

Network Working Group  
Internet-Draft  
Updates: [3981](#) (if approved)  
Intended status: Standards Track  
Expires: September 6, 2007

A. Newton  
VeriSign, Inc.  
March 5, 2007

**XML Pipelining with Chunks for the Information Registry Information  
Service  
draft-ietf-crisp-iris-xpc-06**

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on September 6, 2007.

Copyright Notice

Copyright (C) The IETF Trust (2007).

## Abstract

This document describes a simple TCP transfer protocol for the Internet Registry Information Service (IRIS). Data is transferred between clients and servers using chunks to achieve pipelining.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">3</a>
<a href="#">2.</a>	Document Terminology . . . . .	<a href="#">4</a>
<a href="#">3.</a>	Request Block (RQB) . . . . .	<a href="#">5</a>
<a href="#">4.</a>	Response Blocks . . . . .	<a href="#">6</a>
<a href="#">4.1.</a>	Response Block (RSB) . . . . .	<a href="#">6</a>
<a href="#">4.2.</a>	Connection Response Block (CRB) . . . . .	<a href="#">6</a>
<a href="#">5.</a>	Block Header . . . . .	<a href="#">8</a>
<a href="#">6.</a>	Chunks . . . . .	<a href="#">9</a>
<a href="#">6.1.</a>	No Data Types . . . . .	<a href="#">11</a>
<a href="#">6.2.</a>	Version Information Types . . . . .	<a href="#">11</a>
<a href="#">6.3.</a>	Size Information Types . . . . .	<a href="#">11</a>
<a href="#">6.4.</a>	Other Information Types . . . . .	<a href="#">12</a>
<a href="#">6.5.</a>	SASL Types . . . . .	<a href="#">13</a>
<a href="#">6.6.</a>	Authentication Success Information Types . . . . .	<a href="#">14</a>
<a href="#">6.7.</a>	Authentication Failure Information Types . . . . .	<a href="#">14</a>
<a href="#">6.8.</a>	Application Data Types . . . . .	<a href="#">14</a>
<a href="#">7.</a>	Idle Sessions . . . . .	<a href="#">15</a>
<a href="#">8.</a>	Closing Sessions Due To An Error . . . . .	<a href="#">16</a>
<a href="#">9.</a>	Use over TLS . . . . .	<a href="#">17</a>
<a href="#">10.</a>	Update to <a href="#">RFC 3981</a> . . . . .	<a href="#">18</a>
<a href="#">11.</a>	IRIS Transport Mapping Definitions . . . . .	<a href="#">19</a>
<a href="#">11.1.</a>	URI Scheme . . . . .	<a href="#">19</a>
<a href="#">11.2.</a>	Application Protocol Label . . . . .	<a href="#">19</a>
<a href="#">12.</a>	Internationalization Considerations . . . . .	<a href="#">20</a>
<a href="#">13.</a>	IANA Considerations . . . . .	<a href="#">21</a>
<a href="#">13.1.</a>	XPC URI Scheme Registration . . . . .	<a href="#">21</a>
<a href="#">13.2.</a>	XPCS URI Scheme Registration . . . . .	<a href="#">21</a>
<a href="#">13.3.</a>	S-NAPTR XPC Registration . . . . .	<a href="#">22</a>
<a href="#">13.4.</a>	S-NAPTR XPCS Registration . . . . .	<a href="#">22</a>
<a href="#">13.5.</a>	Well-known TCP Port Registration for XPC . . . . .	<a href="#">22</a>
<a href="#">13.6.</a>	Well-known TCP Port Registration for XPCS . . . . .	<a href="#">23</a>
<a href="#">14.</a>	Security Considerations . . . . .	<a href="#">24</a>
<a href="#">14.1.</a>	Security Mechanisms . . . . .	<a href="#">24</a>
<a href="#">14.2.</a>	SASL Compliance . . . . .	<a href="#">25</a>
<a href="#">15.</a>	Normative References . . . . .	<a href="#">26</a>
<a href="#">Appendix A.</a>	Examples . . . . .	<a href="#">27</a>
<a href="#">Appendix B.</a>	Contributors . . . . .	<a href="#">35</a>
	Author's Address . . . . .	<a href="#">36</a>
	Intellectual Property and Copyright Statements . . . . .	<a href="#">37</a>

Newton

Expires September 6, 2007

[Page 2]

## **1. Introduction**

Using S-NAPTR [5], IRIS has the ability to define the use of multiple application transports (or transfer protocols) for different types of registry services, all at the descretion of the server operator. The TCP transfer protocol defined in this document is completely modular and may be used by any registry types.

This transfer protocol defines simple framing for sending XML in chunks so that XML fragments may be acted upon (or pipelined) before the reception of the entire XML instance. This document calls this XML pipelining with chunks (XPC) and its use with IRIS as IRIS-XPC.

XPC is for use with simple request and response interactions between clients and servers. Clients send a series of requests to a server in data blocks. The server will respond to each data block individually with a corresponding data block, but through the same connection. Request and response data blocks are sent using the TCP SEND function and received using the TCP RECEIVE function.

The lifecycle of an XPC session has the following phases:

1. A client establishes a TCP connection with a server.
2. The server sends a connection response block (CRB).
3. The client sends a request block (RQB). In this request, the client can set a "keep open" flag requesting that the server keep the XPC session open following the response to this request.
4. The server responds with a response block (RSB). In this response, the server can indicate to the client whether or not the XPC session will be closed.
5. If the XPC session is not to be terminated, then the lifecycle repeats from step 3.
6. The TCP connection is closed.



## **2. Document Terminology**

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119](#) [8].

