Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 4, 2018

P. Pillay-Esnault
A. Clemm
Huawei
D. Farinacci
lispers.net
D. Meyer
Brocade

July 3, 2017

Requirements for Generic Identity Services in Identity Enabled Networks
draft-padma-ideas-req-grids-01

Abstract

This document describes requirements for the Generic Identity Services infrastructure for Identity-Enabled Networks.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

This document specifies requirements for Generic Identity Services (GRIDS) that provide a cornerstone of Identity-Enabled Networks. GRIDS includes services to maintain mappings between Identifiers and Locators and to resolve mappings by Identifier. In addition, GRIDS includes services to manage the lifecycle of Identifiers as used in an Identity-Enabled Network, specifically services to register Identifiers.

There are additional services that GRIDS can be extended with. Examples include services to maintain metadata about endpoints that are referenced by Identifiers as well as support for Identity-based network access control. Because those services enable a lot of value-added functionality, important requirements for those services are specified here as well. In order to not overburden GRIDS development, this document focuses on core requirements.

The requirements are rooted in and derived from the problem statement [IDEAS-PS] and use case documents [IDEAS-USE][IDEAS-IDENTITY] for

Identity Enabled Networks. A gap analysis of existing solutions can be found in [IDEAS-GAP].

2. Specification of Requirements

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].

3. Definition of Terms

This document makes use of terms which for the most part have been already defined in the problem statement draft of IDEAS [IDEAS-PS]. They are included here for reader convenience. In case of any discrepancies between the two drafts, the problem statement draft overrides.

- Entity : An entity is a communication endpoint. It can be a device, a node, or a (software) process, that needs to be identified. An entity may have one or multiple Identifiers simultaneously. An entity is reached by the resolution of one or more of its Identifiers to one or more Locators.

- Entity Collection: A set of entities with its own Identifier, e.g., a multicast group, or an ad-hoc vehicular network that needs to be uniquely identified (e.g., a train entity may represent a Closed User Group (CUG) and may contain all the passengers' devices that share the same fate for connectivity).

- Generic Identity Services (GRIDS): GRIDS is a set of services to manage the lifecycle of IDs, to map and resolve Identifiers and Locators, and to associate metadata (META) with entities and entity collections. It is a distributed system that stores the ID, the associated LOC(s), and metadata (META) in the form of tuples (ID, LOC, and META).

- GRIDS-IS (GRIDS Identity Services): The subset of GRIDS that is responsible for managin the lifecycle of Identifiers and Identities.

- GRIDS-MS (GRIDS Mapping Services): The subset of GRIDS that is responsible for mapping and resolving Identifiers and Locators.

- GRIDS-SS (GRIDS Subscription Services): The subset of GRIDS that lets clients subscribe to information updates.

- Identifier (IDf): denotes information to unambiguously identify an entity within a given scope. There is no constraint on the

format, obfuscation or routability of an IDf. The IDf may or may not be present in the packet whose format is defined by ID-based protocols.

- Identifier-based (ID-based): When an entity is only reachable through one or more communication access then a protocol or a solution is said to be ID-based if it uses an ID-LOC decoupling and a mapping system (MS) as base components of the architecture.

- Identity (IDy): the essence of "being" of a specific entity. An Identity is not to be confused with an Identifier: while an Identifier may be used to refer to an entity, an Identifier's lifecycle is not necessarily tied to the lifecycle of the Identity it is referencing. On the other hand, the Identity's lifecycle is inherently tied to the lifecycle of the entity itself.

- Identity-capable (ID-capable): An application is said to be ID-capable if it makes use of an Identifier of an entity to establish communication. For example, an application that initiates its sessions using an ID. An application may use an IP-address as a proxy for an ID if the network resolves this ambiguity. We regard such an application as being ID-capable.

- IDentity Enabled Networks (IDEAS): IDEAS are networks that support the Identifier/Locator decoupling. Reaching an entity is achieved by the resolution of Identifier(s) to Locator(s).

- Locator (LOC): denotes information that is topology-dependent and which is used to forward packets to a given entity attached to a network. An entity can be reached using one or multiple Locators; these Locators may have a limited validity lifetime.

