ICNRG D. Oran
Internet-Draft Network Systems Research and Design
Updates: 8569, 8609 (if approved) D. Kutscher
Intended status: Experimental University of Applied Sciences Emden/Leer
Expires: 4 October 2020 2 April 2020

Reflexive Forwarding for CCNx and NDN Protocols
draft-oran-icnrg-reflexive-forwarding-00

Abstract

Current Information-Centric Networking protocols such as CCNx and NDN have a wide range of useful applications in content retrieval and other scenarios that depend only on a robust two-way exchange in the form of a request and response (represented by an _Interest-Data exchange_ in the case of the two protocols noted above). A number of important applications however, require placing large amounts of data in the Interest message, and/or more than one two-way handshake. While these can be accomplished using independent Interest-Data exchanges by reversing the roles of consumer and producer, such approaches can be both clumsy for applications and problematic from a state management, congestion control, or security standpoint. This specification proposes a _Reflexive Forwarding_ extension to the CCNx and NDN protocol architectures that eliminates the problems inherent in using independent Interest-Data exchanges for such applications. It updates RFC8569 and RFC8609.

Table of Contents

1. Introduction

Current ICN protocols such as CCNx [RFC8569] and [NDN] have a wide range of useful applications in content retrieval and other scenarios that depend only on a robust two-way exchange in the form of a request and response. These ICN architectures use the terms "consumer" and "producer" for the respective roles of the requester and the responder, and the protocols directly capture the mechanics of the two-way exchange through the "Interest message" carrying the request, and the "Data message" carrying the response. Through these constructs, the protocols are heavily biased toward a pure _pull-based_ interaction model where requests are small (carrying little or no user-supplied data other than the name of the requested data object), and responses are relatively large - up to an architecture-defined maximum transmission unit (MTU) on the order of kilobytes or tens of kilobytes.

A number of important applications however require interaction models more complex than individual request/response interactions in the same direction (i.e. between the same consumer and one or more producers). Among these we identify three important classes which are the target of the proposed enhancements defined in this specification. These are described in the following paragraphs.

Remote Method Invocation (RMI, aka RPC): When invoking a remote method, it is common for the method to require arguments supplied by the caller. In conventional TCP/IP style protocols like CORBA or HTTP "Post", these are pushed to the server as part of the message or messages that comprise the request. In ICN-style protocols there is an unattractive choice between inflating the request initiation with pushed arguments, or arranging to have one or more independent request/responses in the opposite direction for the server to fetch the arguments. Both of these approaches have substantial disadvantages. Recently, a viable alternative emerged through the work on RICE [Krol2018] which pioneered the main design elements proposed in this specification.

Phone-Home scenario: Applications in sensing, Internet-of-things (IoT) and other types where data is produced unpredictably and needs to be _pushed_ somewhere create a conundrum for the pure pull-based architectures considered here. If instead one eschews relaxing the size asymmetry between requests and responses, some additional protocol machinery is needed. Earlier efforts in the ICN community have recognized this issue and designed methods to provoke a cooperating element to issue a request to return the

data the originator desires to push, essentially "phoning home" to get the responder to fetch the data. One that has been explored to some extent is the _Interest-Interest-Data_ exchange [Carzaniga2011], where an Interest is sent containing the desired request as encapsulated data. CCNx-1.0 Bidirectional Streams [Mosko2017] are also based on a scheme where an Interest is used to signal a name prefix that a consumer has registered for receiving Interests from a peer in a bidirectional streaming session.

Peer state synchronization: A large class of applications, typified by those built on top of on reliable order-preserving transport protocols, require initial state synchronization between the peers. This is accomplished with a three-way (or longer) handshake, since employing a two-way handshake as provided in the existing NDN and CCNx protocols exposes a number of well-know hazards, such as _half-open connections_. When attempted for security-related operations such as key exchange, additional hazards such as _man-in-the-middle_ attacks become trivial to mount. Existing alternatives, similar to those used in the two examples above, instead utilize either overlapping Interest-Data exchanges in opposite directions (resulting in a four-way handshake) or by adding initialization data to the initial request and employing an Interest-Interest-Data protocol extension as noted in the Phone-home scenarios above.

All of the above application interaction models present interesting challenges, as neither relaxing the architecture to support pushing large amounts of data, nor introducing substantial complexities through multiple independent Interest-Data exchanges is an attractive approach. The following subsections provide further background and justification for why push and/or independent exchanges are problematical.

1.1. Problems with pushing data

There are two substantial problems with the simple approach of just allowing arbitrary amounts of data to be included with requests. These are:

1. In ICN protocols, Interest messages are intended to be small, on the order the size of a TCP ACK, as opposed to the size of a TCP data segment. This is because the hop-by-hop congestion control and forwarder state management requires Interest messages to be buffered in expectation of returning data, and possibly retransmitted hop-by-hop as opposed to end-to-end. In addition, the need to create and manage state on a per-Interest basis is substantially complicated if requests in Interest messages are

larger than a Path MTU (PMTU) and need to be fragmented hop-by-hop.

2. If the payload data of a request is used for invoking a computation (as in the RMI case described above) then substantial bandwidth can be wasted if the computation is either refused or abandoned for any number of reasons, including the requestor failing an authorization check, or the responder not having sufficient resources to execute the associated computation.

These problems also exist in pure datagram transport protocols such as those used for legacy RMI applications like NFS [RFC7530]. More usual are application protocols like HTTP(s) which rely on the TCP or QUIC 3-way handshake to establish a session and then have congestion control and segmentation provided as part of the transport protocol, further allowing sessions to be rejected before large amounts of data are transmitted or significant computational resources expended.

1.2. Problems with utilizing independent exchanges

In order to either complete a three-way handshake, or fetch data via a pull from the original requestor, the role of consumer and producer need to be reversed and an Interest/Data exchange initiated in the direction opposite of the initiating exchange. When done with an independent Interest/Data request and response, a number of complications ensue. Among them are:

1. The originating consumer needs to have a routable name prefix that can be used for the exchange. This means the consumer must arrange to have its name prefix propagated in the ICN routing system with sufficient reach that the producer issuing the interest can be assured it is routed appropriately. While some consumers are generally online and act as application servers, justifying the maintenance of this routing information, many do not. Further, in mobile environments, a pure consumer that does not need to have a routable name prefix can benefit from the inherent consumer mobility support in the CCNx and NDN protocols. By requiring a routable name prefix, extra mobile routing machinery is needed, such as that proposed in KITE [Zhang2018] or MAPME [Auge2018].

2. The consumer name prefix in item (1) above must be communicated to the producer as a payload, name suffix, or other field of the initiating Interest message. Since this name in its entirety is chosen by the consumer, it is highly problematic from a security standpoint, as it can recruit the producer to mount a reflection attack against the consumer's chosen victim.

3. The correlation between the exchanges in opposite directions must be maintained by both the consumer and the producer as independent state, as opposed to being architecturally tied together as would be the case with a conventional 3-way handshake finite state machine. While this can of course be accomplished with care by both parties, experience has shown that it is error prone (for example see the checkered history of interactions between the SIP [RFC3261] and SDP Offer-Answer [RFC6337]) protocols. When employed as the wrapper for a key management protocol such as with TLS [RFC8446] state management errors can be catastrophic for security.

2. Requirements Language

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

3. Overview of the Reflexive Forwarding design

This specification defines a _Reflexive Forwarding_ extension to CCNx and NDN that avoids the problems enumerated in Sections 1.1 and 1.2. It straightforwardly exploits the hop-by-hop state and symmetric routing properties of the current protocols.

