Root Error Status register reports status of error Messages (ERR_COR), ERR_NONFATAL, and ERR_FATAL) received by the Root Complex in IIO, and errors detected by the Root Port itself (which are treated conceptually as if the Root Port had sent an error Message to itself). The ERR_NONFATAL and ERR_FATAL Messages are grouped together as uncorrectable. Each correctable and uncorrectable (Non-fatal and Fatal) error source has a first error bit and a next error bit associated with it respectively. When an error is received by a Root Complex, the respective first error bit is set and the Requestor ID is logged in the Error Source Identification register. A set individual error status bit indicates that a particular error category occurred; software may clear an error status by writing a 1 to the respective bit. If software does not clear the first reported error before another error Message is received of the same category (correctable or uncorrectable), the corresponding next error status bit will be set but the Requestor ID of the subsequent error Message is discarded. The next error status bits may be cleared by software by writing a 1 to the respective bit as well.

Type:	CFG	PortID:	N/A
Bus:	0	Device:	0
Bus:	0	Device:	1
Bus:	0	Device:	2
Bus:	0	Device:	3
Offset:	0x178	Function:	0
		Function:	0-1
		Function:	0-3
		Function:	0-3

Bit	Attr	Default	Description
31:27	RO	0x0	advanced_error_interrupt_message_number: Advanced Error Interrupt Message Number offset between base message data an the MSI message if assigned more than one message number. IIO hardware automatically updates this register to 0x1h if the number of messages allocated to the root port is 2.
6:6	RW1CS	0x0	fatal_error_messages_received: Set when one or more Fatal Uncorrectable error Messages have been received.
5:5	RW1CS	0x0	non_fatal_error_messages_received: Set when one or more Non-Fatal Uncorrectable error Messages have been received.
4:4	RW1CS	0x0	first_uncorrectable_fatal: Set when bit 2 is set (from being clear) and the message causing bit 2 to be set is an ERR_FATAL message.
3:3	RW1CS	0x0	multiple_error_fatal_nonfatal_received: Set when either a fatal or a non-fatal error message is received and Error Fatal/Nonfatal Received is already set, that is, log from the 2nd Fatal or Non fatal error message onwards.
2:2	RW1CS	0x0	error_fatal_nonfatal_received: Set when either a fatal or a non-fatal error message is received and this bit is already not set. i.e. log the first error message. Note that when this bit is set bit 3 could be either set or clear.



Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x178				
Bit	Attr	Default	Description		
1:1	RW1CS	0x0	multiple_correctable_error_received: Set when either a correctable error message is received and Correctable Error Received bit is already set, that is, log from the 2nd Correctable error message onwards .		
0:0	RW1CS	0x0	correctable_error_received: Set when a correctable error message is received and this bit is already not set, that is, log the first error message.		

6.2.81 errsid

Error Source Identification.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x17c				
Bit	Attr	Default	Description		
31:16	ROS_V	0x0	fatal_non_fatal_error_source_id: Requestor ID of the source when an Fatal or Non Fatal error message is received and the Error Fatal/Nonfatal Received bit is not already set, that is, log ID of the first Fatal or Non Fatal error message. Note that when the root port itself is the cause of the received message (virtual message), then a Source ID of CPUBUSNO0:DevNo:0 is logged into this register.		
15:0	ROS_V	0x0	correctable_error_source_id: Requestor ID of the source when a correctable error message is received and the Correctable Error Received bit is not already set, that is, log ID of the first correctable error message. Note that when the root port itself is the cause of the received message (virtual message), then a Source ID of CPUBUSNO0:DevNo:0 is logged into this register.		

6.2.82 perfctrlsts_0

Performance Control and Status Register 0.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x180				
Bit	Attr	Default	Description		
20:16	RW	0x18	outstanding_requests_gen1:		
13:8	RW	0x30	outstanding_requests_gen2:		



Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0x180		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
7:7	RW	0x1	<p>use_allocating_flow_wr:</p> <p>Use Allocating Flows for 'Normal Writes' on VC0 and VCp</p> <p>1: Use allocating flows for the writes that meet the following criteria.</p> <p>0: Use non-allocating flows for writes that meet the following criteria. (TPH=0 OR TPHDIS=1 OR (TPH=1 AND Tag=0 AND CIPCTRL[28]=1)) AND (NS=0 OR NoSnoopOpWrEn=0) AND Non-DCA Write</p> <p>Note: VC1/VCm traffic is not impacted by this bit in Dev#0 When allocating flows are used for the above write types, IIO does not send a Prefetch Hint message. Current recommendation for BIOS is to just leave this bit at default of 1b for all but DMI port. For DMI port when operating in DMI mode, this bit must be left at default value and when operating in PCIe mode, this bit should be set by BIOS. Note there is a coupling between the usage of this bit and bits 2 and 3. TPHDIS is bit 0 of this register NoSnoopOpWrEn is bit 3 of this register</p>		
6:6	RW	0x0	<p>vcp_nosnoopopen:</p> <p>Enables inbound VCp traffic with NS=1 to issue non-snoop IDI/QPI requests.</p>		
5:5	RW	0x0	<p>vc1m_nosnoopopdis — Disables inbound VC1/m traffic with NS=1 from issuing nonsnoop IDI/QPI requests.</p>		
4:4	RW	0x1	<p>read_stream_interleave_size:</p>		
3:3	RW	0x0	<p>nosnoopopwren:</p> <p>Enable No-Snoop Optimization on VC0 writes and VCp writes</p> <p>This applies to writes with the following conditions: NS=1 AND (TPH=0 OR TPHDIS=1)</p> <p>1: Inbound writes to memory with above conditions will be treated as non-coherent (no snoops) writes on Intel QPI</p> <p>0: Inbound writes to memory with above conditions will be treated as allocating or non-allocating writes, depending on bit 4 in this register.</p> <p>If TPH=1 and TPHDIS=0 then NS is ignored and this bit is ignored</p> <p>VC1/VCm writes are not controlled by this bit since they are always non-snoop and can be no other way.</p> <p>Current recommendation for BIOS is to just leave this bit at default of 0b.</p>		
2:2	RW	0x0	<p>nosnoopoprden:</p> <p>Enable No-Snoop Optimization on VC0 reads and VCp reads</p> <p>This applies to reads with the following conditions: NS=1 AND (TPH=0 OR TPHDIS=1)</p> <p>1: When the condition is true for a given inbound read request to memory, it will be treated as non-coherent (no snoops) reads on Intel QPI.</p> <p>0: When the condition is true for a given inbound read request to memory, it will be treated as normal snooped reads from PCIe (which trigger a PCIRdCurrent or DRd.UC on IDI).</p> <p>Notes: If TPH=1 and TPHDIS=0 then NS is ignored and this bit is ignored</p> <p>VC1 and VCm reads are not controlled by this bit and those reads are always non-snoop.</p> <p>Current recommendation for BIOS is to just leave this bit at default of 0b.</p>		
1:1	RW	0x0	<p>read_passing_read_disable:</p> <p>Disable reads bypassing other reads.</p>		



Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x180				
Bit	Attr	Default	Description		
0:0	RW	0x1	read_stream_policy:		

6.2.83 perfctrlsts_1

Performance Control and Status Register 1.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x184				
Bit	Attr	Default	Description		
9:9	RW	0x0	tphdis: TLP Processing Hint Disable When set, writes or reads with TPH=1, will be treated as if TPH=0.		
8:8	RW	0x0	dca_reqid_override: DCA Requester ID Override When this bit is set, Requester ID match for DCA writes is bypassed. All writes from the port are treated as DCA writes and the tag field will convey if DCA is enabled or not and the target information.		
3:3	RW	0x0	max_read_completion_combine_size:		

6.2.84 miscctrlsts_0

MISC Control and Status Register 0.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x188				
Bit	Attr	Default	Description		
31:31	RW	0x0	disable_l0s_on_transmitter: When set, IIO never puts its tx in L0s state, even if OS enables it via the Link Control register.		
30:30	RW_O	0x1	inbound_io_disable:		
29:29	RW	0x1	cfg_to_en: Disables/enables config timeouts, independently of other timeouts.		
28:28	RW	0x0	to_dis: Disables timeouts completely.		



Integrated I/O (IIO) Configuration Registers


Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x188				
Bit	Attr	Default	Description		
27:27	RWS	0x0	system_interrupt_only_on_link_bw_management_status: This bit, when set, will disable generating MSI and Intx interrupts on link bandwidth (speed and/or width) and management changes, even if MSI or INTx is enabled i.e. will disable generating MSI or INTx when LNKSTS bits 15 and 14 are set. Whether or not this condition results in a system event like Intel SMI/PMI/CPEI is dependent on whether this event masked or not in the XPCORERRMSK register.		
24:24	RW	0x0	peer2peer_memory_read_disable: When set, peer-to-peer memory reads are master aborted otherwise they are allowed to progress per the peer-to-peer decoding rules.		
23:23	RW	0x0	phold_disable: Applies only to Dev#0When set, the IIO responds with Unsupported request on receiving assert_phold message from PCH and results in generating a fatal error.		
22:22	RWS	0x0	check_cpl_tc:		
21:21	RW_O	0x0	zero_ob_tc: Forces the TC field to zero for outbound requests. 1: TC is forced to zero on all outbound transactions regardless of the source TC value 0: TC is not altered Note: In DMI mode, TC is always forced to zero and this bit has no effect.		
20:20	RW	0x1	maltp_32baddr64bhdr_en: When set, enables reporting a Malformed packet when the TLP is a 32 bit address in a 4DW header. PCI Express forbids using 4DW header sizes when the address is less than 4 GB, but some cards may use the 4DW header anyway. In these cases, the upper 32 bits of address are all 0.		
18:18	RWS	0x0	max_read_completion_combine_size: When set, all completions are returned without combining. Completions are naturally broken on cacheline boundaries, so all completions will be 64B or less.		
17:17	RO	0x0	force_data_perr: Force Data Parity Error.		
16:16	RO	0x0	force_ep_biterr: Force EP Bit Error (Poison Bit).		
15:15	RWS	0x0	dis_hdr_storage:		
14:14	RWS	0x0	allow_one_np_os:		
13:13	RWS	0x0	tlp_on_any_lane:		
12:12	RWS	0x1	disable_ob_parity_check:		
11:11	RWS	0x1	allow_1nonvc1_after_10vc1s: Allow a non-VC1 request from DMI to go after every ten VC1 request (to prevent starvation of non-VC1). Only available for Device 0 Function 0.		
9:9	RWS	0x0	dispdspolling: Disables gen2 if timeout happens in polling.cfg.		
8:7	RW	0x0	pme2acktoctrl:		



Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0x18c			
Bit	Attr	Default	Description
19:19	RW	0x1	vcm_arb_in_vc1: Only available for Device 0 Function 0.
18:18	RW	0x0	no_vcm_throttle_in_quiesce: Only available for Device 0 Function 0
17:17	RW1CS	0x0	locked_read_timed_out: Indicates that a locked read request incurred a completion time-out on PCI Express/DMI
16:16	RW1C	0x0	received_pme_to_ack: Indicates that IIO received a PME turn off ack packet or it timed out waiting for the packet
9:9	RW	0x0	override_socketid_in_cplid: For TPH/DCA requests, the Completer ID can be returned with SocketID when this bit is set.
6:6	RW	0x0	problematic_port_for_lock_flows: This bit is set by BIOS when it knows that this port is connected to a device that creates Posted-Posted dependency on its In-Out queues. This bit is set on a link if: IIO lock flows depend on the setting of this bit to treat this port in a special way during the flows. Note that if BIOS is setting up the lock flow to be in the "Intel QPI compatible" mode, then this bit must be set to 0. Note: An inbound MSI request can block the posted channel until EOIs are posted to all outbound queues enabled to receive EOI. Because of this, this bit cannot be set unless EOIFD is also set.
4:4	RWS	0x0	formfactor: Indicates what form-factor a particular root port controls 0 - CEM 1 - Express Module This bit is used to interpret bit 6 in the VPP serial stream for the port as either MRL# (CEM) input or EMLSTS# (Express Module) input.
3:3	RW	0x0	override_system_error_on_pcie_fatal_error_enable: When set, fatal errors on PCI Express (that have been successfully propagated to the primary interface of the port) are sent to the IIO core error logic (for further escalation) regardless of the setting of the equivalent bit in the ROOTCTRL register. When clear, the fatal errors are only propagated to the IIO core error logic if the equivalent bit in ROOTCTRL register is set. For Dev#0 in DMI mode and Dev#3/Fn#0, unless this bit is set, DMI link related fatal errors will never be notified to system software.
2:2	RW	0x0	override_system_error_on_pcie_non_fatal_error_enable: When set, non-fatal errors on PCI Express (that have been successfully propagated to the primary interface of the port) are sent to the IIO core error logic (for further escalation) regardless of the setting of the equivalent bit in the ROOTCTRL register. When clear, the non-fatal errors are only propagated to the IIO core error logic if the equivalent bit in ROOTCTRL register is set. For Dev#0 in DMI mode and Dev#3/Fn#0, unless this bit is set, DMI link related non-fatal errors will never be notified to system software.



Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0x18c		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
1:1	RW	0x0	override_system_error_on_pcie_correctable_error_enable: When set, correctable errors on PCI Express (that have been successfully propagated to the primary interface of the port) are sent to the IIO core error logic (for further escalation) regardless of the setting of the equivalent bit in the ROOTCTRL register. When clear, the correctable errors are only propagated to the IIO core error logic if the equivalent bit in ROOTCTRL register is set. For Dev#0 in DMI mode and Dev#3/Fn#0, unless this bit is set, DMI link related correctable errors will never be notified to system software.		
0:0	RW	0x0	acpi_pme_inten: When set, Assert/Deassert_PMEGPE messages are enabled to be generated when ACPI mode is enabled for handling PME messages from PCI Express. When this bit is cleared (from a 1), a Deassert_PMEGPE message is scheduled on behalf of the root port if an Assert_PMEGPE message was sent last from the root port.		



6.2.86 pcie_iou_bif_ctrl

PCIe Port Bifurcation Control.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0
Bus:	0	Device:	1	Function:	0
Bus:	0	Device:	2	Function:	0
Bus:	0	Device:	3	Function:	0
Offset:	0x190				
Bit	Attr	Default	Description		
3:3	WO	0x0	<p>iou_start_bifurcation:</p> <p>When software writes a 1 to this bit, IIO starts the port 0 bifurcation process. After writing to this bit, software can poll the Data Link Layer link active bit in the LNKSTS register to determine if a port is up and running. Once a port bifurcation has been initiated by writing a 1 to this bit, software cannot initiate any more write-1 to this bit (write of 0 is ok).</p> <p>Note: That this bit can be written to a 1 in the same write that changes values for bits 2:0 in this register and in that case, the new value from the write to bits 2:0 take effect. This bit always reads a 0b.</p>		
2:0	RWS RO (Device 0 Function 0)	0x4 0x0 (Device 0 Function 0)	<p>iou_bifurcation_control:</p> <p>To select a IOU bifurcation, software sets this field and then either</p> <ul style="list-style-type: none"> a) sets bit 3 in this register to initiate training OR b) resets the entire Intel® Xeon® Processor E5 v4 product family and on exit from that reset, <p>CPU will bifurcate the ports per the setting in this field.</p> <p>For Device 1 Function 0:</p> <ul style="list-style-type: none"> 000: x4x4 (operate lanes 7:4 as x4, 3:0 as x4) 001: x8 <p>For Device 2 and Device 3 Function 0:</p> <ul style="list-style-type: none"> 000: x4x4x4x4 operate lanes 15:12 as x4, 11:8 as x4, 7:4 as x4 and 3:0 as x4 001: x4x4x8 operate lanes 15:12 as x4, 11:8 as x4 and 7:0 as x8 010: x8x4x4 operate lanes 15:8 as x8, 7:4 as x4 and 3:0 as x4 011: x8x8 operate lanes 15:8 as x8, 7:0 as x8 100: x16 others: Reserved <p>For Device 0 Function 0, read only.</p>		

6.2.87 dmictrl

Type:	CFG	PortID:	N/A	Function:	0 (DMI2 Mode)
Bus:	0	Device:	0		
Offset:	0x1a0				
Bit	Attr	Default	Description		
0:0	RW	0x1	<p>Setting this bit causes IIO to abort all inbound requests on the DMI port. This will be used during specific power state and reset transitions to prevent request from PCH. This bit does not apply in PCI Express mode.</p> <p>Inbound posted requests will be dropped and inbound non-posted requests will be completed with Unsupported Request completion. Completions flowing inbound (from outbound requests) will not be dropped, but will be forwarded normally. This bit will not affect S-state auto-completion, if it is enabled.</p>		



6.2.88 dmists

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	
Offset: 0x1a8		Function: 0 (DMI2 Mode)	
Bit	Attr	Default	Description
0:0	RW1C	0x0	received_cpu_reset_done_ack:

6.2.89 ERRINJCAP

PCI Express Error Injection Capability.

Defines a vendor specific capability for WHEA error injection.

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	
Bus: 0		Device: 1	
Bus: 0		Device: 2	
Bus: 0		Device: 3	
Offset: 0x1d0		Function: 0	
		Function: 0-1	
		Function: 0-3	
		Function: 0-3	
Bit	Attr	Default	Description
31:20	RO	0x250 0x280 (Device 0 Function 0)	nxtptr: Next Capability Offset This field points to the next capability or 0 if there isn't a next capability.
19:16	RO	0x1	capver: Capability Version Set to 2h for this version of the PCI Express specification
15:0	RO	0xb	extcapid: PCI Express Extended Capability ID Vendor Defined Capability

6.2.90 ERRINJHDR

PCI Express Error Injection Capability Header.

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	
Bus: 0		Device: 1	
Bus: 0		Device: 2	
Bus: 0		Device: 3	
Offset: 0x1d4		Function: 0	
		Function: 0-1	
		Function: 0-3	
		Function: 0-3	
Bit	Attr	Default	Description
31:20	RO	0xa	vseclen: Vendor Specific Capability Length Indicates the length of the capability structure, including header bytes.
19:16	RO	0x1	vsecrev: Vendor Specific Capability Revision Set to 1h for this version of the WHEA Error Injection logic.
15:0	RO	0x3	vsecid: Vendor Specific ID Assigned for WHEA Error Injection



6.2.91 ERRINJCON

PCI Express Error Injection Control Register.

Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0x1d8		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
2:2	RW	0x0	cause_ctoerr: Cause a Completion Timeout Error When this bit is written to transition from 0 to 1, one and only one error assertion pulse is produced on the error source signal for the given port. This error will appear equivalent to an actual error assertion because this event is OR'd into the existing error reporting structure. To log another error, this bit must be cleared first, before setting again. Leaving this bit in a 1 state does not produce a persistent error condition. Notes: This bit is used for an uncorrectable error test This bit must be cleared by software before creating another event. This bit is disabled by bit 0 of this register		
1:1	RW	0x0	cause_rcvrr: Cause a Receiver Error When this bit is written to transition from 0 to 1, one and only one error assertion pulse is produced on the error source signal for the given port. This error will appear equivalent to an actual error assertion because this event is OR'd into the existing error reporting structure. To log another error, this bit must be cleared first, before setting again. Leaving this bit in a 1 state does not produce a persistent error condition. Notes: This bit is used for an correctable error test This bit must be cleared by software before creating another event. This bit is disabled by bit 0 of this register		
0:0	RW_0	0x0	errinjdis: Error Injection Disable This bit disables the use of the PCIe error injection bits.		

6.2.92 ctoctrl

Completion Timeout Control.

Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0x1e0		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
9:8	RW	0x0	xp_to_pcie_timeout_select: When OS selects a timeout range of 17s to 64s for XP (that affect NP tx issued to the PCIe/DMI) using the root port's DEVCTRL2 register, this field selects the sub-range within that larger range, for additional controllability. 00 : 17s-30s 01 : 31s-45s 10 : 46s-64s 11 : Reserved		



6.2.93 xpcorerrsts

XP Correctable Error Status

The architecture model for error logging and escalation of internal errors is similar to that of PCI Express AER, except that these internal errors never trigger an MSI and are always reported to the system software. Mask bits mask the reporting of an error and severity bit controls escalation to either fatal or non-fatal error to the internal core error logic. Note that internal errors detected in the PCI Express cluster are not dependent on any other control bits for error escalation other than the mask bit defined in these registers. All these registers are sticky.

Type: CFG		PortID: N/A	
Bus: 0		Device: 0	Function: 0
Bus: 0		Device: 1	Function: 0-1
Bus: 0		Device: 2	Function: 0-3
Bus: 0		Device: 3	Function: 0-3
Offset: 0x200			

Bit	Attr	Default	Description
1:1	RW1CS	0x0	msgd_gt_16dw: This bit is set if the root port receives a message with greater than 16 dwords (64 bytes of data).
0:0	RW1CS	0x0	pci_link_bandwidth_changed_status: This bit is set when the logical OR of LNKSTS[15] and LNKSTS[14] goes from 0 to 1.

6.2.94 xpcorerrmsk

XP Correctable Error Mask.



Type:	CFG	PortID:	N/A	
Bus:	0	Device:	0	Function: 0
Bus:	0	Device:	1	Function: 0-1
Bus:	0	Device:	2	Function: 0-3
Bus:	0	Device:	3	Function: 0-3
Offset:	0x204			

Bit	Attr	Default	Description
0:0	RWS	0x0	pci_link_bandwidth_changed_mask: Masks the BW change event from being propagated to the IIO core error logic as a correctable error

6.2.95 xpuncerrsts

XP Uncorrectable Error Status.

Type:	CFG	PortID:	N/A	
Bus:	0	Device:	0	Function: 0
Bus:	0	Device:	1	Function: 0-1
Bus:	0	Device:	2	Function: 0-3
Bus:	0	Device:	3	Function: 0-3
Offset:	0x208			

Bit	Attr	Default	Description
9:9	RW1CS	0x0	outbound_poisoned_data: Set when outbound poisoned data (from Intel QPI or peer, write or read completion) is received by this port
8:8	RW1CS	0x0	received_msi_writes_greater_than_a_dword_data:
6:6	RW1CS	0x0	received_pcie_completion_with_ur_status:
5:5	RW1CS	0x0	received_pcie_completion_with_ca_status:
4:4	RW1CS	0x0	sent_completion_with_unsupported_request:
3:3	RW1CS	0x0	sent_completion_with_completer_abort:
1:1	RW1CS	0x0	outbound_switch_fifo_data_parity_error_detected:

6.2.96 xpuncerrmsk

XP Uncorrectable Error Mask.

Type:	CFG	PortID:	N/A	
Bus:	0	Device:	0	Function: 0
Bus:	0	Device:	1	Function: 0-1
Bus:	0	Device:	2	Function: 0-3
Bus:	0	Device:	3	Function: 0-3
Offset:	0x20c			

Bit	Attr	Default	Description
9:9	RWS	0x0	outbound_poisoned_data_mask: Masks signaling of stop and scream condition to the core error logic.
8:8	RWS	0x0	received_msi_writes_greater_than_a_dword_data_mask:
6:6	RWS	0x0	received_pcie_completion_with_ur_status_mask:
5:5	RWS	0x0	received_pcie_completion_with_ca_status_mask:
4:4	RWS	0x0	sent_completion_with_unsupported_request_mask:



This register masks PCIe link related uncorrectable errors from causing the associated AER status bit to be set.

Type:	CFG	PortID:	N/A		
Bus:	0	Device:	0	Function:	0
Bus:	0	Device:	1	Function:	0-1
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x218				
Bit	Attr	Default	Description		
21:21	RWS	0x0	acs_violation_detect_mask:		
20:20	RWS	0x0	received_an_unsupported_request_detect_mask:		
19:19	RWS	0x0	ecrc_error_detect_mask:		
18:18	RWS	0x0	malformed_tlp_detect_mask:		
17:17	RWS	0x0	receiver_buffer_overflow_detect_mask:		
16:16	RWS	0x0	unexpected_completion_detect_mask:		
15:15	RWS	0x0	completer_abort_detect_mask:		
14:14	RWS	0x0	completion_time_out_detect_mask:		
13:13	RWS	0x0	flow_control_protocol_error_detect_mask:		
12:12	RWS	0x0	poisoned_tlp_detect_mask:		
5:5	RWS	0x0	surprise_down_error_detect_mask:		
4:4	RWS	0x0	data_link_layer_protocol_error_detect_mask:		

6.2.100 coredmask

Correctable Error Detect Status Mask

This register masks PCIe link related correctable errors from causing the associated status bit in AER status register to be set.

Type:	CFG	PortID:	N/A		
Bus:	0	Device:	0	Function:	0
Bus:	0	Device:	1	Function:	0-1
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x21c				
Bit	Attr	Default	Description		
13:13	RWS	0x0	advisory_non_fatal_error_detect_mask:		
12:12	RWS	0x0	replay_timer_time_out_detect_mask:		
8:8	RWS	0x0	replay_num_rollover_detect_mask:		
7:7	RWS	0x0	bad_dll_detect_mask:		
6:6	RWS	0x0	bad_tlp_detect_mask:		
0:0	RWS	0x0	receiver_error_detect_mask:		

6.2.101 rpedmask

Root Port Error Detect Status Mask



This register masks the associated error messages (received from PCIe link and NOT the virtual ones generated internally), from causing the associated status bits in AER to be set.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x220				

Bit	Attr	Default	Description
2:2	RWS	0x0	fatal_error_detected_status_mask:
1:1	RWS	0x0	non_fatal_error_detected_status_mask:
0:0	RWS	0x0	correctable_error_detected_status_mask:

6.2.102 xpuncedmask

XP Uncorrectable Error Detect Mask

This register masks other uncorrectable errors from causing the associated XPUNCERRSTS status bit to be set.


Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x224				

Bit	Attr	Default	Description
9:9	RWS	0x0	outbound_poisoned_data_detect_mask:
8:8	RWS	0x0	received_msi_writes_greater_than_a_dword_data_detect_mask:
6:6	RWS	0x0	received_pcie_completion_with_ur_detect_mask:
5:5	RWS	0x0	received_pcie_completion_with_ca_detect_mask:
4:4	RWS	0x0	sent_completion_with_unsupported_request_detect_mask:
3:3	RWS	0x0	sent_completion_with_completer_abort_detect_mask:
1:1	RWS	0x0	outbound_switch_fifo_data_parity_error_detect_mask:

6.2.103 xpcoredmask

XP Correctable Error Detect Mask

This register masks other correctable errors from causing the associated XPCORERRSTS status bit to be set.



Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x228				

Bit	Attr	Default	Description
0:0	RWS	0x0	pci_link_bandwidth_changed_detect_mask:

6.2.104 xpglberrsts

XP Global Error Status

This register captures a concise summary of the error logging in AER registers so that sideband system management software can view the errors independent of the main OS that might be controlling the AER errors.


Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x230				

Bit	Attr	Default	Description
2:2	RW1CS	0x0	<p>pci_e_aer_correctable_error:</p> <p>A PCIe correctable error (ERR_COR message received from externally or through a virtual ERR_COR message generated internally) was detected anew. Note that if that error was masked in the PCIe AER, it is not reported in this field. Software clears this bit by writing a 1 and at that stage, only 'subsequent' PCIe unmasked correctable errors will set this bit. Conceptually, per the flow of PCI Express Base Spec 2.0 defined Error message control, this bit is set by the ERR_COR message that is enabled to cause a System Error notification.</p>
1:1	RW1CS	0x0	<p>pci_e_aer_non_fatal_error:</p> <p>A PCIe non-fatal error (ERR_NONFATAL message received from externally or through a virtual ERR_NONFATAL message generated internally) was detected anew. Note that if that error was masked in the PCIe AER, it is not reported in this field. Software clears this bit by writing a 1 and at that stage only 'subsequent' PCIe unmasked non-fatal errors will set this bit again.</p>
0:0	RW1CS	0x0	<p>pci_e_aer_fatal_error:</p> <p>A PCIe fatal error (ERR_FATAL message received from externally or through a virtual ERR_FATAL message generated internally) was detected anew. Note that if that error was masked in the PCIe AER, it is not reported in this field. Software clears this bit by writing a 1 and at that stage, only 'subsequent' PCIe unmasked fatal errors will set this bit.</p>

6.2.105 xpglberrptr

XP Global Error Pointer

Check that the perfmon registers are per "cluster".



Type: CFG Bus: 0 Bus: 0 Bus: 0 Bus: 0 Offset: 0x232		PortID: N/A Device: 0 Device: 1 Device: 2 Device: 3		Function: 0 Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
2:0	ROS_V	0x0	xp_cluster_global_first_error_pointer: This field points to which of the 3 errors indicated in the XPGLBERRSTS register happened first. This field is only valid when the corresponding status bit is set and this field is rearmed to load again when the status bit indicated to by this pointer is cleared by software from 1 to 0. Value of 0x0 corresponds to bit 0 in XPGLBERRSTS register, value of 0x1 corresponds to bit 1, and so forth.		

6.2.106 pxp2cap

Secondary PCI Express Extended Capability Header.

Type: CFG Bus: 0 Bus: 0 Bus: 0 Offset: 0x250		PortID: N/A Device: 1 Device: 2 Device: 3		Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
31:20	RO	0x280	nxtptr: Next Capability Offset. This field contains the offset to the next PCI Express Extended Capability structure or 000h if no other items exist in the linked list of capabilities.		
19:16	RW_O	0x1	version: This field is a PCI-SIG defined version number that indicates the version of the Capability structure present.		
15:0	RW_O	0x19	id: This field is a PCI SIG defined ID number that indicates the nature and format of the Extended Capability. PCI Express Extended Capability ID for the Secondary PCI Express Extended Capability is 0019h.		

6.2.107 Inkcon3

Link Control 3 Register.

Type: CFG Bus: 0 Bus: 0 Bus: 0 Offset: 0x254		PortID: N/A Device: 1 Device: 2 Device: 3		Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
1:1	RW	0x0	Inkeqreqinten: Link Equalization Request Interrupt Enable. When Set, this bit enables the generation of interrupt to indicate that the Link Equalization Request bit has been set.		



Type:	CFG	PortID:	N/A	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x254				

Bit	Attr	Default	Description
0:0	RW	0x0	perfeq: Performance Equalization. When this register is 1b and a 1b is written to the 'Link Retrain' register with 'Target Link Speed' set to 8GTs, the Upstream component must perform Transmitter Equalization.

6.2.108 Inerrsts

Lane Error Status Register

Type:	CFG	PortID:	N/A	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x258				

Bit	Attr	Default	Description
15:0	RW1CS	0x0	lane: A value of 1b in any bit indicates if the corresponding PCIe Express Lane detected lane based error. bit 0 Lane 0 Error Detected bit 1 Lane 1 Error Detected bit 2 Lane 2 Error Detected bit 3 Lane 3 Error Detected bit 4 Lane 4 Error Detected (not used when the link is bifurcated as x4) bit 5 Lane 5 Error Detected (not used when the link is bifurcated as x4) bit 6 Lane 6 Error Detected (not used when the link is bifurcated as x4) bit 7 Lane 7 Error Detected (not used when the link is bifurcated as x4) bit 8 Lane 8 Error Detected (not used when the link is bifurcated as x4 or x8) bit 9 Lane 9 Error Detected (not used when the link is bifurcated as x4 or x8) bit 10 Lane 10 Error Detected (not used when the link is bifurcated as x4 or x8) bit 11 Lane 11 Error Detected (not used when the link is bifurcated as x4 or x8) bit 12 Lane 12 Error Detected (not used when the link is bifurcated as x4 or x8) bit 13 Lane 13 Error Detected (not used when the link is bifurcated as x4 or x8) bit 14 Lane 14 Error Detected (not used when the link is bifurcated as x4 or x8) bit 15 Lane 15 Error Detected (not used when the link is bifurcated as x4 or x8)

6.2.109 In[0:3]eq

Lane 0 through Lane 3 Equalization Control



Type: CFG Bus: 0 Bus: 0 Bus: 0 Offset: 0x25c, 0x25e, 0x260, 0x262		PortID: N/A Device: 1 Device: 2 Device: 3		Function: 0-1 Function: 0-3 Function: 0-3	
Bit	Attr	Default	Description		
14:12	RW_O	0x7	dnrxpreset: Downstream Component Receiver Preset Hint Receiver Preset Hint for Downstream Component with the following encoding. The Upstream component must pass on this value in the EQ TS2'es. 000b: -6 dB 001b: -7 dB 010b: -8 dB 011b: -9 dB 100b: -10 dB 101b: -11 dB 110b: -12 dB 111b: Reserved For a Downstream Component, this field reflects the latest Receiver Preset value requested from the Upstream Component on Lane 0. The default value is 111b.		
11:8	RW_O	0x8	dntxpreset: Downstream Component Transmitter Preset Transmitter Preset for Downstream Component with the following encoding. The Upstream component must pass on this value in the EQ TS2'es. 000b: -6 dB for de-emphasis, 0 dB for preshoot 001b: -3.5 dB for de-emphasis, 0 dB for preshoot 010b: -6 dB for de-emphasis, -3.5 dB for preshoot 011b: -3.5 dB for de-emphasis, -3.5 dB for preshoot 100b: -0 dB for de-emphasis, 0 dB for preshoot 101b: -0 dB for de-emphasis, -3.5 dB for preshoot Others: reserved For a Downstream Component, this field reflects the latest Transmitter Preset requested from the Upstream Component on Lane 0. The default value is 111b.		
6:4	RO	0x7	uprxpreset: Upstream Component Receiver Preset Hint Receiver Preset Hint for Upstream Component. The upstream component uses this hint for receiver equalization. The Root Ports are upstream components. The encodings are defined below. 000b: -6 dB 001b: -7 dB 010b: -8 dB 011b: -9 dB 100b: -10 dB 101b: -11 dB 110b: -12 dB 111b: reserved		



Type:	CFG	PortID:	N/A	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x25c, 0x25e, 0x260, 0x262				

Bit	Attr	Default	Description
3:0	RW_O	0x8	uptxpreset: Upstream Component Transmitter Preset Transmitter Preset for an Upstream Component. The Root Ports are upstream components. The encodings are defined below. 000b: -6 dB for de-emphasis, 0 dB for preshoot 001b: -3.5 dB for de-emphasis, 0 dB for preshoot 010b: -6 dB for de-emphasis, -3.5 dB for preshoot 011b: -3.5 dB for de-emphasis, -3.5 dB for preshoot 100b: -0 dB for de-emphasis, 0 dB for preshoot 101b: -0 dB for de-emphasis, -3.5 dB for preshoot others: reserved

6.2.110 In[4:7]eq

Lane 4 through Lane 7 Equalization Control

This register is unused when the link is configured at x4 in the bifurcation register.

Type:	CFG	PortID:	N/A	Function:	0-1
Bus:	0	Device:	1	Function:	0, 2
Bus:	0	Device:	2	Function:	0, 2
Bus:	0	Device:	3	Function:	0, 2
Offset:	0x264, 0x266, 0x268, 0x26a				

Bit	Attr	Default	Description
14:12	RW_O	0x7	dnrxpreset: Downstream Component Receiver Preset Hint Receiver Preset Hint for Downstream Component with the following encoding. The Upstream component must pass on this value in the EQ TS2'es. 000b: -6 dB 001b: -7 dB 010b: -8 dB 011b: -9 dB 100b: -10 dB 101b: -11 dB 110b: -12 dB 111b: Reserved For a Downstream Component, this field reflects the latest Receiver Preset value requested from the Upstream Component on Lane 0. The default value is 111b.




Type: CFG		PortID: N/A	
Bus: 0		Device: 1	
Bus: 0		Device: 2	
Bus: 0		Device: 3	
Offset: 0x264, 0x266, 0x268, 0x26a		Function: 0-1 0, 2 0, 2	

Bit	Attr	Default	Description
11:8	RW_O	0x8	<p>dntxpreset:</p> <p>Downstream Component Transmitter Preset Transmitter Preset for Downstream Component with the following encoding. The Upstream component must pass on this value in the EQ TS2'es.</p> <p>000b: -6 dB for de-emphasis, 0 dB for preshoot 001b: -3.5 dB for de-emphasis, 0 dB for preshoot 010b: -6 dB for de-emphasis, -3.5 dB for preshoot 011b: -3.5 dB for de-emphasis, -3.5 dB for preshoot 100b: -0 dB for de-emphasis, 0 dB for preshoot 101b: -0 dB for de-emphasis, -3.5 dB for preshoot others: reserved</p> <p>For a Downstream Component, this field reflects the latest Transmitter Preset requested from the Upstream Component on Lane 0. The default value is 111b.</p>
6:4	RO	0x7	<p>uprxpreset:</p> <p>Upstream Component Receiver Preset Hint Receiver Preset Hint for Upstream Component. The upstream component uses this hint for receiver equalization. The Root Ports are upstream components. The encodings are defined below.</p> <p>000b: -6 dB 001b: -7 dB 010b: -8 dB 011b: -9 dB 100b: -10 dB 101b: -11 dB 110b: -12 dB 111b: reserved</p>
3:0	RW_O	0x8	<p>uptxpreset:</p> <p>Upstream Component Transmitter Preset Transmitter Preset for an Upstream Component. The Root Ports are upstream components. The encodings are defined below.</p> <p>000b: -6 dB for de-emphasis, 0 dB for preshoot 001b: -3.5 dB for de-emphasis, 0 dB for preshoot 010b: -6 dB for de-emphasis, -3.5 dB for preshoot 011b: -3.5 dB for de-emphasis, -3.5 dB for preshoot 100b: -0 dB for de-emphasis, 0 dB for preshoot 101b: -0 dB for de-emphasis, -3.5 dB for preshoot others: reserved</p>

6.2.111 In[8:15]eq

Lane 8 though Lane 15 Equalization Control

This register is unused when the link is configured at x4 or x8 in the bifurcation register.



Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x30C				

Bit	Attr	Default	Description
15:15	RO	0x0	mc_ecrc_regen_sup:
13:8	RO	0x0	mc_window_size_req:
5:0	RO	0xF	mc_max_group:

6.2.115 mcast_ctrl

Dualcast Control Register

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x30E				

Bit	Attr	Default	Description
15:15	RW-L	0x0	mc_enable:
5:0	RW-L	0x0	mc_num_group — This field only supports and implements bits [3:0]. Bits [5:4] are not supported.

6.2.116 mcast_base

Dualcast Base Address Register

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	0	Function:	0-1
Bus:	0	Device:	1	Function:	0-3
Bus:	0	Device:	2	Function:	0-3
Bus:	0	Device:	3	Function:	0-3
Offset:	0x310				

Bit	Attr	Default	Description
63:12	RW-L	0x0	mc_base_address:
5:0	RW-L	0x0	mc_index_position:

6.2.117 mcast_rcv

Dualcast Receive Register



Type: CFG PortID: N/A Bus: 0 Device: 0 Function: 0 Bus: 0 Device: 1 Function: 0-1 Bus: 0 Device: 2 Function: 0-3 Bus: 0 Device: 3 Function: 0-3 Offset: 0x318			
Bit	Attr	Default	Description
15:0	RW-L	0x0	mc_rcv:

6.2.118 mcast_blk_all

Dualcast Block All Register

Type: CFG PortID: N/A Bus: 0 Device: 0 Function: 0 Bus: 0 Device: 1 Function: 0-1 Bus: 0 Device: 2 Function: 0-3 Bus: 0 Device: 3 Function: 0-3 Offset: 0x320			
Bit	Attr	Default	Description
15:0	RW-L	0x0	mc_blk_all:

6.2.119 mcast_blk_unt

Dualcast Block Untranslated Register

Type: CFG PortID: N/A Bus: 0 Device: 0 Function: 0 Bus: 0 Device: 1 Function: 0-1 Bus: 0 Device: 2 Function: 0-3 Bus: 0 Device: 3 Function: 0-3 Offset: 0x328			
Bit	Attr	Default	Description
15:0	RW-L	0x0	mc_blk_unt:

6.2.120 mcast_overlay_bar

Dualcast Overlay Bar Register

Type: CFG PortID: N/A Bus: 0 Device: 0 Function: 0 Bus: 0 Device: 1 Function: 0-1 Bus: 0 Device: 2 Function: 0-3 Bus: 0 Device: 3 Function: 0-3 Offset: 0x330			
Bit	Attr	Default	Description
63:6	RW-L	0x0	mc_overlay_addr:
5:0	RW-L	0x0	mc_overlay_size:



6.3 Device 0 Function 0 Region DMIRCBAR

DMI Root Complex Registers Block (RCRB). This block is mapped into memory space, using register DMIRCBAR [Device 0:Function 0, offset 0x50].

Register Name	Offset	Size
dmivc0rcap	0x10	32
dmivc0rctl	0x14	32
dmivc0rst	0x1a	16
dmivc1rcap	0x1c	32
dmivc1rctl	0x20	32
dmivc1rst	0x26	16
dmivcprcap	0x28	32
dmivcprctl	0x2c	32
dmivcprsts	0x32	16
dmivcmrcap	0x34	32
dmivcmrctl	0x38	32
dmivcmrst	0x3e	16
dmivc1cdtthrottle	0x60	32
dmivpcdtthrottle	0x64	32
dmivcmdtthrottle	0x68	32

6.3.1 dmivc0rcap

DMI VC0 Resource Capability

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 0	
Offset: 0x10		Function: 0	
Bit	Attr	Default	Description
31:16	RO	0x0	maxtimeslots: Max Time Slots
15:15	RO	0x0	rejsnpt: Reject Snoop Transactions 0: Transactions with or without the No Snoop bit set within the TLP header are allowed on this VC. 1: Any transaction without the No Snoop bit set within the TLP header will be rejected as an Unsupported Request.

6.3.2 dmivc0rctl

DMI VC0 Resource Control

Controls the resources associated with PCI Express Virtual Channel 0.




Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 0	
Offset: 0x14		Function: 0	
Bit	Attr	Default	Description
31:31	RO	0x1	vc0e: Virtual Channel 0 Enable For VC0 this is hardwired to 1 and read only as VC0 can never be disabled.
26:24	RO	0x0	vc0id: Virtual Channel 0 ID Assigns a VC ID to the VC resource. For VC0 this is hardwired to 0 and read only.
7:7	RO	0x0	tc7vc0m: Traffic Class 7/ Virtual Channel 0 Map Traffic Class 7 is always routed to VCm.
6:1	RW-LB	0x3f	tcvc0m: Traffic Class / Virtual Channel 0 Map Indicates the TCs (Traffic Classes) that are mapped to the VC resource. Bit locations within this field correspond to TC values. For example, when bit 6 is set in this field, TC6 is mapped to this VC resource. When more than one bit in this field is set, it indicates that multiple TCs are mapped to the VC resource. In order to remove one or more TCs from the TC/VC Map of an enabled VC, software must ensure that no new or outstanding transactions with the TC labels are targeted at the given Link.
0:0	RO	0x1	tc0vc0m: Traffic Class 0 / Virtual Channel 0 Map Traffic Class 0 is always routed to VC0.

6.3.3 dmivc0rst

DMI VC0 Resource Status.

Reports the Virtual Channel specific status.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 0	
Offset: 0x1a		Function: 0	
Bit	Attr	Default	Description
1:1	RO-V	0x1	vc0np: Virtual Channel 0 Negotiation Pending 0: The VC negotiation is complete. 1: The VC resource is still in the process of negotiation (initialization or disabling). This bit indicates the status of the process of Flow Control initialization. It is set by default on Reset, as well as whenever the corresponding Virtual Channel is Disabled or the Link is in the DL_Down state. It is cleared when the link successfully exits the FC_INIT2 state. BIOS Requirement: Before using a Virtual Channel, software must check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.



6.3.4 dmivc1rcap

DMI VC1 Resource Capability


Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 0	
Offset: 0x1c		Function: 0	
Bit	Attr	Default	Description
15:15	RO	0x1	rejsnpt: Reject Snoop Transactions 0: Transactions with or without the No Snoop bit set within the TLP header are allowed on this VC. 1: Any transaction without the No Snoop bit set within the TLP header will be rejected as an Unsupported Request.

6.3.5 dmivc1rctl

DMI VC1 Resource Control

Controls the resources associated with PCI Express Virtual Channel 1.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 0	
Offset: 0x20		Function: 0	
Bit	Attr	Default	Description
31:31	RW-LB	0x0	vc1e: Virtual Channel 1 Enable 0: Virtual Channel is disabled. 1: Virtual Channel is enabled. See exceptions below. Software must use the VC Negotiation Pending bit to check whether the VC negotiation is complete. When VC Negotiation Pending bit is cleared, a 1 read from this VC Enable bit indicates that the VC is enabled (Flow Control Initialization is completed for the PCI Express port). A 0 read from this bit indicates that the Virtual Channel is currently disabled. BIOS Requirement: 1. To enable a Virtual Channel, the VC Enable bits for that Virtual Channel must be set in both Components on a Link. 2. To disable a Virtual Channel, the VC Enable bits for that Virtual Channel must be cleared in both Components on a Link. 3. Software must ensure that no traffic is using a Virtual Channel at the time it is disabled. 4. Software must fully disable a Virtual Channel in both Components on a Link before re-enabling the Virtual Channel.
26:24	RW-LB	0x1	vc1id: Virtual Channel 1 ID Assigns a VC ID to the VC resource. Assigned value must be non-zero. This field can not be modified when the VC is already enabled.
7:7	RO	0x0	tc7vc1m: Traffic Class 7/ Virtual Channel 1 Map Traffic Class 7 is always routed to VCm.



Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 0	
Offset: 0x20		Function: 0	
Bit	Attr	Default	Description
6:1	RW-LB	0x0	<p>tcvc1m: Traffic Class / Virtual Channel 1 Map</p> <p>Indicates the TCs (Traffic Classes) that are mapped to the VC resource. Bit locations within this field correspond to TC values. For example, when bit 6 is set in this field, TC6 is mapped to this VC resource. When more than one bit in this field is set, it indicates that multiple TCs are mapped to the VC resource. In order to remove one or more TCs from the TC/VC Map of an enabled VC, software must ensure that no new or outstanding transactions with the TC labels are targeted at the given Link.</p>
0:0	RO	0x0	<p>tc0vc1m: Traffic Class 0 / Virtual Channel 0 Map</p> <p>Traffic Class 0 is always routed to VC0.</p>

6.3.6 dmivc1rst

DMI VC1 Resource Status

Reports the Virtual Channel specific status.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 0	
Offset: 0x26		Function: 0	
Bit	Attr	Default	Description
1:1	RO-V	0x1	<p>vc1np: Virtual Channel 1 Negotiation Pending</p> <p>0: The VC negotiation is complete. 1: The VC resource is still in the process of negotiation (initialization or disabling).</p> <p>This bit indicates the status of the process of Flow Control initialization. It is set by default on Reset, as well as whenever the corresponding Virtual Channel is Disabled or the Link is in the DL_Down state.</p> <p>It is cleared when the link successfully exits the FC_INIT2 state.</p> <p>BIOS Requirement: Before using a Virtual Channel, software must check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.</p>



6.3.7 dmivcprcap

DMI VCP Resource Capability


Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 0	
Offset: 0x1a		Function: 0	
Bit	Attr	Default	Description
15:15	RO	0x0	<p>rejsnpt:</p> <p>Reject Snoop Transactions</p> <p>0: Transactions with or without the No Snoop bit set within the TLP header are allowed on this VC.</p> <p>1: Any transaction without the No Snoop bit set within the TLP header will be rejected as an Unsupported Request.</p>

6.3.8 dmivcprctl

DMI VCP Resource Control

Controls the resources associated with the DMI Private Channel (VCp).

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 0	
Offset: 0x1a		Function: 0	
Bit	Attr	Default	Description
31:31	RW-LB	0x0	<p>vcpe:</p> <p>Virtual Channel Private Enable</p> <p>0: Virtual Channel is disabled.</p> <p>1: Virtual Channel is enabled. See exceptions below.</p> <p>Software must use the VC Negotiation Pending bit to check whether the VC negotiation is complete. When VC Negotiation Pending bit is cleared, a 1 read from this VC Enable bit indicates that the VC is enabled (Flow Control Initialization is completed for the PCI Express port). A 0 read from this bit indicates that the Virtual Channel is currently disabled.</p> <p>BIOS Requirement:</p> <ol style="list-style-type: none"> 1. To enable a Virtual Channel, the VC Enable bits for that Virtual Channel must be set in both Components on a Link. 2. To disable a Virtual Channel, the VC Enable bits for that Virtual Channel must be cleared in both Components on a Link. 3. Software must ensure that no traffic is using a Virtual Channel at the time it is disabled. 4. Software must fully disable a Virtual Channel in both Components on a Link before re-enabling the Virtual Channel.
26:24	RW-LB	0x2	<p>vcpid:</p> <p>Virtual Channel Private ID</p> <p>Assigns a VC ID to the VC resource. This field can not be modified when the VC is already enabled. No private VCs are precluded by hardware and private VC handling is implemented the same way as non-private VC handling.</p>
7:7	RO	0x0	<p>tc7vcpm:</p> <p>Traffic Class 7/ Virtual Channel 0 Map</p> <p>Traffic Class 7 is always routed to VCm.</p>




Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 0	
Offset: 0x1a		Function: 0	
Bit	Attr	Default	Description
6:1	RW-LB	0x0	<p>tcvcpm:</p> <p>Traffic Class / Virtual Channel private Map</p> <p>Indicates the TCs (Traffic Classes) that are mapped to the VC resource. Bit locations within this field correspond to TC values. For example, when bit 6 is set in this field, TC6 is mapped to this VC resource. When more than one bit in this field is set, it indicates that multiple TCs are mapped to the VC resource. In order to remove one or more TCs from the TC/VC Map of an enabled VC, software must ensure that no new or outstanding transactions with the TC labels are targeted at the given Link.</p>
0:0	RO	0x0	<p>tc0vcpm:</p> <p>Traffic Class 0 / Virtual Channel Private Map</p> <p>Traffic Class 0 is always routed to VC0.</p>

6.3.9 dmivcprsts

DMI VCP Resource Status

Reports the Virtual Channel specific status.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 0	
Offset: 0x32		Function: 0	
Bit	Attr	Default	Description
1:1	RO-V	0x1	<p>vcpnp:</p> <p>Virtual Channel Private Negotiation Pending</p> <p>0: The VC negotiation is complete.</p> <p>1: The VC resource is still in the process of negotiation (Initialization or disabling).</p> <p>This bit indicates the status of the process of Flow Control initialization. It is set by default on Reset, as well as whenever the corresponding Virtual Channel is Disabled or the Link is in the DL_Down state.</p> <p>It is cleared when the link successfully exits the FC_INIT2 state.</p> <p>BIOS Requirement: Before using a Virtual Channel, software must check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.</p>



6.3.10 dmivcmrcap

DMI VCM Resource Capability

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 0	
Offset: 0x34		Function: 0	
Bit	Attr	Default	Description
15:15	RO	0x1	rejsnpt: Reject Snoop Transactions 0: Transactions with or without the No Snoop bit set within the TLP header are allowed on this VC. 1: Any transaction without the No Snoop bit set within the TLP header will be rejected as an Unsupported Request.

6.3.11 dmivcmrctl

DMI VCM Resource Control

Controls the resources associated with PCI Express Virtual Channel 0.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 0	
Offset: 0x38		Function: 0	
Bit	Attr	Default	Description
31:31	RW-LB	0x0	vcme: Virtual Channel M Enable 0: Virtual Channel is disabled. 1: Virtual Channel is enabled. See exceptions below. Software must use the VC Negotiation Pending bit to check whether the VC negotiation is complete. When VC Negotiation Pending bit is cleared, a 1 read from this VC Enable bit indicates that the VC is enabled (Flow Control Initialization is completed for the PCI Express port). A 0 read from this bit indicates that the Virtual Channel is currently disabled. BIOS Requirement: 1. To enable a Virtual Channel, the VC Enable bits for that Virtual Channel must be set in both Components on a Link. 2. To disable a Virtual Channel, the VC Enable bits for that Virtual Channel must be cleared in both Components on a Link. 3. Software must ensure that no traffic is using a Virtual Channel at the time it is disabled. 4. Software must fully disable a Virtual Channel in both Components on a Link before re-enabling the Virtual Channel.
26:24	RW-LB	0x0	vcmid: VCm ID
7:7	RO	0x1	tc7vcpm: Traffic Class 7 / Virtual Channel 0 Map Traffic Class 7 is always routed to VCm.
6:1	RO	0x0	tcvcmm: Traffic Class / Virtual Channel M Map No other traffic class is mapped to VCM
0:0	RO	0x0	tc0vcmm: Traffic Class 0 Virtual Channel Map



6.3.12 dmivimrst

DMI VCM Resource Status

Reports the Virtual Channel specific status.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 0	
Offset: 0x3e		Function: 0	
Bit	Attr	Default	Description
1:1	RO-V	1b	<p>vcmnp: Virtual Channel M Negotiation Pending 0: The VC negotiation is complete. 1: The VC resource is still in the process of negotiation (initialization or disabling).</p> <p>This bit indicates the status of the process of Flow Control initialization. It is set by default on Reset, as well as whenever the corresponding Virtual Channel is Disabled or the Link is in the DL_Down state. It is cleared when the link successfully exits the FC_INIT2 state.</p> <p>BIOS Requirement: Before using a Virtual Channel, software must check whether the VC Negotiation Pending fields for that Virtual Channel are cleared in both Components on a Link.</p>

6.3.13 dmivc1cdtthrottle

DMI VC1 Credit Throttle

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 0	
Offset: 0x60		Function: 0	
Bit	Attr	Default	Description
31:24	RWS	0x0	<p>prd: Posted Request Data VC1 Credit Withhold Number of VC1 Posted Data credits to withhold from being reported or used.</p>
21:16	RWS	0x0	<p>prh: Posted Request Header VC1 Credit Withhold Number of VC1 Posted Request credits to withhold from being reported or used.</p>
15:8	RWS	0x0	<p>nprd: Non-Posted Request Data VC1 Credit Withhold Number of VC1 Non-Posted Data credits to withhold from being reported or used.</p>
5:0	RWS	0x0	<p>nprh: Non-Posted Request Header VC1 Credit Withhold Number of VC1 Non-Posted Request credits to withhold from being reported or used.</p>



6.3.14 dmivcpdtthrottle

DMI VCp Credit Throttle

Type:	MEM	PortID:	8'h7e	Function:	0
Bus:	0	Device:	0		
Offset:	0x64				
Bit	Attr	Default	Description		
31:24	RWS	0x0	prd: Posted Request Data VCp Credit Withhold Number of VCp Posted Data credits to withhold from being reported or used.		
21:16	RWS	0x0	prh: Posted Request Header VCp Credit Withhold Number of VCp Posted Request credits to withhold from being reported or used.		
15:8	RWS	0x0	nprd: Non-Posted Request Data VCp Credit Withhold Number of VCp Non-Posted Data credits to withhold from being reported or used.		
5:0	RWS	0x0	nprh: Non-Posted Request Header VCp Credit Withhold Number of VCp Non-Posted Request credits to withhold from being reported or used.		

6.3.15 dmivcmcdtthrottle

DMI VCm Credit Throttle

Type:	MEM	PortID:	8'h7e	Function:	0
Bus:	0	Device:	0		
Offset:	0x68				
Bit	Attr	Default	Description		
31:24	RWS	0x0	prd: Posted Request Data VCm Credit Withhold Number of VCm Posted Data credits to withhold from being reported or used.		
21:16	RWS	0x0	prh: Posted Request Header VCm Credit Withhold Number of VCm Posted Request credits to withhold from being reported or used.		
15:8	RWS	0x0	nprd: Non-Posted Request Data VCm Credit Withhold Number of VCm Non-Posted Data credits to withhold from being reported or used.		
5:0	RWS	0x0	nprh: Non-Posted Request Header VCm Credit Withhold Number of VCm Non-Posted Request credits to withhold from being reported or used.		



6.4 Device 4 Function 0-7

Intel® QuickData Technology DMA Registers.

Register Name	Offset	Size	Function
vid	0x0	16	0-7
did	0x2	16	0-7
pcicmd	0x4	16	0-7
pcists	0x6	16	0-7
rid	0x8	8	0-7
ccr	0x9	24	0-7
clsr	0xc	8	0-7
hdr	0xe	8	0-7
cb_bar	0x10	64	0-7
svid	0x2c	16	0-7
sdid	0x2e	16	0-7
capptr	0x34	8	0-7
intl	0x3c	8	0-7
intpin	0x3d	8	0-7
devcfg	0x60	16	0
msixcapid	0x80	8	0-7
msixnxtptr	0x81	8	0-7
msixmsgctl	0x82	16	0-7
tableoff_bir	0x84	32	0-7
pbaoff_bir	0x88	32	0-7
capid	0x90	8	0-7
nextptr	0x91	8	0-7
expcap	0x92	16	0-7
devcap	0x94	32	0-7
devcon	0x98	16	0-7
devsts	0x9a	16	0-7
devcap2	0xb4	32	0-7
devcon2	0xb8	16	0-7
pmcap	0xe0	32	0-7
pmcsr	0xe4	32	0-7
dmauncerrsts	0x148	32	0
dmauncerrmsk	0x14c	32	0
dmauncerrsev	0x150	32	0
dmauncerrptr	0x154	8	0
dmaglberrptr	0x160	8	0
chanerr_int	0x180	32	0-7
chanerrmsk_int	0x184	32	0-7
chanerrsev_int	0x188	32	0-7
chanerrptr	0x18c	8	0-7



6.4.1 vid


Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x0		Function: 0-7	
Bit	Attr	Default	Description
15:0	RO	0x8086	vendor_identification_number: The value is assigned by PCI-SIG to Intel.

6.4.2 did

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x2		Function: 0-7	
Bit	Attr	Default	Description
15:0	RO	0x2f20 (Function 0) 0x2f21 (Function 1) 0x2f22 (Function 2) 0x2f23 (Function 3) 0x2f24 (Function 4) 0x2f25 (Function 5) 0x2f26 (Function 6) 0x2f27 (Function 7)	device_identification_number: Device ID values vary from function to function.

