On Information Exposure through Named Content

Kostantinos Katsaros, Lorenzo Saino, Ioannis Psaras, George Pavlou

Communications and Information Systems Group
Department of Electrical and Electronics Engineering
University College London
Email: {k.katsaros,l.saino,i.psaras,g.pavlou}@ucl.ac.uk

Q-ICN workshop - August 20th, 2014

Outline

Introduction

Background and motivations

Information exposure

- through content naming
- through name resolution

Handling exposed information

Implementation considerations

Conclusions

Introduction

Introduction

Research on content naming and resolution focused on a variety of different aspects:

- Security
- Routability
- Scalability
- Extensibility

Introduction

Research on content naming and resolution focused on a variety of different aspects:

- Security
- Routability
- Scalability
- Extensibility

We argue however that *information exposure* considerations, i.e. the amount of information leaked by content names and the name resolution process have been overlooked in naming research.

Introduction

Research on content naming and resolution focused on a variety of different aspects:

- Security
- Routability
- Scalability
- Extensibility

We argue however that *information exposure* considerations, i.e. the amount of information leaked by content names and the name resolution process have been overlooked in naming research.

We show that information exposure can enable both desirable and undesirable features.

Information exposure in various networking environments

Information exposure in various networking environments

Content distribution:

- Access logging
- Content neutrality
- Cache purging

Information exposure in various networking environments

Content distribution:

- Access logging
- Content neutrality
- Cache purging

Mobile, opportunistic networks:

- Time and space scoping for efficient usage of scarce network resources

Information exposure in various networking environments

Content distribution:

- Access logging
- Content neutrality
- Cache purging

Mobile, opportunistic networks:

- Time and space scoping for efficient usage of scarce network resources

IoT, smart cities/grids, vehicular networks:

- Time and space scoping to limit spread to interested entities
- Need not to expose sensitive information through content names

Information exposure through content names

A list of possible content characteristics that can be exposed to packet handling network entities are:

Information exposure through content names

A list of possible content characteristics that can be exposed to packet handling network entities are:

- **Service type:** MIME type of traffic associated to the content

Information exposure through content names

A list of possible content characteristics that can be exposed to packet handling network entities are:

- **Service type:** MIME type of traffic associated to the content
- **Ownership:** identity of the content provider

Information exposure through content names

A list of possible content characteristics that can be exposed to packet handling network entities are:

- **Service type:** MIME type of traffic associated to the content
- **Ownership:** identity of the content provider
- **Caching properties:** content cacheability, TTL, etc...

Information exposure through content names

A list of possible content characteristics that can be exposed to packet handling network entities are:

- **Service type:** MIME type of traffic associated to the content
- **Ownership:** identity of the content provider
- **Caching properties:** content cacheability, TTL, etc...
- **Service class:** Class identifying traffic covered by a specific SLA

Information exposure through content names

A list of possible content characteristics that can be exposed to packet handling network entities are:

- **Service type:** MIME type of traffic associated to the content
- **Ownership:** identity of the content provider
- **Caching properties:** content cacheability, TTL, etc...
- **Service class:** Class identifying traffic covered by a specific SLA
- **Scope:** Temporal and geographical scope of a content object

Information exposure through content names

A list of possible content characteristics that can be exposed to packet handling network entities are:

- **Service type:** MIME type of traffic associated to the content
- **Ownership:** identity of the content provider
- **Caching properties:** content cacheability, TTL, etc...
- **Service class:** Class identifying traffic covered by a specific SLA
- **Scope:** Temporal and geographical scope of a content object
- **Content format:** Resolution, codec and other information useful to characterize different versions of the same content.

Exposing information through name resolution

Information can be exposed not only through content naming decisions, but also through the name resolution process.

Exposing information through name resolution

Information can be exposed not only through content naming decisions, but also through the name resolution process.

Example: content access logging via name resolution

- By requiring to resolve a permanent content identifier to an ephemeral (possibly routable) identifier before downloading the content enables content providers or CDNs to log access to requested contents.

