



# **Intel® Open Source HD Graphics and Intel Iris™ Graphics**

## **Programmer's Reference Manual**

For the 2014-2015 Intel Core™ Processors, Celeron™ Processors  
and Pentium™ Processors based on the "Broadwell" Platform

Volume 15: Workarounds

October 2015, Revision 1.1

## Creative Commons License

**You are free to Share** - to copy, distribute, display, and perform the work under the following conditions:

- **Attribution.** You must attribute the work in the manner specified by the author or licensor (but not in any way that suggests that they endorse you or your use of the work).
- **No Derivative Works.** You may not alter, transform, or build upon this work.

## Notices and Disclaimers

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

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

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

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

Implementations of the I2C bus/protocol may require licenses from various entities, including Philips Electronics N.V. and North American Philips Corporation.

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

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

**Copyright © 2015, Intel Corporation. All rights reserved.**



## Table of Contents

<b>Workarounds.....</b>	<b>1</b>
-------------------------	----------



## Workarounds

This table lists all BDW workarounds. Note that the functional area for each item is listed below, and you can search on this value or other content on this page using search (e.g. Ctrl-F).

Functional Area/Component		Workaround Name	Workaround Description
3D	SURFACE_STATE		The R32_FLOAT, R32G32_FLOAT and R8G8_UNORM surface formats are not handled correctly for the Alpha channel for cases where the sample is off-map or out of bounds. The correct behavior should be to force these Alpha values to 1.0, but actual result is for these formats is a blending of the default color with 1.0.
3D	3D Sampler Message Types		When a surface of type surfaceType_Null is accessed by resinfo, the MIPCount returned is undefined instead of 0.
3D	Message Header		offu/offv are calculated in normalized space and hence subject to small truncation error.
3D	3D Sampler Messages and Message Types		offu/offv are calculated in normalized space and hence subject to a small truncation error.
3D	Vertex Fetch Functions		If a 32 bit uscaled or sscaled format is used, then a float format needs to be used so VF will keep the data as is and the kernel needs to convert the format to 32 bit float.
3D		WaDisableAsyncFlipPerf Mode	Async flips hang on MI_WAIT_FOR_EVENT following it. WaDuplicateMiDisplayFlip does not seem to help. To work around this we have enabled MI_SYNC_FLIP always.

Functional Area/Component		Workaround Name	Workaround Description
3D	Patch Header DW0-7 [Project: IVB+]		<p>HW Bug: The Tessellation stage will incorrectly add domain points along patch edges under the following conditions, which may result in conformance failures and/or cracking artifacts:</p> <p>QUAD domain            INTEGER partitioning            All three Tess Factors in a given U or V direction (for example,, V direction: UEQ0, InsideV, UEQ1) are all exactly 1.0            All three Tess Factors in the other direction are &gt; 1.0 and all round up to the same integer value (e.g, U direction: VEQ0 = 3.1, InsideU = 3.7, VEQ1 = 3.4)            The suggested workaround (to be implemented as part of the post-amble to the HS shader in the HS kernel) is:</p> <pre> if (   (TF[UEQ0] &gt; 1.0)      (TF[VEQ0] &gt; 1.0)      (TF[UEQ1] &gt; 1.0)      (TF[VEQ1] &gt; 1.0)      (TF[INSIDE_U] &gt; 1.0)      (TF[INSIDE_V] &gt; 1.0) ) {   TF[INSIDE_U] = (TF[INSIDE_U] == 1.0) ? 2.0 : TF[INSIDE_U];   TF[INSIDE_V] = (TF[INSIDE_V] == 1.0) ? 2.0 : TF[INSIDE_V]; }           </pre>
3D		WaPreventHSTessLevelsI nterference	<p>Inner tessellation levels can interfere with outer tess levels. The proposed SW workaround is to have the HS compiler, upon seeing INTEGER and QUAD, generate instructions to perform the following logic after or when moving the TFs into the Patch Header:</p> <pre> if (   (TF[UEQ0] &gt; 1.0)      (TF[VEQ0] &gt; 1.0)      (TF[UEQ1] &gt; 1.0)      (TF[VEQ1] &gt; 1.0)      (TF[INSIDE_U &gt; INSIDE_U] &gt; 1.0)      (TF[INSIDE_V &gt; INSIDE_V] &gt; 1.0) ) {   TF[INSIDE_U INSIDE_U] = (TF[INSIDE_U INSIDE_U] == 1.0)     ? 2.0: TF[INSIDE_U INSIDE_U];   TF[INSIDE_V INSIDE_V] = (TF[INSIDE_V INSIDE_V] == 1.0)     ? 2.0: TF[INSIDE_V INSIDE_V]; }           </pre> <p>Note that the non-inside TFs are never modified, so this will not impact the HW patch cull testing for 0= or NaN. (This serves as a warning to watch out for NaN comparisons).</p>

Functional Area/Component		Workaround Name	Workaround Description	
3D	GPGPU Indirect Thread Dispatch		The indirect registers are not supposed to be set to 0, but sometimes the kernel computing the value wants no work done and sets them to 0. This does not work correctly, so a workaround in the command stream is needed. The following workaround sets the predicate if IndirectArgs.X=0 or IndirectArgs.Y=0 or IndirectArgs.Z=0	
			<b>Command</b>	<b>Operation</b>
			MI_LOAD_REGISTER_IMM	MI_PREDICATE_RESULT = 0
			MI_LOAD_REGISTER_IMM	MI_PREDICATE_SRC1 = 0
			MI_LOAD_REGISTER_MEM	MI_PREDICATE_SRC0 = [ IndirectArgs ].X
			MI_PREDICATE	LOAD, OR, SOURCES_EQUAL
			MI_LOAD_REGISTER_MEM	MI_PREDICATE_SRC0 = [ IndirectArgs ].Y
			MI_PREDICATE	LOAD, OR, SOURCES_EQUAL
			MI_LOAD_REGISTER_MEM	MI_PREDICATE_SRC0 = [ IndirectArgs ].Z
MI_PREDICATE	LOADINV, OR, SOURCES_EQUAL			
WALKER	Predicate Enable = 1			
3D	TypedUntyped Surface ReadWrite and TypedUntyped Atomic Operation		The Typed Surface Read message returns 0 in all channels for out-of-bounds accesses.	
3D	Shared Local Memory (SLM)		SLM accesses are incorrectly treated as out-of-bounds when the General State Buffer Size in STATE_BASE_ADDRESS is set to zero. The workaround is to set the General State Buffer Size to a non-zero value.	
3D	Media State and Primitive Commands		The MEDIA_STATE_FLUSH command is updated to optionally specify all the resources required for the next thread group via an interface descriptor – if the resources are not available the group cannot start.	
3D	Shared Local Memory (SLM)		When SLM memory is disabled in the L3 configuration, SLM accesses are correctly treated as out-of-bounds but the out-of-bounds data returned may not be zero.	
3D	Byte Scattered ReadWrite		The stateless model is not supported. A workaround is to use the surface model based at General State Base Address, with the additional restriction that surface byte offsets are limited to 2 <sup>31</sup> .	
3D		WaConstantStateStall	No description has been provided yet for this item.	