6.4.3 pcicmd

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x4		Function: 0-7	
Bit	Attr	Default	Description
10:10	RW	0x0	intx_interrupt_disable:
9:9	RO	0x0	fast_back_to_back_enable: Not applicable to PCI Express and is hardwired to 0
8:8	RO	0x0	serre:
7:7	RO	0x0	idsel_stepping_wait_cycle_control: Not applicable to internal devices. Hardwired to 0.
6:6	RO	0x0	perre:
5:5	RO	0x0	vga_palette_snoop_enable: Not applicable to internal devices. Hardwired to 0.
4:4	RO	0x0	mwie:
3:3	RO	0x0	sce:
2:2	RW	0x0	bme:
1:1	RW	0x0	mse:
0:0	RO	0x0	iouse:



6.4.4 pcists

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x6		Function: 0-7	
Bit	Attr	Default	Description
15:15	RW1C	0x0	dpe:
14:14	RO	0x0	sse:
13:13	RO	0x0	rma:
12:12	RO	0x0	rta:
11:11	RW1C	0x0	sta:
10:9	RO	0x0	devsel_timing: Not applicable to PCI Express. Hardwired to 0.
8:8	RW1C	0x0	mdpe:
7:7	RO	0x0	fast_back_to_back: Not applicable to PCI Express. Hardwired to 0.
5:5	RO	0x0	pci66mhz_capable: Not applicable to PCI Express. Hardwired to 0.
4:4	RO	0x1	capabilities_list: This bit indicates the presence of a capabilities list structure
3:3	RO_V	0x0	intxsts:

6.4.5 rid

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x8		Function: 0-7	
Bit	Attr	Default	Description
7:0	RO_V	0x0	revision_id: Reflects the Uncore Revision ID after reset. Reflects the Compatibility Revision ID after BIOS writes 0x69 to any RID register in any Intel® Xeon® Processor E5 v4 product family function.

6.4.6 ccr

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x9		Function: 0-7	
Bit	Attr	Default	Description
23:16	RO_V	0x8	base_class: Generic Device
15:8	RO_V	0x80	sub_class: Generic Device
7:0	RO_V	0x0	register_level_programming_interface: Set to 00h for all non-APIC devices.



6.4.7 clsr

Type:	CFG	PortID:	N/A
Bus:	0	Device:	4
Offset:	0xc	Function:	0-7
Bit	Attr	Default	Description
7:0	RW	0x0	cacheline_size: This register is set as RW for compatibility reasons only. Cacheline size is always 64B.


6.4.8 hdr

Type:	CFG	PortID:	N/A
Bus:	0	Device:	4
Offset:	0xe	Function:	0-7
Bit	Attr	Default	Description
7:7	RO	0x1	multi_function_device: This bit defaults to 1b since all these devices are multi-function
6:0	RO	0x0	configuration_layout: This field identifies the format of the configuration header layout. It is Type 0 for all these devices. The default is 00h, indicating a 'endpoint device'.

6.4.9 cb_bar

Intel QuickData Technology Base Address Register.

Type:	CFG	PortID:	N/A
Bus:	0	Device:	4
Offset:	0x10	Function:	0-7
Bit	Attr	Default	Description
63:14	RW	0x0	bar: This marks the 16 KB aligned 64-bit base address for memory-mapped registers of Intel QuickData Technology-DMA. The BAR register in the 8 functions will be referenced with a logical name of CB_BAR[0:7].
3:3	RO	0x0	prefetchable: The DMA registers are not prefetchable.
2:1	RO	0x2	type: The DMA registers is 64-bit address space and can be placed anywhere within the addressable region of the system.
0:0	RO	0x0	memory_space: This Base Address Register indicates memory space.



6.4.10 svid

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x2c		Function: 0-7	
Bit	Attr	Default	Description
15:0	RW_O	0x8086	vendor_identification_number:

6.4.11 sdid

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x2e		Function: 0-7	
Bit	Attr	Default	Description
15:0	RW_O	0x0	subsystem_identification_number:

6.4.12 capptr

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x34		Function: 0-7	
Bit	Attr	Default	Description
7:0	RO	0x80	capability_pointer: Points to the first capability structure for the device which is the PCIe capability.

6.4.13 intl

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x3c		Function: 0-7	
Bit	Attr	Default	Description
7:0	RW	0x0	interrupt_line: NA for these devices



6.4.14 intpin

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x3d		Function: 0-7	
Bit	Attr	Default	Description
7:0	RW_O	0x1 (Function 0) 0x2 (Function 1) 0x3 (Function 2) 0x4 (Function 3) 0x1 (Function 4) 0x2 (Function 5) 0x3 (Function 6) 0x4 (Function 7)	cb_intpin0: (Function 0) cb_intpin1: (Function 1) cb_intpin2: (Function 2) cb_intpin3: (Function 3) cb_intpin4: (Function 4) cb_intpin5: (Function 5) cb_intpin6: (Function 6) cb_intpin7: (Function 7)

6.4.15 devcfg

This DEVCFG is for Function 0 only

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x60		Function: 0	
Bit	Attr	Default	Description
11:11	RW_O	0x0	f1extop_diden: When set, this bit switches in the Function 1 Device ID that are typically used in storage applications. When clear, the function 1 DID remains at the default value associated with applications (for example, networking). This bit should be written by BIOS prior to enumeration.
10:10	RW_O	0x0	f0extop_diden: When set, this bit switches in the Function 0 Device ID that are typically used in storage applications. When clear, the function 0 DID remains at the default value associated with applications (e.g., networking). This bit should be written by BIOS prior to enumeration.
9:9	RWS	0x0	enable_no_snoop: This bit is akin to the NoSnoop enable bit in the PCI Express capability register, only that this bit is controlled by bios rather than OS. When set, the no snoop optimization is enabled (provided the equivalent bit in the PCI Express DEVCON register is set) on behalf of Intel QuickData Technology DMA otherwise it is not. Note: Due to severe performance degradation, it is not recommended that this bit be set except in debug mode.

6.4.16 msixcapid

MSI-X Capability ID.

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x80		Function: 0-7	
Bit	Attr	Default	Description
7:0	RO	0x11	cb_msixcapid: Assigned by PCI-SIG for MSI-X (Intel QuickData Technology DMA)



6.4.17 msixnxtptr

MSI-X Next Pointer.

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x81		Function: 0-7	
Bit	Attr	Default	Description
7:0	RO	0x90	cb_msixnxtptr: This field is set to 90h for the next capability list (PCI Express capability structure) in the chain.

6.4.18 msixmsgctl

MSI-X Message Control.

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x82		Function: 0-7	
Bit	Attr	Default	Description
15:15	RW	0x0	msi_x_enable: Software uses this bit to select between MSI-X or INTx method for signaling interrupts from the DMA 0: INTx method is chosen for DMA interrupts 1: MSI-X method is chosen for DMA interrupts
14:14	RW	0x0	function_mask: If 1, the 1 vector associated with the dma is masked, regardless of the per-vector mask bit state. If 0, the vector's mask bit determines whether the vector is masked or not. Setting or clearing the MSI-X function mask bit has no effect on the state of the per-vector Mask bit.
10:0	RO	0x0	table_size: Indicates the MSI-X table size which for IIO is 1, encoded as a value of 0h.

6.4.19 tableoff_bir

MSI-X Table Offset and BAR Indicator.

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x84		Function: 0-7	
Bit	Attr	Default	Description
31:3	RO	0x400	table_offset: MSI-X Table Structure is at offset 8K from the Intel QuickData Technology BAR address. See "MSI-X Lower Address Registers (MSGADDR)" for the start of details relating to MSI-X registers.
2:0	RO	0x0	table_bir: Intel QuickData Technology DMA BAR is at offset 10h in the DMA config space and hence this register is 0.



6.4.20 pbaoff_bir

Type:	CFG	PortID:	N/A
Bus:	0	Device:	4
Offset:	0x88	Function:	0-7
Bit	Attr	Default	Description
31:3	RO	0x600	table_offset: MSI-X PBA Structure is at offset 12K from the Intel QuickData Technology BAR address.
2:0	RO	0x0	table_bir: Intel QuickData Technology DMA BAR is at offset 10h in the DMA config space and hence this register is 0.

6.4.21 capid

The PCI Express Capability List register enumerates the PCI Express Capability structure in the PCI 3.0 configuration space

Type:	CFG	PortID:	N/A
Bus:	0	Device:	4
Offset:	0x90	Function:	0-7
Bit	Attr	Default	Description
7:0	RO	0x10	capability_id: Provides the PCI Express capability ID assigned by PCI-SIG.

6.4.22 nextptr

The PCI Express Capability List register enumerates the PCI Express Capability structure in the PCI 3.0 configuration space.

Type:	CFG	PortID:	N/A
Bus:	0	Device:	4
Offset:	0x91	Function:	0-7
Bit	Attr	Default	Description
7:0	RO	0xe0	next_ptr: This field is set to the PCI Power Management capability.



6.4.23 expcap

The PCI Express Capabilities register identifies the PCI Express device type and associated capabilities.

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x92		Function: 0-7	
Bit	Attr	Default	Description
13:9	RO	0x0	interrupt_message_number: N/A
8:8	RO	0x0	slot_implemented: N/A
7:4	RO	0x9	device_port_type: This field identifies the type of device. It is set to for the DMA to indicate root complex integrated endpoint device.
3:0	RO	0x2	capability_version: This field identifies the version of the PCI Express capability structure. Set to 2h for PCI Express and DMA devices for compliance with the extended base registers.

6.4.24 devcap

The PCI Express Device Capabilities register identifies device specific information for the device.

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x94		Function: 0-7	
Bit	Attr	Default	Description
28:28	RWS_O	0x0	flr_supported:
27:26	RO	0x0	captured_slot_power_limit_scale: Does not apply to Intel QuickData Technology DMA
25:18	RO	0x0	captured_slot_power_limit_value: Does not apply to Intel QuickData Technology DMA
15:15	RO	0x1	role_based_error_reporting: IIO is 1.1 compliant and so supports this feature
14:14	RO	0x0	power_indicator_present_on_device: Does not apply to Intel QuickData Technology DMA
13:13	RO	0x0	attention_indicator_present: Does not apply to Intel QuickData Technology DMA
12:12	RO	0x0	attention_button_present: Does not apply to Intel QuickData Technology DMA
11:9	RO	0x0	endpoint_l1_acceptable_latency: N/A
8:6	RO	0x0	endpoint_l0s_acceptable_latency: N/A
5:5	RO	0x0	extended_tag_field_supported:
4:3	RO	0x0	phantom_functions_supported: Intel QuickData Technology DMA does not support phantom functions.




Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x94		Function: 0-7	
Bit	Attr	Default	Description
2:0	RO	0x0	max_payload_size: Intel QuickData Technology DMA supports max 128B on writes to PCI Express

6.4.25 devcon

The PCI Express Device Control register controls PCI Express specific capabilities parameters associated with the device.

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x98		Function: 0-7	
Bit	Attr	Default	Description
15:15	RW	0x0	initiate_flr: Intel QuickData Technology DMA does a reset of that function only per the FLR ECN. This bit always returns 0 when read and a write of 0 has no impact
14:12	RO	0x0	max_read_request_size: N/A to Intel QuickData Technology DMA since it does not issue tx on PCIe
11:11	RW	0x1	enable_no_snoop: For Intel QuickData Technology DMA, when this bit is clear, all DMA transactions must be snooped. When set, DMA transactions to main memory can utilize No Snoop optimization under the guidance of the device driver.
10:10	RO	0x0	auxiliary_power_management_enable: Not applicable to Intel QuickData Technology DMA
9:9	RO	0x0	phantom_functions_enable: Not applicable to Intel QuickData Technology DMA since it never uses phantom functions as a requester.
8:8	RO	0x0	extended_tag_field_enable:
7:5	RO	0x0	max_payload_size: N/A for Intel QuickData Technology DMA
4:4	RW	0x0	enable_relaxed_ordering: For most parts, writes from Intel QuickData Technology DMA are relaxed ordered, except for DMA completion writes. But the fact that Intel QuickData Technology DMA writes are relaxed ordered is not very useful except when the writes are also non-snooped. If the writes are snooped, relaxed ordering does not provide any particular advantage based on IIO uArch. But when writes are non-snooped, relaxed ordering is required to get good BW and this bit is expected to be set. If this bit is clear, NS writes will get terrible performance.
3:3	RO	0x0	unsupported_request_reporting_enable: N/A for Intel QuickData Technology DMA
2:2	RO	0x0	fatal_error_reporting_enable: N/A for Intel QuickData Technology DMA
1:1	RO	0x0	non_fatal_error_reporting_enable: N/A for Intel QuickData Technology DMA
0:0	RO	0x0	correctable_error_reporting_enable: N/A for Intel QuickData Technology DMA



6.4.26 devsts

The PCI Express Device Status register provides information about PCI Express device specific parameters associated with the device.

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x9a		Function: 0-7	
Bit	Attr	Default	Description
5:5	RO	0x0	transactions_pending: 1: indicates that the Intel QuickData Technology DMA device has outstanding Non-Posted Request which it has issued either towards main memory, which have not been completed. 0: Intel QuickData Technology DMA reports this bit cleared only when all Completions for any outstanding Non-Posted Requests it owns have been received.
4:4	RO	0x0	aux_power_detected: Does not apply to IIO
3:3	RO	0x0	unsupported_request_detected: N/A for Intel QuickData Technology DMA
2:2	RO	0x0	fatal_error_detected: N/A for Intel QuickData Technology DMA
1:1	RO	0x0	non_fatal_error_detected: N/A for Intel QuickData Technology DMA
0:0	RO	0x0	correctable_error_detected: N/A for Intel QuickData Technology DMA

6.4.27 devcap2

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0xb4		Function: 0-7	
Bit	Attr	Default	Description
4:4	RO	0x1	completion_timeout_disable_supported:
3:0	RO	0x0	completion_timeout_values_supported: Not Supported

6.4.28 devcon2

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0xb8		Function: 0-7	
Bit	Attr	Default	Description
4:4	RW	0x0	completion_timeout_disable:
3:0	RO	0x0	completion_timeout_value:

6.4.29 pmcap

Power Management Capability.



The Power Management Capabilities Register defines the capability ID, next pointer and other power management related support. The following Power Management registers / capabilities are added for software compliance.

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0xe0		Function: 0-7	
Bit	Attr	Default	Description
26:26	RO	0x0	d2_support: Does not support power management state D2.
25:25	RO	0x0	d1_support: Does not support power management state D1.
24:22	RO	0x0	aux_current:
21:21	RO	0x0	device_specific_initialization:
19:19	RO	0x0	pme_clock: This field is hardwired to 0h as it does not apply to PCI Express.
18:16	RWS_O	0x3	version: This field is set to 3h (Power Management 1.2 compliant) as version number. Bit is RW-O to make the version 2h incase legacy OS'es have any issues.
15:8	RO	0x0	next_capability_pointer: This is the last capability in the chain and hence set to 0.
7:0	RO	0x1	capability_id: Provides the Power Management capability ID assigned by PCI-SIG.

6.4.30 pmcsr

Power Management Control and Status.

This register provides status and control information for Power Management events in the PCI Express port of the IIO.

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0xe4		Function: 0-7	
Bit	Attr	Default	Description
31:24	RO	0x0	data: N/A
23:23	RO	0x0	bus_power_clock_control_enable: N/A
22:22	RO	0x0	b2_b3_support: N/A
15:15	RO	0x0	pme_status: N/A
14:13	RO	0x0	data_scale: N/A
12:9	RO	0x0	data_select: N/A
8:8	RO	0x0	pme_enable: N/A
3:3	RO	0x1	no_soft_reset: Indicates does not reset its registers when transitioning from D3hot to D0.



Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0xe4		Function: 0-7	
Bit	Attr	Default	Description
1:0	RW_V	0x0	<p>power_state:</p> <p>This 2-bit field is used to determine the current power state of the function and to set a new power state as well.</p> <p>00: D0 01: D1 (not supported by IOAPIC) 10: D2 (not supported by IOAPIC) 11: D3_hot</p> <p>If Software tries to write 01 or 10 to this field, the power state does not change from the existing power state which is either (D0 or D3_hot) and nor do these bits[1:0] change value.</p> <p>When in D3_hot state, IOxAPIC will</p> <ul style="list-style-type: none"> a) respond to only Type 0 configuration transactions targeted at the device's configuration space, when in D3_hot state b) will not respond to memory i.e. D3hot state is equivalent to MSE , accesses to MBAR region note: ABAR region access still go through in D3_hot state, if it enabled c) will not generate any MSI writes

6.4.31 dmauncerrsts

DMA Cluster Uncorrectable Error Status.

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x148		Function: 0	
Bit	Attr	Default	Description
12:12	RW1CS	0x0	<p>syndrome:</p> <p>Multiple errors</p>
10:10	RW1CS	0x0	read_address_decode_error_status:
7:7	RW1CS	0x0	rd_cmpl_header_error_status:
3:3	RW1CS	0x0	dma_internal_hw_parity_error_status:
2:2	RW1CS	0x0	received_poisoned_data_from_dp_status:



6.4.32 dmauncerrmsk

DMA Cluster Uncorrectable Error Mask.


Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	4		
Offset:	0x14c				
Bit	Attr	Default	Description		
12:12	RWS	0x0	syndrome: Multiple errors		
10:10	RWS	0x0	read_address_decode_error_mask:		
7:7	RWS	0x0	rd_cmpl_header_error_mask:		
4:4	RWS	0x0	cfg_reg_parity_error_mask:		
3:3	RWS	0x0	dma_internal_hw_parity_error_mask:		
2:2	RWS	0x0	received_poisoned_data_from_dp_mask:		

6.4.33 dmauncerrsev

DMA Cluster Uncorrectable Error Severity.

This register controls severity of uncorrectable DMA unit errors between fatal and non-fatal.

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	4		
Offset:	0x150				
Bit	Attr	Default	Description		
12:12	RWS	0x0	syndrome: Multiple errors		
10:10	RWS	0x0	read_address_decode_error_severity:		
7:7	RWS	0x1	rd_cmpl_header_error_severity:		
4:4	RWS	0x1	cfg_reg_parity_error_severity:		
3:3	RWS	0x1	dma_internal_hw_parity_error_severity:		
2:2	RWS	0x0	received_poisoned_data_from_dp_severity:		



6.4.34 dmauncerrptr

DMA Cluster Uncorrectable Error Pointer..

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x154		Function: 0	
Bit	Attr	Default	Description
4:0	ROS_V	0x0	uncerrptr: Points to the first unmasked uncorrectable error logged in the DMAUNCERRSTS register. This field is only valid when the corresponding error is unmasked and the status bit is set and this register is rearmed to load again once the error pointed by this field in the uncorrectable error status register is cleared. Value of 0x0 corresponds to bit 0 in DMAUNCERRSTS register, value of 0x1 corresponds to bit 1 and so forth.

6.4.35 dmaglberptr

DMA Cluster Global Error Pointer.

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x160		Function: 0	
Bit	Attr	Default	Description
3:0	ROS_V	0x0	global_error_pointer: Points to one of 8 possible sources of uncorrectable errors – DMA channels 0-7. The DMA channel errors are logged in CHANERRx_INT registers. This register is only valid when the register group pointed to by this register has at least one unmasked error status bit set and this register is rearmed to load again once all the unmasked uncorrectable errors in the source pointed to by this field are cleared. Value of 0x0 corresponds to channel#0, value of 0x1 corresponds to channel#1, and value of 0x7 corresponds to channel#7

6.4.36 chanerr_int

Internal DMA Channel Error Status Registers.

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x180		Function: 0-7	
Bit	Attr	Default	Description
18:18	RW1CS (Function 0-1) RO (Function 2-7)	0x0	descnterr: (Function 0-1) The hardware sets this bit when it encounters a base descriptor that requires an extended descriptor (such as an XOR with 8 sources), but DMACount indicates that the Base descriptor is the last descriptor that can be processed. Reserved. (Function 2-7)
17:17	RW1CS (Function 0-1) RO (Function 2-7)	0x0	xorqerr: The hardware sets this bit when the Q validation part of the XOR with Galois Field Multiply Validate operation fails. Reserved. (Function 2-7)
16:16	RW1CS	0x0	crc_xorp_err: The hardware sets this bit when a CRC Test operation or XOR Validity operation fails or when the P validation part of the XOR with Galois Field Multiply Validate operation fails.



Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x180		Function: 0-7	
Bit	Attr	Default	Description
15:15	RO	0x0	unaffil_err: Unaffiliated Error. IIO never sets this bit
14:14	RO	0x0	unused:
13:13	RW1CS	0x0	int_cfg_err: Interrupt Configuration Error. The DMA channel sets this bit indicating that the interrupt registers were not configured properly when the DMA channel attempted to generate an interrupt e.g. interrupt address is not 0xFEE.
12:12	RW1CS	0x0	cmp_addr_err: Completion Address Error. The DMA channel sets this bit indicating that the completion address register was configured to an illegal address or has not been configured.
11:11	RW1CS	0x0	desc_len_err: Descriptor Length Error. The DMA channel sets this bit indicating that the current transfer has an illegal length field value. When this bit has been set, the address of the failed descriptor is in the Channel Status register.
10:10	RW1CS	0x0	desc_ctrl_err: Descriptor Control Error. The DMA channel sets this bit indicating that the current transfer has an illegal control field value. When this bit has been set, the address of the failed descriptor is in the Channel Status register.
9:9	RW1CS	0x0	wr_data_err: Write Data Error. The DMA channel sets this bit indicating that the current transfer has encountered an error while writing the destination data. This error could be because of an internal ram error in the write queue that stores the write data before being written to main memory. When this bit has been set, the address of the failed descriptor is in the Channel Status register.
8:8	RW1CS	0x0	rd_data_err: Read Data Error. The DMA channel sets this bit indicating that the current transfer has encountered an error while accessing the source data. This error could be a read data that is received poisoned. When this bit has been set, the address of the failed descriptor is in the Channel Status register.
7:7	RW1CS	0x0	dma_data_parerr: DMA Data Parity Error. The DMA channel sets this bit indicating that the current transfer has encountered an uncorrectable ECC/parity error reported by the DMA engine.
6:6	RW1CS	0x0	cdata_parerr: Data Parity Error. The DMA channel sets this bit indicating that the current transfer has encountered a parity error. When this bit has been set, the address of the failed descriptor is in the Channel Status register.
5:5	RW1CS	0x0	chancmd_err: CHANCMD Error. The DMA channel sets this bit indicating that a write to the CHANCMD register contained an invalid value (e.g. more than one command bit set).
4:4	RW1CS	0x0	chn_addr_valerr: Chain Address Value Error. The DMA channel sets this bit indicating that the CHAINADDR register has an illegal address including an alignment error (not on a 64-byte boundary).



Type: CFG PortID: N/A Bus: 0 Device: 4 Function: 0-7 Offset: 0x180			
Bit	Attr	Default	Description
3:3	RW1CS	0x0	descriptor_error: The DMA channel sets this bit indicating that the current transfer has encountered an error (not otherwise covered under other error bits) when reading or executing a DMA descriptor. When this bit has been set and the channel returns to the Halted state, the address of the failed descriptor is in the Channel Status register.
2:2	RW1CS	0x0	nxt_desc_addr_err: Next Descriptor Address Error. The DMA channel sets this bit indicating that the current descriptor has an illegal next descriptor address including an alignment error (not on a 64-byte boundary). When this bit has been set and the channel returns to the Halted state, the address of the failed descriptor is in the Channel Status register.
1:1	RW1CS	0x0	dma_xfrer_daddr_err: DMA Transfer Destination Address Error. The DMA channel sets this bit indicating that the current descriptor has an illegal destination address. When this bit has been set, the address of the failure descriptor has been stored in the Channel Status register.
0:0	RW1CS	0x0	dma_trans_saddr_err: DMA Transfer Source Address Error. The DMA channel sets this bit indicating that the current descriptor has an illegal source address. When this bit has been set, the address of the failure descriptor has been stored in the Channel Status register.

6.4.37 chanerrmsk_int

Internal DMA Channel Error Mask Registers.

Type: CFG PortID: N/A Bus: 0 Device: 4 Function: 0-7 Offset: 0x184			
Bit	Attr	Default	Description
18:18	RWS (Function 0-1) RO (Function 2-7)	0x0	mask18: This register is a bit for bit mask for the CHANERR_INT register 0: enable 1: disable
17:17	RWS (Function 0-1) RO (Function 2-7)	0x0	mask17: This register is a bit for bit mask for the CHANERR_INT register 0: enable 1: disable
16:16	RWS	0x0	mask16: This register is a bit for bit mask for the CHANERR_INT register 0: enable 1: disable
15:15	RO	0x0	chanerrintmskro:
13:0	RWS	0x0	mask13_0: This register is a bit for bit mask for the CHANERR_INT register 0: enable 1: disable



6.4.38 chanerrsev_int

Internal DMA Channel Error Severity Registers.

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x188		Function: 0-7	
Bit	Attr	Default	Description
18:18	RWS (Function 0-1) RO (Function 2-7)	0x0	severity18: (Function 0-1) 1: Corresponding error logged in the CHANERR_INT register is escalated as fatal error to the IIO internal core error logic. 0: That error is escalated as non-fatal to the IIO internal core error logic. Reserved. (Function 2-7)
17:17	RWS (Function 0-1) RO (Function 2-7)	0x0	severity17: (Function 0-1) 1: Corresponding error logged in the CHANERR_INT register is escalated as fatal error to the IIO internal core error logic. 0: That error is escalated as non-fatal to the IIO internal core error logic. Reserved. (Function 2-7)
16:16	RWS	0x0	severity16: 1: Corresponding error logged in the CHANERR_INT register is escalated as fatal error to the IIO internal core error logic. 0: That error is escalated as non-fatal to the IIO internal core error logic.
15:14	RO	0x0	chanerrsevro1_0:
13:0	RWS	0x0	severity13_0: 1: Corresponding error logged in the CHANERR_INT register is escalated as fatal error to the IIO internal core error logic. 0: That error is escalated as non-fatal to the IIO internal core error logic.

6.4.39 chanerrptr

DMA Channel Error Pointer.

Type: CFG		PortID: N/A	
Bus: 0		Device: 4	
Offset: 0x18c		Function: 0-7	
Bit	Attr	Default	Description
4:0	ROS_V	0x0	dma_chan_err_pointer: Points to the first uncorrectable, unmasked error logged in the CHANERR_INT register. This register is only valid when the corresponding error is unmasked and its status bit is set and this register is rearmed to load again once the error pointed to by this register, in the CHANERR_INT status register, is cleared.

6.5 Device 4 Function 0 - 7 MMIO Region Intel QuickData Technology BARs

Intel QuickData Technology MMIO Register used to control the DMA functionality. The Intel QuickData Technology BAR register points to the based address to these registers.



All of these registers are accessible from only the processor. The IIO supports accessing the Intel® QuickData Technology device memory-mapped registers via QWORD reads and writes. The offsets indicated in the following table are from the Intel® QuickData Technology BAR value.

Register Name	Offset	Size
chancnt	0x0	8
xfercap	0x1	8
genctrl	0x2	8
intrctrl	0x3	8
attnstatus	0x4	32
cbver	0x8	8
intrdelay	0xc	16
cs_status	0xe	16
dmacapability	0x10	32
dcaoffset	0x14	16
cbprio	0x40	8
chanctrl	0x80	16
dma_comp	0x82	16
chancmd	0x84	8
dmacount	0x86	16
chansts_0	0x88	32
chansts_1	0x8c	32
chainaddr_0	0x90	32
chainaddr_1	0x94	32
chancmp_0	0x98	32
chancmp_1	0x9c	32
chanerr	0xa8	32
chanerrmsk	0xac	32
dcactrl	0xb0	32
dca_ver	0x100	8
dca_reqid_offset	0x102	16
csi_capability	0x108	16
pcie_capability	0x10a	16
csi_cap_enable	0x10c	16
pcie_cap_enable	0x10e	16
apicid_tag_map	0x110	64
dca_reqid0	0x180	32
dca_reqid1	0x184	32
msgaddr	0x2000	32
msgupaddr	0x2004	32
msgdata	0x2008	32
vecctrl	0x200c	32
pendingbits	0x3000	32



6.5.1 chancnt

Channel Count.

The Channel Count register specifies the number of channels that are implemented.

Type:	MEM	PortID:	8'h7e	Function:	0-7
Bus:	0	Device:	4		
Offset:	0x0				
Bit	Attr	Default	Description		
4:0	RO	0x1	num_chan: Number of channels. Specifies the number of DMA channels. The IIO supports 1 DMA Channel per function so this register will always read 1.		

6.5.2 xfercap

Transfer Capacity.

The Transfer Capacity specifies the minimum of the maximum DMA transfer size supported on all channels.

Type:	MEM	PortID:	8'h7e	Function:	0-7
Bus:	0	Device:	4		
Offset:	0x1				
Bit	Attr	Default	Description		
4:0	RO	0x14	trans_size: Transfer size. This field specifies the number of bytes that may be specified in a DMA descriptor's Transfer Size field. This defines the maximum transfer size supported by IIO as a power of 2. CPU will support 1M max.		

6.5.3 genctrl

DMA General Control.

The DMA Control register provides for general control operations.

Type:	MEM	PortID:	8'h7e	Function:	0-7
Bus:	0	Device:	4		
Offset:	0x2				
Bit	Attr	Default	Description		
0:0	RW	0x0	dbgen: Debug Enable		



6.5.4 intrctrl


The Interrupt Control register provides for control of DMA interrupts.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x3		Function: 0-7	
Bit	Attr	Default	Description
3:3	RW	0x0	msix_vecctrl: Intel QuickData Technology DMA ignores this bit
2:2	RO	0x0	intp: Interrupt. This bit is set whenever the channel status bit in the Attention Status register is set and the Master Interrupt Enable bit is set. That is, it is the logical AND of Interrupt Status and Master Interrupt Enable bits of this register. This bit represents the legacy interrupt drive signal (when in legacy interrupt mode). In MSI-X mode, this bit is not used by software and is a don't care.
1:1	RO	0x0	intp_sts: Interrupt Status. This bit is set whenever the bit in the Attention Status register is set. This bit is not used by software in MSI-X mode and is a don't care.
0:0	RW	0x0	mstr_intp_en: Master Interrupt Enable. Setting this bit enables the generation of an interrupt in legacy interrupt mode. This bit is automatically reset each time this register is read. When this bit is cleared, the IIO will not generate a legacy interrupt under otherwise valid conditions. This bit is not used when DMA is in MSI-X mode.

6.5.5 attnstatus

Attention Status.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x4		Function: 0-7	
Bit	Attr	Default	Description
0:0	RO_V	0x0	chanattn: Channel Attention. Represents the interrupt status of the channel. This bit clears when read. Writes have no impact on this bit.



6.5.6 cbver

The Intel® QuickData Technology version register field indicates the version of the Intel QuickData Technology specification that the IIO implements. The most significant 4-bits (range 7:4) are the major version number and the least significant 4-bits (range 3:0) are the minor version number. The IIO implementation for this Intel QuickData Technology version is 3.2 encoded as 0b0011 0010.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x8		Function: 0-7	
Bit	Attr	Default	Description
7:4	RO	0x3	mjrver: Major Version. Specifies Major version of the Intel QuickData Technology implementation. Current value is 2h
3:0	RO	0x2	mnrver: Minor Version. Specifies Minor version of the Intel QuickData Technology implementation. Current value is 0h

6.5.7 intrdelay

Interrupt Delay.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0xc		Function: 0-7	
Bit	Attr	Default	Description
15:15	RO	0x1	interrupt_coalescing_supported: The IIO does support interrupt coalescing by delaying interrupt generation.
13:0	RW	0x0	interrupt_delay_time: Specifies the number of microseconds that the IIO delays generation of an interrupt (legacy or MSI or MSI-X) from the time that interrupts are enabled (That is, Master Interrupt Enable bit in the CSIPINTRCTRL register is set or, for MSI-X when Vector Control bit1, when CHANCTRL: Interrupt Disable for that channel is reset).

6.5.8 cs_status

Chipset Status.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0xe		Function: 0-7	
Bit	Attr	Default	Description
3:3	RO	0x0	address_remapping: This bit reflects the TE bit of the non-VC1 Intel VT-d engine
2:2	RO	0x0	memory_bypass:
1:1	RO	0x0	mmio_restriction:



6.5.9 dmacapability

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x10		Function: 0-7	
Bit	Attr	Default	Description
9:9	RO_V (Function 0-1) RO (Function 2-7)	0x0	<p>xor_raid6:</p> <p>If set, specifies XOR with Galios Field Multiply Parity and Quotient opcodes for RAID5 and RAID6 are supported. The opcodes are:</p> <p>0x89 - XOR with Galios Field Multiply Generation 0x8A - XOR with Galios Field Multiply Validate 0x8B - XOR with Galios Field Multiply Update Generation</p> <p>Note: When this bit is zero, the DMA engine will halt if it encounters a descriptor with these opcodes.</p>
8:8	RO	0x0	<p>xor_raid5:</p> <p>If set, specifies XOR without Galios Field Multiply parity only opcodes for RAID5 are supported. The opcodes are:</p> <p>0x87 - XOR Generation 0x88 - XOR Validate</p> <p>Notes:</p> <p>When this bit is zero, the DMA engine will halt if it encounters a descriptor with these opcodes.</p>
7:7	RO	0x1	<p>extended_apic_id:</p> <p>Set if 32b APIC ID's are supported.</p> <p>1: 32b APIC ID's supported 0: 8b APIC ID's supported</p>
6:6	RO	0x1	<p>block_fill:</p> <p>If set, specifies the Block Fill opcode is supported. The opcode is:</p> <p>0x01 - Block Fill</p> <p>Notes:</p> <p>When this bit is zero, the DMA engine will abort if it encounters a descriptor with these opcodes.</p>
5:5	RO	0x1	<p>move_crc:</p> <p>If set, specifies Move and CRC opcodes are supported. The opcodes are:</p> <p>0x41 - Move and Generate CRC-32 0x42 - Move and Test CRC-32 0x43 - Move and Store CRC-32</p> <p>Note: When this bit is zero, the DMA engine will abort if it encounters a descriptor with these opcodes.</p>
4:4	RW_O	0x1	<p>dca:</p> <p>If set, specifies DMA DCA operations are supported according to the settings in the descriptors.</p> <p>Notes:</p> <p>When this bit is zero, the DMA engine ignores the DCA hints in DMA descriptors.</p> <p>This bit is RW-O to give bios the ability to turn off DCA operation from Intel QuickData Technology DMA.</p>
3:3	RO	0x0	<p>xor:</p> <p>If set, specifies XOR opcodes are supported. Opcodes are:</p> <p>0x85 - original XOR Generation 0x86 - original XOR Validate</p> <p>Notes:</p> <p>These opcodes have been deprecated in Intel QuickData Technology DMA v3.</p> <p>The DMA engine will abort if it encounters a descriptor with these opcodes.</p>



Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x10		Function: 0-7	
Bit	Attr	Default	Description
2:2	RO	0x1	marker_skipping: If set, specifies the Marker Skipping opcode is supported. The opcode is: 0x84 - Marker Skipping Note: When this bit is zero, the DMA engine will abort if it encounters a descriptor with this opcode.
1:1	RO	0x1	crc: If set, specifies CRC Generation opcodes are supported. Opcodes are: 0x81 - CRC-32 Generation 0x82 - CRC-32 Generation & Test 0x83 - CRC-32 Generation & Store Note: When this bit is zero, the DMA engine will abort if it encounters a descriptor with these opcodes.
0:0	RO	0x1	page_break: If set, specifies a transfer crossing physical pages is supported. Note: When this bit is zero, software must not set SPBrk nor DPBrk bits in the DMA descriptor and the DMA engine generates an error if either of those bits are set.


6.5.10 dcaoffset

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x14		Function: 0-7	
Bit	Attr	Default	Description
15:0	RO	0x100	dcaregptr:

6.5.11 cbprio

Intel QuickData Technology DMA Priority Register.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x40		Function: 0-7	
Bit	Attr	Default	Description
7:0	RO	0x0	not_used:



6.5.12 chanctrl

The Channel Control register controls the behavior of the DMA channel when specific events occur such as completion or errors.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x80		Function: 0-7	
Bit	Attr	Default	Description
9:9	RW_L	0x0	<p>cmpwr_dca_enable:</p> <p>When this bit is set, and the DMA engine supports DCA, then completion writes will be directed to the CPU indicated in Target CPU. This field is RW if CHANCNT register is 1 otherwise this register is RO.</p>
8:8	RW_LV	0x0	<p>in_use:</p> <p>In Use. This bit indicates whether the DMA channel is in use. The first time this bit is read after it has been cleared, it will return 0 and automatically transition from 0 to 1, reserving the channel for the first consumer that reads this register. All subsequent reads will return 1 indicating that the channel is in use. This bit is cleared by writing a 0 value, thus releasing the channel. A consumer uses this mechanism to atomically claim exclusive ownership of the DMA channel. This should be done before attempting to program any register in the DMA channel register set. This field is RW if CHANCNT register is 1 otherwise this register is RO.</p>
5:5	RW_L	0x0	<p>desc_addr_snp_ctrl:</p> <p>Descriptor address snoop control.</p> <p>1: When set, this bit indicates that the descriptors are not in coherent space and should not be snooped.</p> <p>0: When cleared, the descriptors are in coherent space and each descriptor address must be snooped on QPI.</p> <p>This field is RW if CHANCNT register is 1 otherwise this register is RO.</p>
4:4	RW_L	0x0	<p>err_int_en:</p> <p>Error Interrupt Enable. This bit enables the DMA channel to generate an interrupt (MSI or legacy) when an error occurs during the DMA transfer. If Any Error Abort Enable (see below) is not set, then unaffiliated errors do not cause an interrupt. This field is RW if CHANCNT register is 1 otherwise this register is RO.</p>
3:3	RW_L	0x0	<p>anyerr_abrt_en:</p> <p>Any Error Abort Enable. This bit enables an abort operation when any error is encountered during the DMA transfer. When the abort occurs, the DMA channel generates an interrupt and a completion update as per the Error Interrupt Enable and Error Completion Enable bits. When this bit is reset, only affiliated errors cause the DMA channel to abort. This field is RW if CHANCNT register is 1 otherwise this register is RO.</p>
2:2	RW_L	0x0	<p>err_cmp_en:</p> <p>Error Completion Enable. This bit enables a completion write to the address specified in the CHANCMP register upon encountering an error during the DMA transfer. If Any Error Abort is not set, then unaffiliated errors do not cause a completion write. This field is RW if CHANCNT register is 1 otherwise this register is RO.</p>
0:0	RW1C	0x0	<p>intp_dis:</p> <p>Interrupt Disable. Upon completing a descriptor, if an interrupt is specified for that descriptor and this bit is reset, then the DMA channel generates an interrupt and sets this bit. The choice between MSI or legacy interrupt mode is determined with the MSICTRL register. Legacy interrupts are further gated through intxDisable in the PCICMD register of the Intel QuickData Technology DMA PCI configuration space. The controlling process can re-enable this channel's interrupt by writing a one to this bit, which clears the bit. Writing a zero has no effect. Thus, each time this bit is reset, it enables the DMA channel to generate one interrupt.</p>



6.5.13 dma_comp

DMA Compatibility Register.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x82		Function: 0-7	
Bit	Attr	Default	Description
2:2	RO	0x1	v3_compatibility: Compatible with version 3 Intel QuickData Technology spec
1:1	RO	0x1	v2_compatibility: Compatible with version 2 Intel QuickData Technology spec
0:0	RO	0x0	v1_compatibility: Not compatible with version 1

6.5.14 chancmd

DMA Channel Command Register.

Setting more than one of these bits with the same write operation will result in an Fatal error affiliated.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x84		Function: 0-7	
Bit	Attr	Default	Description
5:5	RW_LV	0x0	reset_dma: Set this bit to reset the DMA channel. Setting this bit is a last resort to recover the DMA channel from a programming error or other problem such as dead lock from cache coherency protocol. Execution of this command does not generate an interrupt or generate status. This command causes the DMA channel to return to a known state Halted. This field is RW if CHANCNT register is 1 otherwise this register is RO.
2:2	RW_LV	0x0	susp_dma: Suspend DMA. Set this bit to suspend the current DMA transfer. This field is RW if CHANCNT register is 1 otherwise this register is RO.

6.5.15 dmacount

DMA Descriptor Count Register.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x86		Function: 0-7	
Bit	Attr	Default	Description
15:0	RW_L	0x0	numdesc: This is the absolute value of the number of valid descriptors in the chain. The hardware sets this register and an internal counter to zero whenever the CHAINADDR register is written. When this register does not equal the value of the internal register, the DMA channel processes descriptors, incrementing the internal counter each time that it completes (or skips) a descriptor. This register is RW if CHANCNT register is 1 otherwise this register is RO.



6.5.16 chansts_0

Channel Status 0 Register.

The Channel Status Register records the address of the last descriptor completed by the DMA channel.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x88		Function: 0-7	
Bit	Attr	Default	Description
31:6	RO	0x0	cmpdscaddr: This register stores the upper address bits (64B aligned) of the last descriptor processed. The DMA channel automatically updates this register when an error or successful completion occurs. For each completion, the DMA channel over-writes the previous value regardless of whether that value has been read.
2:0	RO	0x3	dma_trans_state: DMA Transfer Status. The DMA engine sets these bits indicating the state of the current DMA transfer. The cause of an abort can be either error during the DMA transfer or invoked by the controlling process via the CHANCMD register.000 - Active 001 - Idle, DMA Transfer Done (no hard errors) 010 - Suspended 011 - Halted, operation aborted 100 - Armed

6.5.17 chansts_1

Channel Status 1 Register.

The Channel Status Register records the address of the last descriptor completed by the DMA channel.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x8c		Function: 0-7	
Bit	Attr	Default	Description
31:0	RO	0x0	cmpdscaddr: This register stores the upper address bits (64B aligned) of the last descriptor processed. The DMA channel automatically updates this register when an error or successful completion occurs. For each completion, the DMA channel over-writes the previous value regardless of whether that value has been read.



6.5.18 chainaddr_0

Descriptor Chain Address 0 Register.

This register is written by the processor to specify the first descriptor to be fetched by the DMA channel.

Type:	MEM	PortID:	8'h7e
Bus:	0	Device:	4
Offset:	0x90	Function:	0-7
Bit	Attr	Default	Description
31:0	RW_L	0x0	dscaddrlo: This 64 bit field marks the address of the first descriptor to be fetched by the DMA channel. The least significant 6 bits must be zero for the address to be valid. This register is RW if CHANCNT register is 1 otherwise this register is RO.

6.5.19 chainaddr_1

Descriptor Chain Address 1 Register.

This register is written by the processor to specify the first descriptor to be fetched by the DMA channel.

Type:	MEM	PortID:	8'h7e
Bus:	0	Device:	4
Offset:	0x94	Function:	0-7
Bit	Attr	Default	Description
31:0	RW_L	0x0	dscaddrhi: This 64 bit field marks the address of the first descriptor to be fetched by the DMA channel. The least significant 6 bits must be zero for the address to be valid. This register is RW if CHANCNT register is 1 otherwise this register is RO.

6.5.20 chancmp_0

Channel Completion Address 0 Register.

This register specifies the address where the DMA channel writes the completion status upon completion or an error condition i.e. it writes the contents of the CHANSTS register to the destination as pointed by the CHANCMP register.

Type:	MEM	PortID:	8'h7e
Bus:	0	Device:	4
Offset:	0x98	Function:	0-7
Bit	Attr	Default	Description
31:3	RW_L	0x0	chcmpladdr_lo: This 64-bit field specifies the address where the DMA engine writes the completion status (CHANSTS). This address can fall within system memory or memory-mapped I/O space but should be 8-byte aligned. This register is RW if CHANCNT register is 1 otherwise this register is RO.



6.5.21 chancmp_1

Channel Completion Address 1 Register.


This register specifies the address where the DMA channel writes the completion status upon completion or an error condition i.e. it writes the contents of the CHANSTS register to the destination as pointed by the CHANCMP register.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x9c		Function: 0-7	
Bit	Attr	Default	Description
31:0	RW_L	0x0	chcmpladdr_hi: This 64-bit field specifies the address where the DMA engine writes the completion status (CHANSTS). This address can fall within system memory or memory-mapped I/O space but should be 8-byte aligned. This register is RW if CHANCNT register is 1 otherwise this register is RO.

6.5.22 chanerr

The Channel Error Register records the error conditions occurring within a given DMA channel.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0xa8		Function: 0-7	
Bit	Attr	Default	Description
18:18	RW1CS (Function 0-1) RO (Function 2-7)	0x0	descntterr: The hardware sets this bit when it encounters a base descriptor that requires an extended descriptor (such as an XOR with 8 sources), but DMACount indicates that the Base descriptor is the last descriptor that can be processed.
17:17	RW1CS (Function 0-1) RO (Function 2-7)	0x0	xorqerr: The hardware sets this bit when the Q validation part of the XOR with Galois Field Multiply Validate operation fails.
16:16	RW1CS	0x0	crc_xorp_err: The hardware sets this bit when a CRC Test operation or XOR Validity operation fails or when the P validation part of the XOR with Galois Field Multiply Validate operation fails.
15:15	RO	0x0	unaffil_err: Unaffiliated Error . IIO never sets this bit
13:13	RW1CS	0x0	int_cfg_err: Interrupt Configuration Error. The DMA channel sets this bit indicating that the interrupt registers were not configured properly when the DMA channel attempted to generate an interrupt. E.g. interrupt address is not 0xFEE.
12:12	RW1CS	0x0	cmp_addr_err: Completion Address Error. The DMA channel sets this bit indicating that the completion address register was configured to an illegal address or has not been configured.
11:11	RW1CS	0x0	desc_len_err: Descriptor Length Error. The DMA channel sets this bit indicating that the current transfer has an illegal length field value. When this bit has been set, the address of the failed descriptor is in the Channel Status register.



Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0xa8		Function: 0-7	
Bit	Attr	Default	Description
10:10	RW1CS	0x0	desc_ctrl_err: Descriptor Control Error. The DMA channel sets this bit indicating that the current transfer has an illegal control field value. When this bit has been set, the address of the failed descriptor is in the Channel Status register.
9:9	RW1CS	0x0	wr_data_err: Write Data Error. The DMA channel sets this bit indicating that the current transfer has encountered an error while writing the destination data. This error could be because of an internal ram error in the write queue that stores the write data before being written to main memory. When this bit has been set, the address of the failed descriptor is in the Channel Status register.
8:8	RW1CS	0x0	rd_data_err: Read Data Error. The DMA channel sets this bit indicating that the current transfer has encountered an error while accessing the source data. This error could be a read data that is received poisoned. When this bit has been set, the address of the failed descriptor is in the Channel Status register.
7:7	RW1CS	0x0	dma_data_parerr: DMA Data Parity Error. The DMA channel sets this bit indicating that the current transfer has encountered an uncorrectable ECC/parity error reported by the DMA engine.
6:6	RW1CS	0x0	cdata_parerr: Chipset Data Parity Error. The DMA channel sets this bit indicating that the current transfer has encountered a parity error reported by the chipset. When this bit has been set, the address of the failed descriptor is in the Channel Status register.
5:5	RW1CS	0x0	chancmd_err: CHANCMD Error. The DMA channel sets this bit indicating that a write to the CHANCMD register contained an invalid value (for example. more than one command bit set).
4:4	RW1CS	0x0	chn_addr_valerr: Chain Address Value Error. The DMA channel sets this bit indicating that the CHAINADDR register has an illegal address including an alignment error (not on a 64-byte boundary).
3:3	RW1CS	0x0	descriptor_error: The DMA channel sets this bit indicating that the current transfer has encountered an error (not otherwise covered under other error bits) when reading or executing a DMA descriptor. When this bit has been set and the channel returns to the Halted state, the address of the failed descriptor is in the Channel Status register.
2:2	RW1CS	0x0	nxt_desc_addr_err: Next Descriptor Address Error. The DMA channel sets this bit indicating that the current descriptor has an illegal next descriptor address including an alignment error (not on a 64-byte boundary). When this bit has been set and the channel returns to the Halted state, the address of the failed descriptor is in the Channel Status register.
1:1	RW1CS	0x0	dma_xfrer_daddr_err: DMA Transfer Destination Address Error. The DMA channel sets this bit indicating that the current descriptor has an illegal destination address. When this bit has been set, the address of the failure descriptor has been stored in the Channel Status register.



Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0xa8		Function: 0-7	
Bit	Attr	Default	Description
0:0	RW1CS	0x0	dma_trans_saddr_err: DMA Transfer Source Address Error. The DMA channel sets this bit indicating that the current descriptor has an illegal source address. When this bit has been set, the address of the failure descriptor has been stored in the Channel Status register.

6.5.23 chanerrmsk

Channel Error Mask Register.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0xac		Function: 0-7	
Bit	Attr	Default	Description
18:18	RWS (Function 0-1) RO (Function 2-7)	0x0	mask18: This register is a bit for bit mask for the CHANERR register 0: enable 1: disable
17:17	RWS (Function 0-1) RO (Function 2-7)	0x0	mask17: This register is a bit for bit mask for the CHANERR register 0: enable 1: disable
16:16	RWS	0x0	mask16: This register is a bit for bit mask for the CHANERR register 0: enable 1: disable
13:0	RWS	0x0	mask13_0: This register is a bit for bit mask for the CHANERR register 0: enable 1: disable

6.5.24 dcactrl

DCA Control.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0xb0		Function: 0-7	
Bit	Attr	Default	Description
15:0	RW_L	0x0	target_cpu: Specifies the APIC ID of the target CPU for Completion Writes. This field is RW if CHANCNT register is 1 otherwise this register is RO.



6.5.25 dca_ver

DCA Version Number Register.

Type:	MEM	PortID:	8'h7e	Function:	0-7
Bus:	0	Device:	4		
Offset:	0x100				
Bit	Attr	Default	Description		
7:4	RO	0x1	major_revision:		
3:0	RO	0x0	minor_revision:		

6.5.26 dca_reqid_offset

DCA Requester ID Offset.

Type:	MEM	PortID:	8'h7e	Function:	0-7
Bus:	0	Device:	4		
Offset:	0x102				
Bit	Attr	Default	Description		
15:0	RO	0x180	dca_reqid_regs: registers are at offset 180h		

6.5.27 csi_capability


Intel QPI Compatibility Register.

Type:	MEM	PortID:	8'h7e	Function:	0-7
Bus:	0	Device:	4		
Offset:	0x108				
Bit	Attr	Default	Description		
0:0	RO	0x1	prefetch_hint: IIO supports Prefetch Hint only method on the coherent interface		

6.5.28 pcie_capability

PCI Express Capability Register.

Type:	MEM	PortID:	8'h7e	Function:	0-7
Bus:	0	Device:	4		
Offset:	0x10a				
Bit	Attr	Default	Description		
0:0	RO	0x1	memwr: IIO supports only memory write method on PCI Express		



6.5.29 csi_cap_enable

Intel QPI Capability Enable Register.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x10c		Function: 0-7	
Bit	Attr	Default	Description
0:0	RW	0x0	enable_prefetch_hint: When set in function 0, DCA on Intel QPI is enabled, else disabled. IIO hardware does not use this bit from functions 1-7. In these functions, this bit is provided primarily for BIOS to communicate to driver that DCA is enabled in the IIO.

6.5.30 pcie_cap_enable

PCI Express Capability Enable Register.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x10e		Function: 0-7	
Bit	Attr	Default	Description
0:0	RW	0x0	enable_memwr_on_pcie: When set in function 0, DCA on PCIe is enabled, else disabled. IIO hardware does not use this bit from functions 1-7. In these functions, this bit is provided primarily for BIOS to communicate to driver that DCA is enabled in the IIO.


6.5.31 apicid_tag_map

APICID to Tag Map Register.

When DCA is disabled, DMA engine uses all 1s in the tag field of the write.

This register is setup by BIOS for the Intel QuickData Technology driver to read. BIOS will map APICID[7:5] to bits Tag[2:0]. BIOS should set Tag[4] to prevent implicit TPH cache target unless it is intended.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x110		Function: 0-7	
Bit	Attr	Default	Description
39:32	RW	0x80	tag_map_4: This field is used by the Intel QuickData Technology DMA engine to populate Tag field bit 4 of the memory write transaction it issues with either 1, 0, or a selected APICID bit. [7:6] 00: Tag[4] = Tag_Map_4[0] 01: Tag[4] = APICID[Tag_Map_4[3:0]] 10: Tag[4] = NOT(APICID [Tag_Map_4[3:0]]) 11: reserved



Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x110		Function: 0-7	
Bit	Attr	Default	Description
31:24	RW	0x80	<p>tag_map_3:</p> <p>This field is used by the Intel QuickData Technology DMA engine to populate Tag field bit 3 of the memory write transaction it issues with either 1, 0, or a selected APICID bit.</p> <p>[7:6]</p> <p>00: Tag[3] = Tag_Map_3[0] 01: Tag[3] = APICID[Tag_Map_3[3:0]] 10: Tag[3] = NOT(APICID[Tag_Map_3[3:0]]) 11: reserved</p>
23:16	RW	0x80	<p>tag_map_2:</p> <p>This field is used by the Intel QuickData Technology DMA engine to populate Tag field bit 2 of the memory write transaction it issues with either 1, 0, or a selected APICID bit.</p> <p>[7:6]</p> <p>00: Tag[2] = Tag_Map_2[0] 01: Tag[2] = APICID[Tag_Map_2[3:0]] 10: Tag[2] = NOT(APICID[Tag_Map_2[3:0]]) 11: reserved</p>
15:8	RW	0x80	<p>tag_map_1:</p> <p>This field is used by the Intel QuickData Technology DMA engine to populate Tag field bit 1 of the memory write transaction it issues with either 1, 0, or a selected APICID bit.</p> <p>[7:6]</p> <p>00: Tag[1] = Tag_Map_1[0] 01: Tag[1] = APICID[Tag_Map_1[3:0]] 10: Tag[1] = NOT(APICID[Tag_Map_1[3:0]]) 11: reserved</p>
7:0	RW	0x80	<p>tag_map_0:</p> <p>This field is used by the Intel QuickData Technology DMA engine to populate Tag field bit 0 of the memory write transaction it issues with either 1, 0, or a selected APICID bit.</p> <p>[7:6]</p> <p>00: Tag[0] = Tag_Map_0[0] 01: Tag[0] = APICID[Tag_Map_0[3:0]] 10: Tag[0] = NOT(APICID[Tag_Map_0[3:0]]) 11: reserved</p>

6.5.32 dca_reqid[0:1]

Global DCA Requester ID Table Registers.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x180, 0x184		Function: 0-7	
Bit	Attr	Default	Description
31:31	RO	0x0	<p>last:</p> <p>This bit is set only in the last RequesterID register for this port. Thus, it identifies that this is the last DCA RequesterID register for this port.</p>
29:29	RW	0x0	<p>valid:</p> <p>when set the requester id programed into bits 15:0 is used by hardware for DCA write identification, otherwise the bits are ignored.</p>



Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x180, 0x184		Function: 0-7	
Bit	Attr	Default	Description
28:28	RW	0x0	ignore_function_number: When set, the function number field in the RequesterID is ignored when authenticating a DCA write, otherwise the function number is included
15:8	RW	0x0	bus_number: PCI bus number of the DCA requester
7:3	RW	0x0	device_number: Device number of the day requester
2:0	RW	0x0	function_number: Function number of the day requester

6.5.33 msgaddr

MSI-X Lower Address Registers.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x2000		Function: 0-7	
Bit	Attr	Default	Description
31:2	RW_V	0x0	chmsgaddr: Specifies the local APIC to which this MSI-X interrupt needs to be sent.
1:0	RO	0x0	chmsgaddr_const:

6.5.34 msgupaddr

MSI-X Upper Address Registers.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x2004		Function: 0-7	
Bit	Attr	Default	Description
31:0	RW_V	0x0	chmsgupaddr_const: Reserved to 0 because does not apply to IA. This field is RW for compatibility reason only.



6.5.35 msgdata

MSI-X Data Registers.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x2008		Function: 0-7	
Bit	Attr	Default	Description
31:0	RW_V	0x0	chmsgdata: Specifies the vector that needs to be used for interrupts from the DMA engine. IIO uses the lower 16 bits of this field to form the data portion of the interrupt on the coherent interface. The upper 16 bits are not used by IIO and left as RW only for compatibility reasons.

6.5.36 vecctrl

MSI-X Vector Control Registers.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x200c		Function: 0-7	
Bit	Attr	Default	Description
31:1	RO	0x0	chvecctrlcnst:
0:0	RW_V	0x1	chmask: When a bit is set, the channel is prohibited from sending a message, even if all other internal conditions for interrupt generation are valid.

6.5.37 pendingbits

MSI-X Interrupt Pending Bits Registers.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 4	
Offset: 0x3000		Function: 0-7	
Bit	Attr	Default	Description
31:1	RO	0x0	chmsipendcnst: unused
0:0	RW_V	0x0	chmsipend: Pending Bit (when set) indicates that the DMA engine has a pending MSI-X message for the DMA Channel. This bit is cleared by hardware as soon as it issues the MSI-X message. Note that a Pending Bit is set only if all internal conditions for generation of an MSIX interrupt (like the Channel Interrupt Disable bit being cleared, etc.) are valid. This does not include the MSI-X Mask bit for the channel and the MSI-X Function Mask bit. Once set, a Pending Bit remains set until: The corresponding MSI-X Mask bit and the MSI-X Function Mask bit are both cleared, at which time the IIO issues the pending message and clears the bit. Pending bit is cleared when the Interrupt Disable bit in the corresponding 'Channel Control Register (CHANCTRL)' transitions from 1b to 0b and there is not another interrupt pending for that channel - no MSI-X message issued. Implementation Note: Implementations can consider an MSI message 'issued to the system', as soon as the message is 'posted' internally in the device.



6.6 Device 5 Function 0

Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d), Address Mapping, System Management, Coherent Interface, Misc Registers.