Octet fields with numeric values are given according to the conventions in [RFC 1166](#) [10]: the left most bit of the whole field is the most significant bit; when a multi-octet quantity is transmitted the most significant octet is transmitted first. Bits signifying flags in an octet are numbered according to the conventions of [RFC 1166](#) [10]: bit 0 is the most significant bit and bit 7 is the least significant bit. When a diagram describes a group of octets, the order of transmission for the octets starts from the left.



### 3. Request Block (RQB)

The format for the request block (RQB) is as follows:

	+-----+	+-----+	+-----+	+-----+
field	header	authority	authority	chunks 1..n
		length		
	+-----+	+-----+	+-----+	+-----+
octets	1	1	0..255	variable

Request Block

These fields have the following meanings:

- o header - as described in [Section 5](#).
- o authority length - the length of the authority field in this request block
- o authority - a string of octets describing the authority against which this request is to be executed. See [\[1\]](#) for the definition and description of an authority. The number of octets in this string MUST be no more and no less than the number specified by the authority length.
- o chunks 1..n - the request data broken into chunks ([Section 6](#)).





## 4. Response Blocks

There are two types of blocks used by a server to respond to a client. The first type is a response block (RSB) defined in [Section 4.1](#). It is used by a server to respond to request blocks (RQB). The second type is a specialized version of a response block called a connection response block (CRB) defined in [Section 4.2](#). It is sent by a server to a client when a connection is established to initiate protocol negotiation. Conceptually, a CRB is a type of RQB; they share the same format, but a CRB is constrained in conveying only specific information and is only sent at the beginning of the session lifecycle.

### 4.1. Response Block (RSB)

The format for the response block (RSB) is as follows:

```

      +-----+-----+
field | header | chunks 1..n |
      |         |         |
      +-----+-----+
octets   1       variable

```

Response Block

These fields have the following meanings:

- o header - as described in [Section 5](#).
- o chunks 1..n - the response data broken into chunks ([Section 6](#)).

Servers SHOULD NOT send an RSB to a client until they have received the entire RQB. Servers that do begin sending an RSB before the reception of the entire RQB must consider that clients will not be expected to start processing the RSB until they have fully sent the RQB, and that the RSB will may fill the clients TCP buffers.

### 4.2. Connection Response Block (CRB)

A connection response block (CRB) is a response block sent by a server to a client in response to the client initiating a session. A connection response block has the same format as a response block (RSB) ([Section 4.1](#)). The only difference is that it is constrained in one of two ways:

1. It contains only one chunk (see [Section 6](#)) containing version information (see [Section 6.2](#)) and the keep-open (or KO) flag in the block header (see [Section 5](#)) has a value of 1 (meaning the



connection is not closing). Servers MUST use this type of CRB to indicate service availability.

2. It contains only one chunk (see [Section 6](#)) containing a system error (see 'system-error' under [Section 6.4](#)) and the keep-open (or KO) flag in the block header (see [Section 5](#)) has a value of 0 (meaning the server will close the connection immediately after sending the CRB). Servers MUST use this type of CRB when they can accept connections but cannot process requests.

## 5. Block Header

Each data block starts with a one octet header called the block header. This header has the same format for both request and response data blocks, though some of the bits in the header only have meaning in one type of data block. The bits are ordered according to the convention given in [RFC 1166](#) [10], where bit 0 is the most significant bit and bit 7 is the least significant bit. Each bit in the block header has the following meaning:

- o bits 0 and 1 - version (V field) - If 0 (both bits are zero), the protocol is the version defined in this document. Otherwise, the rest of the bits in the header and the block may be interpreted as another version. If a server receives a request for a version it does not support, it SHOULD follow the behavior described in [Section 8](#).
- o bits 2 - keep open (KO flag) - This flag is used to request that a connection stay open by a client and to indicate that a connection will stay open by a server, depending on the type of block. In a request block (RQB): a value of 1 indicates that a client is requesting that the server not close the TCP session, and a value of 0 indicates the client will expect the server to close the TCP session immediately after sending the corresponding response. In a response block (RSB) or a connection response block (CRB): a value of 1 indicates that the server expects the client to keep the TCP session open for the server to receive another request, and a value of 0 indicates that the server expects the client to close the TCP session immediately following this block.
- o bit 3, 4, 5, 6, and 7 - reserved - These MUST be 0. If a server receives a request in which any of these bits is set to 1 and the server does not understand the purpose for the value, the server SHOULD follow the behavior described in [Section 8](#).

	+-----+-----+-----+	
field	Version   Keep Open   reserved	
	(V)   (KO)	
	+-----+-----+	
bits	0 and 1	2 3 - 7

Block Header



## 6. Chunks

Request and response blocks break the request and response XML data down into chunks. Request and response blocks MUST always have a minimum of 1 chunk. Each chunk has a one octet descriptor. The first bit of the descriptor determines if chunk is the last chunk in the block.

The bits of the chunk descriptor octet are ordered according to convention given in [RFC 1166](#) [10], where bit 0 is the most significant bit and bit 7 is the least significant bit. The bits of the chunk descriptor octet have the following meaning:

- o bit 0 - last chunk (LC flag) - If 1, this chunk is the last chunk in the block.
- o bit 1 - data complete (DC flag) - If 1, the data in this chunk represents the end of the data for the chunk type given. If this bit is never set to 1 in any chunk descriptor for chunks of the same type in a block, clients and servers MUST NOT assume the data will continue in another block. If the block transitions from one type of chunk to another with out signaling completion of the data, clients and servers MUST assume that the remaining data will not be sent in a remaining chunk.
- o bits 2, 3, and 4 - reserved - These MUST be 0.
- o bit 5, 6, and 7 - chunk type (CT field) - determines the type of data carried in the chunk. These are the binary values for the chunk types:
  - \* 000 - no data or 'nd' type (see [Section 6.1](#))
  - \* 001 - version information or 'vi' type (see [Section 6.2](#))
  - \* 010 - size information or 'si' type (see [Section 6.3](#))
  - \* 011 - other information or 'oi' type (see [Section 6.4](#))
  - \* 100 - SASL data or 'sd' type (see [Section 6.5](#))
  - \* 101 - authentication success information or 'as' type (see [Section 6.6](#))
  - \* 110 - authentication failure information or 'af' type (see [Section 6.7](#))