Ephemeral names

More generally, we argue that requiring a resolution between permanent names for content identification to ephemeral names for routing could provide a number of desirable features:

Ephemeral names

More generally, we argue that requiring a resolution between permanent names for content identification to ephemeral names for routing could provide a number of desirable features:

- Content provider access logging when resolving permanent names to ephemeral names

Ephemeral names

More generally, we argue that requiring a resolution between permanent names for content identification to ephemeral names for routing could provide a number of desirable features:

- Content provider access logging when resolving permanent names to ephemeral names
- Content neutrality

Ephemeral names

More generally, we argue that requiring a resolution between permanent names for content identification to ephemeral names for routing could provide a number of desirable features:

- Content provider access logging when resolving permanent names to ephemeral names
- Content neutrality
- Cache purging

Handling exposed information - I

Handling exposed information - I

Service type

- Can be used by content routers to make informed routing, forwarding and caching decisions with the objective of maximizing QoS depending on traffic characteristics.
- For example, minimize latency for real-time traffic and maximize throughput for bulk data transfer.
- Limited risk of service type misuse as inaccurate assignment degrades performance.

Handling exposed information - I

Service type

- Can be used by content routers to make informed routing, forwarding and caching decisions with the objective of maximizing QoS depending on traffic characteristics.
- For example, minimize latency for real-time traffic and maximize throughput for bulk data transfer.
- Limited risk of service type misuse as inaccurate assignment degrades performance.

Service class

- ISPs can provide preferential treatment for premium traffic. Service class attribute can be used by content providers to identify traffic for preferential treatment by ISPs.
- Differently from service type, there is a more realistic risk of misuse.
- Malicious usage can be mitigated using, for example, algorithmically generated ephemeral names.

Handling exposed information - II

Handling exposed information - II

Ownership

- Ownership information can be exposed to support authenticity verification.
- However, ownership information may hinder *content neutrality*, i.e. ISPs maye be able to deliberately throttle traffic from specific content providers.

Handling exposed information - II

Ownership

- Ownership information can be exposed to support authenticity verification.
- However, ownership information may hinder *content neutrality*, i.e. ISPs maye be able to deliberately throttle traffic from specific content providers.

Caching properties

- Content providers can use this attribute to communicate information that caching nodes can use to improve caching performance.
- These properties may include cacheability information and information to support cache purging operations.
- Content provider based cache purging can be implemented by explicitly labelling each content with the identifiers of content objects it obsoletes.
- This however raises concerns of DoS attacks as malicious providers may attempt to purge content they do not own.

Handling exposed information - III

Handling exposed information - III

Scoping

- Content scoping can be used in push-based applications (e.g. requests for emergency intervention) to limit the spread of information to the region of interest.
- This can be particularly useful for example in the aftermath of a disaster to make efficient use of scarce network resources.
- Use of scoping information may raise concerns of DoS attacks by users maliciously setting larger scopes than needed to maximise impact on network resources.

Handling exposed information - III

Scoping

- Content scoping can be used in push-based applications (e.g. requests for emergency intervention) to limit the spread of information to the region of interest.
- This can be particularly useful for example in the aftermath of a disaster to make efficient use of scarce network resources.
- Use of scoping information may raise concerns of DoS attacks by users maliciously setting larger scopes than needed to maximise impact on network resources.

Content format

- Content format may be used by content providers to distinguish different versions of a content in order to serve most appropriate content version for the requesting client.

Implementation implications

Realising exposure of information has important implementation implications:

Implementation implications

Realising exposure of information has important implementation implications:

- Information exposed in content names may result in excessive header size.

Implementation implications

Realising exposure of information has important implementation implications:

- Information exposed in content names may result in excessive header size.
- Variable lengths of exposed information may hinder line speed operations.

Implementation implications

Realising exposure of information has important implementation implications:

- Information exposed in content names may result in excessive header size.
- Variable lengths of exposed information may hinder line speed operations.
- Utilizing information exposed in content names increases the processing load at in-network devices.

Conclusions

With this work, our intention is to start a discussion about the importance of *information exposure* in the design of naming schemes and name resolution systems.

We showed that information exposure considerations are of great importance as they can lead to both desirable and undesirable features.

We identified a set of information elements whose exposure to network entities can benefit network operations and analysed implementation implications.