- Metadata (META): Metadata is data about an Identity. The metadata may contain information such as the nature of the entity for example.

- Scope: denotes the domain of applicability or usability of an ID. A scope may be limited (e.g., considered local with geographic proximity, or private within an administrative domain) or be global.

- User Equipment (UE): A user equipment is an entity per definition in [IDEAS-PS]

4. Background

Identity-Enabled Networks introduce the concept of Identity into networking. This concept includes an Identity/Identifier split, which complements existing Locator/Identifier separation technologies.

Identity-Enabled Networks are enabled by a set of core services that are provided by common control infrastructure. Both the services and the infrastructure that provides them are referred to as GRIDS, GeneRic IDentity Services. GRIDS comprises several key components. Those components include the following:

- GRIDS-MS (Mapping Services) provides services to maintain and resolve mappings between Identifiers and Locators.

- GRIDS-IS (Identity and Identifier Services) provides services to register Identifiers and maintain bindings between Identifiers and Identities, as we well as manage their overall lifecycle.

- GRIDS-SS (Subscription Services) provides services that let clients subscribe to updates regarding mappings and Identifiers that they are interested in.

- GRIDS-Meta (Metadata Services) provides services to manage metadata about Identities and Identifiers.

The requirements defined in this document do not imply a particular solution. Specifically, they do not imply that infrastructure used to address those requirements would need to be defined or built from scratch. Instead, where possible, existing technologies, components, and services will be used to address the requirements defined in this document. Also, it should be noted that it is possible to introduce additional components that provide value-added functions. One example would be Grouping Services that support groupings of entities and include mechanisms needed to manage Entity Groups.

In the following, requirements are denoted by REQ-xx=n, where "xx" refers to a specific requirements section and "n" refers to the number of the requirement. In some cases, optional requirements are specified and designated as OPT-xx-n. Non-requirements (i.e. aspects that might be considered candidates for requirements, but that are specifically not required to be supported at this point for various reasons) are designated as NON-REQ-xx-n.

5. Requirements for Generic Identity Services (GRIDS)

5.1. Mapping Services

REQ-MS-10: GRIDS-MS needs to maintain mappings between Identifiers and Locators.

REQ-MS-20: GRIDS-MS needs to provide services that allow clients to resolve a Locator for a given Identifier.

REQ-MS-30: GRIDS-MS MUST be able to support different models for authoritative mapping ownership, authorizing only the legitimate owner (or an entity acting on the owner's behalf) to update mapping data. Specifically, GRIDS-MS MUST be able to support (1) a model in which clients of a certain Identity can update mapping data for their Identifier, and (2) a model in which clients with a certain Locator can update mapping data with that Locator.

REQ-MS-40: GRIDS-MS MUST be able to support policy-based authorization for access to mapping services and to mapping information that is associated with specific Identities. Authorization MUST be provided on the basis of the client's identity that is accessing the service, or (in the case of an intermediary client such as a tunnel gateway) on whose behalf the service is being accessed.

Not every client will be entitled to every piece of mapping information. This allows GRIDS to be set up such that information is only available on a "need-to-know" basis to clients, facilitating the protection of private information for systems involved.

5.2. Identity Services and Identifiers

REQ-IS-10: GRIDS MUST support IDfs and IDys with the following characteristics

- Variable length ID

- Fixed length

- Structured

- Unstructured

REQ-IS-20: GRIDS MUST provide proper separation between the concepts of "Identity" and "Identifier".

An Identity is synonymous to the being of an entity that can communicate in an Identity-Enabled Network. Identity information needs to be strongly secured and is generally kept private. Identity is represented by a special type of Identifier that is not expected to ever be exposed over-the-wire in regular data plane communications in the network. An Identifier, on the other hand, is a reference to an Identity respectively associated Entity. An Identifier will in generally be public and constitutes how an Identity will be known to others, including other Entities that wish to communicate with the Entity designated by the Identifier. Identifiers MAY also be included in data plane packets.

An Identity can be associated with multiple Identifiers. It should be noted that Locators are associated with Identifiers, not Identity.