- \* 111 - application data or 'ad' type (see [Section 6.8](#))

	+-----+-----+-----+-----+			
field	Last Chunk   Data Complete   reserved   Chunk Type			
	(LC)   (DC)     (CT)			
	+-----+-----+-----+-----+			
bits	0	1	2 - 4	5 - 7

Chunk Descriptor

A block MAY have multiple types of chunks, but all chunks of the same type MUST be contiguous in a block and MUST be ordered in the block in the order in which their data is to be interpreted. Contiguous chunks must be ordered by type within a block in the following way:

1. authentication related chunks - either SASL data chunks (type 100), authentication success information chunks (type 101) or authentication failure information chunks (type 110), but not more than one type. During the setup of security mechanisms using these chunks, clients MUST NOT send subsequent requests until they have received either an authentication success or failure chunk.
2. data chunks - either no data chunks (type 000) or application data chunks (type 111), but not both.
3. information chunks - either version information (type 001) or other information (type 011), but not both.

A block MUST have at least one type of the above chunks.

The format for a chunk is as follows:

	+-----+-----+-----+		
field	chunk	chunk data	chunk
	descriptor	length	data
	+-----+-----+-----+		
octets	1	2	variable

chunk

These fields have the following meanings:

- o chunk descriptor - as described above.
- o chunk data length - the length of the data of the chunk



- o chunk data - the data of the chunk

### **6.1. No Data Types**

Servers and clients MUST ignore data in chunk types labeled no data. There is no requirement for these types of chunks to be zero length. A client MAY send "no data" to a server, and the server MUST respond with either a chunk of the same type or other information ([Section 6.4](#)).

### **6.2. Version Information Types**

Chunks of this type contain XML conformant to the schema specified in [\[9\]](#) and MUST have the <versions> element as the root element.

In the context of IRIS-XPC, the protocol identifiers for these elements are as follows:

- o <transferProtocol> - the value "iris.xpc1" to indicate the protocol specified in this document.
- o <application> - the XML namespace identifier for IRIS [\[1\]](#).
- o <dataModel> - the XML namespace identifier for IRIS registries.

In the context of IRIS-XPC, the authentication mechanism identifiers are the SASL mechanism names found in the IANA SASL mechanism registry defined by [RFC 4422](#) [\[11\]](#).

This document defines no extension identifiers.

Clients MAY send a block with this type of chunk to a server. These chunks SHOULD be zero length and servers MUST ignore any data in them. When a server receives a chunk of this type, it MUST respond with a chunk of this type. This interchange allows a client to query the version information of a server.

The definition of octet size for the 'requestSizeOctets' and 'responseSizeOctets' attributes of the <transferProtocol> element are defined in [Section 6.3](#).

### **6.3. Size Information Types**

Chunks of this type contain XML conformant to the schema specified in IRIS-COMMON [\[9\]](#) and MUST have the <size> element as the root element.

Octet counts provided by this information are defined as the sum of the count of all chunk data of a particular chunk type. For



instance, if a XML instance is broken up into chunks of 20, 30, and 40 octets, the octet count would be 90 (20 + 30 + 40).

Clients MUST NOT send chunks of this type, and servers MAY close down a session using the procedure in [Section 8](#) if a chunk of this type is received.

#### **6.4. Other Information Types**

Chunks of this type contain XML conformant to the schema specified in IRIS-COMMON [\[9\]](#) and MUST have the <other> element as the root element.

The values for the 'type' attribute of <other> are as follows:

'block-error' - indicates there was an error decoding a block. Servers SHOULD send a block error in the following cases:

1. When a request block is received containing a chunk of this type.
2. When a request block is received containing authentication success (see [Section 6.6](#)) or authentication failure (see [Section 6.7](#)) information.
3. When a request block is received containing size information (see [Section 6.3](#)).
4. When reserved bits in the request block are 1.
5. When a block has not been received in its entirety and the TCP session has been idle for a specific period of time (i.e. a data block has been received but no terminating chunk for the data block has been recieved). Two minutes is RECOMMENDED for this timeout value. Note, there is a difference between an idle condition due to the incomplete reception of a data block and an idle condition between request/response transactions associated with keeping the session open. For the latter, see [Section 7](#).

'data-error' - indicates there was an error parsing data in chunks containing application or SASL data (e.g. XML is not valid in application data).

'system-error' - indicates that the receiver cannot process the request due to a condition not related to this protocol. Servers SHOULD send a system-error when they are capable of responding to requests but not capable of processing requests.



'authority-error' - indicates that the intended authority specified in the corresponding request is not served by the receiver. Servers SHOULD send an authority error when they receive a request directed to an authority other than those they serve.

'idle-timeout' - indicates that an XPC session has been idle for too long. Usage of this value is defined in [Section 7](#). Note, there is a difference between an idle condition due to the incomplete reception of a data block and an idle condition between request/response transactions associated with keeping the session open. For the former, see 'block-error' above.

Clients MUST NOT send chunks of this type, and servers MAY close down a session using the procedure in [Section 8](#) if a chunk of this type is received.

### 6.5. SASL Types

The SASL chunk type allows clients and servers to exchange SASL data.

