



# DATACENTER GPU MANAGER API MANUAL

v1.7 | September 2019

**Reference Manual**



## TABLE OF CONTENTS

<b>Chapter 1. Change Log.....</b>	<b>1</b>
<b>Chapter 2. Modules.....</b>	<b>2</b>
2.1. Administrative.....	3
Init and Shutdown.....	3
Auxiliary information about DCGM engine.....	3
2.1.1. Init and Shutdown.....	3
dcgmlInit.....	3
dcgmShutdown.....	3
dcgmStartEmbedded.....	4
dcgmStopEmbedded.....	4
dcgmConnect.....	5
dcgmConnect_v2.....	6
dcgmDisconnect.....	6
2.1.2. Auxilary information about DCGM engine.....	7
dcgmVersionInfo.....	7
2.2. System.....	7
Discovery.....	7
Grouping.....	7
Field Grouping.....	7
Status handling.....	8
2.2.1. Discovery.....	8
dcgm GetAllDevices.....	8
dcgm GetAllSupportedDevices.....	8
dcgm GetDeviceAttributes.....	9
dcgm GetEntityGroupEntities.....	10
dcgm GetNvLinkLinkStatus.....	10
2.2.2. Grouping.....	11
dcgm GroupCreate.....	11
dcgm GroupDestroy.....	12
dcgm GroupAddDevice.....	13
dcgm GroupAddEntity.....	13
dcgm GroupRemoveDevice.....	14
dcgm GroupRemoveEntity.....	14
dcgm GroupGetInfo.....	15
dcgm GroupGetAllIds.....	16
2.2.3. Field Grouping.....	16
dcgm FieldGroupCreate.....	16
dcgm FieldGroupDestroy.....	17
dcgm FieldGroupGetInfo.....	18
dcgm FieldGroupGetAll.....	18

2.2.4. Status handling.....	19
dcgmStatusCreate.....	19
dcgmStatusDestroy.....	19
dcgmStatusGetCount.....	20
dcgmStatusPopError.....	20
dcgmStatusClear.....	21
2.3. Configuration.....	21
Setup and management.....	21
Manual Invocation.....	21
2.3.1. Setup and management.....	21
dcgmConfigSet.....	22
dcgmConfigGet.....	23
2.3.2. Manual Invocation.....	24
dcgmConfigEnforce.....	24
2.4. Field APIs.....	25
dcgmWatchFields.....	25
dcgmUnwatchFields.....	26
dcgmGetValuesSince.....	26
dcgmGetValuesSince_v2.....	27
dcgmGetLatestValues.....	28
dcgmGetLatestValues_v2.....	29
dcgmGetLatestValuesForFields.....	30
dcgmEntityGetLatestValues.....	30
dcgmEntitiesGetLatestValues.....	31
2.5. Process Statistics.....	32
dcgmWatchPidFields.....	32
dcgmGetPidInfo.....	33
2.6. Job Statistics.....	34
dcgmWatchJobFields.....	34
dcgmJobStartStats.....	35
dcgmJobStopStats.....	35
dcgmJobGetStats.....	36
dcgmJobRemove.....	36
dcgmJobRemoveAll.....	37
2.7. Health Monitor.....	37
dcgmHealthSet.....	38
dcgmHealthGet.....	38
dcgmHealthCheck.....	39
2.8. Policies.....	40
Setup and Management.....	40
Manual Invocation.....	40
2.8.1. Setup and Management.....	40
dcgmPolicySet.....	40

dcgmPolicyGet.....	41
dcgmPolicyRegister.....	42
dcgmPolicyUnregister.....	43
2.8.2. Manual Invocation.....	43
dcgmActionValidate.....	44
dcgmActionValidate_v2.....	44
dcgmRunDiagnostic.....	45
2.9. Topology.....	46
dcgmGetDeviceTopology.....	46
dcgmGetGroupTopology.....	47
2.10. Metadata.....	47
dcgmlntrospectToggleState.....	47
dcgmlntrospectGetFieldsMemoryUsage.....	48
dcgmlntrospectGetHostengineMemoryUsage.....	49
dcgmlntrospectGetFieldsExecTime.....	49
dcgmlntrospectUpdateAll.....	50
2.11. Topology.....	51
dcgmSelectGpusByTopology.....	51
dcgmGetFieldSummary.....	52
2.12. Modules.....	52
dcgmModuleBlacklist.....	52
dcgmModuleGetStatuses.....	53
2.13. Profiling.....	53
dcgmProfGetSupportedMetricGroups.....	53
dcgmProfWatchFields.....	54
dcgmProfUnwatchFields.....	55
2.14. Enums and Macros.....	55
dcgmOperationMode_t.....	56
dcgmOrder_t.....	56
dcgmReturn_t.....	56
dcgmGroupType_t.....	58
dcgmConfigType_t.....	59
dcgmConfigPowerLimitType_t.....	59
DCGM_INT32_BLANK.....	59
DCGM_INT64_BLANK.....	59
DCGM_FP64_BLANK.....	59
DCGM_STR_BLANK.....	60
DCGM_INT32_NOT_FOUND.....	60
DCGM_INT64_NOT_FOUND.....	60
DCGM_FP64_NOT_FOUND.....	60
DCGM_STR_NOT_FOUND.....	60
DCGM_INT32_NOT_SUPPORTED.....	60
DCGM_INT64_NOT_SUPPORTED.....	60

DCGM_FP64_NOT_SUPPORTED.....	60
DCGM_STR_NOT_SUPPORTED.....	60
DCGM_INT32_NOT_PERMISSIONED.....	61
DCGM_INT64_NOT_PERMISSIONED.....	61
DCGM_FP64_NOT_PERMISSIONED.....	61
DCGM_STR_NOT_PERMISSIONED.....	61
DCGM_INT32_IS_BLANK.....	61
DCGM_INT64_IS_BLANK.....	61
DCGM_FP64_IS_BLANK.....	61
DCGM_STR_IS_BLANK.....	62
DCGM_MAX_NUM_DEVICES.....	62
DCGM_NVLINK_MAX_LINKS_PER_GPU.....	62
DCGM_MAX_NUM_SWITCHES.....	62
DCGM_NVLINK_MAX_LINKS_PER_NVSWITCH.....	62
DCGM_MAX_VGPU_INSTANCES_PER_PGPU.....	62
DCGM_MAX_NUM_VGPU_DEVICES.....	62
DCGM_MAX_STR_LENGTH.....	62
DCGM_MAX_CLOCKS.....	62
DCGM_MAX_NUM_GROUPS.....	63
DCGM_MAX_FBC_SESSIONS.....	63
DCGM_VGPU_NAME_BUFFER_SIZE.....	63
DCGM_GRID_LICENSE_BUFFER_SIZE.....	63
DCGM_CONFIG_COMPUTEMODE_DEFAULT.....	63
DCGM_CONFIG_COMPUTEMODE_PROHIBITED.....	63
DCGM_CONFIG_COMPUTEMODE_EXCLUSIVE_PROCESS.....	63
DCGM_HE_PORT_NUMBER.....	63
MAKE_DCGM_VERSION.....	63
DCGM_GROUP_ALL_GPUS.....	63
DCGM_GROUP_MAX_ENTITIES.....	64
2.15. Structure definitions.....	64
dcgmConnectV2Params_v1.....	65
dcgmConnectV2Params_v2.....	65
dcgmGroupInfo_v1.....	65
dcgmGroupEntityPair_t.....	65
dcgmGroupInfo_v2.....	65
dcgmFieldGroupInfo_v1.....	65
dcgmErrorInfo_t.....	65
dcgmClockSet_v1.....	65
dcgmDeviceSupportedClockSets_v1.....	65
dcgmDevicePidAccountingStats_v1.....	65
dcgmDeviceThermals_v1.....	65
dcgmDevicePowerLimits_v1.....	65
dcgmDeviceIdentifiers_v1.....	65

dcgmDeviceMemoryUsage_v1.....	65
dcgmDeviceVgpuUtilInfo_v1.....	65
dcgmDeviceEncStats_v1.....	65
dcgmDeviceFbcStats_y1.....	65
dcgmDeviceFbcSessionInfo_v1.....	65
dcgmDeviceFbcSessions_v1.....	66
dcgmDeviceVgpuEncSessions_v1.....	66
dcgmDeviceVgpuProcessUtilInfo_v1.....	66
dcgmDeviceVgpuids_v1.....	66
dcgmDeviceVgpuTypeInfo_v1.....	66
dcgmDeviceAttributes_v1.....	66
dcgmVgpuDeviceAttributes_v6.....	66
dcgmVgpuInstanceAttributes_v1.....	66
dcgmConfigPerfStateSettings_t.....	66
dcgmConfigPowerLimit_t.....	66
dcgmConfig_v1.....	66
dcgmVgpuConfig_v1.....	66
dcgmPolicyConditionParms_t.....	66
dcgmPolicyViolationNotify_t.....	66
dcgmPolicy_v1.....	66
dcgmPolicyConditionDbe_t.....	66
dcgmPolicyConditionPci_t.....	66
dcgmPolicyConditionMpr_t.....	66
dcgmPolicyConditionThermal_t.....	67
dcgmPolicyConditionPower_t.....	67
dcgmPolicyConditionNvlink_t.....	67
dcgmPolicyConditionXID_t.....	67
dcgmPolicyCallbackResponse_v1.....	67
dcgmFieldValue_v1.....	67
dcgmFieldValue_v2.....	67
dcgmStatSummaryInt64_t.....	67
dcgmStatSummaryInt32_t.....	67
dcgmStatSummaryFp64_t.....	67
dcgmHealthResponse_v1.....	67
dcgmHealthResponse_v2.....	67
dcgmHealthResponse_v3.....	67
dcgmProcessUtilInfo_t.....	67
dcgmProcessUtilSample_t.....	67
dcgmPidSingleInfo_t.....	67
dcgmPidInfo_v1.....	67
dcgmGpuUsageInfo_t.....	67
dcgmJobInfo_v2.....	68
dcgmRunningProcess_v1.....	68

dcgmDiagResponsePerGpu_v1.....	68
dcgmDiagResponse_v3.....	68
dcgmDiagResponse_v4.....	68
dcgmDeviceTopology_v1.....	68
dcgmGroupTopology_v1.....	68
dcgmlntrospectContext_v1.....	68
dcgmlntrospectFieldsExecTime_v1.....	68
dcgmlntrospectFullFieldsExecTime_v1.....	68
dcgmlntrospectMemory_v1.....	68
dcgmlntrospectFullMemory_v1.....	68
dcgmlntrospectCpuUtil_v1.....	68
dcgmNvLinkGpuLinkStatus_t.....	68
dcgmNvLinkNvSwitchLinkStatus_t.....	68
dcgmNvLinkStatus_v1.....	68
dcgmModuleGetStatusesModule_t.....	68
dcgmProfWatchFields_v1.....	68
dcgmProfUnwatchFields_v1.....	69
dcgmVersionInfo_v1.....	69
dcgmPolicyCondition_t.....	69
dcgmPolicyMode_t.....	69
dcgmPolicyIsolation_t.....	69
dcgmPolicyAction_t.....	70
dcgmPolicyValidation_t.....	70
dcgmPolicyFailureResp_t.....	70
dcgmHealthSystems_t.....	70
dcgmHealthWatchResults_t.....	71
dcgmDiagnosticLevel_t.....	71
dcgmDiagResult_t.....	72
dcgmPerGpuTestIndices_t.....	72
dcgmGpuTopologyLevel_t.....	73
dcgmlntrospectLevel_t.....	73
dcgmlntrospectState_t.....	74
dcgmGpuNVLinkErrorType_t.....	74
dcgmNvLinkLinkState_t.....	74
dcgmModuleId_t.....	74
dcgmModuleStatus_t.....	75
dcgmHandle_t.....	75
dcgmGpuGrp_t.....	76
dcgmFieldGrp_t.....	76
dcgmStatus_t.....	76
dcgmConnectV2Params_t.....	76
dcgmGroupInfo_t.....	76
dcgmClockSet_t.....	76

dcgmDeviceSupportedClockSets_t.....	76
dcgmDevicePidAccountingStats_t.....	76
dcgmDeviceThermals_t.....	76
dcgmDevicePowerLimits_t.....	76
dcgmDeviceIdentifiers_t.....	76
dcgmDeviceMemoryUsage_t.....	77
dcgmDeviceVgpuUtilInfo_t.....	77
dcgmDeviceEncStats_t.....	77
dcgmDeviceFbcStats_t.....	77
dcgmDeviceFbcSessionInfo_t.....	77
dcgmDeviceFbcSessions_t.....	77
dcgmDeviceVgpuEncSessions_t.....	77
dcgmDeviceVgpuProcessUtilInfo_t.....	77
dcgmDeviceVgpuids_t.....	77
dcgmDeviceVgpuTypeInfo_t.....	77
dcgmDeviceAttributes_t.....	77
dcgmVgpuDeviceAttributes_t.....	78
dcgmVgpuInstanceAttributes_t.....	78
dcgmConfig_t.....	78
dcgmVgpuConfig_t.....	78
fpRecvUpdates.....	78
dcgmPolicy_t.....	78
dcgmPolicyCallbackResponse_t.....	78
dcgmFieldValueEnumeration_f.....	78
dcgmFieldValueEntityEnumeration_f.....	79
dcgmHealthResponse_t.....	79
dcgmPidInfo_t.....	79
dcgmJobInfo_t.....	79
dcgmRunningProcess_t.....	79
dcgmDiagResponse_t.....	79
dcgmDeviceTopology_t.....	79
dcgmGroupTopology_t.....	79
dcgmlntrospectContext_t.....	79
dcgmlntrospectFieldsExecTime_t.....	80
dcgmlntrospectFullFieldsExecTime_t.....	80
dcgmlntrospectMemory_t.....	80
dcgmlntrospectFullMemory_t.....	80
dcgmlntrospectCpuUtil_t.....	80
dcgmRunDiag_t.....	80
dcgmConnectV2Params_version1.....	80
dcgmConnectV2Params_version2.....	80
dcgmConnectV2Params_version.....	80
dcgmGroupInfo_version1.....	80

dcgmGroupInfo_version2.....	81
dcgmGroupInfo_version.....	81
DCGM_MAX_NUM_FIELD_GROUPS.....	81
DCGM_MAX_FIELD_IDS_PER_FIELD_GROUP.....	81
dcgmFieldGroupInfo_version1.....	81
dcgmFieldGroupInfo_version.....	81
dcgmAllFieldGroup_version1.....	81
dcgmAllFieldGroup_version.....	81
dcgmClockSet_version1.....	81
dcgmClockSet_version.....	82
dcgmDeviceSupportedClockSets_version1.....	82
dcgmDeviceSupportedClockSets_version.....	82
dcgmDevicePidAccountingStats_version1.....	82
dcgmDevicePidAccountingStats_version.....	82
dcgmDeviceThermals_version1.....	82
dcgmDeviceThermals_version.....	82
dcgmDevicePowerLimits_version1.....	82
dcgmDevicePowerLimits_version.....	83
dcgmDeviceIdentifiers_version1.....	83
dcgmDeviceIdentifiers_version.....	83
dcgmDeviceMemoryUsage_version1.....	83
dcgmDeviceMemoryUsage_version.....	83
dcgmDeviceVgpuUtilInfo_version1.....	83
dcgmDeviceVgpuUtilInfo_version.....	83
dcgmDeviceEncStats_version1.....	83
dcgmDeviceEncStats_version.....	84
dcgmDeviceFbcStats_version1.....	84
dcgmDeviceFbcStats_version.....	84
dcgmDeviceFbcSessionInfo_version1.....	84
dcgmDeviceFbcSessionInfo_version.....	84
dcgmDeviceFbcSessions_version1.....	84
dcgmDeviceFbcSessions_version.....	84
dcgmDeviceVgpuEncSessions_version1.....	84
dcgmDeviceVgpuEncSessions_version.....	85
dcgmDeviceVgpuProcessUtilInfo_version1.....	85
dcgmDeviceVgpuProcessUtilInfo_version.....	85
dcgmDeviceVgpuids_version1.....	85
dcgmDeviceVgpuids_version.....	85
dcgmDeviceVgpuTypeInfo_version1.....	85
dcgmDeviceVgpuTypeInfo_version.....	85
dcgmDeviceAttributes_version1.....	85
dcgmDeviceAttributes_version.....	86
DCGM_MAX_VGPU_TYPES_PER_PGPU.....	86

dcgmVgpuDeviceAttributes_version6.....	86
dcgmVgpuDeviceAttributes_version.....	86
DCGM_DEVICE_UUID_BUFFER_SIZE.....	86
dcgmVgpuInstanceAttributes_version1.....	86
dcgmVgpuInstanceAttributes_version.....	86
dcgmConfig_version1.....	86
dcgmConfig_version.....	87
dcgmVgpuConfig_version1.....	87
dcgmVgpuConfig_version.....	87
dcgmPolicy_version1.....	87
dcgmPolicy_version.....	87
dcgmPolicyCallbackResponse_version1.....	87
dcgmPolicyCallbackResponse_version.....	87
DCGM_MAX_BLOB_LENGTH.....	87
dcgmFieldValue_version1.....	87
dcgmFieldValue_version2.....	88
DCGM_FV_FLAG_LIVE_DATA.....	88
DCGM_HEALTH_WATCH_COUNT_V1.....	88
DCGM_HEALTH_WATCH_COUNT_V2.....	88
dcgmHealthResponse_version1.....	88
dcgmHealthResponse_version2.....	88
dcgmHealthResponse_version3.....	88
dcgmHealthResponse_version.....	88
dcgmPidInfo_version1.....	88
dcgmPidInfo_version.....	89
dcgmJobInfo_version2.....	89
dcgmJobInfo_version.....	89
dcgmRunningProcess_version1.....	89
dcgmRunningProcess_version.....	89
dcgmDiagResponse_version3.....	89
dcgmDiagResponse_version4.....	89
dcgmDiagResponse_version5.....	89
dcgmDiagResponse_version.....	89
dcgmDeviceTopology_version1.....	90
dcgmDeviceTopology_version.....	90
dcgmGroupTopology_version1.....	90
dcgmGroupTopology_version.....	90
dcgmlntrospectContext_version1.....	90
dcgmlntrospectContext_version.....	90
dcgmlntrospectFieldsExecTime_version1.....	90
dcgmlntrospectFieldsExecTime_version.....	90
dcgmlntrospectFullFieldsExecTime_version1.....	91
dcgmlntrospectFullFieldsExecTime_version.....	91

dcgmlIntrospectMemory_version1.....	91
dcgmlIntrospectMemory_version.....	91
dcgmlIntrospectFullMemory_version1.....	91
dcgmlIntrospectFullMemory_version.....	91
dcgmlIntrospectCpuUtil_version1.....	91
dcgmlIntrospectCpuUtil_version.....	91
dcgmRunDiag_version1.....	92
dcgmRunDiag_version2.....	92
dcgmRunDiag_version3.....	92
dcgmRunDiag_version4.....	92
dcgmRunDiag_version5.....	92
dcgmRunDiag_version.....	92
DCGM_GEGE_FLAG_ONLY_SUPPORTED.....	92
dcgmNvLinkStatus_version1.....	92
DCGM_MODULE_STATUSES_CAPACITY.....	93
dcgmModuleGetStatuses_version1.....	93
DCGM_PROF_MAX_NUM_GROUPS.....	93
DCGM_PROF_MAX_FIELD_IDS_PER_GROUP.....	93
dcgmProfGetMetricGroups_version2.....	93
dcgmProfWatchFields_version1.....	93
dcgmProfUnwatchFields_version1.....	93
dcgmVersionInfo_version1.....	93
2.16. Field Types.....	94
DCGM_FT_BINARY.....	94
DCGM_FT_DOUBLE.....	94
DCGM_FT_INT64.....	94
DCGM_FT_STRING.....	94
DCGM_FT_TIMESTAMP.....	94
2.17. Field Scope.....	94
DCGM_FS_GLOBAL.....	94
DCGM_FS_ENTITY.....	94
DCGM_FS_DEVICE.....	94
DCGM_CUDA_COMPUTE_CAPABILITY_MAJOR.....	95
DCGM_CLOCKS_THROTTLE_REASON_GPU_IDLE.....	95
DCGM_CLOCKS_THROTTLE_REASON_CLOCKS_SETTING.....	95
DCGM_CLOCKS_THROTTLE_REASON_SW_POWER_CAP.....	95
DCGM_CLOCKS_THROTTLE_REASON_HW_SLOWDOWN.....	95
DCGM_CLOCKS_THROTTLE_REASON_SYNC_BOOST.....	96
DCGM_CLOCKS_THROTTLE_REASON_SW_THERMAL.....	96
DCGM_CLOCKS_THROTTLE_REASON_HW_THERMAL.....	96
DCGM_CLOCKS_THROTTLE_REASON_HW_POWER_BRAKE.....	96
DCGM_CLOCKS_THROTTLE_REASON_DISPLAY_CLOCKS.....	97
2.18. Field Entity.....	97

dcgm_field_entity_group_t.....	97
dcgm_field_eid_t.....	97
<b>2.19. Field Identifiers.....</b>	<b>97</b>
DcgmFieldGetById.....	98
DcgmFieldGetByTag.....	98
DcgmFieldsInit.....	98
DcgmFieldsTerm.....	98
DcgmFieldsGetEntityGroupString.....	99
DCGM_FI_UNKNOWN.....	99
DCGM_FI_DRIVER_VERSION.....	99
DCGM_FI_DEV_COUNT.....	99
DCGM_FI_DEV_NAME.....	99
DCGM_FI_DEV_BRAND.....	99
DCGM_FI_DEV_NVML_INDEX.....	99
DCGM_FI_DEV_SERIAL.....	99
DCGM_FI_DEV_UUID.....	99
DCGM_FI_DEV_MINOR_NUMBER.....	100
DCGM_FI_DEV_OEM_INFOROM_VER.....	100
DCGM_FI_DEV_PCI_BUSID.....	100
DCGM_FI_DEV_PCI_COMBINED_ID.....	100
DCGM_FI_DEV_PCI_SUBSYS_ID.....	100
DCGM_FI_GPU_TOPOLOGY_PCI.....	100
DCGM_FI_GPU_TOPOLOGY_NVLINK.....	100
DCGM_FI_GPU_TOPOLOGY_AFFINITY.....	100
DCGM_FI_DEV_CUDA_COMPUTE_CAPABILITY.....	100
DCGM_FI_DEV_COMPUTE_MODE.....	100
DCGM_FI_DEV_CPU_AFFINITY_0.....	100
DCGM_FI_DEV_CPU_AFFINITY_1.....	101
DCGM_FI_DEV_CPU_AFFINITY_2.....	101
DCGM_FI_DEV_CPU_AFFINITY_3.....	101
DCGM_FI_DEV_ECC_INFOROM_VER.....	101
DCGM_FI_DEV_POWER_INFOROM_VER.....	101
DCGM_FI_DEV_INFOROM_IMAGE_VER.....	101
DCGM_FI_DEV_INFOROM_CONFIG_CHECK.....	101
DCGM_FI_DEV_INFOROM_CONFIG_VALID.....	101
DCGM_FI_DEV_VBIOS_VERSION.....	101
DCGM_FI_DEV_BAR1_TOTAL.....	101
DCGM_FI_SYNC_BOOST.....	101
DCGM_FI_DEV_BAR1_USED.....	102
DCGM_FI_DEV_BAR1_FREE.....	102
DCGM_FI_DEV_SM_CLOCK.....	102
DCGM_FI_DEV_MEM_CLOCK.....	102
DCGM_FI_DEV_VIDEO_CLOCK.....	102

DCGM_FI_DEV_APP_SM_CLOCK.....	102
DCGM_FI_DEV_APP_MEM_CLOCK.....	102
DCGM_FI_DEV_CLOCK_THROTTLE_REASONS.....	102
DCGM_FI_DEV_MAX_SM_CLOCK.....	102
DCGM_FI_DEV_MAX_MEM_CLOCK.....	102
DCGM_FI_DEV_MAX_VIDEO_CLOCK.....	102
DCGM_FI_DEV_AUTOBOOST.....	103
DCGM_FI_DEV_SUPPORTED_CLOCKS.....	103
DCGM_FI_DEV_MEMORY_TEMP.....	103
DCGM_FI_DEV_GPU_TEMP.....	103
DCGM_FI_DEV_POWER_USAGE.....	103
DCGM_FI_DEV_TOTAL_ENERGY_CONSUMPTION.....	103
DCGM_FI_DEV_SLOWDOWN_TEMP.....	103
DCGM_FI_DEV_SHUTDOWN_TEMP.....	103
DCGM_FI_DEV_POWER_MGMT_LIMIT.....	103
DCGM_FI_DEV_POWER_MGMT_LIMIT_MIN.....	103
DCGM_FI_DEV_POWER_MGMT_LIMIT_MAX.....	103
DCGM_FI_DEV_POWER_MGMT_LIMIT_DEF.....	104
DCGM_FI_DEV_ENFORCED_POWER_LIMIT.....	104
DCGM_FI_DEV_PSTATE.....	104
DCGM_FI_DEV_FAN_SPEED.....	104
DCGM_FI_DEV_PCIE_TX_THROUGHPUT.....	104
DCGM_FI_DEV_PCIE_RX_THROUGHPUT.....	104
DCGM_FI_DEV_PCIE_REPLAY_COUNTER.....	104
DCGM_FI_DEV_GPU_UTIL.....	104
DCGM_FI_DEV_MEM_COPY_UTIL.....	104
DCGM_FI_DEV_ACCOUNTING_DATA.....	104
DCGM_FI_DEV_ENC_UTIL.....	105
DCGM_FI_DEV_DEC_UTIL.....	105
DCGM_FI_DEV_MEM_COPY_UTIL_SAMPLES.....	105
DCGM_FI_DEV_GRAPHICS_PIDS.....	105
DCGM_FI_DEV_COMPUTE_PIDS.....	105
DCGM_FI_DEV_XID_ERRORS.....	105
DCGM_FI_DEV_PCIE_MAX_LINK_GEN.....	105
DCGM_FI_DEV_PCIE_MAX_LINK_WIDTH.....	105
DCGM_FI_DEV_PCIE_LINK_GEN.....	105
DCGM_FI_DEV_PCIE_LINK_WIDTH.....	105
DCGM_FI_DEV_POWER_VIOLATION.....	105
DCGM_FI_DEV_THERMAL_VIOLATION.....	106
DCGM_FI_DEV_SYNC_BOOST_VIOLATION.....	106
DCGM_FI_DEV_BOARD_LIMIT_VIOLATION.....	106
DCGM_FI_DEV_LOW_UTIL_VIOLATION.....	106
DCGM_FI_DEV_RELIABILITY_VIOLATION.....	106

DCGM_FI_DEV_TOTAL_APP_CLOCKS_VIOLATION.....	106
DCGM_FI_DEV_TOTAL_BASE_CLOCKS_VIOLATION.....	106
DCGM_FI_DEV_FB_TOTAL.....	106
DCGM_FI_DEV_FB_FREE.....	106
DCGM_FI_DEV_FB_USED.....	106
DCGM_FI_DEV_ECC_CURRENT.....	107
DCGM_FI_DEV_ECC_PENDING.....	107
DCGM_FI_DEV_ECC_SBE_VOL_TOTAL.....	107
DCGM_FI_DEV_ECC_DBE_VOL_TOTAL.....	107
DCGM_FI_DEV_ECC_SBE_AGG_TOTAL.....	107
DCGM_FI_DEV_ECC_DBE_AGG_TOTAL.....	107
DCGM_FI_DEV_ECC_SBE_VOL_L1.....	107
DCGM_FI_DEV_ECC_DBE_VOL_L1.....	107
DCGM_FI_DEV_ECC_SBE_VOL_L2.....	107
DCGM_FI_DEV_ECC_DBE_VOL_L2.....	107
DCGM_FI_DEV_ECC_SBE_VOL_DEV.....	107
DCGM_FI_DEV_ECC_DBE_VOL_DEV.....	108
DCGM_FI_DEV_ECC_SBE_VOL_REG.....	108
DCGM_FI_DEV_ECC_DBE_VOL_REG.....	108
DCGM_FI_DEV_ECC_SBE_VOL_TEX.....	108
DCGM_FI_DEV_ECC_DBE_VOL_TEX.....	108
DCGM_FI_DEV_ECC_SBE_AGG_L1.....	108
DCGM_FI_DEV_ECC_DBE_AGG_L1.....	108
DCGM_FI_DEV_ECC_SBE_AGG_L2.....	108
DCGM_FI_DEV_ECC_DBE_AGG_L2.....	108
DCGM_FI_DEV_ECC_SBE_AGG_DEV.....	108
DCGM_FI_DEV_ECC_DBE_AGG_DEV.....	108
DCGM_FI_DEV_ECC_SBE_AGG_REG.....	109
DCGM_FI_DEV_ECC_DBE_AGG_REG.....	109
DCGM_FI_DEV_ECC_SBE_AGG_TEX.....	109
DCGM_FI_DEV_ECC_DBE_AGG_TEX.....	109
DCGM_FI_DEV_RETIRE_SBE.....	109
DCGM_FI_DEV_RETIRE_DBE.....	109
DCGM_FI_DEV_RETIRE_PENDING.....	109
DCGM_FI_DEV_VIRTUAL_MODE.....	109
DCGM_FI_DEV_SUPPORTED_TYPE_INFO.....	109
DCGM_FI_DEV_CREATABLE_VGPU_TYPE_IDS.....	109
DCGM_FI_DEV_VGPU_INSTANCE_IDS.....	110
DCGM_FI_DEV_VGPU_UTILIZATIONS.....	110
DCGM_FI_DEV_VGPU_PER_PROCESS_UTILIZATION.....	110
DCGM_FI_DEV_ENC_STATS.....	110
DCGM_FI_DEV_FBC_STATS.....	110
DCGM_FI_DEV_FBC_SESSIONS_INFO.....	110

DCGM_FI_DEV_VGPU_VM_ID.....	110
DCGM_FI_DEV_VGPU_VM_NAME.....	110
DCGM_FI_DEV_VGPU_TYPE.....	110
DCGM_FI_DEV_VGPU_UUID.....	110
DCGM_FI_DEV_VGPU_DRIVER_VERSION.....	110
DCGM_FI_DEV_VGPU_MEMORY_USAGE.....	111
DCGM_FI_DEV_VGPU_LICENSE_STATUS.....	111
DCGM_FI_DEV_VGPU_FRAME_RATE_LIMIT.....	111
DCGM_FI_DEV_VGPU_ENC_STATS.....	111
DCGM_FI_DEV_VGPU_ENC_SESSIONS_INFO.....	111
DCGM_FI_DEV_VGPU_FBC_STATS.....	111
DCGM_FI_DEV_VGPU_FBC_SESSIONS_INFO.....	111
DCGM_FI_FIRST_VGPU_FIELD_ID.....	111
DCGM_FI_LAST_VGPU_FIELD_ID.....	111
DCGM_FI_MAX_VGPU_FIELDS.....	111
DCGM_FI_INTERNAL_FIELDS_0_START.....	112
DCGM_FI_INTERNAL_FIELDS_0_END.....	112
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P00.....	112
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P00.....	112
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P00.....	112
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P00.....	112
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P01.....	112
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P01.....	113
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P01.....	113
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P01.....	113
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P02.....	113
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P02.....	113
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P02.....	113
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P02.....	113
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P03.....	114
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P03.....	114
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P03.....	114
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P03.....	114
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P04.....	114
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P04.....	114
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P04.....	114
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P04.....	114
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P05.....	115
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P05.....	115
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P05.....	115
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P05.....	115
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P06.....	115
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P06.....	115

DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P06.....	115
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P06.....	116
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P07.....	116
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P07.....	116
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P07.....	116
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P07.....	116
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P08.....	116
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P08.....	116
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P08.....	117
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P08.....	117
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P09.....	117
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P09.....	117
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P09.....	117
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P09.....	117
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P10.....	117
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P10.....	118
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P10.....	118
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P10.....	118
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P11.....	118
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P11.....	118
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P11.....	118
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P11.....	118
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P12.....	119
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P12.....	119
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P12.....	119
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P12.....	119
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P13.....	119
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P13.....	119
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P13.....	119
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P13.....	119
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P14.....	120
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P14.....	120
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P14.....	120
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P14.....	120
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P15.....	120
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P15.....	120
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P15.....	120
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P15.....	121
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P16.....	121
DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P16.....	121
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P16.....	121
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P16.....	121
DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P17.....	121

DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P17.....	121
DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P17.....	122
DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P17.....	122
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P00.....	122
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P00.....	122
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P01.....	122
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P01.....	122
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P02.....	122
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P02.....	123
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P03.....	123
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P03.....	123
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P04.....	123
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P04.....	123
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P05.....	123
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P05.....	123
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P06.....	123
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P06.....	124
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P07.....	124
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P07.....	124
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P08.....	124
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P08.....	124
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P09.....	124
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P09.....	124
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P10.....	124
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P10.....	125
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P11.....	125
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P11.....	125
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P12.....	125
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P12.....	125
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P13.....	125
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P13.....	125
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P14.....	125
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P14.....	126
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P15.....	126
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P15.....	126
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P16.....	126
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P16.....	126
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P17.....	126
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P17.....	126
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P00.....	127
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P00.....	127
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P01.....	127
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P01.....	127

DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P02.....	127
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P02.....	127
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P03.....	127
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P03.....	127
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P04.....	128
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P04.....	128
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P05.....	128
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P05.....	128
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P06.....	128
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P06.....	128
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P07.....	128
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P07.....	128
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P08.....	129
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P08.....	129
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P09.....	129
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P09.....	129
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P10.....	129
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P10.....	129
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P11.....	129
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P11.....	129
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P12.....	130
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P12.....	130
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P13.....	130
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P13.....	130
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P14.....	130
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P14.....	130
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P15.....	130
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P15.....	130
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P16.....	131
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P16.....	131
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P17.....	131
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P17.....	131
DCGM_FI_DEV_NVSWITCH_FATAL_ERRORS.....	131
DCGM_FI_DEV_NVSWITCH_NON_FATAL_ERRORS.....	131
DCGM_FI_FIRST_NVSWITCH_FIELD_ID.....	131
DCGM_FI_LAST_NVSWITCH_FIELD_ID.....	131
DCGM_FI_MAX_NVSWITCH_FIELDS.....	132
DCGM_FI_PROF_GR_ENGINE_ACTIVE.....	132
DCGM_FI_PROF_SM_ACTIVE.....	132
DCGM_FI_PROF_SM_OCCUPANCY.....	132
DCGM_FI_PROF_PIPE_TENSOR_ACTIVE.....	132
DCGM_FI_PROF_DRAM_ACTIVE.....	132
DCGM_FI_PROF_PIPE_FP64_ACTIVE.....	132