Functional Area/Component		Workaround Name	Workaround Description
3D	Data Port Messages (BDW)		For Atomic Counter OPS other than INC, DEC, and PREDEC, the header is forbidden and not optional as indicated in the table.
3D		WaDividePSInvocationCountBy4	Invocation counter is 4 times actual. WA: SW to divide HW reported PS Invocations value by 4.
3D		WaVSRefCountFullforceMissDisable	Hang; For GS\VS workloads. WA: Set FF_MODE - Thread Mode Register, 0x020A0: DS Reference Count Full Force Miss Enable, bit 19 = 0, VS Reference Count Full Force Miss Enable, bit 15 = 0.
3D		WaHiZAmbiguateRequiredNonAlignedBeforeRender	Corruption. WA: Must do HiZ resolve (ambiguate) on any portion of the HiZ buffer that is not 8x4 aligned before rendering or clearing.
3D	Atomic Counter Operation Message Descriptor		For Atomic Counter OPS other than INC, DEC, or PREDEC, the message header is forbidden and not optional.
3D		WaIndirectDispatchPredicate	GFXDRV: Hang in GT2 WW03d Emulation - GPGPU _WALKER HANG. WA: Use MI_PREDICATE commands to predicate an indirect dispatch whose thread groups in any dimension = 0. GPGPU pipe3d: vfe handle counter underflow for dimx eq 0 followed by curbe workload. Alternatively - can use MI MATH to work around this like we have in the past.
3D	Universal Input Message Phases [BDW]		WA: The following text needs to be maintained so that we can bring back the feature in the next opportunity. Will be used for Field 16x8 Enabled: This field enables 16x8 interlaced-block partitioning for MPEG-2. Note: Enabling Field 16x8 prevents use of sub-partitions types 8x16, 4x4, 4x8 and 8x4, RefAccess and SrcAccess must be 0 and SrcSize must be 16x16 (00). Field8x8 and Field16x8 are mutually exclusive.
3D	Notification Registers		The sub-register numbers for n0.0 and n0.2 are swapped on a write, that is, a destination of n0.0 is required to update n0.2 and n0.2 is required to update n0.0.

Functional Area/Component		Workaround Name	Workaround Description
3D		AccWrEnNotAllowedToAcc1With16bit	Implicit write (AccWrEn) to acc1 (for example, H1, Q3) not allowed with 16-bit data type (for example, hf, w). Acc1 dependency is not set as EUTC doesn't detect it as write to 2nd half accumulator for half float Q3 and takes it as acc0 dest. It is using nib_ctrl[1] to determine which part of acc is being written. this works for float dest but for hf destination nib_ctrl[2] should be used.
3D		WaEnsureMemCoherencyBeforeLoadRegisterMem	MI_STORE_REG_MEM followed by MI_LOAD_REGISTER_MEM does not always stall before the LOAD. Requires that a sync event (like MI_ATOMIC or PIPE_CONTROL with CS stall) be issued between the events to ensure that the memory has been forced to be coherent. See B-Spec for restriction.
3D		WaFastClearZ16K	The definition of the HiZ based fast clear does not specify the ability to clear all the way in X and Y to 16K. The last HiZ in the X and Y would be only partially initialized. Workaround suggested is to clear as four (2x2) sets of 4k x 4k clears. The rectangle size in the HZ_OP packet can only be programmed up to 16k-1 pixels in x and y direction. We need one more bit in each direction to be able to clear the whole thing. Driver already has implemented a software workaround for this in Gen8, so this is not a show-stopper.
3D		WaFlushOpOnCSStall	When a pipecontrol with stall bit is set, an ISP bit is enabled and there is no post sync operation: CS is not doing a flush or dummy atomic while fulsim does and there is a mismatch. Restriction: if stall bit is set, any one of the post sync-operation should be set.
3D		WaRestoreFCandMSGRegistersFromUpperOword	PipeGT Page Faults: Incorrect gpgpu_phase in EUGA causes MEU save cycles to be saved incorrectly.
3D		WaOGLGSVertexReorderingTriStripAdjOnly	GS vtxram read pointer increments incorrectly if control topology value matches tristrip_adj* decoding value...value (6'b001100 , 6'b010101 ). This causes full to not be generated properly and eventually corruption of the GS internal storage structure..
3D		WaScoreboardStallBeforeStateCacheInvalidate	GfxSV PGT: hang on 3DSTATE_CHROMA_KEY instruction; CS, DAPR, PSD, RCZ, I Z, WMFE, SVL. WA: Before PIPE_CONTROL with State Cache Invalidation Enable bit set, PIPE_CONTROL with Stall At Pixel Scoreboard bit set, must be issued.
3D		WADisableWriteCommitForPageFault	Workaround to disable the commit bit when we have Page Faults in Memfence.

Functional Area/Component		Workaround Name	Workaround Description
3D		WaDSRefCountFullforceMissDisable	A BDW bug that needs DS_REFERENCE_COUNT_FULL_FORCE_MISS_ENABLE bit to be disabled in all BDW steppings. Set 0x20a0 masked bit 19.
3D		WaGrfDepClearOnOutstandingSamplerInGpgpuContextSave	<p>Feature: GPGPU Pre-emption Non-page faulting. Outstanding sampler messages to EU, pre-emption occurs before data return. Hang condition.</p> <p>WA: Software workaround in SIP:</p> <ol style="list-style-type: none"> <li>After jumping to save routine, check "fault status" bit in SR0.</li> <li>If it is set, that means there are faulted GRF registers and EU is in page fault mode, so we do not need to wait for them before saving. Jump to normal save routine.</li> <li>If it is not set, that means there are no faulted GRF registers, but there may still be outstanding sampler registers to wait for.</li> <li>Use mov instructions for each GRF register, with null as the destination, to check dependency on each register, to make sure they have returned.</li> <li>Continue with normal save routine.</li> </ol>
3D		WaGuardbandSize	<p>AUTO CLONE from ivbgt2: cref precision loss in Z interpolator (src_cref mismatch causing HIZ memdiff).</p> <p>WA: Limit guardband to -16k to +16k.</p>
3D		WaDisableRSBeforeBTPoolDisable	<p>BT edits are dropped in the R S if the pool is disabled.</p> <p>WA: Disable RS before BT pool is disabled.</p>
3D		WaNoA32ByteScatteredStatelessMessages	<p>The stateless model is not supported.</p> <p>WA: Use the A64 Byte Scattered message, adding the General State Base Address into each address offset in the message. However, the A64 message will not perform A32 stateless bounds checking.</p>
3D		WaForcenon-zeroSBEOutputAttributeCount	<p>WA: Clamp SF number of output attributes to min value of 1 to avoid GS running out of handles.</p>