The format for the data of this type of chunk is as follows:

	+-----+	+-----+	+-----+	+-----+
field	mechanism	mechanism	mechanism	mechanism
	name	name	data	data
	length		length	
	+-----+	+-----+	+-----+	+-----+
octets	1	variable	2	variable

#### SASL Authentication

These fields have the following meaning:

- o mechanism name length - the length of the SASL mechanism name
- o mechanism - the name of the SASL mechanism as registered in the IANA SASL mechanism registry defined by [\[11\]](#).
- o mechanism data length - the length of the SASL data
- o mechanism data - the data used for SASL

These fields MUST NOT span multiple chunks. Therefore it should be noted that SASL data length exceeding the length of the chunk minus the length of SASL profile name minus one is an error.

Depending on the nature of the SASL mechanism being used, SASL data





is sent from clients to servers and from servers to clients and may require multiple request/response transactions to complete. However, once a SASL exchange is complete and a server can determine authentication status, the server MUST send either authentication success information (see [Section 6.6](#)) or authentication failure information (see [Section 6.7](#)).

When used as an initial challenge response for SASL mechanisms that support such a feature, the mechanism data length may be set to a decimal value of 65,535 an absent initial response. A value of 0 indicates an empty initial response.

#### **6.6. Authentication Success Information Types**

Chunks of this type contain XML conformant to the schema specified in IRIS-COMMON [9] and MUST have the <authenticationSuccess> element as the root element.

This type of chunk is only sent from a server to a client. If a client sends it to a server, this will result in a block error (see 'block-error' in [Section 6.4](#)). The usage of this chunk type is defined in [Section 6.5](#). A server MAY close down a session due to reception of this type of chunk using the procedure in [Section 8](#).

SASL mechanisms may use the <data> child element to pass back arbitrary binary data as base 64 binary. The absence of this element indicates the absence of such data, where as the presence of the element with no content indicates an empty data set.

#### **6.7. Authentication Failure Information Types**

Chunks of this type contain XML conformant to the schema specified in IRIS-COMMON [9] and MUST have the <authenticationFailure> element as the root element.

This type of chunk is only sent from a server to a client. If a client sends it to a server, this will result in a block error (see 'block-error' in [Section 6.4](#)). The usage of this chunk type is defined in [Section 6.5](#). A server MAY close down a session due to reception of this type of chunk using the procedure in [Section 8](#).

#### **6.8. Application Data Types**

These chunks contain application data. For IRIS, these are IRIS [1] XML instances.



## **7. Idle Sessions**

An XPC session may become idle between request/response transactions. This can occur when a server honors a client's request to keep the TCP connection running (see the keep-open or KO flag in the block header ([Section 5](#))). Servers are not expected to allow XPC sessions remain idle between requests indefinitely.

Clients **MUST** use the keep-alive feature of TCP to keep the connection active during idle periods.

If a server has not received a request block after sending a response block (either RSB or CRB) and the TCP connection fails to keep-alive, it **SHOULD** do the following:

1. Send an unsolicited response block containing an idle timeout error (see 'idle-timeout' in [Section 6.4](#)) with the keep-open (or KO) flag in the block header ([Section 5](#)) set to a value of 0.
2. Close the TCP connection.



## **8. Closing Sessions Due To An Error**

If a server is to close a session due to an error, it SHOULD do the following:

1. Send a response block containing either a block-error or data-error (see [Section 6.4](#)) or version information (see [Section 6.2](#)) with the keep-open (or K0) flag in the block header ([Section 5](#)) set to a value of 0.
2. Close the TCP connection.

## **9. Use over TLS**

XPC may be tunneled over TLS [\[4\]](#) by establishing a TLS session immediately after a TCP session is opened and before any blocks are to be sent. This type of session is known as XPCS.

When using TLS, a convention must be established to allow a client to authenticate the validity of a server. XPCS uses the same convention as described by IRIS-BEEP [\[2\]](#).

TLS enables authentication and confidentiality.

Implementers should note that while XPC and XPCS have separate URI scheme names and S-NAPTR application protocol labels, both are identified with the same <transferProtocol> value in version information chunks (see [Section 6.2](#)).



**10.** Update to [RFC 3981](#)

[Section 6.2 of RFC 3981](#) [[1](#)] (IRIS-CORE) states that IRIS-BEEP [[2](#)] is the default transport for IRIS. This document revises [RFC 3981](#) and specifies IRIS-XPC as the default transport for IRIS. The TCP well-known port registration is specified in [Section 13.5](#).



## **11. IRIS Transport Mapping Definitions**

This section lists the definitions required by IRIS [[1](#)] for transport mappings.

### **11.1. URI Scheme**

See [Section 13.1](#) and [Section 13.2](#).

### **11.2. Application Protocol Label**

See [Section 13.3](#) and [Section 13.4](#).

## **12. Internationalization Considerations**

XML processors are obliged to recognize both UTF-8 and UTF-16 [[3](#)] encodings. Use of the XML defined by [[9](#)] MUST NOT use any other character encodings other than UTF-8 or UTF-16.

## **13. IANA Considerations**

### **13.1. XPC URI Scheme Registration**

URL scheme name: iris.xpc

URL scheme syntax: defined in [\[1\]](#).

Character encoding considerations: as defined in [RFC2396](#) [\[6\]](#).

Intended usage: identifies IRIS XML using chunks over TCP

Applications using this scheme: defined in IRIS [\[1\]](#).

Interoperability considerations: n/a

Security Considerations: defined in [Section 14](#).

Relevant Publications: IRIS [\[1\]](#).

Contact Information: Andrew Newton <andy@hxr.us>

Author/Change controller: the IESG

### **13.2. XPCS URI Scheme Registration**

URL scheme name: iris.xpcs

URL scheme syntax: defined in [\[1\]](#).

Character encoding considerations: as defined in [RFC2396](#) [\[6\]](#).

Intended usage: identifies IRIS XML using chunks over TLS

Applications using this scheme: defined in IRIS [\[1\]](#).

Interoperability considerations: n/a

Security Considerations: defined in [Section 14](#).

Relevant Publications: IRIS [\[1\]](#).