Figure 1 below illustrates a canonical NDN/CCNx forwarder with its conceptual data structures of the Content Store (CS), Pending Interest Table (PIT) and Forwarding Information Base (FIB). The key observation involves the relation between the PIT and the FIB. Upon arrival of an Interest, a PIT entry is created which contains state recording the incoming interface on which the Interest. If the Interest is not immediately satisfied by cached data in the CS, the forwarder looks up the name in the FIB to ascertain the _next-hop_ to propagate the Interest onward upstream toward the named producer. Therefore, a chain of forwarding state is established during Interest forwarding that couples the PIT entries of the chain of forwarders together conceptually as _breadcrumbs_. These are used to forward the returning Data Message over the inverse path through the chain of forwarders until the Data message arrives at the originating consumer. The state in the PITs is _unwound_ by destroying it as each PIT entry is _satisfied_. This behavior is *critical* to the feasibility of the reflexive forwarding design we propose.

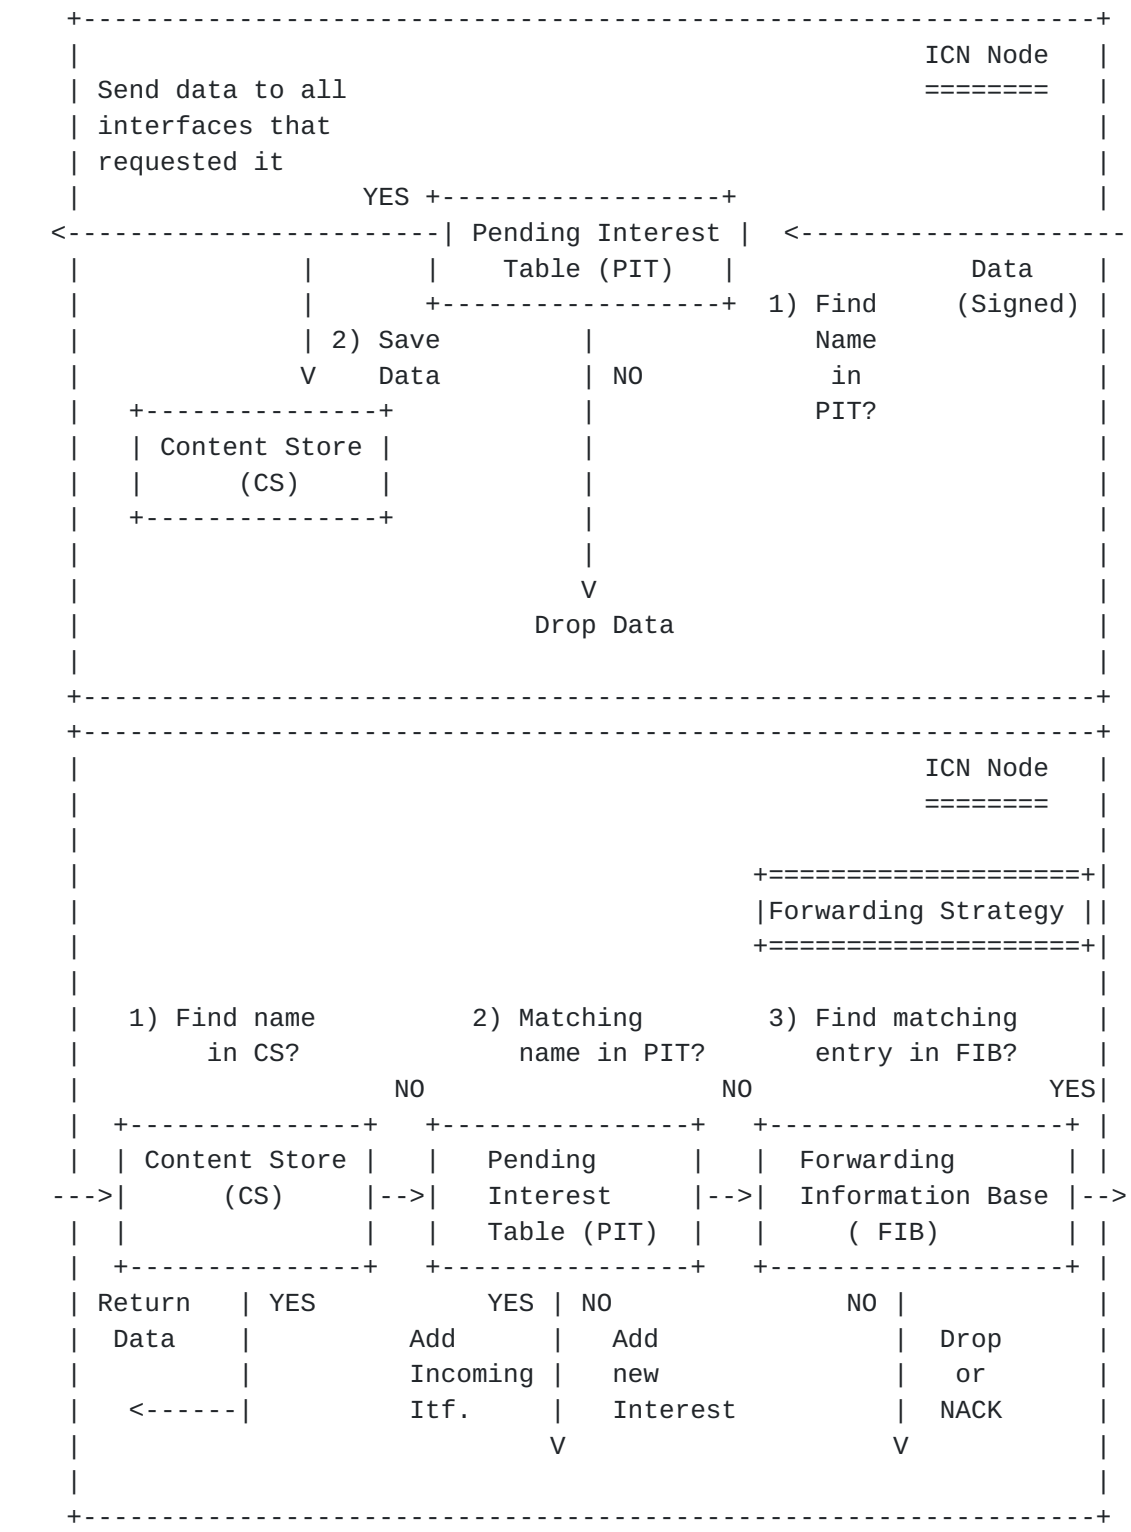

Figure 1: ICN forwarder structure

Given the above forwarding properties for Interests, it should be clear that while an Interest is outstanding and ultimately arrives at

a producer who can respond to it, there is sufficient state in the chain of forwarders to route not just a returning Data message, but potentially another Interest directed through the inverse path to the unique consumer who issued the original Interest. (Section 8.1.3 describes how Interest aggregation interacts with this scheme.) The key question therefore is how to access this state in a way that it can be used to forward Interests.

In order to achieve this _Reflexive Interest_ forwarding on the inverse path recorded in the PIT of each forwarder, we need a few critical design elements. These are as follows:

1. The Reflexive Interest needs to have a Name. This name is what the originating consumer will use to match against the Data object (or objects - more on this later) it wishes the producer to fetch by issuing the Reflexive Interest. This cannot be just any name, but needs to essentially name the state already recorded in the PIT and not allow the consumer to manufacture an arbitrary name and mount a reflection attack as pointed out in Section 1.2, Paragraph 2, Item 2.