DCGM_FI_PROF_PIPE_FP32_ACTIVE.....	132
DCGM_FI_PROF_PIPE_FP16_ACTIVE.....	132
DCGM_FI_PROF_PCIE_TX_BYTES.....	133
DCGM_FI_PROF_PCIE_RX_BYTES.....	133
DCGM_FI_PROF_NVLINK_TX_BYTES.....	133
DCGM_FI_PROF_NVLINK_RX_BYTES.....	133
DCGM_FI_MAX_FIELDS.....	133
2.20. DCGMAPI_Admin_ExecCtrl.....	133
dcgmUpdateAllFields.....	133
dcgmPolicyTrigger.....	134
<b>Chapter 3. Data Structures.....</b>	<b>135</b>
dcgm_field_meta_t.....	137
dcgm_field_output_format_t.....	137
dcgmClockSet_v1.....	137
version.....	137
memClock.....	137
smClock.....	137
dcgmConfig_v1.....	137
version.....	138
gpuld.....	138
eccMode.....	138
computeMode.....	138
perfState.....	138
powerLimit.....	138
dcgmConfigPerfStateSettings_t.....	138
syncBoost.....	139
targetClocks.....	139
dcgmConfigPowerLimit_t.....	139
type.....	139
val.....	139
dcgmConnectV2Params_v1.....	139
version.....	139
persistAfterDisconnect.....	139
dcgmConnectV2Params_v2.....	140
version.....	140
persistAfterDisconnect.....	140
timeoutMs.....	140
addressIsUnixSocket.....	140
dcgmDeviceAttributes_v1.....	140
version.....	141
clockSets.....	141
thermalSettings.....	141
powerLimits.....	141

identifiers.....	141
memoryUsage.....	141
unusedVgpuids.....	141
unusedActiveVgpuInstanceCount.....	141
unusedVgpuInstancelds.....	141
dcmDeviceEncStats_v1.....	141
version.....	142
sessionCount.....	142
averageFps.....	142
averageLatency.....	142
dcmDeviceFbcSessionInfo_v1.....	142
version.....	143
sessionId.....	143
pid.....	143
vgpuld.....	143
displayOrdinal.....	143
sessionType.....	143
sessionFlags.....	143
hMaxResolution.....	143
vMaxResolution.....	143
hResolution.....	143
vResolution.....	143
averageFps.....	144
averageLatency.....	144
dcmDeviceFbcSessions_v1.....	144
version.....	144
sessionCount.....	144
sessionInfo.....	144
dcmDeviceFbcStats_v1.....	144
version.....	145
sessionCount.....	145
averageFps.....	145
averageLatency.....	145
dcmDeviceIdentifiers_v1.....	145
version.....	146
brandName.....	146
deviceName.....	146
pciBusId.....	146
serial.....	146
uuid.....	146
vbios.....	146
inforomImageVersion.....	146
pciDeviceId.....	146

pciSubSystemId.....	146
driverVersion.....	146
virtualizationMode.....	146
dcgmDeviceMemoryUsage_v1.....	147
version.....	147
bar1Total.....	147
fbTotal.....	147
fbUsed.....	147
fbFree.....	147
dcgmDevicePidAccountingStats_v1.....	147
version.....	147
pid.....	147
gpuUtilization.....	147
memoryUtilization.....	148
maxMemoryUsage.....	148
startTimestamp.....	148
activeTimeUsec.....	148
dcgmDevicePowerLimits_v1.....	148
version.....	149
curPowerLimit.....	149
defaultPowerLimit.....	149
enforcedPowerLimit.....	149
minPowerLimit.....	149
maxPowerLimit.....	149
dcgmDeviceSupportedClockSets_v1.....	149
version.....	150
count.....	150
clockSet.....	150
dcgmDeviceThermals_v1.....	150
version.....	150
slowdownTemp.....	150
shutdownTemp.....	150
dcgmDeviceTopology_v1.....	150
version.....	150
cpuAffinityMask.....	150
numGpus.....	151
gpuld.....	151
path.....	151
localNvLinkIds.....	151
dcgmDeviceVgpuEncSessions_v1.....	151
version.....	152
vgpuld.....	152
sessionId.....	152

pid.....	152
codecType.....	152
hResolution.....	152
vResolution.....	152
averageFps.....	152
averageLatency.....	152
dcmDeviceVgpuInfo_v1.....	152
version.....	153
unusedSupportedVgpuTypeCount.....	153
unusedSupportedVgpuTypeIds.....	153
unusedCreatableVgpuTypeCount.....	153
unusedCreatableVgpuTypeIds.....	153
dcmDeviceVgpuProcessUtilInfo_v1.....	153
version.....	154
vgpuld.....	154
vgpuProcessSamplesCount.....	154
pid.....	154
processName.....	154
smUtil.....	154
memUtil.....	154
encUtil.....	154
decUtil.....	154
dcmDeviceVgpuTypeInfo_v1.....	154
version.....	155
vgpuTypeInfo.....	155
vgpuTypeName.....	155
vgpuTypeClass.....	155
vgpuTypeLicense.....	155
deviceid.....	155
subsystemId.....	155
numDisplayHeads.....	155
maxInstances.....	155
frameRateLimit.....	155
maxResolutionX.....	155
maxResolutionY.....	155
fbTotal.....	155
dcmDeviceVgpuUtilInfo_v1.....	156
version.....	156
vgpuld.....	156
smUtil.....	156
memUtil.....	156
encUtil.....	156
decUtil.....	156

dcmDiagResponse_v3.....	156
version.....	157
gpuCount.....	157
blacklist.....	157
nvmlLibrary.....	157
cudaMainLibrary.....	157
cudaRuntimeLibrary.....	157
permissions.....	157
persistenceMode.....	157
environment.....	157
pageRetirement.....	157
inforom.....	157
graphicsProcesses.....	158
perGpuResponses.....	158
systemError.....	158
dcmDiagResponse_v4.....	158
version.....	159
gpuCount.....	159
levelOneTestCount.....	159
levelOneResults.....	159
perGpuResponses.....	159
systemError.....	159
trainingMsg.....	159
dcmDiagResponsePerGpu_v1.....	159
gpuld.....	160
hwDiagnosticReturn.....	160
results.....	160
dcmErrorInfo_t.....	160
gpuld.....	160
fieldId.....	160
status.....	160
dcmFieldGroupInfo_v1.....	160
version.....	161
numFieldIds.....	161
fieldGroupId.....	161
fieldGroupName.....	161
fieldIds.....	161
dcmFieldValue_v1.....	161
version.....	162
fieldId.....	162
fieldType.....	162
status.....	162
ts.....	162

i64.....	162
dbl.....	162
str.....	162
blob.....	162
value.....	162
dcmFieldValue_v2.....	162
version.....	163
entityGroupId.....	163
entityId.....	163
fieldId.....	163
fieldType.....	163
status.....	163
unused.....	163
ts.....	163
i64.....	163
dbl.....	163
str.....	163
blob.....	163
value.....	163
dcmGpuUsageInfo_t.....	164
gpuld.....	165
energyConsumed.....	165
powerUsage.....	165
pcieRxBandwidth.....	165
pcieTxBandwidth.....	165
pcieReplays.....	165
startTime.....	165
endTime.....	165
smUtilization.....	165
memoryUtilization.....	165
eccSingleBit.....	165
eccDoubleBit.....	166
memoryClock.....	166
smClock.....	166
numXidCriticalErrors.....	166
xidCriticalErrorsTs.....	166
numComputePids.....	166
computePidInfo.....	166
numGraphicsPids.....	166
graphicsPidInfo.....	166
maxGpuMemoryUsed.....	166
powerViolationTime.....	166
thermalViolationTime.....	167

reliabilityViolationTime.....	167
boardLimitViolationTime.....	167
lowUtilizationTime.....	167
syncBoostTime.....	167
overallHealth.....	167
system.....	167
health.....	167
dcmgGroupEntityPair_t.....	167
entityGroupId.....	168
entityId.....	168
dcmgGroupInfo_v1.....	168
version.....	168
count.....	168
gpuldList.....	168
groupName.....	168
dcmgGroupInfo_v2.....	168
version.....	169
count.....	169
groupName.....	169
entityList.....	169
dcmgGroupTopology_v1.....	169
version.....	169
groupCpuAffinityMask.....	169
numaOptimalFlag.....	169
slowestPath.....	169
dcmgHealthResponse_v1.....	169
version.....	170
overallHealth.....	170
gpuCount.....	170
gpuld.....	170
incidentCount.....	170
system.....	170
health.....	170
errorString.....	170
dcmgHealthResponse_v2.....	170
version.....	171
overallHealth.....	171
entityCount.....	171
entityGroupId.....	171
entityId.....	171
incidentCount.....	171
system.....	171
health.....	171

errorString.....	171
dcgmHealthResponse_v3.....	171
version.....	172
overallHealth.....	172
entityCount.....	172
entityGroupId.....	172
entityId.....	172
incidentCount.....	172
system.....	172
health.....	172
errors.....	172
errorCount.....	172
dcgmlIntrospectContext_v1.....	172
version.....	173
introspectLvl.....	173
fieldGroupId.....	173
fieldId.....	173
contextId.....	173
dcgmlIntrospectCpuUtil_v1.....	173
version.....	173
total.....	173
kernel.....	173
user.....	173
dcgmlIntrospectFieldsExecTime_v1.....	173
version.....	174
meanUpdateFreqUsec.....	174
recentUpdateUsec.....	174
totalEverUpdateUsec.....	174
dcgmlIntrospectFullFieldsExecTime_v1.....	174
version.....	175
aggregateInfo.....	175
hasGlobalInfo.....	175
globalInfo.....	175
gpuInfoCount.....	175
gpuidsForGpuInfo.....	175
gpuInfo.....	175
dcgmlIntrospectFullMemory_v1.....	175
version.....	176
aggregateInfo.....	176
hasGlobalInfo.....	176
globalInfo.....	176
gpuInfoCount.....	176
gpuidsForGpuInfo.....	176

gpuInfo.....	176
dcgmIntrospectMemory_v1.....	176
version.....	177
bytesUsed.....	177
dcgmJobInfo_v2.....	177
version.....	177
numGpus.....	177
summary.....	177
gpus.....	177
dcgmModuleGetStatusesModule_t.....	177
id.....	178
status.....	178
dcgmNvLinkGpuLinkStatus_t.....	178
entityId.....	178
linkState.....	178
dcgmNvLinkNvSwitchLinkStatus_t.....	178
entityId.....	178
linkState.....	178
dcgmNvLinkStatus_v1.....	178
version.....	179
numGpus.....	179
gpus.....	179
numNvSwitches.....	179
nvSwitches.....	179
dcgmPidInfo_v1.....	179
version.....	180
pid.....	180
numGpus.....	180
summary.....	180
gpus.....	180
dcgmPidSingleInfo_t.....	180
gpuld.....	181
energyConsumed.....	181
pcieRxBandwidth.....	181
pcieTxBandwidth.....	181
pcieReplays.....	181
startTime.....	181
endTime.....	181
processUtilization.....	181
smUtilization.....	181
memoryUtilization.....	181
eccSingleBit.....	182
eccDoubleBit.....	182

memoryClock.....	182
smClock.....	182
numXidCriticalErrors.....	182
xidCriticalErrorsTs.....	182
numOtherComputePids.....	182
otherComputePids.....	182
numOtherGraphicsPids.....	182
otherGraphicsPids.....	182
maxGpuMemoryUsed.....	182
powerViolationTime.....	182
thermalViolationTime.....	183
reliabilityViolationTime.....	183
boardLimitViolationTime.....	183
lowUtilizationTime.....	183
syncBoostTime.....	183
overallHealth.....	183
system.....	183
health.....	183
dcgmPolicy_v1.....	183
version.....	184
condition.....	184
mode.....	184
isolation.....	184
action.....	184
validation.....	184
response.....	184
parms.....	184
dcgmPolicyCallbackResponse_v1.....	184
version.....	185
condition.....	185
dbe.....	185
pci.....	185
mpr.....	185
thermal.....	185
power.....	185
nvlink.....	185
xid.....	185
dcgmPolicyConditionDbe_t.....	185
timestamp.....	186
location.....	186
numerrors.....	186
dcgmPolicyConditionMpr_t.....	186
timestamp.....	186

sbepages.....	186
dbepages.....	186
dcmPolicyConditionNvlink_t.....	186
timestamp.....	187
fieldId.....	187
counter.....	187
dcmPolicyConditionParms_t.....	187
dcmPolicyConditionPci_t.....	187
timestamp.....	187
counter.....	187
dcmPolicyConditionPower_t.....	187
timestamp.....	188
powerViolation.....	188
dcmPolicyConditionThermal_t.....	188
timestamp.....	188
thermalViolation.....	188
dcmPolicyConditionXID_t.....	188
timestamp.....	188
errnum.....	188
dcmPolicyViolationNotify_t.....	188
gpuld.....	189
violationOccurred.....	189
dcmProcessUtilInfo_t.....	189
dcmProcessUtilSample_t.....	189
dcmProfUnwatchFields_v1.....	189
groupId.....	189
flags.....	189
dcmProfWatchFields_v1.....	189
groupId.....	189
numFieldIds.....	190
fieldIds.....	190
updateFreq.....	190
maxKeepAge.....	190
maxKeepSamples.....	190
flags.....	190
dcmRunningProcess_v1.....	190
version.....	191
pid.....	191
memoryUsed.....	191
dcmStatSummaryFp64_t.....	191
minValue.....	191
maxValue.....	191
average.....	191

dcmStatSummaryInt32_t.....	191
minValue.....	192
maxValue.....	192
average.....	192
dcmStatSummaryInt64_t.....	192
minValue.....	192
maxValue.....	192
average.....	192
dcmVersionInfo_v1.....	192
version.....	193
changelist.....	193
platform.....	193
branch.....	193
driverVersion.....	193
buildDate.....	193
dcmVgpuConfig_v1.....	193
version.....	194
gpuld.....	194
eccMode.....	194
computeMode.....	194
perfState.....	194
powerLimit.....	194
dcmVgpuDeviceAttributes_v6.....	194
version.....	195
activeVgpuInstanceCount.....	195
activeVgpuInstancelds.....	195
creatableVgpuTypeCount.....	195
creatableVgpuTypeIds.....	195
supportedVgpuTypeCount.....	195
supportedVgpuTypeInfo.....	195
vgpuUtilInfo.....	195
gpuUtil.....	195
memCopyUtil.....	195
encUtil.....	196
decUtil.....	196
dcmVgpuInstanceAttributes_v1.....	196
version.....	197
vmlId.....	197
vmName.....	197
vgpuTypeld.....	197
vgpuUuid.....	197
vgpuDriverVersion.....	197
fbUsage.....	197

licenseStatus.....	197
frameRateLimit.....	197
<b>Chapter 4. Data Fields.....</b>	<b>198</b>



# Chapter 1. CHANGE LOG

This chapter list changes in API that were introduced to the library.

## 1.3.0

- ▶ Field Groups, GPU Groups, and field watches created with a handle returned from `dcmConnect()` are now cleaned up upon disconnect. `dcmConnect_v2()` can be used to get the old behavior of objects persisting after disconnect.
- ▶ `dcmConnect_v2()` was added as a method for specifying additional connection options when connecting to the host engine.
- ▶ `dcmUnwatchFields()` was added as a method of unwatching fields that were previously watched with `dcmWatchFields()`
- ▶ `dcmActionValidate_v2()` was added to be able to pass more parameters to the DCGM GPU Diagnostic.
- ▶ `dcmDiagResponse_t` was increased from v2 to v3. See `dcmDiagResponse_v3` for details

## 1.2.3

- ▶ No API changes in this version.

## 1.1.1

- ▶ `dcm GetAllSupportedDevices()` was added as a method to get DCGM-supported GPU Ids. `dcm GetAllDevices()` can still be used to get all GPU Ids in the system.

## 1.0.0

- ▶ Initial Release.

# Chapter 2. MODULES

Here is a list of all modules:

- ▶ Administrative
  - ▶ Init and Shutdown
  - ▶ Auxilary information about DCGM engine.
- ▶ System
  - ▶ Discovery
  - ▶ Grouping
  - ▶ Field Grouping
  - ▶ Status handling
- ▶ Configuration
  - ▶ Setup and management
  - ▶ Manual Invocation
- ▶ Field APIs
- ▶ Process Statistics
- ▶ Job Statistics
- ▶ Health Monitor
- ▶ Policies
  - ▶ Setup and Management
  - ▶ Manual Invocation
- ▶ Topology
- ▶ Metadata
- ▶ Topology
- ▶ Modules
- ▶ Profiling
- ▶ Enums and Macros
- ▶ Structure definitions

- ▶ Field Types
- ▶ Field Scope
- ▶ Field Entity
- ▶ Field Identifiers
- ▶ DCGMAPI\_Admin\_ExecCtrl

## 2.1. Administrative

This chapter describes the administration interfaces for DCGM. It is the user's responsibility to call `dcmgInit()` before calling any other methods, and `dcmgShutdown()` once DCGM is no longer being used. The APIs in Administrative module can be broken down into following categories:

### Init and Shutdown

Auxilary information about DCGM engine.

#### 2.1.1. Init and Shutdown

Administrative

Describes APIs to Initialize and Shutdown the DCGM Engine.

`dcmgReturn_t dcmgInit (void)`

**Returns**

- ▶ DCGM\_ST\_OK if DCGM has been properly initialized
- ▶ DCGM\_ST\_INIT\_ERROR if there was an error initializing the library

**Description**

This method is used to initialize DCGM within this process. This must be called before `dcmgStartEmbedded()` or `dcmgConnect()`

\*

`dcmgReturn_t dcmgShutdown (void)`

**Returns**

- ▶ DCGM\_ST\_OK if DCGM has been properly shut down
- ▶ DCGM\_ST\_UNINITIALIZED if the library was not shut down properly

## Description

This method is used to shut down DCGM. Any embedded host engines or remote connections will automatically be shut down as well.

**dcgmReturn\_t dcgmStartEmbedded (dcgmOperationMode\_t opMode,  
dcgmHandle\_t \*pDcgmHandle)**

## Parameters

### opMode

IN : Collect data automatically or manually when asked by the user.

### pDcgmHandle

OUT : DCGM Handle to use for API calls

## Returns

- ▶ DCGM\_ST\_OK if DCGM was started successfully within our process
- ▶ DCGM\_ST\_UNINITIALIZED if DCGM has not been initialized with dcgmInit yet

## Description

Start an embedded host engine agent within this process.

The agent is loaded as a shared library. This mode is provided to avoid any extra jitter associated with an additional autonomous agent needs to be managed. In this mode, the user has to periodically call APIs such as [dcgmPolicyTrigger](#) and [dcgmUpdateAllFields](#) which tells DCGM to wake up and perform data collection and operations needed for policy management.

**dcgmReturn\_t dcgmStopEmbedded (dcgmHandle\_t pDcgmHandle)**

## Parameters

### pDcgmHandle

IN : DCGM Handle of the embedded host engine that came from  
dcgmStartEmbedded

## Returns

- ▶ DCGM\_ST\_OK if DCGM was stopped successfully within our process
- ▶ DCGM\_ST\_UNINITIALIZED if DCGM has not been initialized with dcgmInit or the embedded host engine was not running.
- ▶ DCGM\_ST\_BADPARAM if an invalid parameter was provided
- ▶ DCGM\_ST\_INIT\_ERROR if an error occurred while trying to start the host engine.

## Description

Stop the embedded host engine within this process that was started with dcgmStartEmbedded

**dcgmReturn\_t dcgmConnect (char \*ipAddress, dcgmHandle\_t \*pDcgmHandle)**

## Parameters

### ipAddress

IN : Valid IP address for the remote host engine to connect to. If ipAddress is specified as x.x.x.x it will attempt to connect to the default port specified by DCGM\_HE\_PORT\_NUMBER If ipAddress is specified as x.x.x.x:yyyy it will attempt to connect to the port specified by yyyy

### pDcgmHandle

OUT : DCGM Handle of the remote host engine

## Returns

- ▶ DCGM\_ST\_OK if we successfully connected to the remote host engine
- ▶ DCGM\_ST\_CONNECTION\_NOT\_VALID if the remote host engine could not be reached
- ▶ DCGM\_ST\_UNINITIALIZED if DCGM has not been initialized with dcgmInit.
- ▶ DCGM\_ST\_BADPARAM if pDcgmHandle is NULL or ipAddress is invalid
- ▶ DCGM\_ST\_INIT\_ERROR if DCGM encountered an error while initializing the remote client library
- ▶ DCGM\_ST\_UNINITIALIZED if DCGM has not been initialized with dcgmInit

## Description

This method is used to connect to a stand-alone host engine process. Remote host engines are started by running the nv-hostengine command.

NOTE: dcgmConnect\_v2 provides additional connection options.

```
dcgmReturn_t dcgmConnect_v2 (char *ipAddress,
dcgmConnectV2Params_t *connectParams, dcgmHandle_t
*pDcgmHandle)
```

### Parameters

#### ipAddress

IN : Valid IP address for the remote host engine to connect to. If ipAddress is specified as x.x.x.x it will attempt to connect to the default port specified by DCGM\_HE\_PORT\_NUMBER If ipAddress is specified as x.x.x.x:yyyy it will attempt to connect to the port specified by yyyy

#### connectParams

IN : Additional connection parameters. See [dcgmConnectV2Params\\_t](#) for details.

#### pDcgmHandle

OUT : DCGM Handle of the remote host engine

### Returns

- ▶ DCGM\_ST\_OK if we successfully connected to the remote host engine
- ▶ DCGM\_ST\_CONNECTION\_NOT\_VALID if the remote host engine could not be reached
- ▶ DCGM\_ST\_UNINITIALIZED if DCGM has not been initialized with dcgmInit.
- ▶ DCGM\_ST\_BADPARAM if pDcgmHandle is NULL or ipAddress is invalid
- ▶ DCGM\_ST\_INIT\_ERROR if DCGM encountered an error while initializing the remote client library
- ▶ DCGM\_ST\_UNINITIALIZED if DCGM has not been initialized with dcgmInit

### Description

This method is used to connect to a stand-alone host engine process. Remote host engines are started by running the nv-hostengine command.

[dcgmReturn\\_t dcgmDisconnect \(dcgmHandle\\_t pDcgmHandle\)](#)

### Parameters

#### pDcgmHandle

IN: DCGM Handle that came from dcgmConnect

### Returns

- ▶ DCGM\_ST\_OK if we successfully disconnected from the host engine
- ▶ DCGM\_ST\_UNINITIALIZED if DCGM has not been initialized with dcgmInit
- ▶ DCGM\_ST\_BADPARAM if pDcgmHandle is not a valid DCGM handle

- ▶ DCGM\_ST\_GENERIC\_ERROR if an unspecified internal error occurred

#### Description

This method is used to disconnect from a stand-alone host engine process.

### 2.1.2. Auxilary information about DCGM engine.

Administrative

Describes APIs to get generic information about the DCGM Engine.

**dcgmReturn\_t dcgmVersionInfo (dcgmVersionInfo\_t \*pVersionInfo)**

#### Parameters

##### **pVersionInfo**

OUT : Build environment information

#### Returns

- ▶ DCGM\_ST\_OK if build information is sucessfully obtained
- ▶ DCGM\_ST\_BADPARAM if pVersionInfo is null
- ▶ DCGM\_ST\_VER\_MISMATCH if the expected and provided versions of dcgmVersionInfo\_t do not match

#### Description

This method is used to return information about the build environment where DCGM was built.

## 2.2. System

This chapter describes the APIs used to identify set of GPUs on the node, grouping functions to provide mechanism to operate on a group of GPUs, and status management APIs in order to get individual statuses for each operation. The APIs in System module can be broken down into following categories:

### Discovery

### Grouping

### Field Grouping

## Status handling

### 2.2.1. Discovery

System

The following APIs are used to discover GPUs and their attributes on a Node.

**dcgmReturn\_t dcgmGetAllDevices (dcgmHandle\_t pDcgmHandle, unsigned int gpuidList, int \*count)**

#### Parameters

##### **pDcgmHandle**

IN : DCGM Handle

##### **gpuidList**

OUT : Array reference to fill GPU Ids present on the system.

##### **count**

OUT : Number of GPUs returned in gpuidList.

#### Returns

- ▶ DCGM\_ST\_OK if the call was successful.
- ▶ DCGM\_ST\_BADPARAM if gpuidList or count were not valid.

#### Description

This method is used to get identifiers corresponding to all the devices on the system. The identifier represents DCGM GPU Id corresponding to each GPU on the system and is immutable during the lifespan of the engine. The list should be queried again if the engine is restarted.

The GPUs returned from this function include gpuids of GPUs that are not supported by DCGM. To only get gpuids of GPUs that are supported by DCGM, use dcgmGetAllSupportedDevices().

**dcgmReturn\_t dcgmGetAllSupportedDevices (dcgmHandle\_t pDcgmHandle, unsigned int gpuidList, int \*count)**

#### Parameters

##### **pDcgmHandle**

IN : DCGM Handle

##### **gpuidList**

OUT : Array reference to fill GPU Ids present on the system.

**count**

OUT : Number of GPUs returned in gpuIdList.

**Returns**

- ▶ DCGM\_ST\_OK if the call was successful.
- ▶ DCGM\_ST\_BADPARAM if gpuIdList or count were not valid.

**Description**

This method is used to get identifiers corresponding to all the DCGM-supported devices on the system. The identifier represents DCGM GPU Id corresponding to each GPU on the system and is immutable during the lifespan of the engine. The list should be queried again if the engine is restarted.

The GPUs returned from this function ONLY includes gpuIds of GPUs that are supported by DCGM. To get gpuIds of all GPUs in the system, use [dcgmGetAllDevices\(\)](#).

**dcgmReturn\_t dcgmGetDeviceAttributes (dcgmHandle\_t pDcgmHandle, unsigned int gpuid, dcgmDeviceAttributes\_t \*pDcgmAttr)**

**Parameters****pDcgmHandle**

IN : DCGM Handle

**gpuid**

IN : GPU Id corresponding to which the attributes should be fetched

**pDcgmAttr**

IN/OUT : Device attributes corresponding to gpuid. pDcgmAttr->version should be set to [dcgmDeviceAttributes\\_version](#) before this call.

**Returns**

- ▶ DCGM\_ST\_OK if the call was successful.
- ▶ DCGM\_ST\_VER\_MISMATCH if pDcgmAttr->version is not set or is invalid.

**Description**

Gets device attributes corresponding to the gpuid. If operation is not successful for any of the requested fields then the field is populated with one of DCGM\_BLANK\_VALUES defined in [dcgm\\_structs.h](#).

```
dcgmReturn_t dcgmGetEntityGroupEntities (dcgmHandle_t
dcgmHandle, dcgm_field_entity_group_t entityGroup,
dcgm_field_eid_t *entities, int *numEntities, unsigned int flags)
```

### Parameters

#### **dcgmHandle**

IN: DCGM Handle

#### **entityGroup**

IN: Entity group to list entities of

#### **entities**

OUT: Array of entities for entityGroup

#### **numEntities**

IN/OUT: Upon calling, this should be the number of entities that entityList[] can hold.

Upon return, this will contain the number of entities actually saved to entityList.

#### **flags**

IN: Flags to modify the behavior of this request. See DCGM\_GEGE\_FLAG\_\* defines in dcgm\_structs.h

### Returns

- ▶ DCGM\_ST\_OK if the call was successful.
- ▶ DCGM\_ST\_INSUFFICIENT\_SIZE if numEntities was not large enough to hold the number of entities in the entityGroup. numEntities will contain the capacity needed to complete this request successfully.
- ▶ DCGM\_ST\_NOT\_SUPPORTED if the given entityGroup does not support enumeration.
- ▶ DCGM\_ST\_BADPARAM if any parameter is invalid

### Description

Gets the list of entities that exist for a given entity group. This API can be used in place of [dcgmGetAllDevices](#).

```
dcgmReturn_t dcgmGetNvLinkLinkStatus (dcgmHandle_t
dcgmHandle, dcgmNvLinkStatus_v1 *linkStatus)
```

### Parameters

#### **dcgmHandle**

IN: DCGM Handle

**linkStatus**

OUT: Structure in which to store NvLink link statuses. .version should be set to dcgmNvLinkStatus\_version1 before calling this.

**Returns**

- ▶ DCGM\_ST\_OK if the call was successful.
- ▶ DCGM\_ST\_NOT\_SUPPORTED if the given entityGroup does not support enumeration.
- ▶ DCGM\_ST\_BADPARAM if any parameter is invalid

**Description**

Get the NvLink link status for every NvLink in this system. This includes the NvLinks of both GPUs and NvSwitches. Note that only NvSwitches and GPUs that are visible to the current environment will be returned in this structure.

## 2.2.2. Grouping

### System

The following APIs are used for group management. The user can create a group of entities and perform an operation on a group of entities. If grouping is not needed and the user wishes to run commands on all GPUs seen by DCGM then the user can use DCGM\_GROUP\_ALL\_GPUS or DCGM\_GROUP\_ALL\_NVSWITCHES in place of group IDs when needed.

**dcgmReturn\_t dcgmGroupCreate (dcgmHandle\_t pDcgmHandle,  
dcgmGroupType\_t type, char \*groupName, dcgmGpuGrp\_t  
\*pDcgmGrpId)**

**Parameters****pDcgmHandle**

IN : DCGM Handle

**type**

IN : Type of Entity Group to be formed

**groupName**

IN : Desired name of the GPU group specified as NULL terminated C string

**pDcgmGrpId**

OUT : Reference to group ID

**Returns**

- ▶ DCGM\_ST\_OK if the group has been created

- ▶ DCGM\_ST\_BADPARAM if any of type, groupName, length or pDcgmGrpId is invalid
- ▶ DCGM\_ST\_MAX\_LIMIT if number of groups on the system has reached the max limit DCGM\_MAX\_NUM\_GROUPS
- ▶ DCGM\_ST\_INIT\_ERROR if the library has not been successfully initialized

### Description

Used to create a entity group handle which can store one or more entity Ids as an opaque handle returned in pDcgmGrpId. Instead of executing an operation separately for each entity, the DCGM group enables the user to execute same operation on all the entities present in the group as a single API call.

To create the group with all the entities present on the system, the type field should be specified as DCGM\_GROUP\_DEFAULT or DCGM\_GROUP\_ALL\_NVSWITCHES. To create an empty group, the type field should be specified as DCGM\_GROUP\_EMPTY. The empty group can be updated with the desired set of entities using the APIs [dcgmGroupAddDevice](#), [dcgmGroupAddEntity](#), [dcgmGroupRemoveDevice](#), and [dcgmGroupRemoveEntity](#).

**dcgmReturn\_t dcgmGroupDestroy (dcgmHandle\_t pDcgmHandle, dcgmGpuGrp\_t groupId)**

### Parameters

**pDcgmHandle**  
IN : DCGM Handle  
**groupId**  
IN : Group ID

### Returns

- ▶ DCGM\_ST\_OK if the group has been destroyed
- ▶ DCGM\_ST\_BADPARAM if groupId is invalid
- ▶ DCGM\_ST\_INIT\_ERROR if the library has not been successfully initialized
- ▶ DCGM\_ST\_NOT\_CONFIGURED if entry corresponding to the group does not exists

### Description

Used to destroy a group represented by groupId. Since DCGM group is a logical grouping of entities, the properties applied on the group stay intact for the individual entities even after the group is destroyed.

**dcgmReturn\_t dcgmGroupAddDevice (dcgmHandle\_t pDcgmHandle, dcgmGpuGrp\_t groupId, unsigned int gpuld)**

#### Parameters

##### **pDcgmHandle**

IN : DCGM Handle

##### **groupId**

IN : Group Id to which device should be added

##### **gpuId**

IN : DCGM GPU Id

#### Returns

- ▶ DCGM\_ST\_OK if the GPU Id has been successfully added to the group
- ▶ DCGM\_ST\_INIT\_ERROR if the library has not been successfully initialized
- ▶ DCGM\_ST\_NOT\_CONFIGURED if entry corresponding to the group (groupId) does not exists
- ▶ DCGM\_ST\_BADPARAM if gpuId is invalid or already part of the specified group

#### Description

Used to add specified GPU Id to the group represented by groupId.

**dcgmReturn\_t dcgmGroupAddEntity (dcgmHandle\_t pDcgmHandle, dcgmGpuGrp\_t groupId, dcgm\_field\_entity\_group\_t entityGroupId, dcgm\_field\_eid\_t entityId)**

#### Parameters

##### **pDcgmHandle**

IN : DCGM Handle

##### **groupId**

IN : Group Id to which device should be added

##### **entityGroupId**

IN : Entity group that entityId belongs to

##### **entityId**

IN : DCGM entityId

#### Returns

- ▶ DCGM\_ST\_OK if the entity has been successfully added to the group
- ▶ DCGM\_ST\_INIT\_ERROR if the library has not been successfully initialized

- ▶ DCGM\_ST\_NOT\_CONFIGURED if entry corresponding to the group (groupId) does not exists
- ▶ DCGM\_ST\_BADPARAM if entityId is invalid or already part of the specified group

### Description

Used to add specified entity to the group represented by groupId.

**dcgmReturn\_t dcgmGroupRemoveDevice (dcgmHandle\_t pDcgmHandle, dcgmGpuGrp\_t groupId, unsigned int gpuld)**

### Parameters

#### pDcgmHandle

IN : DCGM Handle

#### groupId

IN : Group ID from which device should be removed

#### gpuld

IN : DCGM GPU Id

### Returns

- ▶ DCGM\_ST\_OK if the GPU Id has been successfully removed from the group
- ▶ DCGM\_ST\_INIT\_ERROR if the library has not been successfully initialized
- ▶ DCGM\_ST\_NOT\_CONFIGURED if entry corresponding to the group (groupId) does not exists
- ▶ DCGM\_ST\_BADPARAM if gpuld is invalid or not part of the specified group

### Description

Used to remove specified GPU Id from the group represented by groupId.

**dcgmReturn\_t dcgmGroupRemoveEntity (dcgmHandle\_t pDcgmHandle, dcgmGpuGrp\_t groupId, dcgm\_field\_entity\_group\_t entityGroupId, dcgm\_field\_eid\_t entityId)**

### Parameters

#### pDcgmHandle

IN : DCGM Handle

#### groupId

IN : Group ID from which device should be removed

#### entityGroupId

IN : Entity group that entityId belongs to

**entityId**

IN : DCGM entityId

**Returns**

- ▶ DCGM\_ST\_OK if the entity has been successfully removed from the group
- ▶ DCGM\_ST\_INIT\_ERROR if the library has not been successfully initialized
- ▶ DCGM\_ST\_NOT\_CONFIGURED if entry corresponding to the group (groupId) does not exists
- ▶ DCGM\_ST\_BADPARAM if entityId is invalid or not part of the specified group

**Description**

Used to remove specified entity from the group represented by groupId.

**dcgmReturn\_t dcgmGroupGetInfo (dcgmHandle\_t pDcgmHandle, dcgmGpuGrp\_t groupId, dcgmGroupInfo\_t \*pDcgmGroupInfo)**

**Parameters****pDcgmHandle**

IN : DCGM Handle

**groupId**

IN : Group ID for which information to be fetched

**pDcgmGroupInfo**

OUT : Group Information

**Returns**

- ▶ DCGM\_ST\_OK if the group info is successfully received.
- ▶ DCGM\_ST\_BADPARAM if any of groupId or pDcgmGroupInfo is invalid.
- ▶ DCGM\_ST\_INIT\_ERROR if the library has not been successfully initialized.
- ▶ DCGM\_ST\_MAX\_LIMIT if the group does not contain the GPU
- ▶ DCGM\_ST\_NOT\_CONFIGURED if entry corresponding to the group (groupId) does not exists

**Description**

Used to get information corresponding to the group represented by groupId. The information returned in pDcgmGroupInfo consists of group name, and the list of entities present in the group.

**dcgmReturn\_t dcgmGroupGetAllIds (dcgmHandle\_t pDcgmHandle, dcgmGpuGrp\_t groupIdList, unsigned int \*count)**

#### Parameters

##### **pDcgmHandle**

IN : DCGM Handle

##### **groupIdList**

OUT : List of Group Ids

##### **count**

OUT : The number of Group ids in the list

#### Returns

- ▶ DCGM\_ST\_OK if the ids of the groups were successfully retrieved
- ▶ DCGM\_ST\_BADPARAM if either of the groupIdList or count is null
- ▶ DCGM\_ST\_GENERIC\_ERROR if an unknown error has occurred

#### Description

Used to get the Ids of all groups of entities. The information returned is a list of group ids in groupIdList as well as a count of how many ids there are in count. Please allocate enough memory for groupIdList. Memory of size MAX\_NUM\_GROUPS should be allocated for groupIdList.

### 2.2.3. Field Grouping

#### System

The following APIs are used for field group management. The user can create a group of fields and perform an operation on a group of fields at once.

**dcgmReturn\_t dcgmFieldGroupCreate (dcgmHandle\_t dcgmHandle, int numFieldIds, unsigned short \*fieldIds, char \*fieldGroupName, dcgmFieldGrp\_t \*dcgmFieldGroupId)**

#### Parameters

##### **dcgmHandle**

IN: DCGM handle

##### **numFieldIds**

IN: Number of field IDs that are being provided in fieldIds[]. Must be between 1 and DCGM\_MAX\_FIELD\_IDS\_PER\_FIELD\_GROUP.

**fieldIds**

IN: Field IDs to be added to the newly-created field group

**fieldGroupName**

IN: Unique name for this group of fields. This must not be the same as any existing field groups.

**dcgmFieldGroupId**

OUT: Handle to the newly-created field group

**Returns**

- ▶ DCGM\_ST\_OK if the field group was successfully created.
- ▶ DCGM\_ST\_BADPARAM if any parameters were bad
- ▶ DCGM\_ST\_INIT\_ERROR if the library has not been successfully initialized.
- ▶ DCGM\_ST\_MAX\_LIMIT if too many field groups already exist

**Description**

Used to create a group of fields and return the handle in `dcgmFieldGroupId`

`dcgmReturn_t dcgmFieldGroupDestroy (dcgmHandle_t dcgmHandle, dcgmFieldGrp_t dcgmFieldGroupId)`

**Parameters****dcgmHandle**

IN: DCGM handle

**dcgmFieldGroupId**

IN: Field group to remove

**Returns**

- ▶ DCGM\_ST\_OK if the field group was successfully removed
- ▶ DCGM\_ST\_BADPARAM if any parameters were bad
- ▶ DCGM\_ST\_INIT\_ERROR if the library has not been successfully initialized.

**Description**

Used to remove a field group that was created with `dcgmFieldGroupCreate`

**dcgmReturn\_t dcgmFieldGroupGetInfo (dcgmHandle\_t dcgmHandle, dcgmFieldGroupInfo\_t \*fieldGroupInfo)**

#### Parameters

##### **dcgmHandle**

IN: DCGM handle

##### **fieldGroupInfo**

IN/OUT: Info about all of the field groups that exist. .version should be set to [dcgmFieldGroupInfo\\_version](#) before this call .fieldGroupId should contain the fieldGroupId you are interested in querying information for.

#### Returns

- ▶ DCGM\_ST\_OK if the field group info was returned successfully
- ▶ DCGM\_ST\_BADPARAM if any parameters were bad
- ▶ DCGM\_ST\_INIT\_ERROR if the library has not been successfully initialized.
- ▶ DCGM\_ST\_VER\_MISMATCH if .version is not set or is invalid.

#### Description

Used to get information about a field group that was created with [dcgmFieldGroupCreate](#).

**dcgmReturn\_t dcgmFieldGroupGetAll (dcgmHandle\_t dcgmHandle, dcgmAllFieldGroup\_t \*allGroupInfo)**

#### Parameters

##### **dcgmHandle**

IN: DCGM handle

##### **allGroupInfo**

IN/OUT: Info about all of the field groups that exist. .version should be set to [dcgmAllFieldGroup\\_version](#) before this call.

#### Returns

- ▶ DCGM\_ST\_OK if the field group info was successfully returned
- ▶ DCGM\_ST\_BADPARAM if any parameters were bad
- ▶ DCGM\_ST\_INIT\_ERROR if the library has not been successfully initialized.
- ▶ DCGM\_ST\_VER\_MISMATCH if .version is not set or is invalid.

**Description**

Used to get information about all field groups in the system.

## 2.2.4. Status handling

**System**

The following APIs are used to manage statuses for multiple operations on one or more GPUs.

### `dcmReturn_t dcgmStatusCreate (dcgmStatus_t *statusHandle)`

**Parameters****statusHandle**

OUT : Reference to handle for list of statuses

**Returns**

- ▶ DCGM\_ST\_OK if the status handle is successfully created
- ▶ DCGM\_ST\_BADPARAM if statusHandle is invalid

**Description**

Creates reference to DCGM status handler which can be used to get the statuses for multiple operations on one or more devices.

The multiple statuses are useful when the operations are performed at group level. The status handle provides a mechanism to access error attributes for the failed operations.

The number of errors stored behind the opaque handle can be accessed using the API `dcgmStatusGetCount`. The errors are accessed from the opaque handle statusHandle using the API `dcgmStatusPopError`. The user can invoke `dcgmStatusPopError` for the number of errors or until all the errors are fetched.

When the status handle is not required any further then it should be deleted using the API `dcgmStatusDestroy`.

### `dcmReturn_t dcgmStatusDestroy (dcgmStatus_t statusHandle)`

**Parameters****statusHandle**

IN : Handle to list of statuses

**Returns**

- ▶ DCGM\_ST\_OK if the status handle is successfully created
- ▶ DCGM\_ST\_BADPARAM if statusHandle is invalid

**Description**

Used to destroy status handle created using [dcgmStatusCreate](#).

**dcgmReturn\_t dcgmStatusGetCount (dcgmStatus\_t statusHandle, unsigned int \*count)**

**Parameters****statusHandle**

IN : Handle to list of statuses

**count**

OUT : Number of error entries present in the list of statuses

**Returns**

- ▶ DCGM\_ST\_OK if the error count is successfully received
- ▶ DCGM\_ST\_BADPARAM if any of statusHandle or count is invalid

**Description**

Used to get count of error entries stored inside the opaque handle statusHandle.

**dcgmReturn\_t dcgmStatusPopError (dcgmStatus\_t statusHandle, dcgmErrorInfo\_t \*pDcgmErrorInfo)**

**Parameters****statusHandle**

IN : Handle to list of statuses

**pDcgmErrorInfo**

OUT : First error from the list of statuses

**Returns**

- ▶ DCGM\_ST\_OK if the error entry is successfully fetched
- ▶ DCGM\_ST\_BADPARAM if any of statusHandle or pDcgmErrorInfo is invalid
- ▶ DCGM\_ST\_NO\_DATA if the status handle list is empty

## Description

Used to iterate through the list of errors maintained behind statusHandle. The method pops the first error from the list of DCGM statuses. In order to iterate through all the errors, the user can invoke this API for the number of errors or until all the errors are fetched.