Register Name	Offset	Size
vid	0x0	16
did	0x2	16
pcicmd	0x4	16
pcists	0x6	16
rid	0x8	8
ccr	0x9	24
clsr	0xc	8
hdr	0xe	8
svid	0x2c	16
sdid	0x2e	16
capptr	0x34	8
intl	0x3c	8
intpin	0x3d	8
pxpcapid	0x40	8
pxpnxtptr	0x41	8
pxpcap	0x42	16
hdrtypectrl	0x80	8
mmcfg_base	0x90	64
mmcfg_limit	0x98	64
tommio_ob	0xa4	32
tseg	0xa8	64
genprotrange1_base	0xb0	64
genprotrange1_limit	0xb8	64
genprotrange2_base	0xc0	64
genprotrange2_limit	0xc8	64
tolm	0xd0	32
tohm	0xd4	64
tommio	0xdc	32
ncmem_base	0xe0	64
ncmem_limit	0xe8	64
mencmem_base	0xf0	64
mencmem_limit	0xf8	64
cpbusno	0x108	32
lmmio_base	0x10c	16
lmmio_limit	0x10e	16
lmmioh_base	0x110	64
lmmioh_limit	0x118	64
genprotrange0_base	0x120	64
genprotrange0_limit	0x128	64



Register Name	Offset	Size
gcfgbus_base	0x134	8
gcfgbus_limit	0x135	8
cipctrl	0x140	32
cipsts	0x144	32
cipdcasad	0x148	32
cipintrc	0x14c	64
cipintrs	0x154	32
vtbar	0x180	32
vtgenctrl	0x184	16
vtisochctrl	0x188	32
vtgenctrl2	0x18c	32
iotlbpartition	0x194	32
vtuncerrsts	0x1a8	32
vtuncerrmsk	0x1ac	32
vtuncerrsev	0x1b0	32
vtuncerrptr	0x1b4	8
iiomiscctrl	0x1c0	64
ltdpr	0x290	32
lcfgbus_base	0x41c	8
lcfgbus_limit	0x41d	8
csipintrs	0x450	32

6.6.1 vid

Type: CFG Bus: 0 Offset: 0x0		PortID: N/A Device: 5 Function: 0	
Bit	Attr	Default	Description
15:0	RO	0x8086	vendor_identification_number: The value is assigned by PCI-SIG to Intel.

6.6.2 did

Type: CFG Bus: 0 Offset: 0x2		PortID: N/A Device: 5 Function: 0	
Bit	Attr	Default	Description
15:0	RO	0x2f28	device_identification_number: Device ID values vary from function to function.



6.6.3 pcicmd

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x4		Function: 0	
Bit	Attr	Default	Description
10:10	RO	0x0	intx_disable: NA for these devices
9:9	RO	0x0	fast_back_to_back_enable: Not applicable to PCI Express and is hardwired to 0
8:8	RO	0x0	serr_enable: This bit has no impact on error reporting from these devices
7:7	RO	0x0	idsel_stepping_wait_cycle_control: Not applicable to internal devices. Hardwired to 0.
6:6	RO	0x0	parity_error_response: This bit has no impact on error reporting from these devices
5:5	RO	0x0	vga_palette_snoop_enable: Not applicable to internal devices. Hardwired to 0.
4:4	RO	0x0	memory_write_and_invalidate_enable: Not applicable to internal devices. Hardwired to 0.
3:3	RO	0x0	special_cycle_enable: Not applicable. Hardwired to 0.
2:2	RO	0x0	bus_master_enable: Hardwired to 0 since these devices don't generate any transactions
1:1	RO	0x0	memory_space_enable: Hardwired to 0 since these devices don't decode any memory BARs
0:0	RO	0x0	io_space_enable: Hardwired to 0 since these devices don't decode any IO BARs

6.6.4 pcists

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x6		Function: 0	
Bit	Attr	Default	Description
15:15	RO	0x0	detected_parity_error: This bit is set when the device receives a packet on the primary side with an uncorrectable data error including a packet with poison bit set or an uncorrectable addresscontrol parity error. The setting of this bit is regardless of the Parity Error Response bit PERRE in the PCICMD register. IIO will never set this bit.
14:14	RO	0x0	signaled_system_error: Hardwired to 0
13:13	RO	0x0	received_master_abort: Hardwired to 0
12:12	RO	0x0	received_target_abort: Hardwired to 0
11:11	RO	0x0	signaled_target_abort: Hardwired to 0
10:9	RO	0x0	devsel_timing: Not applicable to PCI Express. Hardwired to 0.



Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x6		Function: 0	
Bit	Attr	Default	Description
8:8	RO	0x0	master_data_parity_error: Hardwired to 0
7:7	RO	0x0	fast_back_to_back: Not applicable to PCI Express. Hardwired to 0.
5:5	RO	0x0	pci66mhz_capable: Not applicable to PCI Express. Hardwired to 0.
4:4	RO	0x1	capabilities_list: This bit indicates the presence of a capabilities list structure
3:3	RO	0x0	intx_status: Hardwired to 0

6.6.5 rid

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x8		Function: 0	
Bit	Attr	Default	Description
7:0	RO_V	0x0	revision_id: Reflects the Uncore Revision ID after reset. Reflects the Compatibility Revision ID after BIOS writes 0x69 to any RID register in any Intel® Xeon® Processor E5 v4 product family function.

6.6.6 ccr

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x9		Function: 0	
Bit	Attr	Default	Description
23:16	RO_V	0x8	base_class: Generic Device
15:8	RO_V	0x80	sub_class: Generic Device
7:0	RO_V	0x0	register_level_programming_interface: Set to 00h for all non-APIC devices.



6.6.7 clsr

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xc		Function: 0	
Bit	Attr	Default	Description
7:0	RW	0x0	cacheline_size: This register is set as RW for compatibility reasons only. Cacheline size is always 64B.

6.6.8 hdr

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xe		Function: 0	
Bit	Attr	Default	Description
7:7	RO	0x1	multi_function_device: This bit defaults to 1b since all these devices are multi-function
6:0	RO	0x0	configuration_layout: This field identifies the format of the configuration header layout. It is Type 0 for all these devices. The default is 00h, indicating a 'endpoint device'.

6.6.9 svid

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x2c		Function: 0	
Bit	Attr	Default	Description
15:0	RW_O	0x0	subsystem_vendor_identification_number: The default value specifies Intel but can be set to any value once after reset.

6.6.10 sdid

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x2e		Function: 0	
Bit	Attr	Default	Description
15:0	RW_O	0x0	subsystem_device_identification_number: Assigned by the subsystem vendor to uniquely identify the subsystem



6.6.11 capptr

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	5		
Offset:	0x34				
Bit	Attr	Default	Description		
7:0	RO	0x40	capability_pointer: Points to the first capability structure for the device which is the PCIe capability.		

6.6.12 intl

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	5		
Offset:	0x3c				
Bit	Attr	Default	Description		
7:0	RO	0x0	interrupt_line: NA for these devices		

6.6.13 intpin

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	5		
Offset:	0x3d				
Bit	Attr	Default	Description		
7:0	RO	0x0	interrupt_pin: NA since these devices do not generate any interrupt on their own		

6.6.14 pxpcapid

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	5		
Offset:	0x40				
Bit	Attr	Default	Description		
7:0	RO	0x10	capability_id: Provides the PCI Express capability ID assigned by PCI-SIG.		

6.6.15 pxpnxtptr

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	5		
Offset:	0x41				
Bit	Attr	Default	Description		
7:0	RO	0x0	next_ptr: This field is set to the PCI Power Management capability.		



6.6.16 pxpcap

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x42		Function: 0	
Bit	Attr	Default	Description
13:9	RO	0x0	interrupt_message_number_n_a:
8:8	RO	0x0	slot_implemented_n_a:
7:4	RO	0x9	device_port_type: This field identifies the type of device. It is set to for the DMA to indicate root complex integrated endpoint device.
3:0	RO	0x2	capability_version: This field identifies the version of the PCI Express capability structure. Set to 2h for PCI Express and DMA devices for compliance with the extended base registers.

6.6.17 hdrtypectrl

PCI Header Type Control

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x80		Function: 0	
Bit	Attr	Default	Description
2:0	RW	0x0	clr_hdrmfd: When set, function#0 with in the indicated device shows a value of 0 for bit 7 of the HDR register, indicating a single function device. BIOS sets this bit, when only function#0 is visible within the device, either because SKU reasons or BIOS has hidden all functions but function#0 within the device via the DEVHIDE register. Bit 0 is for Device#1 Bit 1 is for Device#2 Bit 3 is for Device#3 Currently this is defined only for devices 1, 2 and 3 because in other devices it is expected that at least 2 functions are visible to OS or the entire device is hidden.

6.6.18 mmcfg_base

MMCFG Address Base

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x90		Function: 0	
Bit	Attr	Default	Description
31:26	RW_LB	0x3f	mmcfg_base_addr: Indicates the base address which is aligned to a 64 MB boundary.



6.6.19 mmcfg_limit

MMCFG Address Limit.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x98		Function: 0	
Bit	Attr	Default	Description
31:26	RW_LB	0x0	mmcfg_limit_addr: Indicates the limit address which is aligned to a 64MB boundary. Any access that decodes to be between $MMCFG.BASE \leq Addr \leq MMCFG.LIMIT$ targets the MMCFG region and is aborted by IIO. Setting the MMCFG.BASE greater than MMCFG.LIMIT, disables this region.

6.6.20 tommiol_ob

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xA4		Function: 0	
Bit	Attr	Default	Description
31:20	RW_LB	0x0FBF	tommiol_ob: This field is used to prevent non-DMI links, along with Intel QuickData Technology/APIC/NTB primary BARs, from claiming outbound addresses starting above this address and ending at 0xffff_ffff. Bits 19:0 are zero and not writable, and are treated as 1's (like TOLM and TOHM). Set this to 0xfff to disable TOMMIOL_OB. This is intended to be set consistently with TOMMIOL, but the two can be different if needed.

6.6.21 tseg

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xa8		Function: 0	
Bit	Attr	Default	Description
63:52	RW_LB	0x0	limit: Indicates the limit address which is aligned to a 1MB boundary. Any access to falls within $TSEG.BASE[31:20] \leq Addr[31:20] \leq TSEG.LIMIT[31:20]$ is considered to target the Tseg region and IIO aborts it. Note that address bits 19:0 are ignored and not compared. The result is that $BASE[19:0]$ is effectively 00000h and $LIMIT$ is effectively FFFFh. Setting the TSEG.BASE greater than the limit, disable this region.
31:20	RW_LB	0xfe0	base: Indicates the base address which is aligned to a 1MB boundary. Bits [31:20] corresponds to A[31:20] address bits.

6.6.22 genprotrange[1:0]_base

Generic Protected Memory Range X Base Address. (X = 1, 0)



Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xb0, 0x120		Function: 0	
Bit	Attr	Default	Description
50:16	RW_LB	0x7fffffff	<p>base_address: [50:16] of generic memory address range that needs to be protected from inbound dma accesses. The protected memory range can be anywhere in the memory space addressable by the processor. Addresses that fall in this range i.e. GenProtRange.Base[63:16] <= Address [63:16] <= GenProtRange.Limit [63:16], are completly aborted by IIO.</p> <p>Setting the Protected range base address greater than the limit address disables the protected memory region. Note that this range is orthogonal to Intel VT-d spec defined protected address range.</p> <p>Since this register provides for a generic range, it can be used to protect any system dram region or MMIO region from DMA accesses. But the expected usage for this range is to abort all PCIe accesses to the PCI-Segments region.</p>



6.6.23 genprotrange[1:0]_limit

Generic Protected Memory Range X Limit Address. (X = 1, 0)

Type: CFG Bus: 0 Offset: 0xb8, 0x128		PortID: N/A Device: 5 Function: 0	
Bit	Attr	Default	Description
50:16	RW_LB	0x0	limit_address: [50:16] of generic memory address range that needs to be protected from inbound dma accesses. The protected memory range can be anywhere in the memory space addressable by the processor. Addresses that fall in this range i.e. GenProtRange.Base[63:16] <= Address [63:16] <= GenProtRange.Limit [63:16], are completler aborted by IIO. Setting the Protected range base address greater than the limit address disables the protected memory region. Note that this range is orthogonal to Intel VT-d spec defined protected address range. This register is programmed once at boot time and does not change after that, including any quiesce flows. Since this register provides for a generic range, it can be used to protect any system dram region from DMA accesses. The expected usage for this range is to abort all PCIe accesses to the PCI-Segments region.

6.6.24 genprotrange2_base

Generic Protected Memory Range 2 Base Address.

Type: CFG Bus: 0 Offset: 0xc0		PortID: N/A Device: 5 Function: 0	
Bit	Attr	Default	Description
50:16	RW_LB	0x7fffffff	base_address: [50:16] of generic memory address range that needs to be protected from inbound dma accesses. The protected memory range can be anywhere in the memory space addressable by the processor. Addresses that fall in this range i.e. GenProtRange.Base[63:16] <= Address [63:16] <= GenProtRange.Limit [63:16], are completler aborted by IIO. Setting the Protected range base address greater than the limit address disables the protected memory region. Note that this range is orthogonal to Intel VT-d spec defined protected address range. This register is programmed once at boot time and does not change after that, including any quiesce flows. This region is expected to be used to protect against PAM region accesses inbound, but could also be used for other purposes, if needed.



6.6.25 genprotrange2_limit

Generic Protected Memory Range 2 Limit Address.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xc8		Function: 0	
Bit	Attr	Default	Description
50:16	RW_LB	0x0	<p>limit_address:</p> <p>[50:16] of generic memory address range that needs to be protected from inbound dma accesses. The protected memory range can be anywhere in the memory space addressable by the processor. Addresses that fall in this range i.e. GenProtRange.Base[63:16] <= Address [63:16] <= GenProtRange.Limit [63:16], are completely aborted by IIO.</p> <p>Setting the Protected range base address greater than the limit address disables the protected memory region.</p> <p>Note that this range is orthogonal to Intel VT-d spec defined protected address range. This register is programmed once at boot time and does not change after that, including any quiesce flows.</p> <p>This region is expected to be used to protect against PAM region accesses inbound, but could also be used for other purposes, if needed.</p>

6.6.26 tolm

Top of Low Memory

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xd0		Function: 0	
Bit	Attr	Default	Description
31:26	RW_LB	0x0	<p>addr:</p> <p>TOLM Address. Indicates the top of low dram memory which is aligned to a 64MB boundary. A 32 bit transaction that satisfies '0 <= Address[31:26] <= TOLM[31:26]' is a transaction towards main memory.</p>

6.6.27 tohm

Top of High Memory.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xd4		Function: 0	
Bit	Attr	Default	Description
63:26	RW_LB	0x0	<p>addr:</p> <p>TOHM Address. Indicates the limit of an aligned 64 MB granular region that decodes >4 GB addresses towards system dram memory. A 64-bit transaction that satisfies '4G <= A[63:26] <= TOHM[63:26]' is a transaction towards main memory. This register is programmed once at boot time and does not change after that, including during quiesce flows.</p>



6.6.28 tommiol

Type:	CFG	PortID:	N/A
Bus:	0	Device:	5
Offset:	0xdc	Function:	0
Bit	Attr	Default	Description
31:20	RW_LB	0x0FBF	tommiol: This field is used to abort inbound MRd/MWr/atomic accesses starting above this address and ending at 0xffff_ffff, exclusive of the interrupt hole (0xfeex_xxxx). Bits 19:0 are zero and not writable, and are treated as 1's (like TOLM and TOHM). Set this to 0xffff to disable TOMMIOL.



6.6.29 ncmem_base

Non-Coherent Memory Base Address.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xe0		Function: 0	
Bit	Attr	Default	Description
63:26	RW_LB	0x3fffffff	<p>addr:</p> <p>Non Coherent memory base address. Describes the base address of a 64MB aligned dram memory region on Intel QPI that is non-coherent. Address bits [63:26] of an inbound address if it satisfies 'NcMem.Base[63:26] <= A[63:26] <= NcMem.Limit[63:26]' is considered to be towards the non-coherent Intel QPI memory region.</p> <p>The range indicated by the Non-coherent memory base and limit registers does not necessarily fall within the low dram or high dram memory regions as described via the corresponding base and limit registers.</p> <p>This register is programmed once at boot time and does not change after that, including any quiesce flows</p>

6.6.30 ncmem_limit

Non-Coherent Memory Limit.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xe8		Function: 0	
Bit	Attr	Default	Description
63:26	RW_LB	0x0	<p>addr:</p> <p>Non Coherent memory limit address. Describes the limit address of a 64 MB aligned dram memory region on Intel QPI that is non-coherent. Address bits [63:26] of an inbound address if it satisfies 'NcMem.Base[63:26] <= A[63:26] <= NcMem.Limit[63:26]' is considered to be towards the non-coherent Intel QPI memory region.</p> <p>The range indicated by the Non-coherent memory base and limit registers does not necessarily fall within the low dram or high dram memory regions as described via the corresponding base and limit registers.</p> <p>This register is programmed once at boot time and does not change after that, including any quiesce flows.</p>

6.6.31 mencmem_base

Intel® Management Engine (Intel® ME) Non-Coherent Memory Base Address.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xf0		Function: 0	
Bit	Attr	Default	Description
63:19	RW_LB	0x1fffffffff	<p>addr:</p> <p>Intel® Management Engine (Intel® ME) UMA Base Address. Indicates the base address which is aligned to a 1MB boundary. Bits [63:19] corresponds to A[63:19] address bits.</p>



6.6.32 mencmem_limit

Intel® Management Engine (Intel® ME) Non-Coherent Memory Base Limit.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xf8		Function: 0	
Bit	Attr	Default	Description
63:19	RW_LB	0x0	<p>addr:</p> <p>Intel ME UMA Limit Address. Indicates the limit address which is aligned to a 1MB boundary. Bits [63:19] corresponds to A[63:19] address bits. Any address that falls within MENCMEMBASE <= Addr <= MENCMEMLIMIT range is considered to target the UMA range. Setting the MENCMEMBASE greater than the MENCMEMLIMIT disables this range.</p> <p>The range indicated by this register must fall within the low dram or high dram memory regions as described via the corresponding base and limit registers.</p>

6.6.33 cpubusno

CPU Internal Bus Numbers.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x108		Function: 0	
Bit	Attr	Default	Description
24:17	RW_LB	0x0	segment:
16:16	RW_LB	0x0	<p>valid:</p> <p>1: IIO claims PCI config accesses if: the bus# matches the value in bits 7:0 of this register and Dev# >= 16 OR the bus# does not match either the value in bits 7:0 or 15:8 of this register</p> <p>0: IIO does not claim PCI config accesses</p>
15:8	RW_LB	0x0	bus1: Is the internal bus# of rest of uncore (not including IIO). All devices are claimed on behalf of this component. Devices that do not exist within this component on this bus number are master aborted.
7:0	RW_LB	0x0	bus0: The internal bus# of IIO and also PCH. Configuration requests that target Devices 16-31 on this bus number must be forwarded to the PCH by the IIO. Devices 0-15 on this bus number are claimed to send to IIO internal registers.



6.6.34 Immiol_base

Local MMIO Low Base.

Type: CFG Bus: 0 Offset: 0x10c		PortID: N/A Device: 5 Function: 0	
Bit	Attr	Default	Description
15:8	RW_LB	0x0	base: Corresponds to A[31:24] of MMIOL base address. An inbound memory address that satisfies 'local MMIOL base[15:8] <= A[31:24] <= local MMIOL limit[15:8]' is treated as a local peer-to-peer transaction that do not cross coherent interface. Note: Setting LMMIOL.BASE greater than LMMIOL.LIMIT disables local MMIOL peer-to-peer. This register is programmed once at boot time and does not change after that, including any quiesce flows.

6.6.35 Immiol_limit

Local MMIO Low Limit.

Type: CFG Bus: 0 Offset: 0x10e		PortID: N/A Device: 5 Function: 0	
Bit	Attr	Default	Description
15:8	RW_LB	0x0	limit: Corresponds to A[31:24] of MMIOL limit. An inbound memory address that satisfies 'local MMIOL base[15:8] <= A[31:24] <= local MMIOL limit[15:8]' is treated as a local peer-to-peer transaction that does not cross the coherent interface. Note: Setting LMMIOL.BASE greater than LMMIOL.LIMIT disables local MMIOL peer-to-peer. This register is programmed once at boot time and does not change after that, including any quiesce flows.



6.6.36 Immioh_base

Local MMIO High Base.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x110		Function: 0	
Bit	Attr	Default	Description
50:26	RW_LB	0x0	<p>base:</p> <p>Corresponds to A[50:26] of MMIOH base. An inbound memory address that satisfies local MMIOH base [50:26] <= A[63:26] <= local MMIOH limit [50:26] is treated as a local peer-to-peer transaction that does not cross the coherent interface.</p> <p>Notes:</p> <p>Setting LMMIOH.BASE greater than LMMIOH.LIMIT disables local MMIOH peer-to-peer.</p> <p>This register is programmed once at boot time and does not change after that, including any quiesce flows.</p>

6.6.37 Immioh_limit

Local MMIO High Limit.


Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x118		Function: 0	
Bit	Attr	Default	Description
50:26	RW_LB	0x0	<p>Local LMMIOH Limit: Address</p> <p>Corresponds to A[50:26] of Local MMIOH Limit (and Base) Address. An inbound memory address that satisfies the Local MMIO Base Address [50:26] <= A[63:26] <= Local MMIOH Limit Address [50:26], with A[63:51] equal to zero, is treated as a local peer2peer transaction that does not cross the coherent interface (ring).</p> <p>Notes:</p> <p>Setting LMMIOH.BASE greater than LMMIOH.LIMIT disables local MMIOH peer-to-peer.</p> <p>This register is programmed once at boot time and does not change after that, including any quiesce flows.</p>



6.6.38 cipctrl

Coherent Interface Protocol Control.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x140		Function: 0	
Bit	Attr	Default	Description
31:31	RW	0x0	flushpendwr: Whenever this bit is written to 1 (regardless what the current value of this bit is), IRP block first clears bit 0 in CIPSTS register and takes a snapshot of the currently pending write transactions to dram in Write Cache, wait for them to complete fully (i.e. deallocate the corresponding Write CacheRRB entry) and then set bit 0 in CIPSTS register.
30:30	RW	0x0	adr_snapshot_req: Whenever this bit is written to 1, this implies wr\$ snapshot request was due to ADR. This is a status indication and does not cause the snapshot to occur.
28:28	RW	0x0	diswrupdtflow: When set, PCIWriteUpdate command is never issued on IDI and the writes that triggered this flow would be treated as 'normal' writes and the rules corresponding to the 'normal writes' apply.
15:15	RW	0x1	rd_merge_enable:
12:12	RW-LB	0x0	dcaen: When clear, PrefetchHint will not be sent on the coherent interface. The CIPDCASAD table is programmed by BIOS and this bit is set when the table is valid.
10:10	RW-LB	0x1	vcp_pri_en: Give VCp transactions high priority in IRP and set pri=3 when issuing VCp transactions to the ring.
9:9	RW-LB	0x1	vc1_pri_en: Give VC1/m transactions high priority in IRP and set pri=3 when issuing VC1/m transactions to the ring.
8:8	RW	0x0	diswrcomb: Disables wr->wr, rd->rd, and rd->wr transfers. This bit is a don't-care if rd_merge_enable==1. Setting diswrcomb==1 and rd_merge_enable==0 disables all entry to entry transfers in IRP (causing a Cbo request for every switch request).
7:4	RW-LB	0x0	numrtids_isoc_pool1: Limits the number of RTIDs used for VC1/VCp/VCm isoch by Home Agent pool 1. BIOS programs value into this register based on SKU. An encoding of 0 in either numrtids_isoc_pool0 or numrtids_isoc_pool1 disables IIO isoch RTID allocation (useful for VCm in non-isoch systems or for debug). 12-15 are illegal values for this register.
3:0	RW-LB	0x0	numrtids_isoc_pool0: Limits the number of RTIDs used for VC1/VCp/VCm isoch by HA pool 0. BIOS programs value into this register based on SKU. An encoding of 0 in either numrtids_isoc_pool0 or numrtids_isoc_pool1 disables IIO isoch RTID allocation (useful for VCm in non-isoch systems or for debug). 12-15 are illegal values for this register.



6.6.39 cipsts

Coherent Interface Protocol Status.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x144		Function: 0	
Bit	Attr	Default	Description
2:2	RO_V	0x1	rrb_non_phold_arb_empty: This indicates that there are no pending requests in the RRB with the exception of ProcLock / Unlock messages to the lock arbiter. 0 - Pending RRB requests 1 - RRB Empty except for any pending Proclock / Unlock This is a live bit and hence can toggle clock by clock. This is provided mostly as a debug visibility feature.
1:1	RO_V	0x1	rrb_empty: This indicates that there are no pending requests in the RRB. 0 - Pending RRB requests 1 - RRB Empty This is a live bit and hence can toggle clock by clock. This is provided mostly as a debug visibility feature.
0:0	RO_V	0x0	flush_pending_writes: This bit gets cleared whenever bit 31 in CIPCTRL is written to 1 by software and gets set by hw when the pending writes in the Write Cache (at the time bit 31 in CIPCTRL is written to 1 by software) complete i.e. the Write Cache/ RRB entry is deallocated for all those writes.

6.6.40 cipdcasad

Coherent Interface Protocol DCA Source Address Decode.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x148		Function: 0	
Bit	Attr	Default	Description
31:28	RW	0x0	dcalt7: For a TPH/DCA request, specifies the target NodeID[3:0] when the inverted Tag[2:0] is 7 NID[2]==1 disables PrefetchHint issue for ST that maps to this entry.
27:24	RW	0x0	dcalt6: For a TPH/DCA request, specifies the target NodeID[3:0] when the inverted Tag[2:0] is 6 NID[2]==1 disables PrefetchHint issue for ST that maps to this entry.
23:20	RW	0x0	dcalt5: For a TPH/DCA request, specifies the target NodeID[3:0] when the inverted Tag[2:0] is 5 NID[2]==1 disables PrefetchHint issue for ST that maps to this entry.
19:16	RW	0x0	dcalt4: For a TPH/DCA request, specifies the target NodeID[3:0] when the inverted Tag[2:0] is 4 NID[2]==1 disables PrefetchHint issue for ST that maps to this entry.
15:12	RW	0x0	dcalt3: For a TPH/DCA request, specifies the target NodeID[3:0] when the inverted Tag[2:0] is 3 NID[2]==1 disables PrefetchHint issue for ST that maps to this entry.



Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x148		Function: 0	
Bit	Attr	Default	Description
11:8	RW	0x0	dcalt2: For a TPH/DCA request, specifies the target NodeID[3:0] when the inverted Tag[2:0] is 2. NID[2]==1 disables PrefetchHint issue for ST that maps to this entry.
7:4	RW	0x0	dcalt1: For a TPH/DCA request, specifies the target NodeID[3:0] when the inverted Tag[2:0] is 1. NID[2]==1 disables PrefetchHint issue for ST that maps to this entry.
3:0	RW	0x0	dcalt0: For a TPH/DCA request, specifies the target NodeID[3:0] when the inverted Tag[2:0] is 0. NID[2]==1 disables PrefetchHint issue for ST that maps to this entry.

6.6.41 cipintrc

Coherent Interface Protocol Interrupt Control.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x14c		Function: 0	
Bit	Attr	Default	Description
25:25	RW	0x0	dis_intx_route2ich:
24:24	RW	0x0	route_nmi2mca: Route NMI to MCA
18:18	RW	0x0	smi_msi_en: Intel SMI MSI Enable
17:17	RW	0x0	init_msi_en: INIT MSI Enable
16:16	RW	0x0	nmi_msi_en: NMI MSI Enable
11:11	RW	0x1	intr_mask: INTR Mask
10:10	RW	0x1	smi_mask: Intel SMI Mask
9:9	RW	0x1	init_mask: INIT Mask
8:8	RW	0x1	nmi_mask: NMI Mask
1:1	RW	0x0	logical:



6.6.42 cipintrs

Coherent Interface Protocol Interrupt Status.

This register is to be polled by BIOS to determine if internal pending system interrupts are drained out of IIO. General usage model is for software to quiesce the source e.g. IOM global error logic of a system event like Intel SMI, then poll this register till this register indicates that the event is not pending inside IIO. One additional read is required from software, after the register first reads 0 for the associated event.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x154		Function: 0	
Bit	Attr	Default	Description
31:31	RW1CS	0x0	Intel SMI: This is set whenever IIO forwards a VLW from PCH that had the Intel SMI bit asserted
30:30	RW1CS	0x0	nmi: This is set whenever IIO forwards a VLW from PCH that had the NMI bit asserted
7:7	RO_V	0x0	mca_ras_evt_pending: MCA RAS Event Pending
6:6	RO_V	0x0	nmi_ras_evt_pending: NMI RAS Event Pending
5:5	RO_V	0x0	smi_ras_evt_pending: Intel SMI RAS Event Pending
4:4	RO_V	0x0	intr_evt_pending: Intel SMI RAS Event Pending
2:2	RO_V	0x0	init_evt_pending: Intel SMI RAS Event Pending
1:1	RO_V	0x0	nmi_evt_pending: Intel SMI RAS Event Pending
0:0	RO_V	0x0	vlw_msgpend: VLW Message Pending, either generated internally or externally

6.6.43 vtbar

Base Address Register for Intel VT-d.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x180		Function: 0	
Bit	Attr	Default	Description
31:13	RW_LB	0x0	vtd_chipset_base_address: Provides an aligned 8K base address for IIO registers relating to Intel VT-d. All inbound accesses to this region are completely aborted by the IIO.
0:0	RW_LB	0x0	vtd_chipset_base_address_enable: Note that accesses to registers pointed to by VTBAR are accessible via message channel, irrespective of the setting of this enable bit i.e. even if this bit is clear, read/write to Intel VT-d registers are completed normally (writes update registers and reads return the value of the register) for accesses from message channel. This bit is RW-LB i.e. lock is determined based on the 'trusted' bit in message channel when VTGENCTRL[15] is set, else it is RO.



6.6.44 vtgenctrl

Intel VT-d General Control.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x184		Function: 0	
Bit	Attr	Default	Description
15:15	RW_O	0x0	lockvtd: When this bit is 0, the VTBAR[0] is RW-LB, else it is RO.
7:4	RW_LB	0xa	hpa_limit: Represents the host processor addressing limit 0000: 2 ³⁶ (i.e. bits 35:0) 0001: 2 ³⁷ (i.e. bits 36:0) ... 1010: 2 ⁴⁶ (i.e. bits 45:0) When Intel VT-d translation is enabled on an Intel VT-d engine, all host addresses (during page walks) that go beyond the limit specified in this register will be aborted by IIO. Note that pass-through and 'translated' ATS accesses carry the host-address directly in the access and are subject to this check as well.
3:0	RW_LB	0x8	gpa_limit: Represents the guest virtual addressing limit for the non-Isch Intel VT-d engine. 0000: 2 ⁴⁰ (i.e. bits 39:0) 0001: 2 ⁴¹ (i.e. bits 40:0) .. 0111: 2 ⁴⁷ 1000: 2 ⁴⁸ Others: Reserved When Intel VT-d translation is enabled, all incoming guest addresses from PCI Express, associated with the non-isoch Intel VT-d engine, that go beyond the limit specified in this register will be aborted by IIO and a UR response returned. This register is not used when translation is not enabled. Note that 'translated' and 'pass-through' addresses are in the 'host-addressing' domain and NOT 'guest-addressing' domain and hence GPA_LIMIT checking on those accesses are bypassed and instead HPA_LIMIT checking applies.

6.6.45 vtgenctrl2

Intel VT-d General Control 2.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x18c		Function: 0	
Bit	Attr	Default	Description
18:12	RW_LB	0x4	tlb_free_entry_limit: Retry prefetch request when number of entries available for allocation in the IOTLB is less than the programmed value. Set this to 0 to disable it.




Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x18c		Function: 0	
Bit	Attr	Default	Description
11:11	RW_LB	0x0	lructrl: Controls what increments the LRU counter that is used to degrade the LRU bits in the IOTLB, L1/L2, and L3 caches. 1: Count Cycles same as TB 0: Count Requests
10:7	RW_LB	0x7	lt: Controls the rate at which the LRU buckets should degrade. If we are in "Request" mode (LRUCTRL = 0), then we will degrade LRU after 16 * N requests where N is the value of this field. If we are in "Cycles" mode (LRUCTRL = 1), then we will degrade LRU after 256 * N cycles where N is the value of this field.
3:3	RW_LB	0x0	ignoreubitleafeviction: Do not use U bit in leaf entry for leaf eviction policy on untranslated DMA requests (AT=00b)
2:2	RW_LB	0x0	evictnonleafat01: Mark non-leaf entries on translation requests with AT=01 for early eviction
1:1	RW_LB	0x0	dontevictleafat01: Do not mark leaf entries with U=0 on translation requests with AT=01 for early eviction

6.6.46 iotlbpartition

IOTLB Partitioning Control.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x194		Function: 0	
Bit	Attr	Default	Description
28:27	RW	0x0	rangesel_dmi_20_22: Range Selection for DMI[20:22]
26:25	RW	0x0	rangesel_iou24_upper_x2: Range Selection for IOU24 upper X2 link
24:23	RW	0x0	rangesel_iou23_upper_x2: Range Selection for IOU23 upper X2 link
14:13	RW	0x0	rangesel_me: Range Selection for ME
12:11	RW	0x0	rangesel_cb: Range Selection for Intel QuickData Technology.
10:9	RW	0x0	rangesel_intr: Range Selection for INTR
0:0	RW_LB	0x0	iotlb_parten: 0: Disabled 1: Enabled



6.6.47 vtuncerrsts

Intel VT-d Uncorrectable Error Status.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x1a8		Function: 0	
Bit	Attr	Default	Description
31:31	RW1CS	0x0	vtderr: When set, this bit is set when an Intel VT-d spec defined error has been detected (and logged in the Intel VT-d fault registers)
8:8	RW1CS	0x0	protmemviol: Protected memory region space violated status
7:7	RW1CS	0x0	miscerrs: This error bit is set when TE is off DMA/INTR request has AT set to nonzero value.
6:6	RW1CS	0x0	unsucc_ci_rdcpl: Unsuccessful status received in the coherent interface read completion status.
5:5	RW1CS	0x0	perr_tlb1: TLB1 Parity Error Status.
4:4	RW1CS	0x0	perr_tlb0: TLB0 Parity Error Status.
3:3	RW1CS	0x0	perr_l3_lookup: Data Parity error while doing a L3 lookup status.
2:2	RW1CS	0x0	perr_l1_lookup: Data Parity error while doing a L1 lookup status. Note the mapping of this register field varies over the mapping in tuncerrmsk and vtuncerrsev.
1:1	RW1CS	0x0	perr_l2_lookup: Data Parity error while doing a L1 lookup status. Note the mapping of this register field varies over the mapping in tuncerrmsk and vtuncerrsev.
0:0	RW1CS	0x0	perr_context_cache:



6.6.48 vtuncerrmsk

Intel VT-d Uncorrectable Error Mask.

Mask out error reporting to IIO. Bit 31 should always be set to 1. It is recommend that the other bits be left as zero so these internal errors are reported out.

Setting bits will not prevent any error collecting inside of Intel VT-d in the Intel VT-d Fault Recording Registers.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x1ac		Function: 0	
Bit	Attr	Default	Description
31:31	RWS	0x1	vtderr_msk: This bit should be set to 1 by BIOS. It is highly recommended that this bit is never set to 0. If Intel VT-d errors are configured to be fatal, leaving this bit set to 0 will cause Fatal errors to be reported when devices send illegal requests. This is generally undesirable.
8:8	RWS	0x0	protmemviol_msk: Protected memory region space violated mask
7:7	RWS	0x0	miscerrm: miscerrm mask Illegal request to 0xFEE, GPAHPA limit error mask
6:6	RWS	0x0	unsucc_ci_rdcpl_msk: Unsuccessful status received in the coherent interface read completion mask.
5:5	RWS	0x0	perr_tlb1_msk: TLB1 Parity Error mask
4:4	RWS	0x0	perr_tlb0_msk: TLB0 Parity Error mask
3:3	RWS	0x0	perr_l3_lookup_msk: Data Parity error while doing a L3 lookup mask
2:2	RWS	0x0	perr_l2_lookup_msk: Data Parity error while doing a L2 lookup mask
1:1	RWS	0x0	perr_l1_lookup_msk: Data Parity error while doing a L1 lookup mask
0:0	RWS	0x0	perr_context_cache_msk: Data Parity error while doing a context cache lookup mask.



6.6.49 vtuncerrsev

Intel VT-d Uncorrectable Error Severity.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x1b0		Function: 0	
Bit	Attr	Default	Description
31:31	RWS	0x0	vtderr_sev: When set, this bit escalates reporting of Intel VT-d spec defined errors, as FATAL errors. When clear, those errors are escalated as Nonfatal errors. Setting this bit to a 1 can allow a guest VM to trigger an unrecoverable FATAL error at the platform. It is HIGHLY recommended that BIOS keep this bit set to 0, as such behavior is generally undesirable.
8:8	RWS	0x1	protmemviol_sev: Protected memory region space violated severity.
7:7	RWS	0x1	miscrrsev: miscrrsev severity. Illegal request to 0xFEE, GPAHPA limit error severity
6:6	RWS	0x0	unsucc_ci_rdcv_sev: Unsuccessful status received in the coherent interface read completion severity.
5:5	RWS	0x1	perr_tlb1_sev: TLB1 Parity Error severity.
4:4	RWS	0x1	perr_tlb0_sev: TLB1 Parity Error severity.
3:3	RWS	0x1	perr_l3_lookup_sev: Data Parity error while doing a L3 lookup severity.
2:2	RWS	0x1	perr_l2_lookup_sev: Data Parity error while doing a L2 lookup severity.
1:1	RWS	0x1	perr_l1_lookup_sev: Data Parity error while doing a L1 lookup severity.
0:0	RWS	0x1	perr_context_cache_sev: Data Parity error while doing a context cache lookup severity.

6.6.50 vtuncerrptr

Intel VT-d Uncorrectable Error Pointer.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x1b4		Function: 0	
Bit	Attr	Default	Description
4:0	ROS_V	0x0	vt_uncerr_ptr: This field points to which of the unmasked uncorrectable errors happened first. This field is only valid when the corresponding error is unmasked and the status bit is set and this field is rearmed to load again when the status bit indicated to by this pointer is cleared by software from 1 to 0. Value of 0x0 corresponds to bit 0 in VTUNCERRSTS register, value of 0x1 corresponds to bit 1 and so forth.



6.6.51 iiomiscctrl

IIO MISC Control.

Type: CFG Bus: 0 Offset: 0x1c0		PortID: N/A Device: 5 Function: 0																															
Bit	Attr	Default	Description																														
41:41	RW	0x0	<p>en_poismsg_spec_behavior:</p> <p>A received poison packet is treated as a Fatal error if it's severity bit is set, but treated as a correctable if the severity bit is cleared and logged in both the UNCERRSTS register and the Advisory Non-Fatal Error bit in the CORERRSTS register.</p> <p>When this bit is clear:</p> <table border="0"> <tr> <td>sev</td> <td>pfen</td> <td>error</td> </tr> <tr> <td>0</td> <td>0</td> <td>non-fatal</td> </tr> <tr> <td>0</td> <td>1</td> <td>correctable</td> </tr> <tr> <td>1</td> <td>0</td> <td>fatal</td> </tr> <tr> <td>1</td> <td>1</td> <td>correctable</td> </tr> </table> <p>When this bit is set:</p> <table border="0"> <tr> <td>sev</td> <td>pfen</td> <td>error</td> </tr> <tr> <td>0</td> <td>0</td> <td>non-fatal</td> </tr> <tr> <td>0</td> <td>1</td> <td>correctable</td> </tr> <tr> <td>1</td> <td>0</td> <td>fatal</td> </tr> <tr> <td>1</td> <td>1</td> <td>fatal</td> </tr> </table>	sev	pfen	error	0	0	non-fatal	0	1	correctable	1	0	fatal	1	1	correctable	sev	pfen	error	0	0	non-fatal	0	1	correctable	1	0	fatal	1	1	fatal
sev	pfen	error																															
0	0	non-fatal																															
0	1	correctable																															
1	0	fatal																															
1	1	correctable																															
sev	pfen	error																															
0	0	non-fatal																															
0	1	correctable																															
1	0	fatal																															
1	1	fatal																															
37:37	RW	0x0	<p>poisfen:</p> <p>Enables poisoned data received inbound (either inbound posted data or completions for outbound reads that have poisoned data) to be forwarded to the destination (DRAM or Cache or PCIe Peer).</p> <p>0: Poison indication is not forwarded with the data (this may result in silent corruption if AER poison reporting is disabled.)</p> <p>1: Poison indication is forwarded with the data (this may result in a conflict with MCA poison reporting if AER poison reporting is enabled)</p>																														
33:33	RWS	0x0	<p>force_6b_mc_group:</p> <p>0 = Use 4 bits for Dualcast group</p> <p>1 = Use 6 bits for Dualcast group</p>																														
25:25	RWS	0x1	<p>cballoccn:</p> <p>When set, use Allocating Flows for non-DCA writes from Intel QuickData Technology DMA. This bit does not affect DCA requests when DCA requests are enabled (bit 21 of this register). A DCA request is identified as matching the DCA requestor ID and having a Tag of non-zero. All DCA requests are always allocating, unless they are disabled, or unless all allocating flows are disabled (bit 24). If all allocating flows are disabled, then DCA requests are also disabled.</p> <p>BIOS is to leave this bit at default of 1b for all but DMI port.</p>																														
24:24	RW	0x0	<p>disable_all_allocating_flows:</p> <p>When this bit is set, IIO will no more issue any new inbound IDI command that can allocate into LLC. Instead, all the writes will use one of the non-allocating commands - PCIWIL/PCIWILF/PCINSWr/PCINSWrF. Software should set this bit only when no requests are being actively issued on IDI. So either a lock/quiesce flow should be employed before this bit is set/cleared or it should be set up before DMA is enabled in system.</p>																														



Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x1c0		Function: 0	
Bit	Attr	Default	Description
19:19	RW	0x0	<p>rvgaen: Remote VGA Enable Enables VGA accesses to be sent to remote node. If set, accesses to the VGA region (A_0000 to B_FFFF) will be forwarded to the CBo where it will determine the NodeID where the VGA region resides. It will then be forwarded to the given remote node. If clear, then VGA accesses will be forwarded to the local PCIe port that has its VGAEN set. If none have their VGAEN set, then the request will be forwarded to the local DMI port, if operating in DMI mode. If it is not operating in DMI mode, then the request will be aborted.</p>
18:18	RW	0x1	<p>disable_inbound_ro_for_vc0: When enabled this mode will treat all inbound write traffic as RO = 0 for VC0. This affects all PCI Express ports and the DMI port.0 - Ordering of inbound transactions is based on RO bit for VC0 1 - RO bit is treated as '0' for all inbound VC0 traffic Note that this pretty much impacts only the NS write traffic because for snooped traffic RO bit is ignored by h/w. When this bit is set, the NS write if enabled BW is going to be generally bad. Note that this bit does not impact VC1 and VCm writes</p>
17:16	RW	0x1	<p>dmi_vc1_write_ordering: Mode is used to control VC1 write traffic from DMI (Intel VT). 00: Reserved 01: Serialize writes on CSI issuing one at a time 10: Pipeline writes on CSI except for writes with Tag value of 0x21 which are issued only after prior writes have all completed and reached global observability 11: Pipeline writes on CSI based on RO bit i.e. if RO = 1, pipeline a write on QPI without waiting for prior write to have reached global observability. If RO0, then it needs to wait till prior writes have all reached global observability.</p>
15:15	RW	0x0	<p>dmi_vc1_vt_d_fetch_ordering: This mode is to allow VC1 Intel VT-d conflicts with outstanding VC0 Intel VT-d reads on IDI to be pipelined. This can occur when Intel VT-d tables are shared between Intel VT (VC1) and other devices. To ensure QoS the Intel VT-d reads from VC1 need to be issued in parallel with non-Isoc accesses to the same cacheline. 0: Serialize all IDI address conflicts to DRAM 1: Pipeline Intel VT-d reads from VC1 with address conflict on IDI Notes: A maximum of 1 VC1 Intel VT-d read and 1 non-VC1 Intel VT-d read to the same address can be outstanding on IDI.</p>
13:13	RW	0x0	<p>vc1_reads_bypass_writes: 0: VC1 Reads push VC1 writes 1: VC1 Reads are allowed to bypass VC1 writes</p>
12:12	RW	0x0	<p>lock_thaw_mode: Mode controls how inbound queues in the south agents (PCIe, DMI) thaw when they are target of a locked read. 0: Thaw only posted requests 1: Thaw posted and non-posted requests. Note that if the lock target is also a 'problematic' port (as indicated by bit 38 in MISCCTRLSTS register), then this becomes meaningless because both posted and non-posted requests are thawed.</p>
8:8	RW	0x0	<p>tocmvalid: Enables the TOCM field.</p>



Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x1c0		Function: 0	
Bit	Attr	Default	Description
7:3	RW	0xe	tocm: Indicates the top of Core physical addressability limit. 00000-00100: Reserved 00101: 2 ³⁷ 00110: 2 ³⁸ ... 1110: 2 ⁴⁶ 01111 -11111: Reserved IIO uses this to abort all inbound transactions that cross this limit.
2:2	RW	0x0	en1k: This bit when set, enables 1K granularity for IO space decode in each of the virtual P2P bridges corresponding to root ports, and DMI ports.
1:1	RWS_O	0x0	uniphy_disable: Place entire UNIPHY in L2 for when no ports are used, as in some multi-socket configurations

6.6.52 Itdpr

Intel TXT DMA Protected Range.

General Description: This register holds the address and size of the DMA protected memory region for Intel® Trusted Execution Technology (Intel® TXT) MP usage.


Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x290		Function: 0	
Bit	Attr	Default	Description
31:20	RO_V	0x0	topofdpr: Top address + 1 of DPR. This is RO, and it is copied by HW from TSEGBASE[31:20].



Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x290		Function: 0	
Bit	Attr	Default	Description
11:4	RW_L	0x0	<p>size:</p> <p>This is the size of memory, in MB, that will be protected from DMA accesses. A value of 0x00 in this field means no additional memory is protected. The maximum amount of memory that will be protected is 255 MB.</p> <p>The amount of memory reported in this field will be protected from all DMA accesses. The top of the protected range is typically the BASE of TSEG -1. BIOS is expected to program that in to bits 31:20 of this register.</p> <p>Notes:</p> <p>If TSEG is not enabled, then the top of this range becomes the base ME stolen space, whichever would have been the location of TSEG, assuming it had been enabled.</p> <p>The DPR range works independently of any other range - Generic Protected ranges, TSEG range, Intel VT-d tables, Intel VT-d protection ranges, MMCFG protection range and is done post any Intel VT-d translation or Intel TXT checks. Therefore incoming cycles are checked against this range after the VTd translation and faulted if they hit this protected range, even if they passed the VTd translation.</p> <p>All the memory checks are OR'ed with respect to NOT being allowed to go to memory. So if either Generic protection range, DPR, Intel VT-d, TSEG range disallows the cycle, then the cycle is not allowed to go to memory. Or in other words, all the above checks must pass before a cycle is allowed to DRAM.</p> <p>DMA remap engines are allowed to access the DPR region without any faulting. It is always legal for any DMA remap engine to read or write into the DPR region, thus DMA remap accesses must not be checked against the DPR range.</p>
2:2	RW_L	0x0	<p>commandbit:</p> <p>Writing a '1' to this bit will enable protection.</p> <p>Writing a '0' to this bit will disable protection.</p>
1:1	RO	0x0	<p>protregsts:</p> <p>IIO sets this bit when the protection has been enabled in hardware and for all practical purposes this should be immediate. When protection is disabled, then this bit is clear</p>
0:0	RW_O	0x0	<p>lock:</p> <p>Bits 19:0 are locked down in this register when this bit is set. Can this be set while other bits are being written to in the same write transaction</p>

6.6.53 Icfgbus_base

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x41c		Function: 0	
Bit	Attr	Default	Description
7:0	RW	0x0	lcfgbus_base:



6.6.54 lcfgbus_limit

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	5		
Offset:	0x41d				
Bit	Attr	Default	Description		
7:0	RW	0x0	lcfgbus_limit:		

6.6.55 csipintrs

Type:	CFG	PortID:	N/A	Function:	0
Bus:	0	Device:	5		
Offset:	0x450				
Bit	Attr	Default	Description		
7:7	RO_V	0x0	mca_ras_evt_pend: MCA event interrupt pending.		
6:6	RO_V	0x0	nmi_ras_evt_pend: NMI RAS event interrupt pending.		
5:5	RO_V	0x0	smi_ras_evt_pend: Intel SMI RAS event interrupt pending.		
4:4	RO_V	0x0	intr_evt_pend: Intr event interrupt pending.		
2:2	RO_V	0x0	init_evt_pend: Init event interrupt pending.		
1:1	RO_V	0x0	nmi_evt_pend: NMI event interrupt pending.		
0:0	RO_V	0x0	smi_evt_pend: Intel SMI event interrupt pending.		

6.7 Device 5 Function 0 MMIO Region VTBAR

Intel VT-d registers are all addressed using aligned dword or aligned qword accesses. Any combination of bits is allowed within a dword or qword access. The Intel VT-d remap engine registers corresponding to the port represented by Device 0, occupy the first 4 K of offset starting from the base address defined by VTBAR register.

Register Name	Offset	Size
vtd0_version	0x0	32
vtd0_cap	0x8	64
vtd0_ext_cap	0x10	64
vtd0_glbcmd	0x18	32
vtd0_glbsts	0x1c	32
vtd0_rootentryadd	0x20	64
vtd0_ctxcmd	0x28	64
vtd0fltsts	0x34	32
nonisoch_fltvctrl	0x38	32



Register Name	Offset	Size
nonisochnv_fltvtdata	0x3c	32
vtd0_fltvtdaddr	0x40	32
vtd0_fltvtdupraddr	0x44	32
vtd0_pmen	0x64	32
vtd0_prot_low_mem_base	0x68	32
vtd0_prot_low_mem_limit	0x6c	32
vtd0_prot_high_mem_base	0x70	64
vtd0_prot_high_mem_limit	0x78	64
vtd0_inv_queue_head	0x80	64
vtd0_inv_queue_tail	0x88	64
vtd0_inv_queue_add	0x90	64
vtd0_inv_comp_status	0x9c	32
nonisochnv_inv_cmp_evtctrl	0xa0	32
nonisochnv_invevtdata	0xa4	32
vtd0_inv_comp_evt_addr	0xa8	32
vtd0_inv_comp_evt_upraddr	0xac	32
vtd0_intr_remap_table_base	0xb8	64
vtd0_fltrec0_gpa	0x100	64
vtd0_fltrec0_src	0x108	64
vtd0_fltrec1_gpa	0x110	64
vtd0_fltrec1_src	0x118	64
vtd0_fltrec2_gpa	0x120	64
vtd0_fltrec2_src	0x128	64
vtd0_fltrec3_gpa	0x130	64
vtd0_fltrec3_src	0x138	64
vtd0_fltrec4_gpa	0x140	64
vtd0_fltrec4_src	0x148	64
vtd0_fltrec5_gpa	0x150	64
vtd0_fltrec5_src	0x158	64
vtd0_fltrec6_gpa	0x160	64
vtd0_fltrec6_src	0x168	64
vtd0_fltrec7_gpa	0x170	64
vtd0_fltrec7_src	0x178	64
vtd0_invaddrreg	0x200	64
vtd0_iotlbinv	0x208	64
vtd1_version	0x1000	32
vtd1_cap	0x1008	64
vtd1_ext_cap	0x1010	64
vtd1_glbcmd	0x1018	32
vtd1_glbsts	0x101c	32
vtd1_rootentryadd	0x1020	64
vtd1_ctxcmd	0x1028	64
vtd1_fltsts	0x1034	32



Register Name	Offset	Size
vtd1_fltvtaddr	0x1040	32
vtd1_fltvtupraddr	0x1044	32
vtd1_pmen	0x1064	32
vtd1_prot_low_mem_base	0x1068	32
vtd1_prot_low_mem_limit	0x106c	32
vtd1_prot_high_mem_base	0x1070	64
vtd1_prot_high_mem_limit	0x1078	64
vtd1_inv_queue_head	0x1080	64
vtd1_inv_queue_tail	0x1088	64
vtd1_inv_queue_add	0x1090	64
vtd1_inv_comp_status	0x109c	32
vtd1_inv_comp_evt_addr	0x10a8	32
vtd1_inv_comp_evt_upraddr	0x10ac	32
vtd1_intr_remap_table_base	0x10b8	64
vtd1_fltrec0_gpa	0x1100	64
vtd1_fltrec0_src	0x1108	64
vtd1_invaddrreg	0x1200	64
vtd1_iotlbinv	0x1208	64

6.7.1 vtd[0:1]_version

Intel VT-d Version Number.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x0, 0x1000		Function: 0	
Bit	Attr	Default	Description
7:4	RO	0x1	major_revision:
3:0	RO	0x0	minor_revision:

6.7.2 vtd[0:1]_cap

Intel VT-d Capabilities.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x8, 0x1008		Function: 0	
Bit	Attr	Default	Description
59:59	RW_O	0x1	posted_interrupts_support: The processor supports posted interrupts
55:55	RO	0x1	dma_read_draining: The processor supports hardware based draining
54:54	RO	0x1	dma_write_draining: The processor supports hardware based write draining




Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x8, 0x1008		Function: 0	
Bit	Attr	Default	Description
53:48	RO	0x12	mamv: The processor support MAMV value of 12h (up to 1G super pages).
47:40	RO	0x7	number_of_fault_recording_registers: The processor supports 8 fault recording registers
39:39	RO	0x1	page_selective_invalidation: Supported in IIO
37:34	RW_O	0x3	super_page_support: 2 MB, 1G supported.
33:24	RO	0x10	fault_recording_register_offset: Fault registers are at offset 100h
23:23	RO	0x0	spatial_separation:
22:22	RO	0x1	zlr: Zero-length DMA requests to write-only pages supported.
21:16	RO_V	0x2f	mgaw: This register is set by the processor-based on the setting of the GPA_LIMIT register. The value is the same for both the Intel VT and non-Intel VT engines. This is because the translation for Intel VT has been extended to be 4-level (instead of 3).
12:8	RO	0x4	sagaw: Supports 4-level walk on both Intel VT and non-Intel VT engines
7:7	RO	0x0	tcm: The processor does not cache invalid pages. This bit should always be set to 0 on HW. It can be set to one when we are doing software virtualization of Intel VT-d.
6:6	RO	0x1	phmr_support: The processor supports protected high memory range
5:5	RO	0x1	plmr_support: The processor supports protected low memory range
4:4	RO	0x0	rwbfi:
3:3	RO	0x0	advanced_fault_logging: The processor does not support advanced fault logging
2:0	RO	0x6	number_of_domains_supported: The processor supports 256 domains with 8 bit domain ID

6.7.3 vtd[0:1]_ext_cap

Extended Intel VT-d Capability.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x10, 0x1010		Function: 0	
Bit	Attr	Default	Description
23:20	RO	0xf	maximum_handle_mask_value: IIO supports all 16 bits of handle being masked. Note IIO always performs global interrupt entry invalidation on any interrupt cache invalidation command and h/w never really looks at the mask value.
17:8	RO	0x20	invalidation_unit_offset: IIO has the invalidation registers at offset 200h




Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x10, 0x1010		Function: 0	
Bit	Attr	Default	Description
7:7	RO	0x1	snoop_control: 0: Hardware does not support 1-setting of the SNP field in the page-table entries. 1: Hardware supports the 1-setting of the SNP field in the page-table entries. IIO supports snoop override only for the non-isoch Intel VT-d engine
6:6	RO	0x1	pass_through: IIO supports pass through.
4:4	RW_O	0x1	ia32_extended_interrupt_mode: IIO supports the extended interrupt mode
3:3	RO	0x1	interrupt_remapping_support: IIO supports this
2:2	RW_O	0x1	device_tlb_support: IIO supports ATS for the non-isoch Intel VT-d engine. This bit is RW-O for non-isoch engine. For VTD[0]_EXT_CAP.Bit[2] the default is 1, but can be programmed to 0. Clarification: For VTD[1]_EXT_CAP.Bit[2] the default is 0
1:1	RO	0x1	queued_invalidation_support: IIO supports this. For VTD[1]_EXT_CAP.Bit[1] the default is 0.
0:0	RW_O	0x0	coherency_support: BIOS can write to this bit to indicate to hardware to either snoop or not-snoop the DMA/Interrupt table structures in memory (root/context/pd/pt/irt). Note that this bit is expected to be always set to 0 for the Intel VT-d engine and programmability is only provided for that engine for debug reasons.

6.7.4 vtd[0:1]_glbcmd

Intel VT-d Global Command.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x18, 0x1018		Function: 0	
Bit	Attr	Default	Description
31:31	RW	0x0	translation_enable: Software writes to this field to request hardware to enable/disable DMA-remapping hardware. 0: Disable DMA-remapping hardware 1: Enable DMA-remapping hardware Hardware reports the status of the translation enable operation through the TES field in the Global Status register. Before enabling (or re-enabling) DMA-remapping hardware through this field, software must: - Setup the DMA-remapping structures in memory - Flush the write buffers (through WBF field), if write buffer flushing is reported as required. - Set the root-entry table pointer in hardware (through S RTP field). - Perform global invalidation of the context-cache and global invalidation of IOTLB - If advanced fault logging supported, setup fault log pointer (through SFL field) and enable advanced fault logging (through EAFL field). There may be active DMA requests in the platform when software updates this field. Hardware must enable or disable remapping logic only at deterministic transaction boundaries, so that any in-flight transaction is either subject to remapping or not at all.



Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x18, 0x1018		Function: 0	
Bit	Attr	Default	Description
30:30	RW_V	0x0	<p>set_root_table_pointer:</p> <p>Software sets this field to set/update the root-entry table pointer used by hardware. The root-entry table pointer is specified through the Root-entry Table Address register. Hardware reports the status of the root table pointer set operation through the RTPS field in the Global Status register. The root table pointer set operation must be performed before enabling or re-enabling (after disabling) DMA remapping hardware.</p> <p>After a root table pointer set operation, software must globally invalidate the context cache followed by global invalidate of IOTLB. This is required to ensure hardware uses only the remapping structures referenced by the new root table pointer, and not any stale cached entries. While DMA-remapping hardware is active, software may update the root table pointer through this field. However, to ensure valid in-flight DMA requests are deterministically remapped, software must ensure that the structures referenced by the new root table pointer are programmed to provide the same remapping results as the structures referenced by the previous root table pointer.</p> <p>Clearing this bit has no effect.</p>
29:29	RO	0x0	set_fault_log_pointer:
28:28	RO	0x0	enable_advanced_fault_logging:
27:27	RO	0x0	write_buffer_flush:
26:26	RW	0x0	<p>queued_invalidation_enable:</p> <p>Software writes to this field to enable queued invalidations.</p> <p>0: Disable queued invalidations. In this case, invalidations must be performed through the Context Command and IOTLB Invalidation Unit registers.</p> <p>1: Enable use of queued invalidations. Once enabled, all invalidations must be submitted through the invalidation queue and the invalidation registers cannot be used till the translation has been disabled. The invalidation queue address register must be initialized before enabling queued invalidations. Also software must make sure that all invalidations submitted prior via the register interface are all completed before enabling the queued invalidation interface.</p>
25:25	RW	0x0	<p>interrupt_remapping_enable:</p> <p>0: Disable Interrupt Remapping Hardware 1: Enable Interrupt Remapping Hardware</p> <p>Hardware reports the status of the interrupt-remap enable operation through the interrupt_remapping_enable field in the Global Status register.</p> <p>Before enabling (or re-enabling) Interrupt-remapping hardware through this field, software must:</p> <ul style="list-style-type: none"> • Setup the interrupt-remapping structures in memory • Set the Interrupt Remap table pointer in hardware (through IRTP field). • Perform global invalidation of IOTLB <p>There may be active interrupt requests in the platform when software updates this field. Hardware must enable or disable remapping logic only at deterministic transaction boundaries, so that any in-flight interrupts are either subject to remapping or not at all. IIO must drain any in-flight translated DMA read/write, MSI interrupt requests queued within the root complex before completing the translation enable command and reflecting the status of the command through the interrupt_remapping_enable field in the VTD[1:0_]GLBSTS. Value returned on read of this field is undefined.</p>



Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x18, 0x1018		Function: 0	
Bit	Attr	Default	Description
24:24	RW_V	0x0	<p>set_interrupt_remap_table_pointer:</p> <p>Software sets this field to set/update the interrupt remapping table pointer used by hardware. The interrupt remapping table pointer is specified through the Interrupt Remapping Table Address register. Hardware reports the status of the interrupt remapping table pointer set operation through the interrupt_remapping_table_pointer_status field in the Global Status register.</p> <p>The interrupt remap table pointer set operation must be performed before enabling or re-enabling (after disabling) interrupt remapping hardware through the interrupt_remapping_enable field.</p> <p>After an interrupt remap table pointer set operation, software must globally invalidate the interrupt entry cache. This is required to ensure hardware uses only the interrupt remapping entries referenced by the new interrupt remap table pointer, and not any stale cached entries.</p> <p>While interrupt remapping is active, software may update the interrupt remapping table pointer through this field. However, to ensure valid in-flight interrupt requests are deterministically remapped, software must ensure that the structures referenced by the new interrupt remap table pointer are programmed to provide the same remapping results as the structures referenced by the previous interrupt remap table pointer. Clearing this bit has no effect. IIO hardware internally clears this field before the 'set' operation requested by software has take effect.</p>
23:23	RW	0x0	<p>cfi:</p> <p>Compatibility Format Interrupt</p> <p>Software writes to this field to enable or disable Compatibility Format interrupts on Intel® 64 platforms. The value in this field is effective only when interrupt-remapping is enabled and Legacy Interrupt Mode is active.</p> <p>0: Block Compatibility format interrupts.</p> <p>1: Process Compatibility format interrupts as pass-through (bypass interrupt remapping).</p> <p>Hardware reports the status of updating this field through the CFIS field in the vtd[0:1]_glbsts register.</p>

6.7.5 vtd[0:1]_glbsts

Intel VT-d Global Status.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x1c, 0x101c		Function: 0	
Bit	Attr	Default	Description
31:31	RO_V	0x0	<p>translation_enable_status:</p> <p>When set, indicates that translation hardware is enabled and when clear indicates the translation hardware is not enabled.</p>
30:30	RO_V	0x0	<p>set_root_table_pointer_status:</p> <p>This field indicates the status of the root- table pointer in hardware. This field is cleared by hardware when software sets the SRTP field in the Global Command register. This field is set by hardware when hardware finishes the set root-table pointer operation (by performing an implicit global invalidation of the context-cache and IOTLB, and setting/updating the root-table pointer in hardware with the value provided in the Root-Entry Table Address register).</p>
29:29	RO	0x0	<p>set_fault_log_pointer_status:</p>
28:28	RO	0x0	<p>advanced_fault_logging_status:</p>
27:27	RO	0x0	<p>write_buffer_flush_status:</p>



Type: MEM		PortID: 8'h7e		Function: 0	
Bus: 0		Device: 5			
Offset: 0x1c, 0x101c					
Bit	Attr	Default	Description		
26:26	RO_V	0x0	queued_invalidation_interface_status: IIO sets this bit once it has completed the software command to enable the queued invalidation interface. Till then this bit is 0.		
25:25	RO_V	0x0	interrupt_remapping_enable_status: OH sets this bit once it has completed the software command to enable the interrupt remapping interface. Till then this bit is 0.		
24:24	RO_V	0x0	interrupt_remapping_table_pointer_status: This field indicates the status of the interrupt remapping table pointer in hardware. This field is cleared by hardware when software sets the SIRTTP field in the Global Command register. This field is set by hardware when hardware completes the set interrupt remap table pointer operation using the value provided in the Interrupt Remapping Table Address register.		
23:23	RO_V	0x0	cfis: Compatibility Format Interrupt Status The value reported in this field is applicable only when interrupt-remapping is enabled and Legacy interrupt mode is active. 0: Compatibility format interrupts are blocked. 1: Compatibility format interrupts are processed as pass-through (bypassing interrupt remapping).		

6.7.6 vtd[0:1]_rootentryadd

Intel VT-d Root Entry Table Address.

Type: MEM		PortID: 8'h7e		Function: 0	
Bus: 0		Device: 5			
Offset: 0x20, 0x1020					
Bit	Attr	Default	Description		
63:12	RW	0x0	root_entry_table_base_address: 4K aligned base address for the root entry table. Software specifies the base address of the root-entry table through this register, and enables it in hardware through the SRTP field in the Global Command register. Reads of this register returns value that was last programmed to it.		



6.7.7 vtd[0:1]_ctxcmd

Intel VT-d Context Command.

Type: MEM Bus: 0 Offset: 0x28, 0x1028		PortID: 8'h7e Device: 5 Function: 0	
Bit	Attr	Default	Description
63:63	RW_V	0x0	<p>icc: Invalidate Context Entry Cache</p> <p>Software requests invalidation of context-cache by setting this field. Software must also set the requested invalidation granularity by programming the CIRG field. Software must read back and check the ICC field to be clear to confirm the invalidation is complete. Software must not update this register when this field is set. Hardware clears the ICC field to indicate the invalidation request is complete. Hardware also indicates the granularity at which the invalidation operation was performed through the CAIG field. Software must not submit another invalidation request through this register while the ICC field is set. Software must submit a context cache invalidation request through this field only when there are no invalidation requests pending at this DMA-remapping hardware unit. Since information from the context-cache may be used by hardware to tag IOTLB entries, software must perform domain-selective (or global) invalidation of IOTLB after the context cache invalidation has completed.</p>
62:61	RW	0x0	<p>cirg: Context Invalidation Request Granularity</p> <p>When requesting hardware to invalidate the context-entry cache (by setting the ICC field), software writes the requested invalidation granularity through this field. Following are the encoding for the 2-bit IRG field.</p> <p>00: Reserved. Hardware ignores the invalidation request and reports invalidation complete by clearing the ICC field and reporting 00 in the CAIG field.</p> <p>01: Global Invalidation request.</p> <p>10: Domain-selective invalidation request. The target domain-id must be specified in the DID field.</p> <p>11: Device-selective invalidation request. The target SID must be specified in the SID field, and the domain-id (programmed in the context-entry for this device) must be provided in the DID field. The processor aliases the h/w behavior for this command to the 'Domain-selective invalidation request'.</p> <p>Hardware indicates completion of the invalidation request by clearing the ICC field. At this time, hardware also indicates the granularity at which the actual invalidation was performed through the CAIG field.</p>
60:59	RO_V	0x0	<p>caig: Context Actual Invalidation Granularity</p> <p>Hardware reports the granularity at which an invalidation request was processed through the CAIG field at the time of reporting invalidation completion (by clearing the ICC field). The following are the encoding for the 2-bit CAIG field.</p> <p>00: Reserved. This is the value on reset.</p> <p>01: Global Invalidation performed. The processor sets this in response to a global invalidation request.</p> <p>10: Domain-selective invalidation performed using the domain-id that was specified by software in the DID field. The processor set this in response to a domain-selective or device-selective invalidation request.</p> <p>11: Device-selective invalidation. The processor never sets this encoding.</p>
33:32	RW	0x0	<p>fm: Function Mask</p> <p>Used by the processor when performing device selective invalidation.</p>
31:16	RW	0x0	<p>source_id: Used by the processor when performing device selective context cache invalidation</p>



Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x28, 0x1028		Function: 0	
Bit	Attr	Default	Description
15:0	RW	0x0	domain_id: Indicates the id of the domain whose context-entries needs to be selectively invalidated. S/W needs to program this for both domain and device selective invalidates. The processor ignores bits 15:8 since it supports only a 8 bit Domain ID.

6.7.8 vtd[0:1]_fltsts

Intel VT-d Fault Status.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x34, 0x1034		Function: 0	
Bit	Attr	Default	Description
15:8	ROS_V	0x0	fault_record_index: This field is valid only when the Primary Fault Pending field is set. This field indicates the index (from base) of the fault recording register to which the first pending fault was recorded when the Primary Fault pending field was set by hardware.
6:6	RW1CS	0x0	invalidation_timeout_error: Hardware detected a Device-IOTLB invalidation completion time-out. At this time, a fault event may be generated based on the programming of the Fault Event Control register.
5:5	RW1CS	0x0	invalidation_completion_error: Hardware received an unexpected or invalid Device-IOTLB invalidation completion. At this time, a fault event is generated based on the programming of the Fault Event Control register.
4:4	RW1CS	0x0	invalidation_queue_error: Hardware detected an error associated with the invalidation queue. For example, hardware detected an erroneous or un-supported Invalidation Descriptor in the Invalidation Queue. At this time, a fault event is generated based on the programming of the Fault Event Control register.
1:1	ROS_V	0x0	primary_fault_pending: This field indicates if there are one or more pending faults logged in the fault recording registers. Hardware computes this field as the logical OR of Fault (F) fields across all the fault recording registers of this DMA-remap hardware unit. 0: No pending faults in any of the fault recording registers 1: One or more fault recording registers has pending faults. The fault recording index field is updated by hardware whenever this field is set by hardware. Also, depending on the programming of fault event control register, a fault event is generated when hardware sets this field.
0:0	RW1CS	0x0	primary_fault_overflow: Hardware sets this bit to indicate overflow of fault recording registers



6.7.9 nonisoch_fltevtctrl


Fault Event Control.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x38		Function: 0	
Bit	Attr	Default	Description
31:31	RW	0x1	<p>fault_nonisoch_msgmsk:</p> <p>1: Hardware is prohibited from issuing interrupt message requests. 0: Software has cleared this bit to indicate interrupt service is available. When a faulting condition is detected, hardware may issue a interrupt request (using the fault event data and fault event address register values) depending on the state of the interrupt mask and interrupt pending bits.</p>
30:30	RO_V	0x0	<p>fault_nonisoch_msi_pend:</p> <p>Hardware sets the IP field whenever it detects an interrupt condition. Interrupt condition is defined as when an interrupt condition occurs when hardware records a fault through one of the Fault Recording registers and sets the PPF field in Fault Status register. - Hardware detected error associated with the Invalidation Queue, setting the IQE field in the Fault Status register.</p> <ul style="list-style-type: none"> - Hardware detected invalidation completion timeout error, setting the ICT field in the Fault Status register. - If any of the above status fields in the Fault Status register was already set at the time of setting any of these fields, it is not treated as a new interrupt condition. <p>The IP field is kept set by hardware while the interrupt message is held pending. The interrupt message could be held pending due to interrupt mask (IM field) being set, or due to other transient hardware conditions.</p> <p>The IP field is cleared by hardware as soon as the interrupt message pending condition is serviced. This could be due to either</p> <ul style="list-style-type: none"> (a) Hardware issuing the interrupt message due to either change in the transient hardware condition that caused interrupt message to be held pending or due to software clearing the IM field. (b) Software servicing all the pending interrupt status fields in the Fault Status register. <ul style="list-style-type: none"> • PPF field is cleared by hardware when it detects all the Fault Recording registers have Fault (F) field clear. • Other status fields in the Fault Status register is cleared by software writing back the value read from the respective fields.
29:0	RO	0x0	fault_nonisoch_msgmsk_const:

6.7.10 nonisoch_fltevtdata

Fault Event Data.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x3c		Function: 0	
Bit	Attr	Default	Description
31:16	RO	0x0	fault_nonisoch_data_const:
15:0	RW	0x0	fault_nonisoch_data:



6.7.11 vtd[0:1]_fltevtaddr

Intel VT-d Fault Event Address.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x40, 0x1040		Function: 0	
Bit	Attr	Default	Description
31:2	RW	0x0	interrupt_address: The interrupt address is interpreted as the address of any other interrupt from a PCI Express port.

6.7.12 vtd[0:1]_fltevtupraddr

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x44, 0x1044		Function: 0	
Bit	Attr	Default	Description
31:0	RW	0x0	address:

6.7.13 vtd[0:1]_pmen

Intel VT-d Protect Memory Enable.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x64, 0x1064		Function: 0	
Bit	Attr	Default	Description
31:31	RW	0x0	protmemen: Enable Protected Memory PROT_LOW_BASE/LIMIT and PROT_HIGH_BASE/LIMIT memory regions. Software can use the protected low/high address ranges to protect both the DMA remapping tables and the interrupt remapping tables. There is no separate set of registers provided for each.
0:0	RO_V	0x0	protregionsts: This bit is set by the processor whenever it has completed enabling the protected memory region per the rules stated in the Intel VT-d spec



6.7.14 vtd[0:1]_prot_low_mem_base

Intel VT-d Protected Memory Low Base.

Type:	MEM	PortID:	8'h7e	Function:	0
Bus:	0	Device:	5		
Offset:	0x68, 0x1068				
Bit	Attr	Default	Description		
31:21	RW	0x0	addr: 16 MB aligned base address of the low protected DRAM region Note that Intel VT-d engine generated reads/writes (page walk, interrupt queue, invalidation queue read, invalidation status) themselves are allowed toward this region, but no DMA accesses (non-translated DMA or ATS translated DMA or pass through DMA, that is, no DMA access of any kind) from any device is allowed toward this region (regardless of whether TE is 0 or 1), when enabled.		

6.7.15 vtd[0:1]_prot_low_mem_limit


Intel VT-d Protected Memory Low Limit.

Type:	MEM	PortID:	8'h7e	Function:	0
Bus:	0	Device:	5		
Offset:	0x6c, 0x106c				
Bit	Attr	Default	Description		
31:21	RW	0x0	addr: 16 MB aligned limit address of the low protected DRAM region Note that Intel VT-d engine generated reads/writes (page walk, interrupt queue, invalidation queue read, invalidation status) themselves are allowed toward this region, but no DMA accesses (non-translated DMA or ATS translated DMA or pass through DMA, that is, no DMA access of any kind) from any device is allowed toward this region (regardless of whether TE is 0 or 1) when enabled.		

6.7.16 vtd[0:1]_prot_high_mem_base

Intel VT-d Protected Memory High Base.

Type:	MEM	PortID:	8'h7e	Function:	0
Bus:	0	Device:	5		
Offset:	0x70, 0x1070				
Bit	Attr	Default	Description		
63:21	RW	0x0	addr: 16 MB aligned base address of the high protected DRAM region Note that Intel VT-d engine generated reads/writes (page walk, interrupt queue, invalidation queue read, invalidation status) themselves are allowed toward this region, but no DMA accesses (non-translated DMA or ATS translated DMA or pass through DMA, that is, no DMA access of any kind) from any device is allowed toward this region (regardless of whether TE is 0 or 1) when enabled.		



6.7.17 vtd[0:1]_prot_high_mem_limit

Intel VT-d Protected Memory High Limit.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x78, 0x1078		Function: 0	
Bit	Attr	Default	Description
63:21	RW	0x0	<p>addr:</p> <p>16 MB aligned limit address of the high protected DRAM region</p> <p>Note that Intel VT-d engine generated reads/writes (page walk, interrupt queue, invalidation queue read, invalidation status) themselves are allowed toward this region, but no DMA accesses (non-translated DMA or ATS translated DMA or pass through DMA, that is, no DMA access of any kind) from any device is allowed toward this region (regardless of whether TE is 0 or 1), when enabled.</p>

6.7.18 vtd[0:1]_inv_queue_head

Intel VT-d Invalidation Queue Header Pointer.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x80, 0x1080		Function: 0	
Bit	Attr	Default	Description
18:4	RO_V	0x0	<p>queue_head:</p> <p>Specifies the offset (128-bit aligned) to the invalidation queue for the command that will be fetched next by hardware. This field is incremented after the command has been fetched successfully and has been verified to be a valid/supported command.</p>

6.7.19 vtd[0:1]_inv_queue_tail

Intel VT-d Invalidation Queue Tail Pointer.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x88, 0x1088		Function: 0	
Bit	Attr	Default	Description
18:4	RW	0x0	<p>queue_tail:</p> <p>Specifies the offset (128-bit aligned) to the invalidation queue for the command that will be written next by software.</p>



6.7.20 vtd[0:1]_inv_queue_add

Intel VT-d Invalidation Queue Address.

Type:	MEM	PortID:	8'h7e	Function:	0
Bus:	0	Device:	5		
Offset:	0x90, 0x1090				
Bit	Attr	Default	Description		
63:12	RW	0x0	invreq_queue_base_address: This field points to the base of size-aligned invalidation request queue.		
2:0	RW	0x0	queue_size: This field specifies the length of the invalidation request queue. The number of entries in the invalidation queue is defined as $2^{(X + 8)}$, where X is the value programmed in this field.		

6.7.21 vtd[0:1]_inv_comp_status

Intel VT-d Invalidation Completion Status.

Type:	MEM	PortID:	8'h7e	Function:	0
Bus:	0	Device:	5		
Offset:	0x9c, 0x109c				
Bit	Attr	Default	Description		
0:0	RW1CS	0x0	invalidation_wait_descriptor_complete: Indicates completion of Invalidation Wait Descriptor with Interrupt Flag (IF) field set. Hardware clears this field whenever it is executing a wait descriptor with IF field set and sets this bit when the descriptor is complete.		

6.7.22 nonisoch_inv_cmp_evtctrl

Invalidation Completion Event Control.

Type:	MEM	PortID:	8'h7e	Function:	0
Bus:	0	Device:	5		
Offset:	0xa0				
Bit	Attr	Default	Description		
31:31	RW	0x1	inval_nonisoch_msgmsk: 0: No masking of interrupt. When a invalidation event condition is detected, hardware issues an interrupt message (using the Invalidation Event Data & Invalidation Event Address register values). 1: This is the value on reset. Software may mask interrupt message generation by setting this field. Hardware is prohibited from sending the interrupt message when this field is set.		



Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0xa0		Function: 0	
Bit	Attr	Default	Description
30:30	RO_V	0x0	inval_nonisoch_msi_pend: Hardware sets the IP field whenever it detects an interrupt condition. Interrupt condition is defined as: - An Invalidation Wait Descriptor with Interrupt Flag (IF) field set completed, setting the IWC field in the Fault Status register. - If the IWC field in the Invalidation Event Status register was already set at the time of setting this field, it is not treated as a new interrupt condition. The IP field is kept set by hardware while the interrupt message is held pending. The interrupt message could be held pending due to interrupt mask (IM field) being set, or due to other transient hardware conditions. The IP field is cleared by hardware as soon as the interrupt message pending condition is serviced. This could be due to either: (a) Hardware issuing the interrupt message due to either change in the transient hardware condition that caused interrupt message to be held pending or due to software clearing the IM field. (b) Software servicing the IWC field in the Fault Status register.
29:0	RO	0x0	inval_nonisoch_msgmsk_const:

6.7.23 nonisoch_invevtdata

Invalidation Event Data.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0xa4		Function: 0	
Bit	Attr	Default	Description
31:16	RO	0x0	inval_nonisoch_data_const:
15:0	RW	0x0	inval_nonisoch_data:

6.7.24 vtd[0:1]_inv_comp_evt_addr

Intel VT-d Invalidation Completion Event Address.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0xa8, 0x10a8		Function: 0	
Bit	Attr	Default	Description
31:2	RW	0x0	interrupt_address:

6.7.25 vtd[0:1]_inv_comp_evt_upraddr

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0xac, 0x10ac		Function: 0	
Bit	Attr	Default	Description
31:0	RW	0x0	address:



6.7.26 vtd[0:1]_intr_remap_table_base


Intel VT-d Interrupt Remapping Table Based Address.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0xb8, 0x10b8		Function: 0	
Bit	Attr	Default	Description
63:12	RW	0x0	intr_remap_base: This field points to the base of page-aligned interrupt remapping table. If the Interrupt Remapping Table is larger than 4 KB in size, it must be size-aligned. Reads of this field returns value that was last programmed to it.
11:11	RW	0x0	ia32_extended_interrupt_enable: 0: IA32 system is operating in legacy IA32 interrupt mode. Hardware interprets only 8-bit APICID in the Interrupt Remapping Table entries. 1: IA32 system is operating in extended IA32 interrupt mode. Hardware interprets 32-bit APICID in the Interrupt Remapping Table entries.
3:0	RW	0x0	size: This field specifies the size of the interrupt remapping table. The number of entries in the interrupt remapping table is $2^{(X+1)}$, where X is the value programmed in this field.

6.7.27 vtd0_fltrec[0:7]_gpa, vtd1_fltrec0_gpa

Intel VT-d Fault Record.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: vtd0: 0x110, 0x120, 0x130, 0x140, 0x150, 0x160, 0x170 vtd1: 0x1100		Function: 0	
Bit	Attr	Default	Description
63:12	ROS_V	0x0	gpa: 4K aligned GPA for the faulting transaction. valid only when F field is set.



6.7.28 vtd0_filtrec[0:7]_src, vtd1_filtrec0_src

Intel VT-d Fault Record.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: vtd0: 0x108, 0x118, 0x128, 0x138, 0x148, 0x158, 0x168, 0x178		Function: 0	
vtd1: 0x1108			
Bit	Attr	Default	Description
63:63	RW1CS	0x0	f: Fault. Hardware sets this field to indicate a fault is logged in this fault recording register. The F field is set by hardware after the details of the fault is recorded in the PADDR, SID, FR and T fields. When this field is set, hardware may collapse additional faults from the same requestor (SID). Software writes the value read from this field to clear it.
62:62	ROS_V	0x0	type: Type of the first faulted DMA request 0: DMA write 1: DMA read request This field is only valid when Fault (F) bit is set.
61:60	ROS_V	0x0	address_type: This field captures the AT field from the faulted DMA request. This field is valid only when the F field is set.
39:32	ROS_V	0x0	fault_reason: Reason for the first translation fault. See Intel VT-d spec for details. This field is only valid when Fault bit is set.
15:0	ROS_V	0x0	source_identifier: Requester ID of the dma request that faulted. Valid only when F bit is set

6.7.29 vtd[0:1]_invaddrreg

Intel VT-d Invalidate Address.

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x200, 0x1200		Function: 0	
Bit	Attr	Default	Description
63:12	RW	0x0	addr: To request a page-specific invalidation request to hardware, software must first write the corresponding guest physical address to this register, and then issue a page-specific invalidate command through the IOTLB_REG.
6:6	RW	0x0	ih: The field provides hint to hardware to preserve or flush the respective non-leaf page-table entries that may be cached in hardware. 0: Software may have modified both leaf and non-leaf page-table entries corresponding to mappings specified in the ADDR and AM fields. On a page-selective invalidation request, IIO must flush both the cached leaf and nonleaf page-table entries corresponding to mappings specified by ADDR and AM fields. IIO performs a domain-level invalidation on non-leaf entries and page-selective-domain-level invalidation at the leaf level. 1: Software has not modified any non-leaf page-table entries corresponding to mappings specified in the ADDR and AM fields. On a page-selective invalidation request, IIO preserves the cached non-leaf page-table entries corresponding to mappings specified by ADDR and AM fields and performs only a page-selective invalidation at the leaf level.
5:0	RW	0x0	am: IIO supports values of 0-9. All other values result in undefined results.



6.7.30 vtd[0:1]_iотlbinv

Intel VT-d IOTLB Invalidate.

Type: MEM Bus: 0 Offset: 0x208, 0x1208		PortID: 8'h7e Device: 5 Function: 0	
Bit	Attr	Default	Description
63:63	RW_V	0x0	<p>Intel VT: Invalidate IOTLB cache Software requests IOTLB invalidation by setting this field. Software must also set the requested invalidation granularity by programming the IIRG field. Hardware clears the Intel VT field to indicate the invalidation request is complete. Hardware also indicates the granularity at which the invalidation operation was performed through the IAIG field. Software must read back and check the CPU field to be clear to confirm the invalidation is complete. When CPU field is set, software must not update the contents of this register (and Invalidate Address register, if it is being used), nor submit new IOTLB invalidation requests.</p>
61:60	RW	0x0	<p>iirg: IOTLB Invalidation Request Granularity When requesting hardware to invalidate the I/OTLB (by setting the Intel VT field), software writes the requested invalidation granularity through this IIRG field. Following are the encoding for the 2-bit IIRG field. 00: Reserved. Hardware ignores the invalidation request and reports invalidation complete by clearing the Intel VT field and reporting 00 in the AIG field. 01: Global Invalidation request. 10: Domain-selective invalidation request. The target domain-id must be specified in the DID field. 11: Page-selective invalidation request. The target address, mask and invalidation hint must be specified in the Invalidate Address register, the domain-id must be provided in the DID field.</p>
58:57	RO_V	0x0	<p>iaig: IOTLB Actual Invalidation Granularity Hardware reports the granularity at which an invalidation request was proceed through the AIG field at the time of reporting invalidation completion (by clearing the Intel VT field). The following are the encoding for the 2-bit IAIG field. 00: Reserved. This indicates hardware detected an incorrect invalidation request and ignored the request. Examples of incorrect invalidation requests include detecting an unsupported address mask value in Invalidate Address register for page-selective invalidation requests or an unsupported/undefined encoding in IIRG. 01: Global Invalidation performed. The processor sets this in response to a global IOTLB invalidation request. 10: Domain-selective invalidation performed using the domain-id that was specified by software in the DID field. The processor sets this in response to a domain selective IOTLB invalidation request. 11: CPU sets this in response to a page selective invalidation request.</p>
49:49	RW	0x0	<p>dr: CPU uses this to drain or not drain reads on an invalidation request.</p>
48:48	RW	0x0	<p>dw: CPU uses this to drain or not drain reads on an invalidation request.</p>
47:32	RW	0x0	<p>did: Domain to be invalidated and is programmed by software for both page and domain selective invalidation requests. CPU ignores the bits 47:40 since it supports only an 8 bit Domain ID.</p>



6.8 Device 5 Function 2

Global System Control and Error Registers.

Register Name	Offset	Size
vid	0x0	16
did	0x2	16
pcicmd	0x4	16
pcists	0x6	16
rid	0x8	8
ccr	0x9	24
clsr	0xc	8
hdr	0xe	8
svid	0x2c	16
sdid	0x2e	16
capptr	0x34	8
intl	0x3c	8
intpin	0x3d	8
pxpcapid	0x40	8
pxpnxtptr	0x41	8
pxpcap	0x42	16
irpperrsv	0x80	64
ioerrsv	0x8c	32
mierrsv	0x90	32
pcierrsv	0x94	32
sysmap	0x9c	32
vppctl	0xb0	64
vppsts	0xb8	32
vppfreq	0xbc	32
gcerrst	0x1a8	32
gcferrst	0x1ac	32
gcerrst	0x1b8	32
gnerrst	0x1c0	32
gferrst	0x1c4	32
gerrctl	0x1c8	32
gsysst	0x1cc	32
gsysctl	0x1d0	32
gfferrst	0x1dc	32
gfnerrst	0x1e8	32
gnferrst	0x1ec	32
gnerrst	0x1f8	32
irpp0errst	0x230	32
irpp0errctl	0x234	32
irpp0ferrst	0x238	32



Register Name	Offset	Size
irpp0fnerrst	0x23c	32
irpp0fferrhd0	0x240	32
irpp0fferrhd1	0x244	32
irpp0fferrhd2	0x248	32
irpp0fferrhd3	0x24c	32
irpp0nferrst	0x250	32
irpp0nnerrst	0x254	32
irpp0nferrhd0	0x258	32
irpp0nferrhd1	0x25c	32
irpp0nferrhd2	0x260	32
irpp0nferrhd3	0x264	32
irpp0errcntsel	0x268	32
irpp0errcnt	0x26c	32
irpp1errst	0x2b0	32
irpp1errctl	0x2b4	32
irpp1fferrst	0x2b8	32
irpp1fnerrst	0x2bc	32
irpp1fferrhd0	0x2c0	32
irpp1fferrhd1	0x2c4	32
irpp1fferrhd2	0x2c8	32
irpp1fferrhd3	0x2cc	32
irpp1nferrst	0x2d0	32
irpp1nnerrst	0x2d4	32
irpp1nferrhd0	0x2d8	32
irpp1nferrhd1	0x2dc	32
irpp1nferrhd2	0x2e0	32
irpp1nferrhd3	0x2e4	32
irpp1errcntsel	0x2e8	32
irpp1errcnt	0x2ec	32
iioerrst	0x300	32
iioerrctl	0x304	32
iiofferrst	0x308	32
iiofferrhd_0	0x30c	32
iiofferrhd_1	0x310	32
iiofferrhd_2	0x314	32
iiofferrhd_3	0x318	32
iiofnerrst	0x31c	32
iionferrst	0x320	32
iionferrhd_0	0x324	32
iionferrhd_1	0x328	32
iionferrhd_2	0x32c	32
iionferrhd_3	0x330	32
iionnerrst	0x334	32



Register Name	Offset	Size
ioerrcntsel	0x33c	32
ioerrcnt	0x340	32
mierrst	0x380	32
mierrctl	0x384	32
mifferrst	0x388	32
mifferrhdr_0	0x38c	32
mifferrhdr_1	0x390	32
mifferrhdr_2	0x394	32
mifferrhdr_3	0x398	32
mifnerrst	0x39c	32
minferrst	0x3a0	32
minferrhdr_0	0x3a4	32
minferrhdr_1	0x3a8	32
minferrhdr_2	0x3ac	32
minferrhdr_3	0x3b0	32
minnerrst	0x3b4	32
mierrcntsel	0x3bc	32
mierrcnt	0x3c0	8

6.8.1 vid

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x0		Function: 2	
Bit	Attr	Default	Description
15:0	RO	0x8086	vendor_identification_number: The value is assigned by PCI-SIG to Intel.

6.8.2 did

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x2		Function: 2	
Bit	Attr	Default	Description
15:0	RO	0x2f2a	device_identification_number: Device ID values vary from function to function.



6.8.3 pcicmd

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x4		Function: 2	
Bit	Attr	Default	Description
10:10	RO	0x0	intx_disable: NA for these devices
9:9	RO	0x0	fast_back_to_back_enable: Not applicable to PCI Express and is hardwired to 0
8:8	RO	0x0	serr_enable: This bit has no impact on error reporting from these devices
7:7	RO	0x0	idsel_stepping_wait_cycle_control: Not applicable to internal devices. Hardwired to 0.
6:6	RO	0x0	parity_error_response: This bit has no impact on error reporting from these devices
5:5	RO	0x0	vga_palette_snoop_enable: Not applicable to internal devices. Hardwired to 0.
4:4	RO	0x0	memory_write_and_invalidate_enable: Not applicable to internal devices. Hardwired to 0.
3:3	RO	0x0	special_cycle_enable: Not applicable. Hardwired to 0.
2:2	RO	0x0	bus_master_enable: Hardwired to 0 since these devices don't generate any transactions
1:1	RO	0x0	memory_space_enable: Hardwired to 0 since these devices don't decode any memory BARs
0:0	RO	0x0	io_space_enable: Hardwired to 0 since these devices don't decode any IO BARs

6.8.4 pcists

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x6		Function: 2	
Bit	Attr	Default	Description
15:15	RO	0x0	detected_parity_error: This bit is set when the device receives a packet on the primary side with an uncorrectable data error including a packet with poison bit set or an uncorrectable addresscontrol parity error. The setting of this bit is regardless of the Parity Error Response bit PERRE in the PCICMD register.
14:14	RO	0x0	signaled_system_error: Hardwired to 0
13:13	RO	0x0	received_master_abort: Hardwired to 0
12:12	RO	0x0	received_target_abort: Hardwired to 0
11:11	RO	0x0	signaled_target_abort: Hardwired to 0
10:9	RO	0x0	devsel_timing: Not applicable to PCI Express. Hardwired to 0.



Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x6		Function: 2	
Bit	Attr	Default	Description
8:8	RO	0x0	master_data_parity_error: Hardwired to 0
7:7	RO	0x0	fast_back_to_back: Not applicable to PCI Express. Hardwired to 0.
5:5	RO	0x0	pci66mhz_capable: Not applicable to PCI Express. Hardwired to 0.
4:4	RO	0x1	capabilities_list: This bit indicates the presence of a capabilities list structure
3:3	RO	0x0	intx_status: Hardwired to 0

6.8.5 rid

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x8		Function: 2	
Bit	Attr	Default	Description
7:0	RO_V	0x0	revision_id: Reflects the Uncore Revision ID after reset. Reflects the Compatibility Revision ID after BIOS writes 0x69 to any RID register in any Intel® Xeon® Processor E5 v4 product family function.

6.8.6 ccr

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x9		Function: 2	
Bit	Attr	Default	Description
23:16	RO_V	0x8	base_class: Generic Device
15:8	RO_V	0x80	sub_class: Generic Device
7:0	RO_V	0x0	register_level_programming_interface: Set to 00h for all non-APIC devices.



6.8.7 clsr

Type:	CFG	PortID:	N/A
Bus:	0	Device:	5
Offset:	0xc	Function:	2
Bit	Attr	Default	Description
7:0	RW	0x0	cacheline_size: This register is set as RW for compatibility reasons only. Cacheline size is always 64B.

6.8.8 hdr


Type:	CFG	PortID:	N/A
Bus:	0	Device:	5
Offset:	0xe	Function:	2
Bit	Attr	Default	Description
7:7	RO	0x1	multi_function_device: This bit defaults to 1b since all these devices are multi-function.
6:0	RO	0x0	configuration_layout: This field identifies the format of the configuration header layout. It is Type 0 for all these devices. The default is 00h, indicating a 'endpoint device'.

6.8.9 svid

Type:	CFG	PortID:	N/A
Bus:	0	Device:	5
Offset:	0x2c	Function:	2
Bit	Attr	Default	Description
15:0	RW_O	0x0	subsystem_vendor_identification_number: The default value specifies Intel but can be set to any value once after reset.

6.8.10 sdid

Type:	CFG	PortID:	N/A
Bus:	0	Device:	5
Offset:	0x2e	Function:	2
Bit	Attr	Default	Description
15:0	RW_O	0x0	subsystem_device_identification_number: Assigned by the subsystem vendor to uniquely identify the subsystem.



6.8.11 capptr

Type: CFG Bus: 0 Offset: 0x34		PortID: N/A Device: 5 Function: 2	
Bit	Attr	Default	Description
7:0	RO	0x40	capability_pointer: Points to the first capability structure for the device which is the PCIe capability.

6.8.12 intl

Type: CFG Bus: 0 Offset: 0x3c		PortID: N/A Device: 5 Function: 2	
Bit	Attr	Default	Description
7:0	RO	0x0	interrupt_line: NA for these devices

6.8.13 intpin

Type: CFG Bus: 0 Offset: 0x3d		PortID: N/A Device: 5 Function: 2	
Bit	Attr	Default	Description
7:0	RO	0x0	interrupt_pin: NA since these devices do not generate any interrupt on their own.

6.8.14 pxpcapid

Type: CFG Bus: 0 Offset: 0x40		PortID: N/A Device: 5 Function: 2	
Bit	Attr	Default	Description
7:0	RO	0x10	capability_id: Provides the PCI Express capability ID assigned by PCI-SIG.

6.8.15 pxpnxtptr

Type: CFG Bus: 0 Offset: 0x41		PortID: N/A Device: 5 Function: 2	
Bit	Attr	Default	Description
7:0	RO	0x0	next_ptr: This field is set to the PCI Power Management capability.




6.8.16 pxpcap

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x42		Function: 2	
Bit	Attr	Default	Description
13:9	RO	0x0	interrupt_message_number_n_a:
8:8	RO	0x0	slot_implemented_n_a:
7:4	RO	0x9	device_port_type: This field identifies the type of device. It is set to for the DMA to indicate root complex integrated endpoint device.
3:0	RO	0x2	capability_version: This field identifies the version of the PCI Express capability structure. Set to 2h for PCI Express and DMA devices for compliance with the extended base registers.

6.8.17 irpperrsv

IRP Protocol Error Severity.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x80		Function: 2	
Bit	Attr	Default	Description
29:28	RWS	0x2	protocol_parity_error: (DB) 00: Error Severity Level 0 (Correctable) 01: Error Severity Level 1 (Recoverable) 10: Error Severity Level 2 (Fatal) 11: Reserved
27:26	RWS	0x2	protocol_qt_overflow_underflow: (DA) 00: Error Severity Level 0 (Correctable) 01: Error Severity Level 1 (Recoverable) 10: Error Severity Level 2 (Fatal) 11: Reserved
21:20	RWS	0x2	protocol_rcvd_unexprsp: (D7) 00: Error Severity Level 0 (Correctable) 01: Error Severity Level 1 (Recoverable) 10: Error Severity Level 2 (Fatal) 11: Reserved
9:8	RWS	0x1	csr_acc_32b_unaligned: (C3) 00: Error Severity Level 0 (Correctable) 01: Error Severity Level 1 (Recoverable) 10: Error Severity Level 2 (Fatal) 11: Reserved
7:6	RWS	0x1	wrcache_uncecc_error: (C2) 00: Error Severity Level 0 (Correctable) 01: Error Severity Level 1 (Recoverable) 10: Error Severity Level 2 (Fatal) 11: Reserved
5:4	RWS	0x1	protocol_rcvd_poison: (C1) 00: Error Severity Level 0 (Correctable) 01: Error Severity Level 1 (Recoverable) 10: Error Severity Level 2 (Fatal) 11: Reserved



Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x80		Function: 2	
Bit	Attr	Default	Description
3:2	RWS	0x0	wrcache_correcc_error: (B4) 00: Error Severity Level 0 (Correctable) 01: Error Severity Level 1 (Recoverable) 10: Error Severity Level 2 (Fatal) 11: Reserved

6.8.18 iioerrsv

IIO Core Error Severity.

This register associates the detected IIO internal core errors to an error severity level. An individual error is reported with the corresponding severity in this register. Software can program the error severity to one of the three severities supported by IIO. This register is sticky and can only be reset by PWRGOOD.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x8c		Function: 2	
Bit	Attr	Default	Description
13:12	RWS_L	0x1	c6_overflow_underflow_error: 00: Error Severity Level 0 (Correctable) 01: Error Severity Level 1 (Recoverable) 10: Error Severity Level 2 (Fatal) 11: Reserved
9:8	RWS_L	0x1	c4_master_abort_address_error: 00: Error Severity Level 0 (Correctable) 01: Error Severity Level 1 (Recoverable) 10: Error Severity Level 2 (Fatal) 11: Reserved
1:0	RWS_L	0x0	c7_multicast_target_error: Multicast target error, indicating a MCAST transaction has targeted more than the number of groups supported. 00: Error Severity Level 0 (Correctable) 01: Error Severity Level 1 (Recoverable) 10: Error Severity Level 2 (Fatal) 11: Reserved



6.8.19 mierrsv

Miscellaneous Error Severity.


Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x90		Function: 2	
Bit	Attr	Default	Description
7:6	RWS	0x0	vpp_err_sts: 00: Error Severity Level 0 (Correctable) 01: Error Severity Level 1 (Recoverable) 10: Error Severity Level 2 (Fatal) 11: Reserved This bit should be programmed to 1.

6.8.20 pcierrsv

PCIe Error Severity Map.

This register allows remapping of the PCIe errors to the IIO error severity.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x94		Function: 2	
Bit	Attr	Default	Description
5:4	RWS	0x2	pciefaterr_map: 10: Map this PCIe error type to Error Severity 2 01: Map this PCIe error type to Error Severity 1 00: Map this PCIe error type to Error Severity 0
3:2	RWS	0x1	pcienonfaterr_map: 10: Map this PCIe error type to Error Severity 2 01: Map this PCIe error type to Error Severity 1 00: Map this PCIe error type to Error Severity 0
1:0	RWS	0x0	pciecorerr_map: 10: Map this PCIe error type to Error Severity 2 01: Map this PCIe error type to Error Severity 1 00: Map this PCIe error type to Error Severity 0



6.8.21 sysmap

System Error Event map.

This register maps the error severity detected by the IIO to one of the system events. When an error is detected by the IIO, its corresponding error severity determines which system event to generate according to this register.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x9c		Function: 2	
Bit	Attr	Default	Description
10:8	RWS	0x1	sev2_map: 010: Generate NMI 001: Generate Intel SMI/PMI 000: No inband message Others: Reserved
6:4	RWS	0x2	sev1_map: 010: Generate NMI 001: Generate SMIPMI 000: No inband message Others: Reserved
2:0	RWS	0x0	sev0_map: 010: Generate NMI 001: Generate SMIPMI 000: No inband message Others: Reserved

6.8.22 vppctl

This register defines the control/command for PCA9555.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xb0		Function: 2	
Bit	Attr	Default	Description
63:60	RO	0x1	vpp_version: Specified the version of this structure for BIOS use. 0: VPPCTL with PCIe ports.
55:55	RWS	0x0	vpp_reset_mode: 0: Power good reset will reset the VPP state machines and hard reset will cause the VPP state machine to terminate at the next 'logical' VPP stream boundary and then reset the VPP state machines 1: Both power good and hard reset will reset the VPP state machines



Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xb0		Function: 2	
Bit	Attr	Default	Description
54:44	RWS	0x0	<p>vpp_en:</p> <p>When set, the VPP function for the corresponding root port is enabled.</p> <p>Enable Root Port</p> <p>[54] Port 3d</p> <p>[53] Port 3c</p> <p>[52] Port 3b</p> <p>[51] Port 3a</p> <p>[50] Port 2d</p> <p>[49] Port 2c</p> <p>[48] Port 2b</p> <p>[47] Port 2a</p> <p>[46] Port 1b</p> <p>[45] Port 1a</p> <p>[44] Port 0 (PCIe mode only)</p>
43:0	RWS	0x0	<p>vpp_enaddr:</p> <p>Assigns the VPP address of the device on the VPP interface and assigns the port address for the ports within the VPP device. There are more address bits than root ports so assignment must be spread across VPP ports.</p> <p>Port Addr Root Port</p> <p>[40] [43:41] Port 3d</p> <p>[36] [39:37] Port 3c</p> <p>[32] [35:33] Port 3b</p> <p>[28] [31:29] Port 3a</p> <p>[24] [27:25] Port 2d</p> <p>[20] [23:21] Port 2c</p> <p>[16] [19:17] Port 2b</p> <p>[12] [15:13] Port 2a</p> <p>[8] [11:9] Port 1a</p> <p>[4] [7:5] Port 1a</p> <p>[0] [3:1] Port 0 (PCIe mode only)</p>

6.8.23 vppsts

This register defines the status from PCA9555

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xb8		Function: 2	
Bit	Attr	Default	Description
0:0	RW1CS	0x0	<p>vpp_error:</p> <p>VPP Port error happened i.e. an unexpected STOP of NACK was seen on the VPP port</p>



6.8.24 vppfreq

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xbc		Function: 2	
Bit	Attr	Default	Description
31:24	RWS	0x1e	vpp_tpf: Pulse Filter should be set to 60 nS. The value used is dependent on the internal clock frequency. In this case, internal clock frequency is 500 MHz, so the default value represents 60 nS at that rate.
23:16	RWS	0x96	vpp_thd_data: Hold time for Data is 300 nS. The default value is set to 300 nS when the internal clock rate is 500 MHz.
11:0	RWS	0x9c4	vpp_tsu_thd: Represents the high time and low time of the SCL pin. It should be set to 5 uS for a 100 kHz SCL clock 5 uS high time and 5 uS low time. The default value represents 5 uS with an internal clock of 500 MHz.

6.8.25 gcerrst

This register indicates the corrected error reported to the IIO global error logic. An individual error status bit that is set indicates that a particular local interface has detected an error.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x1a8		Function: 2	
Bit	Attr	Default	Description
26:26	RV	0x0	MC error Memory Controller Error Status.
25:25	RW	0b	Intel VT-d Error This register indicates the corrected error reported to the Intel VT-d error logic. An individual error status bit that is set indicates that a particular local interface has detected an error.
24:24	RW	0b	Miscellaneous Error
23:23	RW	0b	IIO Core Error
20:20	RW	0b	DMI Error
15:5	RW	0x0	PCIe* Error Bit 5: Port 0 Bit 6: Port 1a Bit 7: Port 1b Bit 8: Port 2a Bit 9: Port 2b Bit 10: Port 2c Bit 11: Port 2d Bit 12: Port 3a Bit 13: Port 3b Bit 14: Port 3c Bit 15: Port 3d
1:1	RW	0x0	IRP1 Error Mask
0:0	RW	0b	IRP0 Error Mask; When set, disables logging of error




6.8.26 gcferrst

Type: CFG		Device: 5	PortID: N/A
Bus: 0		Function: 2	
Offset: 0x1ac			
Bit	Attr	Default	Description
26:26	RV	0x0	MC error Memory Controller Error Status.
25:25	RW	0b	Intel VT-d Error
24:24	RW	0b	Miscellaneous Error
23:23	RW	0b	IIO Core Error
20:20	RW	0b	DMI Error
15:5	RW	0x0	PCIe* Error Bit 5: Port 0 Bit 6: Port 1a Bit 7: Port 1b Bit 8: Port 2a Bit 9: Port 2b Bit 10: Port 2c Bit 11: Port 2d Bit 12: Port 3a Bit 13: Port 3b Bit 14: Port 3c Bit 15: Port 3d
1:1	RW	0x0	IRP1 Error Mask
0:0	RW	0b	IRP0 Error Mask; When set, disables logging of error

6.8.27 gcnerrst

Type: CFG		Device: 5	PortID: N/A
Bus: 0		Function: 2	
Offset: 0x1b8			
Bit	Attr	Default	Description
26:26	RV	0x0	MC error Memory Controller Error Status.
25:25	RW	0b	Intel® VT-d Error
24:24	RW	0b	Miscellaneous Error
23:23	RW	0b	IIO Core Error
20:20	RW	0b	DMI Error
15:5	RW	0x0	PCIe* Error Bit 5: Port 0 Bit 6: Port 1a Bit 7: Port 1b Bit 8: Port 2a Bit 9: Port 2b Bit 10: Port 2c Bit 11: Port 2d Bit 12: Port 3a Bit 13: Port 3b Bit 14: Port 3c Bit 15: Port 3d
1:1	RW	0x0	IRP1 Error Mask
0:0	RW	0b	IRP0 Error Mask; When set, disables logging of error



6.8.28 gnerrst

Global Non-Fatal Error Status.

This register indicates the non-fatal error reported to the IIO global error logic. An individual error status bit that is set indicates that a particular local interface has detected an error.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x1c0		Function: 2	
Bit	Attr	Default	Description
25:25	RW1CS	0x0	vtd: Intel VT-d Error Status This register indicates the non-fatal error reported to the Intel VT-d error logic. An individual error status bit that is set indicates that a particular local interface has detected an error.
24:24	RW1CS	0x0	mi: Miscellaneous Error Status
23:23	RW1CS	0x0	iiio: IIO Core Error Status This bit indicates that IIO core has detected an error
20:20	RW1CS	0x0	dmi: This bit indicates that IIO DMI port 0 has detected an error.
15:15	RW1CS	0x0	pcie10:
14:14	RW1CS	0x0	pcie9:
13:13	RW1CS	0x0	pcie8:
12:12	RW1CS	0x0	pcie7:
11:11	RW1CS	0x0	pcie6:
10:10	RW1CS	0x0	pcie5:
9:9	RW1CS	0x0	pcie4:
8:8	RW1CS	0x0	pcie3:
7:7	RW1CS	0x0	pcie2:
6:6	RW1CS	0x0	pcie1:
5:5	RW1CS	0x0	pcie0:
3:3	RW1CS	0x0	csipro1:
2:2	RW1CS	0x0	csipro0:
1:1	RW1CS	0x0	IRP1 Coherent Interface Error
0:0	RW1CS	0x0	IRP0 Coherent Interface Error



6.8.29 **gferrst**

Global Fatal Error Status.

This register indicates the fatal error reported to the IIO global error logic. An individual error status bit that is set indicates that a particular local interface has detected an error.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x1c4		Function: 2	
Bit	Attr	Default	Description
25:25	RW1CS	0x0	vtd: This register indicates the fatal error reported to the Intel VT-d error logic. An individual error status bit that is set indicates that a particular local interface has detected an error.
24:24	RW1CS	0x0	mi: Miscellaneous Error Status
23:23	RW1CS	0x0	iio: IIO Core Error Status This bit indicates that IIO core has detected an error
20:20	RW1CS	0x0	dmi: This bit indicates that IIO DMI port 0 has detected an error.
15:15	RW1CS	0x0	pcie10:
14:14	RW1CS	0x0	pcie9:
13:13	RW1CS	0x0	pcie8:
12:12	RW1CS	0x0	pcie7:
11:11	RW1CS	0x0	pcie6:
10:10	RW1CS	0x0	pcie5:
9:9	RW1CS	0x0	pcie4:
8:8	RW1CS	0x0	pcie3:
7:7	RW1CS	0x0	pcie2:
6:6	RW1CS	0x0	pcie1:
5:5	RW1CS	0x0	pcie0:
1:1	RW1CS	0x0	IRP1 Coherent Interface Error::
0:0	RW1CS	0x0	IRP0 Coherent Interface Error:

6.8.30 **gerrctl**

Global Error Control.

This register controls/masks the reporting of errors detected by the IIO local interfaces. An individual error control bit that is set masks error reporting of the particular local interface; software may set or clear the control bit. This register is sticky and can only be reset by PWRGOOD. Note that bit fields in this register can become reserved depending on the port configuration. For example, if the PCIe port is configured as 2X8 ports, then only the corresponding PCI-EX8 bit fields are valid; other bits are unused and reserved. Global error control register masks errors reported from the local interface to the global register. If the an error reporting is disabled in this register, all errors from the corresponding local interface will not set any of the global error status bits.




Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x1c8		Function: 2	
Bit	Attr	Default	Description
26:26	RV	0x0	MC error Memory Controller Error Status.
25:25	RW	0x0	vtd_err_msk:
24:24	RW	0x0	mi_err_msk:
23:23	RW	0x0	iio_err_msk:
20:20	RW	0x0	dmi_err_msk: This bit enables/masks the error detected in the DMI[0] Port.
15:15	RW	0x0	pcie_err_msk10:
14:14	RW	0x0	pcie_err_msk9:
13:13	RW	0x0	pcie_err_msk8:
12:12	RW	0x0	pcie_err_msk7:
11:11	RW	0x0	pcie_err_msk6:
10:10	RW	0x0	pcie_err_msk5:
9:9	RW	0x0	pcie_err_msk4:
8:8	RW	0x0	pcie_err_msk3:
7:7	RW	0x0	pcie_err_msk2:
6:6	RW	0x0	pcie_err_msk1:
5:5	RW	0x0	pcie_err_msk0:
3:3	RW	0x0	csip_err_msk1:
2:2	RW	0x0	csip_err_msk0:
1:1	RW	0x0	IRP1 Error Mask:
0:0	RW	0x0	IRP0 Error Mask: When set, disables logging of this error

6.8.31 gsysst

Global System Event Status.

This register indicates the error severity signaled by the IIO global error logic. Setting of an individual error status bit indicates that the corresponding error severity has been detected by the IIO.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x1cc		Function: 2	
Bit	Attr	Default	Description
2:2	ROS_V	0x0	sev2: When set, IIO has detected an error of error severity 2
1:1	ROS_V	0x0	sev1: When set, IIO has detected an error of error severity 1
0:0	ROS_V	0x0	sev0: When set, IIO has detected an error of error severity 0



6.8.32 gsysctl

Global System Event Control.

The system event control register controls/masks the reporting the errors indicated by the system event status register. When cleared, the error severity does not cause the generation of the system event. When set, detection of the error severity generates system events according to system event map register (SYSMAP).

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x1d0		Function: 2	
Bit	Attr	Default	Description
2:2	RW	0x0	sev2_en: When set, the detection of error severity 2 generates system events.
1:1	RW	0x0	sev1_en: When set, the detection of error severity 1 generates system events.
0:0	RW	0x0	sev0_en: When set, the detection of error severity 0 generates system events.

6.8.33 gfferrst, gfnerrst

Global Fatal FERR and NERR Status.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x1dc, 0x1e8		Function: 2	
Bit	Attr	Default	Description
26:0	ROS_V	0x0	log: This field logs the global error status register content when the first fatal error is reported. This has the same format as the global fatal error status register (GFERRST).

6.8.34 gnerrst, gnnerrst

Global Non-Fatal FERR and NERR Status

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x1ec, 0x1f8		Function: 2	
Bit	Attr	Default	Description
26:0	ROS_V	0x0	log: This field logs the global error status register content when the first non-fatal error is reported. This has the same format as the global non-fatal error status register (GNERRST).



6.8.35 irpp[0:1]errst

IRP Protocol Error Status.

This register indicates the error detected by the Coherent Interface.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x230, 0x2b0		Function: 2	
Bit	Attr	Default	Description
14:14	RW1CS	0x0	protocol_parity_error: (DB) Originally used for detecting parity error on coherent interface, however, no parity checks exist. So this logs parity errors on data from the IIO switch on the inbound path.
13:13	RW1CS	0x0	protocol_qt_overflow_underflow: (DA)
10:10	RW1CS	0x0	protocol_rcvd_unexprsp: (D7) A completion has been received from the Coherent Interface that was unexpected.
6:6	RW1CS	0x0	csr_acc_32b_unaligned: (C3)
4:4	RW1CS	0x0	wrcache_uncecc_error1: (C2) A double bit ECC error was detected within the Write Cache in set 1.
3:3	RW1CS	0x0	wrcache_uncecc_error0: (C2) A double bit ECC error was detected within the Write Cache in set 0.
3:3	RW1CS	0x0	protocol_rcvd_poison: (C1) A poisoned packet has been received from the Coherent Interface.
2:2	RW1CS	0x0	wrcache_correcc_error1: (B4) A single bit ECC error was detected and corrected within the Write Cache in set 1.
1:1	RW1CS	0x0	wrcache_correcc_error0: (B4) A single bit ECC error was detected and corrected within the Write Cache in set 0.

6.8.36 irpp[0:1]errctl

IRP Protocol Error Control.

This register enables the error status bit setting for a Coherent Interface detected error. Setting of the bit enables the setting of the corresponding error status bit in IRPPERRST register. If the bit is cleared, the corresponding error status will not be set.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x234, 0x2b4		Function: 2	
Bit	Attr	Default	Description
14:14	RWS	0x0	protocol_parity_error: (DB) 0: Disable error status logging for this error 1: Enable Error status logging for this error
13:13	RWS	0x0	protocol_qt_overflow_underflow: (DA) 0: Disable error status logging for this error 1: Enable Error status logging for this error
10:10	RWS	0x0	protocol_rcvd_unexprsp: (D7) 0: Disable error status logging for this error 1: Enable Error status logging for this error




Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x234, 0x2b4		Function: 2	
Bit	Attr	Default	Description
6:6	RWS	0x0	csr_acc_32b_unaligned: (C3) 0: Disable error status logging for this error 1: Enable Error status logging for this error
3:3	RWS	0x0	wrcache_uncecc_error1: (C2) 0: Disable error status logging for this error 1: Enable Error status logging for this error
4:4	RWS	0x0	wrcache_uncecc_error0: (C2) 0: Disable error status logging for this error 1: Enable Error status logging for this error
3:3	RWS	0x0	protocol_rcvd_poison: (C1) 0: Disable error status logging for this error 1: Enable Error status logging for this error
2:2	RWS	0x0	wrcache_correcc_error1: (B4) 0: Disable error status logging for this error 1: Enable Error status logging for this error.
1:1	RW1CS	0x0	wrcache_correcc_error0: (B4) A single bit ECC error was detected and corrected within the Write Cache in set 0.

6.8.37 irpp[0:1]fferrst, irpp[0:1]fnerrst

IRP Protocol Fatal FERR and NERR Status.