2. There has to be a FIB entry at each forwarder for this name prefix so that when the reflexive interest arrives, the forwarder can forward it downstream toward the originating consumer. This FIB entry points directly to the incoming interface on which the corresponding original Interest arrived. The FIB entry needs to be created as part of the forwarding of the original Interest so that it is available in time to catch any reflexive Interest issued by the producer. It usually makes sense to destroy this FIB entry when the Data message satisfying the original Interest arrives since this avoids any dangling stale state. Given the deign details documented later in this specification, stale FIB state does not represent a correctness hazard and hence can be done lazily if desired in an implementation. See Section 5 for more details on FIB operation considerations.

3. There has to be coupling of the state between the originating Interest-Data exchange and the enclosed Reflexive Interest-Data exchange at both the consumer and the producer. In our design, this accomplished by the way reflexive interest names are chosen.

The following sections provide the normative details on each of these design elements. The overall interaction flow for reflexive forwarding is illustrated below in Figure 2.

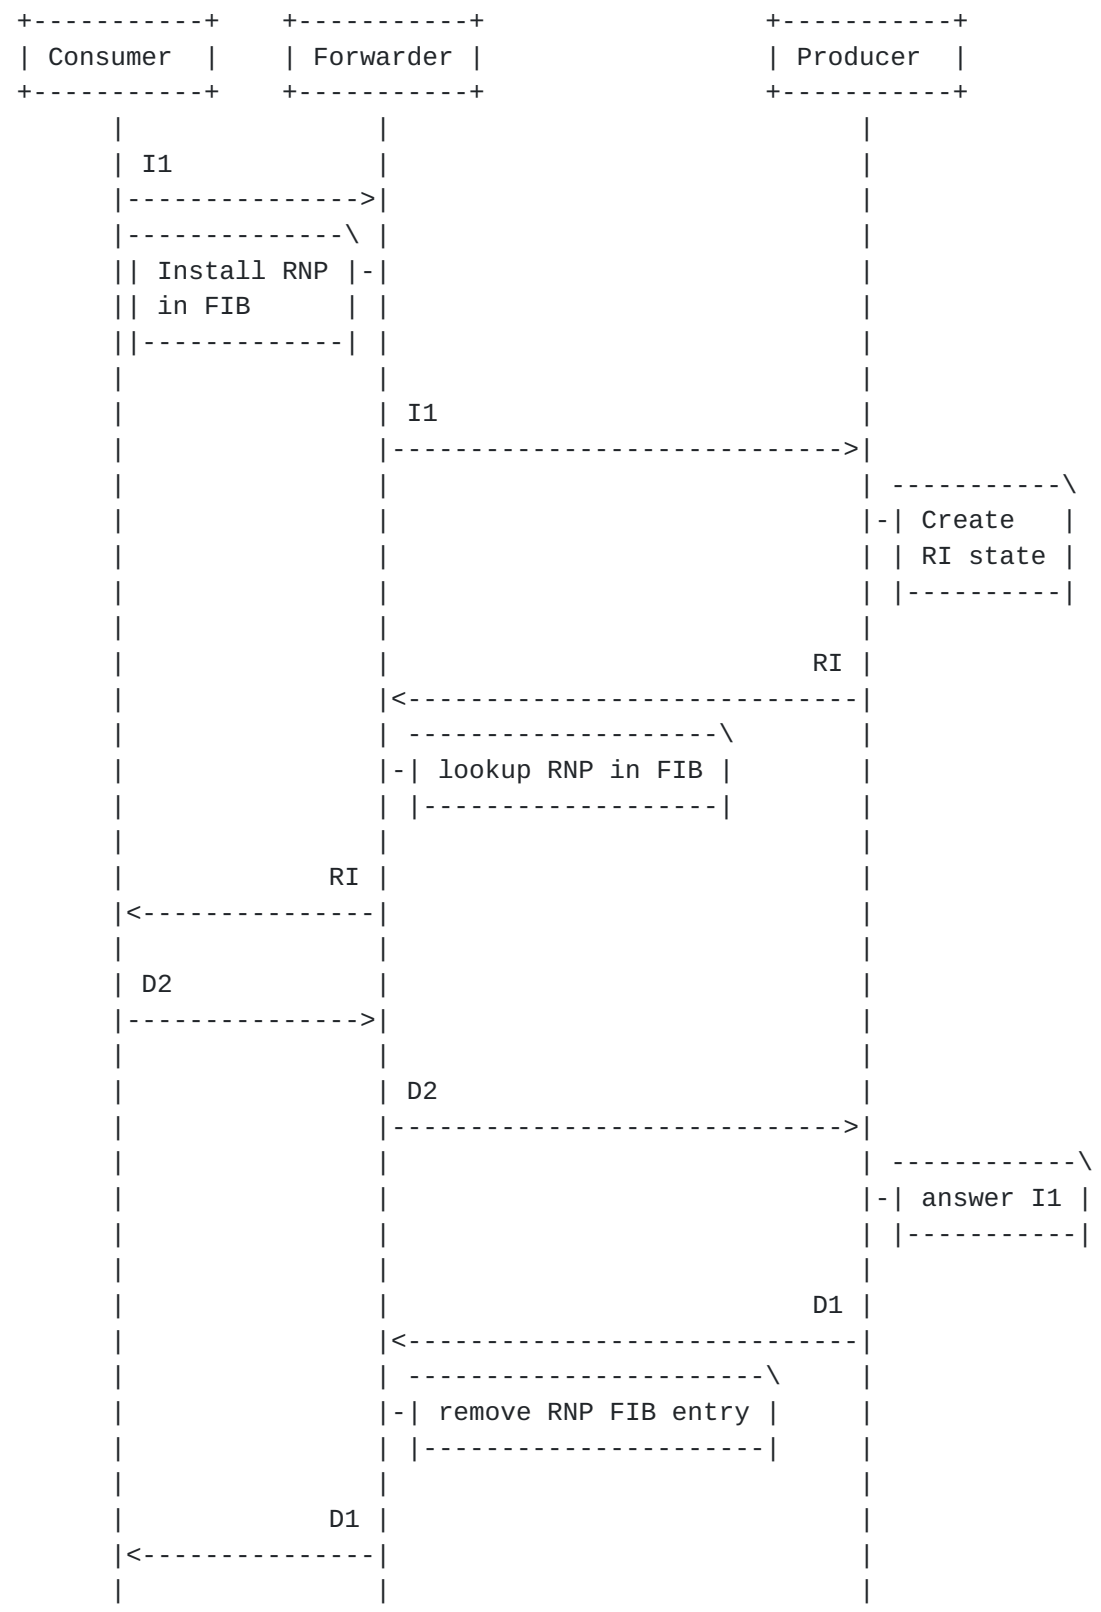
Consumer
Forwarder
Producer
I1
Install RNP in FIB
I1
Create RI state
RI
lookup RNP in FIB
RI
D2
D2
answer I1
D1
remove RNP FIB entry
D1

Legend:
I1: Interest #1 containing the Reflexive Name Prefix TLV
RI: Reflexive Interest with Reflexive Name Prefix Component
RNP: Reflexive Name Prefix
D1: Data message, answering initiating I1 Interest
D2: Data message, answering RI

Figure 2: Message Flow Overview

4. Naming of Reflexive Interests

A consumer may have one or more objects for the producer to fetch, and therefore needs to communicate enough information in their initial Interest to allow the producer to construct properly formed reflexive Interest names. For some applications the set of _full names_ (see [I-D.irtf-icnrg-terminology]) is known a priori, for example through compile time bindings of arguments in interface definitions or by the architectural definition of a simple sensor reading. In other cases the full names of the individual objects must be communicated in the original Interest message. In all cases enough state must be provided by the consumer for the forwarders to construct a FIB entry (as noted in Section 3, Paragraph 6, Item 2). This is accomplished through the following naming construct.

