



LAS SPECIFICATION

VERSION 1.4 – R6

15 August 2011

Approved:
TBD

Published by:
The American Society for Photogrammetry & Remote Sensing
5410 Grosvenor Lane, Suite 210
Bethesda, Maryland 20814-2160
Voice: 301-493-0290
Fax: 301-493-0208
Web: www.asprs.org

All rights reserved.

Copyright © 2002-2011 American Society for Photogrammetry and Remote Sensing (ASPRS). All rights reserved. **Permission to Use:** The copyright owner hereby consents to unlimited use and distribution of this document, or parts thereof, as a specification provided such use references ASPRS as the publisher. This consent does not extend to other uses such as general distribution in any form, including electronic, by any individual or organization whether for advertising or promotional purposes, for creating new collective works, or for resale. For these and all other purposes, reproduction of this publication or any part thereof (excluding short quotations for use in the preparation of reviews and technical and scientific papers) may be made only after obtaining the specific approval of the publisher.

Printed in the United States of America.

LAS FORMAT VERSION 1.4:

1 Purpose, scope, and applicability

The LAS file is intended to contain LIDAR (or other) point cloud data records. The data will generally be put into this format from software (e.g. provided by LIDAR hardware vendors) which combines GPS, IMU, and laser pulse range data to produce X, Y, and Z point data. The intention of the data format is to provide an open format that allows different LIDAR hardware and software tools to output data in a common format.

This document reflects the fourth revision of the LAS format specification since its initial version 1.0 release.

THE ADDITIONS OF LAS 1.4 INCLUDE:

- Extension of offsets and field sizes to support full 64 bit
- Support for up to 15 returns per outgoing pulse
- Extension of the Point Class field to support 256 classes
- Extension of the Scan Angle field to 2 bytes to support finer angle resolution
- Addition of a Sensor Channel bit field to support mobile mapping systems
- Addition of Well Known Text (WKT) definitions for Coordinate Reference Systems
- Addition of an Overlap bit to allow indicating pulses in the overlap region while maintaining the class definition
- Other minor changes

2 Conformance

The data types used in the LAS format definition are conformant to the 1999 ANSI C Language Specification (ANSI/ISO/IEC 9899:1999 ("C99")).

3 Authority

The American Society for Photogrammetry & Remote Sensing (ASPRS) is the owner of the LAS Specification. The standard is maintained by committees within the organization as directed by the ASPRS Board of Directors. Questions related to this standard can be directed to ASPRS at 301-493-0290, by email at asprs@asprs.org, or by mail at 5410 Grosvenor Lane, Suite 210, Bethesda, Maryland 20814-2160.

4 Requirements

LAS FORMAT DEFINITION:

The format contains binary data consisting of a header block, Variable Length Records, and point data.

Table 4.1 – LAS Format Definition

PUBLIC HEADER BLOCK
VARIABLE LENGTH RECORDS
POINT DATA RECORDS

A LAS file that contains waveform data (point record types 4 or 5) would be

Table 4.2 – LAS Format Definition Containing Waveform Data

PUBLIC HEADER BLOCK
VARIABLE LENGTH RECORDS INCLUDING WAVEFORM PACKET DESCRIPTORS (up to 255)
POINT DATA RECORDS
EXTENDED VARIABLE LENGTH RECORD (WAVEFORM DATA PACKETS)

All data are in little-endian format. The header block consists of a public block followed by Variable Length Records. The public block contains generic data such as point numbers and coordinate bounds. The Variable Length Records contain variable types of data including projection information, metadata, waveform packet information and user application data. Waveform Data Packets, if included, comprise the only record that can follow the Point Data Records. It is placed in this position to allow easy "stripping" or externalizing. This record is an Extended Variable Length Record (EVLR). The length of an EVLR is stored in an unsigned long long (8 byte field).

DATA TYPES:

The following data types are used in the LAS format definition. Note that these data types are conformant to the 1999 ANSI C Language Specification (ANSI/ISO/IEC 9899:1999 ("C99")).

- char (1 byte)
- unsigned char (1 byte)
- short (2 bytes)
- unsigned short (2 bytes)
- long (4 bytes)
- unsigned long (4 bytes)
- long long (8 bytes)

- unsigned long long (8 bytes)
- double (8 byte IEEE floating point format)

PUBLIC HEADER BLOCK:

Table 4.3 – Public Header Block

Item	Format	Size	Required
File Signature (“LASF”)	char[4]	4 bytes	*
File Source ID	unsigned short	2 bytes	*
Global Encoding	unsigned short	2 bytes	*
Project ID - GUID data 1	unsigned long	4 bytes	
Project ID - GUID data 2	unsigned short	2 byte	
Project ID - GUID data 3	unsigned short	2 byte	
Project ID - GUID data 4	unsigned char[8]	8 bytes	
Version Major	unsigned char	1 byte	*
Version Minor	unsigned char	1 byte	*
System Identifier	char[32]	32 bytes	*
Generating Software	char[32]	32 bytes	*
File Creation Day of Year	unsigned short	2 bytes	*
File Creation Year	unsigned short	2 bytes	*
Header Size	unsigned short	2 bytes	*
Offset to point data	unsigned long long	8 bytes	*
Number of Variable Length Records	unsigned long	4 bytes	*
Point Data Format ID (0-99 for spec)	unsigned char	1 byte	*
Point Data Record Length	unsigned short	2 bytes	*
Number of point records	unsigned long long	8 bytes	*
Number of points by return	unsigned long long[15]	120 bytes	*
X scale factor	double	8 bytes	*
Y scale factor	double	8 bytes	*
Z scale factor	double	8 bytes	*
X offset	double	8 bytes	*
Y offset	double	8 bytes	*
Z offset	double	8 bytes	*
Max X	double	8 bytes	*
Min X	double	8 bytes	*
Max Y	double	8 bytes	*
Min Y	double	8 bytes	*
Max Z	double	8 bytes	*
Min Z	double	8 bytes	*
Start of Waveform Data Packet Record	unsigned long long	8 bytes	*

Any field in the Public Header Block that is not required and is not used must be zero filled.

File Signature: The file signature must contain the four characters “LASF”, and it is required by the LAS specification. These four characters can be checked by user software as a quick look initial determination of file type.

File Source ID (Flight Line Number if this file was derived from an original flight line): This field should be set to a value between 1 and 65,535, inclusive. A value of zero (0) is interpreted to mean that an ID has not been assigned. In this case, processing software is free to assign any valid number. Note that this scheme allows a LIDAR project to contain up to 65,535 unique sources. A source can be considered an original flight line or it can be the result of merge and/or extract operations.

Global Encoding: This is a bit field used to indicate certain global properties about the file. In LAS 1.2 (the version in which this field was introduced), only the low bit is defined (this is the bit, that if set, would have the unsigned integer yield a value of 1). This bit field is defined as:

Table 4.4 - Global Encoding - Bit Field Encoding

Bits	Field Name	Description
0	GPS Time Type	The meaning of GPS Time in the Point Records 0 (not set) -> GPS time in the point record fields is GPS Week Time (the same as versions 1.0 through 1.2 of LAS) 1 (set) -> GPS Time is standard GPS Time (satellite GPS Time) minus 1×10^9 (Adjusted Standard GPS Time). The offset moves the time back to near zero to improve floating point resolution.
1	Waveform Data Packets Internal	If this bit is set, the waveform data packets are located within this file (note that this bit is mutually exclusive with bit 2)
2	Waveform Data Packets External	If this bit is set, the waveform data packets are located external to this file in an auxiliary file with the same base name as this file and the extension ".wdp". (note that this bit is mutually exclusive with bit 1)
3	Return numbers have been synthetically generated	If set, the point return numbers in the Point Data Records have been synthetically generated. This could be the case, for example, when a composite file is created by combining a First Return File and a Last Return File. In this case, first return data will be labeled "1 of 2" and second return data will be labeled "2 of 2"
4:15	Reserved	Must be set to zero

Project ID (GUID data): The four fields that comprise a complete Globally Unique Identifier (GUID) are now reserved for use as a Project Identifier (Project ID). The field remains optional. The time of assignment of the Project ID is at the discretion of processing software. The Project ID should be the same for all files that are associated with a unique project. By assigning a Project ID and using a File Source ID (defined above) every file within a project and every point within a file can be uniquely identified, globally.

Version Number: The version number consists of a major and minor field. The major and minor fields combine to form the number that indicates the format number of the current specification itself. For example, specification number 1.4 would contain 1 in the major field and 4 in the minor field.

System Identifier: The version 1.0 specification assumed that LAS files are exclusively generated as a result of collection by a hardware sensor. Subsequent versions recognize that files often result from extraction, merging or modifying existing data files. Thus System ID becomes:

Table 4.5 – System Identifier

Generating Agent	System ID
Hardware system	String identifying hardware (e.g. "ALTM 1210" or "ALS50")
Merge of one or more files	"MERGE"
Modification of a single file	"MODIFICATION"

Generating Agent	System ID
Extraction from one or more files	"EXTRACTION"
Reprojection, rescaling, warping, etc.	"TRANSFORMATION"
Some other operation	"OTHER" or a string up to 32 characters identifying the operation

Generating Software: This information is ASCII data describing the generating software itself. This field provides a mechanism for specifying which generating software package and version was used during LAS file creation (e.g. "TerraScan V-10.8", "REALM V-4.2" and etc.). If the character data is less than 16 characters, the remaining data must be null.

File Creation Day of Year: Day, expressed as an unsigned short, on which this file was created. Day is computed as the Greenwich Mean Time (GMT) day. January 1 is considered day 1.

File Creation Year: The year, expressed as a four digit number, in which the file was created.

Header Size: The size, in bytes, of the Public Header Block itself. In the event that the header is extended by a new revision of the LAS specification through the addition of data at the end of the header, the Header Size field will be updated with the new header size. The Public Header Block may not be extended by users.

Offset to point data: The actual number of bytes from the beginning of the file to the first field of the first point record data field. This data offset must be updated if any software adds/removes data to/from the Variable Length Records.