Functional Area/Component		Workaround Name	Workaround Description
3D		WaOCLCoherentLineFlush	In case of OCL buffers with back-to-back walkers with one using binding table index 255 (Coherent) and other using BTI 253 (Non-coherent) using same surface we need flush coherent lines in between these walkers. So the WA requirement is to set bit 21 of 0xb118 MMIO at boot by KMD. With this bit set pipeline flush treat as invalidating even coherent lines along with non coherent lines. KMD needs to set this bit because b118 is not context specific (global) so UMD cannot set it. This WA is not going to impact other UMD components.
3D		WaGrfDependencyClearInGpgpuContextRestore	SVM Page faults.  WA: Software workaround in SIP according to BUN: In Context Restore replicates higher OWORD register part to lower part for flow control registers and message registers as below: <pre> mov (4) r-temp.0:ud r113.4&lt;4;4,1&gt;:ud { NoMask } mov (8) msg0.0:ud r-temp:ud { NoMask }  mov (4) r-temp.0:ud r113.4&lt;4;4,1&gt;:ud { NoMask } mov (8) fc0.0:ud r-temp:ud { NoMask } </pre>
3D	GPGPU Context Switch Workarounds		As a workaround for pre-emption when using ring buffers: An MI_LOAD_REGISTER_IMM writing data 0x0000FFFF to address 0x2054 must be placed at the pre-emption jump address. An MI_LOAD_REGISTER_IMM writing data 0x00000040 (or whatever is needed for the CS IDLE count) to address 0x2054 must be placed after MI_SET_CONTEXT. If MI_SET_CONTEXT is at the pre-emption jump address then this workaround is not needed. This workaround is not needed in execlist mode.
3D		WaSubtract1FromMaxNoOfThreads	Subtract 1 from Max No of Threads per PSD (Per Slice Dispatch) because the PSD RTL incorrectly adds 1 to the max thread programmed. So when programming "MaximumNumberOfThreadsPerPSD" in 3DSTATE_PS command, always subtract 1 from the value to be programmed.
3D		WaUseNonPrivRegisterForMidObjectPreemption	Pre-G0 use GFX_MODE(229c) and G0+ use INSTPM(20c0) MMIO regs to disable Mid-Object pre-emption for draw call.
3D		WaLimitHsUrbEntries	Hang - Tessellation workloads.  WA: Limit the 3DSTATE_URB_HS to 184.

Functional Area/Component		Workaround Name	Workaround Description
3D		WaStateBindingTableOverfetch	HW over-fetches two cache lines of binding table indices. When using the resource streamer, SW needs to pad binding table pointer updates with an additional two cache lines.
3D		WaStoreAcc2to9InAlign16InGpgpuContextSave	Access to acc2 - acc9 arfs in Align16 to prevent corruption in pre-empted workload. During context save/restore, acc2-acc9 are explicit and need to use explicit decoding logic.
3D		WaDisableMidObjectPreemptionForGSLineStripAdj	This is a bug in GS. GS expected vertex count doesn't decode linestrip_adj_cont nor polygon_cont. Only linestrip_adj_cont applies to BDW since polygon is not pre-empted. WA: Disable mid-draw preemption when draw-call is a linestrip_adj and GS is enabled.
3D	Floating-Point Support		When converting from float to int, rounding mode RZ must be used.
3D	OWord Block ReadWrite and Unaligned OWord Block Read		Out-of-Bounds check is disabled.
3D		WaVfPostSyncWrite	Workaround for BDW to set post sync op of write for PIPE_CONTROL when only VF cache invalidate set
3D		WaCsStallBeforenonzeroInstanceCount	Issue CS Stall when going from 3DSTATE_HS zero instance count or HS disabled to 3DSTATE_HS non-zero instance count.
3D		WaCselUnsupported	csel instruction is not supported and must not be used.
3D	Byte Scattered ReadWrite		The stateless model is not supported. A workaround is to use the A64 Byte Scattered message, adding the General State Base Address into each address offset in the message. However, the A64 message will not perform A32 stateless bounds checking.
3D	TypedUntyped Surface ReadWrite and TypedUntyped Atomic Operation		Tile W surfaces must be of format R8_UINT and only support SIMD8. Furthermore, only the RED channel can be enabled.
3D	3D Sampler Message Types		Message sample+killpix cannot be used with Fault & Stream mode.
3D	Stateless Model		A64 Stateless accesses are incorrectly treated as out-of-bounds when the General State Buffer Size in STATE_BASE_ADDRESS is set to zero. The workaround is to set the General State Buffer Size to a non-zero value.

Functional Area/Component		Workaround Name	Workaround Description
3D	Message Mode: M0.2, bit 31:30		BYTE_MASK is not supported.
3D			PSD variable dispatch with per sample 4x and 8x - restriction from previous project has to be removed.  WA: Enable only SIMD8 dispatch mode when in per-sample mode.
3D		WsVfPostSyncWrite	Hang: If the push constant is committed using PIPE_CONTROL with only VF cache invalidation then CS commits the push constant but does not send dirty bits to shaders. Thus the shaders will not see the dirty bits and if there is new push constant command followed by a stalling flush, CS will commit the second push constant and wait for the previous push constant deref. This causes the flush to hang as CS waits for "push constant done" in flush completion. WA:"Post Sync Operation" must be enabled to "Write Immediate Data" or "Write PS Depth Count" or "Write Timestamp".
3D		WaForceMinMaxGSThreadCount	GS being stalled can cause the fftid to go over max threads causing undefined scratch space to be used.  WA: Limit the number of handles to the number of threads, with some GS performance loss. Set min/max threads to 8 for GS. Should be handled in USC/IGC.
3D		WaGrfScoreboardClearInGpgpuContextSave	Need to use stop_done pulse to clear grf scoreboard on save. Logic exists to restore grf scoreboard based on MDE data being restored to MEU.  WA: Software workaround in SIP: State register special handling against page fault issue; change is requested by EU team. In Context save sr0.1 register is stored in temporary register, temporary register is masked and sent to csr buffer, next sr0.1 is cleared. In context restore sr0.1 is restored as one of the last registers (just before r0 restore and exception clear).
3D			WA: A PIPE_CONTROL with CS_STALL must be sent whenever the HS_STATE.InstanceCount changes from 0 (no instancing to > 0 (instancing).
3D		WaAdditionalMovWhenSrc1ModOnMulMach	A source modifier must not be used on src1 for the mul/mach macro operations.  WA: Use extra move instead of src modifier for src1.

Functional Area/Component		Workaround Name	Workaround Description
3D		WaRestoreFC4RegisterDW0fromDW1	GfxSV - [MDT] BDWD0 - GPGPU Pre-emption - Execution Mask not being saved/restored correctly (memdiff). WA: SIP routine has to correct the address while restoring. Flow control register FC4 has to be restored from DW1.
3D		WaScalarAtomic	This is a performance improvement implemented as a W/A. Improves append counter updates from 1/6 clks (L3 limit) to 16/6 clks.
3D		WaRestoreFc0RegistersWithOffset	GfxSV - BDWF0 - GPGPU Pre-emption - EUSTress nested if memdiff. WA: Flow control register FC0 has to be restored with offsets. Restores fc0.4 to fc0.31 register with special offset in pre-emption context restore.
3D	Media State and Primitive Commands		Two MEDIA_STATE_FLUSH commands need to be used to ensure that the flush is complete.
3D	Extended Math Function		When both srcs are NAN, FDIV produces denominator NAN as output.
3D		WaThreadSwitchAfterCall	[MDT]GfxSV - BDW E0 - GPGPU Pre-emption - CALL Instruction Hang. WA: Follow every call by a dummy non-JEU and non-send instruction with a SWITCH for both cases whether a subroutine is taken or not.
3D	MEDIA_STATE_FLUSH		A MEDIA_STATE_FLUSH with no options must be added after a GPGPU_WALKER command which doesn't use either SLM or barriers.
3D		WaNearestFilterLODClamp	DX10.1 LOD clamping VS Max LOD DX case. Workarounds: DX: If ( mipfilter_nearest ) MaxLOD = floor(MaxLOD) MinLOD = floor(MinLOD) OGL: If ( mipfilter_nearest ) lodbias = lodbias - 0.000001b Dx9 - (Not Required for Dx9. Max always set to to 14.0) Mac - Should not be needed.