## `dcmReturn_t dcgmStatusClear (dcgmStatus_t statusHandle)`

### Parameters

#### **statusHandle**

IN : Handle to list of statuses

### Returns

- ▶ DCGM\_ST\_OK if the errors are successfully cleared
- ▶ DCGM\_ST\_BADPARAM if statusHandle is invalid

## Description

Used to clear all the errors in the status handle created by the API `dcgmStatusCreate`. After one set of operation, the statusHandle can be cleared and reused for the next set of operation.

## 2.3. Configuration

This chapter describes the methods that handle device configuration retrieval and default settings. The APIs in Configuration module can be broken down into following categories:

### Setup and management

### Manual Invocation

#### 2.3.1. Setup and management

##### Configuration

Describes APIs to Get/Set configuration on the group of GPUs.

`dcgmReturn_t dcgmConfigSet (dcgmHandle_t pDcgmHandle,  
dcgmGpuGrp_t groupId, dcgmConfig_t *pDeviceConfig,  
dcgmStatus_t statusHandle)`

### Parameters

#### **pDcgmHandle**

IN : DCGM Handle

#### **groupId**

IN : Group ID representing collection of one or more GPUs. Look at [dcgmGroupCreate](#) for details on creating the group.

#### **pDeviceConfig**

IN : Pointer to memory to hold desired configuration to be applied for all the GPU in the group represented by groupId. The caller must populate the version field of pDeviceConfig.

#### **statusHandle**

IN/OUT : Resulting error status for multiple operations. Pass it as NULL if the detailed error information is not needed. Look at [dcgmStatusCreate](#) for details on creating status handle.

### Returns

- ▶ DCGM\_ST\_OK if the configuration has been successfully set.
- ▶ DCGM\_ST\_BADPARAM if any of groupId or pDeviceConfig is invalid.
- ▶ DCGM\_ST\_VER\_MISMATCH if pDeviceConfig has the incorrect version.
- ▶ DCGM\_ST\_GENERIC\_ERROR if an unknown error has occurred.

### Description

Used to set configuration for the group of one or more GPUs identified by groupId.

The configuration settings specified in pDeviceConfig are applied to all the GPUs in the group. Since DCGM group is a logical grouping of GPUs, the configuration settings stays intact for the individual GPUs even after the group is destroyed.

If the user wishes to ignore the configuration of one or more properties in the input pDeviceConfig then the property should be specified as one of DCGM\_INT32\_BLANK, DCGM\_INT64\_BLANK, DCGM\_FP64\_BLANK or DCGM\_STR\_BLANK based on the data type of the property to be ignored.

If any of the properties fail to be configured for any of the GPUs in the group then the API returns an error. The status handle statusHandle should be further evaluated to access error attributes for the failed operations. Please refer to status management APIs at [Status handling](#) to access the error attributes.

To find out valid supported clock values that can be passed to `dcmConfigSet`, look at the device attributes of a GPU in the group using the API `dcmGetDeviceAttributes`.

**`dcmReturn_t dcmConfigGet (dcmHandle_t pDcmHandle, dcmGpuGrp_t groupId, dcmConfigType_t type, int count, dcmConfig_t deviceConfigList, dcmStatus_t statusHandle)`**

### Parameters

#### **pDcmHandle**

IN : DCGM Handle

#### **groupId**

IN : Group ID representing collection of one or more GPUs. Look at `dcmGroupCreate` for details on creating the group.

#### **type**

IN : Type of configuration values to be fetched.

#### **count**

IN : The number of entries that deviceConfigList array can store.

#### **deviceConfigList**

OUT : Pointer to memory to hold requested configuration corresponding to all the GPUs in the group (groupId). The size of the memory must be greater than or equal to hold output information for the number of GPUs present in the group (groupId).

#### **statusHandle**

IN/OUT : Resulting error status for multiple operations. Pass it as NULL if the detailed error information is not needed. Look at `dcmStatusCreate` for details on creating status handle.

### Returns

- ▶ DCGM\_ST\_OK if the configuration has been successfully fetched.
- ▶ DCGM\_ST\_BADPARAM if any of groupId, type, count, or deviceConfigList is invalid.
- ▶ DCGM\_ST\_NOT\_CONFIGURED if the target configuration is not already set.
- ▶ DCGM\_ST\_VER\_MISMATCH if deviceConfigList has the incorrect version.
- ▶ DCGM\_ST\_GENERIC\_ERROR if an unknown error has occurred.

### Description

Used to get configuration for all the GPUs present in the group.

This API can get the most recent target or desired configuration set by `dcmConfigSet`. Set type as DCGM\_CONFIG\_TARGET\_STATE to get target configuration. The target configuration properties are maintained by DCGM and are automatically enforced after a GPU reset or reinitialization is completed.

The method can also be used to get the actual configuration state for the GPUs in the group. Set type as DCGM\_CONFIG\_CURRENT\_STATE to get the actually configuration state. Ideally, the actual configuration state will be exact same as the target configuration state.

If any of the property in the target configuration is unknown then the property value in the output is populated as one of DCGM\_INT32\_BLANK, DCGM\_INT64\_BLANK, DCGM\_FP64\_BLANK or DCGM\_STR\_BLANK based on the data type of the property.

If any of the property in the current configuration state is not supported then the property value in the output is populated as one of DCGM\_INT32\_NOT\_SUPPORTED, DCGM\_INT64\_NOT\_SUPPORTED, DCGM\_FP64\_NOT\_SUPPORTED or DCGM\_STR\_NOT\_SUPPORTED based on the data type of the property.

If any of the properties can't be fetched for any of the GPUs in the group then the API returns an error. The status handle statusHandle should be further evaluated to access error attributes for the failed operations. Please refer to status management APIs at [Status handling](#) to access the error attributes.

### 2.3.2. Manual Invocation

#### Configuration

Describes APIs used to manually enforce the desired configuration on a group of GPUs.

**dcgmReturn\_t dcgmConfigEnforce (dcgmHandle\_t pDcgmHandle, dcgmGpuGrp\_t groupId, dcgmStatus\_t statusHandle)**

#### Parameters

##### pDcgmHandle

IN : DCGM Handle

##### groupId

IN : Group ID representing collection of one or more GPUs. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs.

##### statusHandle

IN/OUT : Resulting error status for multiple operations. Pass it as NULL if the detailed error information is not needed. Look at [dcgmStatusCreate](#) for details on creating status handle.

#### Returns

- ▶ DCGM\_ST\_OK if the configuration has been successfully enforced.
- ▶ DCGM\_ST\_BADPARAM if groupId is invalid.
- ▶ DCGM\_ST\_NOT\_CONFIGURED if the target configuration is not already set.

- ▶ DCGM\_ST\_GENERIC\_ERROR if an unknown error has occurred.

### Description

Used to enforce previously set configuration for all the GPUs present in the group.

This API provides a mechanism to the users to manually enforce the configuration at any point of time. The configuration can only be enforced if it's already configured using the API [dcgmConfigSet](#).

If any of the properties can't be enforced for any of the GPUs in the group then the API returns an error. The status handle statusHandle should be further evaluated to access error attributes for the failed operations. Please refer to status management APIs at [Status handling](#) to access the error attributes.

## 2.4. Field APIs

These APIs are responsible for watching, unwatching, and updating specific fields as defined by DCGM\_FI\_\*

```
dcgmReturn_t dcgmWatchFields (dcgmHandle_t pDcmHandle, dcgmGpuGrp_t groupId, dcgmFieldGrp_t fieldGroupId, long long updateFreq, double maxKeepAge, int maxKeepSamples)
```

### Parameters

#### **pDcmHandle**

IN: DCGM Handle

#### **groupId**

IN: Group ID representing collection of one or more entities. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs or DCGM\_GROUP\_ALL\_NVSWITCHES to perform the operation on all NvSwitches.

#### **fieldGroupId**

IN: Fields to watch.

#### **updateFreq**

IN: How often to update this field in usec

#### **maxKeepAge**

IN: How long to keep data for this field in seconds

#### **maxKeepSamples**

IN: Maximum number of samples to keep. 0=no limit

**Returns**

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if a parameter is invalid

**Description**

Request that DCGM start recording updates for a given field collection.

Note that the first update of the field will not occur until the next field update cycle. To force a field update cycle, call `dcmUpdateAllFields(1)`.

## `dcmReturn_t dcgmUnwatchFields (dcgmHandle_t pDcgmHandle, dcgmGpuGrp_t groupId, dcgmFieldGrp_t fieldGroupId)`

**Parameters****pDcgmHandle**

IN: DCGM Handle

**groupId**

IN: Group ID representing collection of one or more entities. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs or DCGM\_GROUP\_ALL\_NVSWITCHES to to perform the operation on all NvSwitches.

**fieldGroupId**

IN: Fields to unwatch.

**Returns**

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if a parameter is invalid

**Description**

Request that DCGM stop recording updates for a given field collection.

## `dcmReturn_t dcgmGetValuesSince (dcgmHandle_t pDcgmHandle, dcgmGpuGrp_t groupId, dcgmFieldGrp_t fieldGroupId, long long sinceTimestamp, long long`

`*nextSinceTimestamp, dcgmFieldValueEnumeration_f  
enumCB, void *userData)`

#### Parameters

##### **pDcgmHandle**

IN: DCGM Handle

##### **groupId**

IN: Group ID representing collection of one or more GPUs. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs.

##### **fieldGroupId**

IN: Fields to return data for

##### **sinceTimestamp**

IN: Timestamp to request values since in usec since 1970. This will be returned in nextSinceTimestamp for subsequent calls 0 = request all data

##### **nextSinceTimestamp**

OUT: Timestamp to use for sinceTimestamp on next call to this function

##### **enumCB**

IN: Callback to invoke for every field value update. Note that multiple updates can be returned in each invocation

##### **userData**

IN: User data pointer to pass to the userData field of enumCB.

#### Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_NOT\_SUPPORTED if one of the entities was from a non-GPU type
- ▶ DCGM\_ST\_BADPARAM if a parameter is invalid

#### Description

Request updates for all field values that have updated since a given timestamp

This version only works with GPU entities. Use [dcgmGetValuesSince\\_v2](#) for entity groups containing NvSwitches.

**dcgmReturn\_t dcgmGetValuesSince\_v2**  
**(dcgmHandle\_t pDcgmHandle, dcgmGpuGrp\_t**  
**groupId, dcgmFieldGrp\_t fieldGroupId, long long**  
**sinceTimestamp, long long \*nextSinceTimestamp,**

## `dcmFieldValueEntityEnumeration_f enumCB, void *userData)`

### Parameters

#### **pDcgmHandle**

IN: DCGM Handle

#### **groupId**

IN: Group ID representing collection of one or more entities. Look at [dcmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs or DCGM\_GROUP\_ALL\_NVSWITCHES to perform the operation on all NvSwitches.

#### **fieldGroupId**

IN: Fields to return data for

#### **sinceTimestamp**

IN: Timestamp to request values since in usec since 1970. This will be returned in nextSinceTimestamp for subsequent calls 0 = request all data

#### **nextSinceTimestamp**

OUT: Timestamp to use for sinceTimestamp on next call to this function

#### **enumCB**

IN: Callback to invoke for every field value update. Note that multiple updates can be returned in each invocation

#### **userData**

IN: User data pointer to pass to the userData field of enumCB.

### Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if a parameter is invalid

### Description

Request updates for all field values that have updated since a given timestamp

This version works with non-GPU entities like NvSwitches

## `dcmReturn_t dcmGetLatestValues (dcgmHandle_t pDcgmHandle, dcgmGpuGrp_t groupId, dcgmFieldGrp_t`

**fieldGroupId, dcgmFieldValueEnumeration\_f enumCB,  
void \*userData)**

#### Parameters

##### **pDcgmHandle**

IN: DCGM Handle

##### **groupId**

IN: Group ID representing collection of one or more GPUs. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs.

##### **fieldGroupId**

IN: Fields to return data for.

##### **enumCB**

IN: Callback to invoke for every field value update. Note that multiple updates can be returned in each invocation

##### **userData**

IN: User data pointer to pass to the userData field of enumCB.

#### Description

Request latest cached field value for a field value collection

This version only works with GPU entities. Use [dcgmGetLatestValues\\_v2](#) for entity groups containing NvSwitches.

- ▶ `DCGM_ST_OK` if the call was successful
- ▶ `DCGM_ST_NOT_SUPPORTED` if one of the entities was from a non-GPU type
- ▶ `DCGM_ST_BADPARAM` if a parameter is invalid

**dcgmReturn\_t dcgmGetLatestValues\_v2 (dcgmHandle\_t  
pDcgmHandle, dcgmGpuGrp\_t groupId, dcgmFieldGrp\_t  
fieldGroupId, dcgmFieldValueEntityEnumeration\_f  
enumCB, void \*userData)**

#### Parameters

##### **pDcgmHandle**

IN: DCGM Handle

##### **groupId**

IN: Group ID representing collection of one or more entities. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs or DCGM\_GROUP\_ALL\_NVSWITCHES to perform the operation on all NvSwitches.

**fieldGroupId**

IN: Fields to return data for.

**enumCB**

IN: Callback to invoke for every field value update. Note that multiple updates can be returned in each invocation

**userData**

IN: User data pointer to pass to the userData field of enumCB.

**Description**

Request latest cached field value for a field value collection

This version works with non-GPU entities like NvSwitches

- ▶ **DCGM\_ST\_OK** if the call was successful
- ▶ **DCGM\_ST\_NOT\_SUPPORTED** if one of the entities was from a non-GPU type
- ▶ **DCGM\_ST\_BADPARAM** if a parameter is invalid

**dcmReturn\_t dcgmGetLatestValuesForFields**

**(dcgmHandle\_t pDcgmHandle, int gpuld, unsigned short fields, unsigned int count, dcgmFieldValue\_v1 values)**

**Parameters****pDcgmHandle**

IN: DCGM Handle

**gpuId**

IN: Gpu ID representing the GPU for which the fields are being requested.

**fields**

IN: Field IDs to return data for. See the definitions in dcgm\_fields.h that start with **DCGM\_FI\_**.

**count**

IN: Number of field IDs in fields[] array.

**values**

OUT: Latest field values for the fields in fields[].

**Description**

Request latest cached field value for a GPU

**dcmReturn\_t dcgmEntityGetLatestValues**

**(dcgmHandle\_t pDcgmHandle,**

**dcgm\_field\_entity\_group\_t entityGroup, int**

**entityId, unsigned short fields, unsigned int count,  
dcgmFieldValue\_v1 values)**

#### Parameters

##### **pDcgmHandle**

IN: DCGM Handle

##### **entityGroup**

IN: entity\_group\_t (e.g. switch)

##### **entityId**

IN: entity ID representing the rntity for which the fields are being requested.

##### **fields**

IN: Field IDs to return data for. See the definitions in dcgm\_fields.h that start with DCGM\_FI\_.

##### **count**

IN: Number of field IDs in fields[] array.

##### **values**

OUT: Latest field values for the fields in fields[].

#### Description

Request latest cached field value for a group of fields for a specific entity

**dcgmReturn\_t dcgmEntitiesGetLatestValues  
(dcgmHandle\_t pDcgmHandle, dcgmGroupEntityPair\_t  
entities, unsigned int entityCount, unsigned short  
fields, unsigned int fieldCount, unsigned int flags,  
dcgmFieldValue\_v2 values)**

#### Parameters

##### **pDcgmHandle**

IN: DCGM Handle

##### **entities**

IN: List of entities to get values for

##### **entityCount**

IN: Number of entries in entities[]

##### **fields**

IN: Field IDs to return data for. See the definitions in dcgm\_fields.h that start with DCGM\_FI\_.

##### **fieldCount**

IN: Number of field IDs in fields[] array.

**flags**

IN: Optional flags that affect how this request is processed. Pass `DCGM_FV_FLAG_LIVE_DATA` here to retrieve a live driver value rather than a cached value. See that flag's documentation for caveats.

**values**

OUT: Latest field values for the fields requested. This must be able to hold `entityCount * fieldCount` field value records.

**Description**

Request the latest cached or live field value for a list of fields for a group of entities

Note: The returned entities are not guaranteed to be in any order. Reordering can occur internally in order to optimize calls to the NVIDIA driver.

## 2.5. Process Statistics

Describes APIs to investigate statistics such as accounting, performance and errors during the lifetime of a GPU process

`dcmReturn_t dcgmWatchPidFields (dcgmHandle_t pDcgmHandle, dcgmGpuGrp_t groupId, long long updateFreq, double maxKeepAge, int maxKeepSamples)`

**Parameters****pDcgmHandle**

IN: DCGM Handle

**groupId**

IN: Group ID representing collection of one or more GPUs. Look at `dcgmGroupCreate` for details on creating the group. Alternatively, pass in the group id as `DCGM_GROUP_ALL_GPUS` to perform operation on all the GPUs.

**updateFreq**

IN: How often to update this field in usec

**maxKeepAge**

IN: How long to keep data for this field in seconds

**maxKeepSamples**

IN: Maximum number of samples to keep. 0=no limit

**Returns**

- ▶ `DCGM_ST_OK` if the call was successful
- ▶ `DCGM_ST_BADPARAM` if a parameter is invalid

- ▶ DCGM\_ST\_REQUIRES\_ROOT if the host engine is being run as non-root, and accounting mode could not be enabled (requires root). Run "nvidia-smi -am 1" as root on the node before starting DCGM to fix this.

### Description

Request that DCGM start recording stats for fields that can be queried with [dcgmGetPidInfo\(\)](#).

Note that the first update of the field will not occur until the next field update cycle. To force a field update cycle, call `dcgmUpdateAllFields(1)`.

**`dcgmReturn_t dcgmGetPidInfo (dcgmHandle_t pDcgmHandle, dcgmGpuGrp_t groupId, dcgmPidInfo_t *pidInfo)`**

### Parameters

#### **pDcgmHandle**

IN: DCGM Handle

#### **groupId**

IN: Group ID representing collection of one or more GPUs. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs.

#### **pidInfo**

IN/OUT: Structure to return information about pid in. `pidInfo->pid` must be set to the pid in question. `pidInfo->version` should be set to `dcgmPidInfo_version`.

### Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_NO\_DATA if the PID did not run on any GPU

### Description

Get information about all GPUs while the provided pid was running

In order for this request to work, you must first call [dcgmWatchPidFields\(\)](#) to make sure that DCGM is watching the appropriate field IDs that will be populated in `pidInfo`

## 2.6. Job Statistics

The client can invoke DCGM APIs to start and stop collecting the stats at the process boundaries (during prologue and epilogue). This will enable DCGM to monitor all the PIDs while the job is in progress, and provide a summary of active processes and resource usage during the window of interest.

**dcgmReturn\_t dcgmWatchJobFields (dcgmHandle\_t pDcgmHandle, dcgmGpuGrp\_t groupId, long long updateFreq, double maxKeepAge, int maxKeepSamples)**

### Parameters

#### pDcgmHandle

IN: DCGM Handle

#### groupId

IN: Group ID representing collection of one or more GPUs. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs.

#### updateFreq

IN: How often to update this field in usec

#### maxKeepAge

IN: How long to keep data for this field in seconds

#### maxKeepSamples

IN: Maximum number of samples to keep. 0=no limit

### Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if a parameter is invalid
- ▶ DCGM\_ST\_REQUIRES\_ROOT if the host engine is being run as non-root, and accounting mode could not be enabled (requires root). Run "nvidia-smi -am 1" as root on the node before starting DCGM to fix this.

### Description

Request that DCGM start recording stats for fields that are queried with [dcgmJobGetStats\(\)](#)

Note that the first update of the field will not occur until the next field update cycle. To force a field update cycle, call [dcgmUpdateAllFields\(1\)](#).

## `dcmReturn_t dcgmJobStartStats (dcgmHandle_t pDcgmHandle, dcgmGpuGrp_t groupId, char jobId)`

### Parameters

#### **pDcgmHandle**

IN : DCGM Handle

#### **groupId**

IN : Group ID representing collection of one or more GPUs. Look at `dcgmGroupCreate` for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs.

#### **jobId**

IN : User provided string to represent the job

### Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if a parameter is invalid
- ▶ DCGM\_ST\_DUPLICATE\_KEY if the specified jobId is already in use

### Description

This API is used by the client to notify DCGM about the job to be started. Should be invoked as part of job prologue

## `dcmReturn_t dcgmJobStopStats (dcgmHandle_t pDcgmHandle, char jobId)`

### Parameters

#### **pDcgmHandle**

IN : DCGM Handle

#### **jobId**

IN : User provided string to represent the job

### Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if a parameter is invalid
- ▶ DCGM\_ST\_NO\_DATA if jobId is not a valid job identifier.

## Description

This API is used by the clients to notify DCGM to stop collecting stats for the job represented by job id. Should be invoked as part of job epilogue. The job Id remains available to view the stats at any point but cannot be used to start a new job. You must call `dcmWatchJobFields()` before this call to enable watching of job

**dcmReturn\_t dcmJobGetStats (dcgmHandle\_t pDcmHandle, char jobId, dcgmJobInfo\_t \*pJobInfo)**

## Parameters

### pDcmHandle

IN : DCGM Handle

### jobId

IN : User provided string to represent the job

### pJobInfo

IN/OUT : Structure to return information about the job. .version should be set to `dcmJobInfo_version` before this call.

## Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if a parameter is invalid
- ▶ DCGM\_ST\_NO\_DATA if jobId is not a valid job identifier.
- ▶ DCGM\_ST\_VER\_MISMATCH if .version is not set or is invalid.

## Description

Get stats for the job identified by DCGM generated job id. The stats can be retrieved at any point when the job is in process. If you want to reuse this jobId, call `dcmJobRemove` after this call.

**dcmReturn\_t dcmJobRemove (dcgmHandle\_t pDcmHandle, char jobId)**

## Parameters

### pDcmHandle

IN : DCGM Handle

### jobId

IN : User provided string to represent the job

**Returns**

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if a parameter is invalid
- ▶ DCGM\_ST\_NO\_DATA if jobId is not a valid job identifier.

**Description**

This API tells DCGM to stop tracking the job given by jobId. After this call, you will no longer be able to call [dcgmJobGetStats\(\)](#) on this jobId. However, you will be able to reuse jobId after this call.

## **dcgmReturn\_t dcgmJobRemoveAll (dcgmHandle\_t pDcgmHandle)**

**Parameters****pDcgmHandle**

IN : DCGM Handle

**Returns**

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if a parameter is invalid

**Description**

This API tells DCGM to stop tracking all jobs. After this call, you will no longer be able to call [dcgmJobGetStats\(\)](#) on any jobs until you call dcgmJobStartStats again. You will be able to reuse any previously-used jobIds after this call.

## 2.7. Health Monitor

This chapter describes the methods that handle the GPU health monitor.

## dcgmReturn\_t dcgmHealthSet (dcgmHandle\_t pDcgmHandle, dcgmGpuGrp\_t groupId, dcgmHealthSystems\_t systems)

### Parameters

#### pDcgmHandle

IN: DCGM Handle

#### groupId

IN: Group ID representing collection of one or more entities. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs or DCGM\_GROUP\_ALL\_NVSWITCHES to perform operation on all the NvSwitches.

#### systems

IN: An enum representing systems that should be enabled for health checks logically OR'd together. Refer to [dcgmHealthSystems\\_t](#) for details.

### Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if a parameter is invalid

### Description

Enable the DCGM health check system for the given systems defined in [dcgmHealthSystems\\_t](#)

## dcgmReturn\_t dcgmHealthGet (dcgmHandle\_t pDcgmHandle, dcgmGpuGrp\_t groupId, dcgmHealthSystems\_t \*systems)

### Parameters

#### pDcgmHandle

IN: DCGM Handle

#### groupId

IN: Group ID representing collection of one or more entities. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs or DCGM\_GROUP\_ALL\_NVSWITCHES to perform operation on all the NvSwitches.

**systems**

OUT: An integer representing the enabled systems for the given group Refer to [dcgmHealthSystems\\_t](#) for details.

**Returns**

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if a parameter is invalid

**Description**

Retrieve the current state of the DCGM health check system

**dcgmReturn\_t dcgmHealthCheck (dcgmHandle\_t pDcgmHandle, dcgmGpuGrp\_t groupId, dcgmHealthResponse\_t \*results)**

**Parameters****pDcgmHandle**

IN: DCGM Handle

**groupId**

IN: Group ID representing a collection of one or more entities. Refer to [dcgmGroupCreate](#) for details on creating a group

**results**

OUT: A reference to the dcgmHealthResponse\_t structure to populate. results->version must be set to dcgmHealthResponse\_version.

**Returns**

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if a parameter is invalid
- ▶ DCGM\_ST\_VER\_MISMATCH if results->version is not dcgmHealthResponse\_version

**Description**

Check the configured watches for any errors/failures/warnings that have occurred since the last time this check was invoked. On the first call, stateful information about all of the enabled watches within a group is created but no error results are provided. On subsequent calls, any error information will be returned.

## 2.8. Policies

This chapter describes the methods that handle system policy management and violation settings. The APIs in Policies module can be broken down into following categories:

### Setup and Management

#### Manual Invocation

##### 2.8.1. Setup and Management

###### Policies

Describes APIs for setting up policies and registering callbacks to receive notification in case specific policy condition has been violated.

`dcgmReturn_t dcgmPolicySet (dcgmHandle_t pDcgmHandle,  
dcgmGpuGrp_t groupId, dcgmPolicy_t *policy, dcgmStatus_t  
statusHandle)`

###### Parameters

###### **pDcgmHandle**

IN: DCGM Handle

###### **groupId**

IN: Group ID representing collection of one or more GPUs. Look at `dcgmGroupCreate` for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs.

###### **policy**

IN: A reference to `dcgmPolicy_t` that will be applied to all GPUs in the group.

###### **statusHandle**

IN/OUT: Resulting status for the operation. Pass it as NULL if the detailed error information is not needed. Refer to `dcgmStatusCreate` for details on creating a status handle.

###### Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if groupId or policy is invalid
- ▶ DCGM\_ST\_NOT\_SUPPORTED if any non-Tesla GPUs are part of the GPU group specified in groupId

- ▶ DCGM\_ST\_\* a different error has occurred and is stored in statusHandle. Refer to [dcgmReturn\\_t](#)

## Description

Set the current violation policy inside the policy manager. Given the conditions within the [dcgmPolicy\\_t](#) structure, if a violation has occurred, subsequent action(s) may be performed to either report or contain the failure.

This API is only supported on Tesla GPUs and will return DCGM\_ST\_NOT\_SUPPORTED if any non-Tesla GPUs are part of the GPU group specified in groupId.

```
dcgmReturn_t dcgmPolicyGet (dcgmHandle_t pDcgmHandle,
dcgmGpuGrp_t groupId, int count, dcgmPolicy_t *policy,
dcgmStatus_t statusHandle)
```

## Parameters

### **pDcgmHandle**

IN: DCGM Handle

### **groupId**

IN: Group ID representing collection of one or more GPUs. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs.

### **count**

IN: The size of the policy array. This is the maximum number of policies that will be retrieved and ultimately should correspond to the number of GPUs specified in the group.

### **policy**

OUT: A reference to [dcgmPolicy\\_t](#) that will used as storage for the current policies applied to each GPU in the group.

### **statusHandle**

IN/OUT: Resulting status for the operation. Pass it as NULL if the detailed error information for the operation is not needed. Refer to [dcgmStatusCreate](#) for details on creating a status handle.

## Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if groupId or policy is invalid
- ▶ DCGM\_ST\_\* a different error has occurred and is stored in statusHandle. Refer to [dcgmReturn\\_t](#)

## Description

Get the current violation policy inside the policy manager. Given a groupId, a number of policy structures are retrieved.

**dcgmReturn\_t dcgmPolicyRegister (dcgmHandle\_t pDcgmHandle, dcgmGpuGrp\_t groupId, dcgmPolicyCondition\_t condition, fpRecvUpdates beginCallback, fpRecvUpdates finishCallback)**

## Parameters

### pDcgmHandle

IN: DCGM Handle

### groupId

IN: Group ID representing collection of one or more GPUs. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs.

### condition

IN: The set of conditions specified as an OR'd list (see [dcgmPolicyCondition\\_t](#)) for which to register a callback function

### beginCallback

IN: A reference to a function that should be called should a violation occur. This function will be called prior to any actions specified by the policy are taken.

### finishCallback

IN: A reference to a function that should be called should a violation occur. This function will be called after any action specified by the policy are completed.

## Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if groupId, condition, is invalid, beginCallback, or finishCallback is NULL
- ▶ DCGM\_ST\_NOT\_SUPPORTED if any non-Tesla GPUs are part of the GPU group specified in groupId

## Description

Register a function to be called when a specific policy condition (see [dcgmPolicyCondition\\_t](#)) has been violated. This callback(s) will be called automatically when in DCGM\_OPERATION\_MODE\_AUTO mode and only after [dcgmPolicyTrigger](#) when in DCGM\_OPERATION\_MODE\_MANUAL mode. All callbacks are made within a separate thread.

This API is only supported on Tesla GPUs and will return DCGM\_ST\_NOT\_SUPPORTED if any non-Tesla GPUs are part of the GPU group specified in groupId.

## `dcmReturn_t dcgmPolicyUnregister (dcgmHandle_t pDcgmHandle, dcgmGpuGrp_t groupId, dcgmPolicyCondition_t condition)`

### Parameters

#### **pDcgmHandle**

IN: DCGM Handle

#### **groupId**

IN: Group ID representing collection of one or more GPUs. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs.

#### **condition**

IN: The set of conditions specified as an OR'd list (see [dcgmPolicyCondition\\_t](#)) for which to unregister a callback function

### Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if groupId, condition, is invalid or callback is NULL

### Description

Unregister a function to be called for a specific policy condition (see [dcgmPolicyCondition\\_t](#)). This function will unregister all callbacks for a given condition and handle.

## 2.8.2. Manual Invocation

### Policies

Describes APIs which can be used to perform direct actions (e.g. Perform GPU Reset, Run Health Diagnostics) on a group of GPUs.

```
dcgmReturn_t dcgmActionValidate (dcgmHandle_t pDcgmHandle,
dcgmGpuGrp_t groupId, dcgmPolicyValidation_t validate,
dcgmDiagResponse_t *response)
```

### Parameters

#### **pDcgmHandle**

IN: DCGM Handle

#### **groupId**

IN: Group ID representing collection of one or more GPUs. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs.

#### **validate**

IN: The validation to perform after the action.

#### **response**

OUT: Result of the validation process. Refer to [dcgmDiagResponse\\_t](#) for details.

### Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_NOT\_SUPPORTED if running the specified validate is not supported. This is usually due to the Tesla recommended driver not being installed on the system.
- ▶ DCGM\_ST\_BADPARAM if groupId, validate, or statusHandle is invalid
- ▶ DCGM\_ST\_GENERIC\_ERROR an internal error has occurred
- ▶ DCGM\_ST\_GROUP\_INCOMPATIBLE if groupId refers to a group of non-homogeneous GPUs. This is currently not allowed.

### Description

Inform the action manager to perform a manual validation of a group of GPUs on the system

\*\*\*\*\* DEPRECATED \*\*\*\*\*

```
dcgmReturn_t dcgmActionValidate_v2 (dcgmHandle_t pDcgmHandle,
dcgmRunDiag_t *drd, dcgmDiagResponse_t *response)
```

### Parameters

#### **pDcgmHandle**

IN: DCGM Handle

**drd**

IN: Contains the group id, test names, test parameters, struct version, and the validation that should be performed. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs.

**response**

OUT: Result of the validation process. Refer to [dcgmDiagResponse\\_t](#) for details.

**Returns**

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_NOT\_SUPPORTED if running the specified validate is not supported. This is usually due to the Tesla recommended driver not being installed on the system.
- ▶ DCGM\_ST\_BADPARAM if groupId, validate, or statusHandle is invalid
- ▶ DCGM\_ST\_GENERIC\_ERROR an internal error has occurred
- ▶ DCGM\_ST\_GROUP\_INCOMPATIBLE if groupId refers to a group of non-homogeneous GPUs. This is currently not allowed.

**Description**

Inform the action manager to perform a manual validation of a group of GPUs on the system

`dcgmReturn_t dcgmRunDiagnostic (dcgmHandle_t pDcgmHandle,  
dcgmGpuGrp_t groupId, dcgmDiagnosticLevel_t diagLevel,  
dcgmDiagResponse_t *diagResponse)`

**Parameters****pDcgmHandle**

IN: DCGM Handle

**groupId**

IN: Group ID representing collection of one or more GPUs. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs.

**diagLevel**

IN: Diagnostic level to run

**diagResponse**

IN/OUT: Result of running the DCGM diagnostic. .version should be set to [dcgmDiagResponse\\_version](#) before this call.

**Returns**

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_NOT\_SUPPORTED if running the diagnostic is not supported. This is usually due to the Tesla recommended driver not being installed on the system.
- ▶ DCGM\_ST\_BADPARAM if a provided parameter is invalid or missing
- ▶ DCGM\_ST\_GENERIC\_ERROR an internal error has occurred
- ▶ DCGM\_ST\_GROUP\_INCOMPATIBLE if groupId refers to a group of non-homogeneous GPUs. This is currently not allowed.
- ▶ DCGM\_ST\_VER\_MISMATCH if .version is not set or is invalid.

**Description**

Run a diagnostic on a group of GPUs

## 2.9. Topology

**dcmReturn\_t dcgmGetDeviceTopology**  
**(dcgmHandle\_t pDcgmHandle, unsigned int gpuld,**  
**dcgmDeviceTopology\_t \*pDcgmDeviceTopology)**

**Parameters****pDcgmHandle**

IN: DCGM Handle

**gpuld**

IN: GPU Id corresponding to which topology information should be fetched

**pDcgmDeviceTopology**

IN/OUT: Topology information corresponding to gpuld. pDcgmDeviceTopology->version must be set to dcgmDeviceTopology\_version before this call.

**Returns**

- ▶ DCGM\_ST\_OK if the call was successful.
- ▶ DCGM\_ST\_BADPARAM if gpuld or pDcgmDeviceTopology were not valid.
- ▶ DCGM\_ST\_VER\_MISMATCH if pDcgmDeviceTopology->version was not set to dcgmDeviceTopology\_version.

**Description**

Gets device topology corresponding to the gpuld.

**dcgmReturn\_t dcgmGetGroupTopology (dcgmHandle\_t pDcgmHandle, dcgmGpuGrp\_t groupId, dcgmGroupTopology\_t \*pDcgmGroupTopology)**

#### Parameters

##### **pDcgmHandle**

IN: DCGM Handle

##### **groupId**

IN: GroupId corresponding to which topology information should be fetched

##### **pDcgmGroupTopology**

IN/OUT: Topology information corresponding to groupId. pDcgmgroupTopology->version must be set to dcgmGroupTopology\_version.

#### Returns

- ▶ DCGM\_ST\_OK if the call was successful.
- ▶ DCGM\_ST\_BADPARAM if groupId or pDcgmGroupTopology were not valid.
- ▶ DCGM\_ST\_VER\_MISMATCH if pDcgmgroupTopology->version was not set to dcgmGroupTopology\_version.

#### Description

Gets group topology corresponding to the groupId.

## 2.10. Metadata

This chapter describes the methods that query for DCGM metadata.

**dcgmReturn\_t dcgmlntrospectToggleState (dcgmHandle\_t pDcgmHandle, dcgmlntrospectState\_t enabledState)**

#### Parameters

##### **pDcgmHandle**

IN: DCGM Handle

##### **enabledState**

IN: The state to set gathering of introspection data to

**Returns**

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM enabledState is an invalid state for metadata gathering

**Description**

Toggle the state of introspection metadata gathering in DCGM. Metadata gathering will increase the memory usage of DCGM so that it can store the metadata it gathers.

**dcgmReturn\_t dcgmIntrospectGetFieldsMemoryUsage  
(dcgmHandle\_t pDcgmHandle, dcgmIntrospectContext\_t  
\*context, dcgmIntrospectFullMemory\_t \*memoryInfo, int  
waitIfNoData)**

**Parameters****pDcgmHandle**

IN: DCGM Handle

**context**

IN: see [dcgmIntrospectContext\\_t](#). This identifies the level of fields to do introspection for (ex: all fields, field groups) context->version must be set to dcgmIntrospectContext\_version prior to this call.

**memoryInfo**

IN/OUT: see [dcgmIntrospectFullMemory\\_t](#). memoryInfo->version must be set to dcgmIntrospectFullMemory\_version prior to this call.

**waitIfNoData**

IN: if no metadata has been gathered, should this call block until data has been gathered (1), or should this call just return DCGM\_ST\_NO\_DATA (0).

**Returns**

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_NOT\_CONFIGURED if metadata gathering state is DCGM\_INTROSPECT\_STATE\_DISABLED
- ▶ DCGM\_ST\_NO\_DATA if waitIfNoData is false and metadata has not been gathered yet
- ▶ DCGM\_ST\_VER\_MISMATCH if context->version or memoryInfo->version is 0 or invalid.

**Description**

Get the current amount of memory used to store the given field collection.

```
dcgmReturn_t
dcgmlIntrospectGetHostengineMemoryUsage
(dcgmHandle_t pDcgmHandle, dcgmlIntrospectMemory_t
*memoryInfo, int waitIfNoData)
```

#### Parameters

##### **pDcgmHandle**

IN: DCGM Handle

##### **memoryInfo**

IN/OUT: see [dcgmlIntrospectMemory\\_t](#). memoryInfo->version must be set to dcgmlIntrospectMemory\_version prior to this call.

##### **waitIfNoData**

IN: if no metadata is gathered wait till this occurs (!0) or return

DCGM\_ST\_NO\_DATA (0)

#### Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_NOT\_CONFIGURED if metadata gathering state is DCGM\_INTROSPECT\_STATE\_DISABLED
- ▶ DCGM\_ST\_NO\_DATA if waitIfNoData is false and metadata has not been gathered yet
- ▶ DCGM\_ST\_VER\_MISMATCH if memoryInfo->version is 0 or invalid.

#### Description

Retrieve the total amount of memory that the hostengine process is currently using. This measurement represents both the resident set size (what is currently in RAM) and the swapped memory that belongs to the process.

```
dcgmReturn_t dcgmlIntrospectGetFieldsExecTime
(dcgmHandle_t pDcgmHandle, dcgmlIntrospectContext_t
*context, dcgmlIntrospectFullFieldsExecTime_t
*execTime, int waitIfNoData)
```

#### Parameters

##### **pDcgmHandle**

IN: DCGM Handle

**context**

IN: see `dcmIntrospectContext_t`. This identifies the level of fields to do introspection for (ex: all fields, field group ) context->version must be set to `dcmIntrospectContext_version` prior to this call.

