



The XRootD Protocol Version 4.0.0



Andrew Hanushevsky
17-December-2018

©2004-2018 by the Board of Trustees of the Leland Stanford, Jr., University
All Rights Reserved

Produced under contract DE-AC02-76-SFO0515 with the Department of Energy
The protocol specification described in this document falls under BSD license terms.

The specification may be used for any purpose whatsoever.

Use of this specification must cite the original source -- xrootd.org.

Binary definition in header file XProtocol.hh superceed any such definitions in this document.

1 Contents

1	Contents	3
2	Request/Response Protocol	7
2.1	Format of Client-Server Initial Handshake	7
2.2	Data Serialization.....	9
2.3	Client Request Format.....	11
2.3.1	Valid Client Requests	13
2.3.2	Valid Client Paths	14
2.3.3	Client Recovery from Server Failures	15
2.4	Server Response Format	17
2.4.1	Valid Server Response Status Codes.....	18
2.4.2	Server kXR_attn Response Format	19
2.4.2.1	Server kXR_attn Response for kXR_asyncab Client Action.....	21
2.4.2.2	Server kXR_attn Response for kXR_asyncdi Client Action	22
2.4.2.3	Server kXR_attn Response for kXR_asyncgo Client Action.....	23
2.4.2.4	Server kXR_attn Response for kXR_asyncms Client Action.....	24
2.4.2.5	Server kXR_attn Response for kXR_asyncrd Client Action	25
2.4.2.6	Server kXR_attn Response for kXR_asyncresp Client Action	27
2.4.2.7	Server kXR_attn Response for kXR_asyncwt Client Action	29
2.4.3	Server kXR_authmore Response Format.....	30
2.4.4	Server kXR_error Response Format	31
2.4.4.1	Server kXR_error Sub-Codes & Recovery Actions.....	33
2.4.5	Server kXR_ok Response Format	35
2.4.6	Server kXR_oksofar Response Format.....	36
2.4.7	Server kXR_redirect Response Format	37
2.4.8	Server kXR_wait Response Format	40
2.4.9	Server kXR_waitresp Response Format.....	41
2.5	Binary Definitions of Status and Error codes.....	43
2.5.1	Status Codes	43
2.5.2	kXR_attn Status Subcodes	43
2.5.3	Error Codes.....	44
3	Transport Layer Security (TLS) Support	45
3.1	Client-Server interactions to unilaterally use TLS.....	46
3.2	Client-Server interactions to use TLS only when required	46
4	Detailed Protocol Specifications	47
4.1	kXR_admin Request	47
4.2	kXR_auth Request	49
4.3	kXR_bind Request	51
4.3.1	TLS Considerations	52
4.4	kXR_chmod Request	53
4.5	kXR_close Request.....	56
4.6	kXR_decrypt Request.....	57
4.7	kXR_dirlist Request.....	58
4.8	kXR_endsess Request.....	61
4.9	kXR_fattr Request.....	63
4.9.1	Layout of <i>namevec</i>	65
4.9.2	Layout of <i>valuvec</i>	66

4.9.3	kXR_fattr Request – Delete Subcode	67
4.9.4	kXR_fattr Request – Get Subcode	69
4.9.5	kXR_fattr Request – List Subcode.....	71
4.9.6	kXR_fattr Request – Set Subcode	73
4.10	kXR_getfile Request	75
4.11	kXR_locate Request	77
4.12	kXR_login Request	81
4.12.1	Additional Login CGI Tokens.....	84
4.13	kXR_mkdir Request.....	85
4.14	kXR_mv Request.....	87
4.15	kXR_open Request.....	89
4.15.1	Additional Open CGI Tokens	93
4.16	kXR_ping Request	95
4.17	kXR_prepare Request.....	96
4.18	kXR_protocol Request.....	99
4.18.1	Protocol Security Requirements vs Response Implications	107
4.19	kXR_putfile Request.....	109
4.20	kXR_query Request	111
4.20.1	KXR_query Checksum Cancellation Request	115
4.20.2	KXR_query Checksum Request	117
4.20.2.1	Additional Query Checksum CGI Tokens	118
4.20.3	KXR_query Configuration Request.....	119
4.20.3.1	Format for Query Config cms	121
4.20.3.2	Format for Query Config role.....	122
4.20.3.3	Format for Query Config xattrs	122
4.20.4	KXR_query Opaque Request	123
4.20.5	KXR_query Space Request	125
4.20.6	KXR_query Statistics Request	126
4.20.7	KXR_query Visa Request.....	129
4.20.8	KXR_query Xattr Request.....	131
4.21	kXR_read Request.....	133
4.22	kXR_readv Request	137
4.23	kXR_rm Request	141
4.24	kXR_rmdir Request	142
4.25	kXR_set Request.....	143
4.25.1	Valid kXR_Set Values.....	145
4.26	kXR_sigver Request.....	147
4.26.1	Signing a request.....	149
4.26.2	Verifying a signed request.....	150
4.27	kXR_stat Request	151
4.27.1	Additional Stat CGI Tokens	154
4.28	kXR_statx Request	155
4.29	kXR_sync Request	157
4.30	kXR_truncate Request.....	159
4.31	kXR_verifyw Request.....	161
4.32	kXR_write Request	163
4.33	kXR_writew Request.....	165
5	The Security Framework	167

5.1	Framework for Transport Layer Protocols.....	171
5.2	Request Verification	172
6	Document Change History.....	173

2 Request/Response Protocol

2.1 Format of Client-Server Initial Handshake

When a client first connects to the **XRootD** server, it must perform a special handshake. This handshake should determine whether the client is communicating using **XRootD** protocol or another protocol hosted by the server.

The handshake consists of the client sending 20 bytes, as follows:

kXR_int32	0	
kXR_int32	0	
kXR_int32	0	
kXR_int32	4	(network byte order)
kXR_int32	2012	(network byte order)

XRootD protocol, servers should respond, as follows:

<i>streamid:</i>	kXR_char	<i>smid</i> [2]
<i>status:</i>	kXR_unt16	0
<i>msglen:</i>	kXR_int32	<i>rlen</i>
<i>msgval1:</i>	kXR_int32	<i>pval</i>
<i>msgval2:</i>	kXR_int32	<i>flag</i>

Where:

smid is the initial *streamid*. The *smid* for the initial response is always two null characters (i.e., '\0');

rlen is the binary response length (e.g., 8 for the indicated response).

pval is the binary protocol version number.

flag is additional bit-encoded information about the server; as follows:

kXR_DataServer - 0x00 00 00 01 This is a data server.

kXR_LBalServer - 0x00 00 00 00 This is a load-balancing server.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The particular response format was developed for protocol version 2.0 and does not convey all of the information to capture features provided by subsequent protocol versions. In order to provide backward compatibility, this response format has been kept. The recommended mechanism to obtain all of the information that may be needed is to “piggy-back” a **kXR_protocol** Request with the handshake (i.e. send the handshake and the request with a single write).
- 3) All twenty bytes *must* be received by the server at one time. All known TCP implementations should guarantee that the first message is sent intact if all twenty bytes are sent in a single system call. Using multiple system calls for the first message may cause unpredictable results.

2.2 Data Serialization

All data sent and received is serialized (i.e., marshaled) in three ways:

1. Bytes are sent unaligned without any padding,
2. Data type characteristics are predefined (see table below), and
3. All integer quantities are sent in network byte order (i.e, big endian).

XRootD Type	Sign	Bit Length	Bit Alignment	Typical Host Type
kXR_char8	unsigned	8	8	unsigned char
kXR_unt16	unsigned	16	16	unsigned short
kXR_int32	signed	32	32	long ¹
kXR_int64	signed	64	64	long long

Table 1: XRootD Protocol Data Types

Network byte order is defined by the Unix **htons()** and **htonl()** macros for host to network short and host to network long, respectively. The reverse is defined by the **ntohs()** and **ntohl()** macros. Many systems do not define the long long versions of these macros. XRootD protocol requires that the **POSIX** version of long long serialization be used, as defined in the following figures. The OS-dependent **isLittleEndian()** function returns true if the underlying hardware using little endian integer representation.

```

unsigned long long htonll(unsigned long long x)
{
    unsigned long long ret_val;
    if (isLittleEndian())
    {
        *((unsigned long *)(&ret_val) + 1) =
            htonl(*((unsigned long *)(&x)));
        *((unsigned long *)(&ret_val)) =
            htonl(*((unsigned long *)(&x)+1) );
    } else {
        *((unsigned long *)(&ret_val)) =
            htonl(*((unsigned long *)(&x)));
        *((unsigned long *)(&ret_val) + 1) =
            htonl(*((unsigned long *)(&x)+1) );
    }
    return ret_val;
};

```

Figure 1: POSIX Host to Network Byte Order Serialization

¹ As of this writing, the long type has taken on several meanings for 64-bit architectures. Some machines define a long to be 64-bits and int 32-bits while some others reverse the definition.

```

unsigned long long ntohll(unsigned long long x)
{
    unsigned long long ret_val;
    if (isLittleEndian())
    {
        *((unsigned long *)(&ret_val) + 1) =
            ntohl(*(unsigned long *)(&x));
        *((unsigned long *)(&ret_val)) =
            ntohl(*(unsigned long *)(&x)+1);
    }
    else {
        *((unsigned long *)(&ret_val)) =
            ntohl(*(unsigned long *)(&x));
        *((unsigned long *)(&ret_val) + 1) =
            ntohl(*(unsigned long *)(&x)+1);
    }
    return ret_val;
};

```

Figure 2: Network and Host Byte Order Seialization

More compact and efficient, though OS restricted (i.e., Solaris and Linux), versions of 64-bit network byte ordering routines are given in the following figure.

```

#if defined(__sparc) || __BYTE_ORDER==__BIG_ENDIAN
#ifndef htonll
#define htonll(x) x
#endif
#ifndef ntohll
#define ntohll(x) x
#endif
#else
#ifndef htonll
#define htonll(x) __bswap_64(x)
#endif
#ifndef ntohll
#define ntohll(x) __bswap_64(x)
#endif
#endif

```

Figure 3: Network and Host Byte Ordering Macros

2.3 Client Request Format

Requests sent to the server are a mixture of ASCII and binary. All requests, other than the initial handshake request, have the same format, as follows:

kXR_char	<i>streamid</i> [2]
kXR_unt16	<i>requestid</i>
kXR_char	<i>parms</i> [16]
kXR_int32	<i>dlen</i>
kXR_char	<i>data</i> [<i>dlen</i>]

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

requestid

is the binary identifier of the operation to be performed by the server.

parms are parameters specific to the *requestid*.

dlen is the binary length of the *data* portion of the message. If no data is present, then the value is zero.

data are data specific to the *requestid*. Not all requests have associated data. If the request does have data, the length of this field is recorded in the *dlen* field.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) All **XRootD** client requests consist of a standard 24-byte fixed length message. The 24-byte header may then be optionally followed by request specific data.
- 3) Stream id's are arbitrary and are assigned by the client. Typically these id's correspond to logical connections multiplexed over a physical connection established to a particular server.

- 4) The client may send any number of requests to the same server. The order in which requests are performed is undefined. Therefore, each request should have a different *streamid* so that returned results may be paired up with associated requests.
- 5) Requests sent by a client over a single physical connection may be processed in an arbitrary order. Therefore the client is responsible for serializing requests, as needed.

2.3.1 Valid Client Requests

Requestid	Value	Login?	Auth?	Redirect?	Arguments
kXR_admin*	3020	yes	yes	no	<i>args</i>
kXR_auth	3000	y	n	n	<i>authtype, authinfo</i>
KXR_bind	3024	n	n	n	<i>sessid</i>
kXR_chmod	3002	y	y	yes	<i>mode, path</i>
kXR_close	3003	y	-	n	<i>fhandle</i>
kXR_decrypt*	3030	y	y	n	
KXR_dirlist	3004	y	y	y	<i>path</i>
KXR_endsess	3023	y	-	n	<i>sessid</i>
kXR_fattr	3032	y	y	y	<i>Arguments vary by subcode</i>
kXR_getfile*	3005	y	y	y	<i>path</i>
kXR_locate	3027	y	y	y	<i>path</i>
kXR_login	3007	n	n	n	<i>userid, token</i>
kXR_mkdir	3008	y	y	y	<i>mode, path</i>
kXR_mv	3009	y	y	y	<i>old_name, new_name</i>
kXR_open	3010	y	y	y	<i>mode, flags, path</i>
kXR_ping	3011	y	n	n	
kXR_prepare	3021	y	y	n	<i>paths</i>
kXR_protocol	3006	n	n	n	
kXR_putfile*	3012	y	y	y	<i>mode, flags, path</i>
kXR_query	3001	y	y	y	<i>args</i>
kXR_read	3013	y	-	y	<i>fhandle, pathid, length, offset</i>
kXR_readv	3025	y	-	y	<i>fhandle, pathid, length, offset</i>
kXR_rm	3014	y	y	y	<i>path</i>
kXR_rmdir	3014	y	y	y	<i>path</i>
kXR_set	3018	y	y	y	<i>info</i>
kXR_sigver	3029	y	y	n	<i>signature</i>
kXR_stat	3017	y	-	n	<i>fhandle</i>
kXR_stat	3017	y	y	y	<i>path</i>
kXR_statx	3022	y	y	n	<i>pathlist</i>
kXR_sync	3016	y	-	n	<i>fhandle</i>
kXR_truncate	3028	y	-	n	<i>fhandle, length</i>
kXR_truncate	3028	y	-	y	<i>path, length</i>
kXR_verifyw*		y	-	y	<i>fhandle, length, offset, data</i>
kXR_write	3019	y	-	y	<i>fhandle, pathid, length, offset, data</i>
kXR_writev	3031	y	y	n	<i>fhandle, length, offset</i>

Table 2: Valid Client Requests

*

*These requests are currently not supported in the reference implementation.

2.3.2 Valid Client Paths

The **XRootD** server accepts only *absolute paths* where a path may be specified. Relative paths must be resolved by the client interface prior to sending them to **XRootD**. This means that the interface must handle a virtual “current working directory” to resolve relative paths should they arise.

Path names are restricted to the following set of characters:

- Letters (upper or lower case),
- Digits (0-9), and
- Special characters: `!@#%^_+=:./`

In general, paths may not contain shell meta-characters.

Any path may be suffixed by CGI information. The format corresponds to that defined in RFC 3875. However, the protocol does not allow URL encoded characters (i.e. %xx). The meaning of any CGI element that is not specified in this document is implementation specific.

2.3.3 Client Recovery from Server Failures

A server failure should be recognized when the server unexpectedly closes its TCP/IP connection or does not respond for an extended period of time. Should this happen, the client may recover all operations by treating the termination of the connection or unresponsiveness as a redirection request (see page 37) to the initial **XRootD** server for all streams associated with the closed TCP/IP connections.

The initial **XRootD** server is defined as the *first* manager or the *last* meta-manager encountered. In the absence of any manager, the first data server encountered. See the **kXR_protocol** request on how to determine a node's type.

Because many clients are likely to be affected by a server failure, it is important that clients pace their reconnection to the initial **XRootD** server. One effective way to do this is to use the last three bits of the client's IP address as the number of seconds to wait before attempting a reconnection. It is up to the client to determine either the number of times or the time window in which reconnections should be attempted before failure is declared. Typical values are 16 attempts or 3 minutes, whichever is longer.

Note that it may not be possible to recover in this way for files that were opened in update mode. Clients who do not provide proper transactional support generally cannot recover via redirection for any read/write resources.

2.4 Server Response Format

All responses, including the initial handshake response, have the same format, as follows:

kXR_char	<i>streamid</i> [2]
kXR_unt16	<i>status</i>
kXR_int32	<i>dlen</i>
kXR_char	<i>data</i> [<i>dlen</i>]

Where:

streamid

is the binary identifier that is associated with this request stream corresponding to a previous request.

status is the binary status code indicating how the request completed. The next section describes possible status codes.

dlen is the binary length of the *data* portion of the message. If no data is present, then the value is zero.

data are data specific to the *requestid*. Not all responses have associated data. If the response does have data, the length of this field is recorded in the *dlen* field.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Since requests may be completed in any order, the ordering of responses is undefined. The client must appropriately pair responses with requests using the *streamid* value.
- 3) Unsolicited responses are server requests for client configuration changes to make better use of the overall system. Since these responses do not correspond to any request, the *streamid* value has no meaning.
- 4) Unsolicited responses must be immediately acted upon. They should not be paired with any previous request.

2.4.1 Valid Server Response Status Codes

The following table lists all possible responses and their arguments.

Status	Response Data
kXR_attn	<i>Parameters to direct immediate client action</i>
kXR_authmore	<i>Authentication specific data</i>
kXR_error	<i>Error number and corresponding ASCII message text</i>
kXR_ok	<i>Depends on request (this is predefined to be the value 0)</i>
kXR_oksofar	<i>Depends on request</i>
kXR_redirect	<i>Target port number and ASCII host name</i>
kXR_wait	<i>Binary number of seconds and optional ASCII message</i>
kXR_waitresp	<i>Binary number of seconds</i>

Notes

- 1) Any request may receive any of the previous status codes.
- 2) The following sections detail the response format used for each status code.

2.4.2 Server kXR_attn Response Format

kXR_char	<i>pad</i> [2]
kXR_unt16	kXR_attn
kXR_int32	<i>plen</i>
kXR_int32	<i>actnum</i>
kXR_char	<i>parms</i> [<i>plen</i> -4]

Where:

plen is two bytes of padding required by the standard response format. These two bytes can be ignored for this particular response code.

plen is the binary length of the *parms* portion of the message (i.e., the subsequent bytes).

actnum

is the binary action code describing the action that the client is to take.

These are:

- kXR_asyncav** - The file or file(s) the client previously requested to be prepared are now available.
- kXR_asyncab** - The client should immediately disconnect (i.e., close the socket connection) from the server and abort further execution.
- kXR_asyncdi** - The client should immediately disconnect (i.e., close the socket connection) from the server. Parameters indicate when a reconnect may be attempted.
- kXR_asyncgo** - The client may start sending requests. This code is sent to cancel the effects of a previous **kXR_asyncwt** code.
- kXR_asynccms** - The client should send the indicated message to the console. The parameters contain the message text.
- kXR_asyncrd** - The client should immediately disconnect (i.e., close the socket connection) and reconnect to the indicated server.
- kXR_asyncresp** - The client should use the response data in the message to complete the request associated with the indicated streamid.
- kXR_asyncunav** - The file or file(s) the client previously requested to be prepared cannot be made available.

kXR_asyncwt - The client should hold off sending any new requests until the indicated amount of time has passed or until receiving a **kXR_asyncgo** action code.

parms is the parameter data, if any, that is to steer client action.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Servers use the **kXR_attn** response code to optimize overall system performance and to notify clients of any impending events. All responses *except* for **kXR_asynresp**, do not correspond to any client request and should not be paired up with any request.
- 3) When **kXR_attn** is received, the client must perform the requested action and indicated by the *actnum* value.

2.4.2.1 Server **kXR_attn** Response for **kXR_asyncab** Client Action

kXR_char	<i>pad</i> [2]
kXR_unt16	kXR_attn
kXR_int32	<i>mlen</i>
kXR_int32	kXR_asyncab
kXR_char	<i>msg</i> [<i>mlen</i> -4]

Where:

mlen is the binary length of the following action code and message.

msg is the message to be sent to the terminal. The *mlen* value, less four, indicates the length of the message. The ending null byte ('\0') is transmitted and included in the message length.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Servers use the **kXR_attn** response code to optimize overall system performance and to notify clients of any impending events. This response does not correspond to any client request and should not be paired up with any request.
- 3) When **kXR_attn** is received with the **kXR_asyncab** action code, the client should close all physical connections, write the message (*msg*), if any, to standard error, and terminate execution.