Functional Area/Component		Workaround Name	Workaround Description
3D	GPGPU Context Switch Workarounds		After either a MI_SET_CONTEXT or a PIPE_CONTROL with Generic Media State Clear, there must be a MEDIA_VFE_STATE command before any pre-emptable command. The parameter of this MEDIA_VFE_STATE command can be set to default values.
3D	GPGPU Indirect Thread Dispatch		CURBE should be used for the payload when using indirect dispatch rather than indirect payload
3D	Shared Local Memory (SLM)		SLM accesses are incorrectly treated as out-of-bounds when the General State Buffer Size in STATE_BASE_ADDRESS is set to zero. The workaround is to set the General State Buffer Size to a non-zero value. When SLM memory is disabled in the L3 configuration, SLM accesses are correctly treated as out-of-bounds but the out-of-bounds data returned may not be zero.
3D	3D Sampler Message Types		If Surface Format is R10G10B10_SNORM_A2_UNORM and Gather4 Source Channel Select is alpha channel, the returned value may be incorrect.
3D	Programming Media Pipeline - Command Sequence		A MEDIA_STATE_FLUSH needs to be placed right before the MI_BATCH_BUFFER_END of any batch buffer that uses MEDIA_OBJECT.
3D	Depth Buffer Clear		To fast depth clear a full 16k, in the X and/or Y dimensions, the clear operation must be sectioned into rectangles smaller in X and Y than 16k pixels. For example to clear a 16k x 16k surface perform four (2x2) 4k x 4k clears with the proper address offsets.
3D	LOD Information		The LOD is in-bounds if $LOD < MIPCount$ and if $MinLOD + LOD < 15$ . If LOD is not in-bounds then 0 is returned for the width, height, and depth values.
3D	3D Primitives Overview - QUADLIST and 3D Primitives Overview - QUADSTRIP		To work around IVB bug #3665983 the driver must detect the use of a QUADLIST input topology along with the use of primitive ID in the pixel shader, and correspondingly shift right by 1 the primitive ID in the pixel shader.
Blitter	BLT Programming Restrictions		There are two suggested software workarounds to perform coherent overlapping BLTs: (a) The Source and Destination Base Address registers must hold the same value (without alignment restriction) and (b) the Source and Destination Pitch registers (BR11, BR13) must both be a multiple of 64 bytes. Or if (a) is not possible do overlapping source copy BLTs as two blits, using a separate intermediate surface.

Functional Area/Component		Workaround Name	Workaround Description
Blitter		WaUse3dBlitForTiledYSrc Dst	Fallback to 3D blit on blits involving tiledY surfaces.
Blitter		WaUse3dBlitForFBCMMI OProgramming	In this hang, BCS is executing LRI to an external GT address (address 50380, data 00000002). Therefore, instead of sending the cycle through the message channel, it will go through GAB. At the same time, BCS receives a GO=0 message from PM and this will bring the GO flag down which in turn will block the memory cycle from going out. The deadlock here is that the same state machine (in CSCFG) drives the GO ACK response and the external GT message.
Display	Clocks		CDCLK frequency change sequence needs added steps to prevent clock glitches. See North Display Engine Registers, CDCLK Sequence for details.
Display	Clocks		Very intermittent chance of failure of the display spread/bent clock reference.  WA: Adjust the display spread/bent clock reference programming sequence when enabling/disabling HDMI, DVI, and DisplayPort. See ICLKIP Programming.
Display	Clocks		The PCH display clock stops when BIOS enables PCH ISCLK PLL shutdown feature, causing backlight flicker and other problems.  WA: When the ISCLK PLL shutdown feature is enabled, set a clock gate disable C2020h bit 12 before using Aux channel B/C/D, GMBUS, GTC, or panel power sequencing. See South Display Engine Registers for details.
Display	Display		Carry over from IVB. VGA may lockup and give a black screen randomly with some memory configs and/or memory tests.  WA: See VGA_CONTROL for workaround bits since they change with project and stepping.
Display	Display		HDMI and DVI with audio are not supported when HSYNC Start is programmed equal to HBLANK Start.
Display	Display		FDI failures after training causes blank screen or hang.  WA: Transcoder timing generator override bits must be set and cleared at certain point when enabling/disabling PCH transcoder. See Mode Set Sequences for exact workaround details.

Functional Area/Component		Workaround Name	Workaround Description
Display	DisplayPort		<p>There could be a rare case where the eDP modeset does not happen correctly if the link is already in "send normal pixels" when the pipe is enabled.</p> <p>WA: For eDP, do not set Normal Pixel at the end of Link Training. See DisplayPort Enable Sequence.</p>
Display	DisplayPort		<p>DP MST output incorrect for certain M and N and VC payload size values.</p> <p>WA: VC payload must be multiple of 4 in x1 lane config, 2 in x2, 1 in x4. See M/N Values.</p>
Display	DisplayPort		<p>DDIA Aux channel transactions get intermittent NAK errors with some receivers.</p> <p>WA: Increase DDI_AUX_CTL_A bits 27:26 Time out timer value to 600us 01b when doing DDIA aux transactions.</p>
Display	DisplayPort		<p>Big FIFO mode conflicts with enabling of DP Port Sync mode.</p> <p>WA: Set 0x45280 bits 2:1=11b before enabling port sync mode. See DisplayPort Enable Sequence.</p>
Display	DisplayPort		<p>Chance of Aux Channel command failures when using Lynxpoint:H.</p> <p>WA: For the PCH Aux channels (Aux B/C/D) use an aux divider value of 63 decimal (03Fh). If there is a failure, retry at least three times with 63, then retry at least three times with 72 decimal (048h). See South Display Engine Registers, DP_AUX_CTL.</p>
Display	FDI		<p>FDI Rx delay default value incorrect, causing FDI failures.</p> <p>WA: Set FDI delay to 90h before enabling FDI.</p>
Display	FDI		<p>FDI Rx TP1 time default value incorrect, causing FDI training failures.</p> <p>WA: Set TP1 to TP2 time to 48 clocks before enabling.</p>
Display	FDI		<p>Added programming needed to reset FDI mPHY IOSF-SB during mode set sequences.</p> <p>WA: See Mode Set Sequences for details.</p>

Functional Area/Component		Workaround Name	Workaround Description
Display	FDI		Added programming needed to control FDI Rx PwrDn bits during CRT mode set sequences. WA: See Mode Set Sequences for details.
Display	FDI		Added programming needed to setup FDI mPHY registers at boot and in mode set sequences. WA: See Mode Set Sequences.
Display	General		TRANS_CONF Transcoder State may incorrectly show it is enabled before the transcoder has been enabled for the first time after reset or power well disable and enabled. This should not impact the normal usage of this field during the mode set disable sequence.
Display	GTC		LPT:LP Hang when reading GTC_LIVE_SW_CAP with GTC clock disabled. WA: Software needs to prevent reads to GTC_LIVE_SW_CAP or enable the GTC clock (0xC6030 bit 0) when reading it. See South Display Engine Registers, GTC_LIVE_SW_CAP.
Display	GTC		GTC sends incorrect byte count for Aux transaction with receivers, making GTC unusable on LPT-H:A, LPT-H:B 0, LPT-LP:A0. WA: None possible.
Display	GTC		Hang when reading GTC_LIVE_SW_CAP with GTC clock disabled. WA: Software needs to prevent reads to GTC_LIVE_SW_CAP or enable the GTC clock (0xC6030 bit 0) when reading it. See South Display Engine Registers for details.
Display	IPS		It can take up to 42ms for pcode to complete IPS disabling. The driver sequence to disable IPS needs to account for this delay before disabling the last plane on the pipe, otherwise there will be corruption from disabling planes before IPS is completely disabled. See North Display Engine Registers, Intermediate Pixel Storage for details.