An Identity does require a representation itself, which resembles in effect a "special" Identifier. Therefore, one question that is often asked concerns how Identifier and Identity really differ. One way in which to asnwer is that a regular Identifier always refers to another Identifier, whereas the Identity does not. In that sense, the Identity constitutes the root of a "tree" (generally flat with one level of hierarchy only, but not precluding multiple levels) of Identifiers that all belong to and reference the same Identity.

REQ-IS-30: GRIDS MUST support IDfs that refer to User Endpoints of a given Identity.

REQ-IS-40: GRIDS MUST support a model in which multiple Identifiers can be associated with the same Identity. GRIDS-IS MUST NOT have inherent limitations with regards to the number of Identifiers that may be simultaneously associated with the same Identity.

REQ-IS-50: GRIDS-IS MUST support the secure registration of new Identities.

"Secure" refers to mechanisms such as strong encryption and mutual authentication.

REQ-IS-60: GRIDS-IS MUST support the secure unregistration / revocation of an Identity

REQ-IS-70: GRIDS-IS MUST support the registration of new Identifiers (independent of registration of the associated Identity)

REQ-IS-80: GRIDS-IS MUST support the unregistration / revocation of Identifiers (independent of unregistration of the associated Identity)

REQ-IS-90: GRIDS MUST allow for the possibility to support other IDs (i.e. IDs not tied to the Identity of a User Endpoint) in the future, such as Group IDs.

REQ-IS-100: GRIDS-IS MUST support a model in which Identifiers are registered by a client representing the Identity that the IDf is associated with (e.g., a User Endpoint). GRIDS-IS MUST provide mechanisms that prevent usage of identifiers in ways that result in amgibuities with regards to determining an Identifier's associated Identity. To this end, GRIDS-IS MUST either prevent duplicate assignment of IDfs, specifically assignment of the same IDf to multiple Identities, or in case duplicate assignment occurs, ensure that an IDf's associated Identity is clear depending on the context, such as a local scope.

It is to be determined whether GRIDS-IS should prevent recycling of IDfs that had been assigned previously, even if since unregistered, or if it should provide a warning when such an IDf is reassigned.

REQ-IS-110: GRIDS-IS MUST support a model in which Identifiers are assigned and registered by an authority.

REQ-IS-120: GRIDS-IS MUST support the notion of an Identifier preference, providing a service that allows a client to resolve which Identifier it should when directing communication at a given destination. The Identifier used can be simply the same Identfier used by the client to refer to the destination in the resolution request, or it can be an alternative Identifier, such as an ephemeral Identifier. This capability SHOULD be provided in a manner that is integrated with GRIDS-MS, combining the resolution of Identifier with Locator information.

Such a capability is useful to enable anonymization of communciation traffic by obfuscating identifiers. For example, a client could request a Locator for a given, well-known Identifier for a destination, such as an Identifier listed in a public directory. GRIDS could indicate to not use the well-known Identifier, but (for example) an ephemeral Identifier instead, returned (for example) together with a Locator in response to a mapping resolution request.

REQ-IS-130: GRIDS-IS MUST be able to support different models for authoritative ownership of Identifier preferences, authorizing only the legitimate owner (or an entity acting on the owner's behalf) to update preference data. Specifically, GRIDS-IS MUST be able to support a model in which clients of a given Identity can update their own Identity preference data.

5.3. Subscription Services

REQ-SS-10: GRIDS-SS MUST allow clients to subscribe to updates for information that they are entitled to resolve. Specifically, GRIDS-SS MUST provide support for pushing updates about Locators for mappings that are of interest to a client with minimal incurred delay. GRIDS-SS MUST also provide suppport for pushing updates about preferred Identifiers of entities to whose mapping information the client is subscribed to.

5.4. Metadata Support and Services

Metadata can be tremendously useful for Identity-Enabled Networks, as indicated in both Problem Statement and Use Cases. Therefore, GRIDS SHOULD support Metadata Services (GRIDS-Meta) that allow to store and retrieve certain metadata associated with Identities, as well as metadata associated with Identifiers. The metadata supported has several properties in common:

- It is slow changing and does not impose significant requirements with regards to update rates that would have to be supported

- It does not impose significant requirements with regards to latency of propagation of updates

- It is low in size and volume

GRIDS-Meta will have to support requirements that include the following:

- Req-Meta-10: When GRIDS-Meta is supported, it MUST provide support for associating metadata with a given Identity. An example of metadata associated with an Identity is the type of endpoint (e.g. a mobile device, an IoT device, or a compute server).