2.4.2.2 Server `kXR_attn` Response for `kXR_asyncdi` Client Action

<code>kXR_char</code>	<code>pad[2]</code>
<code>kXR_unt16</code>	<code>kXR_attn</code>
<code>kXR_int32</code>	<code>12</code>
<code>kXR_int32</code>	<code>kXR_asyncdi</code>
<code>kXR_int32</code>	<code>wsec</code>
<code>kXR_int32</code>	<code>msec</code>

Where:

`wsec` is the number of seconds the client should wait before attempting to reconnect to the server.

`msec` is the maximum number of seconds the client should wait before declaring reconnect failure.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The `kXR_char` and `kXR_unt16` data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Servers use the `kXR_attn` response code to optimize overall system performance and to notify clients of any impending events. This response does not correspond to any client request and should not be paired up with any request.
- 3) When `kXR_attn` is received with the `kXR_asyncdi` action code, the client should close the physical connection, wait `wsec` seconds, and attempt to reconnect to the server.
- 4) If a server reconnect fails, the client should wait either an additional `wsec` seconds or some other predetermined time and try again. If `msec` seconds have gone since the initial wait and the client has not reconnected to the server, a reconnect failure should be declared.
- 5) When a reconnect failure is declared, the client may either terminate the program or perform an internal redirection to a load-balancing server.
- 6) A reconnect is essentially a delayed redirect to the same server. The actions that must be carried out when reconnecting are identical to those that must be performed when reconnecting to a different server. Refer to the description of the `kXR_asyncrd` action for the set steps that the client must take to successfully reconnect.

2.4.2.3 Server **kXR_attn** Response for **kXR_asyncgo** Client Action

kXR_char	<i>pad</i> [2]
kXR_unt16	kXR_attn
kXR_int32	4
kXR_int32	kXR_asyncgo

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Servers use the **kXR_attn** response code to optimize overall system performance and to notify clients of any impending events. This response does not correspond to any client request and should not be paired up with any request.
- 3) When **kXR_attn** is received with the **kXR_asyncgo** action code, the client may resume sending requests to the server.
- 4) The **kXR_asyncgo** code is sent to cancel the effects of a previously sent **kXR_asyncwt** code. Therefore, if the client is still waiting for the **kXR_asyncwt** interval to expire, the interval should be cancelled.

2.4.2.4 Server **kXR_attn** Response for **kXR_asyncms** Client Action

kXR_char	<i>pad</i> [2]
kXR_unt16	kXR_attn
kXR_int32	<i>mlen</i>
kXR_int32	kXR_asyncms
kXR_char	<i>msg</i> [<i>mlen</i> -4]

Where:

mlen is the binary length of the following action code and message.

msg is the message to be sent to the terminal. The *mlen* value, less four, indicates the length of the message. The ending null byte ('\0') is transmitted and included in the message length.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Servers use the **kXR_attn** response code to optimize overall system performance and to notify clients of any impending events. This response does not correspond to any client request and should not be paired up with any request.
- 3) When **kXR_attn** is received with the **kXR_asyncms** action code, the client should simply write the indicated message to the terminal.

2.4.2.5 Server kXR_attn Response for kXR_asyncrd Client Action

kXR_char	<i>pad</i> [2]
kXR_unt16	kXR_attn
kXR_int32	<i>plen</i>
kXR_int32	kXR_asyncrd
kXR_int32	<i>port</i>
kXR_char	<i>host</i> [? <i>token</i>] [<i>plen</i> -8]

Where:

plen is the binary length of the *parameter* portion of the message (i.e., the subsequent bytes).

port is the binary port number to which the client must connect. If the value is zero, the default **XRootD** port number must be used. If the value is negative, then the text after *port* contains a standard URL that must be used to effect a new connection. This should only occur if the client has indicated that URL redirection responses are acceptable during the most recent kXR_login request to the redirecting server.

host is the ASCII name of the to which the client must connect. The *host* does *not* end with a null (\0) byte. The host should be interpreted as a standard URL if *port* is negative (see above).

token is an optional ASCII token that, when present, must be delivered to the new host during the login phase, if one is needed. The *token*, if present, is separated from the *host* by a single question mark. The *token* does *not* end with a null (\0) byte.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Servers use the **kXR_attn** response code to optimize overall system performance and to notify clients of any impending events. This response does not correspond to any client request and should not be paired up with any request.

- 3) When **kXR_attn** is received with the **kXR_asyncrd** action code, the client should perform the following steps:
 - a. Decompose the response to extract the port number, host name, and possible token value.
 - b. Physically close the connection to the current host, regardless of type.
 - c. Establish a new physical connection with the indicated host at the specified or default port number.
 - d. Perform the initial handshake, login with token (see **kXR_login** description), and authentication (see **kXR_auth** description).
 - e. Re-establish all open files, as needed. Previously opened files may be re-opened all at once or when a request attempts to use the file.
 - f. Re-issue any requests that were sent to the previous server but have not received a response.
- 4) Since **XRootD** allows multiple open files per physical connection, a **kXR_asyncrd** response can become somewhat complicated to handle. The client can re-open files immediately after a new connection is made or can re-open files as they are needed. In either case, the client must:
 - g. Issue a **kXR_open** request using the same file name and options as was originally used.
 - h. Use the returned file handle for all subsequent requests for that file (i.e., substitute the new *fhandle* for the old *fhandle*).
- 5) An **XRootD** server should never redirect a physical connection to a **rootd** server. This differs for logical connections where a logical connection may be so redirected.
- 6) After 256 redirect responses within 10 minutes on the same physical connection, the client should declare an internal system error since it is obvious that effective work is not being performed.

2.4.2.6 Server kXR_attn Response for kXR_asynresp Client Action

kXR_char	<i>pad</i> [2]
kXR_unt16	kXR_attn
kXR_int32	<i>plen</i>
kXR_int32	kXR_asynresp
kXR_char	<i>reserved</i> [4]
kXR_char	<i>streamid</i> [2]
kXR_unt16	<i>status</i>
kXR_int32	<i>dlen</i>
kXR_char	<i>data</i> [<i>dlen</i>]

Where:

rlen is the binary length of the following action code and response.

streamid

is the stream identifier associated with a previously issued request that received a **kXR_waitresp** response.

status is the binary status code indicating how the request completed. The codes definitions are identical as to those described for synchronous responses.

dlen is the binary length of the *data* portion of the message. If no data is present, then the value is zero.

data are data specific to the request. Not all responses have associated data. If the response does have data, the length of this field is recorded in the *dlen* field.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Servers use the **kXR_attn** response code to optimize overall system performance and to notify clients of any impending events.
- 3) Unlike other asynchronous events, this response is associated with a previous request and the response data must be used to complete that request.
- 4) The *rlen-dlen* is *always* 16.

- 5) When **kXR_attn** is received with the **kXR_asynresp** action code, the client should remove the request paired with *streamid* from wait state and complete it using the response data.

2.4.2.7 Server **kXR_attn** Response for **kXR_asyncwt** Client Action

kXR_char	<i>pad</i> [2]
kXR_unt16	kXR_attn
kXR_int32	8
kXR_int32	kXR_asyncwt
kXR_int32	<i>wsec</i>

Where:

wsec is the number of seconds the client should wait before sending any more requests to the server.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Servers use the **kXR_attn** response code to optimize overall system performance and to notify clients of any impending events. This response does not correspond to any client request and should not be paired up with any request.
- 3) When **kXR_attn** is received with the **kXR_asyncwt** action code, the client should queue any new requests (i.e., not send new requests) until *wsec* seconds have elapsed.
- 4) While waiting, the client should still be receiving messages from the server. It is possible for the server to send additional unsolicited responses even after a **kXR_asyncwt** has been sent. For example, the server may send a **kXR_asyncgo** request to cancel the effects of the **kXR_asyncwt** request before the *wsec* interval has gone by.

2.4.3 Server `kXR_authmore` Response Format

<code>kXR_char</code>	<code>streamid[2]</code>
<code>kXR_unt16</code>	<code>kXR_authmore</code>
<code>kXR_int32</code>	<code>dlen</code>
<code>kXR_char</code>	<code>data[dlen]</code>

Where:

streamid

is the binary identifier that is associated with this request stream corresponding to a previous request.

dlen is the binary length of the *data* portion of the message (i.e., the subsequent bytes).

data is the data, if any, required to continue the authentication process.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The `kXR_char` and `kXR_unt16` data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Since requests may be completed in any order, the ordering of responses is undefined. The client must appropriately pair responses with requests using the *streamid* value.
- 3) The `kXR_authmore` response code is issued only for those authentication schemes that require several handshakes in order to complete (e.g., .x500).
- 4) When a `kXR_authmore` response is received, the client must call the appropriate authentication continuation method and pass it *data*, if present. The output of the continuation method should be sent to the server using another `kXR_auth` request. This handshake continues until either the continuation method fails or the server returns a status code of `kXR_error` or `kXR_ok`.
- 5) Refer to the description of the security framework for detailed information.

2.4.4 Server kXR_error Response Format

kXR_char	<i>streamid</i> [2]
kXR_unt16	kXR_error
kXR_int32	<i>dlen</i>
kXR_int32	<i>errnum</i>
kXR_char	<i>errmsg</i> [<i>dlen</i> -4]

Where:

streamid

is the binary identifier that is associated with this request stream corresponding to a previous request.

dlen is the binary length of the *data* portion of the message (i.e., the subsequent bytes).

errnum

is the binary error number indicating the nature of the problem encountered when processing the request.

errmsg

is the human-readable null-terminated message that describes the error. This message may be displayed for informational purposes.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Since the error message is null-terminated, *dlen* includes the null byte in its count of bytes that were sent.
- 3) Since requests may be completed in any order, the ordering of responses is undefined. The client must appropriately pair responses with requests using the *streamid* value.

2.4.4.1 Server kXR_error Sub-Codes & Recovery Actions

The following table lists possible error sub-codes included in the *errnum* field as part of the **kXR_error** response:

Status	Meaning	Redirector Recovery	Server Recovery
kXR_ArgInvalid	A request argument was not valid	n/a	n/a
kXR_ArgMissing	Required request argument was not provided	n/a	n/a
kXR_ArgTooLong	A request argument was too long (e.g., path)	n/a	n/a
kXR_AttrNotFound	The requested file attribute does not exist	n/a	n/a
kXR_BadPayload	The request arguments were malformed	n/a	n/a
kXR_Cancelled	The operation was cancelled by the administrator	n/a	n/a
kXR_ChkLenErr	The close length does not equal the file size	n/a	n/a
kXR_ChkSumErr	The kXR_verifyw checksum does not match	n/a	n/a
kXR_DecryptErr	Data could not be decrypted	n/a	n/a
kXR_FileLocked	File is locked, open request was rejected	n/a	n/a
kXR_FileNotOpen	File if not open for the request (e.g., read)	n/a	n/a
kXR_FSError	The file system indicated an error	n/a	A
kXR_fsReadOnly	The file system is marked read-only.	n/a	n/a
kXR_inProgress	Operation already in progress	B	B
kXR_InvalidRequest	The request code is invalid	n/a	n/a
kXR_IOError	An I/O error has occurred	n/a	A
kXR_isDirectory	Object being opened with kXR_open is a directory	n/a	n/a
kXR_NoMemory	Insufficient memory to complete the request	C	B
kXR_NoSpace	Insufficient disk space to write data	n/a	n/a
kXR_NotAuthorized	Client is not authorized for the request	n/a	n/a
kXR_NotFile	Object being opened with kXR_open is not a file.	n/a	n/a
kXR_NotFound	The requested file was not found	n/a	D
kXR_noserver	There are no servers available to process the request	n/a	n/a
kXR_overQuota	Space quota exceeded	n/a	n/a
kXR_overloaded	Server is overloaded	C	D
kXR_ServerError	An internal server error has occurred	C	A
kXR_SigVerErr	Request signature could not be verified	n/a	n/a
kXR_TLSRequired	Request requires a TLS connection	n/a	n/a
kXR_Unsupported	The request is valid but not supported	n/a	n/a

- A. Go back to the redirector and ask for a different server. **kXR_refresh** *should not* be turned on and "tried=" CGI value should indicate the hostname of the failing server.
- B. Generally, this represents a programming error. However, should an operation subject to a callback response be retried prior to the callback, this status code may be returned. Clients should honor server's callback requests and wait for a callback response. Therefore, this error can be ignored as long as a callback is outstanding. Otherwise, it should be treated as a fatal error.
- C. If the redirector is replicated, a different redirector should be tried. If all redirectors provide the same response, a fatal error should be reported. In the case of intermediate redirectors (i.e., a redirector transferring the request to another redirector), the recovery may be attempted by treating the intermediate as a server and performing the action outline in A.
- D. Go back to the redirector and ask for a different server. **kXR_refresh** *should* be turned on and "tried=" CGI value should indicate the hostname of the failing server. This should normally be done only once.

2.4.5 Server kXR_ok Response Format

kXR_char	<i>streamid</i> [2]
kXR_unt16	kXR_ok
kXR_int32	<i>dlen</i>
kXR_char	<i>data</i> [<i>dlen</i>]

Where:

streamid

is the binary identifier that is associated with this request stream corresponding to a previous request.

dlen is the binary length of the *data* portion of the message (i.e., the subsequent bytes).

data is the result, if any, of the corresponding request.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Since requests may be completed in any order, the ordering of responses is undefined. The client must appropriately pair responses with requests using the *streamid* value.
- 3) The **kXR_ok** response indicates that the request fully completed and no additional responses will be forthcoming.

2.4.6 Server **kXR_oksofar** Response Format

kXR_char	<i>streamid</i> [2]
kXR_unt16	kXR_oksofar
kXR_int32	<i>dlen</i>
kXR_char	<i>data</i> [<i>dlen</i>]

Where:

streamid

is the binary identifier that is associated with this request stream corresponding to a previous request.

dlen is the binary length of the *data* portion of the message (i.e., the subsequent bytes).

data is the result, if any, of the corresponding request.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Since requests may be completed in any order, the ordering of responses is undefined. The client must appropriately pair responses with requests using the *streamid* value.
- 3) The **kXR_oksofar** response indicates that the server is providing partial results and the client should be prepared to receive additional responses on the same stream. This response is primarily used when a read request would transmit more data than the internal server segment size. Refer to the **kXR_getfile** and **kXR_read** requests.
- 4) Sending requests using the same *streamid* when a **kXR_oksofar** status code has been returned may produced unpredictable results. A client must serialize all requests using the *streamid* in the presence of partial results.
- 5) Any status code other than **kXR_oksofar** indicates the end of transmission

2.4.7 Server kXR_redirect Response Format

kXR_char	<i>streamid</i> [2]
kXR_uint16	kXR_redirect
kXR_int32	<i>dlen</i>
kXR_int32	<i>port</i> 0xffffffff
kXR_char	<i>host</i> [? <i>opaque</i>][? <i>token</i>]][<i>dlen-4</i>] <i>url</i>

Where:

streamid

is the binary identifier that is associated with this request stream corresponding to a previous request.

dlen is the binary length of the *data* portion of the message (i.e., the subsequent bytes).

port is the binary port number to which the client must connect. If the value is zero, the default **XRootD** port number must be used. If the value is negative, then the text after *port* contains a standard URL that must be used to effect a new connection. This should only occur if the client has indicated that URL redirection responses are acceptable during the most recent kXR_login request to the redirecting server.

host is the ASCII name of the to which the client must connect. The *host* does *not* end with a null (\0) byte. The *host* should be interpreted as a standard URL if *port* is negative (see above).

opaque is an optional ASCII token that, when present, must be delivered to the new host as opaque information added to the file name² associated with the operation being redirected. The *opaque*, if present, is separated from the *host* by a single question mark. The *opaque* does *not* end with a null (\0) byte but may end with a question mark (see *token* below). Therefore, *opaque* may never contain a question mark.

² In the case of **kXR_mv**, two file names are present. The opaque information must be added to the second of the two file names.

token is an optional ASCII token that, when present, must be delivered to the new host during the login phase, if one is needed (i.e. established connections to the specified host may be re-used without a login). The *token*, if present, is separated from the *host* by a two question marks. The first question mark may be followed by *opaque* information. If none is present, another question mark immediately follows the first one. The *token* does *not* end with a null ($\backslash 0$) byte.

url when a client indicates that it supports multi-protocol redirects, the server may respond with an actual *url*. In this case, the *port* value is set to -1.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Since requests may be completed in any order, the ordering of responses is undefined. The client must appropriately pair responses with requests using the *streamid* value.
- 3) After 256 redirect responses within 10 minutes on the same logical connection, the client should declare an internal system error since it is obvious that effective work is not being performed.
- 4) The client must be prepared to handle a redirect response at any time. A redirect response requires that the client
 - i. Decompose the response to extract the port number, host name, and possible token value.
 - j. Possibly close the connection of the current host, if the current host is a data server and this is the last logical connection to the server. Otherwise, if this is the first load-balancing server encountered in the operation sequence, the connection should remain open since a load-balancing server always responds with a redirect.
 - k. Establish a new logical connection with the indicated host at the specified or default port number. If a physical connection already exists and is session compatible with the new logical connection; the existing physical connection should be reused and the next step (i.e. handshake and login) should be skipped.
 - l. Perform the initial handshake, login with token (see **kXR_login** description), and authentication (see **kXR_auth** description).
 - m. If the redirection occurred for a request using a file handle (i.e., *fhandle*) then a new file handle must be obtained.

- i. A **kXR_open** request must be issued using the same file name and options as was originally used.
 - ii. The returned file handle must be used for the request that is to be re-issued as well as all subsequent requests relating to the file.
 - n. Re-issue the request that was redirected.
- 5) Opaque data must be treated as truly opaque. The client should not inspect nor modify the data in any way.

2.4.8 Server kXR_wait Response Format

kXR_char	<i>streamid</i> [2]
kXR_unt16	kXR_wait
kXR_int32	<i>dlen</i>
kXR_int32	<i>seconds</i>
kXR_char	<i>infomsg</i> [<i>dlen</i> -4]

Where:

streamid

is the binary identifier that is associated with this request stream corresponding to a previous request.

dlen is the binary length of the *data* portion of the message (i.e., the subsequent bytes).

seconds

is the maximum binary number of seconds that the client needs to wait before re-issuing the request.

infomsg

is the human-readable message that describes the reason of why the wait is necessary. The message does *not* end with a null (\0) byte. This message may be displayed for informational purposes.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Since requests may be completed in any order, the ordering of responses is undefined. The client must appropriately pair responses with requests using the *streamid* value.
- 3) The client should wait the indicated number of seconds and retry the request.
- 4) Nothing prohibits the client from waiting for less time than the indicated number of seconds.

2.4.9 Server `kXR_waitresp` Response Format

<code>kXR_char</code>	<code>streamid[2]</code>
<code>kXR_unt16</code>	<code>kXR_waitresp</code>
<code>kXR_int32</code>	<code>4</code>
<code>kXR_int32</code>	<code>seconds</code>

Where:

streamid

is the binary identifier that is associated with this request stream corresponding to a previous request.

seconds

is the *estimated* maximum binary number of seconds that the client needs to wait for the response.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The `kXR_char` and `kXR_unt16` data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Since requests may be completed in any order, the ordering of responses is undefined. The client must appropriately pair responses with requests using the *streamid* value.
- 3) The client should wait the indicated number of seconds for the response. The response should be returned via an unsolicited response (`kXR_attn` with `kXR_asynresp`) at some later time which may be earlier than the time indicated in seconds. When the response arrives, the client must use the response data to complete the request that received the `kXR_waitresp`.
- 4) Nothing prohibits the client from waiting for different time than the indicated number of *seconds*. Generally, if no response is received after at least *seconds* have elapsed; the client should treat the condition as a fatal error.