**execTime**

IN/OUT: see `dcmIntrospectFullFieldsExecTime_t`. execTime->version must be set to `dcmIntrospectFullFieldsExecTime_version` prior to this call.

**waitIfNoData**

IN: if no metadata is gathered, wait until data has been gathered (1) or return `DCGM_ST_NO_DATA` (0)

**Returns**

- ▶ `DCGM_ST_OK` if the call was successful
- ▶ `DCGM_ST_NOT_CONFIGURED` if metadata gathering state is `DCGM_INTROSPECT_STATE_DISABLED`
- ▶ `DCGM_ST_NO_DATA` if `waitIfNoData` is false and metadata has not been gathered yet
- ▶ `DCGM_ST_VER_MISMATCH` if `context->version` or `execTime->version` is 0 or invalid.

**Description**

Get introspection info relating to execution time needed to update the fields identified by context.

## `dcmReturn_t dcmIntrospectUpdateAll (dcgmHandle_t pDcmHandle, int waitForUpdate)`

**Parameters****pDcmHandle**

IN: DCGM Handle

**waitForUpdate**

IN: Whether or not to wait for the update loop to complete before returning to the caller

**Returns**

- ▶ `DCGM_ST_OK` if the call was successful
- ▶ `DCGM_ST_BADPARAM` if `waitForUpdate` is invalid

## Description

This method is used to manually tell the the introspection module to update all DCGM introspection data. This is normally performed automatically on an interval of 1 second.

## 2.11. Topology

This chapter describes the methods that query for DCGM topology information.

```
dcgmReturn_t dcgmSelectGpusByTopology
(dcgmHandle_t pDcgmHandle, uint64_t inputGpuIds,
uint32_t numGpus, uint64_t *outputGpuIds, uint64_t hintFlags)
```

### Parameters

#### **pDcgmHandle**

IN: DCGM Handle

#### **inputGpuIds**

IN: a bitmask of which GPUs DCGM should consider. If some of the GPUs on the system are already in use, they shouldn't be included in the bitmask. 0 means that all of the GPUs in the system should be considered.

#### **numGpus**

IN: the number of GPUs that are desired from inputGpuIds. If this number is greater than the number of healthy GPUs in inputGpuIds, then less than numGpus gpus will be specified in outputGpuIds.

#### **outputGpuIds**

OUT: a bitmask of numGpus or fewer GPUs from inputGpuIds that represent the best placement available from inputGpuIds.

#### **hintFlags**

IN: a bitmask of DCGM\_TOPO\_HINT\_F\_ defines of hints that should be taken into account when assigning outputGpuIds.

### Returns

- ▶ DCGM\_ST\_OK if the call was successful

## Description

Get the best group of gpus from the specified bitmask according to topological proximity: cpuAffinity, NUMA node, and NVLink.

## `dcmReturn_t dcgmGetFieldSummary (dcgmHandle_t pDcgmHandle, dcgmFieldSummaryRequest_t *request)`

### Parameters

#### **pDcgmHandle**

IN: DCGM Handle

#### **request**

IN / OUT: a pointer to the struct detailing the request and containing the response

### Description

Get a summary of the values for a field id over a period of time.

## 2.12. Modules

This chapter describes the methods that query and configure DCGM modules.

## `dcmReturn_t dcgmModuleBlacklist (dcgmHandle_t pDcgmHandle, dcgmModuleId_t moduleId)`

### Parameters

#### **pDcgmHandle**

IN: DCGM Handle

#### **moduleId**

IN: ID of the module to blacklist. Use `dcgmModuleGetStatuses` to get a list of valid module IDs.

### Returns

- ▶ DCGM\_ST\_OK if the module has been blacklisted.
- ▶ DCGM\_ST\_IN\_USE if the module has already been loaded and cannot be blacklisted.
- ▶ DCGM\_ST\_BADPARAM if a parameter is missing or bad.

### Description

Set a module to be blacklisted. This module will be prevented from being loaded if it hasn't been loaded already. Modules are lazy-loaded as they are used by DCGM APIs, so it's important to call this API soon after the host engine has been started. You can also pass --blacklist-modules to the nv-hostengine binary to make sure modules get blacklisted immediately after the host engine starts up.

**dcgmReturn\_t dcgmModuleGetStatuses (dcgmHandle\_t pDcgmHandle, dcgmModuleGetStatuses\_t \*moduleStatuses)**

#### Parameters

##### **pDcgmHandle**

IN: DCGM Handle

##### **moduleStatuses**

OUT: Module statuses. .version should be set to dcgmModuleStatuses\_version upon calling.

#### Returns

- ▶ DCGM\_ST\_OK if the request succeeds.
- ▶ DCGM\_ST\_BADPARAM if a parameter is missing or bad.

#### Description

Get the status of all of the DCGM modules.

## 2.13. Profiling

This chapter describes the methods that watch profiling fields from within DCGM.

**dcgmReturn\_t dcgmProfGetSupportedMetricGroups (dcgmHandle\_t pDcgmHandle, dcgmProfGetMetricGroups\_t \*metricGroups)**

#### Parameters

##### **pDcgmHandle**

IN: DCGM Handle

##### **metricGroups**

IN/OUT: Metric groups supported for metricGroups->groupId. metricGroups->version should be set to dcgmProfGetMetricGroups\_version upon calling.

#### Returns

- ▶ DCGM\_ST\_OK if the request succeeds.
- ▶ DCGM\_ST\_BADPARAM if a parameter is missing or bad.

- ▶ DCGM\_ST\_GROUP\_INCOMPATIBLE if metricGroups->groupId's GPUs are not identical GPUs.
- ▶ DCGM\_ST\_NOT\_SUPPORTED if profiling metrics are not supported for the given GPU group.

## Description

Get all of the profiling metric groups for a given GPU group.

Profiling metrics are watched in groups of fields that are all watched together. For instance, if you want to watch DCGM\_FI\_PROF\_GR\_ENGINE\_ACTIVITY, this might also be in the same group as DCGM\_FI\_PROF\_SM\_EFFICIENCY. Watching this group would result in DCGM storing values for both of these metrics.

Some groups cannot be watched concurrently as others as they utilize the same hardware resource. For instance, you may not be able to watch DCGM\_FI\_PROF\_TENSOR\_OP\_UTIL at the same time as DCGM\_FI\_PROF\_GR\_ENGINE\_ACTIVITY on your hardware. At the same time, you may be able to watch DCGM\_FI\_PROF\_TENSOR\_OP\_UTIL at the same time as DCGM\_FI\_PROF\_NVLINK\_TX\_DATA.

Metrics that can be watched concurrently will have different .majorId fields in their dcgmProfMetricGroupInfo\_t

See [dcgmGroupCreate](#) for details on creating a GPU group See [dcgmProfWatchFields](#) to actually watch a metric group

## **dcgmReturn\_t dcgmProfWatchFields (dcgmHandle\_t pDcgmHandle, dcgmProfWatchFields\_t \*watchFields)**

### Parameters

#### pDcgmHandle

IN: DCGM Handle

#### watchFields

IN: Details of which metric groups to watch for which GPUs. See [dcgmProfWatchFields\\_v1](#) for details of what should be put in each struct member. watchFields->version should be set to dcgmProfWatchFields\_version upon calling.

### Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if a parameter is invalid
- ▶ DCGM\_ST\_NOT\_SUPPORTED if profiling metric group metricGroupTag is not supported for the given GPU group.

- ▶ DCGM\_ST\_GROUP\_INCOMPATIBLE if groupId's GPUs are not identical GPUs.  
Profiling metrics are only support for homogenous groups of GPUs.
- ▶ DCGM\_ST\_PROFILING\_MULTI\_PASS if any of the metric groups could not be watched concurrently due to requiring the hardware to gather them with multiple passes

### Description

Request that DCGM start recording updates for a given list of profiling field IDs.

Once metrics have been watched by this API, any of the normal DCGM field-value retrieval APIs can be used on the underlying fieldIDs of this metric group. See [dcgmGetLatestValues\\_v2](#), [dcgmGetLatestValuesForFields](#), [dcgmEntityGetLatestValues](#), and [dcgmEntitiesGetLatestValues](#).

**dcgmReturn\_t dcgmProfUnwatchFields (dcgmHandle\_t pDcgmHandle, dcgmProfUnwatchFields\_t \*unwatchFields)**

### Parameters

#### pDcgmHandle

IN: DCGM Handle

#### unwatchFields

IN: Details of which metric groups to unwatch for which GPUs. See [dcgmProfUnwatchFields\\_v1](#) for details of what should be put in each struct member. unwatchFields->version should be set to dcgmProfUnwatchFields\_version upon calling.

### Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if a parameter is invalid

### Description

Request that DCGM stop recording updates for all profiling field IDs for all GPUs

## 2.14. Enums and Macros

## enum dcgmOperationMode\_t

Operation mode for DCGM

DCGM can run in auto-mode where it runs additional threads in the background to collect any metrics of interest and auto manages any operations needed for policy management.

DCGM can also operate in manual-mode where its execution is controlled by the user. In this mode, the user has to periodically call APIs such as `dcgmPolicyTrigger` and `dcgmUpdateAllFields` which tells DCGM to wake up and perform data collection and operations needed for policy management.

### Values

`DCGM_OPERATION_MODE_AUTO = 1`

`DCGM_OPERATION_MODE_MANUAL = 2`

## enum dcgmOrder\_t

When more than one value is returned from a query, which order should it be returned in?

### Values

`DCGM_ORDER_ASCENDING = 1`

Data with earliest (lowest) timestamps returned first.

`DCGM_ORDER_DESCENDING = 2`

Data with latest (highest) timestamps returned first.

## enum dcgmReturn\_t

Return values for DCGM API calls.

### Values

`DCGM_ST_OK = 0`

Success.

`DCGM_ST_BADPARAM = -1`

A bad parameter was passed to a function.

`DCGM_ST_GENERIC_ERROR = -3`

A generic, unspecified error.

`DCGM_ST_MEMORY = -4`

An out of memory error occurred.

`DCGM_ST_NOT_CONFIGURED = -5`

Setting not configured.

**DCGM\_ST\_NOT\_SUPPORTED = -6**

Feature not supported.

**DCGM\_ST\_INIT\_ERROR = -7**

DCGM Init error.

**DCGM\_ST\_NVML\_ERROR = -8**

When NVML returns error.

**DCGM\_ST\_PENDING = -9**

Object is in pending state of something else.

**DCGM\_ST\_UNINITIALIZED = -10**

Object is in undefined state.

**DCGM\_ST\_TIMEOUT = -11**

Requested operation timed out.

**DCGM\_ST\_VER\_MISMATCH = -12**

Version mismatch between received and understood API.

**DCGM\_ST\_UNKNOWN\_FIELD = -13**

Unknown field id.

**DCGM\_ST\_NO\_DATA = -14**

No data is available.

**DCGM\_ST\_STALE\_DATA = -15**

Data is considered stale.

**DCGM\_ST\_NOT\_WATCHED = -16**

The given field id is not being updated by the cache manager.

**DCGM\_ST\_NO\_PERMISSION = -17**

Do not have permission to perform the desired action.

**DCGM\_ST\_GPU\_IS\_LOST = -18**

GPU is no longer reachable.

**DCGM\_ST\_RESET\_REQUIRED = -19**

GPU requires a reset.

**DCGM\_ST\_FUNCTION\_NOT\_FOUND = -20**

The function that was requested was not found (bindings only error).

**DCGM\_ST\_CONNECTION\_NOT\_VALID = -21**

The connection to the host engine is not valid any longer.

**DCGM\_ST\_GPU\_NOT\_SUPPORTED = -22**

This GPU is not supported by DCGM.

**DCGM\_ST\_GROUP\_INCOMPATIBLE = -23**

The GPUs of the provided group are not compatible with each other for the requested operation.

**DCGM\_ST\_MAX\_LIMIT = -24**

Max limit reached for the object.

**DCGM\_ST\_LIBRARY\_NOT\_FOUND = -25**

DCGM library could not be found.

**DCGM\_ST\_DUPLICATE\_KEY = -26**

Duplicate key passed to a function.

<b>DCGM_ST_GPU_IN_SYNC_BOOST_GROUP = -27</b>	GPU is already a part of a sync boost group.
<b>DCGM_ST_GPU_NOT_IN_SYNC_BOOST_GROUP = -28</b>	GPU is not a part of a sync boost group.
<b>DCGM_ST_REQUIRES_ROOT = -29</b>	This operation cannot be performed when the host engine is running as non-root.
<b>DCGM_ST_NVVS_ERROR = -30</b>	DCGM GPU Diagnostic was successfully executed, but reported an error.
<b>DCGM_ST_INSUFFICIENT_SIZE = -31</b>	An input argument is not large enough.
<b>DCGM_ST_FIELD_UNSUPPORTED_BY_API = -32</b>	The given field ID is not supported by the API being called.
<b>DCGM_ST_MODULE_NOT_LOADED = -33</b>	This request is serviced by a module of DCGM that is not currently loaded.
<b>DCGM_ST_IN_USE = -34</b>	The requested operation could not be completed because the affected resource is in use.
<b>DCGM_ST_GROUP_IS_EMPTY = -35</b>	This group is empty and the requested operation is not valid on an empty group.
<b>DCGM_ST_PROFILING_NOT_SUPPORTED = -36</b>	Profiling is not supported for this group of GPUs or GPU.
<b>DCGM_ST_PROFILING_LIBRARY_ERROR = -37</b>	The third-party Profiling module returned an unrecoverable error.
<b>DCGM_ST_PROFILING_MULTI_PASS = -38</b>	The requested profiling metrics cannot be collected in a single pass.
<b>DCGM_ST_DIAG_ALREADY_RUNNING = -39</b>	A diag instance is already running, cannot run a new diag until the current one finishes.
<b>DCGM_ST_DIAG_BAD_JSON = -40</b>	The DCGM GPU Diagnostic returned JSON that cannot be parsed.
<b>DCGM_ST_DIAG_BAD_LAUNCH = -41</b>	Error while launching the DCGM GPU Diagnostic.
<b>DCGM_ST_DIAG_VARIANCE = -42</b>	There is too much variance while training the diagnostic.
<b>DCGM_ST_DIAG_THRESHOLD_EXCEEDED = -43</b>	A field value met or exceeded the error threshold.
<b>DCGM_ST_INSUFFICIENT_DRIVER_VERSION = -44</b>	

## enum dcgmGroupType\_t

Type of GPU groups

**Values****DCGM\_GROUP\_DEFAULT = 0**

All the GPUs on the node are added to the group.

**DCGM\_GROUP\_EMPTY = 1**

Creates an empty group.

**DCGM\_GROUP\_DEFAULT\_NVSWITCHES = 2**

All NvSwitches of the node are added to the group.

**enum dcgmConfigType\_t**

Represents the type of configuration to be fetched from the GPUs

**Values****DCGM\_CONFIG\_TARGET\_STATE = 0**

The target configuration values to be applied.

**DCGM\_CONFIG\_CURRENT\_STATE = 1**

The current configuration state.

**enum dcgmConfigPowerLimitType\_t**

Represents the power cap for each member of the group.

**Values****DCGM\_CONFIG\_POWER\_CAP\_INDIVIDUAL = 0**

Represents the power cap to be applied for each member of the group.

**DCGM\_CONFIG\_POWER\_BUDGET\_GROUP = 1**

Represents the power budget for the entire group.

**#define DCGM\_INT32\_BLANK 0x7fffffff0**

Represents value of the field which can be returned by Host Engine in case the operation is not successful Base value for 32 bits integer blank. can be used as an unspecified blank

**#define DCGM\_INT64\_BLANK 0x7fffffffffffff0**

Base value for 64 bits integer blank. can be used as an unspecified blank

**#define DCGM\_FP64\_BLANK 140737488355328.0**

Base value for double blank.  $2^{**} 47$ . FP 64 has 52 bits of mantissa, so 47 bits can still increment by 1 and represent each value from 0-15

```
#define DCGM_STR_BLANK "<<<NULL>>>"
```

Base value for string blank.

```
#define DCGM_INT32_NOT_FOUND (DCGM_INT32_BLANK +1)
```

Represents an error where INT32 data was not found

```
#define DCGM_INT64_NOT_FOUND (DCGM_INT64_BLANK +1)
```

Represents an error where INT64 data was not found

```
#define DCGM_FP64_NOT_FOUND (DCGM_FP64_BLANK +1.0)
```

Represents an error where FP64 data was not found

```
#define DCGM_STR_NOT_FOUND "<<<NOT_FOUND>>>"
```

Represents an error where STR data was not found

```
#define DCGM_INT32_NOT_SUPPORTED  
(DCGM_INT32_BLANK+2)
```

Represents an error where fetching the INT32 value is not supported

```
#define DCGM_INT64_NOT_SUPPORTED  
(DCGM_INT64_BLANK+2)
```

Represents an error where fetching the INT64 value is not supported

```
#define DCGM_FP64_NOT_SUPPORTED  
(DCGM_FP64_BLANK+2.0)
```

Represents an error where fetching the FP64 value is not supported

```
#define DCGM_STR_NOT_SUPPORTED  
"<<<NOT_SUPPORTED>>>"
```

Represents an error where fetching the STR value is not supported

```
#define DCGM_INT32_NOT_PERMISSIONED  
(DCGM_INT32_BLANK+3)
```

Represents and error where fetching the INT32 value is not allowed with our current credentials

```
#define DCGM_INT64_NOT_PERMISSIONED  
(DCGM_INT64_BLANK+3)
```

Represents and error where fetching the INT64 value is not allowed with our current credentials

```
#define DCGM_FP64_NOT_PERMISSIONED  
(DCGM_FP64_BLANK+3.0)
```

Represents and error where fetching the FP64 value is not allowed with our current credentials

```
#define DCGM_STR_NOT_PERMISSIONED  
""><<<NOT_PERM>>>"
```

Represents and error where fetching the STR value is not allowed with our current credentials

```
#define DCGM_INT32_IS_BLANK (((val) >=  
DCGM_INT32_BLANK) ? 1 : 0)
```

Macro to check if a INT32 value is blank or not

```
#define DCGM_INT64_IS_BLANK (((val) >=  
DCGM_INT64_BLANK) ? 1 : 0)
```

Macro to check if a INT64 value is blank or not

```
#define DCGM_FP64_IS_BLANK (((val) >=  
DCGM_FP64_BLANK ? 1 : 0))
```

Macro to check if a FP64 value is blank or not

```
#define DCGM_STR_IS_BLANK (val == strstr(val, "<<<")  
&& strstr(val, ">>>"))
```

Macro to check if a STR value is blank or not Works on (char \*). Looks for <<< at first position and >>> inside string

```
#define DCGM_MAX_NUM_DEVICES 16
```

Max number of GPUs supported by DCGM

```
#define DCGM_NVLINK_MAX_LINKS_PER_GPU 6
```

Number of NvLink links per GPU supported by DCGM This is 6 for Volta and 4 for Pascal

```
#define DCGM_MAX_NUM_SWITCHES 12
```

Max number of NvSwitches supported by DCGM

```
#define DCGM_NVLINK_MAX_LINKS_PER_NVSWITCH 18
```

Number of NvLink links per NvSwitch supported by DCGM

```
#define DCGM_MAX_VGPU_INSTANCES_PER_PGPU 32
```

Maximum number of vGPU instances per physical GPU

```
#define DCGM_MAX_NUM_VGPU_DEVICES  
DCGM_MAX_NUM_DEVICES *  
DCGM_MAX_VGPU_INSTANCES_PER_PGPU
```

Max number of vGPUs supported on DCGM

```
#define DCGM_MAX_STR_LENGTH 256
```

Max length of the DCGM string field

```
#define DCGM_MAX_CLOCKS 256
```

Max number of clocks supported for a device

```
#define DCGM_MAX_NUM_GROUPS 64
```

Max limit on the number of groups supported by DCGM

```
#define DCGM_MAX_FBC_SESSIONS 256
```

Max number of active FBC sessions

```
#define DCGM_VGPU_NAME_BUFFER_SIZE 64
```

Represents the size of a buffer that holds a vGPU type Name or vGPU class type or name of process running on vGPU instance.

```
#define DCGM_GRID_LICENSE_BUFFER_SIZE 128
```

Represents the size of a buffer that holds a vGPU license string

```
#define DCGM_CONFIG_COMPUTEMODE_DEFAULT 0
```

Default compute mode -- multiple contexts per device

```
#define DCGM_CONFIG_COMPUTEMODE_PROHIBITED 1
```

Compute-prohibited mode -- no contexts per device

```
#define  
DCGM_CONFIG_COMPUTEMODE_EXCLUSIVE_PROCESS 2
```

Compute-exclusive-process mode -- only one context per device, usable from multiple threads at a time

```
#define DCGM_HE_PORT_NUMBER 5555
```

Default Port Number for DCGM Host Engine

```
#define MAKE_DCGM_VERSION (unsigned int)  
(sizeof(typeName) | ((ver)<<24))
```

Creates a unique version number for each struct

```
#define DCGM_GROUP_ALL_GPUS 0x7fffffff
```

Identifies for special DCGM groups

```
#define DCGM_GROUP_MAX_ENTITIES 64
```

Maximum number of entities per entity group

## 2.15. Structure definitions

```
struct dcgmConnectV2Params_v1
struct dcgmConnectV2Params_v2
struct dcgmGroupInfo_v1
struct dcgmGroupEntityPair_t
struct dcgmGroupInfo_v2
struct dcgmFieldGroupInfo_v1
struct dcgmErrorInfo_t
struct dcgmClockSet_v1
struct dcgmDeviceSupportedClockSets_v1
struct dcgmDevicePidAccountingStats_v1
struct dcgmDeviceThermals_v1
struct dcgmDevicePowerLimits_v1
struct dcgmDeviceIdentifiers_v1
struct dcgmDeviceMemoryUsage_v1
struct dcgmDeviceVgpuUtilInfo_v1
struct dcgmDeviceEncStats_v1
struct dcgmDeviceFbcStats_v1
struct dcgmDeviceFbcSessionInfo_v1
```

```
struct dcgmDeviceFbcSessions_v1
struct dcgmDeviceVgpuEncSessions_v1
struct dcgmDeviceVgpuProcessUtilInfo_v1
struct dcgmDeviceVgpuls_v1
struct dcgmDeviceVgpuTypeInfo_v1
struct dcgmDeviceAttributes_v1
struct dcgmVgpuDeviceAttributes_v6
struct dcgmVgpuInstanceAttributes_v1
struct dcgmConfigPerfStateSettings_t
struct dcgmConfigPowerLimit_t
struct dcgmConfig_v1
struct dcgmVgpuConfig_v1
struct dcgmPolicyConditionParms_t
struct dcgmPolicyViolationNotify_t
struct dcgmPolicy_v1
struct dcgmPolicyConditionDbe_t
struct dcgmPolicyConditionPci_t
struct dcgmPolicyConditionMpr_t
```

```
struct dcgmPolicyConditionThermal_t
struct dcgmPolicyConditionPower_t
struct dcgmPolicyConditionNvlink_t
struct dcgmPolicyConditionXID_t
struct dcgmPolicyCallbackResponse_v1
struct dcgmFieldValue_v1
struct dcgmFieldValue_v2
struct dcgmStatSummaryInt64_t
struct dcgmStatSummaryInt32_t
struct dcgmStatSummaryFp64_t
struct dcgmHealthResponse_v1
struct dcgmHealthResponse_v2
struct dcgmHealthResponse_v3
struct dcgmProcessUtilInfo_t
struct dcgmProcessUtilSample_t
struct dcgmPidSingleInfo_t
struct dcgmPidInfo_v1
struct dcgmGpuUsageInfo_t
```

```
struct dcgmJobInfo_v2
struct dcgmRunningProcess_v1
struct dcgmDiagResponsePerGpu_v1
struct dcgmDiagResponse_v3
struct dcgmDiagResponse_v4
struct dcgmDeviceTopology_v1
struct dcgmGroupTopology_v1
struct dcgmlntrospectContext_v1
struct dcgmlntrospectFieldsExecTime_v1
struct dcgmlntrospectFullFieldsExecTime_v1
struct dcgmlntrospectMemory_v1
struct dcgmlntrospectFullMemory_v1
struct dcgmlntrospectCpuUtil_v1
struct dcgmNvLinkGpuLinkStatus_t
struct dcgmNvLinkNvSwitchLinkStatus_t
struct dcgmNvLinkStatus_v1
struct dcgmModuleGetStatusesModule_t
struct dcgmProfWatchFields_v1
```

**struct dcgmProfUnwatchFields\_v1****struct dcgmVersionInfo\_v1****enum dcgmPolicyCondition\_t**

Enumeration for policy conditions. When used as part of dcgmPolicy\_t these have corresponding parameters to allow them to be switched on/off or set specific violation thresholds

**Values****DCGM\_POLICY\_COND\_DBE = 0x1**

Double bit errors -- boolean in dcgmPolicyConditionParms\_t.

**DCGM\_POLICY\_COND\_PCI = 0x2**

PCI events/errors -- boolean in dcgmPolicyConditionParms\_t.

**DCGM\_POLICY\_COND\_MAX\_PAGES\_RETIRE = 0x4**

Maximum number of retired pages -- number required in dcgmPolicyConditionParms\_t.

**DCGM\_POLICY\_COND\_THERMAL = 0x8**

Thermal violation -- number required in dcgmPolicyConditionParms\_t.

**DCGM\_POLICY\_COND\_POWER = 0x10**

Power violation -- number required in dcgmPolicyConditionParms\_t.

**DCGM\_POLICY\_COND\_NVLINK = 0x20**

NVLINK errors -- boolean in dcgmPolicyConditionParms\_t.

**DCGM\_POLICY\_COND\_XID = 0x40**

XID errors -- number required in dcgmPolicyConditionParms\_t.

**enum dcgmPolicyMode\_t**

Enumeration for policy modes

**Values****DCGM\_POLICY\_MODE\_AUTOMATED = 0**

automatic mode

**DCGM\_POLICY\_MODE\_MANUAL = 1**

manual mode

**enum dcgmPolicyIsolation\_t**

Enumeration for policy isolation modes

**Values****DCGM\_POLICY\_ISOLATION\_NONE = 0**

no isolation of GPUs on error

**enum dcgmPolicyAction\_t**

Enumeration for policy actions

**Values****DCGM\_POLICY\_ACTION\_NONE = 0**

no action

**DCGM\_POLICY\_ACTION\_GPURESET = 1**

perform a GPU reset on violation

**enum dcgmPolicyValidation\_t**

Enumeration for policy validation actions

**Values****DCGM\_POLICY\_VALID\_NONE = 0**

no validation after an action is performed

**DCGM\_POLICY\_VALID\_SV\_SHORT = 1**

run a short System Validation on the system after failure

**DCGM\_POLICY\_VALID\_SV\_MED = 2**

run a medium System Validation test after failure

**DCGM\_POLICY\_VALID\_SV\_LONG = 3**

run a extensive System Validation test after failure

**enum dcgmPolicyFailureResp\_t**

Enumeration for policy failure responses

**Values****DCGM\_POLICY\_FAILURE\_NONE = 0**

on failure of validation perform no action

**enum dcgmHealthSystems\_t**

Systems structure used to enable or disable health watch systems

**Values****DCGM\_HEALTH\_WATCH\_PCIE = 0x1**

PCIe system watches (must have 1m of data before query).

**DCGM\_HEALTH\_WATCH\_NVLINK = 0x2**  
NVLINK system watches.

**DCGM\_HEALTH\_WATCH\_PMU = 0x4**  
Power management unit watches.

**DCGM\_HEALTH\_WATCH MCU = 0x8**  
Microcontroller unit watches.

**DCGM\_HEALTH\_WATCH\_MEM = 0x10**  
Memory watches.

**DCGM\_HEALTH\_WATCH\_SM = 0x20**  
Streaming multiprocessor watches.

**DCGM\_HEALTH\_WATCH\_INFOROM = 0x40**  
Inforom watches.

**DCGM\_HEALTH\_WATCH\_THERMAL = 0x80**  
Temperature watches (must have 1m of data before query).

**DCGM\_HEALTH\_WATCH\_POWER = 0x100**  
Power watches (must have 1m of data before query).

**DCGM\_HEALTH\_WATCH\_DRIVER = 0x200**  
Driver-related watches.

**DCGM\_HEALTH\_WATCH\_NVSWITCH\_NONFATAL = 0x400**  
Non-fatal errors in NvSwitch.

**DCGM\_HEALTH\_WATCH\_NVSWITCH\_FATAL = 0x800**  
Fatal errors in NvSwitch.

**DCGM\_HEALTH\_WATCH\_ALL = 0xFFFFFFFF**  
All watches enabled.

## enum dcgmHealthWatchResults\_t

Health Watch test results

### Values

**DCGM\_HEALTH\_RESULT\_PASS = 0**  
All results within this system are reporting normal.

**DCGM\_HEALTH\_RESULT\_WARN = 10**  
A warning has been issued, refer to the response for more information.

**DCGM\_HEALTH\_RESULT\_FAIL = 20**  
A failure has been issued, refer to the response for more information.

## enum dcgmDiagnosticLevel\_t

Enumeration for diagnostic levels

**Values****DCGM\_DIAG\_LVL\_INVALID = 0**

Uninitialized.

**DCGM\_DIAG\_LVL\_SHORT = 10**

run a very basic health check on the system

**DCGM\_DIAG\_LVL\_MED = 20**

run a medium-length diagnostic (a few minutes)

**DCGM\_DIAG\_LVL\_LONG = 30**

run a extensive diagnostic (several minutes)

**enum dcgmDiagResult\_t**

Diagnostic test results

**Values****DCGM\_DIAG\_RESULT\_PASS = 0**

This test passed as diagnostics.

**DCGM\_DIAG\_RESULT\_SKIP = 1**

This test was skipped.

**DCGM\_DIAG\_RESULT\_WARN = 2**

This test passed with warnings.

**DCGM\_DIAG\_RESULT\_FAIL = 3**

This test failed the diagnostics.

**DCGM\_DIAG\_RESULT\_NOT\_RUN = 4**

This test wasn't executed.

**enum dcgmPerGpuTestIndices\_t**

Diagnostic per gpu tests - fixed indices for dcgmDiagResponsePerGpu\_t.results[]

**Values****DCGM\_MEMORY\_INDEX = 0**

Memory test index.

**DCGM\_DIAGNOSTIC\_INDEX = 1**

Diagnostic test index.

**DCGM\_PCI\_INDEX = 2**

PCIe test index.

**DCGM\_SM\_PERF\_INDEX = 3**

SM Stress test index.

**DCGM\_TARGETED\_PERF\_INDEX = 4**

Targeted Stress test index.

**DCGM\_TARGETED\_POWER\_INDEX = 5**

Targeted Power test index.

**DCGM\_MEMORY\_BANDWIDTH\_INDEX = 6**

Memory bandwidth test index.

**enum dcgmGpuTopologyLevel\_t**

Represents level relationships within a system between two GPUs. The enums are spaced to allow for future relationships. These match the definitions in nvml.h

**Values****DCGM\_TOPOLOGY\_BOARD = 0x1**

multi-GPU board

**DCGM\_TOPOLOGY\_SINGLE = 0x2**

all devices that only need traverse a single PCIe switch

**DCGM\_TOPOLOGY\_MULTIPLE = 0x4**

all devices that need not traverse a host bridge

**DCGM\_TOPOLOGY\_HOSTBRIDGE = 0x8**

all devices that are connected to the same host bridge

**DCGM\_TOPOLOGY\_CPU = 0x10**

all devices that are connected to the same CPU but possibly multiple host bridges

**DCGM\_TOPOLOGY\_SYSTEM = 0x20**

all devices in the system

**DCGM\_TOPOLOGY\_NVLINK1 = 0x0100**

GPUs connected via a single NVLINK link.

**DCGM\_TOPOLOGY\_NVLINK2 = 0x0200**

GPUs connected via two NVLINK links.

**DCGM\_TOPOLOGY\_NVLINK3 = 0x0400**

GPUs connected via three NVLINK links.

**DCGM\_TOPOLOGY\_NVLINK4 = 0x0800**

GPUs connected via four NVLINK links.

**DCGM\_TOPOLOGY\_NVLINK5 = 0x1000**

GPUs connected via five NVLINK links.

**DCGM\_TOPOLOGY\_NVLINK6 = 0x2000**

GPUs connected via six NVLINK links.

**enum dcgmIntrospectLevel\_t**

Identifies a level to retrieve field introspection info for

**Values****DCGM\_INTROSPECT\_LVL\_INVALID = 0**

Invalid value.

**DCGM\_INTROSPECT\_LVL\_FIELD = 1**

Introspection data is grouped by field ID.

**DCGM\_INTROSPECT\_LVL\_FIELD\_GROUP = 2**

Introspection data is grouped by field group.

**DCGM\_INTROSPECT\_LVL\_ALL\_FIELDS**

Introspection data is aggregated for all fields.

## enum dcgmIntrospectState\_t

State of DCGM metadata gathering. If it is set to DISABLED then "Metadata" API calls to DCGM are not supported.

### Values

**DCGM\_INTROSPECT\_STATE\_DISABLED = 0****DCGM\_INTROSPECT\_STATE\_ENABLED = 1**

## enum dcgmGpuNvLinkErrorType\_t

Identifies a GPU NVLink error type returned by  
DCGM\_FI\_DEV\_GPU\_NVLINK\_ERRORS

### Values

**DCGM\_GPU\_NVLINK\_ERROR\_RECOVERY\_REQUIRED = 1**

NVLink link recovery error occurred.

**DCGM\_GPU\_NVLINK\_ERROR\_FATAL**

NVLink link fatal error occurred.

## enum dcgmNvLinkLinkState\_t

NvLink link states

### Values

**DcgmNvLinkLinkStateNotSupported = 0**

NvLink is unsupported by this GPU (Default for GPUs).

**DcgmNvLinkLinkStateDisabled = 1**

NvLink is supported for this link but this link is disabled (Default for NvSwitches).

**DcgmNvLinkLinkStateDown = 2**

This NvLink link is down (inactive).

**DcgmNvLinkLinkStateUp = 3**

This NvLink link is up (active).

## enum dcgmModuleId\_t

Module IDs

**Values****DcgmModuleIdCore = 0**

Core DCGM - always loaded.

**DcgmModuleIdNvSwitch = 1**

NvSwitch Module.

**DcgmModuleIdVGPU = 2**

VGPU Module.

**DcgmModuleIdIntrospect = 3**

Introspection Module.

**DcgmModuleIdHealth = 4**

Health Module.

**DcgmModuleIdPolicy = 5**

Policy Module.

**DcgmModuleIdConfig = 6**

Config Module.

**DcgmModuleIdDiag = 7**

GPU Diagnostic Module.

**DcgmModuleIdProfiling = 8**

Profiling Module.

**DcgmModuleIdCount**

Always last. 1 greater than largest value above.

**enum dcgmModuleStatus\_t**

Module Status. Modules are lazy loaded, so they will be in status

DcgmModuleStatusNotLoaded until they are used. Once modules are used, they will move to another status.

**Values****DcgmModuleStatusNotLoaded = 0**

Module has not been loaded yet.

**DcgmModuleStatusBlacklisted = 1**

Module has been blacklisted from being loaded.

**DcgmModuleStatusFailed = 2**

Loading the module failed.

**DcgmModuleStatusLoaded = 3**

Module has been loaded.

**typedef void \*dcgmHandle\_t**

Identifier for DCGM Handle.

**typedef void \*dcgmGpuGrp\_t**

Identifier for a group of GPUs. A group can have one or more GPUs.

**typedef void \*dcgmFieldGrp\_t**

Identifier for a group of fields.

**typedef void \*dcgmStatus\_t**

Identifier for list of status codes.

**typedef dcgmConnectV2Params\_t**

Typedef for `dcgmConnectV2Params_v2`

**typedef dcgmGroupInfo\_t**

Typedef for `dcgmGroupInfo_v2`

**typedef dcgmClockSet\_t**

Typedef for `dcgmClockSet_v1`

**typedef dcgmDeviceSupportedClockSets\_t**

Typedef for `dcgmDeviceSupportedClockSets_v1`

**typedef dcgmDevicePidAccountingStats\_t**

Typedef for `dcgmDevicePidAccountingStats_v1`

**typedef dcgmDeviceThermals\_t**

Typedef for `dcgmDeviceThermals_v1`

**typedef dcgmDevicePowerLimits\_t**

Typedef for `dcgmDevicePowerLimits_v1`

**typedef dcgmDeviceIdentifiers\_t**

Typedef for `dcgmDeviceIdentifiers_v1`

**typedef dcgmDeviceMemoryUsage\_t**

Typedef for `dcgmDeviceMemoryUsage_v1`

**typedef dcgmDeviceVgpuUtilInfo\_t**

Typedef for `dcgmDeviceVgpuUtilInfo_v1`

**typedef dcgmDeviceEncStats\_t**

Typedef for `dcgmDeviceEncStats_v1`

**typedef dcgmDeviceFbcStats\_t**

Typedef for `dcgmDeviceFbcStats_v1`

**typedef dcgmDeviceFbcSessionInfo\_t**

Typedef for `dcgmDeviceFbcSessionInfo_v1`

**typedef dcgmDeviceFbcSessions\_t**

Typedef for `dcgmDeviceFbcSessions_v1`

**typedef dcgmDeviceVgpuEncSessions\_t**

Typedef for `dcgmDeviceVgpuEncSessions_v1`

**typedef dcgmDeviceVgpuProcessUtilInfo\_t**

Typedef for `dcgmDeviceVgpuProcessUtilInfo_v1`

**typedef dcgmDeviceVgpuids\_t**

Typedef for `dcgmDeviceVgpuIds_v1`

**typedef dcgmDeviceVgpuTypeInfo\_t**

Typedef for `dcgmDeviceVgpuTypeInfo_v1`

**typedef dcgmDeviceAttributes\_t**

Typedef for `dcgmDeviceAttributes_v1`

**typedef dcgmVgpuDeviceAttributes\_t**

Typedef for `dcgmVgpuDeviceAttributes_v6`

**typedef dcgmVgpuInstanceAttributes\_t**

Typedef for `dcgmVgpuInstanceAttributes_v1`

**typedef dcgmConfig\_t**

Typedef for `dcgmConfig_v1`

**typedef dcgmVgpuConfig\_t**

Typedef for `dcgmVgpuConfig_v1`

**typedef (\*fpRecvUpdates) (void\* userData)**

Represents a callback to receive updates from asynchronous functions. Currently the only implemented callback function is `dcgmPolicyRegister` and the void \* data will be a pointer to `dcgmPolicyCallbackResponse_t`. Ex. `dcgmPolicyCallbackResponse_t *callbackResponse = (dcgmPolicyCallbackResponse_t *) userData;`

**typedef dcgmPolicy\_t**

Typedef for `dcgmPolicy_v1`

**typedef dcgmPolicyCallbackResponse\_t**

Typedef for `dcgmPolicyCallbackResponse_v1`

**typedef (\*dcgmFieldValueEnumeration\_f) (unsigned int gpuld, dcgmFieldValue\_v1\* values, int numValues, void\* userData)**

User callback function for processing one or more field updates. This callback will be invoked one or more times per field until all of the expected field values have been enumerated. It is up to the callee to detect when the field id changes

Returns 0 if OK <0 if enumeration should stop. This allows to callee to abort field value enumeration.

```
typedef (*dcgmFieldValueEntityEnumeration_f)
(dcgm_field_entity_group_t entityIdGroup,
dcgm_field_eid_t entityId, dcgmFieldValue_v1* values,
int numValues, void* userData)
```

User callback function for processing one or more field updates. This callback will be invoked one or more times per field until all of the expected field values have been enumerated. It is up to the callee to detect when the field id changes

Returns 0 if OK <0 if enumeration should stop. This allows to callee to abort field value enumeration.