The error status log indicates which error is causing the report of the first fatal error event.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: irp0: 0x238, 0x23c irp1: 0x2b8, 0x2bc		Function: 2	
Bit	Attr	Default	Description
14:14	ROS_V	0x0	protocol_parity_error: (DB) Originally used for detecting parity error on coherent interface, however, no parity checks exist. So this logs parity errors on data from the IIO switch on the inbound path.
13:13	ROS_V	0x0	protocol_qt_overflow_underflow: (DC)
10:10	ROS_V	0x0	protocol_rcvd_unexprsp: (D7) A completion has been received from the Coherent Interface that was unexpected.
6:6	ROS_V	0x0	csr_acc_32b_unaligned: (C3)
4:4	ROS_V	0x0	wrcache_uncecc_error1: (C2) A double bit ECC error was detected within the Write Cache in set 1.
3:3	ROS_V	0x0	wrcache_uncecc_error0: (C2) A double bit ECC error was detected within the Write Cache in set 0.
3:3	ROS_V	0x0	protocol_rcvd_poison: (C1) A poisoned packet has been received from the Coherent Interface.
2:2	ROS_V	0x0	wrcache_correcc_error1: (B4) A single bit ECC error was detected and corrected within the Write Cache in set 1.



Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x26c, 0x2ec		Function: 2	
Bit	Attr	Default	Description
6:0	RW1CS	0x0	errcnt: This counter accumulates errors that occur when the associated error type is selected in the ERRCNTSEL register. Notes: This register is cleared by writing 7Fh. Maximum counter available is 7Fh

6.8.43 iioerrst

IIO Core Error Status.

This register indicates the IIO internal core errors detected by the IIO error logic. An individual error status bit that is set indicates that a particular error occurred; software may clear an error status by writing a 1 to the respective bit. This register is sticky and can only be reset by PWRGOOD. Clearing of the IIOERRST is done by clearing the corresponding IIOERRST bits.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x300		Function: 2	
Bit	Attr	Default	Description
6:6	RW1CS	0x0	c6: Overflow/Underflow Error Status (C6)
4:4	RW1CS	0x0	c4: Master Abort Error Status (C4)
0:0	RW1CS	0x0	c7_multicast_target_error: Multicast target error indicating a multicast transaction has targeted more than the number of groups supported.



6.8.44 iioerrctl

IIO Core Error Control.

This register controls the reporting of IIO internal core errors detected by the IIO error logic. An individual error control bit that is cleared masks reporting of that a particular error; software may set or clear the respective bit. This register is sticky and can only be reset by PWRGOOD.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x304		Function: 2	
Bit	Attr	Default	Description
6:6	RWS_L	0x0	c6: Overflow/Underflow Error Enable (C6)
4:4	RWS_L	0x0	c4: Master Abort Error Enable (C4)
0:0	RWS_L	0x0	c7_multicast_target_error — Multicast Target Error Enable.

6.8.45 iioferrst, iiofnerrst

IIO Core Fatal FERR and NERR Status.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x308, 0x31c		Function: 2	
Bit	Attr	Default	Description
6:0	ROS_V	0x0	iio_core_error_status_log: The error status log indicates which error is causing the report of the first error event. The encoding indicates the corresponding bit position of the error in the error status register. It has the same field mapping as IIOERRST.

6.8.46 iioferrhd_[0:3]

IIO Core Fatal FERR Header.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x30c, 0x310, 0x314, 0x318		Function: 2	
Bit	Attr	Default	Description
31:0	ROS_V	0x0	iio_core_error_header_log: Logs the respective DWORD of the header on an error condition.



6.8.47 iionferrst, iionnerrst

IIO Core Non-Fatal FERR and NERR Status.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x320, 0x334		Function: 2	
Bit	Attr	Default	Description
6:0	ROS_V	0x0	iio_core_error_status_log: The error status log indicates which error is causing the report of the first error event. The encoding indicates the corresponding bit position of the error in the error status register. It has the same field mapping as IIOERRST.

6.8.48 iionferrhd_[0:3]


IIO Core Non-Fatal FERR Header.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x324, 0x328, 0x32c, 0x330		Function: 2	
Bit	Attr	Default	Description
31:0	ROS_V	0x0	iio_core_error_header_log: Logs the respective DWORD of the header on an error condition. Header log stores the IIO data path header information of the associated IIO core error. The header indicates where the error is originating from and the address of the cycle.

6.8.49 iioerrcntsel

IIO Core Error Counter Selection.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x33c		Function: 2	
Bit	Attr	Default	Description
6:6	RW_L	0x0	c6: Overflow/Underflow Error Count Select
4:4	RW_L	0x0	c4: Master Abort Error Select
1:1	RW_L	0x0	c7_multicast_target_error: Multicast Target Error Select



6.8.50 iioerrcnt

IIO Core Error Counter.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x340		Function: 2	
Bit	Attr	Default	Description
7:7	RW1CS	0x0	errovf: 0: No overflow occurred 1: Error overflow. The error count may not be valid.
6:0	RW1CS	0x0	errcnt: This counter accumulates errors that occur when the associated error type is selected in the ERRCNTSEL register. Notes: This register is cleared by writing 7Fh. Maximum counter available is 7Fh.

6.8.51 mierrst


Miscellaneous Error Status.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x380		Function: 2	
Bit	Attr	Default	Description
3:3	RW1CS	0x0	vpp_err_sts: VPP Hot-plug I/O Extender Port Error Status. I/O module encountered persistent VPP failure. The VPP is unable to operate.

6.8.52 mierrctl

Miscellaneous Error Control.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x384		Function: 2	
Bit	Attr	Default	Description
3:3	RWS	0x0	vpp_err_sts: VPP Error Status Enable.



6.8.53 mifferrst, mifnerrst

Miscellaneous Fatal FERR and NERR Status.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x388, 0x39c		Function: 2	
Bit	Attr	Default	Description
10:0	ROS_V	0x0	mi_err_st_log: There is 1 bit per VPP port to support up to 11 slots. This field only logs VPP errors. Vpp is serial bus that indicates which port (slot) has a hot plug event pending.

6.8.54 mifferrhdr_[0:3]

Miscellaneous Fatal FERR Header Log.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x38c, 0x390, 0x394, 0x398		Function: 2	
Bit	Attr	Default	Description
31:0	ROS_V	0x0	hdr: Logs the respective DWORD of the header on an error condition.

6.8.55 minferrst, minnerrst

Miscellaneous Non-Fatal FERR and NERR Status.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x3a0, 0x3b4		Function: 2	
Bit	Attr	Default	Description
10:0	ROS_V	0x0	mi_err_st_log: There is 1 bit per VPP port to support up to 11 slots. This field only logs VPP errors. Vpp is serial bus that indicates which port (slot) has a hot plug event pending.

6.8.56 minferrhdr_[0:3]

Miscellaneous Non-Fatal FERR Header Log.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x3a4, 0x3a8, 0x3ac, 0x3b0		Function: 2	
Bit	Attr	Default	Description
31:0	ROS_V	0x0	hdr: Logs the respective DWORD of the header on an error condition.



6.8.57 mierrcntsel

Miscellaneous Error Count Select.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x3bc		Function: 2	
Bit	Attr	Default	Description
3:3	RW	0x0	vpp_err_sts: VPP Error Status Count Select.

6.8.58 mierrcnt

Miscellaneous Error Count.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x3c0		Function: 2	
Bit	Attr	Default	Description
7:7	RW1CS	0x0	errovflow: 0: No overflow occurred 1: Error overflow. The error count may not be valid.
6:0	RW1CS	0x0	errcnt: This counter accumulates errors that occur when the associated error type is selected in the ERRCNTSEL register. Notes: This register is cleared by writing 7Fh. Maximum counter available is 127d (7Fh).

6.9 Device 5 Function 4

I/OxAPCI Configuration Space.

Register Name	Offset	Size
vid	0x0	16
did	0x2	16
pcicmd	0x4	16
pcists	0x6	16
rid	0x8	8
ccr	0x9	24
clsr	0xc	8
hdr	0xe	8
mbar	0x10	32
svid	0x2c	16
sid	0x2e	16
capptr	0x34	8
intlin	0x3c	8
intpin	0x3d	8



Register Name	Offset	Size
abar	0x40	16
pxpcap	0x44	32
snapshot_index	0x80	8
snapshot_window	0x90	32
ioapictetpc	0xa0	32
pmcap	0xe0	32
pmcsr	0xe4	32
ioadsels0	0x288	32
iointsrc0	0x2a0	32
iointsrc1	0x2a4	32
ioremintcnt	0x2a8	32
ioremgpecnt	0x2ac	32
FauxGV	0x2c4	32

6.9.1 vid

Type: CFG Bus: 0 Offset: 0x0		PortID: N/A Device: 5		Function: 4
Bit	Attr	Default	Description	
15:0	RO	0x8086	vendor_identification_number:	

6.9.2 did

Type: CFG Bus: 0 Offset: 0x2		PortID: N/A Device: 5		Function: 4
Bit	Attr	Default	Description	
15:0	RO	0x2f2c	device_identification_number:	

6.9.3 pcicmd

Type: CFG Bus: 0 Offset: 0x4		PortID: N/A Device: 5		Function: 4
Bit	Attr	Default	Description	
10:10	RO	0x0	intxdisable:	
9:9	RO	0x0	fb2be:	
8:8	RO	0x0	serre:	
7:7	RO	0x0	idsel:	
6:6	RO	0x0	perrrsp:	
5:5	RO	0x0	vga:	



Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x4		Function: 4	
Bit	Attr	Default	Description
4:4	RO	0x0	memwrinv:
3:3	RO	0x0	spcen:
2:2	RW	0x0	bme:
1:1	RW	0x0	mse:
0:0	RO	0x0	iose:

6.9.4 pcists

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x6		Function: 4	
Bit	Attr	Default	Description
15:15	RO_V	0x0	dpe:
14:14	RO	0x0	sse:
13:13	RO	0x0	rma:
12:12	RO	0x0	rta:
11:11	RW1C	0x0	sta:
10:9	RO	0x0	devselt:
8:8	RO	0x0	medierr:
7:7	RO	0x0	fb2bcap:
5:5	RO	0x0	sixtysixmhzcap:
4:4	RO	0x1	capl:
3:3	RO	0x0	intxst:

6.9.5 rid

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x8		Function: 4	
Bit	Attr	Default	Description
7:0	RO_V	0x0	revision_id: Reflects the Uncore Revision ID after reset. Reflects the Compatibility Revision ID after BIOS writes 0x69 to any RID register in any Intel® Xeon® Processor E5 v4 product family function.



6.9.6 CCR

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x9		Function: 4	
Bit	Attr	Default	Description
23:16	RO_V	0x80	base_class: Generic Device
15:8	RO_V	0x0	sub_class: Generic Device
7:0	RO_V	0x20	interface:

6.9.7 clsr

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xc		Function: 4	
Bit	Attr	Default	Description
7:0	RW	0x0	clsr_reg:

6.9.8 hdr

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xe		Function: 4	
Bit	Attr	Default	Description
7:7	RO	0x1	multi_function_device: This bit defaults to 1b since all these devices are multi-function.
6:0	RO	0x0	configuration_layout: This field identifies the format of the configuration header layout. It is Type 0 for all these devices. The default is 00h, indicating a 'endpoint device'.

6.9.9 mbar

I/OxAPIC Based Address.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x10		Function: 4	
Bit	Attr	Default	Description
31:12	RW	0x0	bar: This marks the 4 KB aligned 32-bit base address for memory-mapped registers of I/OxAPIC. Side note: Any accesses via message channel or JTAG mini port to registers pointed to by the MBAR address, are not gated by MSE bit (in PCICMD register) being set, that is, even if MSE bit is a 0, message channel accesses to the registers pointed to by MBAR address are allowed completed normally. These accesses are accesses from internal ucode/pcode and JTAG and they are allowed to access the registers normally even if this bit is clear.



Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x10		Function: 4	
Bit	Attr	Default	Description
3:3	RO	0x0	prefetchable: The I/OxAPIC registers are not prefetchable.
2:1	RO	0x0	type: The IOAPIC registers can only be placed below 4G system address space.
0:0	RO	0x0	memory_space: This Base Address Register indicates memory space.

6.9.10 svid

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x2c		Function: 4	
Bit	Attr	Default	Description
15:0	RW_O	0x8086	svid_reg: The default value specifies Intel but can be set to any value once after reset.

6.9.11 sid

This value is used to identify a particular subsystem.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x2e		Function: 4	
Bit	Attr	Default	Description
15:0	RW_O	0x0	sid_reg: Assigned by the subsystem vendor to uniquely identify the subsystem.

6.9.12 capptr

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x34		Function: 4	
Bit	Attr	Default	Description
7:0	RO	0x44	capability_pointer: Points to the first capability structure for the device which is the PCIe capability.



6.9.13 intlin

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x3c		Function: 4	
Bit	Attr	Default	Description
7:0	RO	0x0	intlin_reg:


6.9.14 intpin

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x3d		Function: 4	
Bit	Attr	Default	Description
7:0	RO	0x0	intpin_reg:

6.9.15 abar

I/OxAPIC Alternate BAR.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x40		Function: 4	
Bit	Attr	Default	Description
15:15	RW	0x0	abar_enable: When set, the range FECX_YZ00 to FECX_YZFF is enabled as an alternate access method to the I/OxAPIC registers and these addresses are claimed by the IIO's internal I/OxAPIC regardless of the setting the MSE bit in the IOxAPIC config space. Bits 'XYZ' are defined below.
11:8	RW	0x0	base_address_19: 16 (XBAD) These bits determine the high order bits of the I/O APIC address map. When a memory address is recognized by the IIO which matches FECX_YZ00-to-FECX_YZFF, the IIO will respond to the cycle and access the internal I/O APIC.
7:4	RW	0x0	base_address_15: 12 (YBAD) These bits determine the low order bits of the IO APIC address map. When a memory address is recognized by the IIO which matches FECX_YZ00-to-FECX_YZFF, the IIO will respond to the cycle and access the internal I/O APIC.
3:0	RW	0x0	base_address_11: 8 (ZBAD) These bits determine the low order bits of the I/O APIC address map. When a memory address is recognized by the IIO which matches FECX_YZ00-to-FECX_YZFF, the IIO will respond to the cycle and access the internal I/O APIC.



6.9.16 pxpcap

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x44		Function: 4	
Bit	Attr	Default	Description
29:25	RO	0x0	interrupt_message_numnber:
24:24	RO	0x0	slot_implemented:
23:20	RO	0x9	device_port_type: Device type is Root Complex Integrated Endpoint
19:16	RO	0x1	capability_version: PCI Express Capability is Compliant with Version 1.0 of the PCI Express Spec. Note: This capability structure is not compliant with Versions beyond 1.0, since they require additional capability registers to be reserved. The only purpose for this capability structure is to make enhanced configuration space available. Minimizing the size of this structure is accomplished by reporting version 1.0 compliancy and reporting that this is an integrated root port device. As such, only three Dwords of configuration space are required for this structure.
15:8	RO	0xe0	next_ptr: Pointer to the next capability. Set to 0 to indicate there are no more capability structures, else default value.
7:0	RO	0x10	capability_idat: Provides the PCI Express capability ID assigned by PCI-SIG.

6.9.17 snapshot_index

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x80		Function: 4	
Bit	Attr	Default	Description
7:0	RW	0x0	ssidix: When PECI/JTAG wants to read the indirect RTE registers of I/OxAPIC, this register is used to point to the index of the indirect register, as defined in the I/ OxAAPIC indirect memory space. Software writes to this register and then does a read of the RDWINDOW register to read the contents at that index. Note h/w does not preclude software from accessing this register over the coherent interface but that is not what this register is defined for.

6.9.18 snapshot_window

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x90		Function: 4	
Bit	Attr	Default	Description
31:0	RO_V	0x0	sswindow: When SMBUS/JTAG reads this register, the data contained in the indirect register pointed to by the RDINDEX register is returned on the read.



6.9.19 ioapictetpc

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xa0		Function: 4	
Bit	Attr	Default	Description
12:12	RW	0x0	ntb_int: 0: srcint is connected to IOAPIC table entry 16 1: srcint is connected to IOAPIC table entry 23 Notes: NTB interrupt is always mapped to entry 23.
10:10	RW	0x0	port3c_intb: 0: srcint is connected to IOAPIC table entry 21 1: srcint is connected to IOAPIC table entry 19
8:8	RW	0x0	port3a_intb: 0: srcint is connected to IOAPIC table entry 20 1: srcint is connected to IOAPIC table entry 17
6:6	RW	0x0	port2c_intb: 0: srcint is connected to IOAPIC table entry 13 1: srcint is connected to IOAPIC table entry 11
4:4	RW	0x0	port2a_intb: 0: srcint is connected to IOAPIC table entry 12 1: srcint is connected to IOAPIC table entry 9
0:0	RW	0x0	port0_intb: 0: srcint is connected to IOAPIC table entry 1 1: srcint is connected to IOAPIC table entry 3

6.9.20 pmcap

Power Management Capabilities.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xe0		Function: 4	
Bit	Attr	Default	Description
31:27	RO	0x0	pme_support: Bits 31, 30 and 27 must be set to '1' for PCI-PCI bridge structures representing ports on root complexes.
26:26	RO	0x0	d2_support: I/OxAPIC does not support power management state D2
25:25	RO	0x0	d1_support: I/OxAPIC does not support power management state D1
24:22	RO	0x0	aux_current:
21:21	RO	0x0	device_specific_initialization:
19:19	RO	0x0	pme_clock: This field is hardwired to 0h as it does not apply to PCI Express.
18:16	RW_O	0x3	version: This field is set to 3h (Power Management 1.2 compliant) as version number. Bit is RW-O to make the version 2h incase legacy OS'es have any issues.
15:8	RO	0x0	next_pointer: This is the last capability in the chain and hence set to 0.



Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xe0		Function: 4	
Bit	Attr	Default	Description
7:0	RO	0x1	capability_id: Provides the Power Management capability ID assigned by PCI-SIG.

6.9.21 pmcsr

Power Management Control and Status.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0xe4		Function: 4	
Bit	Attr	Default	Description
31:24	RO	0x0	data: Not relevant for I/OxAPIC
23:23	RO	0x0	bpcce: Not relevant for I/OxAPIC
22:22	RO	0x0	b2b3: Not relevant for I/OxAPIC
15:15	RO	0x0	pmests: Not relevant for I/OxAPIC
14:13	RO	0x0	dsc1: Not relevant for I/OxAPIC
12:9	RO	0x0	dse1: Not relevant for I/OxAPIC
8:8	RO	0x0	pmee1: Not relevant for I/OxAPIC
3:3	RO	0x1	rstd3hotd0: Indicates I/OxAPIC does not reset its registers when transitioning from D3hot to D0.
1:0	RW_V	0x0	power_state: This 2-bit field is used to determine the current power state of the function and to set a new power state as well. 00: D0 01: D1 (not supported by IOAPIC) 10: D2 (not supported by IOAPIC) 11: D3_hot If Software tries to write 01 or 10 to this field, the power state does not change from the existing power state (which is either D0 or D3hot) and nor do these bits1:0 change value. When in D3hot state, I/OxAPIC will a) Respond to only Type 0 configuration transactions targeted at the device's configuration space, when in D3hot state. c) Will not respond to memory (that is, D3hot state is equivalent to MSE), accesses to MBAR region (note: ABAR region access still go through in D3hot state, if it enabled). d) Will not generate any MSI writes .



6.9.22 ioadsels0

I/OxAPIC DSELS Register 0.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x288		Function: 4	
Bit	Attr	Default	Description
28:28	RWS	0x0	sw2ipc_aer_negedge_msk: SW2IPC AER Negative Edge Mask
27:27	RWS	0x0	sw2ipc_aer_event_sel: SW2IPC AER Event Select

6.9.23 iointrsrc0

IO Interrupt Source Register 0.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x2a0		Function: 4	
Bit	Attr	Default	Description
31:0	RW_V	0x0	int_src0: bit interrupt source 31: INTD Port 3b 30: INTC Port 3b 29: INTB Port 3b 28: INTA Port 3b 27: INTD Port 3a 26: INTC Port 3a 25: INTB Port 3a 24: INTA Port 3a 23: INTD Port 1b 22: INTC Port 1b 21: INTB Port 1b 20: INTA Port 1b 19: INTD Port 1a 18: INTC Port 1a 17: INTB Port 1a 16: INTA Port 1a 15: INTD Port 2d 14: INTC Port 2d 13: INTB Port 2d 12: INTA Port 2d 11: INTD Port 2c 10: INTC Port 2c 9: INTB Port 2c 8: INTA Port 2c 7: INTD Port 2b 6: INTC Port 2b 5: INTB Port 2b 4: INTA Port 2b 3: INTD Port 2a 2: INTC Port 2a 1: INTB Port 2a 0: INTA Port 2a



6.9.24 iointsrc1

IO Interrupt Source Register 1.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x2a4		Function: 4	
Bit	Attr	Default	Description
20:0	RW_V	0x0	int_src1: bit interrupt source 20: INTA Root Port Core 19: INTB ME KT 18: INTC ME IDE-R 17: INTD ME HECI 16: INTA ME HECI 15: INTD Intel QuickData Technology DMA 14: INTC Intel QuickData Technology DMA 13: INTB Intel QuickData Technology DMA 12: INTA Intel QuickData Technology DMA 11: INTD Port 0 DMI 10: INTC Port 0 DMI 9: INTB Port 0 DMI 8: INTA Port 0 DMI 7: INTD Port 3d 6: INTC Port 3d 5: INTB Port 3d 4: INTA Port 3d 3: INTD Port 3c 2: INTC Port 3c 1: INTB Port 3c 0: INTA Port 3c

6.9.25 ioremintcnt

Remote IO Interrupt Count.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x2a8		Function: 4	
Bit	Attr	Default	Description
31:0	RW_V	0x0	rem_int_cnt: Number of remote interrupts received.



6.9.26 ioremgpecnt

Remote IO GPE Count.

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x2ac		Function: 4	
Bit	Attr	Default	Description
23:16	RW_V	0x0	hpgpe_cnt: Number of remote HPGPEs received.
15:8	RW_V	0x0	pmgpe_cnt: Number of remote PMGPEs received.
7:0	RW_V	0x0	gpe_cnt: Number of remote GPEs received.

6.9.27 FauxGV

Type: CFG		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x2c4		Function: 4	
Bit	Attr	Default	Description
0:0	RWS_L	0x0	FauxGVEn: Enable Fault GV.

6.10 Device 5 Function 4 I/OxAPIC

I/OxAPIC has a direct memory mapped space. An index/data register pair is located within the directed memory mapped region and is used to access the redirection table entries. The offsets shown in the table are from the base address in either ABAR or MBAR or both.

Access to addresses beyond 0x40h return all 0s.

Only addresses up to offset 0xFF can be accessed via the ABAR register whereas offsets up to 0xFFF can be accessed via MBAR.

Only aligned DWORD reads and write are allowed towards the I/OxAPIC memory space. Any other accesses will result in an error.

Register Name	Offset	Size
index	0x0	8
window	0x10	32
eoi	0x40	8

6.10.1 index

The Index Register will select which indirect register appears in the window register to be manipulated by software. Software will program this register to select the desired APIC internal register.



Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x0		Function: 4	
Bit	Attr	Default	Description
7:0	RW_L	0x0	idx: Indirect register to access.

6.10.2 window

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x10		Function: 4	
Bit	Attr	Default	Description
31:0	RW_LV	0x0	window_reg: Data to be written to the indirect registers on writes, and location of read data from the indirect register on reads.

6.10.3 eoi

Type: MEM		PortID: 8'h7e	
Bus: 0		Device: 5	
Offset: 0x40		Function: 4	
Bit	Attr	Default	Description
7:0	RW_L	0x0	eoi_reg: The EOI register is present to provide a mechanism to efficiently convert level interrupts to edge triggered MSI interrupts. When a write is issued to this register, the I/O(x)APIC will check the lower 8 bits written to this register, and compare it with the vector field for each entry in the I/O Redirection Table. When a match is found, the Remote_IRR bit for that I/O Redirection Entry will be cleared. Note that if multiple I/O Redirection entries, for any reason, assign the same vector, each of those entries will have the Remote_IRR bit reset to '0'. This will cause the corresponding I/OxAPIC entries to resample their level interrupt inputs and if they are still asserted, cause more MSI interrupt(s) (if unmasked) which will again set the Remote_IRR bit.

6.11 Device 5 Function 4 Window 0

Register Name	Offset	Size
arbid__window	0x2	32
bcfg__window	0x3	32
rtI0__window	0x10	32
rth0__window	0x11	32
rtI1__window	0x12	32
rth1__window	0x13	32
rtI2__window	0x14	32
rth2__window	0x15	32
rtI3__window	0x16	32



Register Name	Offset	Size
rth3__window	0x17	32
rtl4__window	0x18	32
rth4__window	0x19	32
rtl5__window	0x1a	32
rth5__window	0x1b	32
rtl6__window	0x1c	32
rth6__window	0x1d	32
rtl7__window	0x1e	32
rth7__window	0x1f	32
rtl8__window	0x20	32
rth8__window	0x21	32
rtl9__window	0x22	32
rth9__window	0x23	32
rtl10__window	0x24	32
rth10__window	0x25	32
rtl11__window	0x26	32
rth11__window	0x27	32
rtl12__window	0x28	32
rth12__window	0x29	32
rtl13__window	0x2a	32
rth13__window	0x2b	32
rtl14__window	0x2c	32
rth14__window	0x2d	32
rtl15__window	0x2e	32
rth15__window	0x2f	32
rtl16__window	0x30	32
rth16__window	0x31	32
rtl17__window	0x32	32
rth17__window	0x33	32
rtl18__window	0x34	32
rth18__window	0x35	32
rtl19__window	0x36	32
rth19__window	0x37	32
rtl20__window	0x38	32
rth20__window	0x39	32
rtl21__window	0x3a	32
rth21__window	0x3b	32
rtl22__window	0x3c	32
rth22__window	0x3d	32
rtl23__window	0x3e	32
rth23__window	0x3f	32



6.11.0.1 arbid__window

Tracks the APICID register for compatibility reasons.

Type: MEM		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x2		Function: 4	
Bit	Attr	Default	Description
27:24	RO	0x0	arbitration_id: Tracks the APICID register.

6.11.0.2 bcfg__window

Type: MEM		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x3		Function: 4	
Bit	Attr	Default	Description
0:0	RW	0x1	boot_configuration: This bit is a default1 to indicate FSB delivery mode. A value of 0 has no effect. Its left as RW for software compatibility.

6.11.0.3 rtl[0:23]__window

The information in this register along with Redirection Table High DWORD register is used to construct the MSI interrupt. There is one of these pairs of registers for every interrupt. The first interrupt has the redirection registers at offset 10h. The second interrupt at 12h, third at 14h, etc. until the final interrupt (interrupt 23) at 3Eh.

Type: MEM		PortID: N/A	
Bus: 0		Device: 5	
Offset: 0x10, 0x12, 0x14, 0x16, 0x18, 0x1a, 0x1c, 0x1e, 0x20, 0x22, 0x24, 0x26, 0x28, 0x2a, 0x2c, 0x2e, 0x30, 0x32, 0x34, 0x36, 0x38, 0x3a, 0x3c, 0x3e		Function: 4	
Bit	Attr	Default	Description
17:17	RW	0x0	disable_flushing: This bit has no meaning in IIO. This bit is R/W for software compatibility reasons.



6.12.1 rx_ctle_peak_gen2

This register controls the Continuous Time Linear Equalizer (CTLE) setting for the named receiver bundles on the selected port on the PCIe interface in Gen. 2 mode.

Type: CFG		PortID: N/A	
Bus: 0		Device: 6	
Offset: 0xA78		Function: 0	
Bit	Attr	Default	Description
7:4	RWS_L	0x7	bnd1:
3:0	RWS_L	0x7	bnd0:

6.12.2 rx_ctle_peak_gen2


This register controls the Continuous Time Linear Equalizer (CTLE) setting for the named receiver bundles on the selected port on the PCIe interface in Gen. 2 mode.

Type: CFG		PortID: N/A	
Bus: 0		Device: 6	
Offset: 0xA78		Function: 1	
Bit	Attr	Default	Description
15:12	RWS_L	0x7	bnd3:
11:8	RWS_L	0x7	bnd2:
7:4	RWS_L	0x7	bnd1:
3:0	RWS_L	0x7	bnd0:

6.12.3 rx_ctle_peak_gen3

This register controls the Continuous Time Linear Equalizer (CTLE) setting for the named receiver bundles on the selected port on the PCIe interface in Gen. 3 mode.

Type: CFG		PortID: N/A	
Bus: 0		Device: 6	
Offset: 0xA80		Function: 1	
Bit	Attr	Default	Description
19:15	RWS_L	0x7	bnd3:
14:10	RWS_L	0x7	bnd2:
9:5	RWS_L	0x7	bnd1:
4:0	RWS_L	0x7	bnd0:



6.12.4 rx_ctle_peak_gen2

This register controls the Continuous Time Linear Equalizer (CTLE) setting for the named receiver bundles on the selected port on the PCIe interface in Gen. 2 mode.

Type:	CFG	PortID:	N/A	Function:	3
Bus:	0	Device:	6	Function:	0
Bus:	0	Device:	7		
Offset:	0xA78				
Bit	Attr	Default	Description		
41:37	RWS_L	0x7	bndl7:		
36:32	RWS_L	0x7	bndl6:		
29:25	RWS_L	0x7	bndl5:		
24:20	RWS_L	0x7	bndl4:		
19:15	RWS_L	0x7	bndl3		
14:10	RWS_L	0x7	bndl2		
9:5	RWS_L	0x7	bndl1		
4:0	RWS_L	0x7	bndl0		

6.12.5 rx_ctle_peak_gen3

This register controls the Continuous Time Linear Equalizer (CTLE) setting for the named receiver bundles on the selected port on the PCIe interface in Gen. 3 mode.

Type:	CFG	PortID:	N/A	Function:	3
Bus:	0	Device:	6	Function:	0
Bus:	0	Device:	7		
Offset:	0xA80				
Bit	Attr	Default	Description		
41:37	RWS_L	0x7	bndl7		
36:32	RWS_L	0x7	bndl6		
29:25	RWS_L	0x7	bndl5		
24:20	RWS_L	0x7	bndl4		
19:15	RWS_L	0x7	bndl3		
14:10	RWS_L	0x7	bndl2		
9:5	RWS_L	0x7	bndl1		
4:0	RWS_L	0x7	bndl0		

6.13 Non Transparent Bridge Registers

6.13.1 Configuration Register Map (NTB Primary Side)

This section covers the NTB primary side configuration space registers.

Bus 0, Device 3, Function 0 can function in three modes: PCI Express Root Port, NTB/NTB and NTB/RP. When configured as an NTB there are two sides to discuss for configuration registers. The primary side of the NTB's configuration space is located on Bus 0, Device 3, Function 0 with respect to and a secondary side of the NTB's configuration space is located on some enumerated bus on another system and does not exist as configuration space on the local system anywhere.



Table 6-5. Device 3 Function 0 (Non-Transparent Bridge) Configuration Map Offset 0x00h - 0xFCh

DID		VID		0h	MSIXMSGCTRL	MSIXNXTPT R	MSIXCAPID	80h								
PCISTS		PCICMD		4h	TABLEOFF_BIR			84h								
CCR			RID	8h	PBAOFF_BIR			88h								
BIST	HDR	PLAT	CLSR	Ch				8Ch								
PB01BASE				10h	PXPCAP	PXPNTPTR	PXPCAPID	90h								
				14h	DEVCAP			94h								
PB23BASE				18h	DEVSTS	DEVCTRL		98h								
				1Ch						9Ch						
PB45BASE				20h									A0h			
				24h									A4h			
				28h									A8h			
SDID		SVID		2Ch												ACh
				30h												B0h
					CAPPTR			B4h								
				38h				B8h								
MAXLAT	MINGNT	INTPIN	INTL	3Ch				BCh								
				40h				C0h								
				44h				C4h								
				48h				C8h								
				4Ch				CCh								
				50h				SBAR45SZ	SBAR23SZ	PBAR45SZ	PBAR23SZ	D0h				
				54h							PPD	D4h				
				58h				D8h								
				5Ch				DCh								
MSIMSGCTL		MSINXTPTR	MSICAPID	60h	PMCAP			E0h								
MSGADR				64h	PMCSR			E4h								
MSGDAT				68h				E8h								
MSIMSK				6Ch				ECh								
MISIPENDING				70h				F0h								
				74h				F4h								
				78h				F8h								
				7Ch				FCh								

Table 6-6. Device 3 Function 0 (Non-Transparent Bridge) Configuration Map Offset 0x100h - 0x1FCh

XPREUT_HDR_EXT		100h	PERFCTRLSTS		180h
XPREUT_HDR_CAP		104h			184h
XPREUT_HDR_LEF		108h	MISCCTRLSTS		188h
		10Ch			18Ch
ACSCAPHDR				PCIE_IOW_BIF_CTRL	190h
ACCTRL	ACSCAP	114h	NTBDEVCAP		194h



Table 6-6. Device 3 Function 0 (Non-Transparent Bridge) Configuration Map Offset 0x100h - 0x1FCh

		118h		198h
		11Ch	LNKCAP	
		120h	LNKSTS	LNKCON
		124h	SLTCAP	
		128h	SLTSTS	SLTCON
		12Ch	ROOTCAP	ROOTCON
		130h	ROOTSTS	
		134h	DEVCAP2	
		138h		DEVCTRL2
		13Ch	LNKCAP2	
APICLIMIT	APICBASE	140h	LNKSTS2	LNKCON2
VSECPHDR		144h		
VSHDR		148h		
UNCERRSTS		14Ch		
UNCERRMSK		150h	ERRINJCAP	
UNCERRSEV		154h	ERRINJHDR	
CORERRSTS		158h		ERRINJCON
CORERRMSK		15Ch		
ERRCAP		160h	CTOCTRL	
HDRLOG0		164h		
HDRLOG1		168h		
HDRLOG2		16Ch		
HDRLOG3		170h		
RPERRCMD		174h		
RPERRSTS		178h		
ERRSID		17Ch		

Table 6-7. Device 3 Function 0 (Non-Transparent Bridge) Configuration Map Offset 0x200h - 0x2FCh

XPCORERRSTS		200h		280h
XPCORERRMSK		204h		284h
XPUNCERRSTS		208h		288h
XPUNCERRMSK		20Ch		28Ch
XPUNCERRSEV		210h		290h
	XPUNCERR PTR	214h		294h
UNCEDMASK		218h		298h
COREDMASK		21Ch		29Ch
RPEDMASK		220h		2A0h
XPUNCEDMASK		224h		2A4h
XPCOREDMASK		228h		2A8h
		22Ch		2ACh
XPGLBERRPTR	XPGLBERRSTS	230h		2B0h



Table 6-7. Device 3 Function 0 (Non-Transparent Bridge) Configuration Map Offset 0x200h - 0x2FCh

		234h		2B4h
		238h		2B8h
		23Ch		2BCh
		240h		2C0h
		244h		2C4h
		248h		2C8h
		24Ch		2CCh
PXP2CAP		250h		2D0h
LNKCON3		254h		2D4h
LNERRSTS		258h		2D8h
LN1EQ	LN0EQ	25Ch		2DCh
LN3EQ	LN2EQ	260h		2E0h
LN5EQ	LN4EQ	264h		2E4h
LN7EQ	LN6EQ	268h		2E8h
LN9EQ	LN8EQ	26Ch		2ECh
LN11EQ	LN10EQ	270h		2F0h
LN13EQ	LN12EQ	274h		2F4h
LN15EQ	LN14EQ	278h		2F8h
		27Ch		2FCh

Table 6-8. Device 3 Function 0 (Non-Transparent Bridge) Configuration Map Offset 0x300h - 0x3FCh

mcast_cap_hdr		300h		380h
mcast_cap_ext		304h		384h
		308h		388h
mcast_ctrl	mcast_cap	30Ch		38Ch
mcast_base		310h		390h
		314h		394h
mcast_rcv		318h		398h
		31Ch		39Ch
mcast_blk_all		320h		3A0h
		324h		3A4h
mcast_blk_unt		328h		3A8h
		32Ch		3ACh
mcast_overlay_bar		330h		3B0h
		334h		3B4h
		338h		3B8h
		33Ch		3BCh
		340h		3C0h
		344h		3C4h
		348h		3C8h
		34Ch		3CCh




Table 6-8. Device 3 Function 0 (Non-Transparent Bridge) Configuration Map Offset 0x300h - 0x3FCh

350h	3D0h
354h	3D4h
358h	3D8h
35Ch	3DCh
360h	3E0h
364h	3E4h
368h	3E8h
36Ch	3ECh
370h	3F0h
374h	3F4h
378h	3F8h
37Ch	3FCh

6.13.2 Standard PCI Configuration Space - Type 0 Common Configuration Space

This section covers primary side registers in the 0x0 to 0x3F region that are common to Bus 0, Device 3. Comments at the top of the table indicate what devices/functions the description applies to. Exceptions that apply to specific functions are noted in the individual bit descriptions.

Note: Several registers will be duplicated for device 3 in the three sections discussing the three modes it operates in RP, NTB/NTB, and NTB/RP primary and secondary but are repeated here for readability.

Primary side configuration registers (device 3) can only be read by the local host.

6.13.2.1 VID: Vendor Identification

VID			
Bus: 0 Device: 3 Function: 0 Offset: 0			
Bit	Attr	Default	Description
15:0	RO	8086h	Vendor Identification Number The value is assigned by PCI-SIG to Intel.



6.13.2.2 DID: Device Identification Register

DID Bus: 0 Device: 3 Function: 0 Offset: 2			
Bit	Attr	Default	Description
15:0	RO-V		Device Identification Number This PCI Express Root Port 3.a device ID as follows: 0x2f08: PCI Express Root Port Mode 0x2f0D: Non-Transparent Bridge Primary NTB/NTB mode 0x2f0E: Non-Transparent Bridge Primary NTB/RP mode 0x2f0F: Non-Transparent Bridge Secondary (at BDF = M/N/O accessed from the secondary side) Port3_NTB: Attr: RO-V Default: 2f0Dh

6.13.2.3 PCICMD: PCI Command

This register defines the PCI 3.0 compatible command register values applicable to PCI Express space.

PCICMD Bus: 0 Device: 3 Function: 0 Offset: 4			
Bit	Attr	Default	Description
10	RW	0b	Interrupt Disable Controls the ability of the PCI Express port to generate INTx messages on its own behalf. This bit does not affect the ability of the RP to forward interrupt messages received from the PCI Express port, to the internal I/OxAPIC block. However, this bit controls the internal generation of legacy INTx interrupts for PCI Express RAS events or for INTx interrupts due to Hot Plug/Power Management events or for BW change notification. In NTB mode: 1: Legacy INTx Interrupt mode is disabled 0: Legacy INTx Interrupt mode is enabled and the NTB port can generate INTx interrupts to system Note: If a root port had previously generated an Assert_INTx interrupt when this bit transitions from 0 to 1, then the root port generates a Deassert_INTx message to indicate the interrupt is deasserted.
9	RO	0b	Fast Back-to-Back Enable Not applicable to PCI Express and is hardwired to 0
8	RW	0b	SERR Enable This field enables notifying the internal core error logic of occurrence of an uncorrectable error (fatal or non-fatal) at the NTB port. The internal core error logic of IIO then decides if/how to escalate the error further (pins/message etc.). This bit also controls the propagation of PCI Express ERR_FATAL and ERR_NONFATAL messages received from the port to the internal IIO core error logic. 1: Fatal and Non-fatal error generation and Fatal and Non-fatal error message forwarding is enabled 0: Fatal and Non-fatal error generation and Fatal and Non-fatal error message forwarding is disabled
7	RO	0b	IDSEL Stepping/Wait Cycle Control Not applicable to internal IIO devices. Hardwired to 0.
6	RW	0b	Parity Error Response IIO ignores this bit and always does ECC/parity checking and signaling for data/address of transactions both to and from IIO. This bit though affects the setting of bit 8 in the PCISTS register.
5	RO	0b	VGA palette snoop Enable Not applicable to PCI Express must be hardwired to 0.



PCICMD Bus: 0 Device: 3 Function: 0 Offset: 4			
Bit	Attr	Default	Description
4	RO	0b	Memory Write and Invalidate Enable Not applicable to PCI Express must be hardwired to 0.
3	RO	0b	Special Cycle Enable Not applicable to PCI Express must be hardwired to 0.
2	RW	0b	<p>Bus Master Enable Controls the ability of the PCI Express port in generating and also in forwarding memory (including MSI writes) or I/O transactions (and not messages) or configuration transactions from the secondary side to the primary side.</p> <p>1: Enables the PCI Express port to a) generate MSI writes internally for AER/Hot Plug/Power Management events and also to b) forward memory (including MSI writes from devices south of the RP), config or I/O read/write requests from secondary to primary side</p> <p>0: The Bus Master is disabled. When this bit is 0, IIO root ports will a) treat upstream PCI Express memory writes/reads, IO writes/reads, and configuration reads and writes as unsupported requests (and follow the rules for handling unsupported requests). This behavior is also true towards transactions that are already pending in the IIO root port's internal queues when the BME bit is turned off. b) mask the root port from generating MSI writes internally for AER/Hot Plug/Power Management events at the root port.</p> <p>In NTB mode: When this bit is Set = 1b, the PCIe NTB will forward Memory Requests upstream from the secondary interface to the primary interface. When this bit is Cleared = 0b, the PCIe NTB will not forward Memory Requests from the secondary to the primary interface and will drop all posted memory write requests and will return Unsupported Requests UR for all non-posted memory read requests.</p> <p>Notes: MSI/MSI-X interrupt Messages are in-band memory writes, setting the Bus Master Enable bit = 0b disables MSI/MSI-X interrupt Messages as well. Requests other than Memory or I/O Requests are not controlled by this bit.</p>
1	RW	0b	<p>Memory Space Enable In PCIe mode: 1: Enables a PCI Express port's memory range registers, with the exception of the I/OxAPIC range register ('APICBASE: APIC Base Register (APICBASE)' and 'APICLIMIT: APIC Limit Register (APICLIMIT)'), to be decoded as valid target addresses for transactions from primary side. 0: Disables a PCI Express port's memory range registers, with the exception of the I/OxAPIC range register ('APICBASE: APIC Base Register (APICBASE)' and 'APICLIMIT: APIC Limit Register (APICLIMIT)'), to be decoded as valid target addresses for transactions from primary side.</p> <p>In NTB mode: 1: Enables NTB primary BARs to be decoded as valid target addresses for transactions from primary side. 0: Disables NTB primary BARs to be decoded as valid target addresses for transactions from primary side.</p> <p>Notes: The I/OxAPIC address range of a root port has its own enable bit. This bit is not ever used by hardware to decode transactions from the secondary side of the root port.</p>
0	RO	0b	<p>IO Space Enable 1: Enables the I/O address range, defined in the IOBASE and IOLIM registers of the PCI-to-PCI bridge header, for target decode from primary side 0: Disables the I/O address range, defined in the IOBASE and IOLIM registers of the PCI-to-PCI bridge header, for target decode from primary side</p> <p>Notes: This bit is not ever used by hardware to decode transactions from the secondary side of the root port. NTB does not support I/O space accesses. Hardwired to 0</p>



6.13.2.4 PCISTS: PCI Status

PCISTS Bus: 0 Device: 3 Function: 0 Offset: 6			
Bit	Attr	Default	Description
15	RW1C	0b	<p>Detected Parity Error</p> <p>This bit is set by a device when it receives a packet on the primary side with an uncorrectable data error (i.e. a packet with poison bit set or an uncorrectable data ECC error was detected at the XP-DP interface when ECC checking is done) or an uncorrectable address/control parity error. The setting of this bit is regardless of the Parity Error Response bit (PERRE) in the PCICMD register.</p>
14	RW1C	0b	<p>Signaled System Error</p> <p>1: The root port reported fatal/non-fatal (and not correctable) errors it detected on its PCI Express interface to the IIO core error logic (which might eventually escalate the error through the ERR[2:0] pins or message to cpu core or message to PCH). Note that the SERRE bit in the PCICMD register must be set for a device to report the error the IIO core error logic. Software clears this bit by writing a '1' to it. This bit is also set (when SERR enable bit is set) when a FATAL/NON-FATAL message is forwarded to the IIO core error logic. Note that IIO internal 'core' errors (like parity error in the internal queues) are not reported via this bit.</p> <p>0: The root port did not report a fatal/non-fatal error</p> <p>In NTB mode:</p> <p>1: The device reported fatal/non-fatal (and not correctable) errors it detected on NTB interface. Software clears this bit by writing a '1' to it. Note that IIO internal 'core' errors (like parity error in the internal queues) are not reported via this bit.</p> <p>0: The device did not report a fatal/non-fatal error.</p>
13	RW1C	0b	<p>Received Master Abort</p> <p>This bit is set when a device experiences a master abort condition on a transaction it mastered on the primary interface (IIO internal bus). Note that certain errors might be detected right at the PCI Express interface and those transactions might not 'propagate' to the primary interface before the error is detected (e.g. accesses to memory above TOCM in cases where the PCIe interface logic itself might have visibility into TOCM). Such errors do not cause this bit to be set, and are reported via the PCI Express interface error bits (secondary status register).</p>
12	RW1C	0b	<p>Received Target Abort</p> <p>This bit is set when a device experiences a completer abort condition on a transaction it mastered on the primary interface (uncore internal bus). Note that certain errors might be detected right at the PCI Express interface and those transactions might not 'propagate' to the primary interface before the error is detected (for example, accesses to memory above VTBAR). Such errors do not cause this bit to be set, and are reported via the PCI Express interface error bits (secondary status register).</p> <p>In NTB Mode:</p> <p>Set when a p2p read resulted in CA status</p>
11	RW1C	0b	<p>Signaled Target Abort</p> <p>This bit is set when a root port signals a completer abort completion status on the primary side (internal bus of uncore). This condition includes a PCI Express port forwarding a completer abort status received on a completion from the secondary</p> <p>In NTB Mode:</p> <p>This bit is set when the NTB port forwards a completer abort (CA) completion status from the secondary interface to the primary interface.</p>
10:9	RO	0h	<p>DEVSEL# Timing</p> <p>Not applicable to PCI Express. Hardwired to 0.</p>
8	RW1C	0b	<p>Master Data Parity Error</p> <p>This bit is set if the Parity Error Response bit in the PCI Command register is set and the Requestor receives a poisoned completion on the primary interface or Requestor forwards a poisoned write request (including MSI/MSI-X writes) from the secondary interface to the primary interface.</p>
7	RO	0b	<p>Fast Back-to-Back</p> <p>Not applicable to PCI Express. Hardwired to 0.</p>
5	RO	0b	<p>pci bus 66MHz capable</p> <p>Not applicable to PCI Express. Hardwired to 0.</p>




PCISTS Bus: 0 Device: 3 Function: 0 Offset: 6			
Bit	Attr	Default	Description
4	RO	1b	Capabilities List This bit indicates the presence of a capabilities list structure.
3	RO-V	0b	INTx Status This Read-only bit reflects the state of the interrupt in the PCI Express Root Port. Only when the Interrupt Disable bit in the command register is a 0 and this Interrupt Status bit is a 1, will this device generate INTx interrupt. Setting the Interrupt Disable bit to a 1 has no effect on the state of this bit. This bit does not get set for interrupts forwarded to the root port from downstream devices in the hierarchy. When MSI are enabled, Interrupt status should not be set. The intx status bit should be deasserted when all the relevant events (RAS errors/ Hot Plug/link change status/Power Management) internal to the port using legacy interrupts are cleared by software. In NTB Mode: When Set, indicates that an INTx emulation interrupt is pending internally in the Function. NTB clears this bit when the internal interrupt condition is cleared by software. Note this bit could be set even when INTx assertion is disabled (and INTx mode is enabled though) but an internal interrupt condition is pending.
2:0	RV	0h	Reserved

6.13.2.5 RID: Revision Identification

RID Bus: 0 Device: 3 Function: 0 Offset: 8			
Bit	Attr	Default	Description
7:0	RO	00h	Revision Identification Reflects the Uncore Revision ID after reset. Reflects the Compatibility Revision ID after BIOS writes 0x69 to any RID register in any Processor function.

6.13.2.6 CCR: Class Code

CCR Bus: 0 Device: 3 Function: 0 Offset: 9			
Bit	Attr	Default	Description
23:16	RO	06h	Base Class For PCI Express NTB port this field is hardwired to 06h, indicating it is a 'Bridge Device'.
15:8	RO-V		Sub-Class In NTB mode, this field hardwired to 80h to indicate a 'Other bridge type'. In PCIe mode, it is hardwired to 04h indicating 'PCI-PCI Bridge'. Port3_NTB: Attr: RO-V Default: 80h Port3_PCIe: Attr: RO-V Default: 04h
7:0	RO	00h	Register-Level Programming Interface This field is hardwired to 00h for PCI Express NTB port.



6.13.2.7 CLSR: Cacheline Size

CLSR Bus: 0 Device: 3 Function: 0 Offset: C			
Bit	Attr	Default	Description
7:0	RW	0h	Cacheline Size This register is set as RW for compatibility reasons only. Cacheline size for IIO is always 64B. IIO hardware ignore this setting.

6.13.2.8 HDR: Header Type

HDR Bus: 0 Device: 3 Function: 0 Offset: E			
Bit	Attr	Default	Description
7	RO-V	1b	Multi-function Device This bit defaults to 0 for PCI Express NTB port. BIOS can individually control the value of this bit, based on HDRTPCTRL register. BIOS will write to that register to change this field to 0, if it exposes only function 0 in the device to OS.
6:0	RO		Configuration Layout This field identifies the format of the configuration header layout. It is Type1 for PCI Express and Type0 in NTB mode. The default is 00h, indicating a 'non-bridge function'. Port3_NTB: Attr: RO Default: 00h Port3_PCIe: Attr: RO Default: 01h

6.13.2.9 SVID: Subsystem Vendor ID

Device 3, Function 0, Offset 2Ch. This register exist in both RP and NTB modes. It is documented in [Section 6.2.29](#).

6.13.2.10 SDID: Subsystem Identity

Device 3, Function 0, Offset 2Eh. This register exist in both RP and NTB modes. It is documented in [Section 6.2.30](#).

6.13.2.11 CAPPTR: Capability Pointer

CAPPTR Bus: 0 Device: 3 Function: 0 Offset: 34			
Bit	Attr	Default	Description
7:0	RW-O	60h	Capability Pointer Points to the first capability structure for the device. In NTB mode, capabilities start at a different location.



6.13.2.12 INTL: Interrupt Line

Bus: 0 Device: 3 Function: 0 Offset: 3C			
Bit	Attr	Default	Description
7:0	RW	00h	Interrupt Line This bit is RW for devices that can generate a legacy INTx message and is needed only for compatibility purposes.

6.13.2.13 INTPIN: Interrupt Pin


INTPIN Bus: 0 Device: 3 Function: 0 Offset: 3D			
Bit	Attr	Default	Description
7:0	RW-O	01h	Interrupt Pin This field defines the type of interrupt to generate for the port. 01h: Generate INTA Others: Reserved BIOS can program this to 0 to indicate to OS that the port does not support INTx interrupt.

6.13.3 NTB Port 3A Configured as Primary Endpoint Device

6.13.3.1 PB01BASE: Primary BAR 0/1 Base Address

This register is used to setup the primary side NTB configuration space

PB01BASE Bus: 0 Device: 3 Function: 0 Offset: 10			
Bit	Attr	Default	Description
63:16	RW	0h	Primary BAR 0/1 Base Sets the location of the BAR written by SW on a 64KB alignment
3	RO	1b	Prefetchable BAR points to Prefetchable memory.
2:1	RO	10b	Type Memory type claimed by BAR 0/1 is 64-bit addressable.
0	RO	0b	Memory Space Indicator BAR resource is memory (as opposed to I/O).



6.13.3.2 PB23BASE: Primary BAR 2/3 Base Address

The register is used by the processor on the primary side of the NTB to setup a 64b prefetchable memory window.

PB23BASE Bus: 0 Device: 3 Function: 0 Offset: 18			
Bit	Attr	Default	Description
63:12	RW	0h	Primary BAR 2/3 Base Sets the location of the BAR written by SW NOTE: The number of bits that are writable in this register is dictated by the value loaded into the "PBAR23SZ: Primary BAR 2/3 Size" on page 330 by the BIOS at initialization time (before BIOS PCI enumeration). PBAR23SZ indicates the lowest order bit of this register field that is writeable where valid values are 12-39. If PBAR23SZ is set to 12, all bits are writeable. If set to 39, then bits 38:12 are Read Only and will return values of 0. NOTE: For the special case where PBAR23SZ = '0', bits 63:0 are all RO='0' resulting in the BAR being disabled. NOTE: The lowest order address bit is 12 to enforce a minimum granularity of 4 KB.
3	RO	1b	Prefetchable BAR points to Prefetchable memory.
2:1	RO	10b	Type Memory type claimed by BAR 2/3 is 64-bit addressable.
0	RO	0b	Memory Space Indicator BAR resource is memory (as opposed to I/O).

6.13.3.3 PB45BASE: Primary BAR 4/5 Base Address

The register is used by the processor on the primary side of the NTB to setup a second 64b prefetchable memory window.

PB45BASE Bus: 0 Device: 3 Function: 0 Offset: 20			
Bit	Attr	Default	Description
63:12	RW	0h	Primary BAR 4/5 Base Sets the location of the BAR written by SW NOTE: The number of bits that are writable in this register is dictated by the value loaded into the Section 6.13.3.23, "PBAR45SZ: Primary BAR 4/5 Size" on page 331 by the BIOS at initialization time (before BIOS PCI enumeration). PBAR45SZ indicates the lowest order bit of this register field that is writeable where valid values are 12-39. If PBAR45SZ is set to 12, all bits are writeable. If set to 39, then bits 38:12 are Read Only and will return values of 0. Notes: For the special case where PBAR45SZ = '0', bits 63:0 are all RO='0' resulting in the BAR being disabled. The lowest order address bit is 12 to enforce a minimum granularity of 4 KB.
3	RO	1b	Prefetchable BAR points to Prefetchable memory.
2:1	RO	10b	Type Memory type claimed by BAR 4/5 is 64-bit addressable.
0	RO	0b	Memory Space Indicator BAR resource is memory (as opposed to I/O).

6.13.3.4 MSICAPID: MSI Capability ID

Device 3, Function 0, Offset 60h. This register exist in both RP and NTB modes. It is documented in RP Section 6.2.32.



6.13.3.5 MSINXTPTR: MSI Next Pointer

Device 3, Function 0, Offset 61h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.33](#).

6.13.3.6 MSIMSGCTL: MSI Control

Device 3, Function 0, Offset 62h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.34](#).

6.13.3.7 MSGADR: MSI Address

Device 3, Function 0, Offset 64h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.35](#).

6.13.3.8 MSGDAT: MSI Data Register

Device 3, Function 0, Offset 68h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.36](#).

6.13.3.9 MSIMSK: MSI Mask Bit Register

The Mask Bit register enables software to disable message sending on a per-vector basis.

MSIMSK Bus: 0 Device: 3 Function: 0 Offset: 6Ch			
Bit	Attr	Default	Description
31:2	RV	0h	Reserved
1:0	RW	0b	Mask Bits For each Mask bit that is set, the PCI Express port is prohibited from sending the associated message. NTB supports up to 2 messages. Corresponding bits are masked if set to '1'

6.13.3.10 MISIPENDING: MSI Pending Bit Register

The Mask Pending register enables software to defer message sending on a per-vector basis.

MISIPENDING Bus: 0 Device: 3 Function: 0 Offset: 70h			
Bit	Attr	Default	Description
31:2	RV	0h	Reserved
1:0	RO-V	0h	Pending Bits For each Pending bit that is set, the PCI Express port has a pending associated message. NTB supports up to 2 messages. Corresponding bits are pending if set to '1'



6.13.3.11 MSIXCAPID: MSI-X Capability ID Register

MSIXCAPID Bus: 0 Device: 3 Function: 0 Offset: 80h			
Bit	Attr	Default	Description
7:0	RO	11h	Capability ID Assigned by PCI-SIG for MSI-X.

6.13.3.12 MSIXNXTPTR: MSI-X Next Pointer Register

MSIXNXTPTR Bus: 0 Device: 3 Function: 0 Offset: 81h			
Bit	Attr	Default	Description
7:0	RW-O	90h	Next Ptr This field is set to 90h for the next capability list (PCI Express capability structure) in the chain.

6.13.3.13 MSIXMSGCTRL: MSI-X Message Control Register

MSIXMSGCTRL Bus: 0 Device: 3 Function: 0 Offset: 82h			
Bit	Attr	Default	Description
15	RW	0b	MSI-X Enable Software uses this bit to select between INTx or MSI or MSI-X method for signaling interrupts from the DMA 0: NTB is prohibited from using MSI-X to request service 1: MSI-X method is chosen for NTB interrupts Note: Software must disable INTx and MSI-X for this device when using MSI
14	RW	0b	Function Mask 1: all the vectors associated with the NTB are masked, regardless of the per vector mask bit state. 0: each vector's mask bit determines whether the vector is masked or not. Note: Setting or clearing the MSI-X function mask bit has no effect on the state of the per-vector Mask bit.
13:11	RV	0h	Reserved
10:0	RO-V	003h	Table Size System software reads this field to determine the MSI-X Table Size N, which is encoded as N-1. For example, a returned value of '0000000011' indicates a table size of 4. NTB table size is 4, encoded as a value of 003h

6.13.3.14 TABLEOFF_BIR: MSI-X Table Offset and BAR Indicator

TABLEOFF_BIR Bus: 0 Device: 3 Function: 0 Offset: 84h			
Bit	Attr	Default	Description
31:3	RO	000004 00h	Table Offset MSI-X Table Structure is at offset 8K from the PB01BASE address. See PXCAPID for the start of details relating to MSI-X registers.