Number of Variable Length Records preceding the Point Data Records: This field contains the current number of Variable Length Records that occur in the file preceding the Point Data Records. This number must be updated if the number of Variable Length Records changes at any time.

Point Data Format ID: The point data format ID corresponds to the point data record format type. LAS 1.4 defines types 6 through 10. Previous LAS specification versions defined types 0 through 5, and types 0 through 5 are to be used in LAS files conforming to the LAS 1.3 spec or older.

Point Data Record Length: The size, in bytes, of the Point Data Record. All Point Data Records within a single LAS file must be the same type and hence the same length.

Number of point records: This field contains the total number of point records within the file.

Number of points by return: This field contains an array of the total point records per return. The first unsigned long value will be the total number of records from the first return, the second contains the total number for return two, and so forth up to fifteen returns.

X, Y, and Z scale factors: The scale factor fields contain a double floating point value that is used to scale the corresponding X, Y, and Z long values within the point records. The corresponding X, Y, and Z scale factor must be multiplied by the X, Y, or Z point record value to get the actual X, Y, or Z coordinate. For example, if the X, Y, and Z coordinates are intended to have two decimal point values, then each scale factor will contain the number 0.01.

X, Y, and Z offset: The offset fields should be used to set the overall offset for the point records. In general these numbers will be zero, but for certain cases the resolution of the point data may not be large enough for a given projection system. However, it should always be assumed that these numbers are used. So to scale a given X from the point record, take the point record X multiplied by the X scale factor, and then add the X offset.

$$X_{\text{coordinate}} = (X_{\text{record}} * X_{\text{scale}}) + X_{\text{offset}}$$

$$Y_{\text{coordinate}} = (Y_{\text{record}} * Y_{\text{scale}}) + Y_{\text{offset}}$$

$$Z_{\text{coordinate}} = (Z_{\text{record}} * Z_{\text{scale}}) + Z_{\text{offset}}$$

Max and Min X, Y, Z: The max and min data fields are the actual unscaled extents of the LAS point file data, specified in the coordinate system of the LAS data.

Start of Waveform Data Packet Record: This value provides the offset, in bytes, from the beginning of the LAS file to the first byte of the Waveform Data Package Record. Note that this will be the first byte of the Waveform Data Packet header.

The projection information for the point data is required for all data. The projection information will be placed in the Variable Length Records. Placing the projection information within the Variable Length Records allows for any projection to be defined including custom projections. The Well Known Text (WKT) specification is used for defining the horizontal and vertical reference/projection system information for LAS 1.4 and beyond (www.opengeospatial.org/standards/ct). The GeoTIFF tags will continue to be used as an optional, secondary source of reference/coordinate system information (<http://www.remotesensing.org/geotiff/geotiff.html>).

VARIABLE LENGTH RECORDS:

The Public Header Block is followed by one or more Variable Length Records. (There is one mandatory Variable Length Record, **GeoKeyDirectoryTag**.) The number of Variable Length Records is specified in the “Number of Variable Length Records” field in the Public Header Block. The Variable Length Records must be accessed sequentially since the size of each variable length record is contained in the Variable Length Record Header. Each Variable Length Record Header is 60 bytes in length.

Table 4.6 – Variable Length Record Header

Item	Format	Size	Required
Reserved	unsigned short	2 bytes	
User ID	char[16]	16 bytes	*
Record ID	unsigned short	2 bytes	*
Record Length After Header	unsigned long long	8 bytes	*
Description	char[32]	32 bytes	

User ID: The User ID field is ASCII character data that identifies the user which created the variable length record. It is possible to have many Variable Length Records from different sources with different User IDs. If the character data is less than 16 characters, the remaining data must be null. The User ID must be registered with the LAS specification managing body. The management of these User IDs ensures that no two individuals accidentally use the same User ID. The specification will initially use two IDs: one for globally specified records (LASF_Spec), and another for projection types (LASF_Projection). Keys may be requested at <http://www.asprs.org/lasform/keyform.html>.

Record ID: The Record ID is dependent upon the User ID. There can be 0 to 65535 Record IDs for every User ID. The LAS specification manages its own Record IDs (User IDs owned by the specification), otherwise Record IDs will be managed by the owner of the given User ID. Thus each User ID is allowed to assign 0 to 65535 Record IDs in any manner they desire. Publicizing the meaning of a given Record ID is left to the owner of the given User ID. Unknown User ID/Record ID combinations should be ignored.

Record Length after Header: The record length is the number of bytes for the record after the end of the standard part of the header. Thus the entire record length is 60 bytes (the header size in this version) plus the number of bytes in the variable length portion of the record.

Description: Optional, null terminated text description of the data. Any remaining characters not used must be null.

Note that the record with User ID = LASF_Spec and Record ID = 65535 is the Waveform Packet Data Extended Variable Length Record (EVLRL). Unlike all other Variable Length Records, this VLR (if present) is the only VLR that is placed after the Point Data Records. Thus, if present, it will be the last data record in the LAS file.