## **typedef dcgmHealthResponse\_t**

Typedef for `dcgmHealthResponse_v3`

## **typedef dcgmPidInfo\_t**

Typedef for `dcgmPidInfo_v1`

## **typedef dcgmJobInfo\_t**

Typedef for `dcgmJobInfo_v2`

## **typedef dcgmRunningProcess\_t**

Typedef for `dcgmRunningProcess_v1`

## **typedef dcgmDiagResponse\_v5 dcgmDiagResponse\_t**

Typedef for `dcgmDiagResponse_v4`

## **typedef dcgmDeviceTopology\_t**

Typedef for `dcgmDeviceTopology_v1`

## **typedef dcgmGroupTopology\_t**

Typedef for `dcgmGroupTopology_v1`

## **typedef dcgmIntrospectContext\_t**

Typedef for `dcgmIntrospectContext_v1`

**typedef dcgmIntrospectFieldsExecTime\_t**

Typedef for `dcgmIntrospectFieldsExecTime_t`

**typedef dcgmIntrospectFullFieldsExecTime\_t**

typedef for `dcgmIntrospectFullFieldsExecTime_v1`

**typedef dcgmIntrospectMemory\_t**

Typedef for `dcgmIntrospectMemory_t`

**typedef dcgmIntrospectFullMemory\_t**

typedef for `dcgmIntrospectFullMemory_v1`

**typedef dcgmIntrospectCpuUtil\_t**

Typedef for `dcgmIntrospectCpuUtil_t`

**typedef dcgmRunDiag\_v5 dcgmRunDiag\_t**

Typedef for `dcgmRunDiag_t`

**#define dcgmConnectV2Params\_version1****MAKE\_DCGM\_VERSION(dcgmConnectV2Params\_v1, 1)**

Version 1 for `dcgmConnectV2Params_v1`

**#define dcgmConnectV2Params\_version2****MAKE\_DCGM\_VERSION(dcgmConnectV2Params\_v2, 2)**

Version 2 for `dcgmConnectV2Params_v2`

**#define dcgmConnectV2Params\_version****dcgmConnectV2Params\_version2**

Latest version for `dcgmConnectV2Params_t`

**#define dcgmGroupInfo\_version1****MAKE\_DCGM\_VERSION(dcgmGroupInfo\_v1, 1)**

Version 1 for `dcgmGroupInfo_v1`

```
#define dcgmGroupInfo_version2  
MAKE_DCGM_VERSION(dcgmGroupInfo_v2, 2)
```

Version 2 for dcgmGroupInfo\_v2

```
#define dcgmGroupInfo_version  
dcgmGroupInfo_version2
```

Latest version for dcgmGroupInfo\_t

```
#define DCGM_MAX_NUM_FIELD_GROUPS 64
```

Maximum number of field groups that can exist

```
#define DCGM_MAX_FIELD_IDS_PER_FIELD_GROUP 128
```

Maximum number of field IDs that can be in a single field group

```
#define dcgmFieldGroupInfo_version1  
MAKE_DCGM_VERSION(dcgmFieldGroupInfo_v1, 1)
```

Version 1 for dcgmFieldGroupInfo\_v1

```
#define dcgmFieldGroupInfo_version  
dcgmFieldGroupInfo_version1
```

Latest version for dcgmFieldGroupInfo\_t

```
#define dcgmAllFieldGroup_version1  
MAKE_DCGM_VERSION(dcgmAllFieldGroup_v1, 1)
```

Version 1 for dcgmAllFieldGroup\_v1

```
#define dcgmAllFieldGroup_version  
dcgmAllFieldGroup_version1
```

Latest version for dcgmAllFieldGroup\_t

```
#define dcgmClockSet_version1  
MAKE_DCGM_VERSION(dcgmClockSet_v1, 1)
```

Version 1 for dcgmClockSet\_v1

```
#define dcgmClockSet_version dcgmClockSet_version1
```

Latest version for `dcgmClockSet_t`

```
#define dcgmDeviceSupportedClockSets_version1
```

```
MAKE_DCGM_VERSION(dcgmDeviceSupportedClockSets_v1,  
1)
```

Version 1 for `dcgmDeviceSupportedClockSets_v1`

```
#define dcgmDeviceSupportedClockSets_version
```

```
dcgmDeviceSupportedClockSets_version1
```

Latest version for `dcgmDeviceSupportedClockSets_t`

```
#define dcgmDevicePidAccountingStats_version1
```

```
MAKE_DCGM_VERSION(dcgmDevicePidAccountingStats_v1,  
1)
```

Version 1 for `dcgmDevicePidAccountingStats_v1`

```
#define dcgmDevicePidAccountingStats_version
```

```
dcgmDevicePidAccountingStats_version1
```

Latest version for `dcgmDevicePidAccountingStats_t`

```
#define dcgmDeviceThermals_version1
```

```
MAKE_DCGM_VERSION(dcgmDeviceThermals_v1, 1)
```

Version 1 for `dcgmDeviceThermals_v1`

```
#define dcgmDeviceThermals_version
```

```
dcgmDeviceThermals_version1
```

Latest version for `dcgmDeviceThermals_t`

```
#define dcgmDevicePowerLimits_version1
```

```
MAKE_DCGM_VERSION(dcgmDevicePowerLimits_v1, 1)
```

Version 1 for `dcgmDevicePowerLimits_v1`

```
#define dcgmDevicePowerLimits_version  
dcgmDevicePowerLimits_version1
```

Latest version for `dcgmDevicePowerLimits_t`

```
#define dcgmDeviceIdentifiers_version1  
MAKE_DCGM_VERSION(dcgmDeviceIdentifiers_v1, 1)
```

Version 1 for `dcgmDeviceIdentifiers_v1`

```
#define dcgmDeviceIdentifiers_version  
dcgmDeviceIdentifiers_version1
```

Latest version for `dcgmDeviceIdentifiers_t`

```
#define dcgmDeviceMemoryUsage_version1  
MAKE_DCGM_VERSION(dcgmDeviceMemoryUsage_v1, 1)
```

Version 1 for `dcgmDeviceMemoryUsage_v1`

```
#define dcgmDeviceMemoryUsage_version  
dcgmDeviceMemoryUsage_version1
```

Latest version for `dcgmDeviceMemoryUsage_t`

```
#define dcgmDeviceVgpuUtilInfo_version1  
MAKE_DCGM_VERSION(dcgmDeviceVgpuUtilInfo_v1, 1)
```

Version 1 for `dcgmDeviceVgpuUtilInfo_v1`

```
#define dcgmDeviceVgpuUtilInfo_version  
dcgmDeviceVgpuUtilInfo_version1
```

Latest version for `dcgmDeviceVgpuUtilInfo_t`

```
#define dcgmDeviceEncStats_version1  
MAKE_DCGM_VERSION(dcgmDeviceEncStats_v1, 1)
```

Version 1 for `dcgmDeviceEncStats_v1`

```
#define dcgmDeviceEncStats_version  
dcgmDeviceEncStats_version1
```

Latest version for `dcgmDeviceEncStats_t`

```
#define dcgmDeviceFbcStats_version1  
MAKE_DCGM_VERSION(dcgmDeviceFbcStats_v1, 1)
```

Version 1 for `dcgmDeviceFbcStats_v1`

```
#define dcgmDeviceFbcStats_version  
dcgmDeviceFbcStats_version1
```

Latest version for `dcgmDeviceEncStats_t`

```
#define dcgmDeviceFbcSessionInfo_version1  
MAKE_DCGM_VERSION(dcgmDeviceFbcSessionInfo_v1, 1)
```

Version 1 for `dcgmDeviceFbcSessionInfo_v1`

```
#define dcgmDeviceFbcSessionInfo_version  
dcgmDeviceFbcSessionInfo_version1
```

Latest version for `dcgmDeviceFbcSessionInfo_t`

```
#define dcgmDeviceFbcSessions_version1  
MAKE_DCGM_VERSION(dcgmDeviceFbcSessions_v1, 1)
```

Version 1 for `dcgmDeviceFbcSessions_v1`

```
#define dcgmDeviceFbcSessions_version  
dcgmDeviceFbcSessions_version1
```

Latest version for `dcgmDeviceFbcSessions_t`

```
#define dcgmDeviceVgpuEncSessions_version1  
MAKE_DCGM_VERSION(dcgmDeviceVgpuEncSessions_v1,  
1)
```

Version 1 for `dcgmDeviceVgpuEncSessions_v1`

```
#define dcgmDeviceVgpuEncSessions_version  
dcgmDeviceVgpuEncSessions_version1
```

Latest version for `dcgmDeviceVgpuEncSessions_t`

```
#define dcgmDeviceVgpuProcessUtilInfo_version1  
MAKE_DCGM_VERSION(dcgmDeviceVgpuProcessUtilInfo_v1,  
1)
```

Version 1 for `dcgmDeviceVgpuProcessUtilInfo_v1`

```
#define dcgmDeviceVgpuProcessUtilInfo_version  
dcgmDeviceVgpuProcessUtilInfo_version1
```

Latest version for `dcgmDeviceVgpuProcessUtilInfo_t`

```
#define dcgmDeviceVgpuids_version1  
MAKE_DCGM_VERSION(dcgmDeviceVgpuids_v1, 1)
```

Version 1 for `dcgmDeviceVgpuIds_v1`

```
#define dcgmDeviceVgpuids_version  
dcgmDeviceVgpuids_version1
```

Latest version for `dcgmDeviceVgpuIds_t`

```
#define dcgmDeviceVgpuTypeInfo_version1  
MAKE_DCGM_VERSION(dcgmDeviceVgpuTypeInfo_v1, 1)
```

Version 1 for `dcgmDeviceVgpuTypeInfo_v1`

```
#define dcgmDeviceVgpuTypeInfo_version  
dcgmDeviceVgpuTypeInfo_version1
```

Latest version for `dcgmDeviceVgpuTypeInfo_t`

```
#define dcgmDeviceAttributes_version1  
MAKE_DCGM_VERSION(dcgmDeviceAttributes_v1, 1)
```

Version 1 for `dcgmDeviceAttributes_v1`

```
#define dcgmDeviceAttributes_version  
dcgmDeviceAttributes_version1
```

Latest version for `dcgmDeviceAttributes_t`

```
#define DCGM_MAX_VGPU_TYPES_PER_PGPU 32
```

Maximum number of vGPU types per physical GPU

```
#define dcgmVgpuDeviceAttributes_version6  
MAKE_DCGM_VERSION(dcgmVgpuDeviceAttributes_v6, 1)
```

Version 6 for `dcgmVgpuDeviceAttributes_t`

```
#define dcgmVgpuDeviceAttributes_version  
dcgmVgpuDeviceAttributes_version6
```

Latest version for `dcgmVgpuDeviceAttributes_t`

```
#define DCGM_DEVICE_UUID_BUFFER_SIZE 80
```

Represents the size of a buffer that holds string related to attributes specific to vGPU instance

```
#define dcgmVgpuInstanceAttributes_version1  
MAKE_DCGM_VERSION(dcgmVgpuInstanceAttributes_v1,  
1)
```

Version 1 for `dcgmVgpuInstanceAttributes_t`

```
#define dcgmVgpuInstanceAttributes_version  
dcgmVgpuInstanceAttributes_version1
```

Latest version for `dcgmVgpuInstanceAttributes_t`

```
#define dcgmConfig_version1  
MAKE_DCGM_VERSION(dcgmConfig_v1, 1)
```

Version 1 for `dcgmConfig_t`

```
#define dcgmConfig_version dcgmConfig_version1
```

Latest version for `dcgmConfig_t`

```
#define dcgmVgpuConfig_version1
```

```
MAKE_DCGM_VERSION(dcgmVgpuConfig_v1, 1)
```

Version 1 for `dcgmVgpuConfig_v1`

```
#define dcgmVgpuConfig_version
```

```
dcgmVgpuConfig_version1
```

Latest version for `dcgmVgpuConfig_t`

```
#define dcgmPolicy_version1
```

```
MAKE_DCGM_VERSION(dcgmPolicy_v1, 1)
```

Version 1 for `dcgmPolicy_v1`

```
#define dcgmPolicy_version dcgmPolicy_version1
```

Latest version for `dcgmPolicy_t`

```
#define dcgmPolicyCallbackResponse_version1
```

```
MAKE_DCGM_VERSION(dcgmPolicyCallbackResponse_v1,  
1)
```

Version 1 for `dcgmPolicyCallbackResponse_v1`

```
#define dcgmPolicyCallbackResponse_version
```

```
dcgmPolicyCallbackResponse_version1
```

Latest version for `dcgmPolicyCallbackResponse_t`

```
#define DCGM_MAX_BLOB_LENGTH 4096
```

Set above size of largest blob entry. Currently this is `dcgmDeviceVgpuTypeInfo_v1`.

```
#define dcgmFieldValue_version1
```

```
MAKE_DCGM_VERSION(dcgmFieldValue_v1, 1)
```

Version 1 for `dcgmFieldValue_v1`

```
#define dcgmFieldValue_version2  
MAKE_DCGM_VERSION(dcgmFieldValue_v2, 2)
```

Version 2 for dcgmFieldValue\_v2

```
#define DCGM_FV_FLAG_LIVE_DATA 0x00000001
```

Field value flags used by `dcgmEntitiesGetLatestValues`

```
#define DCGM_HEALTH_WATCH_COUNT_V1 10
```

For iterating through the `dcmHealthSystems_v1` enum.

```
#define DCGM_HEALTH_WATCH_COUNT_V2 12
```

For iterating through the `dcmHealthSystems_v2` enum.

```
#define dcgmHealthResponse_version1
```

```
MAKE_DCGM_VERSION(dcgmHealthResponse_v1, 1)
```

Version 1 for `dcgmHealthResponse_v1`

```
#define dcgmHealthResponse_version2
```

```
MAKE_DCGM_VERSION(dcgmHealthResponse_v2, 2)
```

Version 2 for `dcgmHealthResponse_v2`

```
#define dcgmHealthResponse_version3
```

```
MAKE_DCGM_VERSION(dcgmHealthResponse_v3, 3)
```

Version 3 for `dcgmHealthResponse_v3`

```
#define dcgmHealthResponse_version
```

```
dcgmHealthResponse_version3
```

Latest version for `dcgmHealthResponse_t`

```
#define dcgmPidInfo_version1
```

```
MAKE_DCGM_VERSION(dcgmPidInfo_v1, 1)
```

Version 1 for `dcgmPidInfo_v1`

```
#define dcgmPidInfo_version dcgmPidInfo_version1
```

Latest version for `dcgmPidInfo_t`

```
#define dcgmJobInfo_version2
```

```
MAKE_DCGM_VERSION(dcgmJobInfo_v2, 2)
```

Version 2 for `dcgmJobInfo_v2`

```
#define dcgmJobInfo_version dcgmJobInfo_version2
```

Latest version for `dcgmJobInfo_t`

```
#define dcgmRunningProcess_version1
```

```
MAKE_DCGM_VERSION(dcgmRunningProcess_v1, 1)
```

Version 1 for `dcgmRunningProcess_v1`

```
#define dcgmRunningProcess_version
```

```
dcgmRunningProcess_version1
```

Latest version for `dcgmRunningProcess_t`

```
#define dcgmDiagResponse_version3
```

```
MAKE_DCGM_VERSION(dcgmDiagResponse_v3, 3)
```

Version 3 for `dcgmDiagResponse_v3`

```
#define dcgmDiagResponse_version4
```

```
MAKE_DCGM_VERSION(dcgmDiagResponse_v4, 4)
```

Version 4 for `dcgmDiagResponse_v4`

```
#define dcgmDiagResponse_version5
```

```
MAKE_DCGM_VERSION(dcgmDiagResponse_v5, 5)
```

Version 5 for `dcgmDiagResponse_v5`

```
#define dcgmDiagResponse_version
```

```
dcgmDiagResponse_version5
```

Latest version for `dcgmDiagResponse_t`

```
#define dcgmDeviceTopology_version1  
MAKE_DCGM_VERSION(dcgmDeviceTopology_v1, 1)
```

Version 1 for dcgmDeviceTopology\_t

```
#define dcgmDeviceTopology_version  
dcgmDeviceTopology_version1
```

Latest version for dcgmDeviceTopology\_t

```
#define dcgmGroupTopology_version1  
MAKE_DCGM_VERSION(dcgmGroupTopology_v1, 1)
```

Version 1 for dcgmGroupTopology\_v1

```
#define dcgmGroupTopology_version  
dcgmGroupTopology_version1
```

Latest version for dcgmGroupTopology\_t

```
#define dcgmlIntrospectContext_version1  
MAKE_DCGM_VERSION(dcgmlIntrospectContext_v1, 1)
```

Version 1 for dcgmlIntrospectContext\_t

```
#define dcgmlIntrospectContext_version  
dcgmlIntrospectContext_version1
```

Latest version for dcgmlIntrospectContext\_t

```
#define dcgmlIntrospectFieldsExecTime_version1  
MAKE_DCGM_VERSION(dcgmlIntrospectFieldsExecTime_v1,  
1)
```

Version 1 for dcgmlIntrospectFieldsExecTime\_t

```
#define dcgmlIntrospectFieldsExecTime_version  
dcgmlIntrospectFieldsExecTime_version1
```

Latest version for dcgmlIntrospectFieldsExecTime\_t

```
#define dcgmlIntrospectFullFieldsExecTime_version1  
MAKE_DCGM_VERSION(dcgmlIntrospectFullFieldsExecTime_v1,  
1)
```

Version 1 for `dcgmIntrospectFullFieldsExecTime_t`

```
#define dcgmlIntrospectFullFieldsExecTime_version  
dcgmlIntrospectFullFieldsExecTime_version1
```

Latest version for `dcgmIntrospectFullFieldsExecTime_t`

```
#define dcgmlIntrospectMemory_version1  
MAKE_DCGM_VERSION(dcgmlIntrospectMemory_v1, 1)
```

Version 1 for `dcgmIntrospectMemory_t`

```
#define dcgmlIntrospectMemory_version  
dcgmlIntrospectMemory_version1
```

Latest version for `dcgmIntrospectMemory_t`

```
#define dcgmlIntrospectFullMemory_version1  
MAKE_DCGM_VERSION(dcgmlIntrospectFullMemory_v1, 1)
```

Version 1 for `dcgmIntrospectFullMemory_t`

```
#define dcgmlIntrospectFullMemory_version  
dcgmlIntrospectFullMemory_version1
```

Latest version for `dcgmIntrospectFullMemory_t`

```
#define dcgmlIntrospectCpuUtil_version1  
MAKE_DCGM_VERSION(dcgmlIntrospectCpuUtil_v1, 1)
```

Version 1 for `dcgmIntrospectCpuUtil_t`

```
#define dcgmlIntrospectCpuUtil_version  
dcgmlIntrospectCpuUtil_version1
```

Latest version for `dcgmIntrospectCpuUtil_t`

```
#define dcgmRunDiag_version1  
MAKE_DCGM_VERSION(dcgmRunDiag_v1, 1)
```

Version 1 for dcgmRunDiag\_t

```
#define dcgmRunDiag_version2  
MAKE_DCGM_VERSION(dcgmRunDiag_v2, 2)
```

Version 2 for dcgmRunDiag\_t

```
#define dcgmRunDiag_version3  
MAKE_DCGM_VERSION(dcgmRunDiag_v3, 3)
```

Version 3 for dcgmRunDiag\_t

```
#define dcgmRunDiag_version4  
MAKE_DCGM_VERSION(dcgmRunDiag_v4, 4)
```

Version 4 for dcgmRunDiag\_t

```
#define dcgmRunDiag_version5  
MAKE_DCGM_VERSION(dcgmRunDiag_v5, 5)
```

Version 5 for dcgmRunDiag\_t

```
#define dcgmRunDiag_version dcgmRunDiag_version5
```

Latest version for dcgmRunDiag\_t

```
#define DCGM_GEGE_FLAG_ONLY_SUPPORTED  
0x00000001
```

This mimics the behavior of dcgmGetAllSupportedDevices().

Flags for dcgmGetEntityGroupEntities's flags parameter Only return entities that are supported by DCGM.

```
#define dcgmNvLinkStatus_version1  
MAKE_DCGM_VERSION(dcgmNvLinkStatus_v1, 1)
```

Version 1 of dcgmNvLinkStatus

```
#define DCGM_MODULE_STATUSES_CAPACITY 16
```

This is larger than DcgmModuleIdCount so we can add modules without versioning this request.

```
#define dcgmModuleGetStatuses_version1  
MAKE_DCGM_VERSION(dcgmModuleGetStatuses_v1, 1)
```

Version 1 of dcgmModuleGetStatuses

```
#define DCGM_PROF_MAX_NUM_GROUPS 10
```

Maximum number of metric ID groups that can exist in DCGM.

Structure to return all of the profiling metric groups that are available for the given groupId.

```
#define DCGM_PROF_MAX_FIELD_IDS_PER_GROUP 8
```

Maximum number of field IDs that can be in a single DCGM profiling metric group.

```
#define dcgmProfGetMetricGroups_version2
```

```
MAKE_DCGM_VERSION(dcgmProfGetMetricGroups_v2, 2)
```

Version 1 of dcgmProfGetMetricGroups\_t

```
#define dcgmProfWatchFields_version1
```

```
MAKE_DCGM_VERSION(dcgmProfWatchFields_v1, 1)
```

Version 1 of dcgmProfWatchFields\_v1

```
#define dcgmProfUnwatchFields_version1
```

```
MAKE_DCGM_VERSION(dcgmProfUnwatchFields_v1, 1)
```

Version 1 of dcgmProfUnwatchFields\_v1

```
#define dcgmVersionInfo_version1
```

```
MAKE_DCGM_VERSION(dcgmVersionInfo_v1, 1)
```

Version 1 of dcgmVersionInfo\_v1;

## 2.16. Field Types

Field Types are a single byte.

**#define DCGM\_FT\_BINARY 'b'**

Blob of binary data representing a structure

**#define DCGM\_FT\_DOUBLE 'd'**

8-byte double precision

**#define DCGM\_FT\_INT64 'i'**

8-byte signed integer

**#define DCGM\_FT\_STRING 's'**

Null-terminated ASCII Character string

**#define DCGM\_FT\_TIMESTAMP 't'**

8-byte signed integer usec since 1970

## 2.17. Field Scope

Represents field association with entity scope or global scope.

**#define DCGM\_FS\_GLOBAL 0**

Field is global (ex: driver version)

**#define DCGM\_FS\_ENTITY 1**

Field is associated with an entity (GPU, VGPU...etc)

**#define DCGM\_FS\_DEVICE DCGM\_FS\_ENTITY**

Field is associated with a device. Deprecated. Use DCGM\_FS\_ENTITY

```
#define DCGM_CUDA_COMPUTE_CAPABILITY_MAJOR
((uint64_t)(x) & 0xFFFF0000)
```

DCGM\_FI\_DEV\_CUDA\_COMPUTE\_CAPABILITY is 16 bits of major version followed by 16 bits of the minor version. These macros separate the two.

```
#define DCGM_CLOCKS_THROTTLE_REASON_GPU_IDLE
0x0000000000000001LL
```

DCGM\_FI\_DEV\_CLOCK\_THROTTLE\_REASONS is a bitmap of why the clock is throttled. These macros are masks for relevant throttling, and are a 1:1 map to the NVML reasons documented in nvml.h. The notes for the header are copied below:

Nothing is running on the GPU and the clocks are dropping to Idle state



This limiter may be removed in a later release

```
#define
DCGM_CLOCKS_THROTTLE_REASON_CLOCKS_SETTING
0x0000000000000002LL
```

GPU clocks are limited by current setting of applications clocks

```
#define
DCGM_CLOCKS_THROTTLE_REASON_SW_POWER_CAP
0x0000000000000004LL
```

SW Power Scaling algorithm is reducing the clocks below requested clocks

```
#define
DCGM_CLOCKS_THROTTLE_REASON_HW_SLOWDOWN
0x0000000000000008LL
```

HW Slowdown (reducing the core clocks by a factor of 2 or more) is engaged

This is an indicator of:

- ▶ temperature being too high
- ▶ External Power Brake Assertion is triggered (e.g. by the system power supply)
- ▶ Power draw is too high and Fast Trigger protection is reducing the clocks
- ▶ May be also reported during PState or clock change

- ▶ This behavior may be removed in a later release.

```
#define
DCGM_CLOCKS_THROTTLE_REASON_SYNC_BOOST
0x0000000000000010LL
```

#### Sync Boost

This GPU has been added to a Sync boost group with nvidia-smi or DCGM in order to maximize performance per watt. All GPUs in the sync boost group will boost to the minimum possible clocks across the entire group. Look at the throttle reasons for other GPUs in the system to see why those GPUs are holding this one at lower clocks.

```
#define
DCGM_CLOCKS_THROTTLE_REASON_SW_THERMAL
0x0000000000000020LL
```

#### SW Thermal Slowdown

This is an indicator of one or more of the following:

- ▶ Current GPU temperature above the GPU Max Operating Temperature
- ▶ Current memory temperature above the Memory Max Operating Temperature

```
#define
DCGM_CLOCKS_THROTTLE_REASON_HW_THERMAL
0x0000000000000040LL
```

HW Thermal Slowdown (reducing the core clocks by a factor of 2 or more) is engaged

This is an indicator of:

- ▶ temperature being too high

```
#define
DCGM_CLOCKS_THROTTLE_REASON_HW_POWER_BRAKE
0x0000000000000080LL
```

HW Power Brake Slowdown (reducing the core clocks by a factor of 2 or more) is engaged

This is an indicator of:

- ▶ External Power Brake Assertion being triggered (e.g. by the system power supply)

```
#define
DCGM_CLOCKS_THROTTLE_REASON_DISPLAY_CLOCKS
0x0000000000000000100LL
```

GPU clocks are limited by current setting of Display clocks

## 2.18. Field Entity

Represents field association with a particular entity

### **enum dcgm\_field\_entity\_group\_t**

Enum of possible field entity groups

#### **Values**

**DCGM\_FE\_NONE = 0**

**DCGM\_FE\_GPU**

Field is not associated with an entity. Field scope should be DCGM\_FS\_GLOBAL  
**DCGM\_FE\_VGPU**

Field is associated with a GPU entity

**DCGM\_FE\_SWITCH**

Field is associated with a VGPU entity

**DCGM\_FE\_COUNT**

Field is associated with a Switch entity Number of elements in this enumeration.

Keep this entry last

### **typedef unsigned int dcgm\_field\_eid\_t**

Represents an identifier for an entity within a field entity. For instance, this is the gpuId for DCGM\_FE\_GPU.

## 2.19. Field Identifiers

Field Identifiers

## DcgmFieldGetById (unsigned short fieldId)

### Parameters

#### fieldId

IN: One of the field IDs (DCGM\_FI\_?)

### Returns

0 On Failure > 0 Pointer to field metadata structure if found.

### Description

Get a pointer to the metadata for a field by its field ID. See DCGM\_FI\_? for a list of field IDs.

## DcgmFieldGetByTag (char \*tag)

### Parameters

#### tag

IN: Tag for the field of interest

### Returns

0 On failure or not found > 0 Pointer to field metadata structure if found

### Description

Get a pointer to the metadata for a field by its field tag.

## DcgmFieldsInit (void)

### Returns

0 On success <0 On error

### Description

Initialize the DcgmFields module. Call this once from inside your program

## DcgmFieldsTerm (void)

### Returns

0 On success <0 On error

**Description**

Terminates the DcgmFields module. Call this once from inside your program

```
char *DcgmFieldsGetEntityGroupString  
(dcgm_field_entity_group_t entityGroupId)
```

**Description**

Get the string version of a entityGroupId

Returns Pointer to a string like GPU/NvSwitch..etc Null on error

```
#define DCGM_FI_UNKNOWN 0
```

NULL field

```
#define DCGM_FI_DRIVER_VERSION 1
```

Driver Version

```
#define DCGM_FI_DEV_COUNT 4
```

Number of Devices on the node

```
#define DCGM_FI_DEV_NAME 50
```

Name of the GPU device

```
#define DCGM_FI_DEV_BRAND 51
```

Device Brand

```
#define DCGM_FI_DEV_NVML_INDEX 52
```

NVML index of this GPU

```
#define DCGM_FI_DEV_SERIAL 53
```

Device Serial Number

```
#define DCGM_FI_DEV_UUID 54
```

UUID corresponding to the device

**#define DCGM\_FI\_DEV\_MINOR\_NUMBER 55**  
Device node minor number /dev/nvidia#

**#define DCGM\_FI\_DEV\_OEM\_INFOROM\_VER 56**  
OEM inforom version

**#define DCGM\_FI\_DEV\_PCI\_BUSID 57**  
PCI attributes for the device

**#define DCGM\_FI\_DEV\_PCI\_COMBINED\_ID 58**  
The combined 16-bit device id and 16-bit vendor id

**#define DCGM\_FI\_DEV\_PCI\_SUBSYS\_ID 59**  
The 32-bit Sub System Device ID

**#define DCGM\_FI\_GPU\_TOPOLOGY\_PCI 60**  
Topology of all GPUs on the system via PCI (static)

**#define DCGM\_FI\_GPU\_TOPOLOGY\_NVLINK 61**  
Topology of all GPUs on the system via NVLINK (static)

**#define DCGM\_FI\_GPU\_TOPOLOGY\_AFFINITY 62**  
Affinity of all GPUs on the system (static)

**#define DCGM\_FI\_DEV\_CUDA\_COMPUTE\_CAPABILITY 63**  
Cuda compute capability for the device. The major version is the upper 32 bits and the minor version is the lower 32 bits.

**#define DCGM\_FI\_DEV\_COMPUTE\_MODE 65**  
Compute mode for the device

**#define DCGM\_FI\_DEV\_CPU\_AFFINITY\_0 70**  
Device CPU affinity. part 1/8 = cpus 0 - 63

**#define DCGM\_FI\_DEV\_CPU\_AFFINITY\_1 71**

Device CPU affinity. part 1/8 = cpus 64 - 127

**#define DCGM\_FI\_DEV\_CPU\_AFFINITY\_2 72**

Device CPU affinity. part 2/8 = cpus 128 - 191

**#define DCGM\_FI\_DEV\_CPU\_AFFINITY\_3 73**

Device CPU affinity. part 3/8 = cpus 192 - 255

**#define DCGM\_FI\_DEV\_ECC\_INFOROM\_VER 80**

ECC inforom version

**#define DCGM\_FI\_DEV\_POWER\_INFOROM\_VER 81**

Power management object inforom version

**#define DCGM\_FI\_DEV\_INFOROM\_IMAGE\_VER 82**

Inforom image version

**#define DCGM\_FI\_DEV\_INFOROM\_CONFIG\_CHECK 83**

Inforom configuration checksum

**#define DCGM\_FI\_DEV\_INFOROM\_CONFIG\_VALID 84**

Reads the infoROM from the flash and verifies the checksums

**#define DCGM\_FI\_DEV\_VBIOS\_VERSION 85**

VBIOS version of the device

**#define DCGM\_FI\_DEV\_BAR1\_TOTAL 90**

Total BAR1 of the GPU in MB

**#define DCGM\_FI\_SYNC\_BOOST 91**

Sync boost settings on the node

**#define DCGM\_FI\_DEV\_BAR1\_USED 92**

Used BAR1 of the GPU in MB

**#define DCGM\_FI\_DEV\_BAR1\_FREE 93**

Free BAR1 of the GPU in MB

**#define DCGM\_FI\_DEV\_SM\_CLOCK 100**

SM clock for the device

**#define DCGM\_FI\_DEV\_MEM\_CLOCK 101**

Memory clock for the device

**#define DCGM\_FI\_DEV\_VIDEO\_CLOCK 102**

Video encoder/decoder clock for the device

**#define DCGM\_FI\_DEV\_APP\_SM\_CLOCK 110**

SM Application clocks

**#define DCGM\_FI\_DEV\_APP\_MEM\_CLOCK 111**

Memory Application clocks

**#define DCGM\_FI\_DEV\_CLOCK\_THROTTLE\_REASONS 112**

Current clock throttle reasons (bitmask of DCGM\_CLOCKS\_THROTTLE\_REASON\_\*)

**#define DCGM\_FI\_DEV\_MAX\_SM\_CLOCK 113**

Maximum supported SM clock for the device

**#define DCGM\_FI\_DEV\_MAX\_MEM\_CLOCK 114**

Maximum supported Memory clock for the device

**#define DCGM\_FI\_DEV\_MAX\_VIDEO\_CLOCK 115**

Maximum supported Video encoder/decoder clock for the device

**#define DCGM\_FI\_DEV\_AUTOBOOST 120**

Auto-boost for the device (1 = enabled. 0 = disabled)

**#define DCGM\_FI\_DEV\_SUPPORTED\_CLOCKS 130**

Supported clocks for the device

**#define DCGM\_FI\_DEV\_MEMORY\_TEMP 140**

Memory temperature for the device

**#define DCGM\_FI\_DEV\_GPU\_TEMP 150**

Current temperature readings for the device, in degrees C

**#define DCGM\_FI\_DEV\_POWER\_USAGE 155**

Power usage for the device in Watts

**#define DCGM\_FI\_DEV\_TOTAL\_ENERGY\_CONSUMPTION 156**

Total energy consumption for the GPU in mJ since the driver was last reloaded

**#define DCGM\_FI\_DEV\_SLOWDOWN\_TEMP 158**

Slowdown temperature for the device

**#define DCGM\_FI\_DEV\_SHUTDOWN\_TEMP 159**

Shutdown temperature for the device

**#define DCGM\_FI\_DEV\_POWER\_MGMT\_LIMIT 160**

Current Power limit for the device

**#define DCGM\_FI\_DEV\_POWER\_MGMT\_LIMIT\_MIN 161**

Minimum power management limit for the device

**#define DCGM\_FI\_DEV\_POWER\_MGMT\_LIMIT\_MAX 162**

Maximum power management limit for the device

**#define DCGM\_FI\_DEV\_POWER\_MGMT\_LIMIT\_DEF 163**

Default power management limit for the device

**#define DCGM\_FI\_DEV\_ENFORCED\_POWER\_LIMIT 164**

Effective power limit that the driver enforces after taking into account all limiters

**#define DCGM\_FI\_DEV\_PSTATE 190**

Performance state (P-State) 0-15. 0=highest

**#define DCGM\_FI\_DEV\_FAN\_SPEED 191**

Fan speed for the device in percent 0-100

**#define DCGM\_FI\_DEV\_PCIE\_TX\_THROUGHPUT 200**

PCIe Tx utilization information

**#define DCGM\_FI\_DEV\_PCIE\_RX\_THROUGHPUT 201**

PCIe Rx utilization information

**#define DCGM\_FI\_DEV\_PCIE\_REPLAY\_COUNTER 202**

PCIe replay counter

**#define DCGM\_FI\_DEV\_GPU\_UTIL 203**

GPU Utilization

**#define DCGM\_FI\_DEV\_MEM\_COPY\_UTIL 204**

Memory Utilization

**#define DCGM\_FI\_DEV\_ACCOUNTING\_DATA 205**

Process accounting stats.

This field is only supported when the host engine is running as root unless you enable accounting ahead of time. Accounting mode can be enabled by running "nvidia-smi -am 1" as root on the same node the host engine is running on.

**#define DCGM\_FI\_DEV\_ENC\_UTIL 206**

Encoder Utilization

**#define DCGM\_FI\_DEV\_DEC\_UTIL 207**

Decoder Utilization

**#define DCGM\_FI\_DEV\_MEM\_COPY\_UTIL\_SAMPLES 210**

Memory utilization samples

**#define DCGM\_FI\_DEV\_GRAPHICS\_PIDS 220**

Graphics processes running on the GPU.

**#define DCGM\_FI\_DEV\_COMPUTE\_PIDS 221**

Compute processes running on the GPU.

**#define DCGM\_FI\_DEV\_XID\_ERRORS 230**

XID errors. The value is the specific XID error

**#define DCGM\_FI\_DEV\_PCIE\_MAX\_LINK\_GEN 235**

PCIe Max Link Generation

**#define DCGM\_FI\_DEV\_PCIE\_MAX\_LINK\_WIDTH 236**

PCIe Max Link Width

**#define DCGM\_FI\_DEV\_PCIE\_LINK\_GEN 237**

PCIe Current Link Generation

**#define DCGM\_FI\_DEV\_PCIE\_LINK\_WIDTH 238**

PCIe Current Link Width

**#define DCGM\_FI\_DEV\_POWER\_VIOLATION 240**

Power Violation time in usec

**#define DCGM\_FI\_DEV\_THERMAL\_VIOLATION 241**

Thermal Violation time in usec

**#define DCGM\_FI\_DEV\_SYNC\_BOOST\_VIOLATION 242**

Sync Boost Violation time in usec

**#define DCGM\_FI\_DEV\_BOARD\_LIMIT\_VIOLATION 243**

Board violation limit.

**#define DCGM\_FI\_DEV\_LOW\_UTIL\_VIOLATION 244**

Low utilisation violation limit.

**#define DCGM\_FI\_DEV\_RELIABILITY\_VIOLATION 245**

Reliability violation limit.

**#define DCGM\_FI\_DEV\_TOTAL\_APP\_CLOCKS\_VIOLATION 246**

App clock violation limit.

**#define DCGM\_FI\_DEV\_TOTAL\_BASE\_CLOCKS\_VIOLATION 247**

Base clock violation limit.