Contact Information: Andrew Newton <andy@hxr.us>

Author/Change controller: the IESG



### **13.3. S-NAPTR XPC Registration**

Application Protocol Label (see [\[5\]](#)): iris.xpc

Intended usage: identifies an IRIS server using XPC

Interoperability considerations: n/a

Security Considerations: defined in [Section 14](#).

Relevant Publications: IRIS [\[1\]](#).

Contact Information: Andrew Newton <andy@hxr.us>

Author/Change controller: the IESG

### **13.4. S-NAPTR XPCS Registration**

Application Protocol Label (see [\[5\]](#)): iris.xpcs

Intended usage: identifies an IRIS server using secure XPCS

Interoperability considerations: n/a

Security Considerations: defined in [Section 14](#).

Relevant Publications: IRIS [\[1\]](#).

Contact Information: Andrew Newton <andy@hxr.us>

Author/Change controller: the IESG

### **13.5. Well-known TCP Port Registration for XPC**

Protocol Number: TCP

TCP Port Number: TBD by IANA

Message Formats, Types, Opcodes, and Sequences: defined in [Section 4.2](#), [Section 3](#), and [Section 4.1](#).

Functions: defined in IRIS [\[1\]](#).

Use of Broadcast/Multicast: none

Proposed Name: IRIS over XPC

Short name: iris.xpc



Contact Information: Andrew Newton <andy@hxr.us>

### **13.6. Well-known TCP Port Registration for XPCS**

Protocol Number: TCP

TCP Port Number: TBD by IANA

Message Formats, Types, Opcodes, and Sequences: defined in [Section 9](#), [Section 4.2](#), [Section 3](#), and [Section 4.1](#).

Functions: defined in IRIS [[1](#)].

Use of Broadcast/Multicast: none

Proposed Name: IRIS over XPCS

Short name: iris.xpcs

Contact Information: Andrew Newton <andy@hxr.us>





## **14. Security Considerations**

Implementers should be fully aware of the security considerations given by IRIS [[1](#)] and TLS [[4](#)]. With respect to server authentication with the use of TLS, see [Section 6](#) of IRIS-BEEP [[2](#)].

### **14.1. Security Mechanisms**

Clients SHOULD be prepared to use the following security mechanisms in the following manner:

- o SASL/DIGEST-MD5 - for user authentication without the need of session encryption.
- o SASL/OTP - for user authentication without the need of session encryption.
- o TLS using the TLS\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA cipher - for encryption.
- o TLS using the TLS\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA cipher with client-side certificates - for encryption and user authentication.
- o TLS using the TLS\_RSA\_WITH\_AES\_128\_CBC\_SHA cipher - for encryption. See [[7](#)].
- o TLS using the TLS\_RSA\_WITH\_AES\_128\_CBC\_SHA cipher with client-side certificates - for encryption and user authentication. See [[7](#)].
- o TLS using the TLS\_RSA\_WITH\_AES\_256\_CBC\_SHA cipher - for encryption. See [[7](#)].
- o TLS using the TLS\_RSA\_WITH\_AES\_256\_CBC\_SHA cipher with client-side certificates - for encryption and user authentication. See [[7](#)].

Anonymous client access SHOULD be considered in one of two methods:

1. When no authentication has been used.
2. Using the SASL anonymous profile: SASL/ANONYMOUS

As specified by SASL/PLAIN, clients MUST NOT use the SASL/PLAIN mechanism without first encrypting the TCP session (e.g. such as with TLS). Clients MUST implement SASL/PLAIN and TLS using TLS\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA cipher.



## **14.2. SASL Compliance**

The following list details the compliance of IRIS-XPC for use with SASL, as specified by [RFC 4422](#) [[11](#)], Section 4.

1. The SASL service name to be used by IRIS-XPC is "iris-xpc".
2. [Section 6.2](#) describes the negotiation facility used to determine the available security mechanisms. This facility may be used both before the initiation of SASL exchanges and after the installation of security mechanisms.
3.
  - a) [Section 6.5](#) describes the mechanism to initiate authentication exchanges.
  - b) [Section 6.5](#) describes the mechanism to transfer server challenges and client responses.
  - c) [Section 6.6](#) and [Section 6.7](#) describe the mechanisms to indicate the outcome of an authentication exchange. [Section 6.6](#) describes how additional data may be carried with this message.
4. Non-empty authorization identity strings usign within IRIS-XPC MUST be normalized according to [RFC 4013](#) [[12](#)]. The semantics of the non-empty authorization identity strings is server dependent, and clients MUST use the values for these strings as given by configuration or the user.
5. Clients or servers wishing to abort an ongoing authentication exchange MUST close the connection.
6. After new security layers are negotiated, they take effect on the first octet following the authentication success (as) ([Section 6.6](#)) chunk sent by the server and on the first octet sent after receipt of the authentication success (as) chunk sent by the client.
7. IRIS-XPC can be used with both TLS and SASL. When used in combination, TLS MUST always be applied before any SASL mechanism.
8. IRIS-XPC does not support multiple SASL authentications. However, if TLS is being used in combination with SASL, TLS authentication MUST occur before any SASL authentication.