POINT DATA RECORD

NOTE: Point Data Start Signature was removed in LAS Version 1.1. LAS file I/O software must use the **Offset to Point Data** field in the Public Header Block to locate the starting position of the first Point Data Record. Note that all Point Data Records must be the same type. Point data items that are not 'Required' must be set to the equivalent of zero for the data type (e.g. 0.0 for floating types, null for ACSII, 0 for integers).

NOTE: Point Data Records zero (0) through five (5) have been deprecated in LAS version 1.4 and are no longer supported.

POINT DATA RECORD FORMAT 6:

Table 4.7 – Point Data Record Format 6

Item	Format	Size	Required
X	long	4 bytes	*
Y	long	4 bytes	*
Z	long	4 bytes	*
Intensity	unsigned short	2 bytes	
Return Number	4 bits (bits 0 – 3)	4 bits	*
Number of Returns (given pulse)	4 bits (bits 4 – 7)	4 bits	*
Classification Flags	4 bits (bits 0 – 3)	4 bits	
Scanner Channel	2 bits (bits 4 – 5)	2 bits	*
Scan Direction Flag	1 bit (bit 6)	1 bit	*
Edge of Flight Line	1 bit (bit 7)	1 bit	*
Classification	unsigned char	1 byte	*
Scan Angle	signed short	2 byte	*
User Data	unsigned char	1 byte	
Point Source ID	unsigned short	2 bytes	*
GPS Time	double	8 bytes	*

X, Y, and Z: The X, Y, and Z values are stored as long integers. The X, Y, and Z values are used in conjunction with the scale values and the offset values to determine the coordinate for each point as described in the Public Header Block section.

Intensity: The intensity value is the integer representation of the pulse return magnitude. This value is optional and system specific. However, it should always be included if available. Intensity, when included, is always normalized to a 16 bit, unsigned value by multiplying the value by 65,536/(intensity dynamic range of the sensor). For example, if the dynamic range of the sensor is 10 bits, the scaling value would be (65,536/1,024). If intensity is not included, this value must be set to zero. This normalization is required to ensure that data from different sensors can be correctly merged.

NOTE: The six fields beginning with Return Number (Return Number, Number of Returns, Classification Flags, Scanner Channel, Scan Direction Flag and Edge of Flight Line) are bit fields encoded into two bytes.

Return Number: The Return Number is the pulse return number for a given output pulse. A given output laser pulse can have many returns, and they must be marked in sequence of return. The first return will have a Return Number of one, the second a Return Number of two, and so on up to fifteen returns. The Return Number must be between 1 and the Number of Returns, inclusive.

Number of Returns (for this emitted pulse): The Number of Returns is the total number of returns for a given pulse. For example, a laser data point may be return two (Return Number) within a total number of up to fifteen returns.

Classification Flags: Classification flags are used to indicate special characteristics associated with the point. The bit definitions are:

Table 4.8 - Classification Bit Field Encoding

Bit	Field Name	Description
0	Synthetic	If set then this point was created by a technique other than LIDAR collection such as digitized from a photogrammetric stereo model or by traversing a waveform.
1	Key-point	If set, this point is considered to be a model key-point and thus generally should not be withheld in a thinning algorithm.
2	Withheld	If set, this point should not be included in normal processing. Special routines may explicitly call for the inclusion of this point (i.e., noise analysis, perimeter generation, etc.). Points tagged as "Withheld" are to be retained in the LAS file and included in exports of the LAS file. Withheld is synonymous with "Ignored".
3	Overlap	If set, this point is within the overlap region of two or more swaths or takes. If used, all points in the overlap region must have this bit set. Note that "primary" data in an overlap region may then be identified by using a combination of Point Source ID and Overlap bit. Setting this bit is not mandatory by the LAS specification but may be required by particular delivery specifications.

Note that these bits are treated as flags and can be set or cleared in any combination. For example, a point with bits 0 and 1 both set to one and the Classification field set to 2 (see table below) would be a *ground* point that had been *synthetically* collected and marked as a *model key-point*.

Scanner Channel: Scanner Channel is used to indicate the channel (scanner head) of a multi-channel system. Channel 0 is used for single scanner systems. Up to four channels are supported (0-3).

Scan Direction Flag: The Scan Direction Flag denotes the direction at which the scanner mirror was traveling at the time of the output pulse. A bit value of 1 is a positive scan direction, and a bit value of 0 is a negative scan direction (where positive scan direction is a scan moving from the left side of the in-track direction to the right side and negative the opposite).

Edge of Flight Line: The Edge of Flight Line data bit has a value of 1 only when the point is at the end of a scan. It is the last point on a given scan line before it changes direction. Note that this

field has no meaning for 360° Field of View scanners (such as Mobile LIDAR scanners) and should not be set.

[A note on Bit Fields – The LAS storage format is “Little Endian.” This means that multi-byte data fields are stored in memory from least significant byte at the low address to most significant byte at the high address. Bit fields are always interpreted as bit 0 set to 1 equals 1, bit 1 set to 1 equals 2, bit 2 set to 1 equals 4 and so forth.]