**#define DCGM\_FI\_DEV\_FB\_TOTAL 250**

Total Frame Buffer of the GPU in MB

**#define DCGM\_FI\_DEV\_FB\_FREE 251**

Free Frame Buffer in MB

**#define DCGM\_FI\_DEV\_FB\_USED 252**

Used Frame Buffer in MB

**#define DCGM\_FI\_DEV\_ECC\_CURRENT 300**

Current ECC mode for the device

**#define DCGM\_FI\_DEV\_ECC\_PENDING 301**

Pending ECC mode for the device

**#define DCGM\_FI\_DEV\_ECC\_SBE\_VOL\_TOTAL 310**

Total single bit volatile ECC errors

**#define DCGM\_FI\_DEV\_ECC\_DBE\_VOL\_TOTAL 311**

Total double bit volatile ECC errors

**#define DCGM\_FI\_DEV\_ECC\_SBE\_AGG\_TOTAL 312**

Total single bit aggregate (persistent) ECC errors Note: monotonically increasing

**#define DCGM\_FI\_DEV\_ECC\_DBE\_AGG\_TOTAL 313**

Total double bit aggregate (persistent) ECC errors Note: monotonically increasing

**#define DCGM\_FI\_DEV\_ECC\_SBE\_VOL\_L1 314**

L1 cache single bit volatile ECC errors

**#define DCGM\_FI\_DEV\_ECC\_DBE\_VOL\_L1 315**

L1 cache double bit volatile ECC errors

**#define DCGM\_FI\_DEV\_ECC\_SBE\_VOL\_L2 316**

L2 cache single bit volatile ECC errors

**#define DCGM\_FI\_DEV\_ECC\_DBE\_VOL\_L2 317**

L2 cache double bit volatile ECC errors

**#define DCGM\_FI\_DEV\_ECC\_SBE\_VOL\_DEV 318**

Device memory single bit volatile ECC errors

**#define DCGM\_FI\_DEV\_ECC\_DBE\_VOL\_DEV 319**

Device memory double bit volatile ECC errors

**#define DCGM\_FI\_DEV\_ECC\_SBE\_VOL\_REG 320**

Register file single bit volatile ECC errors

**#define DCGM\_FI\_DEV\_ECC\_DBE\_VOL\_REG 321**

Register file double bit volatile ECC errors

**#define DCGM\_FI\_DEV\_ECC\_SBE\_VOL\_TEX 322**

Texture memory single bit volatile ECC errors

**#define DCGM\_FI\_DEV\_ECC\_DBE\_VOL\_TEX 323**

Texture memory double bit volatile ECC errors

**#define DCGM\_FI\_DEV\_ECC\_SBE\_AGG\_L1 324**

L1 cache single bit aggregate (persistent) ECC errors Note: monotonically increasing

**#define DCGM\_FI\_DEV\_ECC\_DBE\_AGG\_L1 325**

L1 cache double bit aggregate (persistent) ECC errors Note: monotonically increasing

**#define DCGM\_FI\_DEV\_ECC\_SBE\_AGG\_L2 326**

L2 cache single bit aggregate (persistent) ECC errors Note: monotonically increasing

**#define DCGM\_FI\_DEV\_ECC\_DBE\_AGG\_L2 327**

L2 cache double bit aggregate (persistent) ECC errors Note: monotonically increasing

**#define DCGM\_FI\_DEV\_ECC\_SBE\_AGG\_DEV 328**

Device memory single bit aggregate (persistent) ECC errors Note: monotonically increasing

**#define DCGM\_FI\_DEV\_ECC\_DBE\_AGG\_DEV 329**

Device memory double bit aggregate (persistent) ECC errors Note: monotonically increasing

**#define DCGM\_FI\_DEV\_ECC\_SBE\_AGG\_REG 330**

Register File single bit aggregate (persistent) ECC errors Note: monotonically increasing

**#define DCGM\_FI\_DEV\_ECC\_DBE\_AGG\_REG 331**

Register File double bit aggregate (persistent) ECC errors Note: monotonically increasing

**#define DCGM\_FI\_DEV\_ECC\_SBE\_AGG\_TEX 332**

Texture memory single bit aggregate (persistent) ECC errors Note: monotonically increasing

**#define DCGM\_FI\_DEV\_ECC\_DBE\_AGG\_TEX 333**

Texture memory double bit aggregate (persistent) ECC errors Note: monotonically increasing

**#define DCGM\_FI\_DEV\_RETIRED\_SBE 390**

Number of retired pages because of single bit errors Note: monotonically increasing

**#define DCGM\_FI\_DEV\_RETIRED\_DBE 391**

Number of retired pages because of double bit errors Note: monotonically increasing

**#define DCGM\_FI\_DEV\_RETIRED\_PENDING 392**

Number of pages pending retirement

**#define DCGM\_FI\_DEV\_VIRTUAL\_MODE 500**

Virtualization Mode corresponding to the GPU

**#define DCGM\_FI\_DEV\_SUPPORTED\_TYPE\_INFO 501**

Includes Count and Static info of vGPU types supported on a device

**#define DCGM\_FI\_DEV\_CREATABLE\_VGPU\_TYPE\_IDS 502**

Includes Count and currently Creatable vGPU types on a device

**#define DCGM\_FI\_DEV\_VGPU\_INSTANCE\_IDS 503**

Includes Count and currently Active vGPU Instances on a device

**#define DCGM\_FI\_DEV\_VGPU\_UTILIZATIONS 504**

Utilization values for vGPUs running on the device

**#define**

**DCGM\_FI\_DEV\_VGPU\_PER\_PROCESS\_UTILIZATION 505**

Utilization values for processes running within vGPU VMs using the device

**#define DCGM\_FI\_DEV\_ENC\_STATS 506**

Current encoder statistics for a given device

**#define DCGM\_FI\_DEV\_FBC\_STATS 507**

Statistics of current active frame buffer capture sessions on a given device

**#define DCGM\_FI\_DEV\_FBC\_SESSIONS\_INFO 508**

Information about active frame buffer capture sessions on a target device

**#define DCGM\_FI\_DEV\_VGPU\_VM\_ID 520**

VM ID of the vGPU instance

**#define DCGM\_FI\_DEV\_VGPU\_VM\_NAME 521**

VM name of the vGPU instance

**#define DCGM\_FI\_DEV\_VGPU\_TYPE 522**

vGPU type of the vGPU instance

**#define DCGM\_FI\_DEV\_VGPU\_UUID 523**

UUID of the vGPU instance

**#define DCGM\_FI\_DEV\_VGPU\_DRIVER\_VERSION 524**

Driver version of the vGPU instance

**#define DCGM\_FI\_DEV\_VGPU\_MEMORY\_USAGE 525**

Memory usage of the vGPU instance

**#define DCGM\_FI\_DEV\_VGPU\_LICENSE\_STATUS 526**

License status of the vGPU instance

**#define DCGM\_FI\_DEV\_VGPU\_FRAME\_RATE\_LIMIT 527**

Frame rate limit of the vGPU instance

**#define DCGM\_FI\_DEV\_VGPU\_ENC\_STATS 528**

Current encoder statistics of the vGPU instance

**#define DCGM\_FI\_DEV\_VGPU\_ENC\_SESSIONS\_INFO 529**

Information about all active encoder sessions on the vGPU instance

**#define DCGM\_FI\_DEV\_VGPU\_FBC\_STATS 530**

Statistics of current active frame buffer capture sessions on the vGPU instance

**#define DCGM\_FI\_DEV\_VGPU\_FBC\_SESSIONS\_INFO 531**

Information about active frame buffer capture sessions on the vGPU instance

**#define DCGM\_FI\_FIRST\_VGPU\_FIELD\_ID 520**

Starting field ID of the vGPU instance

**#define DCGM\_FI\_LAST\_VGPU\_FIELD\_ID 570**

Last field ID of the vGPU instance

**#define DCGM\_FI\_MAX\_VGPU\_FIELDS**

**DCGM\_FI\_LAST\_VGPU\_FIELD\_ID -**

**DCGM\_FI\_FIRST\_VGPU\_FIELD\_ID**

For now max vGPU field Ids taken as difference of DCGM\_FI\_LAST\_VGPU\_FIELD\_ID and DCGM\_FI\_FIRST\_VGPU\_FIELD\_ID i.e. 50

```
#define DCGM_FI_INTERNAL_FIELDS_0_START 600
```

Starting ID for all the internal fields

```
#define DCGM_FI_INTERNAL_FIELDS_0_END 699
```

Last ID for all the internal fields

NVSwitch entity field IDs start here.

NVSwitch latency bins for port 0

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P00  
700
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P00  
701
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P00  
702
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P00  
703
```

Max latency bin

NVSwitch latency bins for port 1

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P01  
704
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P01  
705
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P01  
706
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P01  
707
```

Max latency bin

NVSwitch latency bins for port 2

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P02  
708
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P02  
709
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P02  
710
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P02  
711
```

Max latency bin

NVSwitch latency bins for port 3

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P03  
712
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P03  
713
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P03  
714
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P03  
715
```

Max latency bin

NVSwitch latency bins for port 4

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P04  
716
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P04  
717
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P04  
718
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P04  
719
```

Max latency bin

NVSwitch latency bins for port 5

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P05  
720
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P05  
721
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P05  
722
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P05  
723
```

Max latency bin

NVSwitch latency bins for port 6

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P06  
724
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P06  
725
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P06  
726
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P06  
727
```

Max latency bin

NVSwitch latency bins for port 7

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P07  
728
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P07  
729
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P07  
730
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P07  
731
```

Max latency bin

NVSwitch latency bins for port 8

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P08  
732
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P08  
733
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P08  
734
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P08  
735
```

Max latency bin

NVSwitch latency bins for port 9

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P09  
736
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P09  
737
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P09  
738
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P09  
739
```

Max latency bin

NVSwitch latency bins for port 10

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P10  
740
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P10  
741
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P10  
742
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P10  
743
```

Max latency bin

NVSwitch latency bins for port 11

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P11  
744
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P11  
745
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P11  
746
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P11  
747
```

Max latency bin

NVSwitch latency bins for port 12

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P12  
748
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P12  
749
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P12  
750
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P12  
751
```

Max latency bin

NVSwitch latency bins for port 13

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P13  
752
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P13  
753
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P13  
754
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P13  
755
```

Max latency bin

NVSwitch latency bins for port 14

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P14  
756
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P14  
757
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P14  
758
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P14  
759
```

Max latency bin

NVSwitch latency bins for port 15

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P15  
760
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P15  
761
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P15  
762
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P15  
763
```

Max latency bin

NVSwitch latency bins for port 16

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P16  
764
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P16  
765
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P16  
766
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P16  
767
```

Max latency bin

NVSwitch latency bins for port 17

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_LOW_P17  
768
```

Low latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MED_P17  
769
```

Medium latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_HIGH_P17  
770
```

High latency bin

```
#define DCGM_FI_DEV_NVSWITCH_LATENCY_MAX_P17  
771
```

Max latency bin

NVSwitch Tx and Rx Counter 0 for each port

By default, Counter 0 counts bytes.

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P00 780
```

NVSwitch Tx Bandwidth Counter 0 for port 0

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P00 781
```

NVSwitch Rx Bandwidth Counter 0 for port 0

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P01 782
```

NVSwitch Tx Bandwidth Counter 0 for port 1

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P01 783
```

NVSwitch Rx Bandwidth Counter 0 for port 1

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P02 784
```

NVSwitch Tx Bandwidth Counter 0 for port 2

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P02 785
```

NVSwitch Rx Bandwidth Counter 0 for port 2

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P03 786
```

NVSwitch Tx Bandwidth Counter 0 for port 3

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P03 787
```

NVSwitch Rx Bandwidth Counter 0 for port 3

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P04 788
```

NVSwitch Tx Bandwidth Counter 0 for port 4

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P04 789
```

NVSwitch Rx Bandwidth Counter 0 for port 4

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P05 790
```

NVSwitch Tx Bandwidth Counter 0 for port 5

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P05 791
```

NVSwitch Rx Bandwidth Counter 0 for port 5

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P06 792
```

NVSwitch Tx Bandwidth Counter 0 for port 6

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P06 793
```

NVSwitch Rx Bandwidth Counter 0 for port 6

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P07 794
```

NVSwitch Tx Bandwidth Counter 0 for port 7

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P07 795
```

NVSwitch Rx Bandwidth Counter 0 for port 7

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P08 796
```

NVSwitch Tx Bandwidth Counter 0 for port 8

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P08 797
```

NVSwitch Rx Bandwidth Counter 0 for port 8

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P09 798
```

NVSwitch Tx Bandwidth Counter 0 for port 9

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P09 799
```

NVSwitch Rx Bandwidth Counter 0 for port 9

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P10 800
```

NVSwitch Tx Bandwidth Counter 0 for port 10

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P10 801
```

NVSwitch Rx Bandwidth Counter 0 for port 10

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P11 802
```

NVSwitch Tx Bandwidth Counter 0 for port 11

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P11 803
```

NVSwitch Rx Bandwidth Counter 0 for port 11

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P12 804
```

NVSwitch Tx Bandwidth Counter 0 for port 12

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P12 805
```

NVSwitch Rx Bandwidth Counter 0 for port 12

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P13 806
```

NVSwitch Tx Bandwidth Counter 0 for port 13

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P13 807
```

NVSwitch Rx Bandwidth Counter 0 for port 13

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P14 808
```

NVSwitch Tx Bandwidth Counter 0 for port 14

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P14 809
```

NVSwitch Rx Bandwidth Counter 0 for port 14

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P15 810
```

NVSwitch Tx Bandwidth Counter 0 for port 15

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P15 811
```

NVSwitch Rx Bandwidth Counter 0 for port 15

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P16 812
```

NVSwitch Tx Bandwidth Counter 0 for port 16

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P16 813
```

NVSwitch Rx Bandwidth Counter 0 for port 16

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_0_P17 814
```

NVSwitch Tx Bandwidth Counter 0 for port 17

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_0_P17 815
```

NVSwitch Rx Bandwidth Counter 0 for port 17

NVSwitch Tx and RX Bandwidth Counter 1 for each port

By default, Counter 1 counts packets.

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P00 820  
NVSwitch Tx Bandwidth Counter 1 for port 0  
  
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P00 821  
NVSwitch Rx Bandwidth Counter 1 for port 0  
  
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P01 822  
NVSwitch Tx Bandwidth Counter 1 for port 1  
  
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P01 823  
NVSwitch Rx Bandwidth Counter 1 for port 1  
  
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P02 824  
NVSwitch Tx Bandwidth Counter 1 for port 2  
  
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P02 825  
NVSwitch Rx Bandwidth Counter 1 for port 2  
  
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P03 826  
NVSwitch Tx Bandwidth Counter 1 for port 3  
  
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P03 827  
NVSwitch Rx Bandwidth Counter 1 for port 3
```

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P04 828
```

NVSwitch Tx Bandwidth Counter 1 for port 4

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P04 829
```

NVSwitch Rx Bandwidth Counter 1 for port 4

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P05 830
```

NVSwitch Tx Bandwidth Counter 1 for port 5

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P05 831
```

NVSwitch Rx Bandwidth Counter 1 for port 5

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P06 832
```

NVSwitch Tx Bandwidth Counter 1 for port 6

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P06 833
```

NVSwitch Rx Bandwidth Counter 1 for port 6

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P07 834
```

NVSwitch Tx Bandwidth Counter 1 for port 7

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P07 835
```

NVSwitch Rx Bandwidth Counter 1 for port 7

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P08 836
```

NVSwitch Tx Bandwidth Counter 1 for port 8

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P08 837
```

NVSwitch Rx Bandwidth Counter 1 for port 8

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P09 838
```

NVSwitch Tx Bandwidth Counter 1 for port 9

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P09 839
```

NVSwitch Rx Bandwidth Counter 1 for port 9

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P10 840
```

NVSwitch Tx Bandwidth Counter 0 for port 10

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P10 841
```

NVSwitch Rx Bandwidth Counter 1 for port 10

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P11 842
```

NVSwitch Tx Bandwidth Counter 1 for port 11

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P11 843
```

NVSwitch Rx Bandwidth Counter 1 for port 11

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P12 844
```

NVSwitch Tx Bandwidth Counter 1 for port 12

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P12 845
```

NVSwitch Rx Bandwidth Counter 1 for port 12

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P13 846
```

NVSwitch Tx Bandwidth Counter 0 for port 13

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P13 847
```

NVSwitch Rx Bandwidth Counter 1 for port 13

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P14 848
```

NVSwitch Tx Bandwidth Counter 1 for port 14

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P14 849
```

NVSwitch Rx Bandwidth Counter 1 for port 14

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P15 850
```

NVSwitch Tx Bandwidth Counter 1 for port 15

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P15 851
```

NVSwitch Rx Bandwidth Counter 1 for port 15

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P16 852
```

NVSwitch Tx Bandwidth Counter 1 for port 16

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P16 853
```

NVSwitch Rx Bandwidth Counter 1 for port 16

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_TX_1_P17 854
```

NVSwitch Tx Bandwidth Counter 1 for port 17

```
#define  
DCGM_FI_DEV_NVSWITCH_BANDWIDTH_RX_1_P17 855
```

NVSwitch Rx Bandwidth Counter 1 for port 17

NVSwitch error counters

```
#define DCGM_FI_DEV_NVSWITCH_FATAL_ERRORS 856
```

NVSwitch fatal error information. Note: value field indicates the specific SXid reported

```
#define DCGM_FI_DEV_NVSWITCH_NON_FATAL_ERRORS  
857
```

NVSwitch non fatal error information. Note: value field indicates the specific SXid reported

```
#define DCGM_FI_FIRST_NVSWITCH_FIELD_ID 700
```

Starting field ID of the NVSwitch instance

```
#define DCGM_FI_LAST_NVSWITCH_FIELD_ID 860
```

Last field ID of the NVSwitch instance

```
#define DCGM_FI_MAX_NVSWITCH_FIELDS
DCGM_FI_LAST_NVSWITCH_FIELD_ID -
DCGM_FI_FIRST_NVSWITCH_FIELD_ID + 1
```

For now max NVSwitch field Ids taken as difference  
of DCGM\_FI\_LAST\_NVSWITCH\_FIELD\_ID and  
DCGM\_FI\_FIRST\_NVSWITCH\_FIELD\_ID + 1 i.e. 200

**#define DCGM\_FI\_PROF\_GR\_ENGINE\_ACTIVE 1001**

Profiling Fields. These all start with DCGM\_FI\_PROF\_\* Ratio of time the graphics engine is active. The graphics engine is active if a graphics/compute context is bound and the graphics pipe or compute pipe is busy.

**#define DCGM\_FI\_PROF\_SM\_ACTIVE 1002**

The ratio of cycles an SM has at least 1 warp assigned (computed from the number of cycles and elapsed cycles)

**#define DCGM\_FI\_PROF\_SM\_OCCUPANCY 1003**

The ratio of number of warps resident on an SM. (number of resident as a ratio of the theoretical maximum number of warps per elapsed cycle)

**#define DCGM\_FI\_PROF\_PIPE\_TENSOR\_ACTIVE 1004**

The ratio of cycles the tensor (HMMA) pipe is active (off the peak sustained elapsed cycles)

**#define DCGM\_FI\_PROF\_DRAM\_ACTIVE 1005**

The ratio of cycles the device memory interface is active sending or receiving data.

**#define DCGM\_FI\_PROF\_PIPE\_FP64\_ACTIVE 1006**

Ratio of cycles the fp64 pipe is active.

**#define DCGM\_FI\_PROF\_PIPE\_FP32\_ACTIVE 1007**

Ratio of cycles the fp32 pipe is active.

**#define DCGM\_FI\_PROF\_PIPE\_FP16\_ACTIVE 1008**

Ratio of cycles the fp16 pipe is active. This does not include HMMA.

## #define DCGM\_FI\_PROF\_PCIE\_TX\_BYTES 1009

The number of bytes of active PCIe tx (transmit) data including both header and payload.

Note that this is from the perspective of the GPU, so copying data from device to host (DtoH) would be reflected in this metric.

## #define DCGM\_FI\_PROF\_PCIE\_RX\_BYTES 1010

The number of bytes of active PCIe rx (read) data including both header and payload.

Note that this is from the perspective of the GPU, so copying data from host to device (HtoD) would be reflected in this metric.

## #define DCGM\_FI\_PROF\_NVLINK\_TX\_BYTES 1011

The number of bytes of active NvLink tx (transmit) data including both header and payload.

## #define DCGM\_FI\_PROF\_NVLINK\_RX\_BYTES 1012

The number of bytes of active NvLink rx (read) data including both header and payload.

## #define DCGM\_FI\_MAX\_FIELDS 1013

1 greater than maximum fields above. This is the 1 greater than the maximum field id that could be allocated

## 2.20. DCGMAPI\_Admin\_ExecCtrl

`dcgmReturn_t dcgmUpdateAllFields (dcgmHandle_t pDcgmHandle, int waitForUpdate)`

### Parameters

#### **pDcgmHandle**

IN: DCGM Handle

#### **waitForUpdate**

IN: Whether or not to wait for the update loop to complete before returning to the caller 1=wait. 0=do not wait.

## Returns

- ▶ DCGM\_ST\_OK if the call was successful
- ▶ DCGM\_ST\_BADPARAM if waitForUpdate is invalid
- ▶ DCGM\_ST\_GENERIC\_ERROR if an unspecified DCGM error occurs

## Description

This method is used to tell the DCGM module to update all the fields being watched.

Note: If the operation mode was set to manual mode (DCGM\_OPERATION\_MODE\_MANUAL) during initialization ([dcgmInit](#)), this method must be caused periodically to allow field value watches the opportunity to gather samples.

## **dcgmReturn\_t dcgmPolicyTrigger (dcgmHandle\_t pDcgmHandle)**

### Parameters

#### pDcgmHandle

IN: DCGM Handle

## Returns

- ▶ DCGM\_ST\_OK If the call was successful
- ▶ DCGM\_ST\_GENERIC\_ERROR The policy manager was unable to perform another iteration.

## Description

Inform the policy manager loop to perform an iteration and trigger the callbacks of any registered functions. Callback functions will be called from a separate thread as the calling function.

Note: The GPU monitoring and management agent must call this method periodically if the operation mode is set to manual mode (DCGM\_OPERATION\_MODE\_MANUAL) during initialization ([dcgmInit](#)).

# Chapter 3. DATA STRUCTURES

Here are the data structures with brief descriptions:

`dcgm_field_meta_t`  
`dcgm_field_output_format_t`  
`dcgmClockSet_v1`  
`dcgmConfig_v1`  
`dcgmConfigPerfStateSettings_t`  
`dcgmConfigPowerLimit_t`  
`dcgmConnectV2Params_v1`  
`dcgmConnectV2Params_v2`  
`dcgmDeviceAttributes_v1`  
`dcgmDeviceEncStats_v1`  
`dcgmDeviceFbcSessionInfo_v1`  
`dcgmDeviceFbcSessions_v1`  
`dcgmDeviceFbcStats_v1`  
`dcgmDeviceIdentifiers_v1`  
`dcgmDeviceMemoryUsage_v1`  
`dcgmDevicePidAccountingStats_v1`  
`dcgmDevicePowerLimits_v1`  
`dcgmDeviceSupportedClockSets_v1`  
`dcgmDeviceThermals_v1`  
`dcgmDeviceTopology_v1`  
`dcgmDeviceVgpuEncSessions_v1`  
`dcgmDeviceVgpuIds_v1`  
`dcgmDeviceVgpuProcessUtilInfo_v1`  
`dcgmDeviceVgpuTypeInfo_v1`  
`dcgmDeviceVgpuUtilInfo_v1`  
`dcgmDiagResponse_v3`  
`dcgmDiagResponse_v4`  
`dcgmDiagResponsePerGpu_v1`  
`dcgmErrorInfo_t`

dcgmFieldGroupInfo\_v1  
dcgmFieldValue\_v1  
dcgmFieldValue\_v2  
dcgmGpuUsageInfo\_t  
dcgmGroupEntityPair\_t  
dcgmGroupInfo\_v1  
dcgmGroupInfo\_v2  
dcgmGroupTopology\_v1  
dcgmHealthResponse\_v1  
dcgmHealthResponse\_v2  
dcgmHealthResponse\_v3  
dcgmIntrospectContext\_v1  
dcgmIntrospectCpuUtil\_v1  
dcgmIntrospectFieldsExecTime\_v1  
dcgmIntrospectFullFieldsExecTime\_v1  
dcgmIntrospectFullMemory\_v1  
dcgmIntrospectMemory\_v1  
dcgmJobInfo\_v2  
dcgmModuleGetStatusesModule\_t  
dcgmNvLinkGpuLinkStatus\_t  
dcgmNvLinkNvSwitchLinkStatus\_t  
dcgmNvLinkStatus\_v1  
dcgmPidInfo\_v1  
dcgmPidSingleInfo\_t  
dcgmPolicy\_v1  
dcgmPolicyCallbackResponse\_v1  
dcgmPolicyConditionDbe\_t  
dcgmPolicyConditionMpr\_t  
dcgmPolicyConditionNvlink\_t  
dcgmPolicyConditionParms\_t  
dcgmPolicyConditionPci\_t  
dcgmPolicyConditionPower\_t  
dcgmPolicyConditionThermal\_t  
dcgmPolicyConditionXID\_t  
dcgmPolicyViolationNotify\_t  
dcgmProcessUtilInfo\_t  
dcgmProcessUtilSample\_t  
dcgmProfUnwatchFields\_v1  
dcgmProfWatchFields\_v1  
dcgmRunningProcess\_v1  
dcgmStatSummaryFp64\_t  
dcgmStatSummaryInt32\_t  
dcgmStatSummaryInt64\_t

`dcgmVersionInfo_v1`  
`dcgmVgpuConfig_v1`  
`dcgmVgpuDeviceAttributes_v6`  
`dcgmVgpuInstanceAttributes_v1`

### 3.1. `dcgm_field_meta_t` Struct Reference

Structure to store meta data for the field

### 3.2. `dcgm_field_output_format_t` Struct Reference

Structure for formating the output for dmon. Used as a member in `dcgm_field_meta_p`

### 3.3. `dcgmClockSet_v1` Struct Reference

Represents a set of memory, SM, and video clocks for a device. This can be current values or a target values based on context

**int `dcgmClockSet_v1::version`**

Version Number (`dcgmClockSet_version`).

**unsigned int `dcgmClockSet_v1::memClock`**

Memory Clock (Memory Clock value OR DCGM\_INT32\_BLANK to Ignore/Use compatible value with `smClk`).

**unsigned int `dcgmClockSet_v1::smClock`**

SM Clock (SM Clock value OR DCGM\_INT32\_BLANK to Ignore/Use compatible value with `memClk`).

### 3.4. `dcgmConfig_v1` Struct Reference

Structure to represent default and target configuration for a device

**unsigned int dcgmConfig\_v1::version**

Version number (dcgmConfig\_version).

**unsigned int dcgmConfig\_v1::gpuld**

GPU ID.

**unsigned int dcgmConfig\_v1::eccMode**

ECC Mode (0: Disabled, 1 : Enabled, DCGM\_INT32\_BLANK : Ignored).

**unsigned int dcgmConfig\_v1::computeMode**

Compute Mode (One of DCGM\_CONFIG\_COMPUTEMODE\_? OR DCGM\_INT32\_BLANK to Ignore).

**struct dcgmConfigPerfStateSettings\_t****dcgmConfig\_v1::perfState**

Performance State Settings (clocks / boost mode).

**struct dcgmConfigPowerLimit\_t****dcgmConfig\_v1::powerLimit**

Power Limits.

## 3.5. dcgmConfigPerfStateSettings\_t Struct Reference

Used to represent Performance state settings

**unsigned int dcgmConfigPerfStateSettings\_t::syncBoost**

Sync Boost Mode (0: Disabled, 1 : Enabled, DCGM\_INT32\_BLANK : Ignored). Note that using this setting may result in lower clocks than targetClocks.

**struct dcgmClockSet\_t****dcgmConfigPerfStateSettings\_t::targetClocks**

Target clocks. Set smClock and memClock to DCGM\_INT32\_BLANK to ignore/use compatible values. For GPUs > Maxwell, setting this implies autoBoost=0.

## 3.6. dcgmConfigPowerLimit\_t Struct Reference

Used to represents the power capping limit for each GPU in the group or to represent the power budget for the entire group

**dcgmConfigPowerLimitType\_t****dcgmConfigPowerLimit\_t::type**

Flag to represent power cap for each GPU or power budget for the group of GPUs.

**unsigned int dcgmConfigPowerLimit\_t::val**

Power Limit in Watts (Set a value OR DCGM\_INT32\_BLANK to Ignore).

## 3.7. dcgmConnectV2Params\_v1 Struct Reference

Connection options for dcgmConnect\_v2 (v1)

NOTE: This version is deprecated. use dcgmConnectV2Params\_v2

**unsigned int dcgmConnectV2Params\_v1::version**

Version number. Use dcgmConnectV2Params\_version.

**unsigned int****dcgmConnectV2Params\_v1::persistAfterDisconnect**

Whether to persist DCGM state modified by this connection once the connection is terminated. Normally, all field watches created by a connection are removed once a connection goes away. 1 = do not clean up after this connection. 0 = clean up after this connection

## 3.8. dcgmConnectV2Params\_v2 Struct Reference

Connection options for dcgmConnect\_v2 (v2)

### **unsigned int dcgmConnectV2Params\_v2::version**

Version number. Use dcgmConnectV2Params\_version.

### **unsigned int**

### **dcgmConnectV2Params\_v2::persistAfterDisconnect**

Whether to persist DCGM state modified by this connection once the connection is terminated. Normally, all field watches created by a connection are removed once a connection goes away. 1 = do not clean up after this connection. 0 = clean up after this connection

### **unsigned int dcgmConnectV2Params\_v2::timeoutMs**

When attempting to connect to the specified host engine, how long should we wait in milliseconds before giving up

### **unsigned int**

### **dcgmConnectV2Params\_v2::addressIsUnixSocket**

Whether or not the passed-in address is a unix socket filename (1) or a TCP/IP address (0)

## 3.9. dcgmDeviceAttributes\_v1 Struct Reference

Represents attributes corresponding to a device

**unsigned int dcgmDeviceAttributes\_v1::version**

Version number (dcgmDeviceAttributes\_version).

**struct dcgmDeviceSupportedClockSets\_t****dcgmDeviceAttributes\_v1::clockSets**

Supported clocks for the device.

**struct dcgmDeviceThermals\_t****dcgmDeviceAttributes\_v1::thermalSettings**

Thermal settings for the device.

**struct dcgmDevicePowerLimits\_t****dcgmDeviceAttributes\_v1::powerLimits**

Various power limits for the device.

**struct dcgmDeviceIdentifiers\_t****dcgmDeviceAttributes\_v1::identifiers**

Identifiers for the device.

**struct dcgmDeviceMemoryUsage\_t****dcgmDeviceAttributes\_v1::memoryUsage**

Memory usage info for the device.

**struct dcgmDeviceVgpulds\_t****dcgmDeviceAttributes\_v1::unusedVgpulds**

Unused Field.

**unsigned int****dcgmDeviceAttributes\_v1::unusedActiveVgpuInstanceCount**

Unused Field.

**unsigned int****dcgmDeviceAttributes\_v1::unusedVgpuInstanceIds**

Unused Field.

## 3.10. dcgmDeviceEncStats\_v1 Struct Reference

Represents current encoder statistics for the given device/vGPU instance

### **unsigned int dcgmDeviceEncStats\_v1::version**

Version Number (dcgmDeviceEncStats\_version).

### **unsigned int dcgmDeviceEncStats\_v1::sessionCount**

Count of active encoder sessions.

### **unsigned int dcgmDeviceEncStats\_v1::averageFps**

Trailing average FPS of all active sessions.

### **unsigned int dcgmDeviceEncStats\_v1::averageLatency**

Encode latency in milliseconds.

## **3.11. dcgmDeviceFbcSessionInfo\_v1 Struct Reference**

Represents information about active FBC session on the given device/vGPU instance

**unsigned int dcgmDeviceFbcSessionInfo\_v1::version**

Version Number (dcgmDeviceFbcSessionInfo\_version).

**unsigned int dcgmDeviceFbcSessionInfo\_v1::sessionId**

Unique session ID.

**unsigned int dcgmDeviceFbcSessionInfo\_v1::pid**

Owning process ID.

**unsigned int dcgmDeviceFbcSessionInfo\_v1::vgpuld**

vGPU instance ID (only valid on vGPU hosts, otherwise zero)

**unsigned int****dcgmDeviceFbcSessionInfo\_v1::displayOrdinal**

Display identifier.

**dcgmFBCSessionType\_t****dcgmDeviceFbcSessionInfo\_v1::sessionType**

Type of frame buffer capture session.

**unsigned int dcgmDeviceFbcSessionInfo\_v1::sessionFlags**

Session flags.

**unsigned int****dcgmDeviceFbcSessionInfo\_v1::hMaxResolution**

Max horizontal resolution supported by the capture session.

**unsigned int****dcgmDeviceFbcSessionInfo\_v1::vMaxResolution**

Max vertical resolution supported by the capture session.

**unsigned int dcgmDeviceFbcSessionInfo\_v1::hResolution**

Horizontal resolution requested by caller in capture call.

**unsigned int dcgmDeviceFbcSessionInfo\_v1::vResolution**

Vertical resolution requested by caller in capture call.

**unsigned int dcgmDeviceFbcSessionInfo\_v1::averageFps**

Moving average new frames captured per second.

**unsigned int****dcgmDeviceFbcSessionInfo\_v1::averageLatency**

Moving average new frame capture latency in microseconds.

## 3.12. dcgmDeviceFbcSessions\_v1 Struct Reference

Represents all the active FBC sessions on the given device/vGPU instance

**unsigned int dcgmDeviceFbcSessions\_v1::version**

Version Number (dcgmDeviceFbcSessions\_version).

**unsigned int dcgmDeviceFbcSessions\_v1::sessionCount**

Count of active FBC sessions.

**struct dcgmDeviceFbcSessionInfo\_t****dcgmDeviceFbcSessions\_v1::sessionInfo**

Info about the active FBC session.

## 3.13. dcgmDeviceFbcStats\_v1 Struct Reference

Represents current frame buffer capture sessions statistics for the given device/vGPU instance

**unsigned int dcgmDeviceFbcStats\_v1::version**

Version Number (dcgmDeviceFbcStats\_version).

**unsigned int dcgmDeviceFbcStats\_v1::sessionCount**

Count of active FBC sessions.

**unsigned int dcgmDeviceFbcStats\_v1::averageFps**

Moving average new frames captured per second.

**unsigned int dcgmDeviceFbcStats\_v1::averageLatency**

Moving average new frame capture latency in microseconds.

## 3.14. dcgmDeviceIdentifiers\_v1 Struct Reference

Represents device identifiers

**unsigned int dcgmDeviceIdentifiers\_v1::version**

Version Number (dcgmDeviceIdentifiers\_version).

**char dcgmDeviceIdentifiers\_v1::brandName**

Brand Name.

**char dcgmDeviceIdentifiers\_v1::deviceName**

Name of the device.

**char dcgmDeviceIdentifiers\_v1::pciBusId**

PCI Bus ID.

**char dcgmDeviceIdentifiers\_v1::serial**

Serial for the device.

**char dcgmDeviceIdentifiers\_v1::uuid**

UUID for the device.

**char dcgmDeviceIdentifiers\_v1::vbios**

VBIOS version.

**char dcgmDeviceIdentifiers\_v1::inforomImageVersion**

Inforom Image version.

**unsigned int dcgmDeviceIdentifiers\_v1::pciDeviceId**

The combined 16-bit device id and 16-bit vendor id.

**unsigned int dcgmDeviceIdentifiers\_v1::pciSubSystemId**

The 32-bit Sub System Device ID.

**char dcgmDeviceIdentifiers\_v1::driverVersion**

Driver Version.

**unsigned int****dcgmDeviceIdentifiers\_v1::virtualizationMode**

Virtualization Mode.

## 3.15. dcgmDeviceMemoryUsage\_v1 Struct Reference

Represents device memory and usage

### **unsigned int dcgmDeviceMemoryUsage\_v1::version**

Version Number (dcgmDeviceMemoryUsage\_version).

### **unsigned int dcgmDeviceMemoryUsage\_v1::bar1Total**

Total BAR1 size in megabytes.

### **unsigned int dcgmDeviceMemoryUsage\_v1::fbTotal**

Total framebuffer memory in megabytes.

### **unsigned int dcgmDeviceMemoryUsage\_v1::fbUsed**

Used framebuffer memory in megabytes.

### **unsigned int dcgmDeviceMemoryUsage\_v1::fbFree**

Free framebuffer memory in megabytes.

## 3.16. dcgmDevicePidAccountingStats\_v1 Struct Reference

Represents accounting data for one process

### **unsigned int dcgmDevicePidAccountingStats\_v1::version**

Version Number. Should match dcgmDevicePidAccountingStats\_version.

### **unsigned int dcgmDevicePidAccountingStats\_v1::pid**

Process id of the process these stats are for.

### **unsigned int**

### **dcgmDevicePidAccountingStats\_v1::gpuUtilization**

Percent of time over the process's lifetime during which one or more kernels was executing on the GPU. Set to DCGM\_INT32\_NOT\_SUPPORTED if is not supported

**unsigned int  
dcgmDevicePidAccountingStats\_v1::memoryUtilization**

Percent of time over the process's lifetime during which global (device) memory was being read or written. Set to DCGM\_INT32\_NOT\_SUPPORTED if is not supported

**unsigned long long  
dcgmDevicePidAccountingStats\_v1::maxMemoryUsage**

Maximum total memory in bytes that was ever allocated by the process. Set to DCGM\_INT64\_NOT\_SUPPORTED if is not supported

**unsigned long long  
dcgmDevicePidAccountingStats\_v1::startTimestamp**

CPU Timestamp in usec representing start time for the process.

**unsigned long long  
dcgmDevicePidAccountingStats\_v1::activeTimeUsec**

Amount of time in usec during which the compute context was active. Note that this does not mean the context was being used. endTimestamp can be computed as startTimestamp + activeTime

## 3.17. dcgmDevicePowerLimits\_v1 Struct Reference

Represents various power limits

**unsigned int dcgmDevicePowerLimits\_v1::version**

Version Number.

**unsigned int dcgmDevicePowerLimits\_v1::curPowerLimit**

Power management limit associated with this device (in W).

**unsigned int****dcgmDevicePowerLimits\_v1::defaultPowerLimit**

Power management limit effective at device boot (in W).

**unsigned int****dcgmDevicePowerLimits\_v1::enforcedPowerLimit**

Effective power limit that the driver enforces after taking into account all limiters (in W).

**unsigned int****dcgmDevicePowerLimits\_v1::minPowerLimit**

Minimum power management limit (in W).

**unsigned int****dcgmDevicePowerLimits\_v1::maxPowerLimit**

Maximum power management limit (in W).

## 3.18. dcgmDeviceSupportedClockSets\_v1 Struct Reference

Represents list of supported clock sets for a device

**unsigned int dcgmDeviceSupportedClockSets\_v1::version**

Version Number (dcgmDeviceSupportedClockSets\_version).

**unsigned int dcgmDeviceSupportedClockSets\_v1::count**

Number of supported clocks.