TABLEOFF_BIR Bus: 0 Device: 3 Function: 0 Offset: 84h			
Bit	Attr	Default	Description
2:0	RO	0h	<p>Table BIR</p> <p>Indicates which one of a function's Base Address registers, located beginning at 10h in Configuration Space, is used to map the function's MSI-X Table into Memory Space.</p> <p>BIR Value Base Address register</p> <p>0: 10h 1: 14h 2: 18h 3: 1Ch 4: 20h 5: 24h 6: Reserved 7: Reserved</p> <p>For a 64-bit Base Address register, the Table BIR indicates the lower DWORD.</p>

6.13.3.15 PBAOFF_BIR: MSI-X Pending Array Offset and BAR Indicator

PBAOFF_BIR Bus: 0 Device: 3 Function: 0 Offset: 88h			
Bit	Attr	Default	Description
31:3	RO	00000600h	<p>Table Offset</p> <p>MSI-X PBA Structure is at offset 12K from the PB01BASE BAR address. See PMSICXPBA register for details.</p>
2:0	RO	0h	<p>PBA BIR</p> <p>Indicates which one of a function's Base Address registers, located beginning at 10h in Configuration Space, is used to map the function's MSI-X Table into Memory Space.</p> <p>BIR Value Base Address register</p> <p>0: 10h 1: 14h 2: 18h 3: 1Ch 4: 20h 5: 24h 6: Reserved 7: Reserved</p> <p>For a 64-bit Base Address register, the Table BIR indicates the lower DWORD.</p>



6.13.3.16 PXPCAPID: PCI Express Capability Identity Register

Device 3, Function 0, Offset 90h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.39](#).

6.13.3.17 PXPNTPTR: PCI Express Next Pointer

Device 3, Function 0, Offset 91h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.40](#).

6.13.3.18 PXPCAP: PCI Express Capabilities Register

Device 3, Function 0, Offset 92h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.41](#).

6.13.3.19 DEVCAP: PCI Express Device Capabilities Register

The PCI Express Device Capabilities register identifies device specific information for the device.

DEVCAP Bus: 0 Device: 3 Function: 0 Offset: 94h			
Bit	Attr	Default	Description
28	RO	0b	Function Level Reset Capability A value of 1b indicates the Function supports the optional Function Level Reset mechanism. NTB does not support this functionality.
27:26	RO	0h	Captured Slot Power Limit Scale Does not apply to RPs or integrated devices This value is hardwired to 00h NTB is required to be able to receive the Set_Slot_Power_Limit message without error but simply discard the Message value. Note: PCI Express Base Specification, Revision 2.0 states Components with Endpoint, Switch, or PCI Express-PCI Bridge Functions that are targeted for integration on an adapter where total consumed power is below the lowest limit defined for the targeted form factor are permitted to ignore Set_Slot_Power_Limit Messages, and to return a value of 0 in the Captured Slot Power Limit Value and Scale fields of the Device Capabilities register
25:18	RO	0h	Captured Slot Power Limit Value Does not apply to RPs or integrated devices This value is hardwired to 00h NTB is required to be able to receive the Set_Slot_Power_Limit message without error but simply discard the Message value. Note: PCI Express Base Specification, Revision 2.0 states Components with Endpoint, Switch, or PCI Express-PCI Bridge Functions that are targeted for integration on an adapter where total consumed power is below the lowest limit defined for the targeted form factor are permitted to ignore Set_Slot_Power_Limit Messages, and to return a value of 0 in the Captured Slot Power Limit Value and Scale fields of the Device Capabilities register
15	RO	1b	Role Based Error Reporting IIO is 1.1 compliant and so supports this feature
14	RO	0b	Power Indicator Present on Device Does not apply to RPs or integrated devices
13	RO	0b	Attention Indicator Present Does not apply to RPs or integrated devices
12	RO	0b	Attention Button Present Does not apply to RPs or integrated devices



DEVCAP Bus: 0 Device: 3 Function: 0 Offset: 94h			
Bit	Attr	Default	Description
11:9	RO	0b	Endpoint L1 Acceptable Latency Does not apply to IIO RCiEP (Link does not exist between host and RCiEP)
8:6	RO	0b	Endpoint L0s Acceptable Latency Does not apply to IIO RCiEP (Link does not exist between host and RCiEP)
5	RO	1b	Extended Tag Field Supported IIO devices support 8-bit tag 1: Maximum Tag field is 8 bits (NTB Mode Only) 0: Maximum Tag field is 5 bits
4:3	RO	0h	Phantom Functions Supported IIO does not support phantom functions.00b = No Function Number bits are used for Phantom Functions
2:0	RO	1h	Max Payload Size Supported IIO supports 256B payloads on PCI Express ports001b = 256 bytes max payload size.

6.13.3.20 DEVCTRL: PCI Express Device Control Register

The PCI Express Device Control register controls PCI Express specific capabilities parameters associated with the device.

DEVCTRL Bus: 0 Device: 3 Function: 0 Offset: 98h			
Bit	Attr	Default	Description
14:12	RO	000b	Max_Read_Request_Size Express/DMI ports in IIO do not generate requests greater than 128B and this field is ignored.
11	RO	0b	Enable No Snoop Not applicable since the NTB is never the originator of a TLP. This bit has no impact on forwarding of NoSnoop attribute on peer requests.
10	RO	0b	Auxiliary Power Management Enable Not applicable to IIO
9	RO	0b	Phantom Functions Enable Not applicable to IIO since it never uses phantom functions as a requester.
8	RO	0h	Extended Tag Field Enable This bit enables the PCI Express port to use an 8-bit Tag field as a requester.
7:5	RW	000b	Max Payload Size This field is set by configuration software for the maximum TLP payload size for the PCI Express port. As a receiver, the IIO must handle TLPs as large as the set value. As a requester (i.e. for requests where IIO's own RequesterID is used), it must not generate TLPs exceeding the set value. Permissible values that can be programmed are indicated by the Max_Payload_Size_Supported in the Device Capabilities register: 000: 128B max payload size 001: 256B max payload size (applies only to standard PCI Express ports and DMI port aliases to 128B) Others: alias to 128B This field is RW for PCI Express ports.
4	RO	0b	Enable Relaxed Ordering When set, NTB will forward RO bit as is from secondary to primary side. When clear, RO bit always cleared on traffic forwarded from secondary to primary



DEVCTRL Bus: 0 Device: 3 Function: 0 Offset: 98h			
Bit	Attr	Default	Description
3	RW	0b	<p>Unsupported Request Reporting Enable</p> <p>Applies only to the PCI Express RP/PCI Express NTB secondary interface/DMI ports. This bit controls the reporting of unsupported requests that IIO itself detects on requests its receives from a PCI Express/DMI port.</p> <p>0: Reporting of unsupported requests is disabled 1: Reporting of unsupported requests is enabled.</p> <p>This bit is hard-wired to 0 in NTB mode.</p> <p>NTB primary side is a RCiEP with no RC event collector.</p> <p>PCI Express Base Specification, Revision 2.0. A Root Complex Integrated Endpoint that is not associated with a Root Complex Event Collector is permitted to hardwire this bit to 0b.</p>
2	RW	0b	<p>Fatal Error Reporting Enable</p> <p>Applies only to the PCI Express RP/PCI Express NTB secondary interface/DMI ports. Controls the reporting of fatal errors that IIO detects on the PCI Express/DMI interface.</p> <p>0: Reporting of Fatal error detected by device is disabled 1: Reporting of Fatal error detected by device is enabled</p> <p>This bit is hard-wired to 0 in NTB mode.</p> <p>NTB primary side is a RCiEP with no RC event collector.</p> <p>PCI Express Base Specification, Revision 2.0 states. A Root Complex Integrated Endpoint that is not associated with a Root Complex Event Collector is permitted to hardwire this bit to 0b.</p>
1	RW	0b	<p>Non Fatal Error Reporting Enable</p> <p>Applies only to the PCI Express RP/PCI Express NTB secondary interface/DMI ports. Controls the reporting of non-fatal errors that IIO detects on the PCI Express/DMI interface.</p> <p>0: Reporting of Non Fatal error detected by device is disabled 1: Reporting of Non Fatal error detected by device is enabled</p> <p>This bit is hard-wired to 0 in NTB mode.</p> <p>NTB primary side is a RCiEP with no RC event collector.</p> <p>PCI Express Base Specification, Revision 2.0 states. A Root Complex Integrated Endpoint that is not associated with a Root Complex Event Collector is permitted to hardwire this bit to 0b.</p>
0	RW	0b	<p>Correctable Error Reporting Enable</p> <p>Applies only to the PCI Express RP/PCI Express NTB secondary interface/DMI ports. Controls the reporting of correctable errors that IIO detects on the PCI Express/DMI interface.</p> <p>0: Reporting of link Correctable error detected by the port is disabled 1: Reporting of link Correctable error detected by port is enabled</p> <p>This bit is hard-wired to 0 in NTB mode.</p> <p>NTB primary side is a RCiEP with no RC event collector.</p> <p>PCI Express Base Specification, Revision 2.0 states. A Root Complex Integrated Endpoint that is not associated with a Root Complex Event Collector is permitted to hardwire this bit to 0b.</p>

6.13.3.21 DEVSTS: PCI Express Device Status Register

The PCI Express Device Status register provides information about PCI Express device specific parameters associated with the device.

DEVSTS Bus: 0 Device: 3 Function: 0 Offset: 9Ah			
Bit	Attr	Default	Description
5	RO	0h	<p>Transactions Pending</p> <p>Does not apply to Root ports, i.e. bit hardwired to 0 for these devices.</p>




DEVSTS Bus: 0 Device: 3 Function: 0 Offset: 9Ah			
Bit	Attr	Default	Description
4	RO	0b	AUX Power Detected Does not apply to IIO.
3	RW1C	0b	Unsupported Request Detected This bit applies only to the root/DMI ports. This bit indicates that the NTB primary detected an Unsupported Request. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control Register. 1: Unsupported Request detected at the device/port. These unsupported requests are NP requests inbound that the RP received and it detected them as unsupported requests (e.g. address decoding failures that the RP detected on a packet, receiving inbound lock reads, BME bit is clear etc.). Note that this bit is not set on peer2peer completions with UR status that are forwarded by the RP to the PCIe link. 0: No unsupported request detected by the RP
2	RW1C	0b	Fatal Error Detected This bit indicates that a fatal (uncorrectable) error is detected by the NTB primary device. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. 1: Fatal errors detected 0: No Fatal errors detected
1	RW1C	0b	Non Fatal Error Detected This bit gets set if a non-fatal uncorrectable error is detected by the NTB primary device. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. 1: Non Fatal errors detected 0: No non-Fatal Errors detected
0	RW1C	0b	Correctable Error Detected This bit gets set if a correctable error is detected by the NTB primary device. Errors are logged in this register regardless of whether error reporting is enabled or not in the PCI Express Device Control register. 1: correctable errors detected 0: No correctable errors detected

6.13.3.22 PBAR23SZ: Primary BAR 2/3 Size

This register contains a value used to set the size of the memory window requested by the 64-bit BAR 2/3 pair for the Primary side of the NTB.

PBAR23SZ Bus: 0 Device: 3 Function: 0 Offset: D0h			
Bit	Attr	Default	Description
7:0	RW-O	00h	Primary BAR 2/3 Size Value indicating the size of 64-bit BAR 2/3 pair on the Primary side of the NTB. This value is loaded by BIOS prior to enumeration. The value indicates the number of bits that will be Read-Only (returning 0 when read regardless of the value written to them) during PCI enumeration. Only legal settings are 12- 39, representing BAR sizes of 212 (4KB) through 239 (512GB) are valid. Note: Programming a value of '0' or any other value other than (12-39) will result in the BAR being disabled.



6.13.3.23 PBAR45SZ: Primary BAR 4/5 Size

This register contains a value used to set the size of the memory window requested by the 64-bit BAR 4/5 pair for the Primary side of the NTB.

PBAR45SZ Bus: 0 Device: 3 Function: 0 Offset: D1h			
Bit	Attr	Default	Description
7:0	RW-O	00h	Primary BAR 4/5 Size Value indicating the size of 64-bit BAR 2/3 pair. This value is loaded by BIOS prior to enumeration. The value indicates the number of bits that will be Read-Only (returning 0 when read regardless of the value written to them) during PCI enumeration. Only legal settings are 12- 39, representing BAR sizes of 212 (4KB) through 239 (512 GB) are valid. Note: Programming a value of '0' or any other value other than (12-39) will result in the BAR being disabled.

6.13.3.24 SBAR23SZ: Secondary BAR 2/3 Size

This register contains a value used to set the size of the memory window requested by the 64-bit BAR 2/3 pair for the Secondary side of the NTB.

SBAR23SZ Bus: 0 Device: 3 Function: 0 Offset: D2h			
Bit	Attr	Default	Description
7:0	RW-O	00h	Secondary BAR 2/3 Size Value indicating the size of 64-bit BAR 2/3 pair on the Secondary side of the NTB. This value is loaded by BIOS prior to enumeration. The value indicates the number of bits that will be Read-Only (returning 0 when read regardless of the value written to them) during PCI enumeration. Only legal settings are 12- 39, representing BAR sizes of 212 (4 KB) through 239 (512 GB) are valid. Note: Programming a value of '0' or any other value other than (12-39) will result in the BAR being disabled.

6.13.3.25 SBAR45SZ: Secondary BAR 4/5 Size

This register contains a value used to set the size of the memory window requested by the 64-bit BAR 4/5 on the secondary side of the NTB.


SBAR45SZ Bus: 0 Device: 3 Function: 0 Offset: D3			
Bit	Attr	Default	Description
7:0	RW-O	00h	Secondary BAR 4/5 Size Value indicating the size of 64-bit BAR 2/3 pair on the Secondary side of the NTB. This value is loaded by BIOS prior to enumeration. The value indicates the number of bits that will be Read-Only (returning 0 when read regardless of the value written to them) during PCI enumeration. Only legal settings are 12- 39, representing BAR sizes of 212 (4 KB) through 239 (512 GB) are valid. Note: Programming a value of '0' or any other value other than (12-39) will result in the BAR being disabled.



6.13.3.26 PPD: PCIe Port Definition

This register defines the behavior of the PCIe port which can be either a RP, NTB connected to another NTB or an NTB connected to a Root Complex. This register is used to set the value in the DID register on the Primary side of the NTB (located at offset 02h). This value is loaded by BIOS prior to running PCI enumeration.

PPD			
Bus: 0 Device: 3 Function: 0 Offset: D4h			
Bit	Attr	Default	Description
6:6	RW-LV	0h	<p>bar45_32bit: If set, BARs 4 and 5 are presented as two 32b non-prefetchable BARs. If clear, BARs 4 and 5 are presented as one 64b BAR. When in this mode the following registers are also presented as 2 32b registers (vs 1 64 bit register): Bus: 0 Device: 3 Function: 0 20h pb4base 24h pb5base D1h pbar4sz D3h sbar4sz D5h pbar5sz D6h sbar5sz MMIO PBBASE32 / SBBASE32 8h pbar4lmt Ch pbar5lmt 18h pbar4xlat 1Ch pbar5xlat 28h sbar4lmt 2Ch sbar5lmt 38h sbar4xlat 3Ch sbar5xlat 50h sbar4base 54h sbar5base 520h sb4base 524h sb5base</p>
5	RW-V	0b	<p>NTB Primary side - MSI-X Single Message Vector This bit when set, causes only a single MSI-X message to be generated if MSI-X is enabled. This bit affects the default value of the MSI-X Table Size field in the Section 6.13.3.13, "MSIXMSGCTRL: MSI-X Message Control Register" on page 325.</p>
4	RO-V	0h	<p>Crosslink Configuration Status This bit is written by hardware and shows the result of the NTBCROSSLINK. 1 - NTB port is configured as USD/DSP 2 - NTB port is configured as DSD/USP</p>
3:2	RW-V	00b	<p>Crosslink Control Directly forces the polarity of the NTB port to be either an Upstream Device (USD) or Downstream Device (DSD). 11 - Force NTB port to USD/DSP; 10 - Force NTB port to DSD/USP; 01 - 00 Reserved Notes: Bits 03:02 of this register only have meaning when bits 01:00 of this same register are programmed as '01'b (NTB/NTB). When configured as NTB/RP hardware directly sets port to DSD/USP so this field is not required. When using crosslink control override, the external strap PECFGSEL[2:0] must be set to '100'b (Wait-on-BIOS). In applications that are DP configuration, and having an external controller set up the crosslink control override through the SMBus master interface. PECFGSEL[2:0] must be set to '100'b (Wait-on-BIOS) on both chipsets. The external controller on the master can then set the crosslink control override field on both chipsets and then enable the ports on both chipsets.</p>



PPD Bus: 0 Device: 3 Function: 0 Offset: D4h			
Bit	Attr	Default	Description
1:0	RW-V	00b	Port Definition Value indicating the value to be loaded into the DID register (offset 02h). 00b - Transparent bridge 01b - 2 NTBs connected back to back 10b - NTB connected to a RP 11b - Reserved Note: When SKU does not support this field becomes RO '00'

6.13.3.27 PMCAP: Power Management Capabilities


Device 3, Function 0, Offset E0h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.59](#).



6.13.3.28 PMCSR: Power Management Control and Status

This register provides status and control information for Power Management events in the PCI Express port of the IIO.

PMCSR Bus: 0 Device: 3 Function: 0 Offset: E4h			
Bit	Attr	Default	Description
31:24	RO	00h	Data Not relevant for IIO
23	RO	0h	Bus Power/Clock Control Enable This field is hardwired to 0h as it does not apply to PCI Express.
22	RO	0h	B2/B3 Support This field is hardwired to 0h as it does not apply to PCI Express.
21:16	RV	0h	Reserved
15	RW1CS	0h	PME Status Applies only to root ports. This PME Status is a sticky bit. This bit is set, independent of the PME Enable bit defined below, on an enabled PCI Express hot-plug event. Software clears this bit by writing a '1' when it has been completed. NTB Mode: This bit is hard-wired to read-only 0, since this function does not support PME# generation from any power state.
14:13	RO	0h	Data Scale Not relevant for IIO
12:9	RO	0h	Data Select Not relevant for IIO
8	RWS	0h	PME Enable Applies only to root ports. This field is a sticky bit and when set, enables a virtual PM_PME message to be generated internally on an enabled PCI Express hot-plug event. This virtual PM_PME message then sets the appropriate bits in the ROOTSTS register (which can then trigger an MSI/INT or cause a _PMEGPE event). 0: Disable ability to send PME messages when an event occurs 1: Enables ability to send PME messages when an event occurs Not used in NTB mode.
7:4	RV	0h	Reserved
3	RW-O	1b	No Soft Reset Indicates IIO does not reset its registers when it transitions from D3hot to D0.
2	RV	0h	Reserved
1:0	RW-V	0h	Power State This 2-bit field is used to determine the current power state of the function and to set a new power state as well. 00: D0 01: D1 (not supported by IIO) 10: D2 (not supported by IIO) 11: D3_hot If Software tries to write 01 or 10 to this field, the power state does not change from the existing power state (which is either D0 or D3hot) and nor do these bits: 0 change value. All devices will respond to only Type 0 configuration transactions when in D3hot state (RP will not forward Type 1 accesses to the downstream link) and will not respond to memory/IO transactions (i.e. D3hot state is equivalent to MSE/IOSE bits being clear) as target and will not generate any memory/IO/configuration transactions as initiator on the primary bus (messages are still allowed to pass through).



6.13.3.29 XPREUT_HDR_EXT: REUT PCIe Header Extended

Device 3, Function 0, Offset 100h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.61](#).

6.13.3.30 XPREUT_HDR_CAP: REUT Header Capability

Device 3, Function 0, Offset 104h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.62](#).

6.13.3.31 XPREUT_HDR_LEF: REUT Header Leaf Capability

Device 3, Function 0, Offset 108h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.63](#).

6.13.3.32 ACSCAPHDR: Access Control Services Extended Capability Header

Device 3, Function 0, Offset 110h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.64](#).

6.13.3.33 ACSCAP: Access Control Services Capability Register

Device 3, Function 0, Offset 114h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.65](#).

6.13.3.34 ACCTRL: Access Control Services Control Register

Device 3, Function 0, Offset 116h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.66](#).

6.13.3.35 APICBASE: APIC Base Register

Device 3, Function 0, Offset 140h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.67](#).

6.13.3.36 APICLIMIT: APIC Limit Register

Device 3, Function 0, Offset 142h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.68](#).

6.13.3.37 VSECPHDR: Vendor Specific Enhanced Capability Header

Device 3, Function 0, Offset 144h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.69](#).

6.13.3.38 VSHDR: Vender Specific Header

This register identifies the capability structure and points to the next structure.

VSHDR Bus: 0 Device: 3 Function: 0 Offset: 148			
Bit	Attr	Default	Description
31:20	RO	03Ch	VSEC Length This field indicates the number of bytes in the entire VSEC structure, including the PCI Express Enhanced Capability header, the Vendor-Specific header, and the Vendor-Specific Registers.



VSHDR Bus: 0 Device: 3 Function: 0 Offset: 148			
Bit	Attr	Default	Description
19:16	RO	1h	VSEC Version Set to 1h for this version of the PCI Express logic
15:0	RO	0004h	VSEC ID Identifies Intel Vendor Specific Capability for AER on NTB

6.13.3.39 UNCERRSTS: Uncorrectable Error Status

Device 3, Function 0, Offset 14Ch. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.72](#).

6.13.3.40 UNCERRMSK: Uncorrectable Error Mask

Device 3, Function 0, Offset 150h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.73](#).

6.13.3.41 UNCERRSEV: Uncorrectable Error Severity

Device 3, Function 0, Offset 154h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.74](#).

6.13.3.42 CORERRSTS: Correctable Error Status

Device 3, Function 0, Offset 158h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.75](#).

6.13.3.43 CORERRMSK: Correctable Error Mask

Device 3, Function 0, Offset 15Ch. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.76](#).

6.13.3.44 ERRRCAP: Advanced Error Capabilities and Control

Device 3, Function 0, Offset 160h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.77](#).


6.13.3.45 HDRLOG[0:3]: Header Log 0

This register contains the header log when the first error occurs. Headers of the subsequent errors are not logged.

HDRLOG[0:3] Bus: 0 Device: 3 Function: 0 Offset: 164, 168, 16C, 170			
Bit	Attr	Default	Description
31:0	ROS-V	000000 00h	Log of Header Dword Logs the respective DWORD of the header on an error condition

6.13.3.46 RPERRCMD: Root Port Error Command

Device 3, Function 0, Offset 174h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.79](#).



6.13.3.47 RPERRSTS: Root Port Error Status

The Root Error Status register reports status of error Messages (ERR_COR, ERR_NONFATAL, and ERR_FATAL) received by the Root Complex in IIO, and errors detected by the RP itself (which are treated conceptually as if the RP had sent an error Message to itself). The ERR_NONFATAL and ERR_FATAL Messages are grouped together as uncorrectable. Each correctable and uncorrectable (Non-fatal and Fatal) error source has a first error bit and a next error bit associated with it respectively. When an error is received by a Root Complex, the respective first error bit is set and the Requestor ID is logged in the Error Source Identification register. A set individual error status bit indicates that a particular error category occurred; software may clear an error status by writing a 1 to the respective bit. If software does not clear the first reported error before another error Message is received of the same category (correctable or uncorrectable), the corresponding next error status bit will be set but the Requestor ID of the subsequent error Message is discarded. The next error status bits may be cleared by software by writing a 1 to the respective bit as well.

RPERRSTS Bus: 0 Device: 3 Function: 0 Offset: 178			
Bit	Attr	Default	Description
31:27	RO	0h	Advanced Error Interrupt Message Number Advanced Error Interrupt Message Number offset between base message data and the MSI/MSI-X message if assigned more than one message number. IIO hardware automatically updates this register to 0x1h if the number of messages allocated to the RP is 2. See bit 6:4 in Section 6.13.5.22, "MSICTRL: MSI Control" on page 369 for details of the number of messages allocated to a RP.
6	RW1CS	0b	Fatal Error Messages Received Set when one or more Fatal Uncorrectable error Messages have been received.
5	RW1CS	0b	Non-Fatal Error Messages Received Set when one or more Non-Fatal Uncorrectable error Messages have been received.
4	RW1CS	0b	First Uncorrectable Fatal Set when bit 2 is set (from being clear) and the message causing bit 2 to be set is an ERR_FATAL message.
3	RW1CS	0b	Multiple Error Fatal/Nonfatal Received Set when either a fatal or a non-fatal error message is received and Error Fatal/Nonfatal Received is already set, i.e log from the 2nd Fatal or No fatal error message onwards
2	RW1CS	0b	Error Fatal/Nonfatal Received Set when either a fatal or a non-fatal error message is received and this bit is already not set. i.e. log the first error message. Note that when this bit is set bit 3 could be either set or clear.
1	RW1CS	0b	Multiple Correctable Error Received Set when either a correctable error message is received and Correctable Error Received bit is already set, i.e log from the 2nd Correctable error message onwards
0	RW1CS	0b	Correctable Error Received Set when a correctable error message is received and this bit is already not set. That is, log the first error message

6.13.3.48 ERRSID: Error Source Identification

Device 3, Function 0, Offset 17Ch. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.81](#).




6.13.3.49 PERFCTRLSTS: Performance Control and Status

PERFCTRLSTS Bus: 0 Device: 3 Function: 0 Offset: 180			
Bit	Attr	Default	Description
45:44	RW	0b	<p>full_cohrd_op: Selects the opcode used on the ring by coherent reads issued on behalf of this root port that need all 64B of data from the requested cacheline.</p> <p>'00 = PCIRdCur '01 = CRd '10 = DRd '11 = RFO</p>
43:42	RW	0b	<p>partial_cohrd_op: Selects the opcode used on the ring by coherent reads issued on behalf of this root port that need less than 64B of data from the requested cacheline.</p> <p>'00 = PCIRdCur '01 = CRd '10 = DRd '11 = RFO</p>
41	RW	0b	<p>TLP Processing Hint Disable When set, writes or reads with TPH=1, will be treated as if TPH=0.</p>
40	RW	0b	<p>DCA Requester ID Override When this bit is set, Requester ID match for DCA writes is bypassed. All writes from the port are treated as DCA writes and the tag field will convey if DCA is enabled or not and the target information.</p>
35	RW	0b	Max read request completion combining size
20:16	RW	18h	<p>Outstanding Requests for Gen1 Number of outstanding RFOs and non-posted requests from a given PCIe port. This register controls the number of outstanding inbound non-posted requests - I/O, Config, Memory - (maximum length of these requests is a single 64B cacheline) that a Gen1 PCI Express downstream port can have. This register provides the value for the port when it is operating in Gen1 mode and for a link width of x4. The value of this parameter for the port when operating in x8/x16 width is obtained by multiplying this register by 2 and 4 respectively. BIOS programs this register based on the read latency to main memory.</p> <p>This register also specifies the number of RFOs that can be kept outstanding on IDI for a given port.</p> <p>The link speed of the port can change during a PCI Express hot-plug event and the port must use the appropriate multiplier.</p> <p>A value of 1 indicates one outstanding pre-allocated request, 2 indicates two outstanding pre-allocated requests, and so on. If software programs a value greater than the buffer size the DMA engine supports, then the maximum hardware supported value is used.</p> <p>Current BIOS recommendation is to leave this field at it's default value.</p>



PERFCTRLSTS Bus: 0 Device: 3 Function: 0 Offset: 180			
Bit	Attr	Default	Description
13:8	RW	30h	<p>Outstanding Requests for Gen2</p> <p>Number of outstanding RFOs and non-posted requests from a given PCIe port. This register controls the number of outstanding inbound non-posted requests - I/O, Config, Memory - (maximum length of these requests is a single 64B cacheline) that a Gen2 PCI Express downstream port can have. This register provides the value for the port when it is operating in Gen2 mode and for a link width of x4. The value of this parameter for the port when operating in x8/x16 width is obtained by multiplying this register by 2 and 4 respectively. BIOS programs this register based on the read latency to main memory. For a port operating in Gen3 mode, a multiplier of x2 is applied.</p> <p>This register also specifies the number of RFOs that can be kept outstanding on IDI for a given port.</p> <p>The link speed of the port can change during a PCI Express hot-plug event and the port must use the appropriate multiplier.</p> <p>A value of 1 indicates one outstanding pre-allocated request, 2 indicates two outstanding pre-allocated requests, and so on. If software programs a value greater than the buffer size the DMA engine supports, then the maximum hardware supported value is used.</p> <p>Current BIOS recommendation is to leave this field at it's default value.</p>
7	RW	1b	<p>Use Allocating Flows for 'Normal Writes'</p> <p>1: Use allocating flows for the writes that meet the following criteria. 0: Use non-allocating flows for writes that meet the following criteria (TPH=0 OR TPHDIS=1 OR (TPH=1 AND Tag=0 AND CIPCTRL[28]=1)) AND (NS=0 OR NoSnoopOpWrEn=0) AND Non-DCA Write</p> <p>Notes: When allocating flows are used for the above write types, IIO does not send a Prefetch Hint message. Current recommendation for BIOS is to just leave this bit at default of 1b. Note there is a coupling between the usage of this bit and bits 2 and 3. TPHDIS is bit 0 of this register NoSnoopOpWrEn is bit 3 of this register</p>
4	RW	1b	Read Stream Interleave Size
3	RW	0b	<p>Enable No-Snoop Optimization on Writes</p> <p>This applies to writes with the following conditions: NS=1 AND (TPH=0 OR TPHDIS=1)</p> <p>1: Inbound writes to memory with above conditions will be treated as non-coherent (no snoops) writes on Intel QPI 0: Inbound writes to memory with above conditions will be treated as allocating or non-allocating writes, depending on bit 4 in this register.</p> <p>Notes: If TPH=1 and TPHDIS=0 then NS is ignored and this bit is ignored Current recommendation for BIOS is to just leave this bit at default of 0b.</p>
2	RW	0b	<p>Enable No-Snoop Optimization on Reads</p> <p>This applies to reads with the following conditions: NS=1 AND (TPH=0 OR TPHDIS=1)</p> <p>1: When the condition is true for a given inbound read request to memory, it will be treated as non-coherent (no snoops) reads on Intel QPI. 0: When the condition is true for a given inbound read request to memory, it will be treated as normal snooped reads from PCIe (which trigger a PCI RdCurrent or DRd.UC on IDI).</p> <p>Notes: If TPH=1 and TPHDIS=0 then NS is ignored and this bit is ignored Current recommendation for BIOS is to just leave this bit at default of 0b.</p>
1	RW	0b	Disable reads bypassing other reads
0	RW	1b	Read Stream Policy



6.13.3.50 MISCTRLSTS: Misc. Control and Status

MISCTRLSTS			
Bus: 0 Device: 3 Function: 0 Offset: 188			
Bit	Attr	Default	Description
49	RW1CS	0b	Locked read timed out Indicates that a locked read request incurred a completion time-out on PCI Express/DMI
48	RW1C	0b	Received PME_TO_ACK Indicates that IIO received a PME turn off ack packet or it timed out waiting for the packet
41	RW	0b	Override SocketID in Completion ID For TPH/DCA requests, the Completer ID can be returned with SocketID when this bit is set.
38	RW	0b	'Problematic Port' for Lock Flows This bit is set by BIOS when it knows that this port is connected to a device that creates Posted-Posted dependency on its In-Out queues. Briefly, this bit is set on a link if: This link is connected to a Processor RP or Processor NTB port on the other side of the link IIO lock flows depend on the setting of this bit to treat this port in a special way during the flows. Note that if BIOS is setting up the lock flow to be in the 'Intel QPI compatible' mode then this bit must be set to 0.
36	RWS	0b	Form-Factor Indicates what form-factor a particular root port controls 0 - CEM 1 - Express Module This bit is used to interpret bit 6 in the VPP serial stream for the port as either MRL# (CEM) input or EMLSTS# (Express Module) input.
35	RW	0b	Override System Error on PCIe Fatal Error Enable When set, fatal errors on PCI Express (that have been successfully propagated to the primary interface of the port) are sent to the IIO core error logic (for further escalation) regardless of the setting of the equivalent bit in the ROOTCTRL register. When clear, the fatal errors are only propagated to the IIO core error logic if the equivalent bit in "ROOTCON: PCI Express Root Control" register is set. For Device #0 in DMI mode and Device #3/Fn#0, unless this bit is set, DMI/NTB link related fatal errors will never be notified to system software.
34	RW	0b	Override System Error on PCIe Non-fatal Error Enable When set, non-fatal errors on PCI Express (that have been successfully propagated to the primary interface of the port) are sent to the IIO core error logic (for further escalation) regardless of the setting of the equivalent bit in the ROOTCTRL register. When clear, the non-fatal errors are only propagated to the IIO core error logic if the equivalent bit in "ROOTCON: PCI Express Root Control" register is set. For Device #0 in DMI mode and Dev#3/Fn#0, unless this bit is set, DMI/NTB link related non-fatal errors will never be notified to system software.
33	RW	0b	Override System Error on PCIe Correctable Error Enable When set, correctable errors on PCI Express (that have been successfully propagated to the primary interface of the port) are sent to the IIO core error logic (for further escalation) regardless of the setting of the equivalent bit in the ROOTCTRL register. When clear, the correctable errors are only propagated to the IIO core error logic if the equivalent bit in "ROOTCON: PCI Express Root Control" register is set. For Dev#0 in DMI mode and Dev#3/Fn#0, unless this bit is set, DMI/NTB link related correctable errors will never be notified to system software.



MISCCTRLSTS			
Bus: 0 Device: 3 Function: 0 Offset: 188			
Bit	Attr	Default	Description
32	RW	0b	ACPI PME Interrupt Enable When set, Assert/Deassert_PMEGPE messages are enabled to be generated when ACPI mode is enabled for handling PME messages from PCI Express. When this bit is cleared (from a 1), a Deassert_PMEGPE message is scheduled on behalf of the root port if an Assert_PMEGPE message was sent last from the root port. When NTB is enabled on Dev#3/Fn#0 this bit is meaningless because PME messages are not expected to be received on the NTB link.
31	RW	0b	Disable LOs on transmitter When set, IIO never puts its tx in L0s state, even if OS enables it via the Link Control register. LOs is not supported.
29	RW	1b	cfg_to_en Disables/enables config timeouts, independently of other timeouts.
28	RW	0b	to_dis Disables timeouts completely.
27	RWS	0b	System Interrupt Only on Link BW/Management Status This bit, when set, will disable generating MSI and Intx interrupts on link bandwidth (speed and/or width) and management changes, even if MSI or INTx is enabled i.e. will disable generating MSI or INTx when LNKSTS bits 15 and 14 are set. Whether or not this condition results in a system event like Intel SMI/PMI/CPEI is dependent on whether this event masked or not in the XPCORERRMSK register. Note that when Dev#3 is operation in NTB mode, this bit still applies and BIOS needs to do the needful if it wants to enable/disable these events from generating MSI/INTx interrupts from the NTB device.
26	RW-LV	0b	EOI Forwarding Disable - Disable EOI broadcast to this PCIe link When set, EOI message will not be broadcast down this PCIe link. When clear, the port is a valid target for EOI broadcast. BIOS must set this bit on a port if it is connected to a another cpu NTB or root port on other end of the link.
24	RW	0b	Peer2peer Memory Read Disable When set, peer2peer memory reads are master aborted otherwise they are allowed to progress per the peer2peer decoding rules.
23	RW	0b	Phold Disable Applies only to Dev#0 When set, the IIO responds with Unsupported request on receiving assert_phold message from PCH and results in generating a fatal error.
22	RWS	0b	check_cpl_tc
21	RW-O	0b	Force Outbound TC to Zero Forces the TC field to zero for outbound requests. 1: TC is forced to zero on all outbound transactions regardless of the source TC value 0: TC is not altered Note: In DMI mode, TC is always forced to zero and this bit has no effect.
20	RWS	0b	maltlp_32baddr64bhdr_en Malformed TLP 32b address in 64b header Enable. When set, enables reporting a Malformed packet when the TLP is a 32 bit address in a 4DW header. PCI Express forbids using 4DW header sizes when the address is less than 4GB, but some cards may use the 4DW header anyway. In these cases, the upper 32 bits of address are all 0.
18	RWS	1b	Max Read Completion Combine Size This bit when set, will enable completion combining to a maximum of 256B (values less than or equal to 256B allowed). When clear, the maximum read completion combining size is 128B (values less than or equal to 256B allowed).
17	RO	0b	Force Data Parity Error
16	RO	0b	Force EP Bit Error Poison Bit
15	RWS	0b	dis_hdr_storage
14	RWS	0b	allow_one_np_os



MISCTRLSTS			
Bus: 0 Device: 3 Function: 0 Offset: 188			
Bit	Attr	Default	Description
13	RWS	0b	tlp_on_any_lane
12	RWS	1b	disable_ob_parity_check
9	RWS	0b	dispdspolling Disables gen2 if timeout happens in polling.cfg.
8:7	RW	0b	PME2ACKTOCTRL
6	RW	0b	Enable timeout for receiving PME_TO_ACK When set, IIO enables the timeout to receiving the PME_TO_ACK
5	RW	0b	Send PME_TURN_OFF message When this bit is written with a 1b, IIO sends a PME_TURN_OFF message to the PCIe link. Hardware clears this bit when the message has been sent on the link.
4	RW	0b	Enable System Error only for AER Applies only to root/NTB ports. For Dev#0 in DMI mode, this bit is to be left at default value always. When this bit is set, the PCI Express errors do not trigger an MSI or Intx interrupt, regardless of the whether MSI or INTx is enabled or not. Whether or not PCI Express errors result in a system event like NMI/Intel SMI/PMI/CPEI is dependent on whether the appropriate system error or override system error enable bits are set or not. When this bit is clear, PCI Express errors are reported via MSI or INTx and/or NMI/Intel SMI/MCA/CPEI. When this bit is clear and if MSI enable bit in the Section 6.13.5.22, "MSICTRL: MSI Control" on page 369 is set (clear), then an MSI (INTx) interrupt is generated for PCI Express errors. When this bit is clear, and 'System Error on Fatal Error Enable' bit in ROOTCON register is set, then NMI/Intel SMI/MCA is (also) generated for a PCI Express fatal error. Similar behavior for non-fatal and corrected errors. Note that this bit applies to Dev#3/Fn#0 in NTB mode as well and BIOS needs to set it up appropriately in that mode.
3	RW	0b	Enable ACPI mode for Hot-plug Applies only to root ports. For Dev#0 in DMI mode, this bit is to be left at default value always. When this bit is set, all Hot Plug events from the PCI Express port are handled via _HPGPE messages to the PCH and no MSI/INTx messages are ever generated for Hot Plug events (regardless of whether MSI or INTx is enabled at the root port or not) at the root port. When this bit is clear, _HPGPE message generation on behalf of root port Hot Plug events is disabled and OS can chose to generate MSI or INTx interrupt for Hot Plug events, by setting the MSI enable bit in root ports. This bit does not apply to the DMI ports. Clearing this bit (from being 1) schedules a Deassert_HPGPE event on behalf of the root port, provided there was any previous Assert_HPGPE message that was sent without an associated Deassert message. Note that this bit applies to Dev#3/Fn#0 in NTB mode as well and BIOS needs to set it up appropriately in that mode.
2	RW	0b	Enable ACPI mode for Power Management Applies only to root ports. For Dev#0 in DMI mode, this bit is to be left at default value always. When this bit is set, all Power Management events at the PCI Express port are handled via _PMEGPE messages to the PCH, and no MSI interrupts are ever generated for Power Management events at the root port (regardless of whether MSI is enabled at the root port or not). When clear, _PMEGPE message generation for Power Management events is disabled and OS can chose to generate MSI interrupts for delivering Power Management events by setting the MSI enable bit in root ports. This bit does not apply to the DMI ports. Clearing this bit (from being 1) schedules a Deassert_PMEGPE event on behalf of the root port, provided there was any previous Assert_PMEGPE message that was sent without an associated Deassert message. Note that this bit applies to Dev#3/Fn#0 in NTB mode as well and BIOS needs to set it up appropriately in that mode.
1	RW-O	0h	inbound_configuration_enable: Enable Inbound Configuration Requests.



6.13.3.51 PCIE_IOU_BIF_CTRL: PCIe IOU Bifurcation Control

PCIE_IOU_BIF_CTRL Bus: 0 Device: 3 Function: 0 Offset: 190			
Bit	Attr	Default	Description
15:4	RV	0h	Reserved
3	WO	0b	<p>IOU Start Bifurcation</p> <p>When software writes a 1 to this bit, IIO starts the port 0 bifurcation process. After writing to this bit, software can poll the Data Link Layer link active bit in the LNKSTS register to determine if a port is up and running. Once a port bifurcation has been initiated by writing a 1 to this bit, software cannot initiate any more write-1 to this bit (write of 0 is ok).</p> <p>Note: That this bit can be written to a 1 in the same write that changes values for bits 2:0 in this register and in that case, the new value from the write to bits 2:0 take effect. This bit always reads a 0b.</p>
2:0	RWS	100b	<p>IOU Bifurcation Control</p> <p>To select a IOU bifurcation, software sets this field and then either</p> <ul style="list-style-type: none"> a) sets bit 3 in this register to initiate training OR b) resets the entire CPU and on exit from that reset, <p>CPU will bifurcate the ports per the setting in this field.</p> <p>For Device 1 Function 0:</p> <p>000: x4x4 (operate lanes 7:4 as x4, 3:0 as x4)</p> <p>001: x8</p> <p>others: Reserved</p> <p>For Device 2 Function 0:</p> <p>000: x4x4x4x4 (operate lanes 15:12 as x4, 11:8 as x4, 7:4 as x4 and 3:0 as x4)</p> <p>001: x4x4x8 (operate lanes 15:12 as x4, 11:8 as x4 and 7:0 as x8)</p> <p>010: x8x4x4 (operate lanes 15:8 as x8, 7:4 as x4 and 3:0 as x4)</p> <p>011: x8x8 (operate lanes 15:8 as x8, 7:0 as x8)</p> <p>100: x16</p> <p>others: Reserved</p>

6.13.3.52 NTBDEVCAP: PCI Express Device Capabilities

The PCI Express Device Capabilities register identifies device specific information for the device.

NTBDEVCAP Bus: 0 Device: 3 Function: 0 Offset: 194h			
Bit	Attr	Default	Description
31:29	RV	0h	Reserved
28	RO	0b	<p>Function Level Reset Capability</p> <p>A value of 1b indicates the Function supports the optional Function Level Reset mechanism. NTB does not support this functionality.</p>
27:26	RO	0h	<p>Captured Slot Power Limit Scale</p> <p>Does not apply to RPs or integrated devices This value is hardwired to 00h NTB is required to be able to receive the Set_Slot_Power_Limit message without error but simply discard the Message value.</p> <p>Note: PCI Express Base Specification, Revision 2.0 states Components with Endpoint, Switch, or PCI Express-PCI Bridge Functions that are targeted for integration on an adapter where total consumed power is below the lowest limit defined for the targeted form factor are permitted to ignore Set_Slot_Power_Limit Messages, and to return a value of 0 in the Captured Slot Power Limit Value and Scale fields of the Device Capabilities register.</p>



NTBDEVCAP Bus: 0 Device: 3 Function: 0 Offset: 194h			
Bit	Attr	Default	Description
25:18	RO	00h	<p>Captured Slot Power Limit Value Does not apply to RPs or integrated devices This value is hardwired to 00h NTB is required to be able to receive the Set_Slot_Power_Limit message without error but simply discard the Message value.</p> <p>Note: PCI Express Base Specification, Revision 2.0 states Components with Endpoint, Switch, or PCI Express-PCI Bridge Functions that are targeted for integration on an adapter where total consumed power is below the lowest limit defined for the targeted form factor are permitted to ignore Set_Slot_Power_Limit Messages, and to return a value of 0 in the Captured Slot Power Limit Value and Scale fields of the Device Capabilities register.</p>
17:16	RV	0h	Reserved
15	RO	1b	<p>Role Based Error Reporting IIO is 1.1 compliant and so supports this feature</p>
14	RO	0b	<p>Power Indicator Present on Device Does not apply to RPs or integrated devices</p>
13	RO	0b	<p>Attention Indicator Present Does not apply to RPs or integrated devices</p>
12	RO	0b	<p>Attention Button Present Does not apply to RPs or integrated devices</p>
11:9	RW-O	110b	<p>Endpoint L1 Acceptable Latency This field indicates the acceptable latency that an Endpoint can withstand due to the transition from L1 state to the L0 state. It is essentially an indirect measure of the Endpoint's internal buffering. Power management software uses the reported L1 Acceptable Latency number to compare against the L1 Exit Latencies reported (see below) by all components comprising the data path from this Endpoint to the Root Complex Root Port to determine whether ASPM L1 entry can be used with no loss of performance.</p> <p>Defined encodings are: 000: Maximum of 1 us 001: Maximum of 2 us 010: Maximum of 4 us 011: Maximum of 8 us 100: Maximum of 16 us 101: Maximum of 32 us 110: Maximum of 64 us 111: No limit</p> <p>Note: BIOS programs this value</p>
8:6	RW-O	000b	<p>Endpoint L0s Acceptable Latency This field indicates the acceptable total latency that an Endpoint can withstand due to the transition from L0s state to the L0 state. It is essentially an indirect measure of the Endpoint's internal buffering. Power management software uses the reported L0s Acceptable Latency number to compare against the L0s exit latencies reported by all components comprising the data path from this Endpoint to the Root Complex Root Port to determine whether ASPM L0s entry can be used with no loss of performance.</p> <p>Defined encodings are: 000: Maximum of 64 ns 001: Maximum of 128 ns 010: Maximum of 256 ns 011: Maximum of 512 ns 100: Maximum of 1 us 101: Maximum of 2 us 110: Maximum of 4 us 111: No limit</p> <p>Note: BIOS programs this value</p>



NTBDEVCAP Bus: 0 Device: 3 Function: 0 Offset: 194h			
Bit	Attr	Default	Description
5	RO	1b	Extended Tag Field Supported IIO devices support 8-bit tag1 = Maximum Tag field is 8 bits 0 = Maximum Tag field is 5 bits
4:3	RO	00b	Phantom Functions Supported IIO does not support phantom functions. 00b = No Function Number bits are used for Phantom Functions
2:0	RO	001b	Max Payload Size Supported IIO supports 256B payloads on PCI Express ports 001b = 256 bytes max payload size

6.13.3.53 LNKCAP: PCI Express Link Capabilities

The Link Capabilities register identifies the PCI Express specific link capabilities. The link capabilities register needs some default values setup by the local host. This register is relocated to the enhanced configuration space region in while in NTB mode.

LNKCAP Bus: 0 Device: 3 Function: 0 Offset: 19Ch			
Bit	Attr	Default	Description
31:24	RW-O	00h	Port Number This field indicates the PCI Express port number for the link and is initialized by software/BIOS. NOTE: This register bit is a RW-O register from the host side. It must be loaded by BIOS in the primary side equivalent register. This register is RO from the secondary side of the NTB.
21	RO	1b	Link Bandwidth Notification Capability A value of 1b indicates support for the Link Bandwidth Notification status and interrupt mechanisms.
20	RO	1b	Data Link Layer Link Active Reporting Capable IIO supports reporting status of the data link layer so software knows when it can enumerate a device on the link or otherwise know the status of the link.
19	RO	0b	Surprise Down Error Reporting Capable IIO supports reporting a surprise down error condition
18	RO	0b	Clock Power Management Does not apply to CPU
17:15	RW-O	010b	L1 Exit Latency This field indicates the L1 exit latency for the given PCI Express port. It indicates the length of time this port requires to complete transition from L1 to L0. 000: Less than 1 us 001: 1 us to less than 2 us 010: 2 us to less than 4 us 011: 4 us to less than 8 us 100: 8 us to less than 16 us 101: 16 us to less than 32 us 110: 32 us to 64 us 111: More than 64 us Notes: This register bit is a RW-O register from the host side. It must be loaded by BIOS in the primary side equivalent register. This register is RO from the secondary side of the NTB.



LNKCAP Bus: 0 Device: 3 Function: 0 Offset: 19Ch			
Bit	Attr	Default	Description
14:12	RW-O	011b	<p>L0s Exit Latency</p> <p>This field indicates the L0s exit latency (i.e L0s to L0) for the PCI Express port.000: Less than 64 ns 001: 64 ns to less than 128 ns 010: 128 ns to less than 256 ns 011: 256 ns to less than 512 ns 100: 512 ns to less than 1 us 101: 1 us to less than 2 us 110: 2 us to 4 us 111: More than 4 us</p> <p>Note: This register bit is a RW-O register from the host side. It must be loaded by BIOS in the primary side equivalent register. This register is RO from the secondary side of the NTB.</p>
11:10	RW-O	11b	<p>Active State Link Power Management Support</p> <p>This field indicates the level of active state power management supported on the given PCI Express port.00: Disabled 01: L0s Entry Supported 10: Reserved 11: L0s and L1 Supported</p> <p>Note: This register bit is a RW-O register from the host side. It must be loaded by BIOS in the primary side equivalent register. This register is RO from the secondary side of the NTB.</p>
9:4	RW-O	4h	<p>Maximum Link Width</p> <p>This field indicates the maximum width of the given PCI Express Link attached to the port.000001: x1 000010: x2 000100: x4 001000: x8 010000: x16 Others - Reserved</p> <p>Note: This register bit is a RW-O register from the host side. It must be loaded by BIOS in the primary side equivalent register. This register is RO from the secondary side of the NTB.</p>
3:0	RW-O	0011b	<p>Maximum Link Speed</p> <p>This field indicates the maximum link speed of this Port.</p> <p>The encoding is the binary value of the bit location in the Supported Link Speeds Vector (in LNKCAP2) that corresponds to the maximum link speed.</p> <p>NTB Device 3 supports a maximum of 8Gbps, unless restricted by SKU.</p> <p>0001b = (2.5Gbps) 0010b = (5Gbps) 0011b = (8Gbps)</p>

6.13.3.54 LNKCON: PCI Express Link Control

The PCI Express Link Control register controls the PCI Express Link specific parameters. The link control register needs some default values setup by the local host. This register is relocated to the enhanced configuration space region in while in NTB mode. In NTB/RP mode RP will program this register. In NTB/NTB mode local host BIOS will program this register

LNKCON Bus: 0 Device: 3 Function: 0 Offset: 1A0h			
Bit	Attr	Default	Description
15:12	RV	0h	Reserved



LNKCON Bus: 0 Device: 3 Function: 0 Offset: 1A0h			
Bit	Attr	Default	Description
11	RW	0b	Link Autonomous Bandwidth Interrupt Enable For root ports, when set to 1b this bit enables the generation of an interrupt to indicate that the Link Autonomous Bandwidth Status bit has been set. For DMI mode on Dev#0, interrupt is not supported and hence this bit is not useful. Expectation is that BIOS will set bit 27 in Section 6.2.84 to notify the system of autonomous BW change event on that port.
10	RW	0b	Link Bandwidth Management Interrupt Enable For root ports, when set to 1b this bit enables the generation of an interrupt to indicate that the Link Bandwidth Management Status bit has been set. For DMI mode on Dev#0, interrupt is not supported and hence this bit is not useful. Expectation is that BIOS will set bit 27 in Section 6.2.84 to notify the system of autonomous BW change event on that port.
9	RW	0b	Hardware Autonomous Width Disable When Set, this bit disables hardware from changing the Link width for reasons other than attempting to correct unreliable Link operation by reducing Link width. Note that IIO does not by itself change width for any reason other than reliability. So this bit only disables such a width change as initiated by the device on the other end of the link.
8	RO	0b	Enable Clock Power Management N/A to CPU
7	RW	0b	Extended Synch This bit when set forces the transmission of additional ordered sets when exiting L0s and when in recovery. See PCI Express Base Specification, Revision 2.0 for details.
6	RW	0b	Common Clock Configuration IIO does nothing with this bit
5	WO	0b	Retrain Link A write of 1 to this bit initiates link retraining in the given PCI Express/DMI port by directing the LTSSM to the recovery state if the current state is [L0, L0s or L1]. If the current state is anything other than L0, L0s, L1 then a write to this bit does nothing. This bit always returns 0 when read. It is permitted to write 1b to this bit while simultaneously writing modified values to other fields in this register. If the LTSSM is not already in Recovery or Configuration, the resulting Link training must use the modified values. If the LTSSM is already in Recovery or Configuration, the modified values are not required to affect the Link training that's already in progress.
4	RW	0b	Link Disable This field controls whether the link associated with the PCI Express/DMI port is enabled or disabled. When this bit is a 1, a previously configured link would return to the 'disabled' state as defined in the PCI Express Base Specification, Revision 2.0. When this bit is clear, an LTSSM in the 'disabled' state goes back to the detect state. 0: Enables the link associated with the PCI Express port 1: Disables the link associated with the PCI Express port
3	RO	0b	Read Completion Boundary Set to zero to indicate IIO could return read completions at 64B boundaries
1:0	RW-V	00b	Active State Link Power Management Control When 01b or 11b, L0s on transmitter is enabled, otherwise it is disabled. 10 and 11 enables L1 ASPM.

6.13.3.55 LNKSTS: PCI Express Link Status Register


The PCI Express Link Status register provides information on the status of the PCI Express Link such as negotiated width, training etc. The link status register needs some default values setup by the local host. This register is relocated to the enhanced configuration space region in while in NTB mode.



LNKSTS Bus: 0 Device: 3 Function: 0 Offset: 1A2h			
Bit	Attr	Default	Description
15	RW1C	0b	<p>Link Autonomous Bandwidth Status</p> <p>This bit is set to 1b by hardware to indicate that hardware has autonomously changed link speed or width, without the port transitioning through DL_Down status, for reasons other than to attempt to correct unreliable link operation. IIO does not, on its own, change speed or width autonomously for non-reliability reasons. IIO only sets this bit when it receives a width or speed change indication from downstream component that is not for link reliability reasons.</p>
14	RW1C	0b	<p>Link Bandwidth Management Status</p> <p>This bit is set to 1b by hardware to indicate that either of the following has occurred without the port transitioning through DL_Down status: a) A link retraining initiated by a write of 1b to the Retrain Link bit has completed b) Hardware has autonomously changed link speed or width to attempt to correct unreliable link operation</p> <p>Note IIO also sets this bit when it receives a width or speed change indication from downstream component that is for link reliability reasons.</p>
13	RO	0b	<p>Data Link Layer Link Active</p> <p>Set to 1b when the Data Link Control and Management State Machine is in the DL_Active state, 0b otherwise. On a downstream port or upstream port, when this bit is 0b, the transaction layer associated with the link will abort all transactions that would otherwise be routed to that link.</p>
12	RW-O	1b	<p>Slot Clock Configuration</p> <p>This bit indicates whether IIO receives clock from the same xtal that also provides clock to the device on the other end of the link.</p> <p>1: indicates that same xtal provides clocks to devices on both ends of the link 0: indicates that different xtals provide clocks to devices on both ends of the link</p> <p>Note: This register bit is a RW-O register from the host side. It must be loaded by BIOS in the primary side equivalent register. This register is RO from the secondary side of the NTB.</p>
11	RO	0b	<p>Link Training</p> <p>This field indicates the status of an ongoing link training session in the PCI Express port0: LTSSM has exited the recovery/configuration state</p> <p>1: LTSSM is in recovery/configuration state or the Retrain Link was set but training has not yet begun.</p> <p>The IIO hardware clears this bit once LTSSM has exited the recovery/configuration state.</p>
9:4	RO	00h	<p>Negotiated Link Width</p> <p>This field indicates the negotiated width of the given PCI Express link after training is completed. Only x1, x2, x4, x8 and x16 link width negotiations are possible in IIO. A value of 0x01 in this field corresponds to a link width of x1, 0x02 indicates a link width of x2 and so on, with a value of 0x10 for a link width of x16. The value in this field is reserved and could show any value when the link is not up. Software determines if the link is up or not by reading bit 13 of this register.</p>
3:0	RO-V	1h	<p>Current Link Speed</p> <p>This field indicates the negotiated Link speed of the given PCI Express Link.</p> <p>0001: 2.5 Gbps 0010: 5 Gbps 0011: 8 Gbps Others: Reserved</p> <p>The value in this field is not defined when the link is not up. Software determines if the link is up or not by reading bit 13 of this register.</p>

6.13.3.56 SLTCAP: PCI Express Slot Capabilities

The Slot Capabilities register identifies the PCI Express specific slot capabilities.



SLTCAP Bus: 0 Device: 3 Function: 0 Offset: 1A4h			
Bit	Attr	Default	Description
31:19	RW-O	0h	Physical Slot Number This field indicates the physical slot number of the slot connected to the PCI Express port and is initialized by BIOS.
18	RO	0h	Command Complete Not Capable IIO is capable of command complete interrupt.
17	RW-O	0h	Electromechanical Interlock Present This bit when set indicates that an Electromechanical Interlock is implemented on the chassis for this slot and that lock is controlled by bit 11 in Slot Control register. BIOS note: this capability is not set if the Electromechanical Interlock control is connected to main slot power control.
16:15	RW-O	0h	Slot Power Limit Scale This field specifies the scale used for the Slot Power Limit Value and is initialized by BIOS. IIO uses this field when it sends a Set_Slot_Power_Limit message on PCI Express. Range of Values: 00: 1.0x 01: 0.1x 10: 0.01x 11: 0.001x
14:7	RW-O	00h	Slot Power Limit Value This field specifies the upper limit on power supplied by slot in conjunction with the Slot Power Limit Scale value defined previously. Power limit (in Watts) = SPLS x SPLV. This field is initialized by BIOS. IIO uses this field when it sends a Set_Slot_Power_Limit message on PCI Express. Design note: IIO can chose to send the Set_Slot_Power_Limit message on the link at first link up condition without regards to whether this register and the Slot Power Limit Scale register are programmed yet by BIOS. IIO must then be designed to discard a received Set_Slot_Power_Limit message without an error.
6	RW-O	0h	Hot-plug Capable This field defines hot-plug support capabilities for the PCI Express port. 0: indicates that this slot is not capable of supporting Hot-plug operations. 1: indicates that this slot is capable of supporting Hot-plug operations This bit is programmed by BIOS based on the system design. This bit must be programmed by BIOS to be consistent with the VPP enable bit for the port.
5	RW-O	0h	Hot-plug Surprise This field indicates that a device in this slot may be removed from the system without prior notification (like for instance a PCI Express cable). 0: indicates that hot-plug surprise is not supported 1: indicates that hot-plug surprise is supported Note that if platform implemented cable solution (either direct or via a SIOM with repeater), on a port, then this could be set. BIOS programs this field with a 0 for CEM/SIOM FFs. This bit is used by IIO hardware to determine if a transition from DL_active to DL_Inactive is to be treated as a surprise down error or not. If a port is associated with a hot pluggable slot and the hot-plug surprise bit is set, then any transition to DL_Inactive is not considered an error.
4	RW-O	0h	Power Indicator Present This bit indicates that a Power Indicator is implemented for this slot and is electrically controlled by the chassis. 0: indicates that a Power Indicator that is electrically controlled by the chassis is not present 1: indicates that Power Indicator that is electrically controlled by the chassis is present BIOS programs this field with a 1 for CEM/SIOM FFs and a 0 for Express cable.