- Req-Meta-20: When GRIDS-Meta is supported, it MUST provide support for associating metadata with a given Identifier (that is not automatically associated with other Identifiers that belong to the same Identity). An example of metadata associated with an Identifier would be information about which Groupings the Identifier belongs to, or whether the Identifier is considered a publicly known Identifier that should, for example, be listed in a public directory.

- Req-Meta-30: When GRIDS-Meta is supported, it MUST provide support that allows a client to retrieve metadata for an Entity as identified by a given Identifier. The retrieved metadata should include both metadata associated with the particular Identifier,

and metadata associated with the Identity that is referred to by the Identifier.

- Req-Meta-50: GRIDS-Meta MUST support for differentiation between public metadata that is generally accessible and can be shared across GRIDS boundaries, and private metadata that is accessible only on a need-to-know basis.

- Example of private metadata includes any metadata that ties an identity to personal information (such as customer data regarding the real-world owner of a communications entity.) Example of public metadata includes metadata such as the type of endpoint (e.g. a mobile device, an IoT device, a compute server), or which Identifier constitutes a publicly known Identfier that should be listed in publicly accessible directories.

- Req-Meta-60: When GRIDS-Meta is supported, it MUST support a notion of ownership of metadata, and give the owner of the metadata full control over security rules that guide who can access that metadata.

- Req-Meta-70: When GRIDS-Meta is supported, it MUST support the definition and enforcement of security policies that guide who is authorized to retrieve metadata, and who is authorized to modify metadata.

5.5. Distribution and Redundancy

REQ-DR-10: GRIDS MUST be robust and very highly available.

REQ-DR-20: Any maintenance or upgrades to GRIDS MUST NOT affect availability of GRIDS services.

REQ-DR-30: GRIDS MUST support implementation using a distributed and redundant architecture. Specifically, failure of individual components MUST NOT bring down GRIDS as a whole.

As this is a requirements document, this document does not mandate a particular implementation architecture. That said, it should be noted that for any mapping system to be successful, it will need to be robust, distributed and provide redundancy. The mapping system design and architecture must avoid being single points of failure and MUST enforce resiliency.

Furthermore, it should be noted that the format of the Identifier may or may not play a role in how any underlying servers used to implement GRIDS might be distributed. It is conceivable that such

distribution and placement of GRIDS components and data maintained by GRIDS will be affected by usage patterns.

5.6. Scale and Performance

REQ-SP-10: GRIDS MUST support very large (Internet-level) scale.

It is anticipated that GRIDS MUST be able to handle from the start billions of distinct Identifiers and mapping entries and allow for substatiantial future growth. While this document makes no statements about GRIDS architecture, it should be noted that GRIDS will likely not be provided by monolithic infrastructure but by means of multiple distributed and interconnected components.

REQ-SP-20: GRIDS MUST scale in a way such that increases in the number of Identifiers and mapping entries do not negatively degrade performance. Performance characteristics SHOULD be independent of scale. If such constant scale performance characteristics cannot be provided, performance MUST NOT degrade in worse than logarithmic manner based on the number of Entities.

REQ-SP-30: A characterization of GRIDS performance at scale, as well as associated GRIDS performance objectives, MUST include the following parameters:

o TR: Time to resolve a Locator by Identifier, in three variants to characterize normal case, performance determinism, and "bottom case" behavior:

   * mean

   * variation

   * bottom percentile

o TM: Time to update a mapping entry (i.e. time until mapping entry first registers with GRIDS), in three variants to characterize normal case, performance determinism, and "bottom case" behavior:

   * mean

   * variation

   * bottom percentile

o TS: Time for mapping entry update to propagate to subscribers of a given mapping (i.e. clients who are subscribed to be notified of mapping updates of a given Identifier), in three variants to

characterize normal case, performance determinism, and "bottom case" behavior:

- mean

- variation

- bottom percentile

o SRT: Sustained resolution throughput for resolution requests, in multiple variants to distinguish overall throughput and throughput as perceived by individual clients:

- overall

- for individual client

o SRM: Sustained mapping update throughput, in multiple variants to distinguish overall throughput and throughput as perceived by individual clients:

- overall

- for individual client

REQ-SP-40: Characterization of performance MUST furthermore include information on scale at which the performance numbers are observed, such as number of Identifiers.