## **15. Normative References**

- [1] Newton, A. and M. Sanz, "Internet Registry Information Service", [RFC 3981](#), January 2004.
- [2] Newton, A. and M. Sanz, "Using the Internet Registry Information Service over the Blocks Extensible Exchange Protocol", [RFC 3983](#), January 2004.
- [3] The Unicode Consortium, "The Unicode Standard, Version 3", ISBN 0-201-61633-5, 2000, <The Unicode Standard, Version 3>.
- [4] Dierks, T., Allen, C., Treese, W., Karlton, P., Freier, A., and P. Kocher, "The TLS Protocol Version 1.0", [RFC 2246](#), January 1999.
- [5] Daigle, L. and A. Newton, "Domain-Based Application Service Location Using SRV RRs and the Dynamic Delegation Discovery Service (DDDS)", [RFC 3958](#), January 2005.
- [6] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifiers (URI): Generic Syntax", [RFC 2396](#), August 1998.
- [7] Chown, P., "Advanced Encryption Standard (AES) Ciphersuites for Transport Layer Security (TLS)", [RFC 3268](#), June 2002.
- [8] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), [BCP 14](#), March 1997.
- [9] Newton, A., "A Common Schema for Internet Registry Information Service Transfer Protocols", [draft-ietf-crips-iris-common-transport-00](#) (work in progress), April 2005.
- [10] Kirkpatrick, S., Stahl, M., and M. Recker, "Internet numbers", [RFC 1166](#), July 1990.
- [11] Melnikov, A. and K. Zeilenga, "Simple Authentication and Security Layer (SASL)", [RFC 4422](#), June 2006.
- [12] Zeilenga, K., "SASLprep: Stringprep Profile for User Names and Passwords", [RFC 4013](#), February 2005.



## [Appendix A](#). Examples

This section gives examples of IRIS-XPC sessions. Lines beginning with "C:" denote data sent by the client to the server, and lines beginning with "S:" denote data sent by the server to the client. Following the "C:" or "S:", the line either contains octet values in hexadecimal notation with comments or XML fragments. No line contains both octet values with comments and XML fragments. Comments are contained within parenthesis.

It should also be noted that flag values of "yes" and "no" reflect binary values 1 and 0.

The following example demonstrates an IRIS client issuing two requests in one XPC session. In the first request, the client is requesting status information for "example.com". This request and its response are transferred with one chunk. In the second request, the client is requesting status information for "milo.example.com", "felix.example.com", and "hobbes.example.com". This request and its response are transferred with three chunks.

```
S:          (connection response block)
S: 0x20      (block header: V=0,K0=yes)
S:          (chunk 1)
S: 0xC1      (LC=yes,DC=yes,CT=vi)
S: 0x01 0xBF (chunk length=447)
S:          (Version Information)
S: <?xml version="1.0"?>
S: <versions xmlns="urn:ietf:params:xml:ns:iris-transport">
S:   <transferProtocol protocolId="iris.xpc1"
S:     authenticationIds="PLAIN EXTERNAL">
S:     <application protocolId="urn:ietf:params:xml:ns:iris1"
S:       extensionIds="http://example.com/SIMPLEBAG">
S:       <dataModel protocolId="urn:ietf:params:xml:ns:dchk1"/>
S:       <dataModel protocolId="urn:ietf:params:xml:ns:dreg1"/>
S:     </application>
S:   </transferProtocol>
S: </versions>
```

```
C:          (request block)
C: 0x20      (block header: V=0,K0=yes)
C: 0x0B      (authority length=11)
C:          (authority="example.com")
C: 0x65 0x78 0x61 0x6D 0x70 0x6C 0x65 0x23 0x63 0x6F 0x6D
C:          (chunk 1)
C: 0xC7      (LC=yes,DC=yes,CT=ad)
C: 0x01 0x53 (chunk length=339)
C:          (IRIS XML request)
```





```
C: <request xmlns="urn:ietf:params:xml:ns:iris1"
C:   xsi:schemaLocation="urn:ietf:params:xml:ns:iris1 iris.xsd" >
C:   <searchSet>
C:     <lookupEntity
C:       registryType="urn:ietf:params:xml:ns:dchk1"
C:       entityClass="domain-name"
C:       entityName="example.com" />
C:   </searchSet>
C: </request>
```

```
S:           (response block)
S: 0x20       (block header: V=0,K0=yes)
S:           (chunk 1)
S: 0xC7       (LC=yes,DC=yes,CT=ad)
S: 0x01 0xE0  (chunk length=480)
S:           (IRIS XML response)
S: <iris:response xmlns:iris="urn:ietf:params:xml:ns:iris1">
S:   <iris:resultSet>
S:     <iris:answer>
S:       <domain authority="example.com" registryType="dchk1"
S:         entityClass="domain-name" entityName="example.com-1"
S:         temporaryReference="true"
S:         xmlns="urn:ietf:params:xml:ns:dchk1">
S:       <domainName>example.com</domainName>
S:       <status>
S:         <assignedAndActive/>
S:       </status>
S:     </domain>
S:   </iris:answer>
S: </iris:resultSet>
S: </iris:response>
```

```
C:           (request block)
C: 0x00       (block header: V=0,K0=no)
C: 0x0B       (authority length=11)
C:           (authority="example.com")
C: 0x65 0x78 0x61 0x6D 0x70 0x6C 0x65 0x23 0x63 0x6F 0x6D
C:           (chunk 1)
C: 0x07       (LC=no,DC=no,CT=ad)
C: 0x01 0x4E  (chunk length=339)
C:           (IRIS XML request)
C: <request xmlns="urn:ietf:params:xml:ns:iris1"
C:   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
C:   xsi:schemaLocation="urn:ietf:params:xml:ns:iris1 iris.xsd" >
C:   <searchSet>
C:     <lookupEntity
C:       registryType="urn:ietf:params:xml:ns:dchk1"
C:       entityClass="domain-name"
```



```
C:      entityName="milo.example.com" />
C:    </searchSet>
C:      (chunk 2)
C: 0x07      (LC=no,DC=no,CT=ad)
C: 0x00 0xA9 (chunk length=169)
C:      (IRIS XML request)
C:    <searchSet>
C:      <lookupEntity
C:        registryType="urn:ietf:params:xml:ns:dchk1"
C:        entityClass="domain-name"
C:        entityName="felix.example.com" />
C:    </searchSet>
C:      (chunk 3)
C: 0xC7      (LC=yes,DC=yes,CT=ad)
C: 0x00 0xB5 (chunk length=181)
C:      (IRIS XML request)
C:    <searchSet>
C:      <lookupEntity
C:        registryType="urn:ietf:params:xml:ns:dchk1"
C:        entityClass="domain-name"
C:        entityName="hobbes.example.com" />
C:    </searchSet>
C:  </request>

S:      (response block)
S: 0x00      (block header: V=0,K0=no)
S:      (chunk 1)
S: 0x07      (LC=no,DC=no,CT=ad)
S: 0x01 0xDA (chunk length=474)
S:      (IRIS XML response)
S: <iris:response xmlns:iris="urn:ietf:params:xml:ns:iris1">
S:   <iris:resultSet>
S:     <iris:answer>
S:       <domain authority="example.com" registryType="dchk1"
S:         entityClass="domain-name" entityName="milo.example.com-1"
S:         temporaryReference="true"
S:         xmlns="urn:ietf:params:xml:ns:dchk1">
S:       <domainName>milo.example.com</domainName>
S:       <status>
S:         <assignedAndActive/>
S:       </status>
S:     </domain>
S:   </iris:answer>
S: </iris:resultSet>
S:      (chunk 2)
S: 0x07      (LC=no,DC=no,CT=ad)
S: 0x01 0xA2 (chunk length=418)
S:      (IRIS XML response)
```



```

S: <iris:resultSet>
S:   <iris:answer>
S:     <domain authority="example.com" registryType="dchk1"
S:       entityClass="domain-name" entityName="felix.example.com-1"
S:       temporaryReference="true"
S:       xmlns="urn:ietf:params:xml:ns:dchk1">
S:       <domainName>felix.example.com</domainName>
S:       <status>
S:         <assignedAndActive/>
S:       </status>
S:     </domain>
S:   </iris:answer>
S: </iris:resultSet>
S:   (chunk 3)
S: 0xC7   (LC=yes, DC=yes, CT=ad)
S: 0x01 0xB5 (chunk length=437)
S:   (IRIS XML response)
S: <iris:resultSet>
S:   <iris:answer>
S:     <domain authority="example.com" registryType="dchk1"
S:       entityClass="domain-name"
S:     entityName="hobbes.example.com-1"
S:       temporaryReference="true"
S:       xmlns="urn:ietf:params:xml:ns:dchk1">
S:       <domainName>hobbes.example.com</domainName>
S:       <status>
S:         <assignedAndActive/>
S:       </status>
S:     </domain>
S:   </iris:answer>
S: </iris:resultSet>
S: </iris:response>