**struct dcgmClockSet\_t****dcgmDeviceSupportedClockSets\_v1::clockSet**

Valid clock sets for the device. Upto count entries are filled.

## 3.19. dcgmDeviceThermals\_v1 Struct Reference

Represents thermal information

**unsigned int dcgmDeviceThermals\_v1::version**

Version Number.

**unsigned int dcgmDeviceThermals\_v1::slowdownTemp**

Slowdown temperature.

**unsigned int dcgmDeviceThermals\_v1::shutdownTemp**

Shutdown temperature.

## 3.20. dcgmDeviceTopology\_v1 Struct Reference

Device topology information

**unsigned int dcgmDeviceTopology\_v1::version**

version number (dcgmDeviceTopology\_version)

**unsignedlong dcgmDeviceTopology\_v1::cpuAffinityMask**

affinity mask for the specified GPU a 1 represents affinity to the CPU in that bit position  
supports up to 256 cores

**unsigned int dcgmDeviceTopology\_v1::numGpus**

number of valid entries in gpuPaths

**unsigned int dcgmDeviceTopology\_v1::gpuId**

gpuId to which the path represents

**dcgmGpuTopologyLevel\_t dcgmDeviceTopology\_v1::path**

path to the gpuId from this GPU. Note that this is a bitmask of DCGM\_TOPOLOGY\_\* values and can contain both PCIe topology and NvLink topology where applicable. For instance: 0x210 = DCGM\_TOPOLOGY\_CPU | DCGM\_TOPOLOGY\_NVLINK2 Use the macros DCGM\_TOPOLOGY\_PATH\_NVLINK and DCGM\_TOPOLOGY\_PATH\_PCI to mask the NvLink and PCI paths, respectively.

**unsigned int dcgmDeviceTopology\_v1::localNvLinkIds**

bits representing the local links connected to gpuId e.g. if this field == 3, links 0 and 1 are connected, field is only valid if NVLINKS actually exist between GPUs

## 3.21. dcgmDeviceVgpuEncSessions\_v1 Struct Reference

Represents information about active encoder sessions on the given vGPU instance

**unsigned int dcgmDeviceVgpuEncSessions\_v1::version**

Version Number (dcgmDeviceVgpuEncSessions\_version).

**unsigned int dcgmDeviceVgpuEncSessions\_v1::vgpuld**

vGPU instance ID

**unsigned int dcgmDeviceVgpuEncSessions\_v1::sessionId**

Unique session ID.

**unsigned int dcgmDeviceVgpuEncSessions\_v1::pid**

Process ID.

**dcgmEncoderType\_t****dcgmDeviceVgpuEncSessions\_v1::codecType**

Video encoder type.

**unsigned int****dcgmDeviceVgpuEncSessions\_v1::hResolution**

Current encode horizontal resolution.

**unsigned int****dcgmDeviceVgpuEncSessions\_v1::vResolution**

Current encode vertical resolution.

**unsigned int****dcgmDeviceVgpuEncSessions\_v1::averageFps**

Moving average encode frames per second.

**unsigned int****dcgmDeviceVgpuEncSessions\_v1::averageLatency**

Moving average encode latency in milliseconds.

## 3.22. dcgmDeviceVgpuls\_v1 Struct Reference

Represents various IDs related to vGPU.

**unsigned int dcgmDeviceVgpuids\_v1::version**

Version Number (dcgmDeviceVgpuIds\_version).

**unsigned int****dcgmDeviceVgpuids\_v1::unusedSupportedVgpuTypeCount**

Unused Field.

**unsigned int****dcgmDeviceVgpuids\_v1::unusedSupportedVgpuTypeIds**

Unused Field.

**unsigned int****dcgmDeviceVgpuids\_v1::unusedCreatableVgpuTypeCount**

Unused Field.

**unsigned int****dcgmDeviceVgpuids\_v1::unusedCreatableVgpuTypeIds**

Unused Field.

### 3.23. dcgmDeviceVgpuProcessUtilInfo\_v1 Struct Reference

Represents utilization values for processes running in vGPU VMs using the device

**unsigned int dcgmDeviceVgpuProcessUtilInfo\_v1::version**

Version Number (dcgmDeviceVgpuProcessUtilInfo\_version).

**unsigned int dcgmDeviceVgpuProcessUtilInfo\_v1::vgpuld**

vGPU instance ID

**unsigned int****dcgmDeviceVgpuProcessUtilInfo\_v1::vgpuProcessSamplesCount**

Count of processes running in the vGPU VM, for which utilization rates are being reported in this cycle.

**unsigned int dcgmDeviceVgpuProcessUtilInfo\_v1::pid**

Process ID of the process running in the vGPU VM.

**char dcgmDeviceVgpuProcessUtilInfo\_v1::processName**

Process Name of process running in the vGPU VM.

**unsigned int dcgmDeviceVgpuProcessUtilInfo\_v1::smUtil**

GPU utilization of process running in the vGPU VM.

**unsigned int****dcgmDeviceVgpuProcessUtilInfo\_v1::memUtil**

Memory utilization of process running in the vGPU VM.

**unsigned int****dcgmDeviceVgpuProcessUtilInfo\_v1::encUtil**

Encoder utilization of process running in the vGPU VM.

**unsigned int****dcgmDeviceVgpuProcessUtilInfo\_v1::decUtil**

Decoder utilization of process running in the vGPU VM.

## 3.24. dcgmDeviceVgpuTypeInfo\_v1 Struct Reference

Represents static info related to vGPUs supported on the device.

**unsigned int dcgmDeviceVgpuTypeInfo\_v1::version**

Version number (dcgmDeviceVgpuTypeIdxStaticInfo\_version).

**dcgmDeviceVgpuTypeInfo\_v1::@2****dcgmDeviceVgpuTypeInfo\_v1::vgpuTypeInfo**

vGPU type ID and Supported vGPU type count

**char dcgmDeviceVgpuTypeInfo\_v1::vgpuTypeName**

vGPU type Name

**char dcgmDeviceVgpuTypeInfo\_v1::vgpuTypeClass**

Class of vGPU type.

**char dcgmDeviceVgpuTypeInfo\_v1::vgpuTypeLicense**

license of vGPU type

**int dcgmDeviceVgpuTypeInfo\_v1::deviceId**

device ID of vGPU type

**int dcgmDeviceVgpuTypeInfo\_v1::subsystemId**

Subsystem ID of vGPU type.

**int dcgmDeviceVgpuTypeInfo\_v1::numDisplayHeads**

Count of vGPU's supported display heads.

**int dcgmDeviceVgpuTypeInfo\_v1::maxInstances**

maximum number of vGPU instances creatable on a device for given vGPU type

**int dcgmDeviceVgpuTypeInfo\_v1::frameRateLimit**

Frame rate limit value of the vGPU type.

**int dcgmDeviceVgpuTypeInfo\_v1::maxResolutionX**

vGPU display head's maximum supported resolution in X dimension

**int dcgmDeviceVgpuTypeInfo\_v1::maxResolutionY**

vGPU display head's maximum supported resolution in Y dimension

**int dcgmDeviceVgpuTypeInfo\_v1::fbTotal**

vGPU Total framebuffer size in megabytes

## 3.25. dcgmDeviceVgpuUtilInfo\_v1 Struct Reference

Represents utilization values for vGPUs running on the device

**unsigned int dcgmDeviceVgpuUtilInfo\_v1::version**

Version Number (dcgmDeviceVgpuUtilInfo\_version).

**unsigned int dcgmDeviceVgpuUtilInfo\_v1::vgpuld**

vGPU instance ID

**unsigned int dcgmDeviceVgpuUtilInfo\_v1::smUtil**

GPU utilization for vGPU.

**unsigned int dcgmDeviceVgpuUtilInfo\_v1::memUtil**

Memory utilization for vGPU.

**unsigned int dcgmDeviceVgpuUtilInfo\_v1::encUtil**

Encoder utilization for vGPU.

**unsigned int dcgmDeviceVgpuUtilInfo\_v1::decUtil**

Decoder utilization for vGPU.

## 3.26. dcgmDiagResponse\_v3 Struct Reference

Global diagnostics result structure

**unsigned int dcgmDiagResponse\_v3::version**

version number (dcgmDiagResult\_version)

**unsigned int dcgmDiagResponse\_v3::gpuCount**

number of valid per GPU results

**dcgmDiagResult\_t dcgmDiagResponse\_v3::blacklist**

test for presence of blacklisted drivers (e.g. nouveau)

**dcgmDiagResult\_t dcgmDiagResponse\_v3::nvmlLibrary**

test for presence (and version) of NVML lib

**dcgmDiagResult\_t****dcgmDiagResponse\_v3::cudaMainLibrary**

test for presence (and version) of CUDA lib

**dcgmDiagResult\_t****dcgmDiagResponse\_v3::cudaRuntimeLibrary**

test for presence (and version) of CUDA RT lib

**dcgmDiagResult\_t dcgmDiagResponse\_v3::permissions**

test for character device permissions

**dcgmDiagResult\_t****dcgmDiagResponse\_v3::persistenceMode**

test for persistence mode enabled

**dcgmDiagResult\_t dcgmDiagResponse\_v3::environment**

test for CUDA environment vars that may slow tests

**dcgmDiagResult\_t****dcgmDiagResponse\_v3::pageRetirement**

test for pending frame buffer page retirement

**dcgmDiagResult\_t dcgmDiagResponse\_v3::inforom**

test for inforom corruption

**dcgmDiagResult\_t**

**dcgmDiagResponse\_v3::graphicsProcesses**

test for graphics processes running

**struct dcgmDiagResponsePerGpu\_v1**

**dcgmDiagResponse\_v3::perGpuResponses**

per GPU test results

**char dcgmDiagResponse\_v3::systemError**

System-wide error reported from NVVS.

## 3.27. dcgmDiagResponse\_v4 Struct Reference

Global diagnostics result structure

**unsigned int dcgmDiagResponse\_v4::version**

version number (dcgmDiagResult\_version)

**unsigned int dcgmDiagResponse\_v4::gpuCount**

number of valid per GPU results

**unsigned int dcgmDiagResponse\_v4::levelOneTestCount**

number of valid levelOne results

**dcgmDiagTestResult\_v1****dcgmDiagResponse\_v4::levelOneResults**

Basic, system-wide test results.

**struct dcgmDiagResponsePerGpu\_v1****dcgmDiagResponse\_v4::perGpuResponses**

per GPU test results

**char dcgmDiagResponse\_v4::systemError**

System-wide error reported from NVVS.

**char dcgmDiagResponse\_v4::trainingMsg**

Training Message.

## 3.28. dcgmDiagResponsePerGpu\_v1 Struct Reference

Per GPU diagnostics result structure

**unsigned int dcgmDiagResponsePerGpu\_v1::gpuld**

ID for the GPU this information pertains.

**unsigned int****dcgmDiagResponsePerGpu\_v1::hwDiagnosticReturn**

Per GPU hardware diagnostic test return code.

**dcgmDiagTestResult\_v1****dcgmDiagResponsePerGpu\_v1::results**

Array with a result for each per-gpu test.

## 3.29. dcgmErrorInfo\_t Struct Reference

Structure to represent error attributes

**unsigned int dcgmErrorInfo\_t::gpuld**

Represents GPU ID.

**short dcgmErrorInfo\_t::fieldId**

One of DCGM\_FI\_?

**int dcgmErrorInfo\_t::status**

One of DCGM\_ST\_?

## 3.30. dcgmFieldGroupInfo\_v1 Struct Reference

Structure to represent information about a field group

**unsigned int dcgmFieldGroupInfo\_v1::version**

Version number (dcgmFieldGroupInfo\_version).

**unsigned int dcgmFieldGroupInfo\_v1::numFieldIds**

Number of entries in fieldIds[] that are valid.

**dcgmFieldGrp\_t dcgmFieldGroupInfo\_v1::fieldGroupID**

ID of this field group.

**char dcgmFieldGroupInfo\_v1::fieldGroupName**

Field Group Name.

**unsigned short dcgmFieldGroupInfo\_v1::fieldIds**

Field ids that belong to this group.

### 3.31. dcgmFieldValue\_v1 Struct Reference

This structure is used to represent value for the field to be queried.

**unsigned int dcgmFieldValue\_v1::version**

version number (dcgmFieldValue\_version1)

**unsigned short dcgmFieldValue\_v1::fieldId**

One of DCGM\_FI\_?

**unsigned short dcgmFieldValue\_v1::fieldType**

One of DCGM\_FT\_?

**int dcgmFieldValue\_v1::status**

Status for the querying the field. DCGM\_ST\_OK or one of DCGM\_ST\_?

**int64\_t dcgmFieldValue\_v1::ts**

Timestamp in usec since 1970 \*/.

**int64\_t dcgmFieldValue\_v1::i64**

Int64 value.

**double dcgmFieldValue\_v1::dbl**

Double value.

**char dcgmFieldValue\_v1::str**

NULL terminated string.

**char dcgmFieldValue\_v1::blob**

Binary blob.

**dcgmFieldValue\_v1::@7 dcgmFieldValue\_v1::value**

Value.

## 3.32. dcgmFieldValue\_v2 Struct Reference

This structure is used to represent value for the field to be queried.

**unsigned int dcgmFieldValue\_v2::version**

version number (dcgmFieldValue\_version2)

**dcgm\_field\_entity\_group\_t****dcgmFieldValue\_v2::entityGroupId**

Entity group this field value's entity belongs to.

**dcgm\_field\_eid\_t dcgmFieldValue\_v2::entityId**

Entity this field value belongs to.

**unsigned short dcgmFieldValue\_v2::fieldId**

One of DCGM\_FI\_?

**unsigned short dcgmFieldValue\_v2::fieldType**

One of DCGM\_FT\_?

**int dcgmFieldValue\_v2::status**

Status for the querying the field. DCGM\_ST\_OK or one of DCGM\_ST\_?

**unsigned int dcgmFieldValue\_v2::unused**

Unused for now to align ts to an 8-byte boundary. \*/.

**int64\_t dcgmFieldValue\_v2::ts**

Timestamp in usec since 1970 \*/.

**int64\_t dcgmFieldValue\_v2::i64**

Int64 value.

**double dcgmFieldValue\_v2::dbl**

Double value.

**char dcgmFieldValue\_v2::str**

NULL terminated string.

**char dcgmFieldValue\_v2::blob**

Binary blob.

**dcgmFieldValue\_v2::@8 dcgmFieldValue\_v2::value**

Value.

## 3.33. dcgmGpuUsageInfo\_t Struct Reference

Info corresponding to the job on a GPU

**unsigned int dcgmGpuUsageInfo\_t::gpuld**

ID of the GPU this pertains to. GPU\_ID\_INVALID = summary information for multiple GPUs.

**long long dcgmGpuUsageInfo\_t::energyConsumed**

Energy consumed in milliwatt-seconds.

**struct dcgmStatSummaryFp64\_t  
dcgmGpuUsageInfo\_t::powerUsage**

Power usage Min/Max/Avg in watts.

**struct dcgmStatSummaryInt64\_t  
dcgmGpuUsageInfo\_t::pcieRxBandwidth**

PCI-E bytes read from the GPU.

**struct dcgmStatSummaryInt64\_t  
dcgmGpuUsageInfo\_t::pcieTxBandwidth**

PCI-E bytes written to the GPU.

**long long dcgmGpuUsageInfo\_t::pcieReplays**

Count of PCI-E replays that occurred.

**long long dcgmGpuUsageInfo\_t::startTime**

User provided job start time in microseconds since 1970.

**long long dcgmGpuUsageInfo\_t::endTime**

User provided job end time in microseconds since 1970.

**struct dcgmStatSummaryInt32\_t  
dcgmGpuUsageInfo\_t::smUtilization**

GPU SM Utilization in percent.

**struct dcgmStatSummaryInt32\_t  
dcgmGpuUsageInfo\_t::memoryUtilization**

GPU Memory Utilization in percent.

**unsigned int dcgmGpuUsageInfo\_t::eccSingleBit**

Count of ECC single bit errors that occurred.

**unsigned int dcgmGpuUsageInfo\_t::eccDoubleBit**

Count of ECC double bit errors that occurred.

**struct dcgmStatSummaryInt32\_t  
dcgmGpuUsageInfo\_t::memoryClock**

Memory clock in MHz.

**struct dcgmStatSummaryInt32\_t  
dcgmGpuUsageInfo\_t::smClock**

SM clock in MHz.

**int dcgmGpuUsageInfo\_t::numXidCriticalErrors**

Number of valid entries in xidCriticalErrorsTs.

**long long dcgmGpuUsageInfo\_t::xidCriticalErrorsTs**

Timestamps of the critical XID errors that occurred.

**int dcgmGpuUsageInfo\_t::numComputePids**

Count of computePids entries that are valid.

**struct dcgmProcessUtilInfo\_t  
dcgmGpuUsageInfo\_t::computePidInfo**

List of compute processes that ran during the job. 0=no process.

**int dcgmGpuUsageInfo\_t::numGraphicsPids**

Count of graphicsPids entries that are valid.

**struct dcgmProcessUtilInfo\_t  
dcgmGpuUsageInfo\_t::graphicsPidInfo**

List of compute processes that ran during the job. 0=no process.

**long long dcgmGpuUsageInfo\_t::maxGpuMemoryUsed**

Maximum amount of GPU memory that was used in bytes.

**long long dcgmGpuUsageInfo\_t::powerViolationTime**

Number of microseconds we were at reduced clocks due to power violation.

**long long dcgmGpuUsageInfo\_t::thermalViolationTime**

Number of microseconds we were at reduced clocks due to thermal violation.

**long long dcgmGpuUsageInfo\_t::reliabilityViolationTime**

Amount of microseconds we were at reduced clocks due to the reliability limit.

**long long****dcgmGpuUsageInfo\_t::boardLimitViolationTime**

Amount of microseconds we were at reduced clocks due to being at the board's max voltage.

**long long dcgmGpuUsageInfo\_t::lowUtilizationTime**

Amount of microseconds we were at reduced clocks due to low utilization.

**long long dcgmGpuUsageInfo\_t::syncBoostTime**

Amount of microseconds we were at reduced clocks due to sync boost.

**dcgmHealthWatchResults\_t****dcgmGpuUsageInfo\_t::overallHealth**

The overall health of the system . dcgmHealthWatchResults\_t.

**dcgmHealthSystems\_t dcgmGpuUsageInfo\_t::system**

system to which this information belongs

**dcgmHealthWatchResults\_t****dcgmGpuUsageInfo\_t::health**

health of the specified system on this GPU

### 3.34. dcgmGroupEntityPair\_t Struct Reference

Represents a entityGroupId + entityId pair to uniquely identify a given entityId inside a group of entities

**dcgm\_field\_entity\_group\_t**  
**dcgmGroupEntityPair\_t::entityGroupId**

Entity Group ID entity belongs to.

**dcgm\_field\_eid\_t dcgmGroupEntityPair\_t::entityId**

Entity ID of the entity.

### 3.35. dcgmGroupInfo\_v1 Struct Reference

Structure to store information for DCGM group

**unsigned int dcgmGroupInfo\_v1::version**

Version Number (use dcgmGroupInfo\_version1).

**unsigned int dcgmGroupInfo\_v1::count**

count of GPU IDs returned in gpuIdList

**unsigned int dcgmGroupInfo\_v1::gpuidList**

List of GPU Ids part of the group.

**char dcgmGroupInfo\_v1::groupName**

Group Name.

### 3.36. dcgmGroupInfo\_v2 Struct Reference

Structure to store information for DCGM group

**unsigned int dcgmGroupInfo\_v2::version**

Version Number (use dcgmGroupInfo\_version2).

**unsigned int dcgmGroupInfo\_v2::count**

count of entityIds returned in entityList

**char dcgmGroupInfo\_v2::groupName**

Group Name.

**struct dcgmGroupEntityPair\_t  
dcgmGroupInfo\_v2::entityList**

List of the entities that are in this group.

### 3.37. dcgmGroupTopology\_v1 Struct Reference

Group topology information

**unsigned int dcgmGroupTopology\_v1::version**

version number (dcgmGroupTopology\_version)

**unsignedlong****dcgmGroupTopology\_v1::groupCpuAffinityMask**

the CPU affinity mask for all GPUs in the group a 1 represents affinity to the CPU in that bit position supports up to 256 cores

**unsigned int dcgmGroupTopology\_v1::numaOptimalFlag**

a zero value indicates that 1 or more GPUs in the group have a different CPU affinity and thus may not be optimal for certain algorithms

**dcgmGpuTopologyLevel\_t****dcgmGroupTopology\_v1::slowestPath**

the slowest path amongst GPUs in the group

### 3.38. dcgmHealthResponse\_v1 Struct Reference

Health Response structure version 1. GPU Only

**unsigned int dcgmHealthResponse\_v1::version**

version number (dcgmHealthResponse\_version)

**dcgmHealthWatchResults\_t****dcgmHealthResponse\_v1::overallHealth**

The overall health of the system. dcgmHealthWatchResults\_t.

overall health of this GPU

**unsigned int dcgmHealthResponse\_v1::gpuCount**

The number of GPUs with warnings/errors.

**unsigned int dcgmHealthResponse\_v1::gpuld**

GPU ID for which this data is valid.

**unsigned int dcgmHealthResponse\_v1::incidentCount**

The number of systems that encountered a warning/error.

**dcgmHealthSystems\_t dcgmHealthResponse\_v1::system**

system to which this information belongs

**dcgmHealthWatchResults\_t****dcgmHealthResponse\_v1::health**

health of the specified system on this GPU

**char dcgmHealthResponse\_v1::errorString**

information about the error(s) or warning(s) flagged

## 3.39. dcgmHealthResponse\_v2 Struct Reference

Health Response structure version 2 - NvSwitch-compatible

**unsigned int dcgmHealthResponse\_v2::version**

version number (dcgmHealthResponse\_version)

**dcgmHealthWatchResults\_t****dcgmHealthResponse\_v2::overallHealth**

The overall health of the system. dcgmHealthWatchResults\_t.

overall health of this entity

**unsigned int dcgmHealthResponse\_v2::entityCount**

The number of entities with warnings/errors.

**dcgm\_field\_entity\_group\_t****dcgmHealthResponse\_v2::entityGroupId**

entity group entityId belongs to

**dcgm\_field\_eid\_t dcgmHealthResponse\_v2::entityId**

entity for which this data is valid

**unsigned int dcgmHealthResponse\_v2::incidentCount**

The number of systems that encountered a warning/error.

**dcgmHealthSystems\_t dcgmHealthResponse\_v2::system**

system to which this information belongs

**dcgmHealthWatchResults\_t****dcgmHealthResponse\_v2::health**

health of the specified system on this entity

**char dcgmHealthResponse\_v2::errorString**

information about the error(s) or warning(s) flagged

## 3.40. dcgmHealthResponse\_v3 Struct Reference

Health Response structure version 3 - NvSwitch-compatible and uses error codes for easier processing

**unsigned int dcgmHealthResponse\_v3::version**

version number (dcgmHealthResponse\_version)

**dcgmHealthWatchResults\_t****dcgmHealthResponse\_v3::overallHealth**

The overall health of the system. dcgmHealthWatchResults\_t.

overall health of this entity

**unsigned int dcgmHealthResponse\_v3::entityCount**

The number of entities with warnings/errors.

**dcgm\_field\_entity\_group\_t****dcgmHealthResponse\_v3::entityGroupId**

entity group entityId belongs to

**dcgm\_field\_eid\_t dcgmHealthResponse\_v3::entityId**

entity for which this data is valid

**unsigned int dcgmHealthResponse\_v3::incidentCount**

The number of systems that encountered a warning/error.

**dcgmHealthSystems\_t dcgmHealthResponse\_v3::system**

system to which this information belongs

**dcgmHealthWatchResults\_t****dcgmHealthResponse\_v3::health**

health of the specified system on this entity

**dcgmDiagErrorDetail\_t dcgmHealthResponse\_v3::errors**

Information about the error(s) and their error codes.

**unsigned int dcgmHealthResponse\_v3::errorCount**

count of errors so far for this system

## 3.41. dcgmIntrospectContext\_v1 Struct Reference

Identifies the retrieval context for introspection API calls.

**unsigned int dcgmlIntrospectContext\_v1::version**

version number (dcgmIntrospectContext\_version)

**dcgmlIntrospectLevel\_t****dcgmlIntrospectContext\_v1::introspectLvl**

Introspect Level dcgmlIntrospectLevel\_t.

**dcgmGpuGrp\_t dcgmlIntrospectContext\_v1::fieldGroupId**

Only needed if introspectLvl is DCGM\_INTROSPECT\_LVL\_FIELD\_GROUP.

**unsigned short dcgmlIntrospectContext\_v1::fieldId**

Only needed if introspectLvl is DCGM\_INTROSPECT\_LVL\_FIELD.

**unsigned long long****dcgmlIntrospectContext\_v1::contextId**

Overloaded way to access both fieldGroupId and fieldId.

## 3.42. dcgmlIntrospectCpuUtil\_v1 Struct Reference

DCGM CPU Utilization information. Multiply values by 100 to get them in %.

**unsigned int dcgmlIntrospectCpuUtil\_v1::version**

version number (dcgmMetadataCpuUtil\_version)

**double dcgmlIntrospectCpuUtil\_v1::total**

fraction of device's CPU resources that were used

**double dcgmlIntrospectCpuUtil\_v1::kernel**

fraction of device's CPU resources that were used in kernel mode

**double dcgmlIntrospectCpuUtil\_v1::user**

fraction of device's CPU resources that were used in user mode

## 3.43. dcgmlIntrospectFieldsExecTime\_v1 Struct Reference

DCGM Execution time info for a set of fields

**unsigned int dcgmlIntrospectFieldsExecTime\_v1::version**

version number (dcgmlIntrospectFieldsExecTime\_version)

**long long****dcgmlIntrospectFieldsExecTime\_v1::meanUpdateFreqUsec**

the mean update frequency of all fields

**double****dcgmlIntrospectFieldsExecTime\_v1::recentUpdateUsec**

the sum of every field's most recent execution time after they have been normalized to [meanUpdateFreqUsec](#)". This is roughly how long it takes to update fields every [meanUpdateFreqUsec](#)

**long long****dcgmlIntrospectFieldsExecTime\_v1::totalEverUpdateUsec**

The total amount of time, ever, that has been spent updating all the fields.

## 3.44. dcgmlIntrospectFullFieldsExecTime\_v1 Struct Reference

Full introspection info for field execution time

**unsigned int**

**dcgmlIntrospectFullFieldsExecTime\_v1::version**

version number (dcgmlIntrospectFullFieldsExecTime\_version)

**struct dcgmlIntrospectFieldsExecTime\_v1**

**dcgmlIntrospectFullFieldsExecTime\_v1::aggregateInfo**

info that includes global and device scope

**int dcgmlIntrospectFullFieldsExecTime\_v1::hasGlobalInfo**

0 means globalInfo is populated, !0 means it's not

**struct dcgmlIntrospectFieldsExecTime\_v1**

**dcgmlIntrospectFullFieldsExecTime\_v1::globalInfo**

info that only includes global field scope

**unsigned short**

**dcgmlIntrospectFullFieldsExecTime\_v1::gpuInfoCount**

count of how many entries in gpuInfo are populated

**unsigned int**

**dcgmlIntrospectFullFieldsExecTime\_v1::gpuidsForGpuInfo**

the GPU ID at a given index identifies which gpu the corresponding entry in `gpuInfo` is from

**struct dcgmlIntrospectFieldsExecTime\_v1**

**dcgmlIntrospectFullFieldsExecTime\_v1::gpuInfo**

info that is separated by the GPU ID that the watches were for

### 3.45. `dcgmlIntrospectFullMemory_v1` Struct Reference

Full introspection info for field memory

**unsigned int dcgmlIntrospectFullMemory\_v1::version**

version number (dcgmlIntrospectFullMemory\_version)

**struct dcgmlIntrospectMemory\_v1**

**dcgmlIntrospectFullMemory\_v1::aggregateInfo**

info that includes global and device scope

**int dcgmlIntrospectFullMemory\_v1::hasGlobalInfo**

0 means globalInfo is populated, !0 means it's not

**struct dcgmlIntrospectMemory\_v1**

**dcgmlIntrospectFullMemory\_v1::globalInfo**

info that only includes global field scope

**unsigned short**

**dcgmlIntrospectFullMemory\_v1::gpuInfoCount**

count of how many entries in gpuInfo are populated

**unsigned int**

**dcgmlIntrospectFullMemory\_v1::gpuidsForGpuInfo**

the GPU ID at a given index identifies which gpu the corresponding entry in **gpuInfo** is from

**struct dcgmlIntrospectMemory\_v1**

**dcgmlIntrospectFullMemory\_v1::gpuInfo**

info that is divided by the GPU ID that the watches were for

## 3.46. dcgmlIntrospectMemory\_v1 Struct Reference

DCGM Memory usage information

**unsigned int dcgmIntrospectMemory\_v1::version**

version number (dcgmIntrospectMemory\_version)

**long long dcgmIntrospectMemory\_v1::bytesUsed**

number of bytes

## 3.47. dcgmJobInfo\_v2 Struct Reference

To store job statistics The following fields are not applicable in the summary info:

- ▶ pcieRxBandwidth (Min/Max)
- ▶ pcieTxBandwidth (Min/Max)
- ▶ smUtilization (Min/Max)
- ▶ memoryUtilization (Min/Max)
- ▶ memoryClock (Min/Max)
- ▶ smClock (Min/Max)
- ▶ processSamples

The average value in the above fields (in the summary) is the average of the averages of respective fields from all GPUs

**unsigned int dcgmJobInfo\_v2::version**

Version of this message (dcgmPidInfo\_version).

**int dcgmJobInfo\_v2::numGpus**

Number of GPUs that are valid in gpus[].

**struct dcgmGpuUsageInfo\_t dcgmJobInfo\_v2::summary**

Summary information for all GPUs listed in gpus[].

**struct dcgmGpuUsageInfo\_t dcgmJobInfo\_v2::gpus**

Per-GPU information for this PID.

## 3.48. dcgmModuleGetStatusesModule\_t Struct Reference

Status of all of the modules of the host engine

**dcgmModuleId\_t dcgmModuleGetStatusesModule\_t::id**

ID of this module.

**dcgmModuleStatus\_t****dcgmModuleGetStatusesModule\_t::status**

Status of this module.

## 3.49. dcgmNvLinkGpuLinkStatus\_t Struct Reference

State of NvLink links for a GPU

**dcgm\_field\_eid\_t dcgmNvLinkGpuLinkStatus\_t::entityId**

Entity ID of the GPU (gpuId).

**dcgmNvLinkLinkState\_t****dcgmNvLinkGpuLinkStatus\_t::linkState**

Per-GPU link states.

## 3.50. dcgmNvLinkNvSwitchLinkStatus\_t Struct Reference

State of NvLink links for a NvSwitch

**dcgm\_field\_eid\_t****dcgmNvLinkNvSwitchLinkStatus\_t::entityId**

Entity ID of the NvSwitch (physicalId).

**dcgmNvLinkLinkState\_t****dcgmNvLinkNvSwitchLinkStatus\_t::linkState**

Per-NvSwitch link states.

## 3.51. dcgmNvLinkStatus\_v1 Struct Reference

Status of all of the NvLinks in a given system

**unsigned int dcgmNvLinkStatus\_v1::version**

Version of this request. Should be dcgmNvLinkStatus\_version1.

**unsigned int dcgmNvLinkStatus\_v1::numGpus**

Number of entries in gpus[] that are populated.

**struct dcgmNvLinkGpuLinkStatus\_t  
dcgmNvLinkStatus\_v1::gpus**

Per-GPU NvLink link statuses.

**unsigned int dcgmNvLinkStatus\_v1::numNvSwitches**

Number of entries in nvSwitches[] that are populated.

**struct dcgmNvLinkNvSwitchLinkStatus\_t  
dcgmNvLinkStatus\_v1::nvSwitches**

Per-NvSwitch link statuses.

## 3.52. dcgmPidInfo\_v1 Struct Reference

To store process statistics

**unsigned int dcgmPidInfo\_v1::version**

Version of this message (dcgmPidInfo\_version).

**unsigned int dcgmPidInfo\_v1::pid**

PID of the process.

**int dcgmPidInfo\_v1::numGpus**

Number of GPUs that are valid in GPUs.

**struct dcgmPidSingleInfo\_t dcgmPidInfo\_v1::summary**

Summary information for all GPUs listed in gpus[].

**struct dcgmPidSingleInfo\_t dcgmPidInfo\_v1::gpus**

Per-GPU information for this PID.

### 3.53. dcgmPidSingleInfo\_t Struct Reference

Info corresponding to single PID

**unsigned int dcgmPidSingleInfo\_t::gpuld**

ID of the GPU this pertains to. GPU\_ID\_INVALID = summary information for multiple GPUs.

**long long dcgmPidSingleInfo\_t::energyConsumed**

Energy consumed by the gpu in milliwatt-seconds.

**struct dcgmStatSummaryInt64\_t****dcgmPidSingleInfo\_t::pcieRxBandwidth**

PCI-E bytes read from the GPU.

**struct dcgmStatSummaryInt64\_t****dcgmPidSingleInfo\_t::pcieTxBandwidth**

PCI-E bytes written to the GPU.

**long long dcgmPidSingleInfo\_t::pcieReplays**

Count of PCI-E replays that occurred.

**long long dcgmPidSingleInfo\_t::startTime**

Process start time in microseconds since 1970.

**long long dcgmPidSingleInfo\_t::endTime**

Process end time in microseconds since 1970 or reported as 0 if the process is not completed.

**struct dcgmProcessUtilInfo\_t****dcgmPidSingleInfo\_t::processUtilization**

Process SM and Memory Utilization (in percent).

**struct dcgmStatSummaryInt32\_t****dcgmPidSingleInfo\_t::smUtilization**

GPU SM Utilization in percent.

**struct dcgmStatSummaryInt32\_t****dcgmPidSingleInfo\_t::memoryUtilization**

GPU Memory Utilization in percent.

**unsigned int dcgmPidSingleInfo\_t::eccSingleBit**

Count of ECC single bit errors that occurred.

**unsigned int dcgmPidSingleInfo\_t::eccDoubleBit**

Count of ECC double bit errors that occurred.

**struct dcgmStatSummaryInt32\_t  
dcgmPidSingleInfo\_t::memoryClock**

Memory clock in MHz.

**struct dcgmStatSummaryInt32\_t  
dcgmPidSingleInfo\_t::smClock**

SM clock in MHz.

**int dcgmPidSingleInfo\_t::numXidCriticalErrors**

Number of valid entries in xidCriticalErrorsTs.

**long long dcgmPidSingleInfo\_t::xidCriticalErrorsTs**

Timestamps of the critical XID errors that occurred.

**int dcgmPidSingleInfo\_t::numOtherComputePids**

Count of otherComputePids entries that are valid.

**unsigned int dcgmPidSingleInfo\_t::otherComputePids**

Other compute processes that ran. 0=no process.

**int dcgmPidSingleInfo\_t::numOtherGraphicsPids**

Count of otherGraphicsPids entries that are valid.

**unsigned int dcgmPidSingleInfo\_t::otherGraphicsPids**

Other graphics processes that ran. 0=no process.

**long long dcgmPidSingleInfo\_t::maxGpuMemoryUsed**

Maximum amount of GPU memory that was used in bytes.

**long long dcgmPidSingleInfo\_t::powerViolationTime**

Number of microseconds we were at reduced clocks due to power violation.

**long long dcgmPidSingleInfo\_t::thermalViolationTime**

Number of microseconds we were at reduced clocks due to thermal violation.

**long long dcgmPidSingleInfo\_t::reliabilityViolationTime**

Amount of microseconds we were at reduced clocks due to the reliability limit.

**long long dcgmPidSingleInfo\_t::boardLimitViolationTime**

Amount of microseconds we were at reduced clocks due to being at the board's max voltage.

**long long dcgmPidSingleInfo\_t::lowUtilizationTime**

Amount of microseconds we were at reduced clocks due to low utilization.

**long long dcgmPidSingleInfo\_t::syncBoostTime**

Amount of microseconds we were at reduced clocks due to sync boost.

**dcgmHealthWatchResults\_t****dcgmPidSingleInfo\_t::overallHealth**

The overall health of the system . dcgmHealthWatchResults\_t.

**dcgmHealthSystems\_t dcgmPidSingleInfo\_t::system**

system to which this information belongs

**dcgmHealthWatchResults\_t dcgmPidSingleInfo\_t::health**

health of the specified system on this GPU

## 3.54. dcgmPolicy\_v1 Struct Reference

Define the structure that specifies a policy to be enforced for a GPU

**unsigned int dcgmPolicy\_v1::version**

version number (dcgmPolicy\_version)

**dcgmPolicyCondition\_t dcgmPolicy\_v1::condition**

Condition(s) to access dcgmPolicyCondition\_t.

**dcgmPolicyMode\_t dcgmPolicy\_v1::mode**

Mode of operation dcgmPolicyMode\_t.

**dcgmPolicyIsolation\_t dcgmPolicy\_v1::isolation**

Isolation level after a policy violation dcgmPolicyIsolation\_t.

**dcgmPolicyAction\_t dcgmPolicy\_v1::action**

Action to perform after a policy violation dcgmPolicyAction\_t action.

**dcgmPolicyValidation\_t dcgmPolicy\_v1::validation**

Validation to perform after action is taken dcgmPolicyValidation\_t.

**dcgmPolicyFailureResp\_t dcgmPolicy\_v1::response**

Failure to validation response dcgmPolicyFailureResp\_t.

**struct dcgmPolicyConditionParms\_t****dcgmPolicy\_v1::parms**

Parameters for the condition fields.

## 3.55. dcgmPolicyCallbackResponse\_v1 Struct Reference

Define the structure that is given to the callback function

**unsigned int dcgmPolicyCallbackResponse\_v1::version**

version number (dcgmPolicyCallbackResponse\_version)

**dcgmPolicyCondition\_t**

**dcgmPolicyCallbackResponse\_v1::condition**

Condition that was violated.

**struct dcgmPolicyConditionDbe\_t**

**dcgmPolicyCallbackResponse\_v1::dbe**

ECC DBE return structure.

**struct dcgmPolicyConditionPci\_t**

**dcgmPolicyCallbackResponse\_v1::pci**

PCI replay error return structure.

**struct dcgmPolicyConditionMpr\_t**

**dcgmPolicyCallbackResponse\_v1::mpr**

Max retired pages limit return structure.

**struct dcgmPolicyConditionThermal\_t**

**dcgmPolicyCallbackResponse\_v1::thermal**

Thermal policy violations return structure.

**struct dcgmPolicyConditionPower\_t**

**dcgmPolicyCallbackResponse\_v1::power**

Power policy violations return structure.