2.5 Binary Definitions of Status and Error codes

2.5.1 Status Codes

Status Code	Value
kXR_ok	0
kXR_oksofar	4000
kXR_attn	4001
kXR_authmore	4002
kXR_error	4003
kXR_redirect	4004
kXR_wait	4005
kXR_waitresp	4006

2.5.2 kXR_attn Status Subcodes

kXR_attn Subcode	Value
kXR_asyncab	5000
kXR_asyncdi	5001
kXR_asyncms	5002
kXR_asyncrd	5003
kXR_asyncwt	5004
kXR_asyncav	5005
kXR_asyncunav	5006
kXR_asyncgo	5007
kXR_asyncresp	5008

2.5.3 Error Codes

Error	Value	Corresponding POSIX errno Value
kXR_ArgInvalid	3000	EINVAL
kXR_ArgMissing	3001	EINVAL
kXR_ArgTooLong	3002	ENAMETOOLONG
kXR_FileLocked	3003	EDEADLK
kXR_FileNotOpen	3004	EBADF
kXR_FSError	3005	EIO
kXR_InvalidRequest	3006	EEXIST
kXR_IOError	3007	EIO
kXR_NoMemory	3008	ENOMEM
kXR_NoSpace	3009	ENOSPC
kXR_NotAuthorized	3010	EACCES
kXR_NotFound	3011	ENOENT
kXR_ServerError	3012	ENOMSG
kXR_Unsupported	3013	ENOTSUP
kXR_noserver	3014	EHOSTUNREACH
kXR_NotFile	3015	ENOTBLK
kXR_isDirectory	3016	EISDIR
kXR_Canceled	3017	ECANCELED
kXR_ChkLenErr	3018	EDOM
kXR_ChkSumErr	3019	EDOM
kXR_inProgress	3020	EINPROGRESS
kXR_overQuota	3021	EDQUOT
kXR_SigVerErr	3022	EILSEQ
kXR_DecryptErr	3023	ERANGE
kXR_Overloaded	3024	EUSERS
kXR_fsReadOnly	3025	EROFS
kXR_BadPayload	3026	EINVAL
kXR_AttrNotFound	3027	ENOATTR
kXR_TLSRequired	3028	EPROTOTYPE

3 Transport Layer Security (TLS) Support

The **XRootD** protocol supports **TLS** mode connections in two explicit ways:

- 1) client request using the **kXR_protocol** request, and
- 2) server request using the **kXR_protocol** response.

This mechanism provides several features:

- A single port can be used for **TLS** and non-**TLS** connections.
- The request channel can be split from the data channel using the **kXR_bind** request so that control information flows on a **TLS** connection while data flows on a non-**TLS** connection. Such an arrangement may significantly improve performance.
- The number of interactions can be reduced when a connection needs to use **TLS**.
- The server may independently enforce **TLS** requirements in for broad categories:
 - logins and all subsequent interactions,
 - all post-login interactions,
 - third party copy requests, and
 - data transfers.

Currently, once a connection switches to **TLS** mode it cannot switch back. This is not a protocol requirement but a practical side-effect of current **TLS** implementations that buffer an indeterminate amount of data making it problematic to deterministically switch modes. However, the **XRootD** protocol is sufficiently open to allow such switches if and when the **TLS** protocol can do so in the future.

A server is not required to support **TLS**. If it does, it should follow the protocol specifications described in the **kXR_protocol** and **kXR_bind** requests.

TLS may be considered a replacement for request signing in most circumstances. However, for certain workflows, request signing may offer better performance. Be ware, that **XRootD** request signing, as defined, does not protect data while **TLS**., when used for data, does so.

TLS

3.1 Client-Server interactions to unilaterally use TLS

- The client should connect to the server using a non-TLS connection and send the handshake packet.
- The client should then send a **kXR_protocol** request indicating that it wants to use **TLS**. For reduced latency, the handshake and the **kXR_protocol** request may be sent together.
- If the server supports **TLS** it should indicate in the **kXR_protocol** response that the connection will be switched to use **TLS** after the response is sent.
- The client should check if the server switched the connection to use **TLS** and do the same if so indicated.
- All communications from then on use **TLS**.

3.2 Client-Server interactions to use TLS only when required

- The client should connect to the server using a non-TLS connection and send the handshake packet.
- The client should then send a **kXR_protocol** request indicating that it is able to use **TLS**. For reduced latency, the handshake and the **kXR_protocol** request may be sent together. In the **kXR_protocol** request the client should also indicate the expected next operation (i.e. login, data transfer, or third party copy).
- If the server supports **TLS** it should indicate in the **kXR_protocol** response that the connection has been switched to use **TLS** if the client's subsequent operation requires **TLS**. Note that it is also possible for the server to indicate that **TLS** is required after the **kXR_login** request (i.e. login does not require **TLS**).
- The client should check if the server switched the connection to use **TLS** and do the same if so indicated. If the next request is a **kXR_login** and the server indicated that **TLS** is not required until after the login, the client should defer switching the connection to **TLS** until after the login and all authentication interactions (i.e. **kXR_auth** requests).

4 Detailed Protocol Specifications

4.1 kXR_admin Request

Purpose: Perform an administrative function.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_admin	kXR_uint16 kXR_ok
kXR_char <i>reserved</i> [16]	kXR_int32 <i>rlen</i>
kXR_int32 <i>rlen</i>	kXR_char <i>resp</i> [<i>rlen</i>]
kXR_char <i>reqs</i> [<i>rlen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

rlen is the binary length of the supplied request, *reqs*.

reqs is the request.

ilen is the binary length of the response, *resp*, that follows *ilen*.

resp is the response to the administrative request.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_uint16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The **kXR_admin** request is only valid for users who have successfully performed a **kXR_login** operation in an administrative role (i.e., logged in as administrator).
- 3) This request type is not currently supported. Use the local socket interface protocol to execute administrative requests.

kXR_admin

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_admin		3020	Perform administrative actions.

kXR_auth

4.2 kXR_auth Request

Purpose: Authenticate client's username to the server.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_auth	kXR_uint16 kXR_ok
kXR_char <i>reserved</i> [12]	kXR_int32 0
kXR_char <i>credtype</i> [4]	
kXR_int32 <i>credlen</i>	
kXR_char <i>cred</i> [<i>credlen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed **kXR_int32** with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

credtype

the first four characters of the protocol name. If the protocol name is less than four characters, the name should be null terminated.

credlen

is the binary length of the supplied credentials, *cred*.

cred are the credentials used to provide authentication information.

kXR_auth

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Authentication credentials may be supplied by many means. The common mechanism used by **XRootD** is to use the classes in the **libXrdSec.so** library. See the "Authentication & Access Control Configuration Reference" for more information.
- 3) Refer to the description of the security framework on how a client authenticates to an **XRootD** server.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_auth		3000	Perform authentication

kXR_bind

4.3 kXR_bind Request

Purpose: Bind a socket to a pre-existing session.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_bind	kXR_unt16 kXR_ok
kXR_char <i>sessid</i> [16]	kXR_int32 1
kXR_int32 0	kXR_char <i>pathid</i>

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

sessid is the session identifier returned by a previous **kXR_login** request.

pathid is the socket identifier associated with this connection. The *pathid* may be used in subsequent **kXR_read**, **kXR_readv**, and **kXR_write** requests to indicate which socket should be used for a response or as a source of data.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The *sessid* value should be treated as opaque data.
- 3) The socket issuing the **kXR_bind** request must neither have a session id (i.e., be logged in) nor be already bound.
- 4) Once a socket is bound to a session, it may only supply data for **kXR_write** requests or receive responses for **kXR_read** and **kXR_readv** requests.
- 5) Each login session is limited to the number of bound sockets. Use the **kXR_Qconfig** sub-request code of **kXR_query** to determine the maximum number of sockets that can be bound to a login session.
- 6) Bound sockets are meant to support parallel data transfer requests across wide-area networks. They are also meant to split control information from data allowing control to flow on a TLS connection while data flows on a non-TLS connection. See [TLS Considerations](#) for more information.

kXR_bind

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_bind		3024	Bind additional sockets to session

4.3.1 TLS Considerations

A server may indicate in the response to the **kXR_protocol** request that all data must flow across a **TLS** connection. The **kXR_protocol** request is normally sent by the client immediately after the handshake. If the the server's response indicates that **TLS** must be used for data then the connection to be bound should be set to **TLS** mode in order for the request to succeed. There are two ways to achieve this.

The client may record whether or not the bound connections must use **TLS**. If bound connections must use **TLS** the **kXR_bind** request should be prefixed by a **kXR_protocol** request indicating that the connection should be switched to **TLS**. To reduce latency, the **kXR_protocol** and **kXR_bind** requests should be sent together. This method is preferred.

Alternatively, the client may always send a **kXR_protocol** request ahead of the **kXR_bind** request indicating in the request that it is able to use **TLS** and the next request will be **kXR_bind**. If the server requires the data to use **TLS** it should respond that the connection will switch to using **TLS** after the **kXR_protocol** response is sent. If the connection was switched to use **TLS** the client should do the same and then send the **kXR_bind** request. Since this involves additional interactions, it is not the preferred method.

The client is also free to switch the connection to use **TLS** whether or not the server requires it to do so.

4.4 kXR_chmod Request

Purpose: Change the access mode on a directory or a file.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_chmod	kXR_uint16 kXR_ok
kXR_char <i>reserved</i> [14]	kXR_int32 0
kXR_int16 <i>mode</i>	
kXR_int32 <i>plen</i>	
kXR_char <i>path</i> [<i>plen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

mode is the access mode to be set for *path*. The access mode is an "or'd" combination of the following values:

Access	Readable	Writable	Executable
Owner	kXR_ur	kXR_uw	<i>not supported</i>
Group	kXR_gr	kXR_gw	<i>not supported</i>
Other	kXR_or	<i>not supported</i>	<i>not supported</i>

plen is the binary length of the supplied path, *path*.

path is the path whose mode is to be set. It may be suffixed with CGI information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_uint16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) No **umask** is applied to the specified mode.

kXR_chmod

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_chmod		3002	Change directory or file permissions
	<i>mode</i>		
	kXR_ur	0x01 00	Owner readable
	kXR_uw	0x00 80	Owner writable
	kXR_ux	0x00 40	Owner searchable (directories)
	kXR_gr	0x00 20	Group readable
	kXR_gw	0x00 10	Group writable
	kXR_gx	0x00 08	Group searchable (directories)
	kXR_or	0x00 04	Other readable
	kXR_ow	0x00 02	Other writable
	kXR_ox	0x00 01	Other searchable (directories)

kXR_close

4.5 kXR_close Request

Purpose: Close a previously opened file, communications path, or path group.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_close	kXR_uint16 kXR_ok
kXR_char <i>fhandle</i> [4]	kXR_int32 0
kXR_char <i>reserved</i> [12]	
kXR_int32 0	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

fhandle

is the file handle value supplied by the successful response to the associated **kXR_open** request.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_uint16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The *fhandle* value should be treated as opaque data.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_close		3003	Close an open file

4.6 kXR_decrypt Request

Purpose: Signal when the data stream is encrypted (placeholder for future enhancement)..

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_decrypt	kXR_uint16 kXR_ok
kXR_char <i>reserved</i> [16]	kXR_int32 0
kXR_int32 0	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is a reserved field and should be set to zero.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_uint16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The **kXR_decrypt** request should be considered as not fully specified. It is currently a place-holder for future enhancement.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_decrypt		3030	Indicate an encrypted data stream

4.7 kXR_dirlist Request

Purpose: Enumerate the contents of a directory.

Request	Normal Response
kXR_char <i>streamid</i> [2] kXR_unt16 kXR_dirlist kXR_char <i>reserved</i> [15] kXR_char <i>options</i> kXR_int32 <i>plen</i> kXR_char <i>path</i> [<i>plen</i>]	kXR_char <i>streamid</i> [2] kXR_unt16 kXR_ok kXR_int32 <i>dlen</i> kXR_char <i>Dirname</i> ₀ \n . . . kXR_char <i>dirname</i> _n kXR_char 0
	Normal Response w/kXR_dstat kXR_char <i>streamid</i> [2] kXR_unt16 kXR_ok kXR_int32 <i>dlen</i> kXR_char ".\n" kXR_char "0 0 0 0\n" kXR_char <i>dirname</i> ₀ \n kXR_char <i>statinfo</i> ₀ \n . . . kXR_char <i>dirname</i> _n \n kXR_char <i>statinfo</i> _n kXR_char 0

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

options

is, optionally, one or more of the following:

kXR_dstat - return stat information with each entry (protocol version 3+).

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

plen is the binary length of the supplied path, *path*.

path is the path of a directory whose entries are to be listed. It may be suffixed with CGI information.

dlen is the binary length of the data that follows *dlen*.

dirname

is an entry in the directory whose listing was requested.

statinfo

the **kXR_stat** information for the preceding *dirname*. Refer to **kXR_stat** for details. The *statinfo* is only returned when **kXR_dstat** is set and the server implements protocol version 3 or higher.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields should be initialized to binary zero.
- 2) A directory may have multiple entries and the response contains all of the entries.
- 3) Each directory entry should be suffixed by a new-line character; except for the last entry which should be suffixed by a null character. When **kXR_dstat** is specified, the last entry is the couplet "*dirname*\n*statinfo*".
- 4) Since more entries may exist than is possible to send at one time, the **kXR_oksofar** protocol may be used to segment the response. Under no circumstances should a directory name be split across a response packet.
- 5) The server should not return the entries "." and ".." except when **kXR_dstat** is specified, in which case only the "." entry is returned.
- 6) An empty directory should return the eight-byte triplet {*streamid*, 0, 0} unless **kXR_dstat** is specified; in which case "{*streamid*,0,8}.\n0 0 0\0" should be returned.
- 7) Clients should always check if the server supports **kXR_dstat**. If the option is supported, the first entry should be a *dot* entry followed the zero stat information.

kXR_dirlist

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_dirlist		3004	List a directory
	<i>options</i>		
	kXR_dstat	0x02	Return stat information with entry
	kXR_online	0x01	Only list online entries

4.8 kXR_endsess Request

Purpose: Terminate a pre-existing session.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_endsess	kXR_unt16 kXR_ok
kXR_char <i>reserved</i> [16]	kXR_int32 0
kXR_int32 0	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

sessid

is the session identifier returned by a previous **kXR_login** request.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The *sessid* value should be treated as opaque data.
- 3) The socket issuing the **kXR_endsess** request must be logged in and, optionally, authenticated.
- 4) If the *sessid* is all binary zeroes, the current session is terminated.
- 5) The server verifies that the process presenting the *sessid* actually received it on a previous **kXR_login**.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_endsess		3023	End previous session

4.9 kXR_fattr Request

The **kXR_fattr** request code is used to delete, list, retrieve, and set file attributes (also known as extended attributes). This is accomplished using request subcodes. File attributes are specific to the file system being exported by the server. The exported file system has its own specific limits on the length of attribute names and the amount of data that may associated with a name. Some even limit the total amount of attribute data that may be associated with a file. Finally, not all file systems support extended attributes. The [kXR_query](#) request using the **kXR_QConfig** subcode with the **xattr** argument may be used to [ascertain limits](#) for any particular server.

The **kXR_fattr** request imposes its own limits on the maximum length of an attribute name (i.e. **kXR_faMaxNlen**, currently 248 bytes) and attributes value (i.e. **kXR_faMaxVlen**, currently 65536 bytes or 64K). Be aware that smaller limits may apply, depending on the underlying file system.

The **kXR_fattr** request supports deleting, retrieving, and setting multiple attributes with one request. However, the operation should not be considered atomic when multiple attributes are specified. A maximum **kXR_faMaxVars** (currently 16) attribute vales may be deleted, set, or retrieved per request.

For delete and retrieve requests, only attribute names are specified. For set requests, the attribute names are followed by the corresponding values (i.e. in 1-to-correspondence to the names) to be used for each attribute. Regardless of the subcode, the first string in each request is the path name of the file to which the request applies; which may be a null string.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_fattr		3032	Perform file attribute function
	kXR_fattrDel	0	Delete one or more attributes
	kXR_fattrGet	1	Get one or more attributes
	kXR_fattrList	2	List file attribute names
	kXR_fattrSet	3	Set one or more attributes
	isNew	0x01	Attribute must not exist
	aData	0x10	Include attribute value

4.9.1 Layout of *namevec*

Subsequent sections refer to *namevec* which is a vector whose elements are laid out as follows:

kXR_unt16	<i>rc</i>
kXR_char	<i>name</i> []
kXR_char	0

Where:

rc as an argument is should be set to zero. In the response, it holds the status code associated with the attribute name. A status code not equal to kXR_ok indicates that the requested operation with respect to the attribute name was not completed.

name is the name of an attribute. The length of each *name*, excluding the null byte, should not be greater than **kXR_faMaxNlen**. Notice that the *name* is followed by a null byte. Attribute names are null terminated strings. These elements are concatenated together to produce a vector of names.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values.
- 2) There is no alignment requirement in the for the *namevec* element. That is, *namevec* elements should be streamed together irrespective of byte boundaries.
- 3) A *namevec* element should not be split across **kXR_oksofar** responses.

kXR_fattr

4.9.2 Layout of *valuvec*

Subsequent sections refer to *valuvec* which is a vector whose elements are laid out as follows:

```
kXR_int32    vlen
kXR_char     value[vlen]
```

Where:

vlen is the length of the subsequent *value*.

value is the value that the attribute is to have when issuing a **kXR_fattrSet** request or the actual value of the attribute when issuing a **kXR_fattrGet** request.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values.
- 2) There is no alignment requirement in the for the *valuvec* element. That is, *valuvec* elements should be streamed together irrespective of byte boundaries.
- 3) A *valuvec* element should not be split across **kXR_oksofar** responses.

4.9.3 kXR_fattr Request – Delete Subcode

Purpose: Delete one or more file attributes.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_fattr	kXR_uint16 kXR_ok
kXR_char <i>fhandle</i> [4]	kXR_int32 <i>rlen</i>
kXR_char kXR_fattrDel	kXR_char <i>nerrs</i>
kXR_char <i>nattr</i>	kXR_char <i>nattr</i>
kXR_char <i>options</i>	kXR_char <i>namevec</i> [<i>nattr</i>]
kXR_char <i>reserved</i> [9]	
kXR_int32 <i>alen</i>	
kXR_char <i>path</i>	
kXR_char <i>namevec</i> [<i>nattr</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

fhandle

is the file handle value supplied by the successful response to the associated **kXR_open** request that is to be used for the request when no path is supplied (i.e, *path* is a null string). If a *path* is supplied, *fhandle* should be ignored.

nattr is the number of attribute names that follow. The value should be one or greater but no more than **kXR_faMaxVars**.

options

reserved for future options.

alen is the binary length of the arguments that follow the request header.

path is the null terminated path. The *path* may be suffixed with CGI information. If *path* is a null string (i.e. only contains a null byte) then *fhandle* should be used to identify the file to which this request applies.

kXR_fattr - del

namevec

is a vector of null terminated attribute names. Each name in the vector is preceded by two bytes of zero. The number of such names concatenated together should equal *nattr*. The length of each name, excluding the null byte, should not be greater than **kXR_faMaxNlen**. The *namevec* layout is [described here](#).

rlen is the binary length of the response that follow the request header.

nerrs is the number of variables in *namevec* that could not be deleted. The two byte field preceding the name contains a status code (i.e. *rc* in *namevec*). When it contains **kXR_OK** then variable was deleted. Otherwise, it should be the error code describing the error encountered when deleting the variable.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The *fhandle* value should be treated as opaque data.
- 3) The server should process the elements in the order specified.
- 4) There are no alignment requirements in the argument or response portions of the request.
- 5) Deletion of extended attributes should be restricted to clients with write access to the target file.