SLTCAP Bus: 0 Device: 3 Function: 0 Offset: 1A4h			
Bit	Attr	Default	Description
3	RW-O	0h	<p>Attention Indicator Present</p> <p>This bit indicates that an Attention Indicator is implemented for this slot and is electrically controlled by the chassis</p> <p>0: indicates that an Attention Indicator that is electrically controlled by the chassis is not present</p> <p>1: indicates that an Attention Indicator that is electrically controlled by the chassis is present</p> <p>BIOS programs this field with a 1 for CEM/SIOM FFs.</p>
2	RW-O	0h	<p>MRL Sensor Present</p> <p>This bit indicates that an MRL Sensor is implemented on the chassis for this slot.</p> <p>0: indicates that an MRL Sensor is not present</p> <p>1: indicates that an MRL Sensor is present</p> <p>BIOS programs this field with a 0 for SIOM/Express cable and with either 0 or 1 for CEM depending on system design.</p>
1	RW-O	0h	<p>Power Controller Present</p> <p>This bit indicates that a software controllable power controller is implemented on the chassis for this slot.</p> <p>0: indicates that a software controllable power controller is not present</p> <p>1: indicates that a software controllable power controller is present</p> <p>BIOS programs this field with a 1 for CEM/SIOM FFs and a 0 for Express cable.</p>
0	RW-O	0h	<p>Attention Button Present</p> <p>This bit indicates that the Attention Button event signal is routed (from slot or on-board in the chassis) to the IIO's hot-plug controller.</p> <p>0: indicates that an Attention Button signal is routed to IIO</p> <p>1: indicates that an Attention Button is not routed to IIO</p> <p>BIOS programs this field with a 1 for CEM/SIOM FFs.</p>

6.13.3.57 SLTCON: PCI Express Slot Control

The Slot Control register identifies the PCI Express specific slot control parameters for operations such as Hot-plug and Power Management.

Warning: Any write to this register will set the Command Completed bit in the SLTSTS register, ONLY if the VPP enable bit for the port is set. If the port's VPP enable bit is set (that is, hot-plug for that slot is enabled), then the required actions on VPP are completed before the Command Completed bit is set in the SLTSTS register. If the VPP enable bit for the port is clear, then the write simply updates this register (see individual bit definitions for details) but the Command Completed bit in the SLTSTS register is not set.

SLTCON Bus: 0 Device: 3 Function: 0 Offset: 1A8h			
Bit	Attr	Default	Description
12	RWS	0b	<p>Data Link Layer State Changed Enable</p> <p>When set to 1, this field enables software notification when Data Link Layer Link Active field is changed</p>
11	RW	0b	<p>Electromechanical Interlock Control</p> <p>When software writes either a 1 to this bit, IIO pulses the EMIL pin per; PCI Express Server/Workstation Module Electromechanical Spec Rev 1.0. Write of 0 has no effect. This bit always returns a 0 when read. If electromechanical lock is not implemented, then either a write of 1 or 0 to this register has no effect.</p>



SLTCON Bus: 0 Device: 3 Function: 0 Offset: 1A8h			
Bit	Attr	Default	Description
10	RWS	1b	<p>Power Controller Control</p> <p>If a power controller is implemented, when written sets the power state of the slot per the defined encodings. Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not executed yet at the VPP, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined.</p> <p>0: Power On 1: Power Off</p>
9:8	RW	3h	<p>Power Indicator Control</p> <p>If a Power Indicator is implemented, writes to this register set the Power Indicator to the written state. Reads of this field must reflect the value from the latest write, even if the corresponding hot-plug command is not executed yet at the VPP, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined.</p> <p>00: Reserved. 01: On 10: Blink (IIO drives 1.5 Hz square wave for Chassis mounted LEDs) 11: Off</p> <p>When this register is written, the event is signaled via the virtual pins of the IIO over a dedicated SMBus port.</p> <p>IIO does not generated the Power_Indicator_On/Off/Blink messages on PCI Express when this field is written to by software.</p>
7:6	RW	3h	<p>Attention Indicator Control</p> <p>If an Attention Indicator is implemented, writes to this register set the Attention Indicator to the written state. Reads of this field reflect the value from the latest write, even if the corresponding hot-plug command is not executed yet at the VPP, unless software issues a write without waiting for the previous command to complete in which case the read value is undefined.</p> <p>00: Reserved. 01: On 10: Blink (The IIO drives 1.5 Hz square wave) 11: Off</p> <p>When this register is written, the event is signaled via the virtual pins of the IIO over a dedicated SMBus port.</p> <p>IIO does not generated the Attention_Indicator_On/Off/Blink messages on PCI Express when this field is written to by software.</p>
5	RW	0h	<p>Hot-plug Interrupt Enable</p> <p>When set to 1b, this bit enables generation of Hot-Plug MSI interrupt (and not wake event) on enabled Hot-Plug events, provided ACPI mode for hot-plug is disabled.</p> <p>0: disables interrupt generation on Hot-plug events 1: enables interrupt generation on Hot-plug events</p>
4	RW	0h	<p>Command Completed Interrupt Enable</p> <p>This field enables the generation of Hot-plug interrupts (and not wake event) when a command is completed by the Hot-plug controller connected to the PCI Express port</p> <p>0: disables hot-plug interrupts on a command completion by a hot-plug Controller 1: Enables hot-plug interrupts on a command completion by a hot-plug Controller</p>
3	RW	0h	<p>Presence Detect Changed Enable</p> <p>This bit enables the generation of hot-plug interrupts or wake messages via a presence detect changed event.</p> <p>0: disables generation of hot-plug interrupts or wake messages when a presence detect changed event happens. 1- Enables generation of hot-plug interrupts or wake messages when a presence detect changed event happens.</p>




SLTCON Bus: 0 Device: 3 Function: 0 Offset: 1A8h			
Bit	Attr	Default	Description
2	RW	0h	MRL Sensor Changed Enable This bit enables the generation of hot-plug interrupts or wake messages via a MRL Sensor changed event. 0: disables generation of hot-plug interrupts or wake messages when an MRL Sensor changed event happens. 1: Enables generation of hot-plug interrupts or wake messages when an MRL Sensor changed event happens.
1	RW	0h	Power Fault Detected Enable This bit enables the generation of hot-plug interrupts or wake messages via a power fault event. 0: disables generation of hot-plug interrupts or wake messages when a power fault event happens. 1: Enables generation of hot-plug interrupts or wake messages when a power fault event happens.
0	RW	0h	Attention Button Pressed Enable This bit enables the generation of hot-plug interrupts or wake messages via an attention button pressed event. 0: disables generation of hot-plug interrupts or wake messages when the attention button is pressed. 1: Enables generation of hot-plug interrupts or wake messages when the attention button is pressed.

6.13.3.58 SLTSTS: PCI Express Slot Status Register

The PCI Express Slot Status register defines important status information for operations such as Hot-plug and Power Management.

SLTSTS Bus: 0 Device: 3 Function: 0 Offset: 1AAh			
Bit	Attr	Default	Description
8	RW1C	0h	Data Link Layer State Changed This bit is set (if it is not already set) when the state of the Data Link Layer Link Active bit in the Link Status register changes. Software must read Data Link Layer Active field to determine the link state before initiating configuration cycles to the hot-plugged device.
7	RO	0h	Electromechanical Latch Status When read this register returns the current state of the Electromechanical Interlock (the EMILS pin) which has the defined encodings as: 0b Electromechanical Interlock Disengaged 1b Electromechanical Interlock Engaged
6	RO	0h	Presence Detect State For ports with slots (where the Slot Implemented bit of the PCI Express Capabilities Registers is 1b), this field is the logical OR of the Presence Detect status determined via an in-band mechanism and sideband Present Detect pins. 0: Card/Module/Cable slot empty or Cable Slot occupied but not powered 1: Card/module Present in slot (powered or unpowered) or cable present and powered on other end For ports with no slots, IIO hardwires this bit to 1b. Note: OS could get confused when it sees an empty PCI Express RP i.e. 'no slots + no presence', since this is now disallowed in the spec. So BIOS must hide all unused RPs devices in IIO config space, via the DEVHIDE register in Intel QPI Configuration Register space.
5	RO	0h	MRL Sensor State This bit reports the status of an MRL sensor if it is implemented. 0: MRL Closed 1: MRL Open



SLTSTS Bus: 0 Device: 3 Function: 0 Offset: 1AAh			
Bit	Attr	Default	Description
4	RW1C	0h	Command Completed This bit is set by the IIO when the hot-plug command has completed and the hot-plug controller is ready to accept a subsequent command. It is subsequently cleared by software after the field has been read and processed. This bit provides no guarantee that the action corresponding to the command is complete.
3	RW1C	0h	Presence Detect Changed This bit is set by the IIO when a Presence Detect Changed event is detected. It is subsequently cleared by software after the field has been read and processed. On-board logic per slot must set the VPP signal corresponding this bit inactive if the FF/system does not support out-of-band presence detect.
2	RW1C	0h	MRL Sensor Changed This bit is set by the IIO when an MRL Sensor Changed event is detected. It is subsequently cleared by software after the field has been read and processed. On-board logic per slot must set the VPP signal corresponding this bit inactive if the FF/system does not support MRL.
1	RW1C	0h	Power Fault Detected This bit is set by the IIO when a power fault event is detected by the power controller. It is subsequently cleared by software after the field has been read and processed. On-board logic per slot must set the VPP signal corresponding this bit inactive if the FF/system does not support power fault detection.
0	RW1C	0h	Attention Button Pressed This bit is set by the IIO when the attention button is pressed. It is subsequently cleared by software after the field has been read and processed. On-board logic per slot must set the VPP signal corresponding this bit inactive if the FF/system does not support attention button. IIO silently discards the Attention_Button_Pressed message if received from PCI Express link without updating this bit.

6.13.3.59 ROOTCON: PCI Express Root Control

Device 3, Function 0, Offset **1ACh**. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.51](#). **Notice the offset differences.**

6.13.3.60 ROOTCAP: PCI Express Root Capabilities

Device 3, Function 0, Offset **1AEh**. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.52](#). **Notice the offset differences.**

6.13.3.61 ROOTSTS: PCI Express Root Status

Device 3, Function 0, Offset **1B0h**. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.53](#). **Notice the offset differences.**

6.13.3.62 DEVCAP2: PCI Express Device Capabilities Register

Device 3, Function 0, Offset **1B4h**. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.54](#). **Notice the offset differences.**

6.13.3.63 DEVCAP2: PCI Express Device Capabilities Register

Device 3, Function 0, Offset **1B8h**. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.55](#). **Notice the offset differences.**



6.13.3.64 DEVCTRL2: PCI Express Device Control 2 Register

DEVCTRL2 Bus: 0 Device: 3 Function: 0 Offset: 1B8h			
Bit	Attr	Default	Description
15:6	RV	0h	Reserved
5	RW	0b	Alternative RID Interpretation Enable When set to 1b, ARI is enabled for the NTB EP. Note: Normally, The 5-bit Device ID is required to be zero in the RID that consists of BDF, but when ARI is enabled, the 8-bit DF is now interpreted as an 8-bit Function Number with the Device Number equal to zero implied.
4	RW-V	0b	Completion Timeout Disable When set to 1b, this bit disables the Completion Timeout mechanism for all NP tx that IIO issues on the PCIe/DMI link and in the case of Intel QuickData Technology DMA, for all NP tx that DMA issues upstream. When 0b, completion timeout is enabled. Software can change this field while there is active traffic in the RP.
3:0	RW-V	0h	Completion Timeout Value on NP Tx that IIO issues on PCIe In Devices that support Completion Timeout programmability, this field allows system software to modify the Completion Timeout range. The following encodings and corresponding timeout ranges are defined: 0000b = 10ms to 50ms 0001b = Reserved (IIO aliases to 0000b) 0010b = Reserved (IIO aliases to 0000b) 0101b = 16 ms to 55 ms 0110b = 65 ms to 210 ms 1001b = 260 ms to 900 ms 1010b = 1 s to 3.5 s 1101b = 4 s to 13 s 1110b = 17 s to 64 s When software selects 17s to 64s range, Section 11.2.92 further controls the timeout value within that range. For all other ranges selected by OS, the timeout value within that range is fixed in IIO hardware. Software can change this field while there is active traffic in the root port. This value will also be used to control PME_TO_ACK Timeout. That is this field sets the timeout value for receiving a PME_TO_ACK message after a PME_TURN_OFF message has been transmitted. The PME_TO_ACK Timeout has meaning only if bit 6 of MISCCTRLSTS register is set to a 1b.

6.13.3.65 LNKCAP2: PCI Express Link Capabilities 2 Register

Device 3, Function 0, Offset 1BCh. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.56](#). **Notice the offset differences.**

6.13.3.66 LNKCON2: PCI Express Link Control 2 Register

Device 3, Function 0, Offset 1C0h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.57](#). **Notice the offset differences.**

6.13.3.67 LNKSTS2: PCI Express Link Status Register 2

Device 3, Function 0, Offset 1C2h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.58](#). **Notice the offset differences.**

6.13.3.68 ERRINJCAP: PCI Express Error Injection Capability

Device 3, Function 0, Offset 1D0h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.89](#).



6.13.3.69 ERRINJHDR: PCI Express Error Injection Capability Header

Device 3, Function 0, Offset D4h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.90](#).

6.13.3.70 ERRINJCON: PCI Express Error Injection Control Register

Device 3, Function 0, Offset 1D8h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.91](#).

6.13.3.71 CTOCTRL: Completion Timeout Control

Device 3, Function 0, Offset 1E0h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.92](#).

6.13.3.72 XPCORERRSTS: XP Correctable Error Status

Device 3, Function 0, Offset 200h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.93](#).

6.13.3.73 XPCORERRMSK: XP Correctable Error Mask

Device 3, Function 0, Offset 204h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.94](#).

6.13.3.74 XPUNCERRSTS: XP Uncorrectable Error Status

Device 3, Function 0, Offset 208h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.95](#).

6.13.3.75 XPUNCERRMSK: XP Uncorrectable Error Mask

Device 3, Function 0, Offset 20Ch. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.96](#).

6.13.3.76 XPUNCERRSEV: XP Uncorrectable Error Severity

Device 3, Function 0, Offset 210h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.97](#).

6.13.3.77 UNCEDMASK: Uncorrectable Error Detect Status Mask

Device 3, Function 0, Offset 218h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.98](#).

6.13.3.78 COREDMASK: Correctable Error Detect Status Mask

Device 3, Function 0, Offset 21Ch. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.100](#).

6.13.3.79 RPEDMASK: Root Port Error Detect Status Mask

Device 3, Function 0, Offset 220h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.101](#).

**6.13.3.80 XPUNCEDMASK: XP Uncorrectable Error Detect Mask**

Device 3, Function 0, Offset 224h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.102](#).

6.13.3.81 XPCOREDMASK: XP Correctable Error Detect Mask

Device 3, Function 0, Offset 228h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.103](#).

6.13.3.82 XPGLBERRSTS: XP Global Error Status

Device 3, Function 0, Offset 230h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.104](#).

6.13.3.83 XPGLBERRPTR: XP Global Error Pointer

Device 3, Function 0, Offset 232h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.105](#).

6.13.3.84 PXP2CAP: Secondary PCI Express Extended Capability Header

Device 3, Function 0, Offset 250h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.106](#)

6.13.3.85 LNKCON3: Link Control 3 Register

Device 3, Function 0, Offset 254h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.107](#).

6.13.3.86 LNERRSTS: Lane Error Status Register

Device 3, Function 0, Offset 258h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.108](#).

6.13.3.87 LN[0:3]EQ: Lane 0 through Lane 3 Equalization Control

Device 3, Function 0, Offset 25Ch, 25Eh, 260h, 262h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.109](#).

6.13.3.88 LN[4:7]EQ: Lane 4 through Lane 7 Equalization Control


Device 3, Function 0, Offset 264h, 266h, 268h, 26Ah. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.110](#).

6.13.3.89 LN[8:15]EQ: Lane 8 though Lane 15 Equalization Control

Device 3, Function 0, Offset 26Ch, 26Eh, 270h, 272h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.111](#).

6.13.3.90 mcast_cap_hdr

Device 3, Function 0, Offset 300h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.112](#).



6.13.3.91 **mcast_cap_ext**

Device 3, Function 0, Offset 304h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.113](#).

6.13.3.92 **mcast_cap**

Device 3, Function 0, Offset 30Ch. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.114](#).

6.13.3.93 **mcast_ctrl**

Device 3, Function 0, Offset 30Eh. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.115](#).

6.13.3.94 **mcast_base**

Device 3, Function 0, Offset 310h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.116](#).

6.13.3.95 **mcast_rcv**

Device 3, Function 0, Offset 318h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.117](#).

6.13.3.96 **mcast_cap_blk_all**

Device 3, Function 0, Offset 320h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.118](#).

6.13.3.97 **mcast_blk_unt**

Device 3, Function 0, Offset 328h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.119](#).

6.13.3.98 **mcast_overlay_bar**

Device 3, Function 0, Offset 330h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.120](#).

6.13.4 **PCI Express Configuration Registers (NTB Secondary Side)**

6.13.5 **Configuration Register Map (NTB Secondary Side)**

This section covers the NTB secondary side configuration space registers.

When configured as an NTB there are two sides to discuss for configuration registers. The primary side of the NTB's configuration space is located on Device 3, Function 0 with respect to the Intel® Xeon® Processor E5 v4 Product Family and a secondary side of the NTB's configuration space is located on some enumerated bus on another system and does not exist as configuration space on the local Intel® Xeon® Processor E5 v4 Product Family-based system anywhere.



Table 6-9. Device 0 Function 0 (Non -Transparent Bridge) Configuration Map 0x00h - 0xFCh

DID		VID		0h	MSIXMSGCTRL	MSIXNXTPT R	MSIXCAPID	80h
PCISTS		PCICMD		4h	TABLEOFF_BIR			84h
CCR			RID	8h	PBAOFF_BIR			88h
BIST	HDR	PLAT	CLSR	Ch				8Ch
SB01BASE				10h	PXPCAP	PXPNTPTR	PXPCAPID	90h
				14h	DEVCAP			94h
SB23BASE				18h	DEVSTS	DEVCTRL		98h
				1Ch	LNKCAP			9Ch
SB45BASE				20h	LNKSTS	LNKCON		A0h
				24h				A4h
				28h				A8h
SID		SUBVID		2Ch				ACh
				30h				B0h
			CAPPTR	34h	DEVCAP2			B4h
				38h		DEVCTRL2		B8h
MAXLAT	MINGNT	INTPIN	INTL	3Ch	LNKCAP2			BCh
				40h	LNKSTS2	LNKCON2		C0h
				44h				C4h
				48h				C8h
				4Ch				CCh
				50h				D0h
				54h			SSCNTL	D4h
				58h				D8h
				5Ch				DCh
MSICTRL		MSINXTPTR	MSICAPID	60h	PMCAP			E0h
MSIAR				64h	PMCSR			E4h
MSIUAR				68h				E8h
MSIDR				6Ch				ECh
MSIMSK				70h				F0h
MSIPENDING				74h				F4h
				78h				F8h
				7Ch				FCh

Table 6-10. Device 0 Function 0 (Non -Transparent Bridge) Configuration Map 0x100h - 0x1FCh (Sheet 1 of 2)

PXP2CAP		100h		180h
LNERRSTS		104h		184h
LN1EQ	LN0EQ	108h		188h
LN3EQ	LN2EQ	10Ch		18Ch
LN5EQ	LN4EQ	110h		190h
LN7EQ	LN6EQ	114h		194h



Table 6-10. Device 0 Function 0 (Non -Transparent Bridge) Configuration Map 0x100h - 0x1FCh (Sheet 2 of 2)

LN9EQ	LN8EQ	118h		198h
LN11EQ	LN10EQ	11Ch		19Ch
LN13EQ	LN12EQ	120h		1A0h
LN15EQ	LN14EQ	124h		1A4h
		128h		1A8h
		12Ch		1ACh
		130h		1B0h
		134h		1B4h
		138h		1B8h
		13Ch		1BCh
		140h		1C0h
		144h		1C4h
		148h		1C8h
		14Ch		1CCh
		150h		1D0h
		154h		1D4h
		158h		1D8h
		15Ch		1DCh
		160h		1E0h
		164h		1E4h
		168h	1E8h	
		16Ch	1ECh	
		170h	1F0h	
		174h	1F4h	
		178h	1F8h	
		17Ch	1FCh	

6.13.5.1 VID: Vendor Identification

VID Bus: M Device: 0 Function: 0 Offset: 0 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 500 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 500			
Bit	Attr	Default	Description
15:0	RO	8086h	Vendor Identification Number The value is assigned by PCI-SIG to Intel.



6.13.5.2 DID: Device Identification

DID Bus: M Device: 0 Function: 0 Offset: 02 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 502 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 502			
Bit	Attr	Default	Description
15:0	RO	2f0Fh	Device Identification Number The value is assigned by Intel to each product.

6.13.5.3 PCICMD: PCI Command

This register defines the PCI 3.0 compatible command register values applicable to PCI Express space.

PCICMD Bus: M Device: 0 Function: 0 Offset: 04 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 504 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 504			
Bit	Attr	Default	Description
15:11	RV	0h	Reserved
10	RW	0b	INTxDisable Interrupt Disable. Controls the ability of the PCI Express port to generate INTx messages. This bit does not affect the ability of Processor to route interrupt messages received at the PCI Express port. However, this bit controls the generation of legacy interrupts to the DMI for PCI Express errors detected internally in this port (e.g. Malformed TLP, CRC error, completion time out etc.) or when receiving RP error messages or interrupts due to Hot Plug/Power Management events generated in legacy mode within Processor. 1: Legacy Interrupt mode is disabled 0: Legacy Interrupt mode is enabled
9	RO	0b	Fast Back-to-Back Enable Not applicable to PCI Express must be hardwired to 0.
8	RO	0b	SERR Enable For PCI Express/DMI ports, this field enables notifying the internal core error logic of occurrence of an uncorrectable error (fatal or non-fatal) at the port. The internal core error logic of IIO then decides if/how to escalate the error further (pins/message etc.). This bit also controls the propagation of PCI Express ERR_FATAL and ERR_NONFATAL messages received from the port to the internal IIO core error logic. 1: Fatal and Non-fatal error generation and Fatal and Non-fatal error message forwarding is enabled 0: Fatal and Non-fatal error generation and Fatal and Non-fatal error message forwarding is disabled
7	RO	0b	IDSEL Stepping/Wait Cycle Control Not applicable to PCI Express must be hardwired to 0.
6	RW	0b	Parity Error Response For PCI Express/DMI ports, IIO ignores this bit and always does ECC/parity checking and signaling for data/address of transactions both to and from IIO. This bit though affects the setting of bit 8 in the Section 6.13.5.4, "PCISTS: PCI Status" on page 361.
5	RO	0b	VGA palette snoop Enable Not applicable to PCI Express must be hardwired to 0.
4	RO	0b	Memory Write and Invalidate Enable Not applicable to PCI Express must be hardwired to 0.



PCICMD Bus: M Device: 0 Function: 0 Offset: 04 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 504 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 504			
Bit	Attr	Default	Description
3	RO	0b	Special Cycle Enable Not applicable to PCI Express must be hardwired to 0.
2	RW	0b	Bus Master Enable 1: When this bit is Set, the PCIe NTB will forward Memory Requests that it receives on its primary internal interface to its secondary external link interface.0: When this bit is Clear, the PCIe NTB will not forward Memory Requests that it receives on its primary internal interface. Memory requests received on the primary internal interface will be returned to requester as an Unsupported Requests UR. Requests other than Memory Requests are not controlled by this bit. Default value of this bit is 0b.
1	RW	0b	Memory Space Enable 1: Enables a PCI Express port's memory range registers to be decoded as valid target addresses for transactions from secondary side. 0: Disables a PCI Express port's memory range registers (including the Configuration Registers range registers) to be decoded as valid target addresses for transactions from secondary side. all memory accesses received from secondary side are UR'ed
0	RO	0b	IO Space Enable Controls a device's response to I/O Space accesses. A value of 0 disables the device response. A value of 1 allows the device to respond to I/O Space accesses. State after RST# is 0. NTB does not support I/O space accesses. Hardwired to 0


6.13.5.4 PCISTS: PCI Status

The PCI Status register is a 16-bit status register that reports the occurrence of various events associated with the primary side of the “virtual” PCI-PCI bridge embedded in PCI Express ports and also primary side of the other devices on the internal IIO bus.

PCISTS Bus: M Device: 0 Function: 0 Offset: 06 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 506 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 506			
Bit	Attr	Default	Description
15	RW1C	0b	Detected Parity Error This bit is set by a device when it receives a packet on the primary side with an uncorrectable data error (that is, a packet with poison bit set or an uncorrectable data ECC error was detected at the XP-DP interface when ECC checking is done) or an uncorrectable address/control parity error. The setting of this bit is regardless of the Parity Error Response bit (PERRE) in the PCICMD register.
14	RO	0b	Signaled System Error 1: The device reported fatal/non-fatal (and not correctable) errors it detected on its PCI Express interface through the ERR[2:0] pins or message to PCH, with SERRE bit enabled. Software clears this bit by writing a '1' to it. For Express ports this bit is also set (when SERR enable bit is set) when a FATAL/NON-FATAL message is forwarded from the Express link to the ERR[2:0] pins or to PCH via a message. Note that IIO internal 'core' errors (like parity error in the internal queues) are not reported via this bit. 0: The device did not report a fatal/non-fatal error



Bit	Attr	Default	Description
PCISTS Bus: M Device: 0 Function: 0 Offset: 06 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 506 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 506			
13	RW1C	0b	Received Master Abort This bit is set when a device experiences a master abort condition on a transaction it mastered on the primary interface (IIO internal bus). Note that certain errors might be detected right at the PCI Express interface and those transactions might not 'propagate' to the primary interface before the error is detected (e.g. accesses to memory above TOCM in cases where the PCIe interface logic itself might have visibility into TOCM). Such errors do not cause this bit to be set, and are reported via the PCI Express interface error bits (secondary status register). Conditions that cause bit 13 to be set, include: Device receives a completion on the primary interface (internal bus of IIO) with Unsupported Request or master abort completion Status. This includes UR status received on the primary side of a PCI Express port on peer-to-peer completions also. Device accesses to holes in the main memory address region that are detected by the Intel QPI source address decoder.
12	RW1C	0b	Received Target Abort This bit is set when a device experiences a completer abort condition on a transaction it mastered on the primary interface (IIO internal bus). Note that certain errors might be detected right at the PCI Express interface and those transactions might not 'propagate' to the primary interface before the error is detected (e.g. accesses to memory above VTCSRBASE). Such errors do not cause this bit to be set, and are reported via the PCI Express interface error bits (secondary status register). Conditions that cause bit 12 to be set, include: Device receives a completion on the primary interface (internal bus of IIO) with completer abort completion Status. This includes CA status received on the primary side of a PCI Express port on peer-to-peer completions also. Accesses to Intel QPI that return a failed completion status
11	RW1C	0b	Signaled Target Abort This bit is set when the NTB port forwards a completer abort (CA) completion status from the primary interface to the secondary interface.
10:9	RO	0h	DEVSEL# Timing Not applicable to PCI Express. Hardwired to 0.
8	RW1C	0b	Master Data Parity Error This bit is set if the Parity Error Response bit in the PCI Command register is set and the Requestor receives a poisoned completion on the secondary interface or Requestor forwards a poisoned write request (including MSI/MSI-X writes) from the primary interface to the secondary interface.
7	RO	0b	Fast Back-to-Back Not applicable to PCI Express. Hardwired to 0.
6	RO	0b	Reserved
5	RO	0b	66 MHz capable Not applicable to PCI Express. Hardwired to 0.
4	RO	1b	Capabilities List This bit indicates the presence of a capabilities list structure
3	RO-V	0b	INTx Status When Set, indicates that an INTx emulation interrupt is pending internally in the Function.



6.13.5.5 RID: Revision Identification

RID Bus: M Device: 0 Function: 0 Offset: 08 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 508 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 508			
Bit	Attr	Default	Description
7:0	RO	00h	Revision_ID Reflects the Uncore Revision ID after reset. Reflects the Compatibility Revision ID after BIOS writes 0x69 to any RID register in any CPU function. Implementation Note: Read and write requests from the host to any RID register in any CPU function are re-directed to the IIO cluster. Accesses to the CCR field are also redirected due to DWORD alignment. It is possible that JTAG accesses are direct, so will not always be redirected.


6.13.5.6 CCR: Class Code

This register contains the Class Code for the device.

CCR Bus: M Device: 0 Function: 0 Offset: 09 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 509 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 509			
Bit	Attr	Default	Description
23:16	RO	06h	Base Class For PCI Express NTB port this field is hardwired to 06h, indicating it is a 'Bridge Device'.
15:8	RO	80h	Sub-Class For PCI Express NTB port, this field hardwired to 80h to indicate an 'Other bridge type'.
7:0	RO	00h	Register-Level Programming Interface This field is hardwired to 00h for PCI Express NTB port.

6.13.5.7 CLSR: Cacheline Size

CLSR Bus: M Device: 0 Function: 0 Offset: 0C Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 50C Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 50C			
Bit	Attr	Default	Description
7:0	RW	0h	Cacheline Size This register is set as RW for compatibility reasons only. Cacheline size for IIO is always 64B. IIO hardware ignore this setting.



6.13.5.8 PLAT: Primary Latency Timer

This register denotes the maximum time slice for a burst transaction in legacy PCI 2.3 on the primary interface. It does not affect/influence PCI Express functionality.

PLAT Bus: M Device: 0 Function: 0 Offset: 0D Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 50D Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 50D			
Bit	Attr	Default	Description
7:0	RO	0h	Prim_Lat_timer Primary Latency Timer Not applicable to PCI Express. Hardwired to 00h.

6.13.5.9 HDR: Header Type

This register identifies the header layout of the configuration space.

HDR Bus: M Device: 0 Function: 0 Offset: 0E Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 50E Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 50E			
Bit	Attr	Default	Description
7	RO	0b	Multi-function Device This bit defaults to 0 for PCI Express NTB port.
6:0	RO	00h	Configuration Layout This field identifies the format of the configuration header layout. It is Type0 for PCI Express NTB port. The default is 00h, indicating a 'non-bridge function'.

6.13.5.10 SB01BASE: Secondary BAR 0/1 Base Address

(PCIe NTB mode) This register is BAR 0/1 for the secondary side of the NTB. This configuration register can be modified via configuration transaction from the secondary side of the NTB and can also be modified from the primary side of the NTB via MMIO transaction to [Section 6.13.7.9, "SBAR0BASE: Secondary BAR 0/1 Base Address"](#) on page 392.

SB01BASE Bus: M Device: 0 Function: 0 Offset: 10 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 510 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 510			
Bit	Attr	Default	Description
63:15	RW	00h	Secondary BAR 0/1 Base This register is reflected into the BAR 0/1 register pair in the Configuration Space of the Secondary side of the NTB written by SW on a 32 KB alignment.
14:4	RO	00h	Reserved Fixed size of 32 KB.
3	RW-O	1b	Prefetchable BAR points to Prefetchable memory (default) BAR points to Non-Prefetchable memory



SB01BASE Bus: M Device: 0 Function: 0 Offset: 10 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 510 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 510			
Bit	Attr	Default	Description
2:1	RO	10b	Type Memory type claimed by BAR 2/3 is 64-bit addressable.
0	RO	0b	Memory Space Indicator BAR resource is memory (as opposed to I/O).

6.13.5.11 SB23BASE: Secondary BAR 2/3 Base Address

(PCIe NTB mode) This register is BAR 2/3 for the secondary side of the NTB. This configuration register can be modified via configuration transaction from the secondary side of the NTB and can also be modified from the primary side of the NTB via MMIO.

SB23BASE Bus: M Device: 0 Function: 0 Offset: 20 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 518 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 518			
Bit	Attr	Default	Description
63:12	RW	0h	Secondary BAR 2/3 Base Sets the location of the BAR written by SWNOTE: The number of bits that are writable in this register is dictated by the value loaded into the SBAR23SZ register Section 6.13.3.24, "SBAR23SZ: Secondary BAR 2/3 Size" on page 331 by the BIOS at initialization time (before BIOS PCI enumeration). SBAR23SZ indicates the lowest order bit of this register field that is writeable where valid values are 12-39. If SBAR23SZ is set to 12, all bits are writeable. If set to 39, then bits 38:12 are Read Only and will return values of 0. Note: For the special case where SBAR23SZ = '0', bits 63:0 are all RO='0' resulting in the BAR being disabled. Note: The lowest order address bit is 12 to enforce a minimum granularity of 4 KB.
11:4	RO	00h	Reserved Granularity must be at least 4 KB.
3	RO	1b	Prefetchable BAR points to Prefetchable memory.
2:1	RO	10b	Type Memory type claimed by BAR 2/3 is 64-bit addressable.
0	RO	0b	Memory Space Indicator BAR resource is memory (as opposed to I/O).



6.13.5.12 SB45BASE: Secondary BAR 4/5 Base Address


This register is BAR 4/5 for the secondary side of the NTB. This configuration register can be modified via configuration transaction from the secondary side of the NTB and can also be modified from the primary side of the NTB via MMIO transaction to “Secondary BAR 4/5 Base Address (SBAR4BASE)”.

SB45BASE				
Bus: M		Device: 0	Function: 0	Offset: 2C
Bus: 0		Device: 3	Function: 0	MMIO BAR: PB01BASE
		Offset: 520		
Bus: 0		Device: 3	Function: 0	MMIO BAR: SB01BASE
		Offset: 520		
Bit	Attr	Default	Description	
63:12	RW	0h	Secondary BAR 4/5 Base Sets the location of the BAR written by SWNOTE: The number of bits that are writable in this register is dictated by the value loaded into the SBAR45SZ register Section 6.13.3.25, “SBAR45SZ: Secondary BAR 4/5 Size” on page 331 by the BIOS at initialization time (before BIOS PCI enumeration). SBAR45SZ indicates the lowest order bit of this register field that is writeable where valid values are 12-39. If SBAR45SZ is set to 12, all bits are writeable. If set to 39, then bits 38:12 are Read Only and will return values of 0. Note: For the special case where SBAR45SZ = ‘0’, bits 63:0 are all RO=‘0’ resulting in the BAR being disabled. Note: The lowest order address bit is 12 to enforce a minimum granularity of 4 KB.	
11:4	RO	00h	Reserved Granularity must be at least 4 KB.	
3	RO	1b	Prefetchable BAR points to Prefetchable memory.	
2:1	RO	10b	Type Memory type claimed by BAR 4/5 is 64-bit addressable.	
0	RO	0b	Memory Space Indicator BAR resource is memory (as opposed to I/O).	

6.13.5.13 SUBVID: Subsystem Vendor ID

This register identifies a particular subsystem.

SUBVID				
Bus: M		Device: 0	Function: 0	Offset: 2C
Bus: 0		Device: 3	Function: 0	MMIO BAR: PB01BASE
		Offset: 52C		
Bus: 0		Device: 3	Function: 0	MMIO BAR: SB01BASE
		Offset: 52C		
Bit	Attr	Default	Description	
15:0	RW-O	0000h	Subsystem Vendor ID This field must be programmed during boot-up to indicate the vendor of the system board. When any byte or combination of bytes of this register is written, the register value locks and cannot be further updated.	



6.13.5.14 SID: Subsystem Identity

This register identifies a particular subsystem.

SID Bus: M Device: 0 Function: 0 Offset: 2E Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 52E Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 52E			
Bit	Attr	Default	Description
15:0	RW-O	0000h	Subsystem ID This field must be programmed during BIOS initialization. When any byte or combination of bytes of this register is written, the register value locks and cannot be further updated.

6.13.5.15 CAPPTR: Capability Pointer

The CAPPTR is used to point to a linked list of additional capabilities implemented by the device. It provides the offset to the first set of capabilities registers located in the PCI compatible space.

CAPPTR Bus: M Device: 0 Function: 0 Offset: 34 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 534 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 534			
Bit	Attr	Default	Description
7:0	RW-O	60h	Capability Pointer Points to the first capability structure for the device.

6.13.5.16 INTL: Interrupt Line

The Interrupt Line register is used to communicate interrupt line routing information between initialization code and the device driver. This register is not used in newer OSes and is just kept as is.

INTL Bus: M Device: 0 Function: 0 Offset: 3C Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 53C Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 53C			
Bit	Attr	Default	Description
7:0	RW	00h	Interrupt Line This bit is RW for devices that can generate a legacy INTx message and is needed only for compatibility purposes.



6.13.5.17 INTPIN: Interrupt Pin

The INTP register identifies legacy interrupts for INTA, INTB, INTC and INTD as determined by BIOS/firmware. These are emulated over the DMI port using the appropriate Assert_Intx commands.


INTPIN Bus: M Device: 0 Function: 0 Offset: 3D Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 53D Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 53D			
Bit	Attr	Default	Description
7:0	RW-O	01h	INTP Interrupt Pin. This field defines the type of interrupt to generate for the PCI Express port.001: Generate INTA 010: Generate INTB 011: Generate INTC 100: Generate INTD Others: Reserved BIOS/configuration Software has the ability to program this register once during boot to set up the correct interrupt for the port. Note: While the PCI spec. defines only one interrupt line (INTA#) for a single function device, the logic for the NTB has been modified to meet customer requests for programmability of the interrupt pin. BIOS should always set this to INTA# for standard OS's.

6.13.5.18 MINGNT: Minimum Grant

MINGNT Bus: M Device: 0 Function: 0 Offset: 3E Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 53E Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 53E			
Bit	Attr	Default	Description
7:0	RO	00h	Minimum Grant This register does not apply to PCI Express. It is hard-coded to '00'h.

6.13.5.19 MAXLAT: Maximum Latency

MAXLAT Bus: M Device: 0 Function: 0 Offset: 3F Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 53F Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 53F			
Bit	Attr	Default	Description
7:0	RO	00h	Maximum Latency This register does not apply to PCI Express. It is hard-coded to '00'h.



6.13.5.20 MSICAPID: MSI Capability ID

MSICAPID Bus: M Device: 0 Function: 0 Offset: 60 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 560 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 560			
Bit	Attr	Default	Description
7:0	RO	05h	Capability ID Assigned by PCI-SIG for MSI.

6.13.5.21 MSINXTPTR: MSI Next Pointer

MSINXTPTR Bus: M Device: 0 Function: 0 Offset: 61 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 561 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 561			
Bit	Attr	Default	Description
7:0	RW-O	80h	Next Ptr This field is set to 80h for the next capability list (PCI Express capability structure) in the chain.

6.13.5.22 MSICTRL: MSI Control

MSICTRL Bus: M Device: 0 Function: 0 Offset: 62 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 562 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 562			
Bit	Attr	Default	Description
8	RO	1b	Per-vector masking capable This bit indicates that PCI Express ports support MSI per-vector masking.
7	RO-V	0b	64-bit Address Capable A PCI Express Endpoint must support the 64-bit Message Address version of the MSI Capability structure 1: Function is capable of sending 64-bit message address 0: Function is not capable of sending 64-bit message address.
6:4	RW	000b	Multiple Message Enable Applicable only to PCI Express ports. Software writes to this field to indicate the number of allocated messages which is aligned to a power of two. When MSI is enabled, the software will allocate at least one message to the device. A value of 000 indicates 1 message. Value Number of Messages Requested 000b = 1 001b = 2 010b = 4 011b = 8 100b = 16 101b = 32 110b = Reserved 111b = Reserved



Bit	Attr	Default	Description
MSICTRL Bus: M Device: 0 Function: 0 Offset: 62 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 562 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 562			
3:1	RO	001b	Multiple Message Capable IOH's PCI Express port supports 16 messages for all internal events. Value Number of Messages Requested 000b = 1 001b = 2 010b = 4 011b = 8 100b = 16 101b = 32 110b = Reserved 111b = Reserved
0	RW	0b	MSI Enable The software sets this bit to select platform-specific interrupts or transmit MSI messages. 0: Disables MSI from being generated. 1: Enables the PCI Express port to use MSI messages for RAS, provided bit 4 in Section 6.13.3.50, "MISCCTRLSTS: Misc. Control and Status" on page 340 is clear and also enables the Express port to use MSI messages for Power Management and Hot Plug events at the root port provided these individual events are not enabled for ACPI handling (see Section 6.13.3.50, "MISCCTRLSTS: Misc. Control and Status" on page 340 for details). Note: Software must disable INTx and MSI-X for this device when using MSI

6.13.5.23 MSIAR: MSI Address

The MSI Address Register (MSIAR) contains the system specific address information to route MSI interrupts from the root ports and is broken into its constituent fields.

Bit	Attr	Default	Description
MSIAR Bus: M Device: 0 Function: 0 Offset: 64 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 564 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 564			
31:20	RW	0h	Address MSB This field specifies the 12 most significant bits of the 32-bit MSI address. This field is R/W.
19:12	RW	00h	Address Destination ID This field is initialized by software for routing the interrupts to the appropriate destination.
3	RW	0h	Address Redirection Hint 0: directed 1: redirectable
2	RW	0h	Address Destination Mode 0: physical 1: logical



6.13.5.24 MSIUAR: Upper Address MSB

If the MSI Enable bit (bit 0 of the MSICTRL) is set, the contents of this register (if non-zero) specify the upper 32-bits of a 64-bit message address (AD[63:32]). If the contents of this register are zero, the function uses the 32 bit address specified by the message address register.

MSIUAR Bus: M Device: 0 Function: 0 Offset: 68 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 568 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 568			
Bit	Attr	Default	Description
31:0	RW	00000000h	MSI Upper Address Register

6.13.5.25 MSIDR: MSI Data

MSIDR Bus: M Device: 0 Function: 0 Offset: 6C Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 568 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 568			
Bit	Attr	Default	Description
31:16	RO	0000h	Reserved.
15	RW	0h	Trigger Mode 0: Edge Triggered 1: Level Triggered Note: IIO does nothing with this bit other than passing it along to Intel QPI
14	RW	0h	Level 0: Deassert 1: Assert Note: IIO does nothing with this bit other than passing it along to Intel QPI
13:12	RW	0h	Don't care for IIO
11:8	RW	0h	Delivery Mode 0000: Fixed: Trigger Mode can be edge or level. 0001: Lowest Priority: Trigger Mode can be edge or level. 0010: Intel SMI/PMI/MCA - Not supported via MSI of root port 0011: Reserved - Not supported via MSI of root port 0100: NMI - Not supported via MSI of root port 0101: INIT - Not supported via MSI of root port 0110: Reserved 0111: ExtINT - Not supported via MSI of root port Others: Reserved
7:0	RW	00h	Interrupt Vector The interrupt vector (LSB) will be modified by the IIO to provide context sensitive interrupt information for different events that require attention from the processor. Only 1 message can be enabled by software, so all events may use any vector.



6.13.5.26 MSIMSK: MSI Mask Bit

The Mask Bit register enables software to disable message sending on a per-vector basis.

MSIMSK Bus: M Device: 0 Function: 0 Offset: 70 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 56C Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 56C			
Bit	Attr	Default	Description
31:1	RV	0h	Reserved
0	RW	0h	Mask Bit For each Mask bit that is set, the PCI Express port is prohibited from sending the associated message. NTB supports up to 1 messages Corresponding bits are masked if set to '1'

6.13.5.27 MSIPENDING: MSI Pending Bit

The Mask Pending register enables software to defer message sending on a per-vector basis.

MSIPENDING Bus: M Device: 0 Function: 0 Offset: 74 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 570 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 570			
Bit	Attr	Default	Description
31:1	RV	0h	Reserved
0	RO	0h	Pending Bits For each Pending bit that is set, the PCI Express port has a pending associated message. NTB supports 1 message. Corresponding bits are pending if set to '1'.

6.13.5.28 MSIXCAPID: MSI-X Capability ID

MSIXCAPID Bus: M Device: 0 Function: 0 Offset: 80 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 580 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 580			
Bit	Attr	Default	Description
7:0	RO	11h	Capability ID Assigned by PCI-SIG for MSI-X.



6.13.5.29 MSIXNXPTR: MSI-X Next Pointer

MSIXNXPTR Bus: M Device: 0 Function: 0 Offset: 81 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 581 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 581			
Bit	Attr	Default	Description
7:0	RO	90h	Next Ptr This field is set to 90h for the next capability list (PCI Express capability structure) in the chain.

6.13.5.30 MSIXMSGCTRL: MSI-X Message Control

MSIXMSGCTRL Bus: M Device: 0 Function: 0 Offset: 82 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 582 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 582			
Bit	Attr	Default	Description
15	RW	0b	MSI-X Enable Software uses this bit to select between INTx or MSI or MSI-X method for signaling interrupts from the NTB 0: NTB is prohibited from using MSI-X to request service 1: MSI-X method is chosen for NTB interrupts Note: Software must disable INTx and MSI for this device when using MSI-X
14	RW	0b	Function Mask If = 1b, all the vectors associated with the NTB are masked, regardless of the per vector mask bit state. If = 0b, each vector's mask bit determines whether the vector is masked or not. Setting or clearing the MSI-X function mask bit has no effect on the state of the per-vector Mask bit.
13:11	RO	0h	Reserved.
10:0	RO	003h	Table Size System software reads this field to determine the MSI-X Table Size N, which is encoded as N-1. For example, a returned value of '0000000011' indicates a table size of 4. NTB table size is 4, encoded as a value of 003h

6.13.5.31 TABLEOFF_BIR: MSI-X Table Offset and BAR Indicator

TABLEOFF_BIR Bus: M Device: 0 Function: 0 Offset: 84 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 584 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 584			
Bit	Attr	Default	Description
31:3	RO	00000800h	Table Offset MSI-X Table Structure is at offset 16K from the SB01BASE BAR address. Section 6.13.8.1, "PMSIXTBL[0:3]: Primary MSI-X Table Address Register 0 - 3" on page 404 for the start of details relating to MSI-X registers. NOTE: Offset placed at 16K so that it can also be visible through the primary BAR for debug purposes.



Bit	Attr	Default	Description
TABLEOFF_BIR Bus: M Device: 0 Function: 0 Offset: 84 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 584 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 584			
2:0	RO	0h	Table BIR Indicates which one of a function's Base Address registers, located beginning at 10h in Configuration Space, is used to map the function's MSI-X Table into Memory Space. BIR Value Base Address register 0 10h 1 14h 2 18h 3 1Ch 4 20h 5 24h 6 Reserved 7 Reserved For a 64-bit Base Address register, the Table BIR indicates the lower DWORD.

6.13.5.32 PBAOFF_BIR: MSI-X Pending Bit Array Offset and BAR Indicator

Bit	Attr	Default	Description
PBAOFF_BIR Bus: M Device: 0 Function: 0 Offset: 88 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 588 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 588			
31:3	RO	00000A00h	Table Offset MSI-X PBA Structure is at offset 20K from the SB01BASE BAR address. See Section 6.13.9.4, "SMSICXPBA: Secondary MSI-X Pending Bit Array" on page 407 for details. Note: Offset placed at 20K so that it can also be visible through the primary BAR for debug purposes.
2:0	RO	0h	PBA BIR Indicates which one of a function's Base Address registers, located beginning at 10h in Configuration Space, is used to map the function's MSI-X Table into Memory Space. BIR Value Base Address register 0 10h 1 14h 2 18h 3 1Ch 4 20h 5 24h 6 Reserved 7 Reserved For a 64-bit Base Address register, the Table BIR indicates the lower DWORD.



6.13.5.33 PXPCAPID: PCI Express Capability Identity

The PCI Express Capability List register enumerates the PCI Express Capability structure in the PCI 3.0 configuration space.

PXPCAPID Bus: M Device: 0 Function: 0 Offset: 90 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 590 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 590			
Bit	Attr	Default	Description
7:0	RO	10h	Capability ID Provides the PCI Express capability ID assigned by PCI-SIG. Required by <i>PCI Express* Base Specification</i> , Revision 2.0 to be this value.

6.13.5.34 PXPNTPTR: PCI Express Next Pointer

The PCI Express Capability List register enumerates the PCI Express Capability structure in the PCI 3.0 configuration space.

PXPNTPTR Bus: M Device: 0 Function: 0 Offset: 91 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 591 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 591			
Bit	Attr	Default	Description
7:0	RW-O	E0h	Next Ptr This field is set to the PCI Power Management capability.

6.13.5.35 PXPCAP: PCI Express Capabilities

The PCI Express Capabilities register identifies the PCI Express device type and associated capabilities.

PXPCAP Bus: M Device: 0 Function: 0 Offset: 92 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 592 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 592			
Bit	Attr	Default	Description
15:14	RV	0h	Reserved
13:9	RO	0h	Interrupt Message Number Applies only to the RPs. This field indicates the interrupt message number that is generated for Power Management/Hot Plug events. When there are more than one MSI interrupt Number, this register field is required to contain the offset between the base Message Data and the MSI Message that is generated when the status bits in the slot status register or RP status registers are set. IIO assigns the first vector for Power Management/Hot Plug events and so this field is set to 0.
8	RW-O	0b	Slot Implemented Applies only to the RPs for NTB this value is kept at 0b. 1: indicates that the PCI Express link associated with the port is connected to a slot. 0: indicates no slot is connected to this port. This register bit is of type 'write once' and is controlled by BIOS/special initialization firmware.



PXPCAP Bus: M Device: 0 Function: 0 Offset: 92 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 592 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 592			
Bit	Attr	Default	Description
7:4	RO	0000b	Device/Port Type This field identifies the type of device. 0000b = PCI Express Endpoint.
3:0	RW-O	2h	Capability Version This field identifies the version of the PCI Express capability structure. Set to 2h for PCI Express devices for compliance with the extended base registers.

6.13.5.36 DEVCAP: PCI Express Device Capabilities Register

The PCI Express Device Capabilities register identifies device specific information for the device.

DEVCAP Bus: M Device: 0 Function: 0 Offset: 94 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 594 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 594			
Bit	Attr	Default	Description
31:29	RV	0h	Reserved
28	RO	0b	Function Level Reset Capability A value of 1b indicates the Function supports the optional Function Level Reset mechanism. NTB does not support this functionality
27:26	RO	0h	Captured Slot Power Limit Scale Does not apply to RPs or integrated devices This value is hardwired to 00h NTB is required to be able to receive the Set_Slot_Power_Limit message without error but simply discard the Message value. <i>Note:</i> Components with Endpoint, Switch, or PCI Express-PCI Bridge Functions that are targeted for integration on an adapter where total consumed power is below the lowest limit defined for the targeted form factor are permitted to ignore Set_Slot_Power_Limit Messages, and to return a value of 0 in the Captured Slot Power Limit Value and Scale fields of the Device Capabilities register.
25:18	RO	00h	Captured Slot Power Limit Value Does not apply to RPs or integrated devices This value is hardwired to 00h NTB is required to be able to receive the Set_Slot_Power_Limit message without error but simply discard the Message value. <i>Note:</i> Components with Endpoint, Switch, or PCI Express-PCI Bridge Functions that are targeted for integration on an adapter where total consumed power is below the lowest limit defined for the targeted form factor are permitted to ignore Set_Slot_Power_Limit Messages, and to return a value of 0 in the Captured Slot Power Limit Value and Scale fields of the Device Capabilities register
15	RO	1b	Role Based Error Reporting IIO is 1.1 compliant and so supports this feature
14	RO	0b	Power Indicator Present on Device Does not apply to RPs or integrated devices
13	RO	0b	Attention Indicator Present Does not apply to RPs or integrated devices
12	RO	0b	Attention Button Present Does not apply to RPs or integrated devices



Bit	Attr	Default	Description
DEVCAP Bus: M Device: 0 Function: 0 Offset: 94 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 594 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 594			
11:9	RO	110b	Endpoint L1 Acceptable Latency This field indicates the acceptable total latency that an Endpoint can withstand due to the transition from L0s state to the L0 state. It is essentially an indirect measure of the Endpoint's internal buffering. Power management software uses the reported L0s Acceptable Latency number to compare against the L0s exit latencies reported by all components comprising the data path from this Endpoint to the Root Complex Root Port to determine whether ASPM L0s entry can be used with no loss of performance. Defined encodings are: 000b Maximum of 64 ns 001b Maximum of 128 ns 010b Maximum of 256 ns 011b Maximum of 512 ns 100b Maximum of 1 us 101b Maximum of 2 us 110b Maximum of 4 us 111b No limit
8:6	RO	000b	Endpoint L0s Acceptable Latency This field indicates the acceptable latency that an Endpoint can withstand due to the transition from L1 state to the L0 state. It is essentially an indirect measure of the Endpoint's internal buffering. Power management software uses the reported L1 Acceptable Latency number to compare against the L1 Exit Latencies reported (see below) by all components comprising the data path from this Endpoint to the Root Complex Root Port to determine whether ASPM L1 entry can be used with no loss of performance. Defined encodings are: 000b Maximum of 1 us 001b Maximum of 2 us 010b Maximum of 4 us 011b Maximum of 8 us 100b Maximum of 16 us 101b Maximum of 32 us 110b Maximum of 64 us 111b No limit
5	RO	1b	Extended Tag Field Supported IIO devices support 8-bit tag1 = Maximum Tag field is 8 bits 0 = Maximum Tag field is 5 bits
4:3	RO	00b	Phantom Functions Supported IIO does not support phantom functions. 00b = No Function Number bits are used for Phantom Functions
2:0	RO	001b	Max Payload Size Supported IIO supports 256B payloads on PCI Express ports 001b = 256 bytes max payload size



6.13.5.37 DEVCTRL: PCI Express Device Control

(PCIe NTB Secondary) The PCI Express Device Control register controls PCI Express specific capabilities parameters associated with the device

Bit	Attr	Default	Description
DEVCTRL Bus: M Device: 0 Function: 0 Offset: 98 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 598 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 598			
15	RV	0h	Reserved
14:12	RO	000b	Max_Read_Request_Size Express/DMI ports in IIO do not generate requests greater than 128B and this field is ignored.
11	RO	0b	Enable No Snoop Not applicable since the NTB is never the originator of a TLP. This bit has no impact on forwarding of NoSnoop attribute on peer requests.
10	RO	0b	Auxiliary Power Management Enable Not applicable to IIO
9	RO	0b	Phantom Functions Enable Not applicable to IIO since it never uses phantom functions as a requester.
8	RW	0h	Extended Tag Field Enable This bit enables the PCI Express/DMI ports to use an 8-bit Tag field as a requester.
7:5	RW	000b	Max Payload Size This field is set by configuration software for the maximum TLP payload size for the PCI Express port. As a receiver, the IIO must handle TLPs as large as the set value. As a requester (i.e. for requests where IIO's own RequesterID is used), it must not generate TLPs exceeding the set value. Permissible values that can be programmed are indicated by the Max_Payload_Size_Supported in the Device Capabilities register: 000: 128B max payload size 001: 256B max payload size (applies only to standard PCI Express ports and DMI port aliases to 128B) others: alias to 128B This field is RW for PCI Express ports.
4	RO	0b	Enable Relaxed Ordering When set, the NTB does not send any outbound traffic with RO bit set, regardless of whether it was forwarded from the local CPU or from a local peer source
3	RW	0b	Unsupported Request Reporting Enable Applies only to the PCI Express/DMI ports. This bit controls the reporting of unsupported requests that IIO itself detects on requests its receives from a PCI Express/DMI port. 0: Reporting of unsupported requests is disabled 1: Reporting of unsupported requests is enabled.
2	RW	0b	Fatal Error Reporting Enable Applies only to the PCI Express RP/PCI Express NTB Secondary interface/DMI ports. Controls the reporting of fatal errors that IIO detects on the PCI Express/DMI interface. 0: Reporting of Fatal error detected by device is disabled 1: Reporting of Fatal error detected by device is enabled For the PCI Express/DMI ports, this bit is not used to control the reporting of other internal component uncorrectable fatal errors (at the port unit) in any way.



DEVCTRL Bus: M Device: 0 Function: 0 Offset: 98 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 598 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 598			
Bit	Attr	Default	Description
1	RW	0b	Non Fatal Error Reporting Enable Applies only to the PCI Express RP/PCI Express NTB Secondary interface/DMI ports. Controls the reporting of non-fatal errors that IIO detects on the PCI Express/DMI interface. 0: Reporting of Non Fatal error detected by device is disabled 1: Reporting of Non Fatal error detected by device is enabled For the PCI Express/DMI ports, this bit is not used to control the reporting of other internal component uncorrectable non-fatal errors (at the port unit) in any way.
0	RW	0b	Correctable Error Reporting Enable Applies only to the PCI Express RP/PCI Express NTB Secondary interface/DMI ports. Controls the reporting of correctable errors that IIO detects on the PCI Express/DMI interface. 0: Reporting of link Correctable error detected by the port is disabled 1: Reporting of link Correctable error detected by port is enabled For the PCI Express/DMI ports, this bit is not used to control the reporting of other internal component correctable errors (at the port unit) in any way.

6.13.5.38 DEVSTS: PCI Express Device Status

The PCI Express Device Status register provides information about PCI Express device specific parameters associated with the device.