```

Figure 7: Example 1

In the following example, an IRIS client requests domain status information for "milo.example.com", "felix.example.com", and "hobbes.example.com" in one request. The request is sent with one chunk, however the answer is returned in three chunks.

```

S:   (connection response block)
S: 0x20   (block header: V=0, K0=yes)
S:   (chunk 1)
S: 0xC1   (LC=yes, DC=yes, CT=vi)
S: 0x01 0xBF (chunk length=447)
S:   (Version Information)
S: <?xml version="1.0"?>
S: <versions xmlns="urn:ietf:params:xml:ns:iris-transport">

```



```
S:  <transferProtocol protocolId="iris.xpc1"
S:    authenticationIds="PLAIN EXTERNAL">
S:    <application protocolId="urn:ietf:params:xml:ns:iris1"
S:      extensionIds="http://example.com/SIMPLEBAG">
S:        <dataModel protocolId="urn:ietf:params:xml:ns:dchk1"/>
S:        <dataModel protocolId="urn:ietf:params:xml:ns:dreg1"/>
S:      </application>
S:    </transferProtocol>
S:  </versions>

C:      (request block)
C: 0x00      (block header: V=0,K0=no)
C: 0x0B      (authority length=11)
C:      (authority="example.com")
C: 0x65 0x78 0x61 0x6D 0x70 0x6C 0x65 0x23 0x63 0x6F 0x6D
C:      (chunk 1)
C: 0xC7      (LC=yes,DC=yes,CT=ad)
C: 0x02 0xAB (chunk length=683)
C:      (IRIS XML request)
C: <request xmlns="urn:ietf:params:xml:ns:iris1"
C:   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
C:   xsi:schemaLocation="urn:ietf:params:xml:ns:iris1 iris.xsd" >
C:   <searchSet>
C:     <lookupEntity
C:       registryType="urn:ietf:params:xml:ns:dchk1"
C:       entityClass="domain-name"
C:       entityName="milo.example.com" />
C:   </searchSet>
C:   <searchSet>
C:     <lookupEntity
C:       registryType="urn:ietf:params:xml:ns:dchk1"
C:       entityClass="domain-name"
C:       entityName="felix.example.com" />
C:   </searchSet>
C:   <searchSet>
C:     <lookupEntity
C:       registryType="urn:ietf:params:xml:ns:dchk1"
C:       entityClass="domain-name"
C:       entityName="hobbes.example.com" />
C:   </searchSet>
C: </request>

S:      (response block)
S: 0x00      (block header: V=0,K0=no)
S:      (chunk 1)
S: 0x07      (LC=no,DC=no,CT=ad)
S: 0x01 0xDA (chunk length=474)
S:      (IRIS XML response)
```





```
S: <iris:response xmlns:iris="urn:ietf:params:xml:ns:iris1">
S:   <iris:resultSet>
S:     <iris:answer>
S:       <domain authority="example.com" registryType="dchk1"
S:         entityClass="domain-name" entityName="milo.example.com-1"
S:         temporaryReference="true"
S:         xmlns="urn:ietf:params:xml:ns:dchk1">
S:         <domainName>milo.example.com</domainName>
S:         <status>
S:           <assignedAndActive/>
S:         </status>
S:       </domain>
S:     </iris:answer>
S:   </iris:resultSet>
S:     (chunk 2)
S: 0x07      (LC=no, DC=no, CT=ad)
S: 0x01 0xA2 (chunk length=418)
S:          (IRIS XML response)
S: <iris:resultSet>
S:   <iris:answer>
S:     <domain authority="example.com" registryType="dchk1"
S:       entityClass="domain-name" entityName="felix.example.com-1"
S:       temporaryReference="true"
S:       xmlns="urn:ietf:params:xml:ns:dchk1">
S:       <domainName>felix.example.com</domainName>
S:       <status>
S:         <assignedAndActive/>
S:       </status>
S:     </domain>
S:   </iris:answer>
S: </iris:resultSet>
S:     (chunk 3)
S: 0xC7      (LC=yes, DC=yes, CT=ad)
S: 0x01 0xB5 (chunk length=437)
S:          (IRIS XML response)
S: <iris:resultSet>
S:   <iris:answer>
S:     <domain authority="example.com" registryType="dchk1"
S:       entityClass="domain-name"
S:       entityName="hobbes.example.com-1"
S:       temporaryReference="true"
S:       xmlns="urn:ietf:params:xml:ns:dchk1">
S:       <domainName>hobbes.example.com</domainName>
S:       <status>
S:         <assignedAndActive/>
S:       </status>
S:     </domain>
S:   </iris:answer>
```



```
S:    </iris:resultSet>
S: </iris:response>
```