4.9.4 kXR_fattr Request – Get Subcode

Purpose: Retrieve one or more file attributes.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_fattr	kXR_unt16 kXR_ok
kXR_char <i>fhandle</i> [4]	kXR_int32 <i>rlen</i>
kXR_char kXR_fattrGet	kXR_char <i>nerrs</i>
kXR_char <i>nattr</i>	kXR_char <i>nattr</i>
kXR_char <i>options</i>	kXR_char <i>namevec</i> [<i>nattr</i>]
kXR_char <i>reserved</i> [9]	kXR_char <i>valuevec</i> [<i>nattr</i>]
kXR_int32 <i>alen</i>	
kXR_char <i>path</i>	
kXR_char <i>namevec</i> [<i>nattr</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

fhandle

is the file handle value supplied by the successful response to the associated **kXR_open** request that is to be used for the request when no path is supplied (i.e, *path* is a null string). If a *path* is supplied, *fhandle* should be ignored.

nattr is the number attribute names that follow. The value should be one or greater but no more than **kXR_faMaxVars**.

options

reserved for future options.

alen is the binary length of the arguments that follow the request header.

path is the null terminated path. The *path* may be suffixed with CGI information. If *path* is a null string (i.e. only contains a null byte) then *fhandle* should be used to identify the file to which this request applies.

kXR_fattr - get

namevec

is a vector of null terminated attribute names. Each name in the vector is preceded by two bytes of zero. The number of such names concatenated together should equal *nattr*. The length of each name, excluding the null byte, should not be greater than **kXR_faMaxNlen**.

The *namevec* is echoed in the response. The two byte header in each name is replaced by the status code associated with retrieving the value (i.e. *rc* in *namevec*). The *namevec* layout is [described here](#).

rlen is the binary length of the response that follows the request header.

nerrs is the number of variables in *namevec* that could not be retrieved. The two byte field preceding the name contains a status code (i.e. *rc* in *namevec*). When it contains **kXR_ok** then variable's value was retrieved. Otherwise, it is the error code describing the error encountered when retrieving the variable.

valuvec

is the value corresponding to the specified attribute name. Values are returned in name specified order (i.e. there should be a 1-to-1 correspondence between *namevec* and *valuvec*). For attribute names that indicate an error the length for the corresponding value should be set to zero. If the attribute, in fact, has no associated value (i.e. it exists but the data is null) then the status code associated with the attribute name should be set to **kXR_ok**. The *valuvec* layout is [described here](#).

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The *fhandle* value should be treated as opaque data.
- 3) The server should process the elements in the order specified.
- 4) Only those variables that can be set via **kXR_fattr** should be returned.
- 5) There are no alignment requirements in the argument or response portions of the request.
- 6) Retrieval of extended attributes should be restricted to clients with read access to the target file.

4.9.5 kXR_fattr Request – List Subcode

Purpose: List file attribute names.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_fattr	kXR_unt16 <i>statok</i>
kXR_char <i>fhandle</i> [4]	kXR_int32 <i>rlen</i>
kXR_char kXR_fattrList	kXR_char <i>names</i> [<i>rlen</i>]
kXR_char <i>reserved</i>	
kXR_char <i>options</i>	Response with ::adata set
kXR_char <i>reserved</i> [9]	kXR_char <i>streamid</i> [2]
kXR_int32 <i>alen</i>	kXR_unt16 <i>statok</i>
kXR_char <i>path</i>	kXR_int32 <i>rlen</i>
	kXR_char { <i>name</i>
	kXR_int32 <i>vlen</i>
	kXR_char <i>value</i> [<i>vlen</i>]
	} []

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

fhandle

is the file handle value supplied by the successful response to the associated **kXR_open** request that is to be used for the request when no path is supplied (i.e, *path* is a null string). If a *path* is supplied, *fhandle* should be ignored.

options

ClientFattrRequest::adata include the attribute value in the response.

alen is the binary length of the arguments that follow the request header.

path is the null terminated path. The *path* may be suffixed with CGI information. If *path* is a null string (i.e. only contains a null byte) then *fhandle* should be used to identify the file to which this request applies. This is should also be the case when *alen* is zero.

kXR_fattr - list

statok is one of two status codes:

kXR_ok

indicates successful completion as a final response.

kXR_oksofar

indicates that a subsequent response should follow with more data. In either case, the response header is followed by one or more null terminated attribute names. Attribute names and optional subsequent values should not be split across response segments.

rlen is the binary length of the of the response data that follows.

names if *rlen* is not zero, then one or more null terminated attribute names forming a list of names (e.g. *name\0[name\0[...]]*).

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The *fhandle* value should be treated as opaque data.
- 3) An attribute name should never be split across multiple responses
- 4) Only settable variables via **kXR_fattr** should be returned.
- 5) When **ClientFattrRequest::adata** is specified, attribute names whose value cannot be retrieved should not be returned.
- 6) There are no alignment requirements in the argument or response portions of the request.
- 7) Listing of extended attributes should be restricted to clients with read access to the target file.

4.9.6 kXR_fattr Request – Set Subcode

Purpose: Set one or more file attributes.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_fattr	kXR_uint16 kXR_ok
kXR_char <i>fhandle</i> [4]	kXR_int32 <i>rlen</i>
kXR_char kXR_fattrList	kXR_char <i>nerrs</i>
kXR_char <i>nattr</i>	kXR_char <i>nattr</i>
kXR_char <i>options</i>	kXR_char <i>namevec</i> [<i>nattr</i>]
kXR_char <i>reserved</i> [9]	
kXR_int32 <i>alen</i>	
kXR_char <i>path</i>	
kXR_char <i>namevec</i> [<i>nattr</i>]	
kXR_char <i>valuevec</i> [<i>nattr</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

fhandle

is the file handle value supplied by the successful response to the associated **kXR_open** request that is to be used for the request when no path is supplied (i.e, *path* is a null string). If a *path* is supplied, *fhandle* should be ignored.

nattr is the number attribute name-value pairs that follow. The value should be one or greater but no more than **kXR_faMaxVars**.

options

is one of the following options:

isNew - the variable should only be set if it does not exist.

alen is the binary length of the arguments that follow the request header.

path is the null terminated path. The *path* may be suffixed with CGI information. If *path* is a null string (i.e. only contains a null byte) then *fhandle* should be used to identify the file to which this request applies.

kXR_fattr - set

namevec

is a vector of null terminated attribute names. Each name in the vector is preceded by two bytes of zero. The number of such names concatenated together should equal *nattr*. The length of each name, excluding the null byte, should not be greater than **kXR_faMaxNlen**.

The *namevec* is echoed in the response. The two byte header in each name is replaced by the status code associated with setting the value (i.e. *rc* in *namevec*). The *namevec* layout is [described here](#).

valuvec

is a vector of attribute values. Each value starts with a four byte length which may be zero to set an attribute without a corresponding value. The length should not be greater than **kXR_faMaxVlen**. The *valuvec* layout is [described here](#).

nerrs is the number of variables in *namevec* that could not be set. The two byte field, *rc*, preceding the name contains a status code. When it contains **kXR_ok** then variable's value was set. Otherwise, it should be the error code describing the error encountered when setting the variable.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The *fhandle* value should be treated as opaque data.
- 3) The server should process the elements in the order specified.
- 4) Attributes set via **kXR_fattr** should be placed in a separate internal namespace to avoid conflicts with other extended attributes.
- 5) There are no alignment requirements in the argument or response portions of the request.
- 6) Setting of extended attributes should be restricted to clients with write access to the target file.

4.10 kXR_getfile Request

Purpose: Retrieve a complete file (placeholder for next generation TPC).

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_getfile	kXR_uint16 kXR_ok
kXR_int32 <i>options</i>	kXR_int32 0
kXR_char <i>reserved</i> [12]	
kXR_int32 <i>plen</i>	
kXR_char <i>path</i> [<i>plen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

options

request options

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

plen is the binary length of the supplied path, *path*.

path is the path of the file to be retrieved. The *path* may be suffixed with CGI information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_uint16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) This request type is not currently supported.
- 3) The **kXR_getfile** request should be considered as not fully specified. It is currently a place-holder for future enhancement and may substantially change in functionality.

kXR_getfile

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_getfile		3005	n/a

4.11 kXR_locate Request

Purpose: Locate a file.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_locate	kXR_uint16 kXR_ok
kXR_int16 <i>options</i>	kXR_int32 <i>rlen</i>
kXR_char <i>reserved</i> [14]	kXR_char <i>info</i> [<i>rlen</i>]
kXR_int32 <i>plen</i>	
kXR_char <i>path</i> [<i>plen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

options

are the options to apply when *path* is opened. The *options* are an "or'd" combination of the following values:

- kXR_nowait** - provide information as soon as possible
- kXR_prefname** - hostname response is preferred
- kXR_refresh** - update cached information on the file's location (see notes)

.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

plen is the binary length of the supplied path, *path*.

path is the path of the file to be located. CGI information appended to the path does not affect the request. Path may also start with an asterisk or be only an asterisk with the following meaning:

- * - return all connected managers and servers
- **path* - return all managers and servers exporting *path*

rlen is the byte length of the response that follows

kXR_locate

info are zero or more node types, IPV6 hybrid addresses, and port numbers of nodes that have the file. The port number is to be used to contact the node.

Node Entry Response Format

```
xy[::aaa.bbb.ccc.ddd.eee]:ppppp  
xyhostname:ppppp
```

Where:

x is a single character that identifies the type of node whose IP address follows. Valid characters are:

- M** - Manager node where the file is online
- m** - Manager node where the file is pending to be online.
- S** - Server node where the file is online
- s** - Server node where the file is pending to be online.

y is a single character that identifies the file access mode at the node whose IP address follows. Valid characters are:

- r** - Read access allowed
- w** - Read and write access allowed.

aaa.bbb.ccc.ddd.eee

is the IPv4 portion of the IPV6 node address, for IPV4 environments. Otherwise, a true IPV6 address is returned.

hostname

is the hostname for the node address. This format may only be returned when kXR_prefname is specified, but does not forbid an address reply.

ppppp is the port number to be used for contacting the node.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Option flags are the same as those defined for the **kXR_open** request.
- 3) The **kXR_refresh** voids the **kXR_nowait** option.
- 4) If the file resides in more than one location, each location is separated by a space.
- 5) The **kXR_nowait** option provides a location as soon as one becomes known. This means that not all locations are necessarily returned. If the file does not exist, a wait is still imposed.
- 6) If available, use the `inet_ntop()` and `inet_pton()` function to convert addresses to suitable format as these accepts traditional IPV4 address as well as IPV6 addresses.
- 7) Nodes identified as **M** or **m**, do not actually hold the file. These are manager nodes that know other locations for the file. To obtain the real file location, the client must contact each **M(m)** node and issue a **kXR_locate** request. The process is iterative, as the response from an **M(m)** node may identify other **M(m)** nodes.
- 8) Clients should guard against circular references by setting an absolute depth limit in the number of **M(m)** to **M(m)** references they will accept before declaring an error. A limit of 4 covers a range of 16,777,216 possible locations.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_locate		3027	Perform location operation
	<i>options</i>		
	kXR_compress	0x00 01	Return unique hosts
	kXR_nowait	0x20 00	Return immediate information
	kXR_prefname	0x01 00	Preferentially return DNS names
	kXR_refresh	0x00 80	Refresh cached information

4.12 kXR_login Request

Purpose: Initialize a server connection.

Request	Normal Response
	server < 2.4.0 client < 1.0
kXR_char streamid[2]	kXR_char streamid[2]
kXR_unt16 kXR_login	kXR_unt16 kXR_ok
kXR_int32 pid	kXR_int32 slen
kXR_char username[8]	kXR_char sec[slen]
kXR_char reserved	server >= 2.4.0 & client >= 1.0
kXR_char ability	kXR_char streamid[2]
kXR_char capver	kXR_unt16 kXR_ok
kXR_char role	kXR_int32 slen+16
kXR_int32 tlen	kXR_char sessid[16]
kXR_char token[tlen]	kXR_char sec[slen]

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

pid is the process number associated with this connection.

username

is the unauthenticated name of the user to be associated with the connection on which the login is sent.

kXR_login

ability are the client's extended capabilities represented as bit flags, as follows:

- 0b00000001** the client accepts full standard URL's in a redirection response. Unless the following ability is set, the protocol in the URL should remain xroot. This bit is also identified as **kXR_fullurl**.
- 0b00000011** the client accepts protocol changes in a full standard URL's in a redirection response. Unless the this ability is set, the protocol in the URL should remain xroot. This bit is also identified as **kXR_multipr**.
- 0b00000100** the client accepts protocol redirects during a **kXR_read** and **kXR_readv** requests. This bit is also identified as **kXR_readrdok**.

capver

is the client's capabilities combined with the binary protocol version number of the client. The capabilities reside in the top-most two bits while the protocol version number is encoded in the lower 6 bits. Currently, for capabilities two values are possible:

- 0b00***vvvvvvv* - client only supports synchronous responses
- 0b10***vvvvvvv* - (**kXR_asyncap**) client supports asynchronous responses

role is the role being assumed for this login. Valid roles are:

- kXR_useradmin** 0x01 - login as an administrator
- kXR_useruser** 0x00 - login as a regular user (the default)

tlen is the binary length of the supplied token, *token*. If no *token* is present, *tlen* is zero.

token is the token supplied by the previous redirection response that has initiated this login request plus other optional elements.

slen is the binary length of the information, *sec*, that follows *slen*.

sessid is the opaque session identifier associated with this login. The *sessid* is always present when the server protocol version is greater than or equal to 2.4.0 and the client protocol version is greater than 0.

sec is the null-terminated security information. The information should be treated as opaque and is meant to be used as input to the security protocol creation routine **XrdSecGetProtocol()**.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) If no security information is returned (i.e., *slen* is zero), the **XRootD** server does not require that the client authenticate.
- 3) If security information is returned, then the client must create the security context allowed by the security information, obtain credentials, and send them using an **kXR_auth** request.
- 4) Authentication must occur prior to any operation that requires authentication. See the table on page 13 for a list of requests that must be authenticated.
- 5) Logging in as an administrator suppresses any redirection attempts and limits the request set to **kXR_auth** and **kXR_admin**.
- 6) A subsequent **kXR_auth** request may revert the login into a normal user login should **XRootD** find that the authenticated user cannot assume the role of administrator.
- 7) Logging in as a normal user prohibits the use of the **kXR_admin** request.
- 8) Sending a **kXR_login** request on a previously authenticated connection destroys the authentication context; requiring that the connection be re-authenticated.
- 9) The *sessid* is used in **kXR_bind** and **kXR_endsess** requests,
- 10) Opaque information must be treated as truly opaque. The client must not inspect nor modify opaque information in any way.

kXR_login

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_login		3007	Perform server login
	<i>ability</i>		
	kXR_fullurl	0x01	Accepts full URL redirect
	kXR_hasipv64	0x08	IPv4 and IPv6 capable
	kXR_multiplr	0x03	Accepts non-root protocol redirects
	kXR_nothing	0x00	No special abilities
	kXR_onlyprv4	0x10	Only accepts private IPv4 addresses
	kXR_onlyprv6	0x20	Only accepts private IPv6 addresses
	<i>capver</i>		
	kXR_asynccap	0x80	Supports asynchronous responses
	kXR_vermask	0x3f	Mask to isolate kXR_vermm
	kXR_ver000	0x00	Predates 2005 protocol
	kXR_ver001	0x01	Implements original 2005 protocol
	kXR_ver002	0x02	Implements above + async responses
	kXR_ver003	0x03	Implements above + 2011 extensions
	kXR_ver004	0x04	Implements above + request signing
	<i>role</i>		
	kXR_useradmin	0x01	Login in administrator mode

4.12.1 Additional Login CGI Tokens

The following table lists additional CGI tokens that may be passed to further identify the client. They are passed in the *token* argument.

Token	Token Value
xrd.cc	the two character country code of the client's location
xrd.if	the client's interface speed in gigabits <i>gggg[.mm]</i>
xrd.ll	the comma separated latitude and longitude of the client in degree [-]DDD[.dddddd] format
xrd.tz	signed timezone relative to UDT of client's location

4.13 kXR_mkdir Request

Purpose: Create a directory.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_mkdir	kXR_uint16 kXR_ok
kXR_char <i>options</i>	kXR_int32 0
kXR_char <i>reserved</i> [13]	
kXR_uint16 <i>mode</i>	
kXR_int32 <i>plen</i>	
kXR_char <i>path</i> [<i>plen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

options

are the options to apply when *path* is created. The *options* are an "or'd" combination of the following values:

kXR_mkdirpath - create directory path if it does not already exist

mode is the access mode to be set for *path*. The access mode is an "or'd" combination of the following values:

Access	Readable	Writeable	Searchable
Owner	kXR_ur	kXR_uw	kXR_ux
Group	kXR_gr	kXR_gw	kXR_gx
Other	kXR_or	<i>not supported</i>	kXR_ox

kXR_mkdir

plen is the binary length of the supplied path, *path*.

path is the path of the of the directory to be created. The *path* may be suffixed with CGI information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) When a directory path is created, as requested by the **kXR_mkdirpath** option, the directory permission specified in *mode* are propagated along the newly created path.
- 3) No **umask** applies to the specified mode.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_mkdir		3008	Create a directory
	<i>mode</i>		
	kXR_ur	0x01 00	Owner readable
	kXR_uw	0x00 80	Owner writable
	kXR_ux	0x00 40	Owner searchable (directories)
	kXR_gr	0x00 20	Group readable
	kXR_gw	0x00 10	Group writable
	kXR_gx	0x00 08	Group searchable (directories)
	kXR_or	0x00 04	Other readable
	kXR_ow	0x00 02	Other writable (not allowed)
	kXR_ox	0x00 01	Other searchable (directories)
	<i>options</i>		
	kXR_mkdirpath	0x01	Create missing directories in <i>path</i>

kXR_mv

4.14 kXR_mv Request

Purpose: Rename a directory or file.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_mv	kXR_unt16 kXR_ok
kXR_char <i>reserved</i> [14]	kXR_int32 0
kXR_unt16 <i>arg1len</i>	
kXR_int32 <i>plen</i>	
kXR_char <i>path</i> [<i>plen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

arg1len

the length of the first component in paths. If *arg1len* is zero, then *paths* is scanned for spaces to delimit the components. See the notes for more information.

plen is the binary length of the supplied old and new paths, *paths*.

path is the old name of the path (i.e., the path to be renamed) followed by a space and then the name that the path is to have. Each *path* string may be suffixed with CGI information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Renames across file systems are not supported.

kXR_open

- 3) Protocol version 3.1.0 introduced *arg1len* in order to specify the actual length of the first component to allow paths to have embedded spaces. When *arg1len* is non-zero then the *paths+arg1len* must point to a space character. All characters before *paths+arg1len* are used as the old name and all characters after *paths+arg1len+1* is taken as the new name.
- 4) When *arg1len* is zero (pre-3.1.0 behaviour), then *paths* is scanned for the first space character and this becomes the breakpoint between the old name and the new name.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_mv		3009	Rename directory or file

4.15 kXR_open Request

Purpose: Open a file or a communications path.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_open	kXR_uint16 kXR_ok
kXR_uint16 <i>mode</i>	kXR_int32 <i>rlen</i>
kXR_uint16 <i>options</i>	kXR_char <i>fhandle</i> [4]
kXR_char <i>reserved</i> [12]	optional addition
kXR_int32 <i>plen</i>	kXR_int32 <i>cpsize</i>
kXR_char <i>path</i> [<i>plen</i>]	kXR_char <i>cptype</i> [4]
	kXR_char <i>info</i> [<i>resplen</i> -12]

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

mode is the advisory mode in which *path* is to be opened. The *mode* is an “or’d” combination of the following values:

Access	Readable	Writeable	Executable
Owner	kXR ur	kXR uw	kXR ux
Group	kXR gr	kXR gw	kXR gx
Other	kXR_or	<i>not supported</i>	kXR_ox

options

are the options to apply when *path* is opened. The *options* are an “or’d” combination of the following values:

- kXR_async** - open the file for asynchronous i/o (see notes)
- kXR_compress** - open a file even when compressed (see notes)
- kXR_delete** - open a new file, deleting any existing file
- kXR_force** - ignore file usage rules
- kXR_mkpath** - create directory path if it does not already exist
- kXR_new** - open a new file only if it does not already exist
- kXR_open_apnd** - open only for appending
- kXR_open_read** - open only for reading
- kXR_open_updt** - open for reading and writing
- kXR_posc** - enable **Persist On Successful Close (POSC)** processing

kXR_open

- kXR_refresh** - update cached information on the file's location (see notes)
- kXR_replica** - the file is being opened for replica creation
- kXR_retstat** - return file status information in the response
- kXR_seqio** - file will be read or written sequentially (see notes)

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

plen is the binary length of the supplied path, *path*.

path is the path of the file to be opened. The *path* may be suffixed with CGI information to provide additional information necessary to properly process the request. See the following section on CGI information for more information.

resplen

is the byte length of the response that follows. At least four bytes should be returned.

fhandle

is the file handle for the associated file. The file handle should be treated as opaque data. It must be used for subsequent **kXR_close**, **kXR_read**, **kXR_sync**, and **kXR_write** requests.

cpsize is the compression page size. The *cpsize* field is returned when the **kXR_compress** or **kXR_retstat** have been specified. Subsequent reads must be equal to this value and read offsets must be an integral multiple of this value. If *cpsize* is zero, the file is not compressed and subsequent reads may use any offset and read length.

cptype is the name of the compression algorithm used to compress the file (e.g. lz4). The *cptype* field is returned when the **kXR_compress** or **kXR_retstat** have been specified. If the file is not compressed, the first byte of the four byte field is a null byte (\0). For compressed files, subsequent reads must use the returned algorithm to decompress each *cpsize* worth of data data.

info is the same information that **kXR_stat** returns for the file. This information is returned only if **kXR_retstat** is set and the server is at protocol version 2.4.0 or greater. The *cpsize* and *cptype* fields are always returned and are only meaningful if **kXR_compress** has been specified. Otherwise, *cpsize* and *cptype* are set to values indicating that the file is not compressed.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Open fails if the path designates a directory.
- 3) No **umask** applies to the specified mode.
- 4) The **kXR_async** option tells the server to overlap file i/o with network requests as much as possible for this file. For instance, read requests may be done in parallel with other read requests sent on the same link. This option is only useful if the client is able to issue multiple requests (i.e., is not serializing the requests-response stream).
- 5) While the **kXR_async** option applies to write operations, as well. Server-side asynchronous opportunities are far more limited. The client needs to perform appropriate multiplexing of write requests with other requests to gain improved parallelism.
- 6) The **kXR_async** option imposes additional overhead on the server and should only be specified when the client can take advantage of request-response parallelism.
- 7) The **kXR_refresh** option imposes additional overhead on the server because it requires that the server obtain the most current information on the file's location before attempting to process the open request. This option should only be used as part of the error recovery process outlined in section "Client Recovery From File Location Failures".
- 8) The **kXR_refresh** option is ignored by any server not functioning as a primary redirecting server.
- 9) When a directory path is created, as requested by the **kXR_mkpath** option, the directory permission of 0775 (i.e., rwxrwxr-x) are propagated along the newly created path.
- 10) Only files may be opened using the **kXR_open** request code.
- 11) The **kXR_retstat** option is meant to eliminate an additional server request for file status information for applications that always need such information.

kXR_open

- 12) The **kXR_seqio** option is meant to be advisory. A server may choose to optimize data layout or access based on this hint. Misusing the hint may lead to degraded performance.
- 13) The **kXR_posc** option requests safe file persistence which persists the file only when it has been explicitly closed.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_open		3010	Open a file
	<i>mode</i>		
	kXR_ur	0x01 00	Owner readable
	kXR_uw	0x00 80	Owner writable
	kXR_ux	0x00 40	Owner searchable (directories)
	kXR_gr	0x00 20	Group readable
	kXR_gw	0x00 10	Group writable
	kXR_gx	0x00 08	Group searchable (directories)
	kXR_or	0x00 04	Other readable
	kXR_ow	0x00 02	Other writable
	kXR_ox	0x00 01	Other searchable (directories)
	<i>options</i>		
	kXR_async	0x00 40	Allow asynchronous I/O
	kXR_compress	0x00 01	Open without inflating files
	kXR_delete	0x00 02	Delete any existing file
	kXR_force	0x00 04	Disregard locking rules
	kXR_mkpath	0x01 00	Create any missing directories
	kXR_new	0x00 08	Create a new file
	kXR_open_apnd	0x02 00	Open only for appending
	kXR_open_read	0x00 10	Open only for reading
	kXR_open_updt	0x00 20	Open for reading and writing
	kXR_open_wrto	0x80 00	Open only for writing
	kXR_posc	0x10 00	Persist on successful close
	kXR_refresh	0x00 80	Refresh cached information
	kXR_replica	0x08 00	Open for replication
	kXR_retstat	0x04 00	Return file stat information
	kXR_seqio	0x40 00	Open for sequential I/O

4.15.1 Additional Open CGI Tokens

The **kXR_open** request allows a client to pass CGI information to properly steer the open. The information may or may not be acted upon, depending on the server's capabilities. The following table lists the defined CGI tokens.

Token	Token Value
<code>ofs.posc</code>	When set to a value of 1 requests "persist on successful close" processing. This is historical as the kXR_posc option should be preferentially used.
<code>oss.asize</code>	The mber of bytes to reserve for a new file.
<code>Oss.cgroup</code>	The desired space name (a.k.a space token).

Notes

- 1) Unrecognized CGI tokens should be ignored.
- 2) Invalid arguments to a recognized CGI token should result in the termination of the request.

Example

```
/tmp/foo&oss.cgroup=index
```


4.16 kXR_ping Request

Purpose: Determine if the server is alive.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_ping	kXR_uint16 kXR_ok
kXR_char <i>reserved</i> [16]	kXR_int32 0
kXR_int32 0	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_uint16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Use the **kXR_ping** request to see if the server is running.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_ping		3011	Send keep alive

kXR_prepare

4.17 kXR_prepare Request

Purpose: Prepare one or more files for access.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_prepare	kXR_uint16 kXR_ok
kXR_char <i>options</i>	kXR_int32 <i>rlen</i>
kXR_char <i>prty</i>	kXR_char <i>resp</i> [<i>rlen</i>]
kXR_uint16 <i>port</i>	
kXR_char <i>reserved</i> [12]	
kXR_int32 <i>plen</i>	
kXR_char <i>plist</i> [<i>plen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

options

are the options to apply to each *path*. The notes explain how these options can be used. The *options* are an “or’d” combination of the following:

- kXR_cancel** - cancel a prepare request
- kXR_coloc** - co-locate staged files, if at all possible
- kXR_fresh** - refresh file access time even when location is known
- kXR_noerrs** - do not send notification of preparation errors
- kXR_notify** - send a message when the file has been processed
- kXR_stage** - stage the file to disk if it is not online
- kXR_wmode** - the file will be accessed for modification

prty is the binary priority the request is to have. Specify a value between 0 (the lowest) and 3 (the highest), inclusive.

port is the binary udp port number in network byte order to which a message is to be sent, as controlled by **kXR_notify** and **kXR_noerrs**. If port is zero and **kXR_notify** is set, notifications are sent via asynchronous messages via the connected server, if possible.

reserved

is an area reserved for future use and must be initialized to null (i.e., '\0').

plen is the binary length of the supplied path list, *plist*.

plist is the list of new-line separated paths that are to be prepared for access. Each *path* may be suffixed with CGI information. If only one path is supplied, it need not be terminated with a new line character (\n). If **kXR_cancel** is specified, then *plist* must be a prepare *locatorid*.

rlen is the binary length of the response, *resp*, that follows *rlen*.

resp is the response to request.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The **kXR_prepare** request attempts to make the indicated files available for access. This may require that the files be brought in from a Mass Storage device.
- 3) The **kXR_prepare** request always executes asynchronously. Therefore, unless there are obvious errors in the request, a successful status code is immediately returned.
- 4) The system makes no guarantees that the files will be made available for access ahead of a future **kXR_open** request. Hence, the **kXR_prepare** request is treated as merely a hint.
- 5) The **kXR_prepare** request should normally be directed to a load-balancing server should one be present.
- 6) The when the prepare request has been accepted in the presence of the **kXR_stage** option, the server returns a request locator (i.e., *locatorid*) as the normal response. This *locatorid* should be treated as an opaque ASCII text string. The *locatorid* can be used to cancel the request at some future time and to pair up asynchronous messages with requests when **kXR_notify** has been set.
- 7) **kXR_coloc** is only meaningful in the presence of **kXR_stage** when more than one file has been specified.
- 8) Co-location of files is not guaranteed. When the **kXR_coloc** and **kXR_stage** options are set, an attempt will be made to co-locate all mentioned files in the request with the first file in the list of files.

kXR_prepare

- 9) Co-location may fail for many reasons, including but not limited to, files already present at different locations, files present in multiple locations, and insufficient space. The success if co-locations is implementation defined.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_prepare		3021	Prepare for future file access
	<i>options</i>		
	kXR_cancel	0x01	Cancel previous prepare request
	kXR_notify	0x02	Send stage completion notification
	kXR_noerrs	0x04	Suppress error notifications
	kXR_stage	0x08	Stage in missing files
	kXR_wmode	0x10	Prepare for writing
	kXR_coloc	0x20	Colocate all specified files
	kXR_fresh	0x40	Update file access time

4.18 kXR_protocol Request

Purpose: Obtain the protocol version number, type of server, and possible security requirements.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_protocol	kXR_uint16 kXR_ok
kXR_int32 <i>clientpv</i>	kXR_int32 <i>dlen</i>
kXR_char <i>options</i>	kXR_int32 <i>pval</i>
kXR_char <i>expect</i>	kXR_int32 <i>flags</i>
kXR_char <i>reserved</i> [10]	Security Requirements
kXR_int32 0	kXR_char 'S'
	kXR_char <i>rsvd</i>
	kXR_char <i>secver</i>
	kXR_char <i>secopt</i>
	kXR_char <i>seclvl</i>
	kXR_char <i>secvsz</i>
	Optional Overrides
	kXR_char { <i>reqidx</i>
	kXR_char <i>reqlvl</i> } [<i>secvsz</i>]
<i>dlen</i> : 8 or 14 + <i>secvsz</i> *2	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

clientpv

the binary protocol version that the client is using. See the usage notes on how to obtain the correct value. The *clientpv* field is recognized only in protocol version 2.9.7 and above.

expect specifies what the client intends to do next and is optional. The presence of an *expect* value allows a server to unilaterally transition the connection to TLS if the subsequent intended operation requires TLS. This avoids an additional roundtrip between the client and server. The expectations are:

kXR_ExpNone no intentions.

kXR_ExpBind a kXR_bind request should be expected.

kXR_ExpLogin a kXR_login should be expected.

kXR_ExpTPC a third party copy should be expected.

kXR_protocol

expect

specifies what the client intends to do next and is optional. The presence of an *expect* value allows a server to unilaterally transition the connection to TLS if the subsequent intended operation requires TLS. This avoids an additional roundtrip between the client and server. The expectations are encoded bits and the server should use the **kXR_ExpMask** to isolate the bits. An invalid expectation setting should be treated as **kXR_ExpNone** (i.e. no expectation). Valid values are:

kXR_ExpNone no intentions.

kXR_ExpBind a kXR_bind request should be expected.

kXR_ExpLogin a kXR_login should be expected.

kXR_ExpTPC a third party copy should be expected.

options

specifies what should be returned. Without any options only the pval and flags should be returned. This is also the case if the server does not support the return option or if no meaningful data exists for the specific request. The options are:

kXR_ableTLS client is TLS capable.

kXR_secreqs return protocol security signing requirements.

kXR_wantTLS client wants to transition the connection to TLS.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

pval is the binary protocol version number the server is using.

flags is additional bit-encoded information about the server. The following flags are returned when *clientpv* is zero (i.e. not specified) or the server's protocol version is 2.9.6 or lower:

- kXR_DataServer** - This is a data server.
- KXR_LBalServer** - This is a load-balancing server.

The following flags are returned when *clientpv* is not zero (i.e. is specified) and the server's protocol version is 2.9.7 or above:

- kXR_isManager** - Has manager role.
- kXR_isServer** - Has server role.
- kXR_attrMeta** - Has the meta attribute (e.g. meta manager).
- kXR_attrProxy** - Has the proxy attribute (e.g. proxy server).
- kXR_attrSuper** - Has the supervisor attribute.

- kXR_anongpf** - Supports anonymous getfile and putfile.
- kXR_supgetf** - Supports **kXR_getfile** request.
- kXR_suppetf** - Supports **kXR_putfile** request.

- kXR_haveTLS** - Supports TLS connections.
- kXR_gotoTLS** - Server has transitioned connection to use TLS.
- kXR_tlsData** - Data must be sent over a TLS connection.
- kXR_tlsLogin** - Login must use a TLS connection.
- kXR_tlsSess** - Connection transitions to TLS after login.
- kXR_tlsTPC** - Third party copy must use a TLS connection.

kXR_protocol

Security Requirements

If the server supports **kXR_secreqs** and the information is meaningful, at least 6 additional bytes are returned:

rsvd a reserved byte that should be set to zero.

secver the controlling security version. Currently, only version 0 is defined so the byte should be set to zero.

secopt security options:

kXR_secOFrce apply signing requirements even if the authentication protocol does not support generic encryption.

seclvl the default security level to be used. The next section defines each of 5 predefined security levels.

secosz the number of security override doublets that follow. Security overrides allow a server to customize the predefined security level specified in *seclvl*. If there are no security overrides, this byte should be set to zero.

Security Overrides

A server may customize any predefined security level by returning alterations needed to the specified predefined security level. The information is contained in a vector of doublets of size *secosz*:

reqidx the request whose security requirements are to be changed. The request code is specified as a request index. Specifically, it is the **kXR** request code minus **kXR_auth** (the lowest numbered request code). Security requirements are explained in the following section.

reqlvl the security requirement that the associated request is to have:

kXR_signIgnore the request need not be signed.

kXR_signLikely a signing requirement is likely and depends on the request's context. If the request modifies data it should be interpreted as **kXR_signNeeded**. Otherwise, it should be interpreted as **kXR_signNone**.

kXR_signNeeded the request must be signed.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The client should not rely on the response data length being 8. In the future, additional information may be returned.
- 3) The protocol version is defined by **kXR_PROTOCOLVERSION** in the header file that defines protocol values and data structures.
- 4) When the client specifies its protocol version in *clientpv*, the server may use that information to tailor responses to be compatible with the stated version. Since any number of **kXR_protocol** requests can be issued, the authoritative protocol version is considered to be the one in effect after the **kXR_login** request succeeds. After that time, the client's protocol version is immutable until the next login.
- 5) For **kXR_bind** requests, the client's protocol version is forced to be the same as that the base login stream to which the bind request refers.
- 6) When testing the bits in *flags* in the protocol response when *clientpv* is specified, the following order should be used:
 - a. **kXR_isManager** -> role manager
 kXR_attrMeta -> role meta manager
 kXR_attrProxy -> role proxy manager
 kXR_attrSuper -> role supervisor
 - b. **kXR_isServer** -> role server
 kXR_attrProxy -> role proxy server
 - c. If none of the above, treat as role manager.
- 7) The protocol specifies that a client must affiliate with the *first* manager or the *last* meta-manager encountered. Client retry requests should be sent to the affiliated [meta] manager established during the connection phase.
- 8) Protocol version 2.9.7 provides for a mechanism to determine whether a connection target is a manager or a meta-manager. Clients using lower versions of the protocol do not have that capability and consequently treat managers and meta-managers identically. While this does not cause functional problems, it markedly reduces efficiency when retrying requests in the presence of multiple meta-managers that control different sets of clusters.

kXR_protocol

- 9) Protocol version 3.1.0 introduced a mechanism to verify that requests came from an authenticated client. Pre 3.1.0 servers should never return security information when requested to do so. Servers that have no security requirements need not return any security information when requested to do so. When security information has not been returned the client should assume that no requirements exist.
- 10) The **kXR_protocol** request should be used to transition a normal connection to one that uses **TLS**. The client requests such a transition by setting the **kXR_wantTLS** flag in *options*. If the option is set and the server supports **TLS**, it should transition the connection to use **TLS** after the response is sent. The server may also unilaterally transition the connection to use **TLS** after sending the response based on the expect flags which indicate what the client intends to do next. When a connection will transition to use **TLS** immediately after the response is sent, the **kXR_gotoTLS** should be set in the response to indicate the client must also transition to **TLS**. The flag should not be set if the connection will not transition to **TLS** mode (e.g. the connection is already in **TLS** mode).
- 11) Clients should indicate whether or not they can use **TLS** on the connection by setting the **kXR_ableTLS** option. Should only the **kXR_wantTLS** option be set, servers should assume the **kXR_ableTLS** option is also in effect.
- 12) If the server does not support **TLS**, none of the **TLS** related flags should be set in *flags*. Otherwise, settings are left up to the implementation, except as noted below.
- 13) The **kXR_tlsLogin** and **kXR_tlsSess** flags are mutually exclusive.
- 14) When **kXR_tlsData** is set, transmission of any file or attribute data should be sent using a **TLS** connection. This applies to the session connection as well as any connections established using **kXR_bind**. This implies that when **kXR_tlsData** is set either **kXR_tlsLogin** or **kXR_tlsSess** should be set as well.
- 15) The **kXR_tlsTPC** flag should cover not only **kXR_getfile** and **kXR_putfile** requests but also **TPC** operations initiated via the **kXR_open** request using the **XRootD TPC** protocol.
- 16) If the server supports **kXR_getfile** requests, the **kXR_supgetf** flag should be set. If the server supports **kXR_putfile** requests, the **kXR_supputf** flag should be set. The client should ignore the **kXR_anongpf** flag if none of the “supp” flags are set.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_protocol		3006	Ascertain server's protocol
	<i>clientpv</i>		Client's protocol version
		0x00 00 02 45	Protocol version 2.4.5 (2005)
		0x00 00 02 96	Protocol version 2.9.6 (2010)
		0x00 00 02 99	Protocol version 2.9.9 (2011)
		0x00 00 03 00	Protocol version 3.0.0 (2013)
		0x00 00 03 10	Protocol version 3.1.0 (2016)
		0x00 00 04 00	Protocol version 4.0.0 (2018)
	<i>expect</i>		
	kXR_expMask	0x0f	Mask to isolate <i>expect</i> encoding
	kXR_ExpNone	0x00	No expectations
	kXR_ExpBind	0x01	Expect a kXR_bind next
	kXR_ExpLogin	0x02	Expect a kXR_login next
	kXR_ExpTPC	0x03	Expect a third party copy request next
	<i>options</i>		
	kXR_ableTLS	0x02	Client is TLS capable
	kXR_secreqs	0x01	Return security requirements
	kXR_wantTLS	0x04	Client wants to transition to TLS
Response	<i>flags</i>		
	kXR_DataServer	0x00 00 00 01	Node is a data server
	kXR_LBalServer	0x00 00 00 00	Node is not a data server
	kXR_isManager	0x00 00 00 02	Node has a manager role
	kXR_isServer	0x00 00 00 01	Node has a server role
	kXR_attrMeta	0x00 00 01 00	Node has a meta role attribute
	kXR_attrProxy	0x00 00 02 00	Node has a proxy role attribute
	kXR_attrSuper	0x00 00 04 00	Node has a supervisor role attribute
	kXR_anongpf	0x00 80 00 00	Supports anonymous get/put file
	kXR_supgetf	0x00 40 00 00	Supports kXR_getfile
	kXR_supputf	0x00 20 00 00	Supports kXR_putfile
	kXR_haveTLS	0x80 00 00 00	Supports TLS connections
	kXR_gotoTLS	0x40 00 00 00	Connection will transition to TLS mode
	kXR_tlsAny	0x0f 00 00 00	Mask to isolate requirement flags
	kXR_tlsData	0x01 00 00 00	All data requires a TLS connection
	kXR_tlsLogin	0x02 00 00 00	kXR_login requires a TLS connection
	kXR_tlsSess	0x04 00 00 00	Connection transition to TLS after login
	kXR_tlsTPC	0x08 00 00 00	TPC requests require a TLS connection
	<i>pval</i>	see <i>clientpv</i>	Server's protocol version
	<i>reqlvl</i>		
	kXR_signIgnore	0x00	Signature is not needed
	kXR_signLikely	0x01	Signature needed when modifying
	kXR_signNeeded	0x02	Signature need in all cases

kXR_protocol

Request	Modifiers	Value	Explanation
	<i>seclvl</i>		
	kXR_secCompatible	0x01	
	kXR_secStandard	0x02	
	kXR_secIntense	0x03	
	kXR_secPedantic	0x04	
	<i>secopt</i>		
	kXR_secOData	0x01	Write data must be signed
	kXR_secOFRCE	0x02	Sign requests even if unencrypted

4.18.1 Protocol Security Requirements vs Response Implications

The **xroot** protocol provides capabilities to verify that a request came from the previously authenticated client. The verification consists of prefixing a request with a **kXR_sigver** request that contains the cryptographic signature of the subsequent request to be verified. The specification of request signature and verification is explained in the **kXR_sigver** section. The **kXR_protocol** request allows a client to determine which requests need to be signed. The table below shows the signing requirements by request for each predefined security level.