Classification: This byte represents the “class” attributes of a point. If a point has never been classified, this byte must be set to zero. Note that the first 64 values (zero through 63) are reserved for ASPRS definition. The remaining values can be freely defined by the user.. Classification must always be used to denote the “class” of the object surface from which the specific return reflected. Class must never be used for denoting other characteristics of a return such as nadir, edge of flight line and so forth.

Classification must adhere to the following standard:

Table 4.9 - ASPRS Standard LIDAR Point Classes

Classification Value	Meaning
0	Created, never classified
1	Unclassified ¹
2	Ground
3	Low Vegetation
4	Medium Vegetation
5	High Vegetation
6	Building
7	Low Point (“low noise”)
8	High Point (typically “high noise”). Note that this value was previously used for Model Key Points. Bit 1 of the Classification Flag must now be used to indicate Model Key Points. This allows the model key point class to be preserved.
9	Water
10	Rail
11	Road Surface
12	Bridge Deck
13	Wire - Guard
14	Wire – Conductor (Phase)
15	Transmission Tower
16	Wire-structure Connector (e.g. Insulator)
17	Reserved
18-63	Reserved
64-255	User definable – The specific use of these classes should be encoded in the Classification lookup VLR.

¹ We are using both 0 and 1 as **Unclassified** to maintain compatibility with current popular classification software such as TerraScan. We extend the idea of classification value 1 to include cases in which data have been subjected to a classification algorithm but emerged in an undefined state. For example, data with class 0 is sent through an algorithm to detect man-made structures – points that emerge without having been assigned as belonging to structures could be remapped from class 0 to class 1.

Scan Angle: The Scan Angle is a signed short that represents the rotational position of the emitted laser pulse with respect to the vertical of the coordinate system of the data. Down in the data coordinate system is the 0.0 position. Each increment represents 0.006 degrees. Counter-Clockwise rotation, as viewed from the rear of the sensor, facing in the along-track (positive trajectory) direction, is positive. The maximum value in the positive sense is 30,000 (180 degrees which is up in the coordinate system of the data). The maximum value in the negative direction is -30,000 which is also directly up.

User Data: This field may be used at the user's discretion. Users are encouraged to document the use of this field via a user defined VLR. If this field is only used for internal processing (e.g. as a temporary storage area during computations), it should be reset to zero in deliverable files.

Point Source ID: This value indicates the file from which this point originated. Valid values for this field are 1 to 65,535 inclusive with zero being used for a special case discussed below. The numerical value corresponds to the File Source ID from which this point originated. Zero is reserved as a convenience to system implementers. A Point Source ID of zero implies that this point originated in this file. This implies that processing software should set the Point Source ID equal to the File Source ID of the file containing this point at some time during processing.

GPS Time: The GPS Time is the double floating point time tag value at which the point was acquired. It is GPS Week Time if the Global Encoding low bit is clear and Adjusted Standard GPS Time if the Global Encoding low bit is set (see *Global Encoding in the Public Header Block description*).

POINT DATA RECORD FORMAT 7:

Point Data Record Format 7 is the same as Point Data Record Format 6 with the addition of three color channels. These fields are used when "colorizing" a LIDAR point using ancillary data, typically from a camera.

Table 4.10 – Point Data Record Format 7

Item	Format	Size	Required
X	long	4 bytes	*
Y	long	4 bytes	*
Z	long	4 bytes	*
Intensity	unsigned short	2 bytes	
Return Number	4 bits (bits 0 – 3)	4 bits	*
Number of Returns (given pulse)	4 bits (bits 4 – 7)	4 bits	*
Classification Flags	4 bits (bits 0 – 3)	4 bits	
Scanner Channel	2 bits (bits 4 – 5)	2 bits	*
Scan Direction Flag	1 bit (bit 6)	1 bit	*
Edge of Flight Line	1 bit (bit 7)	1 bit	*
Classification	unsigned char	1 byte	*
Scan Angle	signed short	2 bytes	*
User Data	unsigned char	1 byte	
Point Source ID	unsigned short	2 bytes	*
GPS Time	double	8 bytes	*
Red	unsigned short	2 bytes	*
Green	unsigned short	2 bytes	*
Blue	unsigned short	2 bytes	*

Red: The Red image channel value associated with this point.

Green: The Green image channel value associated with this point.

Blue: The Blue image channel value associated with this point.

NOTE: Red, Green, Blue values should always be normalized to 16 bit values. For example, when encoding an 8 bit per channel pixel, multiply each channel value by 256 prior to storage in these fields. This normalization allows color values from different camera bit depths to be accurately merged.

POINT DATA RECORD FORMAT 8:

Point Data Record Format 8 is the same as Point Data Record Format 7 with the addition of a NIR (near infrared) channel.