We define a new typed name component, identified by a registered name component type in the IANA registry for [RFC8569]. We call this the _Reflexive Interest Name Component type_. It MUST be the first (i.e. high order) name component of any Reflexive Interest issued by a producer. Its value is a random 64 bit number, assigned by the consumer, which provides the entropy required to uniquely identify the issuing consumer for the duration of any outstanding Interest-Data exchange. The consumer SHOULD choose a different random value for each Interest message it constructs, for two reasons:

1. If stale FIB sate is present, the randomness prevents potential mis-routing of reflexive interests (see Section 8.1.1 below for more details), and

2. Re-use of the same reflexive interest name over multiple interactions might reveal linkability information that could be used by surveillance adversaries for tracking purposes.

This initial name component is either communicated by itself through a _Reflexive Name Prefix TLV_ in the originating Interest, or prepended to any object names the consumer wishes the producer to fetch explicitly where there is more than one object needed by the producer for the current Interest-Data interaction. There are four cases to consider:

1. The reflexive _fullname_ of a single object to fetch.

2. A single reflexive name prefix out of which the producer can (by application-specific means) construct a number of _fullnames_ of the objects it may want to fetch,

3. The reflexive _fullname_ of a FLIC Manifest [I-D.irtf-icnrg-flic] enumerating the suffixes that may be used by the producer to construct the necessary names,

4. Multiple reflexive name TLVs MAY be included in the Interest message if none of the above 3 options covers the desired use case.

The last of the four options above, while not explicitly outlawed, SHOULD NOT be used. This is because it results in a longer Interest message and requires extra FIB resources. Hence, it is more likely a forwarder will reject the Interest for lack of resources. A forwarder MAY optimize for the case of a single Reflexive Name TLV at the expense of those with more than one.

A producer, upon receiving an Interest with one or more Reflexive Name TLVs, may decide it needs the pull the associated data object(s). It therefore can issue one or more Reflexive Interests by appending the necessary name components needed to form valid full names of the associated objects present at the originating consumer. These in fact comprise conventional Interest-Data exchanges, with no alteration of the usual semantics with regard to signatures, caching, expiration, etc. When the producer has retrieved the required objects to complete the original Interest-Data exchange, it can issue its Data response, which unwinds all the established state at the producer, the consumer, and the intermediate forwarders.

5. Forwarder operation for Reflexive Interests

The forwarder operation for CCNx and/or NDN is changed in three respects when supporting Reflexive Interests.

1. The forwarder MUST create short-lifetime FIB entries for any Reflexive Interest Name prefixes communicated in an Interest message. If the forwarder does not have sufficient resources to do so, it MUST reject the Interest with the T_RETURN_NO_RESOURCES error - the same error used if the forwarder were lacking sufficient PIT resources to process the Interest message.

2. Those FIB entries MUST be queried whenever an Interest message arrives whose first name component is of the type _Reflexive Interest Name Component_

3. The FIB entry MUST be removed eventually, after the corresponding Data message has been forwarded. One option would be to remove the FIB directly after the Data message has been forwarded. However, the forwarder MAY do lazy cleanup.

The PIT entry for the Reflexive Interest is consumed per regular Interest/Data message forwarding requirements. The PIT entry for the originating Interest (that communicated the Reflexive Interest Name) is also consumed by a final Data message from the producer to the original consumer.

6. State coupling between producer and consumer

A consumer that wishes to use this scheme MUST utilize one of the reflexive naming options defined in Section 4 and include it in the corresponding Interest message. The Reflexive Name TLV _and_ the full name of the requested data object (that identifies the producer) identify the common state shared by the consumer and the producer. When the producer responds by sending Interests with the Reflexive Name Prefix, the original consumer therefore has sufficient information to map these Interests to the ongoing Interest-Data exchange.

The exchange is finished when the producer who received the original Interest message responds with a Data message (or an Interest Return message in the case of error) answering the original Interest. After sending this Data message, the producer SHOULD destroy the corresponding shared state. It MAY decide to use a timer that will trigger a later state destruction. After receiving this Data message, the originating consumer MUST destroy the corresponding Interest-Data exchange state.

7. Use cases for Reflexive Interests

7.1. Achieving Remote Method Invocation with Reflexive Interests

RICE (Remote Method Invocation in ICN) [Krol2018] uses the Reflexive Interest Forwarding scheme that inspired the design specified in this document.

In RICE, the original Interest denotes the remote method (plus potential parameters) to be invoked at a producer (server). Before committing any computating resources, the server can then request authentication credentials and (optional) parameters using reflexive Interest-Data exchanges.

When the server has obtained the necessary credentials and input parameters, it can decide to commit computing resources, starts the

compute process, and returns a handle ("Thunk") in the final Data message to the original consumer (client).

The client would later request the computation results using a regular Interest-Data exchange (outside the Reflexive-Interest transaction) -- using the Thunk as a name for the computation result.

Figure 3 depicts an abstract message diagram for RICE. In addition to the 4-way Reflexive Forwarding Handshake (see Figure 2 for the details of the interaction), RICE adds another (standard) ICN Interest/Data exchange for transmitting the RMI result. The Thunk name is provided to the consumer in the D1 DATA message (answering the initial I1 Interest).