DEVSTS Bus: M Device: 0 Function: 0 Offset: 9A Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 59A Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 59A			
Bit	Attr	Default	Description
5	RO	0h	Transactions Pending
4	RO	0b	AUX Power Detected Does not apply to IIO
3	RW1C	0b	Unsupported Request Detected This bit applies only to the root/DMI ports. This bit indicates that the NTB secondary detected an Unsupported Request. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control Register. 1: Unsupported Request detected at the device/port. These unsupported requests are NP requests inbound that the RP received and it detected them as unsupported requests (e.g. address decoding failures that the RP detected on a packet, receiving inbound lock reads, BME bit is clear etc.). Note that this bit is not set on peer2peer completions with UR status that are forwarded by the RP to the PCIe link. 0: No unsupported request detected by the RP
2	RW1C	0b	Fatal Error Detected This bit indicates that a fatal (uncorrectable) error is detected by the NTB secondary device. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. 1: Fatal errors detected 0: No Fatal errors detected



DEVSTS				
Bus: M		Device: 0	Function: 0	Offset: 9A
Bus: 0		Device: 3	Function: 0	MMIO BAR: PB01BASE
		Offset: 59A		
Bus: 0		Device: 3	Function: 0	MMIO BAR: SB01BASE
		Offset: 59A		
Bit	Attr	Default	Description	
1	RW1C	0b	Non Fatal Error Detected This bit gets set if a non-fatal uncorrectable error is detected by the NTB secondary device. Errors are logged in this register regardless of whether error reporting is enabled or not in the Device Control register. 1: Non Fatal errors detected 0: No non-Fatal Errors detected	
0	RW1C	0b	Correctable Error Detected This bit gets set if a correctable error is detected by the NTB secondary device. Errors are logged in this register regardless of whether error reporting is enabled or not in the PCI Express Device Control register. 1: correctable errors detected 0: No correctable errors detected	

6.13.5.39 LNKCAP: PCI Express Link Capabilities

The Link Capabilities register identifies the PCI Express specific link capabilities.

LNKCAP				
Bus: M		Device: 0	Function: 0	Offset: 9C
Bus: 0		Device: 3	Function: 0	MMIO BAR: PB01BASE
		Offset: 59C		
Bus: 0		Device: 3	Function: 0	MMIO BAR: SB01BASE
		Offset: 59C		
Bit	Attr	Default	Description	
31:24	RO	00h	Port Number This field indicates the PCI Express port number for the link and is initialized by software/BIOS. Note: This register bit is a RW-O register from the host side. It must be loaded by BIOS in the primary side equivalent register. This register is RO from the secondary side of the NTB.	
21	RO	0b	Link Bandwidth Notification Capability A value of 1b indicates support for the Link Bandwidth Notification status and interrupt mechanisms.	
20	RO	1b	Data Link Layer Link Active Reporting Capable IIO supports reporting status of the data link layer so software knows when it can enumerate a device on the link or otherwise know the status of the link.	
19	RO	1b	Surprise Down Error Reporting Capable IIO supports reporting a surprise down error condition	
18	RO	0b	Clock Power Management Does not apply to IIO.	




Bit	Attr	Default	Description
LNKCAP Bus: M Device: 0 Function: 0 Offset: 9C Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 59C Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 59C			
17:15	RO	010b	L1 Exit Latency This field indicates the L1 exit latency for the given PCI Express port. It indicates the length of time this port requires to complete transition from L1 to L0. 000: Less than 1 us 001: 1 us to less than 2 us 010: 2 us to less than 4 us 011: 4 us to less than 8 us 100: 8 us to less than 16 us 101: 16 us to less than 32 us 110: 32 us to 64 us 111: More than 64 us Note: This register bit is a RW-O register from the host side. It must be loaded by BIOS in the primary side equivalent register. This register is RO from the secondary side of the NTB.
14:12	RO	011b	L0s Exit Latency This field indicates the L0s exit latency (i.e L0s to L0) for the PCI Express port. 000: Less than 64 ns 001: 64 ns to less than 128 ns 010: 128 ns to less than 256 ns 011: 256 ns to less than 512 ns 100: 512 ns to less than 1 is 101: 1 is to less than 2 is 110: 2 is to 4 is 111: More than 4 is Note: This register bit is a RW-O register from the host side. It must be loaded by BIOS in the primary side equivalent register. This register is RO from the secondary side of the NTB.
11:10	RO	11b	Active State Link Power Management Support This field indicates the level of active state power management supported on the given PCI Express port. 00: Disabled 01: L0s Entry Supported 10: Reserved 11: L0s and L1 Supported Note: This register bit is a RW-O register from the host side. It must be loaded by BIOS in the primary side equivalent register. This register is RO from the secondary side of the NTB.
9:4	RO	8h	Maximum Link Width This field indicates the maximum width of the given PCI Express Link attached to the port. 000001: x1 000010: x2 000100: x4 001000: x8 010000: x16 Others: Reserved Note: This register bit is a RW-O register from the host side. It must be loaded by BIOS in the primary side equivalent register. This register is RO from the secondary side of the NTB.
3:0	RO	0011b	Maximum Link Speed This field indicates the maximum link speed of this Port. The encoding is the binary value of the bit location in the Supported Link Speeds Vector (in LNKCAP2) that corresponds to the maximum link speed.



6.13.5.40 LNKCON: PCI Express Link Control

The PCI Express Link Control register controls the PCI Express Link specific parameters.

Bit	Attr	Default	Description
LNKCON Bus: M Device: 0 Function: 0 Offset: A0 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 5A0 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 5A0			
15:10	RV	0h	Reserved
9	RO	0b	Hardware Autonomous Width Disable IIO never changes a configured link width for reasons other than reliability.
8	RO	0b	Enable Clock Power Management N/A to IIO
7	RW-V	0b	Extended Synch This bit when set forces the transmission of additional ordered sets when exiting LOs and when in recovery. See PCI Express Base Specification, Revision 2.0 for details.
6	RW-V	0b	Common Clock Configuration IIO does nothing with this bit
3	RO	0b	Read Completion Boundary Set to zero to indicate IIO could return read completions at 64B boundaries. Note: NTB is not PCIe compliant in this respect. NTB is only capable of 64B RCB. If connecting to non IA IP and the IP does the optional 128B RCB check on received packets, packets will be seen as malformed. This is not an issue with any Intel IP.
1:0	RW	00b	Active State Link Power Management Control When 01b or 11b, LOs on transmitter is enabled, otherwise it is disabled.



6.13.5.41 LNKSTS: PCI Express Link Status

The PCI Express Link Status register provides information on the status of the PCI Express Link such as negotiated width, training, and so forth.

Bit	Attr	Default	Description
LNKSTS Bus: M Device: 0 Function: 0 Offset: A2 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 5A2 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 5A2			
13	RO	0b	Data Link Layer Link Active Set to 1b when the Data Link Control and Management State Machine is in the DL_Active state, 0b otherwise. On a downstream port or upstream port, when this bit is 0b, the transaction layer associated with the link will abort all transactions that would otherwise be routed to that link.
12		1b	Slot Clock Configuration This bit indicates whether IIO receives clock from the same xtal that also provides clock to the device on the other end of the link. 1: indicates that same xtal provides clocks to devices on both ends of the link 0: indicates that different xtals provide clocks to devices on both ends of the link Note: This register bit is a RW-O register from the host side. It must be loaded by BIOS in the primary side equivalent register. This register is RO from the secondary side of the NTB. 0_3_0_PB01BASE: Attr: RO Default: 1b 0_3_0_SB01BASE: Attr: RW-O Default: 1b
11	RO	0b	Link Training This field indicates the status of an ongoing link training session in the PCI Express port. 0: LTSSM has exited the recovery/configuration state 1: LTSSM is in recovery/configuration state or the Retrain Link was set but training has not yet begun. The IIO hardware clears this bit once LTSSM has exited the recovery/configuration state.
9:4	RO	0h	Negotiated Link Width This field indicates the negotiated width of the given PCI Express link after training is completed. Only x1, x2, x4, x8 and x16 link width negotiations are possible in IIO. A value of 0x01 in this field corresponds to a link width of x1, 0x02 indicates a link width of x2 and so on, with a value of 0x16 for a link width of x16. The value in this field is reserved and could show any value when the link is not up. Software determines if the link is up or not by reading bit 13 of this register.
3:0	RO-V	1h	Current Link Speed This field indicates the negotiated Link speed of the given PCI Express Link. 0001: 2.5 Gbps 0010: 5 Gbps 0011: 8Gbps Others: Reserved The value in this field is not defined when the link is not up. Software determines if the link is up or not by reading bit 13 of this register.

6.13.5.42 SSCNTL: Secondary Side Control

This register provides secondary side control of NTB functions.




SSCNTL Bus: M Device: 0 Function: 0 Offset: D4 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 5D4 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 5D4			
Bit	Attr	Default	Description
0	RW	0b	NTB Secondary side MSI-X Single Message Vector: This bit when set, causes only a single MSI-X message to be generated if MSI-X is enabled. This bit affects the default value of the MSI-X Table Size field in the 'SMSIXTBLO-3: Secondary MSI-X Table Address Register 0 - 3.

6.13.5.43 PMCAP: Power Management Capabilities

The Power Management Capabilities Register defines the capability ID, next pointer and other power management related support. The following Power Management registers / capabilities are added for software compliance.

PMCAP Bus: M Device: 0 Function: 0 Offset: E0 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 5E0 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 5E0			
Bit	Attr	Default	Description
31:27	RO	0h	PME Support Indicates the Power Management states within which the function is capable of sending a PME message. NTB secondary side does not forward PME messages. Bit 31 = D3cold Bit 30 = D3hot Bit 29 = D2 Bit 28 = D1 Bit 27 = D0
26	RO	0b	D2 Support IIO does not support power management state D2.
25	RO	0b	D1 Support IIO does not support power management state D1.
24:22	RO	000b	AUX Current Device does not support auxiliary current
21	RO	0b	Device Specific Initialization Device initialization is not required
19	RO	0b	PME Clock This field is hardwired to 0h as it does not apply to PCI Express.
18:16	RO	011b	Version This field is set to 3h as version number for all PCI Express ports.
15:8	RO	00h	Next Capability Pointer This is the last capability in the chain and hence set to 0.
7:0	RO	01h	Capability ID Provides the Power Management capability ID assigned by PCI-SIG.



6.13.5.44 PMCSR: Power Management Control and Status

This register provides status and control information for Power Management events in the PCI Express port of the IIO.

PMCSR Bus: M Device: 0 Function: 0 Offset: E4 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 5E4 Offset: 5E4			
Bit	Attr	Default	Description
31:24	RO	00h	Data Not relevant for IIO
23	RO	0h	Bus Power/Clock Control Enable This field is hardwired to 0h as it does not apply to PCI Express.
22	RO	0h	B2/B3 Support This field is hardwired to 0h as it does not apply to PCI Express.
21:16	RV	0h	Reserved
15	RO	0h	PME Status Applies only to RPs This bit is hard-wired to read-only 0, since this function does not support PME# generation from any power state. This PME Status is a sticky bit. This bit is set, independent of the PMEEN bit defined below, on an enabled PCI Express hot-plug event provided the RP was in D3hot state. Software clears this bit by writing a '1' when it has been completed.
14:13	RO	0h	Data Scale Not relevant for IIO
12:9	RO	0h	Data Select Not relevant for IIO
8	RO	0h	PME Enable Applies only to RPs. 0: Disable ability to send PME messages when an event occurs 1: Enables ability to send PME messages when an event occurs
7:4	RV	0h	Reserved
3	RW-O	1b	Indicates IIO does not reset its registers when it transitions from D3hot to D0
2	RV	0h	Reserved
1:0	RW	0h	Power State This 2-bit field is used to determine the current power state of the function and to set a new power state as well. 00: D0 01: D1 (not supported by IIO) 10: D2 (not supported by IIO) 11: D3_hot If Software tries to write 01 or 10 to this field, the power state does not change from the existing power state (which is either D0 or D3hot) and nor do these bits 1:0 change value. All devices will respond to only Type 0 configuration transactions when in D3hot state (RP will not forward Type 1 accesses to the downstream link) and will not respond to memory/IO transactions (that is, D3hot state is equivalent to MSE/IOSE bits being clear) as target and will not generate any memory/IO/ configuration transactions as initiator on the primary bus (messages are still allowed to pass through).



6.13.5.45 PXP2CAP: Secondary PCI Express Extended Capability Header

PXP2CAP Bus: M Device: 0 Function: 0 Offset: 100			
Bit	Attr	Default	Description
31:20	RO	000h	Next Capability Offset This field contains the offset to the next PCI Express Extended Capability structure or 000h if no other items exist in the linked list of capabilities.
19:16	RO	1h	Capability Version This field is a PCI-SIG defined version number that indicates the version of the Capability structure present. Must be 1h for this version of the specification.
15:0	RO	0000h	PCI Express Extended Capability ID This field is a PCI SIG defined ID number that indicates the nature and format of the Extended Capability. PCI Express Extended Capability ID for the Secondary PCI Express Extended Capability is 0x0019h

6.13.5.46 LNERRSTS: Lane Error Status Register

Device 0, Function 0, Offset 104h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.108](#).

6.13.5.47 LN[0:3]EQ: Lane 0 through Lane 3 Equalization Control

Device 0, Function 0, Offset 108h, 10Ah, 10Ch, 10Eh. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.109](#).

6.13.5.48 LN[4:7]EQ: Lane 4 through Lane 7 Equalization Control

Device 0, Function 0, Offset 110h, 112h, 114h, 116h. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.110](#).

6.13.5.49 LN[8:15]EQ: Lane 8 though Lane 15 Equalization Control

Device 0, Function 0, Offset 118h, 11Ah, 11Ch, 11Eh. This register exist in both RP and NTB modes. It is documented in RP [Section 6.2.111](#).

6.13.6 NTB Shadowed MMIO Space

All shadow registers are visible from the primary side of the NTB. Only some of the shadow registers are visible from the secondary side of the NTB. See each register description for visibility.

Table 6-11. NTB MMIO Shadow Registers (Sheet 1 of 2)

PBAR2LMT	0h	SPAD0	80h
	4h	SPAD1	84h
PBAR4LMT	8h	SPAD2	88h
	Ch	SPAD3	8Ch
PBAR2XLAT	10h	SPAD4	90h
	14h	SPAD5	94h
PBAR4XLAT	18h	SPAD6	98h
	1Ch	SPAD7	9Ch




Table 6-11. NTB MMIO Shadow Registers (Sheet 2 of 2)

SBAR2LMT		20h	SPAD8	A0h
		24h	SPAD9	A4h
SBAR4LMT		28h	SPAD10	A8h
		2Ch	SPAD11	ACh
SBAR2XLAT		30h	SPAD12	B0h
		34h	SPAD13	B4h
SBAR4XLAT		38h	SPAD14	B8h
		3Ch	SPAD15	BCh
SBAR0BASE		40h	SPADSEMA4	C0h
		44h		C4h
SBAR2BASE		48h		C8h
		4Ch		CCh
SBAR4BASE		50h	RSDBMSIXV70	D0h
		54h	RSDBMSIXV158	D4h
NTBCNTL		58h		D8h
CBDFD	SBD	5Ch		DCh
PDBMSK	PDOORBELL	60h		E0h
SDBMSK	SDOORBELL	64h		E4h
		68h		E8h
		6Ch		ECh
USMEMMISS		70h		F0h
		74h		F4h
		78h		F8h
		7Ch		FCh

Table 6-12. NTB MMIO Map (Sheet 1 of 2)

B2BSPAD0	100h		180h
B2BSPAD1	104h		184h
B2BSPAD2	108h		188h
B2BSPAD3	10Ch		18Ch
B2BSPAD4	110h		190h
B2BSPAD5	114h		194h
B2BSPAD6	118h		198h
B2BSPAD7	11Ch		19Ch
B2BSPAD8	120h		1A0h
B2BSPAD9	124h		1A4h
B2BSPAD10	128h		1A8h
B2BSPAD11	12Ch		1ACh
B2BSPAD12	130h		1B0h
B2BSPAD13	134h		1B4h
B2BSPAD14	138h		1B8h
B2BSPAD15	13Ch		1BCh
	B2BDOORBELL	140h	1C0h




Table 6-12. NTB MMIO Map (Sheet 2 of 2)

B2BBAR0XLAT	144h		1C4h
	148h		1C8h
	14Ch		1CCh
	150h		1D0h
	154h		1D4h
	158h		1D8h
	15Ch		1DCh
	160h		1E0h
	164h		1E4h
	168h		1E8h
	16Ch		1ECh
	170h		1F0h
	174h		1F4h
	178h		1F8h
	17Ch		1FCh

6.13.7 NTB Primary/Secondary Host MMIO Registers

6.13.7.1 PBAR2LMT: Primary BAR 2/3 Limit

PBAR2LMT Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 0 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 0			
Bit	Attr	Default	Description
47:12		00000000h	<p>Primary BAR 2/3 Limit</p> <p>Value representing the size of the memory window exposed by Primary BAR 2/3. A value of 00h will disable this register's functionality, resulting in a BAR window equal to that described by the BAR.</p> <p>This register contains a value used to limit the size of the window exposed by 64-bit BAR 2/3 to a size less than the power-of-two expressed in the Primary BAR 2/3 pair. This register is written by the NTB device driver and will contain the formulated sum of the base address plus the size of the BAR. This final value equates to the highest address that will be accepted through this port. Accesses to the memory area above this register will return Unsupported Request.</p> <p>Note: If the value in PBAR2LMT is set to a value less than the value in PB23BASE hardware will force the value in PBAR2LMT to be zero and the full size of the window defined by PBAR23SZ will be used. If the value in PBAR2LMT is set equal to the value in PB23BASE the memory window for PB23BASE is disabled. If the value in PBAR2LMT is set to a value greater than the value in the PB23BASE plus 2^PBAR23SZ hardware will force the value in PBAR2LMT to be zero and the full size of the window defined by PBAR23SZ will be used. If PBAR2LMT is zero the full size of the window defined by PBAR23SZ will be used.</p> <p>This field is RW from PB01BASE (primary side window) and RO from SB01BASE (secondary side window). B01BASE: Attr: RW Default: 00000000h SB01BASE: Attr: RO Default: 00000000h</p>



6.13.7.2 PBAR4LMT: Primary BAR 4/5 Limit

PBAR4LMT Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 8 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 8			
Bit	Attr	Default	Description
47:12		000000000h	<p>Primary BAR 4/5 Limit</p> <p>Value representing the size of the memory window exposed by Primary BAR 4/5. A value of 00h will disable this register's functionality, resulting in a BAR window equal to that described by the BAR.</p> <p>This register is written by the NTB device driver and will contain the formulated sum of the base address plus the size of the BAR. This final value equates to the highest address that will be accepted through this port. Accesses to the memory area above this register will return Unsupported Request.</p> <p>Notes:</p> <p>If the value in PBAR4LMT is set to a value less than the value in PB45BASE hardware will force the value in PBAR4LMT to be zero and the full size of the window defined by PBAR45SZ will be used.</p> <p>If the value in PBAR4LMT is set equal to the value in PB45BASE the memory window for PB45BASE is disabled.</p> <p>If the value in PBAR4LMT is set to a value greater than the value in the PB45BASE plus 2^PBAR45SZ hardware will force the value in PBAR4LMT to be zero and the full size of the window defined by PBAR45SZ will be used.</p> <p>If PBAR4LMT is zero the full size of the window defined by PBAR45SZ will be used.</p> <p>This field is RW from PB01BASE (primary side window) and RO from SB01BASE (secondary side window).</p> <p>PB01BASE: Attr: RW Default: 000000000h SB01BASE: Attr: RO Default: 000000000h</p>

6.13.7.3 PBAR2XLAT: Primary BAR 2/3 Translate

PBAR2XLAT Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 10 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 10			
Bit	Attr	Default	Description
63:12	RW	0000000000000h	<p>Primary BAR 2/3 Translate</p> <p>The aligned base address into Secondary side memory.</p> <p>This register contains a value used to direct accesses into the memory located on the Secondary side of the NTB made from the Primary side of the NTB through the window claimed by BAR 2/3 on the primary side. The register contains the base address of the Secondary side memory window.</p> <p>Notes:</p> <p>There is no hardware enforced limit for this register, care must be taken when setting this register to stay within the addressable range of the attached system.</p> <p>Default is set to 256 GB</p> <p>The number of bits that are writable in this register is dictated by the value loaded into the PBAR23SZ register by the BIOS at initialization time (before BIOS PCI enumeration). PBAR23SZ indicates the lowest order bit of this register field that is writeable where valid values are 12-39. If PBAR23SZ is set to 12, all bits are writeable. If set to 39, then bits 38:12 are Read Only and will return values of 0.</p> <p>For the special case where PBAR23SZ = '0', bits 63:0 are all RO= '0' resulting in the BAR being disabled.</p> <p>The lowest order address bit is 12 to enforce a minimum granularity of 4 KB.</p>



6.13.7.4 PBAR4XLAT: Primary BAR 4/5 Translate

Bit	Attr	Default	Description
PBAR4XLAT Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 18 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 18			
63:12	RW	0000000000000h	<p>Primary BAR 4/5 Translate</p> <p>The aligned base address into Secondary side memory.</p> <p>This register contains a value used to direct accesses into the memory located on the Secondary side of the NTB made from the Primary side of the NTB through the window claimed by BAR 4/5 on the primary side. The register contains the base address of the Secondary side memory window.</p> <p>Notes:</p> <p>There is no hardware enforced limit for this register, care must be taken when setting this register to stay within the addressable range of the attached system.</p> <p>Default is set to 512 GB</p> <p>The number of bits that are writable in this register is dictated by the value loaded into the PBAR45SZ register by the BIOS at initialization time (before BIOS PCI enumeration). PBAR45SZ indicates the lowest order bit of this register field that is writeable where valid values are 12-39. If PBAR45SZ is set to 12, all bits are writeable. If set to 39, then bits 38:12 are Read Only and will return values of 0.</p> <p>For the special case where PBAR45SZ = '0', bits 63:0 are all RO= '0' resulting in the BAR being disabled.</p> <p>The lowest order address bit is 12 to enforce a minimum granularity of 4 KB.</p>

6.13.7.5 SBAR2LMT: Secondary BAR 2/3 Limit

Bit	Attr	Default	Description
SBAR2LMT Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 20 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 20			
63:12	RW-V	0000000000000h	<p>Secondary BAR 2/3 Limit</p> <p>Value representing the size of the memory window exposed by Secondary BAR 2/3. A value of 00h will disable this register's functionality, resulting in a BAR window equal to that described by the BAR.</p> <p>This register contains a value used to limit the size of the window exposed by 64-bit BAR 2/3 to a size less than the power-of-two expressed in the Secondary BAR 2/3 pair. This register is written by the NTB device driver and will contain the formulated sum of the base address plus the size of the BAR. This final value equates to the highest address that will be accepted through this port. Accesses to the memory area above this register will return Unsupported Request.</p> <p>Notes:</p> <p>If the value in SBAR2LMT is set to a value less than the value in SB23BASE hardware will force the value in SBAR2LMT to be zero and the full size of the window defined by SBAR23SZ will be used.</p> <p>If the value in SBAR2LMT is set equal to the value in SB23BASE the memory window for SB23BASE is disabled.</p> <p>If the value in SBAR2LMT is set to a value greater than the value in the SB23BASE plus 2^SBAR23SZ hardware will force the value in SBAR2LMT to be zero and the full size of the window defined by SBAR23SZ will be used.</p> <p>If SBAR2LMT is zero the full size of the window defined by SBAR23SZ will be used.</p>



6.13.7.6 SBAR4LMT: Secondary BAR 4/5 Limit

SBAR4LMT Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 28 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 28			
Bit	Attr	Default	Description
63:12	RW-V	00000000000000h	Secondary BAR 4/5 Limit Value representing the size of the memory window exposed by Secondary BAR 4/5. A value of 00h will disable this register's functionality, resulting in a BAR window equal to that described by the BAR. This register contains a value used to limit the size of the window exposed by 64-bit BAR 4/5 to a size less than the power-of-two expressed in the Secondary BAR 4/5 pair. This register is written by the NTB device driver and will contain the formulated sum of the base address plus the size of the BAR. This final value equates to the highest address that will be accepted through this port. Accesses to the memory area above this register will return Unsupported Request. Notes: If the value in SBAR4LMT is set to a value less than the value in SB45BASE hardware will force the value in SBAR4LMT to be zero and the full size of the window defined by SBAR45SZ will be used. If the value in SBAR4LMT is set equal to the value in SB45BASE the memory window for SB45BASE is disabled. If the value in SBAR4LMT is set to a value greater than the value in the SB45BASE plus 2^SBAR45SZ hardware will force the value in SBAR4LMT to be zero and the full size of the window defined by SBAR45SZ will be used. If SBAR4LMT is zero the full size of the window defined by SBAR45SZ will be used.

6.13.7.7 SBAR2XLAT: Secondary BAR 2/3 Translate

This register contains a value used to direct accesses into the memory located on the Primary side of the NTB made from the Secondary side of the NTB through the window claimed by BAR 2/3 on the secondary side. The register contains the base address of the Primary side memory window.

SBAR2XLAT Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 30 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 30			
Bit	Attr	Default	Description
63:12	RW-L	00000000000000h	Secondary BAR 2/3 Translate The aligned base address into Primary side memory. Notes: Attr will appear as RW to SW The number of bits that are writable in this register is dictated by the value loaded into the SBAR23SZ register by the BIOS at initialization time (before BIOS PCI enumeration). SBAR23SZ indicates the lowest order bit of this register field that is writeable where valid values are 12-39. If SBAR23SZ is set to 12, all bits are writeable. If set to 39, then bits 38:12 are Read Only and will return values of 0. For the special case where SBAR23SZ = '0', bits 63:0 are all RO='0' resulting in the BAR being disabled. The lowest order address bit is 12 to enforce a minimum granularity of 4 KB.



6.13.7.8 SBAR4XLAT: Secondary BAR 4/5 Translate


This register contains a value used to direct accesses into the memory located on the Primary side of the NTB made from the Secondary side of the NTB through the window claimed by BAR 4/5 on the secondary side. The register contains the base address of the Primary side memory window.

Bit	Attr	Default	Description
SBAR4XLAT Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 38 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 38			
63:12	RW-L	0000000000000h	Secondary BAR 4/5Translate The aligned base address into Primary side memory. Attr will appear as RW to SW Notes: The number of bits that are writable in this register is dictated by the value loaded into the SBAR45SZ register by the BIOS at initialization time (before BIOS PCI enumeration). SBAR45SZ indicates the lowest order bit of this register field that is writeable where valid values are 12-39. If SBAR45SZ is set to 12, all bits are writeable. If set to 39, then bits 38:12 are Read Only and will return values of 0. For the special case where SBAR45SZ = '0', bits 63:0 are all RO='0' resulting in the BAR being disabled. The lowest order address bit is 12 to enforce a minimum granularity of 4 KB.

6.13.7.9 SBAR0BASE: Secondary BAR 0/1 Base Address

This register is mirrored from the BAR 0/1 register pair in the Configuration Space of the Secondary side of the NTB. The register is used by the processor on the primary side of the NTB to examine and load the BAR 0/1 register pair on the Secondary side of the NTB.

Bit	Attr	Default	Description
SBAR0BASE Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 38 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 38			
63:13		0000000000000h	Secondary BAR 0/1 Base This register is reflected into the BAR 0/1 register pair in the Configuration Space of the Secondary side of the NTB. O_3_0_PB01BASE: Attr: RW-L Default: 0000000000000h O_3_0_SB01BASE: Attr: RW-L Default: 0000000000000h
3	RW-O	1b	Prefetchable 1: BAR points to Prefetchable memory (default) 0: BAR points to Non-Prefetchable memory
2:1	RO	10b	Type Memory type claimed by BAR 2/3 is 64-bit addressable.
0	RO	0b	Memory Space Indicator BAR resource is memory (as opposed to I/O).



6.13.7.10 SBAR2BASE: Secondary BAR 2/3 Base Address

This register is mirrored from the BAR 2/3 register pair in the Configuration Space of the Secondary side of the NTB. The register is used by the processor on the primary side of the NTB to examine and load the BAR 2/3 register pair on the Secondary side of the NTB.

SBAR2BASE Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 48 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 48			
Bit	Attr	Default	Description
63:12	RW	00000000000000h	Secondary BAR 2/3 Base This register is reflected into the BAR 2/3 register pair in the Configuration Space of the Secondary side of the NTB. Notes: The number of bits that are writable in this register is dictated by the value loaded into the SBAR23SZ register by the BIOS at initialization time (before BIOS PCI enumeration). SBAR23SZ indicates the lowest order bit of this register field that is writeable where valid values are 12-39. If SBAR23SZ is set to 12, all bits are writeable. If set to 39, then bits 38:12 are Read Only and will return values of 0. For the special case where SBAR23SZ = '0', bits 63:0 are all RO='0' resulting in the BAR being disabled. The lowest order address bit is 12 to enforce a minimum granularity of 4 KB.
11:4	RO	00h	Granularity must be at least 4KB.
3	RO	1b	Prefetchable BAR points to Prefetchable memory.
2:1	RO	10b	Type Memory type claimed by BAR 2/3 is 64-bit addressable.
0	RO	0b	Memory Space Indicator BAR resource is memory (as opposed to I/O).

6.13.7.11 SBAR4BASE: Secondary BAR 4/5 Base Address

This register is mirrored from the BAR 4/5 register pair in the Configuration Space of the Secondary side of the NTB. The register is used by the processor on the primary side of the NTB to examine and load the BAR 4/5 register pair on the Secondary side of the NTB.



SBAR4BASE				
Bus: 0		Device: 3	Function: 0	MMIO BAR: PB01BASE
		Offset: 50		
Bus: 0		Device: 3	Function: 0	MMIO BAR: SB01BASE
		Offset: 50		
Bit	Attr	Default	Description	
63:12	RW	000000 000000 0h	<p>Secondary BAR 4/5 Base</p> <p>This register is reflected into the BAR 4/5 register pair in the Configuration Space of the Secondary side of the NTB.</p> <p>Notes:</p> <p>The number of bits that are writable in this register is dictated by the value loaded into the SBAR45SZ register by the BIOS at initialization time (before BIOS PCI enumeration). SBAR45SZ indicates the lowest order bit of this register field that is writable where valid values are 12-39. If SBAR45SZ is set to 12, all bits are writable. If set to 39, then bits 38:12 are Read Only and will return values of 0. For the special case where SBAR45SZ = '0', bits 63:0 are all RO='0' resulting in the BAR being disabled.</p> <p>The lowest order address bit is 12 to enforce a minimum granularity of 4 KB.</p>	
11:4	RO	00h	Granularity must be at least 4 KB.	
3	RO	1b	<p>Prefetchable</p> <p>BAR points to Prefetchable memory.</p>	
2:1	RO	10b	<p>Type</p> <p>Memory type claimed by BAR 4/5 is 64-bit addressable.</p>	
0	RO	0b	<p>Memory Space Indicator</p> <p>BAR resource is memory (as opposed to I/O).</p>	

6.13.7.12 NTBCNTL: NTB Control

This register contains Control bits for the Non-transparent Bridge device.

NTBCNTL				
Bus: 0		Device: 3	Function: 0	MMIO BAR: PB01BASE
		Offset: 58		
Bus: 0		Device: 3	Function: 0	MMIO BAR: SB01BASE
		Offset: 58		
Bit	Attr	Default	Description	
31:11	RV	0h	Reserved	
10		0b	<p>Crosslink SBDF Disable Increment</p> <p>This bit determines if SBDF value on the DSD is incremented or not.</p> <p>0: the DSD will increment SBDF (to SBDF + 1)</p> <p>1: the DSD will leave the SBDF</p> <p>0_3_0_PB01BASE: Attr: RW-V Default: 0b</p> <p>0_3_0_SB01BASE: Attr: RO-V Default: 0b</p>	
9:8		00b	<p>BAR 4/5 Primary to Secondary Snoop Override Control</p> <p>This bit controls the ability to force all transactions within the Primary BAR 4/5 window going from the Primary side to the Secondary side to be snoop/no-snoop independent of the ATTR field in the TLP header.</p> <p>00: All TLP sent as defined by the ATTR field</p> <p>01: Force Snoop on all TLPs: ATTR field overridden to set the 'No Snoop' bit = 0 independent of the setting of the ATTR field of the received TLP.</p> <p>10: Force No-Snoop on all TLPs: ATTR field overridden to set the 'No Snoop' bit = 1 independent of the setting of the ATTR field of the received TLP.</p> <p>11: Reserved</p> <p>0_3_0_PB01BASE: Attr: RW-V Default: 00b</p> <p>0_3_0_SB01BASE: Attr: RO-V Default: 00b</p>	



NTBCNTL Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 58 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 58			
Bit	Attr	Default	Description
7:6		00b	BAR 4/5 Secondary to Primary Snoop Override Control This bit controls the ability to force all transactions within the Secondary BAR 4/5 window going from the Secondary side to the Primary side to be snoop/no-snoop independent of the ATTR field in the TLP header. 00: All TLP sent as defined by the ATTR field 01: Force Snoop on all TLPs: ATTR field overridden to set the 'No Snoop' bit = 0 independent of the setting of the ATTR field of the received TLP. 10: Force No-Snoop on all TLPs: ATTR field overridden to set the 'No Snoop' bit = 1 independent of the setting of the ATTR field of the received TLP. 11: Reserved Notes: This field is RW from PB01BASE (primary side window) and RO from SB01BASE (secondary side window). 0_3_0_PB01BASE: Attr: RW Default: 00b 0_3_0_SB01BASE: Attr: RO Default: 00b
5:4		00b	BAR 2/3 Primary to Secondary Snoop Override Control This bit controls the ability to force all transactions within the Primary BAR 2/3 window going from the Primary side to the Secondary side to be snoop/no-snoop independent of the ATTR field in the TLP header. 00: All TLP sent as defined by the ATTR field 01: Force Snoop on all TLPs: ATTR field overridden to set the 'No Snoop' bit = 0 independent of the setting of the ATTR field of the received TLP. 10: Force No-Snoop on all TLPs: ATTR field overridden to set the 'No Snoop' bit = 1 independent of the setting of the ATTR field of the received TLP. 11: Reserved 0_3_0_PB01BASE: Attr: RW-V Default: 00b 0_3_0_SB01BASE: Attr: RO-V Default: 00b
3:2		00b	BAR 2/3 Secondary to Primary Snoop Override Control This bit controls the ability to force all transactions within the Secondary BAR 2/3 window going from the Secondary side to the Primary side to be snoop/no-snoop independent of the ATTR field in the TLP header. 00: All TLP sent as defined by the ATTR field 01: Force Snoop on all TLPs: ATTR field overridden to set the 'No Snoop' bit = 0 independent of the setting of the ATTR field of the received TLP. 10: Force No-Snoop on all TLPs: ATTR field overridden to set the 'No Snoop' bit = 1 independent of the setting of the ATTR field of the received TLP. 11: Reserved Notes: This field is RW from PB01BASE (primary side window) and RO from SB01BASE (secondary side window). 0_3_0_PB01BASE: Attr: RW Default: 00b 0_3_0_SB01BASE: Attr: RO Default: 00b
1		1b	Secondary Link Disable Control This bit controls the ability to train the link on the secondary side of the NTB. This bit is used to make sure the primary side is up and operational before allowing transactions from the secondary side. 0: Link enabled 1: Link disabled Notes: This bit logically or'd with the LNKCON bit 4 This field is RW from PB01BASE (primary side window) and RO from SB01BASE (secondary side window). 0_3_0_PB01BASE: Attr: RW Default: 1b 0_3_0_SB01BASE: Attr: RO Default: 1b



NTBCNTL Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 58 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 58			
Bit	Attr	Default	Description
0		1b	Secondary Configuration Space Lockout Control This bit controls the ability to modify the Secondary side NTB configuration registers from the Secondary side link partner. 0: Secondary side can read and write secondary registers 1: Secondary side modifications locked out but reads are accepted Notes: This does not block MMIO space. This field is RW from PB01BASE (primary side window) and RO from SB01BASE (secondary side window). 0_3_0_PB01BASE: Attr: RW Default: 1b 0_3_0_SB01BASE: Attr: RO Default: 1b

6.13.7.13 SBDF: Secondary Bus, Device and Function

This register contains the Bus, Device and Function for the secondary side of the NTB when PPD.Port Definition is configured as NTB/NTB [Section 6.13.3.26, “PPD: PCIe Port Definition”](#)

SBDF Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 5C Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 5C			
Bit	Attr	Default	Description
15:8	RW	7Fh	Secondary Bus for the secondary side of the NTB port while in NTB mode Value to be used for the Bus number for ID-based routing. Hardware will leave the default value of 7Fh when this port is USD Hardware will increment the default value to 80h when this port is DSD
7:3	RW	00h	Secondary Device for the secondary side of the NTB port while in NTB mode Value to be used for the Device number for ID-based routing.
2:0	RW	0h	Secondary Function for the secondary side of the NTB port while in NTB mode Value to be used for the Function number for ID-based routing.



6.13.7.14 CBFDF: Captured Bus, Device and Function

Bit	Attr	Default	Description
CBFDF Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 5E Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 5E			
15:8	RO-V	00h	Secondary Bus Value to be used for the Bus number for ID-based routing. This register contains the Bus, Device and Function for the secondary side of the NTB when PPD.Port Definition is configured as NTB/RP. Notes: When configured as a NTB/RP, the NTB must capture the Bus and Device Numbers supplied with all Type 0 Configuration Write Requests completed by the NTB and supply these numbers in the Bus and Device Number fields of the Requester ID for all Requests initiated by the NTB. The Bus Number and Device Number may be changed at run time, and so it is necessary to re-capture this information with each and every Configuration Write Request. When configured as a NTB/RP, if NTB must generate a Completion prior to the initial device Configuration Write Request, 0's must be entered into the Bus Number and Device Number fields This register is only valid when configured as NTB/RP. This register has no meaning when configured as NTB/NTB or RP.
7:3	RO-V	00h	Secondary Device Value to be used for the Device number for ID-based routing.
2:0	RO-V	0h	Secondary Function Value to be used for the Function number for ID-based routing.

6.13.7.15 PDOORBELL: Primary Doorbell

This register contains the bits used to generate interrupts to the processor on the Primary side of the NTB.

Bit	Attr	Default	Description
PDOORBELL Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 60 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 60			
15		0h	Link State Interrupt This bit is set when a link state change occurs on the Secondary side of the NTB (Bit 0 of the NTBSTATUS register). This bit is cleared by writing a 1 from the Primary side of the NTB. Notes: This field is RW1C from PB01BASE (primary side window) and RO from SB01BASE (secondary side window). 0_3_0_PB01BASE: Attr: RW1C Default: 0h 0_3_0_SB01BASE: Attr: RO Default: 0h



PDOORBELL			
Bus: 0		Device: 3	Function: 0
		Offset: 60	MMIO BAR: PB01BASE
Bus: 0		Device: 3	Function: 0
		Offset: 60	MMIO BAR: SB01BASE
Bit	Attr	Default	Description
14:0		0000h	<p>Primary Doorbell Interrupts</p> <p>These bits are written by the processor on the Secondary side of the NTB to cause a doorbell interrupt to be generated to the processor on the Primary side of the NTB if the associated mask bit in the PDBMSK register is not set. A 1 is written to this register from the Secondary side of the NTB to set the bit, and to clear the bit a 1 is written from the Primary side of the NTB.</p> <p>Notes:</p> <p>If both INTx and MSI (NTB PCI CMD bit 10 and NTB MSI Capability bit 0) interrupt mechanisms are disabled software must poll for status since no interrupts of either type are generated.</p> <p>This field is RW1C from PB01BASE (primary side window) and RW1S from SB01BASE (secondary side window).</p> <p>O_3_0_PB01BASE: Attr: RW1C Default: 0000h O_3_0_SB01BASE: Attr: RW1S Default: 0000h</p>

6.13.7.16 PDBMSK: Primary Doorbell Mask

This register is used to mask the generation of interrupts to the Primary side of the NTB.

PDBMSK			
Bus: 0		Device: 3	Function: 0
		Offset: 62	MMIO BAR: PB01BASE
Bus: 0		Device: 3	Function: 0
		Offset: 62	MMIO BAR: SB01BASE
Bit	Attr	Default	Description
15:0		FFFFh	<p>Primary Doorbell Mask</p> <p>This register will allow software to mask the generation of interrupts to the processor on the Primary side of the NTB.</p> <p>0: Allow the interrupt 1: Mask the interrupt</p> <p>Notes:</p> <p>This field is RW from PB01BASE (primary side window) and RO from SB01BASE (secondary side window).</p> <p>O_3_0_PB01BASE: Attr: RW Default: FFFFh O_3_0_SB01BASE: Attr: RO Default: FFFFh</p>



6.13.7.17 SDOORBELL: Secondary Doorbell

This register contains the bits used to generate interrupts to the processor on the Secondary side of the NTB.

SDOORBELL				
Bus: 0		Device: 3	Function: 0	MMIO BAR: PB01BASE
		Offset: 64		
Bus: 0		Device: 3	Function: 0	MMIO BAR: SB01BASE
		Offset: 64		
Bit	Attr	Default	Description	
15:0		0000h	Secondary Doorbell Interrupts These bits are written by the processor on the Primary side of the NTB to cause a doorbell interrupt to be generated to the processor on the Secondary side of the NTB if the associated mask bit in the SDBMSK register is not set. A 1 is written to this register from the Primary side of the NTB to set the bit, and to clear the bit a 1 is written from the Secondary side of the NTB. Notes: If both INTx and MSI (NTB PCI CMD bit 10 and NTB MSI Capability bit 0) interrupt mechanisms are disabled software must poll for status since no interrupts of either type are generated. This field is RW1S from PB01BASE (primary side window) and RW1C from SB01BASE (secondary side window). 0_3_0_PB01BASE: Attr: RW1S Default: 0000h 0_3_0_SB01BASE: Attr: RW1C Default: 0000h	


6.13.7.18 SDBMSK: Secondary Doorbell Mask

This register is used to mask the generation of interrupts to the Secondary side of the NTB.

SDBMSK				
Bus: 0		Device: 3	Function: 0	MMIO BAR: PB01BASE
		Offset: 66		
Bus: 0		Device: 3	Function: 0	MMIO BAR: SB01BASE
		Offset: 66		
Bit	Attr	Default	Description	
15:0	RW-V	0000h	Secondary Doorbell Mask This register will allow software to mask the generation of interrupts to the processor on the Secondary side of the NTB. 0: Allow the interrupt 1: Mask the interrupt Note: This field is RO from PB01BASE (primary side window) and RW from SB01BASE (secondary side window).	

6.13.7.19 USMEMMISS: Upstream Memory Miss

This register is used to keep a rolling count of misses to the memory windows on the upstream port on the secondary side of the NTB. This a rollover counter. This counter can be used as an aid in determining if there are any programming errors in mapping the memory windows in the NTB/NTB configuration.



USMEMMISS Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 70 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 70			
Bit	Attr	Default	Description
15:0	RW-V	0000h	Upstream Memory Miss This register keeps a running count of misses to any of the 3 upstream memory windows on the secondary side of the NTB. The counter does not freeze at max count it rolls over.

6.13.7.20 SPAD[0:15]: Scratchpad Registers 0 - 15


This set of 16 registers, SPAD0 through SPAD15, are shared to both sides of the NTB. They are used to pass information across the bridge.

SPAD[0:15] Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 80, 84, 88, 8C, 90, 94, 98, 9C Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: A0, A4, A8, AC, B0, B4, B8, BC Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 80, 84, 88, 8C, 90, 94, 98, 9C Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: A0, A4, A8, AC, B0, B4, B8, BC			
Bit	Attr	Default	Description
31:0	RW	00h	Scratchpad Register n This set of 16 registers is RW from both sides of the bridge. Synchronization is provided with a hardware semaphore (SPADSEMA4). Software will use these registers to pass a protocol, such as a heartbeat, from system to system across the NTB.

6.13.7.21 SPADSEMA4: Scratchpad Semaphore

This register will allow software to share the Scratchpad registers.

SPADSEMA4 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: C0 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: C0			
Bit	Attr	Default	Description
31:1	RO	00h	Reserved
0	RW-V	0h	Scratchpad Semaphore This bit will allow software to synchronize write ownership of the scratchpad register set. The processor will read the register: If the returned value is 0, the bit is set by hardware to 1 and the reading processor is granted ownership of the scratchpad registers. If the returned value is 1, then the processor on the opposite side of the NTB already owns the scratchpad registers and the reading processor is not allowed to modify the scratchpad registers. To relinquish ownership, the owning processor writes a 1 to this register to reset the value to 0. Ownership of the scratchpad registers is not set in hardware, i.e. the processor on each side of the NTB is still capable of writing the registers regardless of the state of this bit. The attribute of this register is ROTS (Read 0 to Set) and W1TC (Write 1 to Clear)



6.13.7.22 RSDBMSIXV70: Route Secondary Doorbell MSI-X Vector 7 to 0

This register is used to allow flexibility in the SDOORBELL bits 7 to 0 assignments to one of 4 MSI-X vectors. Register is set up to be able to expand to 16 MSI-X vectors in future designs.

Bit	Attr	Default	Description
RSDBMSIXV70 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: D0 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: D0			
29:28	RW	2h	MSI-X Vector assignment for SDOORBELL bit 7
25:24	RW	2h	MSI-X Vector assignment for SDOORBELL bit 6
21:20	RW	1h	MSI-X Vector assignment for SDOORBELL bit 5
17:16	RW	1h	MSI-X Vector assignment for SDOORBELL bit 4
13:12	RW	1h	MSI-X Vector assignment for SDOORBELL bit 3
9:8	RW	1h	MSI-X Vector assignment for SDOORBELL bit 2
5:4	RW	1h	MSI-X Vector assignment for SDOORBELL bit 1
1:0	RW	0h	MSI-X Vector assignment for SDOORBELL bit 0 11 = MSI-X vector allocation 3 10 = MSI-X vector allocation 2 01 = MSI-X vector allocation 1 00 = MSI-X vector allocation 0

6.13.7.23 RSDBMSIXV158: Route Secondary Doorbell MSI-X Vector 15 to 8

This register is used to allow flexibility in the SDOORBELL bits 15 to 8 assignments to one of 4 MSI-X vectors. Register is set up to be able to expand to 16 MSI-X vectors in future designs.

Bit	Attr	Default	Description
RSDBMSIXV158 Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: D4 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: D4			
29:28	RW	3h	MSI-X Vector assignment for SDOORBELL bit 15
25:24	RW	3h	MSI-X Vector assignment for SDOORBELL bit 14
21:20	RW	3h	MSI-X Vector assignment for SDOORBELL bit 13
17:16	RW	3h	MSI-X Vector assignment for SDOORBELL bit 12
13:12	RW	3h	MSI-X Vector assignment for SDOORBELL bit 11
9:8	RW	2h	MSI-X Vector assignment for SDOORBELL bit 10
5:4	RW	2h	MSI-X Vector assignment for SDOORBELL bit 9
1:0	RW	2h	MSI-X Vector assignment for SDOORBELL bit 8 11 = MSI-X vector allocation 3 10 = MSI-X vector allocation 2 01 = MSI-X vector allocation 1 00 = MSI-X vector allocation 0



6.13.7.24 B2BSPAD[0:15]: Back-to-back Scratchpad Registers 0

This set of 16 registers, B2BSPAD0 through B2BSPAD15, is used by the processor on the Primary side of the NTB to generate accesses to the Scratchpad registers on a second NTB whose Secondary side is connected to the Secondary side of this NTB. Writing to these registers will cause the NTB to generate a PCIe packet that is sent to the connected NTB's Scratchpad registers. This mechanism allows inter-system communication through the pair of NTBs. Note that the B2BBAR0XLAT register must be properly configured to point to BAR 0/1 on the opposite NTB for this mechanism to function properly. Note also that this mechanism doesn't require a semaphore because each NTB has a set of Scratchpad registers. The system passing information will always write to the registers on the opposite NTB, and read its own Scratchpad registers to get information from the opposite system.

B2BSPAD[0:15] Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 100, 104, 108, 10C, 110, 114, 118, 11C Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 120, 124, 128, 12C, 130, 134, 138, 13C Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 100, 104, 108, 10C, 110, 114, 118, 11C Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 120, 124, 128, 12C, 130, 134, 138, 13C			
Bit	Attr	Default	Description
31:0		00000000h	Back-to-back Scratchpad Register n This set of 16 registers is written only from the Primary side of the NTB. A write to any of these registers will cause the NTB to generate a PCIe packet which is sent across the link to the opposite NTB's corresponding Scratchpad register. 0_3_0_PB01BASE: Attr: RW Default: 00000000h 0_3_0_SB01BASE: Attr: RO Default: 00000000h

6.13.7.25 B2BDOORBELL: Back-to-back Doorbell

This register is used by the processor on the primary side of the NTB to generate accesses to the PDOORBELL register on a second NTB whose Secondary side is connected to the Secondary side of this NTB. Writing to this register will cause the NTB to generate a PCIe packet that is sent to the connected NTB's PDOORBELL register, causing an interrupt to be sent to the processor on the second system. This mechanism allows inter-system communication through the pair of NTBs. Note that the B2BBAR0XLAT register must be properly configured to point to BAR 0/1 on the opposite NTB for this mechanism to function properly.

B2BDOORBELL Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 140 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 140			
Bit	Attr	Default	Description
15:14	RV	0h	Reserved
13:0		0000h	B2B Doorbell Interrupt These bits are written by the processor on the Primary side of the NTB. Writing to this register will cause a PCIe packet with the same contents as the write to be sent to the PDOORBELL register on a second NTB connected back-to-back with this NTB, which in turn will cause a doorbell interrupt to be generated to the processor on the second NTB. Hardware on the originating NTB clears this register upon scheduling the PCIe packet. 0_3_0_PB01BASE: Attr: RW1S Default: 0000h 0_3_0_SB01BASE: Attr: RO Default: 0000h



6.13.7.26 B2BBAROXLAT: Back-to-back BAR 0/1 Translate


B2BBAROXLAT Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 144 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 144			
Bit	Attr	Default	Description
63:15		000000 000000 0h	B2B translate Base address of Secondary BAR 0/1 on the opposite NTB. This register is used to set the base address where the back-to-back doorbell and scratchpad packets will be sent. This register must match the base address loaded into the BAR 0/1 pair on the opposite NTB, whose Secondary side is linked to the Secondary side of this NTB. Notes: There is no hardware enforced limit for this register, care must be taken when setting this register to stay within the addressable range of the attached system. Primary side MSI-X MMIO registers reached via PB01BASE O_3_0_PB01BASE: Attr: RW Default: 0000000000000h O_3_0_SB01BASE: Attr: RO Default: 0000000000000h
14:0	RO	00h	Reserved Limit register has a granularity of 32 KB (215)

6.13.8 MSI-X MMIO Registers (NTB Primary side)

Primary side MSI-X MMIO registers reached via PB01BASE

Table 6-13. NTB MMIO Map

PMSIXTBLO	2000h	PMSIXPBA				3000h
	2004h					3004h
PMSIXDATA0	2008h					3008h
PMSICXECCNTLO	200Ch					300Ch
PMSIXTBL1	2010h					3010h
	2014h					3014h
PMSIXDATA1	2018h					3018h
PMSICXECCNTL1	201Ch					301Ch
PMSIXTBL2	2020h					3020h
	2024h					3024h
PMSIXDATA2	2028h					3028h
PMSICXECCNTL2	202Ch					302Ch
PMSIXTBL3	2030h					3030h
	2034h					3034h
PMSIXDATA3	2038h					3038h
PMSICXECCNTL3	203Ch					303Ch
	2040h					3040h
	2044h					3044h
	2048h					3048h
	204Ch					304Ch



6.13.8.1 PMSIXTBL[0:3]: Primary MSI-X Table Address Register 0 - 3

PMSIXTBL[0:3] Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 2000, 2010, 2020, 2030 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 2000, 2010, 2020, 2030			
Bit	Attr	Default	Description
63:32	RW	00000000h	MSI-X Upper Address Upper address bits used when generating an MSI.
31:2	RW	00000000h	MSI-X Address System-specified message lower address. For MSI-X messages, the contents of this field from an MSI-X Table entry specifies the lower portion of the DWORD-aligned address (AD[31:02]) for the memory write transaction.
1:0	RO	00b	MSG_ADD10 For proper DWORD alignment, these bits need to be 0's.

6.13.8.2 PMSIXDATA[0:3]: Primary MSI-X Message Data Register 0

PMSIXDATA[0:3] Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 2008, 2018, 2028, 2038 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 2008, 2018, 2028, 2038			
Bit	Attr	Default	Description
31:0	RW	0000h	Message Data System-specified message data.

Table 6-14. MSI-X Vector Handling and Processing by IIO on Primary Side

Number of Messages enabled by Software	Events	IV[7:0]
1	All	xxxxxxx ¹
4	PD[04:00]	xxxxxxx
	PD[09:05]	xxxxxxx
	PD[14:10]	xxxxxxx
	Hot Plug, BW-change, AER, PD[15]	xxxxxxx

Notes:

1. The term "xxxxxx" in the Interrupt vector denotes that software initializes them and IIO will not modify any of the "x" bits



6.13.8.3 PMSICXVECCNTL[0:3]: Primary MSI-X Vector Control Register 0 -3

PMSICXVECCNTL[0:3] Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 200C, 201C, 202C, 203C Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 200C, 201C, 202C, 203C			
Bit	Attr	Default	Description
0	RW	1b	MSI-X Mask When this bit is set, the NTB is prohibited from sending a message using this MSI-X Table entry. However, any other MSI-X Table entries programmed with the same vector will still be capable of sending an equivalent message unless they are also masked.

6.13.8.4 PMSICXPBA: Primary MSI-X Pending Bit Array

Secondary side MSI-X MMIO registers reached via PB01BASE and SB01BASE.

PMSICXPBA Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 3000 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 3000			
Bit	Attr	Default	Description
31:4	RV	0h	Reserved
3	RO-V	0b	MSI-X Table Entry 03 NTB has a Pending Message
2	RO-V	0b	MSI-X Table Entry 02 NTB has a Pending Message
1	RO-V	0b	MSI-X Table Entry 01 NTB has a Pending Message
0	RO-V	0b	MSI-X Table Entry 00 NTB has a Pending Message

6.13.9 MSI-X MMIO registers (NTB Secondary Side)

Secondary side MSI-X MMIO registers reached via PB01BASE and SB01BASE.

These registers are valid when in NTB/RP configuration.

Table 6-15. NTB MMIO Map

SMSIXTBL0	4000h	<table border="1"> <tr> <th colspan="4">SMSIXPBA</th> </tr> <tr><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td></tr> <tr><td></td><td></td><td></td><td></td></tr> </table>	SMSIXPBA																																																												5000h
SMSIXPBA																																																															
SMSIXDATA0	4004h		5004h																																																												
SMSIXVECCNTL0	4008h		5008h																																																												
SMSIXTBL1	400Ch		500Ch																																																												
SMSIXTBL1	4010h		5010h																																																												
SMSIXTBL1	4014h		5014h																																																												
SMSIXDATA1	4018h		5018h																																																												
SMSIXVECCNTL1	401Ch		501Ch																																																												
SMSIXTBL2	4020h		5020h																																																												
SMSIXTBL2	4024h		5024h																																																												
SMSIXDATA2	4028h		5028h																																																												
SMSIXVECCNTL2	402Ch		502Ch																																																												



Table 6-15. NTB MMIO Map

SMSIXTBL3				4030h					5030h
				4034h					5034h
SMSIXDATA3				4038h					5038h
SMSIXVECCNTL3				403Ch					503Ch
				4040h					5040h
				4044h					5044h
				4048h					5048h
				404Ch					504Ch

6.13.9.1 SMSIXTBL[0:3]: Secondary MSI-X Table Address Register 0 - 3

SMSIXTBL[0:3] Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 4000, 4010, 4020, 4030 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 4000, 4010, 4020, 4030			
Bit	Attr	Default	Description
63:32	RW	00000000h	MSI-X Upper Address Upper address bits used when generating an MSI-X.
31:2	RW	00000000h	MSI-X Address System-specified message lower address. For MSI-X messages, the contents of this field from an MSI-X Table entry specifies the lower portion of the DWORD-aligned address (AD[31:02]) for the memory write transaction.
1:0	RO	00b	MSG_ADD10 For proper DWORD alignment, these bits need to be 0's.

6.13.9.2 SMSIXDATA[0:3]: Secondary MSI-X Message Data Register 0 - 3

SDOORBELL bits to MSI-X mapping can be reprogrammed through [Section 6.13.7.22](#) and [Section 6.13.7.23](#).

SMSIXDATA[0:3] Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 4008, 4018, 4028, 4038 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 4008, 4018, 4028, 4038			
Bit	Attr	Default	Description
31:0	RW	0000h	Message Data System-specified message data.

6.13.9.3 SMSIXVECCNTL[0:3]: Secondary MSI-X Vector Control Register 0 - 3



SMSIXVECCNTL[0:3] Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 400C, 401C, 402C, 403C Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 400C, 401C, 402C, 403C			
Bit	Attr	Default	Description
0	RW	1b	MSI-X Mask When this bit is set, the NTB is prohibited from sending a message using this MSI-X Table entry. However, any other MSI-X Table entries programmed with the same vector will still be capable of sending an equivalent message unless they are also masked.

Table 6-16. MSI-X Vector Handling and Processing by IIO on Secondary Side

Number of Messages Enabled by Software	Events	IV[7:0]
1	All	xxxxxxx ¹
4	PD[04:00]	xxxxxxx
	PD[09:05]	xxxxxxx
	PD[14:10]	xxxxxxx
	PD[15]	xxxxxxx

Notes:

1. The term "xxxxxx" in the Interrupt vector denotes that software initializes them and IIO will not modify any of the "x" bits.

6.13.9.4 SMSICXPBA: Secondary MSI-X Pending Bit Array

SMSICXPBA Bus: 0 Device: 3 Function: 0 MMIO BAR: PB01BASE Offset: 5000 Bus: 0 Device: 3 Function: 0 MMIO BAR: SB01BASE Offset: 5000			
Bit	Attr	Default	Description
31:4	RV	0h	Reserved
3	RO-V	0b	MSI-X Table Entry 03 NTB has a Pending Message
2	RO-V	0b	MSI-X Table Entry 02 NTB has a Pending Message
1	RO-V	0b	MSI-X Table Entry 01 NTB has a Pending Message
0	RO-V	0b	MSI-X Table Entry 00 NTB has a Pending Message