Figure 8: Example 2

In the following example, an IRIS client sends a request containing SASL/PLAIN authentication data and a domain status check for "example.com". The server responds with authentication success information and the domain status of "example.com". Note that the client requests that the connection stay open for further requests, but that the server does not honor this request.

```
S:          (connection response block)
S: 0x20      (block header: V=0,K0=yes)
S:          (chunk 1)
S: 0xC1      (LC=yes,DC=yes,CT=vi)
S: 0x01 0xBF (chunk length=447)
S:          (Version Information)
S: <?xml version="1.0"?>
S: <versions xmlns="urn:ietf:params:xml:ns:iris-transport">
S:   <transferProtocol protocolId="iris.xpc1"
S:     authenticationIds="PLAIN EXTERNAL">
S:     <application protocolId="urn:ietf:params:xml:ns:iris1"
S:       extensionIds="http://example.com/SIMPLEBAG">
S:       <dataModel protocolId="urn:ietf:params:xml:ns:dchk1"/>
S:       <dataModel protocolId="urn:ietf:params:xml:ns:dreg1"/>
S:     </application>
S:   </transferProtocol>
S: </versions>

C:          (request block)
C: 0x00      (block header: V=0,K0=no)
C: 0x0B      (authority length=11)
C:          (authority="example.com")
C: 0x65 0x78 0x61 0x6D 0x70 0x6C 0x65 0x23 0x63 0x6F 0x6D
C:          (chunk 1)
C: 0x44      (LC=no,DC=yes,CT=sd)
C: 0x00 0x11 (chunk length=11)
C:          (SASL data)
C: 0x05      (mechanism length=5)
C:          (mechanism name="PLAIN")
C: 0x50 0x4C 0x41 0x49 0x43
C: 0x00 0x0A (sasl PLAIN data length=10)
C:          (sasl PLAIN data: authcid="bob")
C:          (sasl PLAIN data: authzid=NULL)
C:          (sasl PLAIN data: password="kEw1")
C: 0x62 0x6F 0x62 0x20 0x00 0x20 0x6B 0x45 0x77 0x31
C:          (chunk 2)
```



```
C: 0xC7      (LC=yes,DC=yes,CT=ad)
C: 0x01 0x53 (chunk length=339)
C:          (IRIS XML request)
C: <request xmlns="urn:ietf:params:xml:ns:iris1"
C:   xsi:schemaLocation="urn:ietf:params:xml:ns:iris1 iris.xsd" >
C:   <searchSet>
C:     <lookupEntity
C:       registryType="urn:ietf:params:xml:ns:dchk1"
C:       entityClass="domain-name"
C:       entityName="example.com" />
C:   </searchSet>
C: </request>

S:          (response block)
S: 0x00      (block header: V=0,K0=no)
S:          (chunk 1)
S: 0x45      (LC=no,DC=yes,CT=as)
S: 0x00 0xD0 (chunk length=208)
S:          (authentication success response)
S: <?xml version="1.0"?>
S: <authenticationSuccess
S:   xmlns="urn:ietf:params:xml:ns:iris-transport">
S:   <description language="en">
S:     user 'bob' authenticates via password
S:   </description>
S: </authenticationSuccess>
S:          (chunk 2)
S: 0xC7      (LC=yes,DC=yes,CT=ad)
S: 0x01 0xE0 (chunk length=480)
S:          (IRIS XML response)
S: <iris:response xmlns:iris="urn:ietf:params:xml:ns:iris1">
S:   <iris:resultSet>
S:     <iris:answer>
S:       <domain authority="example.com" registryType="dchk1"
S:       entityClass="domain-name" entityName="example.com-1"
S:       temporaryReference="true"
S:       xmlns="urn:ietf:params:xml:ns:dchk1">
S:       <domainName>example.com</domainName>
S:       <status>
S:         <assignedAndActive/>
S:       </status>
S:     </domain>
S:   </iris:answer>
S: </iris:resultSet>
S: </iris:response>
```

Figure 9: Example 3



**Appendix B. Contributors**

Substantive contributions to this document have been provided by the members of the IETF's CRISP Working Group, especially Robert Martin-Legene, Milena Caires, and David Blacka.

Author's Address

Andrew L. Newton  
VeriSign, Inc.  
21345 Ridgetop Circle  
Sterling, VA 20166  
USA

Phone: +1 703 948 3382

Email: [andy@hxr.us](mailto:andy@hxr.us)

URI: <http://www.verisignlabs.com/>



## Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org).

## Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