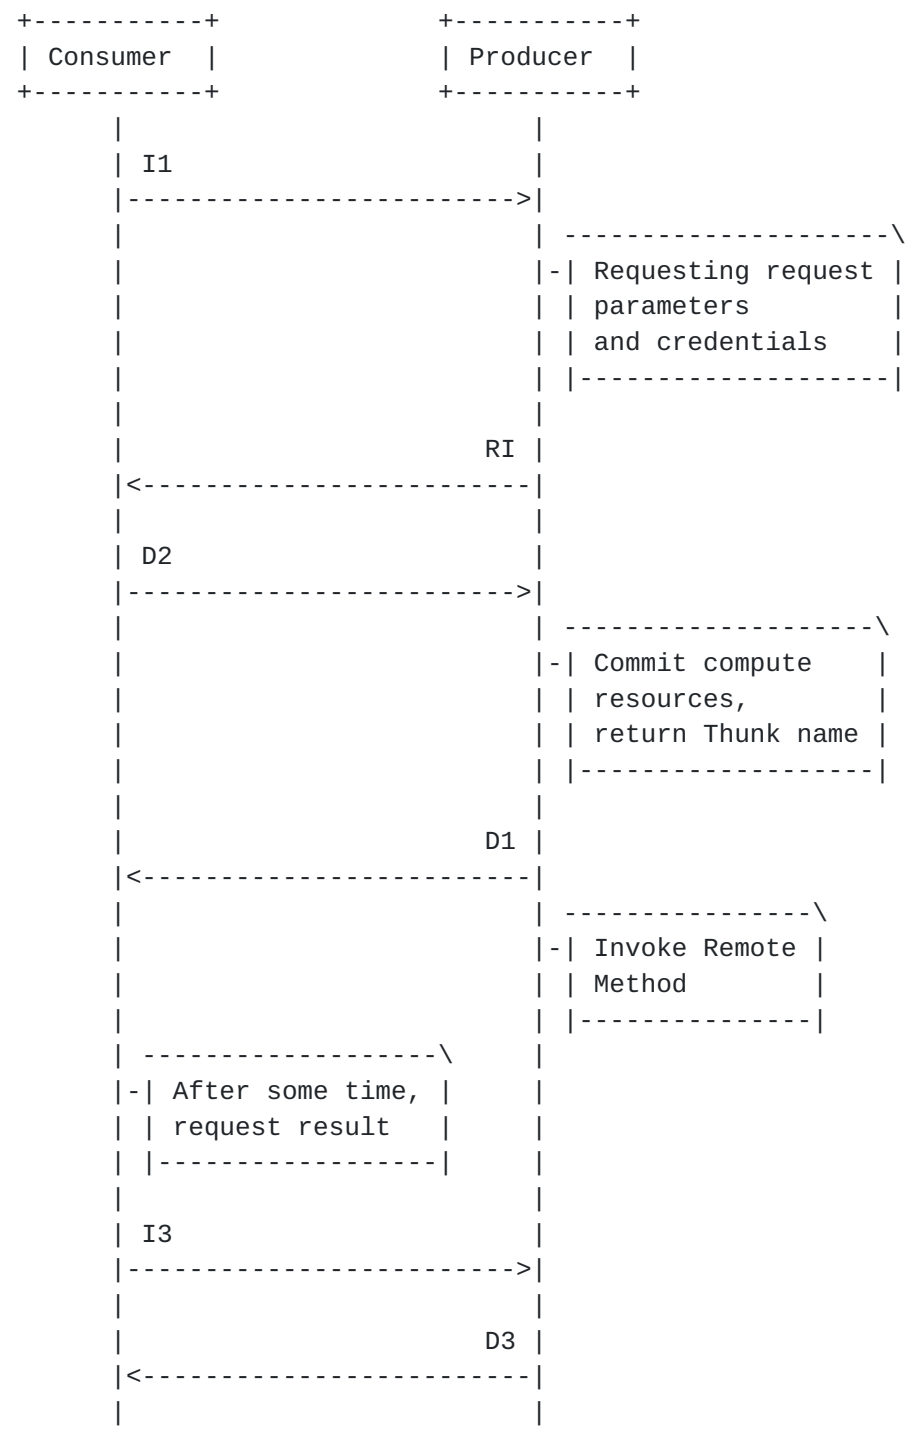

Legend:
I1: Interest #1 containing the Reflexive Name Prefix TLV
D1: Data message, answering initiating I1 Interest, returning Thunk name
D2: Data message, answering RI (parameters, credentials)
I3: Regular Interest for Thunk (compute result)
D3: Data message, answering I3

Figure 3: RICE Message Flow

7.2. RESTful Web Interactions

In todays HTTP-based web, RESTful (Representational State Transfer) web interactions are realized by sending requests in a client/server interaction, where the requests provides the application context (or a reference to it). It has been noted in [Moiseenko2014] that corresponding requests often exceed the response messages in size, and that this raises the problems noted in Section 1.1 when attempting to map such exchanges directly to CCNx/NDN.

Another reason not to include all request parameters in a (possibly encrypted) Interest message is the fact that a server (that is serving thousands of clients) would be obliged to receive, possibly decrypt and parse the complete requests before being able to determine whether the requestor is authorized, whether the request can be served etc. Many non-trivial requests could thus lead to computational overload attacks.

Using Reflexive Interest Forwarding for RESTful Web Interactions would encode the REST request in the Original request, together with a Reflexive Interest Prefix that the server could then use to get back to the client for authentication credentials and request parameters, such as cookies. The request result (response message) could either be transmitted in the Data message answering the original request, or -- in case of dynamic, longer-running computations -- in a seperate Interest/Data exchange, potentially leveraging the Thunk scheme described in section Section 7.1.

Unlike approaches where clients have to signal a globally routable prefix to the network, this approach would not require the client (original consumer) to expose its identity to the network (the network only sees the temporary Reflexive Name Prefix), but it would still be possible to authenticate the client at the server.

7.3. Achieving simple data pull from consumers with reflexive Interests

An oft-cited use case for ICN network architectures is _Internet of Things_ (IoT), where the sources of data are limited-resource sensor/actuators. Many approaches have been tried (e.g. [Baccelli2014], [Lindgren2016], [Gundogan2018]) with varying degrees of success in addressing the issues outlined in Section 1.1. The reflexive forwarding extension may substantially ameliorate the documented difficulties by allowing a different model for the basic interaction of sensors with the ICN network.

Instead of acting as a producer (either directly to the Internet or indirectly through the use of some form of application-layer gateway), the IoT device need only act as a consumer. When it has data to provide, it issues a "phone-home" Interest message to a pre-configured rendezvous name (e.g. an application-layer gateway or ICN Repo [Chen2015]) and provides a reflexive name prefix TLV for the data it wishes to publish. The target producer may then issue the necessary reflexive Interest message(s) to fetch the data. Once fetched, validated, and stored, the producer then responds to the original Interest message with a success indication, possibly containing a Data object if needed to allow the originating device to modify its internal state. Alternatively, the producer might choose to not respond and allow the original Interest to time out, although this is NOT RECOMMENDED except in cases where the extra message transmission bandwith is at a premium compared to the persistence of stale state in the forwarders. We note that this interaction approach mirrors the earlier efforts using Interest-Interest-Data designs.

Figure 4 depicts this interaction with the OPTIONAL D1 message. See Figure 2 for the details of the general Reflexive Forwarding interaction.

Legend:
I1: Interest #1 containing the Reflexive Name Prefix TLV
D1: Data message (OPTIONAL), finalizing interaction
D1: Data message, answering RI, returning IoT data object

Figure 4: "Phone Home" Message Flow

There are two approaches that the IoT device can use for its response to a reflexive Interest. It can simply construct a Data Message

bound through the usual ICN hash name to the reflexive Interest name. Since the scope of any data object bound in this way is only the duration of the enclosing Interest-Data exchange (see Section 8.2) the producer would need to itself construct any persistent Data object, name it, and sign it. This is sometimes the right approach, as for some applications the identity of the originating IoT device is not important from an operational or security point of view; in contrast the identity of the gateway or Repo is what matters.

If alternatively, the persistent Data object should be bound from a naming and security point of view to the originating IoT device, this can be easily accomplished. Instead of directly placing the content in a Data object responding to the reflexive Interest as above, the consumer encapsulates a complete CCNx/NDN Data message (which includes the desired name of the data) as in the response to the reflexive Interest message.

The interaction model described above brings a number potential advantages, some obvious, some less so. We enumerate a few of them as follows:

* By not requiring the IoT device to be actively listening for Interests, it can sleep and only wake up if it has something to communicate. Conversely, parties interested in obtaining data from the device do not need to be constantly polling in order to ascertain if there is new data available.

* No forwarder resources are tied up with state apart from the actual reflexive forwarding interactions. All that is needed is enough routing state in the FIB to be able to forward the "phone home" Interest to an appropriate target producer. While this model does not provide all the richness of a full Pub/Sub system (like that described in [Gundogan2018]) we argue it is adequate for a large subclass of such applications.

* The reflexive interest, through either a name suffix or Interest payload, can give the IoT device useful context from which to craft its Data object in response. One highly useful parameter would be a robust clock value for the device to use as a timestamp of the data, possibly as part of its name to correctly place it in a time seres of sensor readings. This substantially alleviates the need for low-end devices to have a robust time base, as long as they trust the producer they contact to provide it.

8. Implementation Considerations

There are a number of important aspects to the reflexive forwarding design which affect correctness and performance of existing forwarder, consumer, and producer implementations desiring to support it. This section discusses the effect on each of these elements of the CCNx/NDN protocol architecture.

