GNA – Global Network Architecture
Web: http://gna-re.net/

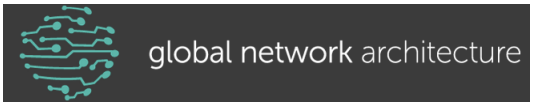

Document name:	GNA Technical Working Group – Technical Notes
Author(s):	Dale Finkelson, Jim Williams, Erik-Jan Bos, Guy Roberts, Dave Reese
Contributor(s):	GNA Technical Group
Date:	January 2017
Version:	v1.0

GNA Technical Working Group – Technical Notes

Table of Contents

Introduction

This 'GNA Technical Working Group - Technical Notes' document is intended to provide policy makers with background information on a set of issues the GNA Technical Working Group feels are important and worthy of discussion. This document is divided into five parts, based around the following subjects:

1. Encryption
2. Security
3. Acceptable Use Policy
4. Privacy
5. Sharing Traffic at Global Exchange Points

It should be noted that in the development of the GNA Technical Working Group set of documents, the working group found that these areas fit both in the technical space and the policy space. While the working group has technical opinions in these areas, the group also understands that these are primarily policy matters and are out of scope of the GNA.

Encryption

Encryption is a service that is agreed between and provided by end-users or networks, rather than by the GNA. The GNA Technical Working Group offers some ideas about how such a service might be provisioned and why it might be useful. Encryption of data can occur both at the application layer and at the link layer.

Encryption at the application layer includes many applications, from e-mail to secure data transfer. We note that application layer encryption is not within the scope of the GNA Technical WG; it is an application that can be provided by the end-user or by the REN directly serving the end-user.

There are two forms of link layer encryption service: authenticated and non-authenticated (aka, opportunistic). Authenticated encryption requires some form of key management, which must by necessity be controlled by the owner of the end equipment. By contrast, opportunistic encryption relies on a TOFU (trust on first use) model where the key of the endpoint is trusted the first time it is encountered (but prevented from being changed by later connections to the same endpoint). In effect, the endpoint is identified by its (often public) key. This model can be applied to several different technology stacks, including MPLS and various VPN technologies.

Link layer encryption can employ specific end-to-end equipment at each network entry/exit, but it can also be implemented using standardized encryption techniques such as IPSec. Specific equipment recommendations or requirements are not within the scope of the GNA. For instance: https://tools.ietf.org/html/draft-farrelll-mpls-opportunistic-encrypt could be a topic for further discussions but is not in scope at this time.

There are examples (e.g., medical data) where security considerations may require (authenticated) link layer encryption, such as between a medical facility and the data repository. These cases may be within NREN services or they may span multiple NRENs; service details must be negotiated as NREN-to-NREN agreements.

Opportunistic link layer encryption can be an effective shield against pervasive monitoring by making it too costly to attack a wide range of targets (forcing the opponent to focus on the "interesting" targets). Pervasive monitoring is described as an attack on the Internet in BCP 188 (RFC 7258). Low cost opportunistic encryption may therefore become a valid use case for the NREN community.

By contrast, the costs of key management in authenticated link layer encryption services often push such applications to high-value use-cases (e.g., medical data). For example, patient-identifiable medical data is collected in the field, transferred to a local collection site and then bulk transferred to an analysis site. Given the sensitive nature of this data, authenticated link layer encryption may be required.

Security

Security is of utmost importance to the Research and Education (R&E) networking community.

This section on security discusses a possible set of services that networks provide to other networks relating to security issues (Network to Network Security Service - NNSS). This section also describes the set of services each network could provide internally (within their network) (Internal Security Service - ISS). This is only one possible structure for deploying security services. Each GNA participating NREN or Global Research and Education Exchange Point (GXP) will configure their security services as best fits their environment.

Security will be further discussed in a comprehensive document being prepared by the Security Working Group (under auspices of the Global R&E Network CEO Forum). That document will greatly expand and enhance or perhaps even obsolete this outlined (NNSS and ISS) structure. This set of notes represents general thoughts of the GNA Technical Working Group, which have been forwarded to the Security Working Group.

Within each network, it is expected that the Internal Security Service (ISS) will assure that all (network and network supporting) equipment is up-to-date with regard to security patches. The ISS is also expected to periodically probe end-user equipment to verify security levels.

Between networks the Network-Network Security Service (NNSS) is responsible for investigating and acting on (if necessary) requests from other networks (such as DDoS attacks/complaints).