Functional Area/Component		Workaround Name	Workaround Description
Display	IPS		<p>Chance of screen corruption if IPS and display planes are not enabled and disabled in a specific sequence.</p> <p>WA: IPS cannot be enabled until at least one plane has been enabled for at least one vertical blank. IPS must be disabled while there is still at least one plane enabled. See North Display Engine Registers, Intermediate Pixel Storage.</p>
Display	Panel fitter	WaPanelFitterDownscale	<p>Not a bug, but good to know. When using panel fitter downscaling (pipe source size is larger than panel fitter window size) the maximum supported pixel rate will be reduced by the downscale amount, and watermarks must be adjusted. Use panel fitter scale amount when calculating maximum pixel rate and watermarks.</p>
Display	Panel power sequencing	WaVDDOverrideT4Power	<p>When software clears the panel power sequencing VDD override bit from 1 to 0 (disable VDD override) it must ensure that T4 power cycle delay is met before setting the bit to 1 again, else panel may be damaged.</p> <p>WA: Use software timers to ensure T4 delay is met or use full panel power enable and not the VDD override.</p>
Display	PSR		<p>PSR single frame update with sprite - Mask the sprite enabled when using single frame update with sprite. See North Display Engine Registers, SRD_CTL register for details.</p>
Display	PSR		<p>PSR single frame update - When single frame update is enabled, the PSR CRC must be disabled for panel compatibility. See North Display Engine Registers SRD_CTL register for details.</p>
Display	PSR		<p>PSR single frame update - Mask register write events when using single frame update. See North Display Engine Registers SRD_CTL register for details.</p>
Display	PSR		<p>PSR power saving - Mask PSR max timeout when PSR CRC is enabled. See North Display Engine Registers, SRD_CTL register for details.</p>
KMD, Media		WaAVCSWHeaderInsertion	<p>Specific to AVC Encode on IVB and HSW. This WA is to make sure that H/W calculates the size of VCL NAL Unit correctly and adds Cabac Zero bytes if required to make the bit-stream AVC spec compliant.</p>
KMD, Media		WaVeboxSliceEnable	<p>Workaround required for checking Single Slice Enable flag in VEBOX State.</p>

Functional Area/Component		Workaround Name	Workaround Description
KMD, Media	Decode	WaJPEGHeightAlignYUV422H2YToNV12	Specific to JPEG decode. If the input is YUV422H_2Y and output is NV12 format, output picture height should be aligned by 16 bytes.
KMD, Media	Decode	WaInvalidSliceStartPosition	Invalidate the slice if the slice start position is out of the picture. Applies to all standards in VLD mode. If the slice position is invalid, skip it -- for example, the vertical or horizontal offset is out of bounds.
KMD, Media	Encode	WaAddMediaStateFlushCmd	Add Media State Flush command after Media Object Walker command.
KMD, Media	Encode	WaEnableVMEReferenceWindowCheck	WA for limiting the minimum downscaled surface dimensions to 48x48.
KMD, Media	Media (Encode/VP)	WaHwWalkerCmdExtraDword	Remove inline data dword 17 from MEDIA_OBJECT_WALKER_CMD_G6 command and send an extra DWORD at the end of this command if this WA is enabled.
KMD, Media	Media (Decode)	WaInsertAVCFrameForFormatSwitchToJPEG	Insert a dummy AVC frame before a JPEG frame to WA format switch issue when Decode format is switched to JPEG.
Memory Views	All Surfaces other than Separate Stencil Buffer.		Sampler MSAA Qpitch will be 4 greater than the value calculated in the equation above, for every other odd Surface Height starting from 1. That is: 1, 5, 9, 13, ....
Memory Views	Fault and STrem and C Context Switch on Demand.		Whole SIMD16 (all channels) being replayed even if non-faulted after context resubmission. If a given message from a thread has split accesses across different pages where one or more of these pages are present and at least one page is not present (i.e. one memory access page faults), the replay of this message does not take into account successful completions. The workaround causes h/w to replay some memory accesses twice, one before the page fault detected and one after replay. Fixed on G-step.
Memory Views	Graphics Memory Address Type: GSM		Due to a workaround, first 4KB of DSM has to be reserved for GFX hardware use during render engine execution.
Memory Views	Memory Interface		32b PPGTT has no implications
Memory Views	Opaque Textures (DXT1_RGB)		The behavior of this format is not compliant with the OGL spec. As a workaround, the Surface Format should be set to BC1 and the Shader Channel Select A should be set to SCS_ONE. This workaround is only available for BDW due to lack of Shader Channel Select support on earlier products.

Functional Area/Component		Workaround Name	Workaround Description
Observability	GPGPU Context Switch Workarounds		In products with slices that have an unequal number of EUs (for example, 1 slice with 23 EUs and 2 slices with 24) there will sometimes be a delay in pre-emption while the hardware waits for threads to exit in the smaller slice.
Display	Power		WA:Disable IPS when dotclock > 95% of cdclk
Display	eDP	N/A	eDP link training failures if package C6 is entered when display on/off power well is disabled. The PCU thinks eDP is completely disabled and the display LCPLL can be shutoff, but eDP is actually partially enabled and needs the LCPLL.  WA:Program SRD_CTL_EDP Max Sleep Time to 0 before starting eDP link training. This will tell the PCU to not turn off the LCPLL. Restore SRD_CTL_EDP Max Sleep Time after enabling TRANS_CONF_EDP.
Display	Dither	N/A	Dither pattern stops as soon as the transcoder configuration register is written to disable transcoder. Dither should continue until the vblank when transcoder actually disables.  WA: A workaround is only needed if it is important to have a correct dither pattern in the final frame as display output is disabled. Any workaround must be specific to the implementation of the operating system and driver.