8.1. Forwarder implementation considerations

8.1.1. Forwarding Information Base (FIB)

The FIB is a performance-critical data structure in any forwarder, as it needs to support relatively expensive longest name prefix match (LNPM) lookup algorithms. A number of well-known FIB data structures are heavily optimized for read access, since for normal Interest message processing the FIB changes slowly - only after topological changes or routing protocol updates. Support for reflexive names changes this, as FIB entries are created and destroyed rapidly as Interest messages containing reflexive name TLVs are processed and the corresponding Data messages come back.

While it may be feasible, especially in low-end forwarders handling a low packet forwarding rate to ignore this problem, for high-speed forwarders there are a number of hazards, including:

1. If the entire FIB needs to be locked in order to insert or remove entries, this could cause inflated forwarding delays or in extreme cases, forwarding performance collapse.

2. A number of high-speed forwarder implementations employ a sharded PIT scheme to better parallelize forwarding across processing cores. The FIB, however, is still a shared data structure which is either read without read locks across cores, or explicitly copied such that there is a separate copy of the FIB for each PIT shard. Clearly, a high update rate without read locks and/or updating many copies of the FIB are unattractive implementation options. (Note: with this reflexive name scheme it is not feasible to force reflexive interests to be hashed or be otherwise directed to the PIT shard holding the original Interest state).

There are any number of alternative FIB implementations that can work well however. The most straightforward is to simply implement a "special" FIB for just reflexive name lookups. This is feasible because reflexive names deterministically contain the distinguished high-order name component type of T_REFLEXIVE_NAME, whose content is a 64-bit value that can be easily hashed to a FIB entry directly,

avoiding the more expensive LNPM lookup. Inserts and deletes then devolve to the well-understood problem of hash table maintenance.

8.1.2. Interactions with Interest Lifetime

If and when a producer decides to fetch data from the consumer using one or more reflexive Interest-Data exchanges, the total latency for the original Interest-Data exchange is inflated, potentially by multiple RTTs. It is difficult for a consumer to predict the inflation factor when issuing the original Interest, and hence there can be a substantial hazard of that Interest lifetime expiring before completion of the full multi-way exchange. This can result in persistent failures, which is obviously highly undesirable.

There is a fairly straightforward technique that can be employed by forwarders to avoid these "false" Interest lifetime expirations. In the absence of a superior alternative technique, it is RECOMMENDED that all forwarders implement the following algorithm.

When processing an Interest containing the reflexive name TLV and creating the necessary FIB entry (see Section 8.1.1 above), the forwarder also creates a _back pointer_ from that FIB entry to the PIT entry for the Interest message that created it. This PIT entry contains the current value of the remaining Interest lifetime or alternatively a value from which the remaining Interest lifetime can be easily computed. Call this value _IL_(t)_.

If and when a reflexive Interest arrives from upstream matching the reflexive FIB entry, the forwarder examines the Interest lifetime of the arriving reflexive Interest. Call this value _IL_(r)_. The forwarder computes MAX(_IL_(t), (IL_(r) * 1.5)_), and replaces _IL_(t)_ with this value. This in effect ensures that the remaining Interest lifetime of the original Interest accounts for the additional 1.5 RTTs that may occur as a result of the reflexive Interest-Data exchange.

If the above algorithm is implemented naively as described above, it may run afoul of a sharded PIT forwarder implementation, since the PIT entry for the reflexive Interest and the PIT entry for the original Interest may be in different shards. Therefore, if the update is done cross-shard on each reflexive Interest arrival, performance may suffer, perhaps dramatically. Instead, the following approach to updating the Interest lifetime after computing the new value is RECOMMMENDED for sharded-PIT forwarders.

When creating the reflexive FIB entry as above in Section 8.1.1, copy the remaining Interest lifetime from the PIT entry. Do the PIT update if and only if this value is about to expire, thus paying the

cross-shard update cost only if the original Interest is about to expire. A further optimization at the cost of modest extra complexity is to instead _queue_ the update to the core holding the shard of the original PIT entry rather than doing the update directly. If the PIT entry expires or is satisfied, instead of removing it the associated core checks the update queue and does the necessary update.

While the above approach of inflating the interest lifetime of the original Interest to accommodate the additional RTTs of reflexive Interest-Data exchanges, this does introduce a new vulnerability that must be dealt with. A Producer, either through a bug or malicious intent, could keep an originating Interest-Data exchange alive by continuing to send reflexive Interests back to the consumer, while the consumer had no way to terminate the enclosing interaction (there is no "cancel Interest" function in either NDN nor CCNx). To eliminate this hazard, if the consumer rejects a reflexive interest with a T_RETURN_PROHIBITED error, the forwarder(s), in addition to satisfying the coresponding PIT entry, MUST also delete the associated reflexive FIB entry, thereby preventing any further reflexive Interests from reaching the consumer. This allows the enclosing Interest-Datsa exchange to either time out or be correctly ended with a Data message or Interest Return from the Producer.

8.1.3. Interactions with Interest aggregation

As with numerous other situations where multiple Interests for the same named object arrive containing different parameters (e.g. Interest Lifetime, QoS, payload hash) the same phenomenon occurs for the reflexive Name TLV. If such Interests collide, the forwarder MUST NOT aggregate these Interest messages and instead MUST create a separate PIT entry for each.

8.2. Consumer Implementation Considerations

8.2.1. Data objects returned by the consumer to reflexive name Interests arriving from a producer

The Data objects returned to the producer in response to a reflexive Interest are normal CCNx/NDN data objects. It is therefore worth noting that the object is bound to the reflexive Interest full name via the hash and hence the scope of the object is under most circumstances meaningful only for the duration of the enclosing Interest-Data interaction. This property is ideal for naming and securing data that is "part of" the enclosing interaction - things like method arguments, authenticators, and key exchange parameters, but not for the creation and naming of objects intended to survive outside the current interaction's scope (c.f. Section 7.3, which

describes how to provide globally-named objects using encapsulation). In general, the consumer should use the following guidelines in creating Data messages in response to reflexive Interest messages from the producer.

(a) Set the recommended cache time (T_CACHETIME) either to zero, or a value no greater than the Interest lifetime (T_INTLIFE) of the original Interest messsage.

(b) Set the payload type (T_PAYLDTYPE) according to the type of object being returned (e.g. object, link, manifest)

(c) Set the expiry time (T_EXPIRY) to a value greater than _now_, and less than or equal to the _now_ + Interest lifetime (T_INTLIFE) of the original Interest messsage.

8.2.2. Terminating unwanted reflexive Interest exchanges

A consumer may wish to stop receiving reflxive Interests due to possible erors or malicious behavior on the part of the producer. Therefore, if the consumer receives an unwanted reflexive Interest, it SHOULD reject that interest with a T_RETURN_PROHIBITED error. This will provoke the forwarders to prevent further reflexive Interests from reaching the consumer, as described above in Section 8.1.2, Paragraph 7.

8.2.3. Interactions with caching

The reflexive named objects provide "local", temporary names that are only defined for one specific interaction between a consumer and a producer. Corresponding Data objects MUST NOT be shared between multiple consumers (violating this would require specail gyrations by the producer since the reflexive Name utilizes per-consumer/per-interaction random values). A producer MUST NOT issue an Interest message for any reflexive name after it has sent the final Data message answering the original Interest.

Forwarders SHOULD still cache reflexive Data objects for retransmissions within a transactions, but they MUST remove them from the content store when they forward the final Data message answering the original Interest.

8.3. Producer Implementation Considerations

Producers receiving an Interest with a Reflexive Name Component, MAY decide to issue Interests for the corresponding Data objects. All Reflexive Interest message that a producer sends MUST be sent over the face that the original Interest was received on.