Request	Compatible	Standard	Intense	Pedantic
kXR_admin	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded
kXR_auth	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore
kXR_bind	kXR_signIgnore	kXR_signIgnore	kXR_signNeeded	kXR_signNeeded
kXR_chmod	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded
kXR_close	kXR_signIgnore	kXR_signIgnore	kXR_signNeeded	kXR_signNeeded
kXR_decrypt	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore
kXR_dirstat	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore	kXR_signNeeded
kXR_endsess	kXR_signIgnore	kXR_signIgnore	kXR_signNeeded	kXR_signNeeded
kXR_fattr	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded
kXR_getfile	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded
kXR_locate	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore	kXR_signNeeded
kXR_login	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore
kXR_mkdir	kXR_signIgnore	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded
kXR_mv	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded
kXR_open	kXR_signLikely	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded
kXR_ping	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore
kXR_prepare	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore	kXR_signNeeded
kXR_protocol	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore
kXR_putfile	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded
kXR_query	kXR_signIgnore	kXR_signIgnore	kXR_signLikely	kXR_signNeeded
kXR_read	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore	kXR_signNeeded
kXR_readv	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore	kXR_signNeeded
kXR_rm	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded
kXR_rmdir	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded
kXR_set	kXR_signLikely	kXR_signLikely	kXR_signNeeded	kXR_signNeeded
kXR_sigver	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore
kXR_stat	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore	kXR_signNeeded
kXR_statx	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore	kXR_signNeeded
kXR_sync	kXR_signIgnore	kXR_signIgnore	kXR_signIgnore	kXR_signNeeded
kXR_truncate	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded	kXR_signNeeded
kXR_verifyw	kXR_signIgnore	kXR_signIgnore	kXR_signNeeded	kXR_signNeeded
kXR_write	kXR_signIgnore	kXR_signIgnore	kXR_signNeeded	kXR_signNeeded

kXR_protocol

A server uses **kXR_protocol** request to specify the security level in effect and any specific overrides. Hence, the protocol provides a framework for, not an absolute definition of, security requirements.

Predefined security levels simplify handling of security requirements. The protocol pre-defines 5 security levels that can be specified in *seclvl*:

kXR_secNone	No security requirements exist.
kXR_secCompatible	A security requirement exists only for potentially destructive requests. (i.e. ones that modify data or metadata).
kXR_secStandard	A security requirement exists for potentially destructive requests. (i.e. ones that modify data or metadata) as well as certain non-destructive requests.
kXR_secIntense	A security requirement exists only for pq wide range of requests that may reveal metadata or modify data.
kXR_secPedantic	Security requirements apply to all requests.

For each request, one of three scenarios exist at each security level:

kXR_signIgnore	The request need not be signed.
kXR_signLikely	The request needs to be signed if it may modify data or metadata.).
kXR_signNeeded	The request must be signed.

The **kXR_signLikely** is the most problematic because it needs to be interpreted the context of what the request is actually doing. Only three requests need to be examined more deeply to determine whether or not they need to be signed.

kXR_open	must be signed if any of the options: kXR_delete , kXR_new , kXR_open_updt , kXR_mkath , and kXR_open_apnd has been specified.
kXR_query	must be signed if any of the options: kXR_Qopque , kXR_qopaquf , and kXR_Qopaqug have been specified.
kXR_set	must be signed if any request options (i.e. a non-default set operation) have been specified.

4.19 kXR_putfile Request

Purpose: Store a complete file (placeholder for next generation TPC).

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_putfile	kXR_unt16 kXR_ok
kXR_char <i>options</i>	kXR_int32 0
kXR_char <i>reserved</i> [12]	
kXR_int32 <i>plen</i>	
kXR_char <i>plist</i> [<i>plen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

options

request options.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0')..

plen is the binary length of the supplied path, *path*.

path is the path of the file to be stored. The *path* may be suffixed with CGI information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) This request type is not currently supported.
- 3) The **kXR_putfile** request should be considered as not fully specified. It is currently a place-holder for future enhancement and may substantially change in functionality.

kXR_putfile

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_putfile		3012	n/a

4.20 kXR_query Request

Purpose: Obtain server information.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_query	kXR_uint16 kXR_ok
kXR_uint16 <i>reqcode</i>	kXR_int32 <i>ilen</i>
kXR_char <i>reserved</i> [2]	kXR_char <i>info</i> [<i>ilen</i>]
kXR_char <i>fhandle</i> [4]	
kXR_char <i>reserved</i> [8]	
kXR_int32 <i>alen</i>	
kXR_char <i>args</i> [<i>alen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

fhandle

is the file handle value supplied by the successful response to the associated **kXR_open** request. Only kXR_Qvisa supports *fhandle*.

reqcode

is the binary code indicating the specific query being made. Valid codes are:

kXR_Qconfig	Query server configuration
kXR_Qckscan	Query file checksum cancellation
kXR_Qcksum	Query file checksum
kXR_Qopaque	Query implementation-dependent information
kXR_Qopaquef	Query implementation-dependent information
kXR_Qopaqueg	Query implementation-dependent information
kXR_QPrep	Query prepare status
kXR_Qspace	Query server logical space statistics
kXR_Qstats	Query server statistics
kXR_Qvisa	Query file visa attributes

kXR_query

kXR_Qxattr Query file extended attributes

alen is the binary length of the supplied arguments, *args*.

args is the arguments to the query, specific to the *reqcode*.

ilen is the binary length of the information, *info*, that follows *ilen*.

info is the requested information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Responses to **kXR_Qspace** and **kXR_Qxattr** requests are documented in the Open File System (**ofs**) and Open Storage System (**oss**) configuration reference. Responses to **kXR_Qopaque** and **kXR_Qopaquf** are implementation dependent. This query type should not be used for portable programs.
- 3) Unstructured data may be passed using the **kXR_Qopaque**. The **kXR_Qopaquf** *reqcode* is meant for structured arguments (i.e., valid path and CGI information).
- 4) The **kXR_waitresp** response is not an error response but merely indicates that the response may take approximately seconds of time to deliver and should be reported using the unsolicited response mechanism (i.e., **kXR_attn** with **kXR_asynresp**). Refer to the description of each server response for detailed handling information.
- 5) A delayed response appears in protocol version 2.5.0 or higher. Earlier protocol versions did not use the delayed response mechanism.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_query		3001	Return information
	<i>reqcode</i>		
	kXR_QStats	0x00 01	Return statistics
	kXR_QPrep	0x00 02	Return prepare status
	kXR_Qcksum	0x00 03	Return checksum
	kXR_qxattr	0x00 04	Return extended space attributes
	kXR_qspace	0x00 05	Return space information
	kXR_qckscan	0x00 06	Return checksum cancellation info
	kXR_Qconfig	0x00 07	Return configuration information
	kXR_Qvisa	0x00 08	Return visa status
	kXR_Qopaque	0x00 10	Return implementation information #1
	kXR_Qopaquf	0x00 20	Return implementation information #2
	kXR_Qopaqug	0x00 40	Return implementation information #3

4.20.1 KXR_query Checksum Cancellation Request

Purpose: Obtain server information.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_query	kXR_unt16 kXR_ok
kXR_unt16 kXR_Qckscan	kXR_int32 0
kXR_char <i>reserved</i> [14]	
kXR_int32 <i>plen</i>	
kXR_char <i>path</i> [<i>plen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

U8

plen is the binary length of the supplied path, *path*.

path is the path of the file whose check sum is to be cancelled. The *path* may be suffixed with CGI information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Only check sums requested by the current client may be cancelled.

4.20.2 KXR_query Checksum Request

Purpose: Obtain server information.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_query	kXR_unt16 kXR_ok
kXR_unt16 kXR_Qcksum	kXR_int32 <i>ilen</i>
kXR_char <i>reserved</i> [14]	kXR_char <i>info</i> [<i>ilen</i>]
kXR_int32 <i>plen</i>	
kXR_char <i>path</i> [<i>plen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

plen is the binary length of the supplied path, *path*.

path is the path of the file whose checksum is to be returned. The *path* may be suffixed with CGI information.

ilen is the binary length of the information, *info*, that follows *ilen*.

info is the requested information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Each installation determines the type of checksum that will be returned. The algorithmic name of the checksum precedes the checksum value.

kXR_Qcksum

Returned Response

The general format for the **kXR_Qcksum** response is:

```
csname csvalue
```

Where:

csname

is the algorithmic name of the checksum algorithm used. This name is selected by the administrator.

csvalue

is the checksum name as a hexadecimal ASCII text string. The format is dependent on the algorithm used to compute the checksum.

4.20.2.1 Additional Query Checksum CGI Tokens

The **kXR_Qcksum** request allows a client to pass CGI information to select a particular checksum should the server support multiple checksums. The information may or may not be acted upon, depending on the server's capabilities. CGI information is passed by suffixing the *path* with a question mark (?) and then coding the cgi information as shown below:

```
path?cks.cktype=arg
```

Where:

arg is the name of the desired checksum.

Notes

- 1) Invalid values or arguments to a recognized directive normally result in termination of the request.

Example

```
/tmp/foo?cks.cktype=md5
```

4.20.3 KXR_query Configuration Request

Purpose: Obtain server information.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_query	kXR_uint16 kXR_ok
kXR_uint16 kXR_Qconfig	kXR_int32 <i>ilen</i>
kXR_char <i>reserved</i> [14]	kXR_char <i>info</i> [<i>ilen</i>]
kXR_int32 <i>qlen</i>	
kXR_char <i>qry</i> [<i>qlen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

qlen is the binary length of the supplied query arguments, *qry*.

qry are the space-separated names of the variables to be returned. Current variables that may be queried are:

bind_max maximum number of sockets that may be bound to login session.

chksum checksum algorithm name supported by the server as "*n:name*" where *n* is the algorithm numeric id and *name* is it's name. If more than one algorithm is supported, they are listed, each separated by a comma.

cid the globally unique cluster identification string.

cms the current dynamic state of the cluster management service configuration. See the next section for the format.

kXR_Qconfig

fattr	get kXR_fattr request limits. The response, if supported, should contain three tokens: " <i>maxattrs maxnlen maxvlen</i> " <i>maxattrs</i> - maximum number of bundled requests <i>maxnlen</i> - maximum length for a name <i>maxvlen</i> - maximum length for attribute data
pio_max	maximum number of requests that may be queued on a bound socket before the session stream must wait.
readv_ior_max	maximum amount of data that may be requested in a single kXR_readv request element.
readv_iov_max	maximum number of elements in a kXR_readv request vector.
role	the configured role. If no role has been configured "none" is returned
sitename	the site name associated with the server. If no sitename is associated, returned value is the token 'sitename'.
ssl_port	the port number to connect to for an ssl encrypted tunnel.
tpc	version number for third party copy protocol. If third party copy protocol is not supported, "tpc" is returned. Otherwise, an integer value is returned.
version	version identification string (implementation dependent).
window	socket buffer size (i.e., window) for the default port.
xattrs	user settable extended attribute limits.

ilen is the binary length of the information, *info*, that follows *ilen*.

info is the requested information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Clients should avoid issuing configuration query requests to a redirector as this may not reflect the actual limits imposed by a server. Instead, configuration requests should be obtained for each server.

Returned Response

The general format for the **kXR_Qconfig** response consists of a list of new-line delimited value in 1-to-1 correspondence to the list of supplied variable:

```
Cvalue\n[Cvalue\n[. . .\n]
```

Where:

Cvalue

is the corresponding value associated with the queried variable. If the variable has no value then the name of the variable is returned as *Cvalue*.

4.20.3.1 Format for Query Config cms

The general format for the **kXR_Qconfig cms** response consists of a space delimited list of host for which outbound connections have been configured and their current state. It should be interpreted in the context of the configured role (ie. **kXR_query** role):

```
state: c | d | s          host:port/state
```

Where:

host:port

is the hostname or IP address of the endpoint to which a connection is to be made. For server roles the *host:port* should always be indicated as localhost:0 to indicate that the clustering service is local.

status is the status of the connection:

- c** - the endpoint is connected.
- d** - the endpoint is disconnected
- s** - the endpoint is connected but is a suspended state.

kXR_Qconfig

4.20.3.2 Format for Query Config role

The general format for the **kXR_Qconfig role** response consists of a space delimited tokens describing the configured role. The role should be one of the following:

- **meta manager**
- **manager**
- **supervisor**
- **server**
- **proxy manager**
- **proxy supervisor**
- **proxy server**

4.20.3.3 Format for Query Config xattrs

The general format for the **kXR_Qconfig xattrs** response consists of two numbers as follows:

<i>maxnsz maxvsz</i>

Where:

maxnsz

is the maximum length a client can specify for an attribute name. If the value is zero, then client settable extended attributes are not allowed.

maxvsz

is the maximum length a client can specify for an attribute value.

4.20.4 KXR_query Opaque Request

Purpose: Obtain implementation-dependent server information.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_query	kXR_uint16 kXR_ok
kXR_uint16 <i>querycode</i>	kXR_int32 <i>ilen</i>
kXR_char <i>reserved</i> [2]	kXR_char <i>info</i> [<i>ilen</i>]
kXR_char <i>fhandle</i>	
kXR_char <i>reserved</i> [8]	
kXR_int32 <i>qlen</i>	
kXR_char <i>qry</i> [<i>qlen</i>]	
<i>querycode</i> : kXR_Qopaque kXR_Qopaquf kXR_Qopaqug	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

querycode

is one of the specified query codes. Each code takes different arguments:

kXR_Qopaque - *qry* is passed to the filesystem plug-in without interpretation.

kXR_Qopaquf - *qry* is interpreted as a path and optional cgi string. If the path is valid, *qry* is passed to the filesystem plug-in.

kXR_Qopaqug - *qry* is passed along without inspection to the file plug-in associated with *fhandle*.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

fhandle

is the file handle value supplied by the successful response to the associated **kXR_open** request. Only **kXR_Qopaqug** uses this field.

qlen is the binary length of the information, *qry*, that follows *qlen*.

qry is the information to be passed to the appropriate plug-in.

kXR_Qstats

ilen is the binary length of the information, *info*, that follows *ilen*.

info is the requested information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) These subcodes provide a mechanism to use special implementation specific features. Use of these subcodes is not portable.

4.20.5 KXR_query Space Request

Purpose: Obtain server information.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_query	kXR_unt16 kXR_ok
kXR_unt16 kXR_Qspace	kXR_int32 <i>ilen</i>
kXR_char <i>reserved</i> [14]	kXR_char <i>info</i> [<i>ilen</i>]
kXR_int32 <i>slen</i>	
kXR_char <i>sname</i> [<i>slen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

slen is the binary length of the supplied path, *sname*.

sname is the logical name of the space whose statistics are to be returned.

ilen is the binary length of the information, *info*, that follows *ilen*.

info is the requested information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Each installation determines the type of logical spaces that exist and the values that can be returned for them.
- 6) The response to the **kXR_Qspace** request is documented in the Open File System (**ofs**) and Open Storage System (**oss**) configuration reference.
- 3) If *sname* is empty, the name "public" is used.

kXR_Qstats

4.20.6 KXR_query Statistics Request

Purpose: Obtain server information.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_query	kXR_unt16 kXR_ok
kXR_unt16 kXR_QStats	kXR_int32 <i>ilen</i>
kXR_char <i>reserved</i> [14]	kXR_char <i>info</i> [<i>ilen</i>]
kXR_int32 <i>alen</i>	
kXR_char <i>args</i> [<i>alen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

alen is the binary length of the supplied arguments, *args*.

args is an optional list of letters, each indicating the statistical components to be returned. Valid letters are:

a - Return all statistics (default)	p - Protocol statistics
b - Buffer usage statistics	s - Scheduling statistics
d - Device polling statistics	u - Usage statistics
i - Server identification	z - Synchronized statistics
l - Connection statistics	

ilen is the binary length of the information, *info*, that follows *ilen*.

info is the requested information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_uint16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Statistical information is returned as an XML text string. The XML schema is outlined below.
- 3) By default, the server returns statistical information that is readily available. The "z" option, informs the server that the information returned must be the accurate in real-time. This requires that the server synchronize activities before gathering information. While this is not this is not a resource intensive activity, it is one that may take considerable amount of elapsed time. The client using "z" option should be ready to wait a significant amount of time for a response.

Returned Response

The general XML schema for the **kXR_Qstats** response is:

```
<statistics tod="time" ver="version">details</statistics>
details: <stats id="sect">details</stats>[details]
```

Where:

time is the Unix **time()** value of when the statistics were generated.

vers is the XRootD version identification string.

sect is the section name assigned to the statistical information. Currently, the following section names should be expected to occur:

id	<i>arg</i>	Information
buff	b	Buffer usage statistics.
Cms	p	Cluster Management Services
info	i	Server identification.
link	l	Connection (i.e., link) statistics.
ofs	p	Open File System layer
oss	p	Open Storage System layer
poll	d	Device polling statistics.

kXR_Qstats

proc **u** - Process usage statistics.
rootd **p** - Protocol statistics for rootd.
sched **s** - Scheduling statistics.
XRootD **p** - Protocol information for XRootD.

Notes

- 1) Each subsection is bracketed by `<stats>` and `</stats>` tags.
- 2) Sections appear in a server-defined order. The sections, corresponding to each requested letter code, are returned.
- 3) The detailed contents of each section beyond what is described here is implementation dependent.

4.20.7 KXR_query Visa Request

Purpose: Obtain server information.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_query	kXR_unt16 kXR_ok
kXR_unt16 kXR_QVisa	kXR_int32 <i>ilen</i>
kXR_char <i>reserved</i> [2]	kXR_char <i>info</i> [<i>ilen</i>]
kXR_char <i>fhandle</i>	
kXR_char <i>reserved2</i> [8]	
kXR_int32 0	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

fhandle

is the file handle value supplied by the successful response to the associated **kXR_open** request.

ilen is the binary length of the information, *info*, that follows *ilen*.

info is the requested information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The response to the **kXR_Qvisa** request is documented in the Bandwidth Manager Configuration reference.

4.20.8 KXR_query Xattr Request

Purpose: Obtain server information.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_query	kXR_unt16 kXR_ok
kXR_unt16 kXR_Qxattr	kXR_int32 <i>ilen</i>
kXR_char <i>reserved</i> [14]	kXR_char <i>info</i> [<i>ilen</i>]
kXR_int32 <i>plen</i>	
kXR_char <i>path</i> [<i>plen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

plen is the binary length of the supplied path, *path*.

path is the path of the file whose extended attributes are to be returned. The *path* may be suffixed with CGI information.

ilen is the binary length of the information, *info*, that follows *ilen*.

info is the requested information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The response to the **kXR_Qxattr** request is documented in the Open File System (**ofs**) and Open Storage System (**oss**) configuration reference.
- 3) The **kXR_Qxattr** request bears no relationship to the **kXR_fattr** request.