**struct dcgmPolicyConditionNvlink\_t**

**dcgmPolicyCallbackResponse\_v1::nvlink**

Nvlink policy violations return structure.

**struct dcgmPolicyConditionXID\_t**

**dcgmPolicyCallbackResponse\_v1::xid**

XID policy violations return structure.

## 3.56. dcgmPolicyConditionDbe\_t Struct Reference

Define the ECC DBE return structure

**long long dcgmPolicyConditionDbe\_t::timestamp**

timestamp of the error

**enum dcgmPolicyConditionDbe\_t::@5**

**dcgmPolicyConditionDbe\_t::location**

location of the error

**unsigned int dcgmPolicyConditionDbe\_t::numerrors**

number of errors

## 3.57. dcgmPolicyConditionMpr\_t Struct Reference

Define the maximum pending retired pages limit return structure

**long long dcgmPolicyConditionMpr\_t::timestamp**

timestamp of the error

**unsigned int dcgmPolicyConditionMpr\_t::sbepages**

number of pending pages due to SBE

**unsigned int dcgmPolicyConditionMpr\_t::dbepages**

number of pending pages due to DBE

## 3.58. dcgmPolicyConditionNvlink\_t Struct Reference

Define the nvlink policy violations return structure

**long long dcgmPolicyConditionNvlink\_t::timestamp**

timestamp of the error

**unsigned short dcgmPolicyConditionNvlink\_t::fieldId**

Nvlink counter field ID that violated policy.

**unsigned int dcgmPolicyConditionNvlink\_t::counter**

Nvlink counter value that violated policy.

## 3.59. dcgmPolicyConditionParms\_t Struct Reference

Structure for policy condition parameters. This structure contains a tag that represents the type of the value being passed as well as a "val" which is a union of the possible value types. For example, to pass a true boolean: tag = BOOL, val.boolean = 1.

## 3.60. dcgmPolicyConditionPci\_t Struct Reference

Define the PCI replay error return structure

**long long dcgmPolicyConditionPci\_t::timestamp**

timestamp of the error

**unsigned int dcgmPolicyConditionPci\_t::counter**

value of the PCIe replay counter

## 3.61. dcgmPolicyConditionPower\_t Struct Reference

Define the power policy violations return structure

**long long dcgmPolicyConditionPower\_t::timestamp**

timestamp of the error

**unsigned int**

**dcgmPolicyConditionPower\_t::powerViolation**

Power value reached that violated policy.

## 3.62. dcgmPolicyConditionThermal\_t Struct Reference

Define the thermal policy violations return structure

**long long dcgmPolicyConditionThermal\_t::timestamp**

timestamp of the error

**unsigned int**

**dcgmPolicyConditionThermal\_t::thermalViolation**

Temperature reached that violated policy.

## 3.63. dcgmPolicyConditionXID\_t Struct Reference

Define the xid policy violations return structure

**long long dcgmPolicyConditionXID\_t::timestamp**

Timestamp of the error.

**unsigned int dcgmPolicyConditionXID\_t::errnum**

The XID error number.

## 3.64. dcgmPolicyViolationNotify\_t Struct Reference

Structure to fill when a user queries for policy violations

**unsigned int dcgmPolicyViolationNotify\_t::gpuld**  
gpu ID

**unsigned int**  
**dcgmPolicyViolationNotify\_t::violationOccurred**  
a violation based on the bit values in dcgmPolicyCondition\_t

## 3.65. dcgmProcessUtilInfo\_t Struct Reference

per process utilization rates

## 3.66. dcgmProcessUtilSample\_t Struct Reference

Internal structure used to get the PID and the corresponding utilization rate

## 3.67. dcgmProfUnwatchFields\_v1 Struct Reference

Structure to pass to dcgmProfUnwatchFields when unwatching profiling metrics

**dcgmGpuGrp\_t dcgmProfUnwatchFields\_v1::groupId**

Version of this request. Should be dcgmProfUnwatchFields\_version

**unsigned int dcgmProfUnwatchFields\_v1::flags**

Group ID representing collection of one or more GPUs. Look at [dcgmGroupCreate](#) for details on creating the group. Alternatively, pass in the group id as DCGM\_GROUP\_ALL\_GPUS to perform operation on all the GPUs. The GPUs of the group must all be identical or DCGM\_ST\_GROUP\_INCOMPATIBLE will be returned by this API.

## 3.68. dcgmProfWatchFields\_v1 Struct Reference

Structure to pass to [dcgmProfWatchFields\(\)](#) when watching profiling metrics

**dcgmGpuGrp\_t dcgmProfWatchFields\_v1::groupId**

Version of this request. Should be dcgmProfWatchFields\_version

## **unsigned int dcgmProfWatchFields\_v1::numFieldIds**

Group ID representing collection of one or more GPUs. Look at `dcgmGroupCreate` for details on creating the group. Alternatively, pass in the group id as `DCGM_GROUP_ALL_GPUS` to perform operation on all the GPUs. The GPUs of the group must all be identical or `DCGM_ST_GROUP_INCOMPATIBLE` will be returned by this API.

## **unsigned short dcgmProfWatchFields\_v1::fieldIds**

Number of field IDs that are being passed in `fieldIds[]`

## **long long dcgmProfWatchFields\_v1::updateFreq**

`DCGM_FI_PROF_?` field IDs to watch

## **double dcgmProfWatchFields\_v1::maxKeepAge**

How often to update this field in usec. Note that profiling metrics may need to be sampled more frequently than this value. See `dcgmProfMetricGroupInfo_t.minUpdateFreqUsec` of the metric group matching `metricGroupTag` to see what this minimum is. If `minUpdateFreqUsec < updateFreq` then samples will be aggregated to `updateFreq` intervals in DCGM's internal cache.

## **int dcgmProfWatchFields\_v1::maxKeepSamples**

How long to keep data for every `fieldId` in seconds

## **unsigned int dcgmProfWatchFields\_v1::flags**

Maximum number of samples to keep for each `fieldId`. 0=no limit

## **3.69. dcgmRunningProcess\_v1 Struct Reference**

Running process information for a compute or graphics process

**unsigned int dcgmRunningProcess\_v1::version**

Version of this message (dcgmRunningProcess\_version).

**unsigned int dcgmRunningProcess\_v1::pid**

PID of the process.

**unsigned long long****dcgmRunningProcess\_v1::memoryUsed**

GPU memory used by this process in bytes.

## 3.70. dcgmStatSummaryFp64\_t Struct Reference

Summary of time series data in double-precision format. Each value will either be set or be a BLANK value. Check for blank with the `DCGM_FP64_IS_BLANK()` macro. See `dcmvalue.h` for the actual values of BLANK values

**double dcgmStatSummaryFp64\_t::minValue**

Minimum value of the samples looked at.

**double dcgmStatSummaryFp64\_t::maxValue**

Maximum value of the samples looked at.

**double dcgmStatSummaryFp64\_t::average**

Simple average of the samples looked at. Blank values are ignored for this calculation.

## 3.71. dcgmStatSummaryInt32\_t Struct Reference

Same as `dcgmStatSummaryInt64_t`, but with 32-bit integer values

**int dcgmStatSummaryInt32\_t::minValue**

Minimum value of the samples looked at.

**int dcgmStatSummaryInt32\_t::maxValue**

Maximum value of the samples looked at.

**int dcgmStatSummaryInt32\_t::average**

Simple average of the samples looked at. Blank values are ignored for this calculation.

## 3.72. dcgmStatSummaryInt64\_t Struct Reference

Summary of time series data in int64 format. Each value will either be set or be a BLANK value. Check for blank with the `DCGM_INT64_IS_BLANK()` macro. See `dcmgvalue.h` for the actual values of BLANK values

**long long dcgmStatSummaryInt64\_t::minValue**

Minimum value of the samples looked at.

**long long dcgmStatSummaryInt64\_t::maxValue**

Maximum value of the samples looked at.

**long long dcgmStatSummaryInt64\_t::average**

Simple average of the samples looked at. Blank values are ignored for this calculation.

## 3.73. dcgmVersionInfo\_v1 Struct Reference

Structure to describe the DCGM build environment

**unsigned int dcgmVersionInfo\_v1::version**

Version of this message.

**char dcgmVersionInfo\_v1::changelist**

Changelist number from which DCGM was built.

**char dcgmVersionInfo\_v1::platform**

Builder platform - uname result without hostname.

**char dcgmVersionInfo\_v1::branch**

Name of the branch where DCGM was built.

**char dcgmVersionInfo\_v1::driverVersion**

The version of NVidia driver DCGM was linked with.

**char dcgmVersionInfo\_v1::buildDate**

Date of the build.

## 3.74. dcgmVgpuConfig\_v1 Struct Reference

Structure to represent default and target vgpu configuration for a device

**unsigned int dcgmVgpuConfig\_v1::version**

Version number (dcgmConfig\_version).

**unsigned int dcgmVgpuConfig\_v1::gpuld**

GPU ID.

**unsigned int dcgmVgpuConfig\_v1::eccMode**

ECC Mode (0: Disabled, 1 : Enabled, DCGM\_INT32\_BLANK : Ignored).

**unsigned int dcgmVgpuConfig\_v1::computeMode**

Compute Mode (One of DCGM\_CONFIG\_COMPUTEMODE\_? OR DCGM\_INT32\_BLANK to Ignore).

**struct dcgmConfigPerfStateSettings\_t****dcgmVgpuConfig\_v1::perfState**

Performance State Settings (clocks / boost mode).

**struct dcgmConfigPowerLimit\_t****dcgmVgpuConfig\_v1::powerLimit**

Power Limits.

## 3.75. dcgmVgpuDeviceAttributes\_v6 Struct Reference

Represents the vGPU attributes corresponding to a physical device

**unsigned int dcgmVgpuDeviceAttributes\_v6::version**

Version number (dcgmVgpuDeviceAttributes\_version).

**unsigned int****dcgmVgpuDeviceAttributes\_v6::activeVgpuInstanceCount**

Count of active vGPU instances on the device.

**unsigned int****dcgmVgpuDeviceAttributes\_v6::activeVgpuInstancelds**

List of vGPU instances.

**unsigned int****dcgmVgpuDeviceAttributes\_v6::creatableVgpuTypeCount**

Creatable vGPU type count.

**unsigned int****dcgmVgpuDeviceAttributes\_v6::creatableVgpuTypelds**

List of Creatable vGPU types.

**unsigned int****dcgmVgpuDeviceAttributes\_v6::supportedVgpuTypeCount**

Supported vGPU type count.

**struct dcgmDeviceVgpuTypeInfo\_t****dcgmVgpuDeviceAttributes\_v6::supportedVgpuTypeInfo**

Info related to vGPUs supported on the device.

**struct dcgmDeviceVgpuUtilInfo\_t****dcgmVgpuDeviceAttributes\_v6::vgpuUtilInfo**

Utilizations specific to vGPU instance.

**unsigned int dcgmVgpuDeviceAttributes\_v6::gpuUtil**

GPU utilization.

**unsigned int****dcgmVgpuDeviceAttributes\_v6::memCopyUtil**

Memory utilization.

**unsigned int dcgmVgpuDeviceAttributes\_v6::encUtil**

Encoder utilization.

**unsigned int dcgmVgpuDeviceAttributes\_v6::decUtil**

Decoder utilization.

## 3.76. dcgmVgpuInstanceAttributes\_v1 Struct Reference

Represents attributes specific to vGPU instance

**unsigned int dcgmVgpuInstanceAttributes\_v1::version**

Version number (dcgmVgpuInstanceAttributes\_version).

**char dcgmVgpuInstanceAttributes\_v1::vmlId**

VM ID of the vGPU instance.

**char dcgmVgpuInstanceAttributes\_v1::vmName**

VM name of the vGPU instance.

**unsigned int****dcgmVgpuInstanceAttributes\_v1::vgpuTypeid**

Type ID of the vGPU instance.

**char dcgmVgpuInstanceAttributes\_v1::vgpuUuid**

UUID of the vGPU instance.

**char dcgmVgpuInstanceAttributes\_v1::vgpuDriverVersion**

Driver version of the vGPU instance.

**unsigned int dcgmVgpuInstanceAttributes\_v1::fbUsage**

Fb usage of the vGPU instance.

**unsigned int****dcgmVgpuInstanceAttributes\_v1::licenseStatus**

License status of the vGPU instance.

**unsigned int****dcgmVgpuInstanceAttributes\_v1::frameRateLimit**

Frame rate limit of the vGPU instance.

# Chapter 4. DATA FIELDS

Here is a list of all documented struct and union fields with links to the struct/union documentation for each field:

## A

### **action**

[dcgmPolicy\\_v1](#)

### **activeTimeUsec**

[dcgmDevicePidAccountingStats\\_v1](#)

### **activeVgpuInstanceCount**

[dcgmVgpuDeviceAttributes\\_v6](#)

### **activeVgpuInstanceIds**

[dcgmVgpuDeviceAttributes\\_v6](#)

### **addressIsUnixSocket**

[dcgmConnectV2Params\\_v2](#)

### **aggregateInfo**

[dcgmIntrospectFullFieldsExecTime\\_v1](#)

[dcgmIntrospectFullMemory\\_v1](#)

### **average**

[dcgmStatSummaryInt64\\_t](#)

[dcgmStatSummaryInt32\\_t](#)

[dcgmStatSummaryFp64\\_t](#)

### **averageFps**

[dcgmDeviceEncStats\\_v1](#)

[dcgmDeviceFbcSessionInfo\\_v1](#)

[dcgmDeviceVgpuEncSessions\\_v1](#)

[dcgmDeviceFbcStats\\_v1](#)

### **averageLatency**

[dcgmDeviceFbcStats\\_v1](#)

[dcgmDeviceFbcSessionInfo\\_v1](#)

[dcgmDeviceVgpuEncSessions\\_v1](#)

`dcmDeviceEncStats_v1`

## B

**bar1Total**

`dcmDeviceMemoryUsage_v1`

**blacklist**

`dcmDiagResponse_v3`

**blob**

`dcmFieldValue_v2`

`dcmFieldValue_v1`

**boardLimitViolationTime**

`dcmPidSingleInfo_t`

`dcmGpuUsageInfo_t`

**branch**

`dcmVersionInfo_v1`

**brandName**

`dcmDeviceIdentifiers_v1`

**buildDate**

`dcmVersionInfo_v1`

**bytesUsed**

`dcmIntrospectMemory_v1`

## C

**changelist**

`dcmVersionInfo_v1`

**clockSet**

`dcmDeviceSupportedClockSets_v1`

**clockSets**

`dcmDeviceAttributes_v1`

**codecType**

`dcmDeviceVgpuEncSessions_v1`

**computeMode**

`dcmConfig_v1`

`dcmVgpuConfig_v1`

**computePidInfo**

`dcmGpuUsageInfo_t`

**condition**

`dcmPolicy_v1`

`dcmPolicyCallbackResponse_v1`

**contextId**

`dcmIntrospectContext_v1`

**count**

`dcmGroupInfo_v1`

```

dcgmGroupInfo_v2
dcgmDeviceSupportedClockSets_v1
counter
  dcgmPolicyConditionNvlink_t
  dcgmPolicyConditionPci_t
cpuAffinityMask
  dcgmDeviceTopology_v1
creatableVgpuTypeCount
  dcgmVgpuDeviceAttributes_v6
creatableVgpuTypeIds
  dcgmVgpuDeviceAttributes_v6
cudaMainLibrary
  dcgmDiagResponse_v3
cudaRuntimeLibrary
  dcgmDiagResponse_v3
curPowerLimit
  dcgmDevicePowerLimits_v1

```

**D**

```

dbe
  dcgmPolicyCallbackResponse_v1
dbepages
  dcgmPolicyConditionMpr_t
dbl
  dcgmFieldValue_v2
  dcgmFieldValue_v1
decUtil
  dcgmDeviceVgpuUtilInfo_v1
  dcgmDeviceVgpuProcessUtilInfo_v1
  dcgmVgpuDeviceAttributes_v6
defaultPowerLimit
  dcgmDevicePowerLimits_v1
deviceId
  dcgmDeviceVgpuTypeInfo_v1
deviceName
  dcgmDeviceIdentifiers_v1
displayOrdinal
  dcgmDeviceFbcSessionInfo_v1
driverVersion
  dcgmDeviceIdentifiers_v1
  dcgmVersionInfo_v1

```

**E****eccDoubleBit**

`dcmPidSingleInfo_t`  
`dcmGpuUsageInfo_t`

**eccMode**

`dcmVgpuConfig_v1`  
`dcmConfig_v1`

**eccSingleBit**

`dcmPidSingleInfo_t`  
`dcmGpuUsageInfo_t`

**encUtil**

`dcmDeviceVgpuProcessUtilInfo_v1`  
`dcmVgpuDeviceAttributes_v6`  
`dcmDeviceVgpuUtilInfo_v1`

**endTime**

`dcmPidSingleInfo_t`  
`dcmGpuUsageInfo_t`

**energyConsumed**

`dcmPidSingleInfo_t`  
`dcmGpuUsageInfo_t`

**enforcedPowerLimit**

`dcmDevicePowerLimits_v1`

**entityCount**

`dcmHealthResponse_v3`  
`dcmHealthResponse_v2`

**entityGroupId**

`dcmGroupEntityPair_t`  
`dcmFieldValue_v2`  
`dcmHealthResponse_v2`  
`dcmHealthResponse_v3`

**entityId**

`dcmNvLinkGpuLinkStatus_t`  
`dcmNvLinkNvSwitchLinkStatus_t`  
`dcmFieldValue_v2`  
`dcmHealthResponse_v2`  
`dcmHealthResponse_v3`  
`dcmGroupEntityPair_t`

**entityList**

`dcmGroupInfo_v2`

**environment**

`dcmDiagResponse_v3`

**errnum**

`dcmPolicyConditionXID_t`

**errorCount**  
  dcgmHealthResponse\_v3  
**errors**  
  dcgmHealthResponse\_v3  
**errorString**  
  dcgmHealthResponse\_v1  
  dcgmHealthResponse\_v2

**F**

**fbFree**  
  dcgmDeviceMemoryUsage\_v1

**fbTotal**  
  dcgmDeviceMemoryUsage\_v1  
  dcgmDeviceVgpuTypeInfo\_v1

**fbUsage**  
  dcgmVgpuInstanceAttributes\_v1

**fbUsed**  
  dcgmDeviceMemoryUsage\_v1

**fieldGroupId**  
  dcgmFieldGroupInfo\_v1  
  dcgmIntrospectContext\_v1

**fieldGroupName**  
  dcgmFieldGroupInfo\_v1

**fieldId**  
  dcgmErrorInfo\_t  
  dcgmPolicyConditionNvlink\_t  
  dcgmFieldValue\_v1  
  dcgmFieldValue\_v2  
  dcgmIntrospectContext\_v1

**fieldIds**  
  dcgmProfWatchFields\_v1  
  dcgmFieldGroupInfo\_v1

**fieldType**  
  dcgmFieldValue\_v1  
  dcgmFieldValue\_v2

**flags**  
  dcgmProfWatchFields\_v1  
  dcgmProfUnwatchFields\_v1

**frameRateLimit**  
  dcgmVgpuInstanceAttributes\_v1  
  dcgmDeviceVgpuTypeInfo\_v1

**G**

**globalInfo**

- dcgmIntrospectFullFieldsExecTime\_v1
- dcgmIntrospectFullMemory\_v1

**gpuCount**

- dcgmDiagResponse\_v3
- dcgmDiagResponse\_v4
- dcgmHealthResponse\_v1

**gpuId**

- dcgmVgpuConfig\_v1
- dcgmPolicyViolationNotify\_t
- dcgmHealthResponse\_v1
- dcgmPidSingleInfo\_t
- dcgmErrorInfo\_t
- dcgmGpuUsageInfo\_t
- dcgmDiagResponsePerGpu\_v1
- dcgmConfig\_v1
- dcgmDeviceTopology\_v1

**gpuIdList**

- dcgmGroupInfo\_v1

**gpuIdsForGpuInfo**

- dcgmIntrospectFullFieldsExecTime\_v1
- dcgmIntrospectFullMemory\_v1

**gpuInfo**

- dcgmIntrospectFullMemory\_v1
- dcgmIntrospectFullFieldsExecTime\_v1

**gpuInfoCount**

- dcgmIntrospectFullFieldsExecTime\_v1
- dcgmIntrospectFullMemory\_v1

**gpus**

- dcgmPidInfo\_v1
- dcgmNvLinkStatus\_v1
- dcgmJobInfo\_v2

**gpuUtil**

- dcgmVgpuDeviceAttributes\_v6

**gpuUtilization**

- dcgmDevicePidAccountingStats\_v1

**graphicsPidInfo**

- dcgmGpuUsageInfo\_t

**graphicsProcesses**

- dcgmDiagResponse\_v3

**groupCpuAffinityMask**

- dcgmGroupTopology\_v1

```

groupId
    dcgmProfUnwatchFields_v1
    dcgmProfWatchFields_v1
groupName
    dcgmGroupInfo_v1
    dcgmGroupInfo_v2

H
hasGlobalInfo
    dcgmIntrospectFullFieldsExecTime_v1
    dcgmIntrospectFullMemory_v1
health
    dcgmHealthResponse_v2
    dcgmHealthResponse_v3
    dcgmHealthResponse_v1
    dcgmGpuUsageInfo_t
    dcgmPidSingleInfo_t
hMaxResolution
    dcgmDeviceFbcSessionInfo_v1
hResolution
    dcgmDeviceFbcSessionInfo_v1
    dcgmDeviceVgpuEncSessions_v1
hwDiagnosticReturn
    dcgmDiagResponsePerGpu_v1

I
i64
    dcgmFieldValue_v1
    dcgmFieldValue_v2
id
    dcgmModuleGetStatusesModule_t
identifiers
    dcgmDeviceAttributes_v1
incidentCount
    dcgmHealthResponse_v1
    dcgmHealthResponse_v3
    dcgmHealthResponse_v2
inforom
    dcgmDiagResponse_v3
inforomImageVersion
    dcgmDeviceIdentifiers_v1
introspectLvl
    dcgmIntrospectContext_v1

```

**isolation**`dcgmPolicy_v1`**K****kernel**`dcgmIntrospectCpuUtil_v1`**L****levelOneResults**`dcgmDiagResponse_v4`**levelOneTestCount**`dcgmDiagResponse_v4`**licenseStatus**`dcgmVgpuInstanceAttributes_v1`**linkState**`dcgmNvLinkGpuLinkStatus_t``dcgmNvLinkNvSwitchLinkStatus_t`**localNvLinkId**`dcgmDeviceTopology_v1`**location**`dcgmPolicyConditionDbe_t`**lowUtilizationTime**`dcgmPidSingleInfo_t``dcgmGpuUsageInfo_t`**M****maxGpuMemoryUsed**`dcgmPidSingleInfo_t``dcgmGpuUsageInfo_t`**maxInstances**`dcgmDeviceVgpuTypeInfo_v1`**maxKeepAge**`dcgmProfWatchFields_v1`**maxKeepSamples**`dcgmProfWatchFields_v1`**maxMemoryUsage**`dcgmDevicePidAccountingStats_v1`**maxPowerLimit**`dcgmDevicePowerLimits_v1`**maxResolutionX**`dcgmDeviceVgpuTypeInfo_v1`**maxResolutionY**`dcgmDeviceVgpuTypeInfo_v1`

**maxValue**  
`dcmStatSummaryInt64_t`  
`dcmStatSummaryInt32_t`  
`dcmStatSummaryFp64_t`

**meanUpdateFreqUsec**  
`dcmIntrospectFieldsExecTime_v1`

**memClock**  
`dcmClockSet_v1`

**memCopyUtil**  
`dcmVgpuDeviceAttributes_v6`

**memoryClock**  
`dcmPidSingleInfo_t`  
`dcmGpuUsageInfo_t`

**memoryUsage**  
`dcmDeviceAttributes_v1`

**memoryUsed**  
`dcmRunningProcess_v1`

**memoryUtilization**  
`dcmDevicePidAccountingStats_v1`  
`dcmPidSingleInfo_t`  
`dcmGpuUsageInfo_t`

**memUtil**  
`dcmDeviceVgpuUtilInfo_v1`  
`dcmDeviceVgpuProcessUtilInfo_v1`

**minPowerLimit**  
`dcmDevicePowerLimits_v1`

**minValue**  
`dcmStatSummaryInt32_t`  
`dcmStatSummaryFp64_t`  
`dcmStatSummaryInt64_t`

**mode**  
`dcmPolicy_v1`

**mp**  
`dcmPolicyCallbackResponse_v1`

**N**

**numaOptimalFlag**  
`dcmGroupTopology_v1`

**numComputePids**  
`dcmGpuUsageInfo_t`

**numDisplayHeads**  
`dcmDeviceVgpuTypeInfo_v1`

**numerrors**  
`dcmPolicyConditionDbe_t`

**numFieldIds**  
`dcmFieldGroupInfo_v1`  
`dcmProfWatchFields_v1`

**numGpus**  
`dcmJobInfo_v2`  
`dcmDeviceTopology_v1`  
`dcmNvLinkStatus_v1`  
`dcmPidInfo_v1`

**numGraphicsPids**  
`dcmGpuUsageInfo_t`

**numNvSwitches**  
`dcmNvLinkStatus_v1`

**numOtherComputePids**  
`dcmPidSingleInfo_t`

**numOtherGraphicsPids**  
`dcmPidSingleInfo_t`

**numXidCriticalErrors**  
`dcmGpuUsageInfo_t`  
`dcmPidSingleInfo_t`

**nvlink**  
`dcmPolicyCallbackResponse_v1`

**nvmlLibrary**  
`dcmDiagResponse_v3`

**nvSwitches**  
`dcmNvLinkStatus_v1`

**O**

**otherComputePids**  
`dcmPidSingleInfo_t`

**otherGraphicsPids**  
`dcmPidSingleInfo_t`

**overallHealth**  
`dcmPidSingleInfo_t`  
`dcmHealthResponse_v3`  
`dcmHealthResponse_v2`  
`dcmHealthResponse_v1`  
`dcmGpuUsageInfo_t`

**P**

**pageRetirement**  
`dcmDiagResponse_v3`

```

parms
    dcgmPolicy_v1
path
    dcgmDeviceTopology_v1
pci
    dcgmPolicyCallbackResponse_v1
pciBusId
    dcgmDeviceIdentifiers_v1
pciDeviceId
    dcgmDeviceIdentifiers_v1
pcieReplays
    dcgmGpuUsageInfo_t
    dcgmPidSingleInfo_t
pcieRxBandwidth
    dcgmPidSingleInfo_t
    dcgmGpuUsageInfo_t
pcieTxBandwidth
    dcgmPidSingleInfo_t
    dcgmGpuUsageInfo_t
pciSubSystemId
    dcgmDeviceIdentifiers_v1
perfState
    dcgmConfig_v1
    dcgmVgpuConfig_v1
perGpuResponses
    dcgmDiagResponse_v3
    dcgmDiagResponse_v4
permissions
    dcgmDiagResponse_v3
persistAfterDisconnect
    dcgmConnectV2Params_v1
    dcgmConnectV2Params_v2
persistenceMode
    dcgmDiagResponse_v3
pid
    dcgmDeviceFbcSessionInfo_v1
    dcgmDevicePidAccountingStats_v1
    dcgmDeviceVgpuEncSessions_v1
    dcgmDeviceVgpuProcessUtilInfo_v1
    dcgmPidInfo_v1
    dcgmRunningProcess_v1
platform
    dcgmVersionInfo_v1

```

**power**  
   dcgmPolicyCallbackResponse\_v1

**powerLimit**  
   dcgmVgpuConfig\_v1  
   dcgmConfig\_v1

**powerLimits**  
   dcgmDeviceAttributes\_v1

**powerUsage**  
   dcgmGpuUsageInfo\_t

**powerViolation**  
   dcgmPolicyConditionPower\_t

**powerViolationTime**  
   dcgmPidSingleInfo\_t  
   dcgmGpuUsageInfo\_t

**processName**  
   dcgmDeviceVgpuProcessUtilInfo\_v1

**processUtilization**  
   dcgmPidSingleInfo\_t

**R**

**recentUpdateUsec**  
   dcgmIntrospectFieldsExecTime\_v1

**reliabilityViolationTime**  
   dcgmPidSingleInfo\_t  
   dcgmGpuUsageInfo\_t

**response**  
   dcgmPolicy\_v1

**results**  
   dcgmDiagResponsePerGpu\_v1

**S**

**sbepages**  
   dcgmPolicyConditionMpr\_t

**serial**  
   dcgmDeviceIdentifiers\_v1

**sessionCount**  
   dcgmDeviceFbcStats\_v1  
   dcgmDeviceFbcSessions\_v1  
   dcgmDeviceEncStats\_v1

**sessionFlags**  
   dcgmDeviceFbcSessionInfo\_v1

**sessionId**  
   dcgmDeviceVgpuEncSessions\_v1

```
dcgmDeviceFbcSessionInfo_v1
sessionInfo
    dcgmDeviceFbcSessions_v1
sessionType
    dcgmDeviceFbcSessionInfo_v1
shutdownTemp
    dcgmDeviceThermals_v1
slowdownTemp
    dcgmDeviceThermals_v1
lowestPath
    dcgmGroupTopology_v1
smClock
    dcgmGpuUsageInfo_t
    dcgmClockSet_v1
    dcgmPidSingleInfo_t
smUtil
    dcgmDeviceVgpuUtilInfo_v1
    dcgmDeviceVgpuProcessUtilInfo_v1
smUtilization
    dcgmPidSingleInfo_t
    dcgmGpuUsageInfo_t
startTime
    dcgmPidSingleInfo_t
    dcgmGpuUsageInfo_t
startTimestamp
    dcgmDevicePidAccountingStats_v1
status
    dcgmModuleGetStatusesModule_t
    dcgmErrorInfo_t
    dcgmFieldValue_v1
    dcgmFieldValue_v2
str
    dcgmFieldValue_v1
    dcgmFieldValue_v2
subsystemId
    dcgmDeviceVgpuTypeInfo_v1
summary
    dcgmJobInfo_v2
    dcgmPidInfo_v1
supportedVgpuTypeCount
    dcgmVgpuDeviceAttributes_v6
supportedVgpuTypeInfo
    dcgmVgpuDeviceAttributes_v6
```

```

syncBoost
    dcgmConfigPerfStateSettings_t
syncBoostTime
    dcgmPidSingleInfo_t
    dcgmGpuUsageInfo_t
system
    dcgmGpuUsageInfo_t
    dcgmHealthResponse_v2
    dcgmHealthResponse_v3
    dcgmPidSingleInfo_t
    dcgmHealthResponse_v1
systemError
    dcgmDiagResponse_v4
    dcgmDiagResponse_v3

```

**T**

```

targetClocks
    dcgmConfigPerfStateSettings_t
thermal
    dcgmPolicyCallbackResponse_v1
thermalSettings
    dcgmDeviceAttributes_v1
thermalViolation
    dcgmPolicyConditionThermal_t
thermalViolationTime
    dcgmPidSingleInfo_t
    dcgmGpuUsageInfo_t
timeoutMs
    dcgmConnectV2Params_v2
timestamp
    dcgmPolicyConditionDbe_t
    dcgmPolicyConditionPci_t
    dcgmPolicyConditionMpr_t
    dcgmPolicyConditionThermal_t
    dcgmPolicyConditionPower_t
    dcgmPolicyConditionNvlink_t
    dcgmPolicyConditionXID_t
total
    dcgmIntrospectCpuUtil_v1
totalEverUpdateUsec
    dcgmIntrospectFieldsExecTime_v1
trainingMsg
    dcgmDiagResponse_v4

```

```

ts
    dcgmFieldValue_v1
    dcgmFieldValue_v2
type
    dcgmConfigPowerLimit_t

U
unused
    dcgmFieldValue_v2
unusedActiveVgpuInstanceCount
    dcgmDeviceAttributes_v1
unusedCreatableVgpuTypeCount
    dcgmDeviceVgpuIds_v1
unusedCreatableVgpuTypeIds
    dcgmDeviceVgpuIds_v1
unusedSupportedVgpuTypeCount
    dcgmDeviceVgpuIds_v1
unusedSupportedVgpuTypeIds
    dcgmDeviceVgpuIds_v1
unusedVgpuIds
    dcgmDeviceAttributes_v1
unusedVgpuInstanceIds
    dcgmDeviceAttributes_v1
updateFreq
    dcgmProfWatchFields_v1
user
    dcgmIntrospectCpuUtil_v1
uuid
    dcgmDeviceIdentifiers_v1

V
val
    dcgmConfigPowerLimit_t
validation
    dcgmPolicy_v1
value
    dcgmFieldValue_v2
    dcgmFieldValue_v1
vbios
    dcgmDeviceIdentifiers_v1
version
    dcgmGroupInfo_v1
    dcgmDeviceIdentifiers_v1

```

dcgmPolicy\_v1  
dcgmPolicyCallbackResponse\_v1  
dcgmDeviceMemoryUsage\_v1  
dcgmFieldValue\_v1  
dcgmFieldValue\_v2  
dcgmGroupInfo\_v2  
dcgmDeviceVgpuUtilInfo\_v1  
dcgmHealthResponse\_v1  
dcgmHealthResponse\_v2  
dcgmDeviceEncStats\_v1  
dcgmHealthResponse\_v3  
dcgmPidInfo\_v1  
dcgmFieldGroupInfo\_v1  
dcgmDeviceFbcStats\_v1  
dcgmVersionInfo\_v1  
dcgmNvLinkStatus\_v1  
dcgmIntrospectCpuUtil\_v1  
dcgmIntrospectFullMemory\_v1  
dcgmIntrospectMemory\_v1  
dcgmJobInfo\_v2  
dcgmIntrospectFieldsExecTime\_v1  
dcgmIntrospectContext\_v1  
dcgmGroupTopology\_v1  
dcgmDeviceTopology\_v1  
dcgmRunningProcess\_v1  
dcgmDeviceFbcSessionInfo\_v1  
dcgmDiagResponse\_v3  
dcgmDiagResponse\_v4  
dcgmClockSet\_v1  
dcgmDeviceFbcSessions\_v1  
dcgmDeviceVgpuEncSessions\_v1  
dcgmConnectV2Params\_v1  
dcgmDeviceSupportedClockSets\_v1  
dcgmDeviceVgpuProcessUtilInfo\_v1  
dcgmIntrospectFullFieldsExecTime\_v1  
dcgmDeviceVgpuIds\_v1  
dcgmVgpuConfig\_v1  
dcgmConfig\_v1  
dcgmVgpuInstanceAttributes\_v1  
dcgmVgpuDeviceAttributes\_v6  
dcgmDevicePidAccountingStats\_v1  
dcgmDeviceVgpuTypeInfo\_v1  
dcgmDeviceAttributes\_v1

**dcmConnectV2Params\_v2**  
**dcmDeviceThermals\_v1**  
**dcmDevicePowerLimits\_v1**  
**vgpuDriverVersion**  
 dcmVgpuInstanceAttributes\_v1  
**vgpuId**  
 dcmDeviceVgpuEncSessions\_v1  
 dcmDeviceVgpuProcessUtilInfo\_v1  
 dcmDeviceVgpuUtilInfo\_v1  
 dcmDeviceFbcSessionInfo\_v1  
**vgpuProcessSamplesCount**  
 dcmDeviceVgpuProcessUtilInfo\_v1  
**vgpuTypeClass**  
 dcmDeviceVgpuTypeInfo\_v1  
**vgpuTypeId**  
 dcmVgpuInstanceAttributes\_v1  
**vgpuTypeInfo**  
 dcmDeviceVgpuTypeInfo\_v1  
**vgpuTypeLicense**  
 dcmDeviceVgpuTypeInfo\_v1  
**vgpuTypeName**  
 dcmDeviceVgpuTypeInfo\_v1  
**vgpuUtilInfo**  
 dcmVgpuDeviceAttributes\_v6  
**vgpuUuid**  
 dcmVgpuInstanceAttributes\_v1  
**violationOccurred**  
 dcmPolicyViolationNotify\_t  
**virtualizationMode**  
 dcmDeviceIdentifiers\_v1  
**vMaxResolution**  
 dcmDeviceFbcSessionInfo\_v1  
**vmId**  
 dcmVgpuInstanceAttributes\_v1  
**vmName**  
 dcmVgpuInstanceAttributes\_v1  
**vResolution**  
 dcmDeviceFbcSessionInfo\_v1  
 dcmDeviceVgpuEncSessions\_v1  
  
**X**  
**xid**  
 dcmPolicyCallbackResponse\_v1

**xidCriticalErrorsTs**

dcgmGpuUsageInfo\_t  
dcgmPidSingleInfo\_t

## **Notice**

THE INFORMATION IN THIS GUIDE AND ALL OTHER INFORMATION CONTAINED IN NVIDIA DOCUMENTATION REFERENCED IN THIS GUIDE IS PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE INFORMATION FOR THE PRODUCT, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the product described in this guide shall be limited in accordance with the NVIDIA terms and conditions of sale for the product.

THE NVIDIA PRODUCT DESCRIBED IN THIS GUIDE IS NOT FAULT TOLERANT AND IS NOT DESIGNED, MANUFACTURED OR INTENDED FOR USE IN CONNECTION WITH THE DESIGN, CONSTRUCTION, MAINTENANCE, AND/OR OPERATION OF ANY SYSTEM WHERE THE USE OR A FAILURE OF SUCH SYSTEM COULD RESULT IN A SITUATION THAT THREATENS THE SAFETY OF HUMAN LIFE OR SEVERE PHYSICAL HARM OR PROPERTY DAMAGE (INCLUDING, FOR EXAMPLE, USE IN CONNECTION WITH ANY NUCLEAR, AVIONICS, LIFE SUPPORT OR OTHER LIFE CRITICAL APPLICATION). NVIDIA EXPRESSLY DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY OF FITNESS FOR SUCH HIGH RISK USES. NVIDIA SHALL NOT BE LIABLE TO CUSTOMER OR ANY THIRD PARTY, IN WHOLE OR IN PART, FOR ANY CLAIMS OR DAMAGES ARISING FROM SUCH HIGH RISK USES.

NVIDIA makes no representation or warranty that the product described in this guide will be suitable for any specified use without further testing or modification. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to ensure the product is suitable and fit for the application planned by customer and to do the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this guide. NVIDIA does not accept any liability related to any default, damage, costs or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this guide, or (ii) customer product designs.

Other than the right for customer to use the information in this guide with the product, no other license, either expressed or implied, is hereby granted by NVIDIA under this guide. Reproduction of information in this guide is permissible only if reproduction is approved by NVIDIA in writing, is reproduced without alteration, and is accompanied by all associated conditions, limitations, and notices.

## **Trademarks**

NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the United States and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

## **Copyright**

© 2013-2019 NVIDIA Corporation. All rights reserved.