9. Operational Considerations

This extension represents a substantial enhancement to the CCNx/NDN protocol architecture and hence has important forward and backward compatibility effects. The most important of these is that correct operation of the scheme requires an unbroken chain of forwarders between the consumer and the desired producer that support the Reflexive Name TLV and the corresponding forwarder capabilities specified in Section 5. When this invariant is not satisfied, some means is necessary to detect and hopefully recover from the error. We have identified three possible approaches to handling the lack of universal deployment of forwarders supporting the reflexive forwarding scheme.

The first approach simply lets the producer detect the error by getting a "no route to destination" error when trying to send an Interest to a reflexive name. This will catch the error, but only after forwarding resources are tied up and the producer has done some work on the original Interest message. Further, the producer would need a bit of smarts to determine that this is a permanent error and not a transient to be retried. In order for the consumer to attempt recovery, there might be a need for some explicit error returned for the original interest to tell the consumer what the likely problem is. This approach does not enable an obvious recovery path for the consumer either, since while we might envision a way to steer a subsequent Interest onto a working path as proposed in [I-D.oran-icnrg-pathsteering], there is no capability to force Interest routing away from an otherwise working path not supporting the reflexive name TLV.

A second approach is to bump the CCNx/NDN protocol version to explicitly indicate the lack of comparability. Such Interests would be rejected by forwarders not supporting this protocol extension. A consumer wishing to use the reflexive name TLV would use the higher protocol version on those Interest messages (but could of course continue to use the current version number on other Interest messages). This is a big hammer, but may be called for in this situation because:

(a) it detects the problem immediately and deterministically, and

(b) one could assume an ICN routing protocol that would only forward to a next hop that supports the updated protocol version number. The supported forwarder protocol versions would have been communicated in the routing protocol ahead of time.

A third option is to, as a precondition utilizing the protocol in a deployment, create and deploy a neighbor capability exchange protocol

which will tell a downstream forwarder if the upstream can handle the new TLV. This might avoid the large hammer of updating the protocol version, but of course this puts a pretty strong dependency on somebody actually designing and publishing such a protocol! On the other hand, a neighbor capability exchange protocol for CCNx/NDN would have a number of other substantial benefits, which makes it worth seriously considering anyway.

10. Mapping to CCNx and NDN packet encodings

10.1. Packet encoding for CCNx

For CCNx[RFC8569] there is one new Name Component TLV type defined in this specification.

Abbrev	Name	Description
T_REFLEXIVE_NAME	Reflexive Name Component	Name component to use as name prefix in Reflexive Interest Message

Table 1: Reflexive Name TLV

10.2. Packet encoding for NDN

TBD based on [NDNTLV]. Suggestions from the NDN team greatly appreciated.

11. IANA Considerations

Please add the T_REFLEXIVE_NAME component TLV to the CCNx Name types TLV types registry of [RFC8609], with Length 9 bytes and type of 64 bit random integer.

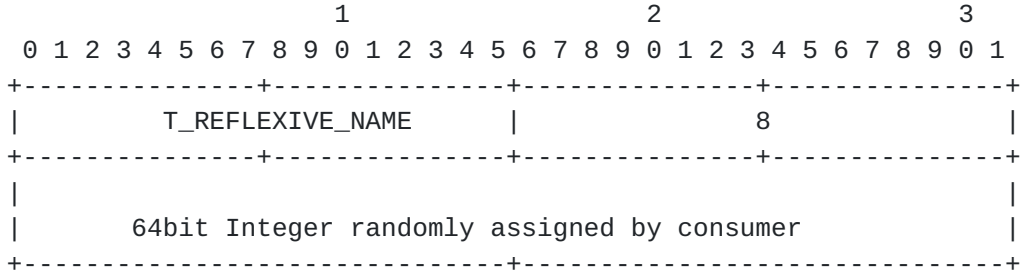

Figure 5: Reflexive Name component type

12. Security Considerations

One of the major motivations for the reflexive forwarding extension specified in this document is in fact to enable better security and privacy characteristics for ICN networks. The main considerations are presented in Section 1, but we briefly recapitulate them here:

* Current approaches to authentication and data transfer often use payloads in Interest messages, which are clumsy to secure (Interest messages must be signed) and as a consequence make it very difficult to ensure consumer privacy. Reflexive forwarding moves all sensitive data to the Data messages sent in response to reflexive Interests, which are secured in the same manner as all other Data messages in the CCNx and NDN protocol designs.

* In many scenarios, consumers are forced to also act as producers so that data may be fetched by either a particular, or arbitrary other party. The means the consumer must arrange to have a routable name prefix and that prefix be disseminated by the routing protocol or other means. This represents both a privacy hazard (by revealing possible important information about the consumer) and a security concern as it opens up the consumer to the full panoply of flooding and crafted Interest denial of service attacks.

* In order to achieve multi-way handshakes, in current designs a consumer wishing a producer to communicate back must inform the producer of what (globally routable) name to use. This gives the consumer a convenient means to mount a variety of reflection attacks by enlisting the producer to send Interests to desired victims.

As a major protocol extension however, this design brings its own potential security issues, which are discussed in the following subsections.

12.1. Collisions of reflexive Interest names

Reflexive Interest names are constructed using 64-bit random numbers. This is intended to ensure an off-path attacker cannot easily manufacture a matching reflexive Interest and either masquerade as the producer, or mount a denial of service attack on the consumer. It also limits tracking through the linkability of Interests containing a re-used random value.

Therefore consumers MUST utilize a robust means of generating these random values, and it is RECOMMENDED that a pseudo-random number

generator (PRNG) approved for use with cryptographic protocols be employed.

12.2. Additional resource pressure on PIT and FIB

Normal Interest message processing in CCNx and NDN needs to consider effect of various resource depletion attacks on the PIT, particularly in the form of Interest flooding attacks (see [Gasti2012] for a good overview of DoS and DDoS mitigation on ICN networks). Interest messages utilizing this reflexive forwarding extension can place additional resource pressure on the PIT, and additionally cause otherwise stable FIB resources to be subject to highly dynamic usage.

While this does not represent a new DoS/DDoS attack vector, the ability of a malicious consumer to utilize this extension in an attack does represent an increased risk of resource depletion, especially if such Interests are given unfair access to PIT and FIB resources. Implementers SHOULD therefore protect PIT and FIB resources by weighing requests for reflexive forwarding resources appropriately relative to other Interests.

12.3. Privacy Considerations

ICN architectures like CCNx and NDN provide a rich tapestry of interesting privacy issues, which have been extensively explored in the research literature. The fundamental tradeoffs for privacy concern the risk of exposing the names of information objects to the forwarding elements of the network, which is a necessary property of any name-based routing and forwarding design. Numerous approaches have been explored with varying degrees of success, such as onion routing ([DiBenedettoGTU12]), name encryption ([Ghali2017]), and name obfuscation ([Arianfar2011]) among others.

Reflexive forwarding does not change the overall landscape of privacy tradeoffs, nor seem to introduce additional hazards. In fact, the privacy exposures are confined to the inverse path of forwarders from the producer to the consumer, through which the original Interest forwarding may have already exposed names on path. Similar name privacy techniques to those cited above may be equally applied to the names in reflexive Interests.

While the individual reflexive Interest-Data exchanges have similar properties to those in any NDN or CCNx exchange, the target usages by applications may have interaction patterns that are subject to relatively straightforward fingerprinting by adversaries. For example, a particular RMI invocation may fingerprint simply through the count of arguments fetched by the producer and their sizes. The

attacker must however be on path, which somewhat ameliorates the exposure hazards.

13. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8569] Mosko, M., Solis, I., and C. Wood, "Content-Centric Networking (CCNx) Semantics", RFC 8569, DOI 10.17487/RFC8569, July 2019, <https://www.rfc-editor.org/info/rfc8569>.