NNSS in coordination with ISS investigates complaints associated with end-user misbehavior.

If possible the ISS and the NNSS should be staffed or contactable 7x24x365, through a well-known and published set of contact details. They should have appropriate software tools and policy authority to take immediate action when required. Each NREN or GXP will have a security structure that best fits its environment.

Acceptable Use Policies

It is desirable and strongly recommended that any Acceptable Use Policies (AUPs) in effect on any network component participating in the GNA should be published and easily found. The GNA itself will not add any AUP to any resource.

Among the AUP issues that networks need to address is the need for consistent access to cloud services, CDNs and other (potentially commercial) services to the R&E community. There is a general recognition that access to these services has become very important to organizations in the pursuit of science, research, and education. It is critical is that documentation exists that specifies how NRENs, link owners, and GXP operators treat this traffic.

Individual NRENs, link owners, and GXP operators can have or impose AUP restrictions within their networks. Organizations are free to restrict internal connections and traffic as they need. However these policies should be clearly published.

Privacy

Background Information

There is a need for operators to collect information about the traffic that is on a circuit or a port. There are several reasons for collecting this information:

1) To understand the utilization of a port or circuit in order to know if and when capacity upgrades are necessary.
2) For diagnostics; specifically, the ability to diagnose a problem reported on a circuit or port. It is useful to be able to determine exactly which circuit is being used by a given individual if they report a problem with an end-to-end connection
3) For security reasons, should there be a Denial of Service attack, or should other activity occur, it is useful to understand the origin and destination of the traffic and possibly gain more in-depth information about the actual content of the traffic (packets).
4) To inform member institutions, funding agencies, or other interested parties about how a facility is being used.

Related to these various uses there are several types of data that can be collected, some of which are less of a privacy concern than others.

1) Simple aggregate utilization data

   This would be simple reporting on the overall amount of traffic that is traversing a port or a circuit. There is no differentiation of the traffic by source, destination or protocol. This is a simple measure of bytes in and bytes out. Generally making this data available is not considered

a privacy concern except in the case where usage of the port is extremely limited. This information is useful in understanding traffic patterns and the need for upgrades.

Utilization data can be broken down into per-VLAN accounting. In this instance it may be possible to use the data to give very rough characterization by organization or by project of the use of the network.

2) Flow Data

   Collection of Flow Data is much more problematic. Here data is collected that contains IP source and destination address of the traffic, protocol used for the traffic, and other information in the header and body of the message. It is this sort of data that presents the greatest privacy challenge.

3) Other Raw Data collections, including "taps" of full packet payload for research, security or other purposes.

Occasionally a full capture (packet capture) is used to debug a particular network problem or to support a network research project. Some network security organizations can request full network captures to support network security research.

It is common practice for the engineering teams of R&E network operators to inspect data on their networks for operational and security purposes. Mature NRENs will have established policy expectations for access to such data. As a result of multi-NREN operations that are common in the R&E community, there are occasions where NOC personnel of NRENs share this data (which could identify individuals or their activities). In some parts of the world, governmental funding agencies require information about how their funded resources are being used.

The above set of operational issues can conflict directly with established privacy principles and expectations. Rather than try to determine a 'point solution' to this problem this document suggests a range of acceptable solutions, a 'solution band'. This provides NRENs with guidance, but does not dictate direction.

Recent revelations of widespread government access to network information, both directly and through commercial partners raises questions as to the role that NRENs should assert as the right position for the NREN community. Should information that could be used to personally identify individuals be shared with third parties? With Governments? Should leading NRENs declare their security and privacy standards so that other NRENs can align their policies (or route around questionable or overly restrictive ones)? As our networks become increasingly interconnected, is a patchwork of privacy policies acceptable?

Some principles:

1) The academic community that is supported by the R&E networks is committed to an open exchange of ideas with a reasonable expectation of privacy among community members.

2) Each NREN should publicly state its privacy policy related to operational and security data sharing on its websites.

3) Each GXP operator participating in the GNA should clearly state/publish its policy regarding sharing data with any third party without the prior written consent of both network parties who exchange data across the GXP.

Internal Operator Privacy Models

There are several operational data collection levels impacting privacy. This document outlines two.