Functional Area/Component		Workaround Name	Workaround Description
KMD		WaRsRestoreWithPerCtxt Bb	<p>To work around a known HW issue, SW must do the below Programming Sequence prior to programming MI_BATCH_BUFFER_END command in BB_PER_CTX_PTR. SW must ensure both MI_LOAD_REGISTER_REG and MI_BATCH_BUFFER_END commands mentioned in the below sequence are placed in the same cacheline of memory.</p> <ol style="list-style-type: none"> <li>1. MI_LOAD_REGISTER_IMM: 0x00800000 -&gt; 0x20C0</li> <li>2. MI_ATOMIC               <ol style="list-style-type: none"> <li>a. Set "CS STALL" (Dword0[17])</li> <li>b. "Return Data Control" enabled (Dword0[16])</li> <li>c. "ATOMIC OPCODE" set to LOAD operation (Dword0[15:8]= 0x4)</li> <li>d. "Memory Address" set to scratch space in GFX memory.</li> <li>e. "Operand1 Data Dword 0" must be programmed to 0x0080_0080</li> </ol> </li> <li>3. MI_LOAD_REGISTER_MEM               <ol style="list-style-type: none"> <li>a. Set "Async Mode Enable" (Dword0[ 21])</li> <li>b. "Memory Address" set to same as in MI_ATOMIC command above.</li> <li>c. "Register Address" set to 0x20C0</li> </ol> </li> <li>4. MI_LOAD_REGISTER_REG: 0x215C -&gt; 0x215C</li> <li>5. MI_BATCH_BUFFER_END // Note that there shouldn't be any commands programmed between step4 &amp; step5 and also these commands must be placed in the same cacheline of memory.</li> </ol>
KMD		WaIdleLiteRestore	<p>SW must always ensure ring buffer head pointer is not equal to tail pointer of a context, whenever it is submitted to HW for execution.</p> <p>WA: Driver should not submit a context with head = tail.</p>
KMD		WaIdleLiteRestore	<p>SW must always ensure "Force Restore Bit" in the context descriptor is set for a preempted context that is getting resubmitted.</p> <p>WA: Driver should not submit a context with head = tail.</p>

Functional Area/Component		Workaround Name	Workaround Description
KMD	Mid-batch Preemption	WaDisableCtxRestoreArbitration	<p>[Render CS Only][Execlist Mode of Scheduling]: SW must ensure arbitration is switched off while context restore is in progress for any given context. This is achieved by disabling arbitration by programming MI_ARB_ON_OFF to "Arbitration Disable" in RCS_INDIRECT_CTX buffer and by enabling back the arbitration by programming MI_ARB_ON_OFF to "Arbitration Enable" as the last command prior to MI_BATCH_END in the BB_PER_CTX_PTR buffer of every context submitted. Note that RCS_INDIRECT_CTX_OFFSET could be set to default value or any other legitimate value as per the programming notes of the register definition.</p> <ul style="list-style-type: none"> <li>Arbitration disable by programming MI_ARB_ON_OFF (Arbitration Disabled) in RCS_INDIRECT_CTX buffer.</li> <li>Arbitration enabled by programming MI_ARB_ON_OFF (Arbitration Enabled) as the last command prior to MI_BATCH_BUFFER_END in BB_PER_CTX_PTR buffer.</li> </ul>
Blitter		WaUse3DBlitForNonCLAlignedLinearSrcDst	Color source and destinations must be restricted to CL alignment for linear addressing.
3D		WaHeadTailNotEqualForExeclist	<p>CS corrupts ctx image when Lite restore with head=tail and per ctx BB/wabb with MI_ARB_OFF/MI_ARB_ON is present</p> <p>W/A: S/W need to make sure for preempted ctx resubmission lite restore should not be attempted. This W/A if applicable till G0 stepping as HSD 1912487 is fixed in H0. s/W Confirmation mail added.</p>
3D		WaGT3ResendDepthStencilAndHierDepthBufferState	<p>Issue is BSPEC specifies X Max and Y Max are "exclusive" when they need to be "inclusive".</p> <p>Here is the BSPEC note for Clear Rectangle X Max in 3DSTATE_WM_HZ_OP:</p> <p>Specifies Xmax value of (exclusive) of clear rectangle with the Depth Buffer, used for clipping. Pixels with X coordinates greater than or equal to Xmax will be not be affected.'</p>

Functional Area/Component		Workaround Name	Workaround Description
Power	FBC		<p>FBC may intermittently fail to update some lines after an image change, causing stale data to appear. The stale data can be prevented by manually invalidating FBC after flips and enabling FBC.</p> <p>Manual invalidation sequence: For (i=0 to 127) { Write 0x50344 = (i«16) }</p> <p>The manual invalidation needs to happen in the frame that the image changes or FBC enables, between the start of vertical blank at the beginning of that frame and the start of vertical blank at the end of that frame.</p> <p>Recommended sequence: After an async flip, sync flip, or enabling FBC, wait for the next start of vblank (sync flips can alternatively wait for flip done), then perform the manual invalidation sequence.</p>
KMD		WaSkipStolenMemoryFirstPage	WA to skip the first page of stolen memory due to sporadic HW write on *CS Idle
KMD		WaForceContextSaveRestoreNonCoherent	To avoid a potential hang condition with TLB invalidation driver should enable masked bit 5 of MMIO 0x7300 at boot.
GMM		WaIommuCCInvalidationHang	When using IOMMU, IOTLB invalidation followed by a Context Cache invalidation can cause a hang. Workaround is to add a NOP (Device TLB invalidation is considered a NOP by HW) every time before doing a context cache invalidation. This is done by writing dev_iotlb_inv_dsc to invalidation queue OR writing to scratch register for register based invalidation. (Scratch register[0x4F104] used for invalidation queue workaround(IommuUncoreUnavailable)) - Refer to PRM for descriptor format - Search "dev_iotlb_inv_dsc".
GMM		Wa32bitGeneralStateOffset	Allocations of type Scratchflat/Instructionflat can only be referenced by 32 bit offset in a Graphics Segment. When Graphics memory size is >2GB(BDW+, win8) 32 bit offsets cannot address the 64 bit aperture segment. This workaround is used to define 32bit Segment to be used only for ScratchFlat/InstructionFlat heaps in >2GB configurations.
GMM		Wa32bitInstructionBaseOffset	Allocations of type Scratchflat/Instructionflat can only be referenced by 32 bit offset in a Graphics Segment. When Graphics memory size is >2GB(BDW+, win8) 32 bit offsets cannot address the 64 bit aperture segment. This workaround is used to define 32bit Segment to be used only for ScratchFlat/InstructionFlat heaps in >2GB configurations.
GMM		WaL3ParitySupportDisable	Disables L3 DPF on BDW. L3 DPF only applies to Shared Local Memory(SLM) use cases and so this flag only affects these use cases. By default L3 DPF is turned off. For other steppings unmask the bit(5) of the register IMR[0x020a8h]

Functional Area/Component		Workaround Name	Workaround Description
KMD		WaForceEnableNonCoherent	Must Force Non-Coherent whenever executing a 3D context. This is a workaround for a possible hang in the unlikely event a TLB invalidation occurs during a PSD flush. Set masked bit 4 in 0x7300 during boot.
KMD		WaDisableSRGBGammaToLinearConv	WA is a carry-over from HSW. WA is to avoid Gamma to Linear conversion when source surface is SRGB and dest is non SRGB as per DX10's Present DDI expectation. In such a case, treat source as non-SRGB that will do bit-for-bit transfer from source to dest. E.g. if source format is B8G8R8A8_UNORM_SRGB and destination format is non-srgb then treat source as B8G8R8A8_UNORM instead by programming the surface state accordingly.
KMD		WaClearRenderResponseMasks	WA is a carry-over from HSW (different programming though for IVB). Need to set the correct render response masks to save power. HW power-on default is not good. Applies only to unmask primary plane flip and sprite flip done on A, B and C pipes. Done by clearing bit 1,2, 9, 10, 15 and 20 of MMIO 0x44050. Must be done at boot and all save/restore paths.
KMD		Wa4x4STCOptimizationDisable	HIZ/STC hang in hawx frames. W/A: Disable 4x4 RCPFE-STC optimization and therefore only send one valid 4x4 to STC on 4x4 interface. This will require setting bit 6 of reg. 0x7004. Must be done at boot and all save/restore paths.
KMD		WaDisableDopClockGating	Euunit - dop_gate_paths (half cycle path) W/A: Disable DOP clock gating for render. This requires setting masked bit 0 of MMIO 0xE4F4 and bit 16 of MMIO 0x9400. Setting bit 0 of 0xE4F4 will disable EU clock gating and setting bit 16 of 0x9400 will disable EUTC clock gating. To be done at boot time.