[RFC8609] Mosko, M., Solis, I., and C. Wood, "Content-Centric Networking (CCNx) Messages in TLV Format", RFC 8609, DOI 10.17487/RFC8609, July 2019, <https://www.rfc-editor.org/info/rfc8609>.

14. Informative References

[Arianfar2011] Arianfar, S., Koponen, T., Raghavan, B., and S. Shenker, "On preserving privacy in content-oriented networks, in ICN '11: Proceedings of the ACM SIGCOMM workshop on Information-centric networking", DOI https://doi.org/10.1145/2018584.2018589, August 2011, <https://dl.acm.org/doi/10.1145/2018584.2018589>.

[Auge2018] Augé, J., Carofiglio, G., Grassi, G., Muscariello, L., Pau, G., and X. Zeng, "MAP-Me: Managing Anchor-Less Producer Mobility in Content-Centric Networks, in IEEE Transactions on Network, Volume 15, Issue 2", DOI 10.1109/TNSM.2018.2796720, June 2018, <https://ieeexplore.ieee.org/document/8267132>.

[Baccelli2014] Baccelli, E., Mehlis, C., Hahm, O., Schmidt, T., and M. Wählisch, "Information centric networking in the IoT: experiments with NDN in the wild, in ACM-ICN '14: Proceedings of the 1st ACM Conference on Information-Centric Networking", DOI 10.1145/2660129.2660144, September 2014, <https://dl.acm.org/doi/abs/10.1145/2660129.2660144>.

[Carzaniga2011] Carzaniga, A., Papalini, M., and A.L. Wolf, "Content-Based

Publish/Subscribe Networking and Information-Centric Networking", DOI 10.1145/2018584.2018599, September 2011, <https://conferences.sigcomm.org/sigcomm/2011/papers/icn/p56.pdf>.

[Chen2015] Chen, S., Cao, J., and L. Zhu, "NDSS: A Named Data Storage System, in International Conference on Cloud and Autonomic Computing", DOI 10.1109/ICCAC.2015.12, September 2014, <https://ieeexplore.ieee.org/document/7312154>.

[DiBenedettoGTU12]
DiBenedetto, S., Gasti, P., Tsudik, G., and E. Uzun, "ANDaNA: Anonymous Named Data Networking Application, in NDSS 2012", DOI https://arxiv.org/abs/1112.2205v2, 2102, <https://www.ndss-symposium.org/ndss2012/andana-anonymous-named-data-networking-application>.

[Gasti2012]
Gasti, P., Tsudik, G., Uzun, Ersin., and L. Zhang, "DoS and DDoS in Named Data Networking, in 22nd International Conference on Computer Communication and Networks (ICCCN)", DOI 10.1109/ICCCN.2013.6614127, August 2013, <https://ieeexplore.ieee.org/document/6614127>.

[Ghali2017]
Tsudik, G., Ghali, C., and C. Wood, "When encryption is not enough: privacy attacks in content-centric networking, in ICN '17: Proceedings of the 4th ACM Conference on Information-Centric Networking", DOI https://doi.org/10.1145/3125719.3125723, September 2017, <https://dl.acm.org/doi/abs/10.1145/3125719.3125723>.

[Gundogan2018]
Gündoğan, C., Kietzmann, P., Schmidt, T., and M. Wählisch, "HoPP: publish-subscribe for the constrained IoT, in ICN '18: Proceedings of the 5th ACM Conference on Information-Centric Networking", DOI 10.1145/3267955.3269020, September 2018, <https://dl.acm.org/doi/abs/10.1145/3267955.3269020>.

[I-D.irtf-icnrg-flic]
Tschudin, C., Wood, C., Mosko, M., and D. Oran, "File-Like ICN Collections (FLIC)", Work in Progress, Internet-Draft, draft-irtf-icnrg-flic-02, 4 November 2019, <https://tools.ietf.org/html/draft-irtf-icnrg-flic-02>.

[I-D.irtf-icnrg-terminology]
Wissingh, B., Wood, C., Afanasyev, A., Zhang, L., Oran, D., and C. Tschudin, "Information-Centric Networking (ICN): CCNx and NDN Terminology", Work in Progress, Internet-Draft, draft-irtf-icnrg-terminology-08, 17 January 2020, <https://tools.ietf.org/html/draft-irtf-icnrg-terminology-08>.

[I-D.oran-icnrg-pathsteering]
Moiseenko, I. and D. Oran, "Path Steering in CCNx and NDN", Work in Progress, Internet-Draft, draft-oran-icnrg-pathsteering-00, 21 October 2019, <https://tools.ietf.org/html/draft-oran-icnrg-pathsteering-00>.

[Krol2018] Krol, M., Habak, K., Oran, D., Kutscher, D., and I. Psaras, "RICE: Remote Method Invocation in ICN, in Proceedings of the 5th ACM Conference on Information-Centric Networking - ICN '18", DOI 10.1145/3267955.3267956, September 2018, <https://conferences.sigcomm.org/acm-icn/2018/proceedings/icn18-final9.pdf>.

[Lindgren2016]
Lindgren, A., Ben Abdessiem, F., Ahlgren, B., Schlelén, O., and A.M. Malik, "Design choices for the IoT in Information-Centric Networks, in 13th IEEE Annual Consumer Communications and Networking Conference (CCNC)", DOI 10.1109/CCNC.2016.7444905, January 2016, <https://ieeexplore.ieee.org/abstract/document/7444905>.

[Moiseenko2014]
Moiseenko, I., Stapp, M., and D. Oran, "Communication patterns for web interaction in named data networking", DOI 10.1145/2660129.2660152, September 2014, <https://dl.acm.org/doi/10.1145/2660129.2660152>.

[Mosko2017]
Mosko, M., "CCNx 1.0 Bidirectional Streams", arXiv 1707.04738, July 2017, <https://arxiv.org/abs/1707.04738>.

[NDN] "Named Data Networking", 2020, <https://named-data.net/project/execsummary/>.

[NDNTLV] "NDN Packet Format Specification", 2016, <http://named-data.net/doc/ndn-tlv/>.

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, DOI 10.17487/RFC3261, June 2002, <https://www.rfc-editor.org/info/rfc3261>.

[RFC6337] Okumura, S., Sawada, T., and P. Kyzivat, "Session Initiation Protocol (SIP) Usage of the Offer/Answer Model", RFC 6337, DOI 10.17487/RFC6337, August 2011, <https://www.rfc-editor.org/info/rfc6337>.

[RFC7530] Haynes, T., Ed. and D. Noveck, Ed., "Network File System (NFS) Version 4 Protocol", RFC 7530, DOI 10.17487/RFC7530, March 2015, <https://www.rfc-editor.org/info/rfc7530>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/info/rfc8446>.

[Zhang2018] Zhang, Y., Xia, Z., Mastorakis, S., and L. Zhang, "KITE: Producer Mobility Support in Named Data Networking, in Proceedings of the 5th ACM Conference on Information-Centric Networking - ICN '18", DOI 10.1145/3267955.3267959, September 2018, <https://conferences.sigcomm.org/acm-icn/2018/proceedings/icn18-final23.pdf>.

Authors' Addresses

Dave Oran
Network Systems Research and Design
4 Shady Hill Square
Cambridge, MA 02138
United States of America

Email: daveoran@orandom.net

Dirk Kutscher
University of Applied Sciences Emden/Leer
Constantiapl. 4
26723 Emden
Germany

Email: ietf@dkutscher.net
URI: https://dirk-kutscher.info