Table 4.11 – Point Data Record Format 8

Item	Format	Size	Required
X	long	4 bytes	*
Y	long	4 bytes	*
Z	long	4 bytes	*
Intensity	unsigned short	2 bytes	
Return Number	4 bits (bits 0 – 3)	4 bits	*
Number of Returns (given pulse)	4 bits (bits 4 – 7)	4 bits	*
Classification Flags	4 bits (bits 0 – 3)	4 bits	
Scanner Channel	2 bits (bits 4 – 5)	2 bits	*
Scan Direction Flag	1 bit (bit 6)	1 bit	*
Edge of Flight Line	1 bit (bit 7)	1 bit	*
Classification	unsigned char	1 byte	*
Scan Angle	signed short	2 bytes	*
User Data	unsigned char	1 byte	
Point Source ID	unsigned short	2 bytes	*
GPS Time	double	8 bytes	*
Red	unsigned short	2 bytes	*
Green	unsigned short	2 bytes	*
Blue	unsigned short	2 bytes	*
NIR	unsigned short	2 bytes	*

NIR: The NIR (near infrared) channel value associated with this point.

NOTE: Red, Green, Blue and NIR values should always be normalized to 16 bit values. For example, when encoding an 8 bit per channel pixel, multiply each channel value by 256 prior to storage in these fields. This normalization allows color values from different camera bit depths to be accurately merged.

POINT DATA RECORD FORMAT 9:

Point Data Record Format 9 adds Wave Packets to Point Data Record Format 6.

Table 4.12 – Point Data Record Format 9

Item	Format	Size	Required
X	long	4 bytes	*
Y	long	4 bytes	*
Z	long	4 bytes	*
Intensity	unsigned short	2 bytes	
Return Number	4 bits (bits 0 – 3)	4 bits	*

Item	Format	Size	Required
Number of Returns (given pulse)	4 bits (bits 4 – 7)	4 bits	*
Classification Flags	4 bits (bits 0 – 3)	4 bits	
Scanner Channel	2 bits (bits 4 – 5)	2 bits	*
Scan Direction Flag	1 bit (bit 6)	1 bit	*
Edge of Flight Line	1 bit (bit 7)	1 bit	*
Classification	unsigned char	1 byte	*
Scan Angle Rank (-90 to +90) – Left side	char	1 byte	*
User Data	unsigned char	1 byte	
Point Source ID	unsigned short	2 bytes	*
GPS Time	double	8 bytes	*
Wave Packet Descriptor Index	unsigned char	1 byte	*
Byte offset to waveform data	unsigned long long	8 bytes	*
Waveform packet size in bytes	unsigned long	4 bytes	*
Return Point Waveform Location	float	4 bytes	*
X(t)	float	4 bytes	*
Y(t)	float	4 bytes	*
Z(t)	float	4 bytes	*

Point Data Record Format 9 is the same as Point Data Record Format 6 with the addition of the waveform packet information.

Wave Packet Descriptor Index: LAS supports up to 255 User Defined Records which describe the waveform packet. This value plus 100 is the Record ID of the Waveform Packet Descriptor User Defined Record that is used to describe the waveform packet data associated with this LIDAR point. Note: A value of zero indicates that there is no waveform data associated with this LIDAR point record.

Byte offset to Waveform Packet Data: The waveform packet data are stored in the LAS file in an Extended Variable Length Record (or, optionally, in an auxiliary file). The Byte Offset represents the location of the start of this LIDAR points' waveform packet within the waveform data variable length record (or external file) relative to the beginning of the Waveform Packet Data header. Note that the absolute location of the beginning of this waveform packet relative to the beginning of the file is given by:

$$\text{Start of Waveform Data Packet Record} + \text{Byte offset to Waveform Packet Data}$$

for waveform packets stored within the LAS file and

$$\text{Byte offset to Waveform Packet Data}$$

for data stored in an auxiliary file

Waveform packet size in bytes: The size, in bytes, of the waveform packet associated with this return. Note that each waveform can be of a different size (even those with the same Waveform Packet Descriptor index) due to packet compression. Also note that waveform packets can be located only via the Byte offset to Waveform Packet Data value since there is no requirement that records be stored sequentially.

Return Point location: The offset in picoseconds (10^{-12}) from the first digitized value to the location within the waveform packet that the associated return pulse was detected.

X(t), Y(t), Z(t): These parameters define a parametric line equation for extrapolating points along the associated waveform. The position along the wave is given by:

$$X = X_0 + X(t)$$

$$Y = Y_0 + Y(t)$$

$$Z = Z_0 + Z(t)$$

where X, Y and Z are the spatial position of the derived point, X_0 , Y_0 , Z_0 are the position of the “anchor” point (the X, Y, Z locations from this point’s data record) and t is the time, in picoseconds, relative to the anchor point (i.e. t = zero at the anchor point). The units of X, Y and Z are the units of the coordinate systems of the LAS data. If the coordinate system is geographic, the horizontal units are decimal degrees and the vertical units are meters.

POINT DATA RECORD FORMAT 10:

Point Data Record Format 10 adds Wave Packets to Point Data Record Format 7.