Functional Area/Component		Workaround Name	Workaround Description
KMD		WaModifyVFESateAfterGPGPUPreemption	<p>GPGPU preemption hang or corruption issue. This WA must be applied before re-submitting a GPGPU preempted workload. SW needs to do two things.</p> <ol style="list-style-type: none"> <li>SW should detect if VFE unit in the context image have LRI commands to 0x54AC and 0x54B0. SW does that by checking if the LRI header at context image offste (0x0E90) &lt;in VFE&gt; is 0x11001003 and the two MMIO offsets in this LRI is 0x54AC and 0x54B0. If so, do this-  D0 (LRI header): 0x1100_1003 -&gt; 0x1100_ 1001  D1 (1st MMIO offset): 0x0000_54AC -&gt; Leave this dword untouched  D2 (1st MMIO offset value): 54AC_DATA -&gt; Leave this dword untouched  D3 (2nd MMIO offset): 0x0000_54B0 -&gt; Overwrite to 0x0  D4 (2nd MMIO offset value): 54B0_DATA -&gt; Overwrite to 0x0</li> <li>Apart from it SW must also increment the Batch Buffer Address by Walker Command Length so that CS moves to the command following it. Size of Walker should be 0xF * sizeof(DWORD) = 0x3C Bytes</li> </ol>
KMD	RTL	WaDisableSamplerPowerBypass	Disable sampler power bypass to prevent MsaaBasic.level16 hang on mi_set_context. This requires setting masked bit 1 of MMIO E184 at boot time.
KMD	RTL	WaBlockMsgChannelDuringGfxReset	WA is a carry-over from HSW. When GPU hangs occurs, inside gfx driver TDR routine the driver writes to the reset register and after that it tries to read back the reset register and GPU never response back so system crashes. As a WA set 0x9424 bits 2,3,4,5,6,9 on boot.
KMD	RTL	WaFlushCoherentL3CacheLinesAtContextSwitch	Coherent L3 cache lines are not getting flushed during context switch which is causing issues like corruption. Need to set bit 21 of MMIO b118, then send PC with DC flush and then reset bit 21 of b118. This programming sequence needs to be part of the indirect context WA BB
KMD		WaDisableSDEUnitClockGating	WA for GPGPU workload hang for which requirement is to disable SDE Unit clock gating. This is done by setting bit 14 of MMIO 9430.
KMD		WaProgramMgsrForCorrectSliceSpecificMmioReads	Slice specific MMIO reads inaccurate so shadow Reg 119 (MGSR) needs to be programmed appropriately to get the correct reads from these slice-related MMIOs. This should be done before any other WAs are applied. Appropriate programming of 119 requires setting its sliceid and subsliceid values as per the slice/sub-slice active on the HW.

Functional Area/Component		Workaround Name	Workaround Description
KMD		WaIncreaseTagClockTimer	Clock gating causes wrong read address from State Array during RO invalidation causing Mem-Diff. So WA requires programming L3 tag clock timing register. At boot, write 0xB10Ch, bits [23:20] programmed to 1000b.
KMD		WaSetVfGuardbandPreemptionVertexCount	Workaround for potential 3d preemption bug that can cause data corruption. Driver should write to register 0x83A4 (preemption vertex count) and set a value of 0x20. At boot, write 0xffff0020 to 0x83a4 (it's a masked register).
KMD		WaPipeControlUpperDwordCorruption	WA to a potential HW issue where the upper DWORD of a PIPECONTROL dword write is corrupted. So split upper dword to a separate pipecontrol.
KMD		WaProgramL3SqcReg1Default	Program the default initial value of the L3SqcReg1 on BDW for performance. This requires writing 0x784000 to MMIO b100.
KMD		WaSkipInvalidSubmitsFromOS	For Invalid submits from OS - simply report fence completion without submitting the DMA buffer to GPU.
KMD		WaTempDisableDOPClkGating	This WA is associated with another WA WaProgramL3SqcReg1Default (which requires doing a cpu write to 0xb100 offset with value 0x784000). Before applying WaProgramL3SqcReg1Default WA, SW should temporarily force disable render GOP clock gating by setting MMIO 0x9424 [bit 0] to 0. After WaProgramL3SqcReg1Default has been applied, S/W should restore 0x9424 with its old value.
3D		WaSendDummyVFEafterPipelineSelect	TSG unit writes null entries into context which can hang restore CS and TSG. WA: Send dummy VFE immediately after GPGPU pipeline select.  We don't use Generic MediaStateClear at all.
3D			Pipeline flush will end up in hang in GPGPU mode when FFDOP is enable. WA: Always use stalling flush in GPGPU.
3D			For indirect dispatch there is handle counter underflow and it leads to hang. WA: GPGPU walker need to use only curbe while doing indirect dispatch.

Functional Area/Component		Workaround Name	Workaround Description
3D			<p>Draw call doesn't send vertices to compute shader because 0-vertex draw call, preempted. Handful of reasons that VS would send no vertices.</p> <p>WA: Pipe control with CS stall only when the UAV is using a vertex shader. (Same as HSW).</p>
3D		WaClearCCStatePriorPipe lineSelect	<p>Architecture hole; on GPGPU context restore, at the end of the context when CS sends a null prim, SVG and SARB does a state prefetch; by the time the data returns from memory, CS gates the FF clock.</p> <p>WA: In GPGPU mode, color cal state should not have valid bits. Before switching pipelines, send null CC state pointers.</p>
3D			<p>FF clocks will be gated, EU kernels are requesting from render cache to have writes from memory.</p> <p>WA: Disable FF CLK gating when GPGPU Kernel require RC access.</p>
3D		WaAvoidPMAStall	<p>PMA bit is getting turned on when color traffic is in flight in RCPFE.</p> <p>WA: Put color flush before changing the PMA bit before stencil write cases.</p>
3D			<p>In BDW, increase the size of sub register from 8 to 16. For MOVI, did not change the code to accommodate. As a result, SW is not able to use this.</p> <p>WA: Use 2 SIMD16 instructions instead of 1 SIMD32. Will have slight performance hit.</p>
3D			<p>Write collision hang condition. When we have both remainder and quotient in same instruction, had write port collision on GRF, drop the dependency clear resulting in hang.</p> <p>WA: Break integer divide into separate quotient and remainder instructions.</p>
3D			<p>When DCS re-allocates handle ID, SDE flops but valids are not used in clock wakeup equations. When reallocation happens, handles lost.</p> <p>WA: Disable SDEunit clock gating. UCGCTL6[14]=1 (09430h[14])</p>