4.21 kXR_read Request

Purpose: Read data from an open file.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_read	kXR_unt16 <i>status</i>
kXR_char <i>fhandle</i> [4]	kXR_int32 <i>dlen</i>
kXR_int64 <i>offset</i>	kXR_char <i>data</i> [<i>dlen</i>]
kXR_int32 <i>rlen</i>	
kXR_int32 <i>alen</i>	
<i>alen</i> > 0: read_args	

read_args	readahead_list
kXR_char <i>pathid</i>	kXR_char <i>fhandle2</i> [4]
<i>alen</i> > 8:	kXR_int31 <i>rlen2</i>
kXR_char <i>reserved</i> [7]	kXR_int64 <i>roffset2</i>
readahead_list [<i>n</i>]	
<i>n</i> = (<i>alen</i> -8)/16 with no remainder	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

status is the ending status of this request. Only the following two status codes indicate a normal ending:

- kXR_ok** - All of the data has been transmitted without error.
- kXR_oksofar** - Partial data has been transmitted without error; additional data should be expected on this stream.

offset is the binary offset from which the data is to be read.

rlen is the binary *maximum* amount of data that is to be read.

alen is the binary length of the arguments that follow the request header. These arguments may include the *pathid* and read-ahead request list, **read_aheadlist**. If no data is to be pre-read, *alen* must be set to less than or equal to eight (typically zero).

kXR_read

pathid is the pathid returned by **kXR_bind**. The response data is sent to this path, if possible.

fhandle2

is the file handle value supplied by the successful response to the associated **kXR_open** request that is to be used for the pre-read request. Each *fhandle2* is treated separately allowing pre-reads to occur from multiple files.

rlen2 is the binary *maximum* amount of data that is to be pre-read. The *rlen2* should correspond to the intended amount of data that will be read at *offset2* in the near future.

offset2 is the binary offset from which the data is to be pre-read. The *offset2* should correspond to the intended offset of data that will be read in the near future.

dlen is the binary length of the of the data, *data*, that was *actually* read.

data is the data that was read.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) If more data is requested than the file contains, the total of all *dlen*'s should be less than *rlen*.
- 3) Reading past the end of file with a valid offset should return a *dlen* of 0.
- 4) The *fhandle* value should be treated as opaque data.
- 5) Since a read may request more data than the allowable internal buffer size, the data is sent in fixed-sized segments until the request is satisfied. This is accomplished using the **kXR_oksofar** status code. Each subsequent data segment is transmitted using a {*streamid*, *status*, *dlen*, *data*} response. The last segment is indicated by a **kXR_ok**, if no error occurred.
- 6) Any status code other than **kXR_oksofar** indicates the end of transmission.
- 7) Sending requests using the same *streamid* when a **kXR_oksofar** status code has been returned may produced unpredictable results. A client must serialize all requests using the *streamid* in the presence of partial results.

- 8) The **kXR_read** request allows you to also schedule the pre-reading of data that you will ask for in the very near future. Pre-reading data may substantially speed up the execution because data will be available in memory when it is actually asked for. On the other hand, requesting data that you will not need will simply cause a general slow-down of the complete system.
- 9) The pre-read request is considered only a hint. The system may or may not honor the pre-read request, depending on the current system load.
- 10) To schedule a pre-read without actually reading any data, issue a **kXR_read** request with *rLen* and *offset* set to zero and **readahead_list** filled out to reflect what data should be pre-read.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_read		3013	Read a file

4.22 kXR_readv Request

Purpose: Read data from one or more open files.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_readv	kXR_uint16 <i>status</i>
kXR_char <i>reserved</i> [15]	kXR_int32 <i>dlen</i>
kXR_char <i>pathid</i>	kXR_char <i>data</i> [<i>dlen</i>]
kXR_int32 <i>alen</i>	
read_list [<i>n</i>]	
<i>n</i> = <i>alen</i> /16 with no remainder	
read_list	
kXR_char <i>fhandle2</i> [4]]
kXR_int31 <i>rlen2</i>	
kXR_int64 <i>roffset2</i>	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

status is the ending status of this request. Only the following two status codes indicate a normal ending:

- kXR_ok** - All of the data has been transmitted without error.
- kXR_oksofar** - Partial data has been transmitted without error; additional data should be expected on this stream.

alen is the binary length of the arguments that follow the request header. These arguments may include the *pathid* and read request list, **struct read_args**. The maximum allowed value for *alen* is 8200. This allows up to 512 read segments.

pathid is the pathid returned by **kXR_bind**. The response data is sent to this path, if possible.

kXR_readv

fhandle

is the file handle value supplied by the successful response to the associated **kXR_open** request that is to be used for the read request. Each *fhandle* is treated separately allowing reads to occur from multiple files.

rlen is the binary *maximum* amount of data that is to be read. Less data will be read if an attempt is made to read past the end of the file.

offset is the binary offset from which the data is to be read..

dlen is the binary length of the of the response data, *data*.

data is the response data. The response data includes **read_list** headers preceding the actual data that was read.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Each **read_list** element represents a read request. All of the read requests are aggregated into one or more responses. Read data is always prefixed by its corresponding **read_list** element. However, the *rlen* value in the element indicates the actual amount of data that was read.
- 3) The **struct read_list** is equivalent to **struct readahead_list**.
- 4) If an element requests more data than the file contains, the returned *rlen* should be smaller than the *rlen* in the request element.
- 5) Reading past the end of file with a valid offset should return a request element whose *rlen* is 0 with no data following.
- 6) The *fhandle* value should be treated as opaque data.
- 7) Since a read may request more data than the allowable internal buffer size, the data is sent in fixed-sized segments until the request is satisfied. This is accomplished using the **kXR_oksofar** status code. Each subsequent data segment is transmitted using a {*streamid*, *status*, *dlen*, *data*} response. The last segment is indicated by a **kXR_ok**, if no error occurred.
- 8) Any status code other than **kXR_oksofar** indicates the end of transmission.
- 9) Sending requests using the same *streamid* when a **kXR_oksofar** status code has been returned may produced unpredictable results. A client must serialize all requests using the *streamid* in the presence of partial results.

10) The server may return the read elements in any order.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_readv		3025	Read file segments

4.23 kXR_rm Request

Purpose: Remove a file.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_rm	kXR_unt16 kXR_ok
kXR_char <i>reserved</i> [16]	kXR_int32 0
kXR_int32 <i>plen</i>	
kXR_char <i>path</i> [<i>plen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

plen is the binary length of the supplied path, *path*.

path is the path of the of the file to be removed. The *path* may be suffixed with CGI information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The kXR_char and kXR_unt16 data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_rm		3014	Remove a file

kXR_rmdir

4.24 kXR_rmdir Request

Purpose: Remove a directory.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_rmdir	kXR_uint16 kXR_ok
kXR_char <i>reserved</i> [16]	kXR_int32 0
kXR_int32 <i>plen</i>	
kXR_char <i>path</i> [<i>plen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

plen is the binary length of the supplied path, *path*.

path is the path of the of the directory to be removed. The *path* may be suffixed with CGI information.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_uint16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The directory must be empty (i.e., no entries other than "." And "..").

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_rmdir		3015	Remove a directory

4.25 kXR_set Request

Purpose: Set server information.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_set	kXR_unt16 kXR_ok
kXR_char <i>reserved</i> [15]	kXR_int32 0
kXR_char <i>modifier</i>	
kXR_int32 <i>dlen</i>	
kXR_char <i>data</i> [<i>dlen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

modifier

set request modifier and should be initialized to zero.

dlen is the binary length of the supplied value, *data*.

data is the value to set.

resp is the response value to the specific set requested.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Set processing takes a command-like string in the *data* field. The following documents valid set arguments.

kXR_set

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_set		3018	Set monitoring information

4.25.1 Valid kXR_Set Values

appid *apmsg*

```
monitor {off | on} [appid] | info [info]}
```

Where:

appid *apmsg*

includes *apmsg* in the server's log. This request is meant to be used to identify the start and stop of certain application processes for rudimentary monitoring purposes. Up to 80 characters should be recorded.

monitor

control monitor settings with respect to the application.

off - turns off monitoring for the application.

appid - includes up to 12 characters of application text in the monitor record.

on - turns on monitoring, if allowed by the configuration.

appid - includes up to 12 characters of application text in the monitor record.

info - insert information into the monitoring record, if information monitoring is active.

info - is up to 1024 characters of information

Response

is the unique four-character identification value that has been assigned to the *info* value.

Notes

- 1) Monitoring is enabled using the **xrootd.monitor** configuration directive. When monitoring is not enabled, the monitor set requests are ignored.
- 2) Use the returned identification value to tag future records in order to correlate related information.

4.26 kXR_sigver Request

Purpose: Provide a signature for the next request.

Request	Normal Response
kXR_char <i>streamid</i> [2]	None (see notes)
kXR_uint16 kXR_sigver	
kXR_uint16 <i>expectrid</i>	
kXR_char <i>version</i>	
kXR_char <i>flags</i>	
kXR_uint64 <i>seqno</i>	
char <i>crypto</i>	
kXR_char <i>reserved</i> [3]	
kXR_int32 <i>dlen</i>	
kXR_char <i>Sigdata</i> [<i>dlen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request. The *streamid* should be identical to the *streamid* in the subsequent request.

expectrid

is the *requestid* of the subsequent request.

version

is version of the signature protocol being used. Currently, there is only one version so *version* should be set to zero.

flags request indicators:

kXR_nodata the data payload is not included in the hash.

seqno is a monotonically increasing sequence number. Each **kXR_sigver** request should have a sequence number that is greater than a previous sequence number used on a particular TCP connection (i.e. login session).

crypto the cryptography used to construct the signature:

kXR_rsaKey the rsa key encrypts the hash.

kXR_SHA256 The hash used is SHA-2.

kXR_sigver

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

dlen is the length of the subsequent signature. This is normally an encrypted hash of the subsequent request.

sigdata

the signature for the subsequent request.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Only an error response is possible to a **kXR_sigver** request. This request simply supplies information used to verify that an authenticated client sent the subsequent request.
- 3) Only SHA-256 hashes and session key encryption are currently defined as supported.
- 4) Encryption should be done using the encryption facilities of the authentication protocol used.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_sigver		3029	Present signature for following request
	<i>crypto</i>		
	kXR_rsaKey	0x80	RSA key encrypts the hash
	kXR_SHA256	0x01	SHA-256 used for the hash
	<i>flags</i>		
	kXR_nodata	0x01	Data payload is not signed

4.26.1 Signing a request

When the response to **kXR_protocol** indicates that a particular request needs to be signed, the client should construct a **kXR_sigver** request and send it immediately ahead of the request that needs to be signed. The server should reject any request that should be signed but has been not signed or incorrectly signed.

A signature consists of a **SHA-256** hash of

- an unsigned 64-bit sequence number,
- the request header, and
- the request payload³;

in that *exact* order. The *crypto* member of **kXR_sigver** should indicate **kXR_SHA256** and the **kXR_rsakey** should not be set.

The client should add one to the sequence number previously used in a **kXR_sigver** request on a particular TCP connection (i.e. login session) before using it in the current request. Sequence numbers should be monotonically increasing on each TCP connection.

Once the hash has been computed, it should be encrypted using the session key developed by the authentication protocol used during the login authentication process. If the authentication protocol does not support generic encryption, an unencrypted hash should be used if the server set the **kXR_secOfrc** option in the **kXR_protocol** response. If the option is not set in this situation the client should not sign the request.

The **kXR_sigver** request should be sent immediately before the request that it signed.

³ When signing a **kXR_write** or **kXR_verifyw** request, the data payload should not be included in the hash and the **kXR_nodata** option should be set in the **kXR_sigver** option field.

kXR_sigver

4.26.2 Verifying a signed request

When the server receives a **kXR_sigver** request, it should assume that it is the signature for the following request. Note that **kXR_sigver** requests cannot be signed so a hold depth of one request is sufficient. Once the next request is received, the following steps should be followed to verify its signature where *sigver* refers to the **kXR_sigver** request and its payload and *thereq* refers to the request whose signature is being verified. If a failure occurs at any step, the request should be rejected with an error sent to the client and the TCP connection can be closed.

- Verify that *sigver.seqno* is greater than the previous *sigver.seqno* the client used on the TCP connection. The saved value should only be updated if the signature is successfully verified.
- Verify that *sigver.streamid* equals *thereq.streamid*.
- Verify that *sigver.expectrid* equals *thereq.requestid*.
- Verify that *sigver.version* matches the version being used.
- Verify that *sigver.crypto* specifies SHA-256.
- Verify that *sigver.crypto* does not specify the use of the rsa key.
- If the hash was encrypted (i.e. the authentication protocol supports generic encryption) successfully decrypt the hash using the session key via the authentication protocol used.
- Generate a new hash using the *sigver.seqno*, *thereq*, and the *thereq* payload⁴ in that exact order.
- Verify that the new hash matches the decrypted hash.
- Upon success update the sequence number used so that the sequence number cannot be reused.

⁴ When verifying a **kXR_write** or **kXR_verifyw** request, the data payload should not be included in the hash (i.e. the **kXR_nodata** option should be set in **the kXR_sigver** option field).

4.27 kXR_stat Request

Purpose: Obtain status information for a path.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_stat	kXR_unt16 kXR_ok
kXR_char <i>opts</i>	kXR_int32 <i>ilen</i>
kXR_char <i>reserved</i> [11]	kXR_char <i>info</i> [<i>ilen</i>]
kXR_char <i>fhandle</i> [4]	
kXR_int32 <i>plen</i>	
kXR_char <i>path</i> [<i>plen</i>]	

Default info: id size flags mtime [ctime atime mode owner group]⁵\0

kXR_vfs *info: nrw frw urw nstg fstg ustg\0*

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

opts

are stat processing options:

kXR_vfs - return virtual file system information for the specified *path*

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

plen

is the binary length of the supplied path, *path*. If *plen* is zero then *fhandle* must hold the file handle value supplied by the successful response to the associated **kXR_open** request. The *fhandle* argument is ignored when a *path* is present or when **kXR_vfs** is specified. In the latter case, a *path* argument is required.

path

is the path whose status information is to be returned. The *path* may be suffixed with CGI information.

⁵ Information in brackets is only returned for servers implementing version 4.0.0 of the protocol.

kXR_stat

ilen is the binary length of the information, *info*, that follows *ilen*.

Default Response

info is the information about the requested path. Note that information in brackets is returned only if the server supports extended stat information as defined in protocol version 4.0.0 or higher.

id is the OS-dependent identifier assigned to this entry. Uniqueness is not guaranteed. The *id* is numeric and convertible to a 64-bit value.

size is the decimal size of the data associated with the path whose information is being returned. The size may represent a number up to $2^{64}-1$ (i.e., a long long).

flags identifies the entry's attributes as a decimal encoded 32-bit string. The entry should be assumed to be a regular file unless one or more of the following bits are set.

- kXR_xset** - Either an executable file or a searchable directory.
- kXR_isDir** - This is a directory.
- kXR_other** - This neither a file nor a directory.
- kXR_offline** - For files, the file is not online (i.e., on disk).
- kXR_pospend** - The file was created with **kXR_posc** and has not yet been successfully closed.
- kXR_readable** - Read access allowed.
- kXR_writable** - Write access allowed.

mtime is the last modification time in Unix time units (i.e., seconds since 00:00:00 UTC, January 1, 1970).

ctime is the creation time in Unix time units (i.e., seconds since 00:00:00 UTC, January 1, 1970).

atime is the last access time in Unix time units (i.e., seconds since 00:00:00 UTC, January 1, 1970).

mode is the octal value of the entries mode bits (i.e. permissions and settings). A leading zero should always appear to indicate this is an octal value.

owner the name of the owner associated with the entry. If the name cannot be determined, the owner's numeric identifier should be returned.

group the name of the group associated with the entry. If the name cannot be determined, the group's numeric identifier should be returned.

Response for kXR_vfs

info the location information about the requested path.

nrw the number of nodes that can provide read/write space.

frw the size, in megabytes, of the largest contiguous area of r/w free space.

urw percent utilization of the partition represented by *frw*.

nstg the number of nodes that can provide staging space.

fstg the size, in megabytes, of the largest contiguous area of staging free space.

ustg percent utilization of the partition represented by *fstg*.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The returned string is compatible to the format returned by the root method `Tsystem::GetPathInfo()`.
- 3) **kXR_stat** requests, without the **kXR_vfs** option, directed to a redirector (i.e., load balancer) referring to a non-file object may result in a non-deterministic response. That is, the response may indicate that the object does not exist when, in fact, it does exist but is not a file. Future versions may resolve the differences between redirectors and file servers.
- 4) **kXR_stat - kXR_vfs** requests need not specify an existing filesystem object. The specified path is used as a path prefix in order to filter out servers and partitions that could not be used to hold objects whose path starts with the specified path prefix.
- 5) **kXR_stat - kXR_vfs** requests directed to a redirector return the space values based on current conditions and should be treated as

kXR_stat

approximations. When the request is directed to an actual server, the server's actual space information is returned.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_stat		3017	Get file information
	<i>opts</i>		
	kXR_vfs	0x01	Return file system information
Response	<i>flags</i>		
	kXR_xset	0x00 00 00 01	Execution/search bit set
	kXR_isDir	0x00 00 00 02	Item is a directory
	kXR_other	0x00 00 00 04	Item neither directory nor file
	kXR_offline	0x00 00 00 08	Item is not online
	kXR_readable	0x00 00 00 10	Item is readable
	kXR_writable	0x00 00 00 20	Item is writable
	kXR_poscpend	0x00 00 00 40	Item is pending a close
	kXR_bkpexist	0x00 00 00 80	Item has a backup

4.27.1 Additional Stat CGI Tokens

The **kXR_stat** request allows a client to pass CGI information to restrict information lookup to online space. The information may or may not be acted upon, depending on the server's capabilities. CGI information is passed by suffixing the *path* with a question mark (?) and then coding the cgi information as shown below:

```
path?oss.lcl=1
```

Example

```
/tmp/foo?oss.lcl=1
```

4.28 kXR_statx Request

Purpose: Obtain type information for one or more paths.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_unt16 kXR_statx	kXR_unt16 kXR_ok
kXR_char <i>reserved</i> [16]	kXR_int32 <i>ilen</i>
kXR_int32 <i>plen</i>	kXR_char <i>info</i> [<i>ilen</i>]
kXR_char <i>path</i> [<i>plen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

plen is the binary length of the supplied path list, *paths*.

paths is the new-line separated path list whose availability information is to be returned. Each *path* in *paths* may be suffixed with CGI information. If a single path is supplied, it need not end with a new line character (\n).

ilen is the binary length of the information, *info*, that follows *ilen*.

info is the information about the requested path consisting of a single binary character flag for each path in *paths*.

flags identifies the entry's attributes as a binary character. The entry should be assumed to be an immediately available regular file unless one or more of the following bits are set.

kXR_xset - Either an executable file or a searchable directory.

kXR_isDir - This is a directory.

kXR_other - This neither a file nor a directory, or does not exist.

kXR_offline - For files, the file is not online (i.e., on disk).

kXR_truncate

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) **kXR_statx** requests directed to a redirector (i.e., load balancer) referring to a non-file object may result in a non-deterministic response. That is, the response may indicate that the object does not exist when, in fact, it does exist but is not a file. Future versions may resolve the differences between redirectors and file servers.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_statx		3022	Get file information
Response	<i>flags</i>		
	kXR_xset	0x00 00 00 01	Execution/search bit set
	kXR_isDir	0x00 00 00 02	Item is a directory
	kXR_other	0x00 00 00 04	Item neither directory nor file
	kXR_offline	0x00 00 00 08	Item is not online
	kXR_readable	0x00 00 00 10	Item is readable
	kXR_writable	0x00 00 00 20	Item is writable
	kXR_poscpend	0x00 00 00 40	Item is pending a close
	kXR_bkpexist	0x00 00 00 80	Item has a backup

4.29 kXR_sync Request

Purpose: Commit all pending writes to an open file.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_sync	kXR_uint16 kXR_ok
kXR_char <i>fhandle</i> [4]	kXR_int32 0
kXR_char <i>reserved</i> [12]	
kXR_int32 0	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

fhandle

is the file handle value supplied by the successful response to the associated **kXR_open** request.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_uint16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The *fhandle* value should be treated as opaque data.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_sync		3016	Synchronize file to disk

4.30 kXR_truncate Request

Purpose: Truncate a file to a particular size.