It is acknowledged that specific implementations may differ in terms of performance characteristics they can accomplish. Specific performance objectives against these parameters MAY be articulated at a later point. It is possible that such objectives will depend on the use case and that such use cases could result in specific qualification requirements imposed on GRIDS implementations for particular deployment scenarios. Furthermore, it is acknowledged that additional performance parameters can be articulated in addition to the ones specified here.

It should be noted that this document does not mandate a particular implementation architecture. However, in order to be able to meet the ambitious performance and scale requirements imposed by GRIDS, we note that an architecture that leverages principles of distribution, hierarchy, and aggregation may help to achieve these goals. Specifically, we note that in order to meet low latency goals, architectural considerations SHOULD include support for predictive and proactive dissemination and caching of data to locations that are close to clients that need to consume and interact with it.

Conceivably, this may also involve application of data analytics and machine learning techniques.

5.7. GRIDS Security

REQ-SEC-10: GRIDS needs to be robust against direct and indirect attacks. If any component of GRIDS is attacked, the system needs to degrade gracefully.

REQ-SEC-20: GRIDS The addition and removal of components of the mapping system must be performed in a secure matter so as to not violate the integrity and operation of the system and service it provides.

REQ-SEC-30: GRIDS MUST authorize any requests directed at it. This includes requests that alter data maintained by GRIDS, as well as requests to retrieve data from GRIDS.

REQ-SEC-40: GRIDS MUST authenticate clients.

REQ-SEC-50: GRIDS MUST support some sort of audit trails. Specifically, GRIDS SHOULD log any requests being served and retain such logs, themselves properly secured, for a minimum (to-be-determined) time interval. In addition, GRIDS SHOULD at a minimum support per-client statistics (such as counter and rate information about resolution requests) and per-Identifier statistics (such as counters for accesses involving a specific Identifier).

REQ-SEC-60: Any Identity information MUST be encrypted. Specifically, Identity information MUST NOT (i.e., must never) be transmitted in the clear between GRIDS and a client. (Note the distinction between "Identity" and "Identifier". While Identity information MUST be protected and highly security sensitive, the same stringent requirements generally do not apply to Identifiers.)

In addition, Identity information MUST NOT be included in dataplane communications.

OPT-SEC-70: Encryption of GRIDS messages is optional. Specifically, it is optional to provide confidentiality of the requesters and the information they are requesting. (Note the exception regarding Identity information; Identity information MUST always be encrypted).

REQ-SEC-80: GRIDS MUST support cryptographic signing of information that it provides to allow clients to verify if the provided information is authentic.

REQ-SEC-90: GRIDS MUST support message rate-limiting and other heuristics must be part of the foundational support of the mapping system to protect GRIDS from sudden overloaded conditions and mitigate the effects of potential attacks.

5.8. Ability to support multiple instances

REQ-MI-10: GRIDS SHOULD be deployable in a private space and provide data isolation. For example, GRIDS MUST NOT require a company to expose all of its IDf as public IDfs if the company does not wish to do so.

Because Identifiers are unique only within a given GRIDS instance, it should be noted that by using multiple isolated instances of GRIDS, it is conceivable that overlapping IDfs can be supported. However this is not encouraged. One way in which this can be avoided is by by allocating private ranges for experimental use in the IDf name space, and requiring GRIDS to not assign any IDfs in an allocated Identifier space.

REQ-MI-20: GRIDS MUST support a distinction between "private" GRIDS data that is refined in scope to a given GRIDS instance, and "public" GRIDS data whose scope can be global. Specifically, private GRIDS data MUST NOT be shared beyond GRIDS boundaries, whereas public GRIDS data can be (and may have to be) shared across multiple GRIDS instances.

For example, some metadata may be private, such as metadata tieing an identity to personal information (such as customer data regarding the real-world owner of a communications entity.) Other metadata may be public, such as the type of endpoint (e.g. a mobile device, an IoT device, a compute server) that is associated with an entity. Likewise, the list of Identifiers that are in use or "claimed" constitute public GRIDS data (but not who those Identifiers are assigned to).