Functional Area/Component		Workaround Name	Workaround Description
3D			<p>When we divide and denominator is DNORM, and DNORM is enabled, should produce infinity, but result is 0.</p> <p>Works OK if DNORMs are disabled, legacy features are fine.</p> <p>WA: Don't use DNORM on FDIV.</p>
3D			<p>GA is sending wrong data to fpu for src0 indirect addressing</p> <p>WA: [BDW]: When an instruction has a source region spanning two registers and a destination regioning contained in one register, one of the following must be true:</p> <p>The destination region is entirely contained in the lower Oword of a register.</p> <p>The destination region is entirely contained in the upper Oword of a register.</p> <p>The destination elements are evenly split between the two OWords of a register AND evenly split between the two source registers.</p>
3D			<p>MsaBasic hang with single subspan dispatch enabled.</p> <p>WA: QtrCtrl must not be used for jmp instruction.</p>
3D			<p>DF --&gt; f format conversion for align16 has wrong emask calculation when the source is immediate.</p> <p>WA: In Align16 mode, format conversion from double-float to floats is not allowed when source is immediate data.</p>
3D			<p>SMOV instruction with byte data type. If this is used, corruption. Data read / write is not as expected.</p> <p>WA: Disable SMOV with byte.</p>
GTI	L3	WaDisableL3ErrorDetectionHangOnError	<p>Model hang in wgf11shader5x store_raw tests.</p> <p>WA: Connected ~SVL[9] to LNCFUNIT Incf_csr_bank_hang_override which is then routed to LBCFUNIT.</p>
3D			<p>Byte Mask Media write have issue with byte enable when block width is less then 32.</p> <p>WA: Restriction for A0 and B0 : byte mask media message cannot be used.</p>

Functional Area/Component		Workaround Name	Workaround Description
3D		WaMSFWithNoWatermarkTSGHang	When CS preempts with MSFLUSH with watermark then TSG is blocking the preemption for indefinite amount of time.  WA: Need to disable watermark bit in MSFLUSH command for Mid thread preemption.
3D			SW not use shader channel select to remap RGB.  WA: Not to use shader channel select.
3D		WaDisableMidObjectPreemptionForQuadStrip	VF preemption on quadstrip topology saving wrong IAV marker value.  WA: To not do preemption when doing quadstrip topology.
3D			When smov src and dest as byte and word then hw detects it as mixed mode instruction and it miscompares.  WA: Smov with byte datatype should not be used.
3D			Group id select is not resetting for media walker during context switch from media to GPGPU mode.  WA: Need to disable media walker with groups.
3D		WaDisableMidObjectPreemptionForLineLoop	VF Stats Counters Missing a vertex when preemption enabled.  WA: Disable mid-draw preemption when the draw uses a lineloop topology. (OGL Only)
3D			In certain cases in 3D workload if CS is preempted in window of MI_RS_CONTROL(OFF) to MI_RS_CONTROL(ON), then CS may start RS for the instruction which are in RS disable window bracketed by MI_RS_CONTROL(OFF) to MI_RS_CONTROL(ON) after resubmission.  WA: Whenever programming MI_RS_CONTROL(OFF), disable all the pools so that RS will not generate any produce after execution instruction from this MI_RS_CONTROL(OFF) to MI_RS_CONTROL(ON) zone.

Functional Area/Component		Workaround Name	Workaround Description
3D		WaDisableIndirectDataAndFlushGPGPUWalker	VFE counter overflow due to missing pending_cntr signal for cntr3. WA: Limit urb entries to 63 and MI_ATOMIC_FLUSH need to be inserted after media curbe load command.
3D			Mid Thread Preemption enabling causes VFE TSG hang in Media Context. WA: MEDIA_STATE_FLUSH need to programmed before MI_BATCH_BUFFER_END of the batch buffer with Media_Object or media object walker command.
GTI		WaSendExtraRSGatherConstantAndRSSStoreImmCmds	RS sends Write and later L3 sends same Read. If GAFM gets GFX fence and do RS flush and stall RS, if RS WR comes, it will stay in FIFO due to STALL after fence, L3 read comes when RS is present in GAFM, L3 gets blocked due to same Address WR present. And gets hang as GAFM cannot respond flush due to this dead lock. WA: Inserting 5+ STDW after RS cycle and once it is out then only send L3 RD cycle.
3D		WaDisableMidObjectPreemptionForTrifanOrPolygon	TriFan miscompare in Execlist Preemption test. Cut index that is on a previous context. End the previous; then resume another context with a tri-fan or polygon, and the vertex count is corrupted. If we preempt again we will cause corruption. WA: Disable mid-draw preemption when draw-call has a tri-fan.
3D		WaDisableRHWOOptimizationForRenderHang	IZ/RCZ hang in 3DM11 frame. WA: Disable RHWO optimization, but will have performance impact.
3D			Issue in Trail mode - rd pointers getting corrupted. WA: Reorder mode bit in 3DSTATE_GS should be always leading.
3D			gpgpu_walker_valid need to be reset when start >= dim to avoid corruption in context image. WA: gpgpu_walker_valid need to be reset when dim=0 or start >= dim to avoid corruption in context image

Functional Area/Component		Workaround Name	Workaround Description
GTI	SQIDI		GFXVTDBAR + x100 shadow register should be RO with value of 050A.  WA: BIOS needs to write to the register prior to use so that shadow will have the correct value.
3D			Preemption protocol of csr_dispatch_done followed by tsg_tdg_preemption is broken on Msflush with flush-to-go.  WA: MSFLUSH without watermark and flush-to-go need to be inserted before MSFLUSH with flush to go command.
3D			VF is corrupting GAFS data when preempted on an instance boundary and replayed with instancing enabled.  WA: Disable preemption when using instancing.
3D			VF sends incorrect object end pulses to VS when autostrip enabled workload is context restored before any loop topologies.  WA (A0/B0): Mid-obj preemption need to be disable with trifan and polygon topology.  WA: Disable autostrip when mid-draw pre-emption is enabled.
3D			Vertex is dropped when the preempted on first vertex of a lineloop. This will cause corruption.  WA: Disable mid-draw preemption when the draw uses a lineloop topology.
Media		WaEnableCB2PipeControlCmdInsertion	Context switch test hangs as I3 is not giving the credits to HDC. WA: Two pipe control commands with cs stall enabled and protection off should be issued when there is change in GT slice state to avoid encoder hang.

Functional Area/Component		Workaround Name	Workaround Description
Media			<p>Two instances of Encoder causing 4 bytes of missing data in PAK output.</p> <p>WA Details:</p> <ul style="list-style-type: none"> <li>• Workaround only needs to be applied to INTRA frames after existing QP modulation algorithm.</li> <li>• For INTRA 16x16 MBs, no restriction for QP adjustment, smaller QP can be used.</li> <li>• For INTRA 8x8 and 4x4 MBs, clamp MB QP to the Min QP based on True INTRA Distortion. <ul style="list-style-type: none"> <li>- Distortion in the range of [0, 128), MinQP=12</li> <li>- Distortion in the range of [128, 384), MinQP=14</li> <li>- Distortion in the range of [384, 768), MinQP=16</li> <li>- Distortion in the range of [768, 65535), MinQP =18</li> <li>- If existing QP modulation increase MB QP above 18, no restrictions.</li> </ul> </li> <li>• If True INTRA Distortion is not available (In Perf mode kernel), final INTRA distortion (with mode cost included) can be used.</li> </ul>