Request (Opened File)	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_truncate	kXR_uint16 kXR_ok
kXR_char <i>fhandle</i> [4]	kXR_int32 0
kXR_int64 <i>size</i>	
kXR_char <i>reserved</i> [4]	
kXR_int32 0	

Request (Closed File)	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_truncate	kXR_uint16 kXR_ok
kXR_char <i>reserved</i> [4]	kXR_int32 0
kXR_int64 <i>size</i>	
kXR_char <i>reserved</i> [4]	
kXR_int32 <i>plen</i>	
kXR_char <i>path</i> [<i>plen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

fhandle

is the file handle value supplied by the successful response to the associated **kXR_open** request. The *fhandle* argument is ignored when a path is present.

size is the binary size that the file is to have.

plen is the binary length of the supplied path, *path*.

path is the path of the of the file to be truncated. The *path* may be suffixed with CGI information.

kXR_truncate

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_uint16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The *fhandle* value should be treated as opaque data.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_truncate		3028	Truncate a file

4.31 kXR_verifyw Request

Purpose: Write data to an open file with checksum validation (placeholder for future enhancement).

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_verifyw	kXR_uint16 kXR_ok
kXR_char <i>fhandle</i> [4]	kXR_int32 0
kXR_int64 <i>offset</i>	
kXR_char <i>pathid</i>	
kXR_char <i>vertype</i>	
kXR_char <i>reserved</i> [2]	
kXR_int32 <i>dlen</i>	
kXR_char <i>data</i> [<i>dlen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

fhandle

is the file handle value supplied by the successful response to the associated **kXR_open** request.

offset is the binary offset to which the data is to be written.

pathid is the *pathid* returned by **kXR_bind**. The actual data is read from this path. If alternate path reading is not wanted, *pathid* should be set to zero.

vertype

identifies the checksum algorithm used; implying how many bytes of checksum data precede the data to be written; as follows:

Option	Bytes	Algorithm
kXR_nocrc	0	no crc computed
kXR_adl32	4	32-bit adler
kXR_crc32	4	32-bit crc

kXR_verifyw

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

dlen is the binary length of the of the data, *data*, to be written.

data is the data to be written.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The *fhandle* value should be treated as opaque data.
- 3) The crc data must immediately precede the data to be written.
- 4) The length of the crc data plus the length of the data to be written must equal the *dlen* value.
- 5) Using an unsupported *vertype* causes the **verifyw** request to fail.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_verifyw		3026	Write wiih verification
	<i>vertype</i>		
	kXR_adl32	0x02	Adler32 checksum supplied
	kXR_crc32	0x01	CRC32 checksum supplied
	kXR_nocrc	0x00	No checksum supplied

4.32 kXR_write Request

Purpose: Write data to an open file.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_write	kXR_uint16 kXR_ok
kXR_char <i>fhandle</i> [4]	kXR_int32 0
kXR_int64 <i>offset</i>	
kXR_char <i>pathid</i>	
kXR_char <i>reserved</i> [3]	
kXR_int32 <i>dlen</i>	
kXR_char <i>data</i> [<i>dlen</i>]	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

fhandle

is the file handle value supplied by the successful response to the associated **kXR_open** request.

offset is the binary offset to which the data is to be written.

pathid is the *pathid* returned by **kXR_bind**. The actual data is read from this path.

reserved

is an area reserved for future use and must be initialized to null characters (i.e., '\0').

dlen is the binary length of the of the data, *data*, to be written.

data is the data to be written.

kXR_write

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_uint16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) The *fhandle* value should be treated as opaque data.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_write		3019	Write into a file

4.33 kXR_writev Request

Purpose: Write data to one or more open files.

Request	Normal Response
kXR_char <i>streamid</i> [2]	kXR_char <i>streamid</i> [2]
kXR_uint16 kXR_readv	kXR_uint16 kXR_ok
kXR_char <i>options</i>	kXR_int32 0
kXR_char <i>reserved</i> [15]	
kXR_int32 <i>alen</i>	
write_list [<i>n</i>]	
<i>n</i> = <i>alen</i> /16 with no remainder	

write_list	
kXR_char <i>fhandle</i> [4]]
kXR_int31 <i>wlen</i>	
kXR_int64 <i>offset</i>	

Where:

streamid

is the binary identifier that is associated with this request stream. This identifier should be echoed along with any response to the request.

options

are the following options or'd together:

kXR_doSync - perform an **fsync()** after data is written.

alen

is the binary length of the arguments that follow the request header. The argument should be a write request list, **struct write_list**. The maximum allowed value for *alen* is 8200. This allows up to 512 write segments.

fhandle

is the file handle value supplied by the successful response to the associated **kXR_open** request that is to be used for the write request. Each *fhandle* is treated separately allowing writes to occur to multiple files. Note that the *fhandle* causing an error is echoed back in an error response.

wlen

is the binary amount of data that is to be written. Elements whose *wlen* is zero are ignored.

kXR_writev

offset is the binary offset to which the data is to be written.

status is the ending status of this request. A **kXR_ok** status indicates successful completion. Any other status code indicates an error and a description follows.

dlen is the binary length of the of the error response data that follows. Note that the first four bytes are the binary filehandle that encountered the error.

msg is a null terminated message describing the error.

Notes

- 1) All binary fields are transmitted in network byte order using an explicit length. The **kXR_char** and **kXR_unt16** data types are treated as **unsigned** values. All reserved fields must be initialized to binary zero.
- 2) Each **write_list** element represents a write request. All of the write requests are aggregated into one response.
- 3) The *fhandle* value should be treated as opaque data.
- 4) Data to be written should follow the **write_list**. The total number of bytes that follow should equal the sum of all the individual *wlen* values.
- 5) The server should process the elements in the order specified.
- 6) In the presence of the **doSync** option, an **fsync()** should be performed after all the elements have been written to the file corresponding to *fhandle*. This may occur at the end of the **write_list** if no other file handles appear in the list or when encountering a new *fhandle*. When a new *fhandle* is encountered, an **fsync()** should be performed for the previous *fhandle* before processing the list for the new *fhandle*.
- 7) When an error occurs, it is unpredictable how many bytes have actually been written to any particular file specified in the **write_list**. Specifically, elements prior to the error may have been partially written while elements subsequent to the element causing the error are ignored.

Binary Definitions

Request	Modifiers	Value	Explanation
kXR_writev		3031	Write file segments
	<i>options</i>		
	kXR_doSync	0x01	Perform kXR_sync afterwards

5 The Security Framework

The **XRootD** protocol provides for a generic security framework in which virtually any security protocol can be used. The **XRootD** protocol neither specifies nor mandates that any particular security protocol be used. However, should a security protocol be implemented, the **XRootD** protocol specifies how the security protocol's data elements are to be framed and how client-server interactions during the security protocol's execution are to be handled.

The first two protocol steps that a client does after connecting to a server are:

1. sends the 20-byte handshake, followed by a
2. **kXR_login** request.

At this point, the server may require that a security protocol be used to authenticate the client. It is the server that initiates the authentication exchange. The client can never force an authentication to occur. The server mandates that authentication occur by responding to the **kXR_login** request with security information. This is an implicit action on the server's behalf.

Authentication is not required when the data length portion of the response to the **kXR_login** is exactly 16 bytes⁶. When the data length is greater than 16 bytes, the data portion of the response contains what is called a security token. This is essentially a list of the security protocols that the server can use to establish the client's identity along with possible configuration information that is specific to each protocol.

⁶ Protocol versions less than 2.4 used a different signifier. Refer to the **kXR_login** request code description for detailed information on deprecated protocol versions.

Security Framework

The format of the security token is:

```
ptoken: &P=protid[,protparms] [ptoken]
```

Where:

protid is a 1- to 7-character protocol name. This name is typically used to locate a shared library that implements the security protocol.

protparms

are optional protocol specific configuration parameters that must be supplied to the protocol's initialization routine when it is instantiated. The comma is required if *protparms* are present. The comma is optional otherwise. The *protparms* may not contain an ampersand (&).

In the reference implementation, the *protid* is specified by the **sec.protocol** configuration directive and *protparms* can come from the **sec.protocol** and the **sec.protparm** directives. The plug-in code that handles the protocol is then assumed to reside in **libXrdSec*protid*.so** which is dynamically loaded by **libXrdSec.so** which itself is loaded by the client when a security token has been returned in response to a login. See the "Scalla Security Configuration Reference" for more information.

Generally, the security token is handled by some class that sequences the authentication process. Only that class should be aware of the token's format. This class is responsible for loading one of the listed protocols and initiating the authentication sequence as defined by the security protocol. For consistency among implementations, it is recommended that protocols be considered from left to right and that protocols be successfully tried until one is found to succeed.

For instance, in the **krb5** security protocol, *protparms* defines the service principal whose ticket must be obtained and sent back to the server to prove the client's identity. The security token would appear as

&P=krb5,*srvname*

with *srvname* being the service principal name. If the service ticket must be forwardable, then the token would be sent as

&P=krb5,*srvname*,fwd****

Each protocol specifies its own *protparms* requirements. Refer to the “Scalla Security Configuration Reference” for more information for each available protocol.

The normal sequence in almost any security protocol is that one side generates data sends it to the other side that either accepts or rejects the data and may respond with other data which the receiver may or may not be required to respond to. This is a generalization of multiple exchanges. The XRootD protocol handles such exchanges without interpretation; as follows:

- 1 After the client-side security manager chooses a protocol and successfully initializes it with the *protparms*; that protocol must return some data that will be sent to the server. The data must be sent to the server as a **kXR_auth** request. This data is known as credentials.
- 2 The first eight characters of the initial credentials should contain the null terminated protocol identifier of the protocol that generated the credentials. Hence, the actual credential data starts eight bytes into the credential data packet in the first credential packet sent to the server. Subsequent packet layout is defined by the security protocol.
- 3 When the server receives the initial **kXR_auth** request; it should attempt to use a protocol handler that matches the protocol identifier contained in the credentials (i.e., in the first eight bytes). When the protocol handler is created, the credentials should be passed to its authentication method. If a match cannot be found or initialization fails, a **kXR_error** response must be sent. The connection should remain opened so that the client may try an alternate protocol without performing a new login. For subsequent **kXR_auth** requests, the same protocol used in the successful handling of the initial **kXR_auth** request must be used.

Security Framework

- 4 After the credentials are processed by the authentication method; three possibilities exist:
 - a) The data is accepted and no more exchanges are needed. In this case, the response to the **kXR_auth** request must be a **kXR_ok**. The client is then considered to be fully authenticated.
 - b) The data is not accepted and authentication failed. In this case, a **kXR_error** response must be sent with the connection remaining open.
 - c) Finally, additional information is needed to complete the authentication. The authentication method supplies the data that data must be sent to the client who is expected to respond with a **kXR_auth** request. The data provided by the authentication method must be used as the body of a **kXR_authmore** response.
- 5 The client's action for each of the three possible responses is:
 - a) Authentication exchanges are stopped when **kXR_ok** is received. The client is considered to be fully authenticated and may issue other requests.
 - b) Authentication exchanges are stopped when a **kXR_error** is received. The client may declare a fatal error at this point.
 - c) An authentication exchange should occur when **kXR_authmore** is received. The protocol's get credentials method should be called with the data in the **kXR_authmore** response. The method may then supply new data that must be sent to the server in a new **kXR_auth** request or indicate failure. In this processing continues with step 3.

5.1 Framework for Transport Layer Protocols

The **XRootD** security framework and the protocol elements that support it are geared to application layer security. Indeed, implementations are highly discouraged to expose the underlying transport to application code. This is necessary in order to allow multiple transports to be used in a transparent way.

Unfortunately, certain commonly available security implementations are either transport based (e.g., **ssl**) or rely on intimate knowledge of the transport out of expediency.

Currently, the **XRootD** protocol elements neither accommodate nor define direct transport layer interactions between the client and server parts of a security protocol; though nothing prevents implementations of doing so out of band or perhaps hijacking the **XRootD** connection for a limited amount of time. However, the reference implementation makes it impossible to do either.

That said, the reference implementation does provide a virtualized transport via a wrapper class called **XrdSecTLayer**. This class can be used to wrap transport layer security protocols implementations and carry out what appear to be transport layer interactions. The wrapper class virtualizes the interactions by appropriately framing all exchanges within the protocol defined in this document.

5.2 Request Verification

The protocol defines a mechanism to verify that a request came from a previously authenticated client. This is done using lightweight cryptographic signing. Signing requirements are completely controlled by the server that the client communicates with. The server uses the **kXR_protocol** request to inform the client of any signing requirements. The client uses the **kXR_sigver** request to sign a request when needed. When the server receives a **kXR_sigver** request it should use the information in the request to verify that the subsequent request actually came from an authenticated client. A server should reject any request where the signature cannot be verified or a request that needs to be signed was not preceded by a **kXR_sigver** request. When a request is reject because it was improperly signed, the server should also close the connection to the client after reporting the error to the client.

6 Document Change History

1 June 2005

- Add **kXR_bind** and **kXR_endsess** request codes.
- Explain how a *sessid* is returned in response to **kXR_login**.
- Add **kXR_open_apnd** and **kXR_retstat** options to **kXR_open**.

28 July 2005

- Document the administrative interface protocol.

16 Aug 2005

- Document the **lsd** administrative command.

25 Jan 2006

- Document the **cj** administrative command.
- Document the **lsj** administrative command.
- Add **kXR_Cancelled** subtype error code.
- Add **kXR_Qckscan** subtype request to **kXR_query**.

25 Jan 2006

- Document **kXR_readv**.
- Complete documentation of **kXR_bind**.
- Redefine the pre-read structure in **kXR_read** to include a *pathid* argument.
- Add a *pathid* to **kXR_write**.

5 Dec 2006

- Document **kXR_Qconfig** subcode of **kXR_query**.
- Document **kXR_unbind**.
- Explain ramification of not using **kXR_unbind** in the **kXR_bind** description.
- Clarify **kXR_open** request with respect **kXR_compress** and **kXR_retstat**.

25 Jan 2007

- Document **pio_max** variable for **kXR_Qconfig** sub-request of **kXR_query**.

26 Feb 2007

- Change **kXR_prepare** to reflect that the priority is really a char.

Change History

1 Aug 2007

- Document the **kXR_verifyw** request.
- Document the **kXR_replica**, **kXR_ulterior**, and **kXR_nowait** options.

26 Sep 2007

- Document the **kXR_locate** request.

15 Nov 2007

- Document the **kXR_nowait** option of the **kXR_locate** request.
- Document the **kXR_vfs** option of the **kXR_stat** request.

13 Mar 2008

- Document the **kXR_qspace** and **kXR_qxattr** options of the **kXR_query** request.

7 Apr 2008

- Document the **kXR_truncate** request.

12 May 2008

- Correct **kXR_Query** documentation w.r.t. the subcode location.
- Document the **kXR_QVisa** variant or **kXR_Query**.

20 Aug 2008

- Correct **kXR_coloc** and **kXR_fresh** options of the **kXR_prepare** request.
- Document the **kXR_Qopaque** and **kXR_Qopaquf** variants of **kXR_Query**.

26 Jan 2009

- Correct description of **lsj** admin command xml output.

8 Apr 2009

- Document **kXR_seqio** option of the **kXR_open** request.
- Add *fhandle* to the **kXR_stat** request to allow getting stat information based on a currently open file.

6 May 2009

- Describe the security framework as related to the protocol.

2 Jun 2009

- Describe the **kXR_posc** open flag and the **kXR_poscpend** stat response flag.

14 Jul 2009

- Alter description of **kXR_query** + **kXR_QStats** to indicate that other than the basic framing of the information, the actual XML package is implementation dependent.

9 Dec 2010

- Document missing field, *credtype*, in the **kXR_auth** request. This field was always there but somehow escaped documentation. Leaving it unset does not change the protocol but also does not allow the client to switch protocols mid-stream.

14 Jul 2011

- Expand description of information that **kXR_protocol** may return when the client optionally specifies its own protocol version number (new extension).

28 Mar 2012

- Correct diagrams and expand on descriptive text for **kXR_query**, **kXR_read**, **kXR_readv**, and **kXR_set**.
- Add missing *pathid* argument to **kXR_readv** request.
- Expand on the text describing responses to **kXR_redirect**.
- Add **tpc** to the list of configuration variables that may be queried.

21 Jun 2012

- Better explain possible error recovery actions.
- Add optional elements that should have been described:
 - zone field in **kXR_login**
 - port field in **kXR_prepare**
 - pathid field in **kXR_readv**

Change History

16 Jul 2013

- Minor corrections on the valid operations table.
- Document the **kXR_dstat** option of **kXR_dirlist**.
- Document the implementation dependent **kXR_query** requests and add **kXR_Qopaqug** to the list.
- Add **cname**, **sitename** and **version** to the list of configuration variables that may be queried.
- Describe full URL redirect responses.
- Document the **kXR_fullurl**, **kXR_multipr**, and **kXR_readrdok** settings in the **kXR_login** request.
- Describe how locate can return a hostname response (i.e. the **kXR_prefname** option in **kXR_locate**).
- Describe the **kXR_locate kXR_addperrs** option.
- Describe optional login tokens.

20 Nov 2013

- Document the **kXR_Qconfig** “cms” and “role” options.

3 Apr 2014

- Add better explanation on how to recover from server failures.

15 Oct 2014

- Correct type **cname** should be **chksum** in query config.

23 Mar 2015

- Document the **cid** option of **kXR_query** config.

10 Feb 2016

- Correct **kXR_dirlist** layout description.
- Document **kXR_sigver** request.
- Document security requirement response to **kXR_protocol**.
- Document extensions to **kXR_mv** to handle names with embedded spaces.

8 Jan 2018

- Document **kXR_fattr** request.
- Document **kXR_writev** request.
- Document the **kXR_query config fattr** request.
- Remove documentation for the **kXR_query config wan_port** and **wan_window** request (no longer useful).

17 Feb 2018

- Specify binary values for all protocol symbols.
- Remove **kXR_unbind** request code.
- Modify description of **kXR_getfile** and **kXR_putfile**.

18 May 2018

- Move local administrative protocol description to a separate document.

28 May 2018

- Correct layout of the **kXR_protocol** request.

20 Jun 2018

- Correct description of **kXR_close** request.
- Remove ambiguities from the **kXR_dirlist** request description.
- Correct description of **kXR_mkdir** request regarding making directory paths.
- Remove ambiguities from the **kXR_open** request description.
- Correct description of **kXR_stat** request.
- Indicate **kXR_verifyw** request is a placeholder.

17 Jul 2018

- Document the **ClientFattrRequest::aData** option for the **kXR_fattrList** request subcode.
- Remove all references to the **kXR_unbind** request code.
- Add caveats to the **kXR_verifyw** request code.

17 Jul 2018

- Document missing error codes.

24 Oct 2018

- Describe extended stat information returned by **kXR_dirlist**, **kXR_open** and **kXR_stat**.
- Describe **CGI** information in greater detail.

17 Dec 2018

- Describe **kXR_protocol** modifications that support **TLS**, **kXR_getfile**, and **kXR_putfile**.
- Add new error code **kXR_TLSRequired**.