Table 4.13 – Point Data Record Format 10

Item	Format	Size	Required
X	long	4 bytes	*
Y	long	4 bytes	*
Z	long	4 bytes	*
Intensity	unsigned short	2 bytes	
Return Number	4 bits (bits 0 – 3)	4 bits	*
Number of Returns (given pulse)	4 bits (bits 4 – 7)	4 bits	*
Classification Flags	4 bits (bits 0 – 3)	4 bits	
Scanner Channel	2 bits (bits 4 – 5)	2 bits	*
Scan Direction Flag	1 bit (bit 6)	1 bit	*
Edge of Flight Line	1 bit (bit 7)	1 bit	*
Classification	unsigned char	1 byte	*
Scan Angle	signed short	2 bytes	*
User Data	unsigned char	1 byte	
Point Source ID	unsigned short	2 bytes	*
GPS Time	double	8 bytes	*
Red	unsigned short	2 bytes	*
Green	unsigned short	2 bytes	*
Blue	unsigned short	2 bytes	*
NIR	unsigned short	2 bytes	*
Wave Packet Descriptor Index	unsigned char	1 byte	*
Byte offset to waveform data	unsigned long long	8 bytes	*
Waveform packet size in bytes	unsigned long	4 bytes	*
Return Point Waveform Location	float	4 bytes	*
X(t)	float	4 bytes	*
Y(t)	float	4 bytes	*
Z(t)	float	4 bytes	*

Point Data Record Format 10 is the same as Point Data Record Format 9 with the addition of RGB and NIR values.

DEFINED VARIABLE LENGTH RECORDS:

Georeferencing Information

As of LAS 1.4, the preferred method for specifying georeferencing information for point data in LAS is the “Well-Known Text Representation of Spatial Reference Systems” (WKT). For

definition of WKT we refer to Open Geospatial Consortium (OGC) specification “OpenGIS Coordinate Transformation Service Implementation Specification” revision 1.00 released 12 January 2001, section 7 (Coordinate Transformation Services Spec).

As there are a few dialects of WKT, please note LAS is *not* using the “ESRI WKT” dialect, which does not include TOWGS84 and authority nodes.

WKT georeferencing information can be specified in two optional Variable Length Records, the OGC Math Transform WKT Record and the OGC Coordinate System WKT Record, as follows. Note that the Math Transform WKT record is added for completeness, and a Coordinate System WKT *may or may not* require a Math Transform WKT Record (a parameterized math transform definition).

OGC Math Transform WKT Record:

User ID: LASF_Projection
Record ID: 2111

This record contains the textual data representing a Math Transform WKT as defined in section 7 of the Coordinate Transformation Services Spec, with the following notes:

- The OGC Math Transform WKT VLR data shall be a null-terminated string.
- The OGC Math Transform WKT VLR data shall be considered UTF-8.
- The OGC Math Transform WKT VLR data shall be considered C locale-based, and no localization of the numeric strings within the WKT should be performed.

OGC Coordinate System WKT:

User ID: LASF_Projection
Record ID: 2112

This record contains the textual data representing a Coordinate System WKT as defined in section 7 of the Coordinate Transformation Services Spec, with the following notes:

- The OGC Math Transform WKT VLR data shall be a null-terminated string.
- The OGC Math Transform WKT VLR data shall be considered UTF-8.
- The OGC Math Transform WKT VLR data shall be considered C locale-based, and no localization of the numeric strings within the WKT should be performed.

As a means of transition, the GeoTIFF records as specified in LAS 1.3 may continue to be used in LAS 1.4. Software reading LAS 1.4 files must be able to consume the GeoTIFF VLR records with user ID “LASF_Projection” and Record IDs 34735 through 34737. It is acceptable to write only GeoTIFF VLR records into an LAS file, while software systems that choose to write OGC WKT VLRs must also write GeoTIFF VLRs for backward compatibility. If both WKT and GeoTIFF VLRs are present, then the WKT information takes precedence

The particulars of these GeoTIFF VLRs are repeated here for specification completeness (same as defined in LAS 1.3).

Georeferencing for the LAS format may use the same robust mechanism that was developed for the GeoTIFF standard. The variable length header records section may contain the same data that would be contained in the GeoTIFF key tags of a TIFF file. Since LAS is not a raster format and each point contains its own absolute location information, only 3 of the 6 GeoTIFF tags are necessary when using GeoTIFF records instead of WKT records. The ModelTiePointTag (33922), ModelPixelScaleTag (33550), and ModelTransformationTag (34264) records can be excluded. The GeoKeyDirectoryTag (34735), GeoDoubleParamsTag (34736), and GeoASCIIParamsTag (34737) records are used.

Only the GeoKeyDirectoryTag record is required when using GEOTIFF records instead of WKT records. The GeoDoubleParamsTag and GeoASCIIParamsTag records may or may not be present, depending on the content of the GeoKeyDirectoryTag record.

GeoKeyDirectoryTag Record:

User ID: LASF_Projection
Record ID: 34735

This record contains the key values that define the coordinate system. A complete description can be found in the GeoTIFF format specification. Here is a summary from a programmatic point of view for someone interested in implementation.