1) Aggregate level accountability. The only data that any NREN NOC can collect is aggregate traffic levels. This means traffic volume as Gbps across time. This is simple and only provides information for the NOC to optimize traffic load/routing at the aggregate level. It provides no information regarding the individual experimenter/flow regarding flow level performance. This is just Utilization Data and as such does not generally present a privacy concern. But it does allow for understanding the overall use of the infrastructure.
    a. However, when the accounting is done at the VLAN level, it is possible to present data on the use made by a specific project which may well correlate to a specific location. While the group using that VLAN may well care to know the information about use, it is a potential privacy concern.
    b. Note that if data is being collected at all edge ports as well as in the core it will be possible to determine the overall use an organization is making of the infrastructure.
2) Flow level accountability. Data is captured at the flow level by some mechanism at access points in the network and is archived for some length of time. Individual flows can be tracked from end-to-end.
    a. This data is of valuable for diagnostic or security issues, but it likely not needed for overall capacity management purposes.
    b. Flow level data present significant privacy concerns due to the detailed level of the data collected, as mentioned above.
        i. This does not mean an organization should not be able to collect such data, only that it needs to have well defined, understood, and published policies about storing, securing and sharing this data.

The critical privacy related questions are: what level of data it is acceptable (or reasonable) to collect and maintain? and, what are the access policies for that data?

It is important to realize that understanding the privacy issues around this data is not a technical problem. If there are technical approaches to making the information in the data anonymous doing so does not address the concerns of data privacy.

Some Suggested Operational Privacy Principles

1) It is acceptable and encouraged for an organization managing a GXP to collect utilization data on the circuit and port. These would be simple counts of bytes in and out or packets in and out as well as error counters—basically high level port (utilization and error) statistics. It is equally acceptable to publish that information on a GXP's website. Aggregated GXP statistics, such as the total volume of traffic through the Exchange, could be used as a marketing tool on the GXP's website without any problem.
2) Upon request by VLAN participants, a GXP operator may collect information on VLAN utilization and report it to the consortia partners for the circuit that VLAN traverses in order for them to make a determination if re-alignment is necessary. But unless specifically authorized by the organizations, this information should not be shared with any party other than the link owners or the GXP's internal operations team. The GXP operator should be transparent and timely with each link operator about every party with access to data.
3) Flow level data is where the challenge occurs. There are certainly times such information is of critical use to both the operators of a circuit or exchange and the end organization. Given the critical nature of flow data, it is necessary for a GXP Operator to collect flow data, whether done per packet or done by sampling packets:
    a. Under no circumstances should this data be made accessible to the public.
    b. This data should only be collected where strictly necessary for operational purposes.
    c. Under no circumstances should this data be retained longer than is reasonable and strictly necessary.
    d. Under no circumstances will this data be shared with any governmental body, commercial vendor or research organization without either the prior consent of *both sides of any peering connection or* required legal demands.
    e. All policies regarding this data should be made available to all participants.
4) Sensitive operational data as described above should be appropriately protected (as should all operational data) and discarded immediately as per 3.c.

Access to and Exchanging Traffic at Global Exchange Points in the GNA

The openness of a GXP means that any organization can request a connection to the GXP fabric. The requester is accepted as a new connector by the owner/operator of the GXP, within the published connection policy, given the requester will obey the conditions that the GXP operates under, including paying the bills. "Open" does not mean that it is "free" to connect. In almost all cases, there is a set-up fee and a monthly recurring fee that is payable to be allowed to connect to the GXP. Although it may vary, the existing GXPs within the GNA charge a fee that is based on the physical interface type. Also, one can expect to incur cost for using the colocation facilities at the GXP.

Having an operational port on the GXP, by no means gives the connector any right to exchange traffic (such as a Layer 2 interconnect or a BGP peering) with any of the other connectors at the GXP: Exchanging traffic between two or more connectors on the GXP requires the parties to explicitly and mutually agree to exchange traffic (bi-lateral peering).

It is in the interest of users and projects in the R&E community that a healthy, resilient, and fair R&E network ecosystem exists, wherein R&E networks do not weaken each other. It is impossible to make rules about who can peer with whom, as there is no way to disallow peering from a governance or policy body, for the simple reason that no such body exists or is likely to exist on a global scale. And even if such rules were written down, enforcing such rules would be close to impossible. Hence, the ecosystem for exchanging traffic is a self-regulating one, and no R&E network that is considered to be part of this ecosystem will purposely try to weaken the ecosystem.