5.9. GRIDS Extensions

GRIDS MUST be designed in such a way to allow future extensions and services.

An example of a future extension concerns grouping services, involving Group IDs that represent groupings of User Endpoints. There are multiple applications as well as multiple types of groupings, for example administrative groupings (used to facilitate management), groupings that represent collections of User Endpoints that temporarily or permanently share the same fate (such as devices in the same railroad car that all use a communications gateway with

the same locator), and groupings that represent multihomed endpoints (which include endpoints that mutually protect each other in case of failures).

The following are examples of requirements that GRIDS will have to support if grouping is to be supported as a feature. It should be noted the following list is incomplete, merely indicative of the types of requirements that will be associated with providing Grouping Services:

- GRIDS SHOULD support group identifiers, used to designate groupings of endpoints.

- GRIDS SHOULD support Group ID (G-ID) Management Services: Adding and removing identifiers from the group, as well as querying group members.

- GRIDS SHOULD support a type of group used to designate a group of endpoints that share the same fate, i.e. that are (temporarily or permanently) assoicated with the same Locator. GRIDS Grouping Services SHALL integrated with GRIDS-MS in such a way that for an Identifier that is part of a group, the Locator of the Group takes precedence over (or determines) the Identfier's "native" Locator (which it would be associated with, if not part of the group).

6. Security Considerations

Due to the sensitivity of Identity tied to Identifier and location data there is a need to pay attention to security ramifications. In particular, the security goals should include confidentiality, possible encryption for integrity of sensitive data and privacy.

7. IANA Considerations

This document has no actions for IANA.

8. Contributors

This present document is based on an extract of the first version of the draft. The authors and their affiliations on the original document are: D. Farinacci (lispers.net), D. Meyer (Brocade), D. Lake (Cisco Systems), T. Herbert (Facebook), M. Menth (University of Tuebingen), Dipenkar Raychaudhuri (Rutgers University), Julius Mueller (ATT) and Padma Pillay-Esnault (Huawei).

There are two companion documents that were extracted from the -00 version of this document: Problem Statement in IDEAS [IDEAS-PS] and GRIDS Requirements [IDEAS-USE] which regroups all the authors above.

Uma Chunduri

Yingzhen Qu

Rutgers University: Parishad Karimi and Shreyasee Mukherjee

9. Acknowledgments

This document was produced using Marshall Rose's xml2rfc tool.

10. References

10.1. 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, <http://www.rfc-editor.org/info/rfc2119>.

10.2. Informative References

[IDEAS-GAP]
Qu, Y., Caballos, A., Moskowitz, R., Liu, B., and A. Stockmayer, "Gap Analysis for Identity Enabled Networks", July 2017, <https://tools.ietf.org/html/draft-xyz-ideas-gap-analysis-00/>.

[IDEAS-IDENTITY]
Chunduri, U., Clemm, A., and M. Menth, "Identity Use Cases in IDEAS", June 2017, <https://tools.ietf.org/html/draft-ccm-ideas-identity-use-cases-00/>.

[IDEAS-PS]
Pillay-Esnault, P., Boucadair, M., Jacquenet, C., Fioccola, G., and A. Nennker, "Problem Statement for Identity Enabled Networks", July 2017, <https://datatracker.ietf.org/doc/draft-padma-ideas-problem-statement/>.

[IDEAS-USE]
Pillay-Esnault, P., Farinacci, D., Herbert, T., Jacquenet, C., Lake, D., Menth, M., Meyer, D., and D. Raychaudhuri, "Use Cases for Identity Enabled Networks", July 2017, <https://datatracker.ietf.org/doc/draft-padma-ideas-use-cases-01/>.

Authors' Addresses

Padma Pillay-Esnault
Huawei
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: padma.ietf@gmail.com

Alexander Clemm
Huawei
2330 Central Expressway
Santa Clara, CA 95050
USA

Email: ludwig@clemm.org

Dino Farinacci
lispers.net
San Jose California
USA

Email: farinacci@gmail.com

Dave Meyer
Brocade

Email: dmm@1-4-5.net