The GeoKeyDirectoryTag is defined as just an array of unsigned short values. But, programmatically, the data can be seen as something like this:

```
struct sGeoKeys
{
    unsigned short wKeyDirectoryVersion;
    unsigned short wKeyRevision;
    unsigned short wMinorRevision;
    unsigned short wNumberOfKeys;
    struct sKeyEntry
    {
        unsigned short wKeyID;
        unsigned short wTIFFTagLocation;
        unsigned short wCount;
        unsigned short wValue_Offset;
    } pKey[1];
};
```

Where:
wKeyDirectoryVersion = 1; // Always
wKeyRevision = 1; // Always
wMinorRevision = 0; // Always
wNumberOfKeys // Number of sets of 4 unsigned shorts to follow

Table 4.14 – GeoKey Four Unsigned Shorts

For each set of 4 unsigned shorts:

Name	Definition
wKeyID	Defined key ID for each piece of GeoTIFF data. IDs contained in the GeoTIFF specification.
wTIFFTagLocation	Indicates where the data for this key is located: 0 means data is in the wValue_Offset field as an unsigned short. 34736 means the data is located at index wValue_Offset of the GeoDoubleParamsTag record. 34737 means the data is located at index wValue_Offset of the GeoAsciiParamsTag record.
wCount	Number of characters in string for values of GeoAsciiParamsTag , otherwise is 1
wValue_Offset	Contents vary depending on value for wTIFFTagLocation above

GeoDoubleParamsTag Record: (optional)

User ID: LASF_Projection
Record ID: 34736

This record is simply an array of doubles that contain values referenced by tag sets in the GeoKeyDirectoryTag record.

GeoAsciiParamsTag Record: (Optional)

User ID: LASF_Projection
Record ID: 34737

This record is simply an array of ASCII data. It contains many strings separated by null terminator characters which are referenced by position from data in the GeoKeyDirectoryTag record.

Classification lookup: (optional)

User ID: LASF_Spec
Record ID: 0
Record Length after Header: 255 recs X 16 byte struct len
struct CLASSIFICATION

```
{
    unsigned char ClassNumber;
    char Description[15];
};
```

Header lookup for flight-lines:

(Removed with Version 1.1 - Point Source ID in combination with Source ID provides the new scheme for directly encoding flight line number. Thus variable Record ID 1 now becomes reserved for future use.)

User ID: LASF_Spec
Record ID: 1

Histogram: (optional)

User ID: LASF_Spec
Record ID: 2

Text area description: (optional)

User ID: LASF_Spec
Record ID: 3

Waveform Packet Descriptor: (required when using Point formats 9 or 10)

User ID: LASF_Spec
Record ID: n

The value of the Record ID, n, must be in the range where n>100 and n<356. Note that the value of the Wave Packet Descriptor Index (in Point Data Records of format 9 or 10) plus 100 equals the Record ID of the applicable Waveform Packet Descriptor

These records contain information that describes the configuration of the waveform packets. Since systems may be configured differently at different times throughout a job, the LAS file supports 255 Waveform Packet Descriptors.

Table 4.15 – Waveform Packet Descriptor User Defined Record

Item	Format	Size	Required
Bits per sample	unsigned char	1 byte	*
Waveform compression type	unsigned char	1 byte	*

Number of samples	unsigned long	4 bytes	*
Temporal Sample Spacing	unsigned long	4 bytes	*
Digitizer Gain	double	8 bits	*
Digitizer Offset	double	8 bits	*

Bits per sample: 2 through 32 bits are supported.

Waveform Compression type: It is expected that in the future standard compression types will be adopted by the LAS committee. This field will indicate the compression algorithm used for the waveform packets associated with this descriptor. A value of 0 indicates no compression. Zero is the only value currently supported.

Number of Samples: The number of samples associated with this waveform packet type. This value always represents the fully decompressed waveform packet.

Temporal Sample Spacing: The temporal sample spacing in picoseconds. Example values might be 500, 1000, 2000 and so on, representing digitizer frequencies of 2 GHz, 1 GHz and 500 MHz respectively.

Digitizer Gain: The gain and offset are used to convert the raw digitized value to an absolute digitizer voltage using the formula: $VOLTS = OFFSET + GAIN * Raw_Waveform_Amplitude$

Digitizer Offset: The gain and voltage offset are used to convert the raw digitized value to a voltage using the formula: $VOLTS = OFFSET + GAIN * Raw_Waveform_Amplitude$

EXTENDED VARIABLE LENGTH RECORD (EVLN)

Extended Variable Length Records occur *after* the Point Data Records.

Table 4.16 – Extended Variable Length Record Header

Item	Format	Size	Required
Reserved	unsigned short	2 bytes	
User ID	char[16]	16 bytes	*
Record ID	unsigned short	2 bytes	*
Record Length After Header	unsigned long long	8 bytes	*
Description	char[32]	32 bytes	

LAS allows only a single EVLN; Waveform Data Packets.

Waveform Data Packets: (required when using Point format 9 or 10)

User ID: LASF_Spec
Record ID: 65,535

The packet of Raw Waveform Amplitude values for all records immediately follow this variable length header.

This is the last Reserved Record for the LASF Specification. This extended variable length record must always be the last record in an LAS file. Unlike all other Variable Length Records, this record and its associated data *follow* the Point Data Records.

NOTE: When using a bit resolution that is not an even increment of 8, the last byte of each waveform packet must be padded such that the next waveform record will start on an even byte boundary.